IN2110: Språkteknologiske metoder

Evaluering + klyngeanalyse

Erik Velldal

Språkteknologigruppen (LTG)

11. februar, 2025



Today

- Evaluation of supervised classifiers
- Unsupervised machine learning for class discovery: Clustering
- k-means clustering
- Recap

Recap – Classification

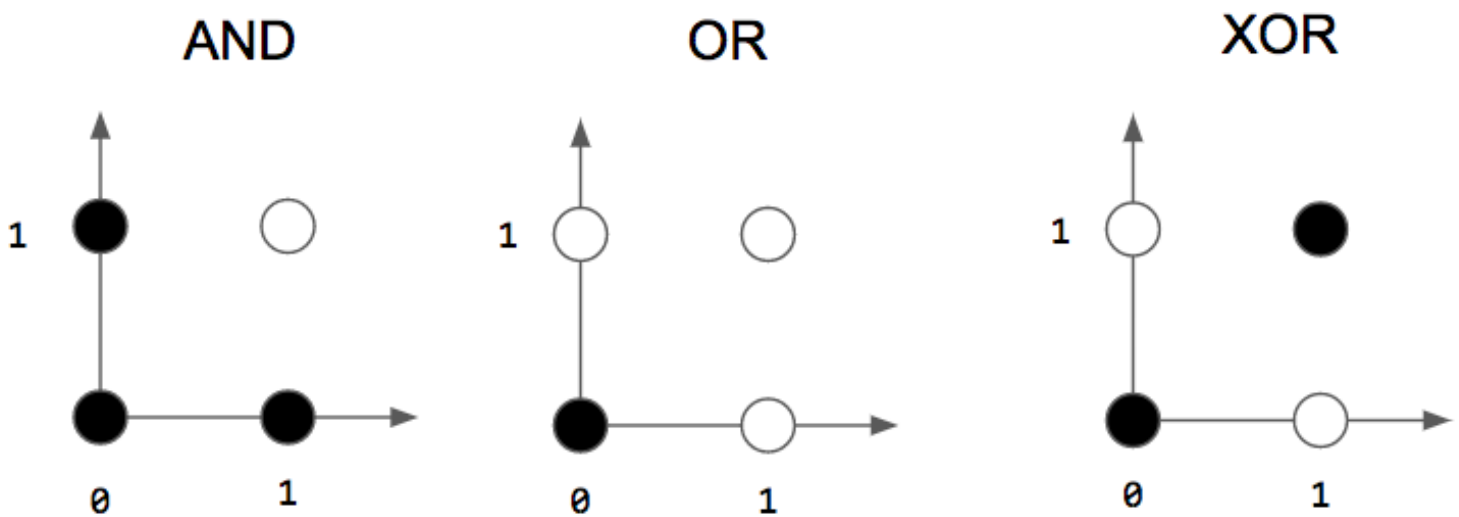
- Supervised vector space classification
- Rocchio
- kNN
- Main difference:

Recap – Classification

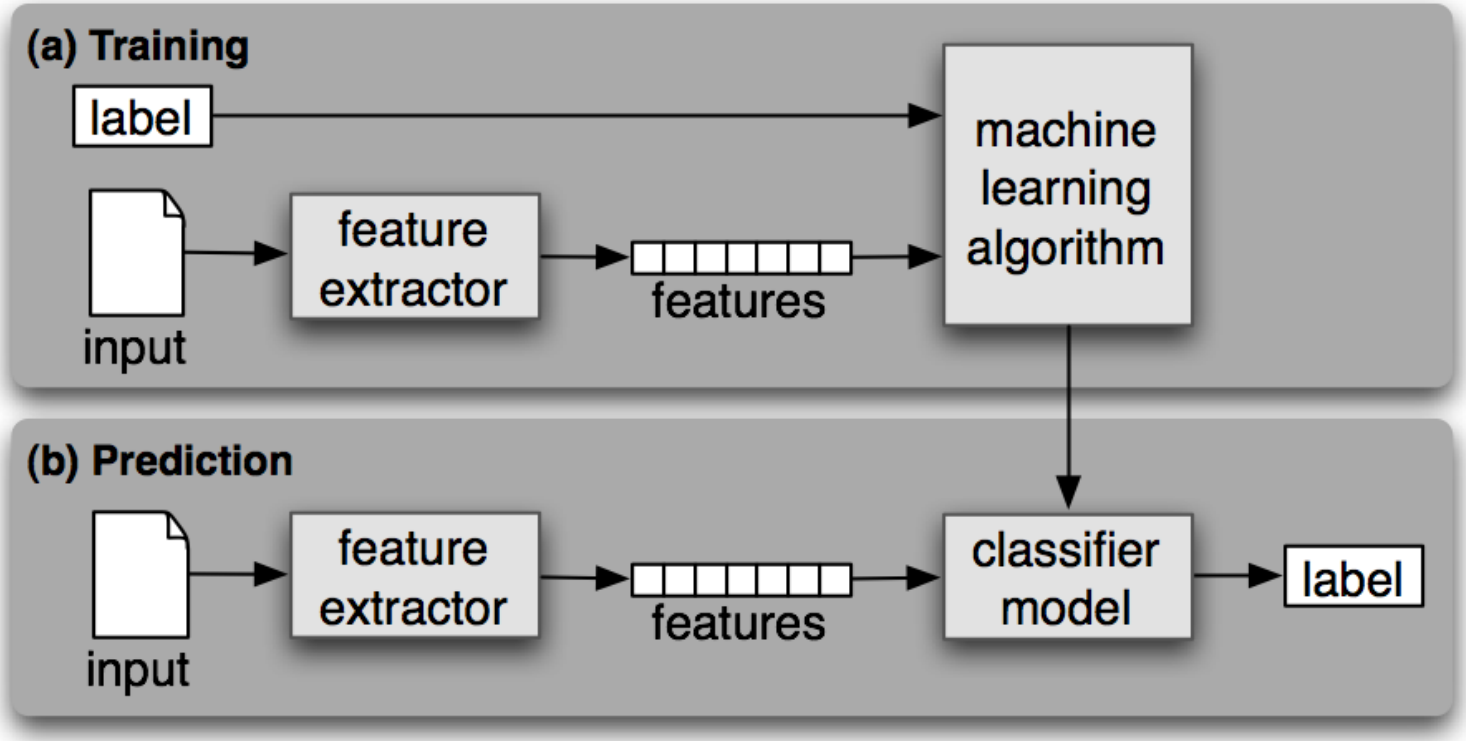
- Supervised vector space classification
- Rocchio
- kNN
- Main difference: Linear vs non-linear decision boundaries.

Recap – Classification

- Supervised vector space classification
- Rocchio
- kNN
- Main difference: Linear vs non-linear decision boundaries.



Classification



(Image: `https://www.nltk.org/book/ch06.html`)

Next topic:



Testing a classifier

- To evaluate a classifier, we measure its number of correct classification predictions on unseen test items.
- The labeled test data is sometimes referred to as the gold standard.
- We evaluate by comparing the predictions made by the model towards the gold labels.
- We will consider different evaluation metrics,
- and the different data splits.
- (Why can't we test on the training data?)

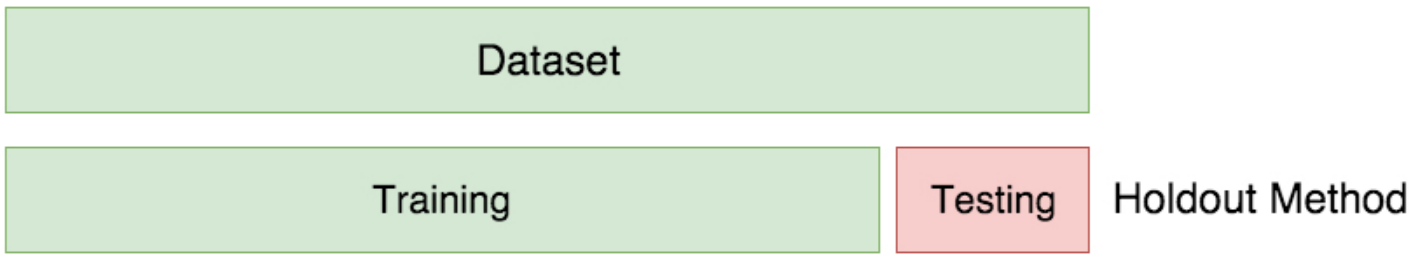
Using data splits

- While tuning our model, estimated from the training set, we repeatedly evaluate towards the development or validation data.
- Or, if we have little data, by n-fold cross-validation.
- Then we evaluate how our *final* model *generalizes* on a held-out test set.



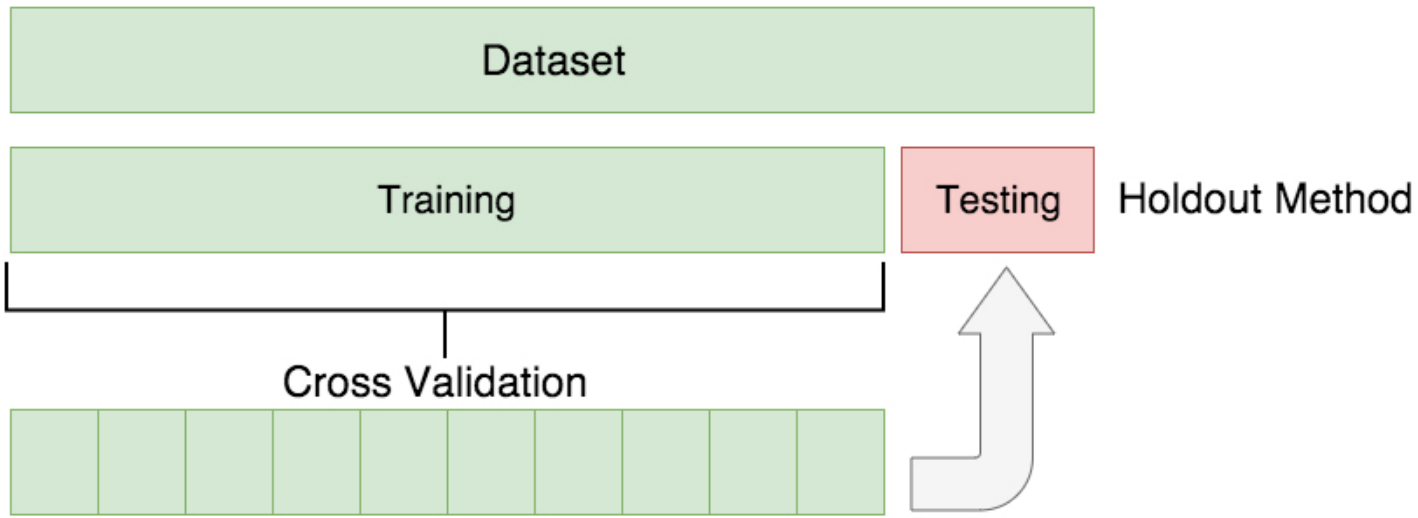
Using data splits

- While tuning our model, estimated from the training set, we repeatedly evaluate towards the development or validation data.
- Or, if we have little data, by n-fold cross-validation.
- Then we evaluate how our *final* model *generalizes* on a held-out test set.



