DATA WAREHOUSING AND DATA MINING

UNIT-I

Data warehousing and OLAP: Data Warehouse basic concepts, Multidimensional Data Model- Stars, Snowflakes and Fact Constellations Schemas, OLAP Operations, A three- Tier Data Warehouse Architecture, Types of OLAP Servers: ROLAP versus MOLAP versus HOLAP and Data Warehouse Implementation.

UNIT-II

Data Mining and its Applications: Introduction- Basic Data Mining tasks, Data mining versus Knowledge Discovery in Databases, Data Mining Issues, Data Mining Metrics. Data Pre- processing, Data cleaning, data integration, data reduction and data transformation.

UNIT-III

Association Analysis: Basic Concepts and Algorithms: Frequent Item set Generation, Rule Generation, Compact Representation of Frequent Item sets, Alternative methods for generating Frequent Item sets, FP Growth Algorithm, Evaluation of Association Patterns.

UNIT-IV

Classification: Methods, improving accuracy of classification: Basics, General approach to solve classification problem, Decision Trees, Rule Based Classifiers, Nearest Neighbour Classifiers. Bayesian Classifiers, Estimating Predictive accuracy of classification methods, Improving accuracy of classification methods, Evaluation criteria for classification methods, Multiclass Problem.

UNIT-V

Clustering Techniques: Overview: Features of cluster analysis, Types of Data and Computing Distance, Types of Cluster Analysis Methods, Partitional Methods, Hierarchical Methods, Density Based Methods, Quality and Validity of Cluster Analysis.

TEXT BOOKS:

1. Jiawei Han and Micheline Kamber: Data Mining - Concepts and Techniques, 2nd Edition, Morgan Kaufmann Publisher, 2006.
2. Margaret H. Dunham: Data Mining- Introductory and Advance Topics, Pearson Edition, 2006.
3. Pang-Ning Tan, Michael Steinbach, Vipin Kumar: Introduction to Data Mining, Addison- Wesley, 2005.

REFERENCE BOOKS:

4. Arun K Pujari: Data Mining Techniques University Press, 2nd Edition, 2009.
5. G. K. Gupta: Introduction to Data Mining with Case Studies, 3rd Edition, PHI, New Delhi, 2009.
6. 6. Alex Berson and Stephen J.Smith: Data Warehousing, Data Mining, and OLAP Computing McGrawHill Publisher, 1997.

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.

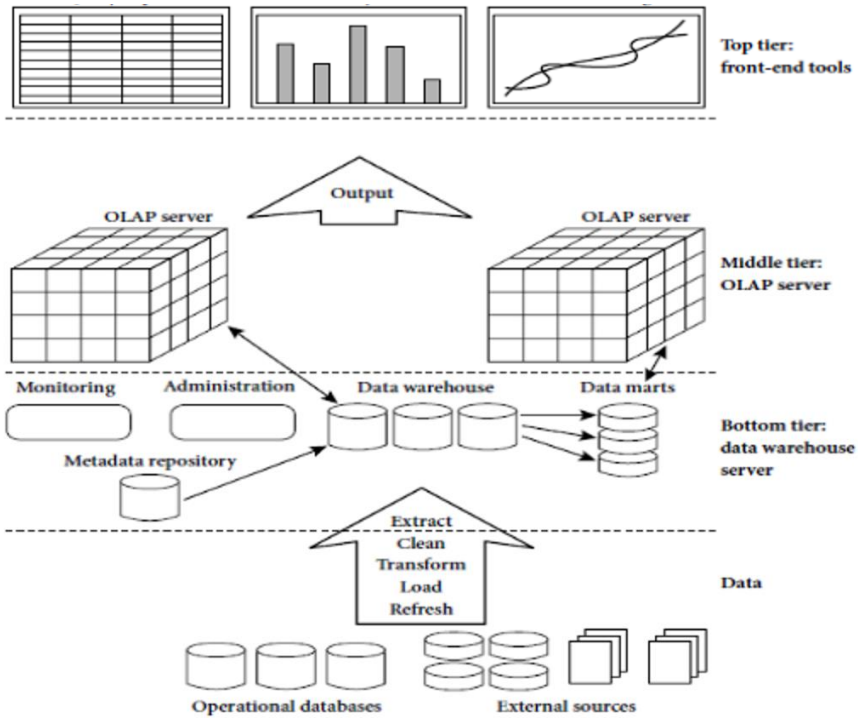
The warehouse design process consists of the following steps:

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:

Query/report Data mining



Middle tier:

OLAP server

Bottom tier:

data warehouse

Data Operational databases

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. e OLAP model is an extended relational DBMS thatmaps operations on multidimensional data to standard relational operations. e 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 a 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. e 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.

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.

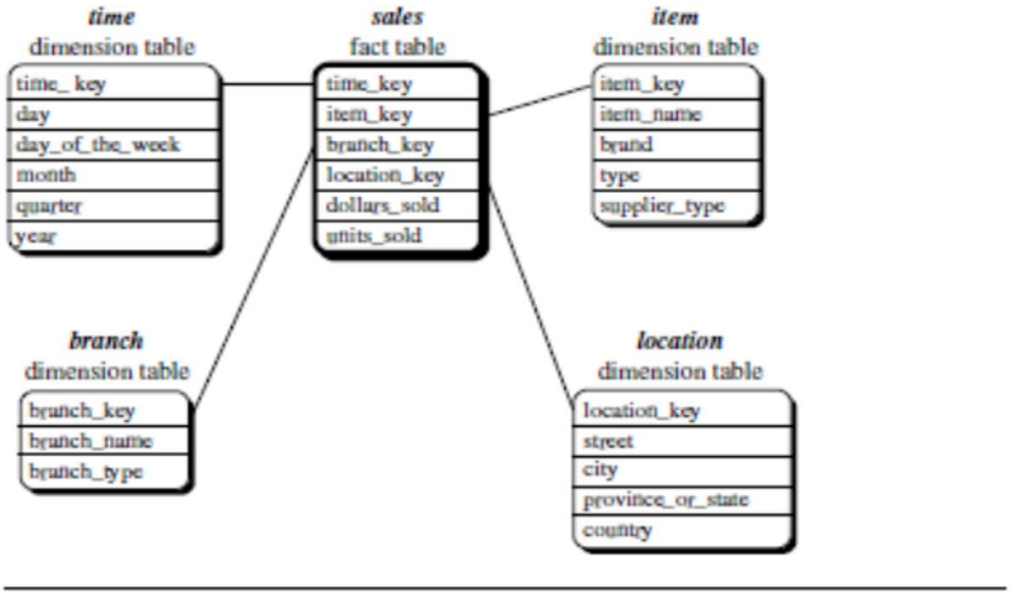
Star schema:

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, (..., Wancouver, 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.

