

Recap: classification

Statistical Methods in NLP 2

ISCL-BA-08

Çağrı Çöltekin
ccolt@infsys.uni-tuebingen.de


University of Tübingen
Seminar für Sprachwissenschaft

Summer Semester 2026

version: 8/27/26 - 00222-01-04

Classification


- Given a training set with (categorical) labels/outcome
- Train a model to predict future data points from the same distribution



© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 1 / 22

The perceptron



$$y = f\left(\sum_{i=1}^n w_i x_i\right)$$

where

$$f(x) = \begin{cases} +1 & \text{if } \sum_{i=1}^n w_i x_i > 0 \\ -1 & \text{otherwise} \end{cases}$$

Note the similarity with linear regression: the separation between the (predicted) classes is a linear line/surface.

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 2 / 22

Learning with perceptron

- We do not update the parameters if classification is correct
- For misclassified examples, we try to minimize

$$E(w) = -\sum_i w x_i y_i$$

where i ranges over all misclassified examples

- Perceptron algorithm updates the weights for misclassified examples

$$w \leftarrow w - \eta \nabla E(w)$$


$$w \leftarrow w + \eta x_i y_i$$

- Perceptron algorithm converges if the classes are *linearly separable*

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 3 / 22

Logistic regression



- Logistic regression is probabilistic, we estimate $p = P(y = 1|x)$
- Note that typical regression would penalize correctly classified examples
- Instead we fit a regression model for $\logit(p) = w^T x + b$
- The probability estimate is the inverse of *logit*, the *logistic* (sigmoid) function


$$p = \frac{1}{1 + e^{-w^T x}}$$

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 4 / 22

Another example with two predictors

Probability assignments are non-linear, but the discriminant function is linear



© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 5 / 22

Fitting a logistic regression model

$$p = \frac{1}{1 + e^{-wx}}$$

The likelihood of the training set is,

$$\mathcal{L}(w) = \prod_i p^{y_i} (1-p)^{1-y_i}$$

In practice, we maximize log likelihood, or minimize '-log likelihood':

$$-\log \mathcal{L}(w) = -\sum_i y_i \log p + (1 - y_i) \log(1 - p)$$

- The derivative / gradient is easy (a good exercise)
- There is no analytic solution for $\nabla -\log \mathcal{L}(w) = 0$
- But the loss function is convex: we can find the global minimum with gradient descent (most of the time)

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 6 / 22

Naive Bayes

- Naive Bayes is another probabilistic classification method
- In any classification task, our aim is to find

$$\hat{y} = \arg \max_y P(y|x)$$

- Sometimes (with some simplifying assumptions) it is easier to predict $P(x|y)$
- Then we use Bayes' formula to invert the conditional probability


$$\hat{y} = \arg \max_y \frac{P(x|y)P(y)}{P(x)} = \arg \max_y P(x|y)P(y)$$

- $P(x|y)$ and $P(y)$ are generally estimated using MLE (with *smoothing*)

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 7 / 22

Maximum-margin methods (e.g., SVMs)




- In perceptron, we stop whenever we found a linear discriminator
- Maximum-margin classifiers seek a discriminator that maximizes the margin
- SVMs have other interesting properties, and they have been one of the best 'out-of-the-box' classifiers for many problems

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 8 / 22

Maximum-margin methods (e.g., SVMs)




- In perceptron, we stop whenever we found a linear discriminator
- Maximum-margin classifiers seek a discriminator that maximizes the margin
- SVMs have other interesting properties, and they have been one of the best 'out-of-the-box' classifiers for many problems

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 9 / 22

Decision trees




- Note that the decision boundary is non-linear

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 10 / 22

Instance/memory based methods



- No training: just memorize the instances
- During test time, decide based on the k nearest neighbors
- Like decision trees, k NN is non-linear
- It can also be used for regression

© Çöltekin, ISL | University of Tübingen

Summer Semester 2026 11 / 22

More than two classes

- Some algorithms can naturally be extended to handle multiple class labels
- Any binary classifier can be turned into a k-way classifier by
 - One-vs-rest or one-vs-all
 - train k classifiers: each learns to discriminate one of the classes from the others
 - at prediction time the classifier with the highest confidence wins
 - needs a confidence score from the base classifiers
 - One-vs-one
 - train $\frac{k(k-1)}{2}$ classifiers: each learns to discriminate a pair of classes
 - decision is made by (weighted) majority vote
 - works without need for confidence scores, but needs more classifiers