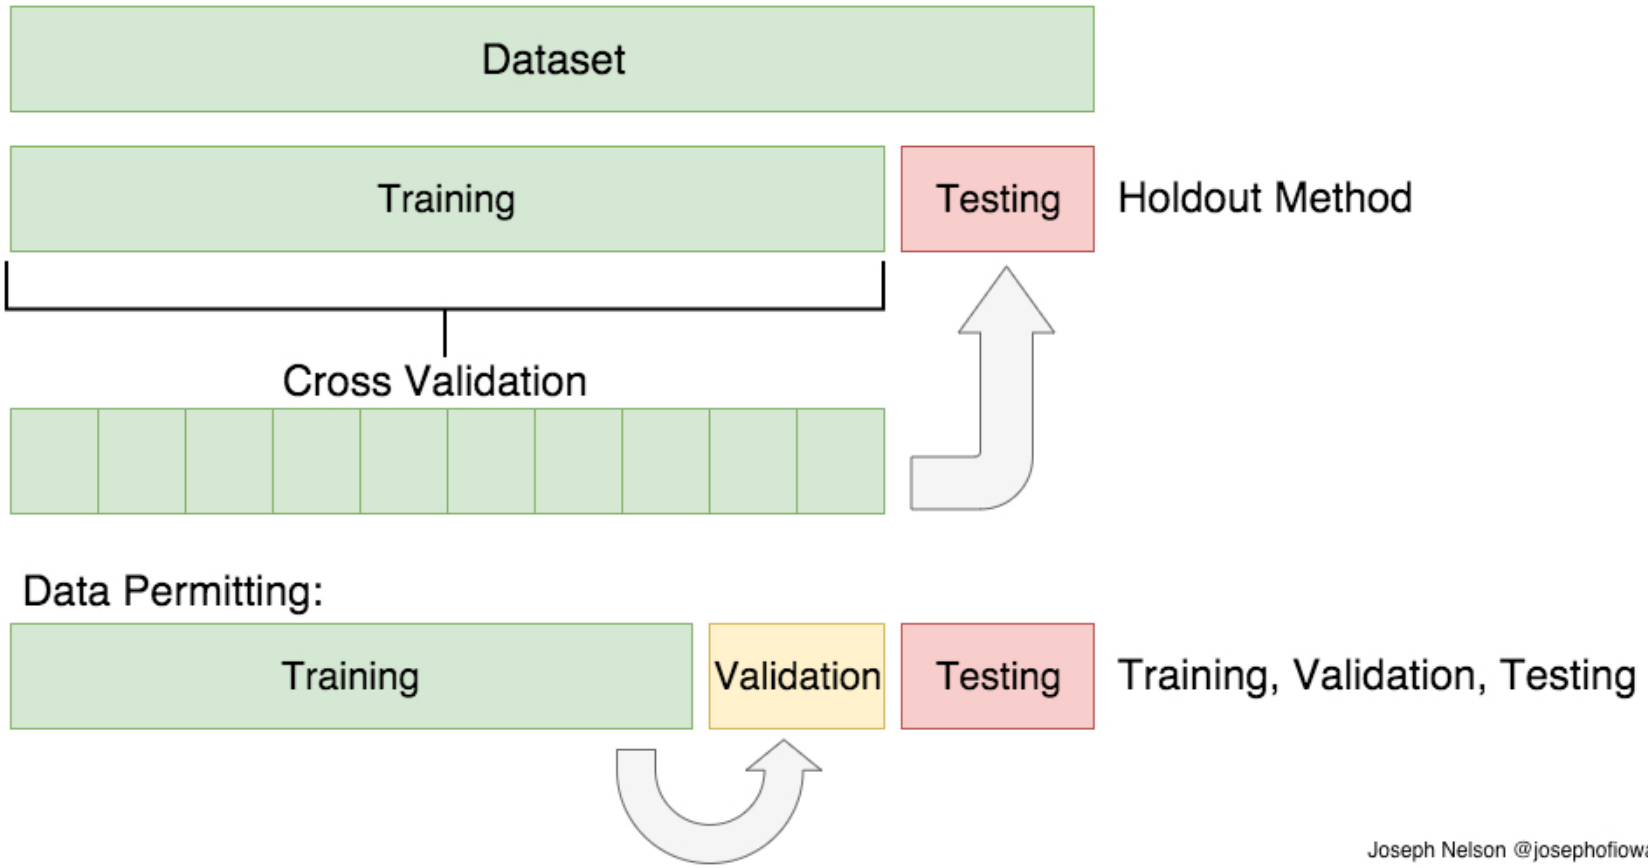
Using data splits

- While tuning our model, estimated from the training set, we repeatedly evaluate towards the development or validation data.
- Or, if we have little data, by n-fold cross-validation.
- Then we evaluate how our *final* model *generalizes* on a held-out test set.

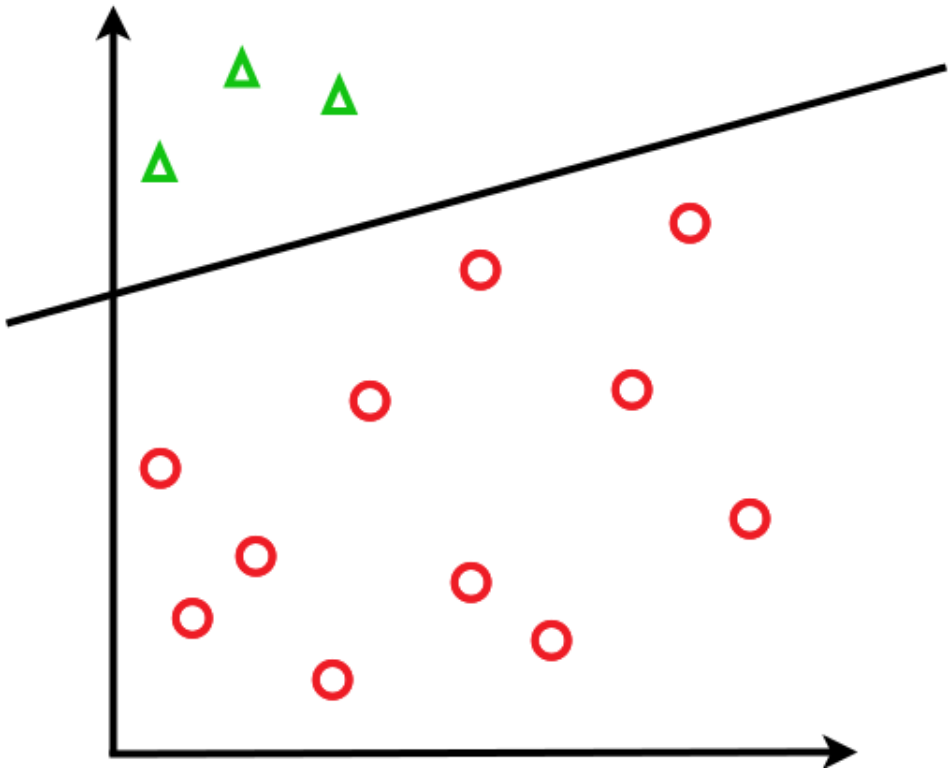


Using data splits

- While tuning our model, estimated from the training set, we repeatedly evaluate towards the development or validation data.
- Or, if we have little data, by n-fold cross-validation.
- Then we evaluate how our *final* model *generalizes* on a held-out test set.

