

Ujjwal Karn's blog

deep learning, computer vision, nlp


A Quick Introduction to Neural Networks

Posted on ~~August 9, 2016~~ August 10, 2016 by ujjwalkarn

An Artificial Neural Network (ANN) is a computational model that is inspired by the way biological neural networks in the human brain process information. Artificial Neural Networks have generated a lot of excitement in Machine Learning research and industry, thanks to many breakthrough results in speech recognition, computer vision and text processing. In this blog post we will try to develop an understanding of a particular type of Artificial Neural Network called the Multi Layer Perceptron.

A Single Neuron

The basic unit of computation in a neural network is the **neuron**, often called a **node** or **unit**. It receives input from some other nodes, or from an external source and computes an output. Each input has an associated **weight** (w), which is assigned on the basis of its relative importance to other inputs. The node applies a function f (defined below) to the weighted sum of its inputs as shown in Figure 1 below:



$$\text{Output of neuron} = Y = f(w1.X1 + w2.X2 + b)$$

Figure 1: a single neuron

The above network takes numerical inputs **X1** and **X2** and has weights **w1** and **w2** associated with those inputs. Additionally, there is another input **1** with weight **b** (called the **Bias**) associated with it. We will learn more details about role of the bias later.

The output **Y** from the neuron is computed as shown in the Figure 1. The function **f** is non-linear and is called the **Activation Function**. The purpose of the activation function is to introduce non-linearity into the output of a neuron. This is important because most real world data is non linear and we want neurons to *learn* these non linear representations.

Every activation function (or *non-linearity*) takes a single number and performs a certain fixed mathematical operation on it [2]. There are several activation functions you may encounter in practice:

- **Sigmoid:** takes a real-valued input and squashes it to range between 0 and 1

$$\sigma(x) = 1 / (1 + \exp(-x))$$

- **tanh:** takes a real-valued input and squashes it to the range [-1, 1]

$$\tanh(x) = 2\sigma(2x) - 1$$

- **ReLU:** ReLU stands for Rectified Linear Unit. It takes a real-valued input and thresholds it at zero (replaces negative values with zero)

$$f(x) = \max(0, x)$$

The below figures [2] show each of the above activation functions.




Figure 2: different activation functions

Importance of Bias: The main function of Bias is to provide every node with a trainable constant value (in addition to the normal inputs that the node receives). See [this link](http://stackoverflow.com/q/2480650/3297280) (<http://stackoverflow.com/q/2480650/3297280>) to learn more about the role of bias in a neuron.

Feedforward Neural Network

The feedforward neural network was the first and simplest type of artificial neural network devised [3]. It contains multiple neurons (nodes) arranged in **layers**. Nodes from adjacent layers have **connections** or **edges** between them. All these connections have **weights** associated with them.

An example of a feedforward neural network is shown in Figure 3.




Figure 3: an example of feedforward neural network

A feedforward neural network can consist of three types of nodes:

1. **Input Nodes** – The Input nodes provide information from the outside world to the network and are together referred to as the “Input Layer”. No computation is performed in any of the Input nodes – they just pass on the information to the hidden nodes.
2. **Hidden Nodes** – The Hidden nodes have no direct connection with the outside world (hence the name “hidden”). They perform computations and transfer information from the input nodes to the output nodes. A collection of hidden nodes forms a “Hidden Layer”. While a feedforward network will only have a single input layer and a single output layer, it can have zero or multiple Hidden Layers.
3. **Output Nodes** – The Output nodes are collectively referred to as the “Output Layer” and are responsible for computations and transferring information from the network to the outside world.

In a feedforward network, the information moves in only one direction – forward – from the input nodes, through the hidden nodes (if any) and to the output nodes. There are no cycles or loops in the network [3] (this property of feed forward networks is different from Recurrent Neural Networks in which the connections between the nodes form a cycle).

