

Chapter 1: Introduction

What is Data Mining?

“The process of discovering meaningful patterns and trends often previously unknown by using some mathematical algorithm on huge amount of stored data”

“Extraction of interesting, non-trivial, implicit, previously unknown and potentially useful information or patterns from data in large database.”


- Data mining is basically concerned with the analysis of data and the use of software techniques for finding patterns and regularities in sets of data.

Two Approaches are:

i. Descriptive Data Mining:

- It characterizes the general properties of data in the database.
- It finds patterns in data the user determinants which ones are important.
- Mostly used during data exploration.
- Typical questions answered by descriptive data mining are:
 - . What is in the data?
 - . What doesn't look like?
 - . Are there any unusual patterns?
 - . What does the data suggest for customer segmentation?
- User may have no idea on which kind of patterns are interesting?
- Functionalities of descriptive data mining are: Clustering, Summarization, Visualization, and Association.

ii. Predictive Data Mining:



X: Vectors of independent variables.

Y: Dependent variables

$Y = f(X)$

- Users don't care about the model, they simply interested in accuracy of predictions.
- Using unknown examples the model is trained and the unknown function is learned from data.

- The more data with known outcomes is available the better is the predictive power of model.
- Used to predict outcomes whose inputs are known but the output values are not realized yet.
- Never 100% accurate.
- The performance of a model on past data is not predicting the known outcomes.
- Suitable for unknown data set.
- Typical questions answered by predictive models are:
 - . Who is likely to respond to next product?
 - . Which customers are likely to leave in the next six months?

Data Mining Process:




Fig: "Data mining process flow"

Problem Definition:

- Focuses on Understanding the project objectives and requirements in terms of business perspective.
Eg: How can I sell more of my product to customer? Which customers are most likely to purchase the product?

Data Gathering and Preparation:

- Data Collection & Exploration.
- Identify data quality, patterns in data.
- Data preparation phase covers all the tasks involved to build the model.
- Data preparation tasks are likely to be performed multiple and not in any prescribed order.

Model Building and Evaluation:

- Various modeling techniques are applied and calibrated the parameters to optimal values.

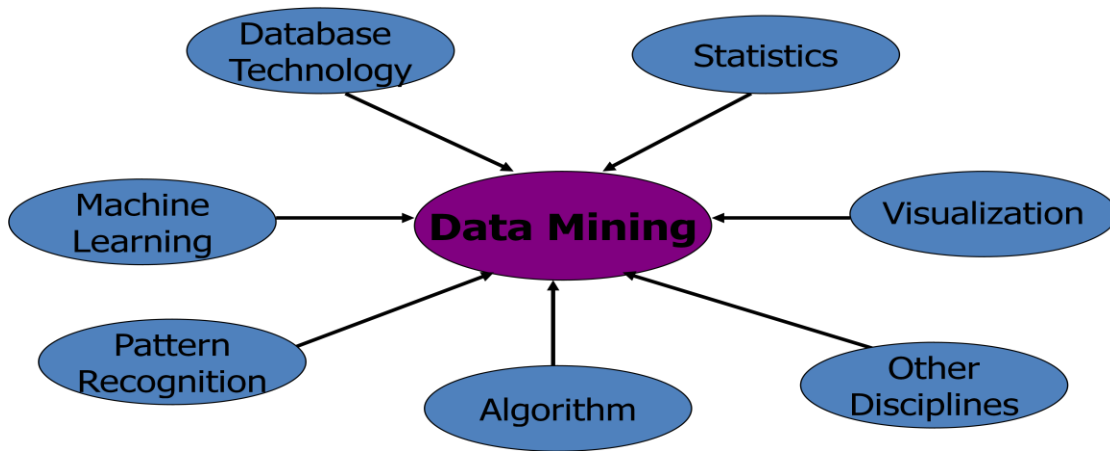
- Evaluate how well the model satisfies the originally stated business goal.

Knowledge Deployment:

- Use data mining within a target environment.
- Insight and actionable information can be derived from data.

Why Data Mining?

Data mining is a combination of multidisciplinary field. It can be applied in many fields and can be done using many algorithm and techniques.