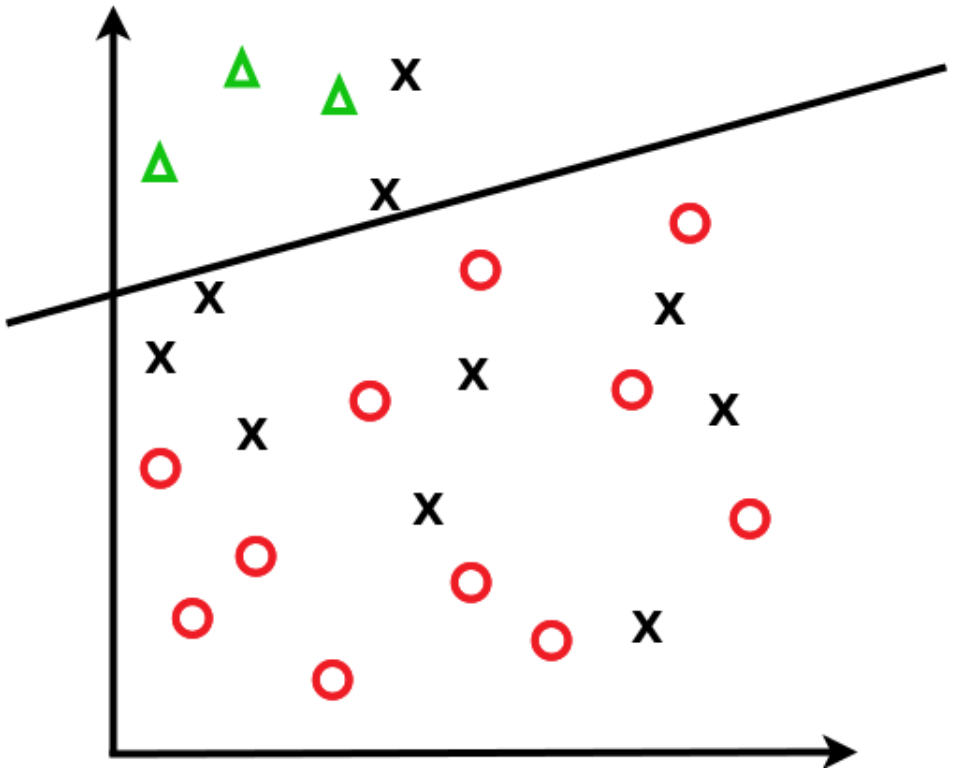


Example: Evaluating classifier decisions



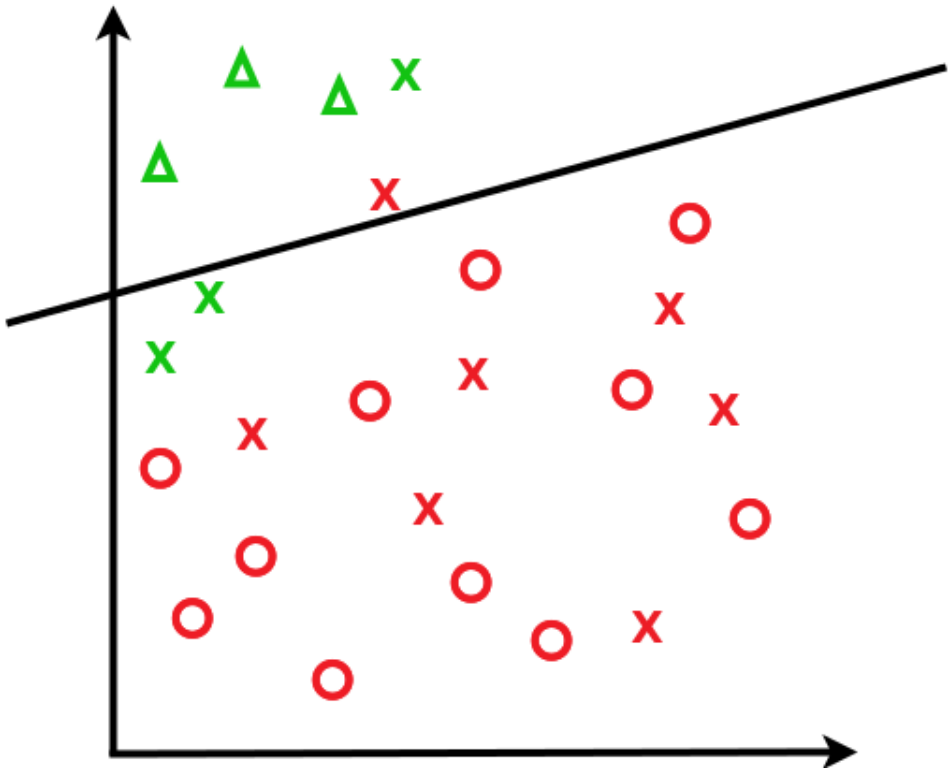
- Predictions for a given class can be wrong or correct in two ways.

Example: Evaluating classifier decisions



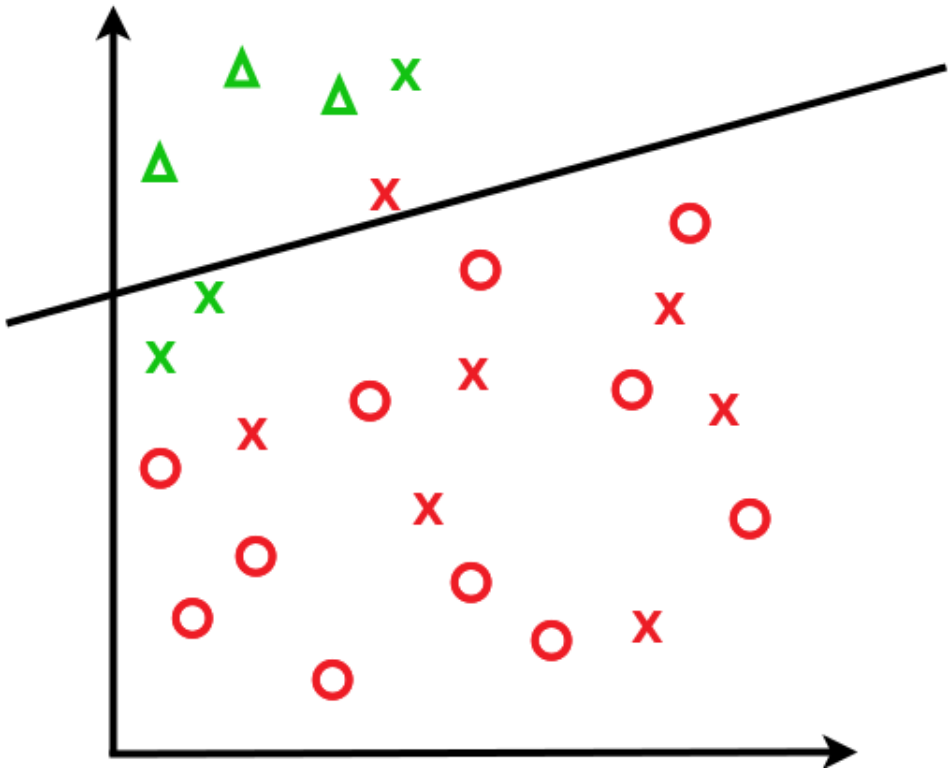
- Predictions for a given class can be wrong or correct in two ways.

Example: Evaluating classifier decisions



- Predictions for a given class can be wrong or correct in two ways.

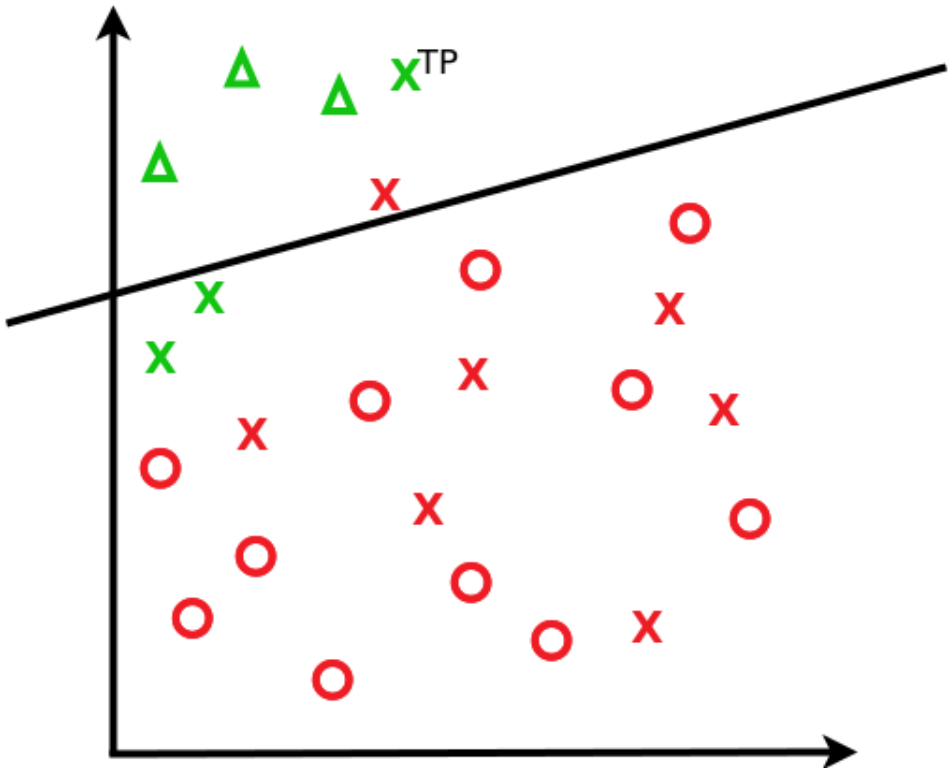
Example: Evaluating classifier decisions



- Predictions for a given class can be wrong or correct in two ways.

	gold = positive	gold = negative
prediction = positive	true positive (TP)	false positive (FP)
prediction = negative	false negative (FN)	true negative (TN)

