

Course Notes: Deep Learning for Visual Computing

Peter Wonka

August 30, 2021

Contents

1	Fully Connected Neural Networks	5
1.1	Literature	6
1.2	NN Compute Complex Functions	7
1.3	Simplification: Vectors	8
1.4	Function Review	9
1.5	Function Notation	10
1.6	Constant Offset in Linear Functions	11
1.7	Non-linear Function: Sigmoid	12
1.8	Non-linear Function: ReLU	13
1.9	Building Complex Functions with Simple Building Blocks	14
1.10	Computational Graph	15
1.11	Computational Graph Drawing Choices	16
1.12	Computational Graph Drawing Choices	17
1.13	Scalar vs. Tensor Nodes	19
1.14	Computational Graph Drawing Choices	20

1.15	Different Types of Variables	21
1.16	Layers	22
1.17	Non-linear Layers	23
1.18	Batch Processing	25
1.19	Fully Connected Neural Network	26
1.20	Types of Neural Networks	27
1.21	Neural Networks Alternative	28
1.22	Building a single Neuron	29
1.23	Single Neuron Drawing	30
1.24	Single Neuron Drawing Alternative	32
1.25	Building a Mini Network	33
1.26	Computing Boolean Functions	36
1.26.1	Building the AND function	39
1.26.2	Building the OR function	41
1.26.3	Building the NOT function	43
1.26.4	Building $(\text{NOT}x_1)$ AND $(\text{NOT} x_2)$	45

1.26.5 Putting Things Together	47
--	----

1 Fully Connected Neural Networks

1.1 Literature

- 2 Andrew NG, coursera, course on machine learning, lecture about neural networks
- 2 CS231n: Convolutional Neural Networks for Visual Recognition, Justin Johnson, Serena Yeung, Fei-Fei Li, youtube

1.2 NN Compute Complex Functions

- 2 A neural network computes a function of an input tensor and maps it to an output tensor
- 2 A neural network defines a **complex** function that maps an input to an output
 - 2 high-dimensional input
 - 2 many parameters
 - 2 possibly high-dimensional output

1.3 Simplification: Vectors

2 Simplification: assume inputs (and outputs) are vectors (for now)

2 Notation:

$$\begin{matrix} & & 0 & 1 \\ & & x_1 & \\ & B & & C \\ x & A & B & x_2 & C & R^n \\ & @ & \dots & A \\ & & & & & x_n \end{matrix}$$

2 We are dealing with

2 images, videos, volumes

2 point clouds

2 graphs, meshes

2 text - not so much in this class

2 n is quite large

1.4 Function Review

2 A function with n inputs and m outputs is denoted by $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$

2 Example: image classification

2 vectorized images as input

2 vector of class probabilities as output

2 $f : \mathbb{R}^{1000000} \rightarrow \mathbb{R}^{1000}$

2 Linear function $f : \mathbb{R}^n \rightarrow \mathbb{R}$:

$$f(x) = w_1 x_1 + w_2 x_2 + \dots + w_n x_n + \text{constant}$$


2 Note: we do not distinguish linear / affine

1.5 Function Notation

- 2 A function has **variables** and **parameters** and **constants**
- 2 Different notations to distinguish between variables and parameters
 - 2 Ignore parameters $f(x) \in \mathbb{R}^{w_1 \times x_1} \dots$
 - 2 Use a semicolon $f(x; w) \in \mathbb{R}^{w_1 \times x_1} \dots$
 - 2 Use a subscript $f_w(x) \in \mathbb{R}^{w_1 \times x_1} \dots$

1.6 Constant 0 set in Linear Functions

- 2 Problem: For a linear function it would be much nicer to write $f(x; w) = w^T x$
- 2 Solution: implicitly add an element to every vector



$$f(x; w) = w^T x = w_0 \cdot 1 + w_1 x_1 + \dots + w_n x_n$$

- 2 It is no longer clear what version is used (with or without 1).
- 2 Typically, authors explicitly consider the bias w_0 in simple equations / figures, but omit bias terms in complex equations / figures