Two examples of feedforward networks are given below:

1. **Single Layer Perceptron** – This is the simplest feedforward neural network [4] and does not contain any hidden layer. You can learn more about Single Layer Perceptrons in [4], [5], [6], [7].
2. **Multi Layer Perceptron** – A Multi Layer Perceptron has one or more hidden layers. We will only discuss Multi Layer Perceptrons below since they are more useful than Single Layer Perceptrons for practical applications today.

Multi Layer Perceptron

A Multi Layer Perceptron (MLP) contains one or more hidden layers (apart from one input and one output layer). While a single layer perceptron can only learn linear functions, a multi layer perceptron can also learn non – linear functions.

Figure 4 shows a multi layer perceptron with a single hidden layer. Note that all connections have weights associated with them, but only three weights (w_0 , w_1 , w_2) are shown in the figure.

Input Layer: The Input layer has three nodes. The Bias node has a value of 1. The other two nodes take X_1 and X_2 as external inputs (which are numerical values depending upon the input dataset). As discussed above, no computation is performed in the Input layer, so the outputs from nodes in the Input layer are 1, X_1 and X_2 respectively, which are fed into the Hidden Layer.

Hidden Layer: The Hidden layer also has three nodes with the Bias node having an output of 1. The output of the other two nodes in the Hidden layer depends on the outputs from the Input layer (1, X_1 , X_2) as well as the weights associated with the connections (edges). Figure 4 shows the output calculation for one of the hidden nodes (highlighted). Similarly, the output from other hidden node can be calculated. Remember that f refers to the activation function. These outputs are then fed to the nodes in the Output layer.