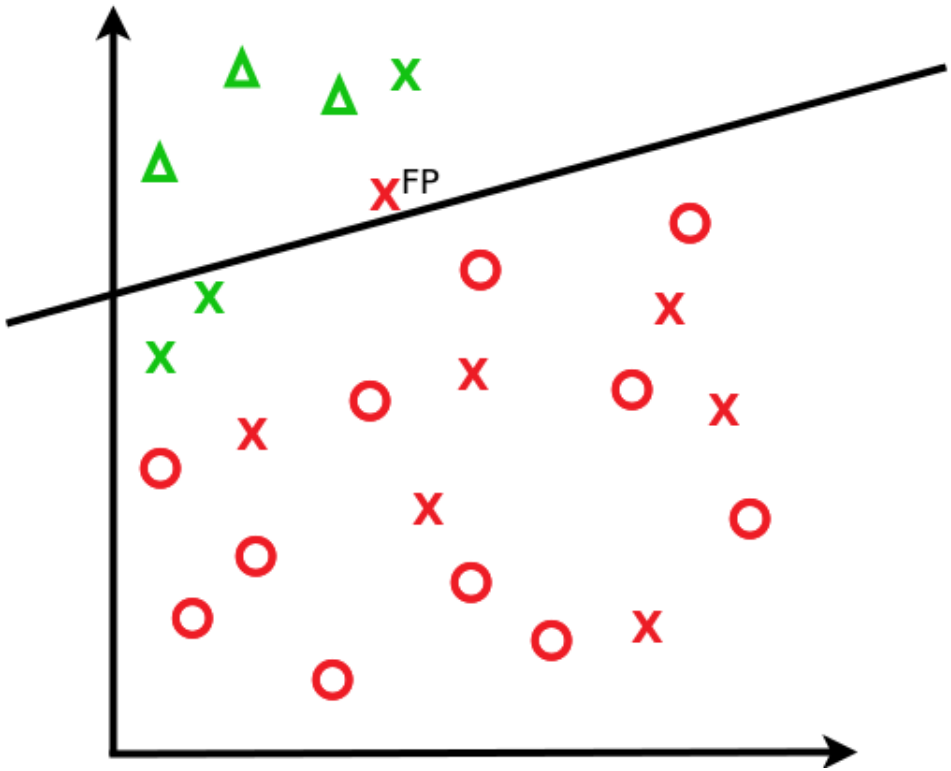
Example: Evaluating classifier decisions



- Predictions for a given class can be wrong or correct in two ways.

	gold = positive	gold = negative
prediction = positive	true positive (TP)	false positive (FP)
prediction = negative	false negative (FN)	true negative (TN)

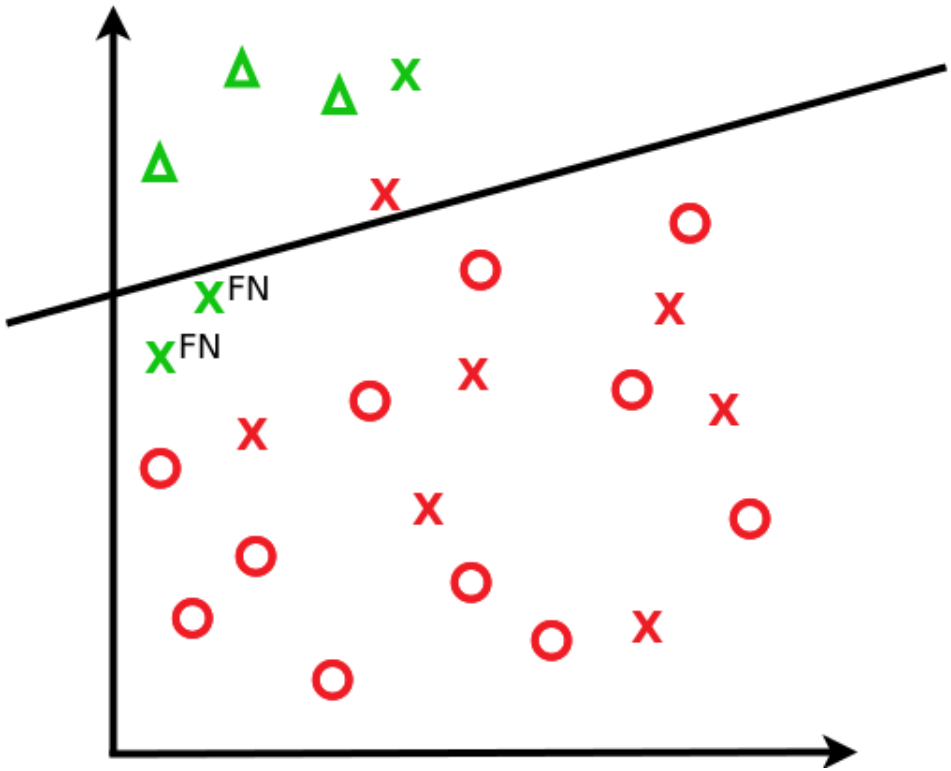
Example: Evaluating classifier decisions



- Predictions for a given class can be wrong or correct in two ways.

	gold = positive	gold = negative
prediction = positive	true positive (TP)	false positive (FP)
prediction = negative	false negative (FN)	true negative (TN)

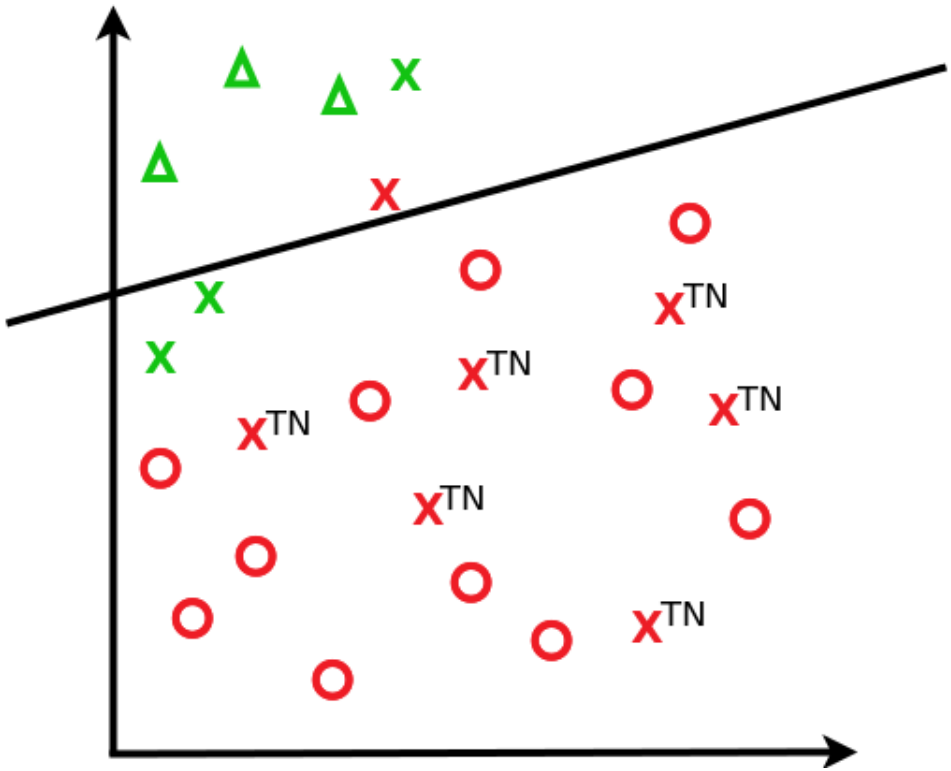
Example: Evaluating classifier decisions



- Predictions for a given class can be wrong or correct in two ways.

	gold = positive	gold = negative
prediction = positive	true positive (TP)	false positive (FP)
prediction = negative	false negative (FN)	true negative (TN)

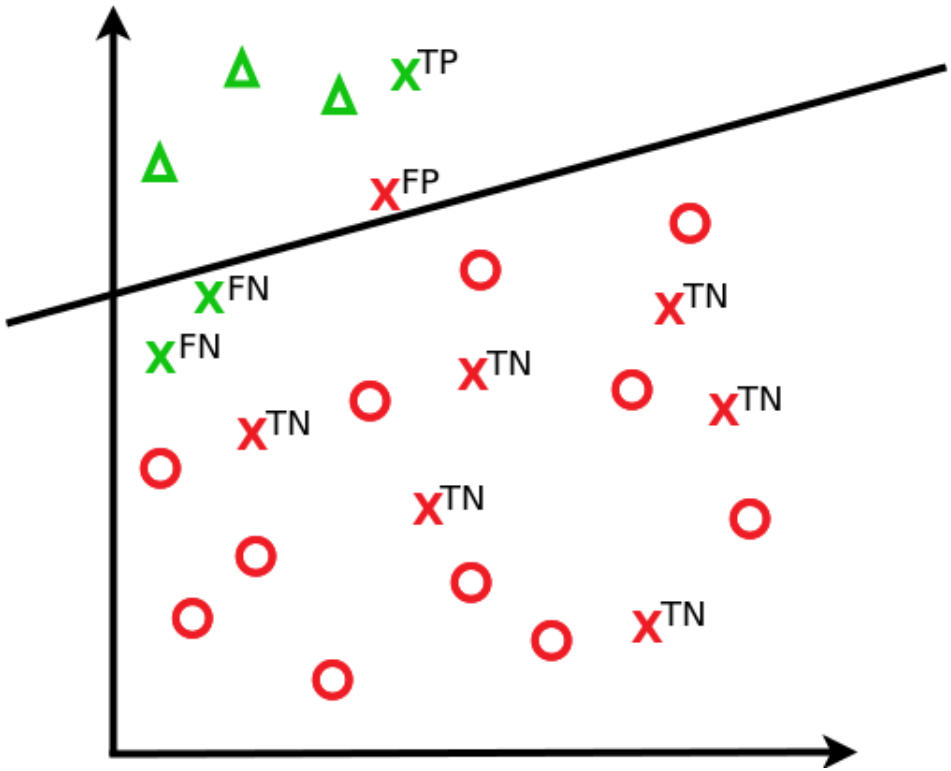
Example: Evaluating classifier decisions



- Predictions for a given class can be wrong or correct in two ways.

	gold = positive	gold = negative
prediction = positive	true positive (TP)	false positive (FP)
prediction = negative	false negative (FN)	true negative (TN)

