



# **ANNAMACHARYA UNIVERSITY**

Estd. under Andhra Pradesh

Private Universities (Establishment and Regulation) Act, 2016

(University listed in UGC as per section 2(f) of the UGC Act, 1956)

## **LECTURE NOTES ON DATA WAREHOUSING AND DATA MINING (23A055CT)**

**B.TECH III Year - I Sem  
(2025-26)**

**DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING**

---



# ANNAMACHARYA UNIVERSITY

Estd. under Andhra Pradesh

Private Universities (Establishment and Regulation) Act, 2016

(University listed in UGC as per section 2(f) of the UGC Act, 1956)

Title of the Course: **Data Warehousing & Data Mining**  
Category: Professional Core  
Course Code: 23A055CT  
Year: III B. Tech  
Semester: I Semester  
Branch: CSE

| Lecture Hours | Tutorial Hours | Practice Hours | Credits |
|---------------|----------------|----------------|---------|
| 3             | -              | -              | 3       |

## Course Objectives:

1. Familiarize with mathematical foundations of data mining tools.
2. Introduce classical models and algorithms in data warehouses and data mining.
3. Investigate the kinds of patterns that can be discovered by association rule mining, classification and clustering.
4. Explore data mining techniques in various applications like social, scientific and environmental context.

## Course Outcomes:

Upon completion of the course, the students should be able to:

1. Design a Data warehouse system and perform business analysis with OLAP tools (L6).
2. Apply suitable pre-processing and visualization techniques for data analysis (L3)
3. Apply frequent pattern and association rule mining techniques for data analysis (L3)
4. Design appropriate classification and clustering techniques for data analysis (L6)
5. Infer knowledge from raw data (L4)

## Unit 1

09

Basic Concepts – Data Warehousing Components – Building a Data Warehouse – Database Architectures for Parallel Processing – Parallel DBMS Vendors – Multidimensional Data Model – Data Warehouse Schemas for Decision Support, Concept Hierarchies -Characteristics of OLAP Systems – Typical OLAP Operations, OLAP and OLTP.

## Unit 2

09

Introduction to Data Mining Systems – Knowledge Discovery Process – Data Mining Techniques – Issues – applications- Data Objects and attribute types, Statistical description of data, Data Preprocessing – Cleaning, Integration, Reduction, Transformation and discretization, Data Visualization, Data similarity and dissimilarity measures.

## Unit 3

08

Mining Frequent Patterns, Associations and Correlations – Mining Methods- Pattern Evaluation Method – Pattern Mining in Multilevel, Multi Dimensional Space – Constraint Based Frequent Pattern Mining, Classification using Frequent Patterns.

**Unit 4**

09

Decision Tree Induction – Bayesian Classification – Rule Based Classification – Classification by Back Propagation – Support Vector Machines — Lazy Learners – Model Evaluation and Selection- Techniques to improve Classification Accuracy. Clustering Techniques – Cluster analysis-Partitioning Methods – Hierarchical Methods – Density Based Methods – Grid Based Methods – Evaluation of clustering – Clustering high dimensional data- Clustering with constraints, Outlier analysis-outlier detection methods.

**Unit 5 WEKA TOOL**

08

Datasets – Introduction, Iris plants database, Breast cancer database, Auto imports database – Introduction to WEKA, The Explorer – Getting started, Exploring the explorer, Learning algorithms, Clustering algorithms, Association–rule learners.

**TEXT BOOK:**

1. Jiawei Han and Micheline Kamber, —Data Mining Concepts and Techniques, Third Edition, Elsevier, 2012.

**REFERENCES:**

1. Alex Berson and Stephen J.Smith, —Data Warehousing, Data Mining & OLAP, Tata McGraw – Hill Edition, 35th Reprint 2016.
2. K.P. Soman, Shyam Diwakar and V. Ajay, —Insight into Data Mining Theory and Practice, Eastern Economy Edition, Prentice Hall of India, 2006.
3. Ian H.Witten and Eibe Frank, —Data Mining: Practical Machine Learning Tools and Techniques, Elsevier, Second Edition.



# ANNAMACHARYA UNIVERSITY

Estd. under Andhra Pradesh

Private Universities (Establishment and Regulation) Act, 2016

(University listed in UGC as per section 2(f) of the UGC Act, 1956)

## DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

### DATA WAREHOUSING AND DATA MINING

#### (23A055CT)

##### INDEX

| Unit | Contents                                                               | Pg.No |
|------|------------------------------------------------------------------------|-------|
| I    | Introduction to Data warehouse & components                            | 2     |
|      | Data warehouse Design and Architecture                                 | 3     |
|      | Database Architectures for Parallel Processing – Parallel DBMS Vendors | 4-7   |
|      | Data warehouse Modelling,                                              | 8     |
|      | Schema Design                                                          | 9     |
|      | Measures                                                               | 10    |
|      | OLAP                                                                   | 11    |
| II   | Fundamentals of data mining                                            | 12    |
|      | Data Mining Functionalities                                            | 13    |
|      | Classification of Data Mining                                          | 16    |
|      | Major Issues in Data Mining                                            | 19    |
|      | Data Preprocessing                                                     | 23    |
| III  | Association Rule Mining                                                | 26    |
|      | Frequent Item set generation                                           | 29    |
|      | Apriori Algorithm                                                      | 30    |
|      | FP growth Algorithm                                                    | 34    |
|      | Compact Representation of Frequent Item set                            | 37    |
| IV   | Classification : General approaches                                    | 43    |
|      | Decision Tree Algorithm                                                | 45    |
|      | Naïve Bayes Classifier                                                 | 49    |
|      | K-Nearest Neighbor classification                                      | 56    |
|      | Prediction: Accuracy & Error Methods                                   | 60    |
|      | Ensemble methods                                                       | 62    |
|      | Clustering Overview                                                    | 64    |
|      | A categorization of major Clustering Methods                           | 67    |
|      | Partitioning clustering_ K-Means Algorithm                             | 71    |

| Unit | Contents                                                              | Pg.No |
|------|-----------------------------------------------------------------------|-------|
| v    | <b>Machine Learning with WEKA</b>                                     |       |
|      | Datasets – Introduction, Iris plants database, Breast cancer database | 3-13  |
|      | Auto imports database                                                 | 3-13  |
|      | Introduction to WEKA, The Explorer                                    | 3-13  |
|      | Getting started, Exploring the explorer                               | 3-13  |
|      | Learning algorithms                                                   | 6-32  |
|      | Clustering algorithms                                                 | 32-35 |
|      | Association-rule learners                                             | 32-35 |

## **UNIT-I**

### **Introduction to Data Warehouse:**

A data warehouse is a subject-oriented, integrated, time-variant and non-volatile collection of data in support of management's decision making process.

**Subject-Oriented:** A data warehouse can be used to analyze a particular subject area. For example, "sales" can be a particular subject.

**Integrated:** A data warehouse integrates data from multiple data sources. For example, source A and source B may have different ways of identifying a product, but in a data warehouse, there will be only a single way of identifying a product.

**Time-Variant:** Historical data is kept in a data warehouse. For example, one can retrieve data from 3 months, 6 months, 12 months, or even older data from a data warehouse. This contrasts with a transactions system, where often only the most recent data is kept. For example, a transaction system may hold the most recent address of a customer, where a data warehouse can hold all addresses associated with a customer.

**Non-volatile:** Once data is in the data warehouse, it will not change. So, historical data in a data warehouse should never be altered.

### **Data Warehouse Design Process:**

A data warehouse can be built using a top-down approach, a bottom-up approach, or a combination of both.


- The top-down approach starts with the overall design and planning. It is useful in cases where the technology is mature and well known, and where the business problems that must be solved are clear and well understood.
- The bottom-up approach starts with experiments and prototypes. This is useful in the early stage of business modeling and technology development. It allows an organization to move forward at considerably less expense and to evaluate the benefits of the technology before making significant commitments.
- In the combined approach, an organization can exploit the planned and strategic nature of the top-down approach while retaining the rapid implementation and opportunistic application of the bottom-up approach.

The warehouse design process consists of the following steps:



- Choose a business process to model, for example, orders, invoices, shipments, inventory, account administration, sales, or the general ledger. If the business process is organizational and involves multiple complex object collections, a data warehouse model should be followed. However, if the process is departmental and focuses on the analysis of one kind of business process, a data mart model should be chosen.
- Choose the grain of the business process. The grain is the fundamental, atomic level of data to be represented in the fact table for this process, for example, individual transactions, individual daily snapshots, and so on.
- Choose the dimensions that will apply to each fact table record. Typical dimensions are time, item, customer, supplier, warehouse, transaction type, and status.
- Choose the measures that will populate each fact table record. Typical measures are numeric additive quantities like dollars sold and units sold.

### A Three Tier Data Warehouse Architecture:



**Tier-1:**

The bottom tier is a warehouse database server that is almost always a relational database system. Back-end tools and utilities are used to feed data into the bottom tier from operational databases or other external sources (such as customer profile information provided by external consultants). These tools and utilities perform data extraction, cleaning, and transformation (e.g., to merge similar data from different sources into a unified format), as well as load and refresh functions to update the data warehouse. The data are extracted using application program interfaces known as gateways. A gateway is supported by the underlying DBMS and allows client programs to generate SQL code to be executed at a server. Examples of gateways include ODBC (Open Database Connection) and OLEDB (Open Linking and Embedding for Databases) by Microsoft and JDBC (Java Database Connection). This tier also contains a metadata repository, which stores information about the data warehouse and its contents.

**Tier-2:**

The middle tier is an OLAP server that is typically implemented using either a relational OLAP (ROLAP) model or a multidimensional OLAP.

- OLAP model is an extended relational DBMS that maps operations on multidimensional data to standard relational operations.
- A multidimensional OLAP (MOLAP) model, that is, a special-purpose server that directly implements multidimensional data and operations.

**Tier-3:**

The top tier is a front-end client layer, which contains query and reporting tools, analysis tools, and/or data mining tools (e.g., trend analysis, prediction, and so on).

# Data Warehouse Models:

There are three data warehouse models.

## 1. Enterprise warehouse:

- An enterprise warehouse collects all of the information about subjects spanning the entire organization.
- It provides corporate-wide data integration, usually from one or more operational systems or external information providers, and is cross-functional in scope.
- It typically contains detailed data as well as summarized data, and can range in size

---

from

few gigabytes to hundreds of gigabytes, terabytes, or beyond.

- An enterprise data warehouse may be implemented on traditional mainframes, computer super servers, or parallel architecture platforms. It requires extensive business modeling and may take years to design and build.

## 2. Data mart:

- A data mart contains a subset of corporate-wide data that is of value to a specific group of users. The scope is confined to specific selected subjects. For example, a marketing data mart may confine its subjects to customer, item, and sales. The data contained in data marts tend to be summarized.
- Data marts are usually implemented on low-cost departmental servers that are UNIX/LINUX- or Windows-based. The implementation cycle of a data mart is more likely to be measured in weeks rather than months or years. However, it may involve complex integration in the long run if its design and planning were not enterprise-wide.
- Depending on the source of data, data marts can be categorized as independent more dependent. Independent data marts are sourced from data captured from one or more operational systems or external information providers, or from data generated locally within a particular department or geographic area. Dependent data marts are source directly

from enterprise data warehouses.

### **3. Virtual warehouse:**

- A virtual warehouse is a set of views over operational databases. For efficient query processing, only some of the possible summary views may be materialized.
- A virtual warehouse is easy to build but requires excess capacity on operational database servers.

## **Meta Data Repository:**

Metadata are data about data. When used in a data warehouse, metadata are the data that define warehouse objects. Metadata are created for the data names and definitions of the given warehouse. Additional metadata are created and captured for time stamping any extracted data, the source of the extracted data, and missing fields that have been added by data cleaning or integration processes.

---

A metadata repository should contain the following:

- A description of the structure of the data warehouse, which includes the warehouse schema, view, dimensions, hierarchies, and derived data definitions, as well as data mart locations and contents.
- Operational metadata, which include data lineage (history of migrated data and the sequence of transformations applied to it), currency of data (active, archived, or purged), and monitoring information (warehouse usage statistics, error reports, and audit trails).
- The algorithms used for summarization, which include measure and dimension definition algorithms, data on granularity, partitions, subject areas, aggregation, summarization, and predefined queries and reports.
- The mapping from the operational environment to the data warehouse, which includes source databases and their contents, gateway descriptions, data partitions, data extraction, cleaning, transformation rules and defaults, data refresh and purging rules, and security (user authorization and access control).

- Data related to system performance, which include indices and profiles that improve data access and retrieval performance, in addition to rules for the timing and scheduling of refresh, update, and replication cycles.
- Business metadata, which include business terms and definitions, data ownership information, and charging policies.

---

## **Introduction to Parallel Databases**

Companies need to handle huge amount of data with high data transfer rate. The client server and centralized system is not much efficient. The need to improve the efficiency gave birth to the concept of Parallel Databases.

Parallel database system improves performance of data processing using multiple resources in parallel, like multiple CPU and disks are used parallelly.

It also performs many parallelization operations like, data loading and query processing.

### **Goals of Parallel Databases**

The concept of Parallel Database was built with a goal to: Improve performance:

The performance of the system can be improved by connecting multiple CPU and disks in parallel. Many small processors can also be connected in parallel.

#### **Improve availability of data:**


Data can be copied to multiple locations to improve the availability of data. For example: if a module contains a relation (table in database) which is unavailable then it is important to make it available from another module.

#### **Improve reliability:**

Reliability of system is improved with completeness, accuracy and availability of data.

### **Provide distributed access of data:**

Companies having many branches in multiple cities can access data with the help of parallel database system.




### **Types of Parallel Database Architecture**

#### **Shared memory system**

Shared memory system uses multiple processors which is attached to a global shared memory via intercommunication channel or communication bus.

Shared memory system have large amount of cache memory at each processors, so referencing of the shared memory is avoided.

If a processor performs a write operation to memory location, the data should be updated or removed from that location.



#### **Shared Memory System in Parallel Databases**

### **Advantages of Shared memory system**

- Data is easily accessible to any processor.
- One processor can send message to other efficiently.

### **Disadvantages of Shared memory system**


- Waiting time of processors is increased due to more number of processors.
- Bandwidth problem.

### **Shared Disk System**

Shared disk system uses multiple processors which are accessible to multiple disks via intercommunication channel and every processor has local memory.

Each processor has its own memory so the data sharing is efficient.

The system built around this system are called as clusters.



### **Shared disk system in Parallel Databases**

#### **Advantages of Shared Disk System**

- Fault tolerance is achieved using shared disk system.
- Fault tolerance: If a processor or its memory fails, the other processor can complete the task. This is called as fault tolerance.

#### **Disadvantage of Shared Disk System**

- Shared disk system has limited scalability as large amount of data travels through the interconnection channel.
- If more processors are added the existing processors are slowed down.

#### **Applications of Shared Disk System**

Digital Equipment Corporation(DEC): DEC cluster running relational databases use


the shared disk system and now owned by Oracle.

### **Shared nothing disk system**

Each processor in the shared nothing system has its own local memory and local disk.

Processors can communicate with each other through intercommunication channel.

Any processor can act as a server to serve the data which is stored on local disk.



### **Shared nothing disk system in Parallel Databases**

#### **Advantages of Shared nothing disk system**

- Number of processors and disk can be connected as per the requirement in share nothing disk system.
- Shared nothing disk system can support for many processor, which makes the system more scalable.

#### **Disadvantages of Shared nothing disk system**

- Data partitioning is required in shared nothing disk system.
- Cost of communication for accessing local disk is much higher.

#### **Applications of Shared nothing disk system**

- Tera data database machine.
- The Grace and Gamma research prototypes.

## Hierarchical System or Non-Uniform Memory Architecture

Hierarchical model system is a hybrid of shared memory system, shared disk system and shared nothing system.

Hierarchical model is also known as Non-Uniform Memory Architecture (NUMA).

In this system each group of processor has a local memory. But processors from other groups can access memory which is associated with the other group in coherent.

NUMA uses local and remote memory(Memory from other group), hence it will take longer time to communicate with each other.

### Advantages of NUMA

Improves the scalability of the system.

Memory bottleneck(shortage of memory) problem is minimized in this architecture.

### Disadvantages of NUMA

The cost of the architecture is higher compared to other architectures.

## Schema Design:


Stars, Snowflakes, and Fact Constellations: Schemas for Multidimensional Databases The entity-relationship data model is commonly used in the design of relational databases, where a database schema consists of a set of entities and the relationships between them. Such a data model is appropriate for on-line transaction processing. A data warehouse, however, requires a concise, subject-oriented schema that facilitates on-line data analysis. The most popular data model for a data warehouse is a multidimensional model. Such a model can exist in the form of a star schema, a snowflake schema, or a fact constellation schema. Let's look at each of these schema types. Star schema: The most common modeling paradigm is the star schema, in which the data warehouse contains (1) a large central table (fact table) containing the bulk of the data, with no redundancy, and (2) a set of smaller attendant tables (dimension tables), one for each dimension. The schema graph resembles a starburst, with the dimension tables displayed in a radial pattern around the central fact table.

## Star schema:

A star schema for AllElectronics sales is shown in Figure. Sales are considered along four dimensions, namely, time, item, branch, and location. The schema contains a central fact table for sales that contains keys to each of the four dimensions, along with two measures: dollars sold and units sold. To minimize the size of the fact table, dimension identifiers (such as time key and item key) are system-generated identifiers. Notice that in the star schema, each dimension is represented by only one table, and each table contains a set of attributes. For example, the location dimension table contains the attribute set {location key, street, city, province or state, country}. This constraint may introduce some redundancy.

For example, “Vancouver” and “Victoria” are both cities in the Canadian province of British Columbia. Entries for such cities in the location dimension table will create redundancy among the attributes province or state and country, that is, (... , Vancouver, British Columbia, Canada) and (... , Victoria, British Columbia, Canada). Moreover, the attributes within a dimension table may form either a hierarchy (total order) or a lattice (partial order).

---




---

Star schema of a data warehouse for sales.

### Snowflake schema.:

A snowflake schema for AllElectronics sales is given in Figure Here, the sales fact table is identical to that of the star schema in Figure . The main difference between the two schemas is in the definition of dimension tables.

The single dimension table for item in the star schema is normalized in the snowflake schema, resulting in new item and supplier tables. For example, the item dimension table now contains the attributes item key, item name, brand, type, and supplier key, where supplier key is linked to the supplier dimension table, containing supplier key and supplier type information. Similarly, the single dimension table for location in the star schema can be normalized into two new tables: location and city. The city key in the new location table links to the city dimension. Notice that further normalization can be performed on province or state and country in the snowflake schema

---




Figure 10.1 Snowflake schema of a data warehouse for sales.


## Fact constellation.

A fact constellation schema is shown in Figure. This schema specifies two fact tables, sales and shipping. The sales table definition is identical to that of the star schema. The shipping table has five dimensions, or keys: item key, time key, shipper key, from location, and to location, and two measures: dollars cost and units shipped.

A fact constellation schema allows dimension tables to be shared between fact tables. For example, the dimensions tables for time, item, and location are shared between both the sales and shipping fact tables.

In data warehousing, there is a distinction between a data warehouse and a data mart.

A data warehouse collects information about subjects that span the entire organization, such as customers, items, sales, assets, and personnel, and thus its scope is enterprise-wide. For data warehouses, the fact constellation schema is commonly used, since it can model multiple, interrelated subjects. A data mart, on the other hand, is a department subset of the data warehouse that focuses on selected subjects, and thus its scope is department wide. For data marts, the star or snowflake schema are commonly used, since both are geared toward modeling single subjects, although the star schema is more popular and efficient.



5 Fact constellation schema of a data warehouse for sales and shipping.

## Measures: Their Categorization and Computation:

“How are measures computed?” To answer this question, we first study how measures can be categorized.<sup>1</sup> Note that a multidimensional point in the data cube space can be defined by a set of dimension-value pairs, for example,  $htime = \text{“Q1”}$ ,  $location = \text{“Vancouver”}$ ,  $item = \text{“computer”}$ . A data cube measure is a numerical function that can be evaluated at each point in the data cube space. A measure value is computed for a given point by aggregating the data corresponding to the respective dimension-value pairs defining the given point.

Measures can be organized into three categories (i.e., distributive, algebraic, holistic), based on the kind of aggregate functions used.

**Distributive:** An aggregate function is distributive if it can be computed in a distributed manner as follows. Suppose the data are partitioned into  $n$  sets. We apply the function to each partition, resulting in  $n$  aggregate values. If the result derived by applying the function to the  $n$  aggregate values is the same as that derived by applying the function to the entire data set (without partitioning), the function can be computed in a distributed manner. For example, `count()` can be computed for a data cube by first partitioning the cube into a set of subcubes, computing `count()` for each subcube, and then summing up the counts obtained for each

subcube. Hence, count() is a distributive aggregate function. For the same reason, sum(), min(), and max() are distributive aggregate functions. A measure is distributive if it is obtained by applying a distributive aggregate function. Distributive measures can be computed efficiently because they can be computed in a distributive manner.

## **OLAP(Online analytical Processing):**

- OLAP is an approach to answering multi-dimensional analytical (MDA) queries swiftly.
- OLAP is part of the broader category of business intelligence, which also encompasses relational database, report writing and data mining.
- OLAP tools enable users to analyze multidimensional data interactively from multiple perspectives.

OLAP consists of three basic analytical operations:

- Consolidation (Roll-Up)
  - Drill-Down
  - Slicing And Dicing
- Consolidation involves the aggregation of data that can be accumulated and computed in one or more dimensions. For example, all sales offices are rolled up to the sales department or sales division to anticipate sales trends.
  - The drill-down is a technique that allows users to navigate through the details. For instance, users can view the sales by individual products that make up a region's sales.
  - Slicing and dicing is a feature whereby users can take out (slicing) a specific set of data of the OLAP cube and view (dicing) the slices from different viewpoints.

## **Types of OLAP:**

### **1. Relational OLAP (ROLAP):**

- ROLAP works directly with relational databases. The base data and the dimension tables are stored as relational tables and new tables are created to hold the aggregated information. It depends on a specialized schema design.
  - This methodology relies on manipulating the data stored in the relational database to
-

give the appearance of traditional OLAP's slicing and dicing functionality. In essence, each action of slicing and dicing is equivalent to adding a "WHERE" clause in the SQL statement. ROLAP tools do not use pre-calculated data cubes but instead pose the query to the standard relational database and its tables in order to bring back the data required to answer the question.

- ROLAP tools feature the ability to ask any question because the methodology does not limit to the contents of a cube. ROLAP also has the ability to drill down to the lowest level of detail in the database.

## **2. Multidimensional OLAP (MOLAP):**

- MOLAP is the 'classic' form of OLAP and is sometimes referred to as just OLAP.
- MOLAP stores this data in an optimized multi-dimensional array storage, rather than in a relational database. Therefore it requires the pre-computation and storage of information in the cube - the operation known as processing.
- MOLAP tools generally utilize a pre-calculated data set referred to as a data cube. The data cube contains all the possible answers to a given range of questions.
- MOLAP tools have a very fast response time and the ability to quickly write back data into the data set.

