

Unsupervised Learning

Notes

- Usually learning can be done in a supervised or unsupervised manner
- There are a strong bias in the machine learning community towards supervised learning
- But a lot of concepts are learnt unsupervisedly
- The discovery of new concepts always is unsupervised

Unsupervised Learning

Notes

Goals:

- To obtain representations that describe an unlabeled dataset
- To discover the concepts inside the data
- The discovery process is biased by a context
- Representation: Unstructured (partitions) or relational (hierarchies)

Numerical Taxonomy

Notes

- Data analysis by the search of predefined structures, fitting of known probability distributions
- It is assumed that the data is embedded in a N-dimensional space that has a distance function defined (similarity/dissimilarity)
- Examples are more related to the nearest examples than to the farthest

Numerical Taxonomy - Algorithms

Notes

Two main strategies:

- Hierarchical algorithms
 - Examples are organized as a binary tree
 - No explicit division in groups
- Partitional algorithms
 - Only a partition of the dataset is obtained

Hierarchical algorithms

Notes

- Based on graph theory
 - The examples form a full connected graph
 - Similarity define the length of the edges
 - The clustering is decided using connectivity criteria
- Based on matrix algebra
 - A distance matrix is calculated from the examples
 - The clustering is computed using the distance matrix
 - The distance matrix is updated after each step (different updating criteria)

Hierarchical algorithms

Notes

- Graphs
 - Single Linkage, Complete Linkage, MST
 - Divisive, Agglomerative
- Matrices
 - Johnson algorithm
 - Different update criteria (S-L, C-L, Centroid, minimum variance)

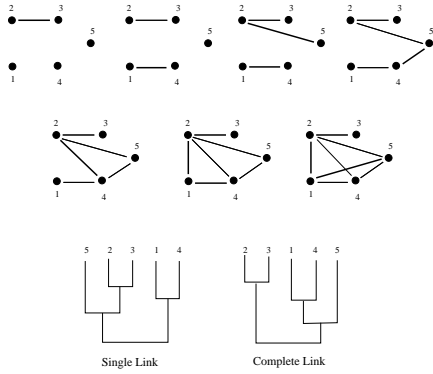
Computational cost

$$O(n_{inst}^3)$$

Hierarchical algorithms - Graphs

Notes

	2	3	4	5
1	6	8	2	7
2		1	5	3
3			10	9
4				4



Single Link

Complete Link

Hierarchical algorithms - Matrices


Notes

	2	3	4	5
1	6	8	2	7
2		1	5	3
3			10	9
4				4
	2,3	4	5	
1	7	2	7	
2,3		7.5	6	
4			4	
	1,4	5		
2,3	7.25	6		
1,4		5.5		
	1,4,5			
2,3	6.725			


Hierarchical algorithms - Example

Notes


Data




Single Link




Complete Link




Median



Centroid



Ward



Hierarchical algorithms - Problems

Notes

- A partition of the data is not given, it has to be decided a posteriori
- Some undesirable and strange behaviours could appear (chaining, inversions)
- Dendrogram is not a practical representation for large amount of data
- Its computational cost is high

Partitional algorithms

Notes

The computational cost to find the optimal partition of N objects in K groups is NP-hard

- Square error minimization (k-means)
- Mixture decomposition (Expectation-Maximization)
- Density estimation
- Graph theory: MST, Relative neighbourhood graph, Gabriel graph
- Neighbourhood relationships
- Fuzzy clustering
- Spectral Clustering
- Unsupervised Neural networks

K-means

Notes

- We assume that the shape of the clusters is hyperspherical
- An iterative algorithm assigns each example to one of K groups (K is a parameter)
- Hill Climbing search
- Optimization criteria (square error, minimize the distance of each example to the centroid of the class)

$$Distorsion = \sum_{k=1}^K \sum_{i \in C_k} \|x_i - \mu_k\|^2$$

- The algorithm converges to a local minima

K-means

Notes

Algorithm: K-means (X : Examples, k :integer)

Generate k prototypes with the k first examples

Assign the $n-k$ examples to its nearest prototype

SumD = Sum of square distances examples-prototypes

repeat

Recalculate prototypes

Reassign examples to its nearest prototype


SumI = SumD

SumD = Sum of square distances examples-prototypes

until $SumI - SumD < \epsilon$

K-means

Notes



K-means - practical problems

Notes

- The algorithm is sensitive to the initialization (to run the algorithm from random initializations could be a good idea)
- Find the value of k is not an easy problem (experimentation with different values is needed)
- You can obtain a solution even if the classes are not hyperspherical (some classes could be splitted)
- No guarantee about the quality of the solution
- The space and computational cost makes it not suitable for big datasets

Mixture Decomposition - EM algorithm

Notes

- We assume that the data are drawn from a mixture of probability distribution functions, we are looking for the parameters of the distributions that explain better the data
- **Model:** The data are from K samples from the same probability distribution function (usually Gaussian)
- Each instance has a probability to belong to a class
- The goal is to estimate the parameters of the distribution that describes each class (i.e.: means and standard deviations)
- The algorithm maximizes the likelihood of the distribution respect the dataset
- It performs iteratively two steps
 - **Expectation:** We calculate a function that assigns a degree of membership to all the instances to any of the K probability distributions
 - **Maximization:** We re-estimate the parameters of the distributions to maximize the membership

EM Algorithm (K Gaussian)

Notes

We suppose that the data are from K n-dimensional Gaussian. The computations depend on the assumptions that we make about the attributes (independents or not, same σ , ...)

- A set of K initial distributions is generated $N(\mu_i, \sigma_i)$, μ_i and σ_i are vectors corresponding to the mean and the variance of each attributes
- We repeat until convergence:
 - 1 **Expectation:** Compute the membership of each instance to each probability distribution. Usually we use the log likelihood function of the distribution
 - Each instance will have a weight depending of the probability assigned by the previous step $w_{x_j, i} = \ln(P(x_j | N(\mu_i, \sigma_i)))$
 - 2 **Maximization:** Recompute the parameters using the weights from the previous steps and obtain the new μ_i and σ_i for each distribution

EM algorithm - Comments

Notes

- K-means is a particular case of this algorithm
- The main advantage is that we obtain a membership as a probability (soft assignments)
- Using different probability distribution we can find different kinds of structures. The problem is to derive the calculations for the iterative updating of the parameters

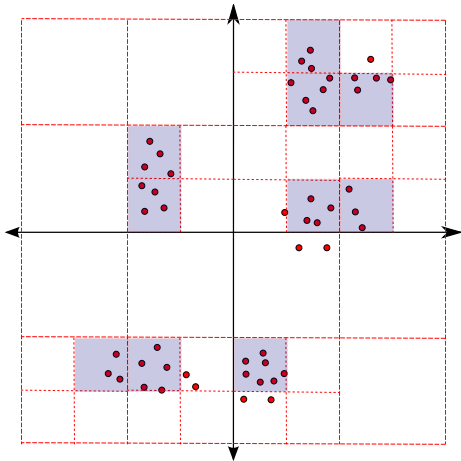
Density estimation

Notes

- The number of groups is not decided beforehand
- We are looking for regions with high density of examples
- We are not limited to predefined shapes (there is no model)
- Different approaches:
 - Space partitioning (exponential number of partitions)
 - Multidimensional histograms (we look for high density regions with less dimensions)
- Usually it is only applied to datasets with low dimensionality


Density estimation

Notes



Density estimation

Notes




Based in graph theory

Notes

- We create different kinds of graphs with the dataset (MST, Voronoi, Delanau, ...)
- We give consistency criteria for the edges of the graph (delete)
- The result is a set of unconnected components
- Two advantages: we do not need to know the number of classes, we do not look for a specific model (any shape is possible)

Based in graph theory

Notes



Incremental algorithms: Neighbourhood relationship

Notes

- The commonality among all the algorithms until this point is that they are not incremental
- Incrementality allows to update a model with new data without starting from scratch
- These algorithms use the neighbourhood relationship defined from a similarity function
- This neighbourhood determines what instances belong to the same group
- Examples: Nearest Neighbour, Mutual Neighbour

Nearest Neighbour

Notes

Algorithm: Nearest-Neighbour (X: Examples, D:double)

Generate a prototype with the first example

while there are examples **do**

 e= current example

 d= distance of e to the the nearest prototype

if $d \leq D$ **then**

 Introduce the example in the class

 Recompute the prototype

else


 Create a new prototype with this example

end

end

Nearest Neighbour

Notes



Disadvantages of Numerical Taxonomy

Notes

From the artificial intelligence perspective

- Are based only on syntactic similarity, no context information is used
- Mainly focused on numerical data
- There are no explicit representation of the groups
- Analysis and search are done only by the data analyst
- No possibility of biasing the search using domain information
- Interpretability is not considered

Cognitive Psychology

Notes

- Its research includes the models of human conceptualization
 - How the concepts that we use are created
 - How concepts are characterized/described/modelled
 - How are organized
- It gives theories about characteristics and properties of the concepts
- It is the basis of some of the algorithms for unsupervised symbolic machine learning

Cognitive Psychology (II)

Notes

Three theories about conceptualization

- Classic view: Disjoint Concepts
- Probabilistic/exemplar view: Similarity
- Theory based view: Context

Classic view

Notes

A concept is defined from a list of necessary and sufficient properties

Problems:

- Is difficult to find this list of necessary and sufficient properties for all concepts
- Grade of membership to a concept of the examples of a concept (Typicality)
- Its difficult to assign some examples to a concept (Ambiguity)
- Preferential level for conceptualization (Basic level)


Probabilistic/exemplar based view

Notes

- Probabilistic
 - A concept is a list of attributes, but they are neither sufficient nor necessary
 - Concepts are organized respect their *family resemblance*
 - There is a prototype that represents the central tendency of the concept
 - Membership to a concept is defined from the number of attributes that are shared with the prototype
- Exemplar based
 - A concept is defined by a list of representative examples
 - Membership to a concept is defined from the number of similar examples
 - The number of similar examples is a measure of the typicality of an example
 - How the representative examples are chosen?

Probabilistic/exemplar based view

Notes



Probabilistic/exemplar based view - Comparison

Notes

- Probabilistic view is more space efficient (just one prototype)
- Probabilistic view has a description of the concepts
- Exemplar based view is biased because of the existence of outliers
- Exemplar based view allows partial comparisons
- Exemplar based view is used by expert individuals, probabilistic view is the categorization used by non experienced individuals
- Exemplar based theory is more accepted as the actual human conceptualization

Probabilistic/exemplar based view - Assumptions

Notes

- Similarity among examples are an increasing function of the attributes shared and decreasing of the attributes not shared
- All attributes are independent
- All attributes are on the same level of abstractness
- A concept is equivalent to a list of properties

Theory based view

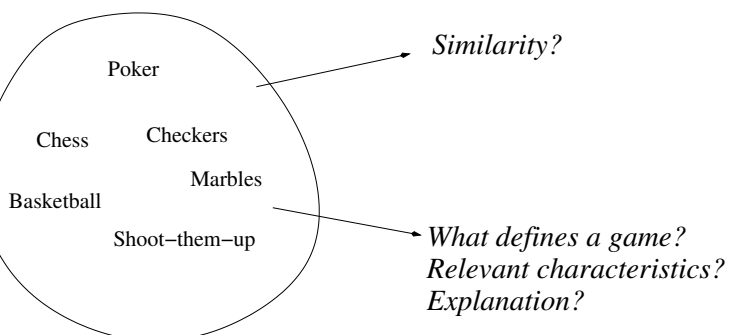
Notes

- Concepts are organised around theories that explain the categorization
- Similarity it is not an absolute, it depends on the context (the perception of an attribute can change)
- Attributes are interrelated (some are redundant)
- Attributes can belong to different levels of abstractness
- Concepts are more complex than a list of attributes

Probabilistic/exemplar based view

Notes

Games



Unsupervised learning

Notes

- We try to extend the view from numerical taxonomy
 - Using qualitative data
 - Allowing the algorithm to do the search in the space of parameters
 - Using background knowledge to reduce the search space
 - Obtaining characterizations of the results and looking for interpretability
- Including the ideas from cognitive psychology about conceptualization

Unsupervised learning

Notes

- Conceptual clustering
 - Non incremental algorithms
 - The result is a partition or a hierarchy
 - Based on numerical taxonomy algorithms
- Concept formation
 - Incremental algorithms
 - The result is a hierarchy
 - Inspired in the ideas from cognitive psychology

Conceptual Clustering

Notes

- Input:
 - A dataset
 - A set of attributes that characterize the dataset
 - Background knowledge (constraints about the concepts, relationships among attributes, relevance, ...)
- Result:
 - A partition or a hierarchy that describes the dataset and characterize the groups

These algorithms combine two tasks


Clustering: We determine the membership of examples to groups that are considered a concept

Characterization: We obtain a description of the groups

Conceptual Clustering

- Similarity is important
- But we also have to take in account the context
- We can only consider a group if there is a matching concept

$$\text{Similarity} = f(\text{Obj}_i, \text{Obj}_j, \text{context}, \text{available concepts})$$



Notes

Conceptual Clustering - CLUSTER/2

[Michalski, Stepp, 1983]

- Characterization using selectors (sets of conjunctive formulas)
- Algorithm similar to K-means
- Optimization of a quality function
- Search is guided by a set of weighted criteria
- Weights are chosen by the user
- Weights are selected depending on the kind of concepts that we are looking for
- Same data could result in different sets of concepts using different weights

Notes

Conceptual Clustering - CLUSTER/2

Lexicographic Evaluation Function (LEF)

- The fit between clusters and data: Descriptions cover the examples
- Simplicity of cluster descriptions: Number of descriptors needed
- Inter-cluster difference: Concepts well separated
- Discrimination index: Descriptions different enough
- Dimensionality reduction: Minimum number of descriptors

Notes

CLUSTER/2

Notes

Algorithm: CLUSTER/2 (X: Examples, k:integer)

Generate k descriptions with the first k examples
 Obtain disjoint descriptions as conjunctive formulas
 Assign each n-k examples to their nearest prototype
 Save the partition

repeat

 Select the more central examples from each group

 Compute LEF

if this partition is better **then**


 substitute last partition with this

end

until LEF is not improved

CLUSTER/2 - Example

Notes



Conceptual Clustering - Autoclass



Notes

[Cheeseman, Stutz, 1996]

- Bayesian approach, uses the EM algorithm as model
- Search in the space of joint probability distributions that better explain the data
- The search in the space of parameters is broader than in EM
 - 1 Number of classes are determined automatically
 - 2 Search for the more suitable probability distribution for each attribute
 - 3 Search in the space of parameters of the joint probability distribution
- The clustering is probabilistic, each example has a grade of membership to each class

Conceptual Clustering - Autoclass

Discover of new classes of stars



Notes

Concept Formation

- Learning has an incremental nature (experience is acquired from continuous observation, not at once)
- Concepts are learnt with their relationships (polithetic hierarchies of concepts)
- Search in the space of hierarchies
- An objective function measures the utility of the learnt structure
- The updating of the structure is performed by a set of conceptual operators
- The result depends on the order of the examples

Notes

Concept Formation - COBWEB

[Fisher, 1989]

- Based on the ideas from cognitive psychology
- Incrementally builds a hierarchy of concepts (top-down)
- The concepts are described probabilistically
- Defines four search operators
- Uses a measure to find the *basic level* (Category utility)

Notes

COBWEB - Category utility

Notes

- Measure defined from a set of categories
- Maximized by the categories in the basic level
- These classes maximize the predictivity of their attributes
- This measures bias the search to categories with high intra-similarity and low inter-similarity

COBWEB - Category utility

Notes

- Intra class similarity : $P(A_i = V_{ij} | C_k)$
- Maximize \rightarrow most of the examples in the class share this value for this attribute
- Inter class similarity: $P(C_k | A_i = V_{ij})$
- Maximize \rightarrow fewer examples from other classes share this value for this attribute
- Maximize the tradeoff between the two measures for a given set of categories:

$$\sum_{k=1}^K P(A_i = V_{ij}) \sum_{i=1}^I \sum_{j=1}^J P(A_i = V_{ij} | C_k) P(C_k | A_i = V_{ij})$$

COBWEB - Category utility

Notes

- Using Bayes theorem:

$$\sum_{k=1}^K P(C_k) \sum_{i=1}^I \sum_{j=1}^J P(A_i = V_{ij} | C_k)^2$$

- $\sum_{i=1}^I \sum_{j=1}^J P(A_i = V_{ij} | C_k)^2$ represents the number of attributes that can be correctly predicted for a class
- We look for a partition that increases this number of attributes compared to a baseline: $\sum_{i=1}^I \sum_{j=1}^J P(A_i = V_{ij})^2$

COBWEB - Category utility

Notes

- Category utility for qualitative attributes for a set of k categories $\{C_1, \dots, C_k\}$

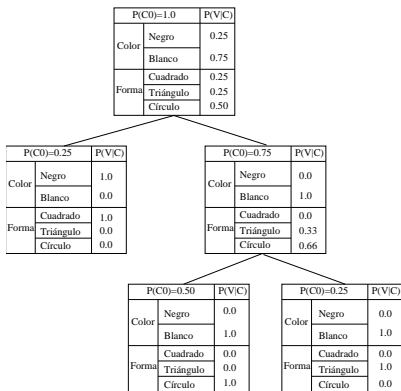
$$\frac{\sum_{k=1}^K P(C_k) \sum_{i=1}^I \sum_{j=1}^J P(A_i = V_{ij} | C_k)^2 - \sum_{i=1}^I \sum_{j=1}^J P(A_i = V_{ij})^2}{K}$$

- Category utility for quantitative attributes (Gaussian distributions)

$$\frac{\sum_{k=1}^K P(C_k) \sum_{i=1}^I \frac{1}{\sigma_{i,k}} - \sum_{i=1}^I \frac{1}{\sigma_{i,p}}}{K}$$

Probabilistic hierarchy

Notes



Algorithm

Notes


- Incremental insertion of each example in the hierarchy
- Look for the path from the root that puts the example in a leaf
- Decide at each level how to modify the hierarchy (which operator apply) to maximize CU and descend recursively the tree

Operators

- **Incorporate:** Put the example inside an existing class
- **New class:** Create a new class at this level
- **Merge:** Two concepts are merge and the example is incorporated inside the new class
- **Divide:** A concept is substituted by its children

Notes

Split - Merge



Notes

COBWEB

Notes

Procedure: Depth-first limited search COBWEB (x: Example, J: Hierarchy)

Update the father with the new example

if we are in a leaf **then**

 Create a new level with this example

else

 Compute **CU** of incorporating the example to each class

 Save the two best **CU**

 Compute **CU** of merging the best two classes

 Compute **CU** of splitting the best class

 Compute **CU** of creating a new class with the example

 Recursive call with the best choice

end