Data Mining Vs. Query Tools

- SQL can find normal queries from the database such as what is an average turnover? Whereas data mining tools find interesting patterns and facts such as what are the important trends in sells?
- Data mining is much more faster than SQL in trend and pattern analysis since it uses algorithm like machine learning, genetic algorithm.
- If we know exactly what we are looking for, we use SQL nut if we know only vaguely what we are looking for we use data mining.
- Hybrid information can't be easily be traced using SQL.

Data Warehouse

In most of the organization, there occur large databases in operation for normal daily transactions called operational database.

A data warehouse is a large database built from the operational database.

A data warehouse should be:

- i. Time – dependent
 - There must be a connection between the information in the warehouse and the time when it was entered.
 - One of the most important aspect of the warehouse as it relates to data mining, because information can then be sourced according to period.

- ii. Non-Volatile
 - Data in a warehouse is never updated, but used only for queries.
 - End-users who want to update data must use operational database.
 - A data warehouse will always be filled with historical data.

- iii. Subject Oriented
 - Not all the information in the operational database is useful for a data warehouse.
 - A data warehouse should be designed especially for decision support and expert system with specific related data.

- iv. Integrated
 - In an operational data, many types of information being used with different names for same entity.
 - In a data warehouse, all entities should be integrated and consistent i.e. only one name must exist to describe each individual entity.

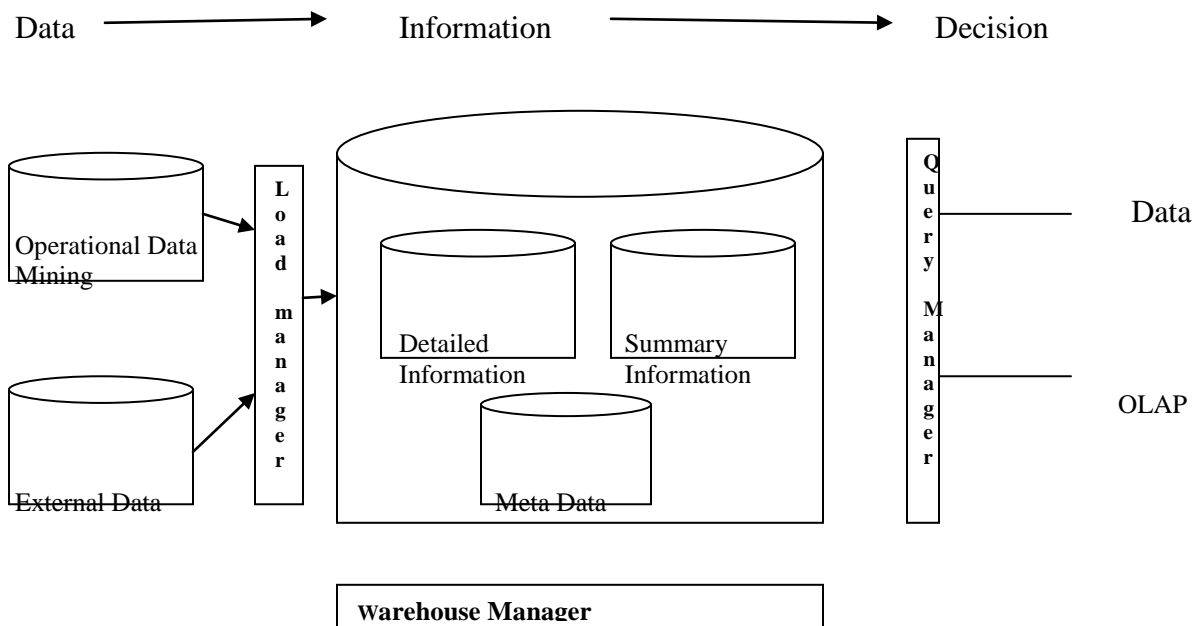


Fig: “Architecture of a Data Warehouse”

Load Manager: The system components that perform all the operations necessary to support the extract and load process. It fast loads the extracted data into a temporary data store and performs simple transformations into a structure similar to the one in the data warehouse.

