

Introduction to artificial neural networks

Statistical Methods in NLP 2

ISCL-BA-08

Çağrı Çöltekin
ccoltekin@fsf.uni-tuebingen.de

University of Tübingen
Seminar für Sprachwissenschaft

Summer Semester 2026

version: iscl-001-000000-00-00

Artificial neural networks

- Artificial neural networks (ANNs) are machine learning models inspired by biological neural networks
- ANNs are powerful non-linear models
- Power comes with a price: there are no guarantees of finding the global minimum of the error function
- ANNs have been used in ML, AI, Cognitive science since 1950's ~ with some ups and downs
- Currently they are the driving force behind the popular 'deep learning' methods


Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 1 / 31

Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

The biological neuron

(showing a picture of a real neuron is mandatory in every ANN lecture)



"Neuron neuron" Wikipedia

Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 2 / 31

Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

Artificial and biological neural networks

- ANNs are *inspired* by biological neural networks
- Similar to biological networks, ANNs are made of many simple processing units
- Despite the similarities, there are many differences: ANNs do not mimic biological networks
- ANNs are practical statistical machine learning methods

Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 3 / 31

Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs


Recap: the perceptron

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

where

$$f(x) = \begin{cases} +1 & \text{if } wx > 0 \\ -1 & \text{otherwise} \end{cases}$$

In ANN-speak $f(\cdot)$ is called an *activation function*.



Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 4 / 31


Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

Recap: logistic regression

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

where

$$f(x) = \frac{1}{1 + e^{-wx}}$$



Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

Linear separability

- A classification problem is said to be *linearly separable* if one can find a linear discriminator
- A well-known counter example is the logical XOR problem



There is no line that can separate positive and negative classes.

Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 5 / 31

Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

Can a linear classifier learn the XOR problem?

- We can use non-linear basis functions

$$w_0 + w_1 x_1 + w_2 x_2 + w_3 \phi(x_1, x_2)$$

is still linear in w for any choice of $\phi(\cdot)$

- For example, adding the product $x_1 x_2$ as an additional feature would allow a solution like: $x_1 + x_2 - 2x_1 x_2$


x_1	x_2	$x_1 + x_2 - 2x_1 x_2$
0	0	0
0	1	1
1	0	1
1	1	0

- Choosing proper basis functions like $x_1 x_2$ is called *feature engineering*

Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

Non-linear basis functions

solution in the original input space



The solution to

$$x_1 + x_2 - 2x_1 x_2 = 0.5 = 0$$

is a (non-linear) discriminant that solves the problem


Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 6 / 31

Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

Non-linear basis functions

solution in the 3D input space



- The additional basis function maps the problem into 3D
- In the new, mapped space, the points are linearly separable

Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 7 / 31

Where do non-linearities come from?

non-linearities are abundant in nature, it is not only the XOR problem

In a linear model, $y = w_0 + w_1 x_1 + \dots + w_n x_n$

- The outcome is *linearly-related* to the predictors
- The effects of the inputs are *additive*

This is not always the case:

- Some predictors affect the outcome in a non-linear way
 - The effect may be strong or positive only in a certain range of the variable (e.g., reaction time change by age)
 - Some effects are periodic (e.g., many measures of time)
- Some predictors interact
 - 'not bad' is not 'not' + 'bad' (e.g., for sentiment analysis)

Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 10 / 31


Introduction Plain-Neurocity MLP Plain-Neurocity and MLP Learning in ANNs

Multi-layer perceptron

- The simplest modern ANN architecture is called multi-layer perceptron (MLP)
- The MLP is a *fully connected, feed-forward* network consisting of perceptron-like units
- Unlike perceptron, the units in an MLP use a continuous activation function
- The MLP can be trained using gradient-based methods
- The MLP can represent many interesting machine learning problems
 - It can be used for both regression and classification


Ç. Çöltekin, ISL | University of Tübingen

Summer Semester 2026 11 / 31

Multi-layer perceptron
the picture

Each unit takes a weighted sum of their input, and applies a (non-linear) activation function.