## **3. Hybrid OLAP (HOLAP):**

- There is no clear agreement across the industry as to what constitutes Hybrid OLAP, except that a database will divide data between relational and specialized storage.
  - For example, for some vendors, a HOLAP database will use relational tables to hold the larger quantities of detailed data, and use specialized storage for at least some aspects of the smaller quantities of more-aggregate or less-detailed data.
  - HOLAP addresses the shortcomings of MOLAP and ROLAP by combining the capabilities of both approaches.
  - HOLAP tools can utilize both pre-calculated cubes and relational data sources.
-

## UNIT-2

### **Fundamentals of Data Mining:**

Data mining refers to extracting or mining knowledge from large amounts of data. The term is actually a misnomer. Thus, data mining should have been more appropriately named as knowledge mining which emphasizes on mining from large amounts of data.

It is the computational process of discovering patterns in large data sets involving methods at the intersection of artificial intelligence, machine learning, statistics, and database systems.

The overall goal of the data mining process is to extract information from a data set and transform it into an understandable structure for further use.

The key properties of data mining are

- Automatic discovery of patterns
- Prediction of likely outcomes
- Creation of actionable information
- Focus on large datasets and databases

### **The Scope of Data Mining**

Data mining derives its name from the similarities between searching for valuable business information in a large database — for example, finding linked products in gigabytes of store scanner data — and mining a mountain for a vein of valuable ore. Both processes require either sifting through an immense amount of material, or intelligently probing it to find exactly where the value resides.

Given databases of sufficient size and quality, data mining technology can generate new business opportunities by providing these capabilities:

**Automated prediction of trends and behaviors.** Data mining automates the process of finding predictive information in large databases. Questions that traditionally required extensive hands-on analysis can now be answered directly from the data — quickly.

A typical example of a predictive problem is targeted marketing. Data mining uses data on past promotional mailings to identify the targets most likely to maximize return on investment in

---

future mailings. Other predictive problems include forecasting bankruptcy and other forms of default, and identifying segments of a population likely to respond similarly to given events.

### **Automated discovery of previously unknown patterns.**

Data mining tools sweep through databases and identify previously hidden patterns in one step. An example of pattern discovery is the analysis of retail sales data to identify seemingly unrelated products that are often purchased together. Other pattern discovery problems include detecting fraudulent credit card transactions and identifying anomalous data that could represent data entry keying errors.

### **Data Mining Functionalities:**

We have observed various types of databases and information repositories on which datamining can be performed. Let us now examine the kinds of data patterns that can be mined. Data mining functionalities are used to specify the kind of patterns to be found in data mining tasks. In general, data mining tasks can be classified into two categories: descriptive and predictive. Descriptive mining tasks characterize the general properties of the data in the database. Predictive mining tasks perform inference on the current data in order to make predictions.

In some cases, users may have no idea regarding what kinds of patterns in their data may be interesting, and hence may like to search for several different kinds of patterns in parallel. Thus it is important to have a data mining system that can mine multiple kinds of patterns to accommodate different user expectations or applications. Furthermore, data mining systems should be able to discover patterns at various granularity (i.e., different levels of abstraction). Data mining systems should also allow users to specify hints to guide or focus the search for interesting patterns. Because some patterns may not hold for all of the data in the database, a measure of certainty or “trustworthiness” is usually associated with each discovered pattern.

Data mining functionalities, and the kinds of patterns they can discover, are described Mining Frequent Patterns, Associations, and Correlations Frequent patterns, as the name suggests, are patterns that occur frequently in data. There are many kinds of frequent patterns, including itemsets, subsequences, and substructures.

A frequent itemset typically refers to a set of items that frequently appear together in a transactional data set, such as milk and bread. A frequently occurring subsequence, such as the pattern that customers tend to purchase first a PC, followed by a digital camera, and then a memory card, is a (frequent) sequential pattern. A substructure can

---

refer to different structural forms, such as graphs, trees, or lattices, which may be combined with itemsets or subsequences. If a substructure occurs frequently, it is called a (frequent) structured pattern. Mining frequent patterns leads to the discovery of interesting associations and correlations within data.


below.

Data mining involves six common classes of tasks:

- **Anomaly detection (Outlier/change/deviation detection)** – The identification of unusual data records, that might be interesting or data errors that require further investigation.
- **Association rule learning (Dependency modelling)** – Searches for relationships between variables. For example a supermarket might gather data on customer purchasing habits. Using association rule learning, the supermarket can determine which products are frequently bought together and use this information for marketing purposes. This is sometimes referred to as market basket analysis.
- **Clustering** – is the task of discovering groups and structures in the data that are in some way or another "similar", without using known structures in the data.
- **Classification** – is the task of generalizing known structure to apply to new data. For example, an e-mail program might attempt to classify an e-mail as "legitimate" or as "spam".
- **Regression** – attempts to find a function which models the data with the least error.
- **Summarization** – providing a more compact representation of the data set, including Visualization and report generation.

# Architecture of Data Mining

A typical data mining system may have the following major components.



## 1. Knowledge Base:

This is the domain knowledge that is used to guide the search or evaluate the interestingness of resulting patterns. Such knowledge can include concept hierarchies, used to organize attributes or attribute values into different levels of abstraction.

Knowledge such as user beliefs, which can be used to assess a pattern's interestingness based on its unexpectedness, may also be included. Other examples of domain knowledge are additional interestingness constraints or thresholds, and metadata (e.g., describing data from multiple heterogeneous sources).

## 2. Data Mining Engine:

This is essential to the data mining system and ideally consists of a set of functional modules for tasks such as characterization, association and correlation analysis, classification,

prediction, cluster analysis, outlier analysis, and evolution analysis.

### **3. Pattern Evaluation Module:**

This component typically employs interestingness measures interacts with the data mining modules so as to focus the search toward interesting patterns. It may use interestingness thresholds to filter out discovered patterns. Alternatively, the pattern evaluation module may be integrated with the mining module, depending on the implementation of the datamining method used. For efficient data mining, it is highly recommended to push the evaluation of pattern interestingness as deep as possible into the mining process as to confine the search to only the interesting patterns.

### **4. User interface:**


This module communicates between users and the data mining system, allowing the user to interact with the system by specifying a data mining query or task, providing information to help focus the search, and performing exploratory datamining based on the intermediate data mining results. In addition, this component allows the user to browse database and data warehouse schemas or data structures, evaluate mined patterns, and visualize the patterns in different forms.

## **Classification of Data Mining Systems**

Data mining is an interdisciplinary field, the confluence of a set of disciplines, including database systems, statistics, machine learning, visualization, and information science .Moreover, depending on the data mining approach used, techniques from other disciplines may be applied, such as neural networks, fuzzy and/or rough set theory, knowledge representation, inductive logic programming, or high-performance computing. Depending on the kinds of data to be mined or on the given data mining application, the data mining system may also integrate techniques from spatial data analysis, information retrieval, pattern recognition, image analysis, signal processing, computer graphics,

Web technology, economics, business, bioinformatics, or psychology. Because of the diversity of disciplines contributing to datamining, datamining research is expected to generate a large variety of data mining systems. Therefore, it is necessary to provide a clear classification of data mining systems, which may help potential users distinguish between such systems and identify those that best match their needs.

---



Data mining systems can be categorized according to various criteria, as follows:

Classification according to the kinds of databases mined: A data mining system can be classified according to the kinds of databases mined. Database systems can be classified according to different criteria (such as data models, or the types of data or applications involved), each of which may require its own data mining technique. Data mining systems can therefore be classified accordingly.

For instance, if classifying according to data models, we may have a relational, transactional, object-relational, or data warehouse mining system. If classifying according to the special types of data handled, we may have a spatial, time-series, text, stream data, multimedia data mining system, or a World Wide Web mining system.

Classification according to the kinds of knowledge mined: Data mining systems can be categorized according to the kinds of knowledge they mine, that is, based on data mining functionalities, such as characterization, discrimination, association and correlation analysis, classification, prediction, clustering, outlier analysis, and evolution analysis. A comprehensive data mining system usually provides multiple and/or integrated data mining functionalities.

Moreover, data mining systems can be distinguished based on the granularity or levels of abstraction of the knowledge mined, including generalized knowledge (at a high level of abstraction), primitive-level knowledge (at a raw data level), or knowledge at multiple levels (considering several levels of abstraction). An advanced data mining system should facilitate the discovery of knowledge at multiple levels of abstraction.

Data mining systems can also be categorized as those that mine data regularities (commonly occurring patterns) versus those that mine data irregularities (such as exceptions, or outliers). In general, concept description, association and correlation analysis, classification, prediction, and clustering mine data regularities, rejecting outliers as noise. These methods may also help detect outliers.

Classification according to the kinds of techniques utilized: Data mining systems can be categorized according to the underlying data mining techniques employed. These techniques can be described according to the degree of user interaction involved (e.g., autonomous systems, interactive exploratory systems, query-driven systems) or the methods of data analysis employed (e.g., database-oriented or data warehouse-oriented techniques, machine learning, statistics, visualization, pattern recognition, neural networks, and so on). A sophisticated data mining system will often adopt multiple data mining techniques or work out an effective, integrated technique that combines the merits of a few individual approaches.

Classification according to the applications adapted: Data mining systems can also be categorized according to the applications they adapt. For example, data mining systems may be tailored specifically for finance, telecommunications, DNA, stock markets, e-mail, and so on. Different applications often require the integration of application-specific methods. Therefore, a generic, all-purpose data mining system may not fit domain-specific mining tasks.

## **Data Mining Process:**

Data Mining is a process of discovering various models, summaries, and derived values from a given collection of data.

The general experimental procedure adapted to data-mining problems involves the following steps:

### **1. State the problem and formulate the hypothesis**

Most data-based modeling studies are performed in a particular application domain. Hence, domain-specific knowledge and experience are usually necessary in order to come up with a meaningful problem statement. Unfortunately, many application studies tend to focus on the data-mining technique at the expense of a clear problem statement. In this step, a modeler usually specifies a set of variables for the unknown dependency and, if possible, a general form of this dependency as an initial hypothesis. There may be several hypotheses formulated for a single problem at this stage.

---

The first step requires the combined expertise of an application domain and a data-mining model. In practice, it usually means a close interaction between the data-mining expert and the application expert. In successful data-mining applications, this cooperation does not stop in the initial phase; it continues during the entire data-mining process.

## **2. Collect the data**

This step is concerned with how the data are generated and collected. In general, there are two distinct possibilities. The first is when the data-generation process is under the control of an expert (modeler): this approach is known as a designed experiment.

The second possibility is when the expert cannot influence the data-generation process: this is known as the observational approach. An observational setting, namely, random data generation, is assumed in most data-mining applications.

Typically, the sampling distribution is completely unknown after data are collected, or it is partially and implicitly given in the data-collection procedure. It is very important, however, to understand how data collection affects its theoretical distribution, since such a priori knowledge can be very useful for modeling and, later, for the final interpretation of results. Also, it is important to make sure that the data used for estimating a model and the data used later for testing and applying a model come from the same, unknown, sampling distribution. If this is not the case, the estimated model cannot be successfully used in a final application of the results.

## **Major Issues In Data Mining:**

- Mining different kinds of knowledge in databases.** - The need of different users is not the same. And Different user may be in interested in different kind of knowledge. Therefore it is necessary for data mining to cover broad range of knowledge discovery task.
  - Interactive mining of knowledge at multiple levels of abstraction.** - The data mining process needs to be interactive because it allows users to focus the search for patterns, providing and refining data mining requests based on returned results.
  - Incorporation of background knowledge.** - To guide discovery process and to express the discovered patterns, the background knowledge can be used. Background knowledge may be used to express the discovered patterns not only in concise terms but at multiple level of abstraction.
-

●**Data mining query languages and ad hoc data mining.** - Data Mining Query language that allows the user to describe ad hoc mining tasks, should be integrated with a data warehouse query language and optimized for efficient and flexible data mining.

●**Presentation and visualization of data mining results.** - Once the patterns are discovered it needs to be expressed in high level languages, visual representations. This representations should be easily understandable by the users.

●**Handling noisy or incomplete data.** - The data cleaning methods are required that can handle the noise, incomplete objects while mining the data regularities. If data cleaning methods are not there then the accuracy of the discovered patterns will be poor.

●**Pattern evaluation.** - It refers to interestingness of the problem. The patterns discovered should be interesting because either they represent common knowledge or lack novelty.

- **Efficiency and scalability of data mining algorithms.** - In order to effectively extract the information from huge amount of data in databases, data mining algorithm must be efficient and scalable.
- **Parallel, distributed, and incremental mining algorithms.** - The factors such as huge size of databases, wide distribution of data, and complexity of data mining methods motivate the development of parallel and distributed data mining algorithms. These algorithms divide the data into partitions which is further processed parallel. Then the results from the partitions are merged. The incremental algorithms, updates the databases without having to mine the data again from the scratch.

## **Data Integration:**


It combines data from multiple sources into a coherent data store, as in data warehousing. These sources may include multiple databases, data cubes, or flat files.

The data integration systems are formally defined as triple  $\langle G, S, M \rangle$

Where G: The global schema

S: Heterogeneous source of schemas

M: Mapping between the queries of source and global schema



## Issues in Data integration:

### 1. Schema integration and object matching:

How can the data analyst or the computer be sure that customer id in one database and customer number in another reference to the same attribute.

### 2. Redundancy:

An attribute (such as annual revenue, for instance) may be redundant if it can be derived from another attribute or set of attributes. Inconsistencies in attribute or dimension naming can also cause redundancies in the resulting data set.

### 3. detection and resolution of data value conflicts:

For the same real-world entity, attribute values from different sources may differ.

## Data Transformation:

In data transformation, the data are transformed or consolidated into forms appropriate for mining.

Data transformation can involve the following:

---

- **Smoothing**, which works to remove noise from the data. Such techniques include binning, regression, and clustering.
- **Aggregation**, where summary or aggregation operations are applied to the data. For example, the daily sales data may be aggregated so as to compute monthly and annual total amounts. This step is typically used in constructing a data cube for analysis of the data at multiple granularities. **Generalization of the data**, where low-level or —primitive (raw) data are replaced by higher-level concepts through the use of concept hierarchies. For example, categorical attributes, like street, can be generalized to higher-level concepts, like city or country.
- **Normalization**, where the attribute data are scaled so as to fall within a small specified range, such as 1:0 to 1:0, or 0:0 to 1:0.
- **Attribute construction** (or feature construction), where new attributes are constructed and added from the given set of attributes to help the mining process.

## Data Reduction:

Data reduction techniques can be applied to obtain a reduced representation of the data set that is much smaller in volume, yet closely maintains the integrity of the original data. That is, mining on the reduced data set should be more efficient yet produce the same (or almost the same) analytical results.

Strategies for data reduction include the following:

- **Data cube aggregation**, where aggregation operations are applied to the data in the construction of a data cube.
- **Attribute subset selection**, where irrelevant, weakly relevant, or redundant attributes or dimensions may be detected and removed.
- **Dimensionality reduction**, where encoding mechanisms are used to reduce the dataset size.
- **Numerosity reduction**, where the data are replaced or estimated by alternative, smaller data representations such as parametric models (which need store only the model parameters instead of the actual data) or nonparametric methods such as clustering, sampling, and the use of histograms.
- **Discretization and concept hierarchy generation**, where raw data values for attributes

are replaced by ranges or higher conceptual levels. Data discretization is a form of numerosity reduction that is very useful for the automatic generation of concept hierarchies. Discretization and concept hierarchy generation are powerful tools for data mining, in that they allow the mining of data at multiple levels of abstraction.

## **Data Preprocessing:**

In the observational setting, data are usually "collected" from the existing databases, data warehouses, and data marts. Data preprocessing usually includes at least two common tasks:

**1. Outlier detection (and removal)** – Outliers are unusual data values that are not consistent with most observations. Commonly, outliers result from measurement errors, coding and recording errors, and, sometimes, are natural, abnormal values. Such non representative samples can seriously affect the model produced later. There are two strategies for dealing with outliers:

- a. Detect and eventually remove outliers as a part of the preprocessing phase, or
- b. Develop robust modeling methods that are insensitive to outliers.

## **2. Scaling, encoding, and selecting features –**

Data preprocessing includes several steps such as variable scaling and different types of encoding. For example, one feature with the range  $[0, 1]$  and the other with the range  $[-100, 1000]$  will not have the same weights in the applied technique; they will also influence the final data-mining results differently. Therefore, it is recommended to scale them and bring both features to the same weight for further analysis.

Also, application-specific encoding methods usually achieve dimensionality reduction by providing a smaller number of informative features for subsequent data modeling. These two classes of preprocessing tasks are only illustrative examples of a large spectrum of preprocessing activities in a data-mining process.

Data-preprocessing steps should not be considered completely independent from other data-mining phases. In every iteration of the data-mining process, all activities, together, could define new and improved data sets for subsequent iterations.

Generally, a good preprocessing method provides an optimal representation for a data-mining

technique by incorporating a priori knowledge in the form of application-specific scaling and encoding.

#### **4. Estimate the model**

The selection and implementation of the appropriate data-mining technique is the main task in this phase. This process is not straightforward; usually, in practice, the implementation is based on several models, and selecting the best one is an additional task.


Interpret the model and draw conclusions

In most cases, data-mining models should help in decision making. Hence, such models need to be interpretable in order to be useful because humans are not likely to base their decisions on complex "black-box" models. Note that the goals of accuracy of the model and accuracy of its interpretation are somewhat contradictory.

Usually, simple models are more interpretable, but they are also less accurate. Modern data-mining methods are expected to yield highly accurate results using high dimensional models. The problem of interpreting these models, also very important, is considered a separate task, with specific techniques to validate the results.

A user does not want hundreds of pages of numeric results. He does not understand them; he cannot summarize, interpret, and use them for successful decision making.

---



## UNIT-III

### Association Rule Mining:

- Association rule mining is a popular and well researched method for discovering interesting relations between variables in large databases.
- It is intended to identify strong rules discovered in databases using different measures of interestingness.
- Based on the concept of strong rules, Rakesh Agrawal et al. introduced association rules.

### Problem Definition:

The problem of association rule mining is defined as:

Let  $I = \{i_1, i_2, \dots, i_n\}$  be a set of  $n$  binary attributes called *items*.

Let  $D = \{t_1, t_2, \dots, t_m\}$  be a set of transactions called the *database*.

Each transaction in  $D$  has a unique transaction ID and contains a subset of the items in  $I$ .

A *rule* is defined as an implication of the form  $X \Rightarrow Y$  where

$X, Y \subseteq I$  and  $X \cap Y = \emptyset$ .

The sets of items (for short *itemsets*)  $X$  and  $Y$  are called *antecedent* (left-hand-side or LHS) and *consequent* (right-hand-side or RHS) of the rule respectively.

### Example:

To illustrate the concepts, we use a small example from the supermarket domain. The set of items is  $I = \{\text{milk, bread, butter, beer}\}$  and a small database containing the items (1 codes presence and 0 absence of an item in a transaction) is shown in the table.

An example rule for the supermarket could be  $\{\text{butter, bread}\} \Rightarrow \{\text{milk}\}$  meaning that if butter and bread are bought, customers also buy milk.

Example database with 4 items and 5 transactions

| Transaction ID | milk | bread | butter | beer |
|----------------|------|-------|--------|------|
| 1              | 1    | 1     | 0      | 0    |
| 2              | 0    | 0     | 1      | 0    |
| 3              | 0    | 0     | 0      | 1    |

|   |  |   |   |   |   |
|---|--|---|---|---|---|
| 4 |  | 1 | 1 | 1 | 0 |
| 5 |  | 0 | 1 | 0 | 0 |

### Important concepts of Association Rule Mining:

- The **support**  $\text{supp}(X)$  of an itemset  $X$  is defined as the proportion of transactions in the data set which contain the itemset. In the example database, the itemset  $\{\text{milk, bread, butter}\}$  has a support of  $1/5 = 0.2$  since it occurs in 20% of all transactions (1 out of 5 transactions).

- The **confidence** of a rule is defined  

$$\text{conf}(X \Rightarrow Y) = \text{supp}(X \cup Y) / \text{supp}(X)$$
For example, the rule  $\{\text{butter, bread}\} \Rightarrow \{\text{milk}\}$  has a confidence of  $0.2/0.2 = 1.0$  in the database, which means that for 100% of the transactions containing butter and bread the rule is correct (100% of the times a customer buys butter and bread, milk is bought as well). Confidence can be interpreted as an estimate of the probability  $P(Y|X)$ , the probability of finding the RHS of the rule in transactions under the condition that these transactions also contain the LHS.

- The **lift** of a rule is defined as  

$$\text{lift}(X \Rightarrow Y) = \frac{\text{supp}(X \cup Y)}{\text{supp}(X) \times \text{supp}(Y)}$$

or the ratio of the observed support to that expected if  $X$  and  $Y$  were independent. The rule

$\{\text{milk, bread}\} \Rightarrow \{\text{butter}\}$  has a lift of  $\frac{0.2}{0.4 \times 0.4} = 1.25$ .

- The **conviction** of a rule is defined as


$$\text{conv}(X \Rightarrow Y) = \frac{1 - \text{supp}(Y)}{1 - \text{conf}(X \Rightarrow Y)}.$$

The rule  $\{\text{milk, bread}\} \Rightarrow \{\text{butter}\}$  has a conviction of  $\frac{1 - 0.4}{1 - .5} = 1.2$ ,

and can be interpreted as the ratio of the expected frequency that X occurs without Y (that is to say, the frequency that the rule makes an incorrect prediction) if X and Y were independent divided by the observed frequency of incorrect predictions.

## Market basket analysis:

This process analyzes customer buying habits by finding associations between the different items that customers place in their shopping baskets. The discovery of such associations can help retailers develop marketing strategies by gaining insight into which items are frequently purchased together by customers. For instance, if customers are buying milk, how likely are they to also buy bread (and what kind of bread) on the same trip to the supermarket. Such information can lead to increased sales by helping retailers do selective marketing and plan their shelfspace.



### Example:

If customers who purchase computers also tend to buy anti virus software at the same time, then placing the hardware display close to the software display may help increase the sales of both items. In an alternative strategy, placing hardware and software at opposite ends of the store may entice customers who purchase such items to pick up other items along the way. For instance, after deciding on an expensive computer, a customer may observe security systems for sale while heading toward the software display to purchase antivirus software and may decide to purchase a

home security system as well. Market basket analysis can also help retailers plan which items to put on sale at reduced prices. If customers tend to purchase computers and printers together, then having a sale on printers may encourage the sale of printers *as well as* computers.

## Frequent Pattern Mining:

Frequent pattern mining can be classified in various ways, based on the following criteria:

### 1. Based on the completeness of patterns to be mined:

- We can mine the complete set of frequent itemsets, the closed frequent itemsets, and the maximal frequent itemsets, given a minimum support threshold.
- We can also mine constrained frequent itemsets, approximate frequent itemsets, near-match frequent itemsets, top-k frequent itemsets and so on.