Warehouse Manager: Performs all the necessary operations to support the warehouse management process. It analyzes the data to perform consistency and referential checks. It also transforms and merges the source data in the temporary data store into the published data warehouse with creating indexes and business views. Update all existing aggregations and back up data in the data warehouse.

Query Manager: Performs all the operations necessary to support the query management process by directing queries to the appropriate tables. In some cases it also stores query profiles to allow the warehouse manager to determine which indexes and aggregations are appropriate.

Detailed Information: Stores all the detailed information to determine the business requirements to analyze the level at which to retain detailed information in the data warehouse.

Summary Information: Stores all the predefined aggregations generated by the warehouse manager. It is a transient area which will change on an ongoing basis in order to respond to changing query profiles. It is essentially a replication to detailed information.

Meta Data: Meta data is data about data which describes how information is structured within a data warehouse. It maps data stores to common view of information with the data warehouse.

Data Mart

- Data Mart is a subset of the information content of a data warehouse that is stored in its own database.
- Data mart may or may not be sourced from an enterprise data warehouse i.e. it could have been directly populated from source data.
- Data mart can improve query performance simply by reducing the volume of data that needs to be scanned to satisfy the query.
- Data marts are created along functional level to reduce the likelihood of queries requiring data outside the mart.
- Data marts may help in multiple queries or tools to access data by creating their own internal database structures.
- Eg: Departmental Store, Banking System.

Chapter 2: Data Preprocessing

What is an Attribute?

- An attribute is a property or characteristic of an object. Examples: eye color of a person, temperature, etc.
- Attribute is also known as variable, field, characteristic, or feature
- A collection of attributes describe an object. Object is also known as record, point, case, sample, entity, or instance.
- Attribute values are numbers or symbols assigned to an attribute
- Same attribute can be mapped to different attribute values. Example: height can be measured in feet or meters.
- Different attributes can be mapped to the same set of values. Example: Attribute values for ID and age are integers but properties of attribute values can be different. ID has no limit but age has a maximum and minimum value.

Types of Attributes

Approach 1:

Attribute Type	Description	Examples
Nominal	The values of a nominal attribute are just different names, i.e., nominal attributes provide only	zip codes, employee ID numbers, eye color.
Ordinal	The values of an ordinal attribute provide enough information to order objects. (<, >)	hardness of minerals, {good, better, best}, grades, street numbers
Interval	For interval attributes, the differences between values are meaningful, i.e., a unit of	calendar dates, temperature in Celsius or Fahrenheit
Ratio	For ratio variables, both differences and ratios are meaningful. (*, /)	temperature in Kelvin, monetary quantities, counts, age, mass, length, electrical current

Approach 2:

Discrete Attribute

- Has only a finite or countably infinite set of values
- Examples: zip codes, counts, or the set of words in a collection of documents
- Often represented as integer variables.
- Note: binary attributes are a special case of discrete attributes

Continuous Attribute

- Has real numbers as attribute values
- Examples: temperature, height, or weight.
- Practically, real values can only be measured and represented using a finite number of digits.
- Continuous attributes are typically represented as floating-point variables.

Approach 3:

Character: values are represented in forms of character or set of characters (string).

Number: values are represented in forms of number. Number may be in form of whole number, decimal number.

Types of data sets

a. Record

- Data that consists of a collection of records, each of which consists of a fixed set of attributes

i. Data Matrix

- If data objects have the same fixed set of numeric attributes, then the data objects can be thought of as points in a multi-dimensional space, where each dimension represents a distinct attribute
- Such data set can be represented by an m by n matrix, where there are m rows, one for each object, and n columns, one for each attribute

Projection of x Load	Projection of y load	Distance	Load	Thickness
10.23	5.27	15.22	2.7	1.2
12.65	6.25	16.22	2.2	1.1

ii. Document Data

- Each document becomes a 'term' vector, each term is a component (attribute) of the vector, the value of each component is the number of times the corresponding term occurs in the document

	team	coach	play	ball	score	game	win	lost	timeout	Season
Document 1	3	0	5	0	2	6	0	2	0	2
Document 2	0	7	0	2	1	0	0	3	0	0
Document 3	0	1	0	0	1	2	2	0	3	0

iii. Transaction Data


- A special type of record data, where each record (transaction) involves a set of items.

