

# CS145: INTRODUCTION TO DATA MINING

## 6: Vector Data: Neural Network

---

**Instructor: Yizhou Sun**

[yzsun@cs.ucla.edu](mailto:yzsun@cs.ucla.edu)


October 7, 2021

# Methods to Learn

	Vector Data	Set Data	Sequence Data/Time Series	Text Data	Graph Data
Classification	Logistic Regression; Decision Tree; <b>NN</b>			Naïve Bayes for Text	Label Propagation
Clustering	K-means; Mixture Models			PLSA	Spectral Clustering
Prediction	Linear Regression; Regression Tree; <b>NN</b> GLM*		AR Model		
Frequent Pattern Mining		Apriori; FP growth	GSP; PrefixSpan		
Similarity Search			DTW		P-PageRank
Ranking					PageRank

# Neural Network

---

- Introduction 
- Connection to Shallow Machine Learning Algorithms
- Multi-Layer Feed-Forward Neural Network
- Deep Learning
- Summary

# A little bit of History

## Beginnings

Thresholded Logic Unit

1943

Perceptron

1957

Adaline

1960

1940

1950

1960





S. McCulloch - W. Pitts



R. Rosenblatt



B. Widrow - M. Hoff



## 1st Neural Winter


XOR Problem

1969

1970



M. Minsky - S. Papert



## 2nd Neural Winter

Multilayer Backprop

1982

CNNs


1986

LSTMs

1989

SVMs

1997



P. Werbos  
D. Rumelhart - G. Hinton - R. Williams



Y. Lecun  
J. Schmidhuber

## GPU Era

Deep Nets

2006

Alex Net

2012

2000

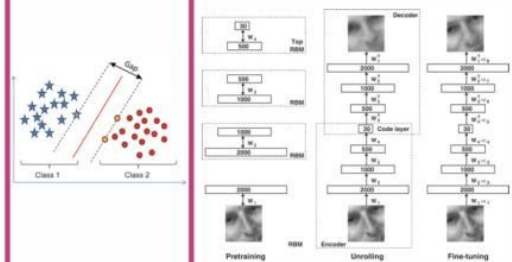
2010




C. Cortes - V. Vapnik



R. Salakhutdinov - J. Hinton - A. Krizhevsky - I. Sutskever



# ImageNet Challenge



<https://semiengineering.com/new-vision-technologies-for-real-world-applications/>

# AlphaGo

## *Master of Go Board Game Is Walloped by Google Computer Program*



By [Choe Sang-Hun](#) and [John Markoff](#)

March 9, 2016



SEOUL, South Korea — Computer, one. Human, zero.


A [Google](#) computer program stunned one of the world's top players on Wednesday in a round of Go, which is believed to be the most complex board game ever created.

# AlphaFold

---

## *A.I. Predicts the Shapes of Molecules to Come*

DeepMind has given 3-D structure to 350,000 proteins, including every one made by humans, promising a boon for medicine and drug design.



Proteins belonging to a fruit fly are given shape by DeepMind's A.I. technology. The company's new database includes the 3-D structures for all proteins expressed by the human genome as well as those appearing in 20 other organisms. DeepMind

# GPT-3



The New York Times

Account ▾

## *Meet GPT-3. It Has Learned to Code (and Blog and Argue).*

The latest natural-language system generates tweets, pens poetry, summarizes emails, answers trivia questions, translates languages and even writes its own computer programs.