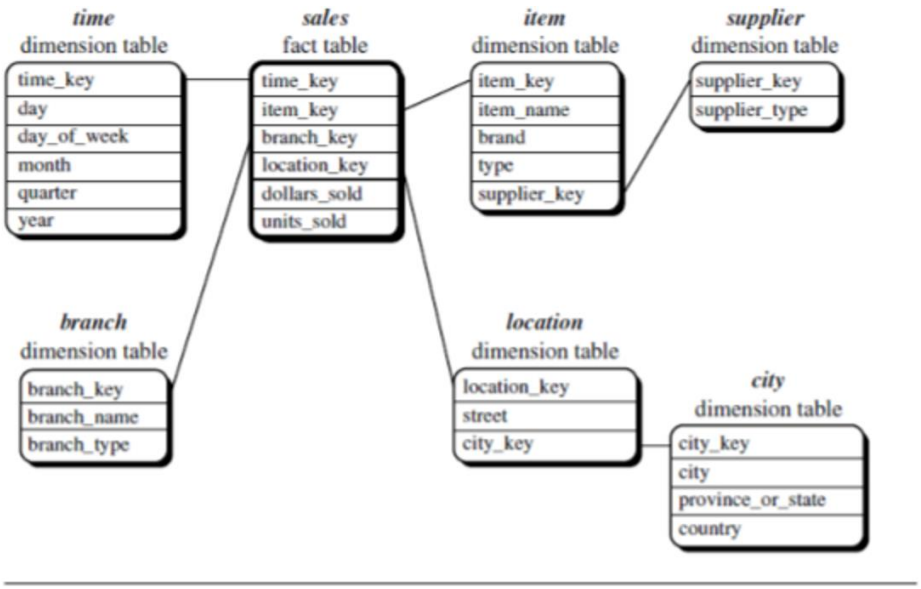
Snowflake schema.:

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

sales item supplier



dimension table dimension table

time_key item_key branch_key location_key

location dimension table

dimension table

city dimension table

city_key

+ Snowflake schema of a data warehouse for sales.

Fact constellation.

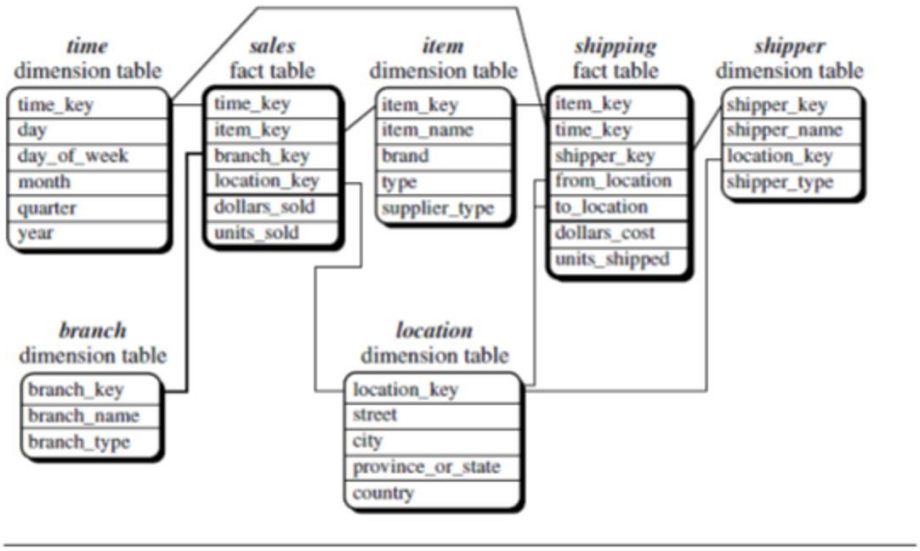
A fact constellation schema is shown in Figure. This schema specifies two fact tables, sales and shipping. The sales table definition is identical to that of the star schema . The shipping table has five dimensions, or keys: item key, time key, shipper key, from location, and to location, and two measures: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



sales item

shipping shipper

time

fact table dimension table

dimension table

fact table dimension table

dimension table

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.1 Note that a multidimensional point in the data cube space can be defined by a set of dimension-value pairs, for example, htime = “Q1”, location = “Vancouver”,item = “computer”i. 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):

e OLAP is an approach to answering multi-dimensional analytical (MDA) queries swiftly. e OLAP is part of the broader category of business intelligence, which also encompasses relational database, report writing and data mining. e 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.

OLAP 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-II

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 emphasis 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".

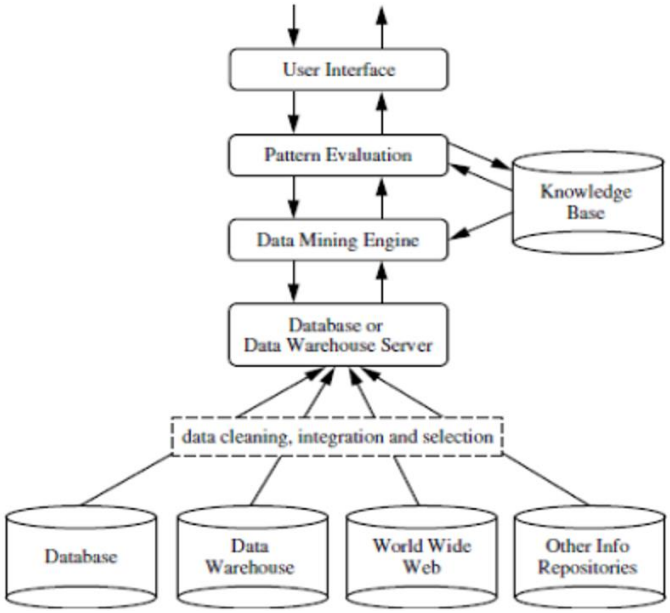
Regression — attempts to find a function which models the data with the least error.

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:

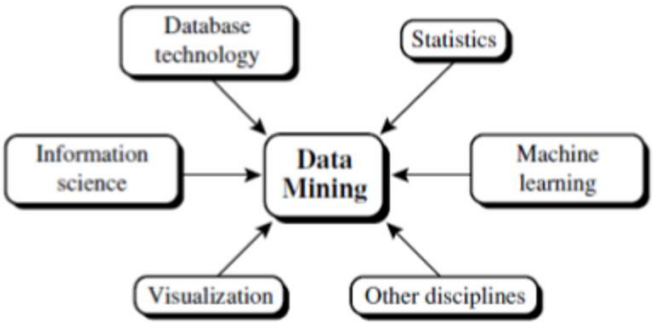
This module communicates between users and the data mining system,allowing the user to

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 aWorldWideWeb 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 providesmultiple 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 highlevel of abstraction),primitive-level knowledge (at a rawdata level), or knowledge atmultiple 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 representation 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 amounts 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 in parallel. Then the results from the partitions are merged. The incremental algorithms, updates the databases without having to mine the data again from scratch.