- For example, consider a grocery store. The set of products purchased by a customer during one shopping trip constitute a transaction, while the individual products that were purchased are the items

<i>TID</i>	<i>Items</i>
1	Bread, Coke, Milk
2	Beer, Bread
3	Beer, Coke, Diaper, Milk
4	Beer, Bread, Diaper, Milk
5	Coke, Diaper, Milk

b. Graph

- Contains nodes and connecting vertices.



Eg: World Wide Web, Molecular Structures

c. Ordered

- Has Sequences of transactions
 - i. Spatial Data

Spatial data, also known as geospatial data, is information about a physical object that can be represented by numerical values in a geographic coordinate system.
 - ii. Temporal Data

A temporal data denotes the evolution of an object characteristic over a period of time. Eg $d=f(t)$.
 - iii. Sequential Data

Data arranged in sequence.

Important Characteristics of Structured Data

a. Dimensionality

- A Data Dimension is a set of data attributes pertaining to something of interest to a business. Dimensions are things like "customers", "products", "stores" and "time".

Curse of Dimensionality

- When dimensionality increases, data becomes increasingly sparse in the space that it occupies.
- Definitions of density and distance between points, which is critical for clustering and outlier detection, become less meaningful

***Purpose:**

- Avoid curse of dimensionality
- Reduce amount of time and memory required by data mining algorithms
- Allow data to be more easily visualized
- May help to eliminate irrelevant features or reduce noise

***Techniques**

- Principle Component Analysis
- Singular Value Decomposition
- Others: supervised and non-linear techniques

Dimensionality Reduction:

i. PCA

- Goal is to find a projection that captures the largest amount of variation in data.
- Find the eigenvectors of the covariance matrix.
- The eigenvectors define the new space.
- Construct a neighborhoods graph
- For each pair of points in the graph, compute the shortest path distances – geodesic distances

ii. Feature Subset Selection

- Another way to reduce dimensionality of data.
- Redundant features.
 - Duplicate much or all of the information contained in one or more other attributes.
 - Example: purchase price of a product and the amount of sales tax paid.
- Irrelevant features
 - Contain no information that is useful for the data mining task at hand
 - Example: students' ID is often irrelevant to the task of predicting students' GPA

Techniques:

- a. Brute-force approach:

- Try all possible feature subsets as input to data mining algorithm
- b. Embedded approaches:
 - Feature selection occurs naturally as part of the data mining algorithm
- c. Filter approaches:
 - Features are selected before data mining algorithm is run
- d. Wrapper approaches:
 - Use the data mining algorithm as a black box to find best subset of attributes.

Feature Creation

- Create new attributes that can capture the important information in a data set much more efficiently than the original attributes.
- Three general methodologies:
 - Feature Extraction: domain-specific
 - Mapping Data to New Space
 - Feature Construction: combining features

b. Sparsity and Density

- Sparsity and density are terms used to describe the percentage of cells in a database table that are not populated and populated, respectively. The sum of the sparsity and density should equal 100.
- Many of the cell combinations might not make sense or the data for them might be missing.
- In the relational world storage of such data is not a problem: we only keep whatever there is. If we want to keep closer to our multidimensional view of the world, we face a dilemma: either store empty space or create an index to keep track of the nonempty cells or search for an alternative solution

c. Resolution

- Scaling of data in different label and classes. Patterns depend on the scale.

Data Quality

- Real world database are highly unprotected from noise, missing and inconsistent data due to their typically huge size and their possible origin from multiple, heterogeneous sources.
- Low quality data will lead to low quality mining results.
- Data pre-processing is required to handle these above mentioned facts.
- The methods for data preprocessing are organized into
 - a. Data Cleaning
 - b. Data Integration
 - c. Data Transformation
 - d. Data Reduction
 - e. Data Discretization

Data Cleaning

- Mostly concern with
 - i. Fill-in missing values
 - ii. Identify outliers and smooth out noisy data
 - iii. Correct inconsistent data
 - iv. Eliminate duplicate data

a. Missing Data

- Data is not always available because many tuples may not have recorded values for several attributes such as age, income.
- Missing data may be due to:
 - . Equipment Malfunction
 - . Inconsistent with other recorded data and thus deleted.
 - . Data not entered due to misunderstanding
 - . Certain data may not be considered important at the time of entry.
 - . No change in recorded data.