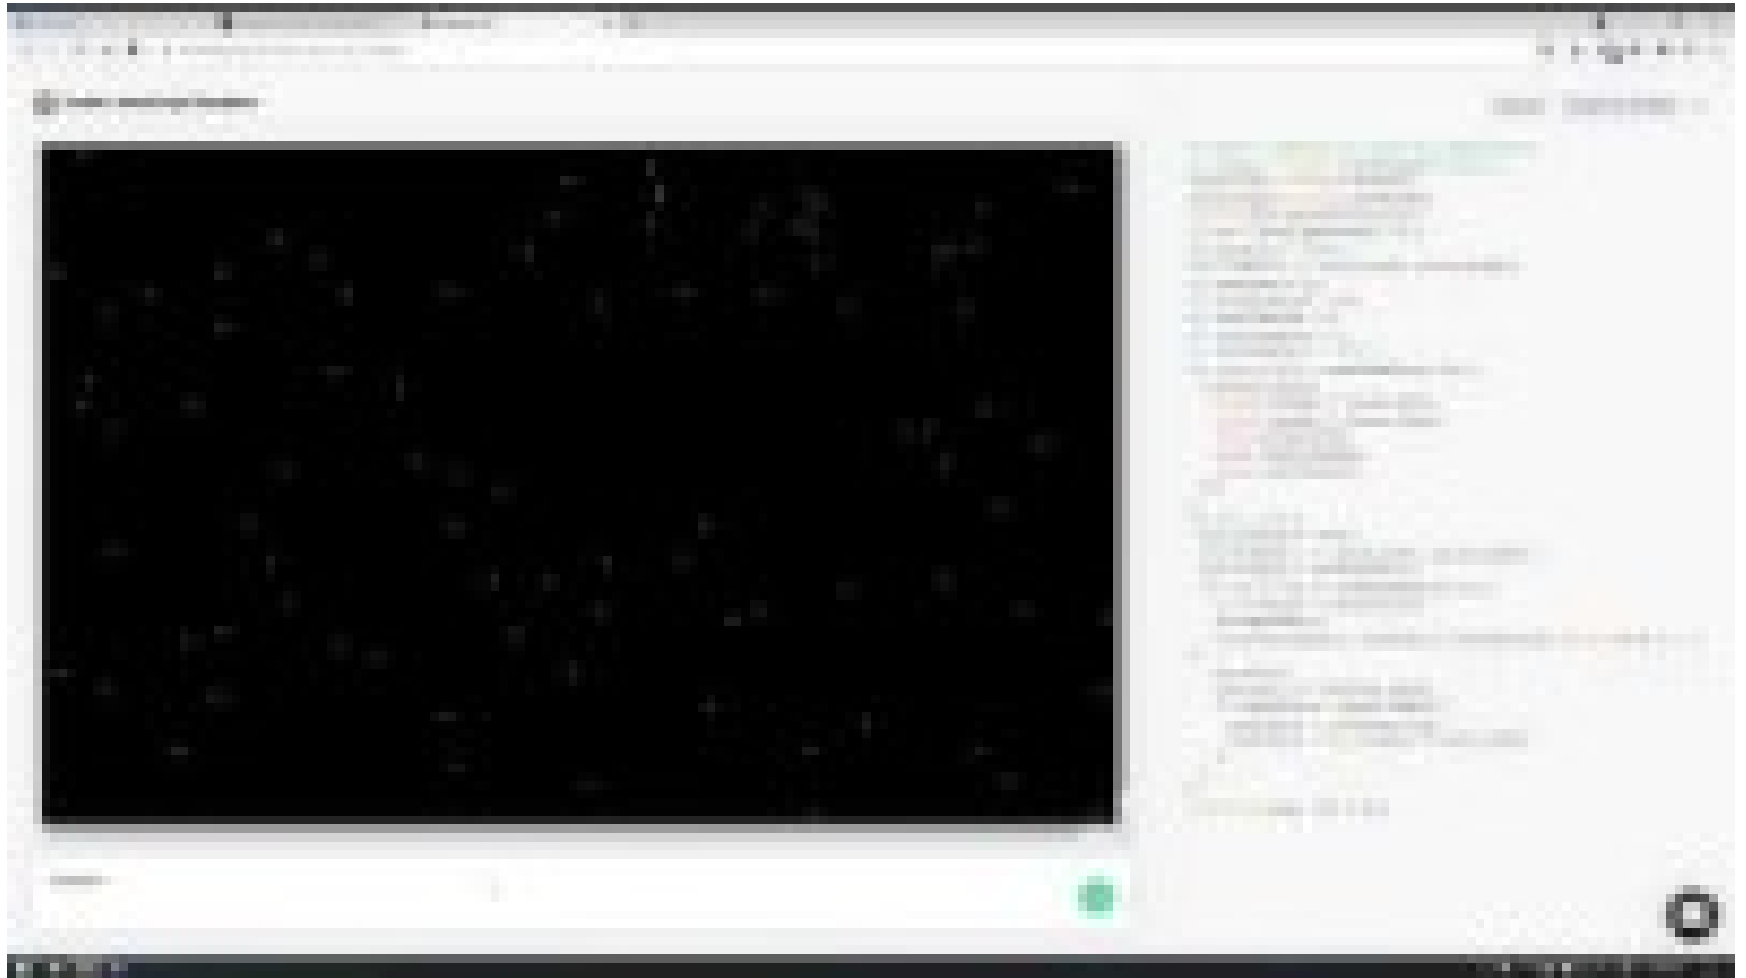
By Cade Metz

Nov. 24, 2020

This summer, an artificial intelligence lab in San Francisco called OpenAI unveiled a technology several months in the making. This new system, GPT-3, had spent the previous months learning the ins and

