

Management Support Systems

DCAP208

Editor
Dr. Kamlesh Lakhwani



L OVELY
P ROFESSIONAL
U NIVERSITY



MANAGEMENT SUPPORT SYSTEMS

Edited By
Dr. Kamlesh Lakhwani

ISBN: 978-93-87034-80-8

Printed by

EXCEL BOOKS PRIVATE LIMITED

Regd. Office: E-77, South Ext. Part-I, Delhi-110049

Corporate Office: 1E/14, Jhandewalan Extension, New Delhi-110055



+91-8800697053, +91-011-47520129



info@excelbooks.com/projects@excelbooks.com

internationalalliance@excelbooks.com



www.excelbooks.com



for

Lovely Professional University
Phagwara

CONTENTS

| | | |
|-----------------|---|-----|
| Unit 1: | Support Systems <i>Mithilesh Kumar Dubey, Lovely Professional University</i> | 1 |
| Unit 2: | Decision Support <i>Rishi Chopra, Lovely Professional University</i> | 15 |
| Unit 3: | Decision Support Systems <i>Mandeep Kaur, Lovely Professional University</i> | 29 |
| Unit 4: | Dialogue Management, Data Management and Model Management for DSS <i>Pallavi Vyas, Lovely Professional University</i> | 41 |
| Unit 5: | Modeling and Analysis <i>Rajni Bhalla, Lovely Professional University</i> | 52 |
| Unit 6: | Data Warehousing <i>Anil Sharma, Lovely Professional University</i> | 71 |
| Unit 7: | Business Analytics and Data Visualization <i>Navpreet Kaur, Lovely Professional University</i> | 94 |
| Unit 8: | Usages, Benefits and Success of Business Analytics and Data Visualization <i>Navpreet Kaur, Lovely Professional University</i> | 116 |
| Unit 9: | Data Mining <i>Anuj Sharma, Lovely Professional University</i> | 130 |
| Unit 10: | Data Mining Tools and Techniques <i>Sukanta Ghosh, Lovely Professional University</i> | 153 |
| Unit 11: | Neural Networks <i>Sukanta Ghosh, Lovely Professional University</i> | 177 |
| Unit 12: | Applications of Neural Network <i>Kamlesh Lakhwani, Lovely Professional University</i> | 193 |
| Unit 13: | Knowledge Management <i>Rajni Bhalla, Lovely Professional University</i> | 210 |
| Unit 14: | Knowledge-based Decision Support <i>Sarabjit Kumar, Lovely Professional University</i> | 234 |

SYLLABUS

Management Support Systems

Objectives: Course objectives include developing an understanding of the purpose, functions, components, and applications of transaction processing systems and management reporting systems in private and public organizations; and describing and evaluating policies for information resource management.

| Sr. No. | Topics |
|---------|--|
| 1. | Support Systems: Changing Business Environment, Managerial Decision Making, Computerised Support, Concept of Decision Support Systems, Major Tools and Techniques for Management Support System |
| 2. | Decision Support: Decision Making |
| 3. | Decision Support Systems: Definition |
| 4. | Modeling and Analysis: Modeling for MSS |
| 5. | Data Warehousing Data Warehousing Definitions and Concepts |
| 6. | Business Analytics and Data Visualization: Overview |
| 7. | Data Mining: Concepts and Applications |
| 8. | Neural Networks: Concept |
| 9. | Knowledge Management: Introduction |
| 10. | Knowledge-based Decision Support: Concepts and Definitions of Artificial Intelligence and Expert Systems |

Unit 1: Support Systems

Notes

CONTENTS

Objectives

Introduction

- 1.1 Changing Business Environment
 - 1.1.1 Results of the Changes
- 1.2 Managerial Decision Making
 - 1.2.1 Types of Management Decisions
- 1.3 Computerized Support
- 1.4 Concept of Decision Support Systems (DSS)
 - 1.4.1 Types of Decision Support Systems
- 1.5 Management Support Systems
- 1.6 Summary
- 1.7 Keywords
- 1.8 Review Questions
- 1.9 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss Changing Business Environment
- Explain Managerial Decision Making
- Discuss the Concept of Computerised Support
- Understand Decision Support Systems
- Explain the Tools and Techniques for Management Support System

Introduction

The vision that is required to maximize control over your operations is provided by Support System module. A graphical user interface consists of various facets of functionality such as online manual, Data browser, Productivity reporting, Workload analyzer. The characteristics of Support System include easy-to-read bar graphs as well as quick access to important data. By means of this system, you can look at a real-time report or “drill down” to the particulars of a specific order, SKU, receipt, etc. The Support System provides you real-time feedback on your operations. This allows you to make rapid and informed decisions. In this unit, we will discuss the concept of management support system.

1.1 Changing Business Environment

Business environments are constantly changing. New developments in technology, politics, customer preferences, and regulations happen all the time. In general, when businesses fail, it is often because they failed to respond adequately to changes in their environments.

Notes



Caution A firm must monitor changes in its environment and share information with key entities in that environment in order to stay in business.

External business environmental forces include: technology and science; economy, international change, and politics. Internal business environmental forces include: customers, suppliers, stockholders, regulations and competitors.

At least three major developments have occurred that have changed dramatically the business environment. First, technology has been developed that has made information preparation and dissemination inexpensive. This technology has taken the form of low cost, high-speed digital and cable video and data transmission, hardware that produces information quickly and easily, and the development of software that makes preparation, data, and communication tools available to individuals who previously did not have access to needed information. With these technology developments, time, space, and other temporal constraints to information have been reduced and, in many cases, eliminated.

A second major development that has significantly impacted business has been globalization. Faster methods of transportation, together with instantaneous information, have allowed the world to become one giant marketplace. Consumers can now buy products from foreign firms as easily as they can from a local store. Organizations such as General Motors have to worry not only about what Chrysler and Ford are doing, but also what Toyota, Volkswagen, and BMW are doing as well. In fact, Chrysler is not just “Chrysler” anymore. It is now a conglomeration of European, North American, and Asian manufacturers known as Daimler Chrysler. Instead of having only two major American competitors, General Motors and all other business organizations now have to compete with similar companies throughout the world. In addition, with the increased availability of inexpensive information, more is known about these competitors and about General Motors than ever before.



Example: If a General Motors product has deficiencies, the world knows about and can act on those problems instantly.

A third major change is the concentration of power in certain market investors, primarily large mutual and pension funds. Mutual funds such as Fidelity and Vanguard, and pensions funds such as CALPERS, for example, now hold major stock positions in many companies. The influence of these major market players is so significant that, if they are displeased, corporate executives will find that their positions within the company are in jeopardy. Armed with easily available and inexpensive information about investees and their competitors, large institutional investors raise the competitive bar very high and shorten the periods over which success is measured.

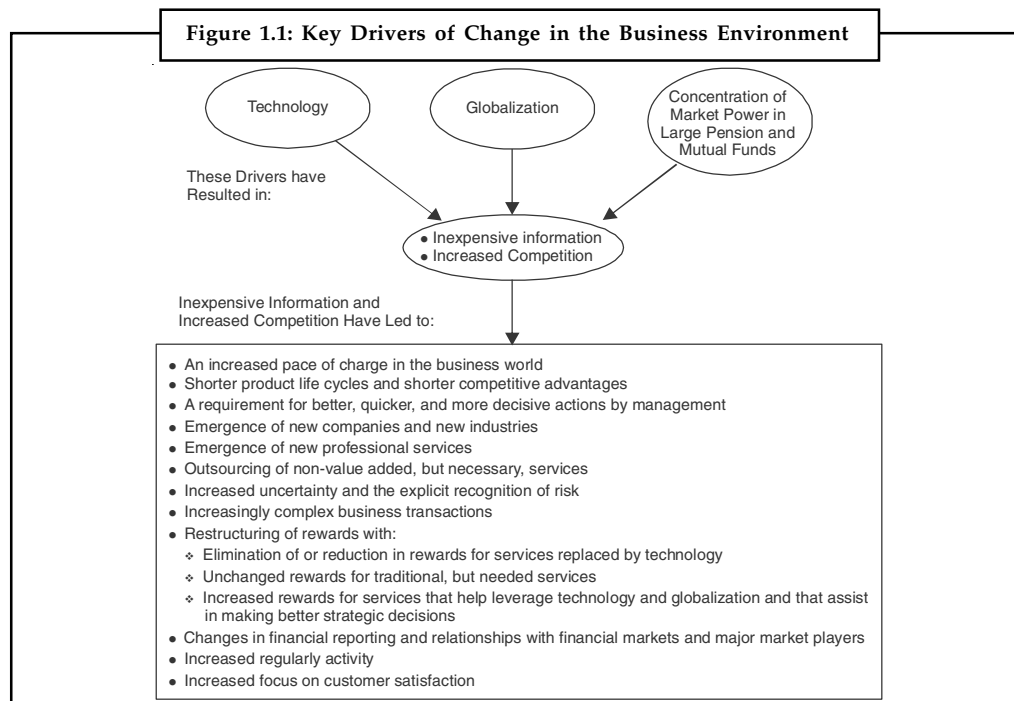
1.1.1 Results of the Changes

While these change drivers have significantly impacted everything we do, including the way we live, they have had two dramatic impacts on business. First, they have eliminated the old model that assumed information is expensive. Today anyone, armed with the right software, can be an “accountant” and produce financial information. Second, they have dramatically increased the level of competition among organizations. Institutional investors want the best performance and they want it now. Global competitors often have different cost structures that can be exploited to render historically successful business models obsolete; and since information about all organizations is widely available, only organizations that are truly the best survive and remain successful.

There have been a number of business developments because of these changes. Some of the most obvious are:

- An increased pace of change in the business world
- Shorter product life cycles and shorter competitive advantages
- A requirement for better, quicker, and more decisive actions by management
- Emergence of new companies and new industries
- Emergence of new professional services
- Outsourcing of non-value-added, but necessary, services
- Increased uncertainty and the explicit recognition of risk
- Increasingly complex business transactions
- Restructuring of rewards with:
 - ❖ Elimination of or reduction in rewards for services replaced by technology
 - ❖ Unchanged rewards for traditional, but needed services
 - ❖ Increased rewards for services that help leverage technology and globalization and that assist in making better strategic decisions
- Changes in financial reporting and relationships with financial markets and major market players
- Increased regulatory activity
- Increased focus on customer satisfaction

The following diagram illustrates the drivers of change, results of these changes, and developments that have occurred because of the changes that have taken place.



Source: <http://aaahq.org/pubs/AESv16/chapter2.pdf>

Notes

Self Assessment

State True or False:

1. Global competitors often have different cost structures that can be exploited to render historically successful business models obsolete.
2. Expensive information and increased competition have led to an decreased pace of change in the business world.

1.2 Managerial Decision Making

All managers have a shortage of knowledge, resources, and time. Working within these parameters, the management process culminates in decisions to implement various actions.

Decision-making is the focal point of all organizational dynamics, and management effectiveness is judged on the basis of the quality of these decisions.

Managerial decisions are deliberate choices made from a range of alternatives. Before making the decision, the manager must evaluate each choice according to its projected outcomes in terms of the organization's resources as well as the amount of information and time available.

Thus, every managerial decision is a best-effort compromise made in an environment of uncertainty.

1.2.1 Types of Management Decisions

From a management perspective there are three types of decisions:

1. Long-term strategic decisions concerning the external environment of the organization.
2. Administrative decisions intended to order the functions of the organization in the most cost-effective way.
3. Operational decisions designed to maximize a firm's profitability through productive procedures.

Strategic Decisions

There are several types of strategic decisions in production and operations management (P/OM):

1. **Product or service strategies:** Management decisions regarding product line market strategies (including design, quality and cost) determine production cost parameters.
2. **Process strategy:** Management decisions regarding process methods are critical in determining technological and organizational production requirements. The process strategy decision is also crucial in determining capital and financial requirements.
3. **Research and Development (R&D) strategy:** R&D is critical for organizational survival in today's rapidly changing marketplace. The R&D strategy includes total resources being devoted to the effort, the type of research to be performed including pure vs. applied research, manufacturing vs. market research, and product development vs. process development.
4. **Location strategy:** Often the success or failure of a business, production, or service is determined by a location decision.

5. **Inventory management strategy:** It is essential to develop a strategy for coordinating production needs with raw material and component inventories. However, the inventory strategy is determined by whether the demand is dependent or independent of the demand for other components. If the demand for one product, such as air conditioners, is independent of another product, such as kitchen chairs, then an independent inventory management strategy is required. However, if the overall component demand is dependent on the demand for the product, then a Material Requirements Planning (MRP) strategy is needed.



Notes MRP is a component manufacturing planning method in which items required for a manufacturing process are indexed to overall product demand. With MRP it is not essential that all inventory items are available at all times, but only when they are required in the production process. Thus, under MRP, inventory needs are coordinated with production needs.

6. **Human resource planning and management strategy:** As a rule of thumb more than 75% of a firm's operating expense is for human resources. Therefore, adequate hiring, training, and utilization of human resources is a critical operational strategy for achieving success.



Example: The management of an organization makes a strategic decision to develop a five-year marketing plan to achieve a competitive advantage through the introduction of a new service.

An automobile manufacturer makes a process strategy decision to offer a standard group of options on its automobiles in order to reduce the variation in its production needs and lower unit costs.

A computer chip manufacturer makes a strategy decision to increase R&D expenditures on an advanced CPU chip design enabling compatibility with multiple computer operating systems.

A firm makes a location strategy decision to conduct a nationwide survey of state industrial development agencies to evaluate where the company could receive the greatest financial and environmental location benefits.

A lawn mower and snow blower manufacturer makes an inventory management decision to use an MRP system to coordinate their need for lawn mower and snow blower components with seasonal manufacturing schedules.

A manufacturer makes a human resource strategy decision to give more responsibility to its employees by creating work teams to assemble entire products rather than components in the belief that it will obtain greater productivity because of job enrichment.

Administrative Decisions

The types of administrative decision are as follows:

1. **Programmed decisions:** Decisions typically made regarding highly routine situations where little discretion is required.
2. **Non-programmed decisions:** Decisions made in unstructured situations where problem conceptualization and original thinking is required.



Example: Management makes a programmed administrative decision to establish a vehicle maintenance schedule.

Notes

Management makes a non-programmed administrative decision to implement an organizational downsizing plan to reduce duplication of services, decrease costs and increase profitability.

Operational Decisions

The different types of operational decisions include:

1. **Quality:** Decision making regarding product and service quality is a vital operations responsibility necessitating comprehensive organizational support. Quality decisions are made in the design stage of the product or service plan and require the creation and maintenance of standards.
2. **Process:** Operational decisions are made regarding the design of the process used in the manufacturing or servicing of a final product. Process decisions normally are long range and cannot easily be reversed.
3. **Capacity:** Operational capacity decisions are concerned with the long-term capability of an organization to produce the required amount of output over time.



Did u know? Capacity planning determines not only the size of an organization's physical productive capability, but also its human resource needs.

4. **Inventory:** Inventory decisions are crucial in fulfilling management's inventory management strategy. The challenge for operations management is to create a balance in inventory between product demand, cost, and supply needs.
5. **Human Resources:** Human resources are an extremely important operational management responsibility. Organizations pay a major portion of their revenues to employees. Therefore, selection, hiring, training, termination, and general management of human resources are critical for the future of the organization.



Example: The franchise management of a fast-food retail chain makes a determination concerning quality standards in terms of the content and temperature of the food when it is served to the customer. It implements a program to ensure the individual franchises meet the quality standards.

The management of a car-washing company makes a process decision to utilize a brushless car-washing facility that requires fewer workers, results in less damage to the car finish, and is more productive.

A seasonal manufacturer of lawn equipment makes a capacity operational decision to hire and train a second shift of employees during peak demand periods rather than increase overall plant capacity. This will make more productive use of existing capacity without increasing long-term overhead costs including plant maintenance and capital financing costs.

The franchise manager of a chain of job printers makes an operational decision to allow the individual store managers to buy their own printing supply inventory as long as they use the franchise's equipment.

A franchise manager makes a human resource operational decision to allow individual franchisees to hire, train, and supervise their own employees. Thus, the individual franchisee has the entire human resource operational responsibility.

The Managerial Decision Making Process can be understood from figure 1.2.

Notes

Figure 1.2: Managerial Decision Making Process



Task Compare and contrast administrative and operational decisions.

Self Assessment

Fill in the blanks:

3. decisions are intended to order the functions of the organization in the most cost-effective way.
4. decisions are the decisions typically made regarding highly routine situations where little discretion is required.
5. decisions are designed to maximize a firm's profitability through productive procedures.
6. Operational decisions are concerned with the long-term capability of an organization to produce the required amount of output over time.

1.3 Computerized Support

Many people use computerized decision support for work and in recent years to aid in personal decision making. Identifying the targeted or intended users for computerized decision support helps differentiate the specific system. Knowing who does or will use a capability provides useful information about how the content and design of the application might or should differ.

Let's go back in time to the "first" DSS developed by Michael S. Scott Morton (1971). That system was designed to support the market planning manager, the production manager and the marketing manager of a consumer product division of a large multi-business firm. "Every month they developed both a production plan and a sales plan for the following twelve months."

In 1978, Keen and Scott Morton described six diverse DSS: a DSS to help investment managers with a stock portfolio, a DSS used by the president of a small manufacturing company to evaluate

Notes

an acquisition prospect, an interactive DSS used by product planners for capacity planning, a model-driven DSS used by a brand marketing manager for making marketing allocations, a DSS (IBM's Geodata Analysis and Display System called GADS) used by police officers and commanders in teams of four to redesign police beats, and also GADS used by school officials to explore and define alternative school district boundaries.

So who uses computerized decision support including analytics and business intelligence systems? Managers, knowledge workers and staff specialists in a wide variety of professions, occupations, industries and disciplines. Decision support users include internal and external stakeholders of an organization. Ultimately, anyone who makes decisions and has access to a computer is a potential user of a computer-based decision aiding applications.

Self Assessment

State True or False:


- 7. Identifying the targeted or intended users for computerized decision support does not help differentiate the specific system.
- 8. Anyone who makes decisions and has access to a computer is a potential user of a computer-based decision aiding applications.

1.4 Concept of Decision Support Systems (DSS)

In the world of business there are so many approaches to decision-making and because of the wide areas where these systems are used, the concept of decision support systems is very broad. A DSS can be any system that supports decision-making. We might say that a DSS is a system for helping people make decisions using a computer.


A Decision Support System is a computerized system that assists in corporate decision making, with a decision being a choice between alternatives based on the estimated values of those alternatives. Supporting a decision means to assist people working alone, or in a group, to gather intelligence, generate alternatives and make choices.

A Decision Support System is a collection of integrated software applications and hardware that form the backbone of an organization's decision making process. Companies across all industries rely on decision support tools, techniques, and models to help them assess and resolve everyday business questions. The decision support system is data-driven, as the entire process feeds off of the collection and availability of data to analyze.



Notes Business Intelligence (BI) reporting tools, processes, and methodologies are key components to any decision support system and provide end users with rich reporting, monitoring, and data analysis.

Generally, decision support systems are interactive, flexible, and adaptable information systems, developed to support the solution of non-structured management problems for improved decision making.

 *Example:* Medical decision making often involves making a diagnosis and selecting an appropriate treatment.

The closer we get to completely automating our lives, the more crucially we depend upon these kinds of decision-making to be correct beyond a shadow of a doubt! Imagine a wrong decision

being made about a medical diagnosis, or even in terms which we would encounter each and every day, decisions about money or our own wealth simply cannot be wrong!

Decision support even extends into the justice system and is therefore another area where mistakes simply cannot be made. (Imagine a judgment being made on an individual, assisted by DSS. If the judgment goes against that individual, the evidence for such a conviction has to be beyond reproach! Fortunately, it is still only a Decision Support System, not a Decision-Making System as yet.)

We should be aware though that, in situations where there are pressures to get something delivered on time, where there are higher stakes, or numerous ambiguities, experts use intuitive decision making rather than structured approaches.

History of Decision Support Systems: It is generally considered that the concept of DSS became an area of research in the mid-1970s, before it gained intensity in the 1980s. Around 1990, data warehousing and OLAP (Online Analytical Processing) began widening the realm of DSS.

1.4.1 Types of Decision Support Systems

Different types of Decision Support Systems are as follows:

- Model-driven DSS puts emphasis on manipulation of a statistical, financial, or simulation model. This type of DSS uses data and parameters provided by users to assist decision makers in analyzing a situation; (they are not necessarily data intensive.) Parameters are provided by users for the analysis of a situation.



Example: Dicosess is an example of an open source model-driven DSS generator. It runs on the Sun Microsystems Java-based Jini-2 engine.

- Communication-driven DSS supports more collaboration on a shared task.



Example: Include integrated tools like Microsoft's NetMeeting or Groove.


- Data-driven DSS emphasizes manipulation of a chronological series of corporate internal data or occasionally, external data.
- Document-driven DSS manages and manipulates unstructured information in from a variety of electronic formats.
- Knowledge-driven DSS provides specialized problem solving expertise stored as facts, rules, procedures, or in similar structures.

Key DSS characteristics and capabilities are:

- Supports decision makers in semi-structured and unstructured problems.
- Supports managers at all levels.
- Supports individuals and groups.
- Supports interdependent or sequential decisions.
- Support intelligence, design, choice and implementation.
- Supports a variety of decision processes and styles.
- Should be adaptable and flexible.
- Should be interactive and easy to use.

Notes

- Benefits exceed cost.
- Complete control by decision-makers.
- Easy modification to suit needs and changing environment.
- Supports modeling and analysis.



Task Analyze the difference between Data-driven DSS and Document-Driven DSS.

Self Assessment

Fill in the blanks:

9. A is a collection of integrated software applications and hardware that form the backbone of an organization’s decision making process.
10. DSS puts emphasis on manipulation of a statistical, financial, or simulation model.
11. DSS supports more collaboration on a shared task.
12. DSS emphasizes manipulation of a chronological series of corporate internal data or occasionally, external data.
13. DSS manages and manipulates unstructured information in from a variety of electronic formats.

1.5 Management Support Systems

Management Support Systems (MSS) are computer-based systems that are supposed to provide information to be used by or at least to support managerial decision making. Management Support Systems is seen as a subset of Management Information System (MIS), which extends the information retrieval capabilities of the end-users with ‘query and analysis functions’ for searching a database, generating ‘what if’ scenarios, and other such purposes. Thus they are collection of all advanced computer technologies for supporting solutions of managerial problems.



Caution Sometimes called Decision Support Systems or Business Intelligence but there are some distinctions.

A major problem in management support systems development is requirements specification. There exist a large number of systems development methods for MSS. Watson et al. (1997) also pointed out that there are differences between traditional SDM and MSS development methods and that the former are not very useful in MSS development. In a study focusing on the MSS development methods used by organizations in the US, Watson et al. (1997) found that only two formal methods were used, namely: the Critical Success Factors (CSF) method and the Strategic Business Objectives (SBO) method. They, as well as other less used methods, focus primarily on specifying managers’ information needs and how an MSS can fulfil information needs.

Although, they can be useful, they have one major limitation. Since they primarily focus on information needs they are not complete in generating MSS requirements. More complete needs requirements specification can be generated by focusing on managerial roles and how an MSS can support a manager’s different organizational roles. We present an MSS design

approach based on a current management theory and model. In doing so, we build on three postulates.

In most cases the MSS features easy-to-read bar graphs and quick access to important data. The system enables you to look at a real-time report or to “drill down” to the details of a particular order, receipt, Stock Keeping Unit (SKU), etc. It also complements decision making by checking current and “what if” scenarios to determine the best use of manpower. With the use of MSS one can examine worker productivity based on comparisons to a pre-set standard and the average worker; refresh your memory on a certain function; and more. The MSS gives you real-time feedback on your operations so you can make quick and informed decisions.

In the broadest sense, a Decision Support Tool (DST) is any guidance, procedure, or analysis tool that can be used to help support a decision. It is important to point out that decision support is not the same as making a decision. Another important point pertaining to decision support is that it can come in the form of written guidance or in the form of software. Written guidance is frequently provided by regulatory agencies as a means of obtaining a standardized, reproducible approach to reaching a decision. Most regulatory agencies view written guidance as an essential part of the approach to contaminated land management.

Guidance documents can be categorized as document-driven DST. Software tools are also developed to assist in the decision process for computationally intensive analysis, (e.g., flow and transport, geostatistical modeling, and multi-criteria analysis) and for mapping the spatial relationship between contamination data and site landmarks (buildings, roads, etc.).



Did u know? Software tools are categorized as data-driven or model-driven DST depending on the output of the tools.

A software DST has six characteristics (Geoffrion, 1983):

1. Explicit design to solve ill-structured problems;
2. Easy-to-use and powerful user interface;
3. Ability to combine analytical models with data;
4. Ability to explore the solution space by building alternatives;
5. Capability of supporting a variety of decision-making styles; and
6. Allowing interactive and recursive problem-solving.

These characteristics lend themselves naturally to contamination problems in environmental systems in which decisions are often made with uncertainty surrounding the data and models used to interpret the data.

The major advantages of using a computerized decision support tool is that it provides improved transparency of the decision process and permits the effects of uncertainty on the decision to be quantitatively addressed.

Self Assessment

Fill in the blanks:

14. are computer-based systems that are supposed to provide information to be used by or at least to support managerial decision making.
15. A is any guidance, procedure, or analysis tool that can be used to help support a decision.

Notes



Case Study

Managerial Decision Making

In May 1994, Gao Feng, a devout Christian, was arrested in Beijing for planning a private worship service and candlelight vigil to commemorate the fifth anniversary of the Tiananmen Square massacre. Gao was a 26-year employee of Beijing Jeep, Chrysler's joint venture with the Chinese government. Gao was accused of violating Chinese laws against the practice of religion outside of a state-authorized venue. Article 36 of the Chinese Constitution nominally provides for freedom of religious belief; however, the government restricts religious practice to government-sanctioned organizations. State Council Regulation 145 requires all worship-places to register with government religious affairs bureaus and thereby to come under the supervision of official 'patriotic' religious organizations. There are almost 85,000 approved venues for religious activities in China. Many religious groups have been reluctant to comply, either out of opposition to state control of religion or due to fear of adverse consequences if they reveal, as the regulations require, the names and addresses of church leaders. The Universal Declaration of Human Rights, endorsed by UN resolution in 1948, states: "Article 18. Everyone has the right to freedom of thought, conscience and religion; this right includes freedom to change his religion or belief, and freedom, either alone or in community with others and in public or private, to manifest his religion or belief in teaching, practice, worship and observance.

Article 19. Everyone has the right to freedom of opinion and expression; this right includes freedom to hold opinions without interference and to seek, receive and impart information and ideas through any media and regardless of frontiers.

Article 20. (1) Everyone has the right to freedom of peaceful assembly and association. (2) No one may be compelled to belong to any association. According to press reports, Gao was detained by the government for 5 weeks, without formal charge. In early July, he returned to work at Beijing Jeep and told his supervisor that the Chinese Public Safety Bureau had imprisoned him for over a month. Chrysler asked Gao to produce proof of his detention. The Chinese police gave Gao a note that said he had been detained for 3 days and then released without trial. Beijing Jeep's general manager was faced with a tough decision. The Chinese joint venture partner was pressuring Chrysler to fire Gao. If the manager did not fire him, millions of dollars of Chrysler's invested capital in China would be put at risk. If, however, Chrysler fired Gao Feng, the company would become complicit in the violation of his rights to religious freedom and political expression. One of the keys to success in the Chinese market is good relations with the Communist Party, which keeps rigid control over the economy. Multinational corporations spend years cultivating good guanxi or connections in China. They are thus extremely vulnerable to retaliation. At the time of the Gao Feng incident, for example, Chrysler was aware that failure to accede to the government's request could result in losing a valuable minivan contract to its German competitor Daimler – Benz.

Question

What would you do if you were in the general manager's position?

Source: <http://besatuva.com/2011/09/05/chrysler-case-study-managerial-decision-making/>

1.6 Summary

- Business environments are constantly changing. New developments in technology, politics, customer preferences, and regulations happen all the time.

- Decision-making is the focal point of all organizational dynamics, and management effectiveness is judged on the basis of the quality of these decisions.
- Many people use computerized decision support for work and in recent years to aid in personal decision making.
- A Decision Support System (DSS) is a computerized system that assists in corporate decision making, with a decision being a choice between alternatives based on the estimated values of those alternatives.
- Decision support even extends into the justice system and is therefore another area where mistakes simply cannot be made.
- Management Support Systems (MSS) are computer-based systems that are supposed to provide information to be used by or at least to support managerial decision making.
- A major problem in management support systems development is requirements specification.
- A Decision Support Tool (DST) is any guidance, procedure, or analysis tool that can be used to help support a decision.

1.7 Keywords

Communication-driven DSS: Communication-driven DSS supports more collaboration on a shared task.

Data-driven DSS: Data-driven DSS emphasizes manipulation of a chronological series of corporate internal data or occasionally, external data.

Document-driven DSS: Document-driven DSS manages and manipulates unstructured information in from a variety of electronic formats.

DSS: A Decision Support System is a collection of integrated software applications and hardware that form the backbone of an organization's decision making process.

DST: A decision support tool is any guidance, procedure, or analysis tool that can be used to help support a decision.

Knowledge-driven DSS: Knowledge-driven DSS provides specialized problem solving expertise stored as facts, rules, procedures, or in similar structures.

Management Support Systems: Management Support Systems are computer-based systems that are supposed to provide information to be used by or at least to support managerial decision making.

Model-driven DSS: Model-driven DSS puts emphasis on manipulation of a statistical, financial, or simulation model.

1.8 Review Questions

1. Discuss the changes developed in external business environment and external business environment.
2. Illustrate the Key Drivers of Change in the Business Environment.
3. Explain the concept of Managerial Decision Making.
4. Describe the different types of strategic and administrative decisions taken by the management.

- Notes**
5. Make distinction between programmed decisions and non-programmed decisions.
 6. Describe the use of computerized decision support.
 7. What are decision support systems? Explain with example.
 8. Explain the different types of Decision Support Systems.
 9. Discuss the concept of management support system.
 10. Illustrate how does a decision support tool helps in supporting a decision?

Answers: Self Assessment

1. True
2. False
3. Administrative
4. Programmed
5. Operational
6. Capacity
7. False
8. True
9. Decision Support System (DSS)
10. Model-driven
11. Communication-driven
12. Data-driven
13. Document-driven
14. Management Support Systems (MSS)
15. decision support tool (DST)

1.9 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

<http://thelawdictionary.org/management-support-system-mss/>

<http://www.anderson.ucla.edu/faculty/art.geoffrion/home/docs/e35.pdf>

<http://www.businessdictionary.com/definition/management-support-system-MSS.html>

<http://www.slideshare.net/Annie05/management-support-systems-presentation>

Unit 2: Decision Support

Notes

CONTENTS

Objectives

Introduction

- 2.1 Decision Making
- 2.2 Definitions of Decision Making
 - 2.2.1 Decision Making Process
 - 2.2.2 Decision Making in a Total Quality Environment
- 2.3 Models
 - 2.3.1 Rational-Economic (or Classical) Model
 - 2.3.2 Administrative (or Behavioral) Model
 - 2.3.3 Implicit Favorite Model
 - 2.3.4 Political Model
- 2.4 Phases of Decision Making
 - 2.4.1 Intelligence Phase
 - 2.4.2 Design Phase
 - 2.4.3 Choice Phase
 - 2.4.4 Implementation Phase
- 2.5 Summary
- 2.6 Keywords
- 2.7 Review Questions
- 2.8 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss the Concept of Decision Making
- Explain the Models in Decision Making
- Discuss the Phases of Decision Making

Introduction

To be an effective top-level manager today, you need automated support. With change now the rule rather than the exception, more decisions must be made more quickly than ever before. The entire decision-making process, from problem recognition to policy implementation, has become so accelerated it's simply impossible to rely on human response alone. You need an automated system that will help you react and adapt to the constantly changing business environment you must relate to. Your internal data processing system, if you have one, can only offer minimum support because it is structured to perform the regular, scheduled operational

Notes

tasks that keep your business running. Decision making can't wait for a schedule. In this unit, we will discuss the concept of decision making.

2.1 Decision Making

Decision making is the act of choosing between two or more courses of action. However, it must always be remembered that there may not always be a correct decision among the available choices. There may have been a better choice that had not been considered, or the right information may not have been available at the time. Because of this, it is important to keep a record of all decisions and the reasons why decisions were made, so that improvements can be made in the future. This also provides justification for any decision taken when something goes wrong. Hindsight might not be able to correct past mistakes, but it will aid improved decision making in the future.


Decision making is a complex situation. To resolve the complexity, the decisions are classified as programmed and non-programmed decisions. If a decision can be based on a rule, method or even guidelines, it is called the programmed decision. If the stock level of an item is 200 numbers, then the decision to raise a purchase requisition for 400 numbers is a programmed-decision-making situation. The decision maker here is told to make a decision based on the instructions or on the rule of ordering a quantity of 400 items when its stock level reaches 200.

If such rules can be developed wherever possible, then the MIS itself can be designed to make a decision and even execute.

The system in such cases plays the role of a decision maker based on a given rule or a method. Since the programmed decision is made through MIS, the effectiveness of the rule can be analyzed and the rule can be revived and modified from time to time for an improvement. The programmed decision making can be delegated to a lower level in the management cadre.

A decision which cannot be made by using a rule or a model is the non-programmed decision.

Such decisions are infrequent but the stakes are usually larger. Therefore, they cannot be delegated to the lower level.



Notes The MIS in the non-programmed-decision situation can help to some extent, in identifying the problem, giving the relevant information to handle the specific decision making situation. The MIS, in other words, can develop decision support systems in the non-programmed decision making situations.

Self Assessment

Fill in the blanks:

1. is the act of choosing between two or more courses of action.
2. If a decision can be based on a rule, method or even guidelines, it is called the decision.
3. A decision which cannot be made by using a rule or a model is the decision.

2.2 Definitions of Decision Making

According to James Stoner:

“Decision making is the process of identifying and selecting a course of action to solve a specific problem.”

According to Trewartha and Newport:

“Decision making involves the selection of a course of action from among two or more possible alternatives in order to arrive at a solution for a given problem.” opportunity

Decision making is the process of making a choice between a number of options and committing to a future course of actions.

Decision making is the thought process of selecting a logical choice from the available options.



Caution When trying to make a good decision, a person must weight the positives and negatives of each option, and consider all the alternatives.

For effective decision making, a person must be able to forecast the outcome of each option as well, and based on all these items, determine which option is the best for that particular situation.

2.2.1 Decision Making Process

Making good decisions is essential to the management process. As discussed, decisions are rational choices among a group of alternatives. Good decisions are the result of a sequential series of analytical steps:

1. **Identify and delineate the problem:** No management action can occur unless there is a need to resolve an issue. Additionally, when identifying a problem, it is necessary to assess the seriousness of the issue. Highly critical issues require more immediate attention and a greater demand on existing resources. Difficulties in identifying problems include:
 - ❖ *Perceptual errors:* Often problems are not identified because of personal biases, which do not allow the individual to perceive that there is a problem needing attention. Preconceived notions of how something should be (as well as personal preferences) will interfere with the ability to identify a problem.
 - ❖ *Insufficient information:* Insufficient research about a specific problem can lead to misleading and unwarranted conclusions regarding the true nature of the problem and its possible solution. Mistaking a symptom as the cause of the problem. An apparent cause of a problem may just mask a systemic cause. Again, further research is essential to find the cause and nature of a problem
2. **Establish decision priorities and goals:** Managers constantly deal with problems. However, all organizations have limited resources. They must assign priorities to problems in terms of their importance relative to the organization's goals. This process results in a matching of organizational resources with priorities and creates a management methodology for administering solutions to problems.
3. **Ascertain the cause of the problem:** In order to develop a solution to a problem, it is essential to understand its cause. This requires a systemic understanding of the dynamics of the situation that has caused the problem.
4. **Develop realistic alternatives:** It is important for the manager to develop a range of alternative realistic solutions. This means doing extensive research into the nature of the problem and discovering what alternatives would be a good fit.
5. **Weigh the best alternative:** This requires extensive evaluation and comparison using a cost benefit analysis. The alternative solutions are developed within the constraints of limited time and resources, and with a degree of uncertainty.

Notes

6. **Choose a solution:** After conducting extensive research, a decision will have to be made regarding an optimal solution. Managers operate within an environment of incomplete information, time deadlines, and limited resources. All solutions represent opportunity choices having limited outcome predictability.



Caution Managers must make decisions within a range of known alternatives having unknown outcomes.

7. **Implement the decision:** This requires developing human resources to carry out the decision. This mandates a high communication level between the manager and the human resource team.
8. **Follow up:** All decisions require constant monitoring. Changes will have to be made over time to ensure optimum results. This requires an effective organizational control and evaluation system for future organizational decisions



Example: A word processing software manufacturer that has been very successful in the text-based operating system market is facing a crucial decision when the industry standard operating system is changed to a Graphical User Interface (GUI) system. The company's text-based word processor is extremely successful and has a large following. The commands used in the text-based word processor are difficult to learn, but once learned, it is a very versatile word processor.

If the word processor is converted to a graphical user interface, then a portion of the installed user base may be lost, and its competitive advantage based on powerful non-intuitive commands may also be compromised. However, failure to convert the word processor to a graphical user interface will mean losing its market share since the major competitors have already released GUI word processors.

After deciding to develop a GUI word processor, the company had to decide whether to do a fundamental rewrite of the program, which could take at least two years, or simply update it and make it GUI compatible. The company decides to release a GUI update to its word processing program with a fundamental GUI rewrite scheduled for a future date.

Decision-making involves managing three major elements:

1. **Decision strategy:** A decision maker develops a plan affecting long-term organizational outcomes utilizing existing organizational resources.
2. **States of nature:** These are elements of the environment over which the manager has little or no control. States of nature include the weather, political environment, the economy, technological developments, etc. They can dramatically affect the outcomes of any decision strategy.
3. **Outcome:** This is the result of the interaction of the implementation of a decision strategy with the states of nature.



Did u know? Because of the many variables within the states of nature, outcomes can be extremely difficult to forecast.

Thus, outcomes of a decision strategy, O , the dependent variable, is a function of the interaction of the two independent variables, D , decision strategies and, S , the states of nature. Figure 2.1 shows a decision matrix. The rows are strategic choices a manager can make while the columns represent decision outcomes. An outcome O_{ij} is a function of a decision strategy D_i ; and a state of nature S_j .

Figure 2.1: Decision Matrix

| Strategies | S ₁ | S ₂ | S _X | S _X | S _X | S _j |
|----------------|-----------------|-----------------|----------------|----------------|----------------|-----------------|
| D ₁ | O ₁₁ | O ₁₂ | * | * | * | O _{1j} |
| D ₂ | O ₂₁ | O ₂₂ | * | * | * | O _{2j} |
| D _x | * | * | | | | * |
| D _x | * | * | | | | * |
| D _i | O _{i1} | O _{i2} | * | * | * | O _{ij} |

Source: http://worldacademyonline.com/article/18/1/managerial_decision_making.html

Mathematically this relationship can be expressed as:

$$O_{ij} = f(D_i, S_j)$$

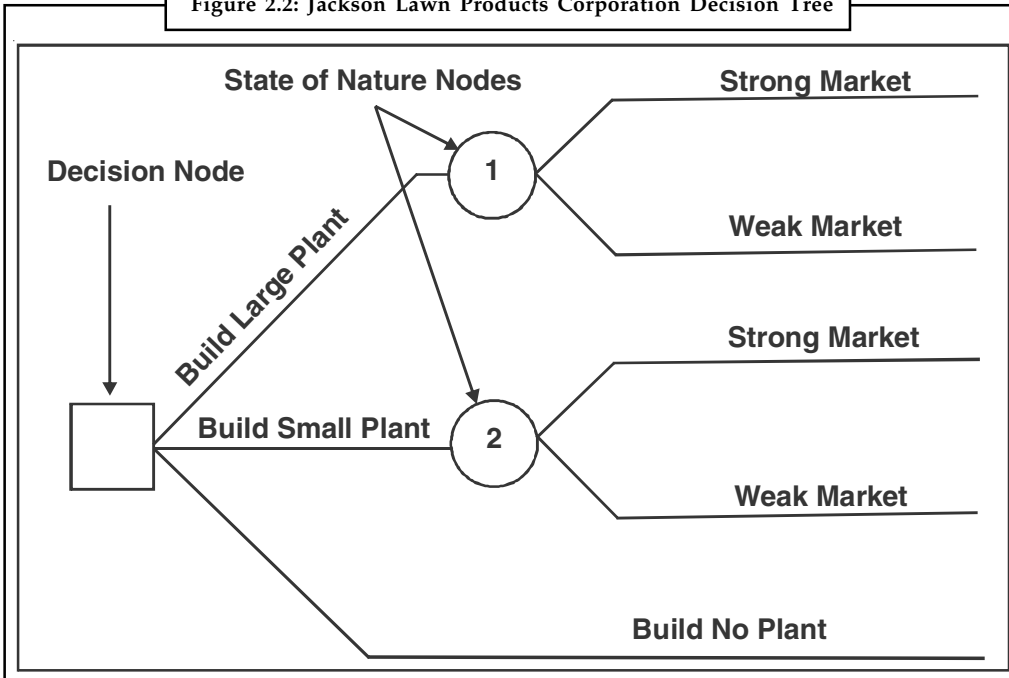
Developing a graphical display is an effective way of mapping the alternatives and probable events that can occur in a complex decision-making environment. Decision trees use symbols consisting of squares and circles. Branches of the decision tree that extend from a square depict an area where several choices can be made while a circle connotes a unique state of nature having certain outcomes.

A decision tree is analyzed in reverse order from right to left going back chronologically. Decision trees are normally accompanied by a payoff or decision table where all the alternatives are listed down the left side of the table with states of nature listed across the top of the table and payoffs stated in the main part of the table.



Example: The Jackson Lawn Products Corporation is studying the possibility of manufacturing a new line of lawn mowers. Since the market for the new mowers is uncertain, the corporation must decide whether to construct a large or small plant, or do nothing. Figure 2.2 presents a decision tree depicting the Jackson Lawn Products decision choices.

Figure 2.2: Jackson Lawn Products Corporation Decision Tree



Source: http://worldacademyonline.com/article/18/1/managerial_decision_making.html

Payoff or decision table can be developed to assist Jackson Lawn Products in determining what type of a lawn mower plant they should build. There is an outcome for each decision and state

Notes

of nature that can be described in units of monetary value. The units of monetary value are also described as conditional values.

Table 2.1 describes all of Jackson’s decision alternatives in the left column of the exhibit, the states of nature across the top, and the payoffs (conditional values) in the main part of the exhibit. In the case of building a large plant, a strong market will produce a \$250,000 net profit, whereas a weak market will result in a \$125,000 net loss. If a small plant is built, a strong market will produce a \$125,000 net profit with a \$60,000 net loss in a weak market.

| Decision | States of Nature | |
|-------------------|------------------|-------------|
| | Strong Market | Weak Market |
| Build Large Plant | \$250,000 | -\$125,000 |
| Build Small Plant | \$125,000 | -\$60,000 |
| Build No Plant | \$0 | \$0 |

2.2.2 Decision Making in a Total Quality Environment

Decision making in a total quality environment essentially involves the elements of a continuous process focusing on three essential components: continuous improvement, assessment management and teamwork.

Implementing decision-making in a total quality environment requires:

- Making quality improvement a central organizational focus.
- Extensive and continuous employee training.
- Total involvement of the employees and management concerning the organization’s mission, goals and operational objectives.
- Continual improvement of organizational processes rather than focusing on employees as the source of quality failures.
- Team decision-making.
- The recognition that the customer defines quality, and that a total quality decision objective is to meet or exceed customer satisfaction standards.

Decision-making in a TQM environment is a shared experience for all employees throughout the organization. Information is an organizational resource essential for making quality decisions. Increased quality leads to increased productivity, lower unit costs, and higher customer satisfaction.



Example: A transmission manufacturing company was machining highly exacting parts to be used in automatic transmissions. The parts consisted of gears, bearings, and assorted spacers and shafts. The parts were engineered to be within the industry standard of + or - .003" of specifications. However, other manufacturers were able to produce the same transmissions at two-thirds of the price and achieve higher productivity and quality.

Management decided to make a 300% improvement in the tolerances of the parts by reducing them to + or -.001" tolerance. The net result was there were fewer returns, lower unit costs, and higher overall customer satisfaction. Consequently, productivity and profits grew substantially.

Self Assessment

Notes

Fill in the blanks:

4. is the result of the interaction of the implementation of a decision strategy with the states of nature.
5. The rows are choices a manager can make while the columns represent decision outcomes.
6. A is analyzed in reverse order from right to left going back chronologically.

2.3 Models

The different types of models are discussed below:

2.3.1 Rational-Economic (or Classical) Model

The Rational-Economic (or Classical) Model consists of the following features:

- It is prescriptive in that it focuses on how decisions ought to be made.
- Assumes the decision maker is completely rational (i.e., seeks to maximize the payoff and utilizes a search process that proceeds in a planned, orderly and consistent fashion) and unbiased.
- Assumes that the decision maker has available all the information needed to make a decision and that all possible alternatives are considered.
- The decision maker selects the optimum or best choice.
- Decision making proceeds through the following sequence of steps: problem identification, development of criteria against which alternative solutions can be evaluated, identification of alternative courses of action, evaluation of alternatives, selection of the best alternative, and implementation.

2.3.2 Administrative (or Behavioral) Model

The Administrative (or Behavioral) Model consists of the following features:

- It is descriptive in that it describes how decisions are actually made.
- Decision makers seek to simplify problems and make them less complex because they are constrained by their individual capabilities (e.g., limited information processing ability) and by organizational conditions (e.g., availability of resources).
- Assumes that decision makers operate with limited (or "bounded") rationality; this means that decision makers are rational within a simplified model which contains fewer components (e.g., fewer decision making criteria, fewer options, etc.).
- Assumes that decision makers identify a limited number of decision making criteria, that they examine a limited range of alternatives (only those which are easy to find, highly visible, have been tried before or are only slightly different from the status quo) and that they do not possess all the information needed to make a decision.
- The decision maker selects a satisficing alternative. This is an alternative that is "good enough" or satisfactory in that it meets the minimum criteria established for a desired solution.

Notes

- Decision making proceeds sequentially: alternatives are examined one at a time and the first satisfactory alternative that is found is selected.

2.3.3 Implicit Favorite Model

The Implicit Favorite Model consists of the following features:

- It is descriptive in that it describes how decisions are actually made.
- The decision maker seeks to simplify the decision making process by identifying an “implicit favorite” before alternatives are evaluated; this often occurs subconsciously.
- The decision maker is neither rational nor objective and unbiased.
- After a “favorite” is selected, the decision maker tries to appear rational and objective by developing decision criteria and by identifying and evaluating various alternatives; however, this is done in a biased way so as to ensure that the favorite appears superior of these criteria and thus, can legitimately be selected as the “best” solution.
- In this model, “decision making” is essentially a process of confirming a choice/decision that has already been made.



Notes The actual decision was made in an intuitive and unscientific fashion.

2.3.4 Political Model

The Political Model consists of the following features:

- It is descriptive in that it describes how decisions are actually made.
- The decision maker is neither rational nor objective and unbiased.
- Since the group members have different agendas, they need to negotiate with each other.
- The process involves a cycle of bargaining among the decision makers in order for each one to try to get his or her perspective to be the one of choice – more specifically, to sway powerful people within the situation to adopt his or her viewpoint and influence the remaining decision makers.
- This model does not involve making full information available, since it is based upon negotiation that is often influenced by power and favors. In fact, information is often withheld in order to better maneuver a given perspective.
- In this model, potential problems and conflict often can be foreseen and minimized. Once powerful people have been swayed to support a particular viewpoint, other group members usually fall in line behind them.
- The nature of bargaining and maneuvering (e.g., withholding information and social pressure) can produce effects that are long-lasting and detrimental. Once they discover it, the individuals involved in the decision may not appreciate the duplicity inherent in the process.



Task Compare and contrast Implicit Favorite Model and Political model.

Self Assessment

Notes

Fill in the blanks:

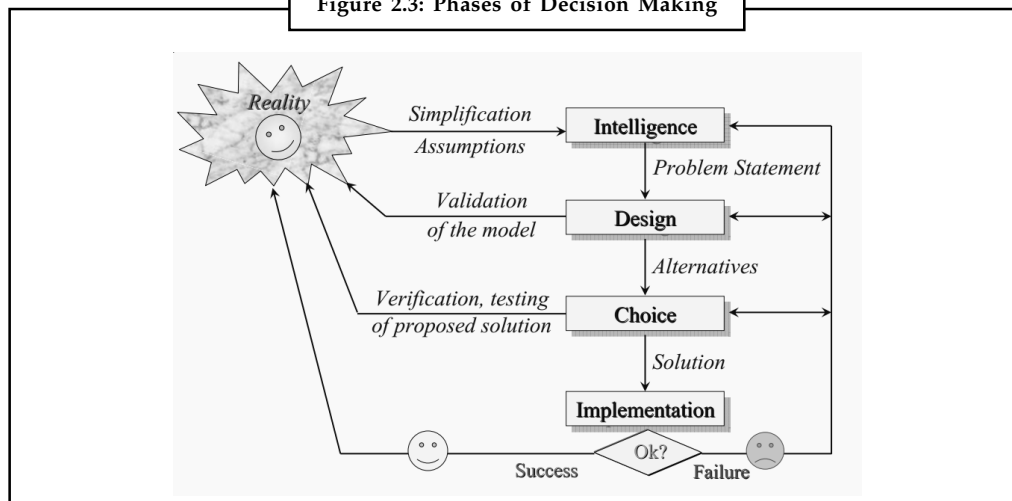
7. The Model is prescriptive in that it focuses on how decisions ought to be made.
8. In case of Model, Decision making proceeds sequentially.
9. In Model, "decision making" is essentially a process of confirming a choice /decision that has already been made.
10. In model, potential problems and conflict often can be foreseen and minimized.

2.4 Phases of Decision Making

There are four major phases:

1. **Intelligence:** Decision maker examines reality and identifies the problem; problem ownership is established as well.
2. **Design:** The model (that represents the system) is constructed making assumptions that simplify reality and writing down the relationships (among variables). The model is validated and criteria are determined (choice of evaluation); identifies alternative solutions.
3. **Choice:** It includes selection of a proposed solution (is tested for viability) to the model. If solution is reasonable we may pass to the next phase else we must return to an earlier phase.
4. **Implementation:** It provides the successful result in solving the real problem.

Figure 2.3: Phases of Decision Making



Source: http://www.cs.ubbcluj.ro/~per/Dss/Dss_2.pdf

These phases are discussed below:

2.4.1 Intelligence Phase

The intelligence phase involves scanning the environment – identifying problem situations (opportunities), monitoring the results of implementation.

Notes

A decision making attempts to determine whether a problem exists, identify its symptom, determine its magnitude and explicitly define it. What is described as a problem may be only a symptom of a problem (real world problems are complicated by many interrelated factors and is difficult to distinguish between the symptoms and the real problem). The measurement of productivity and the construction of a model are based on real data. The collection of data and the estimation of future data is a difficult step in the analysis.

Data may have some inconvenience like:

- data are not available,
- obtaining data may be expensive,
- data estimation may be subjective,
- data may be insecure,
- important data that influence the results may be qualitative,
- there may be too many data – information overload,
- results may occur over an extended period,
- the future data will be similar to historical data.

After the preliminary investigation is possible to determine whether a problem really exists, where is located, how important it is, whether an information system is reporting a problem or only the symptoms of a problem.

Problem Decomposition

The complex problems can be divided into sub-problems. Solving the simpler sub-problems may help in solving the initial one. Poorly structured problems may have highly structured sub-problems. Decomposition facilitates communication among decision makers. Decomposition is one of the most important aspects of the Analytical Hierarchy Process – AHP – which helps decision makers incorporate both qualitative and quantitative factors into the decision-making models.

Problem Ownership

In the intelligence phase it is important to establish problem ownership. A problem exists in an organization only if someone or some group takes on the responsibility of attacking it and if the organization has the ability to solve it.

2.4.2 Design Phase

The design phase involves finding or developing and analyzing possible courses of action: understanding the problem and testing solution for feasibility. A model of the decision-making problem is constructed, tested and validated.

Modeling involves conceptualizing a problem and abstracting it to quantitative and qualitative form – for a mathematical model the variables are identified, their mutual relationship are established. Simplifications are made whenever necessary through assumptions.

A simpler model leads to lower development costs, easier manipulation and a faster solution but is less representative of the real problem and can produce inaccurate results.

The process modeling is art (creativity are required when art determining what simplifying assumptions can work and how to integrate models to obtain valid solutions) and science (there are science many standard model classes available and an analyst can say which one is applicable).

Selection of a Principle of Choice

Notes

A principle of choice is a criteria principle of choice that describes the acceptability of a solution approach. It involves how a person establishes decision making objectives and incorporates them into the model. We must recognize the difference between a criterion and a constraint.

2.4.3 Choice Phase

In the choice phase the actual decision and the commitment to follow a certain course of action is made. The choice phase includes the search for, evaluation of, and recommendation of an appropriate solution to a model. A solution to a model is a specific set of values for the decision variables in a selected alternative. A solution to the model yields a recommended solution to the problem. The problem is solved if the recommended solution is successfully implemented.

Solving a decision-making model involves searching for an appropriate course of action. Search approaches include analytical techniques, algorithms, heuristics, and blind searches. Sensitivity analysis is used to determine the robustness of any given alternative: slight changes in the parameters lead to slight changes in the alternative chosen.



Did u know? Analysis is used to explore major changes in the parameters. Goal seeking helps a manager determine values of the decision variables to meet a specific objective.

2.4.4 Implementation Phase

The implementation of a proposed solution to a problem is the initiation of a new order of things or the introduction of change. The implementation phase involves putting a recommended solution to work, not necessarily implementing a computer system. The primary requirement of decision support for the intelligence phase is the ability to scan internal and external information sources for opportunities and problems and to interpret what the scanning discovers.



Task Analyse the difference between design phase and choice phase.

Self Assessment

Fill in the blanks:

11. The phase involves scanning the environment – identifying problem situations (opportunities), monitoring the results of implementation.
12. The phase involves finding or developing and analyzing possible courses of action: understanding the problem and testing solution for feasibility.
13. In the phase the actual decision and the commitment to follow a certain course of action is made.
14. The of a proposed solution to a problem is the initiation of a new order of things or the introduction of change.
15. is the intelligence applied to the implementation, that is, a form of feed back.

Notes



Case Study

Management Analysis and Decision Making

Times are slow for your company right now and with the rising costs of materials and wages, your profits are at an all-time low. Because of this unfortunate situation, you will need to let some employees go. The senior management team has already compiled the list of people whose employment will be terminated two weeks from today. However, the people on the list will not know until the day of the termination. You have called a meeting of your department managers and supervisors (judges). The managers and supervisors do not know that a list has been created, so you will need to let them know this at some point in the conversation. Also, they will not be able to see the list until the day of the terminations. Obviously, this is a very confidential topic and should not be shared with anybody outside of this meeting. The purpose of your meeting today is to confide in this group and assure them that none of them are on the list. You also want to get their feedback on how the general employee base will react to the news and event in two weeks. Next, you'd like to understand and anticipate any questions that they believe will arise so that appropriate answers can be prepared. Finally, you would like to devise an action plan/transition plan for the day after the event. What you can tell the managers is the number of people they will each be losing, if you find that information important to share. Here is the breakdown:

- Order Processing will lose four of its 12 people
- Human Resources will lose two of its five people
- Production will lose eight of its 40 people. After introductions, you should begin discussing this upcoming event with your managers (judges). Spend as much time on each of the following questions as your group feels is necessary.
 - ❖ Why is this happening?
 - ❖ How will this impact the areas?
 - ❖ How will the department managers plan for this without breaking confidentiality?
 - ❖ How will the departments transition after the event?
- Anything else you feel your managers and supervisors would want to know or need to know.

Each member of your team (individual) should deliver some piece of the conversation to the employees (judges).

Question

Discuss the case.

Source: http://www.fbla-pbl.org/docs/ct/production/PBL/case_study_management_analysis___decision_making.pdf

2.5 Summary

Decision making is the act of choosing between two or more courses of action.

- Decision making is a complex situation. To resolve the complexity, the decisions are classified as programmed and non-programmed decisions.

- Decision making is the process of making a choice between a number of options and committing to a future course of actions.
- A decision maker develops a plan affecting long-term organizational outcomes utilizing existing organizational resources.
- A decision tree is analyzed in reverse order from right to left going back chronologically.
- The intelligence phase involves scanning the environment – identifying problem situations (opportunities), monitoring the results of implementation.
- The design phase involves finding or developing and analyzing possible courses of action: understanding the problem and testing solution for feasibility.
- The choice phase includes the search for, evaluation of, and recommendation of an appropriate solution to a model.
- The implementation phase involves putting a recommended solution to work, not necessarily implementing a computer system.

2.6 Keywords

Administrative Model: Administrative model describes how decisions are actually made.

Choice Phase: The choice phase includes the search for, evaluation of, and recommendation of an appropriate solution to a model.

Decision Making: Decision making is the process of making a choice between a number of options and committing to a future course of actions.

Design Phase: The design phase involves finding or developing and analyzing possible courses of action: understanding the problem and testing solution for feasibility.

Implementation Phase: The implementation phase involves putting a recommended solution to work, not necessarily implementing a computer system.

Intelligence Phase: The intelligence phase involves scanning the environment – identifying problem situations (opportunities), monitoring the results of implementation.

Non-programmed Decision: A decision which cannot be made by using a rule or a model is the non-programmed decision.

Rational-Economic (or Classical) Model: The Rational-Economic (or Classical) Model is prescriptive in that it focuses on how decisions ought to be made.

2.7 Review Questions

1. Explain the concept of decision making with example.
2. Illustrate how to resolve the complexity of decision making.
3. Make distinction between programmed and non-programmed decision.
4. Explain the steps included in the Decision making Process.
5. Describe the decision making process with example.
6. What are the major elements of decision making? Discuss.
7. Discuss the role of Decision making in a total quality environment.

Notes

8. Make distinction between the Rational-Economic (or Classical) Model and Administrative model.
9. Describe the various phases of decision making.
10. Solving a decision-making model involves searching for an appropriate course of action. Comment.

Answers: Self Assessment

- | | |
|----------------------|-----------------------------------|
| 1. Decision making | 2. Programmed |
| 3. Non-programmed | 4. Outcome |
| 5. Strategic | 6. Decision tree |
| 7. Rational-Economic | 8. Administrative (or Behavioral) |
| 9. Implicit Favorite | 10. Political |
| 11. Intelligence | 12. Design |
| 13. Choice | 14. Implementation |
| 15. Monitoring | |

2.8 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

<http://blog.majoriumbusinesspress.com/2011/10/27/the-four-stages-of-effective-decision-making/>

<http://epstein.usc.edu/research/JudicialBehaviorSegal.pdf>

<http://www.robustdecisions.com/decision-making-tools/decision-making-models.php>

http://www.tutorialspoint.com/management_concepts/decision_making_process.htm

Unit 3: Decision Support Systems

Notes

CONTENTS

Objectives

Introduction

- 3.1 Decision Support Systems
- 3.2 Configuration and Characteristics of DSS
 - 3.2.1 Configuration of DSS
 - 3.2.2 Characteristics of DSS
- 3.3 Components of Decision Support Systems
 - 3.3.1 Data Management Component
 - 3.3.2 Model Management Component
 - 3.3.3 User Interface or Dialogue Management Component
 - 3.3.4 Knowledge Management Component
- 3.4 Summary
- 3.5 Keywords
- 3.6 Review Questions
- 3.7 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss the Concept of Decision Support Systems
- Discuss the Configuration of DSS
- Discuss the Characteristics of DSS
- Explain Components of DSS

Introduction

Decision Support Systems have evolved over the past three decades from simple model-oriented systems to advanced multi-function entities. During the 1960s, most Decision Support Systems were fairly based on powerful (and expensive) mainframe computers which provided managers with structured, periodic reports. MIS theory developments during the 1970s saw Decision Support Systems evolve into more elaborate computer-based systems that supported production, promotion, pricing, marketing and some logistical functions. By early 1980s Decision Support Systems enjoyed more interests from academics and the framework for Decision Support Systems was greatly expanded by the end of the decade. It was only during the 1990s that a paradigm shift occurred in Decision Support Systems and more complex systems, which incorporated, advanced database technology and client/server capabilities, were emerging from many areas in business processes. As many organizations started to upgrade their network infrastructure, object-oriented technology and data warehousing started to make its mark on Decision Support Systems.

Notes


The rapid expansion of the Internet provided additional opportunities for the scope of Decision Support Systems and consequently many new innovative systems such as OLAP and other web-drive systems were developed. In this unit, we will discuss the concept of decision support systems.

3.1 Decision Support Systems

Decision support systems are software tools that are developed for the specific purpose of assimilating pertinent data and generating a detail analysis of a given situation. The results of the decision support system, or DSS, can then be utilized in making an informed decision regarding just about any action or set of circumstances.

A decision support system (DSS) is a computer program application that analyzes business data and presents it so that users can make business decisions more easily. It is an “informational application” (to distinguish it from an “operational application” that collects the data in the course of normal business operation).

Typical decision support systems are designed for easy use by end users who may or may not be comfortable with using computer technology as part of the decision making process. Often, the software will be make use of formats such as spreadsheets or databases that work with the use of fields to enter data.



Notes The process of entering data into the program is usually very straightforward, and often includes tips and other forms of assistance as the data is entered into the fields.

Once the information is entered, it is possible to query the system with a series of questions that can allow the user to project not only a range of possible courses of action, but also get some idea on the possible outcomes associated with each option.

In terms of dealing with many managerial and executive decisions that are made daily in the workplace, decision support systems can be a great tool for supervisors, department heads, and site managers. The system can be configured to work within just about any area that involves dealing with labor issues. Human resource personnel can make use of decision support systems that include data about all applicable local and federal laws governing the rights and protections relevant to employees, making the process of administering promotions, dealing with a leave of absence, or developing an equitable approach to disciplinary action much easier. Supervisors can also use decision support systems as part of the identifying strengths and weaknesses among employees in their charge, which can take a lot of the guesswork out of performing period employee evaluations.

Over the last few years, decision support systems have become better known as Business Intelligence systems. However, many persons who have made use of the software for a number of years still tend to make use of the identification of decision support software, which has led many producers of the software to continue use of the DSS identification.

Typical information that a decision support application might gather and present would be:

- Comparative sales figures between one week and the next
- Projected revenue figures based on new product sales assumptions
- The consequences of different decision alternatives, given past experience in a context that is described



Did u know? A decision support system may present information graphically and may include an expert system or artificial intelligence (AI).

A decision support system may be aimed at business executives or some other group of knowledge workers.

Self Assessment

Fill in the blanks:

1. A is a computer program application that analyzes business data and presents it so that users can make business decisions more easily.
2. Over the last few years, decision support systems have become better known as
3. In terms of dealing with many managerial and executive decisions, decision support systems can be a great for supervisors, department heads, and site managers.

3.2 Configuration and Characteristics of DSS

In this section, we will discuss the Configuration and Characteristics of DSS.

3.2.1 Configuration of DSS

The major characteristics of a decision support system utilised in network configuration would be the flexibility to include in the model every significant element of a specific supply chain and every possible option concerning supply chain configuration, regarding facility location, opening or closing of facilities, etc.



Caution It must have the means of integrating managerial decisions concerning service level.

Decision support system Configuration:


- Supports individuals and teams
- Used repeatedly and constantly
- Two major components: data and models
- Web-based
- Uses subjective, personal, and objective data
- Has a simulation model
- Used in public and private sectors
- Has what-if capabilities
- Uses quantitative and qualitative models

Notes

3.2.2 Characteristics of DSS

A decision support system is an internal business element that provides owners, executives, and other employees with information for making informed decisions. Common characteristics of decision support systems include facilitation, interaction, and task-oriented activities within the system. Other characteristics may also exist, such as ancillary options or measures of decision impact. Most effective decision support systems will have these features readily available. Those who create and implement the system are often responsible for measuring its effectiveness and proper application.


- **Facilitation:** Facilitation helps ensure the system provides the correct data and information for upcoming decisions. The decision-making activities and processes must coincide when using a support system. A decision process typically defines the individual activities one must use when making decisions. The support system facilitates information through each individual activity so that the decision maker does not spend too much time gathering data. Multiple individuals can also use the support system to pass information through the pipeline in order to shorten the lead time spent on making decisions.
- **Interaction:** Interaction also represents important characteristics of decision support systems. Many of these systems in current businesses make extensive use of technology, primarily personal computers and decision-making software. When multiple individuals are able to access the software and input information for the decision-making process, interaction occurs.



Notes In some cases, decision makers may have regional restrictions that do not allow visits to all of a company's locations. Therefore, the only interaction during the process is the decision support system and the data contained therein.

- **Task-oriented activities:** Task-oriented activities are also common elements of decision support systems. This element means that the support system can handle more than one activity at a time, which is essential. Many companies provide inputs for multiple outcomes related to one or a few different decisions. Therefore, the multiple tasks in the system must be able to handle all this information. The choice among the various alternatives here can represent the decision one must make for a company.

Companies can define their own characteristics of decision support systems. While some companies may design their own mix of computerized and manual systems, others may use a third-party package. A third-party package often has customizable aspects that allow a company to change various aspects in the system. In short, the characteristics can match the needs and wants of a company.



Task Make a report on the characteristics of decision support systems.

Self Assessment

State True or False:

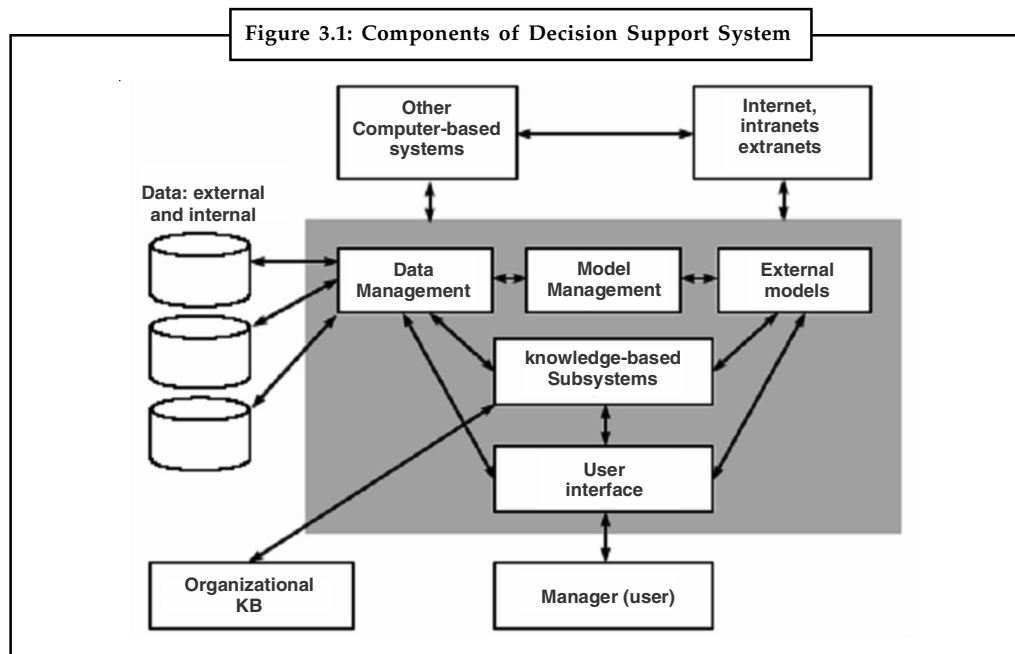
4. Interaction helps ensure the system provides the correct data and information for upcoming decisions.
5. When multiple individuals are able to access the software and input information for the decision-making process, interaction occurs.

6. Task-oriented activities is an element which means that the support system can handle more than one activity at a time, which is essential.
7. A third-party package often has customizable aspects that does not allow a company to change various aspects in the system.

Notes

3.3 Components of Decision Support Systems

A typical Decision support systems has four components: data management, model management, knowledge management and user interface management.



Source: http://www.msit2005.mut.ac.th/msit_media/1_2553/ITEC3614/Lecture/2010074_095946DI.pdf

3.3.1 Data Management Component

The data management component performs the function of storing and maintaining the information that you want your Decision Support System to use. The data management component, therefore, consists of both the Decision Support System information and the Decision Support System database management system. The information you use in your Decision Support System comes from one or more of three sources:

- **Organizational information:** You may want to use virtually any information available in the organization for your Decision Support System. What you use, of course, depends on what you need and whether it is available. You can design your Decision Support System to access this information directly from your company's database and data warehouse. However, specific information is often copied to the Decision Support System database to save time in searching through the organization's database and data warehouses.
- **External information:** Some decisions require input from external sources of information. Various branches of federal government, Dow Jones, and the internet, to mention just a few, can provide additional information for the use with a Decision Support System.
- **Personal information:** You can incorporate your own insights and experience your personal information into your Decision Support System. You can design your Decision

Notes

Support System so that you enter this personal information only as needed, or you can keep the information in a personal database that is accessible by the Decision Support System.