How to Handle Missing Data?

- Ignore the tuple: usually done when class label is missing. Not effective when the percentage of missing values per attribute varies considerably.
- Fill-in missing values manually: Tedious and infeasible task.
- Use a global constant to fill-in missing values.
- Use an attribute mean fill-in missing values belonging to the same class.
- Use the most probable value to fill-in missing value.

b. Noisy Data

- Noisy data is a form of error because of random error in a measured variable.
- Incorrect attribute values may be due to:
 - . Faulty data collection instruments
 - . Data entry problem
 - . Data transmission problem
 - . Technology limitation
 - . Inconsistency in naming convention

How to Handle Noisy Data

- Clustering: Detect and remove outliers
- Regression: Smooth by fitting the data into regression function
- Binning Method: First sort the data and partition into different boundaries with mean, median values.
- Combined computer and human inspection, doing so suspicious values are detected by human

c. Outliers

- Outliers are a set of data points that are considerably dissimilar or inconsistent with the remaining data.
- In most of the cases they are inference of noise while in some cases they may actually carry valuable information.
- Outliers can occur because of:
 - . Transient malfunction of data measurement
 - . Error in data transmission or transcription

- . Changes in system behavior
- . Data contamination from outside the population examined.
- . Flaw in assumed theory

How to Handle Outliers

There are three fundamental approaches to the problem of outlier's detection

- a. Type 1: Determine the outliers with no prior knowledge of data. This is a learning approach analogous to unsupervised learning.
- b. Type 2: Model with normality and abnormality. Analogous to supervised learning.
- c. Type 3: Model with normality. Semi-supervised learning approach.

Data Integration

- Combines data from multiple sources into a coherent store.
- Integrate meta data from different sources (Schema Integration)

Problem: - Entity Identification Problem.

.Different sources have different values for same attributes.

.Data Redundancy

These problems are mainly because of different representation, different scales etc.

How to handle redundant data in data integration?

- Redundant data may be able to be detected by correlation analysis.
- Step-wise and careful integration of data from multiple sources may help to improve mining speed and quality.

Data Transformation

Changing data from one form to another form.

Approaches:

- i. Smoothing: Remove noise from data.
- ii. Aggregation: Summarizations of data
- iii. Generalization: Hierarchy climbing of data
- iv. Normalization: Scaled to fall within a small specified range.
 - a. Min-Max Normalization:

$$V' = ((V - \min) / (\max - \min)) * (\text{new_max} - \text{new_min}) + \text{new_min}$$

- b. Z-Score Normalization:

$$V' = (V - \min) / \text{stand_dev.}$$

- c. Normalization by decimal scaling:

$$V' = V / 10^j \text{ where } j \text{ is the smallest integer such that } \max(|V'|) < 1$$

Data Aggregation:

- Combining two or more attributes (or objects) into a single attribute (or object).

Purpose

Data reduction: Reduce the number of attributes or objects

Change of scale: Cities aggregated into regions, states, countries, etc

More “stable” data: Aggregated data tends to have less variability

Data Reduction:

- Warehouse may store terabytes of data hence complex data mining may take a very long time to run on complete data set.
- Data reduction is the process of obtaining a reduced representation of data set that is much smaller in volume but yet produces the same or almost same analytical results.
- Different methods such as data sampling, dimensionality reduction, data cube, aggregation, discretization and hierarchy are used for data reduction.
- Data compression can also be used mostly in media files or data.

i. Data Sampling:

- It is one of main method for data selection i.e. sampling is the main technique employed for data selection.
- It is often used for both the preliminary investigation of the data and the final data analysis.
- Statisticians sample because obtaining the entire set of data of interest is too expensive or time consuming.
- Sampling is used in data mining because processing the entire set of data of interest is too expensive or time consuming.
- Often used for both preliminary investigation of data and the final data analysis.
- Important since obtaining entire set of data of interest is too expensive or time consuming.
- Sampling should be representative since it must represent approximately the same property as the original set of data.
- Get at least one object from each of 10 groups as sample data.