# OpenAI: Codex

---

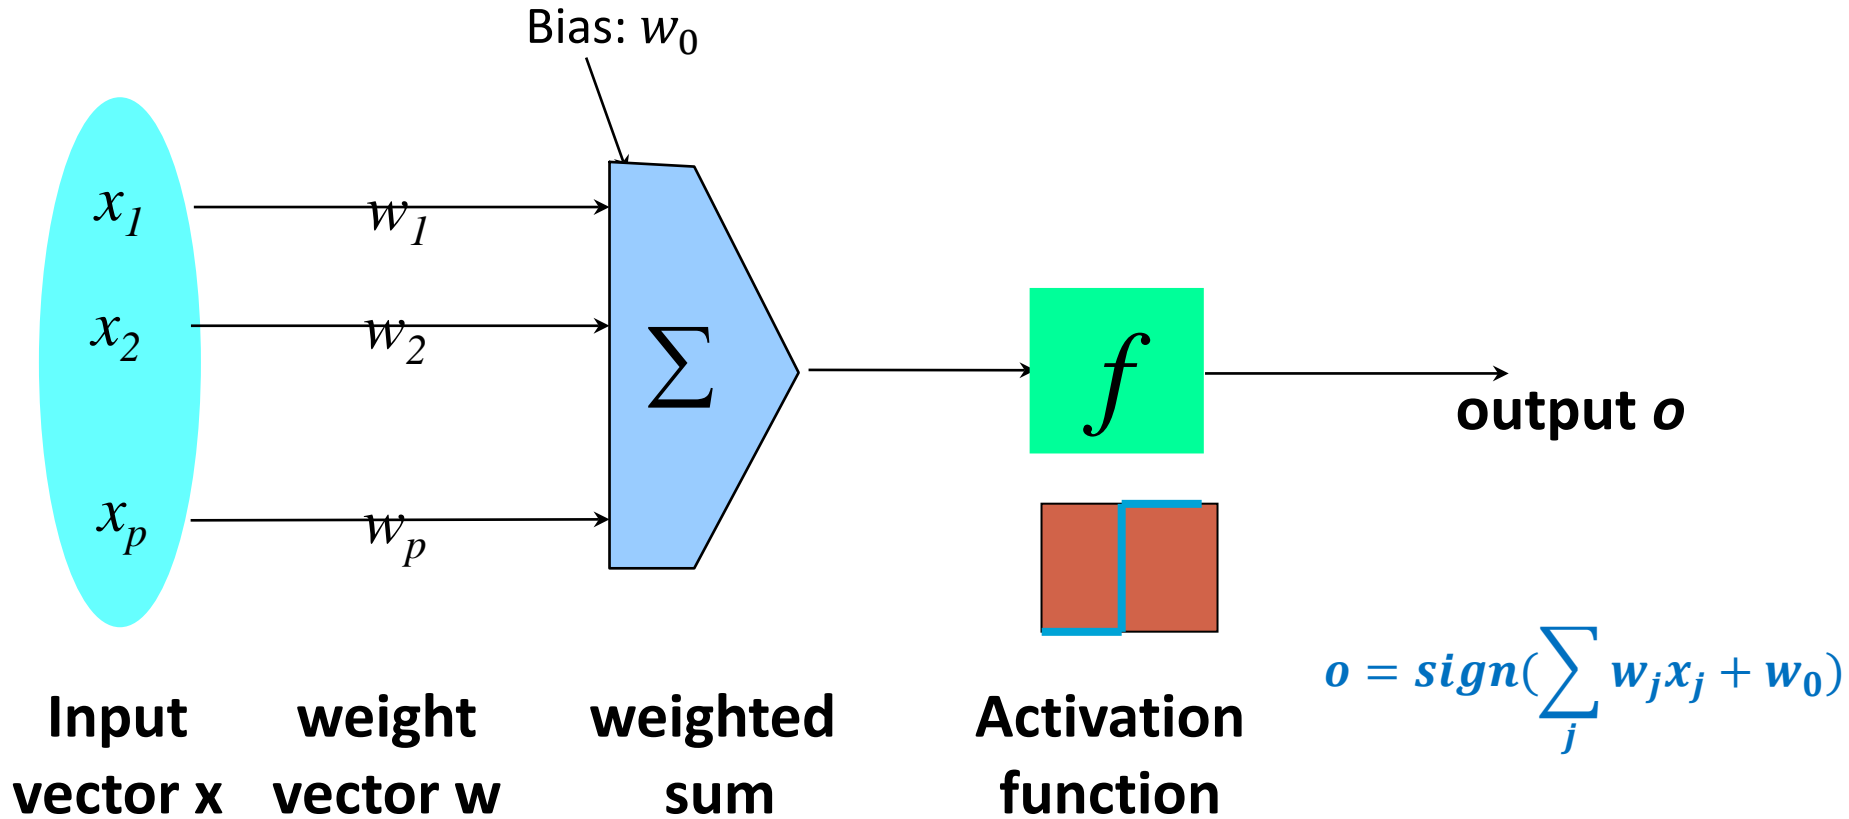


# Artificial Neural Networks

---

- Consider humans:
  - Number of neurons  $\sim 10^{10}$
  - Connections per neuron  $\sim 10^{4-5}$
  - Neuron switching time  $\sim .001$  second
  - Scene recognition time  $\sim .1$  second
  - 100 inference steps doesn't seem like enough -> parallel computation
- Artificial neural networks
  - Many neuron-like threshold switching units
  - Many weighted interconnections among units
  - Highly parallel, distributed process
  - Emphasis on tuning weights automatically

# Single Unit: Perceptron



# Important Concepts

---

- Architecture
- Activation function
- Loss function
- Optimization
- Regularization