### 2. Based on the levels of abstraction involved in the rule set:

Some methods for association rule mining can find rules at differing levels of abstraction.

For example, suppose that a set of association rules mined includes the following rules where X is a variable representing a customer:

$$\text{buys}(X, \text{—computer}\|) \Rightarrow \text{buys}(X, \text{—HP printer}\|) \quad (1)$$

$$\text{buys}(X, \text{—laptop computer}\|) \Rightarrow \text{buys}(X, \text{—HP printer}\|) \quad (2)$$

In rule (1) and (2), the items bought are referenced at different levels of abstraction (e.g., *—computer* is a higher-level abstraction of *—laptop computer*).

### 3. Based on the number of data dimensions involved in the rule:

- If the items or attributes in an association rule reference only one dimension, then it is a single-dimensional association rule.  
 $\text{buys}(X, \text{—computer}\|) \Rightarrow \text{buys}(X, \text{—antivirus software}\|)$
- If a rule references two or more dimensions, such as the dimensions age, income, and buys, then it is a multidimensional association rule. The following rule is an example of a multidimensional rule:

$$\text{age}(X, \text{—30,31...39}\|) \wedge \text{income}(X, \text{—42K,...48K}\|) \Rightarrow \text{buys}(X, \text{—high resolution TV}\|)$$

#### **4. Based on the types of values handled in the rule:**

- If a rule involves associations between the presence or absence of items, it is a Boolean association rule.
- If a rule describes associations between quantitative items or attributes, then it is a quantitative association rule.

#### **5. Based on the kinds of rules to be mined:**

- Frequent pattern analysis can generate various kinds of rules and other interesting relationships.
- Association rule mining can generate a large number of rules, many of which are redundant or do not indicate a correlation relationship among itemsets.
- The discovered associations can be further analyzed to uncover statistical correlations, leading to correlation rules.

#### **6. Based on the kinds of patterns to be mined:**

- Many kinds of frequent patterns can be mined from different kinds of data sets.
- Sequential pattern mining searches for frequent subsequences in a sequence data set, where a sequence records an ordering of events.
- For example, with sequential pattern mining, we can study the order in which items are frequently purchased. For instance, customers may tend to first buy a PC, followed by a digital camera, and then a memory card.
- Structured pattern mining searches for frequent sub structures in a structured data set.
- Single items are the simplest form of structure.
- Each element of an itemset may contain a subsequence, a subtree, and so on.
- Therefore, structured pattern mining can be considered as the most general form of frequent pattern mining.

### **Apriori Algorithm:**

#### **Finding Frequent Itemsets Using Candidate Generation: The Apriori Algorithm**

- Apriori is a seminal algorithm proposed by R. Agrawal and R. Srikant in 1994 for mining frequent itemsets for Boolean association rules.
  - The name of the algorithm is based on the fact that the algorithm uses *prior knowledge* of frequent itemset properties.
-

- Apriori employs an iterative approach known as a *level-wise* search, where  $k$ -itemsets are used to explore  $(k+1)$ -itemsets.
- First, the set of frequent 1-itemsets is found by scanning the database to accumulate the count for each item, and collecting those items that satisfy minimum support. The resulting set is denoted  $L_1$ . Next,  $L_1$  is used to find  $L_2$ , the set of frequent 2-itemsets, which is used to find  $L_3$ , and so on, until no more frequent  $k$ -itemsets can be found.
- The finding of each  $L_k$  requires one full scan of the database.
- A two-step process is followed in Apriori consisting of join and prune action.

**Algorithm: Apriori.** Find frequent itemsets using an iterative level-wise approach based on candidate generation.

**Input:**

- $D$ , a database of transactions;
- $min\_sup$ , the minimum support count threshold.

**Output:**  $L$ , frequent itemsets in  $D$ .

**Method:**

```

(1)  $L_1 = \text{find\_frequent\_1-itemsets}(D)$ ;
(2) for  $(k = 2; L_{k-1} \neq \emptyset; k++)$  {
(3)    $C_k = \text{apriori\_gen}(L_{k-1})$ ;
(4)   for each transaction  $t \in D$  { // scan  $D$  for counts
(5)      $C_t = \text{subset}(C_k, t)$ ; // get the subsets of  $t$  that are candidates
(6)     for each candidate  $c \in C_t$ 
(7)        $c.\text{count}++$ ;
(8)   }
(9)    $L_k = \{c \in C_k \mid c.\text{count} \geq min\_sup\}$ 
(10) }
(11) return  $L = \cup_k L_k$ ;

procedure apriori_gen( $L_{k-1}$ : frequent  $(k-1)$ -itemsets)
(1) for each itemset  $l_1 \in L_{k-1}$ 
(2)   for each itemset  $l_2 \in L_{k-1}$ 
(3)     if  $(l_1[1] = l_2[1]) \wedge (l_1[2] = l_2[2]) \wedge \dots \wedge (l_1[k-2] = l_2[k-2]) \wedge (l_1[k-1] < l_2[k-1])$  then {
(4)        $c = l_1 \bowtie l_2$ ; // join step: generate candidates
(5)       if has_infrequent_subset( $c, L_{k-1}$ ) then
(6)         delete  $c$ ; // prune step: remove unfruitful candidate
(7)       else add  $c$  to  $C_k$ ;
(8)     }
(9) return  $C_k$ ;

procedure has_infrequent_subset( $c$ : candidate  $k$ -itemset;
                                 $L_{k-1}$ : frequent  $(k-1)$ -itemsets); // use prior knowledge
(1) for each  $(k-1)$ -subset  $s$  of  $c$ 
(2)   if  $s \notin L_{k-1}$  then
(3)     return TRUE;
(4) return FALSE;
```

**Example:**


| TID  | List of item IDs |
|------|------------------|
| T100 | I1, I2, I5       |
| T200 | I2, I4           |
| T300 | I2, I3           |

|      |                |
|------|----------------|
| T400 | I1, I2, I4     |
| T500 | I1, I3         |
| T600 | I2, I3         |
| T700 | I1, I3         |
| T800 | I1, I2, I3, I5 |
| T900 | I1, I2, I3     |

There are nine transactions in this database, that is,  $|D| = 9$ .

### Steps:

1. In the first iteration of the algorithm, each item is a member of the set of candidate 1-itemsets,  $C_1$ . The algorithm simply scans all of the transactions in order to count the number of occurrences of each item.
2. Suppose that the minimum support count required is 2, that is,  $\min \text{sup} = 2$ . The set of frequent 1-itemsets,  $L_1$ , can then be determined. It consists of the candidate 1-itemsets satisfying minimum support. In our example, all of the candidates in  $C_1$  satisfy minimum support.
3. To discover the set of frequent 2-itemsets,  $L_2$ , the algorithm uses the join  $L_1$  on  $L_1$  to generate a candidate set of 2-itemsets,  $C_2$ . No candidates are removed from  $C_2$  during the prune step because each subset of the candidates is also frequent.
4. Next, the transactions in  $D$  are scanned and the support count of each candidate itemset in  $C_2$  is accumulated.
5. The set of frequent 2-itemsets,  $L_2$ , is then determined, consisting of those candidate 2-itemsets in  $C_2$  having minimum support.
6. The generation of the set of candidate 3-itemsets,  $C_3$ , From the join step, we first get  $C_3 = L_2 \times L_2 = (\{I1, I2, I3\}, \{I1, I2, I5\}, \{I1, I3, I5\}, \{I2, I3, I4\}, \{I2, I3, I5\}, \{I2, I4, I5\})$ . Based on the Apriori property that all subsets of a frequent itemset must also be frequent, we can determine that the four latter candidates cannot possibly be frequent.
7. The transactions in  $D$  are scanned in order to determine  $L_3$ , consisting of those candidate 3-itemsets in  $C_3$  having minimum support.
8. The algorithm uses  $L_3 \times L_3$  to generate a candidate set of 4-itemsets,  $C_4$ .



Generation of candidate itemsets and frequent itemsets, where the minimum support count is 2.

- (a) Join:  $C_3 = L_2 \bowtie L_2 = \{\{I1, I2\}, \{I1, I3\}, \{I1, I5\}, \{I2, I3\}, \{I2, I4\}, \{I2, I5\}\} \bowtie \{\{I1, I2\}, \{I1, I3\}, \{I1, I5\}, \{I2, I3\}, \{I2, I4\}, \{I2, I5\}\}$   
 $= \{\{I1, I2, I3\}, \{I1, I2, I5\}, \{I1, I3, I5\}, \{I2, I3, I4\}, \{I2, I3, I5\}, \{I2, I4, I5\}\}.$
- (b) Prune using the Apriori property: All nonempty subsets of a frequent itemset must also be frequent. Do any of the candidates have a subset that is not frequent?
- The 2-item subsets of {I1, I2, I3} are {I1, I2}, {I1, I3}, and {I2, I3}. All 2-item subsets of {I1, I2, I3} are members of  $L_2$ . Therefore, keep {I1, I2, I3} in  $C_3$ .
  - The 2-item subsets of {I1, I2, I5} are {I1, I2}, {I1, I5}, and {I2, I5}. All 2-item subsets of {I1, I2, I5} are members of  $L_2$ . Therefore, keep {I1, I2, I5} in  $C_3$ .
  - The 2-item subsets of {I1, I3, I5} are {I1, I3}, {I1, I5}, and {I3, I5}. {I3, I5} is not a member of  $L_2$ , and so it is not frequent. Therefore, remove {I1, I3, I5} from  $C_3$ .
  - The 2-item subsets of {I2, I3, I4} are {I2, I3}, {I2, I4}, and {I3, I4}. {I3, I4} is not a member of  $L_2$ , and so it is not frequent. Therefore, remove {I2, I3, I4} from  $C_3$ .
  - The 2-item subsets of {I2, I3, I5} are {I2, I3}, {I2, I5}, and {I3, I5}. {I3, I5} is not a member of  $L_2$ , and so it is not frequent. Therefore, remove {I2, I3, I5} from  $C_3$ .
  - The 2-item subsets of {I2, I4, I5} are {I2, I4}, {I2, I5}, and {I4, I5}. {I4, I5} is not a member of  $L_2$ , and so it is not frequent. Therefore, remove {I2, I4, I5} from  $C_3$ .
- (c) Therefore,  $C_3 = \{\{I1, I2, I3\}, \{I1, I2, I5\}\}$  after pruning.

Generation and pruning of candidate 3-itemsets,  $C_3$ , from  $L_2$  using the Apriori property.


### **FP-growth (finding frequent itemsets without candidate generation).**

We re-examine the mining of transaction database,  $D$ , of Table 5.1 in Example 5.3 using the frequent pattern growth approach.

The first scan of the database is the same as Apriori, which derives the set of frequent items (1-itemsets) and their support counts (frequencies). Let the minimum support count be 2. The set of frequent items is sorted in the order of descending support count. This resulting set or *list* is denoted  $L$ .

An FP-tree is then constructed as follows. First, create the root of the tree, labeled with “null.” Scan database  $D$  a second time. The items in each transaction are processed in  $L$  order (i.e., sorted according to descending support count), and a branch is created for each transaction. For example, the scan of the first transaction, “T100: I1, I2, I5,” which contains three items (I2, I1, I5 in  $L$  order), leads to the construction of the first branch of the tree with three nodes, hI2: 1i, hI1:1i, and hI5: 1i, where I2 is linked as a child of the root, I1 is linked to I2, and I5 is linked to I1. The second transaction, T200, contains the items I2 and I4 in  $L$  order, which would result in a branch where I2 is linked to the root and I4 is linked to I2. However, this branch would share a common prefix, I2, with the existing path for T100. Therefore, we instead increment the count of the I2 node by 1, and create a new node, hI4: 1i, which is linked as a child of hI2: 2i. In general, when considering the branch to be added for a transaction, the count of each node along a common prefix is incremented by 1, and nodes for the items following the prefix are created and linked accordingly.

To facilitate tree traversal, an item header table is built so that each item points to its occurrences in the tree via a chain of node-links. The tree obtained after scanning all of the transactions is shown in Figure 5.7 with the associated node-links. In this way, the problem of mining frequent patterns in databases is transformed to that of mining the FP-tree.



An FP-tree registers compressed, frequent pattern information.

Mining the FP-tree by creating conditional (sub-)pattern bases.

| Item | Conditional Pattern Base                  | Conditional FP-tree                                   | Frequent Patterns Generated                       |
|------|-------------------------------------------|-------------------------------------------------------|---------------------------------------------------|
| I5   | $\{\{I2, I1: 1\}, \{I2, I1, I3: 1\}\}$    | $\langle I2: 2, I1: 2 \rangle$                        | $\{I2, I5: 2\}, \{I1, I5: 2\}, \{I2, I1, I5: 2\}$ |
| I4   | $\{\{I2, I1: 1\}, \{I2: 1\}\}$            | $\langle I2: 2 \rangle$                               | $\{I2, I4: 2\}$                                   |
| I3   | $\{\{I2, I1: 2\}, \{I2: 2\}, \{I1: 2\}\}$ | $\langle I2: 4, I1: 2 \rangle, \langle I1: 2 \rangle$ | $\{I2, I3: 4\}, \{I1, I3: 4\}, \{I2, I1, I3: 2\}$ |
| I1   | $\{\{I2: 4\}\}$                           | $\langle I2: 4 \rangle$                               | $\{I2, I1: 4\}$                                   |

The FP-tree is mined as follows.

Start from each frequent length-1 pattern (as an initial suffix pattern), construct its conditional pattern base (a “subdatabase,” which consists of the set of prefix paths in the FP-tree co-occurring with the suffix pattern), then construct its (conditional) FP-tree, and perform mining recursively on such a tree. The pattern growth is achieved by the concatenation of the suffix pattern with the frequent patterns generated from a conditional FP-tree.

Mining of the FP-tree is summarized in Table 5.2 and detailed as follows. We first consider I5, which is the last item in L, rather than the first. The reason for starting at the end of the list will become apparent as we explain the FP-tree mining process. I5 occurs in two branches of the FP-tree of Figure 5.7. (The occurrences of I5 can easily be found by following its chain of node-links.) The paths formed by these branches are hI2, I1, I5: 1i and hI2, I1, I3, I5: 1i.

Therefore, considering I5 as a suffix, its corresponding two prefix paths are hI2, I1: 1i and hI2, I1, I3: 1i, which form its conditional pattern base. Its conditional FP-tree contains only a single path, hI2: 2, I1: 2i; I3 is not included because its support count of 1 is less than the minimum support count.

The single path generates all the combinations of frequent patterns: fI2, I5: 2g, fI1, I5: 2g, fI2, I1, I5: 2g.

Generating Association Rules from Frequent Itemsets:

Once the frequent itemsets from transactions in a database  $D$  have been found, it is straightforward to

generate strong association rules from them.

$$\text{confidence}(A \Rightarrow B) = P(B|A) = \frac{\text{support\_count}(A \cup B)}{\text{support\_count}(A)}.$$

The conditional probability is expressed in terms of itemset support count, where  $\text{support\_count}(A \cup B)$  is the number of transactions containing the itemsets  $A \cup B$ , and  $\text{support\_count}(A)$  is the number of transactions containing the itemset  $A$ . Based on this equation, association rules can be generated as follows:


- For each frequent itemset  $l$ , generate all nonempty subsets of  $l$ .
- For every nonempty subset  $s$  of  $l$ , output the rule " $s \Rightarrow (l - s)$ " if  $\frac{\text{support\_count}(l)}{\text{support\_count}(s)} \geq \text{min\_conf}$ , where  $\text{min\_conf}$  is the minimum confidence threshold.

## **Compact Representation of Frequent Item Set:**

- For many applications, it is difficult to find strong associations among data items at low or primitive levels of abstraction due to the sparsity of data at those levels.
- Strong associations discovered at high levels of abstraction may represent commonsense knowledge.
- Therefore, data mining systems should provide capabilities for mining association rules at multiple levels of abstraction, with sufficient flexibility for easy traversal among different abstraction spaces.
- Association rules generated from mining data at multiple levels of abstraction are called multiple-level or multilevel association rules.
- Multilevel association rules can be mined efficiently using concept hierarchies under a support-confidence framework.
- In general, a top-down strategy is employed, where counts are accumulated for the calculation of frequent itemsets at each concept level, starting at the concept level 1 and working downward in the hierarchy toward the more specific concept levels, until no more frequent itemsets can be found.

A concept hierarchy defines a sequence of mappings from a set of low-level concepts to higher level, more general concepts. Data can be generalized by replacing low-level concepts within the data by their higher-level concepts, or ancestors, from a concept hierarchy.

| <i>TID</i> | <i>Items Purchased</i>                                            |
|------------|-------------------------------------------------------------------|
| T100       | IBM-ThinkPad-T40/2373, HP-Photosmart-7660                         |
| T200       | Microsoft-Office-Professional-2003, Microsoft-Plus!-Digital-Media |
| T300       | Logitech-MX700-Cordless-Mouse, Fellowes-Wrist-Rest                |
| T400       | Dell-Dimension-XPS, Canon-PowerShot-S400                          |
| T500       | IBM-ThinkPad-R40/P4M, Symantec-Norton-Antivirus-2003              |
| ...        | ...                                                               |



A concept hierarchy for *AllElectronics* computer items.


The concept hierarchy has five levels, respectively referred to as levels 0 to 4, starting with level 0 at the root node for all.

- Here, Level 1 includes computer, software, printer&camera, and computer accessory.
- Level 2 includes laptop computer, desktop computer, office software, antivirus software
- Level 3 includes IBM desktop computer, . . . , Microsoft office software, and so on.
- Level 4 is the most specific abstraction level of this hierarchy.

## 2.5.1 Approaches For Mining Multilevel Association Rules:

### 1. Uniform Minimum Support:

- The same minimum support threshold is used when mining at each level of abstraction.
- When a uniform minimum support threshold is used, the search procedure is simplified.
- The method is also simple in that users are required to specify only one minimum support threshold.
- The uniform support approach, however, has some difficulties. It is unlikely that items at lower levels of abstraction will occur as frequently as those at higher levels of abstraction.
- If the minimum support threshold is set too high, it could miss some meaningful associations occurring at low abstraction levels. If the threshold is set too low, it may generate many uninteresting associations occurring at high abstraction levels.




---

Multilevel mining with uniform support.

### 2. Reduced Minimum Support:

- Each level of abstraction has its own minimum support threshold.

- The deeper the level of abstraction, the smaller the corresponding threshold is.
- For example, the minimum support thresholds for levels 1 and 2 are 5% and 3%, respectively. In this way,  $\text{—computer}$ ,  $\text{—laptop computer}$ , and  $\text{—desktop computer}$  are all considered frequent.



### 3. Group-Based Minimum Support:

- Because users or experts often have insight as to which groups are more important than others, it is sometimes more desirable to set up user-specific, item, or group based minimal support thresholds when mining multilevel rules.
- For example, a user could set up the minimum support thresholds based on product price, or on items of interest, such as by setting particularly low support thresholds for laptop computers and flash drives in order to pay particular attention to the association patterns containing items in these categories.

## Mining Multidimensional Association Rules from Relational Databases and Data Warehouses:

- Single dimensional or intra dimensional association rule contains a single distinct predicate (e.g., buys) with multiple occurrences i.e., the predicate occurs more than once within the rule.  

$$\text{buys}(X, \text{—digital camera}) \Rightarrow \text{buys}(X, \text{—HP printer})$$
- Association rules that involve two or more dimensions or predicates can be referred to as multidimensional association rules.  

$$\text{age}(X, "20...29") \wedge \text{occupation}(X, "student") \Rightarrow \text{buys}(X, "laptop")$$
- Above Rule contains three predicates (age, occupation, and buys), each of which occurs only once in the rule. Hence, we say that it has no repeated predicates.

- Multidimensional association rules with no repeated predicates are called interdimensional association rules.
- We can also mine multidimensional association rules with repeated predicates, which contain multiple occurrences of some predicates. These rules are called hybrid-dimensional association rules. An example of such a rule is the following, where the predicate buys is repeated:

$age(X, [20...29]) \wedge buys(X, [laptop]) \Rightarrow buys(X, [HP\ printer])$

### Mining Quantitative Association Rules:

- Quantitative association rules are multidimensional association rules in which the numeric attributes are *dynamically* discretized during the mining process so as to satisfy some mining criteria, such as maximizing the confidence or compactness of the rules mined.
- In this section, we focus specifically on how to mine quantitative association rules having two quantitative attributes on the left-hand side of the rule and one categorical attribute on the right-hand side of the rule. That is

$A_{quan1} \wedge A_{quan2} \Rightarrow A_{cat}$

Where  $A_{quan1}$  and  $A_{quan2}$  are tests on quantitative attribute interval

$A_{cat}$  tests a categorical attribute from the task-relevant data.

- Such rules have been referred to as two-dimensional quantitative association rules, because they contain two quantitative dimensions.
- For instance, suppose you are curious about the association relationship between pairs of quantitative attributes, like customer age and income, and the type of television (such as *high-definition TV*, i.e., *HDTV*) that customers like to buy.

An example of such a 2-D quantitative association rule is

$age(X, [30...39]) \wedge income(X, [42K...48K]) \Rightarrow buys(X, [HDTV])$

### From Association Mining to Correlation Analysis:

- A correlation measure can be used to augment the support-confidence framework for association rules. This leads to *correlation rules* of the form

$A \Rightarrow B$  [support, confidence, correlation]

- That is, a correlation rule is measured not only by its support and confidence but also by the correlation between itemsets  $A$  and  $B$ . There are many different correlation measures from which to choose. In this section, we study various correlation measures to determine which would be good for mining large data sets.
- Lift is a simple correlation measure that is given as follows. The occurrence of itemset  $A$  is independent of the occurrence of itemset  $B$  if  $P(A \cup B) = P(A)P(B)$ ; otherwise, itemsets  $A$  and  $B$  are dependent and correlated as events. This definition can easily be extended to more than two itemsets.

The lift between the occurrence of  $A$  and  $B$  can be measured by computing

$$\text{lift}(A, B) = \frac{P(A \cup B)}{P(A)P(B)}.$$

- If the  $\text{lift}(A, B)$  is less than 1, then the occurrence of  $A$  is negatively correlated with the occurrence of  $B$ .
- If the resulting value is greater than 1, then  $A$  and  $B$  are positively correlated, meaning that the occurrence of one implies the occurrence of the other.
- If the resulting value is equal to 1, then  $A$  and  $B$  are independent and there is no correlation between them.

## **UNIT-IV**

### **Classification and Prediction:**

- Classification and prediction are two forms of data analysis that can be used to extract models describing important data classes or to predict future data trends.
- Classification predicts categorical (discrete, unordered) labels, *prediction* models continuous valued functions.
- For example, we can build a classification model to categorize bank loan applications as either safe or risky, or a prediction model to predict the expenditures of potential customers on computer equipment given their income and occupation.
- A predictor is constructed that predicts a continuous-valued function, or ordered value, as opposed to a categorical label.
- Regression analysis is a statistical methodology that is most often used for numeric prediction.
- Many classification and prediction methods have been proposed by researchers in machine learning, pattern recognition, and statistics.
- Most algorithms are memory resident, typically assuming a small data size. Recent data mining research has built on such work, developing scalable classification and prediction techniques capable of handling large disk-resident data.