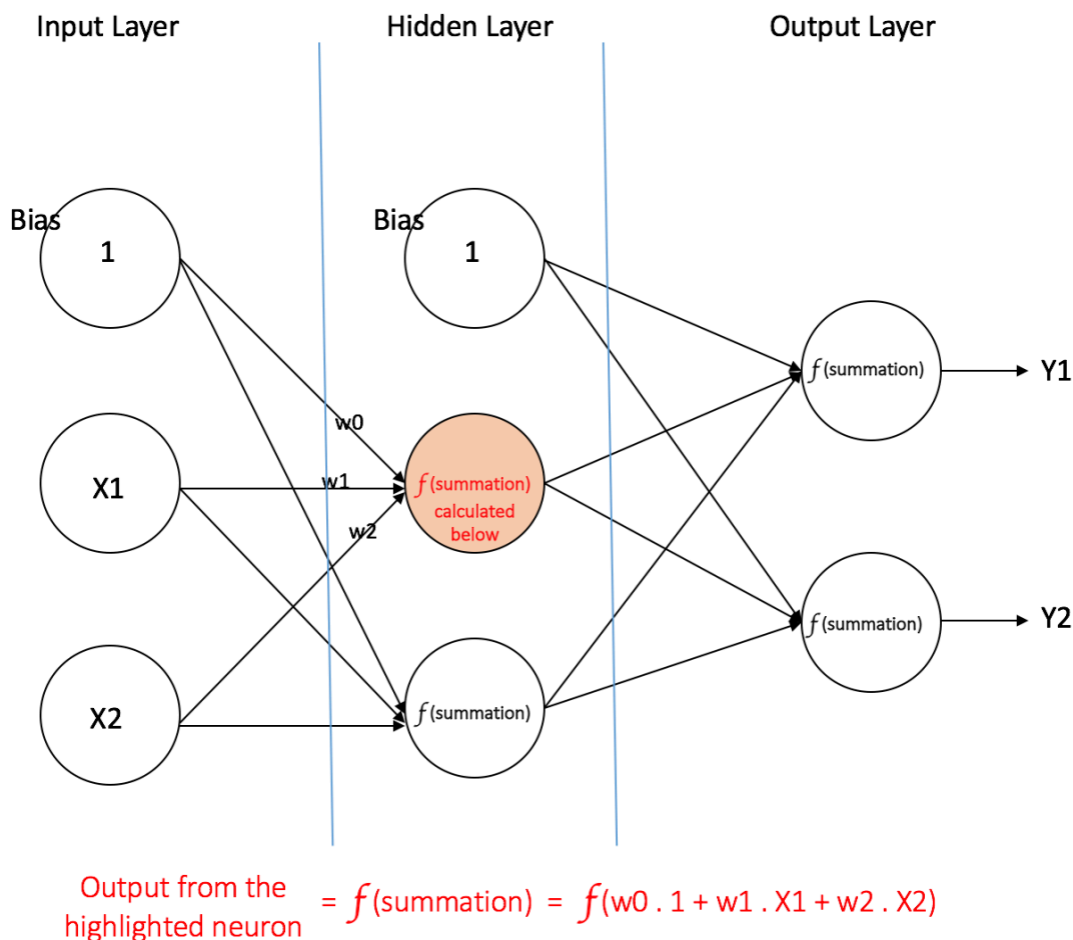


Figure 4: a multi layer perceptron having one hidden layer

Output Layer: The Output layer has two nodes which take inputs from the Hidden layer and perform similar computations as shown for the highlighted hidden node. The values calculated (Y_1 and Y_2) as a result of these computations act as outputs of the Multi Layer Perceptron.

Given a set of features $X = (x_1, x_2, \dots)$ and a target y , a Multi Layer Perceptron can learn the relationship between the features and the target, for either classification or regression.

Lets take an example to understand Multi Layer Perceptrons better. Suppose we have the following student-marks dataset:

Hours Studied	Mid Term Marks	Final Term Result
35	67	1 (Pass)
12	75	0 (Fail)
16	89	1 (Pass)
45	56	1 (Pass)
10	90	0 (Fail)

The two input columns show the number of hours the student has studied and the mid term marks obtained by the student. The Final Result column can have two values 1 or 0 indicating whether the student passed in the final term. For example, we can see that if the student studied 35 hours and had obtained 67 marks in the mid term, he / she ended up passing the final term.

Now, suppose, we want to predict whether a student studying 25 hours and having 70 marks in the mid term will pass the final term.

Hours Studied	Mid Term Marks	Final Term Result
25	70	?

This is a binary classification problem where a multi layer perceptron can learn from the given examples (training data) and make an informed prediction given a new data point. We will see below how a multi layer perceptron learns such relationships.

Training our MLP: The Back-Propagation Algorithm

The process by which a Multi Layer Perceptron learns is called the Backpropagation algorithm. I would recommend reading [this Quora answer by Hemanth Kumar \(https://www.quora.com/How-do-you-explain-back-propagation-algorithm-to-a-beginner-in-neural-network/answer/Hemanth-Kumar-Mantri\)](https://www.quora.com/How-do-you-explain-back-propagation-algorithm-to-a-beginner-in-neural-network/answer/Hemanth-Kumar-Mantri) (quoted below) which explains Backpropagation clearly.

Backward Propagation of Errors, often abbreviated as BackProp is one of the several ways in which an artificial neural network (ANN) can be trained. It is a supervised training scheme, which means, it learns from labeled training data (there is a supervisor, to guide its learning).

To put in simple terms, BackProp is like “**learning from mistakes**”. The supervisor *corrects* the ANN whenever it makes mistakes.

An ANN consists of nodes in different layers; input layer, intermediate hidden layer(s) and the output layer. The connections between nodes of adjacent layers have “weights” associated with them. The goal of learning is to assign correct weights for these edges. Given an input vector, these weights determine what the output vector is.

In supervised learning, the training set is labeled. This means, for some given inputs, we know the desired/expected output (label).

BackProp Algorithm:

Initially all the edge weights are randomly assigned. For every input in the training dataset, the ANN is activated and its output is observed. This output is compared with the desired output that we already know, and the error is “propagated” back to the previous layer. This error is noted and the weights are “adjusted” accordingly. This process is repeated until the output error is below a predetermined threshold.

Once the above algorithm terminates, we have a “learned” ANN which, we consider is ready to work with “new” inputs. This ANN is said to have learned from several examples (labeled data) and from its mistakes (error propagation).