# Architecture

---


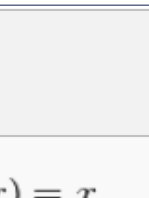
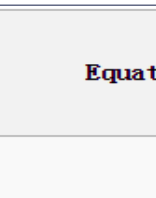
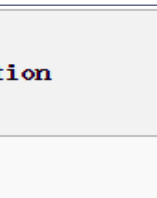
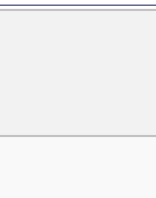


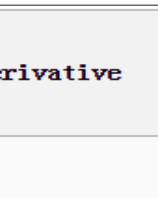
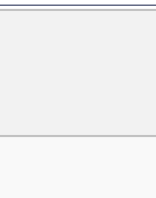
- Decide the **network topology**:
  - Specify # of units in the *input layer*, # of *hidden layers* (if  $> 1$ ), # of units in *each hidden layer*, the unit types, connection between layers, and # of units in the *output layer*
- Architecture specifies the function that maps input to output, which contains parameters to be learned

# Activation function

---

- An activation function  $f(\cdot)$  in the output layer can control the nature of the output (e.g., probability value in  $[0, 1]$ )
- Activation functions bring **nonlinearity** into hidden layers, which increases the complexity of the model.
- Good activation functions should be **differentiable** for optimization purpose

# Examples of Activation Functions

Name	Plot	Equation	Derivative
Identity		$f(x) = x$	$f'(x) = 1$
Binary step		$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} 0 & \text{for } x \neq 0 \\ ? & \text{for } x = 0 \end{cases}$
Logistic (a.k.a. Soft step)		$f(x) = \frac{1}{1 + e^{-x}}$	$f'(x) = f(x)(1 - f(x))$
TanH		$f(x) = \tanh(x) = \frac{2}{1 + e^{-2x}} - 1$	$f'(x) = 1 - f(x)^2$
ArcTan		$f(x) = \tan^{-1}(x)$	$f'(x) = \frac{1}{x^2 + 1}$
Rectified Linear Unit (ReLU)		$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$
Parameteric Rectified Linear Unit (PReLU) [2]		$f(x) = \begin{cases} \alpha x & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} \alpha & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$
Exponential Linear Unit (ELU) [3]		$f(x) = \begin{cases} \alpha(e^x - 1) & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} f(x) + \alpha & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$
SoftPlus		$f(x) = \log_e(1 + e^x)$	$f'(x) = \frac{1}{1 + e^{-x}}$

# Loss Functions

---

- How good are the outputs compared with the labels (target)?
  - Empirical risk
    - $\mathcal{L}(\mathbf{w}) = \frac{1}{n} \sum_i l(y^{(i)}, \hat{y}^{(i)})$ , where  $\hat{y}^{(i)} = f(\mathbf{x}^{(i)}, \mathbf{w})$
    - $\mathbf{w}$ : parameters in the model
  - Loss function: difference between actual value and predicted value
    - $l(y, \hat{y})$

# Examples of Loss Functions

---

- Squared error
  - $l(y, \hat{y}) = (y - \hat{y})^2$
- (Binary) cross entropy loss
  - $l(y, \hat{y}) = -y \log \hat{y} - (1 - y) \log(1 - \hat{y})$
  - $y \in \{0, 1\}, \hat{y} \in [0, 1]$
- Hinge loss
  - $l(y, \hat{y}) = \max(0, 1 - y\hat{y})$
  - $y \in \{-1, 1\}, \hat{y} \in (-\infty, +\infty)$

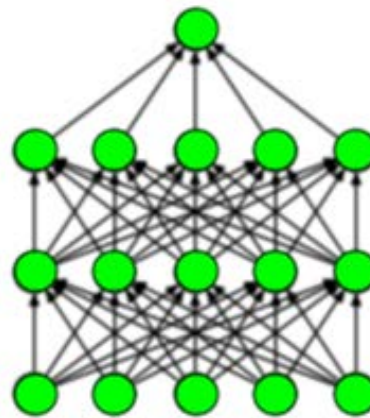
# Optimization

---


- Given a training dataset, minimize the empirical risk
  - Find  $\mathbf{w}$ , such that  $\mathcal{L}(\mathbf{w}) = \frac{1}{n} \sum_i l(y^{(i)}, \hat{y}^{(i)})$  is minimized
- Solution:
  - Stochastic gradient descent + chain rule = backpropagation

# Regularization

- Avoid overfitting
- Techniques
  - L2/L1 regularization
  - Dropout
  - Early stopping
  - ...



(a) Standard Neural Net




(b) After applying dropout.