3.3.2 Model Management Component

The model management component consists of both the Decision Support System models and the Decision Support System model management system. A model is a representation of some event, fact, or situation. As it is not always practical, or wise, to experiment with reality, people build models and use them for experimentation. Models can take various forms.

Businesses use models to represent variables and their relationships.



Example: You would use a statistical model called analysis of variance to determine whether newspaper, TV, and billboard advertising are equally effective in increasing sales.

Decision Support Systems help in various decision-making situations by utilizing models that allow you to analyze information in many different ways. The models you use in a Decision Support System depend on the decision you are making and, consequently, the kind of analysis you require.



Example: You would use what-if analysis to see what effect the change of one or more variables will have on other variables, or optimization to find the most profitable solution given operating restrictions and limited resources.



Did u know? Spreadsheet software such as excel can be used as a Decision Support System for what-if analysis.

The model management system stores and maintains the Decision Support System's models. Its function of managing models is similar to that of a database management system. The model management component cannot select the best model for you to use for a particular problem that requires your expertise but it can help you create and manipulate models quickly and easily.

3.3.3 User Interface or Dialogue Management Component

The user interface management component allows you to communicate with the Decision Support System. It consists of the user interface management system. This is the component that allows you to combine your know-how with the storage and processing capabilities of the computer.

The user interface is the part of the system you see through it when enter information, commands, and models.



Caution This is the only component of the system with which you have direct contract.

If you have a Decision Support System with a poorly designed user interface, if it is too rigid or too cumbersome to use, you simply won't use it no matter what its capabilities. The best user interface uses your terminology and methods and is flexible, consistent, simple, and adaptable.



Example: The components of a Decision Support System, let's consider the Decision Support System that Land's End has tens of millions of names in its customer database. It sells a

wide range of women's, men's, and children's clothing, as well various household wares. To match the right customer with the catalog, Land's End has identified 20 different specialty target markets. Customers in these target markets receive catalogs of merchandise that they are likely to buy, saving Land's End the expense of sending catalogs of all products to all 20 million customers. To predict customer demand, Land's End needs to continuously monitor buying trends. And to meet that demand, Land's End must accurately forecast sales levels. To accomplish these goals, it uses a Decision Support System which performs three tasks:

- **Data management:** The Decision Support System stores customer and product information. In addition to this organizational information, Land's End also needs external information, such as demographic information and industry and style trend information.
- **Model management:** The Decision Support System has to have models to analyze the information. The models create new information that decision makers need to plan product lines and inventory levels.



Example: Land's End uses a statistical model called regression analysis to determine trends in customer buying patterns and forecasting models to predict sales levels.

- **User interface management:** A user interface enables Land's End decision makers to access information and to specify the models they want to use to create the information they need.

3.3.4 Knowledge Management Component

The knowledge management component, like that in an expert system, provides information about the relationship among data that is too complex for a database to represent. It consists of rules that can constrain possible solution as well as alternative solutions and methods for evaluating them.



Example: When analyzing the impact of a price reduction, a Decision Support System should signal if the forecasted volume of activity exceeds the volume that the projected staff can service. Such signaling requires the Decision Support System to incorporate some rules-of-thumb about an appropriate ratio of staff to sales volume. Such rules-of-thumb, also known as heuristics, make up the knowledge base.

These components are discussed in detail in the next unit.



Task Compare and contrast data management and model management model.

Self Assessment

Fill in the blanks:

8. The component performs the function of storing and maintaining the information that you want your Decision Support System to use.
9. A is a representation of some event, fact, or situation.
10. The component consists of both the Decision Support System models and the Decision Support System model management system.
11. The management component allows you to communicate with the Decision Support System.

Notes

12. The is the part of the system you see through it when enter information, commands, and models.
13. The component provides information about the relationship among data that is too complex for a database to represent.
14. The function of managing models in model management is similar to that of a
15. Businesses use models to represent and their relationships.



Case Study

An Intelligent Decision Support System

Information technology applications that support decision-making processes and problem-solving activities have proliferated and evolved over the past few decades. In the 1970s, these applications were simple and based on spreadsheet software. During the 1980s, decision support systems incorporated optimization models, which originated in the operations research and management science communities. In the 1990s, these systems were further enhanced with components from artificial intelligence and statistics. This evolution led to many different types of decision support systems with somewhat confusing names, including management information systems, intelligent information systems, expert systems, management support systems, and knowledge-based systems. Because businesses realized that data was a precious asset, they often based these “intelligent” systems on data warehousing and online analytical processing technologies. They gathered and stored a lot of data, assuming valuable assets were implicitly coded in it. Raw data, however, is rarely beneficial. Its value depends on a user’s ability to extract knowledge that is useful for decision support. Thousands of “business intelligence” companies thus emerged to provide such services. After analyzing a corporation’s operational data, for example, these companies might return intelligence (in the form of tables, graphs, charts, and so on) stating that, say, 57 percent of the corporation’s customers are between 40 and 50, or product Q sells much better in Florida than in Georgia.

Many businesses have realized, however, that the return on investment for pure “business intelligence” is much smaller than initially thought. The “discovery” that 57 percent of your customers are between 40 and 50 doesn’t directly lead to decisions that increase profit or market share. Moreover, we live in a dynamic environment where everything is in flux.

Interest rates change, new fraud patterns emerge, weather conditions vary, the stock markets rise and fall, new regulations and policies surface, and so on. These economic and environmental changes render some data obsolete and make other data—which might have been useless just six weeks ago—suddenly meaningful. We developed a software system to address these complexities and implemented it on a real distribution problem for a large car manufacturer. The system detects data trends in a dynamic environment, incorporates optimization modules to recommend a near-optimum decision, and includes self-learning modules to improve future recommendations. As figure 1 shows, such a system lets enterprises monitor business trends, evolve and adapt quickly as situations change, and make intelligent decisions based on uncertain and incomplete information.

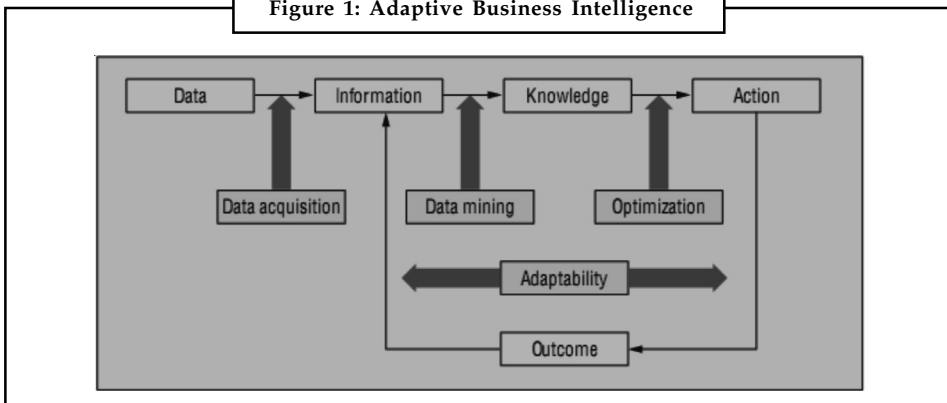
This diagram shows the flow from data acquisition to recommended action, including an adaptive feedback loop.

Contd....

Problem Overview

We developed the system for a US-based car manufacturer that has more than 1 million cars returned from leases or rentals each year. The manufacturer owns the cars, and the problem is how to best distribute these cars among hundreds of auction sites around the United States. The cars vary by make and model, mileage, options, wear and tear, and so on. These characteristics, along with others, influence the car's sale price at each particular auction. Our central challenge was to achieve the "best" possible distribution among these auction sites—that is, the distribution that maximizes the net proceeds from all sales. The process of making optimal recommendations involves many considerations, ranging from price prediction for various car types at different locations, to price depreciation and volume effects, to transportation issues. One million cars per year corresponds to approximately 4,000 cars per working day. So, each day, a remarketing team must make 4,000 decisions regarding which auction site will maximize the sales price of each car. Further, due to volume effects, assigning cars to auctions is highly interrelated, and therefore it's not possible to process these cars sequentially.

Figure 1: Adaptive Business Intelligence



Say, for example, that a company uses 50 auction sites and processes a mere 1,000 cars per day. This results in a mind-boggling 501,000 distribution choices! No computer can check out all these possible combinations in a human lifetime. Nevertheless, the manufacturer requires decisions on all of the cars today.

Problem Complexity

To illustrate the task, we'll use a silver, four door 2002 Toyota Corolla with 34,983 miles, a sun roof, automatic transmission, power windows, power seats, and many other options. At the moment, the car sits at a dealership in Virginia, and we must decide where to send it. At first glance, this looks easy. We might be tempted to simply look up the car's average sales price at each auction using one of many guides, such as the Black Book, Kelley Blue Book, or Manheim Auction Report. After adjusting the price based on the car's mileage, options, and so on, and estimating transportation costs—both manageable calculations—we might simply decide to target the auction with the highest current average sales price.

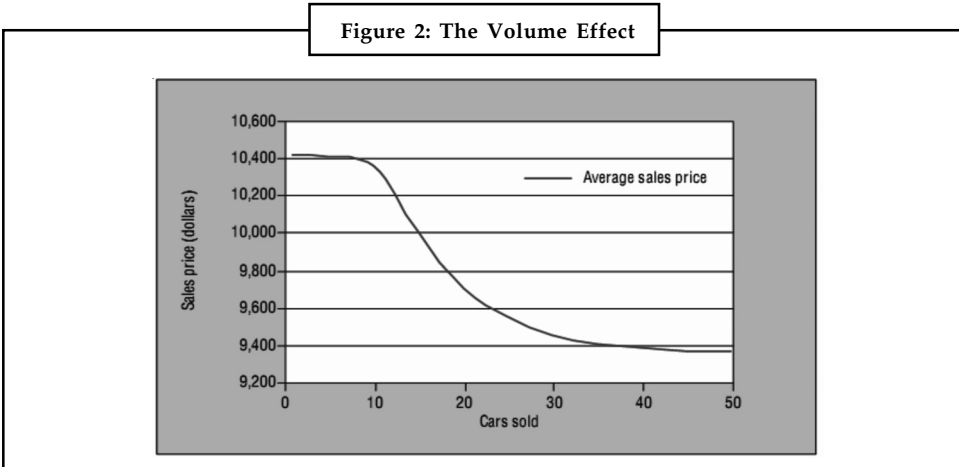
So, what's the problem? In a word: volume. Per car, it's cheaper to ship a truckload of cars from one place to another than it is to ship one or a few cars at a time.

For more than 14 cars, we usually calculate transportation costs by determining the fee for transporting 14 cars (or multiples of 14 cars), then calculate the remainder at the applicable rate. For example, transporting 20 cars would cost us \$85 per car for the first 14 cars, and then \$120 per car for the remaining 6 cars, for a total cost of \$1,910. Beyond transportation

Contd....

Notes

costs, we must also address other considerations, such as the volume effect, auction and transportation scheduling, depreciation, and insurance and risk factors (cars might be damaged or stolen in transit). The volume effect kicks in as the number of similar cars for sale grows. If we send many similar cars to a single auction site—which is reasonable, assuming it offers the best net price—the volume will result in less money per car (see figure 2).



Past a certain point, selling many similar cars at the same auction reduces the price per car.

Say, for example, the current average sale price for a 2002 Toyota Corolla on a particular auction site is \$7,200. We would likely get this price if we ship up to seven cars to that location. However, transporting 30 similar cars to that site would drop the average price to \$6,900. To complicate this further, “similar” doesn’t mean the same make, model, and color. Even though the makes and models might differ, shipping 30 silver sedans to the same auction gives buyers more options, thereby depressing the average sales price per car. Also, the volume effect’s curve is different for different types of cars. With Toyota Corollas, for example, the volume effect is significant, whereas for Porsche 911s, it is moderate.

Scheduling is also a major issue. Every auction has a typical sales day, such as at 11 a.m. every second Friday. So, let’s say we have 20 cars that we’d like to ship to an auction site, the transport time is 10 days, and the next auction is 11 days away. If there is even a slight delay in the delivery of these 20 cars, then we might miss the auction. The cars would then have to sit in the auction’s parking lot for almost two weeks. This is bad—not only because the company wants them sold as soon as possible, but because the cars would lose value each day. The price depending on distribution volume. In contrast, if a smart decision-support system improves the daily car distribution and thereby lifts net sales by, say, \$200 per car (which is only a 1.33 percent increase in the price of an average \$15,000 off-lease car), the leasing company would increase its annual profits by hundreds of million of dollars. Exploring this possibility is clearly worthwhile.

Solution

To address this problem, we developed an intelligent system comprising several building blocks, including prediction, optimization, and adaptation modules. All three modules involve research challenges, which we briefly review in the “Research Issues in Dynamic Optimization” sidebar.

Question

Discuss the concept of adaptive business intelligence.

Source: <http://cs.adelaide.edu.au/~zbyszek/Papers/IEEE.pdf>

3.4 Summary

- A decision support system (DSS) is a computer program application that analyzes business data and presents it so that users can make business decisions more easily.
- In terms of dealing with many managerial and executive decisions that are made daily in the workplace, decision support systems can be a great tool for supervisors, department heads, and site managers.
- A decision support system is an internal business element that provides owners, executives, and other employees with information for making informed decisions.
- Companies can define their own characteristics of decision support systems.
- A typical Decision support systems has four components: data management, model management, knowledge management and user interface management.
- The data management component performs the function of storing and maintaining the information that you want your Decision Support System to use.
- The model management component consists of both the Decision Support System models and the Decision Support System model management system.
- The user interface management component allows you to communicate with the Decision Support System.

3.5 Keywords

Data Management: The data management component performs the function of storing and maintaining the information that you want your Decision Support System to use.

DSS: A decision support system (DSS) is a computer program application that analyzes business data and presents it so that users can make business decisions more easily.

Knowledge Management: The knowledge management component provides information about the relationship among data that is too complex for a database to represent.

Model Management: The model management component consists of both the Decision Support System models and the Decision Support System model management system.

Model: A model is a representation of some event, fact, or situation.

Task-oriented Activities: Task-oriented activities means that the support system can handle more than one activity at a time, which is essential.

User Interface Management: The user interface management component allows you to communicate with the Decision Support System.

User Interface: The user interface is the part of the system you see through it when enter information, commands, and models.

3.6 Review Questions

1. Explain the concept of decision support system.
2. Illustrate how decision support systems can be a great tool for supervisors, department heads, and site managers.
3. Discuss the Configuration of DSS.

Notes

4. Discuss the Characteristics of decision support system.
5. Interaction also represents important characteristics of decision support systems. Comment.
6. What are the components of Decision Support Systems.
7. Discuss the functions performed by data management component.
8. Explain the use of Model Management Component with example.
9. Which component is used to communicate with the Decision Support System? Discuss.
10. Make distinction between user interface component and knowledge management component.

Answers: Self Assessment

- | | |
|----------------------------------|----------------------------------|
| 1. Decision Support System (DSS) | 2. Business Intelligence systems |
| 3. Tool | 4. False |
| 5. True | 6. True |
| 7. False | 8. Data management |
| 9. Model | 10. Model management |
| 11. User Interface | 12. User interface |
| 13. Knowledge management | 14. Database management system |
| 15. Variables | |

3.7 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

<http://anale-informatica.tibiscus.ro/download/lucrari/7-1-33-Stanciu.pdf>

<http://dspace.upce.cz/bitstream/10195/32436/1/CL585.pdf>

http://mis.postech.ac.kr/class/MEIE580_DSS/paper/Decision%20Support%20Systems_A%20Summary,%20Problems,%20and%20Future%20Trends.pdf

http://www.cadrc.calpoly.edu/pdf/decision_brochure.pdf

Unit 4: Dialogue Management, Data Management and Model Management for DSS

Notes

CONTENTS

Objectives

Introduction

- 4.1 Dialogue Management for DSS
 - 4.1.1 Support for Model Construction and Model Analysis
 - 4.1.2 Support for Reasoning about the Problem Structure in Addition to Numerical Calculations
 - 4.1.3 Support for both Choice and Optimization of Decision Variables
 - 4.1.4 Graphical Interface
- 4.2 Data Management for DSS
- 4.3 Models Management for DSS
 - 4.3.1 Model Base
 - 4.3.2 Model Building Blocks and Routines
 - 4.3.3 Model Directory
- 4.4 Summary
- 4.5 Keywords
- 4.6 Review Questions
- 4.7 Further Readings

Objectives

After studying this unit, you will be able to:

- Explain Dialogue management component
- Discuss data management component
- Discuss model management component

Introduction

A DSS is an interactive, flexible, and adaptable CBIS, specially developed for supporting the solution of a non-structured management problem for improved decision making. It utilizes data, it provides easy user interface, and it allows for the decision maker's own insights. DSS may utilize models, is built by an interactive process (frequently by end-users), supports all the phases of the decision making, and may include a knowledge component. In this unit, we will discuss the various components of decision support system.

4.1 Dialogue Management for DSS

While the quality and reliability of modeling tools and the internal architectures of DSSs are important, the most crucial aspect of DSSs is, by far, their user interface. Systems with user interfaces that are cumbersome or unclear or that require unusual skills are rarely useful and accepted in practice. The most important result of a session with a DSS is insight into the decision problem. In addition, when the system is based on normative principles, it can play a tutoring role; one might hope that users will learn the domain model and how to reason with it over time, and improve their own thinking.

A good user interface to DSSs should support model construction and model analysis, reasoning about the problem structure in addition to numerical calculations and both choice and optimization of decision variables.

The term “process” in the context of user interface process refers to the flow of information (a) from the user to the system and (b) from the system to the user. It is handled by the User Interface Management System (UIMS). The UIMS processes user commands, issued in whatever action language it requires, and passes them on to the data and model management subsystems. In the reverse direction, it presents information from those subsystems to the user.



Notes Increasingly, the action language is based on Web or operating system GUI concepts. It may also incorporate natural language processing capabilities.

4.1.1 Support for Model Construction and Model Analysis

User interface is the vehicle for both model construction (or model choice) and for investigating the results. Even if a system is based on a theoretically sound reasoning scheme, its recommendations will be as good as the model they are based on. Furthermore, even if the model is a very good approximation of reality and its recommendations are correct, they will not be followed if they are not understood. Without understanding, the users may accept or reject a system’s advice for the wrong reasons and the combined decision-making performance may deteriorate even below unaided performance. A good user interface should make the model on which the system’s reasoning is based transparent to the user.

Modeling is rarely a one-shot process, and good models are usually refined and enhanced as their users gather practical experiences with the system recommendations.



Caution It is important to strike a careful balance between precision and modeling efforts; some parts of a model need to be very precise while others do not.

A good user interface should include tools for examining the model and identifying its most sensitive parts, which can be subsequently elaborated on. Systems employed in practice will need their models refined, and a good user interface should make it easy to access, examine, and refine its models.

4.1.2 Support for Reasoning about the Problem Structure in Addition to Numerical Calculations

While numerical calculations are important in decision support, reasoning about the problem structure is even more important. Often when the system and its model are complex it is insightful

for the decision maker to realize how the system variables are interrelated. This is helpful in designing creative decision options but also in understanding how a policy decision will impact the objective. Graphical models, such as those used in decision analysis or in equation-based and hybrid systems, are particularly suitable for reasoning about structure. Under certain assumptions, a directed graphical model can be given a causal interpretation. This is especially convenient in situations where the DSS autonomously suggests decision options; given a causal interpretation of its model, it is capable of predicting effects of interventions. A causal graph facilitates building an effective user interface.



Did u know? The system can refer to causal interactions during its dialogue with the user, which is known to enhance user insight.

4.1.3 Support for both Choice and Optimization of Decision Variables

Many DSSs have an inflexible structure in the sense that the variables that will be manipulated are determined at the model-building stage. This is not very suitable for planning of the strategic type when the object of the decision-making process is identifying both the objectives and the methods of achieving them.

For example, changing policy variables in a spreadsheet-based model often requires that the entire spreadsheet be rebuilt. If there is no support for that, few users will consider it as an option. This closes the world of possibilities for flexible reframing of a decision problem in the exploratory process of searching for opportunities.

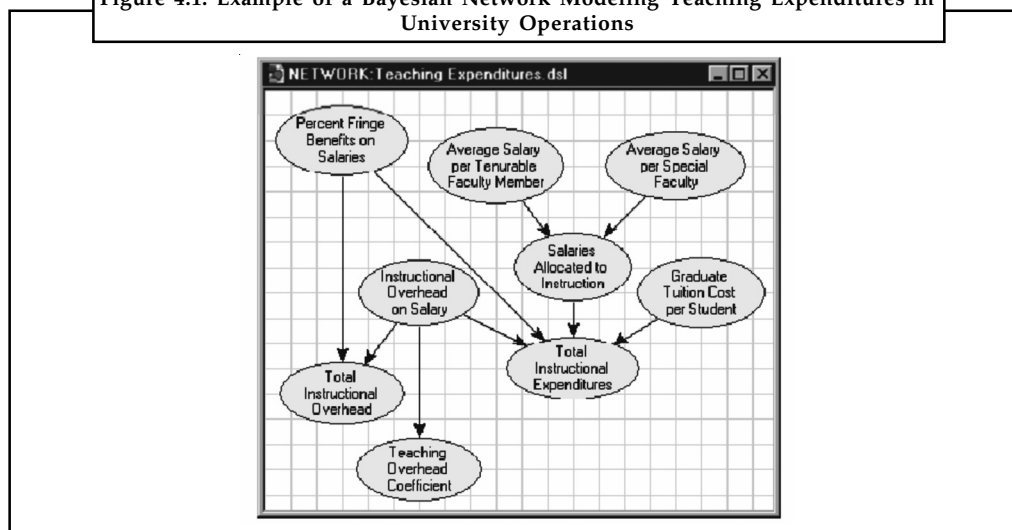
Support for both choice and optimization of decision variables should be an inherent part of DSSs.

4.1.4 Graphical Interface

Insight into a model can be increased greatly at the user interface level by a diagram representing the interactions among its components.

For example, a drawing of a graph on which a model is based is shown in Figure 4.1.

Figure 4.1: Example of a Bayesian Network Modeling Teaching Expenditures in University Operations

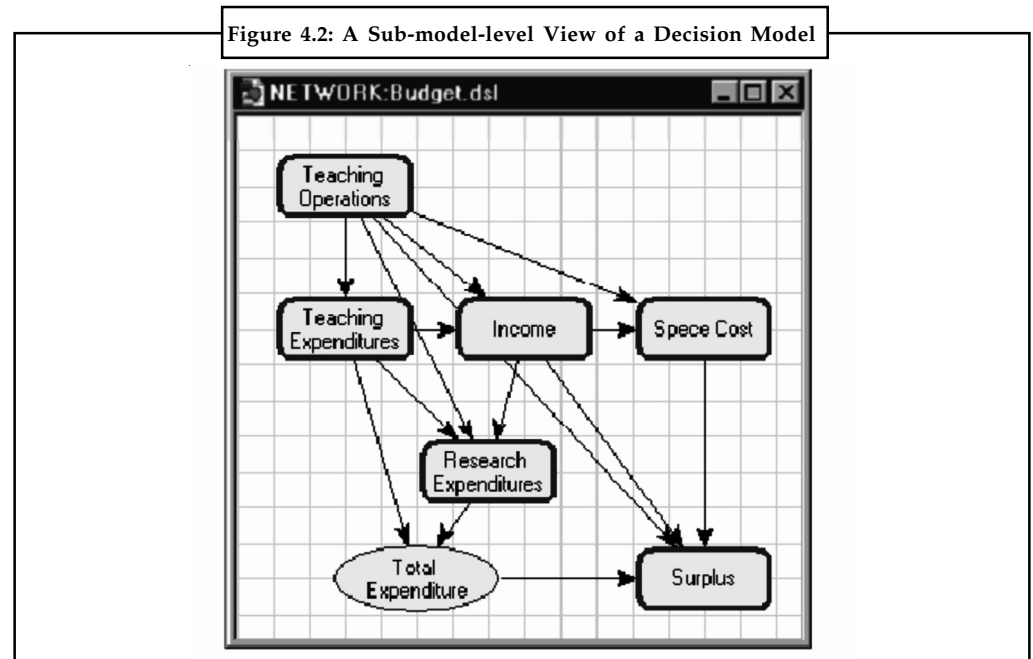


Source: <http://www.pitt.edu/~druzdel/psfiles/dss.pdf>

Notes

This graph is a qualitative, structural explanation of how information flows from the independent variables to the dependent variables of interest.

As models may become very large, it is convenient to structure them into sub-models, groups of variables that form a subsystem of the modeled system. Such sub-models can be again shown graphically with interactions among them, increasing simplicity and clarity of the interface. Figure 4.2 shows a sub-model-level view of a model developed in our ESP project. Note that the graph in Figure 4.1 is an expanded version of the Teaching Expenditures sub-model in Figure 4.2.



Source: <http://www.pitt.edu/~druzdzl/psfiles/dss.pdf>



Notes The user can navigate through the hierarchy of the entire model in her quest for insight, opening and closing sub-models on demand.

Self Assessment

Fill in the blanks:

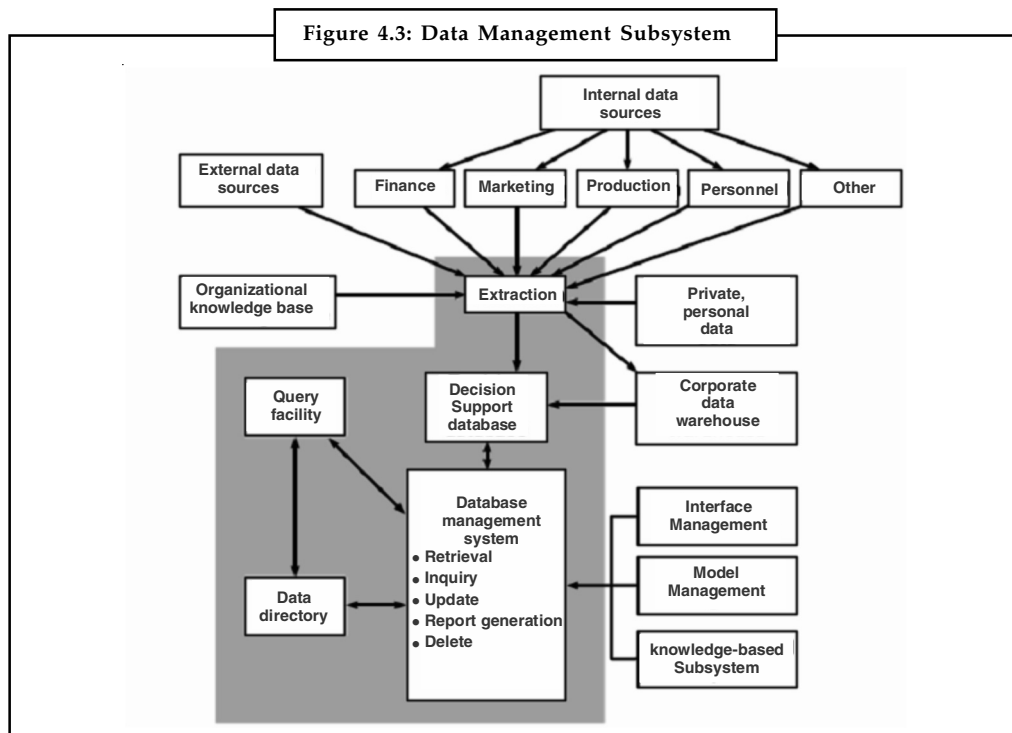
1. The term “.....” in the context of user interface process refers to the flow of information (a) from the user to the system and (b) from the system to the user.
2. The processes user commands, issued in whatever action language it requires, and passes them on to the data and model management subsystems.
3. is the vehicle for both model construction (or model choice) and for investigating the results.
4. A good user interface should include for examining the model and identifying its most sensitive parts, which can be subsequently elaborated on.
5. As models may become very large, it is convenient to structure them into....., groups of variables that form a subsystem of the modeled system.

4.2 Data Management for DSS

Notes

The data management subsystem is composed of:

- DSS database
- DBMS
- Data directory
- Query facility



These are discussed below:


- **Database:** Internal data come mainly from the organization's transaction processing system. External data include industry data, market research data, census data, regional employment data, government regulations, tax rate schedules, and national economic data. Private data can include guidelines used by specific decision makers and assessments of specific data and/or situations.
- **Data extraction:** The process of capturing data from several sources, synthesizing them, summarizing them, determining which of them are relevant, and organizing them, resulting in their effective integration.
- **DBMS:** Software for establishing, updating, and querying (e.g., managing) a database
- **Query Facility:** The (database) mechanism that accepts requests for data, accesses them, manipulates them, and queries them
- **Directory:** A catalog of all the data in a database or all the models in a model base.

Key database and database management system issues include:

- Data quality
- Data integration

Notes

- Scalability
- Data security



Task Compare and contrast internal data and external data.

Self Assessment

Fill in the blanks:

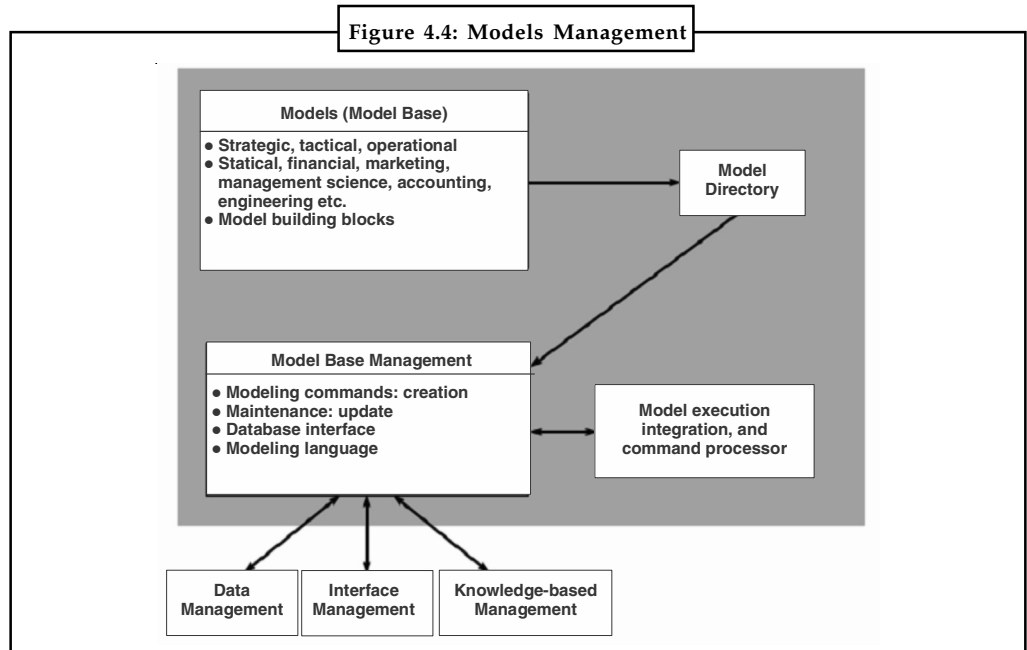
6. data come mainly from the organization’s transaction processing system.
7. data can include guidelines used by specific decision makers and assessments of specific data and/or situations.
8. is the process of capturing data from several sources, synthesizing them, summarizing them, determining which of them are relevant, and organizing them, resulting in their effective integration.
9. is a catalog of all the data in a database or all the models in a model base.

4.3 Models Management for DSS

Model management is one of the most important areas and is widely recognized as a key component of decision support systems. One of the prime tasks of model management is to perform mod integration which consists of combining existing DSS models and components. It involves tasks such as connecting models, formulating composite models from existing models, collection of models an model reuse and analysis.



Did u know? Model integration, over the past two decades, has emerged as one of the most widely researched topics in the field of MOD management.



4.3.1 Model Base

Notes

A collection of pre-programmed quantitative models (e.g., statistical, financial, optimization) organized as a single unit. Four categories of models with the model base:

- Strategic models
- Tactical models
- Operational models
- Analytical models

Strategic Models: Models that represent problems for the strategic level (i.e., executive level) of management.

Tactical Models: Models that represent problems for the tactical level (i.e., mid-level) of management.

Operational Models: Models that represent problems for the operational level of management.

Analytical Models: Mathematical models into which data are loaded for analysis.



Task Make distinction between Tactical models and Operational models.

4.3.2 Model Building Blocks and Routines

Model building blocks are Pre-programmed software elements that can be used to build computerized models.



Example: A random-number generator can be employed in the construction of a simulation model.

Model Base Management System

MBMS software has four main functions:

- Model creation, using programming languages, DSS tools and/or subroutines, and other building blocks
- Generation of new routines and reports
- Model updating and changing
- Model data manipulation

4.3.3 Model Directory

- Model execution is the process of controlling the actual running of the model
- Model integration involves combining the operations of several models when needed
- A model command processor is used to accept and interpret modeling instructions from the user interface component and route them to the MBMS, model execution, or integration functions

Notes

Self Assessment

Fill in the blanks:

10. One of the prime tasks of is to perform mod integration which consists of combining existing DSS models and components.
11. is a collection of pre-programmed quantitative models (e.g., statistical, financial, optimization) organized as a single unit.
12. Models that represent problems for the strategic level (i.e., executive level) of management are called
13. are the mathematical models into which data are loaded for analysis.
14. are Pre-programmed software elements that can be used to build computerized models.
15. is the process of controlling the actual running of the model.



Case Study

KOC Seismic Data Management System

Challenge: Develop and implement a master database for seismic data, new group workflows, and standards for managing and handling physical seismic data.

Solution: Schlumberger and KOC E&P Information Management planned for a master database, containing all post-stack seismic data. This involved creation of new workflows, standards, and conventions related to seismic data across several KOC groups, including research and technology, exploration, and field development. In addition, the plan addressed management of physical seismic tapes and reports.

Result:

- Dedicated, comprehensive seismic database.
- Effective corporate workflows.
- Increased confidence in data quality.
- Barcoded field tapes and seismic reports, stored and loaded to the eSearch physical data management systems.
- Easily interrogated repository, data is found in minutes, rather than days.

Schlumberger Information Management consultants conducted a site assessment to investigate the existing data management infrastructure, workflows, and systems. Data management procedures were not clearly documented or universally understood. Further, there were no master databases for seismic data storage. This translated to a lack of trust in the quality of the available data. In addition, the need to improve physical asset management and capture interpretation results was identified.

Based on the findings of this assessment, the project was initiated. The Schlumberger team used ProSource E&P corporate information management system—based on Seabed advanced E&P data store system—to build the seismic database. After installation and configuration, standard naming conventions were defined. Seismic data was then collated from users and stored tapes, and exported from existing data management systems. The

Contd....

loading and quality control procedures were defined and seismic trace and navigation data was loaded to the Seabed system.

To successfully manage the physical data items, the team created inventories for documents and tapes. The physical inventory data were then loaded to the eSearch system for powerful, searchable storage. A gap analysis was undertaken, per asset, for any missing tapes and associated reports. The team then recommended suitable workflows to manage new seismic data going forward.

Productivity and Data Confidence

KOC now benefits from a dedicated and comprehensive seismic database. This promotes confidence in data quality, when previously there was little. Over 45,000 field tapes have been barcoded, stored, and loaded to the eSearch database. Seismic reports were also processed in this way and stored in an exploration library. Users now benefit from an easily interrogated repository, and can request any tape or report through the system.

The Schlumberger team ran a survey to define efficiency and productivity improvements realized through the deployment. Users reported significant time savings in obtaining seismic data; tasks which previously took days now take minutes using the master database.

Question

Discuss the concept of developing and implementing a master database for seismic data.

Source: http://www.slb.com/resources/case_studies/software/cs_prosource_koc.aspx

4.4 Summary

- The term “process” in the context of user interface process refers to the flow of information (a) from the user to the system and (b) from the system to the user.
- User interface is the vehicle for both model construction (or model choice) and for investigating the results.
- Many DSSs have an inflexible structure in the sense that the variables that will be manipulated are determined at the model-building stage.
- As models may become very large, it is convenient to structure them into sub-models, groups of variables that form a subsystem of the modeled system.
- The data management subsystem is composed of DSS database, DBMS, Data directory, Query facility.
- One of the prime tasks of model management is to perform mod integration which consists of combining existing DSS models and components.
- Model base is a collection of pre-programmed quantitative models (e.g., statistical, financial, optimization) organized as a single unit.
- Model building blocks are Pre-programmed software elements that can be used to build computerized models.

4.5 Keywords

Data Extraction: It is the process of capturing data from several sources, synthesizing them, summarizing them, determining which of them are relevant, and organizing them, resulting in their effective integration.

Notes

Directory: It is a catalog of all the data in a database or all the models in a model base.

Model Base: It is a collection of pre-programmed quantitative models (e.g., statistical, financial, optimization) organized as a single unit.

Model Building Blocks: Model building blocks are Pre-programmed software elements that can be used to build computerized models.

Process: The term “process” in the context of user interface process refers to the flow of information (a) from the user to the system and (b) from the system to the user.

Query Facility: It is the (database) mechanism that accepts requests for data, accesses them, manipulates them, and queries them.

User Interface Management System (UIMS): The UIMS processes user commands, issued in whatever action language it requires, and passes them on to the data and model management subsystems.

User Interface: User interface is the vehicle for both model construction (or model choice) and for investigating the results.

4.6 Review Questions

1. What is user interface process? Discuss.
2. Explain the concept of Dialogue Management for DSS.
3. Describe the concept of Graphical Interface for DSS.
4. What are the components of data management subsystem? Discuss.
5. Discuss the concept of data extraction.
6. Explain the issues related to database management system.
7. Illustrate the use of Models Management for DSS.
8. What is Model base? Discuss the different categories of models with model base.
9. Discuss the functions of Model base management system.
10. Explain the concept of model directory.

Answers: Self Assessment

- | | |
|-----------------------|--|
| 1. Process | 2. User Interface Management System (UIMS) |
| 3. User interface | 4. Tools |
| 5. Sub-models | 6. Internal |
| 7. Private | 8. Data extraction |
| 9. Directory | 10. Model management |
| 11. Model base | 12. Strategic models |
| 13. Analytical models | 14. Model building blocks |
| 15. Model execution | |

4.7 Further Readings

Notes



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

http://www.cs.ubbcluj.ro/~per/Dss/Dss_3.pdf

<http://www.isqa.unomaha.edu/pietron/oat/Chapt15.pdf>

<http://www.recordit.org/workshop/Present%20Claudia.pdf>

http://www.umsl.edu/~sauterv/analysis/488_f02_papers/dss.html

Unit 5: Modeling and Analysis

CONTENTS

Objectives

Introduction

5.1 Modeling for MSS

5.1.1 Some Major Modeling Issues

5.2 Static and Dynamic Models

5.3 Certainty, Uncertainty and Risk

5.3.1 Decision Making under Certainty

5.3.2 Decision Making under Uncertainty

5.3.3 Decision Making under Risk

5.4 MSS Modeling in Spreadsheets

5.5 Simulation

5.5.1 Characteristics of Simulation

5.5.2 Advantages of Simulation

5.5.3 Disadvantages of Simulation

5.5.4 Methodology of Simulation

5.5.5 Simulation Types

5.5.6 Visual Interactive Simulation

5.5.7 Conventional Simulation Inadequacies

5.5.8 Visual Interactive Simulation

5.5.9 Visual Interactive Models

5.6 Optimization Via Mathematical Programming

5.7 Summary

5.8 Keywords

5.9 Review Questions

5.10 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss Modeling for MSS
- Explain Static and Dynamic Models
- Discuss the Concept of Certainty, Uncertainty and Risk
- Describe MSS Modeling in Spreadsheets
- Explain the Concept of Simulation
- Discuss Optimization via Mathematical Programming

Introduction

Notes

Modeling is a key element in most DSS and a necessity in a model-based DSS. In this unit, we will discuss Modeling for MSS, explain static and dynamic models, discuss the concept of certainty, uncertainty, and risk, describe MSS modeling in spreadsheets, explain the concept of Simulation, discuss optimization via Mathematical Programming.

5.1 Modeling for MSS

There are many classes of models, and there are often many specialized techniques for solving each one.

Simulation is a common modeling approach, but there are several others:

- A general model, based on an algorithm (for example to make transportation cost estimates). This model is programmed directly into the DSS.
- A demand-forecasting model (statistics based).
- A distribution center location model. This model uses aggregated data (a special modeling technique) and is solved with a standard linear/integer optimization package.
- A transportation model (i.e. a specialization of a linear programming model) to determine the best shipping option from sources to distribution centers (fed to it from the previous model) and hence to customers. This model is solved using commercial software and is loosely integrated with the distribution location model.



Caution The DSS must interface with commercial software and integrate the model.


- A financial and risk simulation model that takes into consideration some qualitative factors that require important human judgment.
- A geographical information system (GIS; effectively a graphical model of the data) for a user interface.

5.1.1 Some Major Modeling Issues

1. **Identification of the Problem and Environmental Analysis:** An important aspect of the environmental analysis is environmental scanning and analysis, which is the monitoring, scanning, and interpretation of collected information. No decision is made in a vacuum. It is important to analyze the scope of the domain and the forces and dynamics of the environment. A decision maker needs to identify the organizational culture and the corporate decision-making process. It is entirely possible that environmental factors have created the current problem. The problem must be understood, and everyone involved should share the same frame of understanding because the problem will ultimately be represented by the model in one form or another. Otherwise, the model will not help the decision maker.
2. **Variable identification:** Identification of a model's variables is critical, as are relationships of the variables. Influence diagrams, which are graphical models of mathematical models, can facilitate the identification process. A more general form of an influence diagram, a cognitive map, can help a decision maker develop a better understanding of a problem, especially of variables and their interactions.
3. **Forecasting – Predictive Analytics:** Forecasting is predicting the future. This form of predictive analytics is essential for construction and manipulating models because when

Notes

a decision is implemented, the results usually occur in the future. Whereas DSS are typically designed to determine what will be, traditional MIS report what is or what was. There is no point in running a what-if (sensitivity) analysis on the past because decisions made then have no impact in the future. Forecasting is getting easier as software vendors automate many of the complications of developing such models.



Notes Forecasting system that incorporates its predictive analytics technology is ideally for retailers. This software is more automated than most other forecasting packages.

4. **Multiple Models:** A DSS can include several models which represents a different part of the decision-making problem.



Example: P&G supply-chain DSS includes a location model to locate distribution center, a product-strategy model, a demand-forecasting model, a cost generation model, a financial-and risk-simulation model, and even a GIS model.

5. **Model Categories:** There are seven groups of DSS models. Each category can be applied to either a static or a dynamic model, which can be constructed under assumed environments of certainty, uncertainty, or risk. To expedite model construction, we can use special decision analysis systems that have modeling languages and capabilities embedded in them. These includes spreadsheets, data mining systems, OLAP systems, and even fourth-generation languages (formerly financial planning languages).

Table 5.1: Model Categories

| Category | Process and Objectives | Representative Techniques |
|--|--|---|
| Optimization of problems with few alternatives | Find the best solution from a small number of alternatives | Decision tables, decision tree |
| Optimization via algorithm | Find the best solution from a large number of alternatives, using a step-by-step improvement process | Linear programming models, network models |
| Optimization via an analytic formula | Find the best solution in one step, using a formula | Some inventory models |
| Simulation | Finding a good enough solution or the best among the alternatives checked, using experimentation | Several types of simulation |
| Heuristics | Find a good enough solution, using rules | Heuristic programming, expert systems |
| Predictive models | Predict the future for a given scenario | Forecasting models, Markov analysis |
| Other models | Solve a what-if case, using a formula | Financial modeling, waiting lines |

6. **Model Management:** Models, like data, must be managed to maintain their integrity and thus their applicability. Such management is done with the aid of the model base management systems (MBMS), which are analogous to MBMS database management systems (DBMS).
7. **Knowledge-based Modeling:** DSS uses mostly quantitative models, whereas expert systems use qualitative, knowledge-based models in their applications. Some knowledge is necessary to construct solvable (and therefore usable) models.

8. **Current Trends in Modeling:** One recent trend in modeling involves the development of model libraries and solution technique libraries. Some of these codes can be run directly on the owner's web server for free, and others can be downloaded and run on an individual's PC, Unix machine, or server. The availability of these codes means that powerful optimization and simulation packages are available to decision makers who may have only experienced these tools from the perspective of classroom problems.

There is a clear trend toward developing and using Web tools and software to access and even run software to perform modeling, optimization, simulation, and so on. This has, in many ways, simplified the application of many models to real-world problems. Another trend, involves the lack of understanding of what models and their solutions can do in the real world.

There is a continuing trend toward making MSS models completely transparent to decision maker.



Example: Multidimensional analysis (modeling) involves data analysis in several dimensions.

There is also a trend to model a model to help in its analysis. An influence diagram is a graphical representation of a model; that is, it is a model of a model. Some influence diagram software packages are capable of generating and solving the resultant model.

Self Assessment

Fill in the blanks:

- is a key element in most DSS and a necessity in a model-based DSS.
- An is a graphical representation of a model; that is, it is a model of a model.

5.2 Static and Dynamic Models

DSS models can be classified as static or dynamic models.

Static Models

Static Models describe a single interval of a situation. A static model takes a single snapshot of a situation. During this snapshot, everything occurs in a single interval.



Example: A decision about whether to make or buy a product is static in nature.

Dynamic Models

Dynamic models represent scenarios that change over time.

For example, a 5 years profit and loss projections in which the input data (costs, prices, and quantities) change from year to year.

Dynamic models are time dependent. Dynamic models use, represent, or generate trends and patterns over time, show averages per periods, moving averages, and comparative analyses. When a static model is constructed to describe a given situation, it can be expanded to represent the dynamic nature of the problem.



Task Compare and contrast static and dynamic models.

Notes

Self Assessment

Fill in the blanks:

3. models describe a single interval of a situation.
4. models represent scenarios that change over time.

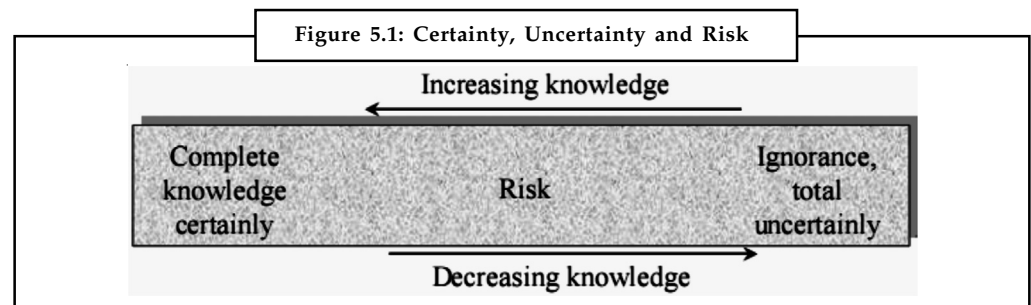
5.3 Certainty, Uncertainty and Risk

Decision making process involves evaluating and comparing alternatives. During the process, it is necessary to predict the future outcome of each proposed alternative. Decision situation are often classified on the basis of what the decision maker knows or believes about the forecasted result. We customarily classify this knowledge into three categories, ranging from complete knowledge to total ignorance:

- Certainty
- Risk
- Uncertainty



Did u know? When we are develop models, any of these conditions can occur, and different kind of models are appropriate for each case.



5.3.1 Decision Making under Certainty

In decision making under certainty, it is assumed that complete knowledge is assumed available – the decision maker knows exactly what the outcome of each course of action will be (a deterministic environment). The decision maker is viewed as a perfect predictor on the future. Certainly models are relatively easy to develop and solve, and they can yield optimal solution.

5.3.2 Decision Making under Uncertainty

In decision making under uncertainty, the decision maker considers situations in which several outcomes are possible for each course of action. The decision maker does not know, or cannot estimate, the probability of occurrence of the possible outcomes. The problem, in this case, is more difficult because there is insufficient information – modeling involves assessment of the decision maker’s attitude toward risk.

5.3.3 Decision Making under Risk

In decision made under risk (a probabilistic or stochastic decision-making situation), is one in which the decision maker must consider several possible outcomes for each alternative, each

with a given probability of occurrence. The decision maker can assess the degree of risk associated with each alternative (calculated risk).

Risk analysis can be performed by calculating the expected value of each performed alternative and selecting the one with the best expected value.

Self Assessment

Fill in the blanks:

- In decision making under, it is assumed that complete knowledge is assumed available – the decision maker knows exactly what the outcome of each course of action will be.
- In decision making under, the decision maker considers situations in which several outcomes are possible for each course of action.
- can be performed by calculating the expected value of each performed alternative and selecting the one with the best expected value.

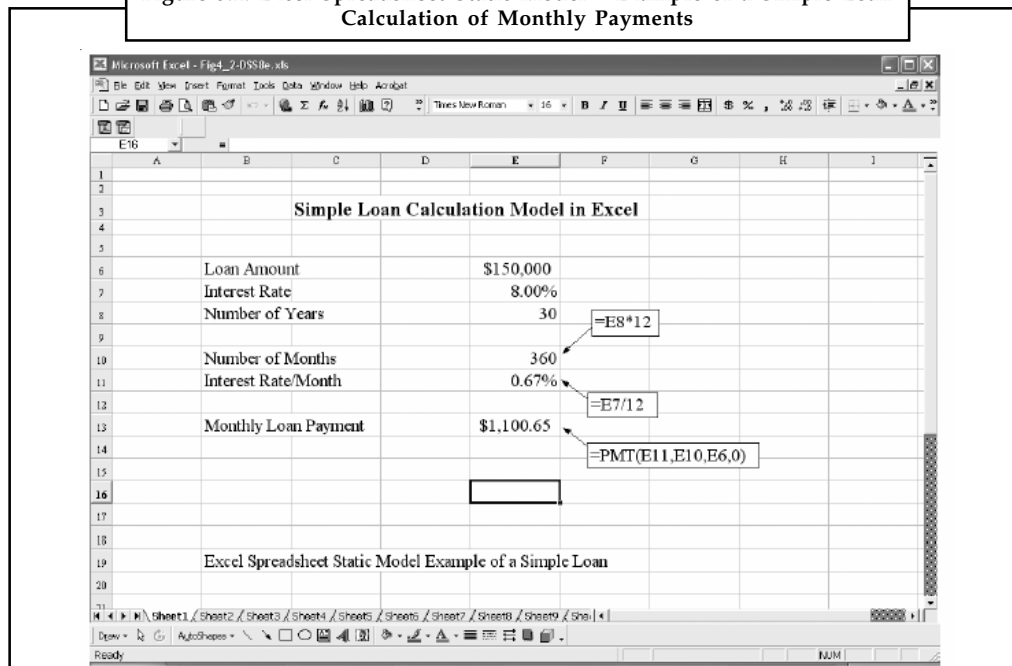
5.4 MSS Modeling in Spreadsheets

Models can be developed and implemented in a variety of programming languages and systems: from third, fourth, and fifth-generation to Computer Aided Software Engineering (CASE) systems and other systems that automatically generate usable software. We focus primarily on spreadsheets, model languages, and transparent data analysis tools – the most popular end-user modeling tool. Spreadsheets include extensive statistical, forecasting, and other modeling and database management capabilities, functions and routines.



Example: Examples of excel spread sheet static model and dynamic model are shown in figure 5.2 and figure 5.3 respectively.

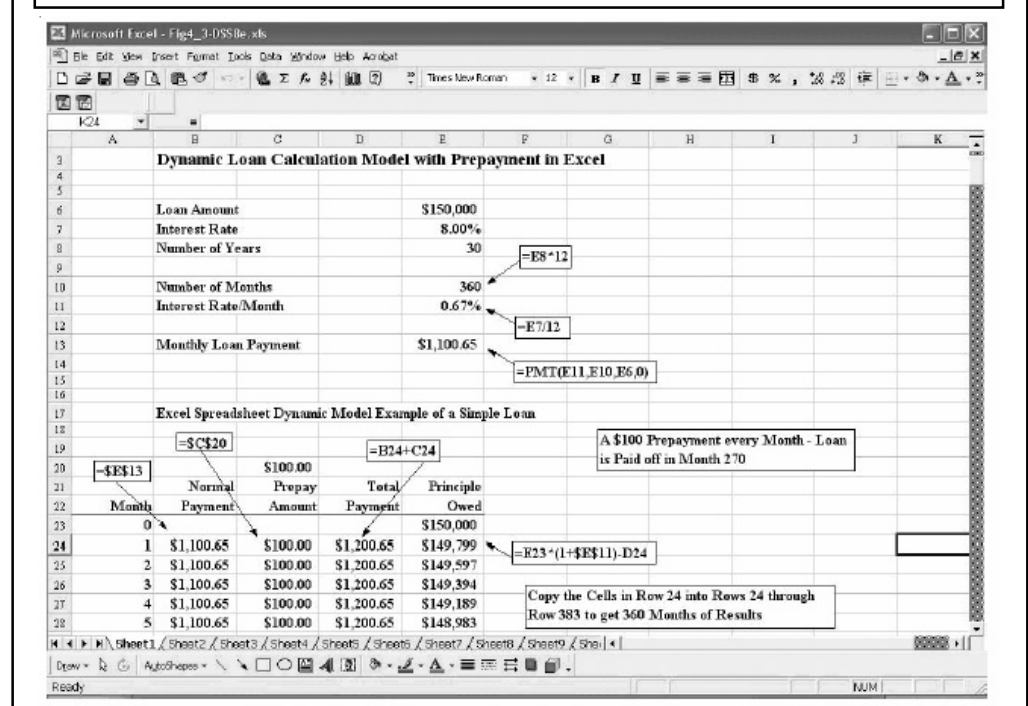
Figure 5.2: Excel Spreadsheet Static Model – Example of a Simple Loan Calculation of Monthly Payments



Notes

- Other important spreadsheet features include what-if analysis, goal seeking, trial and error, optimization, data management, and programmability
- Most spreadsheet packages provide fairly seamless integration because they read and write common file structures and easily interface with databases and other tools
- Static or dynamic models can be built in a spreadsheet.

Figure 5.3: Excel Spreadsheet Dynamic Model – Example of a Simple Loan Calculation of Monthly Payments and the Effects of Prepayment



Self Assessment

State True or False:

8. Spreadsheets include extensive statistical, forecasting, and other modeling and database management capabilities, functions, and routines.
9. Most spreadsheet packages does not provide fairly seamless integration.

5.5 Simulation

“Simulation is the appearance of reality. In MSS is a technique for Simulation conducting experiments with a computer on a model of a management system. Because DSS deals with semi-structured or unstructured situations, reality is complex, which may not be easily represented by optimization or other models but can often be handled by simulation. Simulation is one of the most commonly used DSS methods.”

5.5.1 Characteristics of Simulation

The Characteristics of Simulation are:

- Models represent reality, simulation imitates it.

- Simulation is a technique for conducting experiments.
- Simulation is a descriptive method – there is no automatic search for an optimal solution.
- A simulation model describes or predicts the characteristics of a given system under different conditions.
- The simulation process usually repeats an experiment many times to obtain an estimate (and a variance) of the overall effect of certain action.
- Simulation is used when the problem is too complex to be treated using numerical optimization techniques.

5.5.2 Advantages of Simulation

Advantages of simulation are as follows:

- The theory is fairly straightforward.
- A great amount of time compression can be attained.
- Simulation is descriptive – allows the manager to pose what-if questions → trial-and-error approach to problem solving and can do so faster, at less expense, more accurately, and with less risk.
- The model is built from the manager's perspective.
- The simulation model is built for one particular problem and typically cannot solve any other problem – every component in the model corresponds to part of the real system.
- Simulation is often the only DSS modeling method that can readily handle relatively unstructured problems.

5.5.3 Disadvantages of Simulation

Disadvantages of simulation are as follows:

- The optimal solution cannot be guaranteed, but relatively good ones are generally found.
- Simulation model construction can be a slow and costly process, although newer modeling systems are easier to use than ever.
- Solutions and inferences from a simulation study are usually not transferable to other problems because the model incorporates unique problem factors.
- Simulation is sometimes so easy to explain to managers that analytic methods are often overlooked.
- Simulation software sometimes requires special skills because of the complexity of the formal solution method.

5.5.4 Methodology of Simulation

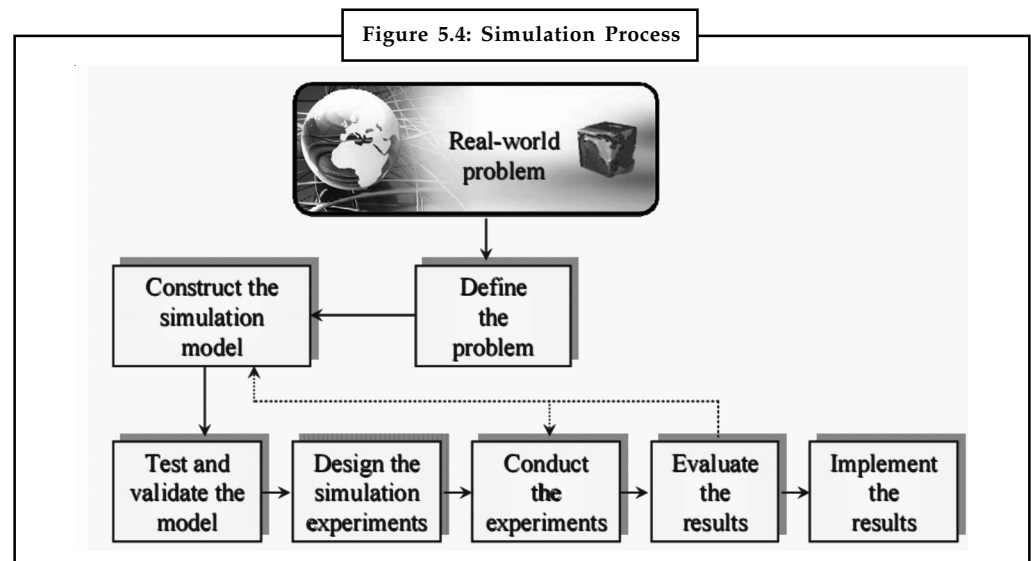
Simulation involves setting up a model of a real system and conducting repetitive experiments on it. The methodology consists of the following steps:

1. **Define the problem:** We examine and classify the real-world problem – we specify why a simulation approach is appropriate.
2. **Construct the simulation model:** This step involves determination of the variables and their relationships, as well as data gathering.

Notes

3. **Test and validate the model:** The simulation model must represent the system being studied.
4. **Design the simulation experiments:** When the model has been proven valid, an experiment is designed – how long to run the simulation (accuracy and cost?).
5. **Conduct the experiments:** involves issues ranging from random-number : generation to result presentation.
6. **Evaluate the results:** The results must be interpreted – sensitivity analysis.
7. **Implement the results:** The chances of success are better because the manager is usually more involved with the simulation process than other models.

The Process of Simulation is shown in the figure 5.4:



Source: http://www.cs.ubbcluj.ro/~per/Dss/Dss_11.pdf

5.5.5 Simulation Types

The types of Simulation are:

1. **Probabilistic Simulation:** The independent variables are probabilistic:
 - ❖ *Discrete distributions:* A situation with a limited number of events (variables) that can take on only a finite number of values.
 - ❖ *Continuous distributions:* Unlimited number of possible events that follow density functions, such as the normal distribution.
2. **Time Dependent Simulation:** It is important to know the precise time of arrival.
3. **Time Independent Simulation:** It is not important to know exactly when the event occurred.
4. **Object-oriented Simulation:** SIMPROCESS: object-oriented process modeling tool that lets the user create a simulation model by using screen-based objects. UML could be used in practice for modeling complex, real-time systems.
5. **Visual Simulation:** The graphical display of computerized results, which may include animation, is one of the most successful developments in computer-human interaction and problem solving.



Task Make distinction between Time Dependent Simulation and Time Independent Simulation.

5.5.6 Visual Interactive Simulation

We examine methods that show a decision maker a representation of the decision-making situation in action as it runs through scenarios of the decision maker's choices of alternatives. "These powerful methods overcome some of the inadequacies of conventional methods and help build trust in the solution attained because they can be visualized directly".

5.5.7 Conventional Simulation Inadequacies

Simulation is a descriptive and mathematics-based method for gaining insight into complex decision-making situations. Simulation does not usually allow decision makers to see how a solution to a complex problem evolves over time, nor can decision makers interact with the simulation. Simulation generally reports statistical results at the end of experiments.



Caution If the simulation results do not match the intuition or judgment of the decision maker, a confidence gap in the results can occur.

5.5.8 Visual Interactive Simulation

Visual Interactive Simulation (VIS) also known as Visual Interactive Modeling (VIM) and Visual Interactive Problem Solving is a simulation method that lets decision makers see what the model is doing and how it interacts with the decisions made, as they are made. The user can employ his knowledge to determine and try different decision strategies while interacting with the model. Decision maker also contribute to model validation, and support and trust their results. VIS uses animated computer graphics display to present the impact of different decisions. It differs from regular graphics in that the user can adjust the decision-making process and see the results of the intervention. A visual model is a graphic used as an integral part of decision making or problem solving. VIS displays the effects of different decisions in graphic form on a computer screen.



Did u know? The simulation system identified the most important input factors that significantly affected performance.

5.5.9 Visual Interactive Models

Visual Interactive Simulation can represent static or dynamic systems:


- Static models display a visual image of the result of one decision alternative at a time.
- Dynamic models display systems that evolve over time, and the evolution is represented by animation.

The latest visual simulation technology has been coupled with the concept of virtual reality, where an artificial world is created for a number of purposes, from training to entertainment to viewing data in an artificial landscape. VIM in DSS has been used in several operations

Notes

management decision. The method consists of priming a visual interactive model of a plant (or company) with its current status. The model then runs rapidly on a computer, allowing managers to observe how a plant is likely to operate in the future.

The VIM approach can also be used in conjunction with artificial intelligence.



Notes Integration of the two techniques adds several capabilities that range from the ability to build systems graphically to learning about the dynamics of the system.

Self Assessment

Fill in the blanks:

10. In MSS is a technique for conducting experiments with a computer on a model of a management system.
11. is a situation with a limited number of events (variables) that can take on only a finite number of values.
12. is the graphical display of computerized results, which may include animation.
13. Visual Interactive Simulation (VIS) is also known as

5.6 Optimization Via Mathematical Programming

Linear programming (LP) is the best-know technique in family of optimization tools called mathematical programming. In LP, all relationships among the LP variables are linear.

Mathematical programming is a family of tools designed to help solve managerial problems in which the decision maker must allocate scarce resources among competing activities to optimize a measurable goal.

Every LP problem is composed of:

- Decision variable – whose values are unknown and are searched for an objective function – a linear mathematical function that relates the decision variables to the goal, measures goal attainment, and is to be optimized.
- Objective function coefficients – unit profit or cost coefficients indicating the contribution to the objective of one unit of a decision variable.
- Constraints – expressed in the form of linear inequalities or equalities that limit resources or requirements; these relate the variables through linear relationships.
- Capacities – which describe the upper and sometimes lower limits on the constraints and variables, and
- Input/output (technology) coefficients – which indicate resource utilization for a decision variable.



Example:

An example of modeling in LP is shown below:

- The decision variables:
 X_1 = units of CC-7

X_2 = units of CC-8

- The result variable:

Total profit = Z

- The objective is to maximize:

$$Z = 8,000X_1 + 12,000X_2$$

- The uncontrollable variables (constraints):

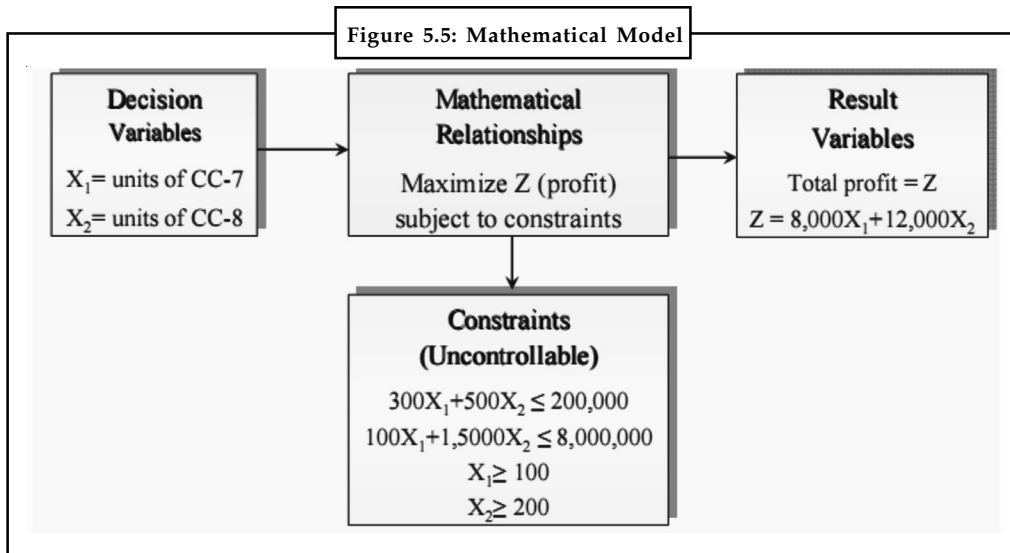
Labor constrain: $300X_1 + 500X_2 \leq 200,000$

Budget constrain: $10,000X_1 + 15,000X_2 \leq 8,000,000$

Marketing requirement for CC-7: $X_1 \geq 100$

Marketing requirement for CC-8: $X_2 \geq 200$

The Mathematical Model is:



Source: http://www.cs.ubbcluj.ro/~per/Dss/Dss_10.pdf

The Excel Solver Solution is:

Figure 5.6: Excel Solver Solution

| | X_1 | X_2 | | |
|----------------------------|---------------|---------------|-----------------|----------|
| Decision Variables: | 333.33 | 200.00 | | |
| Total Profit: | 8.00 | 12.00 | 5,066.67 | |
| Labor: | 0.30 | 0.50 | 200.00 | 200.00 |
| Budget: | 10.00 | 15.00 | 6,333.33 | 8,000.00 |
| X1 lower: | 1.00 | 0.00 | 333.33 | 100.00 |
| X2 lower: | 0.00 | 1.00 | 200.00 | 200.00 |

Source: http://www.cs.ubbcluj.ro/~per/Dss/Dss_10.pdf

Notes

The Excel Solver Answer Report is:

Figure 5.7: Excel Solver Answer Report

| Target Cell (Max) | | | | | | |
|-------------------|--------------|----------------|-------------|--|--|--|
| Cell | Name | Original Value | Final Value | | | |
| \$E\$7 | Total Profit | 5,066.67 | 5,066.67 | | | |

| Adjustable Cells | | | | | | |
|------------------|------------------------|----------------|-------------|--|--|--|
| Cell | Name | Original Value | Final Value | | | |
| \$C\$6 | Decision Variables: X1 | 333.33 | 333.33 | | | |
| \$D\$6 | Decision Variables: X2 | 200.00 | 200.00 | | | |

| Constraints | | | | | | |
|-------------|-----------|------------|------------------|-------------|-------------|--|
| Cell | Name | Cell Value | Formula | Status | Slack | |
| \$E\$8 | Labor: | 200.00 | \$E\$8<=\$F\$8 | Binding | 0 | |
| \$E\$9 | Budget: | 6,333.33 | \$E\$9<=\$F\$9 | Not Binding | 1666.666667 | |
| \$E\$10 | X1 lower: | 333.33 | \$E\$10>=\$F\$10 | Not Binding | 233.33 | |
| \$E\$11 | X2 lower: | 200.00 | \$E\$11>=\$F\$11 | Binding | 0.00 | |

Source: http://www.cs.ubbcluj.ro/~per/Dss/Dss_10.pdf

The Excel Solver Sensitivity Report is:

Figure 5.8: Excel Solver Sensitivity Report

| Adjustable Cells | | | | |
|------------------|------------------------|-------------|------------------|--|
| Cell | Name | Final Value | Reduced Gradient | |
| \$C\$6 | Decision Variables: X1 | 333.33 | 0.00 | |
| \$D\$6 | Decision Variables: X2 | 200.00 | 0.00 | |

| Constraints | | | | |
|-------------|-----------|-------------|---------------------|--|
| Cell | Name | Final Value | Lagrange Multiplier | |
| \$E\$8 | Labor: | 200.00 | 26.67 | |
| \$E\$9 | Budget: | 6,333.33 | 0.00 | |
| \$E\$10 | X1 lower: | 333.33 | 0.00 | |
| \$E\$11 | X2 lower: | 200.00 | -1.33 | |

Source: http://www.cs.ubbcluj.ro/~per/Dss/Dss_10.pdf

The Excel Solver Limits Report is:

Figure 5.9: Excel Solver Limits Report

| Target | | | | | | |
|--------|--------------|----------|--|--|--|--|
| Cell | Name | Value | | | | |
| \$E\$7 | Total Profit | 5,066.67 | | | | |

| Adjustable | | | Lower | Target | Upper | Target |
|------------|------------------------|--------|--------|----------|--------|----------|
| Cell | Name | Value | Limit | Result | Limit | Result |
| \$C\$6 | Decision Variables: X1 | 333.33 | 100.00 | 3,200.00 | 333.33 | 5,066.67 |
| \$D\$6 | Decision Variables: X2 | 200.00 | 200.00 | 5,066.67 | 200.00 | 5,066.67 |

Source: http://www.cs.ubbcluj.ro/~per/Dss/Dss_10.pdf

The most common optimization models can be solved by a variety of mathematical programming methods, including the following:

Notes

- Assignment (best matching of objects)
- Dynamic programming
- Goal programming
- Investment (maximizing rate of return)
- Linear and integer programming
- Network models for planning and scheduling
- Non-linear programming
- Replacement (capital budgeting)
- Simple inventory models (e.g., economic order quantity)
- Transportation (minimize cost of shipments)

Self Assessment

Fill in the blanks:

14. is the best-know technique in family of optimization tools called mathematical programming in which all relationships among the LP variables are linear.
15. is a family of tools designed to help solve managerial problems in which the decision maker must allocate scarce resources among competing activities to optimize a measurable goal.