Now that we have an idea of how Backpropagation works, lets come back to our student-marks dataset shown above.

The Multi Layer Perceptron shown in Figure 5 (adapted from Sebastian Raschka’s [excellent visual explanation of the backpropagation algorithm \(https://github.com/rasbt/python-machine-learning-book/blob/master/faq/visual-backpropagation.md\)](https://github.com/rasbt/python-machine-learning-book/blob/master/faq/visual-backpropagation.md)) has two nodes in the input layer (apart from the Bias node) which take the inputs ‘Hours Studied’ and ‘Mid Term Marks’. It also has a hidden layer with two nodes (apart from the Bias node). The output layer has two nodes as well – the upper node outputs the probability of ‘Pass’ while the lower node outputs the probability of ‘Fail’.

In classification tasks, we generally use a [Softmax function \(http://cs231n.github.io/linear-classify/#softmax\)](http://cs231n.github.io/linear-classify/#softmax) as the Activation Function in the Output layer of the Multi Layer Perceptron to ensure that the outputs are probabilities and they add up to 1. The Softmax function takes a vector of arbitrary real-valued scores and squashes it to a vector of values between zero and one that sum to one. So, in this case,

$$\text{Probability (Pass)} + \text{Probability (Fail)} = 1$$

Step 1: Forward Propagation

All weights in the network are randomly assigned. Lets consider the hidden layer node marked **V** in Figure 5 below. Assume the weights of the connections from the inputs to that node are w_1 , w_2 and w_3 (as shown).

The network then takes the first training example as input (we know that for inputs 35 and 67, the probability of Pass is 1).

- Input to the network = [35, 67]
- Desired output from the network (target) = [1, 0]

Then output V from the node in consideration can be calculated as below (f is an activation function such as sigmoid):

$$V = f(1*w1 + 35*w2 + 67*w3)$$

Similarly, outputs from the other node in the hidden layer is also calculated. The outputs of the two nodes in the hidden layer act as inputs to the two nodes in the output layer. This enables us to calculate output probabilities from the two nodes in output layer.

Suppose the output probabilities from the two nodes in the output layer are 0.4 and 0.6 respectively (since the weights are randomly assigned, outputs will also be random). We can see that the calculated probabilities (0.4 and 0.6) are very far from the desired probabilities (1 and 0 respectively), hence the network in Figure 5 is said to have an 'Incorrect Output'.




Figure 5: forward propagation step in a multi layer perceptron

Step 2: Back Propagation and Weight Updation

We calculate the total error at the output nodes and propagate these errors back through the network using Backpropagation to calculate the *gradients*. Then we use an optimization method such as *Gradient Descent* to 'adjust' **all** weights in the network with an aim of reducing the error at the output layer. This is shown in the Figure 6 below (ignore the mathematical equations in the figure for now).

Suppose that the new weights associated with the node in consideration are $w4$, $w5$ and $w6$ (after Backpropagation and adjusting weights).

Backpropagation
+
Weights Adjusted

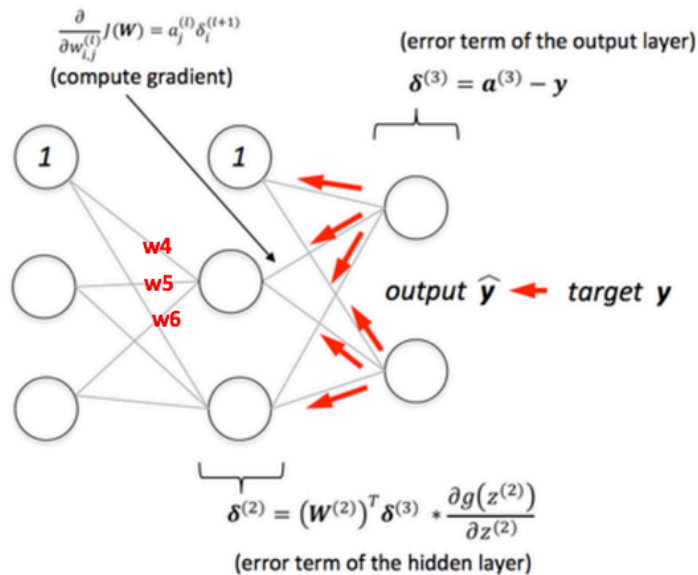


Figure 6: backward propagation and weight updation step in a multi layer perceptron