### **Classification General Approaches:**

#### **1. Preparing the Data for Classification and Prediction:**

The following preprocessing steps may be applied to the data to help improve the accuracy, efficiency, and scalability of the classification or prediction process.

##### **(i) Data cleaning:**

- This refers to the preprocessing of data in order to remove or reduce *noise* (by applying smoothing techniques) and the treatment of *missing values* (e.g., by replacing a missing value with the most commonly occurring value for that attribute, or with the most probable
-

value based on statistics).

- Although most classification algorithms have some mechanisms for handling noisy or missing data, this step can help reduce confusion during learning.

**(ii) Relevance analysis:**

- Many of the attributes in the data may be *redundant*.
- Correlation analysis can be used to identify whether any two given attributes are statistically related.
- For example, a strong correlation between attributes A1 and A2 would suggest that one of the two could be removed from further analysis.
- A database may also contain *irrelevant* attributes. Attribute subset selection can be used in these cases to find a reduced set of attributes such that the resulting probability distribution of the data classes is as close as possible to the original distribution obtained using all attributes.
- Hence, relevance analysis, in the form of correlation analysis and attribute subset selection, can be used to detect attributes that do not contribute to the classification or prediction task.
- Such analysis can help improve classification efficiency and scalability.

**(iii) Data Transformation And Reduction**

- The data may be transformed by normalization, particularly when neural networks or methods involving distance measurements are used in the learning step.
  - Normalization involves scaling all values for a given attribute so that they fall within a small specified range, such as -1 to +1 or 0 to 1.
  - The data can also be transformed by *generalizing* it to higher-level concepts. Concept hierarchies may be used for this purpose. This is particularly useful for continuous valued attributes.
  - For example, numeric values for the attribute *income* can be generalized to discrete ranges, such as *low*, *medium*, and *high*. Similarly, categorical attributes, like *street*, can be generalized to higher-level concepts, like *city*.
  - Data can also be reduced by applying many other methods, ranging from wavelet transformation and principle components analysis to discretization techniques, such
-

as binning, histogram analysis, and clustering.

## **Comparing Classification and Prediction Methods:**

### ➤ **Accuracy:**

- The accuracy of a classifier refers to the ability of a given classifier to correctly predict the class label of new or previously unseen data (i.e., tuples without class label information).
- The accuracy of a predictor refers to how well a given predictor can guess the value of the predicted attribute for new or previously unseen data.

### ➤ **Speed:**

This refers to the computational costs involved in generating and using the given classifier or predictor.

### ➤ **Robustness:**

This is the ability of the classifier or predictor to make correct predictions given noisy data or data with missing values.

### ➤ **Scalability:**


This refers to the ability to construct the classifier or predictor efficiently given large amounts of data.

### ➤ **Interpretability:**

- This refers to the level of understanding and insight that is provided by the classifier or predictor.
- Interpretability is subjective and therefore more difficult to assess.

## **Decision Tree Algorithm:**

- Decision tree induction is the learning of decision trees from class-labeled training tuples.
- A decision tree is a flowchart-like tree structure, where
  - Each internal node denotes a test on an attribute.
  - Each branch represents an outcome of the test.
  - Each leaf node holds a class label.
  - The topmost node in a tree is the root node.



- The construction of decision tree classifiers does not require any domain knowledge or parameter setting, and therefore is appropriate for exploratory knowledge discovery.
- Decision trees can handle high dimensional data.
- Their representation of acquired knowledge in tree form is intuitive and generally easy to assimilate by humans.
- The learning and classification steps of decision tree induction are simple and fast.
- In general, decision tree classifiers have good accuracy.
- Decision tree induction algorithms have been used for classification in many applications.

**Algorithm: Generate\_decision\_tree.** Generate a decision tree from the training tuples of data partition  $D$ .

**Input:**

- Data partition,  $D$ , which is a set of training tuples and their associated class labels;
- *attribute\_list*, the set of candidate attributes;
- *Attribute\_selection\_method*, a procedure to determine the splitting criterion that “best” partitions the data tuples into individual classes. This criterion consists of a *splitting\_attribute* and, possibly, either a *split point* or *splitting subset*.

**Output:** A decision tree.

**Method:**

- (1) create a node  $N$ ;
- (2) **if** tuples in  $D$  are all of the same class,  $C$  **then**
- (3)     return  $N$  as a leaf node labeled with the class  $C$ ;
- (4) **if** *attribute\_list* is empty **then**
- (5)     return  $N$  as a leaf node labeled with the majority class in  $D$ ; // majority voting
- (6) apply *Attribute\_selection\_method*( $D$ , *attribute\_list*) to find the “best” *splitting\_criterion*;
- (7) label node  $N$  with *splitting\_criterion*;
- (8) **if** *splitting\_attribute* is discrete-valued and multiway splits allowed **then** // not restricted to binary trees
- (9)     *attribute\_list* ← *attribute\_list* – *splitting\_attribute*; // remove *splitting\_attribute*
- (10) **for each** outcome  $j$  of *splitting\_criterion*
- // partition the tuples and grow subtrees for each partition
- (11)     let  $D_j$  be the set of data tuples in  $D$  satisfying outcome  $j$ ; // a partition
- (12)     **if**  $D_j$  is empty **then**
- (13)         attach a leaf labeled with the majority class in  $D$  to node  $N$ ;
- (14)     **else** attach the node returned by *Generate\_decision\_tree*( $D_j$ , *attribute\_list*) to node  $N$ ;
- (15) **endfor**
- (16) return  $N$ ;

areas, such as medicine, manufacturing and production, financial analysis, astronomy, and molecular biology.

### **Algorithm For Decision Tree Induction:**

The algorithm is called with three parameters:

- Data partition
  - Attribute list
  - Attribute selection method
- The parameter attribute list is a list of attributes describing the tuples.
  - Attribute selection method specifies a heuristic procedure for selecting the attribute that —best discriminates the given tuples according to class.
  - The tree starts as a single node,  $N$ , representing the training tuples in  $D$ .
  - If the tuples in  $D$  are all of the same class, then node  $N$  becomes a leaf and is labeled with that class.
  - All of the terminating conditions are explained at the end of the algorithm.
  - Otherwise, the algorithm calls Attribute selection method to determine the splitting criterion.
  - The splitting criterion tells us which attribute to test at node  $N$  by determining the —best way to separate or partition the tuples in  $D$  into individual classes.

There are three possible scenarios. Let  $A$  be the splitting attribute.  $A$  has  $v$  distinct values,  $\{a_1, a_2, \dots, a_v\}$ , based on the training data.

#### **1 $A$ is discrete-valued:**

- In this case, the outcomes of the test at node  $N$  correspond directly to the known values of  $A$ .
- A branch is created for each known value,  $a_j$ , of  $A$  and labeled with that value.
- $A$  need not be considered in any future partitioning of the tuples.

#### **2 $A$ is continuous-valued:**


In this case, the test at node  $N$  has two possible outcomes, corresponding to the conditions

$A \leq \text{split point}$  and  $A > \text{split point}$ , respectively where split point is the split-point

returned by Attribute selection method as part of the splitting criterion.

### 3 A is discrete-valued and a binary tree must be produced:

The test at node N is of the form— $A \in S_A?$ .  $S_A$  is the splitting subset for A, returned by Attribute selection method as part of the splitting criterion. It is a subset of the known values of A.



(a) If A is Discrete valued (b) If A is continuous valued (c) If A is discrete-valued and a binary tree must be produced:

## Bayesian Classification:

- Bayesian classifiers are statistical classifiers.
- They can predict class membership probabilities, such as the probability that a given tuple belongs to a particular class.
- Bayesian classification is based on Bayes' theorem.

## Bayes' Theorem:

- Let X be a data tuple. In Bayesian terms, X is considered —evidence.‖and it is described by measurements made on a set of n attributes.

- Let  $H$  be some hypothesis, such as that the data tuple  $X$  belongs to a specified class  $C$ .
- For classification problems, we want to determine  $P(H|X)$ , the probability that the hypothesis  $H$  holds given the —evidence— or observed data tuple  $X$ .
- $P(H|X)$  is the posterior probability, or a posteriori probability, of  $H$  conditioned on  $X$ .
- Bayes' theorem is useful in that it provides a way of calculating the posterior probability,  $P(H|X)$ , from  $P(H)$ ,  $P(X|H)$ , and  $P(X)$ .

$$P(H|X) = \frac{P(X|H)P(H)}{P(X)}.$$

## Naïve Bayesian Classifier:

The naïve Bayesian classifier, or simple Bayesian classifier, works as follows:

1. Let  $D$  be a training set of tuples and their associated class labels. As usual, each tuple is represented by an  $n$ -dimensional attribute vector,  $X = (x_1, x_2, \dots, x_n)$ , depicting  $n$  measurements made on the tuple from  $n$  attributes, respectively,  $A_1, A_2, \dots, A_n$ .
2. Suppose that there are  $m$  classes,  $C_1, C_2, \dots, C_m$ . Given a tuple,  $X$ , the classifier will predict that  $X$  belongs to the class having the highest posterior probability, conditioned on  $X$ .

That is, the naïve Bayesian classifier predicts that tuple  $X$  belongs to the class  $C_i$  if and only if

$$P(C_i|X) > P(C_j|X) \quad \text{for } 1 \leq j \leq m, j \neq i.$$

Thus we maximize  $P(C_i|X)$ . The class  $C_i$  for which  $P(C_i|X)$  is maximized is called the maximum posteriori hypothesis. By Bayes' theorem

$$P(C_i|X) = \frac{P(X|C_i)P(C_i)}{P(X)}.$$

3. As  $P(X)$  is constant for all classes, only  $P(X|C_i)P(C_i)$  need be maximized. If the class prior probabilities are not known, then it is commonly assumed that the classes are equally likely, that is,  $P(C_1) = P(C_2) = \dots = P(C_m)$ , and we would therefore maximize  $P(X|C_i)$ . Otherwise, we maximize  $P(X|C_i)P(C_i)$ .

4. Given data sets with many attributes, it would be extremely computationally expensive to compute  $P(X|C_i)$ . In order to reduce computation in evaluating  $P(X|C_i)$ , the naive assumption of class conditional independence is made. This presumes that the values of the attributes are conditionally independent of one another, given the class label of the tuple. Thus,

$$\begin{aligned} P(X|C_i) &= \prod_{k=1}^n P(x_k|C_i) \\ &= P(x_1|C_i) \times P(x_2|C_i) \times \cdots \times P(x_n|C_i). \end{aligned}$$

We can easily estimate the probabilities  $P(x_1|C_i)$ ,  $P(x_2|C_i)$ ,  $\dots$ ,  $P(x_n|C_i)$  from the training tuples. For each attribute, we look at whether the attribute is categorical or continuous-valued. For instance, to compute  $P(X|C_i)$ , we consider the following:

- If  $A_k$  is categorical, then  $P(x_k|C_i)$  is the number of tuples of class  $C_i$  in  $D$  having the value  $x_k$  for  $A_k$ , divided by  $|C_{i,D}|$  the number of tuples of class  $C_i$  in  $D$ .
- If  $A_k$  is continuous-valued, then we need to do a bit more work, but the calculation is pretty straightforward.

A continuous-valued attribute is typically assumed to have a Gaussian distribution with a mean  $\mu$  and standard deviation  $\sigma$ , defined by

$$g(x, \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}},$$

$$P(x_k|C_i) = g(x_k, \mu_{C_i}, \sigma_{C_i}).$$

5. In order to predict the class label of  $X$ ,  $P(X|C_i)P(C_i)$  is evaluated for each class  $C_i$ .


The classifier predicts that the class label of tuple  $X$  is the class  $C_i$  if and only if

$$P(X|C_i)P(C_i) > P(X|C_j)P(C_j) \quad \text{for } 1 \leq j \leq m, j \neq i.$$

## A Multilayer Feed-Forward Neural Network:

- The back propagation algorithm performs learning on a multilayer feed-forward neural network.
- It iteratively learns a set of weights for prediction of the class label of tuples. A multilayer feed-forward neural network consists of an input layer, one or more hidden layers, and an output layer.

Example:



- The inputs to the network correspond to the attributes measured for each training tuple. The inputs are fed simultaneously into the units making up the input layer. These inputs pass through the input layer and are then weighted and fed simultaneously to a second layer known as a hidden layer.
- The outputs of the hidden layer units can be input to another hidden layer, and so on. The number of hidden layers is arbitrary.
- The weighted outputs of the last hidden layer are input to units making up the output layer, which emits the network's prediction for given tuples

### Classification by Backpropagation:

- Back propagation is a neural network learning algorithm.
- A neural network is a set of connected input/output units in which each connection has a weight associated with it.
- During the learning phase, the network learns by adjusting the weights so as to be able to predict the correct class label of the input tuples.
- Neural network learning is also referred to as connectionist learning due to the connections between units.

- Neural networks involve long training times and are therefore more suitable for applications where this is feasible.
- Back propagation learns by iteratively processing a data set of training tuples, comparing the network's prediction for each tuple with the actual known target value.
- The target value may be the known class label of the training tuple (for classification problems) or a continuous value (for prediction).
- For each training tuple, the weights are modified so as to minimize the mean squared error between the network's prediction and the actual target value. These modifications are made in the —backwards direction, that is, from the output layer, through each hidden layer down to the first hidden layer hence the name is back propagation.
- Although it is not guaranteed, in general the weights will eventually converge, and the learning process stops.

#### **Advantages:**

- It includes their high tolerance of noisy data as well as their ability to classify patterns on which they have not been trained.
- They can be used when you may have little knowledge of the relationships between attributes and classes.
- They are well-suited for continuous-valued inputs and outputs, unlike most decision tree algorithms.
- They have been successful on a wide array of real-world data, including handwritten character recognition, pathology and laboratory medicine, and training a computer to pronounce English text.
- Neural network algorithms are inherently parallel; parallelization techniques can be used to speed up the computation process.

#### **Process:**

##### **Initialize the weights:**

The weights in the network are initialized to small random numbers ranging from -1.0 to 1.0, or -0.5 to 1.5. Each unit has a *bias* associated with it. The biases are similarly initialized to small random numbers.

Each training tuple,  $X$ , is processed by the following steps.

### Propagate the inputs forward:

First, the training tuple is fed to the input layer of the network. The inputs pass through the input units, unchanged. That is, for an input unit  $j$ , its output,  $O_j$ , is equal to its input value,  $I_j$ . Next, the net input and output of each unit in the hidden and output layers are computed. The net input to a unit in the hidden or output layers is computed as a linear combination of its inputs.

Each such unit has a number of inputs to it that are, in fact, the outputs of the units connected to it in the previous layer. Each connection has a weight. To compute the net input to the unit, each input connected to the unit is multiplied by its corresponding weight, and this is summed.


$$I_j = \sum_i w_{ij} O_i + \theta_j,$$

Where  $w_{i,j}$  is the weight of the connection from unit  $i$  in the previous layer to unit

$j$ ;  $O_i$  is the output of unit  $i$  from the previous layer

$\theta_j$  is the bias of the unit & it acts as a threshold in that it serves to vary the activity of the unit.

Each unit in the hidden and output layers takes its net input and then applies an activation function to it.



## Backpropagate the error:

The error is propagated backward by updating the weights and biases to reflect the error of the network's prediction. For a unit  $j$  in the output layer, the error  $Err_j$  is computed by

$$Err_j = O_j(1 - O_j)(T_j - O_j)$$

Where  $O_j$  is the actual output of unit  $j$ , and  $T_j$  is the known target value of the given training tuple.

The error of a hidden layer unit  $j$  is

$$Err_j = O_j(1 - O_j) \sum_k Err_k w_{jk}$$

Where  $w_{jk}$  is the weight of the connection from unit  $j$  to a unit  $k$  in the next higher layer, and  $Err_k$  is the error of unit  $k$ .

Weights are updated by the following equations, where  $\Delta w_{ij}$  is the change in weight  $w_{ij}$ :

$$\Delta w_{ij} = (l)Err_j O_i$$

$$w_{ij} = w_{ij} + \Delta w_{ij}$$

Biases are updated by the following equations below

$$\Delta \theta_j = (l)Err_j$$

$$\theta_j = \theta_j + \Delta \theta_j$$

### Algorithm: k-Nearest-Neighbor Classifier:

#### Input:

- $D$ , a data set consisting of the training tuples and their associated target values;
- $l$ , the learning rate;
- $network$ , a multilayer feed-forward network.

**Output:** A trained neural network.

#### Method:

```
(1) Initialize all weights and biases in network;  
(2) while terminating condition is not satisfied {  
(3)   for each training tuple  $X$  in  $D$  {  
(4)     // Propagate the inputs forward:  
(5)     for each input layer unit  $j$  {  
(6)        $O_j = I_j$ ; // output of an input unit is its actual input value  
(7)     for each hidden or output layer unit  $j$  {  
(8)        $I_j = \sum_i w_{ij} O_i + \theta_j$ ; // compute the net input of unit  $j$  with respect to the  
        previous layer,  $i$   
(9)        $O_j = \frac{1}{1 + e^{-I_j}}$ ; // compute the output of each unit  $j$   
(10)    // Backpropagate the errors:  
(11)    for each unit  $j$  in the output layer  
(12)       $Err_j = O_j(1 - O_j)(T_j - O_j)$ ; // compute the error  
(13)    for each unit  $j$  in the hidden layers, from the last to the first hidden layer  
(14)       $Err_j = O_j(1 - O_j) \sum_k Err_k w_{jk}$ ; // compute the error with respect to the  
        next higher layer,  $k$   
(15)    for each weight  $w_{ij}$  in network {  
(16)       $\Delta w_{ij} = (l)Err_j O_i$ ; // weight increment  
(17)       $w_{ij} = w_{ij} + \Delta w_{ij}$ ; // weight update  
(18)    for each bias  $\theta_j$  in network {  
(19)       $\Delta \theta_j = (l)Err_j$ ; // bias increment  
(20)       $\theta_j = \theta_j + \Delta \theta_j$ ; // bias update  
(21)    } }
```

- Nearest-neighbor classifiers are based on learning by analogy, that is, by comparing a given test tuple with training tuples that are similar to it.
- The training tuples are described by  $n$  attributes. Each tuple represents a point in an  $n$ -dimensional space. In this way, all of the training tuples are stored in an  $n$ -dimensional pattern space. When given an unknown tuple, a  $k$ -nearest-neighbor classifier searches the pattern space for the  $k$  training tuples that are closest to the unknown tuple. These  $k$  training tuples are the  $k$  nearest neighbors of the unknown tuple.
- Closeness is defined in terms of a distance metric, such as Euclidean distance.
- The Euclidean distance between two points or tuples, say,  $X_1 = (x_{11}, x_{12}, \dots, x_{1n})$  and  $X_2 = (x_{21}, x_{22}, \dots, x_{2n})$ , is

$$dist(X_1, X_2) = \sqrt{\sum_{i=1}^n (x_{1i} - x_{2i})^2}.$$

In other words, for each numeric attribute, we take the difference between the corresponding values of that attribute in tuple  $X_1$  and in tuple  $X_2$ , square this difference, and accumulate it. The square root is taken of the total accumulated distance count.

Min-Max normalization can be used to transform a value  $v$  of a numeric attribute  $A$  to  $v_0$  in the range  $[0, 1]$  by computing

$$v' = \frac{v - \min_A}{\max_A - \min_A},$$

Where  $\min_A$  and  $\max_A$  are the minimum and maximum values of attribute  $A$

- For  $k$ -nearest-neighbor classification, the unknown tuple is assigned the most common class among its  $k$  nearest neighbors.
- When  $k = 1$ , the unknown tuple is assigned the class of the training tuple that is closest to it in pattern space.
- Nearest neighbor classifiers can also be used for prediction, that is, to return a real-valued prediction for a given unknown tuple.
- In this case, the classifier returns the average value of the real-valued labels associated with the  $k$  nearest neighbors of the unknown tuple.

## ***k*-Nearest-Neighbor Classifiers**

The *k*-nearest-neighbor method was first described in the early 1950s. The method is labor intensive when given large training sets, and did not gain popularity until the 1960s when increased computing power became available. It has since been widely used in the area of pattern recognition.

Nearest-neighbor classifiers are based on learning by analogy, that is, by comparing a given test tuple with training tuples that are similar to it. The training tuples are described by *n* attributes. Each tuple represents a point in an *n*-dimensional space. In this way, all of the training tuples are stored in an *n*-dimensional pattern space. When given an unknown tuple, a *k*-nearest-neighbor classifier searches the pattern space for the *k* training tuples that are closest to the unknown tuple. These *k* training tuples are the *k* “nearest neighbors” of the unknown tuple.

“Closeness” is defined in terms of a distance metric, such as Euclidean distance.

The Euclidean distance between two points or tuples, say,  $X_1 = (x_{11}, x_{12}, \dots, x_{1n})$  and  $X_2 = (x_{21}, x_{22}, \dots, x_{2n})$ , is

$$dist(X_1, X_2) = \sqrt{\sum_{i=1}^n (x_{1i} - x_{2i})^2}.$$

In other words, for each numeric attribute, we take the difference between the corresponding values of that attribute in tuple  $X_1$  and in tuple  $X_2$ , square this difference, and accumulate it. The square root is taken of the total accumulated distance count. Typically, we normalize the values of each attribute before using Equation This helps prevent attributes with initially large ranges (such as income) from outweighing attributes with initially smaller ranges (such as binary attributes). Min-max normalization, for example, can be used to transform a value *v* of a numeric attribute *A* to *v'* in the range [0, 1] by computing

$$v' = \frac{v - \min_A}{\max_A - \min_A},$$

where  $\min_A$  and  $\max_A$  are the minimum and maximum values of attribute *A*. For *k*-nearest-neighbor classification, the unknown tuple is assigned the most common class among its *k* nearest neighbors. When *k* = 1, the unknown

tuple is assigned the class of the training tuple that is closest to it in pattern space. Nearestneighbor classifiers can also be used for prediction, that is, to return a real-valued prediction for a given unknown tuple. In this case, the classifier returns the average value of the real-valued labels associated with the  $k$  nearest neighbors of the unknown tuple. “But how can distance be computed for attributes that not numeric, but categorical, such as color?” The above discussion assumes that the attributes used to describe the tuples are all numeric.

For categorical attributes, a simple method is to compare the corresponding value of the attribute in tuple  $X_1$  with that in tuple  $X_2$ . If the two are identical (e.g., tuples  $X_1$  and  $X_2$  both have the color blue), then the difference between the two is taken as 0.