Data Integration:

It combines data from multiple sources into a coherent data store, as in data warehousing. These

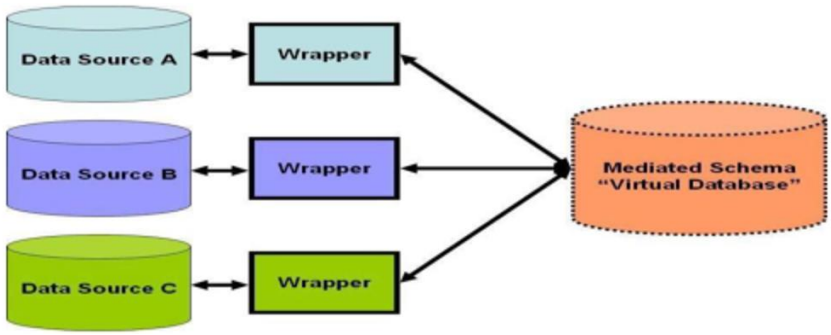
sources may include multiple databases, data cubes, or flat files. The data integration systems are formally defined as triple<G,S,M>

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 datavalue 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 —primitivel (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 datamining, 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 andrecording 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.

1. 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.

2. Estimate the model

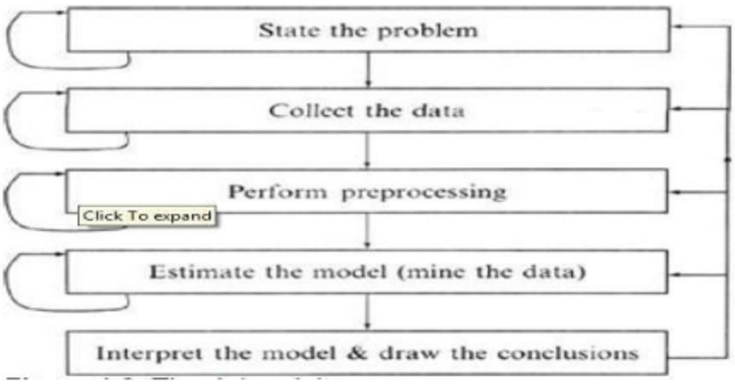
The selection and implementation of the appropriate data-mining technique is the main task in this phase. This process is not straightforward; usually, in practice, the implementation is based on several models, and selecting the best one is an additional task.

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:

e 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. e Based on the concept of strong rules, RakeshAgrawal et al. introduced association rules.

Problem Definition:

The problem of association rule mining is defined as:

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

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

tet J = {t1,t2,.--; Inde aset of Mbinary attributes called items. Let D = {t1,to,..-, tm }oe a set of transactions called the database. Each transaction in Dhas a unique transaction ID and contains a subset of the items in J.

A rule is defined as an implication of the form X => Y

where X,¥ € Tana X NY = Oh

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

The sets of items (for short itemsets) Avand 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 1 = {milk, bread, butter, beer}ang a small database containing the items (1

items is $I = \{\text{milk}, \text{bread}, \text{butter}, \text{beer}\}$ 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 {butter, bread} > {milk } meaning that if

An example rule for the supermarket could be $\{\text{butter}, \text{bread}\} \Rightarrow \{\text{milk}\}$ 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

Transaction ID | milk | bread | butter | beer

1 |i]! 0 | 0 2 | 9 | 9 1 | 0 o |!

4	1	1	1	0
5	0	1	0	0

Important concepts of Association Rule Mining:

The support supp(X)ofan itemset Xis defined as the proportion of transactions in the data set which contain the itemset. In the example database, the itemset

{milk, bread, butter} has a support of 1/5 = 0.2since it occurs in 20% of all

transactions (1 out of 5 transactions). The confidenceof a rule is defined

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

$$\text{conf}(X \Rightarrow Y) = \text{supp}(X \cup Y)/\text{supp}(X).$$

For example, the rule $\{\text{butter}, \text{bread}\} \Rightarrow \{\text{milk}\}$ has a confidence of

conf(X = Y) = supp(X U Y)/supp(X), For example, the rule {butter, bread} > {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

under the condition that these transactions also contain the LHS'

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

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

The Jiftof a rule is defined as) . supp(X UY

(supp(X) x supp(Y)

lift(x + 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 $\frac{0.2}{0.4 \times 0.4} = 1.25$.

0.2

rule {milk, bread} = {butter}pas a tit of04x04 125 The conviction of a rule is defined as

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

x Y)= 1 — supp(Y) conv(X > Y) = 7 conf(X > VY). 1-04 _

1.2

The rule {milk, bread} > {butter}nas aconviction of 1 —.5 ~

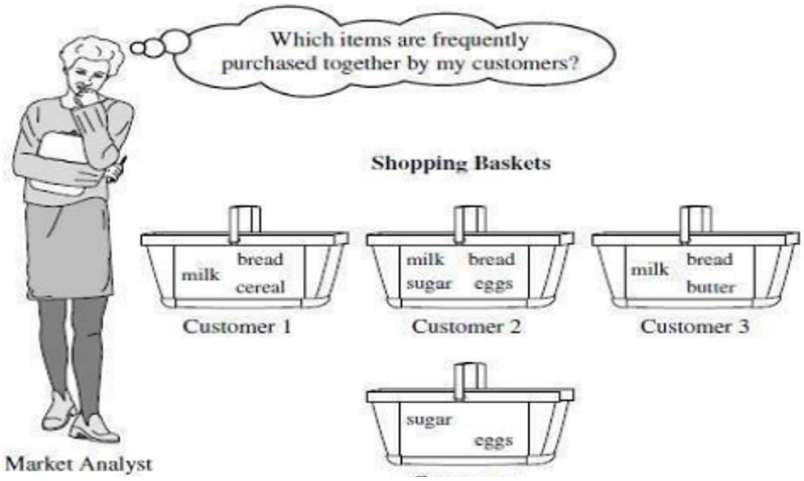
and can be interpreted as the ratio of the expected frequency that X occurs without Y (that is to say, the frequency that the rule makes an incorrect prediction) if X and Y were

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 shelf space.



Which items are frequently purchased together by my customers?

Shopping Baskets

milk bread bread milk

Sugar eggs | buuer |

Customer 2 Customer 3

Market Analyst ————— te Customer n

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 andsoftware at opposite ends of the store may entice customers who purchase such items topick 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 displayto 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:

e Wecan mine the complete set of frequent itemsets, the closed frequent itemsets, and the maximal frequent itemsets, given a minimum support threshold. e Wecan 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:

buys(X, —computer!))=>buys(X, —HP printerl) (1)

buys(X, —laptop computer|)) =>buys(X, —HP printer|) (2)

In rule (1) and (2), the items bought are referenced at different levels ofabstraction (e.g.,

—*computer*l is a higher-level abstraction of —*laptop computer*l).

—computerl is a higher-level abstraction of —laptop computer|).

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

e Ifthe items or attributes in an association rule reference only one dimension, then it is a single-dimensional association rule.

buys(X, —computerl))=>buys(X, —antivirus softwarel)

e Ifarule 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 exampleof a

multidimensional rule:

age(X, —30,31...391) * income(X, —42K,...48K1))=>buys(X, —high resolution TVI)

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

Ifa rule involves associations between the presence or absence of items, it is a Boolean association rule. Ifa 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 /evel-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 L1.Next, L1 is used to find L2, the set of frequent 2-itemsets, which is used to find £3, and so on, until no more frequent k-itemsets can be found.
- The finding of each Lirequires one full scan of the database.
- A two-step process is followed in Apriori consisting of join and prune action.
- Algorithm: Apriori. Find frequent itenmpets using an iterative level-wise approach: based on candidate generation.