If we now input the same example to the network again, the network should perform better than before since the weights have now been adjusted to minimize the error in prediction. As shown in Figure 7, the errors at the output nodes now reduce to [0.2, -0.2] as compared to [0.6, -0.4] earlier. This means that our network has learnt to correctly classify our first training example.

Correct Output

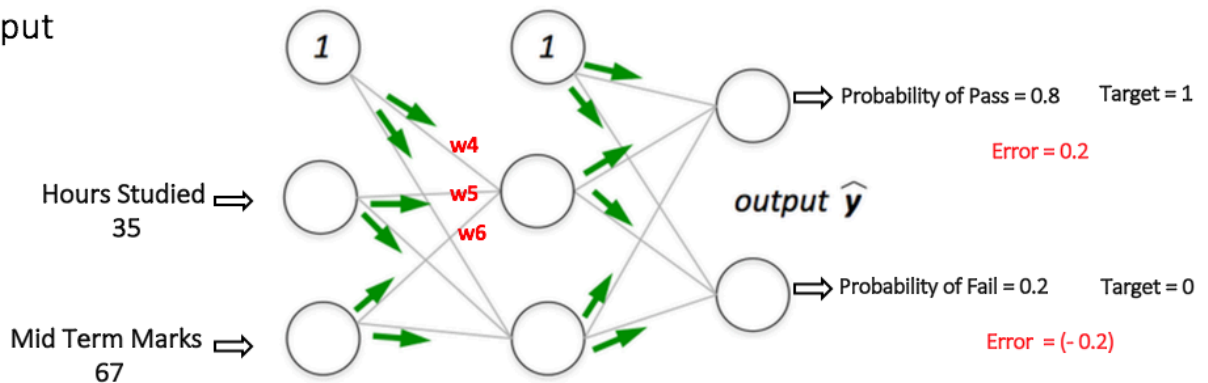


Figure 7: the MLP network now performs better on the same input

We repeat this process with all other training examples in our dataset. Then, our network is said to have *learnt* those examples.

If we now want to predict whether a student studying 25 hours and having 70 marks in the mid term will pass the final term, we go through the forward propagation step and find the output probabilities for Pass and Fail.

I have avoided mathematical equations and explanation of concepts such as 'Gradient Descent' here and have rather tried to develop an intuition for the algorithm. For a more mathematically involved discussion of the Backpropagation algorithm, refer to [this link](http://home.agh.edu.pl/~vlsi/AI/backp_t_en/backprop.html) (http://home.agh.edu.pl/~vlsi/AI/backp_t_en/backprop.html).

3d Visualization of a Multi Layer Perceptron

Adam Harley has created a [3d visualization](http://scs.ryerson.ca/~aharley/vis/fc/) (<http://scs.ryerson.ca/~aharley/vis/fc/>) of a Multi Layer Perceptron which has already been trained (using Backpropagation) on the MNIST Database of handwritten digits.

The network takes 784 numeric pixel values as inputs from a 28 x 28 image of a handwritten digit (it has 784 nodes in the Input Layer corresponding to pixels). The network has 300 nodes in the first hidden layer, 100 nodes in the second hidden layer, and 10 nodes in the output layer (corresponding to the 10 digits) [15].

Although the network described here is much larger (uses more hidden layers and nodes) compared to the one we discussed in the previous section, all computations in the forward propagation step and backpropagation step are done in the same way (at each node) as discussed before.

Figure 8 shows the network when the input is the digit '5'.