If the two are different (e.g., tuple  $X_1$  is blue but tuple  $X_2$  is red), then the difference is considered to be 1. Other methods may incorporate more sophisticated schemes for differential grading (e.g., where a larger difference score is assigned, say, for blue and white than for blue and black). “What about missing values?” In general, if the value of a given attribute  $A$  is missing in tuple  $X_1$  and/or in tuple  $X_2$ , we assume the maximum possible difference. Suppose that each of the attributes have been mapped to the range  $[0, 1]$ . For categorical attributes, we take the difference value to be 1 if either one or both of the corresponding values of  $A$  are missing. If  $A$  is numeric and missing from both tuples  $X_1$  and  $X_2$ , then the difference is also taken to be 1. If only one value is missing and the other (which we’ll call  $v_0$ ) is present and normalized, then we can take the difference to be either  $|1 - v_0|$  or  $|0 - v_0|$  (i.e.,  $1 - v_0$  or  $v_0$ ), whichever is greater. “How can I determine a good value for  $k$ , the number of neighbors?” This can be determined experimentally. Starting with  $k = 1$ , we use a test set to estimate the error rate of the classifier.

This process can be repeated each time by incrementing  $k$  to allow for one more neighbor. The  $k$  value that gives the minimum error rate may be selected. In general, the larger the number of training tuples is, the larger the value of  $k$  will be (so that classification and prediction decisions can be based on a larger portion of the stored tuples). As the number of training tuples approaches infinity and  $k = 1$ , the error rate can be no worse than twice the Bayes error rate (the latter being the theoretical minimum). If  $k$  also approaches infinity, the error rate approaches the Bayes error rate.

Nearest-neighbor classifiers use distance-based comparisons that intrinsically assign equal weight to each attribute. They therefore can suffer from poor accuracy when given noisy or irrelevant attributes. The method, however, has been modified to incorporate attribute weighting and the pruning of noisy data tuples. The choice of a distance

metric can be critical. The Manhattan (city block) distance ,or other distance measurements, may also be used.

### **Prediction :**

What if we would like to predict a continuous value, rather than a categorical label?” Numeric prediction is the task of predicting continuous (or ordered) values for given input. For example, we may wish to predict the salary of college graduates with 10 years of work experience, or the potential sales of a new product given its price. By far, the most widely used approach for numeric prediction (hereafter referred to as prediction) is regression, a statistical methodology that was developed by Sir Frances Galton (1822– 1911), a mathematician who was also a cousin of Charles Darwin. In fact, many texts use the terms “regression” and “numeric prediction” synonymously. However, as we have seen, some classification techniques (such as back propagation, support vector machines, and k-nearest-neighbor classifiers) can be adapted for prediction. In this section, we discuss the use of regression techniques for prediction.

Regression analysis can be used to model the relationship between one or more independent or predictor variables and a dependent or response variable (which is continuous-valued). In the context of data mining, the predictor variables are the attributes of interest describing the tuple (i.e., making up the attribute vector). In general, the values of the predictor variables are known. (Techniques exist for handling cases where such values may be missing.) The response variable is what we want to predict—it is what we referred to in Section 6.1 as the predicted attribute. Given a tuple described by predictor variables, we want to predict the associated value of the response variable.

Regression analysis is a good choice when all of the predictor variables are continuous valued as well. Many problems can be solved by linear regression, and even more can be tackled by applying transformations to the variables so that a nonlinear problem can be converted to a linear one. For reasons of space, we cannot give a fully detailed treatment of regression. Instead, this section provides an intuitive introduction to the topic.

Section 6.11.1 discusses straight-line regression analysis (which involves a single predictor variable) and multiple linear regression analysis (which involves two or more predictor variables). Section 6.11.2 provides some pointers on dealing with nonlinear regression. Section 6.11.3 mentions other regression-based methods, such as generalized linear models, Poisson regression, log-linear models, and regression trees.

---

## Linear Regression

Straight-line regression analysis involves a response variable,  $y$ , and a single predictor variable,  $x$ . It is the simplest form of regression, and models  $y$  as a linear function of  $x$ . That is,

$$y = b + wx;$$

where the variance of  $y$  is assumed to be constant, and  $b$  and  $w$  are regression coefficients specifying the  $Y$ -intercept and slope of the line, respectively. The regression coefficients,  $w$  and  $b$ , can also be thought of as weights, so that we can equivalently write

$$y = w_0 + w_1x;$$

These coefficients can be solved for by the method of least squares, which estimates the best-fitting straight line as the one that minimizes the error between the actual data and the estimate of the line. Let  $D$  be a training set consisting of values of predictor variable,  $x$ , for some population and their associated values for response variable,  $y$ . The training set contains  $|D|$  data points of the form  $(x_1, y_1), (x_2, y_2), \dots, (x_{|D|}, y_{|D|})$ .<sup>12</sup> The regression coefficients can be estimated using this method with the following equations:

$$w_1 = \frac{\sum_{i=1}^{|D|} (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^{|D|} (x_i - \bar{x})^2}$$


<sup>12</sup>Note that earlier, we had used the notation  $(X_i, y_i)$  to refer to training tuple  $i$  having associated class label  $y_i$ , where  $X_i$  was an attribute (or feature) vector (that is,  $X_i$  was described by more than one attribute). Here, however, we are dealing with just one predictor variable. Since the  $X_i$  here are one-dimensional, we use the notation  $x_i$  over  $X_i$  in this case.

$$w_0 = \bar{y} - w_1\bar{x}$$

where  $\bar{x}$  is the mean value of  $x_1, x_2, \dots, x_{|D|}$ , and  $\bar{y}$  is the mean value of  $y_1, y_2, \dots, y_{|D|}$ . The coefficients  $w_0$  and  $w_1$  often provide good approximations to otherwise complicated regression equations.

Salary data.

| <i>x</i> years experience | <i>y</i> salary (in \$1000s) |
|---------------------------|------------------------------|
| 3                         | 30                           |
| 8                         | 57                           |
| 9                         | 64                           |
| 13                        | 72                           |
| 3                         | 36                           |
| 6                         | 43                           |
| 11                        | 59                           |
| 21                        | 90                           |
| 1                         | 20                           |
| 16                        | 83                           |



## Accuracy and Error Measures

Now that you may have trained a classifier or predictor, there may be many questions going through your mind. For example, suppose you used data from previous sales to train a classifier to predict customer purchasing behavior. You would like an estimate of how accurately the classifier can predict the purchasing behavior of future customers, that is, future customer data on which the classifier has not been trained. You may even have tried different methods to build more than one classifier (or predictor) and now wish to compare their accuracy. But what is accuracy? How can we estimate it? Are there strategies for increasing the accuracy of a learned model? These questions are addressed in the next few sections. describes measures for computing classifier accuracy. Predictor error measures are given can use these measures in techniques for accuracy estimation, such as the holdout, random subsampling, k-fold cross-validation, and bootstrap methods .

| Classes                    | <i>buys_computer = yes</i> | <i>buys_computer = no</i> | Total  | Recognition (%) |
|----------------------------|----------------------------|---------------------------|--------|-----------------|
| <i>buys_computer = yes</i> | 6,954                      | 46                        | 7,000  | 99.34           |
| <i>buys_computer = no</i>  | 412                        | 2,588                     | 3,000  | 86.27           |
| Total                      | 7,366                      | 2,634                     | 10,000 | 95.52           |

A confusion matrix for the classes *buys\_computer = yes* and *buys\_computer = no*, where an entry is row *i* and column *j* shows the number of tuples of class *i* that were labeled by the classifier as class *j*. Ideally, the nondiagonal entries should be zero or close to zero.

## Classifier Accuracy Measures

Using training data to derive a classifier or predictor and then to estimate the accuracy of the resulting learned model can result in misleading overoptimistic estimates due to overspecialization of the learning algorithm to the data. (We'll say more on this in a moment!) Instead, accuracy is better measured on a test set consisting of class-labeled tuples that were not used to train the model. The accuracy of a classifier on a given test set is the percentage of test set tuples that are correctly classified by the classifier. In the pattern recognition literature, this is also referred to as the overall recognition rate of the classifier, that is, it reflects how well the classifier recognizes tuples of the various classes. We can also speak of the error rate or misclassification rate of a classifier,  $M$ , which is simply  $1 - \text{Acc}(M)$ , where  $\text{Acc}(M)$  is the accuracy of  $M$ . If we were to use the training set to estimate the error rate of a model, this quantity is known as the resubstitution error. This error estimate is optimistic of the true error rate (and similarly, the corresponding accuracy estimate is optimistic) because the model is not tested on any samples that it has not already seen.

## Predictor Error Measures

"How can we measure predictor accuracy?" Let  $DT$  be a test set of the form  $(X_1, y_1), (X_2, y_2), \dots, (X_d, y_d)$ , where the  $X_i$  are the  $n$ -dimensional test tuples with associated known values,  $y_i$ , for a response variable,  $y$ , and  $d$  is the number of tuples in  $DT$ . Since predictors return a continuous value rather than a categorical label, it is difficult to say exactly whether the predicted value,  $y_0 i$ , for  $X_i$  is correct. Instead of focusing on whether  $y_0 i$  is an "exact" match with  $y_i$ , we instead look at how far off the predicted value is from the actual known value. Loss functions measure the error between  $y_i$  and the predicted value,  $y_0 i$ .

## Evaluating the Accuracy of a Classifier or Predictor

---

How can we use the above measures to obtain a reliable estimate of classifier accuracy (or predictor accuracy in terms of error)? Holdout, random subsampling, cross validation, and the bootstrap are common techniques for assessing accuracy based on

### Ensemble Methods—Increasing the Accuracy

Bagging and boosting are two techniques. They are examples of ensemble methods, or methods that use a combination of models. Each combines a series of  $k$  learned models (classifiers or predictors),  $M_1, M_2, \dots, M_k$ , with the aim of creating an improved composite model. Both bagging and boosting can be used for classification as well as prediction.

#### Bagging

We first take an intuitive look at how bagging works as a method of increasing accuracy. For ease of explanation, we will assume at first that our model is a classifier. Suppose that you are a patient and would like to have a diagnosis made based on your symptoms. Instead of asking one doctor, you may choose to ask several. If a certain diagnosis occurs more than any of the others, you may choose this as the final or best diagnosis. That is, the final diagnosis is made based on a majority vote, where each doctor gets an equal vote. Now replace each doctor by a classifier, and you have the basic idea behind bagging. Intuitively, a majority vote made by a large group of doctors may be more reliable than a majority vote made by a small group.

#### Boosting

We now look at the ensemble method of boosting. As in the previous section, suppose that as a patient, you have certain symptoms. Instead of consulting one doctor, you choose to consult several. Suppose you assign weights to the value or worth of each doctor's diagnosis, based on the accuracies of previous diagnoses they have made. The final diagnosis is then a combination of the weighted diagnoses. This is the essence behind boosting.

In boosting, weights are assigned to each training tuple. A series of  $k$  classifiers is iteratively learned. After a classifier  $M_i$  is learned, the weights are updated to allow the subsequent classifier,  $M_{i+1}$ , to “pay more attention” to the training tuples that were misclassified by  $M_i$ . The final boosted classifier,  $M$ —————, combines the votes of each individual classifier, where the weight of each classifier's vote is a function of its

---

accuracy. The boosting algorithm can be extended for the prediction of continuous values.

Adaboost is a popular boosting algorithm. Suppose we would like to boost the accuracy of some learning method. We are given  $D$ , a data set of  $d$  class-labeled tuples,  $(X_1, y_1), (X_2, y_2), \dots, (X_d, y_d)$ , where  $y_i$  is the class label of tuple  $X_i$ . Initially, Adaboost assigns each training tuple an equal weight of  $1/d$ . Generating  $k$  classifiers for the ensemble requires  $k$  rounds through the rest of the algorithm. In round  $i$ , the tuples from  $D$  are sampled to form a training set,  $D_i$ , of size  $d$ . Sampling with replacement is used—the same tuple may be selected more than once. Each tuple's chance of being selected is based on its weight. A classifier model,  $M_i$ , is derived from the training tuples of  $D_i$ . Its error is then calculated using  $D_i$  as a test set. The weights of the training tuples are then adjusted according to how they were classified. If a tuple was incorrectly classified, its weight is increased. If a tuple was correctly classified, its weight is decreased. A tuple's weight reflects how hard it is to classify—the higher the weight, the more often it has been misclassified. These weights will be used to generate the training samples for the classifier of the next round. The basic idea is that when we build a classifier, we want it to focus more on the misclassified tuples of the previous round.

## Clustering Overview:

- The process of grouping a set of physical or abstract objects into classes of similar objects is called clustering.
- A cluster is a collection of data objects that are similar to one another within the same cluster and are dissimilar to the objects in other clusters.
- A cluster of data objects can be treated collectively as one group and so may be considered as a form of data compression.
- Cluster analysis tools based on k-means, k-medoids, and several methods have also been built into many statistical analysis software packages or systems, such as S-Plus, SPSS, and SAS.

## Applications:

- Cluster analysis has been widely used in numerous applications, including market research, pattern recognition, data analysis, and image processing.
  - In business, clustering can help marketers discover distinct groups in their customer bases and characterize customer groups based on purchasing patterns.
  - In biology, it can be used to derive plant and animal taxonomies, categorize genes with similar functionality, and gain insight into structures inherent in populations.
  - Clustering may also help in the identification of areas of similar land use in an earth observation database and in the identification of groups of houses in a city according to house type, value, and geographic location, as well as the identification of groups of automobile insurance policy holders with a high average claim cost.
  - Clustering is also called data segmentation in some applications because clustering partitions large data sets into groups according to their *similarity*.
-

- Clustering can also be used for outlier detection, Applications of outlier detection include the detection of credit card fraud and the monitoring of criminal activities in electronic commerce.

## **Typical Requirements Of Clustering In Data Mining:**

### **➤ Scalability:**

Many clustering algorithms work well on small data sets containing fewer than several hundred data objects; however, a large database may contain millions of objects. Clustering on a sample of a given large data set may lead to biased results.

Highly scalable clustering algorithms are needed.

### **➤ Ability to deal with different types of attributes:**

Many algorithms are designed to cluster interval-based (numerical) data. However, applications may require clustering other types of data, such as binary, categorical (nominal), and ordinal data, or mixtures of these data types.

### **➤ Discovery of clusters with arbitrary shape:**

Many clustering algorithms determine clusters based on Euclidean or Manhattan distance measures. Algorithms based on such distance measures tend to find spherical clusters with similar size and density.

However, a cluster could be of any shape. It is important to develop algorithms that can detect clusters of arbitrary shape.

### **➤ Minimal requirements for domain knowledge to determine input parameters:**

Many clustering algorithms require users to input certain parameters in cluster analysis (such as the number of desired clusters). The clustering results can be quite sensitive to input parameters. Parameters are often difficult to determine, especially for data sets containing high-dimensional objects. This not only burdens users, but it also makes the quality of clustering difficult to control.

### **➤ Ability to deal with noisy data:**

Most real-world databases contain outliers or missing, unknown, or erroneous data. Some clustering algorithms are sensitive to such data and may lead to clusters of poor quality.

➤ **Incremental clustering and insensitivity to the order of input records:**

Some clustering algorithms cannot incorporate newly inserted data (i.e., database updates) into existing clustering structures and, instead, must determine a new clustering from scratch. Some clustering algorithms are sensitive to the order of input data.

That is, given a set of data objects, such an algorithm may return dramatically different clusterings depending on the order of presentation of the input objects.

It is important to develop incremental clustering algorithms and algorithms that are insensitive to the order of input.

➤ **High dimensionality:**

A database or a data warehouse can contain several dimensions or attributes. Many clustering algorithms are good at handling low-dimensional data, involving only two to three dimensions. Human eyes are good at judging the quality of clustering for up to three dimensions. Finding clusters of data objects in high dimensional space is challenging, especially considering that such data can be sparse and highly skewed.

➤ **Constraint-based clustering:**

Real-world applications may need to perform clustering under various kinds of constraints. Suppose that your job is to choose the locations for a given number of new automatic banking machines (ATMs) in a city. To decide upon this, you may cluster households while considering constraints such as the city's rivers and highway networks, and the type and number of customers per cluster. A challenging task is to find groups of data with good clustering behavior that satisfy specified constraints.

➤ **Interpretability and usability:**

Users expect clustering results to be interpretable, comprehensible, and usable. That is, clustering may need to be tied to specific semantic interpretations and applications. It is important to study how an application goal may influence the selection of clustering features and methods.

## **A Categorization of Major Clustering Methods:**

- Partitioning Methods
- Hierarchical Methods
- Density-Based Methods
- Grid-Based Methods
- Model-Based Methods

### **Partitioning Methods:**

A partitioning method constructs  $k$  partitions of the data, where each partition represents a cluster and  $k \leq n$ . That is, it classifies the data into  $k$  groups, which together satisfy the following requirements:

- Each group must contain at least one object, and
- Each object must belong to exactly one group.

A partitioning method creates an initial partitioning. It then uses an iterative relocation technique that attempts to improve the partitioning by moving objects from one group to another.

The general criterion of a good partitioning is that objects in the same cluster are close or related to each other, whereas objects of different clusters are far apart or very different.

### **Hierarchical Methods:**

A hierarchical method creates a hierarchical decomposition of the given set of data objects. A hierarchical method can be classified as being either agglomerative or divisive, based on how the hierarchical decomposition is formed.

- ❖ The agglomerative approach, also called the bottom-up approach, starts with each object forming a separate group. It successively merges the objects or groups that are close to one another, until all of the groups are merged into one or until a termination condition holds.
  - ❖ The divisive approach, also called the top-down approach, starts with all of the objects in the same cluster. In each successive iteration, a cluster is split up into smaller clusters,
-

until eventually each object is in one cluster, or until a termination condition holds.

Hierarchical methods suffer from the fact that once a step (merge or split) is done, it can never be undone. This rigidity is useful in that it leads to smaller computation costs by not having to worry about a combinatorial number of different choices.

There are two approaches to improving the quality of hierarchical clustering:

- ❖ Perform careful analysis of object —linkages‖ at each hierarchical partitioning, such as in Chameleon, or
- ❖ Integrate hierarchical agglomeration and other approaches by first using a hierarchical agglomerative algorithm to group objects into micro clusters, and then performing macro clustering on the microclusters using another clustering method such as iterative relocation.

### **Density-based methods:**

- ❖ Most partitioning methods cluster objects based on the distance between objects. Such methods can find only spherical-shaped clusters and encounter difficulty at discovering clusters of arbitrary shapes.
- ❖ Other clustering methods have been developed based on the notion of density. Their general idea is to continue growing the given cluster as long as the density in the neighborhood exceeds some threshold; that is, for each data point within a given cluster, the neighborhood of a given radius has to contain at least a minimum number of points. Such a method can be used to filter out noise (outliers) and discover clusters of arbitrary shape.
- ❖ DBSCAN and its extension, OPTICS, are typical density-based methods that grow clusters according to a density-based connectivity analysis. DENCLUE is a method that clusters objects based on the analysis of the value distributions of density functions.

### **Grid-Based Methods:**

- ❖ Grid-based methods quantize the object space into a finite number of cells that form a grid structure.
-

- ❖ All of the clustering operations are performed on the grid structure i.e., on the quantized space. The main advantage of this approach is its fast processing time, which is typically independent of the number of data objects and dependent only on the number of cells in each dimension in the quantized space.
- ❖ STING is a typical example of a grid-based method. Wave Cluster applies wavelet transformation for clustering analysis and is both grid-based and density-based.

### **Model-Based Methods:**

- ❖ Model-based methods hypothesize a model for each of the clusters and find the best fit of the data to the given model.
- ❖ A model-based algorithm may locate clusters by constructing a density function that reflects the spatial distribution of the data points.
- ❖ It also leads to a way of automatically determining the number of clusters based on standard statistics, taking —noise or outliers into account and thus yielding robust clustering methods.

### **Tasks in Data Mining:**

- Clustering High-Dimensional Data
- Constraint-Based Clustering

#### **Clustering High-Dimensional Data:**

- It is a particularly important task in cluster analysis because many applications require the analysis of objects containing a large number of features or dimensions.
- For example, text documents may contain thousands of terms or keywords as features, and DNA micro array data may provide information on the expression levels of thousands of genes under hundreds of conditions.
- Clustering high-dimensional data is challenging due to the curse of dimensionality.
- Many dimensions may not be relevant. As the number of dimensions increases, the data become increasingly sparse so that the distance measurement between pairs of points become meaningless and the average density of points anywhere in the data is likely to be low. Therefore, a different clustering methodology needs to be

developed for high-dimensional data.

- CLIQUE and PROCLUS are two influential subspace clustering methods, which search for clusters in subspaces of the data, rather than over the entire data space.
- Frequent pattern-based clustering, another clustering methodology, extracts distinct frequent patterns among subsets of dimensions that occur frequently. It uses such patterns to group objects and generate meaningful clusters.

### **Constraint-Based Clustering:**

- It is a clustering approach that performs clustering by incorporation of user-specified or application-oriented constraints.
- A constraint expresses a user's expectation or describes properties of the desired clustering results, and provides an effective means for communicating with the clustering process.
- Various kinds of constraints can be specified, either by a user or as per application requirements.
- Spatial clustering employs with the existence of obstacles and clustering under user-specified constraints. In addition, semi-supervised clustering employs for pairwise constraints in order to improve the quality of the resulting clustering.

### **Classical Partitioning Methods:**

The most well-known and commonly used partitioning methods are

- ❖ The  $k$ -Means Method
- ❖  $k$ -Medoids Method

#### **Partitioning Clustering: The $K$ -Means Method:**

The  $k$ -means algorithm takes the input parameter,  $k$ , and partitions a set of  $n$  objects into  $k$  clusters so that the resulting intracluster similarity is high but the intercluster similarity is low. Cluster similarity is measured in regard to the mean value of the objects in a cluster, which can be viewed as the cluster's centroid or center of gravity.

The  $k$ -means algorithm proceeds as follows.

---

- First, it randomly selects  $k$  of the objects, each of which initially represents a cluster mean or center.
- For each of the remaining objects, an object is assigned to the cluster to which it is the most similar, based on the distance between the object and the cluster mean.
- It then computes the new mean for each cluster.
- This process iterates until the criterion function converges.

Typically, the square-error criterion is used, defined as

$$E = \sum_{i=1}^k \sum_{p \in C_i} |p - m_i|^2,$$

Where  $E$  is the sum of the square error for all objects in the data set

$p$  is the point in space representing a given object

$M_i$  is the mean of cluster  $C_i$

### **The k-means partitioning algorithm:**

The  $k$ -means algorithm for partitioning, where each cluster's center is represented by the mean value of the objects in the cluster.


#### **Input:**

- $k$ : the number of clusters,
- $D$ : a data set containing  $n$  objects.