1.7 Non-linear Function: Sigmoid

2 Two popular options for non-linear functions are **Sigmoid** and **ReLU**

2 Sigmoid: $\text{sigmoid}(x) = \frac{1}{1 + e^{-x}}$

± $\text{sigmoid}(10) \approx 0.9999$ and $\text{sigmoid}(-10) \approx 0.0001$

1.8 Non-linear Function: ReLU

2 **ReLU**: $\text{ReLU}(x) = \max(x, 0)$

2 Reminder: vector version $\max(x, 0)$ computes the \max elementwise

1.9 Building Complex Functions with Simple Building Blocks


- 2 We can use simple building blocks to build a complex function
- 2 A building block is a simpler function, e.g. $f_i(x)$ or $Y = f_i(X)$
- 2 Countless building blocks have been proposed
- 2 One important criteria is that the building block has to be differentiable

1.10 Computational Graph

- 2 A NN is a type of computational graph
- 2 A **Computational Graph** is a directed graph that has nodes that are **operations** (**function**) or **variables**,
 - 2 Data travels on the edges


1.11 Computational Graph Drawing Choices


- 2 Example: $f : \mathbb{R}^2 \rightarrow \mathbb{R}, y \in f(x), y \in x_1 + 3x_2$
- 2 Unusual drawing where output variable is drawn as a separate node



1.12 Computational Graph Drawing Choices


- 2 Drawing variables as separate node or not
 - 2 Typically only input variables are drawn as nodes
 - 2 Variables can just be drawn on edges
 - 2 Output node is typically not drawn at all
 - 2 Confusing either way






1.13 Scalar vs. Tensor Nodes

- 2 Traditional deep learning: each node computes a single scalar value
- 2 Overkill for drawing, manual computation, and programming
 - 2 E.g. $f : \mathbb{R}^4 \rightarrow \mathbb{R}^2$ with two **hidden layers**




1.14 Computational Graph Drawing Choices

- 2 Much easier with inputs and outputs as tensors
- 2 Example $z = \text{AE}f(A, x, b)$




1.15 Different Types of Variables

- 2 There are two types of variables
 - 2 Example $z \in f(x; W, b) \in W x \dot{A} b$
 - 2 Input variables that describe the input data,
 - 2 **Parameters** or **weights** of the neural network W, b
 - 2 Constants are fixed once
- 2 Runtime: new input data determines variables W, b are not changed.
- 2 Training: input data x is taken from the training set; W, b are variables we optimize
- 2 Simplification: not drawing network parameters




1.16 Layers

- 2 A computational node is also called **layer** of the network
- 2 How complex should a layer be? The complete network could be a single layer
- 2 **Linear Layer**: computes $y \in \mathbb{R}^b$ linear $(x; W, b) \in \mathbb{R}^W \times \mathbb{R}^b$
- 2 A layer is often drawn as a rectangle (the height and width of the rectangle can be proportional to the size of the output tensor)




1.17 Non-linear Layers

- 2 Non-linear layers work element-wise, e.g. $\text{ReLU}(x)$ or $\text{sigmoid}(x)$
- 2 Traditional term: **activation function**
- 2 Programming: separate layer that is implemented, e.g. on top of a linear layer
- 2 Visualization: omitted, or drawn directly next to another layer



$$y \in \text{ReLU}(Wx + b)$$



$$y \in Wx + b$$

$$z \in \text{ReLU}(y)$$




1.18 Batch Processing

- 2 The neural network needs to process data **batches** (**mini-batches**)
- 2 If the input to a layer is a single vector, a sequence of vectors can be stored as a matrix (rank-2 tensor)
- 2 The layer has to process this rank-2 tensor
 - 2 Data can be in columns or rows
 - 2 Careful with transpose! E.g. $\text{ReLU}(WX \dot{\wedge} b)$? $\text{ReLU}(X^T W \dot{\wedge} b^T)$?
 - 2 Careful with broadcasting! E.g. $WX \dot{\wedge} b$. Is b added to columns or rows?
 - 2 Broadcasting is not allowed in mathematics, only in computer science