Input:

- *D*, a database of transactions;
- *min_sup*, the minimum support count threshold.

Method:

(1) Ly — find_frequent_1-itemsets(D);

(2) for (k —2:a a A o:k++) {

```
(2)    for (k = 2; L_{k-1} ≠ φ; k++) {
(3)        C_k = apriori_gen(L_{k-1});
(4)        for each transaction t ∈ D { // scan D for counts
(5)            C_t = subset(C_k, t); // get the subsets of t that are candidates
(6)            for each candidate c ∈ C_t
(7)                c.count++;
(8)        }
(9)        L_k = {c ∈ C_k | c.count ≥ min_sup}
(10)   }
(11)   return L = ∪_k L_k;

procedure apriori_gen(L_{k-1}:frequent (k-1)-itemsets)
(1)    for each itemset l_1 ∈ L_{k-1}
(2)        for each itemset l_2 ∈ L_{k-1}
(3)            if (l_1[1] = l_2[1]) ∧ (l_1[2] = l_2[2]) ∧ ... ∧ (l_1[k-2] = l_2[k-2]) ∧ (l_1[k-1] < l_2[k-1]) then {
(4)                c = l_1 ⋈ l_2; // join step: generate candidates
(5)                if has_infrequent_subset(c, L_{k-1}) then
```

() ; =apriort gen(l, 1)s

(4) tor each transaction € D { // scan D tor counts (5) C, — subset(C,, 1); // get the subsets of ¢ that are candidates (6) for each candidate © € CG

@) ©.count+—+s

(8)

(9) Ly = {e CQ le.count > minsup}

ao) }

(11) return L — UyL,?

procedure apriori_gen(Z,_;:frequent (A — 1)-itemsets)

a) for each tremset fy HS 4

(~ for each itemset f9 € 1,4

(3) f(A MAUNA =h2YA..AK ZB=hk AACA 1 <b[k 1) then { (4) c= & Hb; // juin step: generate Candidates

(&) if hasinfroquont_subsot(c, 4, ;) then

(6) delete c; // prune step: remove untruitful candidate

(7) else add c to Ci;

(8) }

```
(9)    return C_k;

procedure has_infrequent_subset(c: candidate k-itemset;
          L_{k-1}: frequent (k-1)-itemsets); // use prior knowledge
```

(9) return (;;

procedure lias_infreyguentsubset(c: candidate A-itemmsets 1, 1° frequent (k — 1)-tremsets); // nse prior knowledge

(1) for each (k — |)-subset s of ¢ (2) ifs ZLy_ then (3) return TRUE;

(4) return FALSE:

Example:

TID | List of item IDs |

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 candidatel- itemsets, C1. 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 sup = 2. The set of frequent 1-itemsets, L1, can then be determined. It consists of the candidate 1-itemsets satisfying minimum support. In our example, all of the candidates in C1 satisfy minimum support.

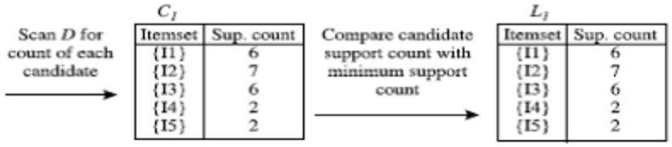
3. To discover the set of frequent 2-itemsets, L2, the algorithm uses the join L1 on LI to generate a candidate set of 2-itemsets, C2.No candidates are removed fromC2 during the prune step because each subset of the candidates is also frequent.

4. Next, the transactions in D are scanned and the support count of each candidate itemsetInC2 is accumulated.
5. The set of frequent 2-itemsets, L2, is then determined, consisting of those candidate2- itemsets in C2 having minimum support.

6. The generation of the set of candidate 3-itemsets,C3, From the join step, we first getC3 =L2x L2 = ({I1, 12,13), {11, 12, 15}, (11, 13, 15}, (12, 13, 14}, {12, 13, 15}, (12, 14, 15}. 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 L3, consisting of those candidate 3-itemsets in C3 having minimum support.

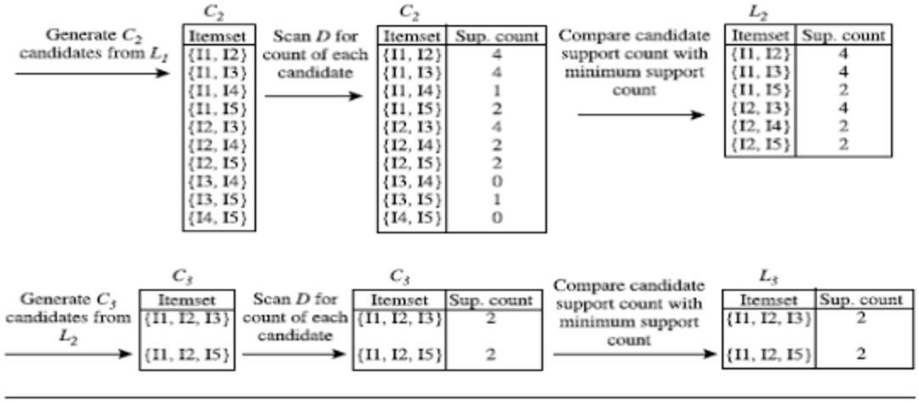
8. The algorithm uses L3x L3 to generate a candidate set of 4-itemsets, C4.



Scan Dfor — [Ttemset] Sup. count | Compare candidate _ [Itemset] Sup. count |

count of cach 6 support count with 6 candidate 7 minimum support 7

2 2



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

(a) Join: C3 = Lz x Lz = {{11, 12}, {11, 13}, {11, 15}, {12, 13}, {12, 14}, {12, 15}} » {{11, 12}, {11,13}, {11, 15}, {12, 13}, {12, 14}, {12, 15}}

= {{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?

= {{I1, 12, 13}, {11, 12, 15}, {11, 13, 15}, {12, 13, 14}, {12, 13, 15}, {12, 14, 15} }.

(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 {11, 12, 13} are {11, 12}, {11,13}, and {12, 13}. All 2-item subsets of {11, 12, 13} are members of La. Therefore, keep {11, 12, 13} in C3.

The 2-item subsets of {11, 12, 15} are {11, 12}, {I1, 15}, and {12, 15}. All 2-item subsets of {11, 12, 15} are members of La. Therefore, keep {11, 12, 15} in C3.

The 2-item subsets of {11, 13, 15} are {11,13}, {11,15}, and {13, 15}. {13, 15} is not a member of L2,

- 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.

and so it is not frequent. Therefore, remove {11, 13, 15} from Cs.

The 2-item subsets of {12, 13, 14} are {12,13}, {I2, 14}, and {13, 14}. {13, 14} is not a member of L2, and so it is not frequent. Therefore, remove {12, 13, 14} from C3.

The 2-item subsets of {12, 13, 15} are {12,13}, {12,15}, and {13, 15}. {13, 15} is not a member of L2, and so it is not frequent. Therefore, remove {12, 13, 15} from C3.