Case Study

Multi-objectives Mathematical Programming using "Payoff Technique" for Andhra Pradesh

An Approximation of the Multi-objective Programming Problem

Multi-Objective Programming of vector optimization technique tackle the problem of simultaneous optimization of several objectives subject to a set of constraints, usually linear. As an optimum solution is undefined for several simultaneous objectives, MOP seeks to identify the set that contains efficient (non-dominated or Pareto optimal) solutions.

The elements of this efficient set are feasible solutions such that there are no other feasible solutions that can achieve the same or better performance for all the objectives and strictly better for at least one objective.

Given that the purpose of MOP is to generate the efficient set, the general nature of problem can be stated as:

$$\text{Eff } z(x) = [z_1(x), z_2(x), \dots, z_q(x)]$$

Subject to: $x \in F$

Where Eff means the search for the efficient solutions (in a minimizing or maximizing sense) and F represents the feasible set.

Contd....

Notes

It is necessary to resort to multi-objective techniques to generate or at least to approximate the efficient set. There are basically three approaches to undertake that task: (a) the constraint method, (b) the weighting method and (c) the multi-objective simplex method. Of these three approaches only the last one can obtain an exact representation of the efficient set.

The Pay-Off Matrix in Mop

One way of obtaining the initial and useful information of an MOP problem is to optimize each of the objectives separately over the efficient set and then to compute the value of each objective at each of the optimal solutions. In this way a square matrix, called the 'Pay-off Matrix' is obtained.

The Constraint Method

The basic idea of this method is to optimize one of the objectives while the others are specified as constraints. The efficient set is then generated by parameterize the right hand side of the objectives treated as restraints. Thus, for a MOP problem with q objectives to be maximized, the constraint method leads to the following mathematical programming problem:

$$\text{Maximize } Z_k(x)$$

$$\text{Subject to: } x \in F$$

$$Z_j(x) \geq L_j, j = 1, 2, \dots, k-1, k+1, \dots, q$$

Being $Z_j(x)$ the objective to be optimized. Through parametric variations of the right hand sides L_j the efficient set is generated.

The Weighting Method

The basic idea of this method is to combine all the objectives in to a single objective function. Each objective function is given a weight before all the objectives are added. Subsequently, the efficient set is generated through parametric variation of weights. Zadeh (1963) was first to propose this method. Thus, for a MOP problem with q objectives to be maximized, the weighting method leads to the following mathematical programming structure:

$$\text{Maximize } W_1 Z_1(x) + W_2 Z_2(x) + \dots + W_q Z_q(x)$$

$$\text{Subject to } x \in F$$

$$W > 0$$

Through parametric variations of the weights W the efficient set can be generated. It should be noted that the weighting method guarantees efficient solutions only when the weights are larger than zero ($W > 0$). If one of the weights is zero, and there are alternative optimal solutions, then the corresponding optimal solution provided by the weighting method can be inferior or non efficient. Further the weighting method can only generate extreme efficient points and not both the extreme and inferior one as the constraint method does.

India is predominantly an agricultural economy and thus prosperity to its people depends upon the progress of its agricultural sector. Agriculture and allied activities contribute 29.1 percent to the gross domestic product of India as compared to 2 percent in U.S., France, Norway, Japan, 5 percent in Korea and 49 percent in Ethiopia. Indian Agriculture employees 69 percent of total work force as compared to 2 percent in Tanzania, 93 percent in Nepal and is a major source of poverty alleviation and empowerment of the agrarian folk. It is also the major source of sustaining life for majority of its population by providing them

Contd....

food stuff, which is one of the basic needs for human existence. The classical economists have assigned a central place to food in their growth model. Starting from Malthus, the question of feeding the ever growing population has been engaging the attention of politicians and academicians alike. Food grains occupies an important place in human food so as self-sufficiency in food grain production was the derived goal of planned development.

For instance, a subsistence farmer may be interested in securing adequate food supplies for the family, maximizing cash income, increasing leisure, avoiding risk, etc. but not necessarily in that order. Similarly a commercial farmer may wish to maximize gross margin, minimize his indebtedness, acquire more land, reduce fixed costs, etc.

Technique of Analysis

The objective functions are optimized simultaneously in the multiple objective programming farm planning model. First, the pay-off matrix has been constructed using 'ideal point' which represents the optimum values of the objectives under consideration. In fact, this ideal point is not feasible because the objectives are in conflict, we select the efficient farm plans closest to it or best compromise by using compromise programming techniques. The worst element from each column of the pay-off matrix will be the 'Anti-Ideal Point'. Among the different techniques to generate the efficient set, a variant of the weighting method has been chosen known as Non-Inferior Set Estimation (NISE) method, as the most suitable multiple objective programming technique for generating the efficient set (Cohan, Church and Steer, 1979).

For some value of and different values of P different compromise solution for distant function Lp were obtained and the farmer/nation can chose any one solution for a given preferences of the different objectives out of the various compromise solutions. However the distance function Lp is usually used for P=1 and P2=α which shows a longest and the chebyshev distance in the geometric sense respectively (greater weight is given to the largest deviation). Therefore, maximum of the individual deviations is minimized with p = α For different values of P and j, we can generate different compromise solutions. The alternate with the lowest value for the distance function will be the best compromise solution because it is the nearest solution with respect to the ideal point. For L1 metric (P=1), the best compromise solution to the ideal point can be obtained by solving the following linear programming problems.

$$\text{i.e. Min } L_1 = \sum \delta_j \frac{Z_j^* - Z_j(\underline{x})}{Z_j^* - Z_j^*} \text{ subject to } (\underline{X}) \in \underline{F}$$

For L α matrix (P=α), the minimised by minimum of the individual deviation is minimized by solving the following linear programming model.

$$\text{Min } L_e = d_e$$

Such that

$$\delta_1 = \frac{Z_j^* - Z_j(\underline{x})}{Z_j^* - Z_j^*} \leq d_e$$

$$\delta_2 = \frac{Z_j^* - Z_j(\underline{x})}{Z_j^* - Z_j^*} \leq d_e$$

:
:
:

Contd....

Notes

$$\delta_k = \frac{Z_j^* - Z_j(\underline{x})}{Z_j^* - Z_j^*} \leq d_e$$

$(X) \in F$

Question

How to search for the efficient solutions? Discuss.

Source: <http://indianresearchjournals.com/pdf/IJSSIR/2012/July/3.pdf>

5.7 Summary

- Modeling is a key element in most DSS and a necessity in a model-based DSS.
- A static model takes a single snapshot of a situation. During this snapshot, everything occurs in a single interval.
- Dynamic models represent scenarios that change over time.
- In decision making under certainty, it is assumed that complete knowledge is assumed available – the decision maker knows exactly what the outcome of each course of action will be.
- In decision making under uncertainty, the decision maker considers situations in which several outcomes are possible for each course of action.
- In decision made under risk (a probabilistic or stochastic decision-making situation), is one in which the decision maker must consider several possible outcomes for each alternative, each with a given probability of occurrence.
- Spreadsheets include extensive statistical, forecasting, and other modeling and database management capabilities, functions, and routines.
- Simulation is the appearance of reality. In MSS is a technique for Simulation conducting experiments with a computer on a model of a management system.
- Mathematical programming is a family of tools designed to help solve managerial problems in which the decision maker must allocate scarce resources among competing activities to optimize a measurable goal.

5.8 Keywords

Dynamic Model: Dynamic models represent scenarios that change over time.

Linear Programming: Linear programming (LP) is the best-know technique in family of optimization tools in which all relationships among the LP variables are linear.

Mathematical Programming: Mathematical programming is a family of tools designed to help solve managerial problems in which the decision maker must allocate scarce resources among competing activities to optimize a measurable goal.

Modeling: Modeling is a key element in most DSS and a necessity in a model-based DSS.

Risk Analysis: Risk analysis is a decision-making method that analyzes the risk associated with different alternatives.

Simulation: Simulation is the appearance of reality. In MSS is a technique for Simulation conducting experiments with a computer on a model of a management system.

Static Model: A static model takes a single snapshot of a situation. During this snapshot, everything occurs in a single interval.

Visual Simulation: Visual Simulation is the graphical display of computerized results, which may include animation.

5.9 Review Questions

1. Describe the concept of modeling for MSS. Also discuss the various modeling issues.
2. Discuss the classification of Decision situation.
3. Make distinction between the Decision making under certainty and Decision making under uncertainty.
4. Explain the concept of MSS Modeling in Spreadsheets with example.
5. What is simulation? Discuss the characteristics of simulation.
6. Discuss the advantages and disadvantages of simulation.
7. Discuss the uses of Visual Interactive Simulation.
8. Illustrate the steps included in the simulation process.
9. Describe the process of Optimization via Mathematical Programming. Give example.
10. Analyse the difference between Discrete distributions and Continuous distributions.

Answers: Self Assessment

- | | |
|---------------------------------|-----------------------------|
| 1. Modeling | 2. Influence diagram |
| 3. Static | 4. Dynamic |
| 5. Certainty | 6. Uncertainty |
| 7. Risk analysis | 8. True |
| 9. False | 10. Simulation |
| 11. Discrete distributions | 12. Visual Simulation |
| 13. Visual Interactive Modeling | 14. Linear programming (LP) |
| 15. Mathematical programming | |

5.10 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.

Notes



Online links

<http://freepdfdb.com/ppt/spreadsheet-modeling-decision-analysis-pdf>

<http://fti.uajm.ac.id/ajar/Sistem%20Penunjang%20Keputusan/modelling%20and%20analysis.pdf>

http://www.mines.unr.edu/dtaylor/courses/361/Ragsdale_ch01.pdf

Unit 6: Data Warehousing

Notes

CONTENTS

Objectives

Introduction

6.1 Definitions and Concepts

6.1.1 Other Definitions

6.1.2 Concepts

6.2 Process Overview

6.3 Data Warehouse Architecture

6.3.1 Single-Layer Architecture

6.3.2 Two-Layer Architecture

6.3.3 Three-Layer Architecture

6.3.4 An Additional Architecture Classification

6.4 Data Warehouse Development

6.4.1 Business Requirements Definition

6.4.2 Data Acquisition

6.4.3 Architecture

6.4.4 Data Quality

6.4.5 Warehouse Administration

6.4.6 Metadata Management

6.4.7 Data Access

6.4.8 Database Design and Build

6.4.9 Documentation

6.4.10 Testing

6.4.11 Training

6.4.12 Transition

6.4.13 Post-Implementation Support

6.5 Administration and Security Issues

6.6 Summary

6.7 Keywords

6.8 Review Questions

6.9 Further Readings

Notes

Objectives

After studying this unit, you will be able to:

- Discuss the Definitions and Concepts of Data Warehousing
- Explain Process Overview
- Explain Architecture of Data Warehouse
- Discuss Data Warehousing Development
- Discuss Administration and Security Issues

Introduction

A fundamental concept of a data warehouse is the distinction between data and information. Data is composed of observable and recordable facts that are often found in operational or transactional systems. At Rutgers, these systems include the registrar's data on students (widely known as the SRDB), human resource and payroll databases, course scheduling data, and data on financial aid. In a data warehouse environment, data only comes to have value to end-users when it is organized and presented as information. Information is an integrated collection of facts and is used as the basis for decision-making. For example, an academic unit needs to have diachronic information about its extent of instructional output of its different faculty members to gauge if it is becoming more or less reliant on part-time faculty. In this unit, we will discuss the concept of data warehousing.

6.1 Definitions and Concepts

The data warehouse is that portion of an overall Architected Data Environment that serves as the single integrated source of data for processing information. The data warehouse has specific characteristics that include the following:

- **Subject-Oriented:** Information is presented according to specific subjects or areas of interest, not simply as computer files. Data is manipulated to provide information about a particular subject.



Example: The SRDB is not simply made accessible to end-users, but is provided structure and organized according to the specific needs.

- **Integrated:** A single source of information for and about understanding multiple areas of interest. The data warehouse provides one-stop shopping and contains information about a variety of subjects. Thus the OIRAP data warehouse has information on students, faculty and staff, instructional workload, and student outcomes.
- **Non-Volatile:** Stable information that doesn't change each time an operational process is executed. Information is consistent regardless of when the warehouse is accessed.
- **Time-Variant:** Containing a history of the subject, as well as current information. Historical information is an important component of a data warehouse.
- **Accessible:** The primary purpose of a data warehouse is to provide readily accessible information to end-users.
- **Process-Oriented:** It is important to view data warehousing as a process for delivery of information.

The maintenance of a data warehouse is ongoing and iterative in nature.

6.1.1 Other Definitions

Other definitions of data warehouse are discussed below.

Data Warehouse: A data structure that is optimized for distribution. It collects and stores integrated sets of historical data from multiple operational systems and feeds them to one or more data marts. It may also provide end-user access to support enterprise views of data.

Data Mart: A data structure that is optimized for access. It is designed to facilitate end-user analysis of data. It typically supports a single, analytic application used by a distinct set of workers.

Staging Area: Any data store that is designed primarily to receive data into a warehousing environment.

Operational Data Store: A collection of data that addresses operational needs of various operational units. It is not a component of a data warehousing architecture, but a solution to operational needs.

OLAP (On-Line Analytical Processing): A method by which multidimensional analysis occurs.

Multidimensional Analysis: The ability to manipulate information by a variety of relevant categories or “dimensions” to facilitate analysis and understanding of the underlying data. It is also sometimes referred to as “drilling-down”, “drilling-across” and “slicing and dicing”.

Hypercube: A means of visually representing multidimensional data.

Star Schema: A means of aggregating data based on a set of known dimensions. It stores data multidimensionally in a two dimensional Relational Database Management System (RDBMS), such as Oracle.

Snowflake Schema: An extension of the star schema by means of applying additional dimensions to the dimensions of a star schema in a relational environment.

Multidimensional Database: Also known as MDDB or MDDBS. A class of proprietary, non-relational database management tools that store and manage data in a multidimensional manner, as opposed to the two dimensions associated with traditional relational database management systems.

OLAP Tools: A set of software products that attempt to facilitate multidimensional analysis. Can incorporate data acquisition, data access, data manipulation, or any combination thereof.

6.1.2 Concepts

The definition of data warehousing presented here is intentionally generic; it gives you an idea of the process but does not include specific features of the process. To understand the role and the useful properties of data warehousing completely, you must first understand the needs that brought it into being. In 1996, R. Kimball efficiently summed up a few claims frequently submitted by end users of classic information systems:

- “We have heaps of data, but we cannot access it!” This shows the frustration of those who are responsible for the future of their enterprises but have no technical tools to help them extract the required information in a proper format.
- “How can people playing the same role achieve substantially different results?” In midsize to large enterprises, many databases are usually available, each devoted to a specific business area. They are often stored on different logical and physical media that are not conceptually integrated. For this reason, the results achieved in every business area are likely to be inconsistent.

Notes

- “We want to select, group, and manipulate data in every possible way!” Decision-making processes cannot always be planned before the decisions are made. End users need a tool that is user-friendly and flexible enough to conduct ad hoc analyses. They want to choose which new correlations they need to search for in real time as they analyze the information retrieved.
- “Show me just what matters!” Examining data at the maximum level of detail is not only useless for decision-making processes, but is also self-defeating, because it does not allow users to focus their attention on meaningful information.
- “Everyone knows that some data is wrong!” This is another sore point. An appreciable percentage of transactional data is not correct—or it is unavailable. It is clear that you cannot achieve good results if you base your analyses on incorrect or incomplete data.

We can use the previous list of problems and difficulties to extract a list of key words that become distinguishing marks and essential requirements for a data warehouse process, a set of tasks that allow us to turn operational data into decision-making support information:

- accessibility to users not very familiar with IT and data structures;
- integration of data on the basis of a standard enterprise model;
- query flexibility to maximize the advantages obtained from the existing information;
- information conciseness allowing for target-oriented and effective analyses;
- multidimensional representation giving users an intuitive and manageable view of information;
- correctness and completeness of integrated data.



Did u know? Data warehouses are placed right in the middle of this process and act as repositories for data. They make sure that the requirements set can be fulfilled.

Data warehouses are subject-oriented because they hinge on enterprise-specific concepts, such as customers, products, sales, and orders. On the contrary, operational databases hinge on many different enterprise-specific applications.

We put emphasis on integration and consistency because data warehouses take advantage of multiple data sources, such as data extracted from production and then stored to enterprise databases, or even data from a third party’s information systems. A data warehouse should provide a unified view of all the data. Generally speaking, we can state that creating a data warehouse system does not require that new information be added; rather, existing information needs rearranging. This implicitly means that an information system should be previously available.

Operational data usually covers a short period of time, because most transactions involve the latest data. A data warehouse should enable analyses that instead cover a few years. For this reason, data warehouses are regularly updated from operational data and keep on growing. If data were visually represented, it might progress like so: A photograph of operational data would be made at regular intervals. The sequence of photographs would be stored to a data warehouse, and results would be shown in a movie that reveals the status of an enterprise from its foundation until present.

Fundamentally, data is never deleted from data warehouses and updates are normally carried out when data warehouses are off-line. This means that data warehouses can be essentially

viewed as read-only databases. This satisfies the users' need for a short analysis query response time and has other important effects. First, it affects data warehouse-specific database management system (DBMS) technologies, because there is no need for advanced transaction management techniques required by operational applications. Second, data warehouses operate in read-only mode, so data warehouse-specific logical design solutions are completely different from those used for operational databases. For instance, the most obvious feature of data warehouse relational implementations is that table normalization can be given up to partially denormalize tables and improve performance.

Other differences between operational databases and data warehouses are connected with query types. Operational queries execute transactions that generally read/write a small number of tuples from/to many tables connected by simple relations.



Example: This applies if you search for the data of a customer in order to insert a new customer order. This kind of query is an OLTP query. On the contrary, the type of query required in data warehouses is OLAP. It features dynamic, multidimensional analyses that need to scan a huge amount of records to process a set of numeric data summing up the performance of an enterprise.



Notes OLTP systems have an essential workload core "frozen" in application programs, and ad hoc data queries are occasionally run for data maintenance. Conversely, data warehouse interactivity is an essential property for analysis sessions, so the actual workload constantly changes as time goes by.

The distinctive features of OLAP queries suggest adoption of a multidimensional representation for data warehouse data. Basically, data is viewed as points in space, whose dimensions correspond to many possible analysis dimensions. Each point represents an event that occurs in an enterprise and is described by a set of measures relevant to decision making processes.

Table 6.1 summarizes the main differences between operational databases and data warehouses.

Table 6.1: Differences between Operational Databases and Data Warehouses

| System | Description |
|--------------------------|--|
| Passive DSS | Supports decision-making processes, but it does not offer explicit suggestions on decisions or solutions. |
| Active DSS | Offers suggestions and solutions. |
| Collaborative DSS | Operates interactively and allows decision-makers to modify, integrate, or refine suggestions given by the system. Suggestions are sent back to the system for validation. |
| Model-driven DSS | Enhances management of statistical, financial, optimization, and simulation models. |
| Communication-driven DSS | Supports a group of people working on a common task. |
| Data-driven DSS | Enhances the access and management of time series of corporate and external data. |
| Document-driven DSS | Manages and processes nonstructured data in many formats. |
| Knowledge-driven DSS | Provides problem-solving features in the form of facts, rules, and procedures. |

Notes

Benefits of Data Warehouses

The core benefits include:

- Historical information for comparative and competitive analysis
- Enhanced data quality and completeness
- Supplementing disaster recovery plans with another data back up source.

Among the greatest benefits of a data warehouse is the ability to analyze and execute business decisions based on data from multiple sources.



Example: An organization has collected valuable data and stored it in 30 databases.

A data warehouse is not only a convenient way to analyze and compare data in all the databases, but it can also give historical data and perspective. Thus data warehouse is a one-stop shop, but it is also a one-stop shop from an historical perspective as well. Using data warehouse, one can look at past trends, whether they be product sales or customers or whatever and may be do some predictions of what is going to happen in the future.

Also data retrieved from multiple databases is not constrained by the tables in each of those databases. A data warehouse receives application neutral data. Whatever database application is supplying the information to the data warehouse is not preconditioning the data to be presented in a way the originator of the data requires. That means, the data from the inventory system, the financial system, or the sales system is sent to the data warehouse for processing as application neutral data that is not formatted to answer only queries from an inventory database, finance database, or sales database program. If not for application-neutral data, the data warehouse would be nothing more than a collection of data marts.

A data warehouse by itself does not create value, but value comes from the use of the data in the warehouse. In support of a low cost strategy, the data warehouse can provide savings in billing processes, reduce fraud losses, and reduce the cost of reporting. The data warehouses can provide analysts with precalculated reports and graphs. This increases the productivity of business analysts.

Most companies can benefit from a data warehouse when the proper tools are in place and users are trained in analysis of results.

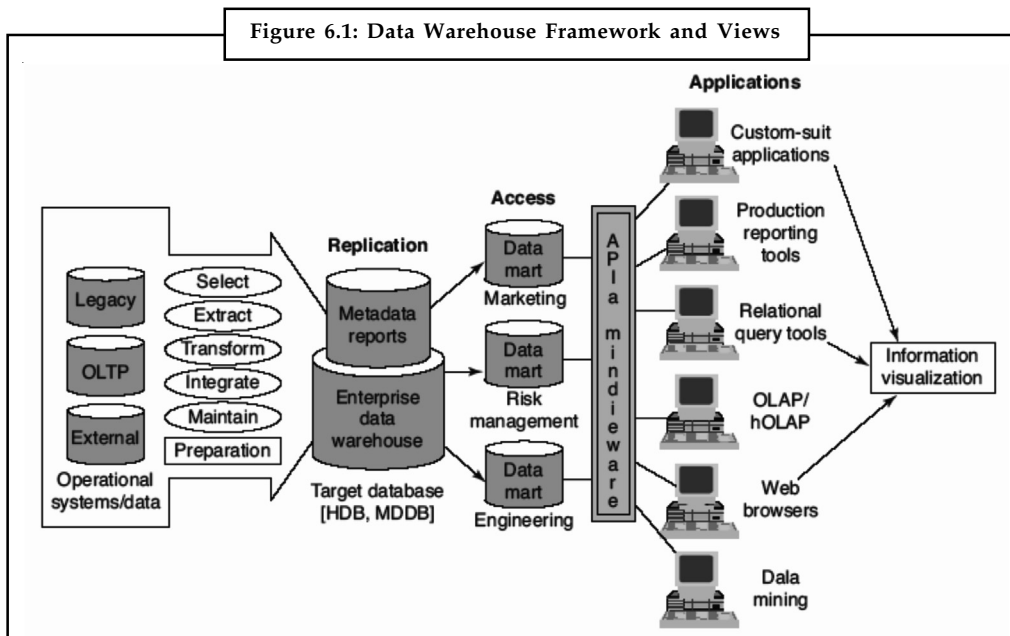
Self Assessment

Fill in the blanks:

1. The is that portion of an overall Architected Data Environment that serves as the single integrated source of data for processing information.
2. Any data store that is designed primarily to receive data into a warehousing environment is called
3. is a means of aggregating data based on a set of known dimensions.

6.2 Process Overview

Organizations continuously collect data, information, and knowledge at an increasingly accelerated rate and store them in computerized systems. The number of users needing to access the information continues to increase as a result of improved reliability and availability of network access, especially the Internet.



The major components of a data warehousing process are:

- Data sources (legacy systems, ERP, OLTP, web logs)
- Data extraction
- Data loading
- Comprehensive database (EDW)
- Metadata
- Middleware tools (enable access to the DW)

ETL (Extract, Transform and Load) is a process in data warehousing responsible for pulling data out of the source systems and placing it into a data warehouse.

ETL Technology is an important component of the Data Warehousing Architecture. It is used to copy data from Operational Applications to the Data Warehouse Staging Area, from the DW Staging Area into the Data Warehouse and finally from the Data Warehouse into a set of conformed Data Marts that are accessible by decision makers.

The ETL software extracts data, transforms values of inconsistent data, cleanses “bad” data, filters data and loads data into a target database. The scheduling of ETL jobs is critical. Should there be a failure in one ETL job, the remaining ETL jobs must respond appropriately.

ETL involves the following tasks:

- extracting the data from source systems (SAP, ERP, other operational systems), data from different source systems is converted into one consolidated data warehouse format which is ready for transformation processing.
- transforming the data may involve the following tasks:
 - ❖ applying business rules



Example: So-called derivations, such as calculating new measures and dimensions.

Notes

- ❖ Cleaning



Example: Mapping NULL to 0 or “Male” to “M” and “Female” to “F” etc.

- ❖ filtering



Example: Selecting only certain columns to load.

- ❖ splitting a column into multiple columns and vice versa,
- ❖ joining together data from multiple sources (e.g., lookup, merge),
- ❖ transposing rows and columns,
- ❖ applying any kind of simple or complex data validation



Example: If the first three columns in a row are empty then reject the row from processing.

- loading the data into a data warehouse or data repository other reporting applications

Self Assessment

Fill in the blanks:

4. is a process in data warehousing responsible for pulling data out of the source systems and placing it into a data warehouse.
5. ETL process loads the data into a data warehouse or

6.3 Data Warehouse Architecture

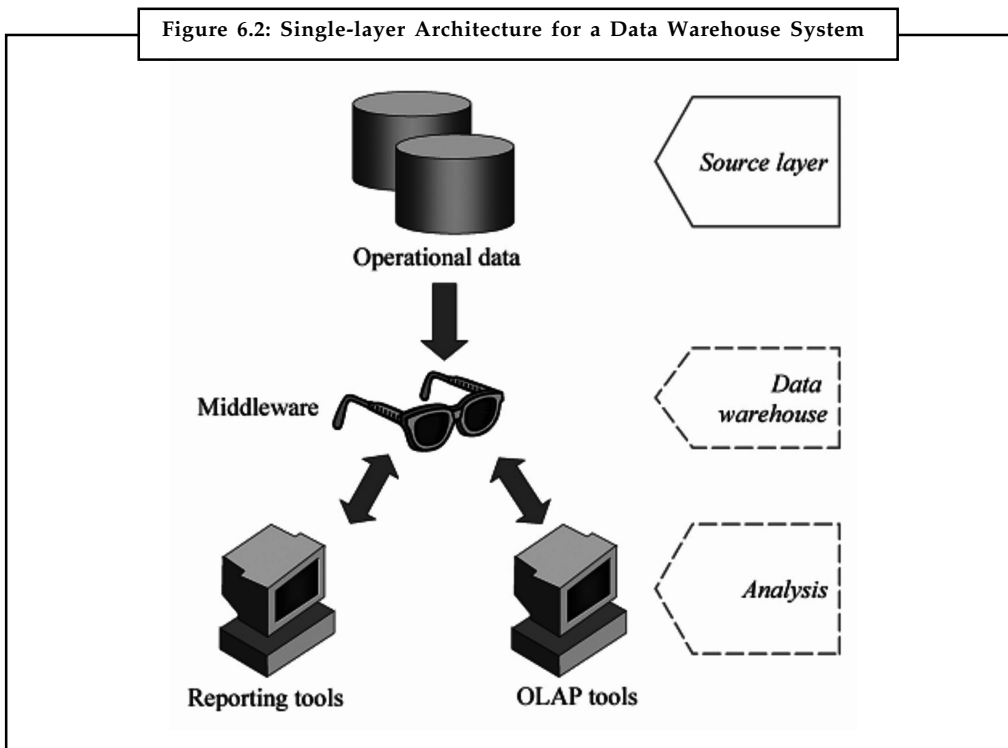
The following architecture properties are essential for a data warehouse system:

- **Separation:** Analytical and transactional processing should be kept apart as much as possible.
- **Scalability:** Hardware and software architectures should be easy to upgrade as the data volume, which has to be managed and processed, and the number of users’ requirements, which have to be met, progressively increase.
- **Extensibility:** The architecture should be able to host new applications and technologies without redesigning the whole system.
- **Security:** Monitoring accesses is essential because of the strategic data stored in data warehouses.
- **Administerability:** Data warehouse management should not be overly difficult.

Two different classifications are commonly adopted for data warehouse architectures. The first classification is a structure-oriented one that depends on the number of layers used by the architecture. The second classification depends on how the different layers are employed to create enterprise-oriented or department-oriented views of data warehouses.

6.3.1 Single-Layer Architecture

A single-layer architecture is not frequently used in practice. Its goal is to minimize the amount of data stored; to reach this goal, it removes data redundancies. Figure 6.2 shows the only layer physically available: the source layer. In this case, data warehouses are virtual.



Source: http://www.mhprofessional.com/downloads/products/0071610391/0071610391_chap01.pdf

This means that a data warehouse is implemented as a multidimensional view of operational data created by specific middleware, or an intermediate processing layer. The weakness of this architecture lies in its failure to meet the requirement for separation between analytical and transactional processing. Analysis queries are submitted to operational data after the middleware interprets them. In this way, the queries affect regular transactional workloads. In addition, although this architecture can meet the requirement for integration and correctness of data, it cannot log more data than sources do. For these reasons, a virtual approach to data warehouses can be successful only if analysis needs are particularly restricted and the data volume to analyze is huge.

6.3.2 Two-Layer Architecture

The requirement for separation plays a fundamental role in defining the typical architecture for a data warehouse system, as shown in Figure 6.3. Although it is typically called a two-layer architecture to highlight a separation between physically available sources and data warehouses, it actually consists of four subsequent data flow stages:

1. **Source layer:** A data warehouse system uses heterogeneous sources of data. That data is originally stored to corporate relational databases or legacy databases, or it may come from information systems outside the corporate walls.
2. **Data staging:** The data stored to sources should be extracted, cleansed to remove inconsistencies and fill gaps, and integrated to merge heterogeneous sources into one common schema. The so-called Extraction, Transformation, and Loading tools (ETL) can merge heterogeneous schemata, extract, transform, cleanse, validate, filter, and load source data into a data warehouse. Technologically speaking, this stage deals with problems that are typical for distributed information systems, such as inconsistent data management and incompatible data structures.

Notes

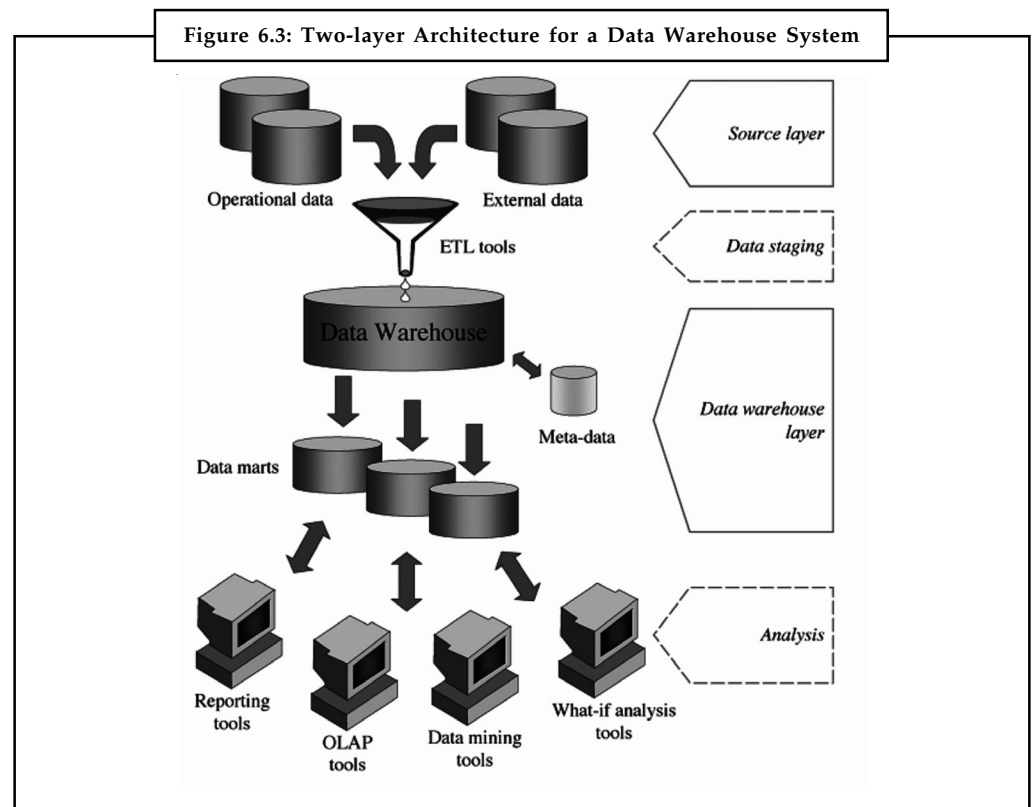
3. **Data warehouse layer:** Information is stored to one logically centralized single repository: a data warehouse. The data warehouse can be directly accessed, but it can also be used as a source for creating data marts, which partially replicate data warehouse contents and are designed for specific enterprise departments. Metadata repositories store information on sources, access procedures, data staging, users, data mart schemata, and so on.
4. **Analysis:** In this layer, integrated data is efficiently and flexibly accessed to issue reports, dynamically analyze information, and simulate hypothetical business scenarios. Technologically speaking, it should feature aggregate data navigators, complex query optimizers, and user-friendly GUIs.

The architectural difference between data warehouses and data marts needs to be studied closer. The component marked as a data warehouse in Figure 6.3 is also often called the primary data warehouse or corporate data warehouse.



Did u know? Primary data warehouse acts as a centralized storage system for all the data being summed up.

Data marts can be viewed as small, local data warehouses replicating (and summing up as much as possible) the part of a primary data warehouse required for a specific application domain.



Source: http://www.mhprofessional.com/downloads/products/0071610391/0071610391_chap01.pdf

The data marts populated from a primary data warehouse are often called dependent. Although data marts are not strictly necessary, they are very useful for data warehouse systems in midsize to large enterprises because:

- they are used as building blocks while incrementally developing data warehouses;

- they mark out the information required by a specific group of users to solve queries;
- they can deliver better performance because they are smaller than primary data warehouses.

Sometimes, mainly for organization and policy purposes, you should use a different architecture in which sources are used to directly populate data marts. These data marts are called independent. If there is no primary data warehouse, this streamlines the design process, but it leads to the risk of inconsistencies between data marts. To avoid these problems, you can create a primary data warehouse and still have independent data marts. In comparison with the standard two-layer architecture of Figure 6.3, the roles of data marts and data warehouses are actually inverted. In this case, the data warehouse is populated from its data marts, and it can be directly queried to make access patterns as easy as possible.

The following list sums up all the benefits of a two-layer architecture, in which a data warehouse separates sources from analysis applications:

- In data warehouse systems, good quality information is always available, even when access to sources is denied temporarily for technical or organizational reasons.
- Data warehouse analysis queries do not affect the management of transactions, the reliability of which is vital for enterprises to work properly at an operational level.
- Data warehouses are logically structured according to the multidimensional model, while operational sources are generally based on relational or semi-structured models.
- A mismatch in terms of time and granularity occurs between OLTP systems, which manage current data at a maximum level of detail, and OLAP systems, which manage historical and summarized data.
- Data warehouses can use specific design solutions aimed at performance optimization of analysis and report applications.



Task Analyze the use of primary data warehouse.

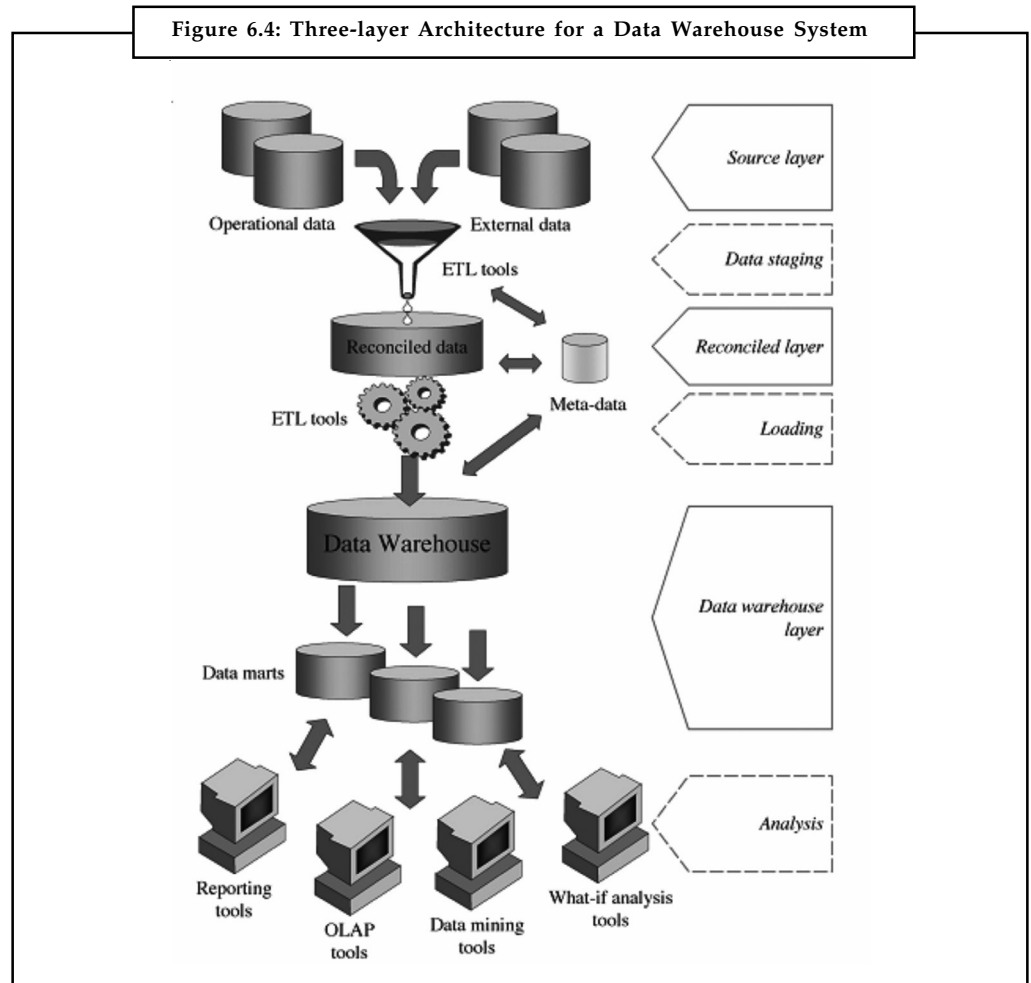
6.3.3 Three-Layer Architecture

In this architecture, the third layer is the reconciled data layer or operational data store. This layer materializes operational data obtained after integrating and cleansing source data. As a result, those data are integrated, consistent, correct, current, and detailed. Figure 6.4 shows a data warehouse that is not populated from its sources directly, but from reconciled data. The main advantage of the reconciled data layer is that it creates a common reference data model for a whole enterprise. At the same time, it sharply separates the problems of source data extraction and integration from those of data warehouse population. Remarkably, in some cases, the reconciled layer is also directly used to better accomplish some operational tasks, such as producing daily reports that cannot be satisfactorily prepared using the corporate applications, or generating data flows to feed external processes periodically so as to benefit from cleaning and integration. However, reconciled data leads to more redundancy of operational source data.



Notes Note that we may assume that even two-layer architectures can have a reconciled layer that is not specifically materialized, but only virtual, because it is defined as a consistent integrated view of operational source data.

Notes



Source: http://www.mhprofessional.com/downloads/products/0071610391/0071610391_chap01.pdf

Finally, let's consider a supplementary architectural approach, which provides a comprehensive picture. This approach can be described as a hybrid solution between the single-layer architecture and the two/three-layer architecture. This approach assumes that although a data warehouse is available, it is unable to solve all the queries formulated. This means that users may be interested in directly accessing source data from aggregate data (drill-through). To reach this goal, some queries have to be rewritten on the basis of source data (or reconciled data if it is available). This type of architecture is implemented in a prototype by Cui and Widom, 2000, and it needs to be able to go dynamically back to the source data required for queries to be solved (lineage).

6.3.4 An Additional Architecture Classification

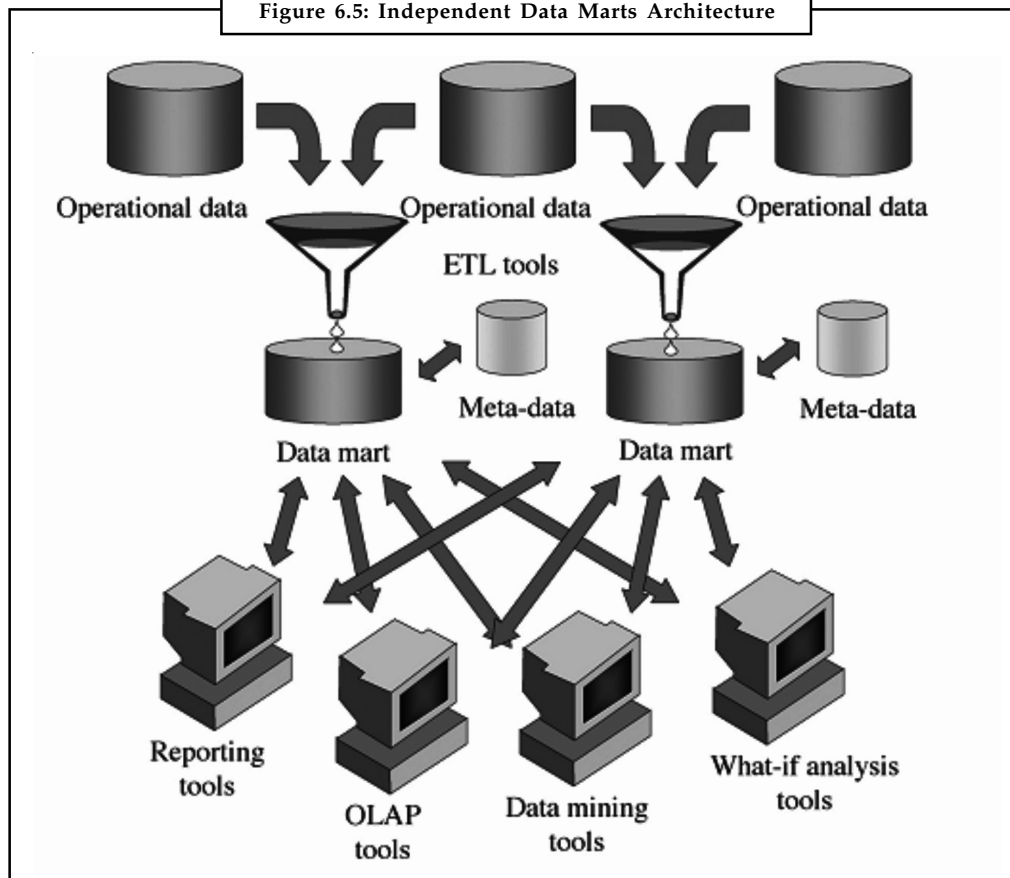
The scientific literature often distinguishes five types of architecture for data warehouse systems, in which the same basic layers mentioned in the preceding paragraphs are combined in different ways.

In independent data marts architecture, different data marts are separately designed and built in a non-integrated fashion (Figure 6.5). This architecture can be initially adopted in the absence of a strong sponsorship toward an enterprise-wide warehousing project, or when the organizational divisions that make up the company are loosely coupled. However, it tends to be soon replaced by other architectures that better achieve data integration and cross-reporting.

The bus architecture, recommended by Ralph Kimball, is apparently similar to the preceding architecture, with one important difference. A basic set of conformed dimensions (that is, analysis dimensions that preserve the same meaning throughout all the facts they belong to), derived by a careful analysis of the main enterprise processes, is adopted and shared as a common design guideline. This ensures logical integration of data marts and an enterprise-wide view of information.

In the hub and spoke architecture, one of the most used in medium to large contexts, there is much attention to scalability and extensibility, and to achieving an enterprise-wide view of information. Atomic, normalized data is stored in a reconciled layer that feeds a set of data marts containing summarized data in multidimensional form (Figure 6.6). Users mainly access the data marts, but they may occasionally query the reconciled layer.

Figure 6.5: Independent Data Marts Architecture

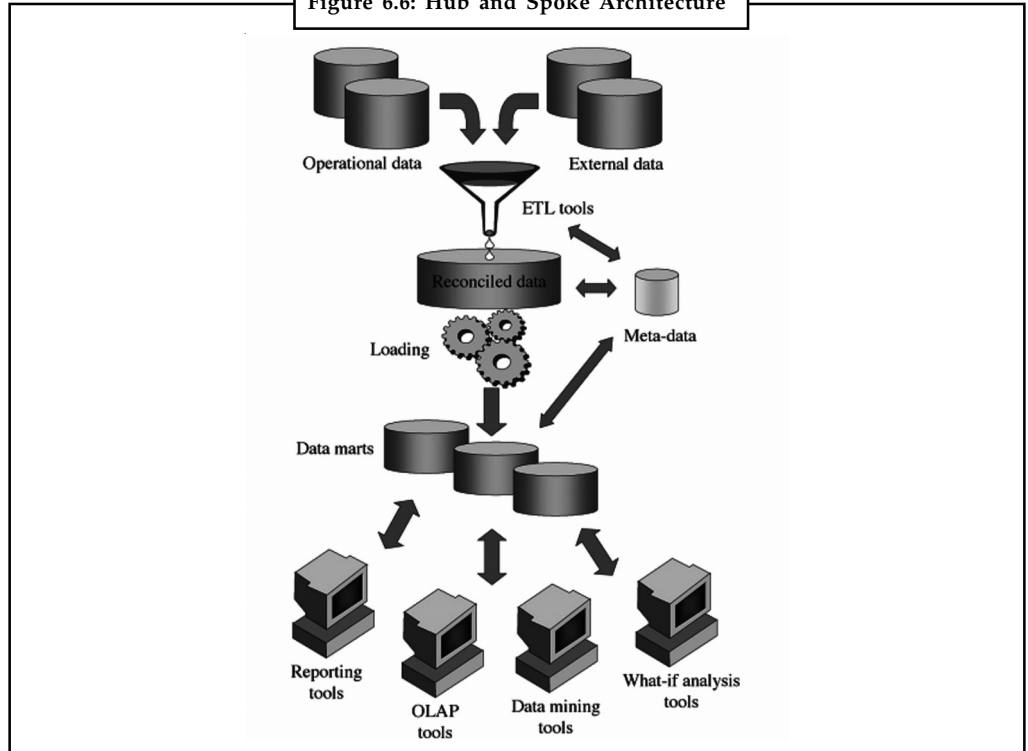


Source: http://www.mhprofessional.com/downloads/products/0071610391/0071610391_chap01.pdf

The centralized architecture, recommended by Bill Inmon, can be seen as a particular implementation of the hub and spoke architecture, where the reconciled layer and the data marts are collapsed into a single physical repository. The federated architecture is sometimes adopted in dynamic contexts where preexisting data warehouses/data marts are to be non-invasively integrated to provide a single, cross organization decision support environment (for instance, in the case of mergers and acquisitions). Each data warehouse/data mart is either virtually or physically integrated with the others, leaning on a variety of advanced techniques such as distributed querying, ontologies, and meta-data interoperability (Figure 6.7).

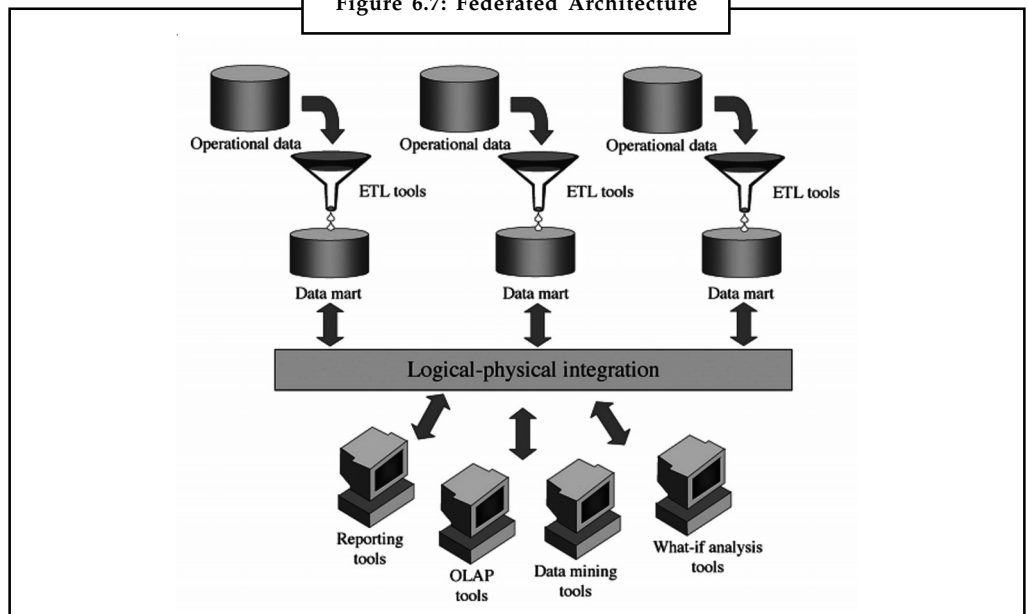
Notes

Figure 6.6: Hub and Spoke Architecture



Source: http://www.mhprofessional.com/downloads/products/0071610391/0071610391_chap01.pdf

Figure 6.7: Federated Architecture



Source: http://www.mhprofessional.com/downloads/products/0071610391/0071610391_chap01.pdf

The following list includes the factors that are particularly influential when it comes to choosing one of these architectures:

- The amount of interdependent information exchanged between organizational units in an enterprise and the organizational role played by the data warehouse project sponsor may

lead to the implementation of enterprise-wide architectures, such as bus architectures, or department-specific architectures, such as independent data marts.

- An urgent need for a data warehouse project, restrictions on economic and human resources, as well as poor IT staff skills may suggest that a type of “quick” architecture, such as independent data marts, should be implemented.
- The minor role played by a data warehouse project in enterprise strategies can make you prefer an architecture type based on independent data marts over a hub and spoke architecture type.
- The frequent need for integrating preexisting data warehouses, possibly deployed on heterogeneous platforms, and the pressing demand for uniformly accessing their data can require a federated architecture type.

Self Assessment

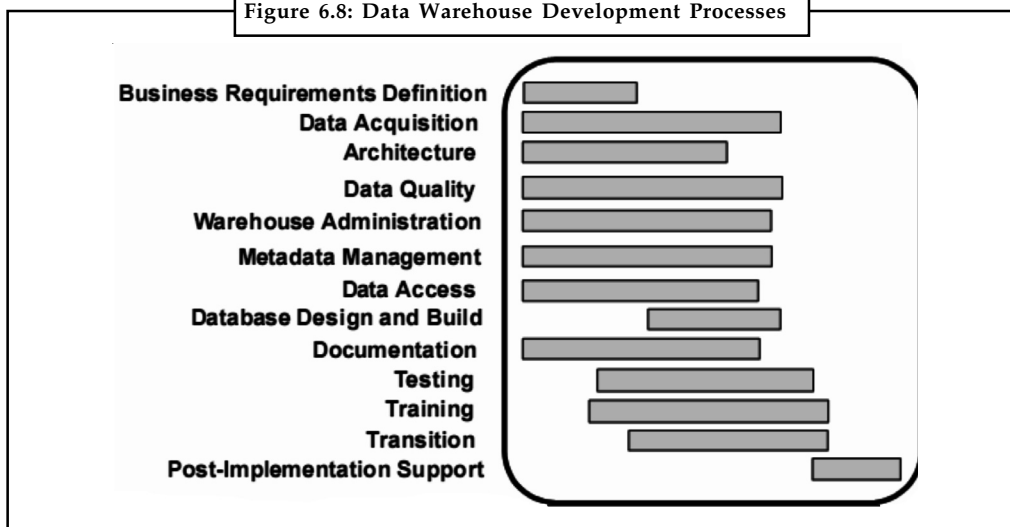
Fill in the blanks:

6. The goal of architecture is to minimize the amount of data stored; to reach this goal, it removes data redundancies.
7. data warehouse acts as a centralized storage system for all the data being summed up.
8. The populated from a primary data warehouse are often called dependent.
9. In data marts architecture, different data marts are separately designed and built in a nonintegrated fashion.

6.4 Data Warehouse Development

The sections below provide perspective on each of the critical processes in a data warehouse development effort. In addition to the details discussed in each section, numerous control mechanisms and management techniques to facilitate the success of the overall project support each process.

Figure 6.8: Data Warehouse Development Processes



Source: http://www.grtcorp.com/sites/grtcorp.com/files/DW_Development_Approach.pdf

6.4.1 Business Requirements Definition

The Business Requirements Definition process defines the requirements, clarifies the scope, and establishes the implementation roadmap for the data warehouse. With the direction of the client organization, strategic business goals and initiatives are established and used to direct the strategies, purpose and goals for each phase of the data warehouse solution. Early in the process, the focus is on the enterprise aspects of the data warehouse such as enterprise information requirements, subject areas, an implementation roadmap, and a business case for the data warehouse. As the process continues, Business Requirements Definition focuses on determining the specifics of the solution to be developed and delivered, identifying the client's information needs, and modeling the requirements.

6.4.2 Data Acquisition

The objective of the Data Acquisition process is to identify, extract, transform, and transport the various source data necessary for the operation of the data warehouse. Data Acquisition is performed between several components of the warehouse, including operational and external data sources to data warehouse, data warehouse to data mart, and data mart to individual marts. Early in the Data Acquisition process, data sources are identified and evaluated against the subject areas, and a gap analysis is conducted to verify that data is available to support the information requirements.

The Data Acquisition Strategy is also developed to outline the approach for extraction, transformation, and transportation of the source data to the data warehouse. The strategy includes selecting a tool or set of tools as the data pump or defining the specifications of one that must be built. If tools are to be utilized, high-level tool requirements, tool evaluations, and tool recommendations are also addressed.

6.4.3 Architecture

The Data Warehouse Architecture provides an integrated data warehouse environment while delivering incremental solutions. The architectural design focuses on the application of a centralized data warehouse, data marts, individual marts, metadata repositories, and incremental solution architectures. As the process continues, the development and execution of the integration plans are completed and the compliance of incremental solutions with the strategic architecture is validated.

6.4.4 Data Quality

The objective of the Data Quality process is to provide consistent, reliable, and accurate data in the warehouse. The Data Quality Strategy is developed based upon a clear understanding of which data cleansing and integrity functions meet the needs of the customer. To facilitate timely and reliable resolution of data issues, Data Management Procedures are also defined. Early in the process, data quality tools are also evaluated and recommended. The Data Quality process identifies the business rules for error exception handling, data cleansing, and audit and control. In addition, any variations in business rules for error handling between initial loads and subsequent updates to the data warehouse are defined. Utilizing the strategy, procedures, and tools, modules are developed to support the requirements for data quality. In order to populate the data warehouse with reliable data, the Data Quality modules are integrated with the Data Acquisition modules to check that the quality functions are properly sequenced within the overall transformation of source data to the target environment.

6.4.5 Warehouse Administration

The Warehouse Administration process specifies the strategy and requirements for the maintenance, use and ongoing updates to the data warehouse. Early in the process, the Warehouse Administration Strategy is established specifically outlining areas including version control, scheduling, data warehouse usage, security, audit and data governing.

The warehouse administration workflow, tool requirements, evaluation, and testing are also addressed. As the process continues, modules are designed and built for version control, scheduling, backup and recovery, archiving, security, audit, and data governing. In addition, several administration and monitoring tasks are addressed during the process. These include authorization to access appropriate levels of data, monitoring usage, governing queries, identifying repetitive queries, calculating metrics, defining access thresholds, adding or removing users, and updating access authority. To provide successful ongoing support and maintenance of the warehouse, this process focuses on the automation of the administration tasks, wherever possible.

6.4.6 Metadata Management

The Metadata Management process determines the Metadata Strategy, and defines requirements for metadata types, the metadata repository, metadata integration, and metadata access. The process addresses the integration of metadata for both the incremental and enterprise data warehouse solutions. A primary objective for this process is to provide technical and business views of the various aspects of warehouse metadata.

6.4.7 Data Access

The Data Access process focuses on the identification, selection, and design of tools that support end-user access to the warehouse data. Early in the process, a strategy is established and requirements are defined as a framework for the data access environment. Tools are evaluated, tested and recommended.

As the process continues, the user profiles are defined based on the level of data required to support their analysis, decision-making requirements, and skill level. Detailed requirements are also collected for the user interface style, queries and reports. With the user profiles in place, functional requirements, the levels of data to be accessed, and tool criteria are established for each data access component. In most cases, data access is supported by a variety of tools, rather than one tool to support every type of user.

6.4.8 Database Design and Build

The goal of the Database Design and Build process is to determine how the database objects are designed to support the data requirements and provide efficient access to the data. In addition, the logical and physical database designs are created and validated. Database designs are created for the relational and multidimensional database objects. During this process, physical data partitioning, segmentation, and data placement is evaluated against business and user requirements versus operational constraints. In addition, indexes and key definitions are determined. This process also generates the database Data Definition Language (DDL), and builds and implements the development, testing, and production data warehouse database objects.

6.4.9 Documentation

The Documentation process centers on producing high quality textual deliverables. All user and technical documentation for the data warehouse is developed, including references, user and system operations guides, and online help.

To facilitate active and successful use of the warehouse, the Metadata Reference describes the contents of the data warehouse in business terms and provides a navigational roadmap to the contents of the data warehouse. In addition, the Warehouse Administration Reference outlines the workflow and the manual and automated administration procedures. The New Features Guide highlights any new enhancements to warehouse functionality that result from the implementation of the solution.

6.4.10 Testing

The Testing process is an integrated approach to testing the quality of the various components of the data warehouse. Initially, the Testing Strategy is developed and approved by the client, followed by the creation of system, system integration, and module test plans, scripts, and scenarios. Each test is performed, including the volume tests on the physical designs for the database objects and regression testing of the enhanced warehouse against the current warehouse. Data Acquisition Modules, Data Access functions, canned queries and reports also undergo thorough module and modules integration testing.



Caution The testing strategy must address the various components of the solution, including the ad hoc access processes.

Regression testing is performed, allowing changes to the data warehouse to be tested against a baseline, ensuring past functionality works when an enhancement is added. Volume testing is conducted on the production platform to check that performance meets the established objectives. Preparation of the acceptance environment and support for acceptance testing is also performed during the Testing process.

6.4.11 Training

The Training process defines the development and end-user training requirements, identifies the technical and business personnel requiring training, and establishes timeframes for executing the training plans. During this process, the training plan and training materials are designed and developed, and the user and technical training is conducted.

The objective is to provide both users and administrators with adequate training to take on the tasks of operating, maintaining and using the data warehouse. Training focuses on tool training as well as the ways by which business value is generated from the information in the data warehouse. The team also trains the client's maintenance personnel and the acceptance test team.

6.4.12 Transition

The Transition process focuses on the tasks needed to perform the cutover to the production data warehouse. It includes tasks to create the installation plan and prepare the client maintenance and production environments. During this process, the warehouse administration workflow is implemented, and the production data warehouse is available.

6.4.13 Post-Implementation Support

Notes

The Post-Implementation Support process addresses the ongoing administration of the warehouse and provides an opportunity to evaluate and review the implemented solution.

During this process, use of the data warehouse is evaluated by accessing metadata and evaluating queries and reports run against the warehouse. This information assists with the management of standard queries or reports, the end-user layer, and the identification of potential indexes.

The process also focuses on several key warehouse administration functions including refreshing the warehouse, monitoring and responding to system problems, correcting errors, and conducting performance and tuning activities for the various components of the data warehouse. This includes change control for information requirements, roll-out of metadata, queries, reports, filters, and conditions, the library of shared objects, security, incorporation of new users, and the distribution of data marts and catalogs. During this process, responsibility for the data warehouse may be transferred to the client organization.



Task Compare and contrast testing and training.

Self Assessment

Fill in the blanks:

10. The objective of the process is to identify, extract, transform, and transport the various source data necessary for the operation of the data warehouse.
11. The process centers on producing high quality textual deliverables.
12. The process is an integrated approach to testing the quality of the various components of the data warehouse.
13. The process focuses on the tasks needed to perform the cutover to the production data warehouse.

6.5 Administration and Security Issues

Data warehouse administrator (DWA) is a person responsible for the administration and management of a data warehouse.



Caution Information security is generally a fundamental requirement for a system, and it should be carefully considered in software engineering at every project development stage from requirement analysis through implementation to maintenance.

Security is particularly relevant to data warehouse projects, because data warehouses are used to manage information crucial for strategic decision-making processes. Furthermore, multidimensional properties and aggregation cause additional security problems similar to those that generally arise in statistic databases, because they implicitly offer the opportunity to infer information from data. Finally, the huge amount of information exchange that takes place in data warehouses in the data-staging phase causes specific problems related to network security.

Appropriate management and auditing control systems are important for data warehouses. Management control systems can be implemented in front-end tools or can exploit operating

Notes

system services. As far as auditing is concerned, the techniques provided by DBMS servers are not generally appropriate for this scope. For this reason, you must take advantage of the systems implemented by OLAP engines. From the viewpoint of users profile-based data access, basic requirements are related to hiding whole cubes, specific cube slices, and specific cube measures. Sometimes you also have to hide cube data beyond a given detail level.

Effective security in a data warehouse should focus on four main areas:

- Establishing effective corporate and security policies and procedures
- Implementing logical security procedures and techniques to restrict access
- Limiting physical access to the data center environment
- Establishing an effective internal control review process with an emphasis on security and privacy

Self Assessment

Fill in the blanks:

14. is a person responsible for the administration and management of a data warehouse.
15. Security is particularly relevant to data warehouse projects, because data warehouses are used to manage information crucial for strategic processes.



Case Study

Large Data Warehouse

Knoware was instrumental in developing one of the first major data warehouse in New Zealand, starting in 1996. It is still one of the largest data warehouses in New Zealand. It takes data from multiple source systems, supports over 100 analysts in five different government departments, and provides customised reporting to several thousand users through an Intranet.

It records a full unit record history of all events for all systems. With this history it can recreate reports as at any given point in time, and track all events, for all clients, through any period of time and across all data systems.

Verification processes regularly prove that the data in the warehouse is correct. Metadata is generated and maintained by the data warehouse system to support the work of the business analysts. It documents the rules on how to use the data, and guarantees the quality of the information. Security controls appropriate access to the data to ensure that analysts see only the data they need to do their work.

Front line staff around the country can access information without the need to learn complex software, data or systems. They can point and click, slice and dice, drill down, and explore. For the client, this has created a whole new way of managing and operating their business.

Portals provides secure access to case load and case management information for around 4,000 frontline staff in service centres around the country. Staff can generate standard reports when they like; or customize and filter their own information.

Contd...

Notes

Knowledge transfer means that the data warehouse is now fully operated by the client. And like any successful data warehouse project, it continues to grow. "It is very satisfying to know that this data warehouse continues to provide the sound infrastructure that's so important to high quality information delivery." Says Clare Somerville, Knoware Director.

Question

Discuss the concept of Verification processes in data warehousing.

Source: <http://www.knoware.co.nz/large-data-warehouse-case-study>

6.6 Summary

- Data Warehouse collects and stores integrated sets of historical data from multiple operational systems and feeds them to one or more data marts.
- Data warehouses are subject-oriented because they hinge on enterprise-specific concepts, such as customers, products, sales, and orders.
- Among the greatest benefits of a data warehouse is the ability to analyze and execute business decisions based on data from multiple sources.
- Organizations continuously collect data, information, and knowledge at an increasingly accelerated rate and store them in computerized systems.
- ETL (Extract, Transform and Load) is a process in data warehousing responsible for pulling data out of the source systems and placing it into a data warehouse.
- The Data Warehouse Architecture provides an integrated data warehouse environment while delivering incremental solutions.
- A single-layer architecture minimize the amount of data stored; to reach this goal, it removes data redundancies.
- Data Warehouse Administrator (DWA) is a person responsible for the administration and management of a data warehouse.

6.7 Keywords

Business Requirements Definition: The Business Requirements Definition process defines the requirements, clarifies the scope, and establishes the implementation roadmap for the data warehouse.

Data Acquisition: Data Acquisition process identifies, extracts, transforms, and transports the various source data necessary for the operation of the data warehouse.

Data Mart: It is a data structure that is optimized for access which is designed to facilitate end-user analysis of data.

Data Warehouse: The data warehouse is that portion of an overall Architected Data Environment that serves as the single integrated source of data for processing information.

DWA: Data Warehouse Administrator (DWA) is a person responsible for the administration and management of a data warehouse.

ETL: ETL (Extract, Transform and Load) is a process in data warehousing responsible for pulling data out of the source systems and placing it into a data warehouse.

Notes

Multidimensional Analysis: It is the ability to manipulate information by a variety of relevant categories or “dimensions” to facilitate analysis and understanding of the underlying data.

OLAP Tools: It is a set of software products that attempt to facilitate multidimensional analysis.

6.8 Review Questions

1. Explain the concept of data warehousing.
2. Discuss the characteristics of data warehouse.
3. Discuss the use of data mart and operational data store.
4. Analyze the differences between operational databases and data warehouses.
5. Elucidate the major components of a data warehousing process.
6. Describe the Architecture of Data Warehouse.
7. Discuss the steps included in Two-Layer Architecture of data warehouse.
8. Discuss the factors that are to be considered when choosing one of the architectures of data warehouse.
9. Illustrate the steps included in the development of data warehouse.
10. Discuss the security issues involved in data warehousing.

Answers: Self Assessment

- | | |
|---------------------|--|
| 1. Data warehouse | 2. Staging Area |
| 3. Star Schema | 4. ETL (Extract, Transform and Load) |
| 5. Data repository | 6. Single-layer |
| 7. Primary | 8. Data marts |
| 9. Independent | 10. Data Acquisition |
| 11. Documentation | 12. Testing |
| 13. Transition | 14. Data warehouse administrator (DWA) |
| 15. Decision-making | |

6.9 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.

Notes



Online links

[http://63.209.69.107/see.php?q=data%20warehousing%20 architecture&affid=err1&subid=1&p=2&r=0](http://63.209.69.107/see.php?q=data%20warehousing%20architecture&affid=err1&subid=1&p=2&r=0)

http://it.toolbox.com/wiki/index.php/Data_Warehouse_Administration

<http://www.1keydata.com/datawarehousing/datawarehouse.html>

http://www.sba-research.org/wp-content/uploads/publications/weippl_arh_securityDWH.pdf

Notes

Unit 7: Business Analytics and Data Visualization

CONTENTS

Objectives

Introduction

7.1 An Overview of Business Analytics and Data Visualization

7.1.1 Business Analytics

7.1.2 Data Visualization

7.2 Online Analytical Processing (OLAP)

7.2.1 Historical Background

7.2.2 How OLAP Works

7.2.3 Types of OLAP

7.2.4 Features of OLAP

7.2.5 OLAP Cubes Support Business Intelligence

7.3 Geographical Information Systems (GIS)

7.3.1 History of GIS

7.3.2 Components of GIS

7.3.3 Types of GIS

7.3.4 Multimedia and Geographical Information System (GIS)

7.4 GIS Subsystems

7.4.1 Data Input Subsystem

7.4.2 Data Storage, Editing and Retrieval Subsystem

7.4.3 Data Manipulation and Analysis Subsystem

7.4.4 Data Output and Display Subsystem

7.5 Applications of Multimedia GIS

7.5.1 Education

7.5.2 Map-making

7.5.3 Land Information

7.5.4 Infrastructure and Utilities

7.5.5 Environmental

7.5.6 Archaeology

7.5.7 Natural Hazards

7.5.8 Forestry

7.5.9 Military GIS

7.5.10 Oceanography

Contd....

| | |
|--------|-----------------------------|
| 7.5.11 | Water Resources |
| 7.5.12 | GIS in Agriculture and Soil |
| 7.6 | Summary |
| 7.7 | Keywords |
| 7.8 | Review Questions |
| 7.9 | Further Readings |

Objectives

After studying this unit, you will be able to:

- Discuss the Concept of Business Analytics
- Explain Data Visualization
- Discuss the Concept of Online Analytical Processing (OLAP)
- Explain the Concept of GIS

Introduction

Business analytics is a term that refers to the applications, practices, skills, and technologies that are necessary for a complete investigation of a company's past business performance. Data visualization is a general term used to describe any technology that lets corporate executives and other end users "see" data in order to help them better understand the information and put it in a business context. Data analysis is a process that requires the reviewing of data to determine trends and abnormalities for a business. This is a daunting task because large data sets are typically presented in a spread sheet format. Data visualization is the process of converting standard textual data into pictures that are easily understood by a general audience. These pictures typically include graphs, shapes, and abstract objects that are color coded based on the significances of the data. In this unit, we will discuss the concept of business analytics and data visualization.