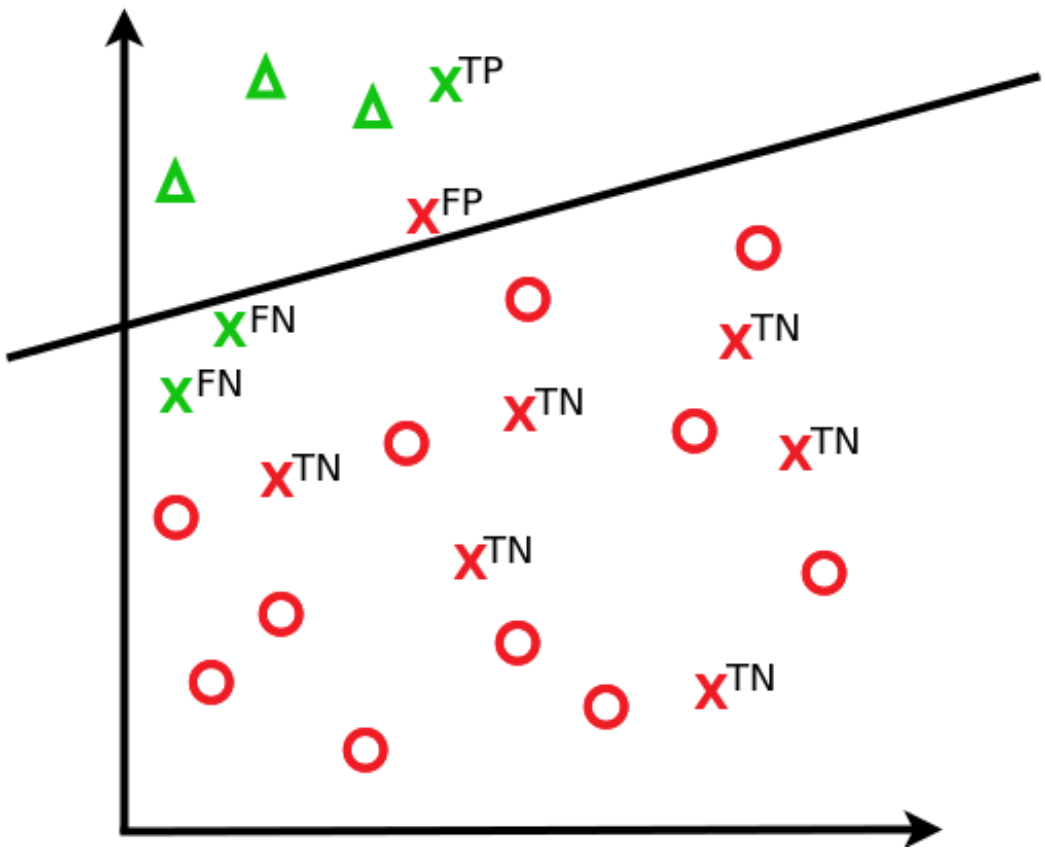
Example: Evaluating classifier decisions



- Predictions for a given class can be wrong or correct in two ways.

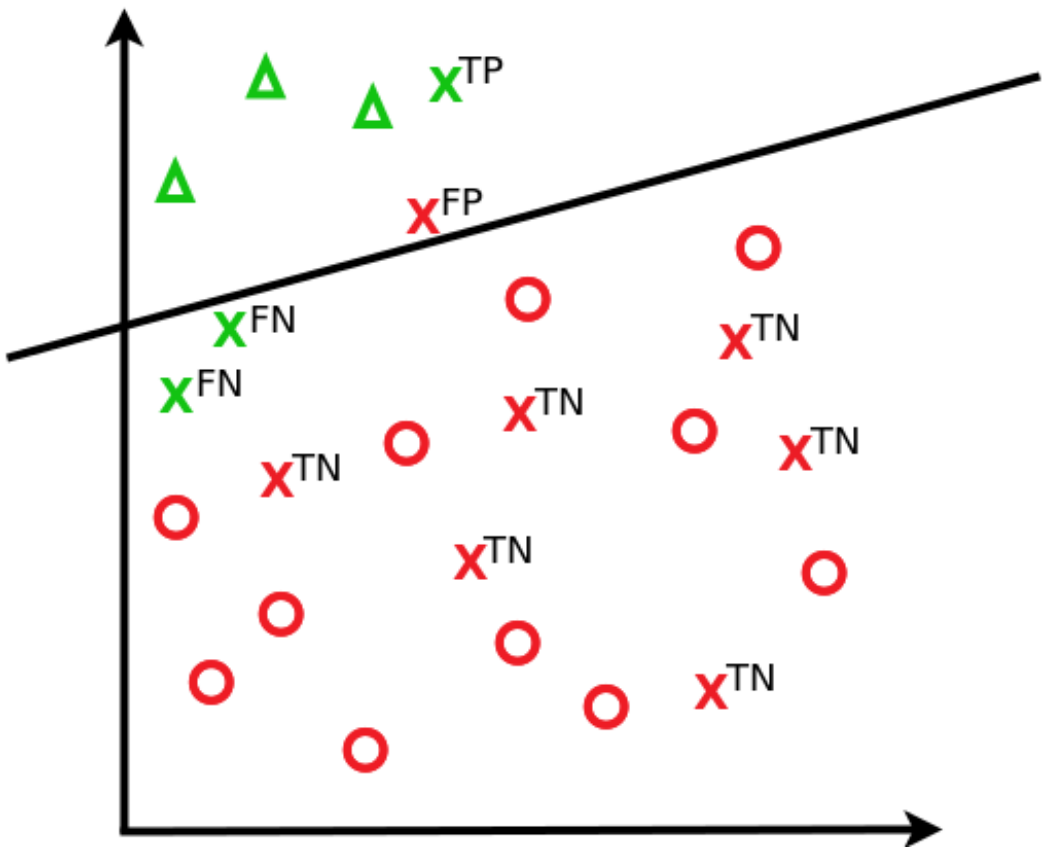
	gold = positive	gold = negative
prediction = positive	true positive (TP)	false positive (FP)
prediction = negative	false negative (FN)	true negative (TN)