The 2-item subsets of {12, 14, 15} are {12, 14}, {12, 15}, and {14, 15}. {14, 15} is not a member of Lz, and so it is not frequent. Therefore, remove {12, 14, 15} from C3.

(c) Therefore, C; = {{I1, 12, 13}, {I1, 12, 15}} after pruning.

Generation and pruning of candidate 3-itemsets, Cz, from L2 using the Apriori property.

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

We re-examinethe mining of transaction database, D, of Table 5.1 in Example 5.3 using the frequentpattern growth approach.

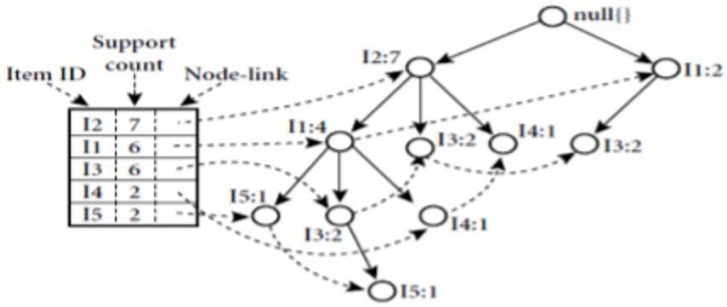
The first scan of the database is the same as Apriori, which derives the set of frequent items (1-itemsets) and their support counts (frequencies). Let the minimum support count be 2. The set of frequent items is sorted in the order of descending support count. This resulting set or /ist is denoted L.

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, 12, 15,” which contains three items (12, I1, IS in Z order), leads to the construction of the first branch of the tree with three nodes, hI2: 1i, hI1:1i, and hIS: 1i, where 12 is linked as a child of the root, I1 is linked to I2, and I5 is linked to Il. The second transaction, T200, contains the items I2 and I4 in L order, which would result in a branch where 12 is linked to the root and I4 is linked to 12. However, this branch would share a common prefix, 12, with the existing path for T100. Therefore,we instead increment the count of the I2 node by 1, and create a newnode, hI4:

1i,which is linked as a child of hI2: 2i. In general,when considering the branch to be added for a transaction, the

li,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
15	{{12, 11: 1}, {12, 11, 13: 1}}	⟨12: 2, 11: 2⟩	{12, 15: 2}, {11, 15: 2}, {12, 11, 15: 2}

Item Conditional Pattern Base Conditional FP-tree Frequent Patterns Generated

1S {{12, Ut: bp, (12, 11, 13: 1} } (12: 2, 11: 2) {12, 1S: 2}, {11, 1S: 2}, {12, 11, 15: 2} 14 {{12, 11: 2}, {12s 1} } (12: 2) {12, 14: 2}

13) {{12, 11: 2}, {12:2}, {0s 2h} (N2s 4, Ube 2), Cs 2) {12, 13: 4}, (11, 13: 4}, {12, 11, 13: 2} 11 {{12: 4}} (12: 4) {12, 11: 4}

The FP-tree is mined as follows.

Start fromeach 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 15, which is the last itemin 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. IS 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, IS: li and h12, I1, 13, IS: 1i.

Therefore, considering IS as a suffix, its corresponding two prefix paths are hI2, I1: li and h12, I1, 13: li, which form its conditional pattern base. Its conditional FP-tree contains only a single path, hI2: 2, 11: 2i; I3 is not included

because its support count of | is less than the minimum support count. The single path generates all the combinations of frequent patterns: f12, 15: 2g, f11, 15: 2g, f12, 11, IS: 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.

support_count(A U B)

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

confidence(A => B) = P(B\|A) = mneuesrouitA)

The conditional probability is expressed in terms of itemset support count, where support_count(A UB) is the number of transactions containing the itemsets A UB, and 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 /, generate all nonempty subsets of /.

For every nonempty subset s of /, output the rule “s = (/—s)” if ae >

- For every nonempty subset s of l, output the rule “$s \Rightarrow (l - s)$” if $\frac{support_count(l)}{support_count(s)} \geq$

min_conf, where 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 arecalled 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 | 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.

TID Items Purchased

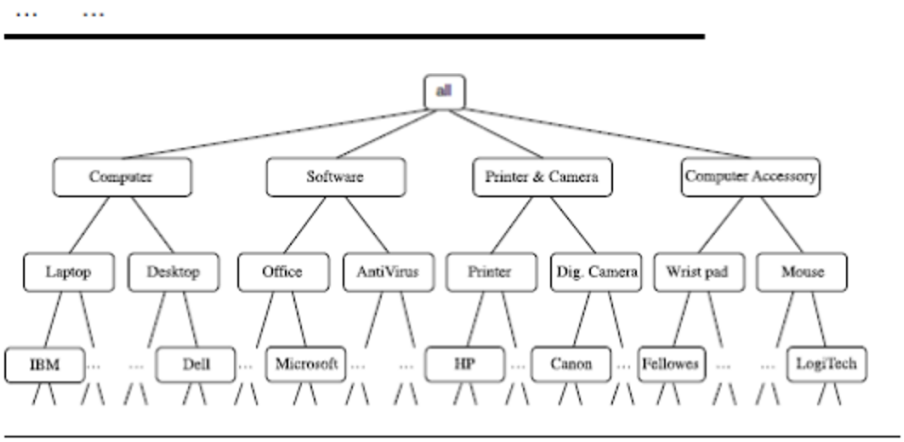
T100 IBM-ThinkPad-T40/2373, HP-Photosmart-7660

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 Aconcept hierarchy for AllElectronicscomputeritems for AllElectronics computer items.

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

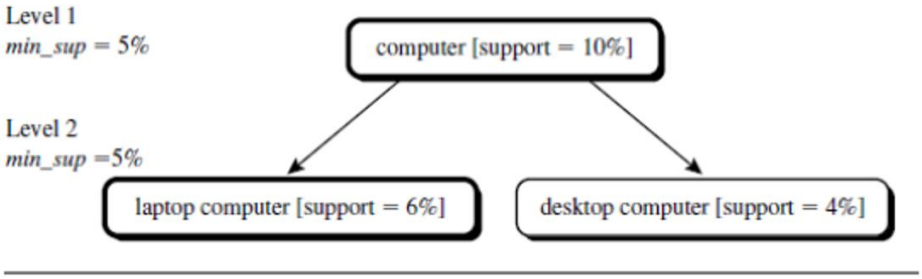
Here,

- 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.
- lf 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.



Level 1 min_sup = 5% computer [support = 10%] Level 2

min_sup =5% laptop computer [support = 6%] desktop computer [support = 4%]

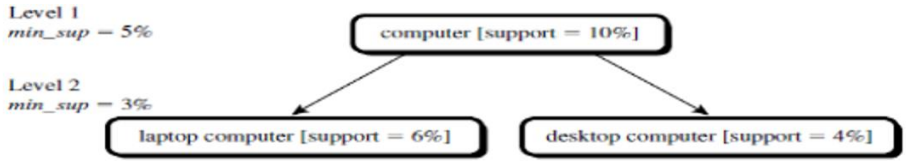
Multilevel mining with uniform support.