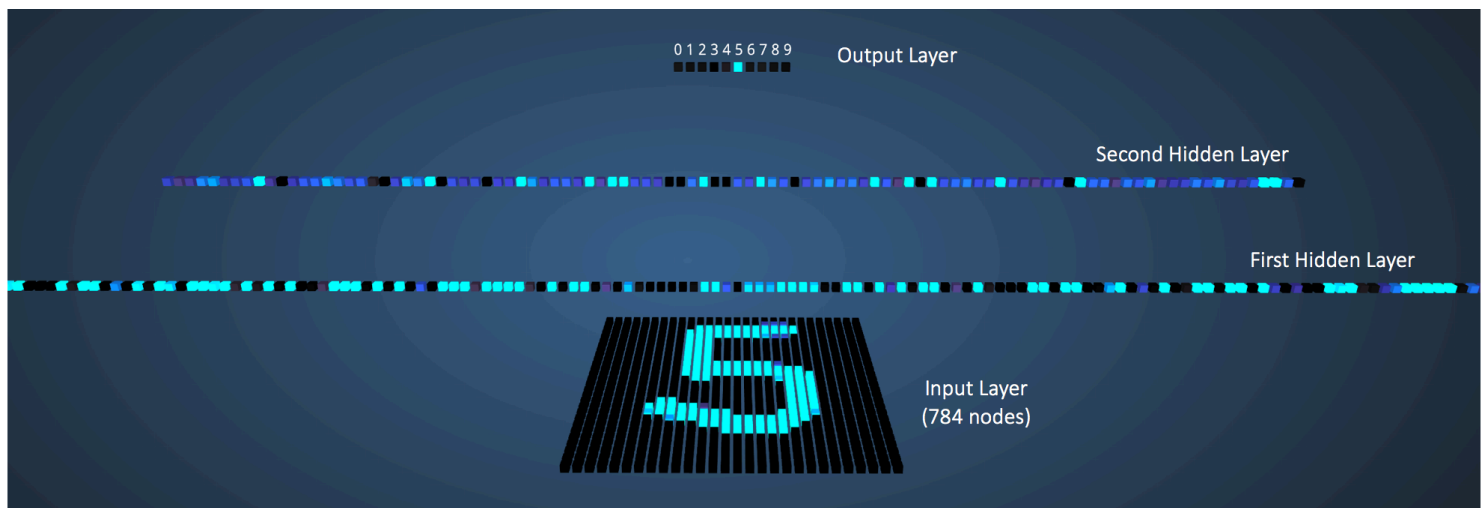


Figure 8: visualizing the network for an input of '5'

A node which has a higher output value than others is represented by a brighter color. In the Input layer, the bright nodes are those which receive higher numerical pixel values as input. Notice how in the output layer, the only bright node corresponds to the digit 5 (it has an output probability of 1, which is higher than the other nine nodes which have an output probability of 0). This indicates that the MLP has correctly classified the input digit. I highly recommend playing around with this visualization and observing connections between nodes of different layers.

Deep Neural Networks

1. [What is the difference between deep learning and usual machine learning?](https://github.com/rasbt/python-machine-learning-book/blob/master/faq/difference-deep-and-normal-learning.md)
(<https://github.com/rasbt/python-machine-learning-book/blob/master/faq/difference-deep-and-normal-learning.md>).
2. [What is the difference between a neural network and a deep neural network?](http://stats.stackexchange.com/questions/182734/what-is-the-difference-between-a-neural-)
(<http://stats.stackexchange.com/questions/182734/what-is-the-difference-between-a-neural->

[network-and-a-deep-neural-network?rq=1](#)).

3. [How is deep learning different from multilayer perceptron?](https://www.quora.com/How-is-deep-learning-different-from-multilayer-perceptron) (<https://www.quora.com/How-is-deep-learning-different-from-multilayer-perceptron>).