Artificial neurons



- The unit calculates a weighted sum of the inputs


$$\sum_j w_j x_j = \mathbf{W}\mathbf{x}$$

- result is a linear transformation
- a non-linear activation function $f(\cdot)$ applied to the result
- Output is

$$\mathbf{y} = \mathbf{f}(\mathbf{W}\mathbf{x})$$

Artificial neurons

an example



- A common activation function is (was?) the logistic sigmoid function

$$f(x) = \frac{1}{1 + e^{-x}}$$

- The output of the network becomes

$$\mathbf{y} = \frac{1}{1 + e^{-\mathbf{W}\mathbf{x}}}$$

Activation functions in ANNs

hidden units

- The activation functions in MLP are typically continuous (differentiable) functions
- For hidden units common choices are



Sigmoid (logistic)

$$\frac{1}{1 + e^{-x}}$$



Hyperbolic tangent (tanh)

$$\frac{e^{2x} - 1}{e^{2x} + 1}$$

Rectified linear unit (relu) $\max(0, x)$

Activation functions in ANNs

output units


- The activation functions of the output units depends on the task. Common choices are
 - For regression, the identity function ($y = x$)
 - For binary classification, logistic sigmoid

$$P(y = 1 | x) = \frac{1}{1 + e^{-\mathbf{W}\mathbf{x}}} = \frac{e^{\mathbf{W}\mathbf{x}}}{1 + e^{\mathbf{W}\mathbf{x}}}$$

- For multi-class classification, softmax

$$P(y = k | x) = \frac{e^{\mathbf{W}_k \mathbf{x}}}{\sum_j e^{\mathbf{W}_j \mathbf{x}}}$$

MLP: a simple example




$$\mathbf{h}_1 = f\left(\sum_i w_{1i}^{(1)} x_i\right)$$

$$\mathbf{y}_k = g\left(\sum_j w_{jk}^{(2)} \mathbf{h}_j\right)$$

$$\mathbf{y}_k = g\left(\sum_j w_{jk}^{(2)} f\left(\sum_i w_{ij}^{(1)} x_i\right)\right)$$

MLP: a simple example



- Alternatively, we can write the computations in matrix form

$$\mathbf{h} = f(\mathbf{W}^{(1)} \mathbf{x})$$

$$\mathbf{y} = g(\mathbf{W}^{(2)} \mathbf{h})$$

$$= g\left(\mathbf{W}^{(2)} f(\mathbf{W}^{(1)} \mathbf{x})\right)$$

- This corresponds to a series of linear transformations followed by elementwise (non-linear) function applications


Solving non-linear problems with ANNs

a solution to XOR problem




Solving non-linear problems with ANNs

a solution to XOR problem




Solving non-linear problems with ANNs

a solution to XOR problem





Solving non-linear problems with ANNs

a solution to XOR problem



Solving non-linear problems with ANNs


a solution to XOR problem



Is this different from non-linear basis functions?

Non-linear activation functions are necessary

Without non-linear activation functions, an ANN with any number of layers is equivalent to a linear model.



$$\begin{aligned}h_1 &= ax_1 + cx_2 \\h_2 &= bx_1 + dx_2 \\y &= eh_1 + fh_2 \\&= (ea + fb)x_1 + (ec + fd)x_2\end{aligned}$$

y is still a linear function of x_1


Gradient descent: a refresher

- The general idea is to approach a minimum of the error function in small (but not so small) steps

$$\mathbf{w} \leftarrow \mathbf{w} - \eta \nabla J(\mathbf{w})$$

- ∇J is the gradient of the loss function, it points to the direction of the maximum increase
 - η is the learning rate
- The updates can be performed
 - batch: for the complete training set
 - on-line: after every training instance
 - this is known as *stochastic gradient descent* (SGD)
 - mini-batch: after small fixed-sized batches

ANN objectives are not convex



Error functions in ANN training

depend on the task

- For regression, a natural choice is minimizing the sum of squared error

$$E(\mathbf{w}) = \sum_i (y_i - \hat{y}_i)^2$$

- For binary classification, we use *cross entropy*

$$E(\mathbf{w}) = - \sum_i y_i \log \hat{y}_i + (1 - y_i) \log (1 - \hat{y}_i)$$

- Similarly, for multi-class classification, also cross entropy

$$E(\mathbf{w}) = - \sum_{i,k} y_{i,k} \log \hat{y}_{i,k}$$


In practice, the ANN loss functions will not be convex.

Learning in ANNs

- ANNs implement complex functions: we need to use optimization methods (e.g., gradient descent) to train them
- Typically error functions for ANNs are not convex, gradient descent will find a local minimum
- Optimization requires updating multiple layers of weights
- Assigning credit (or blame) to each weight during learning is not trivial
- An effective solution to the last problem is the *backpropagation* algorithm

Learning in multi-layer networks: the problem


Introduction Non-linearity and MLP Non-linearity and MLP Learning in ANNs



We want a way to (efficiently) update non-final weights based on final error.