**Output:** A set of  $k$  clusters.

#### **Method:**

- (1) arbitrarily choose  $k$  objects from  $D$  as the initial cluster centers;
- (2) repeat
- (3) (re)assign each object to the cluster to which the object is the most similar,  
based on the mean value of the objects in the cluster;
- (4) update the cluster means, i.e., calculate the mean value of the objects for  
each cluster;
- (5) until no change;



Clustering of a set of objects based on the  $k$ -means method

### The $k$ -Medoids Method:

- The  $k$ -means algorithm is sensitive to outliers because an object with an extremely large value may substantially distort the distribution of data. This effect is particularly exacerbated due to the use of the square-error function.
- Instead of taking the mean value of the objects in a cluster as a reference point, we can pick actual objects to represent the clusters, using one representative object per cluster. Each remaining object is clustered with the representative object to which it is the most similar.
- The partitioning method is then performed based on the principle of minimizing the sum of the dissimilarities between each object and its corresponding reference point. That is, an absolute-error criterion is used, defined as

$$E = \sum_{j=1}^k \sum_{p \in C_j} |p - o_j|,$$

where  $E$  is the sum of the absolute error for all objects in the data set

$p$  is the point in space representing a given object in cluster  $C_j$

$o_j$  is the representative object of  $C_j$

- The initial representative objects are chosen arbitrarily. The iterative process of replacing representative objects by non representative objects continues as long as the quality of the resulting clustering is improved.
- This quality is estimated using a cost function that measures the average dissimilarity between an object and the representative object of its cluster.
- To determine whether a non representative object,  $o_j$  random, is a good replacement for a current representative object,  $o_j$ , the following four cases are examined for each of the non representative objects.

**Case 1:**

'p' currently belongs to representative object,  $o_j$ . If  $o_j$  is replaced by  $o_{\text{random}}$  as a representative object and p is closest to one of the other representative objects,  $o_i, i \neq j$ , then p is reassigned to  $o_i$ .

**Case 2:**


'p' currently belongs to representative object,  $o_j$ . If  $o_j$  is replaced by  $o_{\text{random}}$  as a representative object and p is closest to  $o_{\text{random}}$ , then p is reassigned to  $o_{\text{random}}$ .

**Case 3:**

'p' currently belongs to representative object,  $o_i, i \neq j$ . If  $o_j$  is replaced by  $o_{\text{random}}$  as a representative object and p is still closest to  $o_i$ , then the assignment does not change.

**Case 4:**

'p' currently belongs to representative object,  $o_i, i \neq j$ . If  $o_j$  is replaced by  $o_{\text{random}}$  as a representative object and p is closest to  $o_{\text{random}}$ , then p is reassigned to  $o_{\text{random}}$ .



Four cases of the cost function for  $k$ -medoids clustering

### The $k$ -Medoids Algorithm:

The  $k$ -medoids algorithm for partitioning based on medoid or central objects.

#### Input:

- $k$ : the number of clusters,
- $D$ : a data set containing  $n$  objects.

**Output:** A set of  $k$  clusters.

#### Method:

- (1) arbitrarily choose  $k$  objects in  $D$  as the initial representative objects or seeds;
- (2) **repeat**
- (3) assign each remaining object to the cluster with the nearest representative object;
- (4) randomly select a nonrepresentative object,  $o_{random}$ ;
- (5) compute the total cost,  $S$ , of swapping representative object,  $o_j$ , with  $o_{random}$ ;
- (6) **if**  $S < 0$  **then** swap  $o_j$  with  $o_{random}$  to form the new set of  $k$  representative objects;
- (7) **until** no change;

The  $k$ -medoids method is more robust than  $k$ -means in the presence of noise and outliers, because a medoid is less influenced by outliers or other extreme values than a mean. However, its processing is more costly than the  $k$ -means method.

## **Hierarchical Clustering Methods:**

- A hierarchical clustering method works by grouping data objects into a tree of clusters.
- The quality of a pure hierarchical clustering method suffers from its inability to perform adjustment once a merge or split decision has been executed. That is, if a particular merge or split decision later turns out to have been a poor choice, the method cannot backtrack and correct it.

Hierarchical clustering methods can be further classified as either agglomerative or divisive, depending on whether the hierarchical decomposition is formed in a bottom-up or top-down fashion.

### **Agglomerative hierarchical clustering:**

- This bottom-up strategy starts by placing each object in its own cluster and then merges these atomic clusters into larger and larger clusters, until all of the objects are in a single cluster or until certain termination conditions are satisfied.
- Most hierarchical clustering methods belong to this category. They differ only in their definition of intercluster similarity.

### **Divisive hierarchical clustering:**

- This top-down strategy does the reverse of agglomerative hierarchical clustering by starting with all objects in one cluster.
- It subdivides the cluster into smaller and smaller pieces, until each object forms a cluster on its own or until it satisfies certain termination conditions, such as a desired number of clusters is obtained or the diameter of each cluster is within a certain threshold.

## Constraint-Based Cluster Analysis:

Constraint-based clustering finds clusters that satisfy user-specified preferences or constraints. Depending on the nature of the constraints, constraint-based clustering may adopt rather different approaches.

There are a few categories of constraints.

- **Constraints on individual objects:**

We can specify constraints on the objects to be clustered. In a real estate application, for example, one may like to spatially cluster only those luxury mansions worth over a million dollars. This constraint confines the set of objects to be clustered. It can easily be handled by preprocessing after which the problem reduces to an instance of unconstrained clustering.

- **Constraints on the selection of clustering parameters:**

A user may like to set a desired range for each clustering parameter. Clustering parameters are usually quite specific to the given clustering algorithm. Examples of parameters include  $k$ , the desired number of clusters in a  $k$ -means algorithm; or  $e$  the radius and the minimum number of points in the DBSCAN algorithm. Although such user-specified parameters may strongly influence the clustering results, they are usually confined to the algorithm itself. Thus, their fine tuning and processing are usually not considered a form of constraint-based clustering.

- **Constraints on distance or similarity functions:**

We can specify different distance or similarity functions for specific attributes of the objects to be clustered, or different distance measures for specific pairs of objects. When clustering sportsmen, for example, we may use different weighting schemes for height, body weight, age, and skill level. Although this will likely change the mining results, it may not alter the clustering process per se. However, in some cases, such changes may make the evaluation of the distance function nontrivial, especially when it is tightly intertwined with the clustering process.

---

➤ **User-specified constraints on the properties of individual clusters:**

A user may like to specify desired characteristics of the resulting clusters, which may strongly influence the clustering process.

➤ **Semi-supervised clustering based on partial supervision:**

The quality of unsupervised clustering can be significantly improved using some weak form of supervision. This may be in the form of pairwise constraints (i.e., pairs of objects labeled as belonging to the same or different cluster). Such a constrained clustering process is called semi-supervised clustering.

## **Outlier Detection:**

- There exist data objects that do not comply with the general behavior or model of the data. Such data objects, which are grossly different from or inconsistent with the remaining set of data, are called outliers.
- Many data mining algorithms try to minimize the influence of outliers or eliminate them all together. This, however, could result in the loss of important hidden information because one person's noise could be another person's signal. In other words, the outliers may be of particular interest, such as in the case of fraud detection, where outliers may indicate fraudulent activity. Thus, outlier detection and analysis is an interesting data mining task, referred to as outlier mining.
- It can be used in fraud detection, for example, by detecting unusual usage of credit cards or telecommunication services. In addition, it is useful in customized marketing for identifying the spending behavior of customers with extremely low or extremely high incomes, or in medical analysis for finding unusual responses to various medical treatments.

Outlier mining can be described as follows: Given a set of  $n$  data points or objects and  $k$ , the expected number of outliers, find the top  $k$  objects that are considerably dissimilar, exceptional, or inconsistent with respect to the remaining data. The outlier mining problem can be viewed as two subproblems:

- Define what data can be considered as inconsistent in a given data set, and
-

Find an efficient method to mine the outliers so defined.

### Types of outlier detection:

- ☐ Statistical Distribution-Based Outlier Detection
- ☐ Distance-Based Outlier Detection
- ☐ Density-Based Local Outlier Detection
- ☐ Deviation-Based Outlier Detection

### Statistical Distribution-Based Outlier Detection:

The statistical distribution-based approach to outlier detection assumes a distribution or probability model for the given data set (e.g., a normal or Poisson distribution) and then identifies outliers with respect to the model using a discordancy test. Application of the test requires knowledge of the data set parameters knowledge of distribution parameters such as the mean and variance and the expected number of outliers.

A statistical discordancy test examines two hypotheses:

- A working hypothesis
- An alternative hypothesis

A working hypothesis,  $H$ , is a statement that the entire data set of  $n$  objects comes from an initial distribution model,  $F$ , that is,

$$H : o_i \in F, \text{ where } i = 1, 2, \dots, n.$$

The hypothesis is retained if there is no statistically significant evidence supporting its rejection. A discordancy test verifies whether an object,  $o_i$ , is significantly large (or small) in relation to the distribution  $F$ . Different test statistics have been proposed for use as a discordancy test, depending on the available knowledge of the data. Assuming that some statistic,  $T$ , has been chosen for discordancy testing, and the value of the statistic for object  $o_i$  is  $v_i$ , then the distribution of  $T$  is constructed. Significance probability,  $SP(v_i) = \text{Prob}(T > v_i)$ , is evaluated. If  $SP(v_i)$  is sufficiently small, then  $o_i$  is discordant and the working hypothesis is rejected.

An alternative hypothesis,  $H$ , which states that  $o_i$  comes from another distribution model,  $G$ , is adopted. The result is very much dependent on which model  $F$  is chosen because  $o_i$  may be an outlier under one model and a perfectly valid value under another. The

alternative distribution is very important in determining the power of the test, that is, the probability that the working hypothesis is rejected when  $o_i$  is really an outlier.

There are different kinds of alternative distributions.

- **Inherent alternative distribution:**

In this case, the working hypothesis that all of the objects come from distribution  $F$  is rejected in favor of the alternative hypothesis that all of the objects arise from another distribution,  $G$ :

$H : o_i \in G$ , where  $i = 1, 2, \dots, n$

$F$  and  $G$  may be different distributions or differ only in parameters of the same distribution.

There are constraints on the form of the  $G$  distribution in that it must have potential to produce outliers. For example, it may have a different mean or dispersion, or a longer tail.

- **Mixture alternative distribution:**

The mixture alternative states that discordant values are not outliers in the  $F$  population, but contaminants from some other population,

$G$ . In this case, the alternative hypothesis is

$$\bar{H} : o_i \in (1 - \lambda)F + \lambda G, \text{ where } i = 1, 2, \dots, n.$$

- **Slippage alternative distribution:**

This alternative states that all of the objects (apart from some prescribed small number) arise independently from the initial model,  $F$ , with its given parameters, whereas the remaining objects are independent observations from a modified version of  $F$  in which the parameters have been shifted.

There are two basic types of procedures for detecting outliers:

**Block procedures:**

In this case, either all of the suspect objects are treated as outliers or all of them are accepted as consistent.

**Consecutive procedures:**

An example of such a procedure is the *insideout* procedure. Its main idea is that the object that is least likely to be an outlier is tested first. If it is found to be an outlier, then all of

the more extreme values are also considered outliers; otherwise, the next most extreme object is tested, and so on. This procedure tends to be more effective than block procedures.

### **Distance-Based Outlier Detection:**

The notion of distance-based outliers was introduced to counter the main limitations imposed by statistical methods. An object,  $o$ , in a data set,  $D$ , is a distance-based (DB) outlier with parameters  $pct$  and  $dmin$ , that is, a  $DB(pct; dmin)$ -outlier, if at least a fraction,  $pct$ , of the objects in  $D$  lie at a distance greater than  $dmin$  from  $o$ .

In other words, rather than relying on statistical tests, we can think of distance-based outliers as those objects that do not have enough neighbors, where neighbors are defined based on distance from the given object.

In comparison with statistical-based methods, distance based outlier detection generalizes the ideas behind discordancy testing for various standard distributions. Distance-based outlier detection avoids the excessive computation that can be associated with fitting the observed distribution into some standard distribution and in selecting discordancy tests.

For many discordancy tests, it can be shown that if an object,  $o$ , is an outlier according to the given test, then  $o$  is also a  $DB(pct, dmin)$ -outlier for some suitably defined  $pct$  and  $dmin$ .

For example, if objects that lie three or more standard deviations from the mean are considered to be outliers, assuming a normal distribution, then this definition can be generalized by a  $DB(0.9988, 0.13s)$  outlier.

Several efficient algorithms for mining distance-based outliers have been developed.

#### **Index-based algorithm:**

Given a data set, the index-based algorithm uses multidimensional indexing structures, such as R-trees or k-d trees, to search for neighbors of each object  $o$  within radius  $dmin$  around that object.

Let  $M$  be the maximum number of objects within the  $dmin$ -neighborhood of an outlier. Therefore, once  $M+1$  neighbors of object  $o$  are found, it is clear that  $o$  is not an outlier.

This algorithm has a worst-case complexity of  $O(n^2k)$ , where  $n$  is the number of objects in the data set and  $k$  is the dimensionality. The index-based algorithm scales well as  $k$  increases.

However, this complexity evaluation takes only the search time into account, even though the task of building an index in itself can be computationally intensive.

#### **Nested-loop algorithm:**

The nested-loop algorithm has the same computational complexity as the index-based algorithm but avoids index structure construction and tries to minimize the number of I/Os.

It divides the memory buffer space into two halves and the data set into several logical blocks. By carefully choosing the order in which blocks are loaded into each half, I/O efficiency can be achieved.

#### **Cell-based algorithm:**

To avoid  $O(n^2)$  computational complexity, a cell-based algorithm was developed for memory-resident data sets. Its complexity is  $O(c^k + n)$ , where  $c$  is a constant depending on the number of cells and  $k$  is the dimensionality.

In this method, the data space is partitioned into cells with a side length equal to  $\frac{d_{min}}{2\sqrt{k}}$ . Each cell has two layers surrounding it. The first layer is one cell thick, while the second is cells thick, rounded up to the closest integer.

The algorithm counts outliers on a cell-by-cell rather than an object-by-object basis. For a given cell, it accumulates three counts—the number of objects in the cell, in the cell and the first layer together, and in the cell and both layers together. Let's refer to these counts as cell count, cell + 1 layer count, and cell + 2 layers count, respectively.

Let  $M$  be the maximum number of outliers that can exist in the  $d_{min}$ -neighborhood of an outlier.

- An object,  $o$ , in the current cell is considered an outlier only if cell + 1 layer count is less than or equal to  $M$ . If this condition does not hold, then all of the objects in the cell can be removed from further investigation as they cannot be outliers.
- If cell + 2 layers count is less than or equal to  $M$ , then all of the objects in the cell are considered outliers. Otherwise, if this number is more than  $M$ , then it is possible that some of the objects in the cell may be outliers. To detect these outliers, object-by-object processing is used where, for each object,  $o$ , in the cell, objects in the second layer of  $o$  are examined. For objects in the cell, only those objects having no more than  $M$  points in their  $d_{min}$ -neighborhoods are outliers. The  $d_{min}$ -neighborhood of an object consists of the object's cell, all of its first layer, and some of its second layer.

A variation to the algorithm is linear with respect to  $n$  and guarantees that no more than three passes over the data set are required. It can be used for large disk-resident data sets, yet does not scale well for high dimensions.

### **Density-Based Local Outlier Detection:**

Statistical and distance-based outlier detection both depend on the overall or global distribution of the given set of data points,  $D$ . However, data are usually not uniformly distributed. These methods encounter difficulties when analyzing data with rather different density distributions.

To define the local outlier factor of an object, we need to introduce the concepts of  $k$ -distance,  $k$ -distance neighborhood, reachability distance,<sup>13</sup> and local reachability density.

These are defined as follows:

The  $k$ -distance of an object  $p$  is the maximal distance that  $p$  gets from its  $k$ -nearest neighbors. This distance is denoted as  $k\text{-distance}(p)$ . It is defined as the distance,  $d(p, o)$ , between  $p$  and an object  $o \in D$ , such that for at least  $k$  objects,  $o_0 \in D$ , it holds that  $d(p, o_0) \leq d(p, o)$ . That is, there are at least  $k$  objects in  $D$  that are as close as or closer to  $p$  than  $o$ , and for at most  $k-1$  objects,  $o' \in D$ , it holds that  $d(p, o') < d(p, o)$ .

That is, there are at most  $k-1$  objects that are closer to  $p$  than  $o$ . You may be wondering at this point how  $k$  is determined. The LOF method links to density-based clustering in that it sets  $k$  to the parameter  $rMinPts$ , which specifies the minimum number of points for use in identifying clusters based on density.

Here,  $MinPts$  (as  $k$ ) is used to define the local neighborhood of an object,  $p$ .

The  $k$ -distance neighborhood of an object  $p$  is denoted  $N_{k\text{-distance}(p)}(p)$ , or  $N_k(p)$  for short. By setting  $k$  to  $MinPts$ , we get  $N_{MinPts}(p)$ . It contains the  $MinPts$ -nearest neighbors of  $p$ . That is, it contains every object whose distance is not greater than the  $MinPts$ -distance of  $p$ .

The reachability distance of an object  $p$  with respect to object  $o$  (where  $o$  is within the  $MinPts$ -nearest neighbors of  $p$ ), is defined as reach

$$\text{distMinPts}(p, o) = \max\{\text{MinPtsdistance}(o), d(p, o)\}.$$

Intuitively, if an object  $p$  is far away, then the reachability distance between the two is simply their actual distance. However, if they are sufficiently close (i.e., where  $p$  is within the  $MinPts$ -distance neighborhood of  $o$ ), then the actual distance is replaced by the  $MinPts$ -distance of  $o$ . This helps to significantly reduce the statistical fluctuations of  $d(p, o)$  for all of the  $p$  close to  $o$ .

The higher the value of  $MinPts$  is, the more similar is the reachability distance for objects within the same neighborhood.

Intuitively, the local reachability density of  $p$  is the inverse of the average reachability density based on the  $MinPts$ -nearest neighbors of  $p$ . It is defined as

$$lrd_{MinPts}(p) = \frac{|N_{MinPts}(p)|}{\sum_{o \in N_{MinPts}(p)} reach\_dist_{MinPts}(p, o)}$$

The local outlier factor (LOF) of  $p$  captures the degree to which we call  $p$  an outlier. It is defined as

$$LOF_{MinPts}(p) = \frac{\sum_{o \in N_{MinPts}(p)} \frac{lrd_{MinPts}(o)}{lrd_{MinPts}(p)}}{|N_{MinPts}(p)|}$$

It is the average of the ratio of the local reachability density of  $p$  and those of  $p$ 's  $MinPts$ -nearest neighbors. It is easy to see that the lower  $p$ 's local reachability density is, and the higher the local reachability density of  $p$ 's  $MinPts$ -nearest neighbors are, the higher  $LOF(p)$  is.

### Deviation-Based Outlier Detection:

Deviation-based outlier detection does not use statistical tests or distance-based measures to identify exceptional objects. Instead, it identifies outliers by examining the main characteristics of objects in a group. Objects that deviate from this description are considered outliers. Hence, in this approach the term deviations is typically used to refer to outliers. In this section, we study two techniques for deviation-based outlier detection. The first sequentially compares objects in a set, while the second employs an OLAP data cube approach.

### **Sequential Exception Technique:**

The sequential exception technique simulates the way in which humans can distinguish unusual objects from among a series of supposedly like objects. It uses implicit redundancy of the data. Given a data set,  $D$ , of  $n$  objects, it builds a sequence of subsets,  $\{D_1, D_2, \dots, D_m\}$ , of these objects with  $2 \leq m \leq n$  such that

$$D_{j-1} \subset D_j, \text{ where } D_j \subseteq D.$$

Dissimilarities are assessed between subsets in the sequence. The technique introduces the following key terms.

#### **Exception set:**

This is the set of deviations or outliers. It is defined as the smallest subset of objects whose removal results in the greatest reduction of dissimilarity in the residual set.

#### **Dissimilarity function:**

This function does not require a metric distance between the objects. It is any function that, if given a set of objects, returns a low value if the objects are similar to one another. The greater the dissimilarity among the objects, the higher the value returned by the function. The dissimilarity of a subset is incrementally computed based on the subset prior to it in the sequence. Given a subset of  $n$  numbers,  $\{x_1, \dots, x_n\}$ , a possible dissimilarity function is the variance of the numbers in the set, that is,

$$\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2,$$

where  $\bar{x}$  is the mean of the  $n$  numbers in the set. For character strings, the dissimilarity function may be in the form of a pattern string (e.g., containing wildcard characters that is used to cover all of the patterns seen so far. The dissimilarity increases when the pattern covering all of the strings in  $D_{j-1}$  does not cover any string in  $D_j$  that is not in  $D_{j-1}$ .

#### **Cardinality function:**

This is typically the count of the number of objects in a given set.

#### **Smoothing factor:**

This function is computed for each subset in the sequence. It assesses how much the dissimilarity can be reduced by removing the subset from the original set of objects.

## Unit-5

### **Machine Learning with WEKA WEKA Explorer Tutorial for WEKA Version 3.4.3**

#### **TABLE OF CONTENTS**

|                                                  |           |
|--------------------------------------------------|-----------|
| <b>1. INTRODUCTION .....</b>                     | <b>2</b>  |
| <b>2. LAUNCHING WEKA EXPLORER.....</b>           | <b>2</b>  |
| <b>3. PREPROCESSING DATA .....</b>               | <b>3</b>  |
| 3.1. FILE CONVERSION .....                       | 4         |
| 3.2. OPENING FILE FROM A LOCAL FILE SYSTEM ..... | 5         |
| 3.3. OPENING FILE FROM A WEB SITE .....          | 7         |
| 3.4. READING DATA FROM A DATABASE .....          | 8         |
| 3.5. PREPROCESSING WINDOW .....                  | 9         |
| 3.6. SETTING FILTERS.....                        | 13        |
| <b>4. BUILDING “CLASSIFIERS” .....</b>           | <b>16</b> |
| 4.1. CHOOSING A CLASSIFIER .....                 | 17        |
| 4.2. SETTING TEST OPTIONS .....                  | 17        |
| 4.3. ANALYZING RESULTS .....                     | 21        |
| 4.4. VISUALIZATION OF RESULTS.....               | 22        |
| <i>Classification Exercise</i> .....             | 25        |
| <b>5. CLUSTERING DATA .....</b>                  | <b>25</b> |
| 5.1. CHOOSING CLUSTERING SCHEME.....             | 26        |
| 5.2. SETTING TEST OPTIONS .....                  | 27        |
| 5.3. ANALYZING RESULTS .....                     | 29        |
| 5.4. VISUALIZATION OF RESULTS.....               | 30        |
| <i>Clustering Exercise</i> .....                 | 32        |
| <b>6. FINDING ASSOCIATIONS.....</b>              | <b>32</b> |
| 6.1. CHOOSING ASSOCIATION SCHEME .....           | 32        |
| 6.2. SETTING TEST OPTIONS .....                  | 33        |
| 6.3. ANALYZING RESULTS .....                     | 35        |
| <i>Association Rules Exercise</i> .....          | 35        |
| <b>7. ATTRIBUTE SELECTION.....</b>               | <b>35</b> |
| 7.1. SELECTING OPTIONS.....                      | 36        |
| 7.2. ANALYZING RESULTS .....                     | 37        |
| 7.3. VISUALIZING RESULTS.....                    | 37        |
| <b>8. DATA VISUALIZATION .....</b>               | <b>39</b> |
| 8.1. CHANGING THE VIEW.....                      | 40        |
| 8.2. SELECTING INSTANCES.....                    | 41        |
| <b>9. CONCLUSION .....</b>                       | <b>43</b> |
| <b>10. REFERENCES .....</b>                      | <b>44</b> |

## 1. Introduction


WEKA is a data mining system developed by the University of Waikato in New Zealand that implements data mining algorithms. WEKA is a state-of-the-art facility for developing machine learning (ML) techniques and their application to real-world data mining problems. It is a collection of machine learning algorithms for data mining tasks. The algorithms are applied directly to a dataset. WEKA implements algorithms for data preprocessing, classification, regression, clustering, association rules; it also includes a visualization tools. The new machine learning schemes can also be developed with this package. WEKA is open source software issued under the GNU General Public License [3].