<https://medium.com/analytics-vidhya/a-simple-introduction-to-dropout-regularization-with-code-5279489dda1e>

# Neural Network

---

- Introduction
- Connection to Shallow Machine Learning Algorithms 
- Multi-Layer Feed-Forward Neural Network
- Deep Learning
- Summary

# Perceptron

---

- Architecture:
  - A single neuron
- Activation function
  - Training: identity function
  - Inference: sign function/step function
- Loss function
  - $l(y, \hat{y}) = \max(0, -y\hat{y})$
- Optimization
  - $\mathbf{w} \leftarrow \mathbf{w} + \eta y^{(i)} \mathbf{x}^{(i)}$ , for a misclassified training data point  $(\mathbf{x}^{(i)}, y^{(i)})$ , i.e.,  $y^{(i)} \mathbf{w}^T \mathbf{x}^{(i)} \leq 0$
  - $\eta$ : learning rate


# Example: 1 for “Y” and -1 for “N”;

$$\eta = 0.9$$

x0	x1	x2	true label	w before update	predicted label	w after update
1	0	1	Y	(0.0, 0.0, 0.0)	N	(0.9, 0.0, 0.9)
1	1	1	N	(0.9, 0.0, 0.9)	Y	(0.0, -0.9, 0.0)
1	0	0	Y	(0.0, -0.9, 0.0)	N	(0.9, -0.9, 0.0)
1	1	0	Y	(0.9, -0.9, 0.0)	N	(1.8, 0.0, 0.0)
1	0	1	Y	(1.8, 0.0, 0.0)	Y	(1.8, 0.0, 0.0)
1	1	1	N	(1.8, 0.0, 0.0)	Y	(0.9, -0.9, -0.9)
1	0	0	Y	(0.9, -0.9, -0.9)	Y	(0.9, -0.9, -0.9)
1	1	0	Y	(0.9, -0.9, -0.9)	N	(1.8, 0.0, -0.9)
1	0	1	Y	(1.8, 0.0, -0.9)	Y	(1.8, 0.0, -0.9)
1	1	1	N	(1.8, 0.0, -0.9)	Y	(0.9, -0.9, -1.8)
1	0	0	Y	(0.9, -0.9, -1.8)	Y	(0.9, -0.9, -1.8)
1	1	0	Y	(0.9, -0.9, -1.8)	N	(1.8, 0.0, -1.8)

# Linear SVM

- Architecture:
  - A single neuron
- Activation function
  - Training: identity function
  - Inference: sign function/step function
- Loss function
  - $l(y, \hat{y}) = \max(0, 1 - y\hat{y})$
- Regularization
  - L2, i.e.,  $\frac{1}{2} \lambda \|\mathbf{w}\|^2$
- Optimization
  - $\mathbf{w} \leftarrow \mathbf{w} - \eta(\lambda \mathbf{w} - y^{(i)} \mathbf{x}^{(i)})$ , for a misclassified or barely correct training data point  $(\mathbf{x}^{(i)}, y^{(i)})$ , i.e.,  $y^{(i)} \mathbf{w}^T \mathbf{x}^{(i)} < 1$
  - $\mathbf{w} \leftarrow \mathbf{w} - \eta \lambda \mathbf{w}$ , for a confidently correct training data point
  - $\eta$ : learning rate




# Perceptron V.S. Linear SVM

- Source:

<http://www.charuaggarwal.net/neural.htm>


MARGINALLY  
CORRECT PREDICTION



$$\bar{W} \cdot \bar{X} = 0$$

OPTIMAL SOLUTION  
FOUND BY PERCEPTRON

LOSS FUNCTION DISCOURAGES  
MARGINALLY CORRECT PREDICTIONS



$$\bar{W} \cdot \bar{X} = 0$$

OPTIMAL SOLUTION  
FOUND BY SVM

# Logistic Regression

---

- Architecture:
  - A single neuron
- Activation function
  - Sigmoid function
- Loss function
  - $l(y, \hat{y}) = -y \log \hat{y} - (1 - y) \log(1 - \hat{y})$
  - Note  $\hat{y}$  is the predicted probability of taking class 1
- Optimization
  - $\mathbf{w} \leftarrow \mathbf{w} + \eta (y^{(i)} - \sigma(\mathbf{w}^T \mathbf{x}^{(i)})) \mathbf{x}^{(i)}$ , for a training data point  $(\mathbf{x}^{(i)}, y^{(i)})$
  - $\eta$ : learning rate


# Question

---

- How about multivariate logistic regression?


# Neural Network

---

- Introduction
- Connection to Shallow Machine Learning Algorithms
- Multi-Layer Feed-Forward Neural Network 
- Deep Learning
- Summary

# A Multi-Layer Feed-Forward Neural Network

A two-layer network



$$\mathbf{o} = g(W^{(2)}\mathbf{h} + b^{(2)})$$

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

Bias term

Weight matrix

Nonlinear transformation,  
e.g. sigmoid transformation

# How A Multi-Layer Neural Network Works

---

- The **inputs** to the network correspond to the attributes measured for each training tuple
- Inputs are fed simultaneously into the units making up the **input layer**
- They are then weighted and fed simultaneously to a **hidden layer**
  - The number of hidden layers is arbitrary
- The weighted outputs of the last hidden layer are input to units making up the **output layer**, which emits the network's prediction
- The network is **feed-forward**: None of the weights cycles back to an input unit or to an output unit of a previous layer
- From a math point of view, networks perform **nonlinear regression**: **Given enough hidden units and enough training samples, they can closely approximate any continuous function**