Example: Evaluating classifier decisions



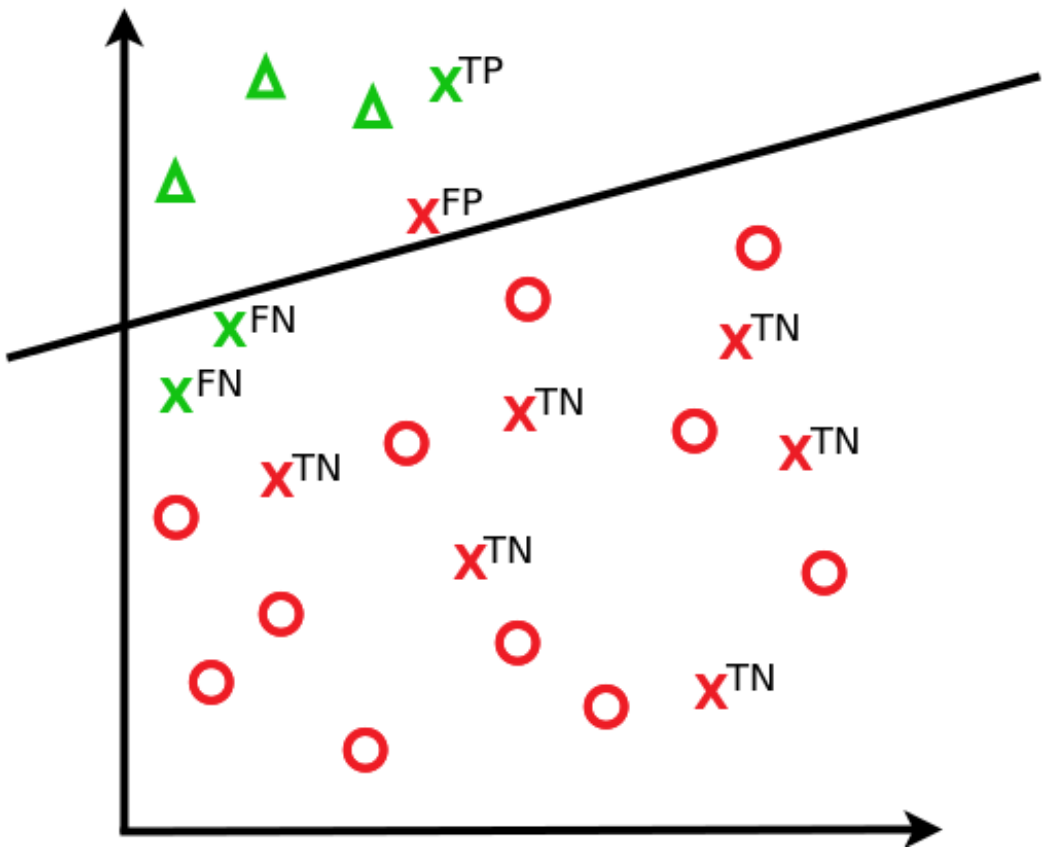
$$\text{Accuracy} = \frac{TP+TN}{N}$$

Example: Evaluating classifier decisions



Accuracy $= \frac{TP+TN}{N}$
$= \frac{1+6}{10} = 0.7$

Example: Evaluating classifier decisions

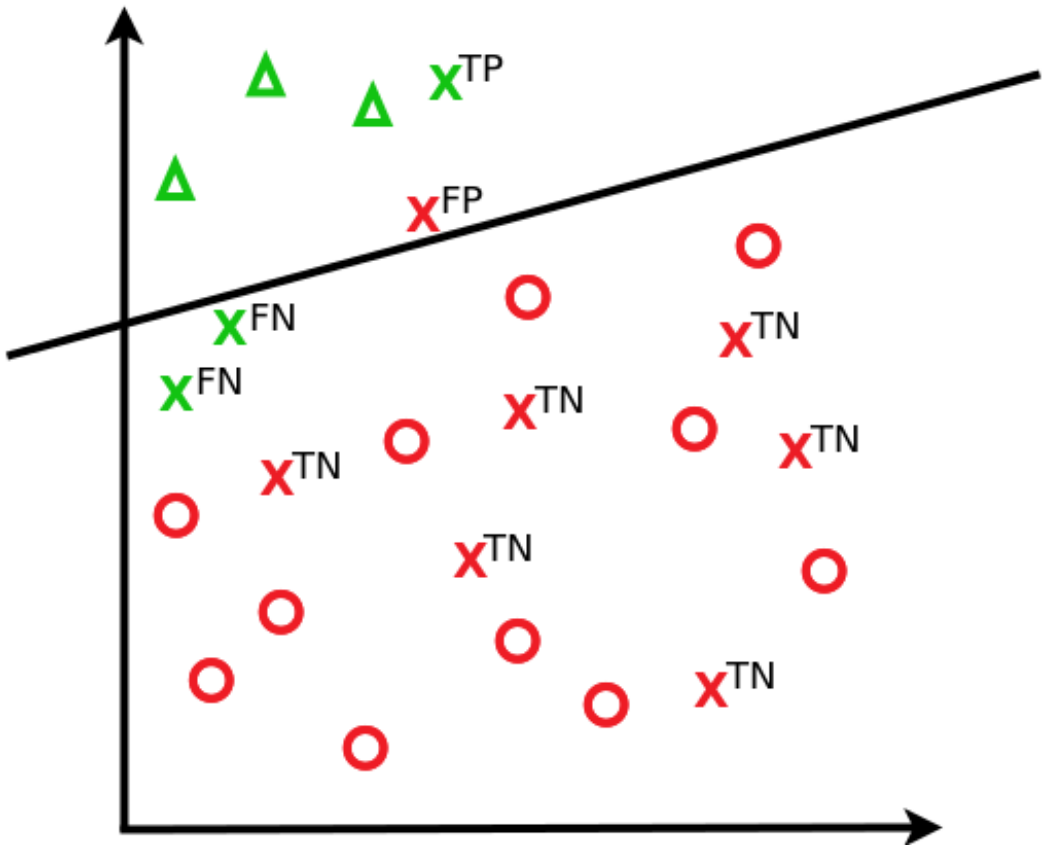


Accuracy $= \frac{TP+TN}{N}$
$= \frac{1+6}{10} = 0.7$

Precision $= \frac{TP}{TP+FP}$

Recall $= \frac{TP}{TP+FN}$

Example: Evaluating classifier decisions

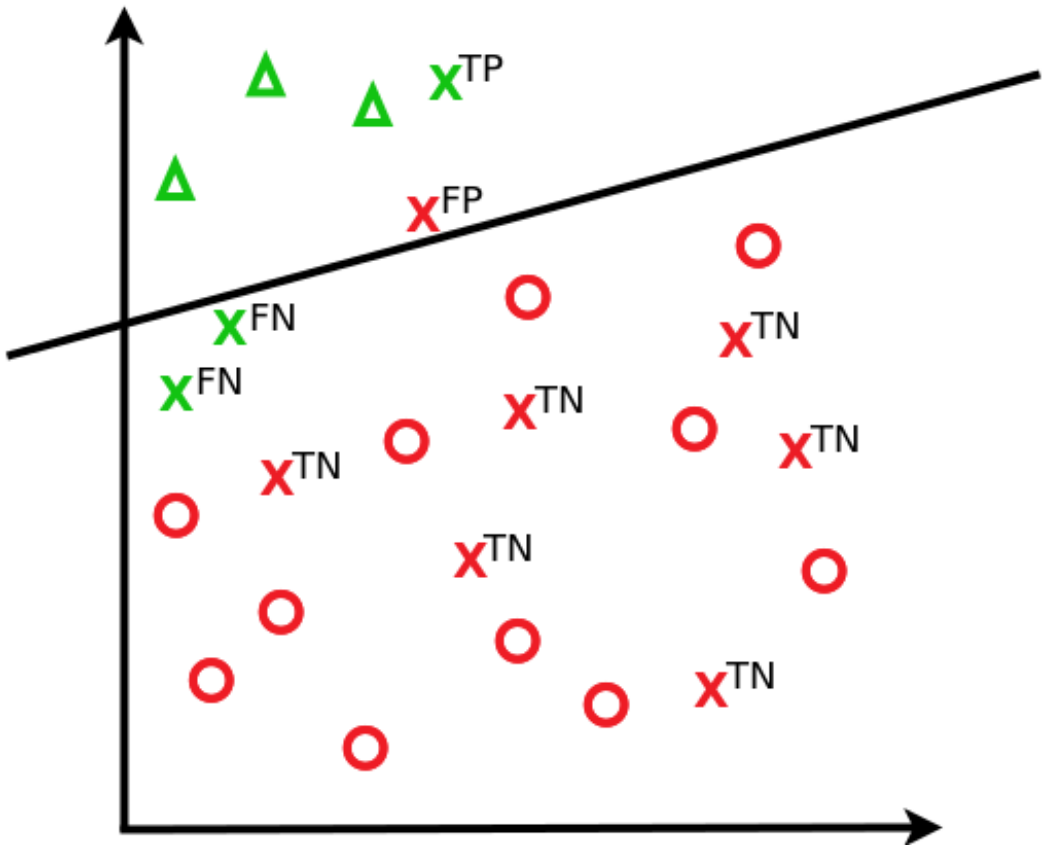


$$\text{Accuracy} = \frac{TP+TN}{N} = \frac{1+6}{10} = 0.7$$

$$\text{Precision} = \frac{TP}{TP+FP} = \frac{1}{1+1} = 0.5$$

$$\text{Recall} = \frac{TP}{TP+FN}$$

Example: Evaluating classifier decisions

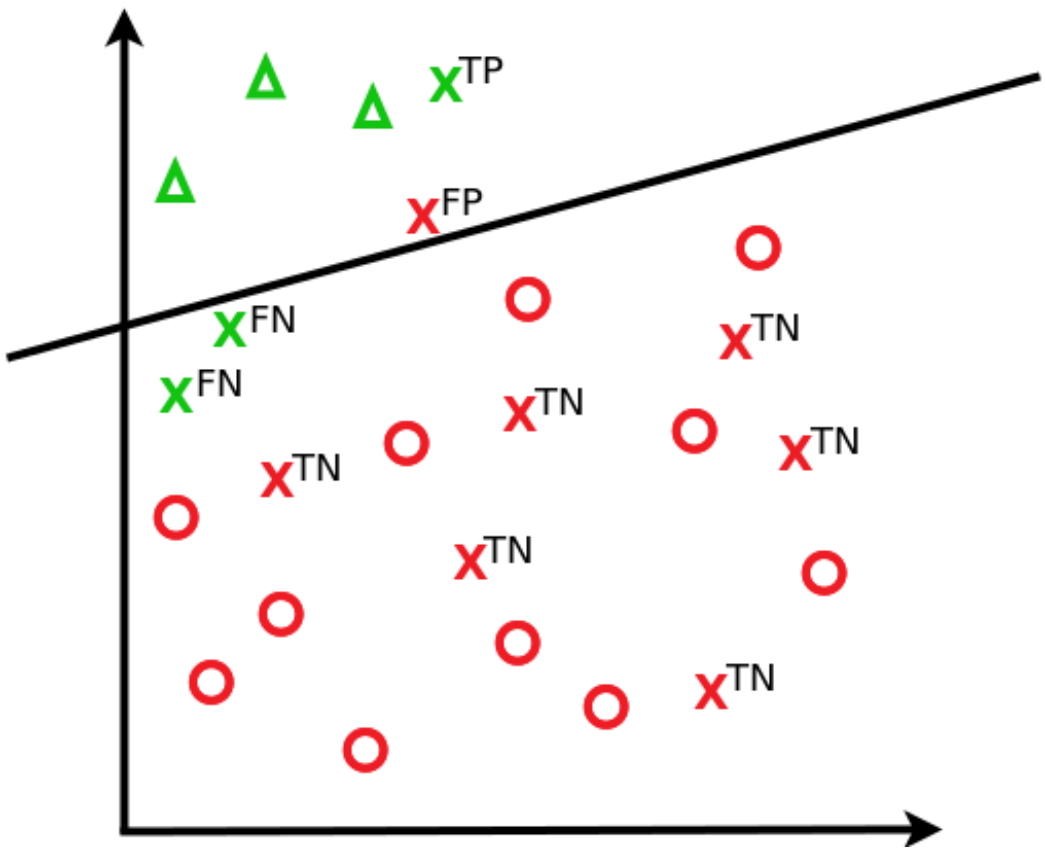


Accuracy $= \frac{TP+TN}{N}$
$= \frac{1+6}{10} = 0.7$

Precision $= \frac{TP}{TP+FP}$
$= \frac{1}{1+1} = 0.5$

Recall $= \frac{TP}{TP+FN}$
$= \frac{1}{1+2} = 0.33$

Example: Evaluating classifier decisions



Accuracy $= \frac{TP+TN}{N}$
$= \frac{1+6}{10} = 0.7$

Precision $= \frac{TP}{TP+FP}$
$= \frac{1}{1+1} = 0.5$

Recall $= \frac{TP}{TP+FN}$
$= \frac{1}{1+2} = 0.33$

F-score
$= 2 \times \frac{precision \times recall}{precision+recall} = 0.4$

Evaluation measures

- Accuracy $= \frac{TP+TN}{N} = \frac{TP+TN}{TP+TN+FP+FN}$
  - The ratio of correct predictions.
  - Not suitable for unbalanced numbers of positive / negative examples.
- Precision $= \frac{TP}{TP+FP}$
  - The number of detected class members that were correct.
- Recall $= \frac{TP}{TP+FN}$
  - The number of actual class members that were detected.
  - Trade-off: Positive predictions for all examples would give 100% recall but (typically) terrible precision.
- F-score $= 2 \times \frac{precision \times recall}{precision+recall}$
  - Balanced measure of precision and recall (harmonic mean).
- Some degree of manual error analysis also advised!

Evaluating multi-class predictions

Micro-averaging

- Sum TPs, FPs, and FNs for all points/objects across all classes, and then compute global metric based on this.
- The micro average will be highly influenced by the large classes.

Evaluating multi-class predictions

Micro-averaging

- Sum TPs, FPs, and FNs for all points/objects across all classes, and then compute global metric based on this.
- The micro average will be highly influenced by the large classes.

Macro-averaging

- First compute metric for each class, and then compute the global average of these.
- The macro average will be highly influenced by the small classes.