Calculating gradient on a neural network

(with some simplification)



- We need to calculate the gradient:

$$\nabla E = \left(\frac{\partial E}{\partial a}, \frac{\partial E}{\partial b}, \frac{\partial E}{\partial c}, \frac{\partial E}{\partial d}, \frac{\partial E}{\partial e}, \frac{\partial E}{\partial f} \right)$$


- we can use gradient descent directly
- $\frac{\partial E}{\partial a}$ and $\frac{\partial E}{\partial b}$ is easy, they do not depend on other variables

- We factor others using chain rule

$$\frac{\partial E}{\partial a} = \frac{\partial h_1}{\partial a} \frac{\partial E}{\partial h_1} \quad \text{and} \quad \frac{\partial E}{\partial c} = \frac{\partial h_1}{\partial c} \frac{\partial E}{\partial h_1}$$

Backpropagation

Introduction Non-linearity and MLP Non-linearity and MLP Learning in ANNs



- So far, it is just math

$$\frac{\partial E}{\partial a} = \frac{\partial h_1}{\partial a} \frac{\partial E}{\partial h_1} \quad \text{and} \quad \frac{\partial E}{\partial c} = \frac{\partial h_1}{\partial c} \frac{\partial E}{\partial h_1}$$

- But a naive implementation does many repeated calculations
- Backpropagation is an efficient (dynamic programming) algorithm that avoids repeated calculations
- Backpropagation works for any computation graph without cycles

Preventing overfitting in neural networks

- As in linear models, we can use L1 and L2 regularization by adding a regularization term to the error function (known as *weight decay*). For example,

$$J(\mathbf{w}) = E(\mathbf{w}) + \|\mathbf{w}\|$$

- There are other ways to fight overfitting
 - With *early stopping*, one stops the training before it reaches to the smallest training error
 - With *dropout*, random units (with all of their connections) are dropped during training
 - Injecting noise at the output, as a way to (implicitly) model the noise in the target classes/values

How many layers, units

Introduction Non-linearity and MLP Non-linearity and MLP Learning in ANNs

- A network with single hidden layer is said to be a *universal approximator*: it can approximate any continuous function with arbitrary precision
- However, in practice multiple interconnected layers are useful and commonly used in modern ANN models
- The choice of layers, in general the architecture of the system, depends on the application

A bit of history

1950-60 ANNs (perceptron) became popular:

lots of excitement in AI, cognitive science

1970s Not much interest

criticism on perceptron: linear separability

1980s ANNs became popular again

– backpropagation algorithm

– multi-layer networks

1990s ANNs had again fallen 'out of fashion'

– Engineering: other algorithms (such as SVMs) performed generally better

– From the cognitive science perspective: ANNs are difficult to interpret

present ANNs (again) enjoy a renewed popularity with the name 'deep learning'

Summary

Introduction Non-linearity and MLP Non-linearity and MLP Learning in ANNs

- ANNs are powerful non-linear learners
 - based on some inspiration from biological NNs
 - using many simple processing units
 - built on linear models (logistic regression)
 - For non-linear problems we need non-linear activation functions, and at least one hidden layer
 - ANNs can be used for both regression and classification
 - In general, ANN loss functions are not convex, what we find is a local minimum
 - They (typically) are trained with *backpropagation* algorithm
 - Reading: Jurafsky and Martin (2026, Chapter 7)
- Next:
- Models for sequential data, reading: Jurafsky and Martin (2026, Chapter 17)

Additional reading, references, credits

- Hastie, Tibshirani, and Friedman (2009, Chapter. 11) also includes an accessible introduction
- For a review of use of ANNs in NLP, including more advanced topics, see Goldberg, 2016

Additional reading, references, credits (cont.)

Goldberg, Yoav (2016). "A primer on neural network models for natural language processing". In: *Journal of Artificial Intelligence Research* 57, pp. 343–420.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Second. Springer series in statistics. Springer-Verlag New York, Inc. ISBN 9780387848587. URL: <http://web.stanford.edu/~hastie/ElemStatLearn/>

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/>

References & further reading

Goldberg, Yoav (2016). "A primer on neural network models for natural language processing". In: *Journal of Artificial Intelligence Research* 57, pp. 343–420.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Second. Springer series in statistics. Springer-Verlag New York, Inc. ISBN 9780387848587. URL: <http://web.stanford.edu/~hastie/ElemStatLearn/>

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/>