7.1 An Overview of Business Analytics and Data Visualization

In this section, we will discuss an overview of Business Analytics and Data Visualization.

7.1.1 Business Analytics

Business Analytics (BA) is the practice of iterative, methodical exploration of an organization's data with emphasis on statistical analysis. Business analytics is used by companies committed to data-driven decision making.

BA is used to gain insights that inform business decisions and can be used to automate and optimize business processes. Data-driven companies treat their data as a corporate asset and leverage it for competitive advantage. Successful business analytics depends on data quality, skilled analysts who understand the technologies and the business and an organizational commitment to data-driven decision making.

Examples of BA uses include:

- Exploring data to find new patterns and relationships (data mining)

Notes

- Explaining why a certain result occurred (statistical analysis, quantitative analysis)
- Experimenting to test previous decisions (A/B testing, multivariate testing)
- Forecasting future results (predictive modeling, predictive analytics)

This discipline helps business owners and executives to obtain valuable insight about how their particular businesses are performing, on average. It also helps them to determine how to plan successfully for improvement.

Once the business goal of the analysis is determined, an analysis methodology is selected and data is acquired to support the analysis. Data acquisition often involves extraction from one or more business systems, cleansing, and integration into a single repository such as a data warehouse or data mart. The analysis is typically performed against a smaller sample set of data. Analytic tools range from spreadsheets with statistical functions to complex data mining and predictive modeling applications. As patterns and relationships in the data are uncovered, new questions are asked and the analytic process iterates until the business goal is met. Deployment of predictive models involves scoring data records (typically in a database) and using the scores to optimize real-time decisions within applications and business processes. BA also supports tactical decision making in response to unforeseen events, and in many cases the decision making is automated to support real-time responses.

The process of business analytics focuses specifically on understanding a company's overall business performance. Through these intricate activities, new ideas can develop that may help to better prepare a business for future, or continued, success against its competitors. People who work in the field of business analytics rely heavily on a wide variety of data to help them with their daily tasks. Analysts who work in these types of jobs make regular use of quantitative and statistical analyses, and they are also usually quite involved in predictive modeling, which is essentially the process of predicting the likelihood of a particular outcome.

Practices of business analytics may be used to address issues of human activities and decisions, or may be more focused on automated functions. People who work in this field are generally quite good at coming up with answers to a wide range of questions that executives and company owners are likely to ask, in an effort to determine the best course of action on a company-wide scale. Analytics solutions usually involve the study of data in large quantities, so most people in this profession are comfortable working with numbers on a grand scale.



Caution In order for analytics practices to be the most successful and beneficial for companies, there must generally be massive quantities of data available for analysis.

When only small amounts of data are available, analytics tends to be less worthwhile for businesses. Essentially, this means that in some cases, business analysts are quite likely to be dealing with data that spans at least a few years. These positions, as a result, are most suited to individuals who work well under pressure and who thrive on deadlines.

For the most part, people involved in business analytics are also responsible for anticipating the changing needs of customers, which requires an insightful mind and shrewd deductive skills. Professionals within this field are usually capable of solving difficult problems fairly quickly. They are also usually able to suggest unique, innovative ideas to company executives for consistent, regular, positive change regarding business practices.

7.1.2 Data Visualization

Data visualization is a pretty literal term that means, quite simply, the visual representation of quantitative data.

Data visualization is the presentation of data in a pictorial or graphical format. For centuries, people have depended on visual representations such as charts and maps to understand information more easily and quickly. As more and more data is collected and analyzed, decision makers at all levels welcome data visualization software that enables them to see analytical results presented visually, find relevance among the millions of variables, communicate concepts and hypotheses to others, and even predict the future.

Because of the way the human brain processes information, it is faster for people to grasp the meaning of many data points when they are displayed in charts and graphs rather than poring over piles of spreadsheets or reading pages and pages of reports.

Visualization tools go beyond the standard charts and graphs used in Excel spreadsheets, displaying data in more sophisticated ways such as dials and gauges, geographic maps, time-series charts, spark lines, heat maps, tree maps and detailed bar, pie and fever charts. Patterns, trends and correlations that might go undetected in text-based data can be exposed and recognized easier with data visualization software.

Visualized data is frequently displayed in business intelligence (BI) dashboards and performance scorecards that provide users with high-level views of corporate information, metrics and Key Performance Indicators (KPIs). The images may include interactive capabilities, enabling users to manipulate them or drill into the data for querying and analysis. Indicators designed to alert users when data has been updated or predefined conditions occur, can also be included.



Notes Most business intelligence software vendors embed data visualization tools into their products, either developing the visualization technology themselves or sourcing it from companies that specialize in visualization.

Interactive Visualization

Interactive data visualization goes a step further – moving beyond the display of static graphics and spreadsheets to using computers and mobile devices to drill down into charts and graphs for more details, and interactively (and immediately) changing what data you see and how it is processed.

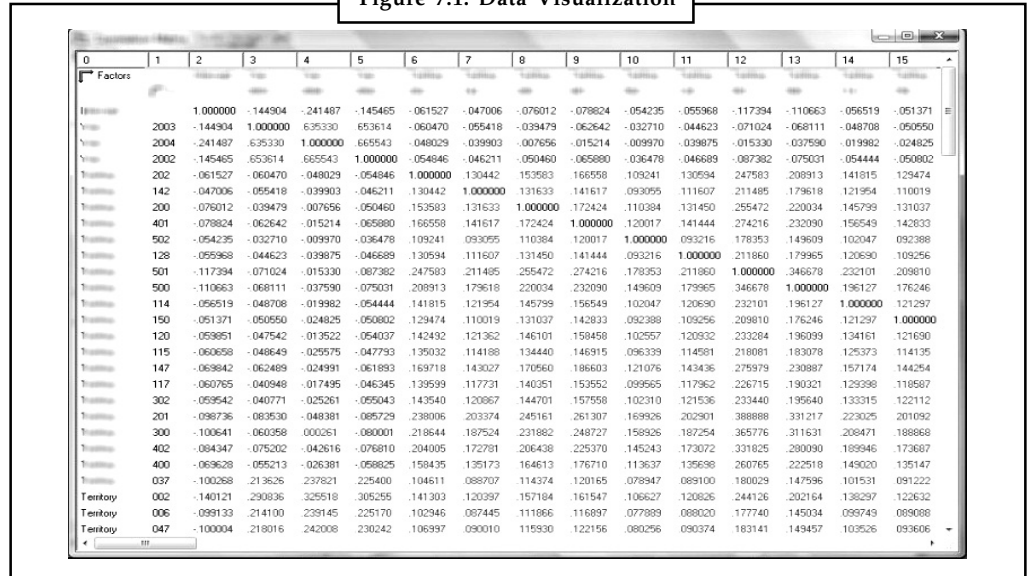
Importance of Data Visualization

Visualizations help people see things that were not obvious to them before. Even when data volumes are very large, patterns can be spotted quickly and easily. Visualizations convey information in a universal manner and make it simple to share ideas with others. It lets people ask others, “Do you see what I see?” And it can even answer questions like “What would happen if we made an adjustment to that area?”

Consider the manufacturing director of product reliability for an international company that produces small vibrating cell phone motors. One of the director’s principal responsibilities is to determine how reliable the cell phone motors will be with each year of age. If the product’s reliability falls short of the standards set forth by the cell phone manufacturers who use the motors, his company could lose major contracts.

Notes

Figure 7.1: Data Visualization



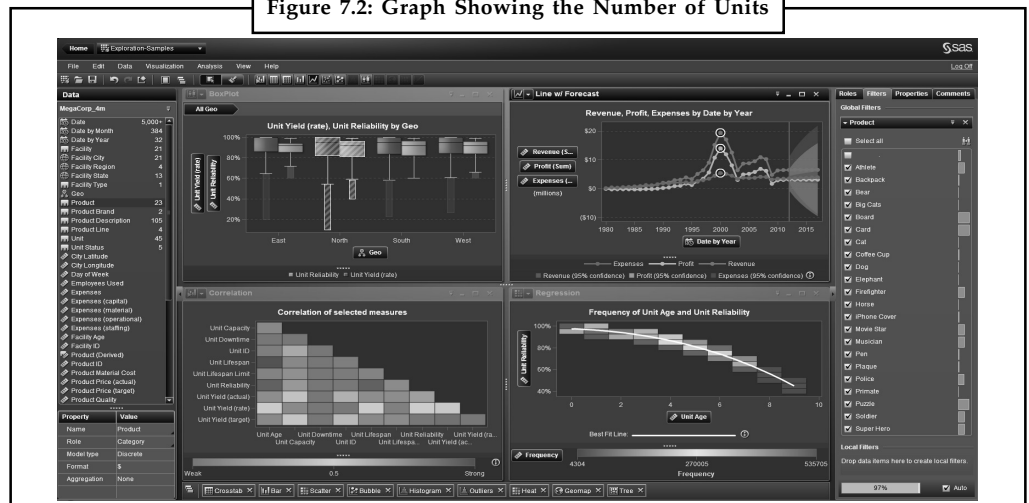
Source: <http://www.sas.com/data-visualization/overview.html>

Because of the amount of data collected on the age and reliability of the cell phone motors, a traditional electronic spreadsheet cannot visually represent the information due to data presentation limitations. And, if printed out, the spreadsheets would be a humongous pile of paper on the director's desk. In both cases, the director would spend hours searching among thousands of rows and columns of data with still no concrete answer to the original question about the relationship between the motor's age and its reliability.

Data visualization presents the data in a way that the director can easily interpret, saving time and energy.

For example, the graph below shows the number of units that correspond to each age (represented by the color gradient) as well as the reliability as the age of a unit increases. In a matter of seconds, the director can see that units approaching 10 years of age are approximately 40 percent reliable. This visual simplifies the data, instantly clarifying the factors affecting the reliability of the cell phone motors.

Figure 7.2: Graph Showing the Number of Units



Source: <http://www.sas.com/data-visualization/overview.html>

Interactive charts and graphs like the ones shown above make it easier for decision makers across all organizations to:

Notes

- Identify areas that need attention or improvement.
- Understand what factors influence your customers' behavior.
- Know which products to place where.
- Predict sales volumes.
- Discover how to increase revenues or reduce expenses.

Common Techniques

There are a few basic concepts that can help you generate the best visuals for displaying your data:

- Understand the data you are trying to visualize, including its size and cardinality (the uniqueness of data values in a column).
- Determine what you are trying to visualize and what kind of information you want to communicate.
- Know your audience and understand how it processes visual information.
- Use a visual that conveys the information in the best and simplest form for your audience.

Data visualization is an art and a science unto itself, and there are many graphical techniques that can be used to help people understand the story their data is telling.

Self Assessment

Fill in the blanks:

1. is the practice of iterative, methodical exploration of an organization's data with emphasis on statistical analysis.
2. is the presentation of data in a pictorial or graphical format.
3. go beyond the standard charts and graphs used in Excel spreadsheets, displaying data in more sophisticated ways.

7.2 Online Analytical Processing (OLAP)

Online analytical processing is defined as "The dynamic synthesis, analysis, and consolidation large volumes of multi-dimensional data."

Online analytical processing is a method of using multidimensional databases to support quick reporting, frequently involving trend-analysis. The primary query language for OLAP is called Multidimensional Expressions (MDX). Its name is derived from the program class known as online transactional processing (OLTP).



Did u know? Online analytical processing is a technique of data analysis used in the business intelligence (BI) field.

BI involves using technology to analyze an organization's internal processes and data to support its decision making. When using online analytical processing for BI, historic data is often the

Notes

subject of the analysis, but BI can also encompass analysis of current and future states. Along with OLAP, other data management techniques that fall into the realm of BI include data mining, reporting, operational performance management, and predictive analytics.

Online Analytical Processing is used to answer the complex queries posted on data warehouse. In order to solve the queries of nature 'who?' and 'what?' We can use the simple tools but to answer the advanced queries like 'what if?' and 'why?', we require special tool that can support online analytical processing (OLAP).

OLAP is a term that describes a technology that uses a multi-dimensional view of aggregate data to provide quick access to strategic information for the purposes of advanced analysis. OLAP enables users to gain a deeper understanding and knowledge about various aspects of their corporate data through fast, consistent, interactive access to a wide variety of possible views of the data.

OLAP enables decision-making about future actions. Atypical OLAP calculation can be more complex than simply aggregating data, for example, 'What would be the effect on property sales in the different regions of Punjab if legal costs went up by 3.5% and Government taxes went down by 1.5% for properties over ₹ 100,000?'

Analytical Queries per Minute (AQM) is used as a standard benchmark for comparison of performances of different OLAP tools. OLAP systems should as much possible hide users from the syntax of complex queries and provide consistent response times for all queries no matter how complex.

Online analytical processing is frequently used for ad hoc reporting, and typically generates reports in a pivot or matrix format. Departments that may make use of OLAP include finance, operations, sales, and marketing. Types of uses can include budgeting and forecasting.

One of the defining characteristics of online analytical processing is the OLAP cube. The concept of the cube correlates the elements known as measures and dimensions, which describe the various measures' metadata. A relational database's snowflake or star schema tables may be the source of the metadata.

An example of a cube is using a business' individual accounts receivable amount as a measure, with a due date as a dimension.

OLAP uses databases that are designed with multiple dimensions. These databases may be smaller than those needed for the data warehousing capabilities that are often used for business intelligence. Compared to other types of analysis, fewer details of transaction are usually needed in online analytical processing. Not only are the OLAP databases often smaller than data warehouses, accessing the OLAP databases is often faster than accessing relational databases.

There are various specialties of online transaction processing. Several of the more frequently used specialties include multidimensional, relational, and hybrid. Multidimensional OLAP stores data in multidimensional arrays, relational OLAP uses relational databases, and hybrid OLAP uses a combination of the relational and specialized tables.

Though online transactional processing is an important technique in BI, more sophisticated tools or improvements to OLAP may be required for organizations that are interested in predictive analysis and business analytics. Predictive analysis is frequently used to forecast events such as customer buying behavior. Business performance data is usually the target of business analytics.

7.2.1 Historical Background

The first fully functional online analytical system was introduced in 1970 by Express and later on in 1995 the Oracle acquired the release for the resource of information in 2007 the official

ceremony for acquisition of OLAP was held in 2007. Oracle also released its own system called Essbase using the OLAP theoretical background and functionality. In the year of 1998 the Microsoft stepped in for upgrading and advancement in the OLAP technology. Microsoft worked on the mainstream idea and developed highly advanced online analytical system that is deployed in many large organizations today.

7.2.2 How OLAP Works

Online analytical processing use multidimensional model for processing an executing the given queries online the entire database hierarchical system is derived from the attached resources such as relational database from record keeping. It follows they cubic structure created by the star schema. The output is communicated in the form of the matrix first to the system and then to the ad hoc query makers. The programming structure and processing of online analytical systems are complexes the whole working. Typical knowledge is required fro dealing with these kinds of systems.

7.2.3 Types of OLAP

There are several types of the online analytical systems each with the different properties according to the level of use.

Multi-dimensional OLAP

Multi-dimensional OLAP is also called as MOLAP in short it's a classic type of leading online analytical systems. Instead of storing the data in relational form, it has the ability to store data in the multi-dimensional array that is highly optimized. This makes the data storage more efficient and quick. For keeping the data in the multi-dimensional array the MOLAP require pre processing and certain computational functions built in the internal cubic structure this computation is known as processing.


Relational OLAP

Another type of the online analytical processing models is relational OLAP which is also known as ROLAP in short. This works similar to the relational database as and store the information in the form of rows and columns in the particular sequence serialized by address. Base tables are created at the deep database and the new tables which are created by the users are aggregated to connect the data in meaningful way. This type of OLAP is used in the situations where data manipulation is frequently required which may involve slicing and dicing of the stored information. All the functions and structured language queries are applicable in relational OLAP. Data modifications can be performed by using WHERE clause in combination with sql statements.

Hybrid OLAP

The third type of the online analytical processing is the hybrid OLAP systems that aggregates the properties of both relational and multi dimensional OLAPs. It is also known as HOLAP in short. Hybrid OLAP divides the total data between the relational database and the specialized storage area for storing important data elements. This type is typically used by the large organization particularly the vendor and manufacturer. For instance for some vendors HOLAP serves as the central systems for partitioning the large amount of data into relational form and then use small aggregation functions to connect the smaller data stored in the special storage area. The data stored in the special storage is less detailed.

Notes

 *Task* Compare and contrast Multidimensional OLAP and Relational OLAP.

7.2.4 Features of OLAP

There are the following key features of OLAP:

- Multi-dimensional views of data
- Support for complex calculations
- Time intelligence

Multi-dimensional Views of Data

A multi-dimensional view of data provides the basis for analytical processing through flexible access to corporate data. It enables users to analyze data across any dimension at any level of aggregation with equal functionality and ease.

Support for Complex Calculations

OLAP software must provide a range of powerful computational methods such as that required by sales forecasting such as moving averages and percentage growth.

Time Intelligence

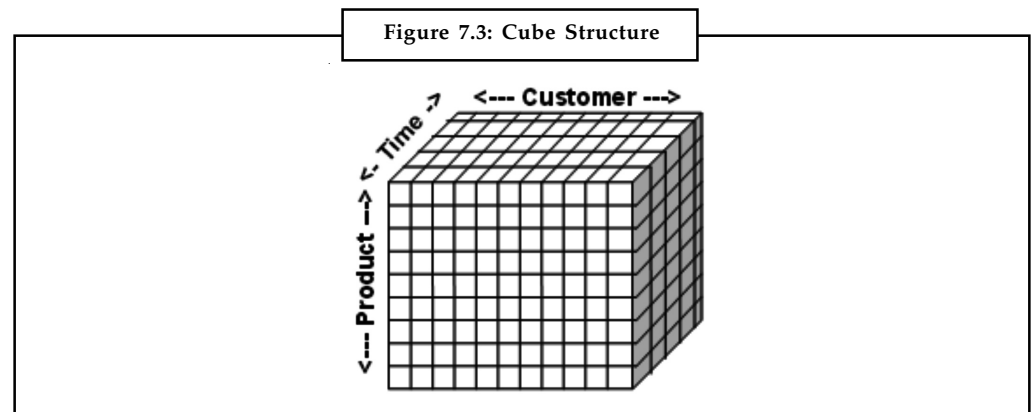
Time intelligence is used to judge the performance of almost any analytical application over time.

For example, this month versus last month or this month versus the same month last year or a user may require to view, the sales of the month of May or the sales for the first five months of 2007.

Concepts such as year-to-date and period-over-period comparisons should be easily defined in an OLAP system.

7.2.5 OLAP Cubes Support Business Intelligence

OLAP helps to visualize data as cube structures.



Source: <http://infogoal.com/datawarehousing/olap.htm>

A cube is a multi-dimensional structure consisting of dimensions and measurements. Cells are the points where dimensions intersect and contain the measurements.

Dimensions provide the context for analysis which are used for labels on reports and selection criteria for queries. Dimensions answer questions like:

- Who (customers, employees, partners, ...)
- When (year, quarter, month, ...)
- What (products, contracts, ...)
- Where (state, zip code, territory, ...)
- How (method, process, formula, ...)

Cells supply quantitative information. Cells answers questions like:

- How many (customer count, inventory count, ...)
- How much (revenue amount, budget amount, ...)

An atomic measurement is one that is stored at the lowest level, such as an individual sale or a single receipt of goods. The benefit of atomic data is that it supports detailed analysis and atomic data can be summed as needed. The drawbacks of atomic data are that it takes more space to store and it requires time to aggregate into totals for analysis.

Aggregated data is a summation of atomic data.



Example: Sales by quarter and rejects by month are aggregations.

The benefit is that query and analysis time are reduced. The drawbacks are that analysis detail can be lost and it is difficult to predict which aggregations the analyst will want to use.

Self Assessment

Fill in the blanks:

4. is a method of using multidimensional databases to support quick reporting, frequently involving trend-analysis.
5. OLAP works similar to the relational database as and store the information in the form of rows and columns in the particular sequence serialized by address.
6. OLAP divides the total data between the relational database and the specialized storage area for storing important data elements.
7. is used to judge the performance of almost any analytical application over time.
8. An is one that is stored at the lowest level, such as an individual sale or a single receipt of goods.

7.3 Geographical Information Systems (GIS)

GIS refers to three integrated parts.

- **Geographic:** Of the real world; the spatial realities, the geography.
- **Information:** Data and information; their meaning and use.
- **Systems:** The computer technology and support infrastructure.

Notes

GIS therefore refers to a set of three aspects of our modern world, and offers new ways to deal with them. Geographical Information Systems (GIS) are computer-based systems that enable users to collect, store, process, analyze and present spatial data.

It provides an electronic representation of information, called spatial data, about the Earth's natural and man-made features. A GIS references these real-world spatial data elements to a coordinate system. These features can be separated into different layers. A GIS system stores each category of information in a separate "layer" for ease of maintenance, analysis, and visualization.



Example: Layers can represent terrain characteristics, census data, demographics information, environmental and ecological data, roads, land use, river drainage and flood plains, and rare wildlife habitats. Different applications create and use different layers.

A GIS can also store attribute data, which is descriptive information of the map features. This attribute information is placed in a database separate from the graphics data but is linked to them.

A GIS allows the examination of both spatial and attribute data at the same time. Also, a GIS lets users search the attribute data and relate it to the spatial data. Therefore, a GIS can combine geographic and other types of data to generate maps and reports, enabling users to collect, manage, and interpret location-based information in a planned and systematic way.




Notes In short, a GIS can be defined as a computer system capable of assembling, storing, manipulating, and displaying geographically referenced information.

GIS is used to display and analyze spatial data which are tied to a relational database. This connection is what gives GIS its power: maps can be drawn from the database and data can be referenced from the maps. When a database is updated, the associated map can be dynamically updated as well. GIS databases include a wide variety of information: geographic, social, political, environmental, and demographic.

GIS uses layers, called "themes," to overlay different types of information, much as some static maps use mylar overlays to add tiers of information to a geographic background. Each theme represents a category of information, such as roads or forest cover.



Caution As with the old mylar maps, the layers which are underneath remain visible while additional themes are placed above.



Source: <http://www.mohaminfo.com/ITServices/Services/gis.aspx>

7.3.1 History of GIS

Notes

Spatial analysis has been around as long as maps have been used for navigation. Traditional GIS has included paper maps, and manual tools that were utilized before the computer age. There are many theories on how and where the original GIS began and took hold in the information-based world we live in today. There is little doubt that the Harvard Laboratory for Computer Graphics, later renamed by adding “and Spatial Analysis” to its title, was a key institution in the birth and early development of GIS in the United States. A large number of the founders of GIS in the academic and commercial sectors spent time at the Harvard Lab in the 1960s and 1970s, as staff, students, or visitors, exchanging ideas that would shape the field.


Most would acknowledge that the two dominant commercial GIS producers in the United States are Environmental Systems Research Institute (ESRI) and Intergraph. Estimates usually indicate that together they produce at least half of the GIS software in the country.

7.3.2 Components of GIS

The infrastructure to any GIS is made up of four components whose interrelation is shown on the illustration below. These components can be listed in order of importance as follows:

- **Methods and People:** The most important part of a GIS infrastructure. Although GIS is a powerful tool, it will not work without some well-adapted methods and trained people.
- **Data:** The needed raw material to be processed by the system. The focus and attention should be on data. In fact, most work will be devoted to data input and management. Data (information) is the foundation of GIS applications.
- **Software:** The computer programs needed to run GIS. There are many GIS programs available, from low-cost and low-performance packages to expensive and very powerful ones. This also includes support programs, such as statistical, word processing, graphing, and others.
- **Hardware:** The machinery on which GIS operates – computers, printers, plotters, digitizers, and other types of equipment.

Figure 7.5: Components of a GIS



Source: <http://igre.emich.edu/wsatraining/TManual/Chapter1/Chap1.pdf>

Notes

Unfortunately, GIS is often developed in the wrong order and more attention is given to the dazzling hardware and software, rather than trying to fit the data to the computer technology. The statement ‘garbage in; garbage out’ can be applied to the data collection process for the organization and the people who utilize and implement that data. The greatest piece of software, and the computers with the greatest amount of memory and speed, will not improve the information from data that is poorly gathered and inputted into the system.

7.3.3 Types of GIS

The following GIS types are not necessarily mutually exclusive and a GIS application can be always classified under more than one type.

Four-dimensional GIS

While spatio-temporal geo-representations can handle two dimensions of space and one of time, four-dimensional GIS are designed for three dimensions of space and one of time.

Multimedia/Hypermedia GIS


Multimedia/hypermedia GIS allow the user to access a wide range of geo-referenced multimedia data (e.g., simulations, sounds and videos) by selecting resources from a geo-referenced image map base. A map serving as the primary index to multimedia data in a multimedia geo-representation is termed a hypermap. Multimedia and virtual geo-representations can be stored either in extended relational databases, object databases or in application-specific data stores. GIS systems are dynamic and permit rapid updating, analysis, and display. They use data from many diverse sources such as satellite imagery, aerial photos, maps, ground surveys, and Global Positioning Systems (GPS).

Web GIS

Widespread access to the Internet, the ubiquity of browsers and the explosion of commodified geographic information has made it possible to develop new forms of multimedia geo-representations on the Web. Many current geomatics solutions are Web-based overtaking the traditional Desktop environment and most future ones are expected to follow the same direction.

Virtual Reality GIS

Virtual Reality GIS have been developed to allow the creation, manipulation and exploration of geo-referenced virtual environments, e.g., using VRML modelling (Virtual Reality Modelling Language). Virtual Reality GIS can be also Web-based. Applications include 3D simulation for planning (to experiment with different scenarios).



Task Make distinction between Four-dimensional GIS and Multimedia/hypermedia GIS.

7.3.4 Multimedia and Geographical Information System (GIS)

Multimedia is a technology that encompasses various types of data and presents them in an integrated form. There are several types of data that are used by the technology, including text,

graphics, hyperlinks, images, sound, digital and analogue video and animation. Although many GIS have been successfully implemented, it has become quite clear that two-dimensional maps cannot precisely present multidimensional and dynamic spatial phenomena. Moreover, there is a growing need towards accessing spatial data. It seems that merging GIS and Multimedia is a way to deal with these issues.

The latest advances in computer industry especially in hardware have led to the development of the Multimedia and Geographical Information System (GIS) technologies. Multimedia provides communications using text, graphics, animation, and video. Multimedia GIS systems is a way to overcome the limitations displayed by the technologies when they are used separately. Multimedia can extend GIS capabilities of presenting geographic and other information.

The combination of several media often results in a powerful and richer presentation of information and ideas to stimulate interest and enhance information retention. They can also make GIS more friendly and easier to use. On the other hand, multimedia can benefit from GIS by gaining an environment which facilitates the use and analysis of spatial data. The result is a system, which has the advantages of both worlds without retaining most of their disadvantages.

Self Assessment

Fill in the blanks:

9. GIS uses layers, called “.....” to overlay different types of information.
10. A map serving as the primary index to multimedia data in a multimedia geo-representation is termed a
11. is a technology that encompasses various types of data and presents them in an integrated form.

7.4 GIS Subsystems

A GIS has four main functional subsystems. These are:

7.4.1 Data Input Subsystem

A Data Input subsystem allows the user to capture, collect, and transform spatial and thematic data into digital form. The data inputs are usually derived from a combination of hard copy maps, aerial photographs, remotely sensed images, reports, survey documents, etc.

GIS Data Types

These data will contain maps of different detail levels (maps of the county, its main cities and villages, maps of the archaeological and historical sites etc.), photos of places and monuments, video images, text (in many languages), music and sound. For more complex applications, multimedia data can be remotely sensed imagery, scanned maps, digitized video clips, DTMs, one or more dimensional measurements, simulation model outputs and others. Most of them are complicated objects, which have large data volumes, intensive processing requirements and rich semantics.

The basic data types in a GIS reflect traditional data found on a map. Accordingly, GIS technology utilizes two basic types of data. These are:

- **Spatial data:** Spatial data describes the absolute and relative location of geographic features.

Notes

- **Attribute data:** Attribute data describes characteristics of the spatial features. These characteristics can be quantitative and/or qualitative in nature. Attribute data is often referred to as tabular data.

The coordinate location of a forestry stand would be spatial data, while the characteristics of that forestry stand, e.g. cover group, dominant species, crown closure, height, etc., would be attribute data. Other data types, in particular image and multimedia data, are becoming more prevalent with changing technology. Depending on the specific content of the data, image data may be considered either spatial, e.g. photographs, animation, movies, etc., or attribute, e.g. sound, descriptions, narrations, etc.

Sources of Data

A wide variety of data sources exist for both spatial and attribute data. The most common general sources for spatial data are:

- Hard copy maps
- Aerial photographs
- Remotely-sensed imagery
- Point data samples from surveys
- Existing digital data files

This spatial data is usually in analog form and needs to be converted to digital form before it can be used. Maps can be digitized, or hand traced with a computer mouse, to collect the coordinates of features.

Attribute data has an even wider variety of data sources. Any textual or tabular data that can be referenced to a geographic feature, e.g. a point, line, or area, can be input into a GIS.



Did u know? Attribute data is usually input by manual keying or via a bulk loading utility of the DBMS software.

Data Editing and Quality Assurance

Data editing and verification is in response to the errors that arise during the encoding of spatial and non-spatial data. The editing of spatial data is a time consuming, interactive process that can take as long, if not longer, than the data input process itself. Several kinds of errors can occur during data input. They can be classified as:

- **Incompleteness of the spatial data:** This includes missing points, line segments, and/or polygons.
- **Locational placement errors of spatial data:** These types of errors usually are the result of careless digitizing or poor quality of the original data source.
- **Distortion of the spatial data:** This kind of error is usually caused by base maps that are not scale-correct over the whole image, e.g. aerial photographs.
- **Incorrect linkages between spatial and attribute data:** This type of error is commonly the result of incorrect unique identifiers (labels) being assigned during manual key in or digitizing. This may involve the assigning of an entirely wrong label to a feature, or more than one label being assigned to a feature.

- **Attribute data is wrong or incomplete:** Often the attribute data does not match exactly with the spatial data. This is because they are frequently from independent sources and often different time periods. Missing data records or too many data records are the most common problems.

7.4.2 Data Storage, Editing and Retrieval Subsystem

The second necessary component for a GIS is the data storage and retrieval subsystem. The Data Storage and retrieval subsystem organizes the data, spatial and attribute, in a form, which permits it to be quickly retrieved by the user for analysis, and permits rapid and accurate updates to be made to the database. This component usually involves use of a database management system (DBMS) for maintaining attribute data. Spatial data is usually encoded and maintained in a proprietary file format.

Organizing Data for Analysis

Most GIS software organizes spatial data in a thematic approach that categorizes data in vertical layers. The definition of layers is fully dependent on the organization's requirements.

Typical layers used in natural resource management agencies or companies include forest cover, soil classification, elevation, road network (access), ecological areas, hydrology, etc.

Editing and Updating of Data

Perhaps the primary function in the data storage and retrieval subsystem involves the editing and updating of data. Frequently, the following data editing capabilities are required:

- Interactive editing of spatial data
- Interactive editing of attribute data
- The ability to add, manipulate, modify, and delete both spatial features and attributes (independently or simultaneously)
- Ability to edit selected features in a batch-processing mode.

Data Retrieval and Querying

The ability to retrieve data is based on the unique structure of the DBMS and command interfaces are commonly provided with the software. Most GIS software also provides a programming subroutine library, or macro language, so the user can write their own specific data retrieval routines if required.

Querying is the capability to retrieve data, usually a data subset, based on some user-defined formula. These data subsets are often referred to as logical views. Often the querying is closely linked to the data manipulation and analysis subsystem. Querying can be either by example or by content.

7.4.3 Data Manipulation and Analysis Subsystem

The Data Manipulation and Analysis subsystem allows the user to define and execute spatial and attribute procedures to generate derived information. This subsystem is commonly thought of as the heart of a GIS, and usually distinguishes it from other database information systems and Computer-Aided Drafting (CAD) systems.

Notes

Manipulation and Transformations of Spatial Data

The maintenance and transformation of spatial data concerns the ability to input, manipulate, and transform data once it has been created. Some specific functions are:

- Coordinate thinning: involves the reduction of the coordinate pairs (X and Y) from arcs.
- Geometric Transformations
- Map Projection Transformations
- Edge Matching
- Interactive Graphic Editing

Analytical Functions in a GIS

The primitive analytical functions that must be provided by any GIS are:

- Retrieval, Reclassification, and Generalization
- Topological Overlay Techniques
- Neighbourhood Operations
- Connectivity Functions

7.4.4 Data Output and Display Subsystem

The Data Output subsystem allows the user to generate graphic displays, normally maps, and tabular reports representing derived information products. This subsystem conveys the results of analysis to the people who make decisions about resources. Wall maps and other graphics can be generated, allowing the viewer to visualize and thereby understand the results of analyses or simulations of potential events.

Self Assessment

Fill in the blanks:

12. A subsystem allows the user to capture, collect, and transform spatial and thematic data into digital form.
13. describes the absolute and relative location of geographic features.

7.5 Applications of Multimedia GIS

Applications of multimedia GIS are discussed below:

7.5.1 Education

Education is a field where integration of multimedia and GIS can bring enormous benefits. Students will learn faster and more efficiently. In addition, it will be possible to individualize learning and tune it to particular preferences of each student. In this model a teacher becomes a guide rather than a repository of facts. It is the computer that takes on a role of "an infinitely patient teacher."

7.5.2 Map-making

GIS can use and combine all layers that are available for an area, in order to produce an overlay that can be analyzed by using the same GIS. Such overlays and their analysis radically change decision-making process that include, among others:

- Site selection
- Simulation of environmental effects



Example: Creating perspective views of a terrain before and after mining.

- Emergency response planning



Example: Combining road network and earth science information to analyze the effects of a potential earthquake.

7.5.3 Land Information

GIS has aided management of land information by enabling easy creation and maintenance of data for land records, land planning and land use. GIS makes input, updates, and retrieval of data such as tax records, land-use plan, and zoning codes much easier than during the paper-map era. Typical uses of GIS in land information management include managing land registry for recording titles to land holdings, preparing land-use plan and zoning maps, cadastral mapping, etc. Input of data into a land information GIS includes: political and administrative boundaries, transportation, and soil cover.

7.5.4 Infrastructure and Utilities

GIS technologies are also widely applied to the planning and management of public utilities. Typical uses include management of the following services: electric, gas, water, roads, telecommunication, storm sewers, TV/FM transmitting facilities, hazards analysis, and dispatch and emergency services. Typical data input includes street network, topographic data, demographic data and local government administration boundary.

7.5.5 Environmental

The environmental field has long used GIS for a variety of applications that range from simple inventory and query, to map analysis and overlay, to complex spatial decision-making systems.

Examples include: forest modeling, air/water quality modeling and monitoring, environmentally sensitive zone mapping, analysis of interaction between economic, meteorological, and hydrological & geological change.

Typical data input into an environmental GIS include: elevation, forest cover, soil quality and hydrogeology coverage.

7.5.6 Archaeology

Archaeology, as a spatial discipline, has used GIS in a variety of ways. At the simplest level, GIS has found applications as database management for archaeological records, with the added benefit of being able to create instant maps. It has been implemented in cultural resource management contexts, where archaeological site locations are predicted using statistical models

Notes

based on previously identified site locations. It has also been used to simulate diachronic changes in past landscapes, and as a tool in intra-site analysis.

7.5.7 Natural Hazards

Areas vulnerable to earthquakes, floods, cyclones, storms, drought, fire, volcano, land slides, soil erosion can be used to accurately predict future disasters.

7.5.8 Forestry

GIS has been emerging as a strong tool for many areas of forestry, from harvesting schedules to urban forestry.

7.5.9 Military GIS

GIS offers a virtually unique ability to aggregate, automate, integrate and analyze geographical data, which further enhance the intelligence base for defense operations.

7.5.10 Oceanography

GIS enables study of sea level change, marine population, sea surface temperature, and coral reef ecosystem.

7.5.11 Water Resources

GIS enables spatial representation of ground water resources, waste quality, watershed management, surface water management, and water pollution.

7.5.12 GIS in Agriculture and Soil

Data includes information on the country's land resources including physiography, soils, climate, hydrology, cropping systems and crop suitability.

Self Assessment

State True or False:

- 14. GIS makes input, updates, and retrieval of data much difficult than during the paper-map era.
- 15. Archaeology, as a spatial discipline, has used GIS in a variety of ways.



Case Study

Sun Banks on Data Visualization to Win Customers

Sun National, a subsidiary of Sun Bancorp, offers consumer and business services to customers in more than 70 community banking centers. The bank's portfolio includes electronic banking, retail and business accounts, and business, personal and wholesale commercial loans. The bank sought a way to differentiate itself because "the financial industry is dominated by huge banks with \$28 billion in assets, and we're a \$3-billion

Contd....

asset bank," says Louise O'Donnell, Sun National's vice president of management information services.

Financial institutions generally offer the same products, so the bank wanted to understand its customers better as a way to set itself apart from the competition. "A picture says a thousands words," and finding the best key performance indicators for analysts would help them evaluate their businesses, she says.

The bank then chose data visualization software as its best bet, going through a 2-year study to find the best product that fit in its environment and within its budget. "We needed something to get the biggest bang for our buck," O'Donnell says. Persuading upper management of the need for a data visualization app was a challenge with the bank's budget constraints, so the adoption process was methodical. O'Donnell and her IT staff conducted in-house case studies with software including Business Objects, Cognos and Advizor Solutions, dedicating about 45 days to each software product. The testing cycle consisted of bringing software in, building a demo and sharing it with different user groups.

As part of its criteria, the MIS department wanted data visualization software that fit into Microsoft SQL Server, so it selected Advizor's Analyst Edition software for its dashboard and analysis capabilities, which would allow users to monitor their departments through key performance indicators, drill down to underlying data, then publish their views and findings. Sun National also chose Microsoft Analysis Services and Reporting Services. As part of its push to data visualization, Sun National built a data warehouse to centralize its data, "so there's one version of the truth," O'Donnell says. To populate the SQL Server data warehouse and Analysis Services data marts, the bank uses Microsoft Integration Services to extract data from core systems, apply business rules, and aggregate and load the data.

The data marts are organized by business unit, including customer, retail and lending, and function, such as ATM. Analysts then select a dashboard to see their information, rather than viewing data in stand-alone spreadsheets and systems. They can now identify trends and anomalies across customer and sales information, as well as monitor cash flow projections and profitability trends. "They can create 65 reports which can be sorted 65 different ways," she says. Analysts only have to drag their mouse and highlight certain areas to drill down. "There's a lot of trending and [data visualization] complements the knowledge we really want to know about-particular products, customers and behaviors."

The bank's analysts aren't the only ones who have benefited from data visualization software. O'Donnell and her staff have been spared from producing reports for almost 25 departments, including deposit operations, accounting, small business loans and human resources. MIS – a three-member staff – would write queries when an analyst made a request and send data to him or her. If the analyst wanted to see more details, MIS would have to run the query again, a process that could take about several hours. The data visualization makes it possible to fulfill such requests in minutes. O'Donnell claims Advizor data visualization software has reduced its workload "easily by 20 to 30 percent." Prior to the Advizor implementation, trending was difficult to track, but now between Advizor and the data warehouse, Sun National has 3 years worth of information for trending. "Life is so much easier now to fulfill our informational needs," she says.

Question

Discuss the use of data visualization software.

Source: <http://www.advizorsolutions.com/markets/ADT%20Sun%20Banks%20on%20Data%20Visualization%20to%20Win%20Customers.pdf>

7.6 Summary

- Business analytics (BA) is the practice of iterative, methodical exploration of an organization's data with emphasis on statistical analysis.
- Once the business goal of the analysis is determined, an analysis methodology is selected and data is acquired to support the analysis.
- Data visualization is the presentation of data in a pictorial or graphical format.
- Visualizations convey information in a universal manner and make it simple to share ideas with others.
- Online analytical processing (OLAP) is defined as "The dynamic synthesis, analysis, and consolidation large volumes of multi-dimensional data."
- Online analytical processing (OLAP) is a method of using multidimensional databases to support quick reporting, frequently involving trend-analysis.
- Geographical Information Systems (GIS) are computer-based systems that enable users to collect, store, process, analyze and present spatial data.
- Multimedia is a technology that encompasses various types of data and presents them in an integrated form.

7.7 Keywords

Business Analytics: Business analytics (BA) is the practice of iterative, methodical exploration of an organization's data with emphasis on statistical analysis.

Data Input Subsystem: A Data Input subsystem allows the user to capture, collect, and transform spatial and thematic data into digital form.

Data Visualization: Data visualization is the presentation of data in a pictorial or graphical format.

GIS: Geographical Information Systems (GIS) are computer-based systems that enable users to collect, store, process, analyze and present spatial data.

Hybrid OLAP: Hybrid OLAP divides the total data between the relational database and the specialized storage area for storing important data elements.

Multimedia: Multimedia is a technology that encompasses various types of data and presents them in an integrated form.

OLAP: Online analytical processing (OLAP) is a method of using multidimensional databases to support quick reporting, frequently involving trend-analysis.

Virtual Reality GIS: Virtual Reality GIS allows the creation, manipulation and exploration of geo-referenced virtual environments.

7.8 Review Questions

1. Explain the concept of business analytics with example.
2. Describe the importance of Data visualization.
3. What is Interactive Visualization? Illustrate with example.
4. Explain the concept of OLAP with example.

5. Describe the working of online analytical processing.
6. Discuss the different types of online analytical processing (OLAP).
7. Illustrate how OLAP helps in visualizing data as cube structures.
8. Explain the basic data types in a GIS.
9. Discuss the various components of GIS.
10. Discuss various applications of GIS.

Notes

Answers: Self Assessment

- | | |
|----------------------------|--|
| 1. Business analytics (BA) | 2. Data visualization |
| 3. Visualization tools | 4. Online analytical processing (OLAP) |
| 5. Relational | 6. Hybrid |
| 7. Time intelligence | 8. Atomic Measurement |
| 9. Themes | 10. Hypermap |
| 11. Multimedia | 12. Data Input |
| 13. Spatial data | 14. False |
| 15. True | |

7.9 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

<http://linfiniti.com/dla/AGentleIntroductionToGIS.pdf>

http://www.cs.ubbcluj.ro/~per/Dss/Dss_12.pdf

http://www.sas.com/resources/asset/8Essentials_BusinessAnalytics_SpecialReport.pdf

<http://www.westminster.edu/staff/athrock/GIS/GIS.pdf>

Notes

Unit 8: Usages, Benefits and Success of Business Analytics and Data Visualization

CONTENTS

Objectives

Introduction

- 8.1 Usages of Business Analytics and Data Visualization
 - 8.1.1 Usages of Business Analytics
 - 8.1.2 Usage of Data Visualization
- 8.2 Benefits of Business Analytics and Data Visualization
 - 8.2.1 Benefits of Business Analytics
 - 8.2.2 Benefits of Data Visualization
- 8.3 Success Factors of Business Analytics and Data Visualization
 - 8.3.1 Success Factors of Business Analytics
 - 8.3.2 Success Factors of Data Visualization
- 8.4 Summary
- 8.5 Keywords
- 8.6 Review Questions
- 8.7 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss the Usages of Business Analytics and Data Visualization
- Discuss the Benefits of Business Analytics and Data Visualization
- Discuss Success Factors of Business Analytics and Data Visualization

Introduction

Focus on business analytics has increased steadily over the past decade as evidenced by the continuously growing business analytics software market. Business analytics is reaching more organizations and extends to a wider range of users, from executives and line of business managers to analysts and other knowledge workers, within organizations. In an environment of increasingly faster growing data volumes where operating on intuition is no longer an option, business analytics provide the means to both optimize the organization internally and at the same time maintain flexibility to face unexpected external forces. Visualized data is frequently displayed in business intelligence (BI) dashboards and performance scorecards that provide users with high-level views of corporate information, metrics and Key Performance Indicators (KPIs). In this unit, we will discuss the usages, benefits, and success of business analytics and data visualization.

8.1 Usages of Business Analytics and Data Visualization

In this section, we will discuss the usages of Business Analytics and Data Visualization.

8.1.1 Usages of Business Analytics

Leading banks use business analytics to predict and prevent credit fraud, saving millions. Retailers use business analytics to predict the best location for stores and how to stock them. Pharmaceutical firms use it to get life-saving drugs to market more quickly. Even sports teams are getting in on the action, using business analytics to determine both game strategy and optimal ticket prices.

But these advanced business applications tell only part of the story. What's going on inside these market-leading companies that sets them apart? They have committed to deploying their people, technologies and business processes in new ways. They have committed to a culture that is based on fact-based decisions – which helps them anticipate and solve complex business problems throughout the organization. By embracing an analytical approach, these companies identify their most profitable customers, accelerate product innovation, optimize supply chains and pricing, and identify the true drivers of financial performance.

And you can too. Get started with business analytics by taking these eight essential actions:

1. **Improve the flow and flexibility of data:** High-quality data must be integrated and accessible across your organization. It should also be structured in a flexible way that allows your analysts to discover new insights and provide leaders the information they need to adjust strategies quickly.



Notes Strengthening and flexing the data backbone of your enterprise will pay off when you need to change business processes quickly in response to market shifts, regulatory or stakeholder demands.

2. **Get the right technology in place:** Take an enterprise approach to data management and analytics to effect better decisions. Remove disconnected silos of data, technology or expertise. Your technology portfolio should include:
 - ❖ Optimized data stores to support core business processes and discovery.
 - ❖ Data integration and data quality software.
 - ❖ Analytical software with the means to effectively deploy, explore and share results in a meaningful way.
 - ❖ Integrated analytical applications designed to solve defined issues quickly.

When selecting technologies, consider “risk-to-value”: Can the technology be applied to help reduce costs and increase revenue? And getting the right technology in place doesn't have to mean a complete overhaul.

3. **Develop the talent you need:** Develop or recruit analytic thinkers who seek and explore the right data to make discoveries. To make analytics work, analysts must also be able to communicate effectively with leaders and link analytics to key decisions and the bottom line.
4. **Demand fact-based decisions:** An analytical company makes a wide range of decisions. Some are ad hoc; some are automated; some are transformative. The common thread? Evidence backs them all. Managers encourage asking the right questions of the data to get

Notes

maximum insight. How results are deployed is also important – through operation systems such as customer relationship management applications or real-time fraud applications to interactive dashboards, data movies, in databases – wherever needed to ensure decision makers have the information they need when they need it (and in the way they can best consume it).

5. **Keep the process transparent:** Transparency implies openness, communication and accountability; it is key to successful business analytics projects.



Caution The value delivered from an investment in business analytics must be visible and measureable.

Who the analysts are and what they're seeking to accomplish should be clearly communicated to the business, as should their findings.

6. **Develop an analytical center of excellence:** Create a centralized team approach – an analytical center of excellence (ACE) – which promotes the use of analytics and associated best practices. Your implementation of an ACE will depend on your organization's maturity and requirements, but the most effective implementations address all elements of the organization's analytic infrastructure: people, process, technology and culture to support the business' strategy and operations.
7. **Transform the culture:** A strong analytical culture has executive sponsorship and encourages creativity. Experimentation should be seen as part of learning, and employees should be given permission to fail as they learn from trying new things.
8. **Revise your strategies – often:** Your competitors will often duplicate your analytical initiatives. Staying ahead requires continuous review of strategy and development of new skills and capabilities.

8.1.2 Usage of Data Visualization

The modes of information communication evolve constantly in order to improve its efficiency, clarity and aesthetic appeal. Generally, there are no wrong ways to communicate information but the traditional textual forms are slowly giving their way to data visualisation. Whether the latter is really the best way to communicate information or not depends on several factors including the type of data you want to present and the target audience to which you are communicating information

If you are having difficulties deciding between communication of data in the "raw" form or using visualisation methods such as graphs, dials, charts, etc. instead, the following overview of advantages of data visualisation may help:

- **Clarity:** It is a lot easier to understand a dial or graph than numbers. The viewer understands what you are trying to say at a first sight.
- **Saving time:** Since a "picture is worth a million words", using data visualisation helps the audience quickly absorb and interpret the presented data. As a result, data visualisation enables you to present a considerably larger amount of data in comparison to the textual format which often requires repetition in order to help the audience understand the information.
- **Less confusion:** It is not difficult to get confused when dealing with lots of numbers as you actually need to memorize them to be able to understand the communicated information. Using visual presentation of numbers, however, dramatically reduces confusion because the audience does not need to process the numbers to be able to see where you are going.

Visualisations look better and attract more attention than the textual format. They are also more likely to keep the audience interested in your presentation.

Notes

Although data visualisations are easier to understand and look more attractive to the audience, it is crucial to achieve a perfect balance between visual appeal and functionality. Data visualisation is in the first place used to improve efficiency of the communicated information. A beautiful presentation which, however, fails to emphasise relevant data or is not clear enough is of little value. For that reason it is highly important to make sure that the presented data are clear and understandable, and only then focus on aesthetically appealing and attention drawing design. Just as important is to support the communicated information with additional materials such as official statistic data, facts, examples, etc., if you want the audience to accept your view/ interpretation.

In the end, it is necessary to mention that the use of data visualisation does not necessarily exclude the textual format or vice versa. If you are dealing with numbers, you cannot avoid them completely no matter how sophisticated visualisations you use. Nevertheless, it is a lot easier to make a point and help the audience understand the importance of the communicated information if you also use visual presentation along the numbers.

Self Assessment

State True or False:

1. High-quality data must be integrated and accessible across your organization.
2. Take an enterprise approach to data management and analytics to effect better decisions.
3. It is a very difficult to understand a dial or graph than numbers.
4. If you are dealing with numbers, you cannot avoid them completely no matter how sophisticated visualisations you use.
5. Transparency is not necessary for successful business analytics projects.

8.2 Benefits of Business Analytics and Data Visualization

In this section, we will discuss the benefits of Business Analytics and Data Visualization.

8.2.1 Benefits of Business Analytics

Business analytics is a term that encompasses many things. To put it simply, studying previous business performance to gain insight into what future business performance will be is business analytics. Business analysis employs many different methods to help them with this including data and applications. You may have also heard of the term 'business intelligence' which is very similar to business analytics in that there is also the collection and study of data. However, the difference between the two is that business analytics strives to find new and better ways while business intelligence is about using tried and true methods. Check out Transportation Consulting & Management Services, for more details.

As you can imagine, for one medium sized company, there would be a substantial amount of data and analysis needed to make good decisions. And when analyzing a company, you need to analyze each part of the business, e.g. production, human resources, environmental etc.



Example: With a company that manufactures products, just the getting the finished products to the customers is a big undertaking that would result in a lot of data that needed to be processed.

Notes

Each step that is required to move product from supplier to customer is referred to as the supply chain strategy. To provide business analysis of a supply chain an individual would not only have to collect data from the manufacture but also from the trucking companies and retailers. And there are other the areas like production that would need a separate analysis as well.

The majority of businesses are not capable of conducting their own analysis. A company interested in improving how their products get from them to their customers would need to seek the advice of supply chain consultants. Supply chain consultants are experts who provide their expertise to businesses and individuals in order to help them become successful. Supply chain consultants have what it takes to investigate past procedures and data to make the informed decision to make things better.

It can be very costly for a company to hire a supply chain consulting firm or any other kind of consulting firm. But the value of business analytics far outweigh the cost. Thorough analysis means more profits. An experiences consulting firm can help a company increase their profits threefold. If you are starting a business, surely you have thought about all of the normal cost: supplies, staff, utilities, etc. You should also include business analytics as well. The number one reason companies aren't as successful as they could be is because they don't put enough emphasis on analysis and consulting.



Caution Business analytics is crucial to the success of any company and if you have a business it is suggested to add that to your to-do list.

The five most popular benefits include:

1. Improving the decision-making process.
2. Speeding up the decision-making process.
3. Better alignment of resources with strategies.
4. Realizing cost efficiencies.
5. Responding to user needs for availability of data on a timely basis.

8.2.2 Benefits of Data Visualization

Data visualization is the practice of representing data in graphic or abstract form in order to concisely demonstrate information and results. Data visualization greatly enhances not only data comprehension, it allows for a greater coverage of trends and patterns within data.

For example, billing and medical debt collectors have been using data visualization in order to best understand how to impact billing operations and increase efficiency.

David Stodder explains and clarifies the numerous benefits of using data visualization through the article. One of the top benefits is that data is easier to understand and is more accessible in this format such that people can better interact with the data and analyze it. Normally data is used for mining and reporting. A few ad-hoc data visualization tools, however, can take your data analysis to a whole new level.



Example: According to Stodder, it is possible to quickly understand trends, patterns, and outliers that would not have been noticed with traditional reporting.

A huge benefit of having greater visibility and accessibility allows more better insight into your business.

Some of the benefits are discussed below:

- **Improve Understanding:** The data purpose is clear and you can dig into the details directly on the screen to quickly obtain awareness about any problems and discover a new perception.
- **Maintain Focus:** Data visualizations can be strikingly appealing. The mixture of data, visuals, and interactivity creates a method that can inform, engage, and influence viewers. Not many people enjoy staring at enormous data sets, or rows and rows of numbers, and fewer can make sense of them without some type of data visualization tool, even if it's only a pivot table.
- **Tackle the growing volume of data:** With huge amounts of incoming data, businesses need to transform the data into simple visuals for effective evaluation. Interaction with large data sets accelerates the analysis process.
- **Improve Decision-Making:** By generating and analyzing data visualizations, you can quickly identify dependencies, trends, and abnormalities that otherwise may go unnoticed. Overlooking important pieces of data could be a disaster for your business, but data visualization offers accurate views of your information to pinpoint specific data and therefore speed up the decision-making process.

As we continue to collect growing amounts of data, possessing the tools and the ability to take it and break it down for everyone to understand will only become more critical. In the past, data visualization was viewed as a key analytical tool for researchers; it is now recognized as a vital phase of successful research communication.

Self Assessment

State True or False:

6. Business intelligence' is different from business analytics.
7. Business analytics is crucial to the success of any company and if you have a business it is suggested to add that to your to-do list.
8. The majority of businesses are not capable of conducting their own analysis.
9. Data visualization is the practice of representing data in graphic or abstract form in order to concisely demonstrate information and results.
10. Data visualization is not considered a very important phase of successful research communication.

8.3 Success Factors of Business Analytics and Data Visualization

In this section, we will discuss the success of Business Analytics and Data Visualization.

8.3.1 Success Factors of Business Analytics

While most large companies have some unit deeply engaged in data mining, predictive analysis or forecasting, very few have embedded analytic thinking throughout the culture.




Did u know? Rising midsize and smaller companies often struggle over where to start.

The advice to begin small and build on successes still remains true. Yet there is nothing more painful than watching a company delay an analytics project over and over because they are

Notes

trying to get just the right data set or they want to understand just one more thing about the process before they jump in. Please go ahead and dive in. And as you analyze the results, consider the Do's and Don'ts list for refining the process.

- **Don't second-guess your analytic results:** This happens time and again. A company will invest in analytics but not trust the results. This often occurs when organizations fail to get executive buy-in prior to rolling out an extensive analytics initiative. But failing to stick to what the analysis suggests renders your efforts moot. If all you do is say "When it matches the hypothesis we will run with it. When it doesn't we will override it," then you are not using analytics. And yes, sometimes it is hard to stick to your guns – particularly if the recommendation is a little uncomfortable or different from what your organization traditionally does. The key is this: When you can't follow recommendations that go against traditional thinking, analytics just becomes a layer that reinforces conventional wisdom – not something that helps your organization grow.
- **Do respect the creative elements of analytics:** An organization can go too far in assuming that analytics is a pure science. It's not. There is science involved in building a model, but questions like "What is the right thing to predict?" and "What factors are needed to build the model?" require the artistic and creative efforts of business users who think about these problems daily. When we talk to companies about model building, we emphasize the need to bring the scientists together with the non-scientific business people. Doing that ensures that the analytics address the right problem in the best possible way.
- **Don't be afraid of new data sources:** We've leapt from having a household file with demographics to factoring in transactions and other data points. Leading-edge companies are adding Web interaction information like what the customers are looking at on the Web, what reviews they are reading, what product pictures they are zooming in on and what search terms they used to get to the site. This data is important to understand what is going on inside a customer's head before they make a purchase. In addition, sensors and RFID data are critical new data sources, particularly relating to supply chains, transportation and manufacturing. There are countless other data sources arising across all industries. In fact, there are so many that the term "big data" has become popular these days as a catchall term for the wealth of new, large data sources. The more you can take every one of these pieces and stitch them together, the more you'll know about your customers and processes. Organizations that do this will be far more successful than those sitting back frightened about incorporating new data sources.
- **Do stay on the cutting edge:** Along with embracing new data sources, organizations need to embrace new ways of looking at data. This might include looking at cost-effective ways to speed model processing, pursuing additional modeling techniques or improving the way analytic results are distributed to users. The last thing you want is to have your team training a new person and starting the conversation with "Here's how we do it here. This is how you'll do it too."



Notes It is important to have standard procedures and approaches, but you also need to regularly challenge them and ensure there isn't room for improvement. After all, you won't be the leader if you are simply copying what everyone else is doing.

- **Don't expect one person to lead the charge:** We have seen companies – typically the industry laggards in using analytics – decide that they are going to hire one person to handle all their analytic needs. After a year or two, when this poor, beleaguered soul has not single-handedly transformed their business, they decide that "predictive analytics doesn't work."

- **Do look internally for analytic talent:** Large companies often have areas within their firms with mature analytics users. We say cross-pollinate and take advantage of the skills you've already built up in addition to looking outside. With so much demand for analytics talent, you need to keep your people challenged and engaged so they stick with you. Giving them a new internal opportunity is a great way to do that.



Task Conduct a research and analyze the application areas of business analytics.

8.3.2 Success Factors of Data Visualization

Demand for data visualization tools is rising sharply, partly as a result of more companies seeking to gain valuable business insights through “big data” analytics initiatives. But achieving success with data visualization often requires fresh thinking about how to present information to business users, especially in big-data environments, according to data management analysts.

Data visualization — which enables users to create graphical and often interactive representations of data sets big and small — can contribute greatly to improvements in corporate business intelligence (BI) efforts and organizational productivity.

Informed decision making is the foundation upon which successful businesses are built. As a decision maker for your business, you need access to highly visual business intelligence tools that can help you make the right decisions quickly. As your organization grows, so does the amount of collected information. If this data is delivered to you in spreadsheets or tabular reports, it becomes more and more challenging to find the patterns, trends and correlations necessary to perform your job well.

Effective data visualization is an important tool in the decision making process. It allows business decision makers to quickly examine large amounts of data, expose trends and issues efficiently, exchange ideas with key players, and influence the decisions that will ultimately lead to success.

The practice of representing information visually is nothing new. Scientists, students, and analysts have been using data visualization for centuries to track everything from astrological phenomena to stock prices. Only recently, with the adoption of more sophisticated BI technology in the corporate world and the ever-increasing practice of data collection and data mining activities, has data visualization in the form of dashboards been used as an important presentation tool in business analysis. As a result, the use of dashboards in making quick and accurate business decisions has become an essential requirement for remaining competitive.

Common Forms of Data Visualization

- **Basic Charts:** The most recognizable and utilized form of data visualization is the basic chart. Line, bar, area and pie charts represent the most common types of this form. The first function of a good chart is to allow decision makers to examine the data and reduce the time required to extract key information.
- **Status Indicators:** In addition to basic charts that visualize a set or sets of data, status indicators are also a commonly used visualization to indicate the business condition of a particular measure or unit of data. These indicators can take on many forms, including gauges, traffic lights or symbols. Status indicators become even more effective when they incorporate contextual metrics, such as targets and thresholds, because they can provide quick feedback as to whether a specific measure is good or bad, high or low, below or above target.

Notes

- **Advanced Data Visualizations:** More advanced examples of data visualization include scatter graphs, bubble charts, spark line charts, geographical maps, tree maps, Pareto charts, and many others. These more sophisticated visualizations are designed to display data in ways tailored to a specific function or industry.
- **Quick Analysis:** Successful visuals that depict measurable, actionable data allow decision makers to easily pinpoint and examine outliers. They also allow quick analysis to expose patterns, correlations, business conditions and trends.

Analysts who do not know what the target should be or who do not have the background information to assist them, will interpret this gauge differently than someone who has additional knowledge of the situation. This leads to confusion, missed opportunities and loss of time.

However, if you add context to the gauge in the form of a target and adjust the scale of the gauge so that the start and end points are more in line with that target, you can clearly see that the number of hits of this landing page is clearly lower than desired.



Did u know? Context allows a story to be told by the data without the risk of misinterpretation and allows everyone to come to the same conclusion.

- **Take Action:** Decision makers need to interact with their data to expose trends, highlight opportunities and raise red flags quickly and accurately. Their data should answer key questions and provide insight into issues that contribute directly to the decision making process. Presenting this data visually and adding contextual information to complement the analysis process not only makes it quicker and easier to pinpoint areas of opportunities and concern, but also enables decision makers to take action with their data. Successful data visualization provides the ability to expose problem areas and communicate those problems universally. Not being able to clearly identify and share your discoveries to back up your decisions can mean the difference between taking appropriate and decisive action and losing momentum or failing to act.

Using data visualization to display large amounts of data is nothing new. However, its value and use in making business decisions is often overlooked or poorly implemented. The key to success in using data visualization is ensuring that: the best and most appropriate types of visualizations are used; the data is always put into perspective with contextual information allowing for the information to be universally understood; and that the data being measured within the data visualization enables the user to take action based on the observations being made. With a good set of visuals that keep these key success factors in mind, decisions can be made more quickly and with more confidence so that your business can continue to grow.



Task Compare and contrast line chart and pie chart.

Self Assessment

State True or False:

11. Informed decision making is the foundation upon which successful businesses are built.
12. Data visualization enables users to create graphical and often interactive representations of data sets big and small.

13. By means of effective data visualization, business decision makers cannot quickly examine large amounts of data.
14. The first function of a good chart is to allow decision makers to examine the data and reduce the time required to extract key information.
15. Status indicators are commonly used visualization to indicate the business condition of a particular measure or unit of data.

Notes



Case Study

A Case Study in Business Analytics

Procter & Gamble

P&G's has 127,000 employees and 300 brands sold in 180 countries. P&G averages about 4 billion transactions daily. P&G CEO Bob McDonald has staked out a strategy to "digitize" the company's processes from end to end, and Business Sufficiency, Business Sphere and Decision Cockpits is enabler of that agenda.

P&G is building analytics expertise at a time when P&G is cutting costs in other areas, including eliminating 1,600 non-manufacturing jobs. The company's IT organization itself has cut \$900 million in total spending over the past nine years.

P&G is investing in analytics talent, even as the company cuts in other areas, to speed up business decision making. CIO Filippo Passerini says he plans to increase fourfold the number of company staff with expertise in business analytics.

CIO Passerini who leads the Global Business Services (GBS organization) is investing in analytics expertise (almost like a business competency center or center of excellence) because the model for using data to run a company is changing.

Business Sphere and Decision Cockpits

P&G has made available to 38,000 users analytical solutions called Business Sphere and Decision Cockpits. The Business Sphere was developed in partnership with BOI, Cisco, HP, SAP, Nielsen and TIBCO Spotfire.

The first project, launched in 2010, is the Business Sufficiency program, which gives executives predictions about P&G market share and other performance stats six to 12 months into the future. At its core is a series of analytic models designed to reveal what's happening in the business now, why it's happening, and what actions P&G can take. The "what" models focus on data such as shipments, sales, and market share. The "why" models highlight sales data down to the country, territory, product line, and store levels, as well as drivers such as advertising and consumer consumption, factoring in region- and country-specific economic data. The "actions" analysis look at levers P&G can pull, such as pricing, advertising, and product mix, and provide estimates on what they deliver.

Business Sphere is the further integration of technology, visualization, and information enables leaders to drill-down into data to get answers in real-time. To answer a set of questions, the program analyzes and connects as much as 200 terabytes of data (equal to the amount of information contained in 200,000 copies of Encyclopedia Britannica), allowing for unprecedented granularity and customization.

The way the data is presented uncovers insights, trends, and opportunities for the business leaders and prompts them to ask different and very focused business questions. If one question elicits a follow-up question, it can be addressed with data on-the-spot.

Contd....

Notes

The visualization helps people to “see” the data in ways they would not have been able to with just numbers and spreadsheets. It challenges assumptions while simultaneously presenting the data in different ways, revealing potential solutions that previously may have not been apparent.

Evolution of the P&G Decision Making Model

The old IT model was to figure out which reports people wanted, capture the data, and deliver it to the key people weeks or days after the fact. “That model is an obsolete model,” Passerini says.

The new model Passerini envisions is something of a virtual, instant-on war room, where people huddle in person or by video around the needed data, pulling in the right experts to fix a problem the moment it arises.

This decision-making environment requires better collaboration via easy-to-use video, more real-time data, and business analytics expertise.

A new building block is high-quality video-conferencing, because people solve hard problems faster and better when they can see one another, Passerini maintains. P&G has been an avid user for several years of room-sized Cisco telepresence systems. The video is used as part of a collaboration environment P&G calls Business Sphere, which CEO Bob McDonald and his executive council use to collaborate with colleagues worldwide. It combines video with large screens that display data visualizations on sales, market share, ad spending and the like, so everyone in the meeting is seeing the same information. In the past year, P&G added 50 smaller Business Sphere systems around the world, giving more people access to the technology.

Passerini’s team is working on a video platform that broadens access even more by letting people join in regardless of the video system they’re using, whether it’s Cisco telepresence or WebEx or FaceTime. That would mean a key team member can video in from an iPad, Droid smartphone, or PC if need be. In terms of data, this strategy needs the right real-time data.

What’s real time? The goal P&G’s working toward is that as soon as data is collected, it’s available for use, Passerini says. P&G isn’t after new data types; it still wants to share and analyze point-of-sale, inventory, ad spending, and shipment data. What’s new is the higher frequency and speed at which P&G gets that data, and the finer granularity. Passerini says P&G has about two-thirds of the real-time data it needs.

Passerini talks about the what, why, and how of a problem. “What” is the problem itself — is market share stable or has it shrunk two points? He thinks P&G has beaten the what problem by giving 58,000 employees business intelligence “cockpits,” which are dashboards that link to common data sources so people spend little time arguing over whose data to use. “Why” is the cause of a problem — was it a bad TV ad, out-of-stock shelves, or a competitor’s new product or price cut that caused a problem? Right now, the P&G IT team is working on automating analysis of the why, so employees get alerts when key events like a supply chain snafu or rival product launch happen.

If P&G can eliminate “what” discussions and some of the “why,” and decision-makers can jump right to how to solve a problem, “that radically increases the pace at which they do business,” Passerini says.

The final piece is bringing in that business analytics expertise. These are people “at the intersection of business and IT,” Passerini says. They need to be as well versed in P&G business issues as a marketing pro. And they need to be skilled in finding information,

Contd....

building data models, and creating simulations. For example, when CEO Bob McDonald and his executive committee meet each Monday, one data slice they look at is the “top 503 – combinations of products and country markets (Brazil hair care, perhaps, or U.S. pet care) that are the company’s 50 largest, making up about 60% of sales. Data visualizations show at a glance if sales or share are moving materially. If they are, and executives want to drill down, Jeffrey Goldman is the business analyst in that key meeting who delivers those insights, delivering analysis in real time on screens that all the executives see. Is a sales dip in detergent in France because of one retailer, so that’s where to focus? Is that retailer buying less only in France, or across Europe? Did P&G cut promotions or raise prices, letting a rival grab share, or is the category overall losing sales? Goldman delivers this kind of data so executives can decide how to respond. Passerini pictures analytics experts like Goldman sitting in on more meetings to make sure the “how” to solve problems gets sorted out right then and there, not postponed until everyone gets more information. The old model would mean “let’s get back to this in two weeks,” he says. “You need to be able to answer that question immediately.” Passerini describes the video and data collaboration efforts, and the role of the business analysts, as “harmonizing” how people do business across P&G. It’s the opposite of creating standard reports. It’s about creating a standard environment with the right tools, then it’s up to the experts in that room to use whatever data they need to make the right decisions.