Types:

- a. **Simple Random Sampling:** Equal probability of selecting any particular item.
- b. **Sampling without replacement:** As each item is selected, it is removed from population.
- c. **Sampling with replacement:** Objects are not removed from the population as they are selected from the sample. The same objects can be picked-up more than once.
- d. **Stratified Sampling:** Split the data into several partitions, then draw random samples from each partition.

ii. Dimensionality Reduction:

- Dimensionality Reduction is about converting data of very high dimensionality into data of much lower dimensionality such that each of the lower dimensions conveys much more information.
- This is typically done while solving data mining/machine learning problems to get better features for a classification or regression task.

Data Discretization:

- Convert continuous data into discrete data.

- Partition data into different classes.
- Two approaches are:
 - a. Equal width (distance) partitioning:**
 - It divides the range into N intervals of equal size.
 - If A and B are the lowest and the highest values of the attribute, the width of interval will be
 - $W = (A - B)/N$.
 - The most straight forward approach for data discretization.
 - b. Equal depth (frequency) partitioning:**
 - It divides the range into N intervals, each containing approximately same number of samples.
 - Good data scaling
 - Managing categorical attributes can be tricky.

OLAP Tool

- OLAP stands for On-Line Analytical Processing.
- An OLAP cube is a data structure that allows fast analysis of data.
- OLAP tools were developed to solve multi-dimensional data analysis which stores their data in a special multi-dimensional format (data cube) with no updating facility.
- An OLAP tool doesn't learn, it creates no new knowledge and they can't reach new solutions.
- Information of multi-dimension nature can't be easily analyzed when the table has the standard 2-D representation.
- A table with n- independent attributes can be seen as an n-dimensional space.
- It is required to explore the relationships between several dimensions and standard relational databases are not very good for this.

	Product A	Product B	Product C	Product D	Product E
Chicago	20M	2M	12M	2M	21M
Cincinnati	30M	4M	10M	8M	26M
Dallas	14M	3M	14M	9M	24M
Louisville	16M	5M	11M	4M	23M

OLAP Operations:

- i. Slicing:** A slice is a subset of multi-dimensional array corresponding to a single value for one or more members of the dimensions. Eg: Product A sales.
- ii. Dicing:** Dicing operation is the slice on more than two dimensions of data cube. (More than two consecutive slice). Eg: Product A sales in 2004.

- iii. **Drill-Down:** Drill-down is specific analytical technique where the user navigates among levels of data ranging from the most summarized to the most detailed i.e. it navigates from less detailed data to more detailed data. Eg: Product A sales in Chicago in 2004.
- iv. **Roll-Up:** Computing of all the data relationship for more than one or more dimensions i.e. summarization of data to one or more dimensions. Eg: Total Product.
- v. **Pivoting:** Pivoting is also called rotate operation. It rotates the data in order to provide an alternative presentation of data.

OLTP (Online Transaction Processing)

- Used to carry out day to day business functions such as ERP (Enterprise Resource Planning), CRM (Customer Relationship Planning)
- OLTP system solved a critical business problem of automating daily business functions and running real time report and analysis.

OLAP Vs OLTP

Facts	OLTP	OLAP
Source of Data	Operational Data	Data warehouse (From various database)
Purpose of data	Control and run fundamental business tasks	For planning, problem solving and decision support
Queries	Simple queries	Complex queries and algorithms
Processing Speed	Typically very fast	Depends on data size, techniques and algorithms
Space requirements	Can be relatively small	Larger due to aggregated databases
Database Design	Highly Normalized with many tables.	Typically denormalized with fewer tables. Use of star or snowflake schema.

Similarity and Dissimilarity

Similarity

- Numerical measure of how alike two data objects are.
- Is higher when objects are more alike.
- Often falls in the range [0,1]

Dissimilarity

- Numerical measure of how different are two data objects
- Lower when objects are more alike
- Minimum dissimilarity is often 0
- Upper limit varies

Similarity Measure Methods: