



Machine learning in computer vision

Lesson 10



Nonhierarchical methods

the space is divided into one set of clusters



K-means clustering

K clusters

Minimizes total intra-cluster scatter (within sum of squares - WSS)

$$W = \sum_{k=1}^K \sum_{x_i \in C_k} \sum_{x_j \in C_k} \|x_i - x_j\|^2 = \sum_{k=1}^K 2N_k \sum_{x_i \in C_k} \|x_i - m_k\|^2 = \sum_{k=1}^K WSS_k$$

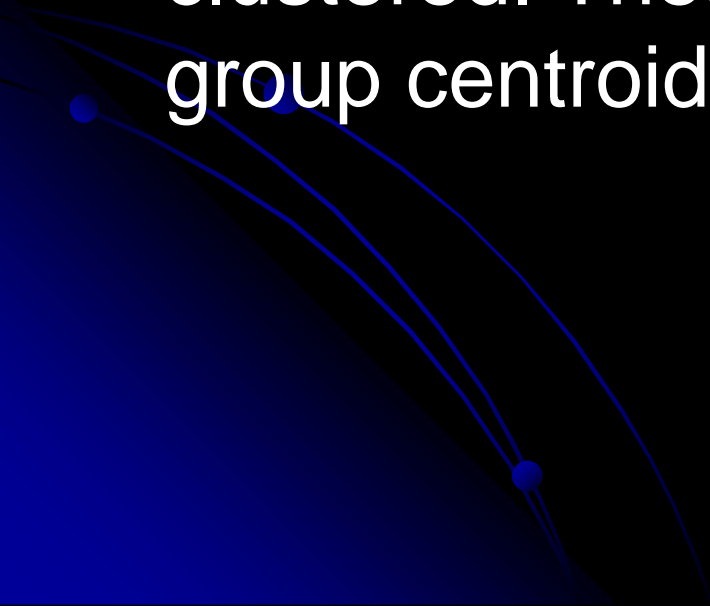
m_k centroid of cluster k

N_k number of points in cluster k

K-Means Algorithm

Initialization:

Randomly place K points into the space represented by the objects that are being clustered. These points represent initial group centroids.

A decorative graphic in the bottom-left corner of the slide. It features three blue dots of varying sizes and two thin, curved blue lines that sweep upwards and to the right, creating a sense of motion or flow.

K-Means Algorithm

Assign objects to the group that has the closest centroid


$$C(x) = \arg \min_k \|x - m_k\|^2$$

When all objects have been assigned, recalculate the positions of the K centroids


$$m_k = \frac{\sum_{x:C(x)=k} x}{N_k}, k = 1, \dots, K$$

Repeat until the stopping criteria is met.
(MSE < threshold, or no change in clustering)


K-Means Clustering



K-Means Clustering




K-Means Clustering




Assign objects

K-Means Clustering




Recalculate centroids

K-Means Clustering



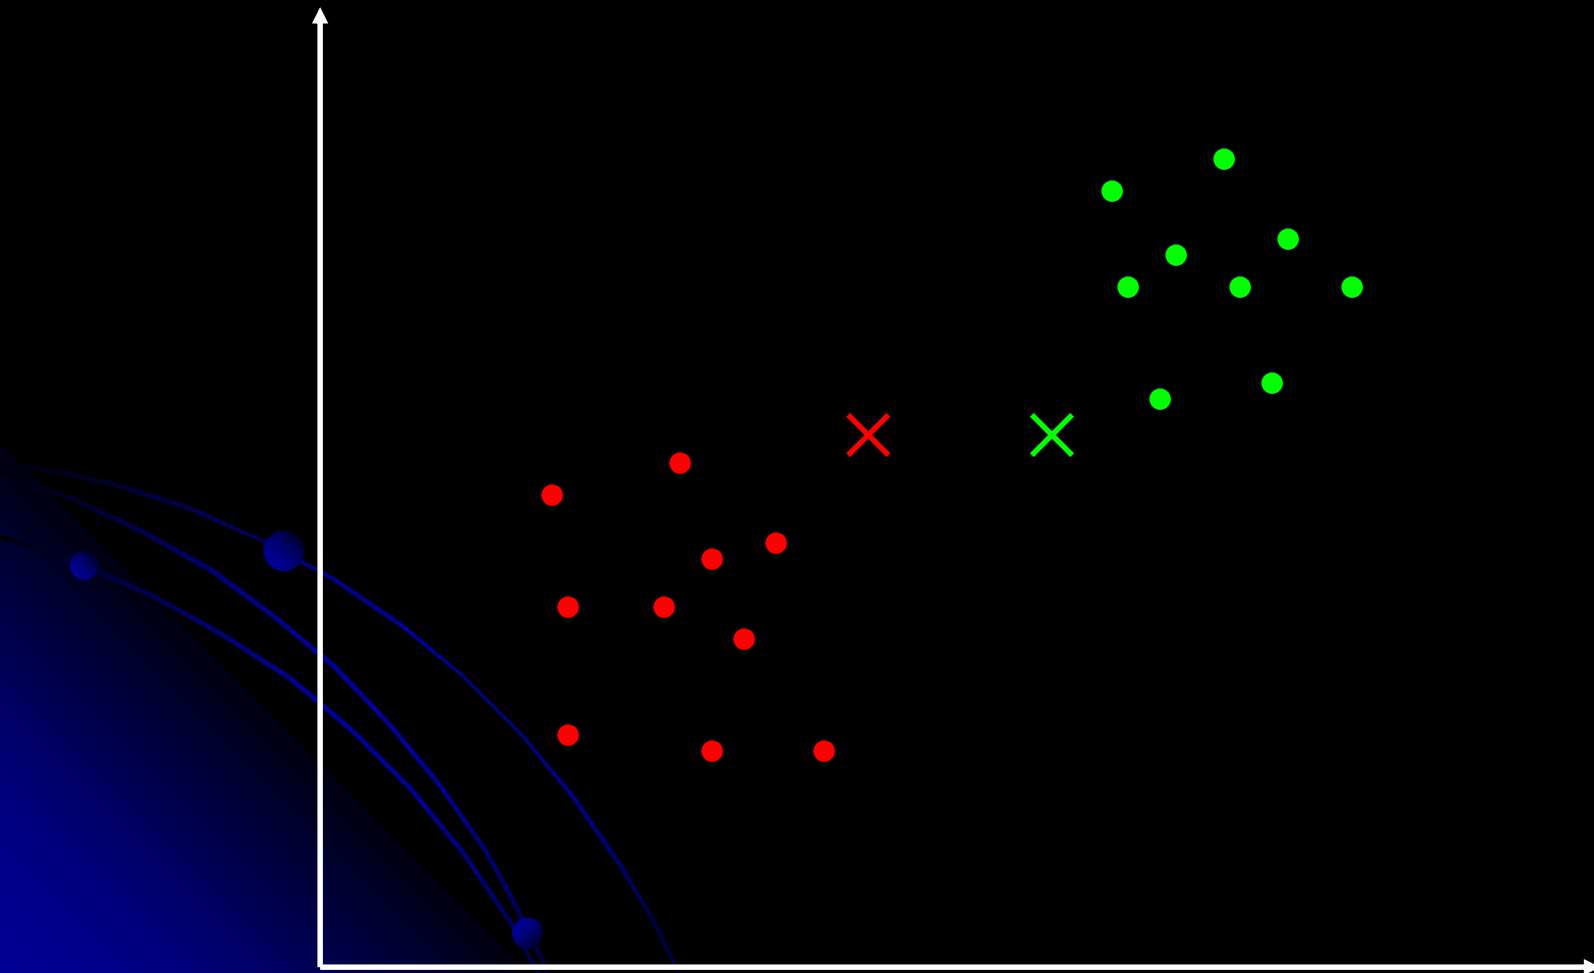
Assign objects

K-Means Clustering



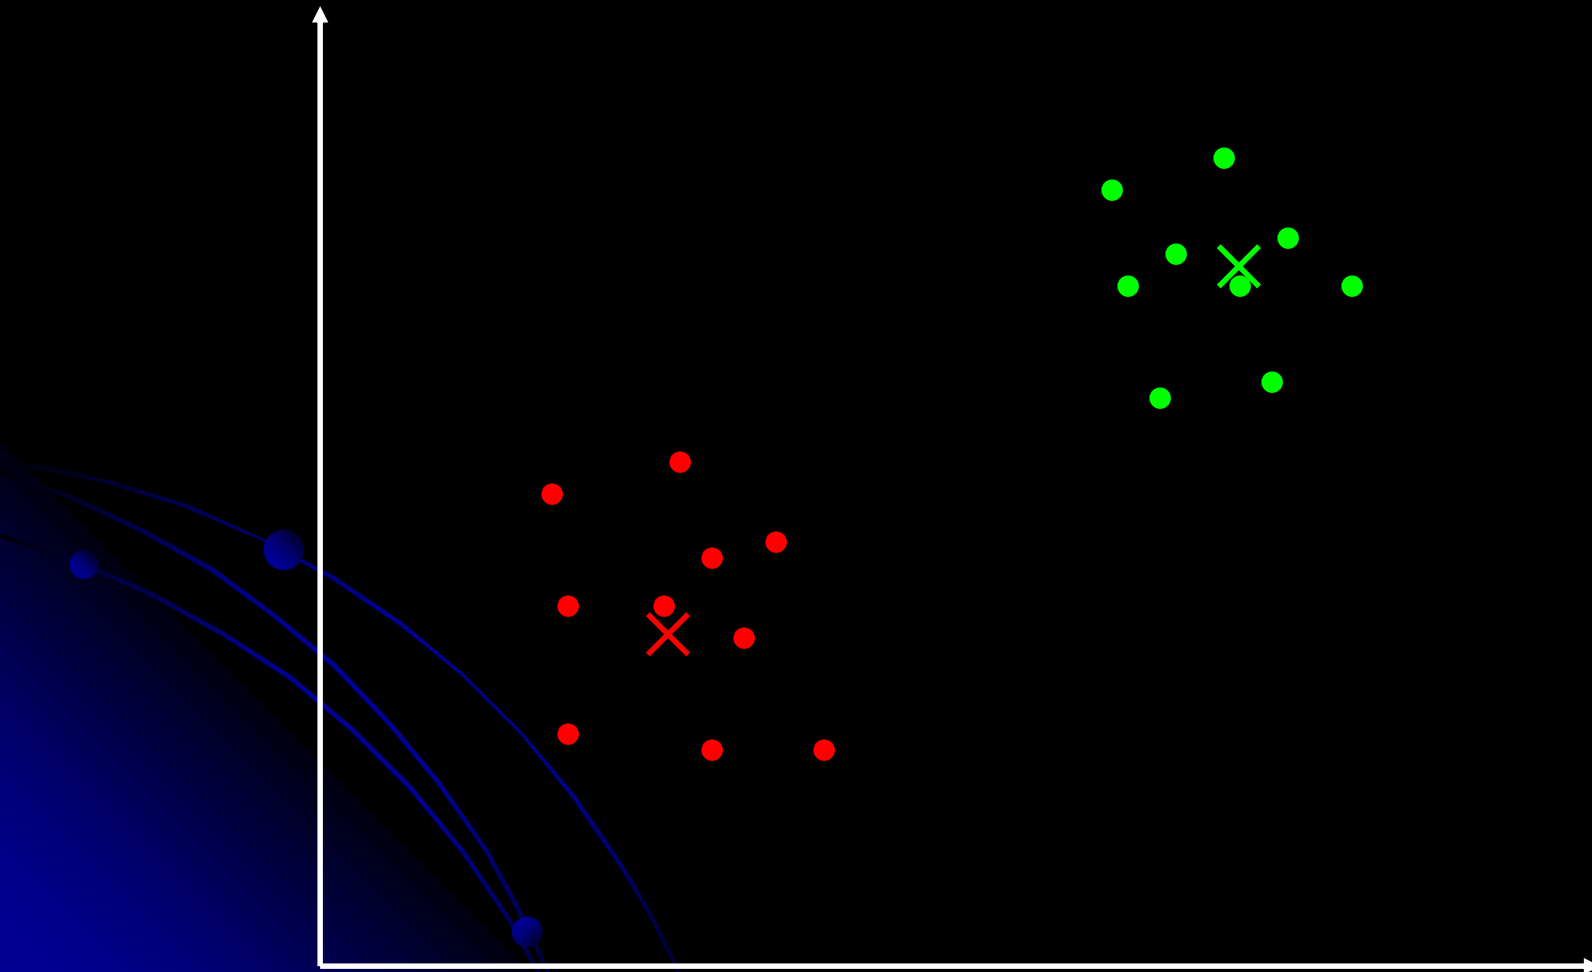
Recalculate centroids

K-Means Clustering



Assign objects

K-Means Clustering



Recalculate centroids

END

K-Means Clustering

Guaranteed to converge

Guaranteed to achieve local optimum, not necessarily global optimum

Sensitive to noise and outlier data points


Clusters are sensitive to initial assignment of centroids (not a deterministic algorithm)

Clusters can be inconsistent from one run to another


Initialization




Initial Cluster Assignments and Centroids




Final Cluster Assignments and Centroids




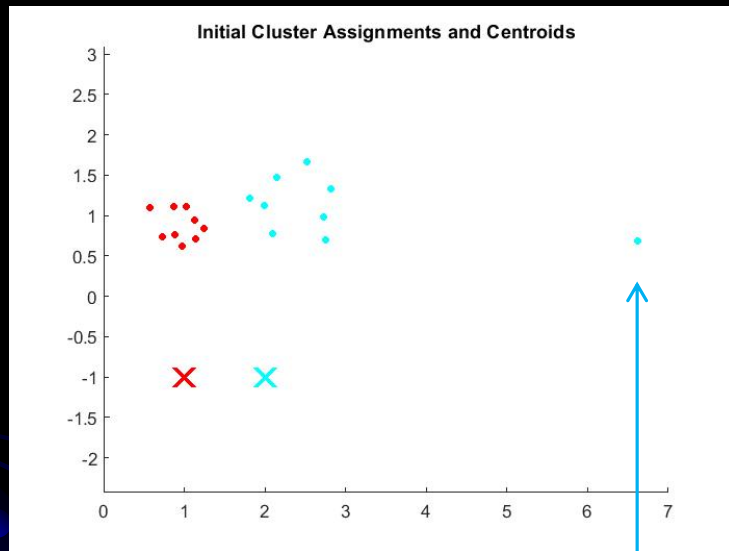
Initial Cluster Assignments and Centroids



Final Cluster Assignments and Centroids



Outliers




Outlier

Choice of K

Many methods

Elbow point: compute total WSS




Choice of K

Gap value

$$\text{Gap}_N(K) = E_N^*\{W_K\} - \log W_K$$

$$W_K = \sum_{k=1}^K \frac{1}{2N_k} WSS_k$$



Choice of K

The silhouette value:

a measure of how similar a point is to points in its own cluster, when compared to points in other clusters

$$S(i) = \frac{b(i) - a(i)}{\max(a(i), b(i))}$$

$a(i)$ – average distance from the i -th point to the other points in the same cluster as i


$b(i)$ – minimum average distance from the i -th point to points in a different cluster, minimized over clusters

Try different K s, compute average silhouette

Choice of K

The silhouette plot

Silhouette analysis for KMeans clustering on sample data with $n_clusters = 6$

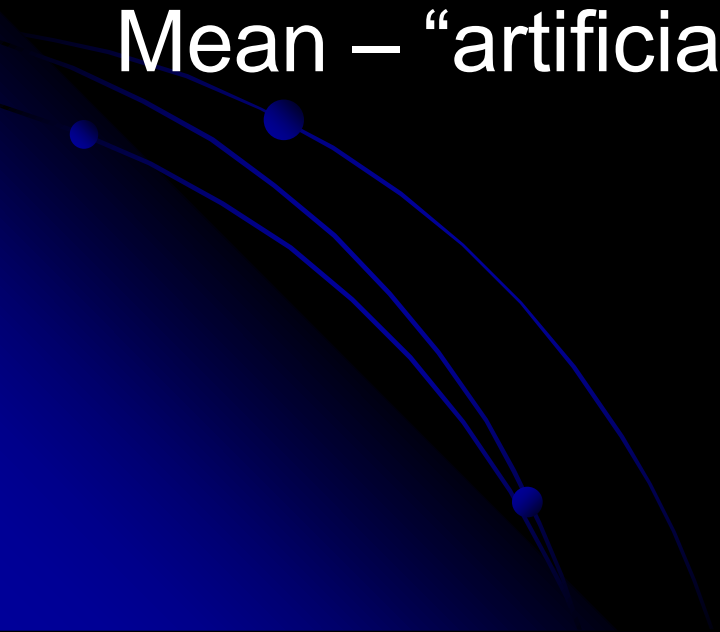


K-medoids

instead of means, use medoids of each cluster

Medoid – object already in the set (e.g. existing color)


Mean – “artificial” object



K-medians

instead of means, use medians in each dimension


object might not be in the set



Fuzzy C-means

Soft membership function

$$\sum_{k=1}^K \sum_{x_i \in C_k} w_{ik}^f \|x_i - m_k\|^2$$



Buckshot

hierarchical agglomerative clustering (HAC) and K-Means clustering

First randomly take a sample of instances of size \sqrt{n}

Run group-average HAC on this sample, which takes only $O(n)$ time

- Use the results of HAC as initial seeds for K-means

Overall algorithm is $O(n)$ and avoids problems of bad seed selection

Bisecting K-means

Divisive hierarchical clustering method using K-means

For $I=1$ to $K-1$ do {

 Pick a leaf cluster C to split

 For $J=1$ to ITER do {

 Use K-means to split C into two sub-clusters, C_1
 and C_2


 }

 Choose the best of the above splits and make it
 permanent

}

Self-Organizing Maps


SOM – Kohonen nets




Self-Organizing Maps

Neurons form a lattice


Input data connected with all neurons



Two spaces of SOM



SOM lattice
Topological structure




Weight space
Same dimensionality as input space

SOM learning

Winner-takes-all algorithm:
The closest node is updated

Algorithm:

1. Randomize the map's nodes weight
2. Select randomly one input vector




SOM learning

3. Find the closest node: best matching unit


$$i^* = \arg \min_i \|\mathbf{x} - \mathbf{w}_i\|$$

Closest using Euclid distance



4. The weight of this node is updated

Winner-takes-all




SOM learning

5. The weights of the adjacent nodes are also updated, by not to the same degree

$$\mathbf{w}_i(t+1) = \mathbf{w}_i(t) + \eta(t)O(i, i^*, t)\|\mathbf{x} - \mathbf{w}_i(t)\|$$

$O(i, i^*, t)$ – neighborhood specification

$\eta(t)$ – learning rate




SOM learning

$$\mathbf{w}_i(t+1) = \mathbf{w}_i(t) + \eta(t)O(i, i^*, t)\|\mathbf{x} - \mathbf{w}_i(t)\|$$

Cooperation phase

This is what ensures the
similarity of weights between
contiguous nodes



SOM learning

6. Reduce the intensity of the update progressively



Adaptation phase

At first, high learning rate, move quickly to the solution; at the end, small learning rate, to avoid oscillations.




7. Repeat 1 to 6 for T_{\max} iterations

Neighborhood specification


$$O(i, i^*, t) = e^{-\frac{\|\mathbf{r}_{i^*}(t) - \mathbf{r}_i(t)\|^2}{2\sigma^2(t)}}$$
$$\sigma(t) = \sigma_0 e^{\frac{-t}{T_{max}}}$$




Neighborhood specification




SOM progress




Weights
random numbers from
 $\langle -0.5, 0.5 \rangle \times \langle -0.5, 0.5 \rangle$




Weights
After 100 iterations




Weights
After 200 iterations



Weights
After 600 iterations






Weights
After 3000 iterations



Weights
After 7000 iterations


Input vectors: Uniform random numbers from $\langle -1, 1 \rangle \times \langle -1, 1 \rangle$.

SOM progress




Classification


U-matrix (unified distance matrix):
visualizes the distances between the neurons




SOM Neighbor Weight Distances




Hits




SOM Weight Positions



Classification



Structural (syntactic) recognition



Structural pattern recognition


Patterns can contain structural and relational information that are difficult or impossible to quantify in feature vector form

Statistical



Number of segments: 4
Number of horizontal segments: 2
Number of vertical segments: 2
Number of diagonal segments: 0



Number of segments: 3
Number of horizontal segments: 1
Number of vertical segments: 0
Number of diagonal segments: 2



Structural



Structural pattern recognition

Structure quantification and description are mainly done using:

- Formal grammars

- Relational descriptions (principally graphs)


Recognition and classification are done using:

- Parsing (for formal grammars)

- Relational graph matching (for relational descriptions)

Applications


- a) Classification of time data (e.g. ECG)
- b) Object recognition described by structural codes (e.g. Freeman code, signature...)
- c) Scene recognition, scene represented as a graph of primitive objects



Time data



Line approximation of ECG:

0 /\ 0 /\ / 00 /\ /\ 0





Structural description of objects

d, b, a, b, c, b, a, b, d, b, a, b, c, b, a, b





221171076555443



Structural scene description

Hierarchical tree structure



Recognition

Theory of formal languages

A grammar generates a (possibly infinite) set of strings (objects)

If we can design a grammar which generates a class of strings, then we can build a machine which will recognize any string in that class

Formal languages

Alphabet is a finite set of symbols, $V = \{x_1, x_2, \dots, x_n\}$

Word over V is a finite string of ordered symbols from V

Example: $V = \{a, b, c\}$, valid words are “abcab”, “abba”, “aaa”, null

V^* set of all words over V

Language is an arbitrary subset L of V^*

Example: $V = \{0, 1\}$, then $L_1 = \{001, 110, 111, 0, \text{null}\}$ is a finite language

$L_2 = \{s \mid s = 1^n 0^2 1^m, n \geq 1, 1 \leq m \leq 10\}$ is an infinite language

Recognition

Objects from one class – words from the language of this class

Classification – decide whether a word belongs to the language of a class

- Finite language – check all words

- Infinite language – use the language grammar or automaton to check

Grammars

Grammar $G = \{V_T, V_N, P, S\}$

V_T is a set of terminal symbols

V_N is a set of *non-terminal* symbols

$V_T \cap V_N = \emptyset$;

P is the set of production rules

$$(V_T \cup V_N)^* V_N (V_T \cup V_N)^* \rightarrow (V_T \cup V_N)^*$$

S is the starting symbol or the root; S belongs to V_N

$L(G)$ is a formal language generated by the grammar G

Each string is composed of only terminals

Each string can be derived from S using the production rules P


Example: $V_T = \{a, b\}$, $V_N = \{S\}$; $P = \{S \rightarrow aSb, S \rightarrow ab\} \Rightarrow L(G) : a^n b^n, n \geq 1$

Inference

Derive grammar based on training set or domain knowledge

Not unique solutions

No general method, usually user interaction is required



Example

Consider,

a: 0° horizontal unit length

b: 120° unit length

c: 240° unit length

$$L = \{a^n b^n c^n; 1 \leq n \leq 3\}$$

L(G) represents the class of equilateral triangles

What is the grammar?

Example

Type 3 Grammar solution

$$V_T = \{a, b, c\}$$

$$V_N = \{S, A, B, C, D, E, F, G, H, I, J, K\}$$

$$S \mapsto aA \quad C \mapsto bI \quad H \mapsto bK$$

$$S \mapsto aC \quad D \mapsto bF \quad I \mapsto c$$

$$A \mapsto aB \quad F \mapsto bJ \quad J \mapsto cI$$

$$A \mapsto aD \quad E \mapsto bG \quad K \mapsto cJ$$

$$B \mapsto aE \quad G \mapsto bH$$

Type 2 Grammar solution

$$V_T = \{a, b, c\}$$

$$V_N = \{S, A, B, C, D, E, F\}$$

$$S \mapsto aAF \quad A \mapsto aBF \quad D \mapsto bc$$

$$A \mapsto b \quad B \mapsto aEF \quad C \mapsto b$$

$$A \mapsto aDF \quad E \mapsto bD \quad F \mapsto c$$

Inference from training set

input: $T = \{x_1, \dots, x_t\}$

output: regular grammar $G = (V_N, V_T, S, P)$

Step 1

Find all terminals in $T \rightarrow$ create V_T

Step 2

For each word $x_i = a_{i1} \dots a_{in}$ ($x_i \in T$) create rules

$S \rightarrow a_{i1} Z_{i1}$

$Z_{i1} \rightarrow a_{i2} Z_{i2}$

.....

$Z_{i,n-2} \rightarrow a_{i,n-1} Z_{i,n-1}$

$Z_{i,n-1} \rightarrow a_{in}$

every Z_{ij} is a new non-terminal

Example

Regular grammar G^* unknown


$G^* = (\{S, A, B, C\}, \{a, b\}, S, P)$

$S \rightarrow aA \mid bB$

$A \rightarrow a \mid aS \mid bC$

$B \rightarrow b \mid bS \mid aC$

$C \rightarrow aB \mid bA$



Finite automaton of G^*

Training set

$T = \{abab, bbaa, baba, aabb\}$


Example

Inference: $VT = \{a, b\}$

$VN = \{S, Z_1, Z_2, Z_3, Z_4, Z_5, Z_6, Z_7, Z_8, Z_9, Z_{10}, Z_{11}, Z_{12}\}$

$S \rightarrow aZ_1$	$Z_1 \rightarrow bZ_2$	$Z_2 \rightarrow aZ_3$	$Z_3 \rightarrow b$
$S \rightarrow bZ_4$	$Z_4 \rightarrow bZ_5$	$Z_5 \rightarrow aZ_6$	$Z_6 \rightarrow a$
$S \rightarrow bZ_7$	$Z_7 \rightarrow aZ_8$	$Z_8 \rightarrow bZ_9$	$Z_9 \rightarrow a$
$S \rightarrow aZ_{10}$	$Z_{10} \rightarrow aZ_{11}$	$Z_{11} \rightarrow bZ_{12}$	$Z_{12} \rightarrow b$

Training set
 $T = \{abab, bbaa, baba, aabb\}$



Example

$L(G^*)$ – infinite

$L(G)$ – finite

$$L(G) = T \subseteq L(G^*)$$




Many non-terminals, some equivalent

Generates only words from the training set

Recognition

1st step – check the terminal symbols

2nd step – try to derive the word from
compliant grammars: top-down, bottom-up



Syntactic deformations

Errors, noise, ...

Structural deformations

Search for most similar word

e.g. Levenshtein distance (number of transformations needed to transform word A to B)

Grammars contains deformation rules (insertion, deletion, substitution)

Summary


The classifier for a structural pattern recognition system consists of a set of grammars, one for each class

The main difficulty lies in grammar inference

Applications – mainly user-constructed grammars



END OF TERM:



FREEDOOOOOM!!!

makeameme.org