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 %,respectively.

In this way, —computer,| —laptop computer,| and —desktop computer| are all considered

In this way, —computer,! —laptop computer, and —desktop computer! are all considered frequent.



Level 1

min_sup — 5% computer [support — 10%] Level 2

min_sup — 3%

laptop computer [support = 6%] desktop computer [support = 4%]

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

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. buys(X, — digital cameral)=>buys(X, —HP printer)
- Association rules that involve two or more dimensions or predicates can be referred to as multidimensional association rules.

  age(X, “20...29")°occupation(X, “student”)=>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 arecalled inter dimensional 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!)*buys(X, —laptopl)=>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 Aquan| “Aquan2 =>Acat Where Aquan! and Aquan2 are tests on quantitative attribute interval Acat 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...391)income(X, —42K...48K1)=>buys(X, —HDTVI|)

From Association Mining to Correlation Analysis:

- Acorrelation measure can be used to augment the support-confidence framework for

association rules. This leads to correlation rules of the form A=>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(AUB) = 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 item sets. The lift between the occurrence of A and B can be measured by computing lift(A, B) = PR

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

  - If the lift(A,B) is less than 1, then the occurrence of A is negatively correlated with

- If the lift(A,B) is less than 1, then the occurrence of A is negatively correlated with the occurrence of B. e 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 missingvalues (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 Al 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. e 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. e 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.

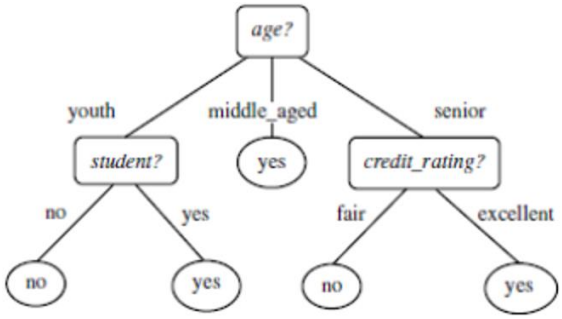
> Interpretability:

- This refers to the level of understanding and insight that is providedby 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 I 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

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 application

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 application

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 attributetist is empty then «s) return as a leaf node labeled with the majority class in D;m// majority voting

(6) apply Attribute_selection_methoed(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_Jist «— attributeJlist splitting_attribute; // remove splitting_attribute

(10) for each outcome / of splitting_criterion // partition the tuples and grow subtrees for each partition

(11) let D, be the set of data tuples in D satisfying outcome jj; // a 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;
```

(12) if D; 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,, attributc_list) to node N; endfor

(15) 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 giventuples 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, {a1, a2, ... ,av}, based on the training data.

1 A is discrete-valued:

e In this case, the outcomes of the test at node N correspond directly to the known values of A. e A branch is created for each known value, aj, of A and labeled with that value. e 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

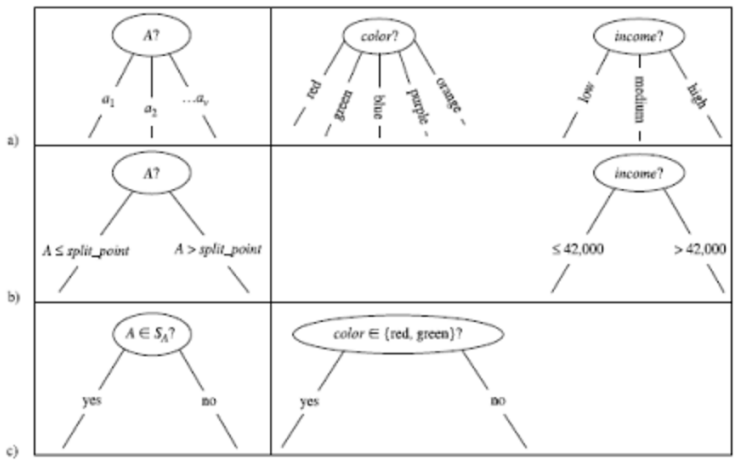
A <=split point and A >split point, respectively where split point is the split-point

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€SA?I. SA 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.

Partiboning Scenarios Examples



(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.land 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 Hholds given the —evidencel 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(X|H)P(H)
    - P(H|K) = P(X)

Naive Bayesian Classifier:

The naive 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 = (xl, x2, ...,xn), depicting n

measurements made on the tuple from n attributes, respectively, Al, A2, ..., An.

2. Suppose that there are m classes, Cl, C2, ..., Cm. 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 naive Bayesian classifier predicts that tuple X belongs to the class Ci if and only if

P(C\X) > P(C|X) for 1< j<m.j Fi.

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

Thus we maximize P(CijX). The class Ci for which P(CijX) is maximized is called the maximum posteriori hypothesis. By Bayes' theorem

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

P(X|Ci)P(Ci)

P(C|X) = P(X)

3. As P(X) is constant for all classes, only P(X|Ci)P(Ci) 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(C1) = P(C2) = ...= P(Cm), and we would therefore maximize P(X|Ci). Otherwise, we maximize P(X|Ci)P(Ci).

4.Given data sets with many attributes, it would be extremely computationally expensive to compute P(X|Ci). In order to reduce computation in evaluating P(X|Ci), 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,

$$P(X|C_i) = \prod^{n} P(x_k|C_i)$$

T] Pala)

P(X|G)

kel

$$= P(x_1|C_i) \times P(x_2|C_i) \times \cdots \times P(x_n|C_i).$$

P(xiGi) x P(x2|Ci) x «++ x Plxn|Ci)-

We can easily estimate the probabilities P(x1|Ci), P(x2|Ci), : : : , P(an|Ci) from the training tuples. For each attribute, we look at whether the attribute is categorical or continuous-

valued. For instance, to compute P(X|Ci), we consider the following:

- If Axis categorical, then P(x:|Ci) is the number of tuples of class Ci in D having the value x«for Ax, divided by |Ci,p| the number of tuples of class Ciin D.
- If A;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 4 and standard deviation , defined by

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

- In order to predict the class label of X, P(XjCi)P(Ci) is evaluated for each class Ci.
- The classifier predicts that the class label of tuple X is the class Ci if and only if

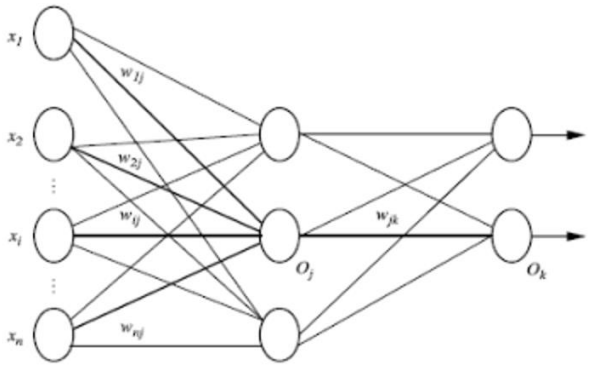
$$P(X|C_i)P(C_i) > P(X|C_j)P(C_j) \quad \text{for } 1 \le j \le m, j \ne i.$$

A Multilayer Feed-Forward Neural Network:

e@ The back propagation algorithm performs learning on a multilayer feed- forward neural network. e 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:

Input Hidden Output layer layer layer



- 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 hddkn 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 include 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, Oj, is equal to its input value, /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.

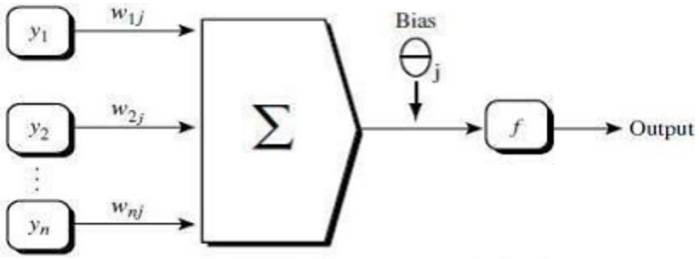
$$I_j = \sum_i w_{ij} O_i + \theta_j,$$

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

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.

Weights



Inputs Weighted Activation (outputs from sum function previous layer)

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 jis computed by

Err; = 0j(1—0j)(T;—0))

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

Where Qj is the actual output of unit j, and Tjis the known target value of the given training tuple. The error of a hidden layer unit 7 is

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

Err; =Oj(1— 0) > Errewjx.

Where wixis the weight of the connection from unit j to a unit k in the next higher layer, and Erryis the error of unit k.

Weights are updated by the following equations,

where Dwi j is the change in weight wij:

Awij = (DErrjO;

wij = wij + Aw;;

Biases are updated by the following equations below

AO; = (D)Err; 6; =6;+A0;

Algorithm: k-Nearest-Neighbor Classifier:

Input:

DP, a data set consisting of the training tuples and their associated target values:

1, 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 ¥ 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 = Σ_i w_ij O_i + θ_j; //compute the net input of unit j with respect to the
                previous layer, i
(9)         O_j = 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

(15)    for each weight w_ij in network {
(16)        Δw_ij = (l)Err_j O_i; // weight increment
(17)        w_ij = w_ij + Δw_ij; } // weight update
```