The goal of this Tutorial is to help you to learn WEKA Explorer. The tutorial will guide you step by step through the analysis of a simple problem using WEKA Explorer preprocessing, classification, clustering, association, attribute selection, and visualization tools. At the end of each problem there is a representation of the results with explanations side by side. Each part is concluded with the exercise for individual practice. By the time you reach the end of this tutorial, you will be able to analyze your data with WEKA Explorer using various learning schemes and interpret received results.

Before starting this tutorial, you should be familiar with data mining algorithms such as C4.5 (C5), ID3, K-means, and Apriori. All working files are provided. For better performance, the archive of all files used in this tutorial can be downloaded or copied from CD to your hard drive as well as a printable version of the lessons. A trial version of Weka package can be downloaded from the University of Waikato website at <http://www.cs.waikato.ac.nz/~ml/weka/index.html>.

## 2. Launching WEKA Explorer


You can launch Weka from C:\Program Files directory, from your desktop selecting  icon, or from the Windows task bar 'Start' → 'Programs' → 'Weka 3-4'. When 'WEKA GUI Chooser' window appears on the screen, you can select one of the four options at the bottom of the window [2]:




1. **Simple CLI** provides a simple command-line interface and allows direct execution of Weka commands.

2. **Explorer** is an environment for exploring data.
3. **Experimenter** is an environment for performing experiments and conducting statistical tests between learning schemes.
4. **KnowledgeFlow** is a Java-Beans-based interface for setting up and running machine learning experiments.

For the exercises in this tutorial you will use 'Explorer'. Click on 'Explorer' button in the 'WEKA GUI Chooser' window.



'WEKA Explorer' window appears on a screen.



### 3. Preprocessing Data


At the very top of the window, just below the title bar there is a row of tabs. Only the first tab, 'Preprocess', is active at the moment because there is no dataset open. The first three

buttons at the top of the preprocess section enable you to load data into WEKA. Data can be imported from a file in various formats: ARFF, CSV, C4.5, binary, it can also be read from a URL or from an SQL database (using JDBC) [4]. The easiest and the most common way of getting the data into WEKA is to store it as Attribute-Relation File Format (ARFF) file.

You've already been given "weather.arff" file for this exercise; therefore, you can skip section 3.1 that will guide you through the file conversion.

### 3.1. File Conversion


We assume that all your data stored in a Microsoft Excel spreadsheet "weather.xls".




The screenshot shows a Microsoft Excel window titled "Microsoft Excel - weather.xls". The spreadsheet contains a dataset with 5 columns: outlook, temperature, humidity, windy, and play. The data is organized into rows, with the first row serving as a header. The following table represents the data shown in the spreadsheet:

|    | A        | B          | C        | D     | E    |
|----|----------|------------|----------|-------|------|
| 1  | outlook  | temperatur | humidity | windy | play |
| 2  |          |            |          |       |      |
| 3  | sunny    | 85         | 85       | f     | no   |
| 4  | sunny    | 80         | 90       | t     | no   |
| 5  | overcast | 83         | 86       | f     | yes  |
| 6  | rainy    | 70         | 96       | f     | yes  |
| 7  | rainy    | 68         | 80       | f     | yes  |
| 8  | rainy    | 65         | 70       | t     | no   |
| 9  | overcast | 64         | 65       | t     | yes  |
| 10 | sunny    | 72         | 95       | f     | no   |
| 11 | sunny    | 69         | 70       | f     | yes  |
| 12 | rainy    | 75         | 80       | f     | yes  |
| 13 | sunny    | 75         | 70       | t     | yes  |
| 14 | overcast | 72         | 90       | t     | yes  |
| 15 | overcast | 81         | 75       | f     | yes  |
| 16 | rainy    | 71         | 91       | t     | no   |
| 17 |          |            |          |       |      |

WEKA expects the data file to be in Attribute-Relation File Format (ARFF) file. Before you apply the algorithm to your data, you need to convert your data into comma-separated file into ARFF format (into the file with .arff extension) [1]. To save you data in comma-separated format, select the 'Save As...' menu item from Excel 'File' pull-down menu. In the ensuing dialog box select 'CSV (Comma Delimited)' from the file type pop-up menu, enter a name of the file, and click 'Save' button. Ignore all messages that appear by clicking 'OK'. Open this file with Microsoft Word. Your screen will look like the screen below.




The rows of the original spreadsheet are converted into lines of text where the elements are separated from each other by commas. In this file you need to change the first line, which holds the attribute names, into the header structure that makes up the beginning of an ARFF file. Add a `@relation` tag with the dataset's name, an `@attribute` tag with the attribute information, and a `@data` tag as shown below.




Choose 'Save As...' from the 'File' menu and specify 'Text Only with Line Breaks' as the file type. Enter a file name and click 'Save' button. Rename the file to the file with extension `.arff` to indicate that it is in ARFF format.

### 3.2. Opening file from a local file system


Click on 'Open file...' button.



It brings up a dialog box allowing you to browse for the data file on the local file system, choose “weather.arff” file.




Some databases have the ability to save data in CSV format. In this case, you can select CSV file from the local filesystem. If you would like to convert this file into ARFF format, you can click on ‘Save’ button. WEKA automatically creates ARFF file from your CSV file.




### 3.3. Opening file from a web site

A file can be opened from a website. Suppose, that “weather.arff” is on the following website:




The URL of the web site in our example is <http://gaia.ecs.csus.edu/~aksenovs/>. It means that the file is stored in this directory, just as in the case with your local file system. To open this file, click on ‘Open URL...’ button, it brings up a dialog box requesting to enter source URL.



Enter the URL of the web site followed by the file name, in this example the URL is <http://gaia.ecs.csus.edu/~aksenovs/weather.arff>, where weather.arff is the name of the file you are trying to load from the website.


### 3.4. Reading data from a database

Data can also be read from an SQL database using JDBC. Click on 'Open DB...' button, 'GenericObjectEditor' appears on the screen.



To read data from a database, click on 'Open' button and select the database from a filesystem.

### 3.5. Preprocessing window



At the bottom of the window there is 'Status' box. The 'Status' box displays messages that keep you informed about what is going on. For example, when you first opened the 'Explorer', the message says, "Welcome to the Weka Explorer". When you loading "weather.arff" file, the 'Status' box displays the message "Reading from file...". Once the file is loaded, the message in the 'Status' box changes to say "OK". Right-click anywhere in 'Status box', it brings up a menu with two options:

1. **Available Memory** that displays in the log and in 'Status' box the amount of memory available to WEKA in bytes.
2. **Run garbage collector** that forces Java garbage collector to search for memory that is no longer used, free this memory up and to allow this memory for new tasks.

To the right of 'Status box' there is a 'Log' button that opens up the log. The log records every action in WEKA and keeps a record of what has happened. Each line of text in the log contains time of entry. For example, if the file you tried to open is not loaded, the log will have record of the problem that occurred during opening.

To the right of the 'Log' button there is an image of a bird. The bird is WEKA status icon. The number next to 'X' symbol indicates a number of concurrently running processes. When you loading a file, the bird sits down that means that there are no processes running. The number of processes besides symbol 'X' is zero that means that the system is idle. Later, in classification problem, when generating result look at the bird, it gets up and start moving that indicates that a process started. The number next to 'X' becomes 1 that means that there is one process running, in this case calculation.




If the bird is standing and not moving for a long time, it means that something has gone wrong. In this case you should restart WEKA Explorer.

### Loading data

Lets load the data and look what is happening in the 'Preprocess' window.

The most common and easiest way of loading data into WEKA is from ARFF file, using 'Open file...' button (section 3.2). Click on 'Open file...' button and choose "weather.arff" file from your local filesystem. Note, the data can be loaded from CSV file as well because some databases have the ability to convert data only into CSV format.



Once the data is loaded, WEKA recognizes attributes that are shown in the 'Attribute' window. Left panel of 'Preprocess' window shows the list of recognized attributes:

**No.** is a number that identifies the order of the attribute as they are in data file, **Selection tick boxes** allow you to select the attributes for working relation, **Name** is a name of an attribute as it was declared in the data file.

The 'Current relation' box above 'Attribute' box displays the base relation (table) name and the current working relation (which are initially the same) - "weather", the number of instances - 14 and the number of attributes - 5.

During the scan of the data, WEKA computes some basic statistics on each attribute. The following statistics are shown in 'Selected attribute' box on the right panel of 'Preprocess' window:

**Name** is the name of an attribute,

**Type** is most commonly Nominal or Numeric, and

**Missing** is the number (percentage) of instances in the data for which this attribute is unspecified,


**Distinct** is the number of different values that the data contains for this attribute, and

**Unique** is the number (percentage) of instances in the data having a value for this attribute that no other instances have.

An attribute can be deleted from the 'Attributes' window. Highlight an attribute you would like to delete and hit Delete button on your keyboard.

By clicking on an attribute, you can see the basic statistics on that attribute. The frequency for each attribute value is shown for categorical attributes. Min, max, mean, standard deviation (StdDev) is shown for continuous attributes.

Click on attribute Outlook in the 'Attribute' window.




Outlook is nominal. Therefore, you can see the following frequency statistics for this attribute in the 'Selected attributes' window:

Missing = 0 means that the attribute is specified for all instances (no missing values), Distinct = 3 means that Outlook has three different values: sunny, overcast, rainy, and Unique = 0 means that other instances do not have the same value as Outlook has.

Just below these values there is a table displaying count of instances of the attribute Outlook. As you can see, there are three values: sunny with 5 instances, overcast with 4 instances, and rainy with 5 instances. These numbers match the numbers of instances in the base relation and table "weather.xls".

Lets take a look at the attribute Temperature.




|    | A        | B           | C        | D     | E    |
|----|----------|-------------|----------|-------|------|
| 1  | outlook  | temperature | humidity | windy | play |
| 2  |          |             |          |       |      |
| 3  | sunny    | 85          | 85       | f     | no   |
| 4  | sunny    | 80          | 90       | t     | no   |
| 5  | overcast | 83          | 86       | f     | yes  |
| 6  | rainy    | 70          | 96       | f     | yes  |
| 7  | rainy    | 68          | 80       | f     | yes  |
| 8  | rainy    | 65          | 70       | t     | no   |
| 9  | overcast | 64          | 65       | t     | yes  |
| 10 | sunny    | 72          | 95       | f     | no   |
| 11 | sunny    | 69          | 70       | f     | yes  |
| 12 | rainy    | 75          | 80       | f     | yes  |
| 13 | sunny    | 75          | 70       | t     | yes  |
| 14 | overcast | 72          | 90       | t     | yes  |
| 15 | overcast | 81          | 75       | f     | yes  |
| 16 | rainy    | 71          | 91       | t     | no   |

Temperature is a numeric value; therefore, you can see min, max, means, and standard deviation in 'Selected Attribute' window.


Missing = 0 means that the attribute is specified for all instances (no missing values), Distinct = 12 means that Temperature has twelve different values, and Unique = 10 means that other attributes or instances have the same 10 value as Temperature has.

Temperature is a Numeric value; therefore, you can see the statistics describing the distribution of values in the data - Minimum, Maximum, Mean and Standard Deviation. Minimum = 64 is the lowest temperature, Maximum = 85 is the highest temperature, mean and standard deviation. Compare the result with the attribute table "weather.xls"; the numbers in WEKA match the numbers in the table.

You can select a class in the 'Class' pull-down box. The last attribute in the 'Attributes' window is the default class selected in the 'Class' pull-down box.



You can Visualize the attributes based on selected class. One way is to visualize selected attribute based on class selected in the 'Class' pull-down window, or visualize all attributes by clicking on 'Visualize All' button.




### 3.6. Setting Filters


Pre-processing tools in WEKA are called "filters". WEKA contains filters for discretization, normalization, resampling, attribute selection, transformation and combination of attributes [4]. Some techniques, such as association rule mining, can only be performed on categorical data. This requires performing discretization on numeric or continuous attributes [5]. For classification example you do not need to transform the data. For you practice, suppose you need to perform a test on categorical data. There are two attributes that need to be converted: 'temperature' and 'humidity'. In other words, you will keep all of the values for these attributes in the data. This means you can discretize by removing the keyword "numeric" as the type for the

'temperature' attribute and replace it with the set of "nominal" values. You can do this by applying a filter.


In 'Filters' window, click on the 'Choose' button.




This will show pull-down menu with a list of available filters. Select Supervised → Attribute → Discretize and click on 'Apply' button. The filter will convert Numeric values into Nominal.




When filter is chosen, the fields in the window changes to reflect available options.



As you can see, there is no change in the value Outlook. Select value Temperature, look at the 'Selected attribute' box, the 'Type' field shows that the attribute type has changed from Numeric to Nominal. The list has changed as well: instead of statistical values there is count of instances, and the count of it is 14 that means that there are 14 instances of the value Temperature.



Note, when you right-click on filter, a 'GenericObjectEditor' dialog box comes up on your screen. The box lets you to choose the filter configuration options. The same box can be used for classifiers, clusters and association rules. Clicking on 'More' button brings up an 'Information' window describing what the different options can do.




At the bottom of the editor window there are four buttons. 'Open' and 'Save' buttons allow you to save object configurations for future use. 'Cancel' button allows you to exit without saving changes. Once you have made changes, click 'OK' to apply them.


#### 4. Building "Classifiers"

Classifiers in WEKA are the models for predicting nominal or numeric quantities. The learning schemes available in WEKA include decision trees and lists, instance-based classifiers, support vector machines, multi-layer perceptrons, logistic regression, and bayes' nets. "Meta"-classifiers include bagging, boosting, stacking, error-correcting output codes, and locally weighted learning [4].

Once you have your data set loaded, all the tabs are available to you. Click on the 'Classify' tab.




'Classify' window comes up on the screen.



Now you can start analyzing the data using the provided algorithms. In this exercise you will analyze the data with C4.5 algorithm using J48, WEKA's implementation of decision tree learner. The sample data used in this exercise is the weather data from the file "weather.arff". Since C4.5 algorithm can handle numeric attributes, in contrast to the ID3 algorithm from which C4.5 has evolved, there is no need to discretize any of the attributes. Before you start this exercise, make sure you do not have filters set in the 'Preprocess' window. Filter exercise in section 3.6 was just a practice.

#### 4.1. Choosing a Classifier

Click on 'Choose' button in the 'Classifier' box just below the tabs and select C4.5 classifier WEKA → Classifiers → Trees → J48.




#### 4.2. Setting Test Options

Before you run the classification algorithm, you need to set test options. Set test options in the 'Test options' box. The test options that available to you are [2]:


1. **Use training set.** Evaluates the classifier on how well it predicts the class of the instances it was trained on.
2. **Supplied test set.** Evaluates the classifier on how well it predicts the class of a set of instances loaded from a file. Clicking on the 'Set...' button brings up a dialog allowing you to choose the file to test on.
3. **Cross-validation.** Evaluates the classifier by cross-validation, using the number of folds that are entered in the 'Folds' text field.
4. **Percentage split.** Evaluates the classifier on how well it predicts a certain percentage of the data, which is held out for testing. The amount of data held out depends on the value entered in the '%' field.

In this exercise you will evaluate classifier based on how well it predicts 66% of the tested data. Check 'Percentage split' radio-button and keep it as default 66%. Click on 'More options...' button.



Identify what is included into the output. In the 'Classifier evaluation options' make sure that the following options are checked [2]:


1. **Output model.** The output is the classification model on the full training set, so that it can be viewed, visualized, etc.
2. **Output per-class stats.** The precision/recall and true/false statistics for each class output.
3. **Output confusion matrix.** The confusion matrix of the classifier's predictions is included in the output.
4. **Store predictions for visualization.** The classifier's predictions are remembered so that they can be visualized.
5. Set '**Random seed for Xval / % Split**' to 1. This specifies the random seed used when randomizing the data before it is divided up for evaluation purposes.




The remaining options that you do not use in this exercise but that available to you are:

6. **Output entropy evaluation measures.** Entropy evaluation measures are included in the output.
7. **Output predictions.** The classifier's predictions are remembered so that they can be visualized.

Once the options have been specified, you can run the classification algorithm. Click on 'Start' button to start the learning process. You can stop learning process at any time by clicking on 'Stop' button.



When training set is complete, the 'Classifier' output area on the right panel of 'Classify' window is filled with text describing the results of training and testing. A new entry appears in the 'Result list' box on the left panel of 'Classify' window.



### 4.3. Analyzing Results

=== Run information ===

Scheme: weka.classifiers.trees.J48 -C 0.25 -M 2  
Relation: weather  
Instances: 14  
Attributes: 5  
outlook  
temperature  
humidity  
windy  
play

Test mode: split 66% train, remainder test

=== Classifier model (full training set) ===

J48 pruned tree

-----

outlook = sunny  
| humidity <= 75: yes (2.0)  
| humidity > 75: no (3.0)  
outlook = overcast: yes (4.0)  
outlook = rainy  
| windy = t: no (2.0)  
| windy = f: yes (3.0)

Number of Leaves : 5

Size of the tree : 8

Time taken to build model: 0.06 seconds

=== Evaluation on test split ===

=== Summary ===

|                                  |          |    |   |
|----------------------------------|----------|----|---|
| Correctly Classified Instances   | 2        | 40 | % |
| Incorrectly Classified Instances | 3        | 60 | % |
| Kappa statistic                  | -0.3636  |    |   |
| Mean absolute error              | 0.6      |    |   |
| Root mean squared error          | 0.7746   |    |   |
| Relative absolute error          | 126.9231 | %  |   |
| Root relative squared error      | 157.6801 | %  |   |
| Total Number of Instances        | 5        |    |   |

=== Detailed Accuracy By Class ===

| TP Rate | FP Rate | Precision | Recall | F-Measure | Class |
|---------|---------|-----------|--------|-----------|-------|
| 0.667   | 1       | 0.5       | 0.667  | 0.571     | yes   |
| 0       | 0.333   | 0         | 0      | 0         | no    |

=== Confusion Matrix ===

a b <-- classified as  
2 1 | a = yes  
2 0 | b = no

Run Information gives you the following information:

- the algorithm you used - J48
  - the relation name – “weather”
  - number of instances in the relation – 14
  - number of attributes in the relation – 5 and the list of the attributes: outlook, temperature, humidity, windy, play
- 
- the test mode you selected: split=66%

Classifier model is a pruned decision tree in textual form that was produced on the full training data. As you can see, the first split is on the ‘outlook’ attribute, at the second level, the splits are on ‘humidity’ and ‘windy’.

In the tree structure, a colon represents the class label that has been assigned to a particular leaf, followed by the number of instances that reach that leaf.

Below the tree structure, there is a number of leaves (which is 5), and the number of nodes in the tree - size of the tree (which is 8). The program gives a time it took to build the model, which is 0.06 seconds.

Evaluation on test split. This part of the output gives estimates of the tree’s predictive performance, generated by WEKA’s evaluation module. It outputs the list of statistics summarizing how accurately the classifier was able to predict the true class of the instances under the chosen test module. The set of measurements is derived from the training data.


In this case only 40% of 14 training instances have been classified correctly. This indicates that the results obtained from the training data are not optimistic compared with what might be obtained from the independent test set from the same source. In addition to classification error, the evaluation output measurements derived from the class probabilities assigned by the tree. More specifically, it outputs mean output error (0.6) of the probability estimates, the root mean squared error (0.77) is the square root of the quadratic loss. The mean absolute error calculated in a similar way by using the absolute instead of squared difference. The reason that the errors are not 1 or 0 is because not all training instances are classified correctly.

Detailed Accuracy By Class demonstrates a more detailed per-class break down of the classifier’s prediction accuracy.


From the Confusion matrix you can see that one instance of a class ‘yes’ have been assigned to a class ‘no’, and two of class ‘no’ are assigned to class ‘yes’.

## 4.4. Visualization of Results


After training a classifier, the result list adds an entry.




WEKA lets you to see a graphical representation of the classification tree. Right-click on the entry in 'Result list' for which you would like to visualize a tree. It invokes a menu containing the following items:




Select the item 'Visualize tree'; a new window comes up to the screen displaying the tree.




WEKA also lets you to visualize classification errors. Right-click on the entry in 'Result list' again and select 'Visualize classifier errors' from the menu:



'Weka Classifier Visualize' window displaying graph appears on the screen.



On the 'Weka Classifier Visualize' window, beneath the X-axis selector there is a drop-down list, 'Colour', for choosing the color scheme. This allows you to choose the color of points based on the attribute selected. Below the plot area, there is a legend that describes what values the colors correspond to. In your example, red represents 'no', while blue represents 'yes'. For better visibility you should change the color of label 'yes'. Left-click on 'yes' in the 'Class colour' box and select lighter color from the color palette.



To the right of the plot area there are series of horizontal strips. Each strip represents an attribute, and the dots within it show the distribution values of the attribute. You can choose what axes are used in the main graph by clicking on these strips (left-click changes X-axis, right-click changes Y-axis).

Change X - axis to 'Outlook' attribute and Y - axis to 'Play'. The instances are spread out in the plot area and concentration points are not visible. Keep sliding 'Jitter', a random displacement given to all points in the plot, to the right, until you can spot concentration points.



On the plot you can see the results of classification. Correctly classified instances are represented as crosses, incorrectly classified once represented as squares. In this example in the left lower corner you can see blue cross indicating correctly classified instance: if Outlook = 'sunny' → play = 'yes'.

Look to the upper left corner of the graph, there are two red squares in this corner. The square represents incorrectly classified instance. The following is not correct: if Outlook = 'sunny' → play = 'no'.


### **Classification Exercise**

Use ID3 algorithm to classify weather data from the “weather.arff” file. Perform initial preprocessing and create a version of the initial dataset in which all numeric attributes should be converted to categorical data.

## **5. Clustering Data**

WEKA contains “clusterers” for finding groups of similar instances in a dataset. The clustering schemes available in WEKA are *k*-Means, EM, Cobweb, X-means, FarthestFirst. Clusters can be visualized and compared to “true” clusters (if given). Evaluation is based on log likelihood if clustering scheme produces a probability distribution [4].