1.19 Fully Connected Neural Network

- 2 **Fully Connected Neural Network** : A NN that consists of alternating linear layers and non-linear layers
- 2 Alternative term: **Multi-layer Perceptron**
- 2 As final layer, there is typically another type of non-linear layer, e.g. softmax for classification



1.20 Types of Neural Networks

- 2 Feedforward Neural Networks vs. Recurrent Neural Networks:
 - 2 Feedforward Neural Networks are also called Deep Feedforward Networks
 - 2 Feedforward networks have no feedback connections.
 - 2 Recurrent Neural Networks (RNNs) have feedback connections.
- 2 Types of blocks
 - 2 A Fully Connected Network also called Multi-layer Perceptron (MLP) uses (only) fully connected layers (and non-linear layers)
 - 2 A fully connected layer is also called linear layer in PyTorch
 - 2 A Convolutional Neural Network (CNN) uses predominantly convolutional layers

1.21 Neural Networks Alternative

- ² Traditionally, many people built neural networks using neurons that output a single scalar value
- ² We now look at this alternative variation

1.22 Building a single Neuron

- 2 Neuron combines a linear function with a nonlinear function
- 2 Neuron has many inputs, but only one output
- 2 Example:


$$\frac{1}{1 + e^{-\mathbf{w}^T \mathbf{x}}}$$
$$\max(0, \mathbf{w}^T \mathbf{x})$$

- 2 Terms: the non-linear function is called **activation function**


1.23 Single Neuron Drawing

2 Example:

0 1
 $x_0 \in \mathbb{R}^1$
neuron $y \in f(\mathbf{x}; \mathbf{w})$
 x_1
 x_2
 x_3




2 Variables




- 2 The bias unit x_0 can be drawn as separate input or not
- 2 The variables w are called **weights** or **parameters** of the network

1.24 Single Neuron Drawing Alternative

- 2 Drawings inspired by transfer from neurology
 - 2 taken from another context with slightly different variables



1.25 Building a Mini Network



² Terms:

² Layer 0 is also called **Input Layer**

² Layer 1 is a **Hidden Layer**

2 Layer 2 is the **Output Layer**

2 $a_i^{(j)}$ = **activation** of unit i in layer j

2 $W^{(j)}$ = matrix of weights controlling function mapping from layer j to layer $j+1$

2 Computation

2 $a_1^{(2)} \in \text{nl}(w_{10}^{(1)} x_0 \dot{+} w_{11}^{(1)} x_1 \dot{+} w_{12}^{(1)} x_2 \dot{+} w_{13}^{(1)} x_3)$

2 $a_2^{(2)} \in \text{nl}(w_{20}^{(1)} x_0 \dot{+} w_{21}^{(1)} x_1 \dot{+} w_{22}^{(1)} x_2 \dot{+} w_{23}^{(1)} x_3)$

2 $a_3^{(2)} \in \text{nl}(w_{30}^{(1)} x_0 \dot{+} w_{31}^{(1)} x_1 \dot{+} w_{32}^{(1)} x_2 \dot{+} w_{33}^{(1)} x_3)$

2 $output \in a_1^{(3)} \in \text{nl}(w_{10}^{(2)} a_0^{(2)} \dot{+} w_{11}^{(2)} a_1^{(2)} \dot{+} w_{12}^{(2)} a_2^{(2)} \dot{+} w_{13}^{(2)} a_3^{(2)})$

2 If network has s_j units in layer j and s_{j+1} units in layer $j+1$, then $W^{(j)}$ will be of dimension $s_{j+1} \times (s_j + 1)$.