# Learning by Backpropagation

---

- Iteratively process a set of training tuples & compare the network's prediction with the actual known target value
- For each training tuple, the weights are modified to **minimize the loss function** between the network's prediction and the actual target value, say **mean squared error**
  - Stochastic gradient descent + chain rule
- Modifications are made in the “**backwards**” direction: from the output layer, through each hidden layer down to the first hidden layer, hence “**backpropagation**”

# Recap: Chain Rule

---

**The Chain Rule** If  $f$  and  $g$  are both differentiable and  $F = f \circ g$  is the composite function defined by  $F(x) = f(g(x))$ , then  $F$  is differentiable and  $F'$  is given by the product

$$F'(x) = f'(g(x))g'(x)$$


In Leibniz notation, if  $y = f(u)$  and  $u = g(x)$  are both differentiable functions, then

$$\frac{dy}{dx} = \frac{dy}{du} \frac{du}{dx}$$

[http://faculty.wlc.edu/buelow/Web%20Library/chain\\_rule.htm](http://faculty.wlc.edu/buelow/Web%20Library/chain_rule.htm)

# Example

- Loss function:  $\mathcal{L} = \frac{1}{2} \|\mathbf{y} - \hat{\mathbf{y}}\|^2$



$$\mathbf{z}^{(l)} = W^{(l)} \mathbf{x}^{(l-1)} + \mathbf{b}^{(l)}$$

$$\mathbf{x}^{(l)} = f^{(l)}(\mathbf{z}^{(l)})$$

$$\mathcal{L} = \frac{1}{2} \|\mathbf{y} - \mathbf{x}^{(3)}\|^2$$

# Gradient for Layer 3 (Last Layer)


- Stochastic gradient for  $W_{ij}^{(3)}$  and  $b_i^{(3)}$

- Recall:*

- $\mathcal{L} = \frac{1}{2} \|\mathbf{y} - \mathbf{x}^{(3)}\|^2 = \frac{1}{2} \sum_i (y_i - x_i^{(3)})^2$

- $x_i^{(3)} = f^{(3)}(z_i^{(3)})$

- $z_i^{(3)} = \sum_j W_{ij}^{(3)} x_j^{(2)} + b_i^{(3)}$



- $\frac{\partial \mathcal{L}}{\partial W_{ij}^{(3)}} = \frac{\partial \mathcal{L}}{\partial x_i^{(3)}} \frac{\partial x_i^{(3)}}{\partial z_i^{(3)}} \frac{\partial z_i^{(3)}}{\partial W_{ij}^{(3)}} = \underbrace{-(y_i - x_i^{(3)}) f'^{(3)}(z_i^{(3)})}_{\delta_i^{(3)}} x_j^{(2)}$


- $\frac{\partial \mathcal{L}}{\partial b_i^{(3)}} = \frac{\partial \mathcal{L}}{\partial x_i^{(3)}} \frac{\partial x_i^{(3)}}{\partial z_i^{(3)}} \frac{\partial z_i^{(3)}}{\partial b_i^{(3)}} = \underbrace{-(y_i - x_i^{(3)}) \delta_i^{(3)}}_{\delta_i^{(3)}} f'^{(3)}(z_i^{(3)})$

# Gradient for Layer 2

- Stochastic gradient for  $W_{ij}^{(2)}$  and  $b_i^{(2)}$

- Recall:*

- $\mathcal{L} = \frac{1}{2} \|\mathbf{y} - \mathbf{x}^{(3)}\|^2 = \frac{1}{2} \sum_i (y_i - x_i^{(3)})^2$
- $x_i^{(3)} = f^{(3)}(z_i^{(3)}); z_i^{(3)} = \sum_j W_{ij}^{(3)} x_j^{(2)} + b_i^{(3)}$
- $x_j^{(2)} = f^{(2)}(z_j^{(2)}); z_j^{(2)} = \sum_k W_{jk}^{(2)} x_k^{(1)} + b_j^{(2)}$




$$\begin{aligned}
 \bullet \frac{\partial \mathcal{L}}{\partial W_{jk}^{(2)}} &= \sum_i \frac{\partial \mathcal{L}}{\partial x_i^{(3)}} \frac{\partial x_i^{(3)}}{\partial z_i^{(3)}} \frac{\partial z_i^{(3)}}{\partial x_j^{(2)}} \frac{\partial x_j^{(2)}}{\partial z_j^{(2)}} \frac{\partial z_j^{(2)}}{\partial W_{jk}^{(2)}} \\
 &= \sum_i \underbrace{-(y_i - x_i^{(3)}) f'^{(3)}(z_i^{(3)})}_{\delta_i^{(3)}} \underbrace{W_{ij}^{(3)} f'^{(2)}(z_j^{(2)})}_{\delta_j^{(2)}} x_k^{(1)}
 \end{aligned}$$

# Gradient for Layer 1


- Stochastic gradient for  $W_{ks}^{(1)}$

- $$\frac{\partial \mathcal{L}}{\partial W_{ks}^{(1)}} = \sum_j \delta_j^{(2)} W_{jk}^{(2)} f'^{(1)} \left( z_k^{(1)} \right) x_s^{(0)}$$