Evaluating multi-class predictions

Micro-averaging

- Sum TPs, FPs, and FNs for all points/objects across all classes, and then compute global metric based on this.
- The micro average will be highly influenced by the large classes.

Macro-averaging

- First compute metric for each class, and then compute the global average of these.
- The macro average will be highly influenced by the small classes.

Weighted

- Like macro avg. but taking class imbalance into account:
- Weights the score of each class by number of (gold) examples.

Two approaches to evaluation

Intrinsic

- Focuses on how well a model performs on its own task,
- using metrics that directly measure performance on a labeled dataset.
- (What we have assumed so far.)

Two approaches to evaluation

Intrinsic

- Focuses on how well a model performs on its own task,
- using metrics that directly measure performance on a labeled dataset.
- (What we have assumed so far.)

Extrinsic

- Measures impact on a larger application or end-to-end task.
- Example for topic classification; do the predictions improve user satisfaction in a product recommendation system?

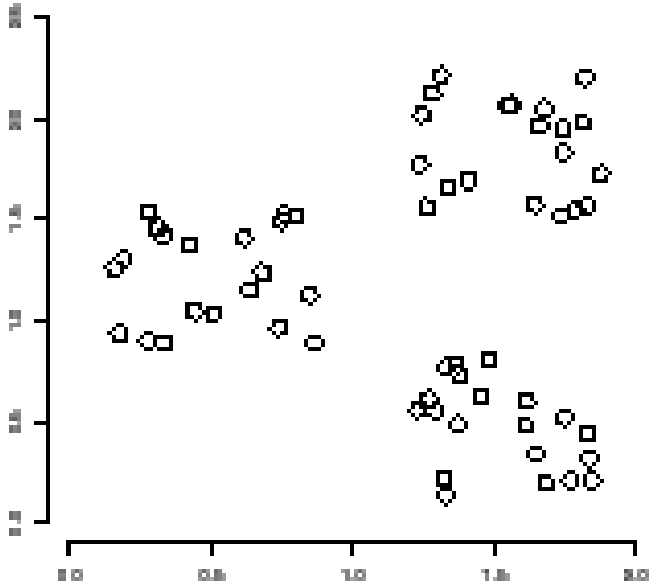
Classification:



(We will return to classification later in the term.)

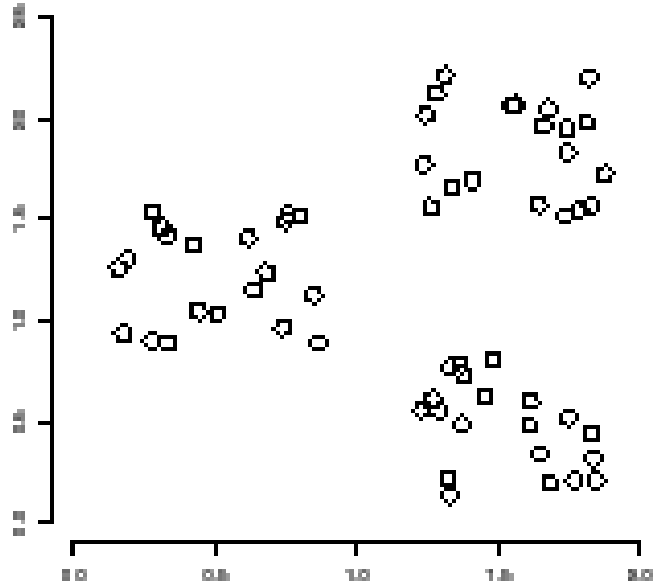
Cluster analysis (*klyngeanalyse*)

- A cluster: 'A group of similar things or people positioned or occurring closely together.' *Oxford Dictionaries*
- Cf. the contiguity hypothesis in classification



Cluster analysis (*klyngeanalyse*)

- A cluster: 'A group of similar things or people positioned or occurring closely together.' *Oxford Dictionaries*
- Cf. the contiguity hypothesis in classification



Clustering or cluster analysis

- Unsupervised learning from unlabeled data.
- Automatically group similar objects together into k categories.
- No pre-defined classes:
- We only specify the features and similarity measure (and k, usually).

Example applications of cluster analysis

- Visualization and exploratory data analysis.

Example applications of cluster analysis

- Visualization and exploratory data analysis.
- Many applications within IR, e.g.:
  - Speed up search: First retrieve the most relevant cluster, then retrieve documents from within the cluster.
  - Presenting the search results: Instead of ranked lists, organize the results as clusters.

Example applications of cluster analysis

- Visualization and exploratory data analysis.
- Many applications within IR, e.g.:
  - Speed up search: First retrieve the most relevant cluster, then retrieve documents from within the cluster.
  - Presenting the search results: Instead of ranked lists, organize the results as clusters.
- Dimensionality reduction: class-based features.

Example applications of cluster analysis

- Visualization and exploratory data analysis.
- Many applications within IR, e.g.:
  - Speed up search: First retrieve the most relevant cluster, then retrieve documents from within the cluster.
  - Presenting the search results: Instead of ranked lists, organize the results as clusters.
- Dimensionality reduction: class-based features.
- Product recommendations, demographic analysis, news aggregation, topic directories, ...

Main types of clustering methods

Hierarchical

- Creates a tree structure of hierarchically nested clusters.

Flat

- Tries to directly decompose the data into a set of clusters.
- What we will focus on.

Flat clustering

- Given a set of objects $O = \{o_1, \ldots, o_n\}$, construct a set of clusters $C = \{c_1, \ldots, c_k\}$, where each object o_i is assigned to a cluster c_j.
- $=$ a partition.
- Parameters:
  - The cardinality k (the number of clusters).
  - The similarity function s.

Flat clustering

- Given a set of objects $O = \{o_1, \ldots, o_n\}$, construct a set of clusters $C = \{c_1, \ldots, c_k\}$, where each object o_i is assigned to a cluster c_j.
- $=$ a partition.
- Parameters:
  - The cardinality k (the number of clusters).
  - The similarity function s.
- Formally defined as an optimization problem:
- We want to find an assignment $\gamma : O \rightarrow C$ that optimizes some objective function $F_s(\gamma)$.
- In general terms, we want to optimize for:
  - High intra-cluster similarity
  - Low inter-cluster similarity

Flat clustering (cont'd)

Optimization problems are search problems:

- There's a finite number of possible partitionings of O.
- Naive solution: enumerate all possible assignments $\Gamma = \{\gamma_1, \ldots, \gamma_m\}$ and choose the best one,

$$\hat{\gamma} = \arg\min_{\gamma \in \Gamma} F_s(\gamma)$$

Flat clustering (cont'd)

Optimization problems are search problems:

- There's a finite number of possible partitionings of O.
- Naive solution: enumerate all possible assignments $\Gamma = \{\gamma_1, \ldots, \gamma_m\}$ and choose the best one,

$$\hat{\gamma} = \arg\min_{\gamma \in \Gamma} F_s(\gamma)$$

- Problem: Exponentially many possible partitions.
- Approximate the solution by iteratively improving on an initial (possibly random) partition until some stopping criterion is met.

k-means

- Unsupervised variant of the Rocchio classifier.
- Goal: Partition the n observed objects into k clusters C so that each point $\boldsymbol{x}_j$ belongs to the cluster c_i with the nearest centroid $\boldsymbol{\mu}_i$.
- Typically assumes Euclidean distance as the similarity function s.

k-means

- Unsupervised variant of the Rocchio classifier.
- Goal: Partition the n observed objects into k clusters C so that each point $\boldsymbol{x}_j$ belongs to the cluster c_i with the nearest centroid $\boldsymbol{\mu}_i$.
- Typically assumes Euclidean distance as the similarity function s.
- The optimization problem: For each cluster, minimize the *within-cluster sum of squares*, $F_s = \mathrm{WCSS}$:

$$\mathrm{WCSS} = \sum_{c_i \in C} \sum_{\boldsymbol{x}_j \in c_i} \|\boldsymbol{x}_j - \boldsymbol{\mu}_i\|^2$$

- Equivalent to minimizing the average squared distance between objects and their cluster centroids (since n is fixed) – a measure of how well each centroid represents the members assigned to the cluster.

k-means (cont'd)

- Goal: Partition the n observed objects into k clusters C so that each point $\boldsymbol{x}_j$ belongs to the cluster c_i with the nearest centroid $\boldsymbol{\mu}_i$.

Algorithm

Initialize: Randomly select k centroid seeds.

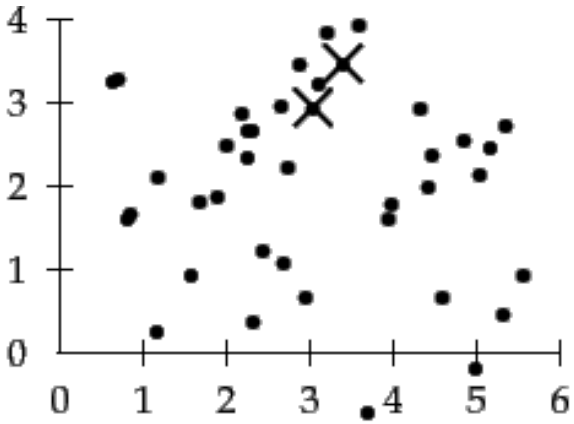
Iterate:

- Assign each object to the cluster with the nearest centroid.
- Compute new centroids for the clusters.

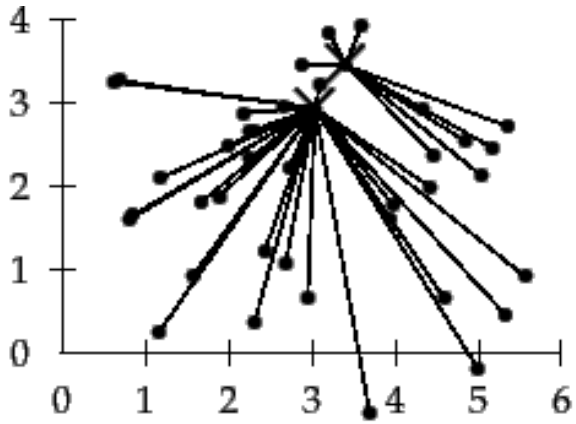
Terminate: When stopping criterion is satisfied.