Question

Discuss the Evolution of the P&G Decision Making model.

Source: <http://www.saama.com/blog/bid/79545/Procter-Gamble-A-Case-Study-in-Business-Analytics>

8.4 Summary

- Take an enterprise approach to data management and analytics to effect better decisions.
- Develop or recruit analytic thinkers who seek and explore the right data to make discoveries.
- Generally, there are no wrong ways to communicate information but the traditional textual forms are slowly giving their way to data visualisation.
- Data visualisation enables you to present a considerably larger amount of data in comparison to the textual format which often requires repetition in order to help the audience understand the information.
- Business analysis employs many different methods to help them with this including data and applications.
- Data visualization is the practice of representing data in graphic or abstract form in order to concisely demonstrate information and results.
- Paralysis is a common condition when it comes to organizations and analytics.
- Informed decision making is the foundation upon which successful businesses are built.

8.5 Keywords

Basic Charts: Basic charts allow decision makers to examine the data and reduce the time required to extract key information.

Business Analytics: Business analytics (BA) is the practice of iterative, methodical exploration of an organization’s data with emphasis on statistical analysis.

Notes

Data Visualization: Data visualization is a general term used to describe any technology that lets corporate executives and other end users “see” data in order to help them better understand the information and put it in a business context.

Informed Decision Making: Informed decision making is the foundation upon which successful businesses are built.

Status Indicators: Status indicators are a commonly used visualization to indicate the business condition of a particular measure or unit of data.

8.6 Review Questions

1. Discuss the usages of Business Analytics.
2. What are the benefits of Business Analytics? Discuss.
3. Explain the Success Factors of Business Analytics.
4. Discuss the uses of Data Visualization.
5. What are the benefits of Data Visualization? Discuss.
6. Discuss the success factors of Data Visualization.
7. Discuss the essential actions taken in business analytics.
8. The majority of businesses are not capable of conducting their own analysis. Comment.
9. Explain the role of data visualization in improving decision making.
10. Discuss the Do’s and Don’ts of business analytics.

Answers: Self Assessment

- | | |
|-----------|-----------|
| 1. True | 2. True |
| 3. False | 4. True |
| 5. False | 6. False |
| 7. True | 8. True |
| 9. True | 10. False |
| 11. True | 12. True |
| 13. False | 14. True |
| 15. True | |

8.7 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

Notes

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

http://informationashvins.files.wordpress.com/2012/04/varshney_icassp2012.pdf

http://mines.humanoriented.com/classes/2009/fall/csci568/visualizing_data.pdf

<http://www.bentley.edu/centers/sites/www.bentley.edu.centers/files/csbig/wang2.pdf>

http://www.cs.ubbcluj.ro/~per/Dss/Dss_12.pdf

Unit 9: Data Mining

CONTENTS

Objectives

Introduction

9.1 Concepts of Data Mining

9.1.1 Types of Information

9.1.2 Data Mining and Knowledge Discovery

9.1.3 Types of Data

9.1.4 Data Mining Functionalities

9.1.5 Working of Data Mining

9.1.6 Categorization of Data Mining Systems

9.1.7 Issues in Data Mining

9.2 Applications of Data Mining

9.3 Summary

9.4 Keywords

9.5 Review Questions

9.6 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss the Concepts of Data Mining
- Explain Data Mining and Knowledge Discovery
- Discuss Applications of Data Mining

Introduction

Data mining, the extraction of hidden predictive information from large databases, is a powerful new technology with great potential to help companies focus on the most important information in their data warehouses. Data mining tools predict future trends and behaviors, allowing businesses to make proactive, knowledge-driven decisions. The automated, prospective analyses offered by data mining move beyond the analyses of past events provided by retrospective tools typical of decision support systems. Most companies already collect and refine massive quantities of data.

9.1 Concepts of Data Mining

Data mining uses a relatively large amount of computing power operating on a large set of data to determine regularities and connections between data points. Algorithms that employ techniques from statistics, machine learning and pattern recognition are used to search large databases automatically. Data mining is also known as Knowledge-Discovery in Databases (KDD).

Like the term artificial intelligence, data mining is an umbrella term that can be applied to a number of varying activities. In the corporate world, data mining is used most frequently to determine the direction of trends and predict the future. It is employed to build models and decision support systems that give people information they can use. Data mining takes a frontline role in the battle against terrorism. It was supposedly used to determine the leader of the 9/11 attacks.

Data miners are statisticians who use techniques with names like near-neighbor models, k-means clustering, holdout method, k-fold cross validation, the leave-one-out method, and so on. Regression techniques are used to subtract irrelevant patterns, leaving only useful information. The term Bayesian is seen frequently in the field, referring to a class of inference techniques that predict the likelihood of future events by combining prior probabilities and probabilities based on conditional events.



Notes Spam filtering is arguably a form of data mining, which automatically brings relevant messages to the surface from a chaotic sea of phishing attempts and Viagra pitches.

Decision trees are used to filter mountains of data. In a decision tree, all data passes through an entrance node, where it faces a filter that separates the data into streams depending on its characteristics.



Example: Data about consumer behavior is likely to be filtered based on demographic factors.

Data mining is not primarily about fancy graphs and visualization techniques, but it does employ them to show what it has found. It is known that we can absorb more statistical information visually than verbally and this format for presentation can be very persuasive and powerful if used in the right context.

As our civilization becomes increasingly data-saturated and sensors are distributed en masse into our local environments, we will inadvertently discover things that might be missed on the first pass over. Data mining will let us correct these mistakes and discover new insights based on past data, giving us more bang for our data storage buck.

9.1.1 Types of Information

We have been collecting a myriad of data, from simple numerical measurements and text documents, to more complex information such as spatial data, multimedia channels, and hypertext documents. Here is a non-exclusive list of a variety of information collected in digital form in databases and in flat files.

- **Business transactions:** Every transaction in the business industry is (often) “memorized” for perpetuity. Such transactions are usually time related and can be inter-business deals such as purchases, exchanges, banking, stock, etc., or intra-business operations such as management of in-house wares and assets. Large department stores, for example, thanks to the widespread use of bar codes, store millions of transactions daily representing often terabytes of data. Storage space is not the major problem, as the price of hard disks is continuously dropping, but the effective use of the data in a reasonable time frame for competitive decision-making is definitely the most important problem to solve for businesses that struggle to survive in a highly competitive world.

Notes

- **Scientific data:** Whether in a Swiss nuclear accelerator laboratory counting particles, in the Canadian forest studying readings from a grizzly bear radio collar, on a South Pole iceberg gathering data about oceanic activity, or in an American university investigating human psychology, our society is amassing colossal amounts of scientific data that need to be analyzed. Unfortunately, we can capture and store more new data faster than we can analyze the old data already accumulated.
- **Medical and personal data:** From government census to personnel and customer files, very large collections of information are continuously gathered about individuals and groups. Governments, companies and organizations such as hospitals, are stockpiling very important quantities of personal data to help them manage human resources, better understand a market, or simply assist clientele. Regardless of the privacy issues this type of data often reveals, this information is collected, used and even shared. When correlated with other data this information can shed light on customer behaviour and the like.
- **Surveillance video and pictures:** With the amazing collapse of video camera prices, video cameras are becoming ubiquitous. Video tapes from surveillance cameras are usually recycled and thus the content is lost. However, there is a tendency today to store the tapes and even digitize them for future use and analysis.
- **Satellite sensing:** There is a countless number of satellites around the globe: some are geostationary above a region, and some are orbiting around the Earth, but all are sending a non-stop stream of data to the surface. NASA, which controls a large number of satellites, receives more data every second than what all NASA researchers and engineers can cope with. Many satellite pictures and data are made public as soon as they are received in the hopes that other researchers can analyze them.
- **Games:** Our society is collecting a tremendous amount of data and statistics about games, players and athletes. From hockey scores, basketball passes and car-racing lapses, to swimming times, boxers pushes and chess positions, all the data are stored. Commentators and journalists are using this information for reporting, but trainers and athletes would want to exploit this data to improve performance and better understand opponents.
- **Digital media:** The proliferation of cheap scanners, desktop video cameras and digital cameras is one of the causes of the explosion in digital media repositories. In addition, many radio stations, television channels and film studios are digitizing their audio and video collections to improve the management of their multimedia assets. Associations such as the NHL and the NBA have already started converting their huge game collection into digital forms.
- **CAD and Software engineering data:** There are a multitude of Computer Assisted Design (CAD) systems for architects to design buildings or engineers to conceive system components or circuits. These systems are generating a tremendous amount of data. Moreover, software engineering is a source of considerable similar data with code, function libraries, objects, etc., which need powerful tools for management and maintenance.
- **Virtual Worlds:** There are many applications making use of three-dimensional virtual spaces. These spaces and the objects they contain are described with special languages such as VRML. Ideally, these virtual spaces are described in such a way that they can share objects and places. There is a remarkable amount of virtual reality object and space repositories available. Management of these repositories as well as content-based search and retrieval from these repositories are still research issues, while the size of the collections continues to grow.
- **Text reports and memos (e-mail messages):** Most of the communications within and between companies or research organizations or even private people, are based on reports and

memos in textual forms often exchanged by e-mail. These messages are regularly stored in digital form for future use and reference creating formidable digital libraries.

- **The World Wide Web repositories:** Since the inception of the World Wide Web in 1993, documents of all sorts of formats, content and description have been collected and interconnected with hyperlinks making it the largest repository of data ever built. Despite its dynamic and unstructured nature, its heterogeneous characteristic, and its very often redundancy and inconsistency, the World Wide Web is the most important data collection regularly used for reference because of the broad variety of topics covered and the infinite contributions of resources and publishers. Many believe that the World Wide Web will become the compilation of human knowledge.

9.1.2 Data Mining and Knowledge Discovery

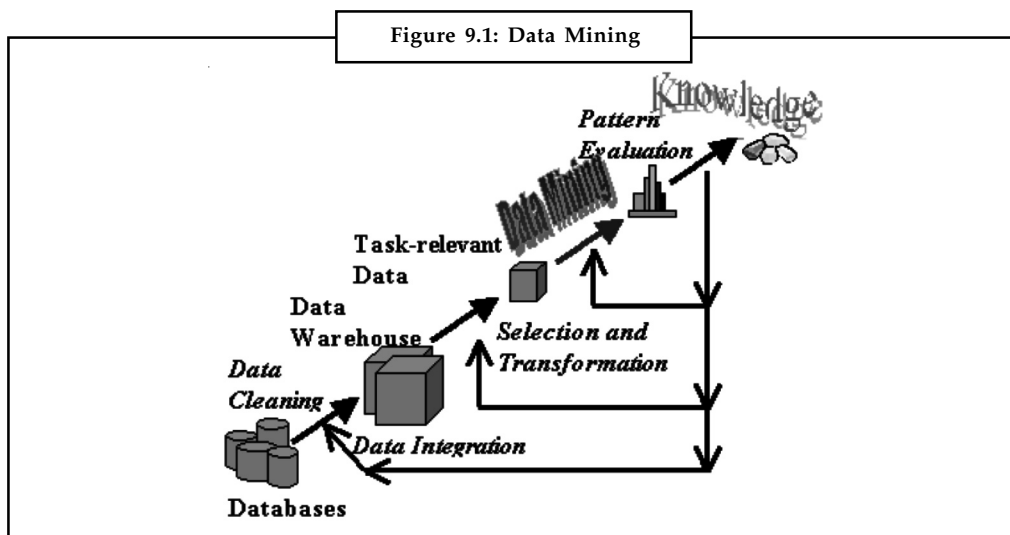
With the enormous amount of data stored in files, databases, and other repositories, it is increasingly important, if not necessary, to develop powerful means for analysis and perhaps interpretation of such data and for the extraction of interesting knowledge that could help in decision-making.

Data Mining, also popularly known as *Knowledge Discovery in Databases (KDD)*, refers to the nontrivial extraction of implicit, previously unknown and potentially useful information from data in databases.



Did u know? While data mining and knowledge discovery in databases (or KDD) are frequently treated as synonyms, data mining is actually part of the knowledge discovery process.

The figure 9.1 shows data mining as a step in an iterative knowledge discovery process.



Source: <http://webdocs.cs.ualberta.ca/~zaiane/courses/cmput690/notes/Chapter1/>

The Knowledge Discovery in Databases process comprises of a few steps leading from raw data collections to some form of new knowledge. The iterative process consists of the following steps:

- **Data cleaning:** It is also known as data cleansing, it is a phase in which noise data and irrelevant data are removed from the collection.

Notes

- **Data integration:** At this stage, multiple data sources, often heterogeneous, may be combined in a common source.
- **Data selection:** At this step, the data relevant to the analysis is decided on and retrieved from the data collection.
- **Data transformation:** It is also known as data consolidation, it is a phase in which the selected data is transformed into forms appropriate for the mining procedure.
- **Data mining:** It is the crucial step in which clever techniques are applied to extract patterns potentially useful.
- **Pattern evaluation:** In this step, strictly interesting patterns representing knowledge are identified based on given measures.
- **Knowledge representation:** It is the final phase in which the discovered knowledge is visually represented to the user. This essential step uses visualization techniques to help users understand and interpret the data mining results.

It is common to combine some of these steps together. For instance, *data cleaning* and *data integration* can be performed together as a pre-processing phase to generate a data warehouse. *Data selection* and *data transformation* can also be combined where the consolidation of the data is the result of the selection, or, as for the case of data warehouses, the selection is done on transformed data.

The KDD is an iterative process. Once the discovered knowledge is presented to the user, the evaluation measures can be enhanced, the mining can be further refined, new data can be selected or further transformed, or new data sources can be integrated, in order to get different, more appropriate results.

Data mining derives its name from the similarities between searching for valuable information in a large database and mining rocks for a vein of valuable ore. Both imply either sifting through a large amount of material or ingeniously probing the material to exactly pinpoint where the values reside. It is, however, a misnomer, since mining for gold in rocks is usually called "gold mining" and not "rock mining", thus by analogy, data mining should have been called "knowledge mining" instead. Nevertheless, data mining became the accepted customary term, and very rapidly a trend that even overshadowed more general terms such as knowledge discovery in databases (KDD) that describe a more complete process.



Caution Other similar terms referring to data mining are: data dredging, knowledge extraction and pattern discovery.

9.1.3 Types of Data

In principle, data mining is not specific to one type of media or data. Data mining should be applicable to any kind of information repository. However, algorithms and approaches may differ when applied to different types of data. Indeed, the challenges presented by different types of data vary significantly. Data mining is being put into use and studied for databases, including relational databases, object-relational databases and object-oriented databases, data warehouses, transactional databases, unstructured and semi-structured repositories such as the World Wide Web, advanced databases such as spatial databases, multimedia databases, time-series databases and textual databases, and even flat files. Here are some examples in more detail:

- **Flat files:** Flat files are actually the most common data source for data mining algorithms, especially at the research level. Flat files are simple data files in text or binary format with a structure known by the data mining algorithm to be applied. The data in these files can be transactions, time-series data, scientific measurements, etc.
- **Relational Databases:** Briefly, a relational database consists of a set of tables containing either values of entity attributes, or values of attributes from entity relationships. Tables have columns and rows, where columns represent attributes and rows represent tuples. A tuple in a relational table corresponds to either an object or a relationship between objects and is identified by a set of attribute values representing a unique key.



Example: In Figure 9.2 we present some relations Customer, Items, and Borrow representing business activity in a fictitious video store OurVideoStore. These relations are just a subset of what could be a database for the video store and is given as an example.

Figure 9.2: Fragments of Some Relations from a Relational Database for OurVideoStore

| Customer | | | | | | | |
|------------|------------|-----------------|----------|-----------|---------------|---------|-----|
| customerID | name | address | password | birthdate | family_income | group | ... |
| C1234 | John Smith | 120 main street | | Marty | 1965/10/10 | \$45000 | |
| ... | | | | | | | |

| Borrow | | | |
|------------|----------|--------|---|
| customerID | date | itemID | # |
| C1234 | 99/09/06 | 98765 | |
| ... | | | |

| Items | | | | | | | |
|--------|-------|---------|-------|----------|---------|---|-----|
| itemID | type | title | media | category | Value | # | ... |
| 98765 | Video | Titanic | DVD | Drama | \$15.00 | 2 | |
| ... | | | | | | | |

Source: <http://webdocs.cs.ualberta.ca/~zaiane/courses/cmput690/notes/Chapter1/>

The most commonly used query language for relational database is SQL, which allows retrieval and manipulation of the data stored in the tables, as well as the calculation of aggregate functions such as average, sum, min, max and count. For instance, an SQL query to select the videos grouped by category would be:

```
SELECT count(*) FROM Items WHERE type=video GROUP BY category.
```

Data mining algorithms using relational databases can be more versatile than data mining algorithms specifically written for flat files, since they can take advantage of the structure inherent to relational databases. While data mining can benefit from SQL for data selection, transformation and consolidation, it goes beyond what SQL could provide, such as predicting, comparing, detecting deviations, etc.

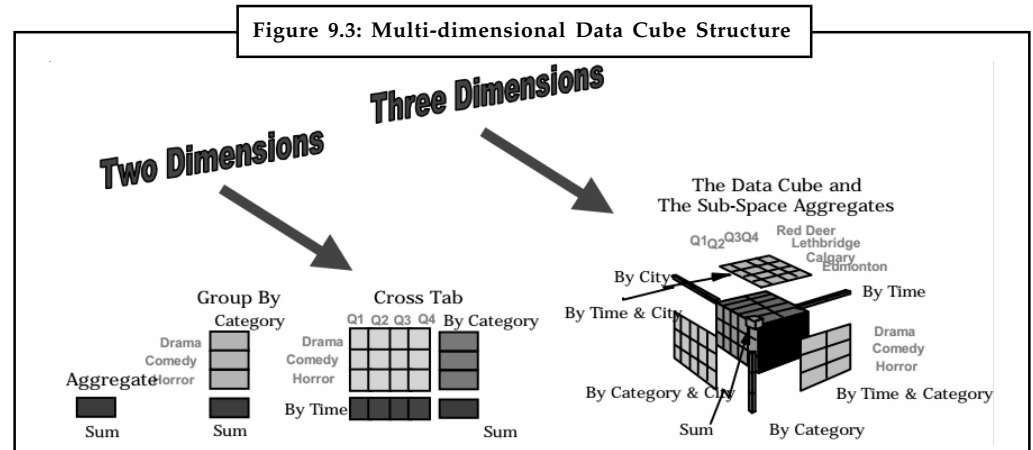
- **Data Warehouses:** A data warehouse as a storehouse, is a repository of data collected from multiple data sources (often heterogeneous) and is intended to be used as a whole under the same unified schema. A data warehouse gives the option to analyze data from different sources under the same roof. Let us suppose that OurVideoStore becomes a franchise in North America. Many video stores belonging to OurVideoStore company may have different databases and different structures. If the executive of the company wants to access the data from all stores for strategic decision-making, future direction, marketing,

Notes

etc., it would be more appropriate to store all the data in one site with a homogeneous structure that allows interactive analysis. In other words, data from the different stores would be loaded, cleaned, transformed and integrated together. To facilitate decision-making and multi-dimensional views, data warehouses are usually modeled by a multi-dimensional data structure.



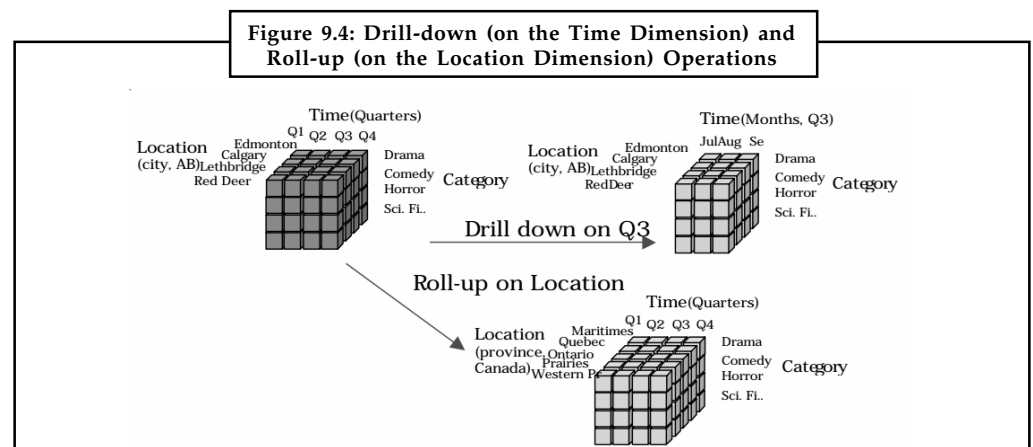
Example: Figure 9.3 shows an example of a three dimensional subset of a data cube structure used for OurVideoStore data warehouse.



Source: <http://webdocs.cs.ualberta.ca/~zaiane/courses/cmput690/notes/Chapter1/>

The figure shows summarized rentals grouped by film categories, then a cross table of summarized rentals by film categories and time (in quarters). The data cube gives the summarized rentals along three dimensions: category, time and city. A cube contains cells that store values of some aggregate measures (in this case rental counts), and special cells that store summations along dimensions. Each dimension of the data cube contains a hierarchy of values for one attribute.

Because of their structure, the pre-computed summarized data they contain and the hierarchical attribute values of their dimensions, data cubes are well suited for fast interactive querying and analysis of data at different conceptual levels, known as On-Line Analytical Processing (OLAP). OLAP operations allow the navigation of data at different levels of abstraction, such as drill-down, roll-up, slice, dice, etc. Figure 9.4 illustrates the drill-down (on the time dimension) and roll-up (on the location dimension) operations.



Source: <http://webdocs.cs.ualberta.ca/~zaiane/courses/cmput690/notes/Chapter1/>

- **Transaction Databases:** A transaction database is a set of records representing transactions, each with a time stamp, an identifier and a set of items. Associated with the transaction files could also be descriptive data for the items.



Example: In the case of the video store, the rentals table such as shown in Figure 9.5, represents the transaction database. Each record is a rental contract with a customer identifier, a date, and the list of items rented (i.e. video tapes, games, VCR, etc.).

Since relational databases do not allow nested tables (i.e. a set as attribute value), transactions are usually stored in flat files or stored in two normalized transaction tables, one for the transactions and one for the transaction items. One typical data mining analysis on such data is the so-called market basket analysis or association rules in which associations between items occurring together or in sequence are studied.

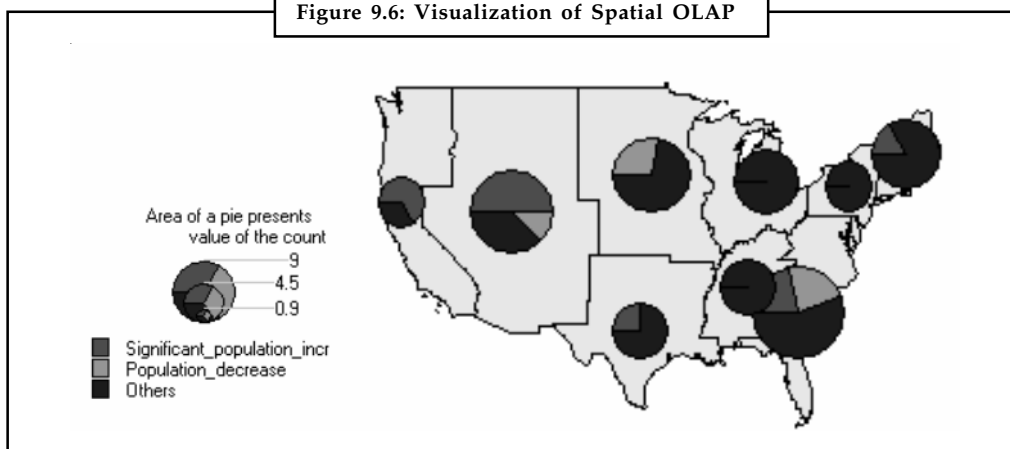
Figure 9.5: Fragment of a Transaction Database for the Rentals at OurVideoStore

| Rentals | | | | |
|---------------|----------|-------|------------|------------------------|
| transactionID | date | time | customerID | itemList |
| T12345 | 99/09/06 | 19:38 | C1234 | {I2, I6, I10, I45 ...} |
| ... | | | | |

Source: <http://webdocs.cs.ualberta.ca/~zaiane/courses/cmput690/notes/Chapter1/>

- **Multimedia Databases:** Multimedia databases include video, images, audio and text media. They can be stored on extended object-relational or object-oriented databases, or simply on a file system. Multimedia is characterized by its high dimensionality, which makes data mining even more challenging. Data mining from multimedia repositories may require computer vision, computer graphics, image interpretation, and natural language processing methodologies.
- **Spatial Databases:** Spatial databases are databases that, in addition to usual data, store geographical information like maps, and global or regional positioning. Such spatial databases present new challenges to data mining algorithms.

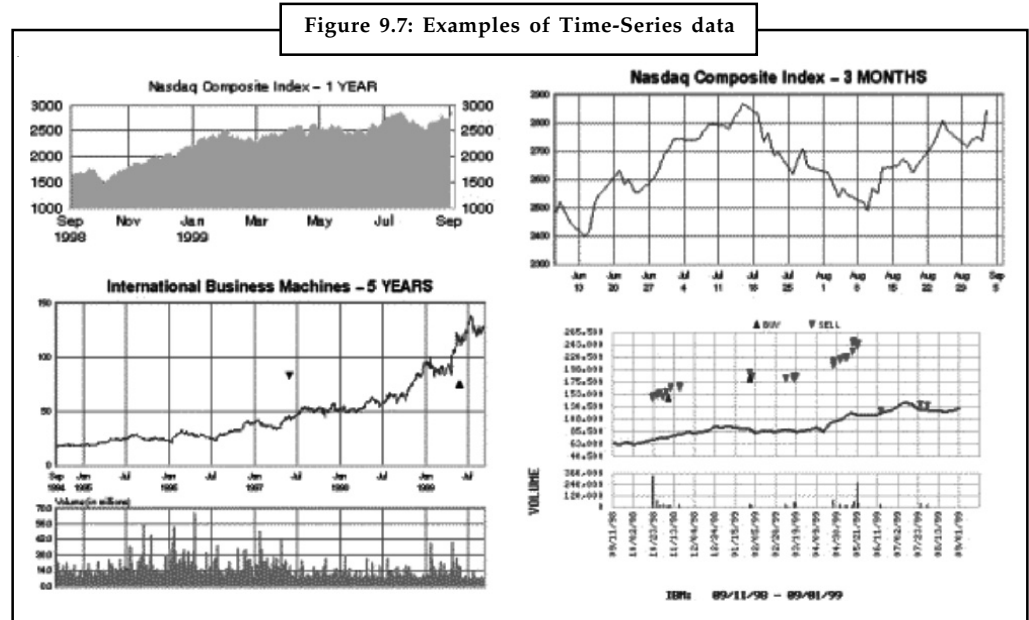
Figure 9.6: Visualization of Spatial OLAP



Source: <http://webdocs.cs.ualberta.ca/~zaiane/courses/cmput690/notes/Chapter1/>

Notes

- Time-Series Databases:** Time-series databases contain time related data such stock market data or logged activities. These databases usually have a continuous flow of new data coming in, which sometimes causes the need for a challenging real time analysis. Data mining in such databases commonly includes the study of trends and correlations between evolutions of different variables, as well as the prediction of trends and movements of the variables in time. Figure 9.7 shows some examples of time-series data.



Source: <http://webdocs.cs.ualberta.ca/~zaiane/courses/cmput690/notes/Chapter1/>

- World Wide Web:** The World Wide Web is the most heterogeneous and dynamic repository available. A very large number of authors and publishers are continuously contributing to its growth and metamorphosis, and a massive number of users are accessing its resources daily.



Caution Data in the World Wide Web is organized in inter-connected documents. These documents can be text, audio, video, raw data, and even applications.

Conceptually, the World Wide Web is comprised of three major components: the content of the Web, which encompasses documents available; the structure of the Web, which covers the hyperlinks and the relationships between documents; and the usage of the web, describing how and when the resources are accessed. A fourth dimension can be added relating the dynamic nature or evolution of the documents. Data mining in the World Wide Web, or web mining, tries to address all these issues and is often divided into web content mining, web structure mining and web usage mining.

9.1.4 Data Mining Functionalities

The kinds of patterns that can be discovered depend upon the data mining tasks employed. By and large, there are two types of data mining tasks: descriptive data mining tasks that describe the general properties of the existing data, and predictive data mining tasks that attempt to do predictions based on inference on available data.

The data mining functionalities and the variety of knowledge they discover are briefly presented in the following list:

- **Characterization:** Data characterization is a summarization of general features of objects in a target class, and produces what is called *characteristic rules*. The data relevant to a user-specified class are normally retrieved by a database query and run through a summarization module to extract the essence of the data at different levels of abstractions.



Example: One may want to characterize the OurVideoStore customers who regularly rent more than 30 movies a year. With concept hierarchies on the attributes describing the target class, the *attribute-oriented induction n method* can be used, for example, to carry out data summarization. Note that with a data cube containing summarization of data, simple OLAP operations fit the purpose of data characterization.

- **Discrimination:** Data discrimination produces what are called *discriminant rules* and is basically the comparison of the general features of objects between two classes referred to as the *target class* and the *contrasting class*.



Example: One may want to compare the general characteristics of the customers who rented more than 30 movies in the last year with those whose rental account is lower than 5.

The techniques used for data discrimination are very similar to the techniques used for data characterization with the exception that data discrimination results include comparative measures.

- **Association analysis:** Association analysis is the discovery of what are commonly called *association rules*. It studies the frequency of items occurring together in transactional databases, and based on a threshold called *support*, identifies the frequent item sets. Another threshold, *confidence*, which is the conditional probability than an item appears in a transaction when another item appears, is used to pinpoint association rules. Association analysis is commonly used for market basket analysis.



Example: It could be useful for the OurVideoStore manager to know what movies are often rented together or if there is a relationship between renting a certain type of movies and buying popcorn or pop.

The discovered association rules are of the form: $P \rightarrow Q [s,c]$, where P and Q are conjunctions of attribute value-pairs, and s (for support) is the probability that P and Q appear together in a transaction and c (for confidence) is the conditional probability that Q appears in a transaction when P is present.



Example: The hypothetical association rule: $RentType(X, "game") \text{ AND } Age(X, "13-19") \rightarrow Buys(X, "pop") [s=2\%, c=55\%]$ would indicate that 2% of the transactions considered are of customers aged between 13 and 19 who are renting a game and buying a pop, and that there is a certainty of 55% that teenage customers who rent a game also buy pop.

- **Classification:** Classification analysis is the organization of data in given classes. Also known a supervised *classification*, the classification uses given class labels to order the objects in the data collection. Classification approaches normally use a *training set* where all objects are already associated with known class labels. The classification algorithm learns from the training set and builds a model. The model is used to classify new objects.

Notes



Example: After starting a credit policy, the OurVideoStore managers could analyze the customers behaviours vis-a-vis their credit, and label accordingly the customers who received credits with three possible labels “safe”, “risky” and “very risky”.

The classification analysis would generate a model that could be used to either accept or reject credit requests in the future.

- **Prediction:** Prediction has attracted considerable attention given the potential implications of successful forecasting in a business context. There are two major types of predictions: one can either try to predict some unavailable data values or pending trends, or predict a class label for some data. The latter is tied to classification. Once a classification model is built based on a training set, the class label of an object can be foreseen based on the attribute values of the object and the attribute values of the classes. Prediction is however more often referred to the forecast of missing numerical values, or increase/decrease trends in time related data. The major idea is to use a large number of past values to consider probable future values.
- **Clustering:** Similar to classification, clustering is the organization of data in classes. However, unlike classification, in clustering, class labels are unknown and it is up to the clustering algorithm to discover acceptable classes. Clustering is also called *unsupervised classification*, because the classification is not dictated by given class labels. There are many clustering approaches all based on the principle of maximizing the similarity between objects in a same class (*intra-class similarity*) and minimizing the similarity between objects of different classes (*inter-class similarity*).
- **Outlier analysis:** Outliers are data elements that cannot be grouped in a given class or cluster. Also known as *exceptions* or *surprises*, they are often very important to identify. While outliers can be considered noise and discarded in some applications, they can reveal important knowledge in other domains, and thus can be very significant and their analysis valuable.
- **Evolution and deviation analysis:** Evolution and deviation analysis pertain to the study of time related data that changes in time. Evolution analysis models evolutionary trends in data, which consent to characterizing, comparing, classifying or clustering of time related data. Deviation analysis, on the other hand, considers differences between measured values and expected values, and attempts to find the cause of the deviations from the anticipated values.

It is common that users do not have a clear idea of the kind of patterns they can discover or need to discover from the data at hand. It is therefore important to have a versatile and inclusive data mining system that allows the discovery of different kinds of knowledge and at different levels of abstraction. This also makes interactivity an important attribute of a data mining system.



Task Compare and contrast characterization and discrimination.

9.1.5 Working of Data Mining

How exactly is data mining able to tell you important things that you didn't know or what is going to happen next? The technique that is used to perform these feats in data mining is called modeling. Modeling is simply the act of building a model in one situation where you know the answer and then applying it to another situation that you don't. For instance, if you were

looking for a sunken Spanish galleon on the high seas the first thing you might do is to research the times when Spanish treasure had been found by others in the past.

Notes



Notes You might note that these ships often tend to be found off the coast of Bermuda and that there are certain characteristics to the ocean currents, and certain routes that have likely been taken by the ship's captains in that era. You note these similarities and build a model that includes the characteristics that are common to the locations of these sunken treasures.

With these models in hand you sail off looking for treasure where your model indicates it most likely might be given a similar situation in the past. Hopefully, if you've got a good model, you find your treasure.

This act of model building is thus something that people have been doing for a long time, certainly before the advent of computers or data mining technology. What happens on computers, however, is not much different than the way people build models. Computers are loaded up with lots of information about a variety of situations where an answer is known and then the data mining software on the computer must run through that data and distill the characteristics of the data that should go into the model. Once the model is built it can then be used in similar situations where you don't know the answer.



Example: Say that you are the director of marketing for a telecommunications company and you'd like to acquire some new long distance phone customers. You could just randomly go out and mail coupons to the general population – just as you could randomly sail the seas looking for sunken treasure. In neither case would you achieve the results you desired and of course you have the opportunity to do much better than random – you could use your business experience stored in your database to build a model.

As the marketing director you have access to a lot of information about all of your customers: their age, sex, credit history and long distance calling usage. The good news is that you also have a lot of information about your prospective customers: their age, sex, credit history etc. Your problem is that you don't know the long distance calling usage of these prospects (since they are most likely now customers of your competition). You'd like to concentrate on those prospects who have large amounts of long distance usage. You can accomplish this by building a model. Table 9.1 illustrates the data used for building a model for new customer prospecting in a data warehouse.

Table 9.1: Data Mining for Prospecting

| | Customers | Prospects |
|--|-----------|-----------|
| General information (e.g. demographic data) | Known | Known |
| Proprietary information (e.g. customer transactions) | Known | Target |

The goal in prospecting is to make some calculated guesses about the information in the lower right hand quadrant based on the model that we build going from Customer General Information to Customer Proprietary Information. For instance, a simple model for a telecommunications company might be:

98% of my customers who make more than \$60,000/year spend more than \$80/month on long distance.

This model could then be applied to the prospect data to try to tell something about the proprietary information that this telecommunications company does not currently have access to. With this model in hand new customers can be selectively targeted.

Notes

Test marketing is an excellent source of data for this kind of modeling. Mining the results of a test market representing a broad but relatively small sample of prospects can provide a foundation for identifying good prospects in the overall market. Table 9.2 shows another common scenario for building models: predict what is going to happen in the future.

| | Yesterday | Today | Tomorrow |
|---|-----------|-------|----------|
| Static information and current plans (e.g. demographic data, marketing plans) | Known | Known | Known |
| Dynamic information (e.g. customer transactions) | Known | Known | Target |

If someone told you that he had a model that could predict customer usage how would you know if he really had a good model? The first thing you might try would be to ask him to apply his model to your customer base - where you already knew the answer. With data mining, the best way to accomplish this is by setting aside some of your data in a vault to isolate it from the mining process. Once the mining is complete, the results can be tested against the data held in the vault to confirm the model's validity. If the model works, its observations should hold for the vaulted data.

9.1.6 Categorization of Data Mining Systems

There are many data mining systems available or being developed. Some are specialized systems dedicated to a given data source or are confined to limited data mining functionalities, other are more versatile and comprehensive. Data mining systems can be categorized according to various criteria among other classification are the following:

- **Classification according to the type of data source mined:** This classification categorizes data mining systems according to the type of data handled such as spatial data, multimedia data, time-series data, text data, World Wide Web, etc.
- **Classification according to the data model drawn on:** This classification categorizes data mining systems based on the data model involved such as relational database, object-oriented database, data warehouse, transactional, etc.
- **Classification according to the kind of knowledge discovered:** This classification categorizes data mining systems based on the kind of knowledge discovered or data mining functionalities, such as characterization, discrimination, association, classification, clustering, etc. Some systems tend to be comprehensive systems offering several data mining functionalities together.
- **Classification according to mining techniques used:** Data mining systems employ and provide different techniques. This classification categorizes data mining systems according to the data analysis approach used such as machine learning, neural networks, genetic algorithms, statistics, visualization, database-oriented or data warehouse-oriented, etc. The classification can also take into account the degree of user interaction involved in the data mining process such as query-driven systems, interactive exploratory systems, or autonomous systems. A comprehensive system would provide a wide variety of data mining techniques to fit different situations and options, and offer different degrees of user interaction.

9.1.7 Issues in Data Mining

Data mining algorithms embody techniques that have sometimes existed for many years, but have only lately been applied as reliable and scalable tools that time and again outperform

older classical statistical methods. While data mining is still in its infancy, it is becoming a trend and ubiquitous. Before data mining develops into a conventional, mature and trusted discipline, many still pending issues have to be addressed. Some of these issues are discussed below.



Did u know? These issues are not exclusive and are not ordered in any way.

Security and Social Issues

Security is an important issue with any data collection that is shared and/or is intended to be used for strategic decision-making. In addition, when data is collected for customer profiling, user behaviour understanding, correlating personal data with other information, etc., large amounts of sensitive and private information about individuals or companies is gathered and stored. This becomes controversial given the confidential nature of some of this data and the potential illegal access to the information. Moreover, data mining could disclose new implicit knowledge about individuals or groups that could be against privacy policies, especially if there is potential dissemination of discovered information. Another issue that arises from this concern is the appropriate use of data mining. Due to the value of data, databases of all sorts of content are regularly sold, and because of the competitive advantage that can be attained from implicit knowledge discovered, some important information could be withheld, while other information could be widely distributed and used without control.

User Interface Issues

The knowledge discovered by data mining tools is useful as long as it is interesting, and above all understandable by the user. Good data visualization eases the interpretation of data mining results, as well as helps users better understand their needs. Many data exploratory analysis tasks are significantly facilitated by the ability to see data in an appropriate visual presentation. There are many visualization ideas and proposals for effective data graphical presentation. However, there is still much research to accomplish in order to obtain good visualization tools for large datasets that could be used to display and manipulate mined knowledge. The major issues related to user interfaces and visualization are “screen real-estate”, information rendering, and interaction. Interactivity with the data and data mining results is crucial since it provides means for the user to focus and refine the mining tasks, as well as to picture the discovered knowledge from different angles and at different conceptual levels.

Mining Methodology Issues

These issues pertain to the data mining approaches applied and their limitations. Topics such as versatility of the mining approaches, the diversity of data available, the dimensionality of the domain, the broad analysis needs (when known), the assessment of the knowledge discovered, the exploitation of background knowledge and metadata, the control and handling of noise in data, etc. are all examples that can dictate mining methodology choices. For instance, it is often desirable to have different data mining methods available since different approaches may perform differently depending upon the data at hand. Moreover, different approaches may suit and solve user’s needs differently.

Most algorithms assume the data to be noise-free. This is of course a strong assumption. Most datasets contain exceptions, invalid or incomplete information, etc., which may complicate, if not obscure, the analysis process and in many cases compromise the accuracy of the results. As a consequence, data preprocessing (data cleaning and transformation) becomes vital. It is often seen as lost time, but data cleaning, as time-consuming and frustrating as it may be, is one of the

Notes

most important phases in the knowledge discovery process. Data mining techniques should be able to handle noise in data or incomplete information.


More than the size of data, the size of the search space is even more decisive for data mining techniques. The size of the search space is often depending upon the number of dimensions in the domain space. The search space usually grows exponentially when the number of dimensions increases. This is known as the *curse of dimensionality*. This “curse” affects so badly the performance of some data mining approaches that it is becoming one of the most urgent issues to solve.

Performance Issues

Many artificial intelligence and statistical methods exist for data analysis and interpretation. However, these methods were often not designed for the very large data sets data mining is dealing with today. Terabyte sizes are common. This raises the issues of scalability and efficiency of the data mining methods when processing considerably large data. Algorithms with exponential and even medium-order polynomial complexity cannot be of practical use for data mining. Linear algorithms are usually the norm. In same theme, sampling can be used for mining instead of the whole dataset. However, concerns such as completeness and choice of samples may arise. Other topics in the issue of performance are *incremental updating*, and parallel programming. There is no doubt that parallelism can help solve the size problem if the dataset can be subdivided and the results can be merged later. Incremental updating is important for merging results from parallel mining, or updating data mining results when new data becomes available without having to re-analyze the complete dataset.

Data Source Issues

There are many issues related to the data sources, some are practical such as the diversity of data types, while others are philosophical like the data glut problem. We certainly have an excess of data since we already have more data than we can handle and we are still collecting data at an even higher rate. If the spread of database management systems has helped increase the gathering of information, the advent of data mining is certainly encouraging more data harvesting. The current practice is to collect as much data as possible now and process it, or try to process it, later. The concern is whether we are collecting the right data at the appropriate amount, whether we know what we want to do with it, and whether we distinguish between what data is important and what data is insignificant. Regarding the practical issues related to data sources, there is the subject of heterogeneous databases and the focus on diverse complex data types. We are storing different types of data in a variety of repositories. It is difficult to expect a data mining system to effectively and efficiently achieve good mining results on all kinds of data and sources. Different kinds of data and sources may require distinct algorithms and methodologies. Currently, there is a focus on relational databases and data warehouses, but other approaches need to be pioneered for other specific complex data types. A versatile data mining tool, for all sorts of data, may not be realistic. Moreover, the proliferation of heterogeneous data sources, at structural and semantic levels, poses important challenges not only to the database community but also to the data mining community.

| | |
|--|---|
|  <i>Task</i> | Make a report on various issues related to data mining. |
|--|---|

Self Assessment

Fill in the blanks:

1. Data mining is also known as in Databases.

2. is arguably a form of data mining, which automatically brings relevant messages to the surface from a chaotic sea of phishing attempts and Viagra pitches.
3. is a phase in which noise data and irrelevant data are removed from the collection.
4. is a phase in which the selected data is transformed into forms appropriate for the mining procedure.
5. is the final phase in which the discovered knowledge is visually represented to the user.
6. are actually the most common data source for data mining algorithms, especially at the research level.
7. A database consists of a set of tables containing either values of entity attributes, or values of attributes from entity relationships.
8. A database is a set of records representing transactions, each with a time stamp, an identifier and a set of items.
9. databases are databases that, in addition to usual data, store geographical information like maps, and global or regional positioning.
10. databases contain time related data such stock market data or logged activities.
11. Data characterization is a summarization of general features of objects in a target class, and produces what is called
12. is the discovery of what are commonly called *association rules*.
13. is the organization of data in given classes.

9.2 Applications of Data Mining

Data mining is a relatively new technology that has not fully matured. Despite this, there are a number of industries that are already using it on a regular basis. Some of these organizations include retail stores, hospitals, banks and insurance companies.

Many of these organizations are combining data mining with such things as statistics, pattern recognition, and other important tools. Data mining can be used to find patterns and connections that would otherwise be difficult to find. This technology is popular with many businesses because it allows them to learn more about their customers and make smart marketing decisions.

There are a number of applications that data mining has. The first is called market segmentation. With market segmentation, you will be able to find behaviors that are common among your customers. You can look for patterns among customers that seem to purchase the same products at the same time. Another application of data mining is called customer churn. Customer churn will allow you to estimate which customers are the most likely to stop purchasing your products or services and go to one of your competitors. In addition to this, a company can use data mining to find out which purchases are the most likely to be fraudulent.



Example: By using data mining a retail store may be able to determine which products are stolen the most. By finding out which products are stolen the most, steps can be taken to protect those products and detect those who are stealing them.

While direct mail marketing is an older technique that has been used for many years, companies who combine it with data mining can experience fantastic results.

Notes



Example: You can use data mining to find out which customers will respond favorably to a direct mail marketing strategy. You can also use data mining to determine the effectiveness of interactive marketing. Some of your customers will be more likely to purchase your products online than off-line, and you must identify them.

While many businesses use data mining to help increase their profits, many of them don't realize that it can be used to create new businesses and industries. One industry that can be created by data mining is the automatic prediction of both behaviors and trends. Imagine for a moment that you were the owner of a fashion company, and you were able to precisely predict the next big fashion trend based on the behavior and shopping patterns of your customers? It is easy to see that you could become very wealthy within a short period of time. You would have an advantage over your competitors. Instead of simply guessing what the next big trend will be, you will determine it based on statistics, patterns, and logic.

Another example of automatic prediction is to use data mining to look at your past marketing strategies. Which one worked the best? Why did it work the best? Who were the customers that responded most favorably to it? Data mining will allow you to answer these questions, and once you have the answers, you will be able to avoid making any mistakes that you made in your previous marketing campaign.

Data mining can allow you to become better at what you do. It is also a powerful tool for those who deal with finances. A financial institution such as a bank can predict the number of defaults that will occur among their customers within a given period of time, and they can also predict the amount of fraud that will occur as well.

Another potential application of data mining is the automatic recognition of patterns that were not previously known. Imagine if you had a tool that could automatically search your database to look for patterns which are hidden. If you had access to this technology, you would be able to find relationships that could allow you to make strategic decisions.

Because your decisions are based on logic, you would increase the chances of being successful. While data mining is a very valuable tool, it is important to realize that it is not a panacea. Even if an automated technology should be invented, it will not guarantee the success of you or your company. However, it will tip the odds in your favor.

Two critical factors for success with data mining are: a large, well-integrated data warehouse and a well-defined understanding of the business process within which data mining is to be applied (such as customer prospecting, retention, campaign management, and so on).

Some successful application areas include:

- A pharmaceutical company can analyze its recent sales force activity and their results to improve targeting of high-value physicians and determine which marketing activities will have the greatest impact in the next few months. The data needs to include competitor market activity as well as information about the local health care systems. The results can be distributed to the sales force via a wide-area network that enables the representatives to review the recommendations from the perspective of the key attributes in the decision process. The ongoing, dynamic analysis of the data warehouse allows best practices from throughout the organization to be applied in specific sales situations.
- A credit card company can leverage its vast warehouse of customer transaction data to identify customers most likely to be interested in a new credit product. Using a small test mailing, the attributes of customers with an affinity for the product can be identified. Recent projects have indicated more than a 20-fold decrease in costs for targeted mailing campaigns over conventional approaches.

- A diversified transportation company with a large direct sales force can apply data mining to identify the best prospects for its services. Using data mining to analyze its own customer experience, this company can build a unique segmentation identifying the attributes of high-value prospects. Applying this segmentation to a general business database such as those provided by Dun & Bradstreet can yield a prioritized list of prospects by region.
- A large consumer package goods company can apply data mining to improve its sales process to retailers. Data from consumer panels, shipments, and competitor activity can be applied to understand the reasons for brand and store switching. Through this analysis, the manufacturer can select promotional strategies that best reach their target customer segments.

Each of these examples have a clear common ground. They leverage the knowledge about customers implicit in a data warehouse to reduce costs and improve the value of customer relationships. These organizations can now focus their efforts on the most important (profitable) customers and prospects, and design targeted marketing strategies to best reach them.

Self Assessment

Fill in the blanks:

14. With, you will be able to find behaviors that are common among your customers.
15. will allow you to estimate which customers are the most likely to stop purchasing your products or services and go to one of your competitors.



Case Study

Jaeger Uses Data Mining to Reduce Losses from Crime and Waste

Lamb is the most stolen item at Iceland. Thieves also like cheese, bacon and coffee. With the UK in recession, shoplifters appear to be switching their sights from alcohol, electric toothbrushes and perfume to food. Tesco, Marks & Spencer and Iceland have all reported an increase in shoplifting since the economy began to contract in the second quarter of 2008. Tesco alone caught some 43,000 would-be thieves in the first half of 2008, up 36% from the same period in 2007.

The impact of the recession on retailers is yet to be reflected in any of the major surveys of shoplifting. The Centre for Retail Research's Retail Theft Barometer only has figures up to the end of 2007. Those figures show that shrinkage – losses from crime and waste – cost retailers 1.3% of sales in 2007, down from 1.34% in 2006. Even though 2007 was the peak of the boom, the losses were still huge. Customers stole some £1.6bn and employees another £1.3bn. Suppliers took £209m fraudulently. Some £73m was lost through card fraud and another £39m through robberies or burglaries. Retailers lost £666m through waste. The systems and security guards intended to reduce losses cost £785m. The total bill was £4.6bn.

Retailers have used a variety of technologies to reduce their losses. Closed-Circuit Television (CCTV) and Electronic Article Surveillance (EAS) – the tags attached to individual items – are all visible in stores. Some retailers, however, are using a different type of technology to reduce losses data mining.

Contd....

Notes

During the summer of 2008, British clothing chain Jaeger went live with a data mining application in an attempt to identify where it was losing money. The application, which is located centrally, interrogates data held on different systems throughout the business, including both the head office and the company's 90-plus stores and concessions.

In common with every other data mining application, Jaeger's LossManager system, from IDM Software, uses a feed from the company's electronic point of sale (Epos) system to spot potential fraud such as excessive discounting by a single member of staff.

"It's a centralised system, but every single store is feeding into it. We have got time and attendance feeding into it as well," says Steve Hearn, head of safety and security at Jaeger.

Jaeger had used a data mining application from SFR, a small British supplier, before it took the decision to implement the IDM system.

"We had SFR Storescan, but it had long since been defunct and we were not using it because our till architecture had changed. I started to wonder if that was the right piece of software for us. It was quite cumbersome," says Hearn.

Jaeger does not disclose its net profits because it is privately owned. However, it is a mid-sized clothing retailer lacking the colossal IT budgets available to, say, Tesco or Sainsbury's. Hearn's first choice supplier was too expensive for his budget.

"I would have gone with an IntelliQ product [another British software supplier], but for the price," he says. IDM Software is a start-up aimed at mid-sized retailers and Jaeger was its first retail customer.

Unlike CCTV and EAS, which are designed to catch thieving customers, data mining applications are supposed to catch thieving employees. "We do an awful lot of work on internal fraud," says Hearn.

"The data mining system was put in to separate the usual from the unusual," he adds.

Jaeger set up an audit team when the system went live in June. The team's job is to use the new application to identify losses wherever they occur – from dishonest employees to working practices that waste stock.

Like other data mining applications, LossManager generates exception reports. However, it would be misleading to rely solely on these reports, according to IDM's chief executive officer Khuram Kirmani.

"It's very easy to get swamped with false positives," Kirmani says.

The employees responsible for loss prevention (in Jaeger's case, the audit team) use their data mining application to generate exception reports as usual. Then they continue to use the application to ask more questions of the data so that they can understand whether the system is reporting a false positive or a genuine loss.

"Each question is based on the answer to the previous question," says David Snocken, IDM's commercial director.

Any project's success is limited by the user's willingness to extract as much value as possible.

"It depends on the amount of effort the retailer has put in," says Kirmani. IDM says its system has reduced losses as a percentage of sales below Global Retail Theft Barometer's 1.3% average for UK retailers.

Contd....

Notes

Although Jaeger has only had the system since June, it already expects a return on investment in its first financial year. Hearn says, "Data mining is widely accepted as having one of the fastest returns on investment of any technology. We are still in the early days in terms of assessing the benefits, but we are almost double-counting our results to check they are right."

One of the earliest discoveries was that theft by employees was only a small part of total losses at Jaeger.

"We have not gone out en masse and started arresting staff members for fraud, but we have identified considerable numbers of erroneous transactions. That is not to say that they are all fraud," explains Hearn.

Data mining is helping the clothing retailer to manage its stock, thereby reducing the need for markdowns when items go out of season and reducing the number of items that go missing altogether.

In a recession that has already claimed the scalps of established retailers such as Woolworths and MFI, any initiative that helps a retailer conserve cash will receive management support.

"Data mining is even more important now in terms of being able to understand margin erosion. Shrinkage is the last free margin on the table. We have got to keep the stock current," says Hearn.

At the start of the data mining project, Jaeger forecast that it would make a return on investment within six to nine months of the project going live. That target will be met. Jaeger now expects both a significant improvement in margins and a substantial benefit to its net profits.

"The sheer opportunities to improve margin - it's not just about fraud, it's about putting the wrong stock in the wrong place at the wrong time. As a result, the decision to go with data mining was very quick. I had no resistance from Jaeger," Hearn says.

In Jaeger's case, the difficulty with implementing its data mining application did not come from the management it came from the complexity of setting up data feeds between Jaeger's existing store applications and its new centralised system. The company decided to buy a data mining application in the summer of 2007.

"It was nearly a year," says Hearn. "It was nothing to do with IDM, but to do with Jaeger. Our data was very complicated because we have had so much in-house development of our systems. For instance, at just one meeting, we had to review at line level the data we used in over 800 fields."

Jaeger's data mining project will make a positive contribution to profits at the most important part of the business cycle. As the recession worsens in 2009, retailers will need to develop similar projects that produce rapid returns on investment those that make sustained improvements to net profits year after year will stand the best chance of winning management approval. As money strains lead more customers and employees to steal from retailers, applications that can reduce theft will become increasingly important.

How Data Mining Gathers Information?

A data mining application becomes more powerful if it uses a greater number of feeds from the retailer's other systems. LossManager was built in the Microsoft Development Environment and was written in C++ so it can be used to accept feeds from as many different systems as possible.

Contd....

Notes

“We can take feeds from almost anything. We can use that information to ask if there is a correlation between a store that loses a lot of product and EAS deactivations and alarms. One of the departures from previous approaches is that for an application to be truly effective, we have to integrate multiple sets of data,” says Hearn.

Several data mining applications already use video feed from CCTV cameras to make sense of Epos data. Many retailers would like to use the two technologies together, but they are unable to do so because their CCTV cameras use analogue rather than digital film. For most retailers, the cost of replacing analogue cameras with digital cameras far exceeds the financial benefits that they expect to gain from reducing their losses.

Retailers with radio frequency identification (RFID) projects could even use the information from tagged pallets or individually tagged items within their data mining applications. Unfortunately for advocates of RFID technology, the only retailer with a public RFID project in the UK is Marks & Spencer, which tags different ranges of clothing in most of its major stores.

Audit teams use a type of network theory called link analysis to understand the patterns between data on different systems. Auditors look for symmetric patterns between two sets of data, or more likely asymmetric patterns, to understand the relationships between different types of information.

Retailers are not the only organisations that use data mining to look for correlating information. Governments have used data mining to sift through huge amounts of data to identify potential terrorist attacks. In 2002, the Pentagon started a secret project called Total Information Awareness in an attempt to identify terrorists. Total Information Awareness was a data mining project on a massive scale. In 2003, it was cancelled after Congress removed funding over fears that it was too intrusive.

Question

Discuss how data mining is used to reduce losses from crime and waste.

Source: <http://www.computerweekly.com/feature/Case-study-Jaeger-uses-data-mining-to-reduce-losses-from-crime-and-waste>

9.3 Summary

- Data mining uses a relatively large amount of computing power operating on a large set of data to determine regularities and connections between data points.
- Data Mining, also popularly known as *Knowledge Discovery in Databases* (KDD), refers to the nontrivial extraction of implicit, previously unknown and potentially useful information from data in databases.
- Flat files are actually the most common data source for data mining algorithms, especially at the research level.
- A relational database consists of a set of tables containing either values of entity attributes, or values of attributes from entity relationships.
- A data warehouse as a storehouse, is a repository of data collected from multiple data sources (often heterogeneous) and is intended to be used as a whole under the same unified schema.
- A transaction database is a set of records representing transactions, each with a time stamp, an identifier and a set of items.

- Multimedia databases include video, images, audio and text media.
- There are two types of data mining tasks: descriptive data mining tasks that describe the general properties of the existing data, and predictive data mining tasks that attempt to do predictions based on inference on available data.

9.4 Keywords

Association Analysis: Association analysis is the discovery of what are commonly called *association rules*.

Data Warehouse: A data warehouse as a storehouse, is a repository of data collected from multiple data sources and is intended to be used as a whole under the same unified schema.

Flat Files: Flat files are simple data files in text or binary format with a structure known by the data mining algorithm to be applied.

KDD: Knowledge Discovery in Databases (KDD) refers to the nontrivial extraction of implicit, previously unknown and potentially useful information from data in databases.

Multimedia Database: Multimedia databases include video, images, audio and text media.

Relational Database: A relational database consists of a set of tables containing either values of entity attributes, or values of attributes from entity relationships.

Spatial Databases: Spatial databases are databases that, in addition to usual data, store geographical information like maps, and global or regional positioning.

Time-series Databases: Time-series databases contain time related data such stock market data or logged activities.

Transaction Database: A transaction database is a set of records representing transactions, each with a time stamp, an identifier and a set of items.

9.5 Review Questions

1. Explain the concept of data mining with example.
2. Discuss the different types of information collected in digital form in databases and in flat files.
3. Discuss the concept of Data Mining and Knowledge Discovery.
4. Describe the steps included in the Knowledge Discovery process.
5. Make distinction between relational database and transaction database.
6. Explain the functionalities of data mining.
7. Discuss association analysis with example.
8. Illustrate the categorization of data mining systems.
9. Describe various issues in Data Mining.
10. Explain some application areas of data mining.

Answers: Self Assessment

1. Knowledge-Discovery
2. Spam filtering

Notes

- | | |
|-----------------------------|--------------------------|
| 3. Data cleaning | 4. Data transformation |
| 5. Knowledge representation | 6. Flat files |
| 7. Relational | 8. Transaction |
| 9. Spatial | 10. Time-series |
| 11. Characteristic rules | 12. Association analysis |
| 13. Association analysis | 14. Market segmentation |
| 15. Customer churn | |

9.6 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

<http://dbdmg.polito.it/wordpress/teaching/data-mining-concepts-and-algorithms/>

http://making.csie.ndhu.edu.tw/course/2006/Fall/Data_mining/06.pdf

<http://www.cs.sfu.ca/~han/dmbook>

<http://www.laits.utexas.edu/~anorman/BUS.FOR/course.mat/Alex/>

Unit 10: Data Mining Tools and Techniques

Notes

CONTENTS

Objectives

Introduction

10.1 Data Mining Tools

10.2 Data Mining Techniques

10.2.1 Statistics

10.2.2 Nearest Neighbor

10.2.3 Clustering

10.2.4 Decision Trees

10.2.5 Neural Networks

10.2.6 Rule Induction

10.3 Text Mining

10.4 Web Mining

10.5 Summary

10.6 Keywords

10.7 Review Questions

10.8 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss Various Data Mining Tools
- Explain Data Mining Techniques such as Decision Tree, Neural Network, etc.

Introduction

Data Mining can be defined as a technique for extracting the “meaning” contained in information to allow the understanding needed by a user to make a “right” decision. It is Data Mining that allows a computer to digest the constant stream of data being generated by the computerized sensors and monitors of the plant, and then extract from that information that has some meaning content. Data mining tools and techniques can be used for rationalizing the data so as to reduce the overload that tends to occur and make it simple for the personnel to make a right decision in textile industry. In this unit, we will discuss various data mining tools and techniques.

10.1 Data Mining Tools

Data mining tools collect data and model the data to represent the reality. The model will represent and describe the data relationship and pattern. Based on orientation process, data mining activities divide into three categories which include discovery, predictive modeling and forensic analysis. Discovery is the process of finding the hidden patterns in a database without

Notes

gives idea and hypothesis on what the patterns might be. While predictive modeling is the process of using the pattern gather from the database and use the data to predict future. The third categories are the forensic analysis.



Did u know? Forensic analysis is the process of implementing the extracted patterns to determine differences or non-standardized data.

Data mining automates the process relevant patterns of current and historical data in the database to be analyzed to forecast the future. Through the ability of data mining tools to predict and analyze behaviors of data in the databases, it will be able to guide the organization to produce proactive and efficient decision making and answer question that is urgently need to be solve in a little time.

It is very difficult for data mining software companies to create tools which are geared towards businesses. The reason for this is because many of the people who are responsible for this technology will place an emphasis on computer algorithms.

To most business owners, algorithms are not important. Data mining is a technology that is now being used mostly be large corporations, and because of this, the focus cannot be algorithms. Many popular data mining programs have algorithms that only compose about 10% of their structure. The question that many developers must as themselves is where should the emphasis be placed on the other 90 percent?

The first place that data mining developers can focus on is database integration. The data mining tools that are created must be able to function with data warehouses. When the files are flat, this will not allow the tool to work with many databases, and this will cause problems. Fortunately, many data mining developers have taken this advice, and are designing their tools in a way that allows them to work seamlessly with the data warehouses of many companies. However, there are still some developers that are not doing this. The next area that is important for data mining tools is called automatic model scoring.



Caution Scoring is one of the most tedious aspects of data mining.

There are a number of contemporary data mining programs that cannot score the models that they create. If you are using any of these programs, you will have to develop your own scoring system. This is tedious, time consuming, and unnecessary. In addition to this, when you have to manually produce a scoring system, it is likely that you will have many errors. The scoring system will often have to be done by the information technology department, and they don't do it correctly, there could be a number of problems. To solve these problems, developers will want to create data mining tools that automate the process of scoring models that have been created.

By automating the process of scoring data mining models, companies could become more efficient and less prone to errors. Another area where data mining programs need to improve is exporting models between different software programs. Once a model has been generated, it is important for other programs to be able to understand it. By doing this, the process of scoring can be much more efficient, the models can be used by numerous tools. In addition to this, it is important for data mining tools to begin using more business templates. The goal of a company is to solve a business problem rather than a statistical issue. Developers will want to calibrate the data mining tools in a way that makes them more relevant to business users.

Users should also be given more control over the data mining programs they use.



Example: A user should be able to set a value which will determine the speed and detail of models that are generated.

For example, if a user needs to be able to create a rough model quickly that will give them a general idea of how to solve a problem, this should be available. However, if the customer needs a detailed model which may take an extended period of time to generate, this function should be available as well.

Data mining developers will want to allow users to set their own parameters.

Because most businesses will naturally want to deal with financial issues, this should be integrated into data mining programs. The goal of a business owner is to increase their profits rather than deal with technical issues which are not relevant for their operation. Because of this, finance issues should play an important role in the development of data mining tools.

There are various types of data mining available in the market. Each tool comes with its own advantage and weaknesses. Information personal have to keep update with the different type of data mining tools and suggest to purchase the right tools that support the best need of the organization. Data mining tools can be classified in to three main categories which is dashboard, text mining tools and traditional data mining tools.

- Traditional data mining tools use complex algorithms and technique to establish data trends and patterns. To monitor data, trends and captures information that not in the database, these tools should be installed in the desktop. Most of the tools are compatible with both Windows and UNIX version.
- The second categories of data mining are dashboard. Basically organization will install these tools to monitor the data changes, information contained in the database and onscreen update. Basically this tool comes in the form of table and chart to allow the user to get better seeing of the business performance. Beside that, dashboard also allow user to refer historical data. This will enable user to find changes on the data. Beside easy to use, this function makes dashboard interesting and easier for the manager to view the company's overall performance.
- Text mining tools is the third type of data mining. This tool has the ability to mine data in various kind of text such as Microsoft words and acrobat PDF. The ability of this tools to scan and convert data into the right format that suitable with the tool's database has brings easy and convenient data access to the user. By the use of this tool, user does not need to open different application for every different data format. The data scanned may contain structured or unstructured data. This input captured will gives organization a wealth of information which can be mined to determine attitudes, trend and concept. The origins of data mining began on the first storage of data in the computers and continue with the progress in data access, until nowadays technology that allows users to browse through data in actual time.

Best way in applying advanced data mining techniques is should have interactive and flexible data mining tools which is directly integrated with the organization's data warehouse. It is the best practice to integrate data mining to data warehouse. This allows organization to simplify the application and mining result implementation. Besides, if the data warehouse grows larger, organization can mine best practice continually and apply for the future decision making. In contrast, with using outside mining tools that is not efficient and time consuming where by, few extra mining steps are required. The techniques are discussed in the next section.

Notes

Self Assessment

Fill in the blanks:

1. Data mining collect data and model the data to represent the reality.
2. is used to monitor the data changes, information contained in the database and onscreen update.
3. tool has the ability to mine data in various kind of text such as Microsoft words and acrobat PDF.
4. is the process of implementing the extracted patterns to determine differences or non-standardized data.

10.2 Data Mining Techniques

These techniques have been divided into classical techniques and next generation techniques. This division is based on when the data mining technique was developed and when it became technically mature enough to be used for business, especially for aiding in the optimization of customer relationship management systems.

- **Classical Techniques:** Statistics, Neighborhoods and Clustering are the techniques that have classically been used for decades the next section represents techniques that have only been widely used since the early 1980s. The main techniques that we will discuss here are the ones that are used 99.9% of the time on existing business problems. There are certainly many other ones as well as proprietary techniques from particular vendors – but in general the industry is converging to those techniques that work consistently and are understandable and explainable.
- **Next Generation Techniques:** Trees, Networks and Rules are data mining techniques represent the most often used techniques that have been developed over the last two decades of research. They also represent the vast majority of the techniques that are being spoken about when data mining is mentioned in the popular press. These techniques can be used for either discovering new information within large databases or for building predictive models. Though the older decision tree techniques such as CHAID are currently highly used the new techniques such as CART are gaining wider acceptance.

10.2.1 Statistics

By strict definition “statistics” or statistical techniques are not data mining. They were being used long before the term data mining was coined to apply to business applications. However, statistical techniques are driven by the data and are used to discover patterns and build predictive models. And from the users perspective you will be faced with a conscious choice when solving a “data mining” problem as to whether you wish to attack it with statistical methods or other data mining techniques.

Statistics is a branch of mathematics concerning the collection and the description of data. Usually statistics is considered to be one of those scary topics in college right up there with chemistry and physics. However, statistics is probably a much friendlier branch of mathematics because it really can be used every day. Statistics was in fact born from very humble beginnings of real world problems from business, biology, and gambling!

Knowing statistics in your everyday life will help the average business person make better decisions by allowing them to figure out risk and uncertainty when all the facts either aren’t known or can’t be collected. Even with all the data stored in the largest of data warehouses

business decisions still just become more informed guesses. The more and better the data and the better the understanding of statistics the better the decision that can be made.

Statistics has been around for a long time easily a century and arguably many centuries when the ideas of probability began to gel. It could even be argued that the data collected by the ancient Egyptians, Babylonians, and Greeks were all statistics long before the field was officially recognized. Today data mining has been defined independently of statistics though “mining data” for patterns and predictions is really what statistics is all about.



Notes Some of the techniques that are classified under data mining such as CHAID and CART really grew out of the statistical profession more than anywhere else, and the basic ideas of probability, independence and causality and overfitting are the foundation on which both data mining and statistics are built.

10.2.2 Nearest Neighbor

Clustering and the Nearest Neighbor prediction technique are among the oldest techniques used in data mining. Most people have an intuition that they understand what clustering is – namely that like records are grouped or clustered together. Nearest neighbor is a prediction technique that is quite similar to clustering – its essence is that in order to predict what a prediction value is in one record look for records with similar predictor values in the historical database and use the prediction value from the record that it “nearest” to the unclassified record.

A Simple Example of Clustering

A simple example of clustering would be the clustering that most people perform when they do the laundry – grouping the permanent press, dry cleaning, whites and brightly colored clothes is important because they have similar characteristics. And it turns out they have important attributes in common about the way they behave (and can be ruined) in the wash. To “cluster” your laundry most of your decisions are relatively straightforward. There are of course difficult decisions to be made about which cluster your white shirt with red stripes goes into (since it is mostly white but has some color and is permanent press). When clustering is used in business the clusters are often much more dynamic – even changing weekly to monthly and many more of the decisions concerning which cluster a record falls into can be difficult.