# Question: General form

- layer  $l - 1 \rightarrow$  layer  $l \rightarrow$  layer  $l + 1$ 
  - Index:  $k \rightarrow j \rightarrow i$



- $\delta_j^{(l)} = \sum_i \delta_i^{(l+1)} W_{ij}^{(l+1)} f'^{(l)}(z_j^{(l)})$
- $\frac{\partial \mathcal{L}}{\partial W_{jk}^{(l)}} = \delta_j^{(l)} x_k^{(l-1)}$

# Backpropagation Steps to Learn Weights

---

- Initialize weights to small random numbers, associated with biases
- **Repeat** until terminating condition meets
  - **For** each training example
    - **Propagate the inputs forward** (by applying activation function)
      - For layer  $l = 1: L$ 
        - Calculate  $\mathbf{z}^{(l)} = W^{(l)}\mathbf{x}^{(l-1)} + \mathbf{b}^{(l)}$
        - Calculate  $\mathbf{x}^{(l)} = f^{(l)}(\mathbf{z}^{(l)})$  (elementwise activation)
      - **Backpropagate the error** (by updating weights and biases)
        - Calculate  $\delta^{(L)}$ , Update  $W^{(L)}$  and  $\mathbf{b}^{(L)}$  based on  $\frac{\partial \mathcal{L}}{\partial W^{(L)}} = \delta^{(L)}(\mathbf{x}^{(L-1)})^T$  and  $\frac{\partial \mathcal{L}}{\partial \mathbf{b}^{(L)}} = \delta^{(L)}$
        - For layer  $l = L-1: 1$ 
          - Calculate  $\delta^{(l)} = (W^{(l+1)})^T \delta^{(l+1)} \circ f'^{(l)}(\mathbf{z}^{(l)})$
          - Update  $W^{(l)}$  and  $\mathbf{b}^{(l)}$  based on  $\frac{\partial \mathcal{L}}{\partial W^{(l)}} = \delta^{(l)}(\mathbf{x}^{(l-1)})^T$  and Update  $\frac{\partial \mathcal{L}}{\partial \mathbf{b}^{(l)}} = \delta^{(l)}$
  - Terminating condition (convergence, max iteration, etc.)

# Neural Network as a Classifier

---

- Weakness

- Long training time
- Require a number of hyper-parameters typically best determined empirically, e.g., the network topology or “structure.”
- Poor interpretability: Difficult to interpret the symbolic meaning behind the learned weights and of “hidden units” in the network

- Strength

- High tolerance to noisy data
- Successful on an array of real-world data, e.g., hand-written letters
- Algorithms are inherently parallel
- Techniques have recently been developed for the extraction of rules from trained neural networks
- Deep neural network is powerful


# Digits Recognition Example

---

- Obtain sequence of digits by segmentation




- Recognition (our focus)




# Digits Recognition Example

- The architecture of the used neural network




- What each neurons are doing?





# Neural Network

---

- Introduction
- Connection to Shallow Machine Learning Algorithms
- Multi-Layer Feed-Forward Neural Network
- Deep Learning\* 
- Summary

# Towards Deep Learning

## Deep neural network




# Popular Deep NNs

---

- Deep layers and parameter sharing
  - Convolutional Neural Network
    - Applications to: Images, NLP, Graph
  - Recurrent Neural Network
    - Applications to: Sequence data

# Convolutional Neural Network

- The LeNet Architecture (By Yann LeCun et al.)



Source: <https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/>

# Convolution

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

Input

\*

1	0	1
0	1	0
1	0	1

Filter

Convolution operator: weighted sum where weights are from filter matrix

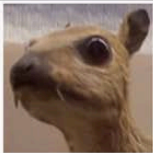




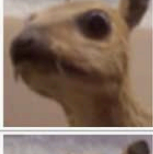
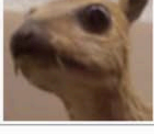
1 <sub>x1</sub>	1 <sub>x0</sub>	1 <sub>x1</sub>	0	0
0 <sub>x0</sub>	1 <sub>x1</sub>	1 <sub>x0</sub>	1	0
0 <sub>x1</sub>	0 <sub>x0</sub>	1 <sub>x1</sub>	1	1
0	0	1	1	0
0	1	1	0	0

Image

4		


Convolved  
Feature

# Effects of Different Filters

Operation	Filter	Convolved Image
<b>Identity</b>	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	
<b>Edge detection</b>	$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$	
	$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$	
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	
<b>Sharpen</b>	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
<b>Box blur</b> (normalized)	$\frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$	
<b>Gaussian blur</b> (approximation)	$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$	

Source:  
[https://en.wikipedia.org/wiki/Kernel\\_\(image\\_processing\)](https://en.wikipedia.org/wiki/Kernel_(image_processing))