- In short, we iteratively reassign memberships and recompute centroids until the configuration stabilizes.

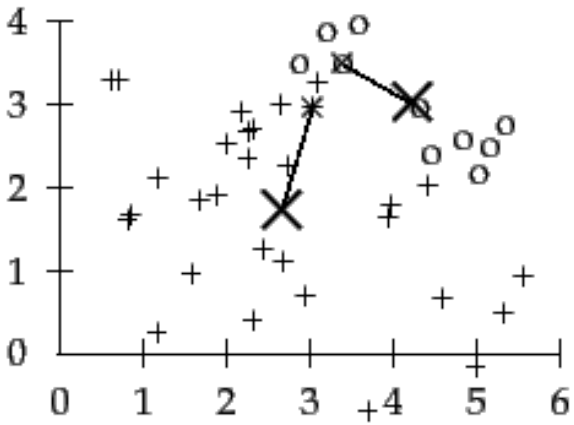
k-means example for $k = 2$ in R^2 (Manning, Raghavan & Schütze 2008)



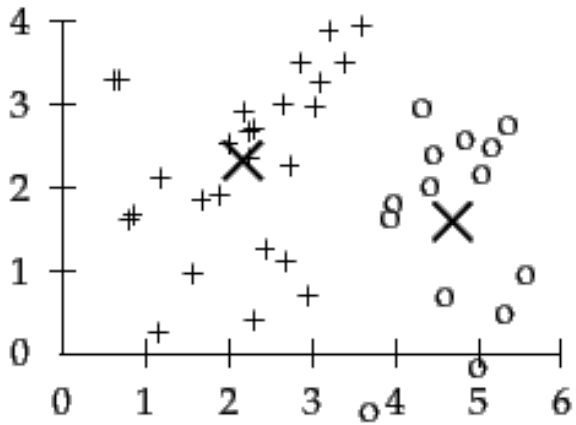
selection of seeds



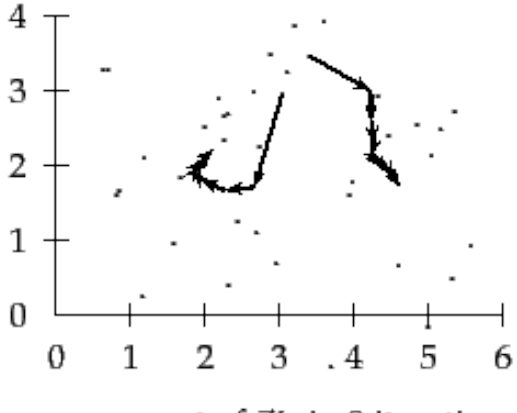
assignment of documents (iter. 1)



recomputation/movement of $\vec{\mu}$'s (iter. 1)



$\vec{\mu}$'s after convergence (iter. 9)



movement of $\vec{\mu}$'s in 9 iterations

Comments on k-means

Possible termination criteria

- Fixed number of iterations
- Clusters or centroids are unchanged between iterations.
- Threshold on the decrease of the objective function (absolute or relative to previous iteration)

Comments on k-means

Possible termination criteria

- Fixed number of iterations
- Clusters or centroids are unchanged between iterations.
- Threshold on the decrease of the objective function (absolute or relative to previous iteration)

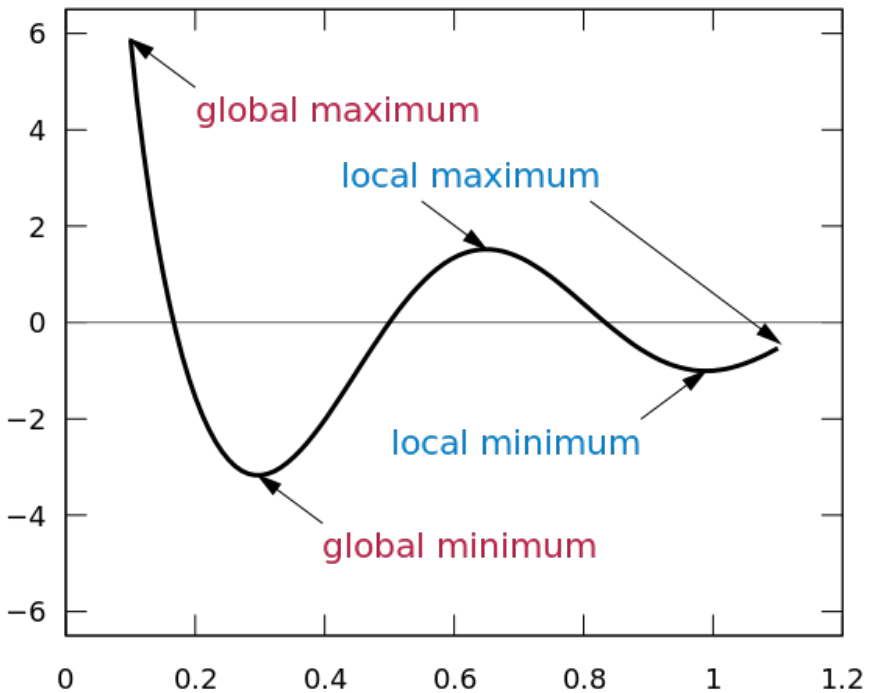
Some close relatives of k-means

- k-medoids: Like k-means but uses medoids instead of centroids to represent the cluster centers.
- Fuzzy c-means (FCM): Like k-means but assigns soft memberships in $[0, 1]$, where membership is a function of the centroid distance.
  - The computations of both WCSS and centroids are weighted by the membership function.

Properties of k-means

- The time complexity is linear, $\mathrm{O}(kn)$.
- WCSS is monotonically decreasing (or unchanged) for each iteration.

- Guaranteed to converge but not to find the global minimum.
- Possible solution: multiple random initializations.
- (k-means is non-deterministic)



Comments on k-means

'Seeding'

- We initialize the algorithm by choosing random seeds that we use to compute the first set of centroids, e.g:
  - pick k random objects from the collection;
  - pick k random points in the space;
  - pick k sets of m random points and compute centroids for each set; etc.
- The seeds can have a large impact on the resulting clustering.
- Outliers are troublemakers.

Flat Clustering: The good and the bad



Pros

- Conceptually simple, and easy to implement.
- Efficient. Typically linear in the number of objects.

Cons

- The dependence on random seeds as in k-means makes the clustering non-deterministic.
- The number of clusters k must be pre-specified. Often no principled means of *a priori* specifying k.
- Not as informative as the more structured clusterings produced by hierarchical methods.

- In general; often difficult to evaluate clustering.

Intrinsic evaluation of clustering

Evaluation With labeled gold data

- Cluster the examples in the gold test set (without using the labels).
- Compare the resulting clusters to the gold classes.
- Several available measures that focuses on the internal quality of the clusters, e.g. *purity*, measuring whether clusters tend to contain mostly items from a single class.

Intrinsic evaluation of clustering

Evaluation With labeled gold data

- Cluster the examples in the gold test set (without using the labels).
- Compare the resulting clusters to the gold classes.
- Several available measures that focuses on the internal quality of the clusters, e.g. *purity*, measuring whether clusters tend to contain mostly items from a single class.

Without

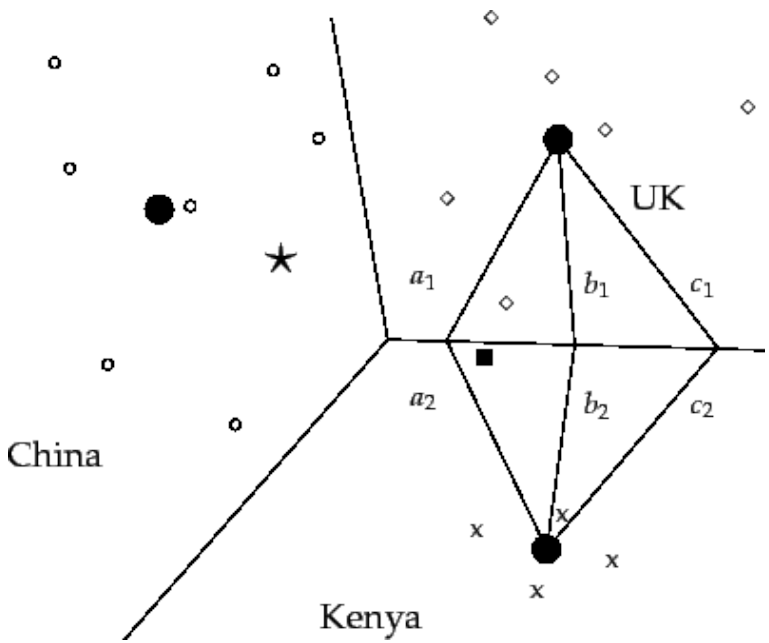
- Several measures that aim to measure the quality of a partitioning.
- E.g. Silhouette Coefficient – measures how similar an object is to its own cluster (cohesion) compared to other clusters (separation).
- But the scores are often sensitive to the particular dataset and difficult to interpret.

Extrinsic evaluation of clustering

- See which (or whether) clustering improves results in a downstream task
- E.g., for different versions of a recommender system, using topic clusters as features, measure
  - User satisfaction
  - How many recommended articles they read, or click on
  - Improvement in sales

Connecting the dots

- We have seen how Rocchio classification can be thought of as a 1-Nearest-Neighbor classification with respect to the centroids.
- Note how k-means clustering can be thought of as performing Rocchio classification in each iteration.



Next lecture

- Focus on *words* rather than *documents*.
- Distributional models of word meaning (lexical semantics).
- Semantic spaces: Vector space models of word meaning.
- Example tasks for evaluating word vectors.