A Simple Example of Nearest Neighbor

A simple example of the nearest neighbor prediction algorithm is that if you look at the people in your neighborhood (in this case those people that are in fact geographically near to you). You may notice that, in general, you all have somewhat similar incomes. Thus if your neighbor has an income greater than \$100,000 chances are good that you too have a high income. Certainly the chances that you have a high income are greater when all of your neighbors have incomes over \$100,000 than if all of your neighbors have incomes of \$20,000. Within your neighborhood there may still be a wide variety of incomes possible among even your “closest” neighbors but if you had to predict someone’s income based on only knowing their neighbors you’re best chance of being right would be to predict the incomes of the neighbors who live closest to the unknown person.

The nearest neighbor prediction algorithm works in very much the same way except that “nearness” in a database may consist of a variety of factors not just where the person lives. It may, for instance, be far more important to know which school someone attended and what

Notes

degree they attained when predicting income. The better definition of “near” might in fact be other people that you graduated from college with rather than the people that you live next to.

Nearest Neighbor techniques are among the easiest to use and understand because they work in a way similar to the way that people think - by detecting closely matching examples. They also perform quite well in terms of automation, as many of the algorithms are robust with respect to dirty data and missing data. Lastly they are particularly adept at performing complex ROI calculations because the predictions are made at a local level where business simulations could be performed in order to optimize ROI. As they enjoy similar levels of accuracy compared to other techniques the measures of accuracy such as lift are as good as from any other.

How to Use Nearest Neighbor for Prediction

One of the essential elements underlying the concept of clustering is that one particular object (whether they be cars, food or customers) can be closer to another object than can some third object. It is interesting that most people have an innate sense of ordering placed on a variety of different objects. Most people would agree that an apple is closer to an orange than it is to a tomato and that a Toyota Corolla is closer to a Honda Civic than to a Porsche. This sense of ordering on many different objects helps us place them in time and space and to make sense of the world. It is what allows us to build clusters – both in databases on computers as well as in our daily lives. This definition of nearness that seems to be ubiquitous also allows us to make predictions.

The nearest neighbor prediction algorithm simply stated is:

Objects that are “near” to each other will have similar prediction values as well. Thus if you know the prediction value of one of the objects you can predict it for it’s nearest neighbors.

Where has the Nearest Neighbor Technique been Used in Business?

One of the classical places that nearest neighbor has been used for prediction has been in text retrieval. The problem to be solved in text retrieval is one where the end user defines a document that is interesting to them and they solicit the system to “find more documents like this one”. Effectively defining a target of: “this is the interesting document” or “this is not interesting”. The prediction problem is that only a very few of the documents in the database actually have values for this prediction field (namely only the documents that the reader has had a chance to look at so far). The nearest neighbor technique is used to find other documents that share important characteristics with those documents that have been marked as interesting.

Using Nearest Neighbor for Stock Market Data

As with almost all prediction algorithms, nearest neighbor can be used in a variety of places. Its successful use is mostly dependent on the pre-formatting of the data so that nearness can be calculated and where individual records can be defined. In the text retrieval example this was not too difficult - the objects were documents. This is not always as easy as it is for text retrieval. Consider what it might be like in a time series problem - say for predicting the stock market. In this case the input data is just a long series of stock prices over time without any particular record that could be considered to be an object. The value to be predicted is just the next value of the stock price.

The way that this problem is solved for both nearest neighbor techniques and for some other types of prediction algorithms is to create training records by taking, for instance, 10 consecutive stock prices and using the first 9 as predictor values and the 10th as the prediction value. Doing things this way, if you had 100 data points in your time series you could create 10 different training records.

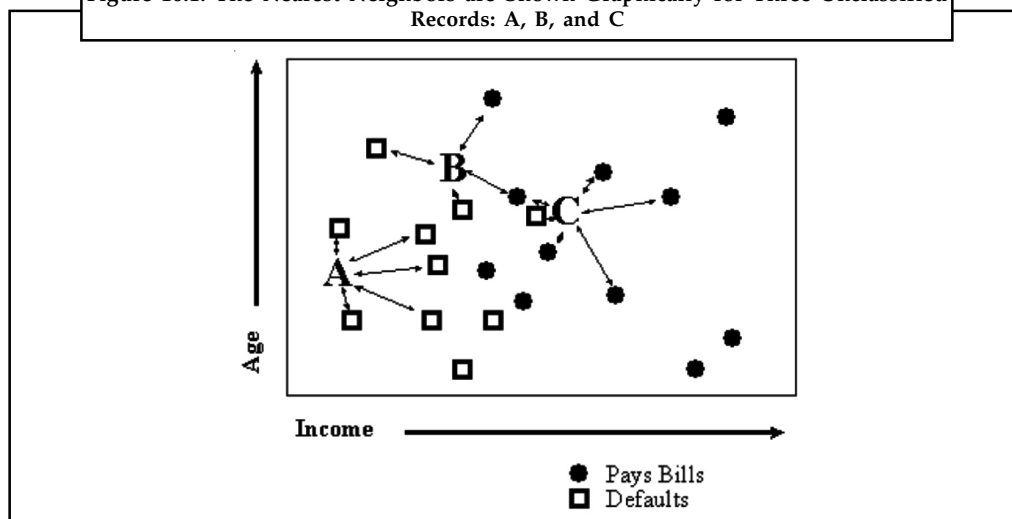
You could create even more training records than 10 by creating a new record starting at every data point. For instance you could take the first 10 data points and create a record. Then you could take the 10 consecutive data points starting at the second data point, then the 10 consecutive data point starting at the third data point. Even though some of the data points would overlap from one record to the next the prediction value would always be different. In our example of 100 initial data points 90 different training records could be created this way as opposed to the 10 training records created via the other method.

Why Voting is Better – K Nearest Neighbors

One of the improvements that is usually made to the basic nearest neighbor algorithm is to take a vote from the “K” nearest neighbors rather than just relying on the sole nearest neighbor to the unclassified record.

Example: In Figure 10.1 we can see that unclassified example C has a nearest neighbor that is a defaulter and yet is surrounded almost exclusively by records that are good credit risks. In this case the nearest neighbor to record C is probably an outlier - which may be incorrect data or some non-repeatable idiosyncrasy. In either case it is more than likely that C is a non-defaulter yet would be predicted to be a defaulter if the sole nearest neighbor were used for the prediction.

Figure 10.1: The Nearest Neighbors are Shown Graphically for Three Unclassified Records: A, B, and C



In cases like these a vote of the 9 or 15 nearest neighbors would provide a better prediction accuracy for the system than would just the single nearest neighbor. Usually this is accomplished by simply taking the majority or plurality of predictions from the K nearest neighbors if the prediction column is a binary or categorical or taking the average value of the prediction column from the K nearest neighbors.

Another important aspect of any system that is used to make predictions is that the user be provided with, not only the prediction, but also some sense of the confidence in that prediction (e.g. the prediction is defaulter with the chance of being correct 60% of the time). The nearest neighbor algorithm provides this confidence information in a number of ways.

The distance to the nearest neighbor provides a level of confidence. If the neighbor is very close or an exact match then there is much higher confidence in the prediction than if the nearest record is a great distance from the unclassified record.

The degree of homogeneity amongst the predictions within the K nearest neighbors can also be used. If all the nearest neighbors make the same prediction then there is much higher confidence

Notes

in the prediction than if half the records made one prediction and the other half made another prediction.

10.2.3 Clustering

Clustering is the method by which like records are grouped together. Usually this is done to give the end user a high level view of what is going on in the database. Clustering is sometimes used to mean segmentation – which most marketing people will tell you is useful for coming up with a birds eye view of the business. Two of these clustering systems are the PRIZM™ system from Claritas corporation and MicroVision™ from Equifax corporation. These companies have grouped the population by demographic information into segments that they believe are useful for direct marketing and sales. To build these groupings they use information such as income, age, occupation, housing and race collect in the US Census. Then they assign memorable “nicknames” to the clusters. Some examples are shown in Table 10.1.


Table 10.1: Some Commercially Available Cluster Tags

| Name | Income | Age | Education | Vendor |
|----------------------|-------------|---------------------|----------------|-------------------------|
| Blue Blood Estates | Wealthy | 35-54 | College | Claritas Prizm™ |
| Shotguns and Pickups | Middle | 35-64 | High School | Claritas Prizm™ |
| Southside City | Poor | Mix | Grade School | Claritas Prizm™ |
| Living Off the Land | Middle-Poor | School Age Families | Low | Equifax MicroVision™ |
| University USA | Very low | Young - Mix | Medium to High | Equifax MicroVision™ |
| Sunset Years | Medium | Seniors | Medium | Equifax MicroVision™ |

This clustering information is then used by the end user to tag the customers in their database. Once this is done the business user can get a quick high level view of what is happening within the cluster. Once the business user has worked with these codes for some time they also begin to build intuitions about how these different customers clusters will react to the marketing offers particular to their business. For instance some of these clusters may relate to their business and some of them may not. But given that their competition may well be using these same clusters to structure their business and marketing offers it is important to be aware of how your customer base behaves in regard to these clusters.

Difference between Clustering and Nearest Neighbor Prediction

The main distinction between clustering and the nearest neighbor technique is that clustering is what is called an unsupervised learning technique and nearest neighbor is generally used for prediction or a supervised learning technique. Unsupervised learning techniques are unsupervised in the sense that when they are run there is not particular reason for the creation of the models the way there is for supervised learning techniques that are trying to perform prediction. In prediction, the patterns that are found in the database and presented in the model are always the most important patterns in the database for performing some particular prediction. In clustering there is no particular sense of why certain records are near to each other or why they all fall into the same cluster. Some of the differences between clustering and nearest neighbor prediction can be summarized in Table 10.2.



Task Conduct a research and analyze some Cluster Tags used in market.

Notes

Table 10.2: Some of the Differences between the Nearest-Neighbor Data Mining Technique and Clustering

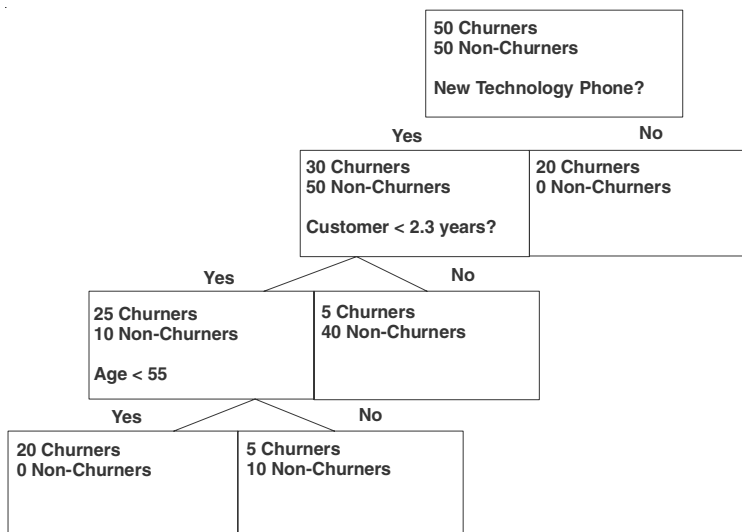
| Nearest Neighbor | Clustering |
|---|---|
| Used for prediction as well as consolidation. | Used mostly for consolidating data into a high-level view and general grouping of records into like behaviors. |
| Space is defined by the problem to be solved (supervised learning). | Space is defined as default n-dimensional space, or is defined by the user, or is a predefined space driven by past experience (unsupervised learning). |
| Generally only uses distance metrics to determine nearness. | Can use other metrics besides distance to determine nearness of two records - for example linking two points together. |

10.2.4 Decision Trees

A decision tree is a predictive model that, as its name implies, can be viewed as a tree. Specifically each branch of the tree is a classification question and the leaves of the tree are partitions of the dataset with their classification.

For example, if we were going to classify customers who churn (don't renew their phone contracts) in the Cellular Telephone Industry a decision tree might look something like that found in Figure 10.2.

Figure 10.2: A Decision Tree is a Predictive Model that makes a Prediction on the Basis of a Series of Decision much Like the Game of 20 Questions



You may notice some interesting things about the tree:

- It divides up the data on each branch point without losing any of the data (the number of total records in a given parent node is equal to the sum of the records contained in its two children).
- The number of churners and non-churners is conserved as you move up or down the tree
- It is pretty easy to understand how the model is being built (in contrast to the models from neural networks or from standard statistics).

Notes

- It would also be pretty easy to use this model if you actually had to target those customers that are likely to churn with a targeted marketing offer.

You may also build some intuitions about your customer base. E.g. “customers who have been with you for a couple of years and have up to date cellular phones are pretty loyal”.

Viewing Decision Trees as Segmentation with a Purpose

From a business perspective decision trees can be viewed as creating a segmentation of the original dataset (each segment would be one of the leaves of the tree). Segmentation of customers, products, and sales regions is something that marketing managers have been doing for many years. In the past this segmentation has been performed in order to get a high level view of a large amount of data - with no particular reason for creating the segmentation except that the records within each segmentation were somewhat similar to each other.

In this case the segmentation is done for a particular reason – namely for the prediction of some important piece of information. The records that fall within each segment fall there because they have similarity with respect to the information being predicted – not just that they are similar – without similarity being well defined. These predictive segments that are derived from the decision tree also come with a description of the characteristics that define the predictive segment. Thus the decision trees and the algorithms that create them may be complex, the results can be presented in an easy to understand way that can be quite useful to the business user.

Applying Decision Trees to Business

Because of their tree structure and ability to easily generate rules decision trees are the favored technique for building understandable models. Because of this clarity they also allow for more complex profit and ROI models to be added easily in on top of the predictive model. For instance once a customer population is found with high predicted likelihood to attract a variety of cost models can be used to see if an expensive marketing intervention should be used because the customers are highly valuable or a less expensive intervention should be used because the revenue from this sub-population of customers is marginal.

Because of their high level of automation and the ease of translating decision tree models into SQL for deployment in relational databases the technology has also proven to be easy to integrate with existing IT processes, requiring little preprocessing and cleansing of the data, or extraction of a special purpose file specifically for data mining.

Use of Decision Trees

Decision trees are data mining technology that has been around in a form very similar to the technology of today for almost twenty years now and early versions of the algorithms date back in the 1960s. Often times these techniques were originally developed for statisticians to automate the process of determining which fields in their database were actually useful or correlated with the particular problem that they were trying to understand. Partially because of this history, decision tree algorithms tend to automate the entire process of hypothesis generation and then validation much more completely and in a much more integrated way than any other data mining techniques. They are also particularly adept at handling raw data with little or no pre-processing. Perhaps also because they were originally developed to mimic the way an analyst interactively performs data mining they provide a simple to understand predictive model based on rules (such as “90% of the time credit card customers of less than three months who max out their credit limit are going to default on their credit card loan.”).

Because decision trees score so highly on so many of the critical features of data mining they can be used in a wide variety of business problems for both exploration and for prediction. They

have been used for problems ranging from credit card attrition prediction to time series prediction of the exchange rate of different international currencies. There are also some problems where decision trees will not do as well. Some very simple problems where the prediction is just a simple multiple of the predictor can be solved much more quickly and easily by linear regression.



Notes Usually the models to be built and the interactions to be detected are much more complex in real world problems and this is where decision trees excel.

Using Decision Trees for Exploration

The decision tree technology can be used for exploration of the dataset and business problem. This is often done by looking at the predictors and values that are chosen for each split of the tree. Often times these predictors provide usable insights or propose questions that need to be answered. For instance, if you ran across the following in your database for cellular phone churn you might seriously wonder about the way your tele-sales operators were making their calls and maybe change the way that they are compensated: "IF customer lifetime < 1.1 years AND sales channel = tele-sales THEN chance of churn is 65%."

Using Decision Trees for Data Preprocessing

Another way that the decision tree technology has been used is for preprocessing data for other prediction algorithms. Because the algorithm is fairly robust with respect to a variety of predictor types (e.g. number, categorical, etc.) and because it can be run relatively quickly decision trees can be used on the first pass of a data mining run to create a subset of possibly useful predictors that can then be fed into neural networks, nearest neighbor and normal statistical routines - which can take a considerable amount of time to run if there are large numbers of possible predictors to be used in the model.

Decision Trees for Prediction

Although some forms of decision trees were initially developed as exploratory tools to refine and preprocess data for more standard statistical techniques like logistic regression. They have also been used and more increasingly often being used for prediction. This is interesting because many statisticians will still use decision trees for exploratory analysis effectively building a predictive model as a by product but then ignore the predictive model in favor of techniques that they are most comfortable with. Sometimes veteran analysts will do this even excluding the predictive model when it is superior to that produced by other techniques. With a host of new products and skilled users now appearing this tendency to use decision trees only for exploration now seems to be changing.

The First Step is Growing the Tree

The first step in the process is that of growing the tree. Specifically the algorithm seeks to create a tree that works as perfectly as possible on all the data that is available. Most of the time it is not possible to have the algorithm work perfectly. There is always noise in the database to some degree (there are variables that are not being collected that have an impact on the target you are trying to predict).

The name of the game in growing the tree is in finding the best possible question to ask at each branch point of the tree. At the bottom of the tree you will come up with nodes that you would like to be all of one type or the other. Thus the question: "Are you over 40?" probably does not

Notes

sufficiently distinguish between those who are churners and those who are not - let's say it is 40%/60%. On the other hand, there may be a series of questions that do quite a nice job in distinguishing those cellular phone customers who will churn and those who won't. Maybe the series of questions would be something like: "Have you been a customer for less than a year, do you have a telephone that is more than two years old and were you originally landed as a customer via tele-sales rather than direct sales?" This series of questions defines a segment of the customer population in which 90% churn. These are then relevant questions to be asking in relation to predicting churn.

If the decision tree algorithm just continued growing the tree like this it could conceivably create more and more questions and branches in the tree so that eventually there was only one record in the segment. To let the tree grow to this size is both computationally expensive but also unnecessary. Most decision tree algorithms stop growing the tree when one of three criteria are met:

- The segment contains only one record. (There is no further question that you could ask which could further refine a segment of just one.)
- All the records in the segment have identical characteristics. (There is no reason to continue asking further questions segmentation since all the remaining records are the same.)
- The improvement is not substantial enough to warrant making the split.

10.2.5 Neural Networks

When data mining algorithms are talked about these days most of the time people are talking about either decision trees or neural networks. Of the two neural networks have probably been of greater interest through the formative stages of data mining technology. As we will see neural networks do have disadvantages that can be limiting in their ease of use and ease of deployment, but they do also have some significant advantages. Foremost among these advantages is their highly accurate predictive models that can be applied across a large number of different types of problems.

To be more precise with the term "neural network" one might better speak of an "artificial neural network". True neural networks are biological systems (a k a brains) that detect patterns, make predictions and learn. The artificial ones are computer programs implementing sophisticated pattern detection and machine learning algorithms on a computer to build predictive models from large historical databases. Artificial neural networks derive their name from their historical development which started off with the premise that machines could be made to "think" if scientists found ways to mimic the structure and functioning of the human brain on the computer. Thus historically neural networks grew out of the community of Artificial Intelligence rather than from the discipline of statistics. Despite the fact that scientists are still far from understanding the human brain let alone mimicking it, neural networks that run on computers can do some of the things that people can do.

It is difficult to say exactly when the first "neural network" on a computer was built. During World War II a seminal paper was published by McCulloch and Pitts which first outlined the idea that simple processing units (like the individual neurons in the human brain) could be connected together in large networks to create a system that could solve difficult problems and display behavior that was much more complex than the simple pieces that made it up. Since that time much progress has been made in finding ways to apply artificial neural networks to real world prediction problems and in improving the performance of the algorithm in general. In many respects the greatest breakthroughs in neural networks in recent years have been in their application to more mundane real world problems like customer response prediction or fraud detection rather than the loftier goals that were originally set out for the techniques such as overall human learning and computer speech and image understanding.

Because of the origins of the techniques and because of some of their early successes the techniques have enjoyed a great deal of interest. To understand how neural networks can detect patterns in a database an analogy is often made that they “learn” to detect these patterns and make better predictions in a similar way to the way that human beings do. This view is encouraged by the way the historical training data is often supplied to the network - one record (example) at a time. Neural networks do “learn” in a very real sense but under the hood the algorithms and techniques that are being deployed are not truly different from the techniques found in statistics or other data mining algorithms. It is for instance, unfair to assume that neural networks could outperform other techniques because they “learn” and improve over time while the other techniques are static. The other techniques in fact “learn” from historical examples in exactly the same way but often times the examples (historical records) to learn from are processed all at once in a more efficient manner than neural networks which often modify their model one record at a time.

A common claim for neural networks is that they are automated to a degree where the user does not need to know that much about how they work, or predictive modeling or even the database in order to use them. The implicit claim is also that most neural networks can be unleashed on your data straight out of the box without having to rearrange or modify the data very much to begin with.

Just the opposite is often true. There are many important design decisions that need to be made in order to effectively use a neural network such as:

- How should the nodes in the network be connected?
- How many neuron like processing units should be used?
- When should “training” be stopped in order to avoid overfitting?

There are also many important steps required for preprocessing the data that goes into a neural network - most often there is a requirement to normalize numeric data between 0.0 and 1.0 and categorical predictors may need to be broken up into virtual predictors that are 0 or 1 for each value of the original categorical predictor. And, as always, understanding what the data in your database means and a clear definition of the business problem to be solved are essential to ensuring eventual success. The bottom line is that neural networks provide no short cuts.

Applying Neural Networks to Business

Neural networks are very powerful predictive modeling techniques but some of the power comes at the expense of ease of use and ease of deployment. As we will see in this section, neural networks, create very complex models that are almost always impossible to fully understand even by experts. The model itself is represented by numeric values in a complex calculation that requires all of the predictor values to be in the form of a number. The output of the neural network is also numeric and needs to be translated if the actual prediction value is categorical (e.g. predicting the demand for blue, white or black jeans for a clothing manufacturer requires that the predictor values blue, black and white for the predictor color to be converted to numbers).

Because of the complexity of these techniques much effort has been expended in trying to increase the clarity with which the model can be understood by the end user. These efforts are still in their infancy but are of tremendous importance since most data mining techniques including neural networks are being deployed against real business problems where significant investments are made based on the predictions from the models (e.g. consider trusting the predictive model from a neural network that dictates which one million customers will receive a \$1 mailing).

There are two ways that these shortcomings in understanding the meaning of the neural network model have been successfully addressed:

Notes

- The neural network is package up into a complete solution such as fraud prediction. This allows the neural network to be carefully crafted for one particular application and once it has been proven successful it can be used over and over again without requiring a deep understanding of how it works.
- The neural network is package up with expert consulting services. Here the neural network is deployed by trusted experts who have a track record of success. Either the experts are able to explain the models or they are trusted that the models do work.

The first tactic has seemed to work quite well because when the technique is used for a well defined problem many of the difficulties in preprocessing the data can be automated (because the data structures have been seen before) and interpretation of the model is less of an issue since entire industries begin to use the technology successfully and a level of trust is created. There are several vendors who have deployed this strategy (e.g. HNC's Falcon system for credit card fraud prediction and Advanced Software Applications ModelMAX package for direct marketing).

Packaging up neural networks with expert consultants is also a viable strategy that avoids many of the pitfalls of using neural networks, but it can be quite expensive because it is human intensive. One of the great promises of data mining is, after all, the automation of the predictive modeling process. These neural network consulting teams are little different from the analytical departments many companies already have in house. Since there is not a great difference in the overall predictive accuracy of neural networks over standard statistical techniques the main difference becomes the replacement of the statistical expert with the neural network expert. Either with statistics or neural network experts the value of putting easy to use tools into the hands of the business end user is still not achieved.

Where to Use Neural Networks

Neural networks are used in a wide variety of applications. They have been used in all facets of business from detecting the fraudulent use of credit cards and credit risk prediction to increasing the hit rate of targeted mailings. They also have a long history of application in other areas such as the military for the automated driving of an unmanned vehicle at 30 miles per hour on paved roads to biological simulations such as learning the correct pronunciation of English words from written text.

Neural Networks for Clustering

Neural networks of various kinds can be used for clustering and prototype creation. The Kohonen network described in this section is probably the most common network used for clustering and segmentation of the database. Typically the networks are used in a unsupervised learning mode to create the clusters. The clusters are created by forcing the system to compress the data by creating prototypes or by algorithms that steer the system toward creating clusters that compete against each other for the records that they contain, thus ensuring that the clusters overlap as little as possible.

Neural Networks for Outlier Analysis

Sometimes clustering is performed not so much to keep records together as to make it easier to see when one record sticks out from the rest. For instance:

Most wine distributors selling inexpensive wine in Missouri and that ship a certain volume of product produce a certain level of profit. There is a cluster of stores that can be formed with these characteristics. One store stands out, however, as producing significantly lower profit. On closer

examination it turns out that the distributor was delivering product to but not collecting payment from one of their customers.

A sale on men's suits is being held in all branches of a department store for southern California. All stores with these characteristics have seen at least a 100% jump in revenue since the start of the sale except one. It turns out that this store had, unlike the others, advertised via radio rather than television.

Neural Networks for Feature Extraction

One of the important problems in all of data mining is that of determining which predictors are the most relevant and the most important in building models that are most accurate at prediction. These predictors may be used by themselves or they may be used in conjunction with other predictors to form "features".

Example: A simple example of a feature in problems that neural networks are working on is the feature of a vertical line in a computer image.

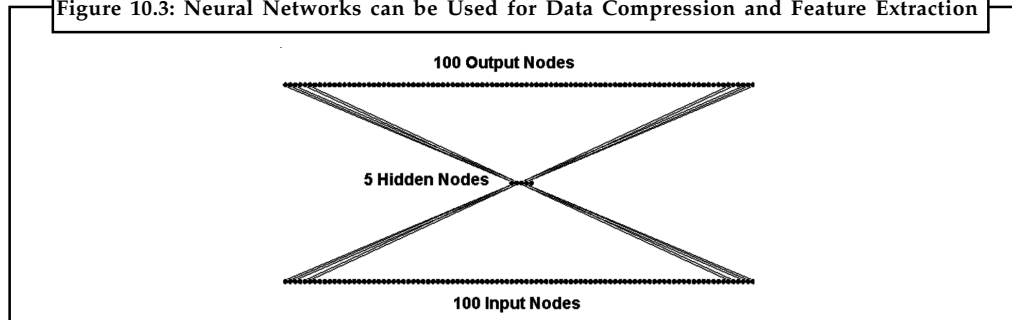
The predictors, or raw input data are just the colored pixels that make up the picture. Recognizing that the predictors (pixels) can be organized in such a way as to create lines, and then using the line as the input predictor can prove to dramatically improve the accuracy of the model and decrease the time to create it.

Some features like lines in computer images are things that humans are already pretty good at detecting, in other problem domains it is more difficult to recognize the features.

One novel way that neural networks have been used to detect features is the idea that features are sort of a compression of the training database. For instance you could describe an image to a friend by rattling off the color and intensity of each pixel on every point in the picture or you could describe it at a higher level in terms of lines, circles – or maybe even at a higher level of features such as trees, mountains etc. In either case your friend eventually gets all the information that they need in order to know what the picture looks like, but certainly describing it in terms of high level features requires much less communication of information than the "paint by numbers" approach of describing the color on each square millimeter of the image.

If we think of features in this way, as an efficient way to communicate our data, then neural networks can be used to automatically extract them. The neural network shown in Figure 10.3 is used to extract features by requiring the network to learn to recreate the input data at the output nodes by using just 5 hidden nodes. Consider that if you were allowed 100 hidden nodes, that recreating the data for the network would be rather trivial - simply pass the input node value directly through the corresponding hidden node and on to the output node. But as there are fewer and fewer hidden nodes, that information has to be passed through the hidden layer in a more and more efficient manner since there are less hidden nodes to help pass along the information.

Figure 10.3: Neural Networks can be Used for Data Compression and Feature Extraction



Notes

In order to accomplish this the neural network tries to have the hidden nodes extract features from the input nodes that efficiently describe the record represented at the input layer. This forced “squeezing” of the data through the narrow hidden layer forces the neural network to extract only those predictors and combinations of predictors that are best at recreating the input record. The link weights used to create the inputs to the hidden nodes are effectively creating features that are combinations of the input nodes values.

The concepts of neural networks are discussed in detail in the next unit.

10.2.6 Rule Induction

Rule induction is one of the major forms of data mining and is perhaps the most common form of knowledge discovery in unsupervised learning systems. It is also perhaps the form of data mining that most closely resembles the process that most people think about when they think about data mining, namely “mining” for gold through a vast database. The gold in this case would be a rule that is interesting – that tells you something about your database that you didn’t already know and probably weren’t able to explicitly articulate (aside from saying “show me things that are interesting”).

Rule induction on a data base can be a massive undertaking where all possible patterns are systematically pulled out of the data and then an accuracy and significance are added to them that tell the user how strong the pattern is and how likely it is to occur again. In general these rules are relatively simple such as for a market basket database of items scanned in a consumer market basket you might find interesting correlations in your database such as:

- If bagels are purchased then cream cheese is purchased 90% of the time and this pattern occurs in 3% of all shopping baskets.
- If live plants are purchased from a hardware store then plant fertilizer is purchased 60% of the time and these two items are bought together in 6% of the shopping baskets.



Did u know? The rules that are pulled from the database are extracted and ordered to be presented to the user based on the percentage of times that they are correct and how often they apply.

The bane of rule induction systems is also its strength - that it retrieves all possible interesting patterns in the database. This is a strength in the sense that it leaves no stone unturned but it can also be viewed as a weakness because the user can easily become overwhelmed with such a large number of rules that it is difficult to look through all of them. You almost need a second pass of data mining to go through the list of interesting rules that have been generated by the rule induction system in the first place in order to find the most valuable gold nugget amongst them all. This overabundance of patterns can also be problematic for the simple task of prediction because all possible patterns are culled from the database there may be conflicting predictions made by equally interesting rules. Automating the process of culling the most interesting rules and of combing the recommendations of a variety of rules are well handled by many of the commercially available rule induction systems on the market today and is also an area of active research.

Self Assessment

Fill in the blanks:

5. is a branch of mathematics concerning the collection and the description of data.

6. is a technique that classifies each record in a dataset based on a combination of the classes of the k record(s) most similar to it in a historical dataset (where k 1).
7. is the method by which like records are grouped together.
8. A is a predictive model that, as its name implies, can be viewed as a tree.
9. Artificial are computer programs implementing sophisticated pattern detection and machine learning algorithms on a computer.
10. is the form of data mining that most closely resembles the process that most people think about when they think about data mining, namely “mining” for gold through a vast database.

Notes

10.3 Text Mining

Text mining is the process of using computer technology to sift through text documents for the purposes of research and analysis. It is often considered very similar to the process known as data mining, but it relies on special programming to look in uncategorized text and find meaning or patterns instead of analyzing pre-categorized database information. Text mining has many applications in areas like science, marketing, and data organization.

The complexity involved in organizing words into language is much too extreme for computers to handle, but scientists have worked hard to improve this kind of programming. Many methods have been developed that let scientists identify phrases and discover facts about text. This is generally not the same as fully deciphering the meaning, but it allows for shortcuts that achieve many of the same goals. Text mining takes advantage of some of these techniques, and as this technology improves, text mining is generally expected to improve as well.

Experts use text information analysis primarily to do research into written documents. Large amounts of written data can be hard to analyze because of the tremendous amount of time required. Computers can go through this text much quicker, but they can’t understand it. Text mining techniques allow computers to find useful trends in text, presenting the data in a way that may reveal new facts or allow experts to make discoveries.

An example of a use for this technology would be market research.

Experts could analyze search results on a product name and have the program look for phrases that express user sentiment. In this way, they may find out how people really feel about their product in a very detailed way. They could also simply look for their product and see which phrases were popping up most often, and this might help them develop new ideas about how to please their customers.

Another use for mining text is analyzing scientific papers on similar subjects looking for new trends or agreements. This has allowed some scientists to make predictive assumptions that have proven useful in fields like protein analysis. Some experts think these sorts of applications may eventually provide unexpected discoveries.

A process called data mining is actually quite similar to the mining of text, but it is generally less complex to do because it relies on text that’s already been formatted into categories.

For example, the software could go through all the information for job applicants in a database, looking for trends.



Caution Text mining is more difficult for computers to do because pure text is harder to analyze than data with categories.

Notes

Self Assessment

State True or False:

- 11. Text mining is more difficult for computers to do.
- 12. A process called data mining is different from the mining of text.

10.4 Web Mining

Web mining is a form of information harvesting that applies to data gathered from online sources. Data collection from sources across the Internet allows users to aggregate large volumes of information for analysis to make key business decisions in an online environment.



Example: A researcher might use web mining to collect information regarding the use of specific keywords in web content.

Alternatively, retailers and other marketing professionals use online data mining to spot trends in web traffic, the conversion of site visitors to buyers, and other web usage.

In terms of gathering, sorting, and analyzing data, web mining mimics traditional data mining activities. Comparatively, web mining activities focus on web-based information, rather than a large cross section of information sources such as off-line computer databases, customer records, or hard copy accounting data, as typically occurs with traditional data mining. Focusing solely on data collection from online sources provides targeted analysis needed for online marketing strategies, website structure decisions, and similar electronic commerce-related decision making. Collecting data via web mining also provides the added benefit of a broad international demographic, since websites from all over the world are available to researchers and information gatherers.

Professionally, web mining is divided into three specific categories: web structure mining, usage mining, and web content mining. Each area focuses on specific information such as the structure and hyperlinks of a particular website, server log information regarding visitor usage, and specific content available online.

Website analytic software packages and services are a prime example of web usage mining, providing webmasters with information regarding visitor traffic, search results used, links clicked, and time spent interacting with specific pages. Structure mining, on the other hand, provides detailed information about a specific website's internal structure, including hyperlinks, databases, and query functions.

To the marketing professional, web mining offers a wealth of uses relative to marketing activities. Knowing how site visitors use a particular website, how competitors set up a competing site, and what content is already online is valuable information. Such information helps key decision makers craft a marketing strategy based on previously proven techniques and documented information.

Colleges and universities also utilize web mining via software that verifies student papers are unique and not plagiarized. Using web content mining principles, such grading aides search the entirety of the Internet for like content. Instructors upload the text of a student document and then instruct the plagiarism software to check the Internet for similar phrases or copied text online. Results are often expressed as percentage of matching text. Links to any similar results are provided to allow instructors the ability to visit sites to determine if matches are indeed plagiarized.



Task Compare and contrast text mining and web mining.

Notes

Self Assessment

Fill in the blanks:

13. is a form of information harvesting that applies to data gathered from online sources.
14. Web mining activities focus on information, rather than a large cross section of information sources.
15. provides detailed information about a specific website's internal structure.



Case Study

Data Mining in Mobile Communication

Application Background

Mobile communication data analysis has been often used as a background application to motivate many technical problems in data mining research, such as mining frequent patterns and clusters on data streams, social network analysis, collaborative filtering and recommendation. However, very few data mining researchers have a chance to see a working data mining system on real mobile communication data. The lack of this experience prevents those researchers from deeply understanding the business application scenarios in mobile communication as well as the successes and the limitations of the existing techniques.

We are developing MobileMiner, a data mining tool for mobile data analysis and business strategy development. Built on the state-of-the-art data mining techniques, MobileMiner presents a real case study on how to integrate data mining techniques into a business solution. In a large mobile communication company like China Mobile Communication Corporation, there are many analytical tasks where data mining can help to address the business interests of the company. Clearly, a system cannot cover all aspects. MobileMiner starts with customer relation management, the core component of mobile communication business. In this demo, we focus on two tasks, mobile user segmentation and community discovery from user calling networks.

MobileMiner provides a platform for the analytical tasks, where user profiles are extracted continuously from users' moving and calling records. The profiles are extremely important and valuable in business. Based on the profile mining platform, various data mining tasks can be effectively performed using different features of the profiles. The mobile user segmentation task tries to group customers by their frequent moving patterns. The features used for grouping are obtained by mining users' moving records continuously on the profile mining platform. Knowing the moving patterns for different customer groups, a service provider can dynamically deploy resources to improve the service quality (e.g., adjusting the angles of antennas or re-positioning a mobile station). For example, in Beijing Olympic period, many people are moving from Bird Nest around 9pm to Olympic Village around 11pm. It is interesting to find the clusters of customers in terms of service areas and time.

Contd....

Notes

The community discovery task aims to discover coherent calling communities. Based on the profile mining platform, a social network can be constructed using calls between customers and the calling frequencies. Communities in the social network capture the connectivity and similarity among customers. By considering the properties of the communities, effective market campaign can be designed for targeted customers. For example, the customers with broad social connections should be taken care specially.

We emphasize the following points in our demo. First, we present how we solve the business tasks in mobile communication using novel data mining techniques. Second, we use MobileMiner on real data to elaborate what can be done and how the data mining techniques can be integrated in a business-driven model. Third, we show some examples of what still cannot be done satisfactorily using the current data mining techniques, which may motivate future research and development.

Technology and Novelty

Customer records are collected by the mobile communication base stations and fed into MobileMiner as data streams, including customer moving trajectories and calling records. A base station serves the cell phones in a specific region, and can detect a mobile customer once she turns on her phone. Once the records are imported into the system, profile mining is performed to generate user profiles for the upper layer data mining tasks.

Specifically, in the profile mining part, customers' moving profiles or their frequent moving patterns are constructed based on their moving records continuously. The core of this task is to mine sequential patterns on data streams, which is challenging since there can be many customers and the sliding window can be large. A customer's moving profile is formed using the set of closed sequential patterns that match the customer's trajectory and the profile is incrementally maintained. We developed a novel algorithm to mine and incrementally maintain on fast data streams closed sequential patterns, which are non-redundant representation of sequential patterns. An effective data structure is designed to keep close sequential patterns in memory and various strategies are proposed to prune search space aggressively. Based on the experiments on both real and synthetic databases, our algorithm outperforms the best existing algorithms by a large margin. The details of the techniques can be found in.

The mobile user segmentation module clusters customers according to their profiles. The goal is to partition customers into groups such that the customers in a group are similar to each other in moving patterns. Importantly, timestamps should be considered. Since each point in a customer trajectory is associated with a timestamp, two trajectories are similar only if they are close to each other in time dimension. The problem is formulated as clustering trajectories in both space and time. The spatio-temporal patterns of clusters are very useful for the company to allocate base stations effectively for specific customer groups. Some related work clusters spatio-temporal patterns in bio-informatics. Here, we adapt the algorithm to group 2-dimensional trajectories in different time stamps. The main idea is to find biclusters with low mean squared residue through effectively iterative search. The mean squared residue captures the variance of the set of trajectories in a bicluster over time.

In mobile communication business, the social relationship among customers often plays a significant role in marketing. For example, losing some customers with broad social connections may cause customer churning. A social network among customers is constructed. Each customer is represented by a node in the network. An edge is drawn to connect two customers if they call each other over a certain number of times in the current sliding window. A social community in the network is a set of nodes such that they are

Contd....

relatively well connected to each other and much less connected to the other nodes in the network. Some previous work discovers communities in a network. In this application, the connection weights on edges in graphs should be considered. The core set is a set of customers whom are frequently called by other customers. The affiliated customers are the customers surrounding the core with different layers. We use the calling frequencies as the weights in the process of finding core customers and ranking affiliated customers. To control the granularity of the discovering communities, a merging schema is used to merge similar communities to get coarser results.

Demo Scenario

Our demo consists of three parts. First, we showcase how we integrate the state-of-the-art data mining techniques into a business framework and building MobileMiner as a business solution. Second, we illustrate how the underlying data mining techniques affect business analysis. Last, we present some interesting observations found from real data which unfortunately still cannot be handled by the existing techniques. Such case studies may motivate novel data mining research and development.

Techniques Meeting Business Requirements

We will demonstrate some common business analysis tasks in mobile communication companies, including customer segmentation for mobile service bases deployment, and calling community discovery for marketing campaign design. For example, the user interface of the mobile user segmentation module is not a simple list of the users grouped in each cluster. MobileMiner visualizes the user groups by showing their moving patterns, each group in a different color. Moreover, the moving patterns are shown in temporal order with a local map as the background. With this information, analysts can make informative decisions about how to deploy mobile base stations more effectively. We will also show how calling community discovery techniques help companies to design marketing campaign. The graph mining results are presented properly in a business driven way. Based on the discovered knowledge, business analysts can identify targeted customers in an effective way.

Sharpening Business Analysis by Tuning Techniques

Data mining techniques need to be tuned to make business analysis effective. To understand how well the data mining techniques in MobileMiner work in practice, we use a real mobile communication data set to show some interesting mining results.

To demonstrate the tuning needs, we will show how the parameters of our sequential pattern mining algorithms may affect the mining results. Moreover, it is important that the user interface can help business analysts to tune the underlying data mining methods. A business analyst can interact with the social community visualization to tune the parameters of the social network construction such as the call frequency threshold and the time window.

Opportunities for Future Research

Mobile communication is a fast growing industry. We demonstrate some patterns found yet from real data by human analysts but cannot be found using the data mining techniques. For example, a new service of low calling charge by the company may negatively effect the sales of another service such as monthly SMS. It is critical to analyze whether such a new service overall improves the business and thus whether it should be introduced. Usually, this decision is based on the experiential analysis on both potential profits and potential customers. This task can be modeled as a hypothesis mining problem, which is highly demanded in business but has not been systematically studied in a practical setting.

Contd....

Notes

Question

Discuss the techniques used for fulfilling Business Requirements.

Source: <http://www.cs.sfu.ca/~jpei/publications/mobileminer-sigmod09.pdf>

10.5 Summary

- Data mining tools collect data and model the data to represent the reality.
- Basically organization install dashboard to monitor the data changes, information contained in the database and onscreen update.
- Text mining tools has the ability to mine data in various kind of text such as Microsoft words and acrobat PDF.
- Nearest neighbor method is a technique that classifies each record in a dataset based on a combination of the classes of the k record(s) most similar to it in a historical dataset.
- Clustering is the method by which like records are grouped together.
- A decision tree is a predictive model that, as its name implies, can be viewed as a tree.
- Artificial neural networks derive their name from their historical development which started off with the premise that machines could be made to “think” if scientists found ways to mimic the structure and functioning of the human brain on the computer.
- Rule induction is perhaps the form of data mining that most closely resembles the process that most people think about when they think about data mining, namely “mining” for gold through a vast database.
- Text mining is the process of using computer technology to sift through text documents for the purposes of research and analysis.
- Web mining is a form of information harvesting that applies to data gathered from online sources.

10.6 Keywords

Clustering: Clustering is the method by which like records are grouped together.

Dashboard: Dashboard is used to monitor the data changes, information contained in the database and onscreen update.

Decision Tree: A decision tree is a predictive model that, as its name implies, can be viewed as a tree.

Nearest Neighbor Method: Nearest neighbor method is a technique that classifies each record in a dataset based on a combination of the classes of the k record(s) most similar to it in a historical dataset (where k > 1).

Rule Induction: Rule induction is perhaps the form of data mining that most closely resembles the process that most people think about when they think about data mining, namely “mining” for gold through a vast database.

Text Mining Tool: Text mining tools has the ability to mine data in various kind of text such as Microsoft words and acrobat PDF.

Text Mining: Text mining is the process of using computer technology to sift through text documents for the purposes of research and analysis.

Notes

Web Mining: Web mining is a form of information harvesting that applies to data gathered from online sources.

10.7 Review Questions

1. Explain different data mining tools available in the market.
2. Discuss the advantages and disadvantages of dashboard.
3. Make distinction between data mining classical techniques and next generation techniques.
4. Explain the concept of the Nearest Neighbor technique with example.
5. Describe how to use Nearest Neighbor for Prediction.
6. Discuss the use of clustering technique with example.
7. Explain the use of decision tree with example.
8. Illustrate the use of decision trees for Data Preprocessing.
9. Discuss the significant steps required for preprocessing the data that goes into a neural network.
10. Illustrate the concept of web mining with example.

Answers: Self Assessment

- | | |
|----------------------|----------------------------|
| 1. Tools | 2. Dashboard |
| 3. Text mining | 4. Forensic analysis |
| 5. Statistics | 6. Nearest neighbor method |
| 7. Clustering | 8. Decision tree |
| 9. Neural networks | 10. Rule induction |
| 11. True | 12. False |
| 13. Web mining | 14. Web-based |
| 15. Structure mining | |

10.8 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.

Notes



Online links

<http://www.anderson.ucla.edu/faculty/jason.frand/teacher/technologies/palace/datamining.htm>

<http://www.ijcaonline.org/volume13/number4/pxc3872425.pdf>

<http://www.ijcaonline.org/volume13/number4/pxc3872425.pdf>

http://www.ncbi.nlm.nih.gov/About/tools/restable_data.html

Unit 11: Neural Networks

Notes

CONTENTS

Objectives

Introduction

11.1 Concept of Neural Network

11.1.1 Artificial Neural Network (ANN)

11.1.2 Biological and Artificial Neural Networks

11.1.3 Elements of ANN

11.1.4 Network Information Processing

11.1.5 Architecture

11.1.6 Learning in ANN

11.2 Development of Neural Network-based System

11.2.1 Data Collection and Preparation

11.2.2 Selection of Network Structure

11.2.3 Learning Algorithm Selection

11.2.4 Network Training

11.2.5 Testing

11.2.6 Implementation of an ANN

11.3 Summary

11.4 Keywords

11.5 Review Questions

11.6 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss the Concept of Neural Networks
- Explain the Elements of ANN
- Discuss the Development of Neural Network Based Mining

Introduction

An Artificial Neural Network (ANN), often just called a “neural network” (NN), is a mathematical model or computational model based on biological neural networks, in other words, is an emulation of biological neural system. It consists of an interconnected group of artificial neurons and processes information using a connectionist approach to computation. In most cases an ANN is an adaptive system that changes its structure based on external or internal information that flows through the network during the learning phase. In this unit, we will discuss the concept of neural networks.

11.1 Concept of Neural Network

Neural networks represent a brain metaphor for information processing. These models are biologically inspired rather than an exact replica of how the brain actually functions. Neural networks have been shown to be very promising systems in many forecasting applications and business classification applications due to their ability to “learn” from the data, their nonparametric nature (i.e., no rigid assumptions), and their ability to generalize. Neural computing refers to a pattern recognition methodology for machine learning. The resulting model from neural computing is often called an artificial neural network (ANN) or a neural network.

11.1.1 Artificial Neural Network (ANN)

The human brain possesses bewildering capabilities for information processing and problem solving that modern computers cannot compete with in many aspects. It has been postulated that a model or a system that is enlightened and supported by the results from brain research, with a structure similar to that of biological neural networks, could exhibit similar intelligent functionality. Based on this bottom-up postulation, ANN (also known as connectionist models, parallel distributed processing models, neuromorphic systems, or simply neural networks) have been developed as biologically inspired and plausible models for various tasks.

Biological neural networks are composed of many massively interconnected primitive biological neurons. Each neuron possesses axons and dendrites, finger-like projections that enable the neuron to communicate with its neighboring neurons by transmitting and receiving electrical and chemical signals. More or less resembling the structure of their counterparts, ANN are composed of interconnected, simple processing elements called artificial neurons. In processing information, the processing elements in an ANN operate concurrently and collectively in a similar fashion to biological neurons. ANN possess some desirable traits similar to those of biological neural networks, such as the capabilities of learning, self-organization, and fault tolerance.

The formal study of ANN began with the pioneering work of McCulloch and Pitts in 1943. Stimulated by results of biological experiments and observations, McCulloch and Pitts (1943) introduced a simple model of a binary artificial neuron that captures some functions of a living neuron. Considering information processing machines as a means for modeling the brain, McCulloch and Pitts built their neural networks model using a large number of interconnected binary artificial neurons. Led by a school of researchers, neural network research was quite popular in the late 1950s and early 1960s.



Notes After a thorough analysis of an early neural network model (called the perceptron, which used no hidden layer) as well as a pessimistic evaluation of the research potential by Minsky and Papert in 1969, the interest in neural networks diminished.

During the past two decades, there has been an exciting resurgence in the studies of ANN due to the introduction of new network topologies, new activation functions, and new learning algorithms, as well as progress in neuroscience and cognitive science. On the one hand, advances in theory and methodology have overcome many obstacles that hindered neural network research a few decades ago. Evidenced by the appealing results of numerous studies, neural networks are gaining acceptance and popularity. On the other hand, as complex problems solvers, ANN have been applied to solve numerous problems in a variety of application settings. The desirable features in neural information processing make neural networks attractive for solving complex


problems. The initial success in neural network applications has inspired renewed interest from industry and business.

11.1.2 Biological and Artificial Neural Networks

The human brain is composed of special cells called neurons. These cells do not die when a human is injured (all other cells reproduce to replace themselves and then die). This phenomenon may explain why we retain information. Information storage spans sets of neurons. The estimated number of neurons in a human brain is 50 to 150 billion, of which there are more than 100 different kinds. Neurons are partitioned into groups called networks. Each network contains several thousand highly interconnected neurons. Thus, the brain can be viewed as a collection of neural networks. The ability to learn and react to changes in our environment requires intelligence. The brain and the central nervous system control thinking and intelligent behavior. People who suffer brain damage have difficulty learning and reacting to changing environments. Even so, undamaged parts of the brain can often compensate with new learning.

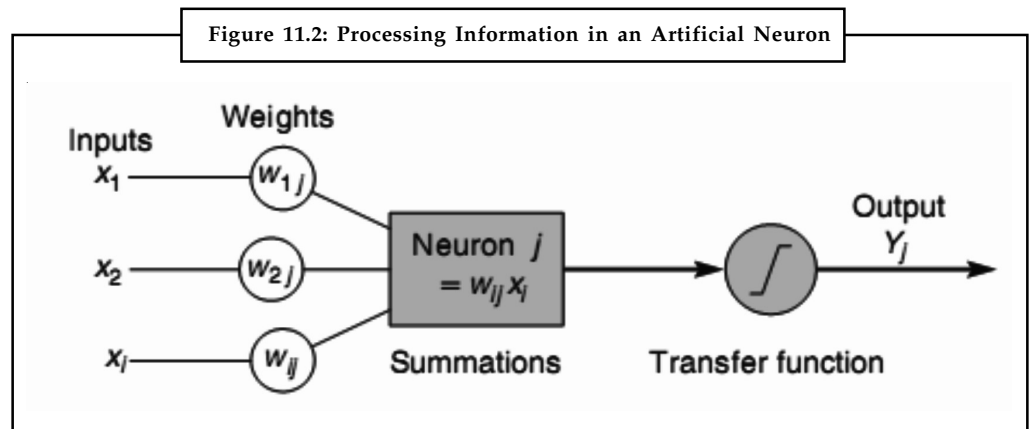
A portion of a network composed of two cells is shown in Figure 11.1. The cell itself includes a nucleus (the central processing portion of the neuron). To the left of cell 1, the dendrites provide input signals to the cell. To the right, the axon sends output signals to cell 2 via the axon terminals. These axon terminals merge with the dendrites of cell 2. Signals can be transmitted unchanged, or they can be altered by synapses. A synapse is able to increase or decrease the strength of the connection from neuron to neuron and cause excitation or inhibition of a subsequent neuron. This is where information is stored.

An ANN model emulates a biological neural network. Neural computing actually uses a very limited set of concepts from biological neural systems. It is more of an analogy to the human brain than an accurate model of it. Neural concepts are usually implemented as software simulations of the massively parallel processes that involve processing elements (also called artificial neurons, or neurodes) interconnected in a network architecture. The artificial neuron receives inputs analogous to the electrochemical impulses the dendrites of biological neurons receive from other neurons. The output of the artificial neuron corresponds to signals sent out from a biological neuron over its axon. These artificial signals can be changed by weights in a manner similar to the physical changes that occur in the synapses (see Figure 11.2).



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

Notes



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

Several ANN paradigms have been proposed for applications in a variety of problem domains. Perhaps the easiest way to differentiate between the various models is on the basis of how these models structurally emulate the human brain, the way in which the neural model processes information and how the neural models learn to perform their designated tasks. As they are biologically inspired, the main processing elements of a neural network are individual neurons, analogous to the brain’s neurons. These artificial neurons receive the sum “information” from other neurons or external input stimuli, perform a transformation on the inputs, and then pass on the transformed information to other neurons or external outputs. This is similar to how it is presently thought that the human brain works. Passing information from neuron to neuron can be thought of as a way to activate, or trigger a response from certain neurons based on the information or stimulus received.

Thus, how information is processed by a neural network is inherently a function of its structure. Neural networks can have one or more layers of neurons. These neurons can be highly or fully interconnected, or only certain layers can be connected together. Connections between neurons have an associated weight. In essence, the “knowledge” possessed by the network is encapsulated in these interconnection weights. Each neuron calculates a weighted sum of the incoming neuron values, transforms this input, and passes on its neural value as the input to subsequent neurons.



Did u know? Typically, although not always, this input/output transformation process at the individual neuron level is done in a nonlinear fashion.

11.1.3 Elements of ANN

A neural network is composed of processing elements organized in different ways to form the network’s structure. The basic processing unit is the neuron. A number of neurons are organized into a network. There are many ways to organize neurons; they are referred to as topologies. One popular approach, known as the feedforward back propagation paradigm (or simply back propagation), allows all neurons to link the output in one layer to the input of the next layer, but it does not allow any feedback linkage. This is the most commonly used paradigm.

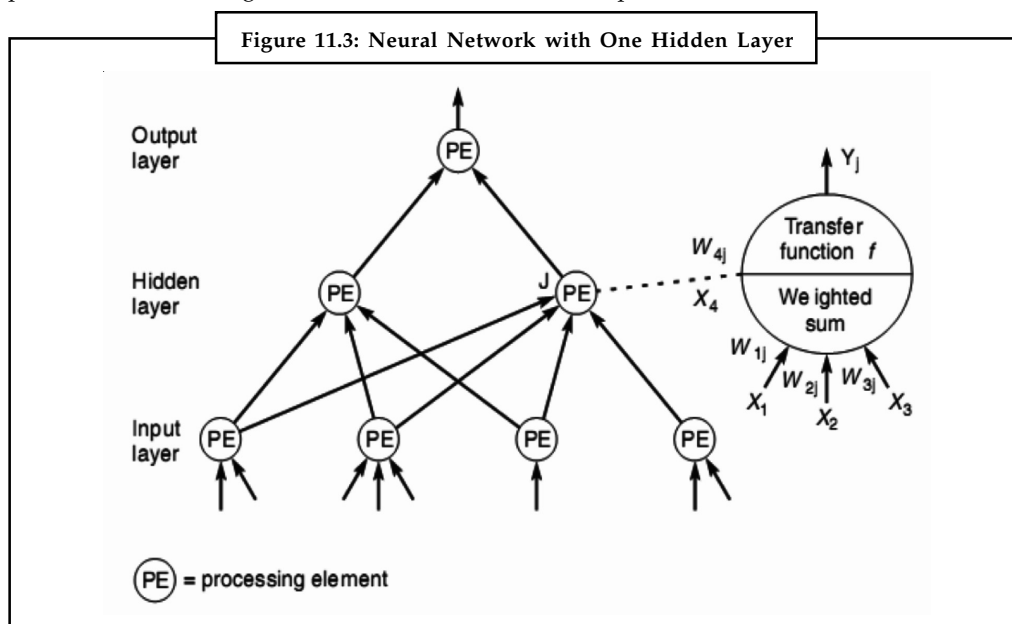
Processing Elements (PE)

The processing elements of an ANN are artificial neurons. Each of the neurons receives inputs, processes them, and delivers a single output, as shown in Figure 11.2. The input can be raw input data or the output of other processing elements. The output can be the final result (e.g., 1 means yes, 0 means no), or it can be inputs to other neurons.

Network Structure

Notes

Each ANN is composed of a collection of neurons, grouped in layers. A typical structure is shown in Figure 11.3. Note the three layers: input, intermediate (called the hidden layer), and output. A hidden layer is a layer of neurons that takes input from the previous layer and converts those inputs into outputs for further processing. Several hidden layers can be placed between the input and output layers, although it is quite common to use only one hidden layer. In that case, the hidden layer simply converts inputs into a nonlinear combination and passes the transformed inputs to the output layer. The most common interpretation of the hidden layer is as a feature extraction mechanism. That is, the hidden layer converts the original inputs in the problem into some higher level combinations of such inputs.



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

Like a biological network, an ANN can be organized in several different ways (i.e., topologies or architectures); that is, the neurons can be interconnected in different ways. Therefore, ANN appear in many configurations called architectures. When information is processed, many of the processing elements perform their computations at the same time. This parallel processing resembles the way the brain works, and it differs from the serial processing of conventional computing.

11.1.4 Network Information Processing

Once the structure of a neural network is determined, information can be processed. We now present the major concepts related to the processing.

Inputs: Each input corresponds to a single attribute.



Example: If the problem is to decide on approval or disapproval of a loan, some attributes could be the applicant's income level, age, and home ownership.

The numeric value, or representation, of an attribute is the input to the network. Several types of data, such as text, pictures, and voice, can be used as inputs. Preprocessing may be needed to convert the data to meaningful inputs from symbolic data or to scale the data.

Notes

Outputs: The outputs of a network contain the solution to a problem.



Example: In the case of a loan application, the outputs can be yes or no.

The ANN assigns numeric values to the outputs, such as 1 for yes and 0 for no. The purpose of the network is to compute the values of the output. Often, post-processing of the outputs is required because some networks use two outputs: one for yes and another for no. It is common to have to round the outputs to the nearest 0 or 1.

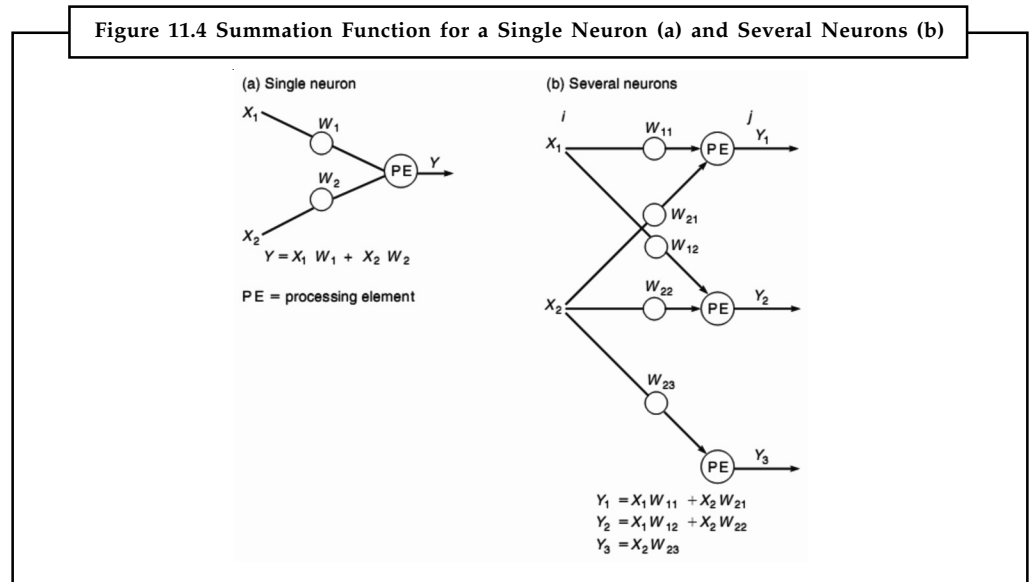
Connection Weights: Connection weights are the key elements in an ANN. They express the relative strength (or mathematical value) of the input data or the many connections that transfer data from layer to layer. In other words, weights express the relative importance of each input to a processing element and, ultimately, the outputs weights are crucial in that they store learned patterns of information. It is through repeated adjustments of weights that a network learns.

Summation Function: The summation function computes the weighted sums of all the input elements entering each processing element. A summation function multiplies each input value by its weight and totals the values for a weighted sum Y . The formula for n inputs in one processing element (see Figure 11.4a) is:

$$Y = \sum_{i=1}^n X_i W_i$$

For the j th neuron of several processing neurons in a layer (see Figure 11.4b), the formula is:

$$Y_j = \sum_{i=1}^n X_i W_{ij}$$



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

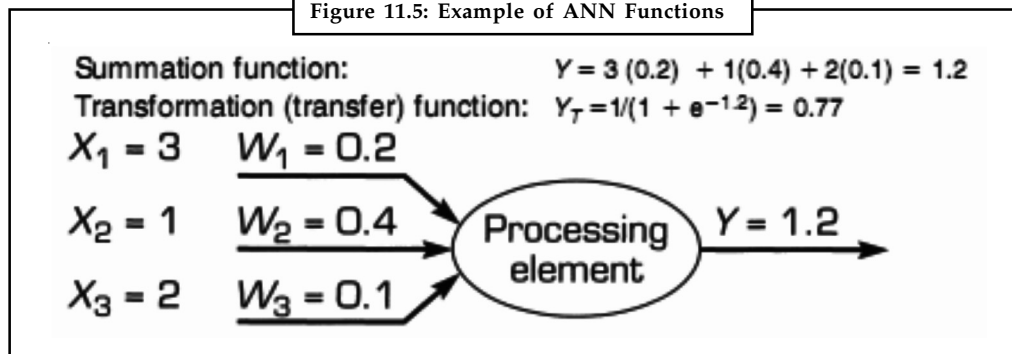
Transformation (Transfer) Function: The summation function computes the internal stimulation, or activation level, of the neuron. Based on this level, the neuron may or may not produce an output. The relationship between the internal activation level and the output can be linear or

nonlinear. The relationship is expressed by one of several types of transformation (transfer) functions. The transformation (transfer) function combines (i.e., adds up) the inputs coming into a neuron from other neurons/sources and then produces an output based on the choice of the transfer function. Selection of the specific function affects the network's operation. The sigmoid (logical activation) function (or sigmoid transfer function) is an S-shaped transfer function in the range of 0 to 1, and it is a popular as well as useful nonlinear transfer function:

$$Y_T = 1/(1 + e^{-Y})$$

where Y_T is the transformed (i.e., normalized) value of Y (see Figure 11.5).

Figure 11.5: Example of ANN Functions



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

The transformation modifies the output levels to reasonable values (typically between 0 and 1). This transformation is performed before the output reaches the next level. Without such a transformation, the value of the output becomes very large, especially when there are several layers of neurons. Sometimes, instead of a transformation function, a threshold value is used. A threshold value is a hurdle value for the output of a neuron to trigger the next level of neurons. If an output value is smaller than the threshold value, it will not be passed to the next level of neurons.



Example: Any value of 0.5 or less becomes 0, and any value above 0.5 becomes 1.

A transformation can occur at the output of each processing element, or it can be performed only at the final output nodes.

Hidden Layers: Complex practical applications require one or more hidden layers between the input and output neurons and a correspondingly large number of weights. Many commercial ANN include three and sometimes up to five layers, with each containing 10 to 1,000 processing elements. Some experimental ANN use millions of processing elements. Because each layer increases the training effort exponentially and also increases the computation required, the use of more than three hidden layers is rare in most commercial systems.

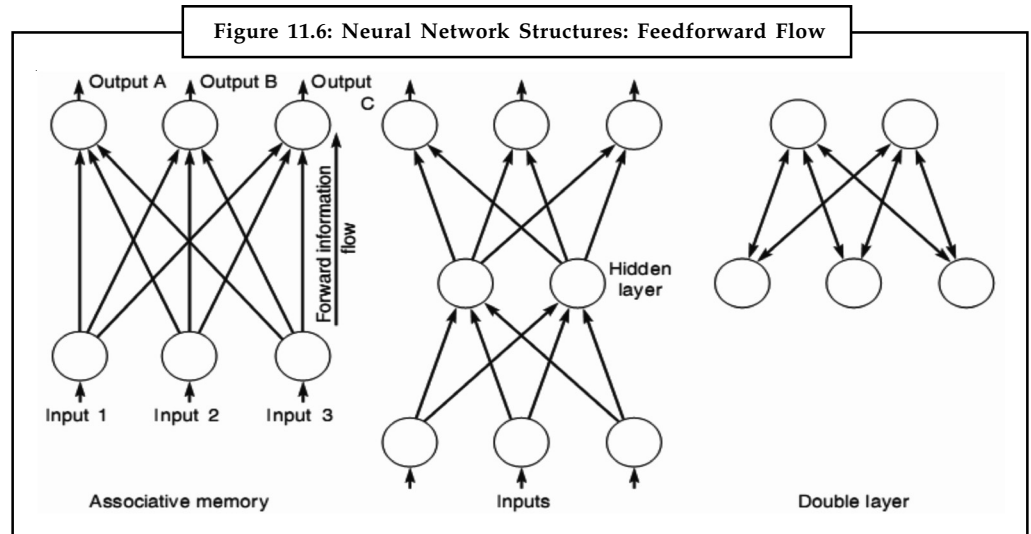


Task Compare and contrast summation function and transformation function.

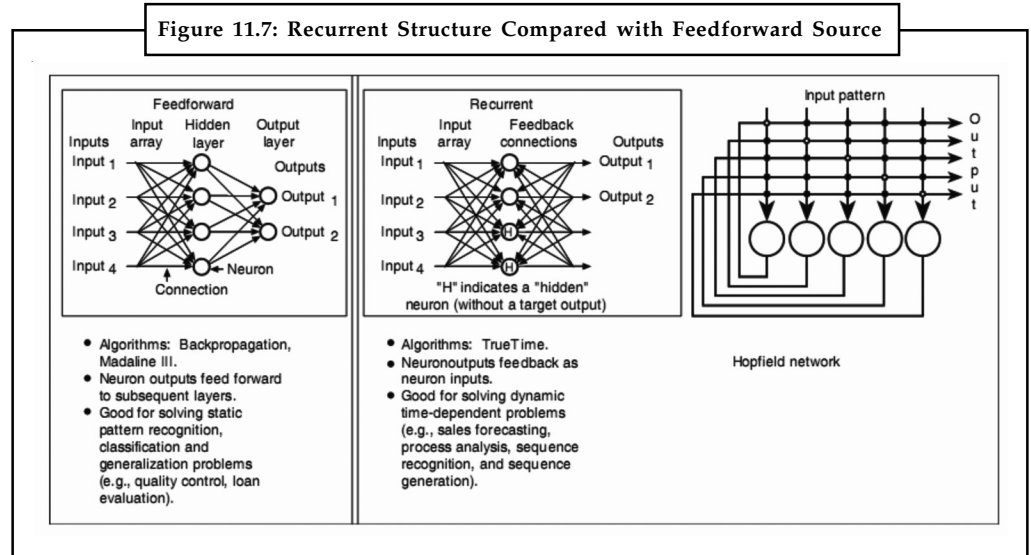
11.1.5 Architecture

There are several effective neural network models and algorithms. Some of the most common are back propagation (or feedforward), associative memory, and the recurrent network.

Notes



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

Ultimately, the operation of the entire neural network model is driven by the task it is designed to address. For instance, neural network models have been used as classifiers, as forecasting tools, and as general optimizers. Neural network classifiers are typically multilayer models in which information is passed from one layer to the next, with the ultimate goal of mapping an input to the network to a specific category, as identified by an output of the network. A neural model used as an optimizer, on the other hand, can be a single layer of neurons, highly interconnected, and can compute neuron values iteratively until the model converges to a stable state. This stable state would then represent an optimal solution to the problem under analysis.

Finally, how a network is trained to perform its desired task is another identifying model characteristic. Neural network learning can occur in either a supervised or unsupervised mode. In supervised learning, a sample training set is used to "teach" the network about its problem domain. This training set of exemplar cases (input and the desired output[s]) is iteratively presented to the neural network. Output of the network in its present form is calculated and compared to the desired output. The learning algorithm is the training procedure that an ANN uses.



Notes The learning algorithm being used determines how the neural interconnection weights are corrected due to differences in the actual and desired output for a member of the training set.

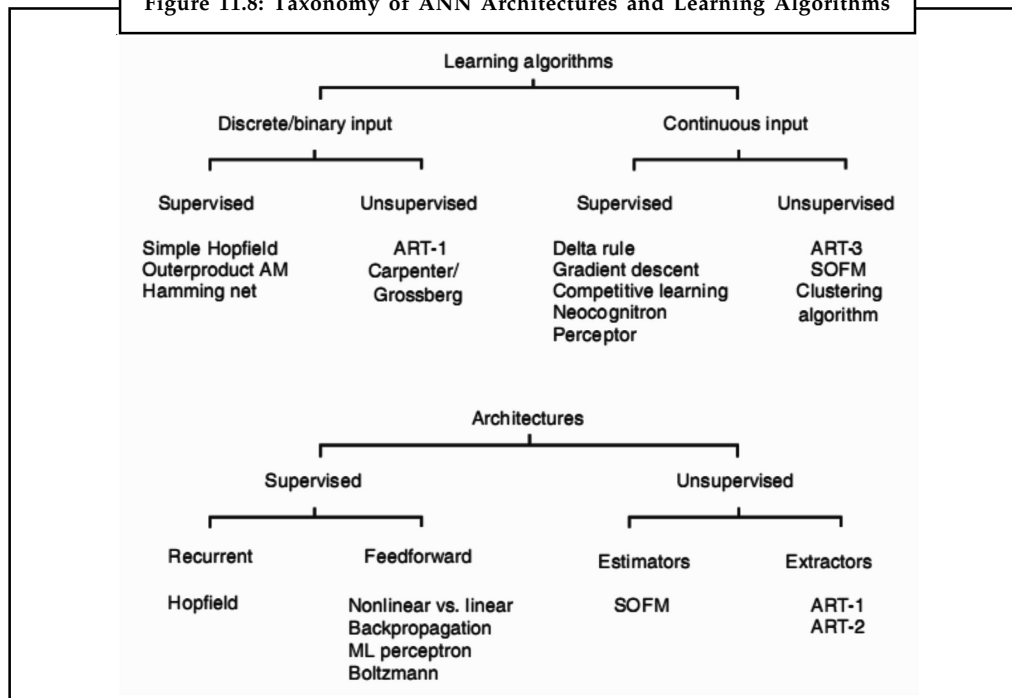
Notes

Updating of the network's interconnection weights continues until the stopping criteria of the training algorithm are met (e.g., all cases must be correctly classified within a certain tolerance level). Alternatively, in unsupervised learning, there are no target answers that the network tries to learn. Instead, the neural network learns a pattern through repeated exposure. Thus, this kind of learning can be envisioned as the neural network appropriately self-organizing or clustering its neurons related to the specific desired task. Multilayer, feedforward neural networks are a class of models that show promise in classification and forecasting problems. As the name implies, these models structurally consist of multiple layers of neurons. Information is passed through the network in one direction, from the input layers of the network, through one or more hidden layers, toward the output layer of neurons. Neurons of each layer are connected only to the neurons of the subsequent layer.