Nearest-neighbor classifiers are based on learning by analogy, that is, by comparing a

piven 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 & 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, $X_1 = (x_{11}, x_{12}, \ldots, x_{1n})$ and $X_2 = (x_{21}, x_{22}, \ldots, x_{2n})$, is

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

- The Euclidean distance between two points or tuples, say, X) = (211, X12, ... , Xin) and X27 = (x21, X22, -.. X2n),
- In other words, for each numeric attribute, we take the difference between the corresponding values of that attribute in tuple Xiand in tuple X2, 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 vo in the range [0, 1] by computing v= v—ming max, — ming Where minaand maxaare the minimum and maximum values of attribute A
- For k-nearest-neighbor classification, the unknown tuple is assigned the mostcommon 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 tuplewith 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, $\boldsymbol{X}1 = (x11, x12, : : : , x1n)$ and $\boldsymbol{X}2 = (x21, x22, : : : , x2n)$, is

The Euclidean distance between two points or tuples, say, X1 = (x11, x12, : 2:2, xn) and X2 = (x21, .x22,:::, x2n), is

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

In other words, for each numeric attribute, we take the difference between the corresponding values of that attribute in tuple X1 and in tuple X2, 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 transforma value v of a numeric attribute A to vO in the range [0, 1] by computing P v—ming

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

y¥ =—____+_ max, — min,'

where minA and maxA 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 X1 with that in tuple X2. If the two are identical (e.g., tuples X1 and X2 both have the color blue), then the difference between the two is taken as 0.

If the two are different (e.g., tuple X1 is blue but tuple X2 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, ifthe value of a given attribute A is missing in tuple X1 and/or in tuple X2, 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 | if either one or both of the corresponding values of A are missing. If A is numeric and missing fromboth tuples X1 and X2, then the difference is also taken to be 1. If only one value is missing and the other (which we’ll call vO) is present and normalized, then we can take the difference to be either jl || vOj or jO\!v0j

other (which we’ll call v0) is present and normalized, then we can take the difference to be either j1 □v0j or j0□v0j

(i.e., 1□v0 or v0), whichever is greater. “How can I determine a good value for k, the number of neighbors?” This

(i.e., 1-1 vO or vO), 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 then 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 frompoor 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 (c ity 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 = w0+w Ix:

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 jDj data points of the form(x1, y1), (x2, y2),:::, (xjDj, yjDj).12 The

variable, y. The training set contains jDj data points of the form(x1, yl), (x2, y2).::, (xjDj, yjDj).12 The

regression coefficients can be estimated using this method with the following equations:

\

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

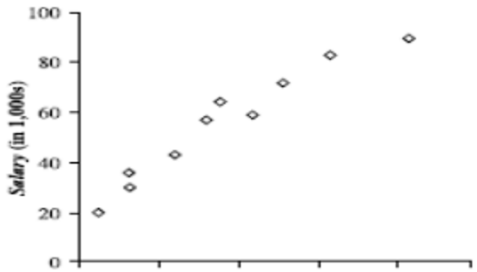
12Note that earlier, we had used the notation (Xi, yi) to refer to training tuple i having associated class label yi, where Xi was an attribute (or feature) vector (that is, Xi was described by more than one attribute). Here, however, we are dealing with just one predictor variable. Since the Xi here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

w0 = y-w1x where x is the mean value of x1, x2, ::: , xjDj, and y is the mean value of yl, y2,::: , yjDj. The coefficients

w0 and w! often provide good approximations to otherwise complicated regression equations.

Salary data.

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



x years experience

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	buys_computer = yes	buys_computer = no	Total	Recognition (%)
buys_computer = yes	6,954	46	7,000	99.34
buys_computer = no	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

> 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 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 Acc(M), where Acc(M) is the accuracy ofM. 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 (X1, yl), (X2.y2), : :: , (Xd, yd), where the Xi are the n-dimensional test tuples with associated known values, yi, 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, y0 i, for Xi is correct. Instead of focusing on whether y0 i is an "exact" match with yi, we instead look at how far off the predicted value is from the actual known value. Loss functions measure the error between yi and the

predicted value, y0 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 onEnsemble 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), M1, M2, :::, Mk, 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 Mi is learned, the weights are updated to allow the subsequent classifier,Mi+1, to "pay more attention" to the training tuples that were misclassified by Mi. 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 apopular 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, (X1, yl), (X2, y2), : : :, (Xd, yd),where yi is the class label of tupleXi. Initially, Adaboost assigns each training tuple an equal weight of 1=d. Generating k classifiers for the

of tupleXi. 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 forma

ensemble requires k rounds through the rest of the algorithm. In round i, the tuples from D are sampled to forma training set,Di, 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, Mi, is derived from the training tuples ofDi. Its error is then calculated using Di as a test set. The weights of the training tuples are then adjusted according to how they were classified. Ifa 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.