# Pooling




Rectified Feature Map

Source: <http://cs231n.github.io/convolutional-networks/>

# Recurrent Neural Network

- Model sequential information

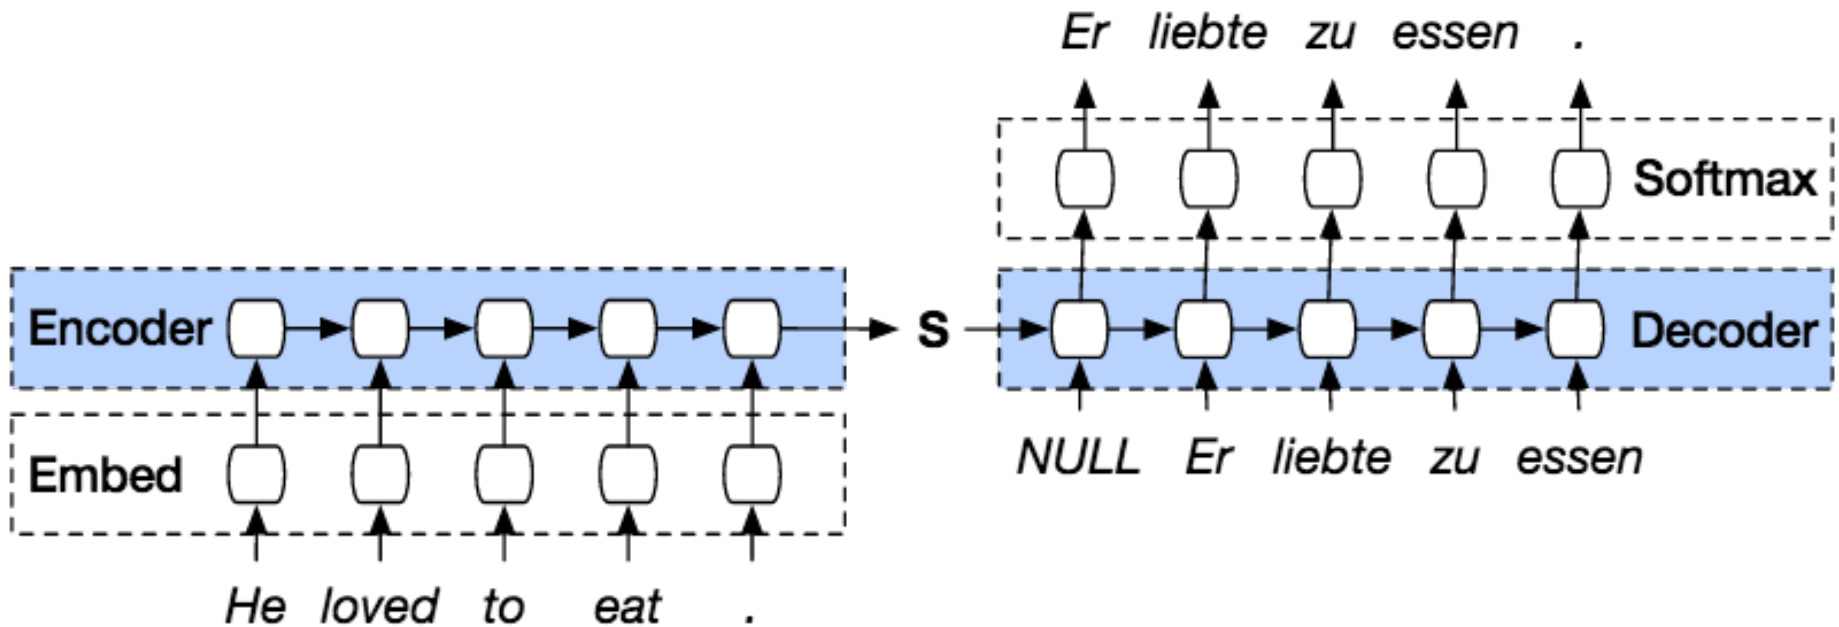


$x_t$ : input vector at time  $t$   
 $s_t$ : hidden state at time  $t$   
 $o_t$ : output vector at time  $t$

$$s_t = f(Ws_{t-1} + Ux_t)$$
$$o_t = g(Vs_t)$$


# Application of Machine Translation

- Seq2Seq model



# Neural Network

---

- Introduction
- Connection to Shallow Machine Learning Algorithms
- Multi-Layer Feed-Forward Neural Network
- Deep Learning
- Summary 

# Summary

---

- Neural Network
  - Architecture; activation function; loss function; backpropagation
- Existing shallow machine learning algorithms can be represented in NN
- Multilayer feedforward NN
- Deep learning

# Further References

---

- 3Blue1Brown NN series:

[https://www.youtube.com/watch?v=aircAruvnKk&list=PLZHQObOWTQDNU6R1\\_67000Dx\\_ZCJB-3pi](https://www.youtube.com/watch?v=aircAruvnKk&list=PLZHQObOWTQDNU6R1_67000Dx_ZCJB-3pi)

- Deep Learning

- <http://neuralnetworksanddeeplearning.com/>
- <http://www.deeplearningbook.org/>
- <http://www.charuaggarwal.net/neural.htm>
- <http://d2l.ai/index.html>