11.1.6 Learning in ANN

An important consideration in an ANN is the use of an appropriate learning algorithm (or training algorithm). Learning algorithms specify the process by which a neural network learns the underlying relationship between input and outputs, or just among the inputs. There are hundreds of them. Learning algorithms in ANN can also be classified as supervised learning and unsupervised learning (see Figure 11.8).

Figure 11.8: Taxonomy of ANN Architectures and Learning Algorithms



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

Supervised learning uses a set of inputs for which the appropriate (i.e., desired) outputs are known.

Notes



Example: A historical set of loan applications with the success or failure of the individual to repay the loan has a set of input parameters and presumed known outputs.

In one type, the difference between the desired and actual outputs is used to calculate corrections to the weights of the neural network. A variation of this approach simply acknowledges for each input trial whether the output is correct as the network adjusts weights in an attempt to achieve correct results.

Examples of this type of learning are backpropagation and the Hopfield network.

In unsupervised learning, only input stimuli are shown to the network. The network is self-organizing; that is, it organizes itself internally so that each hidden processing element responds strategically to a different set of input stimuli (or groups of stimuli). No knowledge is supplied about which classifications (i.e., outputs) are correct, and those that the network derives may or may not be meaningful to the network developer (this is useful for cluster analysis). However, by setting model parameters, we can control the number of categories into which a network classifies the inputs. Regardless, a human must examine the final categories to assign meaning and determine the usefulness of the results.

Examples of this type of learning are Adaptive Resonance Theory (ART) (i.e., a neural network architecture that is aimed at being brain-like in unsupervised mode) and Kohonen self-organizing feature maps (i.e., neural network models for machine learning).

The backpropagation learning algorithm is the standard way of implementing supervised training of feedforward neural networks. It is an iterative gradient-descent technique designed to minimize an error function between the actual output of the network and its desired output, as specified in the training set of data. Adjustment of the interconnection weights, which contain the mapping function per se, starts at the output node where the error measure is initially calculated and is then propagated back through the layers of the network, toward the input layer.

Self Assessment

Fill in the blanks:

1. represent a brain metaphor for information processing.
2. ANN are composed of interconnected, simple processing elements called
3. Neurons are partitioned into groups called
4. A is able to increase or decrease the strength of the connection from neuron to neuron.
5. The propagation paradigm, allows all neurons to link the output in one layer to the input of the next layer, but it does not allow any feedback linkage.
6. A is a layer of neurons that takes input from the previous layer and converts those inputs into outputs for further processing.
7. The of a network contain the solution to a problem.
8. express the relative strength of the input data or the many connections that transfer data from layer to layer.
9. The function computes the weighted sums of all the input elements entering each processing element.

10. The function combines the inputs coming into a neuron from other neurons/sources and then produces an output based on the choice of the transfer function.
11. A is a hurdle value for the output of a neuron to trigger the next level of neurons.
12. The learning algorithm is the standard way of implementing supervised training of feedforward neural networks.

11.2 Development of Neural Network-based System

Although the development process of ANN is similar to the structured design methodologies of traditional computer-based information systems, some phases are unique or have some unique aspects. In the process described here, we assume that the preliminary steps of system development, such as determining information requirements, conducting a feasibility analysis, and gaining a champion in top management for the project, have been completed successfully. Such steps are generic to any information system.

As shown in Figure 11.9, the development process for an ANN application includes nine steps. In step 1, the data to be used for training and testing the network are collected. Important considerations are that the particular problem is amenable to neural network solution and that adequate data exist and can be obtained. In step 2, training data must be identified, and a plan must be made for testing the performance of the network.

In steps 3 and 4, a network architecture and a learning method are selected. The availability of a particular development tool or the capabilities of the development personnel may determine the type of neural network to be constructed. Also, certain problem types have demonstrated high success rates with certain configurations (e.g., multilayer feedforward neural networks for bankruptcy prediction). Important considerations are the exact number of neurons and the number of layers. Some packages use genetic algorithms to select the network design.

There are parameters for tuning the network to the desired learning-performance level. Part of the process in step 5 is the initialization of the network weights and parameters, followed by the modification of the parameters as training-performance feedback is received. Often, the initial values are important in determining the efficiency and length of training. Some methods change the parameters during training to enhance performance.

Step 6 transforms the application data into the type and format required by the neural network. This may require writing software to preprocess the data or performing these operations directly in an ANN package.



Caution Data storage and manipulation techniques and processes must be designed for conveniently and efficiently retraining the neural network, when needed.

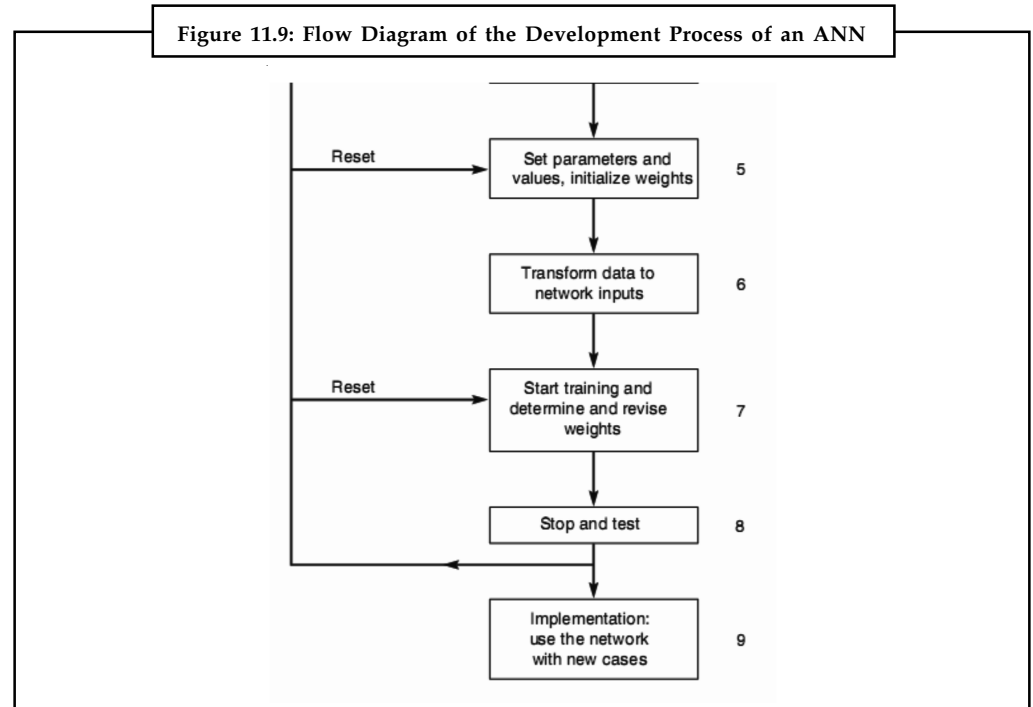
The application data representation and ordering often influence the efficiency and possibly the accuracy of the results.

In steps 7 and 8, training and testing are conducted iteratively by presenting input and desired or known output data to the network. The network computes the outputs and adjusts the weights until the computed outputs are within an acceptable tolerance of the known outputs for the input cases. The desired outputs and their relationships to input data are derived from historical data (i.e., a portion of the data collected in step 1).

In step 9, a stable set of weights is obtained. Now the network can reproduce the desired outputs, given inputs such as those in the training set. The network is ready for use as a stand-alone

Notes

system or as part of another software system where new input data will be presented to it and its output will be a recommended decision. In the following sections, we examine these steps in more detail.



Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

11.2.1 Data Collection and Preparation

The first two steps in the ANN development process involve collecting data and separating them into a training set and a testing set. The training cases are used to adjust the weights, and the testing cases are used for network validation. The data used for training and testing must include all the attributes that are useful for solving the problem. The system can only learn as much as the data can tell.



Did u know? Collection and preparation of data is the most critical step in building a good system.

11.2.2 Selection of Network Structure

After the training and testing data sets are identified, the next step is to design the structure of the neural networks. This includes the selection of a topology and determination of

- input nodes,
- output nodes,
- number of hidden layers, and
- number of hidden nodes.

The multilayer feedforward topology is often used in business applications, although other network models are beginning to find some business use as well.

11.2.3 Learning Algorithm Selection

Notes

After the network structure is chosen, we need to find a learning algorithm to identify a set of connection weights that best cover the training data and have the best predictive accuracy. For the feedforward topology we chose for the bankruptcy-prediction problem, a typical approach is to use the backpropagation algorithm. Because many commercial packages are available on the market, there is no need to implement the learning algorithm by ourselves. Instead, we can choose a suitable commercial package to analyze the data.

11.2.4 Network Training

Training of ANN is an iterative process that starts from a random set of weights and gradually enhances the fitness of the network model and the known data set. The iteration continues until the error sum is converged to below a preset acceptable level. In the backpropagation algorithm, two parameters, learning rate and momentum, can be adjusted to control the speed of reaching a solution. These determine the ratio of the difference between the calculated value and the actual value of the training cases. Some software packages may have their own parameters in their learning heuristics to speed up the learning process.



Caution It is important to read carefully when using this type of software.

11.2.5 Testing

Recall that in step 2 of the development process shown in Figure 11.9, the available data are divided into training and testing data sets. When the training has been completed, it is necessary to test the network. Testing (step 8) examines the performance of the derived network model by measuring its ability to classify the testing data correctly. Black-box testing (i.e., comparing test results to historical results) is the primary approach for verifying that inputs produce the appropriate outputs. Error terms can be used to compare results against known benchmark methods.



Task Analyze the importance of testing.

11.2.6 Implementation of an ANN

Implementation of an ANN (step 9) often requires interfaces with other computer based information systems and user training. Ongoing monitoring and feedback to the developers are recommended for system improvements and long-term success. It is also important to gain the confidence of users and management early in the deployment to ensure that the system is accepted and used properly.

Self Assessment

Fill in the blanks:

13. of ANN is an iterative process that starts from a random set of weights and gradually enhances the fitness of the network model and the known data set.
14. is the primary approach for verifying that inputs produce the appropriate outputs.

Notes

15. of an ANN often requires interfaces with other computer based information systems and user training.



Case Study

Neural Networks for Breast Cancer Diagnosis

ANN have proven to be a useful tool in pattern recognition and classification tasks in diverse areas, including clinical medicine. Despite the wide applicability of ANN, the large amount of data required for training makes using them an unsuitable classification technique when the available data are scarce. Magnetic Resonance Spectroscopy (MRS) plays a pivotal role in the investigation of cell biochemistry and provides a reliable method for detection of metabolic changes in breast tissue. The scarcity of data and the complexity of interpretation of relevant physiological information impose extra demands that prohibit the applicability of most statistical and machine learning techniques developed.

Knowledge-based artificial neural networks (KBANN) help to prevail over such difficulties and complexities. A KBANN combines knowledge from a domain, in the form of simple rules, with connectionist learning. This combination trains the network through the use of small sets of data (as is typical of medical diagnosis tasks). The primary structure is based on the dependencies of a set of known domain rules, and it is necessary to refine those rules through training. The KBANN process consists of two algorithms.

One is the Rules-to-Network algorithm, in which the main task is the translation process between a knowledge base containing information about a domain theory and the initial structure of a neural network. This algorithm maps the structure of an approximately correct domain theory, with all the rules and their dependencies, into a neural network structure. The defined network is then trained using the backpropagation learning algorithm.

Feedback mechanisms, which inhibit or stimulate the growth of normal cells, control the division and replacement of cells in normal tissues. In the case of tumors, that process is incapable of controlling the production of new cells, and the division is done without any regard to the need for replacement, disturbing the structure of normal tissue. Changes observed in phospholipid metabolite concentrations, which are associated with differences in cell proliferation in malignant tissues, have served as the basic inputs for the identification of relevant features present in malignant or cancerous tissues but not in normal tissues. The abnormal levels of certain phospholipid characteristics are considered indicators of tumors. These include several parameters, such as PDE, PME, Pi, PCr, γ ATP, α ATP, and β ATP. KBANN produced an accurate tumor classification of 87 percent from a set of 26, with an average pattern error of 0.0500 and a standard deviation of 0.0179.

Question

Discuss the role of Magnetic resonance spectroscopy (MRS) in the investigation of cell biochemistry.

Source: http://www70.homepage.villanova.edu/matthew.liberatore/Mgt2206/turban_online_ch06.pdf

11.3 Summary

- Neural networks represent a brain metaphor for information processing.
- The resulting model from neural computing is often called an artificial neural network (ANN) or a neural network.

- More or less resembling the structure of their counterparts, ANN are composed of interconnected, simple processing elements called artificial neurons.
- An ANN model emulates a biological neural network. Neural computing actually uses a very limited set of concepts from biological neural systems.
- A neural network is composed of processing elements organized in different ways to form the network's structure.
- Complex practical applications require one or more hidden layers between the input and output neurons and a correspondingly large number of weights.
- There are several effective neural network models and algorithms. Some of the most common are backpropagation (or feedforward), associative memory, and the recurrent network.
- Learning algorithms specify the process by which a neural network learns the underlying relationship between input and outputs, or just among the inputs.
- Although the development process of ANN is similar to the structured design methodologies of traditional computer-based information systems, some phases are unique or have some unique aspects.

11.4 Keywords

ANN: An artificial neural network (ANN), often just called a “neural network” (NN), is a mathematical model or computational model based on biological neural networks.

Backpropagation: The backpropagation learning algorithm is the standard way of implementing supervised training of feedforward neural networks.

Connection Weights: Connection weights express the relative strength of the input data or the many connections that transfer data from layer to layer.

Feedforward Neural Network: Feedforward neural networks are a class of models that show promise in classification and forecasting problems.

Hidden Layer: A hidden layer is a layer of neurons that takes input from the previous layer and converts those inputs into outputs for further processing.

Neurons: The human brain is composed of special cells called neurons.

Summation Weights: The summation function computes the weighted sums of all the input elements entering each processing element.

Transformation Function: The transformation (transfer) function combines the inputs coming into a neuron from other neurons/sources and then produces an output based on the choice of the transfer function

11.5 Review Questions

1. Explain the concept of ANN.
2. Describe the following terms: neuron, axon and synapse.
3. How do weights function in an ANN? Discuss.
4. What is the role of inputs and outputs in neural networks?
5. Discuss the steps in developing a neural network.
6. Describe the Elements of ANN.

Notes

7. Explain the concept of hidden layers in neural networks.
8. Discuss the major concepts related to Network Information Processing.
9. Illustrate the backpropagation architecture of neural network.
10. Each network contains several thousand highly interconnected neurons. Comment.

Answers: Self Assessment

- | | |
|---------------------|-----------------------|
| 1. Neural networks | 2. Neural networks |
| 3. Networks | 4. Synapse |
| 5. Feedforward back | 6. Hidden layer |
| 7. Outputs | 8. Connection weights |
| 9. Summation | 10. Transformation |
| 11. Threshold value | 12. Backpropagation |
| 13. Training | 14. Black-box testing |
| 15. Implementation | |

11.6 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

<http://nordbotten.com/articles/NN.pdf>

<http://web.iitd.ac.in/~sumeet/Jain.pdf>

<http://www.jatit.org/volumes/research-papers/Vol5No1/1Vol5No6.pdf>

<http://www.uncg.edu/ism/ism611/neuralnet.pdf>

Unit 12: Applications of Neural Network

Notes

CONTENTS

Objectives

Introduction

12.1 Applications of ANN

12.2 Other Applications

12.2.1 Pattern Recognition

12.2.2 Evolution of Neural Networks for Control of Pursuit & Evasion

12.2.3 Learning the Distribution of Object Trajectories for Event Recognition

12.2.4 Radiosity for Virtual Reality Systems (ROVER)

12.2.5 Using HMM's for Audio-to-Visual Conversion

12.2.6 Artificial Life: Galapagos

12.2.7 Speechreading (Lipreading)

12.2.8 Detection and Tracking of Moving Targets

12.2.9 Real-time Target Identification for Security Applications

12.2.10 A Three Layer Feedforward Neural Network

12.2.11 Artificial Life for Graphics, Animation, Multimedia and Virtual Reality

12.3 Summary

12.4 Keywords

12.5 Review Questions

12.6 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss Applications of Neural networks
- Discuss Various Application such as Pattern Recognition, Galapagos, etc.

Introduction

ANN have been applied in many domains. There have been several tests of neural networks in financial markets. Collard (1990) stated that his neural network model for commodity trading would have resulted in significant profits over other trading strategies. Kamijo and Tanigawa (1990) used a neural network to chart Tokyo Stock Exchange data. They found that the results of the model would beat a "buy and hold" strategy. Finally, a neural model for predicting percentage change in the S&P 500 five days ahead, using a variety of economic indicators, was developed. It is claimed that the model has provided more accurate prediction than alleged experts in the field using the same indicators. In this unit, we will discuss various applications of neural networks.

12.1 Applications of ANN

Neural networks have been successfully trained to determine whether loan applications should be approved. It has also been shown that neural networks can predict mortgage applicant solvency better than mortgage writers. Predicting rating of corporate bonds and attempting to predict their profitability is another area where neural networks have been successfully applied. Neural networks outperformed regression analysis and other mathematical modeling tools in predicting bond rating and profitability. The main conclusion reached was that neural networks provided a more general framework for connecting financial information of a firm to the respective bond rating.

Fraud prevention is another area of neural network application in business. Chase Manhattan Bank successfully used neural networks in dealing with credit card fraud (Rochester, 1990), with the neural network models outperforming traditional regression approaches. Also, neural networks have been used in the validation of bank signatures. These networks identified forgeries significantly better than any human expert. Another significant area of statistical application of neural networks is in time series forecasting. Several studies have attempted to use neural networks for time-series prediction.

Examples include Fozzard et al. (1989), Tang et al. (1991), and Hill et al. (1994).

The general conclusion is that neural networks appear to do at least as well as the Box-Jenkins forecasting technique.

Because neural networks have been a subject of intense study since late 1980s, there have been many applications as well as experiments with applications. Other recent reports include live intrusion tracking , Web content filtering, exchange rate prediction, and hospital bed allocation. Newer applications are emerging in health care and medicine.

In general, ANN are suitable for problems whose inputs are both categorical and numeric, and where the relationships between inputs and outputs are not linear or the input data are not normally distributed. In such cases, classical statistical methods may not be reliable enough. Because ANN do not make any assumptions about the data distribution, their power is less affected than traditional statistical methods when data are not properly distributed. Finally, there are cases in which the neural networks simply provide one more way of building a predictive model for the situation at hand.



Did u know? Given the ease of experimentation using the available software tools, it is certainly worth exploring the power of neural networks in any data modeling situation.

Self Assessment

State True or False:

1. Neural networks outperformed regression analysis and other mathematical modeling tools in predicting bond rating and profitability.
2. Neural networks are not used in dealing with credit card fraud .
3. We can use neural networks for time-series prediction.
4. ANN are not suitable for problems whose inputs are both categorical and numeric.

12.2 Other Applications

In this section, we will discuss various applications of ANN.

12.2.1 Pattern Recognition


Notes

Pattern Recognition can be defined as the act of taking in raw data and taking an action based on the category of the pattern. Pattern recognition aims to classify data (patterns) based either on a priori knowledge or on statistical information extracted from the patterns. The patterns to be classified are usually groups of measurements or observations, denning points in an appropriate multidimensional space. This is in contrast to pattern matching, where the pattern is rigidly specified.

Artificial neural nets can be used for pattern recognition (a task in which humans excel machines!) The first step is to take an image and convert it to a pixel map (a map of squares with different intensity of gray color). This pixel map is a two dimensional function with 2 input variables: pixel position and 1 output variable: intensity of the pixel. We will use standard ANN techniques to train this 2D function.

Any image can be represented as two-dimensional array, where every element of that array contains color information for one pixel.


Figure 12.1: Image Represented as Two-dimensional Array



Source: <http://www.teor.mi.infn.it/~rojo/teaching/Milano-NNet-Course-Lecture4.pdf>

This information is the input patterns for the artificial neural network training.

Figure 12.2: Artificial Neural Network



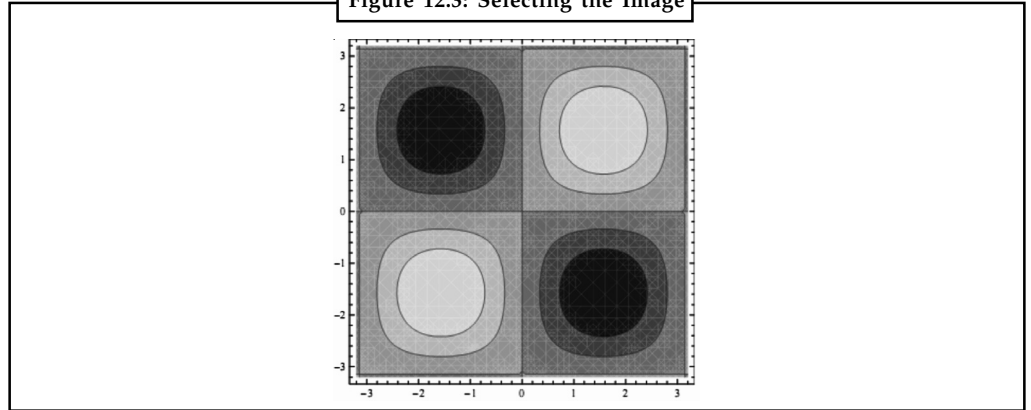
Source: <http://www.teor.mi.infn.it/~rojo/teaching/Milano-NNet-Course-Lecture4.pdf>

Notes



Example: Select the image which we want the NN to learn. For simplicity start from a simple geometrical image.

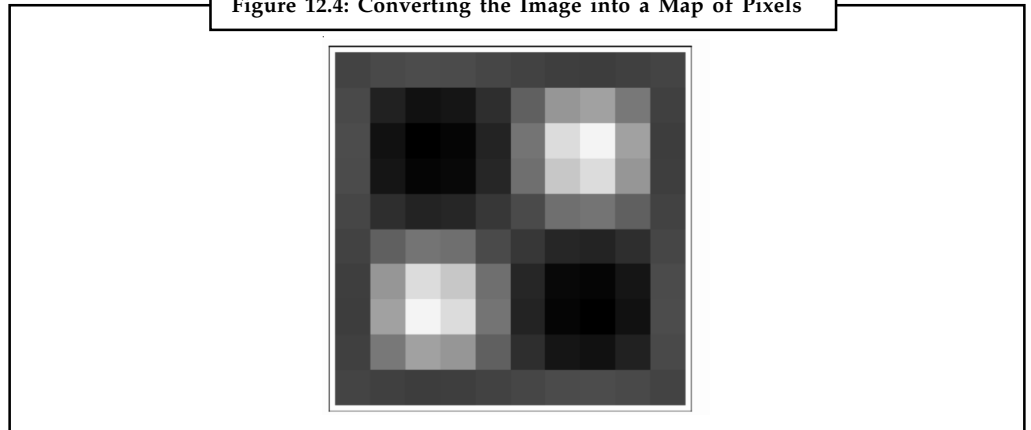
Figure 12.3: Selecting the Image



Source: <http://www.teor.mi.infn.it/~rojo/teaching/Milano-NNet-Course-Lecture4.pdf>

Convert the image into a map of pixels (with a given resolution).


Figure 12.4: Converting the Image into a Map of Pixels



Source: <http://www.teor.mi.infn.it/~rojo/teaching/Milano-NNet-Course-Lecture4.pdf>

Initialize at random the neural network.

Figure 12.5: Initialize at Random



Source: <http://www.teor.mi.infn.it/~rojo/teaching/Milano-NNet-Course-Lecture4.pdf>

Now train the neural network with the error function:

Notes

$$E[\omega, \theta] \equiv \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} (o(x_i, y_j) - z_{ij})^2$$

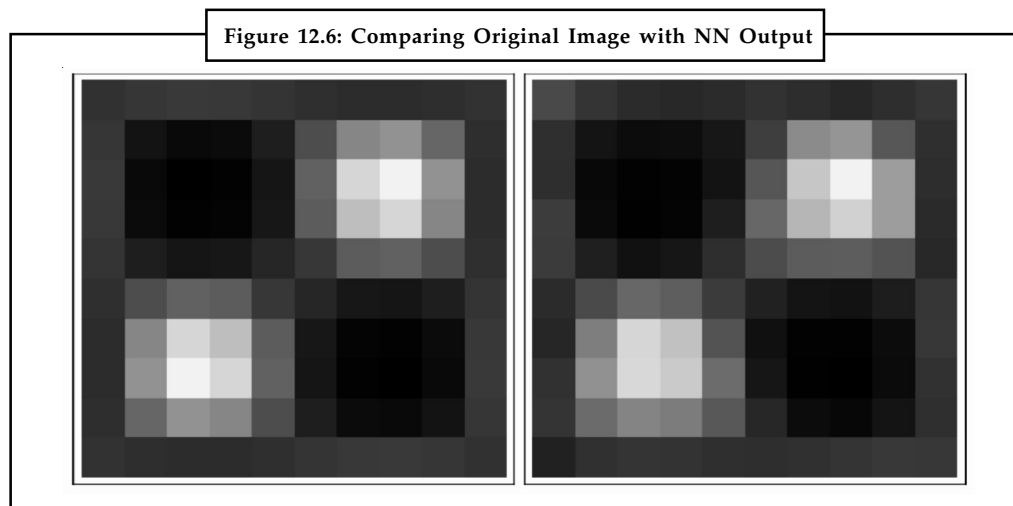
Where,

- n_x, n_y are the dimensions (number of pixels) of the image in the x, y directions
- Each pixel is specified by x_i, y_j, z_{ij} where z_{ij} is the gray intensity of the pixel
- $o(x_i, y_j)$ is the output of the neural network when the input is the position of the pixel i, j .

More complex images will require larger architectures.

Check the results of the trained neural network.

Compare original (left) with NN output (right):



Source: <http://www.teor.mi.infn.it/~rojo/teaching/Milano-NNet-Course-Lecture4.pdf>

Using Image Recognition

Once the NN has been trained, we can use it for image recognition. This can be done by comparing the error function between the trained ANN and another different image → Below a given threshold (to be validated) the two images are assigned to be identical.



Caution The trained ANN works also as compressed version of the image, where all the information of the image is now encoded into the weights and thresholds.


12.2.2 Evolution of Neural Networks for Control of Pursuit & Evasion

The following MPEG movie sequences illustrate behaviour generated by dynamical recurrent neural network controllers co-evolved for pursuit and evasion capabilities. From an initial population of random network designs, successful designs in each generation are selected for reproduction with recombination, mutation, and gene duplication. Selection is based on measures of how well each controller performs in a number of pursuit-evasion contests. In each contest a pursuer controller and an evader controller are pitched against each other, controlling simple

Notes

“visually guided” 2-dimensional autonomous virtual agents. Both the pursuer and the evader have limited amounts of energy, which is used up in movement, so they have to evolve to move economically. Each contest results in a time-series of position and orientation data for the two agents.



These time-series are then fed into a custom 3-D movie generator.



Notes It is important to note that, although the chase behaviors are genuine data, the 3D structures, surface physics, and shading are all purely for illustrative effect.

12.2.3 Learning the Distribution of Object Trajectories for Event Recognition

This research work is about the modelling of object behaviours using detailed, learnt statistical models. The techniques being developed will allow models of characteristic object behaviours to be learnt from the continuous observation of long image sequences. It is hoped that these models of characteristic behaviours will have a number of uses, particularly in automated surveillance and event recognition, allowing the surveillance problem to be approached from a lower level, without the need for high-level scene/behavioural knowledge. Other possible uses include the random generation of realistic looking object behaviour for use in Virtual Reality, and long-term prediction of object behaviours to aid occlusion reasoning in object tracking.



Source: <http://travex.com/pub/nap/#CoEvolution> of Neural Networks for Control of Pursuit & Evasion

In figure 12.7, the model is learnt in an unsupervised manner by tracking objects over long image sequences, and is based on a combination of a neural network implementing Vector Quantization and a type of neuron with short-term memory capabilities.

In figure 12.8, Models of the trajectories of pedestrians have been generated and used to assess the typicality of new trajectories (allowing the identification of ‘incidents of interest’ within the scene), predict future object trajectories, and randomly generate new trajectories.

12.2.4 Radiosity for Virtual Reality Systems (ROVER)

The synthesis of actual and computer generated photo-realistic images has been the aim of artists and graphic designers for many decades. Some of the most realistic images were generated using radiosity techniques. Unlike ray tracing, radiosity models the actual interaction between the lights and the environment. In photo realistic Virtual Reality (VR) environments, the need for quick feedback based on user actions is crucial.



Did u know? It is generally recognised that traditional implementation of radiosity is computationally very expensive and therefore not feasible for use in VR systems where practical data sets are of huge complexity.

In the original thesis, we introduce two new methods and several hybrid techniques to the radiosity research community on using radiosity in VR applications.

For example, in figure 12.9, flyby, walkthrough and a virtual space are first introduced and in figure 12.10, we showcase one of the two novel methods which was proposed using Neural Network technology.

Figure 12.9: Introduction to Flyby, Walkthrough and Virtual Space





Figure 12.10: ROVER



Source: [http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion](http://tralvex.com/pub/nap/#CoEvolution%20of%20Neural%20Networks%20for%20Control%20of%20Pursuit%20&%20Evasion)


12.2.5 Using HMM's for Audio-to-Visual Conversion

One emerging application which exploits the correlation between audio and video is speech-driven facial animation. The goal of speech-driven facial animation is to synthesize realistic video sequences from acoustic speech. Much of the previous research has implemented this audio-to-visual conversion strategy with existing techniques such as vector quantization and neural networks. Here, they examine how this conversion process can be accomplished with hidden Markov models (HMM).

- (a) **Tracking Demo:** The parabolic contour is fit to each frame of the video sequence using a modified deformable template algorithm. The height between the two contours, and the width between the corners of the mouth can be extracted from the templates to form our visual parameter sets.

Notes

Figure 12.11: Tracking



Source: <http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion>

- (b) **Morphing Demo:** Another important piece of the speech-driven facial animation system is a visual synthesis module. Here we are attempting to synthesize the word “wow” from a single image. Each frame in the video sequence is morphed from the first frame shown below. The parameters used to morph these images were obtained by hand.

Figure 12.12: Morphing



Source: <http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion>



Task Compare and contrast tracking demo and morphing demo.

12.2.6 Artificial Life: Galapagos

Galapagos is a fantastic and dangerous place where up and down have no meaning, where rivers of iridescent acid and high-energy laser mines are beautiful but deadly artifacts of some other time. Through spatially twisted puzzles and bewildering cyber-landscapes, the artificial creature called Mendel struggles to survive, and you must help him.

Mendel is a synthetic organism that can sense infrared radiation and tactile stimulus. His mind is an advanced adaptive controller featuring Non-stationary Entropic Reduction Mapping — a new form of artificial life technology developed by Anark. He can learn like your dog, he can adapt to hostile environments like a cockroach, but he can't solve the puzzles that prevent his escape from Galapagos.

Galapagos features rich, 3D texture-mapped worlds, with continuous-motion graphics and 6 degrees of freedom. Dramatic camera movement and incredible lighting effects make your passage through Galapagos breathtaking. Explosions and other chilling effects will make you fear for your synthetic friend. Active panning 3D stereo sound will draw you into the exotic worlds of Galapagos.

12.2.7 Speechreading (Lipreading)

Notes

As part of the research program Neuroinformatik the IPVR develops a neural speechreading system as part of a user interface for a workstation. The three main parts of the system include a face tracker (done by Marco Sommerau), lip modeling and speech processing (done by Michael Vogt) and the development and application of SNNS for neural network training (done by Günter Mamier).

Automatic speechreading is based on a robust lip image analysis. In this approach, no special illumination or lip make-up is used. The analysis is based on true color video images. The system allows for realtime tracking and storage of the lip region and robust off-line lip model matching. The proposed model is based on cubic outline curves. A neural classifier detects visibility of teeth edges and other attributes. At this stage of the approach the edge between the closed lips is automatically modeled if applicable, based on a neural network's decision.


To achieve high flexibility during lip-model development, a model description language has been defined and implemented. The language allows the definition of edge models (in general) based on knots and edge functions. Inner model forces stabilize the overall model shape. User defined image processing functions may be applied along the model edges. These functions and the inner forces contribute to an overall energy function.



Caution Adaptation of the model is done by gradient descent or simulated annealing like algorithms.

The figure shows one configuration of the lip model, consisting of an upper lip edge and a lower lip edge. The model edges are defined by Bezier-functions. Outer control knots stabilize the position of the corners of the mouth.

Figure 12.13: Configuration of the Lip Model



Source: [http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion](http://tralvex.com/pub/nap/#CoEvolution%20of%20Neural%20Networks%20for%20Control%20of%20Pursuit%20&%20Evasion)

The model interpreter enables a permanent measurement of model knot positions and color blends along model edges during adaptation to an utterance.

The resulting parameters may be used for speech recognition tasks in further steps.



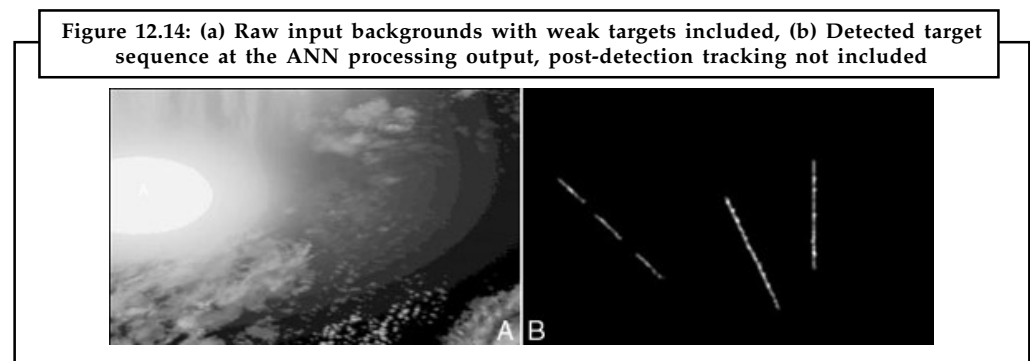
Task Analyze the use of Bezier-functions.

12.2.8 Detection and Tracking of Moving Targets

The moving target detection and track methods here are “track before detect” methods. They correlate sensor data versus time and location, based on the nature of actual tracks. The track statistics are “learned” based on Artificial Neural Network (ANN) training with prior real or simulated data. Effects of different clutter backgrounds are partially compensated based on space-time-adaptive processing of the sensor inputs, and further compensated based on the ANN training. Specific processing structures are adapted to the target track statistics and sensor characteristics of interest. Fusion of data over multiple wavelengths and sensors is also supported.

Compared to conventional fixed matched filter techniques, these methods have been shown to reduce false alarm rates by up to a factor of 1000 based on simulated SBIRS data for very weak ICBM targets against cloud and nuclear backgrounds, with photon, quantization, and thermal noise, and sensor jitter included.

Examples of the backgrounds, and processing results, are given below.



Source: [http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion](http://tralvex.com/pub/nap/#CoEvolution%20of%20Neural%20Networks%20for%20Control%20of%20Pursuit%20&%20Evasion)

The methods are designed to overcome the weaknesses of other advanced track-before-detect methods, such as 3+-D (space, time, etc.) matched filtering, dynamic programming (DP), and multi-hypothesis tracking (MHT). Loosely speaking, 3+-D matched filtering requires too many filters in practice for long-term track correlation; DP cannot realistically exploit the non-Markovian nature of real tracks, and strong targets mask out weak targets; and MHT cannot support the low pre-detection thresholds required for very weak targets in high clutter. They have developed and tested versions of the above (and other) methods in their research, as well as Kalman-filter probabilistic data association (KF/PDA) methods, which they use for post-detection tracking.

Space-time-adaptive methods are used to deal with correlated, non-stationary, non-Gaussian clutter, followed by a multi-stage filter sequence and soft-thresholding units that combine current and prior sensor data, plus feed back of prior outputs, to estimate the probability of target presence. The details are optimized by adaptive “training” over very large data sets, and special methods are used to maximize the efficiency of this training.

12.2.9 Real-time Target Identification for Security Applications

The system localises and tracks peoples’ faces as they move through a scene. It integrates the following techniques:

- Motion detection
- Tracking people based upon motion
- Tracking faces using an appearance model

Faces are tracked robustly by integrating motion and model-based tracking.

Notes

(a) Tracking in low resolution and poor lighting conditions.

Figure 12.15: Tracking in Low Resolution



Source: [http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion](http://tralvex.com/pub/nap/#CoEvolution%20of%20Neural%20Networks%20for%20Control%20of%20Pursuit%20&%20Evasion)

(b) Tracking two people simultaneously: lock is maintained on the faces despite unreliable motion-based body tracking.

Figure 12.16: Tracking Two People Simultaneously




Source: [http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion](http://tralvex.com/pub/nap/#CoEvolution%20of%20Neural%20Networks%20for%20Control%20of%20Pursuit%20&%20Evasion)

12.2.10 A Three Layer Feedforward Neural Network

A three layer feedforward neural network with two input nodes and one output node is trained with backpropagation using some sample points inside a circle in the 2D plane. The evolution of decision regions formed during the training are shown in the following MPEG movies.

Figure 12.17: Sequence 3, Sequence 10




Source: [http://tralvex.com/pub/nap/#CoEvolution of Neural Networks for Control of Pursuit & Evasion](http://tralvex.com/pub/nap/#CoEvolution%20of%20Neural%20Networks%20for%20Control%20of%20Pursuit%20&%20Evasion)

Notes

12.2.11 Artificial Life for Graphics, Animation, Multimedia and Virtual Reality

Some graphics researchers have begun to explore a new frontier—a world of objects of enormously greater complexity than is typically accessible through physical modeling alone—objects that are alive. The modeling and simulation of living systems for computer graphics resonates with the burgeoning field of scientific inquiry called Artificial Life. Conceptually, artificial life transcends the traditional boundaries of computer science and biological science. The natural synergy between computer graphics and artificial life can be potentially beneficial to both disciplines. As some of the demos here demonstrate, potential is becoming fulfillment.

The demos demonstrate and elucidate new models that realistically emulate a broad variety of living things—both plants and animals—from lower animals all the way up the evolutionary ladder to humans. Typically, these models inhabit virtual worlds in which they are subject to physical laws. Consequently, they often make use of physics-based modeling techniques. More significantly, however, they must also simulate many of the natural processes that uniquely characterize living systems—such as birth and death, growth, natural selection, evolution, perception, locomotion, manipulation, adaptive behavior, intelligence and learning.



Notes The challenge is to develop sophisticated graphics models that are self-creating, self-evolving, self-controlling, and/or self-animating by simulating the natural mechanisms fundamental to life.

Self Assessment

Fill in the blanks:

5. can be defined as the act of taking in raw data and taking an action based on the category of the pattern.
6. The is a two dimensional function with 2 input variables: pixel position and 1 output variable: intensity of the pixel.
7. Any image can be represented as array, where every element of that array contains color information for one pixel.
8. Selection is based on measures of how well each performs in a number of pursuit-evasion contests.
9. Unlike ray tracing, models the actual interaction between the lights and the environment.
10. In photo realistic environments, the need for quick feedback based on user actions is crucial.
11. The goal of facial animation is to synthesize realistic video sequences from acoustic speech.
12. In case of, each frame in the video sequence is morphed from the first frame.
13. is a fantastic and dangerous place where up and down have no meaning, where rivers of iridescent acid and high-energy laser mines are beautiful but deadly artifacts of some other time.
14. is a synthetic organism that can sense infrared radiation and tactile stimulus.
15. Automatic is based on a robust lip image analysis.



Case Study

Artificial Neural Network (ANN) Approach for an Intelligent System

Artificial Neural Networks and Cognition

ANNs have massive parallelism, high connectivity, which tries to emulate the Biological neurons and their synapses. NNs are especially useful for classification and function approximation/mapping problems. The networks use Input/Output layers and none or few hidden layers (HLs). Every hidden and output unit has its own bias term Learning algorithms can be supervised or unsupervised. In supervised learning, the correct results (desired outputs) are known and are given to the NN during training so that the NN can adjust its weights to try match its outputs to the target values. After training, NN is tested by giving only input values, to see how closely it predicts target values. LM, Momentum Step are common learning algorithms. Performances of different network topologies can be compared by evaluating the error function. Various networks are trained by minimization of an appropriate error function defined with respect to a training data set. Performance of networks are compared by evaluating the error function using an independent cross-validation (CV) set.

Network having smallest error with respect to CV set is selected (hold out method). Performance of the selected network is confirmed using a third independent set of data called a test set to avoid overfitting.

Multilayer Perceptron (MLP)

MLP has one or more HLs with a linear combination function is the inner product of inputs, weights and a bias. Activation function is logistic or tanh function. Inputs are fully connected to the first HL, each HL is fully connected to the next, and the last HL is fully connected to the outputs. MLPs use supervised learning or backpropagation. Designing and training an MLP requires (i) selecting number of HLs, (ii) number of neurons to be used in each HL, (iii) avoiding local minima, (iv) converging to an optimal solution in a reasonable period of time, (v) validating NN to test for overfitting. MLPs are used in classification problems, in fitness approximations.

Linear-Regression

Linear regression assumes that expected value of the output given an input, is linear. BAYES rule is used to obtain posterior distribution for given data. Maximum likelihood estimation helps in predictions.

Tonic-Shift-Property

CCM is a highly scientific and evolved form of classical music which uses a heptatonic scale of seven notes symbolically given as S, R, G, M, P, D, N in ascending order of frequencies. It is not equi-tempered scale like the Western, since notes have varying ratios with their successive notes. Two notes S and P are inherently stable with no variants. The remaining 5 notes R, G, M, D, N have variants. M has 2 variants -M1 and M2. They result in 12 semi-tones as shown in TABLE 1. In CCM a full scale (heptatonic) consists of all the 7 notes in both ascending and descending called a 'melakartha' ('Janaka') raga if i) all the 7 notes in ascending and descending are identical.

Contd...

Notes

Table 1: Relative Frequencies of the 12 Semi-tones [6]-[8],[15]

| S | R1 | R2 | G2 | G3 | M1 | M2 | P | D1 | D2 | N2 | N3 |
|---|---------|-----|---------|-------|-----|---------|-----|--------|-------|------|---------|
| 1 | 256/243 | 9/8 | 192/162 | 81/64 | 4/3 | 729/512 | 3/2 | 128/81 | 27/16 | 16/9 | 243/128 |

Frequencies Occur in Order

Exactly 72 such scales are possible. They are given characteristic names and are serially arranged in a table. Beauty of a raga delineation takes a new dimension when performer sings or plays another raga by staying in the domain of the original raga. This is possible by temporarily shifting the reference note to a new note existing in the original raga. Hence all the remaining semi-tones will take different relative values leading to a new raga. This is called ‘tonic-shift’. Every such shift may not give a new raga. For example, 29th melakartha raga called ‘Dheera Shankarabharana’ has the scale S, R2, G3, M1, P, D2, N3., i.e., relative frequencies are 1 9/8 81/64 4/3 3/2 27/16 and 243/128. Shifting reference from S to R2 makes the frequency ratios of R2, G3, M1, P, D2, N3 and upper octave S,=1, 9/8 192/162 4/3 3/2 27/16 and 16/9 respectively. This corresponds to another new raga. Continuing the process of shifting to various notes give rise to 4 other ragas using G3, M1, P and D2. When reference is shifted to N3, relative frequencies give rise to both M1 and M2 in the new scale which does not correspond to any valid scale.

Formulation of the Problem

CCM being highly scientific, shows very interesting phenomena which are intriguing researchers from a long time. One such phenomenon known as modal tonic shifting exhibited by ragas of CCM have been investigated in the present paper using ANN. MLP and LR models were constructed using inputs of frequencies present in all possible heptatonic ragas of CCM. Applying TS on them, frequencies of the base pitch was shifted to every note present in a given scale. The new scales generated were verified to evaluate if a new valid raga was obtained. Since there are 72 scales, TS to remaining 6 notes of each scale theoretically gives 432 combinations. But only 122 of these are valid scales. Hence a total of 194 exemplars were used as inputs.

100% accurate results were obtained with MLP and LR models, using LM and momentum learning, 2 HLs. Sensitivity analysis was performed to study the effects of the inputs on target values. MLP used here was hetero-associative, supervised learning since correct results (desired outputs) were known, so that during training the NN could adjust its weights to match its outputs to target values. After training, NN was tested giving only input values LR method allows user to test a network on a chosen data set Best network weights were used to minimise CV error.

Experiments and Results

The present paper concerns with the study of an intelligent system to analyse the performance of MLP and LR for a classification and regression problem. A case study was taken up to study the unique phenomenon of tonic shifts in CCM. Input data consisted of relative frequency ratios of the notes in heptatonic scales and their TS. MLP neural network was constructed with one and two HLs and studied for online/batch processing. LM and momentum learning rules were used. About 194 exemplars were used out of which 70% was used for training, 10% for (CV), 20% for testing. 1000 epochs (iterations) were used. Classification and Regression reports were generated. Regression gave a plot of network output and desired output for each value and correlation coefficient. Classification Report gave the meansquared error (MSE), normalized mean-squared error (NMSE), mean absolute error (MAE), minimum absolute error, maximum absolute error (MAE), correlation coefficient (r) for each output and percent correct for each class. Bread boards generated for

Contd....

Notes

each model are shown. Increasing number of HLs show an improvement in performance, as seen in plots for MLP- 1-B-L and MLP-2_B-L. Plots for best performing network, MLP-2-O-M is also shown. Very good correlation was obtained for this network Highest accuracy of 97-100% was obtained for LR-0-B-L/MLP- 2-O-M NNs.

A sensitivity analysis (SA) was performed on the results. SA is a valid tools for characterizing uncertainty associated with a model. It orders by importance the strength and relevance of inputs to determine the variation in output.

Question

Discuss the process of designing and training an MLP.

Source: <http://psrcentre.org/images/extraimages/0112244.pdf>

12.3 Summary

- Neural networks outperformed regression analysis and other mathematical modeling tools in predicting bond rating and profitability.
- In general, ANN are suitable for problems whose inputs are both categorical and numeric, and where the relationships between inputs and outputs are not linear or the input data are not normally distributed.
- Pattern recognition aims to classify data (patterns) based either on a priori knowledge or on statistical information extracted from the patterns.
- Artificial neural nets can be used for pattern recognition.
- Although the chase behaviors are genuine data, the 3D structures, surface physics, and shading are all purely for illustrative effect.
- In photo realistic Virtual Reality (VR) environments, the need for quick feedback based on user actions is crucial.
- The goal of speech-driven facial animation is to synthesize realistic video sequences from acoustic speech.
- Galapagos is a fantastic and dangerous place where up and down have no meaning, where rivers of iridescent acid and high-energy laser mines are beautiful but deadly artifacts of some other time.

12.4 Keywords

Galapagos: Galapagos is a fantastic and dangerous place where up and down have no meaning, where rivers of iridescent acid and high-energy laser mines are beautiful but deadly artifacts of some other time.

Mendel: Mendel is a synthetic organism that can sense infrared radiation and tactile stimulus.

Neural Classifier: A neural classifier detects visibility of teeth edges and other attributes.

Pattern Recognition: Pattern Recognition can be defined as the act of taking in raw data and taking an action based on the category of the pattern.

Pixel Map: The pixel map is a two dimensional function with two input variables: pixel position and one output variable.

Radiosity: Radiosity models the actual interaction between the lights and the environment.

Notes

Speech-driven Facial Animation: The goal of speech-driven facial animation is to synthesize realistic video sequences from acoustic speech.

Speechreading: Neural speechreading system as part of a user interface for a workstation.

12.5 Review Questions

1. Explain the use of artificial neural networks in financial sector.
2. Describe the use of neural networks in preventing frauds.
3. What is pattern recognition? Discuss the objective of pattern recognition.
4. Illustrate with example the use of Artificial neural nets for pattern recognition .
5. Discuss Radiosity techniques used for Virtual Reality Systems.
6. Describe the use of NN or Audio-to-Visual Conversion. Illustrate with example.
7. Illustrate the application of Galapagos in neural networks.
8. Explain the use of ANN in Speechreading.
9. Faces are tracked robustly by integrating motion and model-based tracking. Comment.
10. Illustrate the concept of three layer feedforward neural network.

Answers: Self Assessment

- | | |
|------------------------|--------------------------|
| 1. True | 2. False |
| 3. True | 4. False |
| 5. Pattern Recognition | 6. Pixel map |
| 7. Two-dimensional | 8. Controller |
| 9. Radiosity | 10. Virtual Reality (VR) |
| 11. Speech-driven | 12. Morphing |
| 13. Galapagos | 14. Mendel |
| 15. Speechreading | |

12.6 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.

Notes



Online links

<http://www.nd.com/apps/>

http://www.sans.org/reading_room/whitepapers/detection/application-neural-networks-intrusion-detection_336

<http://www-cs-faculty.stanford.edu/~eroberts/courses/soco/projects/neural-networks/Applications/index.html>

Unit 13: Knowledge Management

CONTENTS

Objectives

Introduction

13.1 Knowledge Management Activities

13.1.1 Knowledge Identification

13.1.2 Knowledge Acquisition

13.1.3 Knowledge Application

13.1.4 Knowledge Sharing

13.1.5 Knowledge Development

13.1.6 Knowledge Creation

13.1.7 Knowledge Preservation

13.1.8 Knowledge Measurement

13.1.9 Conceptual Framework for KM Activities

13.2 Knowledge Management Approaches

13.2.1 Knowledge Elicitation

13.2.2 Tacit Knowledge Management

13.3 Information Technology

13.4 Role of People

13.5 Success in KM Processes

13.5.1 Water Lily Strategy

13.5.2 Change factors

13.6 Summary

13.7 Keywords

13.8 Review Questions

13.9 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss the Concept of Knowledge Management
- Explain Knowledge Management Activities and Approaches
- Analyze the Role of Information Technology and People in Knowledge Management
- Discuss the Success Factors of Knowledge Management

Introduction

Notes

In simple terms, knowledge management refers to the process of acquiring, organizing, storing, sharing and using knowledge by organizations. Knowledge management is a branch of management that aims at attaining the optimum business performance through the synergy of people, processes and technology in creating and sharing relevant knowledge. To succeed in any venture, including business ventures, sufficient knowledge is required. Knowledge management has become necessary as we have moved from a society where information was scarce to a society where there is a glut of information. The problem today is generally not about procuring information but deciding which information to use.



Caution Organizations should be careful in acquiring relevant data and processing those in an ever-changing business landscape.

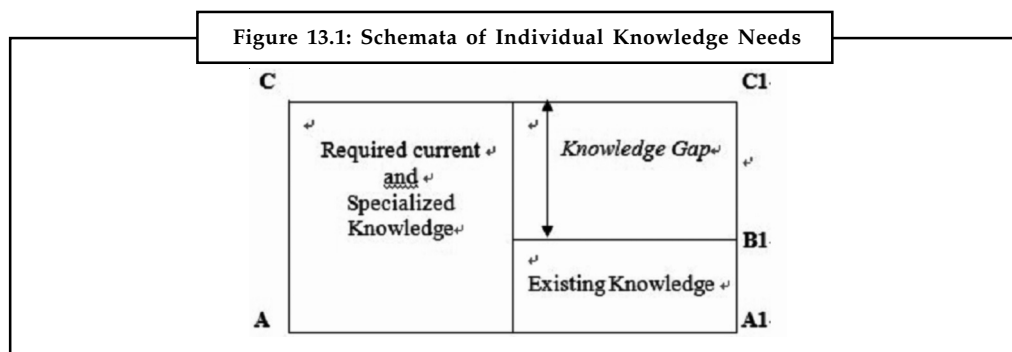
It is equally important for businesses to discard obsolete information and acquire the latest information to survive and stay competitive. In this unit, we will discuss the concept of knowledge management.

13.1 Knowledge Management Activities

Knowledge management (KM) is a process that helps organizations identify, select, organize, disseminate, and transfer important information and expertise that are part of the organization's memory. The term knowledge management often leads to confusion, resulting in knowledge management activities being individually identified rather than grouped as part of a knowledge management program to eliminate confusion and help ensure buy-in and support from the organisation. Knowledge management activities in an organization include knowledge identification, knowledge acquisition, knowledge application, knowledge sharing, knowledge development, knowledge creation, knowledge preservation and knowledge measurement. These activities are discussed below.

13.1.1 Knowledge Identification

An enterprise must state its business strategies and objectives. The knowledge requirements have to be identified to meet these goals. The difference between what the enterprise requires and what it currently has is what is called the knowledge gap. Besides that, it can be used to identify the knowledge gap of the individual employees which is shown in Figure 13.1. A to C shows the current and specialized knowledge that employee requires in an organization. If the individual employee has the existing knowledge from A1 to B1, there will be a knowledge gap in the employee which is shown from B1 to C1.



Source: <http://www.ipcsit.com/vol36/004-ICIIM2012-M20038.pdf>

Notes



Task Illustrate the process of identifying knowledge with example.

13.1.2 Knowledge Acquisition

After determining the knowledge gap, the next step is to close the knowledge gap by applying knowledge acquisition or knowledge import. Companies import a substantial part of their knowledge from outside sources. Relationships with customers, suppliers, competitors and partners in co-operative ventures have considerable potential to provide knowledge – a potential that is seldom fully utilized. Firms can also buy knowledge which they could not develop themselves by recruiting experts or acquiring other particularly innovative companies.

Knowledge acquisition is quite simply the process of acquiring knowledge that is available somewhere. For an organization, this might entail capturing knowledge from existing documents. It could mean capturing the tacit knowledge of its people into its repositories. Or it might mean identifying external sources of either process/technology expertise or market intelligence so that this knowledge can be purchased.

13.1.3 Knowledge Application

After acquiring knowledge, it must be applied in an organization. Knowledge application means making knowledge more active and relevant for an organization in creating values. Knowledge application deals with the fact that employees continually apply their knowledge to their working situation. We have to make local knowledge usable in global application. Organizational knowledge needs to be employed into a company's products, processes and services. If an organization does not find it easy to locate the right kind of knowledge in the right form, the organization may find it difficult to sustain its competitive advantage.

13.1.4 Knowledge Sharing

After applying the knowledge in an organization, the next step is knowledge sharing. The sharing and distribution of knowledge within an organization are a vital precondition for turning isolated information or experiences into something which the whole organization can use. The most important step is to analyze the transition of knowledge from the individual to group or organization. In knowledge based economy, knowledge itself is not power, when knowledge is shared, it become power.

13.1.5 Knowledge Development

Knowledge development is a building block which complements knowledge acquisition. Its focus is on generating new skills, new products, better ideas and more efficient processes. Knowledge development includes all management efforts consciously aimed at producing capabilities which are not yet present within the organization or which do not exist either inside or outside it.



Notes Traditionally, knowledge development was anchored in the company's market research and in its research and development department; however, important knowledge can also spring from any other part of the organization. This can provide the company general ways of dealing with new ideas and utilizing the creativity of the employees.

13.1.6 Knowledge Creation

Knowledge creation is the key focus about creating new knowledge or innovating existing knowledge for the organization. Nonaka and Takeuchi (1995) suggested a KM model in terms of knowledge creation perspectives based on four kinds of process knowledge conversions. The spiral – type conversions between explicit knowledge and tacit knowledge, i.e. the SECI model.

They are as follows:

- From tacit knowledge to tacit knowledge (Socialization)
- From tacit knowledge to explicit knowledge (Externalization)
- From explicit knowledge to explicit knowledge (Combination)
- From explicit knowledge to tacit knowledge (Internalization)

13.1.7 Knowledge Preservation

Competencies once acquired are not automatically available for all time. The selective retention of information, documents and experience requires management. Organizations commonly complain that recognition has cost them part of the memory.



Caution The processes for selecting, storing and regularly updating knowledge of potential future value must be carefully structured.

Knowledge once acquired needs to be preserved. Storing or preserving is certainly not about putting it somewhere and forgetting all about it. Unless knowledge is constantly updated and kept relevant, it ceases to be knowledge. Obsolete knowledge can be most dangerous.

13.1.8 Knowledge Measurement

The next activity is knowledge measurement; this is to measure the impact and effects after implementing knowledge management in an organization.




Example: Customer satisfaction, efficiency, productivity, quality, etc.

This must be developed to link actions to strategies, monitor changes in intellectual capital and encourage value-creating work. Besides that, a comprehensive survey of the German TOP 1000 and European TOP 200 companies showed that KM helps to achieve the goals of a company. KM can best be used to increase innovation ability, increase of product quality, reduction of goals, increase of effectiveness and customer satisfaction.

13.1.9 Conceptual Framework for KM activities

The conceptual framework for KM activities is proposed in Figure 13.2.



Source: <http://www.ipcsit.com/vol36/004-ICIIM2012-M20038.pdf>

Notes

This concluded that KM activities consist of knowledge identification, knowledge acquisition, knowledge application, knowledge sharing, knowledge development, knowledge creation, knowledge preservation and knowledge measurement. The most significant contribution is to provide a KM activities framework.

Self Assessment

Fill in the blanks:

1. The difference between what the enterprise requires and what it currently has is what is called the
2. is quite simply the process of acquiring knowledge that is available somewhere.
3. deals with the fact that employees continually apply their knowledge to their working situation.
4. is a building block which complements knowledge acquisition.
5. is the key focus about creating new knowledge or innovating existing knowledge for the organization.
6. is to measure the impact and effects after implementing knowledge management in an organization.

13.2 Knowledge Management Approaches

The Knowledge Capital in the enterprise is an intangible capital, which is not visible in the organization; its content even remained elusive. One can agree to now say that the content of that capital is both hidden and scattered in two essential components of the enterprise:

- **The human and social capital:** The quintessence of the knowledge in the enterprise (core knowledge) is in the head of its employees. It is anyway the ultimate place before its operational use. It is in a tacit form, and hardly expressible (or even not at all), according to the adage: “we know more that what we can tell”. This tacit knowledge capital, tightly linked to the human capital, is at a time collective and individual and lives through knowledge networks of the organization (knowledge workers) that produce and use a precious and operational knowledge constantly.
- **The information capital:** Enterprises stored, since decades, huge information masses, that they stock and distribute with more and more sophisticated information. Employees of the enterprise constantly use this gigantic system of information for their operational activity, they acquire information, attach a precise sense in their operational context to transform it into useful knowledge to their profession. Unfortunately, to find the right information, at the right moment for the right person becomes a difficult task, regarding the available amount of information. The potential knowledge is then really buried in masses of information, and is not easily accessible.

Therefore, managing a knowledge capital is problematic, because of its two-headed character two-headed human capital/information capital and its character hidden tacit/buried.

To approach this problematic two main approaches are possible:

- The first is to face (partially) the hidden character of knowledge, by making it explicit, either from the tacit knowledge in knowledge networks, either from knowledge buried in information systems.
- The second is to keep to the available knowledge its hidden character (tacit, especially) and to manage them by managing networks of the knowledge themselves.

13.2.1 Knowledge Elicitation

Notes

According to whether one is interested in tacit knowledge in networks in the organisation or knowledge buried in information systems, elicitation of the Knowledge Capital uses two distinct approaches.

Tacit Knowledge Elicitation

- **Knowledge Transcription:** A first type of approach is what one will call the knowledge transcription: some tacit knowledge can be elicited simply, by transcribing them, in a more or less structured manner. It is the case in the setting up of systems quality (of which the first rule is “write what you have to do”), or in return on experience files, or in the writing of publications. It is also the case of the “secondary documents” that synthesize knowledge contained in given documents.
- **Knowledge Engineering:** Knowledge Engineering is a manner more sophisticated than the transcription to capture parcels of tacit knowledge. It appeared with expert systems (or knowledge-based systems). These systems were supposed to replicate reasoning of experts on domains of specific knowledge. One perceived quickly that if powerful technologies were available to design such systems, the essential difficulty resided in the capacity to transfer knowledge of one or several human experts into a computer program. Knowledge Engineering put methods therefore in place to collect the knowledge, the most often from interviews, and to structure it, in general from models.

These methods can be used therefore with profit to clarify, from interviews with knowledge holders, a part of the capital of tacit knowledge of the organization.

An example typical of this evolution is the MASK method that was method to specify experts systems and became a method of knowledge capitalization, integrated then in a general KM problematic.

Knowledge Extraction

- **Knowledge Extraction from data:** All enterprises detain big quantities of data, resulting of their production activity. Those data are very diverse (technical, management, marketing...), their mass doesn't quit to grow (it doubles every 20 months on average). Moreover there are other considerable masses of data called non-structured or semi-structured, that are all textual data (and others media) that correspond to the production of texts, cards, reports and other documents of any kind. This informational capital is probably a wealth of the enterprise, for its production needs, but that could be reused efficiently – a posteriori – for other needs. However, it proves to be that hardly 10% of this capital is exploited. It can be explained by the difficulty to reuse information that has been structured for objectives different from capitalization and reuse. However, big efforts are made currently to valorise these layers of information accumulated by data processing for production.



Did u know? The objective is about producing, from these layers, new information that are useful to the action in the enterprise, in other words producing “operational knowledge”, in the KM meaning.

It is a manipulation of information in an objective of knowledge discovery, called knowledge extraction from data (or Knowledge Discovery from Data or KDD), linked also to rather equivalent concepts as “Text Mining”, or “Data Mining” or “Data Warehouse”.

Notes

- **Knowledge extraction from texts:** Knowledge extraction from texts (Text Mining) is the discovery of useful information from hidden patterns buried in large corpus of texts (sometimes called non structured or semi-structured information). Research engines now process those types or texts, more and more abundant, included web pages that are a growing source of knowledge.

Supporting Technologies for Explicit (Elicited) Knowledge Management

After having elicited a part of the knowledge capital, one has an informational corpus susceptible to be used for transferring or operating knowledge. To put this corpus in action, one can elaborate two kinds of systems:

- **Knowledge servers:** These are systems, usually included in the intranet of the company, from which one can browse (in the most possible intelligent and ergonomic manner) the elicited knowledge. These systems don't solve problems directly for the users, but give means, in a rich and flexible way, for retrieving knowledge that may be useful to solve an operational problem.
- **Knowledge based systems:** These are computer-based systems which operate the elicited knowledge, as, for example, expert systems. They use the elicited data and structures to solve a precise high level problem: decision support, process supervision, diagnostic, resource planning, design support.



Task Analyze the uses of knowledge servers.

13.2.2 Tacit Knowledge Management

Knowledge elicitation is an approach that may be not chosen by some organisations, for various reasons: difficulties to set up such processes that may be long and time consuming, direct cost too high, confidentiality problems, problems with people, with the knowledge networks.

Another possible approach can be derived from the way in which knowledge is produced in organisations, more precisely from the different forms of groups and functions which participate: networks, communities ... Knowledge is there seen as the result of a cooperative process in a collective action. The problem is then not to elicit this knowledge, but to foster its creation, its sharing by managing the cooperative work of a community of people. One then don't manage knowledge, but the community which creates it. This knowledge may then remain tacit within the community, while being shared and operational. One may then talk of "cooperative Knowledge Management".

The tacit/explicit approaches are not opposed but complementary. It is, for instance, useful that a knowledge community which manages its own knowledge produces visible and tangible records; and on the other hand, an elicited corpus of knowledge needs a knowledge community to operate it and make it evolve.

Cooperative Knowledge Management requires four key points:

- Identification of knowledge communities
- Exchange mechanisms that allows knowledge transfer in knowledge communities
- Principles of managing and supervising cooperation
- Technologies supporting cooperative Knowledge Management

Most of times, there already are communities in the company (for specific jobs, specific profession, for projects, for given practices, specific interests ...). One must analyse them and include them in a strategic schema, in order to optimize their performance, in terms of knowledge sharing and creation. The problem is to map the communities, existing or to be developed, to define for each community, the issue, the type of the community, the nature of its need or the problem to solve.

The Functioning Modes of Knowledge Communities

There are three types of way for working, exchanging through a community:

- **The classic type of the "labour division":** A worker realises a limited number of tasks, with reduced autonomy in designing and realising the job. This is the most common way, even in new and modern forms due to the introduction of information technology. Dedicated structures analyse, design and codify the work to be done. This type of work is supported by formalised links (planning of activities, performance supervision ...), which don't facilitate cooperation. Cooperation links are hidden, implicit and limited.
- **The coordination type ("soft" cooperation):** It relies on a better coordination of activities, with information sharing and tasks synchronising. It is compatible with the Taylor point of view. The only problem is to do better and quicker. The consequence is the implementation of a global information system with a fast communication function and a generalised access to information (the archetype is company intranet and email).
- **The cooperation type ("strong" cooperation):** This is the work in synergy, not only coordination between separated tasks. Work together is to establish a community with sound bases for:
 - ❖ nature of problems and of knowledge to develop,
 - ❖ objectives identification,
 - ❖ convergence of reasons why people work together

Knowledge Community Supervision

Implementation of knowledge communities may be performed through their life cycle. This leads to the definition of:

- Management commitment at the different phases in the cycle
- Indicators for supervision and control
- Adequate supporting tools

The different phases in a community development are:

1. **Emergence:** A community is created on a knowledge domain. This creation arises from aggregation according to affinities. It is spontaneous (a community cannot be prescribed). Management is not an actor, maybe an observer. The main need is connection between actors.
2. **Structuring:** The emergent community begins to collaborate. It is consolidated by peer co-optation. Information sharing leads to coordination, but every actor is autonomous. The community is still informal, driven by the actors themselves. Management is still an observer. The needs are for coordination, for mutual information.
3. **Becoming Official:** The community is structured and active enough to become visible for the organisation (if it is too premature, it stops). The community becomes coherent through

Notes

a shared common goal. Management interferes to give sense in the company strategic objectives and to give an official recognition, together with resources. The needs are then for collaborative work.

4. **Consolidation:** The community works and produces on a regular basis, and this work is integrated in the collective action. Management evaluates periodically its action and its production. The needs are then for sharing and capitalisation.
5. **Dissociation:** The well anchored communities tend to withdraw into themselves, and some outside signs are often there to show it (apparition of a jargon, gregarious mind,). The management must clearly observe these phenomena and must facilitate cooperation and transfer of knowledge between communities.
6. **Know how to stop a community, as one can stop a project:** This dissociation maybe generated in intern by actors blocked in their creative capacity and who will find in the emergence of another community the necessary rebirth to the expression of their knowledge need.

To pilot cooperation it is necessary to define:

- **Indicators of cooperation:** How to recognize that there is a development of cooperation within the different types of collective work (team, project, process and network)?
- **Indicators of results of the cooperation:** What are effects of the cooperation on the knowledge capital? Several criteria seem useful to recognize that members of a network cooperate:
 - ❖ They construct, from their individual representations, of shared representations of problems to solve, of objectives to reach, of goals to achieve.
 - ❖ They communicate efficiently, using a common language, understanding the language of the others and sharing the point of view of others.
 - ❖ They cross their domain (discipline, sectors, geographical...) by interdisciplinary actions
 - ❖ They accept the existence of conflicts (of criteria, of points of view, of priorities...), managing them in appropriate time, and proceeding to arbitrations.
 - ❖ They put some new applicable and evolutionary organization in place.

Supporting Technologies for Knowledge Communities

The new types of work described above are tightly linked to rapid adoption and dissemination of systems and Technologies for Information and Communication (ICT) within the organisation. The rapid increase and development of new ICT has a considerable influence on Knowledge Management. There is sometimes confusion (maybe on purpose) between ICT and KM.