For this exercise we will use customer data [6] that is contained in “customers.arff” file and analyze it with *k*-means clustering scheme.



|    | A      | B   | C        | D              | E               |
|----|--------|-----|----------|----------------|-----------------|
| 1  | Income | Age | Children | Marital Status | Education       |
| 2  |        |     |          |                |                 |
| 3  | 25000  | 35  | 3        | single         | high school     |
| 4  | 15000  | 25  | 1        | married        | high school     |
| 5  | 20000  | 40  | 0        | single         | high school     |
| 6  | 30000  | 20  | 0        | divorced       | high school     |
| 7  | 20000  | 25  | 3        | divorced       | college         |
| 8  | 70000  | 60  | 0        | married        | college         |
| 9  | 90000  | 30  | 0        | married        | graduate school |
| 10 | 200000 | 45  | 5        | married        | graduate school |
| 11 | 100000 | 50  | 2        | divorced       | college         |
| 12 |        |     |          |                |                 |


An international online catalog company wishes to group its customers based on common features. Company management does not have any predefined labels for these groups. Based on the outcome of the grouping, they will target marketing and advertising campaigns to the different groups. The information they have about the customers includes income, age, number of children, marital status, and education. For our exercise we will use a part of the database for customers in US. Depending on the type of advertising, not all attributes are important. For example, suppose the advertising is for a special sale on children’s clothes. We will target the advertising only to the persons with young children. The clustering that you will perform in this exercise is as follows. The first group of people has young children and a high school degree, the second group does not have children but has high school degree. The third group has both children and a college degree. The fourth group has higher income and at least a college degree. The fifth group has children and higher degree. Different clustering would have been found by examining either age or marital status.

In ‘Preprocess’ window click on ‘Open file...’ button and select “customers.arff” file. Click ‘Cluster’ tab at the top of WEKA Explorer window.




## 5.1. Choosing Clustering Scheme

In the 'Clusterer' box click on 'Choose' button. In pull-down menu select WEKA → Clusterers, and select the cluster scheme 'SimpleKMeans'. Some implementations of K-means only allow numerical values for attributes; therefore, we do not need to use a filter.




Once the clustering algorithm is chosen, right-click on the algorithm, "weka.gui.GenericObjectEditor" comes up to the screen. Set the value in "numClusters" box to 5 (instead of default 2) because you have five clusters in your .arff file. Leave the value of 'seed' as is. The seed value is used in generating a random number, which is used for making the initial assignment of instances to clusters. Note that, in general, K-means is quite sensitive to how clusters are initially assigned. Thus, it is often necessary to try different values and evaluate the results.




## 5.2. Setting Test Options


Before you run the clustering algorithm, you need to choose 'Cluster mode'. Click on 'Classes to cluster evaluation' radio-button in 'Cluster mode' box and select 'marital\_status' in the pull-down box below. It means that you will compare how well the chosen clusters match up with a pre-assigned class ('marital\_status') in the data.



Once the options have been specified, you can run the clustering algorithm. Click on the 'Start' button to execute the algorithm.



When training set is complete, the 'Cluster' output area on the right panel of 'Cluster' window is filled with text describing the results of training and testing. A new entry appears in the 'Result list' box on the left of the result. These behave just like their classification counterparts.



## 5.3. Analyzing Results

=== Run information ===

Scheme: weka.clusterers.SimpleKMeans -N 5 -S 10  
 Relation: customers  
 Instances: 9  
 Attributes: 5  
     income  
     age  
     children  
     education  
 Ignored: marital\_status  
 Test mode: Classes to clusters evaluation on training data

=== Clustering model (full training set) ===  
 kMeans  
 =====

Number of iterations: 4  
 Within cluster sum of squared errors: 3.449558299853908

Cluster centroids:

| Cluster   | Mean/Mode | Std Devs | Income | Age             | Children | Education |
|-----------|-----------|----------|--------|-----------------|----------|-----------|
| Cluster 0 | 22500     | 30       | 3      | high_school     | N/A      |           |
| Cluster 1 | 145000    | 37.5     | 0      | graduate_school | N/A      |           |
| Cluster 2 | 85000     | 55       | 0      | college         | N/A      |           |
| Cluster 3 | 15000     | 25       | 1      | high_school     | N/A      |           |
| Cluster 4 | 25000     | 30       | 0      | high_school     | N/A      |           |

=== Evaluation on training set ===

kMeans  
 =====

Number of iterations: 4  
 Within cluster sum of squared errors: 6.899116599707816

Cluster centroids:

| Cluster   | Mean/Mode | Std Devs | Income | Age             | Children | Education |
|-----------|-----------|----------|--------|-----------------|----------|-----------|
| Cluster 0 | 22500     | 30       | 3      | high_school     | N/A      |           |
| Cluster 1 | 145000    | 37.5     | 0      | graduate_school | N/A      |           |
| Cluster 2 | 85000     | 55       | 0      | college         | N/A      |           |
| Cluster 3 | 15000     | 25       | 1      | high_school     | N/A      |           |
| Cluster 4 | 25000     | 30       | 0      | high_school     | N/A      |           |

Clustered Instances

|   |         |
|---|---------|
| 0 | 2 (22%) |
| 1 | 2 (22%) |
| 2 | 2 (22%) |
| 3 | 1 (11%) |
| 4 | 2 (22%) |

Class attribute: marital\_status  
 Classes to Clusters:

| 0 | 1 | 2 | 3 | 4 | <-- assigned to cluster |
|---|---|---|---|---|-------------------------|
| 1 | 0 | 0 | 0 | 1 | single                  |
| 0 | 2 | 1 | 1 | 0 | married                 |
| 1 | 0 | 1 | 0 | 1 | divorced                |

Cluster 0 <-- No class  
 Cluster 1 <-- married  
 Cluster 2 <-- divorced  
 Cluster 3 <-- No class  
 Cluster 4 <-- single

Incorrectly clustered instances : 5.0 55.5556 %

'Run Information' gives you the following information:

- the clustering scheme used: SimpleKMeans with 5 clusters
- the relation name "customers"
- number of instances in the relation – 9
- number of attributes in the relation – 6
- list of attributes used in clustering
- the ignored cluster 'marital\_status' is an attribute the clustering is performed on.

The clustering model shows the centroid of each cluster and statistics on the number and percentage of instances assigned to different clusters. Cluster centroids are the mean vectors for each cluster; so, each dimension value and the centroid represents the mean value for that dimension in the cluster.

Thus, centroids can be used to characterize the clusters. WEKA generated clusters are:

Cluster 0 shows that this is a segment of cases representing 25 and 35 year old, either single or divorced, people with income \$22,500 in average, who have 3 children.

In cluster 1 there are 30 and 45 year old married people who do not have children.

In cluster 2 there are 50 and 60 year old married and divorced people with higher income college degree and no children.

Cluster 3 represents 25 year old married people with one child lower income and high school degree.

Cluster 4 represents 20 and 40 year old single and divorced people with lower income, high school degree and no children.

Sum of errors within the clusters is recalculated.

'Cluster Instances' section shows the number of instances in each new cluster.

For example, cluster 3 has 1 instance: people of age 25 who have one child.


Cluster 4 has 2 instances: people of age 30 in average (including 20 and 40 y.o.), whose average income is \$25,000, with high school education and no children.

'Classes to Clusters' represents class ('marital-status') assigned to clusters.


The last line displays the you have 5 number incorrectly classified instances, which is 55.5 %.

## 5.4. Visualization of Results

Another way of representation of results of clustering is through visualization. Right-click on the entry in the 'Result list' and select 'Visualize cluster assignments' in the pull-down window.



This brings up the 'Weka Clusterer Visualize' window.




On the 'Weka Clusterer Visualize' window, beneath the X-axis selector there is a drop-down list, 'Colour', for choosing the color scheme. This allows you to choose the color of points based on the attribute selected. Below the plot area, there is a legend that describes what values the colors correspond to. In your example, seven different colors represent seven numbers (number of children). For better visibility you should change the color of label '3'. Left-click on '3' in the 'Class colour' box and select lighter color from the color palette.

To the right of the plot area there are series of horizontal strips. Each strip represents an attribute, and the dots within it show the distribution values of the attribute. You can choose what axes are used in the main graph by clicking on these strips (left-click changes X-axis, right-

click changes Y-axis). Set X - axis to 'Cluster' attribute, Y - axis to 'Age'. Select 'Children' as the color dimension. You can see the result in a visual rendering of the relationship within each cluster. For instance, you can note that 'cluster 0' represents a group of people of age 25 and 35, who have 3 children, 'cluster 1' represents a group of people of age 30 and 45 who do not have children, 'cluster 2' represents 50 and 60 year old people with no children, 'cluster 3' represents 25 year old married people with one child, and 'cluster 4' represents 20 and 40 year old people without children.

The initially correctly clustered instances are represented by crosses, incorrectly clustered once represented as squares. By changing the color dimension to other attributes, you can see their distribution within each of the clusters.

You may want to save the resulting data set, which included each instance along with its assigned cluster. To do so, click 'Save' button in the visualization window and save the result as the file "customers\_kmeans.arff".



As you can see, there is a new attribute appeared in the file – 'cluster' that was added by WEKA. This attribute represents the clustering done by WEKA.

```

relation customers_clustered

@attribute Instance_number numeric
@attribute income numeric
@attribute age numeric
@attribute children {0,1,2,3,4,5}
@attribute marital_status {single,married,divorced}
@attribute education {high_school,college,graduate_school}
@attribute Cluster {cluster0,cluster1,cluster2,cluster3,cluster4}

@data
0,25000,35,3,single,high_school,cluster0
1,15000,25,1,married,high_school,cluster3
2,20000,40,0,single,high_school,cluster4
3,30000,20,0,divorced,high_school,cluster4
4,20000,25,3,divorced,college,cluster0
5,70000,60,0,married,college,cluster2
6,90000,30,0,married,graduate_school,cluster1
7,200000,45,5,married,graduate_school,cluster1
8,100000,50,2,divorced,college,cluster2
  
```

### **Clustering Exercise**


Use *k-means* algorithm to bank data from the “bank.arff” file. Perform initial preprocessing and create a version of the initial data set in which the ID field should be removed and the “children” attribute should be converted to categorical data.

## **6. Finding Associations**


WEKA contains an implementation of the Apriori algorithm for learning association rules. This is the only currently available scheme for learning associations in WEKA. It works only with discrete data and will identify statistical dependencies between groups of attributes, milk, peanut butter and bread, jelly, beer and diapers, with confidence 40% and support 30%. Apriori can compute all rules that have a given minimum support and exceed a given confidence.

### **6.1. Choosing Association Scheme**

Click ‘Associate’ tab at the top of ‘WEKA Explorer’ window. It brings up interface for the Apriori algorithm.




The association rule scheme cannot handle numeric values; therefore, for this exercise you will use grocery store data from the “grocery.arff” file where all values are nominal. Go back to ‘Preprocessing’ section described in part 4 and open “grocery.arff” file.




## 6.2. Setting Test Options


Check the text field in the 'Associator' box at the top of the window. As you can see, there are no other associators to choose and no extra options for testing the learning scheme.



Right-click on the 'Associator' box, 'GenericObjectEditor' appears on your screen. In the dialog box, change the value in 'minMetric' to 0.4 for confidence = 40%. Make sure that the default value of rules is set to 100. The upper bound for minimum support 'upperBoundMinSupport' should be set to 1.0 (100%) and 'lowerBoundMinSupport' to 0.1. Apriori in WEKA starts with the upper bound support and incrementally decreases support (by delta increments, which by default is set to 0.05 or 5%). The algorithm halts when either the specified number of rules is generated, or the lower bound for minimum support is reached. The 'significanceLevel' testing option is only applicable in the case of confidence and is (-1.0) by default (not used).



Once the options have been specified, you can run Apriori algorithm. Click on the 'Start' button to execute the algorithm.



## 6.3. Analyzing Results

```
=== Run information ===

Scheme:   weka.associations.Apriori -N 10 -T 0 -C 0.4 -D 0.05 -U 1.0 -M 0.1 -S -
1.0 -A false -c -1
Relation: grocery_store
Instances: 5
Attributes: 5
    bread
    jelly
    peanut_butter
    milk
    beer

=== Associator model (full training set) ===

Apriori
=====

Minimum support: 0.3
Minimum metric <confidence>: 0.4
Number of cycles performed: 14

Generated sets of large itemsets:

Size of set of large itemsets L(1): 5
Size of set of large itemsets L(2): 7
Size of set of large itemsets L(3): 2

Best rules found:

1. peanut_butter=yes 3 ==> bread=yes 3   conf:(1)
2. jelly=yes 1 ==> bread=yes 1   conf:(1)
3. jelly=yes 1 ==> peanut_butter=yes 1   conf:(1)
4. jelly=yes peanut_butter=yes 1 ==> bread=yes 1   conf:(1)
5. bread=yes jelly=yes 1 ==> peanut_butter=yes 1   conf:(1)
6. jelly=yes 1 ==> bread=yes peanut_butter=yes 1   conf:(1)
7. peanut_butter=yes milk=yes 1 ==> bread=yes 1   conf:(1)
8. bread=yes milk=yes 1 ==> peanut_butter=yes 1   conf:(1)
9. bread=yes 4 ==> peanut_butter=yes 3   conf:(0.75)
10. milk=yes 2 ==> bread=yes 1   conf:(0.5)
```

Run Information gives you the following information:

- the scheme for learning association we used - Apriori
- the relation name – “grocery\_store”
- number of instances in the relation – 5
- number of attributes in the relation – 4 and the list of attributes

The results for Apriori algorithm are the following:

First, the program generated the sets of large itemsets found for each support size considered. In this case five item sets of three items were found to have the required minimum support.

By default, Apriori tries to generate ten rules. It begins with a minimum support of 100% of the data items and decreases this in steps of 5% until there are at least ten rules with the required minimum confidence, or until the support has reached a lower bound of 10% whichever occurs first. The minimum confidence is set 0.4 (40%). As you can see, the minimum support decreased to 0.3 (30%), before the required number of rules can be generated. Generation of the required number of rules involved a total of 14 iterations.

The last part gives the association rules that are found. The number preceding = => symbol indicates the rule's support, that is, the number of items covered by its premise. Following the rule is the number of those items for which the rule's consequent holds as well. In the parentheses there is a confidence of the rule.


### **Association Rules Exercise**

*Use Apriori algorithm to generate association rules for Iris data from the “iris.arff” file. Perform initial preprocessing and create a version of the initial data set in which the numeric attributes should be converted to categorical data.*

## 7. Attribute Selection

Attribute selection searches through all possible combinations of attributes in the data and finds which subset of attributes works best for prediction [1]. Attribute selection methods contain two parts: a search method such as best-first, forward selection, random, exhaustive, genetic algorithm, ranking, and an evaluation method such as correlation-based, wrapper, information gain, chi-squared. Attribute selection mechanism is very flexible - WEKA allows (almost) arbitrary combinations of the two methods [4].

For this exercise you will use weather data from the “weather.arff” file. To begin an attribute selection, click ‘Select attributes’ tab.



## 7.1. Selecting Options


To search through all possible combinations of attributes in the data and find which subset of attributes works best for prediction, make sure that you set up attribute evaluator to 'CfsSubsetEval' and a search method to 'BestFirst'. The evaluator will determine what method to use to assign a worth to each subset of attributes. The search method will determine what style of search to perform.

The options that you can set for selection in the 'Attribute Selection Mode' box are [2]:

1. **Use full training set.** The worth of the attribute subset is determined using the full set of training data.
2. **Cross-validation.** The worth of the attribute subset is determined by a process of cross-validation. The 'Fold' and 'Seed' fields set the number of folds to use and the random seed used when shuffling the data.

Specify which attribute to treat as the class in the drop-down box below the test options.

Once all the test options are set, you can start the attribute selection process by clicking on 'Start' button.



When it is finished, the results of selection are shown on the right part of the window and entry is added to the 'Result list'.

## 7.2. Analyzing Results

=== Run information ===

Evaluator: weka.attributeSelection.CfsSubsetEval  
Search: weka.attributeSelection.BestFirst -D 1 -N 5  
Relation: weather  
Instances: 14  
Attributes: 5  
    outlook  
    temperature  
    humidity  
    windy  
    play  
Evaluation mode: evaluate on all training data

=== Attribute Selection on all input data ===

Search Method:  
    Best first.  
    Start set: no attributes  
    Search direction: forward  
    Stale search after 5 node expansions  
    Total number of subsets evaluated: 11  
    Merit of best subset found: 0.196  
  
Attribute Subset Evaluator (supervised, Class (nominal): 5  
play):  
    CFS Subset Evaluator  
  
Selected attributes: 1 : 1  
    outlook

Run Information gives you the following information:

- the evaluator we used – CfsSubsetEval
- the search method - BestFit
- the relation name – “weather”
- number of instances in the relation – 14
- number of attributes in the relation – 5 and the list of attributes


The search method selected is the Best Fit. The software started search with no attributes, and it is forward search. We evaluated 11 subsets and the merit of the best subset is 0.196.

The attribute evaluator used is CFS Subset Evaluator. We used supervised learning with labels in the attribute 'play'.

The selected attribute for prediction is 'outlook'.

## 7.3. Visualizing Results

Right-click on the entry in the 'Result list'. From the pull-down menu select 'Visualize reduced data'.




In the window below you can see a prediction for 'play' depending on the 'outlook'. For better visibility the color of label 'yes' was changed to the lighter one and 'Jitter' was slid to the right to see concentration points.

In the WEKA visualization window, beneath the X-axis selector there is a drop-down list, 'Colour', for choosing the color scheme. This allows you to choose the color of points based on the attribute selected. Below the plot area, there is a legend that describes what values the colors correspond to. In your example, red represents 'no', while blue represents 'yes'. For better visibility you should change the color of label 'yes'. Left-click on 'yes' in the 'Class colour' box and select lighter color from the color palette.

To the right of the plot area there are series of horizontal strips. Each strip represents an attribute, and the dots within it show the distribution values of the attribute. You can choose what axes are used in the main graph by clicking on these strips (left-click changes X-axis, right-click changes Y-axis).

Change X - axis to 'Outlook' attribute and Y - axis to 'Play'. The instances are spread out in the plot area and concentration points are not visible. Keep sliding 'Jitter', a random displacement given to all points in the plot, to the right, until you can spot concentration points.




The prediction is as follows: if the 'outlook' is sunny, play = 'yes', and if the 'outlook' is 'rainy', play = 'no', which is very likely to happen. There are few instances displayed in the window that

may or may not happen: if 'outlook' = 'sunny', 'play' = 'no' and if 'outlook' = 'rainy', 'play' = 'yes'. Note, in this section there are no correctly or incorrectly classified symbols in the graph because the result is based on probability.


## 8. Data Visualization

WEKA's visualization allows you to visualize a 2-D plot of the current working relation. Visualization is very useful in practice, it helps to determine difficulty of the learning problem. WEKA can visualize single attributes (1-d) and pairs of attributes (2-d), rotate 3-d visualizations (Xgobi-style). WEKA has "Jitter" option to deal with nominal attributes and to detect "hidden" data points [4].


To open Visualization screen, click 'Visualize' tab.



Select a square that corresponds to the attributes you would like to visualize. For example, let's choose 'outlook' for X – axis and 'play' for Y – axis. Click anywhere inside the square that corresponds to 'play' on the left and 'outlook' at the top.



A 'Visualizing weather' window appears on the screen.




## 8.1. Changing the View

In the visualization window, beneath the X-axis selector there is a drop-down list, 'Colour', for choosing the color scheme. This allows you to choose the color of points based on the attribute selected. Below the plot area, there is a legend that describes what values the colors correspond to. In your example, red represents 'no', while blue represents 'yes'. For better visibility you should change the color of label 'yes'. Left-click on 'yes' in the 'Class colour' box and select lighter color from the color palette.

To the right of the plot area there are series of horizontal strips. Each strip represents an attribute, and the dots within it show the distribution values of the attribute. You can choose what axes are used in the main graph by clicking on these strips (left-click changes X-axis, right-click changes Y-axis).

The software sets X - axis to 'Outlook' attribute and Y - axis to 'Play'. The instances are spread out in the plot area and concentration points are not visible. Keep sliding 'Jitter', a random displacement given to all points in the plot, to the right, until you can spot concentration points.


The results are shown below. But on this screen we changed 'Colour' to temperature. Besides 'outlook' and 'play', this allows you to see the 'temperature' corresponding to the 'outlook'. It will affect your result because if you see 'outlook' = 'sunny' and 'play' = 'no' to explain the result, you need to see the 'temperature' – if it is too hot, you do not want to play. Change 'Colour' to 'windy', you can see that if it is windy, you do not want to play as well.




## 8.2. Selecting Instances

Sometimes it is helpful to select a subset of the data using visualization tool. A special case is the 'UserClassifier', which lets you to build your own classifier by interactively selecting instances. Below the Y – axis there is a drop-down list that allows you to choose a selection method. A group of points on the graph can be selected in four ways [2]:


1. **Select Instance.** Click on an individual data point. It brings up a window listing attributes of the point. If more than one point will appear at the same location, more than one set of attributes will be shown.




2. **Rectangle.** You can create a rectangle by dragging it around the points.



3. **Polygon.** You can select several points by building a free-form polygon. Left-click on the graph to add vertices to the polygon and right-click to complete it.



4. **Polyline.** To distinguish the points on one side from the once on another, you can build a polyline. Left-click on the graph to add vertices to the polyline and right-click to finish.



Once the area has been selected it is colored gray. You can click on 'Submit' button to remove the points outside the gray area. To erase selected (gray) area without affecting the graph, click on 'Clear' button. When you clicked on 'Submit' button, it changes to 'Reset' button. By clicking on 'Reset' button, you can undo all changes and restore the original graph. To save all your currently visible instances to ARFF file, click on 'Save' button.

## 9. Conclusion

This concludes WEKA Explorer Tutorial. You have covered a lot of material since the Tutorial Introduction. There is a lot more to learn about WEKA than what you have covered in these seven exercises. But you have already learned enough to be able to analyze your data using preprocessing, classification, clustering, and association rule tools. You have learned how to visualize the result and select attributes. This knowledge will prove invaluable to you. If you plan to do any complicated data analysis, which require software flexibility, I recommend you to use WEKA's 'Simple CLI' interface. So, are you ready yet? Probably not. You have few new tools, but practice makes perfect. Good luck with your data analysis.

## 10. References

1. Witten, E. Frank, Data Mining, Practical Machine Learning Tools and Techniques with Java Implementation, Morgan Kaufmann Publishers, 2000.
2. R. Kirkby, WEKA Explorer User Guide for version 3-3-4, University of Weikato, 2002.
3. Weka Machine Learning Project, <http://www.cs.waikato.ac.nz/~ml/index.html>.
4. E.Frank, Machine Learning With WEKA, University of Waikato, New Zealand.
5. B. Mobasher, Data Preparation and Mining with WEKA, [http://maya.cs.depaul.edu/~classes/ect584/WEKA/association\\_rules.html](http://maya.cs.depaul.edu/~classes/ect584/WEKA/association_rules.html), DePaul University, 2003.
6. M. H. Dunham, Data Mining, Introductory and Advanced Topics, Prentice Hall, 2002.