UNIT-V

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 ona 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 <= 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 o fthe 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 icroclusters 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

++ 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 —noisel 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:

e 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.

e Clustering high-dimensional data is challenging due to the curse of dimensionality.

e 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.

e 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.

e 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:

e It is a clustering approach that performs clustering by incorporation of user-specified or application-oriented constraints.

e 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.

e = 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

"* 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.

e First, it randomly selects k of the objects, each of which initially represents a cluster mean or center.

e 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.

e It then computes the new mean for each cluster.

e This process iterates until the criterion function converges.

Typically, the square-error criterion is used, defined as k 2

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

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

E=> YD p—mil, f=IpeC; 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

Miis the mean of cluster Cj

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 & clusters.

Method:

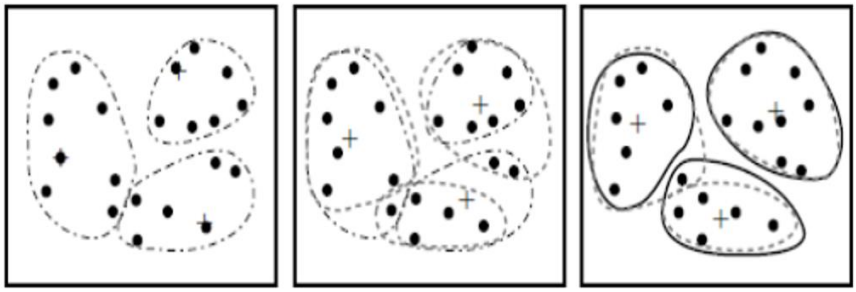
(1) arbitrarily choose k objects from D as the initial cluster centers;

(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:

e 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.

e 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.

e 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

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

E= y ¥ p—al,

jal pec; where£ is the sum of the absolute error for all objects in the data set pis the point inspace representing a given object in clusterC;

ojis the representative object of Cj

e 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.

e This quality is estimated using a cost function that measures the average dissimilarity between an object and the representative object of its cluster.

e To determine whether a non representative object, oj random, is a good replacement for a current representative object, oj, 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_{random}asa representative object and p is closest to one of the other representative objects, o_i,i≠j, then p is reassigned to o_i.

'p' currently belongs to representative object, 0}. If ojis replaced by O;sandomasa representative object and p is closest to one of the other representative objects, 0i,i¢j, then p is reassigned to 0).

Case 2:

'p' currently belongs to representative object, oj. If ojis replaced by Orandomasa representative object and p is closest tO Oyandom, then p is reassigned tO Oyandom-

Case 3:

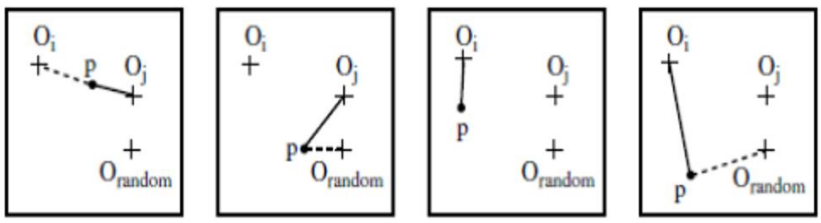
'p' currently belongs to representative object, o_i, i≠j. If ojis replaced by o_{random}as a representative object and p is still closest to o_i, then the assignment does not change.

'p' currently belongs to representative object, oj, i¢j. If ojis replaced by OyandomaS a representative object and p is still closest to 0;, then the assignment does not change.

Case 4:

'p' currently belongs to representative object, oi, i≠j. If $o_j$is replaced byo_{random}as a representative object and p is closest to o_{random}, then p is reassigned

'p' currently belongs to representative object, oi, &j. If ojis replaced byO;andomas a representative object and p is closest tO O;andom, then p is reassigned lOOsandom.



1. ReassignedtoO; — 2. Reassigned to 3. No change 4, Reassigned to random Orandom

O_{random} O_{random}

• data object

+ cluster center — before swapping +--+ after swapping

Four cases of the cost function for k-medoids clustering

Thek-MedoidsAlgorithm:

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. e 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 orconstraints. Depending on the nature of the constraints, constraint-based clusteringmay 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:

- Aworking 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,

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

The hypothesis is retained if there is no statistically significant evidence supporting its rejection. A discordancy test verifies whether an object, oi, 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 oi is vi, then the distribution of T is constructed. Significance probability, SP(vi)=Prob(T > vi), is evaluated. If SP(vi) is sufficiently small, then oi is discordant and the working hypothesis is rejected.

An alternative hypothesis, H, which states that oj comes from another distribution model, G, is adopted. The result is very much dependent on which model F is chosen because Oi 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 0i 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 :oi € G, where i= 1, 2,...,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 H:0;€(1—-A)F+AG, wherei=1,2,...,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, 0, 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, pet, of the objects in D lie at a distance greater than dmin from o.

In other words, rather that 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, 0, 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, onceM+1 neighbors of object 0 are found, it is clear that 9 is not an outlier.

This algorithm has a worst-case complexity of O(n2k), 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 complexityas 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*+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 $\frac{dmin}{2\sqrt{k}}$.

dmin

IQ dts-nigthod, the data space is partitioned into cells with a side length equal to 2%" 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 dmin-neighborhood of an

outlier.

An object, 0, 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, 0, in the cell, objects in the second layer of 0 are examined. For objects in the cell, only those objects having no more than M points in their dmin-neighborhoods are outliers. The dmin-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, 13 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-distance(p). It is defined as the distance, d(p, 0), between p and an object o 2 D, such that for at least k objects, 002 D, it holds that d(p,

o·)_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, o00 2 D, it holds that d(p;o'') <d(p, o).

o:)_d(p, 0). That is, there are at least k objects in D that are as close as or closer to p than 0,

and for at most k-1 objects, 000 2 D, it holds that d(p:o"') <d(p, 0).

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 $N_{kdistance(p)}(p)$, or $N_k(p)$for short. By

The k-distance neighborhood of an object p is denoted Nixdistancepyp), or Nx(p)for short. By setting k to MinPts, we get Nminpis(p). It contains the MinPts- nearestneighbors 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

distMinPts(p, 0) = max{ MinPtsdistance(o), d(p, 0)}.

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 0), then the actual distance is replaced by the MinPts- distance of 0. This helps to significantly reduce the statistical fluctuations of d(p, 0) for all of the p close to o.

The higher thevalue of MinPts is, the more similar is the reachability distance for objects withinthe 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

|Nutinpes(P) |

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

(rd yiness(p) = : :

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-basedmeasures to identify exceptional objects. Instead, it identifies outliers by examining themain characteristics of objects in a group.Objects that —deviatel fromthisdescription areconsidered outliers. Hence, in this approach the term deviations is typically used to referto 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 OLAPdata 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,{D1, D2,

...,Dm}, of these objects with 2<=m <= n such that Dj_; C Dj, where Dj CD.

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

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 m numbers, {x1, ...%n}, a possible dissimilarity function is the variance of the numbers in the set, that is,

Ss (x; —x)*,

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

where 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 Dj-; does not cover any string in D; that is not in Dj.;.

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.