In classical information systems (data bases, data banks ...) functionalities are well known: an information system is by definition a system to process, store and present information. Coming from computer networks, communication technology had the clear functionality of data transmission. With the introduction of new types of tools, especially web based, those distinctions are no longer clear; they are more complex and difficult to understand deeply from the user's point of view. In fact the new ICT products must be analysed regarding the notion of services, and those services justify their use in Knowledge Management problems, especially cooperative work, in the sense defined above.

There are four types of services attached to ICT:

- Communication

- Coordination
- Resources sharing
- Information sharing

Among those services, one can distinguish the basic services limited to communication software (communication and coordination), and services with added value based on the information system of the company. But this distinction communication tools, as on a website. This set of services help the collaborators in an organisation for knowledge sharing and for working in a cooperative way, especially through knowledge communities. But one must not forget that if a technical platform is necessary, it is neither a prerequisite, nor a goal. Cooperative Knowledge Management is really a problem of organisation and method becomes fuzzy in the new usages of ICT (e-mail may be used as an information basis, information bases maybe linked to communication tools, as on a website).

This set of services help the collaborators in an organisation for knowledge sharing and for working in a cooperative way, especially through knowledge communities. But one must not forget that if a technical platform is necessary, it is neither a prerequisite, nor a goal. Cooperative Knowledge Management is really a problem of organisation and method.

Self Assessment

Fill in the blanks:

7. is a manner more sophisticated than the transcription to capture parcels of tacit knowledge.
8. Knowledge extraction from is the discovery of useful information from hidden patterns buried in large corpus of texts.
9. are systems, usually included in the intranet of the company, from which one can browse the elicited knowledge.

13.3 Information Technology

The information technology infrastructure should provide a seamless “pipeline” for the flow of explicit knowledge to enable

- capturing knowledge,
- defining, storing, categorizing, indexing and linking digital objects corresponding to knowledge units,
- searching for (“pulling”) and subscribing to (“pushing”) relevant content,
- presenting content with sufficient flexibility to render it meaningful and applicable across multiple contexts of use.

Information technologies such as the World Wide Web and Lotus Notes offer a potentially useful environment within which to build a multimedia repository for rich, explicit knowledge. Input is captured by forms for assigning various labels, categories, and indices to each unit of knowledge. The structure is flexible enough to create knowledge units, indexed and linked using categories that reflect the structure of the contextual knowledge and the content of factual knowledge of the organization, displayed as flexible subsets via dynamically customizable views.

Effective use of information technology to communicate knowledge requires an organization to share an interpretive context. The more that communicators share similar knowledge, background

Notes

and experience, the more effectively knowledge can be communicated via electronically mediated channels. At one extreme, the dissemination of explicit, factual knowledge within a stable community having a high degree of shared contextual knowledge can be accomplished through access to a central electronic repository. However, when interpretive context is moderately shared, or the knowledge exchanged is less explicit, or the community is loosely affiliated, then more interactive modes such as electronic mail or discussion databases are appropriate.



Notes When context is not well shared and knowledge is primarily tacit, communication and narrated experience is best supported with the richest and most interactive modes such as video conferencing or face-to-face conversation.

Self Assessment

State True or False:

- 10. Information technologies offer a potentially useful environment to build a multimedia repository for rich, explicit knowledge.
- 11. Effective use of information technology to communicate knowledge does not requires an organization to share an interpretive context.

13.4 Role of People

Human experts do not have static memories. They can change their internal classification systems when their conception of something changes, or when their needs for retrieval changes. People change their focus or their interests and the things they think about and remember change as well. For the most part, such changes are not conscious. People do not typically know the internal categorization scheme that they use. They can do this without even realizing they have done it. This is what a dynamic memory is all about – getting smarter over time without realizing it. The acquisition of new knowledge actually makes experts smarter, while it often just makes knowledge management (KM) systems slower.

People seem to be able to cope with new information with ease. We can readily find a place to store new information in our memories, although we don't know where or what that location is. This is all handled unconsciously. We can also find old information, but again we don't know where we found it and we can't really say what the look-up procedure might have been. Our memories change dynamically in the way they store information by abstracting significant generalizations from our experiences and storing the exceptions to those generalizations. As we have more experiences, we alter our generalizations and categorizations of information to meet our current needs and account for our new experiences.

Despite constant changes in organization, we continue to be able to call up relevant memories without consciously considering where we have stored them. People are not aware of their own internal categorization schemes — they are just capable of using them.

The question for KM is how to make systems more like those of people. Human memories dynamically adjust to reflect new experiences. A dynamic memory is one that can change its own organization when new experiences demand it. A dynamic memory is by nature a learning system. No KM system learns. But they need to learn in order to actually work properly.

The underlying question is how knowledge is structured. People structure knowledge when they build any KM system by inventing a set of categories to put documents in.

But, what categories do people use inside their own heads? People use knowledge structures, ways of organizing information into a coherent whole, in order to process what goes on around them. Knowledge structures help us make sense of the world around us. What knowledge structures does an expert have and how do they acquire them?

Understanding how knowledge structures are acquired helps us understand what kinds of entities they are. Learning depends upon knowledge and knowledge depends upon learning. A script is a simple knowledge structure that organizes knowledge we all know about event sequences in situations like restaurants, air travel, hotel check in, and so on. We know what to expect and interpret events in light of our expectations. But do we use this kind of structure to manage knowledge?

If something odd happens to us in a restaurant, how do we recall it later? Let me count the ways. We would recall it if we entered the same restaurant at another time or if we had the same waitress at a different restaurant, or if we ate with the same dinner companions (assuming we ate with them rarely). It is clear that an incident in memory is indexed in many ways. One set of indices is the “people, props and places” that appeared in an incident and are associated specifically with that incident.

But there is more abstract indexing method that goes beyond the “people, props, and places” type of index. Those indices are about actions, results of actions, and lessons learned from actions. These indices matter greatly in a KM system. If they do not exist no one will learn anything from a description of an event that has lessons in it beyond those about the people and places of that event.

One set of abstractions about actions involves roles and tasks. Organizing information around role and tasks allows events to be easily accessed if one has an implicit understanding of the roles and tasks involved in a given situation.

Beyond using an organizational scheme involving roles and tasks, human experts can do something that is quite significant. They can abstract up a level to organize information around plans and goals. To put this another way, if the waitress dumped spaghetti on the head of someone who offended her, you should get reminded of that event by the 3 Ps and also if you should happen to witness this event in some other setting. But, far more importantly, you should get reminded of this event if you witness the same kind of event another time. The question is what does it mean to be the same kind of event? Whatever this means, it would mean different things to different people. One person might see it as an instance of “female rage” and another as an instance of “justifiable retribution.” Another might see it as a kind of art.

The key issue is to learn from it. Any learning that takes place involves placing the new memory in a place in memory whereby it adds to and expands upon what is already in that place. So, it might tell us more about that waitress, or waitresses in general, or women in general, or about that restaurant and so on, depending upon what we previously believed to be true of all those things. New events modify existing beliefs by adding data to what we already know or by contradicting what we already know and forcing us to new conclusions. Either way, learning is more than simply adding new information. Since information helps us form a point of view, when we add new information it changes the information we already have.

The question is: how do we find the information we already have? This answer depends upon how it was indexed (or categorized) in the first place. Learning depends upon the initial categorization of what we know and may involve changing those categorizations in order to get smarter. A KM system that does not do this will never be very smart and will fail to absorb new information in significant ways.

Notes

What are the Knowledge Structures that Human Memory Uses?

Higher level knowledge structures, about more complex issues than those dealt with by scripts, hinge upon the notion of abstraction and generalization. Scripts are specific sets of information associated with specific situations that frequently repeat themselves. Scripts are a source of information, naturally acquired by having undergone an experience many times, yielding the notion of a script as a very specific set of sequential facts about a very specific situation. Scripts allow one to organize low level sequential knowledge in a KM. Higher level, more abstract notions, enable learning because they allow sharing of knowledge across script boundaries.

Consider "fixing something." Is this a script? Not really. There might be many scripts each dependent upon what you are fixing. But, if we treat each instance of fixing something as different from every other one, how will we get smarter as a result of having fixed something when we encounter something different that has to be fixed? It is reasonable to assume that people who fix things do get smarter about the process each time. They learn. The question is "how?"

In any knowledge-based understanding system, any given set of materials can be stored in either a script-organized or plan/goal-based form. If we choose to give up generalizability, then we can use a script based organization for a KM system. This is more efficient in the short term. Knowing a great deal about a discrete set of roles and tasks in a human organization is a good way to organize information within that organization.

But information organized by scripts never transcends the boundaries of those scripts. Any knowing system must be able to know a great deal about one script without losing the power to apply generalizations drawn from that knowledge to a different set of issues within the system.



Did u know? A script, and any other memory structure, should be part of a dynamic memory, changeable as a result of incorporating new experiences.

Any structure proposed for organizing information within a KM must be capable of self-modification. Such modification comes about as a result of new events differing in some way from the normative events that a script describes.

When new information is placed in a dynamically organized KM, if that information simply amplifies or clarifies the script to which it belongs, then it is simply added to it. On the other hand, when a new event modifies the script, its difference must be noted. And its difference must be applied to all structures to which it is relevant.

By utilizing general structures to encode what we know in memory, a system can learn. Given enough modifications from a script, an intelligent system would begin to create new general structures that account for those differences. That is, just because an episode is new once, it does not follow that it should be seen as forever new. Eventually we will recognize what was once novel as "old hat." To do this, we must be constantly modifying our general structures, which, as we have said, is what we mean by a dynamic memory.

All this depends on detailed information, like scripts, that describe the processes in a situation. To put this another way, there cannot be, at least for quiet some time, generalized KM systems that work for every industry or every enterprise. The backbone of intelligence is knowledge of situations. You don't ask a novice to captain a ship, nor do you ask a beginning salesman to call on your biggest account. An organization knows quite a bit about the processes involved in these situations. We can say what a ship captain does, not simply, but we can say it. And, we can say what a salesman does. Further we can say the experiences a company has had with selling into a particular company or with a particular client or with clients who are like a prospective client.

There is a lot of knowledge in an enterprise that can be used to organize new knowledge that is coming in. People understand new knowledge in terms of what they already know. A smart KM system must know a lot of about an industry and a particular enterprise before it starts up. This is hard but by no means impossible. And it is the future of software – namely software that really knows a great deal about your business.

Simply put, any business could use someone who knew all about every job, and every person doing that job and every experience the company had had in the past and what its goals and plans were at the moment and could use all that knowledge to know what to do with new information it has just received. That senior citizen of the company used to exist, and he can exist again. Only this time he will be a computer equipped with a very new kind of KM system one that is not about document retrieval but about delivering just in time advice to those who need it.

Self Assessment

Fill in the blanks:

12. A memory is one that can change its own organization when new experiences demand it.
13. If we choose to give up generalizability, then we can use a based organization for a KM system.

13.5 Success in KM Processes

A KM process is put in place to optimize knowledge use and evolution in an enterprise. It concerns in crucial manner knowledge actors: experts, specialists, competent people without whom the knowledge capital would not have added value. As it is an element of strategic management, management must be strongly implicated to many hierarchical levels. Here resides both strength and weakness Knowledge Management of: the cohabitation of two fundamental networks: the knowledge network and the power network. It is why a KM project doesn't look like a classic project, because two poles of decision, completely different, must cohabit to construct a common system: a Knowledge Management system.

The setting up of a Knowledge Management in an enterprise is therefore delicate, and a certain number of unusual factors for a classic project must be taken in account. We enumerate some of them here.

13.5.1 Water Lily Strategy

One key changes factor in a Knowledge Management project is the strategy of project deployment. There is a radical change compared to a classic project.

Let's take the example of an Information System project. The project deployment strategy is composed by a needs analysis, a specification and a roadmap that foresees a progressive and linear implementation, integrating progressively the users, though, in most of cases, a training program. One now knows advantages and inconveniences of this project deployment type: more mistakes are upstream the process, more they are costly; risks of failure near the end of the process are far from being negligible, early endings are not exceptional, etc.

Project management is centralized in general by dedicated professional in information systems, it is linear (for example, the cost of realization of the last half is the half of the total cost), it is planned, users are implied in a sporadic manner (sometimes in beginning of project, and in the end for training) etc. Those types of "roadmap strategies" for a Knowledge Management strategy (except if, of course, it is reduced to an information system project) are revealed to be often a failure. It is too linear and don't take enough in account incentive factors.

Notes

A “socio-technical” type strategy is necessary. One can build a new type of project strategy, radically opposite to the “roadmap” one. It may be called the “water lily strategy”. This term expresses the idea of steady organic growth, that designates the cell growing process (meiosis in morphogenesis), that processes by successive cellular divisions to create some viable and complex structures. The water lily evokes a well known childish riddle: a water lily doubles its surface every day, if it already took forty days to cover half of the pond, how many days are needed to cover the entire pond? An imprudent and too fast answer would be eighty days while considering that as much time is necessary to cover one half that the other. This answer neglects the water lily nature and the initial hypothesis! If one transcribes this riddle in term of project management, it may be: a KM project is expensive, if the project costs 40 M€ to cover the needs for managing half of the Knowledge Capital, how much will be needed for managing the entire Knowledge Capital? The linear extrapolation of costs will answer undoubtedly rather toward 80 M€ than 41M€! One doesn't have the habit to evoke the hypothesis of the water lily in this kind of problem.

It is yet “water lily” strategy that is the most adequate for a Knowledge Management project. It is the best way to include the change process, incentive factors, and the emergence phenomenon in complex systems. The water lily strategy is a strategy of constant effort and cumulative effect, which perfectly suits the cumulative economic nature of knowledge. It indicates notably that if efforts can be considerable to start the project (from a human or other resources point of views), they will not be necessarily multiplied during the project spreading. Incentive factors being essential, one benefits from a leadership effect. The material costs are not necessarily the most important; they don't grow therefore linearly etc. One of the big interests of this project management type is to minimize costs and risks. The first water lilies are not too expensive. Stops, often unavoidable on medium-term projects, are not damaging. Failures don't necessarily imply to completely revise the entire approach.

The water lily strategy usually comprises three phases. These phrases are discussed below:

The Pilot Project(s)

The first water lilies are persuasive projects. They are autonomous projects, concerning a restricted part of the knowledge capital. It is often possible to identify such projects that already emerged spontaneously (often long time ago) in the enterprise: it's indeed true that the knowledge is (fortunately!) a common and pragmatic practice for a long time. These pilot projects, concerned with knowledge capitalization, sharing or creation can be revealed or initialized by knowledge cartography. Some of their features are:

- They concern a small knowledge domain (too small), but meaningful.
- They are performed by people convinced of the necessity to act on their knowledge, and that are therefore susceptible to have a leadership action in a general approach.
- They are settled and achieved locally in units. They are not necessarily supported by the management (it is necessary however, to get an “understanding neutrality” from the management). This allows avoiding the considerable effort to get, from the first beginning, the “top management commitment”, that is often a prerequisite for strategic projects, but that is sometimes a delaying or even blocking factor.
- They show real life examples as for the possibility to manage the knowledge capital.

The pilot project(s) are the ferments of a global and strategic Knowledge Management approach. They are significant of a bottom up approach where people who produce and use knowledge in their daily practice (Knowledge Workers) prove the need and the possibility to capitalize, to share or to make evolve this knowledge in their context. A pilot project must show that it is useful to work on knowledge, and to show an original added value compared with a more classic project (quality, documentation, data processing...). It is one concrete and pragmatic

starting point to communicate to management and the other people in the company, contributing to the awareness in this domain.

The KM Projects Federation

The first water lilies grow from emergence, which is a typical phenomenon due to the complexity of a knowledge system. This emergence phenomenon must be carefully managed in the first phase. Coordination tasks must be undertaken to gather and to federate the different projects, the strategic and global dimension is not still major, but begins to appear. Cooperation tasks then take place, giving an official status to this federation of projects organised as a network (the KM network of the enterprise). Several points are addressed:

- The definition of a common objective that gives sense to the set of projects, and a unified shared view. It allows to become persuasive and initiate a leadership for the rest of the enterprise
- The official status legitimating through actions of communication and explanation by the concerned hierarchies. It allows to give a formal structure to the network, and means to support it, as well as a strong recognition of the project
- A regular and coherent communication on Knowledge Management of knowledge inside (maybe outside) the Enterprise.

Project Deployment

The network of water lilies being now in place, one must assure its “steady organic growth” in order to cover the set of needs of the enterprise for the management of its knowledge capital. This project deployment may be usefully initialised by critical knowledge cartography (cf. infra). The set of critical domains and corresponding actions can thus be defined thus and included in a global plan. KM processes (cf. infra) must be put in place, with assessment tools. Global supervising tools must be put also in place (as balanced scorecards or the Intellectual Capital Navigator). Communication and incentives must be especially studied, because a global project must mobilize a large number of people to enrich and to make live the KM system. This last point often implies a deep change in habits and beliefs. Incentives are not always of classic type, and rely often more on intellectual or immaterial satisfactions than on material rewards. The key factors of success are also (and maybe especially) in these incentives that are the essential drivers for change in this type of project. This is the topic of the following paragraph.

13.5.2 Change Factors

The setting up of a Knowledge Management system goes through a certain number of delicate phases that must taken into account according to the bivalent nature of the project. These phases are key factors of success to which it is necessary to pay a careful attention. Among these factors there are:

- **The “Mirror Step”:** A Knowledge Management system is built with knowledge holders. One of the first factors of incentive and acceptance for the system is that these people recognize themselves in the implemented knowledge. It must a structured and valorising picture of them. This step is essential. If knowledge holders don’t recognize themselves in the system, this one doesn’t have any chance to be validated thereafter. This first step allows the concerned “knowledge network» to become support of the project.
- **Consensus:** Knowledge implemented in the system must consensual, reflecting the individual knowledge that is part of it. Consensus is not natural, especially in a company. It requires a specific and sensible process.

Notes

- **Legitimizing:** Consensus between participants is not sufficient to validate the knowledge put in the system. This one cannot be validated in the classic sense, because it deals with personal and collective knowledge that is very particular to the enterprise, to the highest level, and for which external model of reference never exists. The only possible validation is obtaining a consensus from a “peer committee”, in general constituted by other actors of the knowledge network or related actors.
- **Approval:** The knowledge, once legitimized, must be endorsed by the hierarchy that gives it an official and productive status in the organisation. This stage is the starting point for project dissemination and use in the company.
- **Appropriation:** It is clear that a Knowledge Management system doesn't have an interest if there is no appropriation and no evolution of the knowledge in the enterprise. It is an ending step that initializes a “virtuous knowledge cycle”. Up to now, one knows very few about how to correctly perform this phase.

This set of key factors show, if still needed, that the Knowledge Management in an enterprise is a complex process, that must be implemented progressively and in a rather long perspective. In this domain, as in many others, there are no miracle tools or methods capable to solve any problems. In this domain, more that elsewhere, we deal with knowledge, a matter that concerns human being since the beginning of mankind!

Self Assessment

Fill in the blanks:

14. The are the ferments of a global and strategic Knowledge Management approach.
15. The strategy is a strategy of constant effort and cumulative effect, which perfectly suits the cumulative economic nature of knowledge.



Case Study

Knowledge Management

Practice of Knowledge Management

Wal-Mart is an American multinational corporation running a chain of large discount departmental stores and warehouse in several countries. Wal-Mart inc. was started in 1962 and incorporated officially in 1969 and began trading stock as a publicly held company in 1970. At current, Wal-Mart international experiences a rapid growth and it has acquired 4, 573 stores along with 730,000 associates evident in 14 countries that are out of the US continent. The company has wholly owned operations in Brazil, Argentina, Canada and the United Kingdom. In addition to these wholly owned international operations, Wal-Mart has several joint ventures in China as well as a number of majority owned subsidiaries. Wal-Mart operates under different banners in most of its international stores but they all work towards the same goal; enabling their customers save money and live a better life.

Knowledge management refers to the practices and strategies that a company uses in an attempt to create distribute and enable adoption of strategic insights and specific experiences (O'Leary, 2002). This knowledge can be embodied in either individual or organizational practices or processes. Successful business firms and companies have committed hefty investments towards development of internal knowledge management efforts. One such

Contd....

company is Wal-Mart Stores Inc. which through this effort has emerged the world's largest corporation.

Knowledge has to be organized for it to be useful and valuable. Depending on what use will be made of it, knowledge should be organized differently. Addicott et al, (2006, identified four dimensions to knowledge management. These are connectedness, completeness, perspective and congruency.

Creation of knowledge is brought about through recognition of gaps existing in an organization's knowledge management practices. Choo, (2002), says that knowledge gaps can stand in the way of problem solving and decision making and therefore hinder development of new products. Choo developed a model for decision making which aids in knowledge management. This model involves identifying and evaluating alternatives by processing all information and knowledge collected up to date.

In its efforts at developing and nurturing insightful knowledge management skills, Wal-Mart focuses on its organizational objectives which specifically include improved business performance, innovation, competitive advantage as well as continuous improvement of business operations and processes. The company's knowledge management incorporates such important issues as formal apprenticeship, discussion forums, corporate libraries, mentoring programs and professional training. The company also employs specific adaptations of such technologies as expert systems, knowledge bases, group decision support systems, intranets, knowledge repositories and also computer supported cooperative work.

Importance of Knowledge Management at the Wal-Mart Stores

Wal-Mart stores have created an indelible mark in the retail business industry. The giant chain store has an overwhelmingly effective knowledge management program that has enabled it to retain its competitive advantage even at times of turbulent economic situations. The corporation's goals are being achieved simultaneously by strictly adhering to business aims and balancing operations with the economic growth. Accordingly, economic gain as a result of positive knowledge management strategies is achieved throughout. The corporations' knowledge management strategies have enabled it to focus its strategic business operations into cutting operational costs and building up a value for its shareholders.

Knowledge can be classified into three broad forms, namely public, shared and personal knowledge (Alavi, 2001). Public knowledge can be accessed through public domains such as internet or books. Shared knowledge on the other hand refers to knowledge that is exclusively held by employees and is only used in work. Personal knowledge is the least accessible knowledge and is used mainly in work and daily life.

Through focused differentiation, the Wal-Mart has succeeded in cultivating enormous productivity and making. The corporation is now the leading of its type in not only the USA but the whole world itself. Nevertheless, the Wal-Mart stores have encountered numerous challenges in their attempt at introducing various knowledge management strategies. Change management analysis has greatly overridden the corporation's organizational growth (Alavi, Maryam, Leidner & Dorothy, 2001). Effective knowledge management at the Wal-Mart stores is largely contributed by its ushering in of big financial gains on annual production. This important and positive trend has reigned over all major stores associated with the corporation all over the world. The Wal-Mart stores are now ideal for businesses, organizations and units at the modern global age. Improved employee skills as one way of enhancing effective knowledge management by the corporation has

Contd....

Notes

proved extremely successful in countering challenges encountered threatening its overall mission (O'Leary, 2002).

The stores effective and reliable knowledge management skills have enhanced efficiency in service delivery and customer care as well as development of great competition. The corporation now seems success or failure of its business as a deliverable of staff productivity. The management's intellectuality focuses on strengthening the staff by way of transformational leadership. The company plans language to its employee's loyalty, diligence and good conduct and this has gone along way into ensuring long service in duty by the employees. Apt knowledge management in the company has enabled it to satisfactorily solve all issues related to organizational practices and conduct as well as make decisions and discover such issues as animosity.

How the Stores Apply Knowledge Management in Human Resource Management

Human resource strategies refer to the ways in organizations and business firms select, recruit, train and nurture their workforces for effective performance and delivery of services. Different human resource strategies are implemented by different organization and much as this would be different, they all achieve similar or almost same goals. Effective human resource strategies have enabled Wal-Mart to implement policies for management of its vast human resource. The application of effective human resource management strategies in the company is not a new phenomenon and this has enabled it receive much recognition in the recent years all over the world.

As an integral part of knowledge management, Wal-Mart's human resource management strategies are used chiefly to enable the organization and its vast array of branches manage their managerial activities through coordination of employee related processes. The effective use of these strategies has resulted to improved return on investments and has minimized the company's financial risks. This has largely been achieved by integrating skilled workforce with the firm's existing workforce. This is done in line with the organization's future business plans and requirements for securing its future success and survival. For that case, the organization implements human resource strategies taking into account the current state of its business operations as well as its future goals and objectives (O'Leary, 2002).

As Wolpert, (2001) says, the success of an organization more especially in the field of knowledge management is determined mainly by the kind of Human Resource management strategies displayed. Accordingly, Human Resource management strategies involve the power of an individual to persuade and inspire others towards attaining the objectives and goals of the organization. Managers across all Wal-Mart's branches motivate their workers to achieve the objectives and goals of the organization and outdo their rivals. As leaders, they daily interact with workers freely and offer directions personally and within teams, various departments as well as divisions. The effective and sound management of Wal-Mart is a big strength to the organization in achieving its objectives and remaining competitive.

Effective human resource management has contributed immensely to redefining knowledge management at Wal-Mart stores. In its efforts to remain competitive, the company has made substantial progress in relation to restructuring its long term objectives so that it can stand a better chance of achieving its goals.

Just like any other business, the Wal-Mart Stores Inc. has undertaken tremendous transformations with regard to adoption and implementation of the best knowledge management strategies. As an organization, it has found it useful for it to learn to manage

Contd....

its people and business operations in a well planned and coherent framework which is well reflected in the business' strategy. This way, all its branches have been able to ensure that all aspects of personnel management are reinforced for the purpose of improving and developing the performance of the business.

Knowledge Management with Regard to Communication and Feedback at the Wal-Mart

At the Wal-Mart stores, knowledge management is made into use with regard to conflict management and as such, a complex system of solving conflicts has been created. The company's intellectual negotiation capability enables it to differentiate between principles and discover prone areas and discard them. Ascending principles and self-esteem are two key issues that have been addressed appropriately through subjective application of knowledge management in enhancing effective communication in company.

Managers at the Wal-Mart have succeeded primarily due to their ability to manoeuvre the staff into undertaking positive transformations. The managers are very intellectual in tackling company targets and are able to dictate the decision making conclusions without necessarily having to court grievances from some quarters of the staff. The management is always ready to guide and protect the staff by providing articulate resources through which the staff and the management can share their knowledge management experiences and the management is always ready for unexpected issues raised by the staff. The environment at the Wal-Mart stores encourages team work with great focus on future growth.

Through the use of knowledge management practices, the Wal-Mart stores has managed to usher in a new form of employee motivation that focuses on employee retention till retirement and gives annual rewards to best performing employees with great emphasis on appraisals as well as potentiality in performance and productivity. For this reason, knowledge management at the Wal-Mart has greatly been enhanced by managers who have always ensure that the staff are well maintained and can get time offs and cash rewards which includes gift vouchers prioritizing family needs. Production and performance has been achieved by the collective responsibility of the staff and the company is now able to balance management roles with global trends in the same industry.

The Extent of Knowledge Management

Enhancing knowledge and talent management at the Wal-Mart's intellect for quick requirement has really widened its scope for global gains with respect to the corporation's targeted goals and tasks. At the Wal-Mart stores, assignments and projects are monitored continuously. Checking operation progress thoroughly ensures high productivity and performance for all staff. In the Wal-Mart jobs are always checked fully to ensure that staff and employers blend for one target of productivity and targeted tasks. Continuous need for productivity is essential for upheavals identification. Help in performance for staff has significantly elevated the appraisal times thereby enabling constant blending between managers and staff.

In enhancing the stores' acceleration of performance, it has been extremely essential for it to define its objectives. The company's staff is very competent throughout the branches and subsidiaries and thus, they can accept any new applications with respect to knowledge management. Performing required duties on all necessary management issues will result in excellence in the firm thus motivating the corporation's all staff to go up the ladder.

Promotions for staff annually by "on the Job Training" achieve a greater scope in implementation of the productivity and performance of staff and manager. At any moment, appraisal granting it necessary to staff by pay grade increases moral and motivation of all

Contd....

Notes

employees. An updated current summation of staff will comfort staff due to their frequent performances. It is quite satisfying for the staff at the Wal-Mart when they receive Awards and cash gifts from the manager to enhance their morals. Perfect performance is derived by the Manager of Wal-Mart selecting the best employee of the year annually to lift up the staff spirit and motivation by issuing gift vouchers as rewards. The conclusion for the Wal-Mart is achieved by a thorough management checked by performance. Need to check out behavioural perfection for all components blending in together.

At the Wal-Mart stores, staffs are a priority with union factors being checked out carefully, plus added value to stakeholders for stability. There should be no interference with staff due to less staff levelling. The recruitment should administer to correct quality and quantity, who will be accountants, admits, producing items, selling and managerial. Unions thus should be more co-operational to staff and managers. Staff priorities family even with perfect packages given with benefits, whose expectation and goal is deemed important to the firm. The staff are motivated properly and always focused for positivist throughout, and not just qualifications stashed away in office metal cabinets.

Wal-Mart's Concept Map

The need to have a sound knowledge management system is one key driving force behind Wal-Mart's success. The corporation maintains and acts upon awareness in a potentially changing scope and as such, the directions of many of its investment projects has been the major compelling factors in determining which strategy to best adopt. In this regard, Wal-Mart regularly has established and regularly assesses the scope of its knowledge management Endeavour. The corporation's experience has shown that failure to keep track of changes in knowledge management can lead to adverse elicitation effort that may well go into directions that do not serve the goals and aspirations of the business (Terreberry, 1968).

Wal-Mart's actual elicitation of knowledge takes different forms and these can be categorized broadly into those which are direct and the indirect. Direct methods of eliciting knowledge occur in collaboration with expert programs and are very efficient in knowledge management. These are chosen from a vast array of techniques some of which are based upon interviews and analysis of familiar techniques. Indirect methods of eliciting knowledge are on the other hand enhanced through study or use of relevant information resources as can be identified by knowledge engineer. Knowledge management and elicitation process as such leads to creation of concept maps that are included in the corporation's knowledge models.

Conclusion

Although knowledge management activities are all over the map in as far as Wal-Mart stores is concerned particularity in regard to technology development and its subsequent implementation, the stores have not yet fully reaped all the benefits that accrue from its implementation. Recent analyses on the corporation's pursuit of knowledge management strategies have demonstrated on the relationship between technological investment and improvement in business performance. All in all, the corporation seems to be faring quite well with regard to the way in which it is implementing knowledge management not only as a way of enhancing profitability and financial growth but also as away of remaining competitive in a turbulent world of businesses.

As of now, the company is contend with its leading position in the retail business on the global scale but still has more work ahead to ensure that it retains this commanding position. The realization of this will no doubt depend on the corporation's ability to

Contd....

bolster up its knowledge management strategies and also focus on offering services and products that meet its customer's diverse needs. Wal-Mart expects to enhance its future operations through the application of modern knowledge management techniques. This will to a great extent result in improved employee training, innovation, apt responsiveness and also a positive oriented organizational culture.

Question

Discuss the significance of knowledge management at the Wal-Mart stores.

Source: <http://seo-kisumu.hubpages.com/hub/knowledge-management-a-case-study-of-wallmart>

13.6 Summary

- Knowledge management (KM) is a process that helps organizations identify, select, organize, disseminate, and transfer important information and expertise that are part of the organization's memory.
- Knowledge acquisition is quite simply the process of acquiring knowledge that is available somewhere.
- Knowledge application means making knowledge more active and relevant for an organization in creating values.
- Knowledge development is a building block which complements knowledge acquisition. Its focus is on generating new skills, new products, better ideas and more efficient processes.
- Knowledge creation is the key focus about creating new knowledge or innovating existing knowledge for the organization.
- Knowledge Engineering is a manner more sophisticated than the transcription to capture parcels of tacit knowledge.
- Information technologies offer a potentially useful environment within which to build a multimedia repository for rich, explicit knowledge.
- A KM process is put in place to optimize knowledge use and evolution in an enterprise. It concerns in crucial manner knowledge actors: experts, specialists, competent people without whom the knowledge capital would not have added value.

13.7 Keywords

Knowledge Acquisition: Knowledge acquisition is quite simply the process of acquiring knowledge that is available somewhere.

Knowledge Application: Knowledge application means making knowledge more active and relevant for an organization in creating values.

Knowledge Creation: Knowledge creation is the key focus about creating new knowledge or innovating existing knowledge for the organization.

Knowledge Engineering: Knowledge Engineering is a manner more sophisticated than the transcription to capture parcels of tacit knowledge.

Knowledge Gap: It is the difference between what the enterprise requires and what it currently has.

Knowledge Management (KM): Knowledge management (KM) is a process that helps organizations identify, select, organize, disseminate, and transfer important information and expertise that are part of the organization's memory.

Notes

Knowledge Measurement: This is to measure the impact and effects after implementing knowledge management in an organization.

Knowledge Servers: These are systems, usually included in the intranet of the company, from which one can browse (in the most possible intelligent and ergonomic manner) the elicited knowledge.

13.8 Review Questions

1. Explain the concept of knowledge management.
2. Explain the various knowledge management activities.
3. Make distinction between knowledge acquisition and knowledge sharing.
4. Illustrate the conceptual framework for KM activities.
5. Describe the approaches for Knowledge Management.
6. Differentiate between Tacit knowledge elicitation and Knowledge Extraction.
7. The tacit/explicit approaches are not opposed but complementary. Comment.
8. Explain the different phases in a community development.
9. Explain the role of information technology and people in knowledge management.
10. Describe the success factors of knowledge management.

Answers: Self Assessment

- | | |
|--------------------------|--------------------------|
| 1. Knowledge gap | 2. Knowledge acquisition |
| 3. Knowledge application | 4. Knowledge development |
| 5. Knowledge creation | 6. Knowledge measurement |
| 7. Knowledge Engineering | 8. Texts |
| 9. Knowledge servers | 10. True |
| 11. False | 12. Dynamic |
| 13. Script | 14. Pilot project(s) |
| 15. Water lily | |

13.9 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.

Notes



Online links

<http://polaris.umuc.edu/~busilane/tman636/articles/12prin.pdf>

http://www.kyield.com/images/cko_competency_chapter_with_graphics.pdf

<http://www.sveiby.com/articles/KnowledgeManagement.html>

http://www.uky.edu/~gmswan3/575/IT_as_KM_Enabler.pdf

Unit 14: Knowledge-based Decision Support

CONTENTS

Objectives

Introduction

14.1 Artificial Intelligence

14.1.1 Definitions of Artificial Intelligence

14.1.2 Areas of Artificial Intelligence

14.2 Expert System

14.2.1 Rule-based System

14.2.2 Components of Rule-based Expert System

14.2.3 Theory of Rule-based Systems

14.2.4 Direction of Searching

14.2.5 Techniques of Searching

14.2.6 Creation of Expert Systems

14.2.7 Benefits of Expert System

14.2.8 Problems with Expert Systems

14.2.9 Limitations of Expert System

14.2.10 Success Factors

14.3 Summary

14.4 Keywords

14.5 Review Questions

14.6 Further Readings

Objectives

After studying this unit, you will be able to:

- Discuss the Concept of Artificial Intelligence
- Explain the Concept of Expert Systems
- Discuss Benefits, Problems and Limitations of Expert Systems
- Discuss the Success Factors of Expert System

Introduction

A knowledge-based decision support system is a system that can undertake intelligent tasks in a specific domain that is normally performed by highly skilled people. Typically, the success of such a system relies on the ability to represent the knowledge for a particular subject. Computerized decision support systems can be used by project participants to help make more informed decisions regarding the management of variation orders in projects by providing

access to useful, organized and timely information. It is important to understand that the KBDSS for the management of variation orders is not designed to make decisions for users, but rather it provides pertinent information in an efficient and easy-to-access format that allows users to make more informed decisions. The model contains two main components, i.e., a knowledge-base and a decision support shell, for selecting appropriate potential controls for variation orders for educational buildings. The database is developed through collecting data from source documents of the 79 educational projects, questionnaire survey, literature review and interview sessions with the professionals. The knowledge-base was developed through initial sieving and organization of data from the database. The segment that contained information pertinent to possible effects and controls of the causes of variation orders for educational buildings was integrated with a decision support shell. The decision support shell provided decision support through a structured process consisting of building the hierarchy among the main criteria and the suggested controls, rating the controls, and techniques. The KBDSS is developed in the MS Excel environment using numerous macros for developing the user-interface that carry out stipulated functions. These are incorporated within a decision support shell. In this unit, we will discuss the concepts of artificial intelligence and expert system.

14.1 Artificial Intelligence

Artificial Intelligence (AI) is a branch of Science which deals with helping machines finding solutions to complex problems in a more human-like fashion. This generally involves borrowing characteristics from human intelligence, and applying them as algorithms in a computer friendly way. A more or less flexible or efficient approach can be taken depending on the requirements established, which influences how artificial the intelligent behaviour appears.

AI is generally associated with Computer Science, but it has many important links with other fields such as Maths, Psychology, Cognition, Biology and Philosophy, among many others. Our ability to combine knowledge from all these fields will ultimately benefit our progress in the quest of creating an intelligent artificial being.

AI is one of the newest disciplines. It was formally initiated in 1956, when the name was coined, although at that point work had been under way for about five years. However, the study of intelligence is one of the oldest disciplines. For over 2000 years, philosophers have tried to understand how seeing, learning, remembering, and reasoning could, or should, be done. The advent of usable computers in the early 1950s turned the learned but armchair speculation concerning these mental faculties into a real experimental and theoretical discipline. Many felt that the new "Electronic Super-Brains" had unlimited potential for intelligence. "Faster than Einstein" was a typical headline. But as well as providing a vehicle for creating artificially intelligent entities, the computer provides a tool for testing theories of intelligence, and many theories failed to withstand the test. AI has turned out to be more difficult than many at first imagined, and modern ideas are much richer, more subtle, and more interesting as a result.

AI currently encompasses a huge variety of subfields, from general-purpose areas such as perception and logical reasoning, to specific tasks such as playing chess, proving mathematical theorems, writing poetry, and diagnosing diseases. Often, scientists in other fields move gradually into artificial intelligence, where they find the tools and vocabulary to systematize and automate the intellectual tasks on which they have been working all their lives. Similarly, workers in AI can choose to apply their methods to any area of human intellectual endeavour. In this sense, it is truly a universal field.

14.1.1 Definitions of Artificial Intelligence


It is often difficult to construct a definition of a discipline that is satisfying to all of its practitioners. AI research encompasses a spectrum of related topics. Broadly, AI is the computer-based

Notes

exploration of methods for solving challenging tasks that have traditionally depended on people for solution. Such tasks include complex logical inference, diagnosis, visual recognition, comprehension of natural language, game playing, explanation, and planning.

We shall begin our study of AI, by considering a number of alternative definitions of this Topic.

- AI is the study of how to make computers do things which at the moment people do better. This is ephemeral as it refers to the current state of computer science and it excludes a major area; problems that cannot be solved well either by computers or by people at the moment.
- AI is a field of study that encompasses computational techniques for performing tasks that apparently require intelligence when performed by humans.
- AI is the branch of computer science that is concerned with the automation of intelligent behaviour. AI is based upon the principles of computer science namely data structures used in knowledge representation, the algorithms needed to apply that knowledge and the languages and programming techniques used in their implementation. These definitions avoid philosophic discussions as to what is meant by artificial or intelligence.
- AI is the field of study that seeks to explain and emulate intelligent behaviour in terms of computational processes.
- AI is about generating representations and procedures that automatically or autonomously solve problems heretofore solved by humans.
- AI is the part of computer science concerned with designing intelligent computer systems, that is, computer systems that exhibit the characteristics we associate with intelligence in human behaviour such as understanding language, learning, reasoning and solving problems.
- AI is the study of mental faculties through the use of computational models.
- AI is the study of the computations that make it possible to perceive, reason, and act.
- AI is the exciting new effort to make computers think machines with minds, in the full and literal sense.



Notes AI is concerned with developing computer systems that can store knowledge and effectively use the knowledge to help solve problems and accomplish tasks. This brief statement sounds a lot like one of the commonly accepted goals in the education of humans. We want students to learn (gain knowledge) and to learn to use this knowledge to help solve problems and accomplish tasks.

The above definitions give us four possible goals to pursue in artificial intelligence:

- Systems that think like humans
- Systems that act like humans
- Systems that think rationally. (A system is rational if it does the right thing.)
- Systems that act rationally

Historically, all four approaches have been followed. As one might expect, a tension exists between approaches centered around humans and approaches centered around rationality. We should point out that by distinguishing between human and rational behaviour, we are not suggesting that humans are necessarily “irrational” in the sense of “emotionally unstable” or

“insane.” One merely need note that we often make mistakes; we are not all chess grandmasters even though we may know all the rules of chess; and unfortunately, not everyone gets an A on the exam.



Caution A human-centered approach must be an empirical science, involving hypothesis and experimental confirmation.

A rationalist approach involves a combination of mathematics and engineering. People in each group sometimes cast aspersions on work done in the other groups, but the truth is that each direction has yielded valuable insights.

Acting Humanly: The Turing Test Approach

The Turing Test, proposed by Alan Turing, was designed to provide a satisfactory operational definition of intelligence. Turing defined intelligent behaviour as the ability to achieve human-level performance in all cognitive tasks, sufficient to fool an interrogator. Roughly speaking, the test proposed is that the computer should be interrogated by a human via a teletype, and passes the test if the interrogator cannot tell if there is a computer or a human at the other end. Programming a computer to pass the test provides plenty to work on. The computer would need to possess the following capabilities:

- Natural language processing to enable it to communicate successfully in English (or some other human language);
- Knowledge representation to store information provided before or during the interrogation;
- Automated reasoning to use the stored information to answer questions and to draw new conclusions;
- Machine learning to adapt to new circumstances and to detect and extrapolate patterns.

Turing’s test deliberately avoided direct physical interaction between the interrogator and the computer, because physical simulation of a person is unnecessary for intelligence. However, the so-called total Turing Test includes a video signal so that the interrogator can test the subject’s perceptual abilities, as well as the opportunity for the interrogator to pass physical objects “through the hatch.”

Thinking Humanly: The Cognitive Modelling Approach

If we are going to say that a given program thinks like a human, we must have some way of determining how humans think. We need to get inside the actual workings of human minds. There are two ways to do this: through introspection (trying to catch our own thoughts as they go by) and through psychological experiments.

Once we have a sufficiently precise theory of the mind, it becomes possible to express the theory as a computer program. If the program’s input/output and timing behavior matches human behavior, that is evidence that some of the program’s mechanisms may also be operating in humans.



Example: Newell and Simon, who developed GPS, the “General Problem Solver” (Newell and Simon, 1961), were not content to have their program correctly solve problems. They were more concerned with comparing the trace of its reasoning steps to traces of human subjects solving the same problems. This is in contrast to other researchers of the same time (such as

Notes

Wang (1960)), who were concerned with getting the right answers regardless of how humans might do it. The interdisciplinary field of cognitive science brings together computer models from AI and experimental techniques from psychology to try to construct precise and testable theories of the workings of the human mind.

Thinking Rationally: The Laws of Thought Approach

The Greek philosopher Aristotle was one of the first to attempt to codify “right thinking,” that is, irrefutable reasoning processes. His famous syllogisms provided patterns for argument structures that always gave correct conclusions given correct premises.



Example: “Socrates is a man; all men are mortal; therefore Socrates is mortal.” These laws of thought were supposed to govern the operation of the mind, and initiated the field of logic.

The development of formal logic in the late nineteenth and early twentieth centuries, provided a precise notation for statements about all kinds of things in the world and the relations between them. (Contrast this with ordinary arithmetic notation, which provides mainly for equality and inequality statements about numbers.) By 1965, programs existed that could, given enough time and memory, take a description of a problem in logical notation and find the solution to the problem, if one exists. (If there is no solution, the program might never stop looking for it.) The so-called logicist tradition within artificial intelligence hopes to build on such programs to create intelligent systems.

There are two main obstacles to this approach. First, it is not easy to take informal knowledge and state it in the formal terms required by logical notation, particularly when the knowledge is less than 100% certain. Second, there is a big difference between being able to solve a problem “in principle” and doing so in practice. Even problems with just a few dozen facts can exhaust the computational resources of any computer unless it has some guidance as to which reasoning steps to try first. Although both of these obstacles apply to any attempt to build computational reasoning systems, they appeared first in the logicist tradition because the power of the representation and reasoning systems are well-defined and fairly well understood.

Acting Rationally: The Rational Agent Approach

Acting rationally means acting so as to achieve one’s goals, given one’s beliefs. An agent is just something that perceives and acts. In this approach, AI is viewed as the study and construction of rational agents.

In the “laws of thought” approach to AI, the whole emphasis was on correct inferences. Making correct inferences is sometimes part of being a rational agent, because one way to act rationally is to reason logically to the conclusion that a given action will achieve one’s goals, and then to act on that conclusion. On the other hand, correct inference is not all of rationality, because there are often situations where there is no provably correct thing to do, yet something must still be done. There are also ways of acting rationally that cannot be reasonably said to involve inference.



Example: Pulling one’s hand off of a hot stove is a reflex action that is more successful than a slower action taken after careful deliberation.

All the “cognitive skills” needed for the Turing Test are there to allow rational actions. Thus, we need the ability to represent knowledge and reason with it because this enables us to reach good decisions in a wide variety of situations. We need to be able to generate comprehensible sentences in natural language because saying those sentences helps us get by in a complex society.

We need learning not just for erudition, but because having a better idea of how the world works enables us to generate more effective strategies for dealing with it. We need visual perception not just because seeing is fun, but in order to get a better idea of what an action might achieve.

14.1.2 Areas of Artificial Intelligence

Various application areas of artificial intelligence are:

- **Robotics:** Although industrial robots have been expensive, robot hardware can be cheap: Radio Shack has sold a working robot arm and hand for \$15. The limiting factor in application of robotics is not the cost of the robot hardware itself. What is needed is perception and intelligence to tell the robot what to do; “blind” robots are limited to very well-structured tasks (like spray painting car bodies).
- **Planning:** Planning attempts to order actions to achieve goals. Planning applications include logistics, manufacturing scheduling, planning manufacturing steps to construct a desired product. There are huge amounts of money to be saved through better planning.
- **Expert Systems:** Expert Systems attempt to capture the knowledge of a human expert and make it available through a computer program. There have been many successful and economically valuable applications of expert systems. Expert systems provide the following benefits:
 - ❖ Reducing skill level needed to operate complex devices.
 - ❖ Diagnostic advice for device repair.
 - ❖ Interpretation of complex data.
 - ❖ “Cloning” of scarce expertise.
 - ❖ Capturing knowledge of expert who is about to retire.
 - ❖ Combining knowledge of multiple experts.
 - ❖ Intelligent training.
- **Theorem Proving:** Proving mathematical theorems might seem to be mainly of academic interest. However, many practical problems can be cast in terms of theorems. A general theorem prover can therefore be widely applicable.



Example:

- ❖ Automatic construction of compiler code generators from a description of a CPU’s instruction set.
- ❖ J Moore and colleagues proved correctness of the floating-point division algorithm on AMD CPU chip.
- **Symbolic Mathematics:** Symbolic mathematics refers to manipulation of formulas, rather than arithmetic on numeric values.
 - ❖ Algebra
 - ❖ Differential and Integral Calculus

Symbolic manipulation is often used in conjunction with ordinary scientific computation as a generator of programs used to actually do the calculations. Symbolic manipulation programs are an important component of scientific and engineering workstations.

Notes

- **Game Playing:** Games are good vehicles for research because they are well formalized, small, and self-contained. They are therefore easily programmed. Games can be good models of competitive situations, so principles discovered in game-playing programs may be applicable to practical problems.

Self Assessment

Fill in the blanks:

1. is a field of study that encompasses computational techniques for performing tasks that apparently require intelligence when performed by humans.
2. The, proposed by Alan Turing, was designed to provide a satisfactory operational definition of intelligence.
3. In the approach to AI, the whole emphasis was on correct inferences.
4. refers to manipulation of formulas, rather than arithmetic on numeric values.

14.2 Expert System

Expert system is programs that attempt to perform the duty of an expert in the problem domain in which it is defined. Expert systems are computer programs that have been constructed (with the assistance of human experts) in such a way that they are capable of functioning at the standard of (and sometimes even at a higher standard than) human experts in given fields that embody a depth and richness of knowledge that permit them to perform at the level of an expert.

14.2.1 Rule-based System

Using a set of assertions, which collectively form the ‘working memory’, and a set of rules that specify how to act on the assertion set, a rule-based system can be created. Rule-based systems are fairly simplistic, consisting of little more than a set of if-then statements, but provide the basis for so-called “expert systems” which are widely used in many fields. The concept of an expert system is this: the knowledge of an expert is encoded into the rule set. When exposed to the same data, the expert system will perform in a similar manner as the expert.

14.2.2 Components of Rule-based Expert System

Rule-based systems are a relatively simple model that can be adapted to any number of problems. To create a rule-based system for a given problem, you must have (or create) the following:

- A set of facts to represent the initial working memory. This should be anything relevant to the beginning state of the system.
- A set of rules. This should encompass any and all actions that should be taken within the scope of a problem, but nothing irrelevant. The number of rules in the system can affect its performance, so you don’t want any that aren’t needed.
- A condition that determines that a solution has been found or that none exists. This is necessary to terminate some rule-based systems that find themselves in infinite loops otherwise.

In fact, there are three essential components to a fully functional rule based expert system: the knowledge base, the working memory and the inference engine.

The Knowledge Base

Notes

The knowledge base is the store in which the knowledge in the particular domain is kept. The knowledge base stores information about the subject domain. However, this goes further than a passive collection of records in a database. Rather it contains symbolic representations of experts' knowledge, including definitions of domain terms, interconnections of component entities, and cause-effect relationships between these components. The knowledge in the knowledge based is expressed as a collection of fact and rule. Each fact expresses relationship between two or more object in the problem domain and can be expressed in term of predicates IF condition THEN conclusion where the condition or conclusion are fact or sets of fact connected by the logical connectives NOT, AND, OR.



Caution We need to create a variable name list to help to deal with long clumsy name and to help writing the rule.

| Variable Name | Meaning |
|----------------|---|
| Interest rise | interest rate rise |
| Interest fall | interest rate fall |
| Stock rise | market stock fall |
| Stock fall | market stock fall |
| Naira rise | exchange rate rise |
| Naira fall | exchange rate fall |
| Inflation rise | cost of market product rise |
| Fedmont add | increment of federal reserve money in circulation |
| Tax add | taxation on market product increase |
| Tax fall | taxation on market product decrease |

Using the above variable name, the following set of rule can then be constructed.

Rule 1: IF interest rise THEN stock fall

Rule 2: IF interest fall THEN stock rise

Rule 3: IF naira fall THEN interest fall

Rule 4: IF naira rise THEN interest rise

Rule 5 IF stock fall THEN inflation rise

Rule 6 IF fedmont add AND tax fall THEN naira rise

IF THEN rules are treated very differently from similar constructs in a natural programming language.

While natural programming languages treats IF THEN construct as part of a sequence of instructions, to be considered in order, the rule based system treats each rule as an independent chunk of knowledge, to be invoked when needed under the control of the interpreter. The rules are more like implication in logic (e.g. naira rise \rightarrow interest rise).

Notes

The Working Memory

The working memory is a temporal store that hold the fact produced during processing and possibly awaiting further processing produced by the Inference engine during its activities.



Did u know? Note that the working memory contains only facts and these fact are those produced during the searching process.

The Inference Engine

The core of any expert system is its inference engine. This is the part of expert system that manipulates the knowledge based to produce new fact in order to solve the given problem. An inference engine consists of search and reasoning procedures to enable the system to find solutions, and, if necessary, provide justifications for its answers. In this process it can used either forward or backward searching as a direction of search while applying some searching technique such as depth first search, breath first search, etc.

The roles of inference engine are:

- It identified the rule to be fired. The rule selected is the one whose conditional part is the same as the fact been considered in the case of forward chaining or the one whose conclusion part is the one as the fact been considered in the case of backward chaining.
- It resolve conflict when more than one rule satisfy the matching this is called conflict resolution which is based on certain criteria mentioned further.
- It recognizes the goal state. When the goal state is reached it report the conclusion of searching.

14.2.3 Theory of Rule-based Systems

The rule-based system itself uses a simple technique: It starts with a knowledge-base, which contains all of the appropriate knowledge encoded into IF-THEN rules, and a working memory, which may or may not initially contain any data, assertions or initially known information. The system examines all the rule conditions (IF) and determines a subset, the conflict set, of the rules whose conditions are satisfied based on the working memory. Of this conflict set, one of those rules is triggered (fired). Which one is chosen is based on a conflict resolution strategy. When the rule is fired, any actions specified in its THEN clause are carried out. These actions can modify the working memory, the rule base itself, or do just about anything else the system programmer decides to include. This loop of firing rules and performing actions continues until one of two conditions is met: there are no more rules whose conditions are satisfied or a rule is fired whose action specifies the program should terminate.

Which rule is chosen to fire is a function of the conflict resolution strategy. Which strategy is chosen can be determined by the problem or it may be a matter of preference. In any case, it is vital as it controls which of the applicable rules are fired and thus how the entire system behaves. There are several different strategies, but here are a few of the most common:

- **First Applicable:** If the rules are in a specified order, firing the first applicable one allows control over the order in which rules fire. This is the simplest strategy and has a potential for a large problem: that of an infinite loop on the same rule. If the working memory remains the same, as does the rule-base, then the conditions of the first rule have not changed and it will fire again and again. To solve this, it is a common practice to suspend

a fired rule and prevent it from re-firing until the data in working memory, that satisfied the rule's conditions, has changed.

- **Random:** Though it doesn't provide the predictability or control of the first-applicable strategy, it does have its advantages. For one thing, its unpredictability is an advantage in some circumstances (such as games). A random strategy simply chooses a single random rule to fire from the conflict set. Another possibility for a random strategy is a fuzzy rule-based system in which each of the rules has a probability such that some rules are more likely to fire than others.
- **Most Specific:** This strategy is based on the number of conditions of the rules. From the conflict set, the rule with the most conditions is chosen. This is based on the assumption that if it has the most conditions then it has the most relevance to the existing data.
- **Least Recently Used:** Each of the rules is accompanied by a time or step stamp, which marks the last time it was used. This maximizes the number of individual rules that are fired at least once.



Did u know? If all rules are needed for the solution of a given problem, this is a perfect strategy.

- **Best rule:** For this to work, each rule is given a 'weight,' which specifies how much it should be considered over the alternatives. The rule with the most preferable outcomes is chosen based on this weight.

14.2.4 Direction of Searching

There are two broad kinds of direction of searching in a rule-based system: forward chaining systems, and backward chaining systems. In a forward chaining system you start with the initial facts, and keep using the rules to draw new conclusions (or take certain actions) given those facts. In a backward chaining system you start with some hypothesis (or goal) you are trying to prove, and keep looking for rules that would allow you to conclude that hypothesis, perhaps setting new sub-goals to prove as you go. Forward chaining systems are primarily data-driven, while backward chaining systems are goal-driven. We'll look at both, and when each might be useful.

Forward Chaining Systems

In a forward chaining system the facts in the system are represented in a working memory which is continually updated as rules are invoked. Rules in the system represent possible actions to take when specified conditions hold on items in the working memory – they are sometimes called condition-action rules. The conditions are usually patterns that must match items in the working memory, while the actions usually involve adding or deleting items from the working memory.


The inference engine controls the application of the rules, given the working memory, thus controlling the system's activity. It is based on a cycle of activity sometimes known as a recognize-act cycle. The system first checks to find all the rules whose conditions hold, given the current state of working memory. It then selects one and performs the actions in the action part of the rule. (The selection of a rule to fire is based on fixed strategies, known as conflict resolution strategies.) The actions will result in a new working memory, and the cycle begins again. This cycle will be repeated until either no rules fire, or some specified goal state is satisfied.

Notes

Backward Chaining Systems


So far we have looked at how rule-based systems can be used to draw new conclusions from existing data, adding these conclusions to a working memory. This approach is most useful when you know all the initial facts, but don't have much idea what the conclusion might be. If you DO know what the conclusion might be, or have some specific hypothesis to test, forward chaining systems may be inefficient. You could keep on forward chaining until no more rules apply or you have added your hypothesis to the working memory. But in the process the system is likely to do a lot of irrelevant work, adding uninteresting conclusions to working memory.

This can be done by backward chaining from the goal state (or on some state that we are interested in). Given a goal state to try and prove (e.g., inflation rise) the system will first check to see if the goal matches the initial facts given. If it does, then that goal succeeds. If it doesn't the system will look for rules whose conclusions (previously referred to as actions) match the goal. One such rule will be chosen, and the system will then try to prove any facts in the preconditions of the rule using the same procedure, setting these as new goals to prove.



Notes A backward chaining system does not need to update a working memory. Instead it needs to keep track of what goals it needs to prove its main hypothesis.

In principle we can use the same set of rules for both forward and backward chaining. However, in practice we may choose to write the rules slightly differently if we are going to be using them for backward chaining. In backward chaining we are concerned with matching the conclusion of a rule against some goal that we are trying to prove. So the 'then' part of the rule is usually not expressed as an action to take but as a state which will be true if the premises are true.



Task Compare and contrast forward chaining system and backward chaining system.

14.2.5 Techniques of Searching

The order that rules fire may be crucial, especially when rules may result in items being deleted from working memory. The system must implement a searching technique that is used to process the knowledge base. Some of these technique are:

Depth First Search

In depth first search technique, the most recently fact added to the working memory is first selected for processing. We can thus implement it using stack so that the rule we have recently added to the working memory will be the one to be selected for the next cycle.

Breadth First Search

In the breath first search technique, the fact selected in the working memory for processing are selected in the order in which they were added in the working memory. We can use queue data structure to implement it. Since the rule to be processed will be selected in the front of the queue and the new fact are added at the rear of the queue.

14.2.6 Creation of Expert Systems

Notes

The best way to do this is to use an expert system shell. An expert system shell can be viewed as an expert system minus the domain knowledge (the analogy would be the difference between a database tool and a database system). It allows knowledge of a domain to be encoded in a specific format and put into the system to create expert systems for different domains. The advantage of using a shell is that it avoids the need for computer programming and allows the developer to focus only on the domain knowledge. This enables even non-computer professionals to create expert systems. For instance, one could create a system by just keying in the rules which were given. The shell would provide the interface, the inference engine and the explanation system.



Task Analyze the steps for creating expert systems.

14.2.7 Benefits of Expert System

The main benefits of expert systems are:-

- **Permanence:** Expert systems do not forget which may happen with human experts. Actually it is accumulation of experts knowledge.
- **Reproductive:** Many copies of an expert system can be made, but training new human experts is time-consuming and expensive.
- **Complex Problem Solving & Decision Making:** If there is a maze of rules (e.g. tax and auditing), then the expert system can work out the maze easily.
- **Productivity:** Expert Systems can increase output & productivity which may result in decreased personnel costs.
- **Inexpensive maintenance:** Although expert systems are expensive to build and maintain, they are inexpensive to operate. Development and maintenance costs can be spread over many users. The overall cost can be quite reasonable when compared to expensive and scarce human experts.
- **Consistency:** With expert systems similar transactions are handled in the same way. The system will make comparable recommendations for like situations whereas humans are influenced by recency effects (most recent information having a disproportionate impact on judgment) & primacy effects (early information dominates the judgment).
- **Documentation:** An expert system can provide permanent documentation of the decision process.
- **Completeness:** An expert system can review all the transactions where a human expert can only review a sample.
- **Timeliness:** Fraud and/or errors can be prevented. Information is available sooner for decision making.
- **Breadth:** The knowledge of multiple human experts can be combined to give a system more breadth than a single person is likely to achieve.

14.2.8 Problems with Expert Systems

On the technical side, there is the problem of the size of the database and using it efficiently. If the system consists of several thousand rules, it takes a very powerful control program to produce any conclusions in a reasonable amount of time. If the system also has a large quantity of information in the working memory, this will also slow things down unless you have a very good indexing and search system.

A second problem that comes from a large database is that as the number of rules increases the conflict set also becomes large so a good conflict resolving algorithm is needed if the system is to be usable.

Another problem that appears is that of responsibility.



Example: A system used by a doctor that is designed to administer drugs to patients according to their needs and that it must first determine what is wrong with them, very much like the prescribing work of a GP. If the system causes someone to take the wrong medicine and the person is harmed, who is legally responsible? Some would say the health authority who allowed the doctor to use the system, others would say the doctor, others the suppliers of the Expert System. A problem is produced that is not at all a trivial one. Think about the implications of using Expert Systems in other scenarios.

A more obvious problem is that of gathering the rules. Human experts are expensive and are not extremely likely to want to sit down and write out a large number of rules as to how they come to their conclusions. More to the point, they may not be able to. Although they will usually follow a logical path to their conclusions, putting these into a set of IF ... THEN rules may actually be very difficult and maybe impossible.

It is quite possible that many human experts, though starting off in their professions with a set of rules, learn to do their job through experiential knowledge and 'just know' what the correct solution is. Again they may have followed a logical path, but mentally they may have 'skipped some steps' along the way to get there. An Expert System cannot do this and needs to know the rules very clearly.

What may be a way round this problem is to enable Expert Systems to learn as they go, starting off with a smaller number of rules but given the ability to deduce new rules from what they know and what they 'experience'. This leads us very nicely into the field of Computer Learning.

14.2.9 Limitations of Expert System

However, Expert Systems suffer from following limitations:

- **Common sense:** In addition to a great deal of technical knowledge, human experts have common sense. It is not yet known how to give expert systems common sense.
- **Creativity:** Human experts can respond creatively to unusual situations whereas expert systems cannot.
- **Learning:** Human experts automatically adapt to changing environments; expert systems must be explicitly updated. Case-based reasoning and neural networks are methods that can incorporate learning.
- **Sensory Experience:** Human experts have available to them a wide range of sensory experience; expert systems are currently dependent on symbolic input.

14.2.10 Success Factors

Notes

Two major of success factors of expert system include:

- Champion in Management
- User Involvement and Training

Other success factors are:

- The level of knowledge must be sufficiently high
- There must be (at least) one cooperative expert
- The problem to be solved must be qualitative (fuzzy) not quantitative
- The problem must be sufficiently narrow in scope
- The ES shell must be high quality, and naturally store and manipulate the knowledge
- A friendly user interface
- The problem must be important and difficult enough
- Need knowledgeable and high quality system developers with good people skills
- The impact of ES as a source of end-users' job improvement must be favorable. End user attitudes and expectations must be considered
- Management support must be cultivated.

Expert systems need end-user training programs. The organizational environment should favor new technology adoption.

Self Assessment

Fill in the blanks:

5. is programs that attempt to perform the duty of an expert in the problem domain in which it is defined.
6. are fairly simplistic, consisting of little more than a set of if-then statements, but provide the basis for so-called "expert systems".
7. The is the store in which the knowledge in the particular domain is kept.
8. The in the knowledge base is expressed as a collection of fact and rule.
9. The is a temporal store that hold the fact produced during processing.
10. is the part of expert system that manipulates the knowledge based to produce new fact in order to solve the given problem.
11. In a system the facts in the system are represented in a working memory which is continually updated as rules are invoked.
12. A system does not need to update a working memory.
13. In technique, the most recently fact added to the working memory is first selected for processing.
14. In the technique, the fact selected in the working memory for processing are selected in the order in which they were added in the working memory.
15. The inference engine is based on a cycle of activity sometimes known as a

Notes



Case Study

Developing Healthcare Rule-based Expert Systems

The use of expert systems to generate automated alerts and patient instructions based on telemonitoring data could enable increased self-care and improve clinical management. However, of great importance is the development of the rule set to ensure safe and clinically relevant alerts and instructions are sent. The purpose of this work was to develop a rule-based expert system for a heart failure mobile phone-based telemonitoring system, to evaluate the expert system, and to generalize the lessons learned from the development process for use in other healthcare applications.

Methods

Semi-structured interviews were conducted with 10 heart failure clinicians to inform the development of a draft heart failure rule set for alerts and patient instructions. The draft rule set was validated and refined with 9 additional interviews with heart failure clinicians. Finally, the clinical champion of the project vetted the rule set. The concerns voiced by the clinicians during the interviews were noted, and methods to mitigate these concerns were employed. The rule set was then evaluated as part of a 6-month randomized controlled trial of a mobile phone-based heart failure telemonitoring system (n=50 for each of the telemonitoring and control groups).

Results

The developed expert system generated alerts and instructions based on the patient's weight, blood pressure, heart rate, and symptoms. During the trial, 1620 alerts were generated, which led to various clinical actions including 105 medication changes/instructions. The findings from the trial indicated the rule set was associated with improved quality of life and self-care.

Conclusions

A rule set was developed with extensive input by heart failure clinicians. The results from the trial indicated the rule set was associated with significantly increased self-care and improved the clinical management of heart failure. The developed rule set can be used as a basis for other heart failure telemonitoring systems, but should be validated and modified as necessary. In addition, the process used to develop the rule set can be generalized and applied to create robust and complete rule sets for other healthcare expert systems.

Question

Discuss the development of a rule set.

Source: <http://www.ncbi.nlm.nih.gov/pubmed/22465288>

14.3 Summary

- AI is a field of study that encompasses computational techniques for performing tasks that apparently require intelligence when performed by humans.
- The Turing Test, proposed by Alan Turing, was designed to provide a satisfactory operational definition of intelligence.
- Expert system is programs that attempt to perform the duty of an expert in the problem domain in which it is defined.

- Rule-based systems are fairly simplistic, consisting of little more than a set of if-then statements, but provide the basis for so-called “expert systems” which are widely used in many fields.
- The knowledge base is the store in which the knowledge in the particular domain is kept.
- The working memory is a temporal store that hold the fact produced during processing and possibly awaiting further processing produced by the Inference engine during its activities.
- An inference engine consists of search and reasoning procedures to enable the system to find solutions, and, if necessary, provide justifications for its answers.
- There are two broad kinds of direction of searching in a rule-based system: forward chaining systems, and backward chaining systems.

14.4 Keywords

Artificial Intelligence: AI is a field of study that encompasses computational techniques for performing tasks that apparently require intelligence when performed by humans.

BFS: In the breath first search technique, the fact selected in the working memory for processing are selected in the order in which they were added in the working memory.

Condition-action Rules: Rules in the system represent possible actions to take when specified conditions hold on items in the working memory – they are sometimes called condition-action rules.

DFS: In depth first search technique, the most recently fact added to the working memory is first selected for processing.

Expert System: Expert system is programs that attempt to perform the duty of an expert in the problem domain in which it is defined.

Inference Engine: An inference engine consists of search and reasoning procedures to enable the system to find solutions, and, if necessary, provide justifications for its answers.

Knowledge Base: The knowledge base is the store in which the knowledge in the particular domain is kept.

Working Memory: The working memory is a temporal store that hold the fact produced during processing and possibly awaiting further processing produced by the Inference engine during its activities.

14.5 Review Questions

1. Explain the concept of artificial intelligence.
2. Discuss the use of Turing Test in defining intelligent behavior.
3. Describe various applications areas of artificial intelligence.
4. What is an expert system? Discuss with example.
5. Explain the components of expert system.
6. Discuss the different strategies used in rule-based systems.
7. What are the different types of direction of searching in a rule-based system? Discuss.
8. Discuss the advantages and disadvantages of expert system.

Notes

9. Illustrate the problems associated with expert system.
10. In depth first search technique, the most recently fact added to the working memory is first selected for processing. Comment.

Answers: Self Assessment

- | | |
|----------------------------|-------------------------|
| 1. Artificial Intelligence | 2. Turing Test |
| 3. Laws of thought'' | 4. Symbolic mathematics |
| 5. Expert system | 6. Rule-based systems |
| 7. Knowledge base | 8. Knowledge |
| 9. Working memory | 10. Inference engine |
| 11. Forward chaining | 12. Backward chaining |
| 13. Depth first search | 14. Breath first search |
| 15. Recognize-act cycle | |

14.6 Further Readings



Books

Daniel Power, 2002, *Decision Support Systems: Concepts and Resources for Managers*, Greenwood Publishing Group

Efraim Turban, 1995, *Decision Support and Expert Systems: Management Support Systems*, Prentice Hall

Harry Katzan, 1984, *Management Support Systems*, Van Nostrand Reinhold Company

K. Sarukesi, 2004, *Decision Support Systems*, PHI Learning Pvt. Ltd.



Online links

<http://herkules.oulu.fi/isbn9514282051/isbn9514282051.pdf>

<http://www.csetube.in/2/post/2013/05/cs2057-knowledge-based-decision-support-system-introduction-notes-updated-8th-semester.html>

<http://www.jstor.org/discover/10.2307/40397838?uid=3738256&uid=2129&uid=2&uid=70&uid=4&sid=21102209397171>

<http://www.nrs.fs.fed.us/pubs/gtr/other/gtr-nc205/pdffiles/p84.PDF>

LOVELY PROFESSIONAL UNIVERSITY

Jalandhar-Delhi G.T. Road (NH-1)
Phagwara, Punjab (India)-144411
For Enquiry: +91-1824-521360
Fax.: +91-1824-506111
Email: odl@lpu.co.in