Measuring success in classification

Accuracy, precision, recall, F-score

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{precision} = \frac{TP}{TP + FP}$$

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

$$F_1\text{-score} = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$$

	predicted	
	positive	negative
true value	pos. TP	FN
	neg. FP	TN

Multi-class evaluation

- For multi-class problems, it is common to report average precision/recall/f-score
- For C classes, averaging can be done two ways:

$$\text{precision}_{\text{M}} = \frac{\sum_{i=1}^C \frac{TP_i}{FP_i + TP_i}}{C} \quad \text{recall}_{\text{M}} = \frac{\sum_{i=1}^C \frac{TP_i}{TP_i + FN_i}}{C}$$

$$\text{precision}_{\mu} = \frac{\sum_{i=1}^C TP_i}{\sum_{i=1}^C TP_i + FP_i} \quad \text{recall}_{\mu} = \frac{\sum_{i=1}^C TP_i}{\sum_{i=1}^C TP_i + FN_i}$$

- (M = macro, μ = micro)
- The averaging can also be useful for binary classification, if there is no natural positive class

Confusion matrix

- A confusion matrix is often useful for multi-class classification tasks

	predicted		
	negative	neutral	positive
true value	negative 10	2	0
	neutral 3	12	7
	positive 4	8	7

- Are the classes balanced?
- What is the accuracy?
- What is per-class, and averaged precision/recall?

Overfitting & Underfitting

We want our models to *generalize*, perform well on unseen data.


- Overfitting* occurs when the model learns the idiosyncrasies of the training data
- Underfitting* occurs when the model is not flexible enough for solving the problem at hand

We want simpler models, but not too simple for the task at hand.

Causes of overfitting

(1) Model complexity

Quadratic 6th degree




Causes of overfitting

(1) Noise

With noise

No noise




Causes of overfitting

(1) Lack of training data

5 training instances

10 training instances



Preventing overfitting

- Add more training data
- Increase data quality
- Lower model complexity
 - Simpler models
 - Regularization
 - For iterative methods, stop early
- In most cases, selecting the best model requires *hyperparameter* tuning / optimization.
- Always check the model performance on a held-out (development / validation) set.
- K-fold cross-validation is a method for efficient use of available labeled data

Bias and variance

Bias of an estimate is the difference between the value being estimated, and the expected value of the estimate

$$B(\hat{w}) = E[\hat{w}] - w$$

- An unbiased estimator has 0 bias
- Variance* of an estimate is, simply its variance, the value of the squared deviations from the mean estimate

$$\text{var}(\hat{w}) = E[(\hat{w} - E[\hat{w}])^2]$$

w is the parameter (vector) that defines the model

Bias-variance relationship is a trade-off: models with low bias result in high variance.

ML evaluation in general

The first principle is that you must not fool yourself and you are the easiest person to fool. — Richard P. Feynman

- We want models with low bias and low variance, but this is a trade-off
- Estimators/models with high variance (hence, low/no bias) are likely to *overfit*
- Estimators/models with low variance (hence, high bias) are likely to *underfit*
- Evaluating ML system requires special care:
 - Tuning your system on a development set
 - Cross-validation allows efficient use of labeled data during tuning
 - A test set is often used when comparing results obtained by different models
- In most cases, the scores are not meaningful by themselves, we need to compare with *baselines*

Final remarks

- Most NLP problems we try to solve are classification problems
- We reviewed some of the 'traditional' classification methods
- Understanding them will help understanding more 'modern' methods
- The models we review can serve as baselines for more complex models, and sometimes their performance may surprise you

Next:

Mon/Wed Gradient descent, reading: Jurafsky and Martin (2026, Section 4.6)
Fri Lab: numpy tutorial

Some sources of information

- Any ML textbook covers most of the methods reviewed (and more), here are a few: James et al. (2024), Bishop (2006), and MacKay (2003)

- 1 Bishop, Christopher M. (2006). *Pattern Recognition and Machine Learning*. Springer. isbn: 978-0387-31073-2.
- 2 James, Gareth, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor (2024). *An introduction to statistical learning*. Springer. isbn: 9783031391897. url: <https://www.statlearning.com/>
- 3 Jurafsky, Daniel and James H. Martin (2026). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models*. 3rd. Online manuscript released January 6, 2026. url: <https://web.stanford.edu/~jurafsky/slp3/>.

Some sources of information (cont.)

- 4 MacKay, David J. C. (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge University Press. isbn: 978-05-2164-298-9. url: <http://www.inference.phy.cam.ac.uk/itprn/book.html>.