2 In the example above, $W^{(1)}$ has $3 \times 4 = 12$ entries.

2 Terms:

2 The inputs x_i could also be considered as activations $a_i^{(1)}$

² Since all activations from layer j connect to all neurons of layer $j+1$ this network is called a **fully connected** neural network.

² Vectorized Implementation:

² $\mathbf{a}^{(1)} \in \mathbf{x}$

² add $a_0^{(1)} \in 1$

² $lin^{(2)} \in \mathbf{W}^{(1)} \mathbf{a}^{(1)}$

² $\mathbf{a}^{(2)} \in nl(lin^{(2)})$


² add $a_0^{(2)} \in 1$

² $lin^{(3)} \in \mathbf{W}^{(2)} \mathbf{a}^{(2)}$


² $h_w(\mathbf{x}) \in \mathbf{a}^{(3)} \in nl(lin^{(3)})$

1.26 Computing Boolean Functions

² We are interested in computing Boolean functions as examples



- 2 This example shows the XNOR function
- 2 $y \in x_1 XNOR x_2$
- 2 $y \in NOT(x_1 XOR x_2)$
- 2 $x_1, x_2 \in \{0, 1\}$
- 2 red X means $y \in 1$, blue circle means $y \in 0$
- 2 Problem: there is no 2 layer (input and output layer) network that can build the XNOR function
- 2 This could be imagined as a simplified version of a more complex classification example



1.26.1 Building the AND function

$$y = x_1 \text{ AND } x_2$$

	+1	
	-30	
x_1	20	$h(\mathbf{x}; \mathbf{w})$
x_2	20	

2 Verification:

± reminder: $\text{sigmoid}(10) \hat{=} 0.9999$ and $\text{sigmoid}(-10) \hat{=} 0.0001$

± $0 \text{ AND } 0 = 0$ $\hat{=} 20 \hat{=} 0$ $\hat{=} 20 \hat{=} 1$ $\hat{=} (-30) \hat{=} -30 \hat{=} 0$

$\pm 1 \text{ AND } 0 \text{ } \mathcal{E} 20_j \text{ } 30 \text{ } \mathcal{E} j \text{ } 10 \text{ } \frac{1}{4} 0$

$\pm 0 \text{ AND } 1 \text{ } \frac{1}{4} 0$

$\pm 1 \text{ AND } 1 \text{ } \frac{1}{4} 1$

1.26.2 Building the OR function

$$y \approx x_1 \text{ OR } x_2$$

$$+1 \quad -10$$

$$x_1 \quad 20$$

$$x_2 \quad 20$$

$$h(\mathbf{x}; \mathbf{w})$$

2 Verification:

± reminder: $\text{sigmoid}(10) \hat{=} 0.9999$ and $\text{sigmoid}(-10) \hat{=} 0.0001$

± 0 OR 0 = 0


$\pm 1OR0 \frac{1}{4}1$

$\pm 0OR1 \frac{1}{4}1$

$\pm 1OR1 \frac{1}{4}1$

1.26.3 Building the NOT function

$$y \neq \text{NOT } x_1$$



Verification:

$$\neq \text{NOT } 0 \neq 1$$

\pm *NOT* 1 $\frac{1}{4}$ 0

1.26.4 Building (NOT x_1) AND (NOT x_2)

$$y \in (\text{NOT } x_1) \text{ AND } (\text{NOT } x_2)$$

$$+1 \quad 10$$

$$x_1 \quad -20$$

$$x_2 \quad -20$$

$$h(\mathbf{x}; \mathbf{w})$$

2 Verification:

$$2 \quad 0 \text{ op } 0 \text{ } \frac{1}{4} 1$$

$$2 \quad 1 \text{ op } 0 \text{ } \frac{1}{4} 0$$

² $0op1 \frac{1}{4}0$

² $1op1 \frac{1}{4}0$

² op is short for the operation $(NOT x_1) AND (NOT x_2)$