Conclusion

I have skipped important details of some of the concepts discussed in this post to facilitate understanding. I would recommend going through [Part1](http://cs231n.github.io/neural-networks-1/) (<http://cs231n.github.io/neural-networks-1/>), [Part2](http://cs231n.github.io/neural-networks-2/) (<http://cs231n.github.io/neural-networks-2/>), [Part3](http://cs231n.github.io/neural-networks-3/) (<http://cs231n.github.io/neural-networks-3/>) and [Case Study](http://cs231n.github.io/neural-networks-case-study/) (<http://cs231n.github.io/neural-networks-case-study/>) from Stanford's Neural Network tutorial for a thorough understanding of Multi Layer Perceptrons.

Let me know in the comments below if you have any questions or suggestions!

References

1. [Artificial Neuron Models](https://www.willamette.edu/~gorr/classes/cs449/ann-overview.html) (<https://www.willamette.edu/~gorr/classes/cs449/ann-overview.html>)
2. [Neural Networks Part 1: Setting up the Architecture](http://cs231n.github.io/neural-networks-1/) (Stanford CNN Tutorial) (<http://cs231n.github.io/neural-networks-1/>).
3. [Wikipedia article on Feed Forward Neural Network](https://en.wikipedia.org/wiki/Feedforward_neural_network) (https://en.wikipedia.org/wiki/Feedforward_neural_network).
4. [Wikipedia article on Perceptron](https://en.wikipedia.org/wiki/Perceptron) (<https://en.wikipedia.org/wiki/Perceptron>).
5. [Single-layer Neural Networks](http://computing.dcu.ie/~humphrys/Notes/Neural/single.neural.html) (Perceptrons) (<http://computing.dcu.ie/~humphrys/Notes/Neural/single.neural.html>).
6. [Single Layer Perceptrons](http://www.cs.stir.ac.uk/courses/ITNP4B/lectures/kms/2-Perceptrons.pdf) (<http://www.cs.stir.ac.uk/courses/ITNP4B/lectures/kms/2-Perceptrons.pdf>).
7. [Weighted Networks – The Perceptron](http://page.mi.fu-berlin.de/rojas/neural/chapter/K3.pdf) (<http://page.mi.fu-berlin.de/rojas/neural/chapter/K3.pdf>).
8. [Neural network models \(supervised\) \(scikit learn documentation\)](http://scikit-learn.org/dev/modules/neural_networks_supervised.html) (http://scikit-learn.org/dev/modules/neural_networks_supervised.html).
9. [What does the hidden layer in a neural network compute?](http://stats.stackexchange.com/a/63163/53914) (<http://stats.stackexchange.com/a/63163/53914>).
10. [How to choose the number of hidden layers and nodes in a feedforward neural network?](http://stats.stackexchange.com/a/1097/53914) (<http://stats.stackexchange.com/a/1097/53914>).
11. [Crash Introduction to Artificial Neural Networks](http://ulcar.uml.edu/~iag/CS/Intro-to-ANN.html) (<http://ulcar.uml.edu/~iag/CS/Intro-to-ANN.html>).
12. [Why the BIAS is necessary in ANN? Should we have separate BIAS for each layer?](http://stackoverflow.com/questions/7175099/why-the-bias-is-necessary-in-ann-should-we-have-separate-bias-for-each-layer) (<http://stackoverflow.com/questions/7175099/why-the-bias-is-necessary-in-ann-should-we-have-separate-bias-for-each-layer>).
13. [Basic Neural Network Tutorial – Theory](https://takinginitiative.wordpress.com/2008/04/03/basic-neural-network-tutorial-theory/) (<https://takinginitiative.wordpress.com/2008/04/03/basic-neural-network-tutorial-theory/>).
14. [Neural Networks Demystified \(Video Series\): Part 1, Welch Labs @ MLconf SF](https://www.youtube.com/watch?v=5MXp9UUKSmc) (<https://www.youtube.com/watch?v=5MXp9UUKSmc>).
15. A. W. Harley, "An Interactive Node-Link Visualization of Convolutional Neural Networks," in ISVC, pages 867-877, 2015 ([link](http://scs.ryerson.ca/~aharley/vis/harley_vis_isvc15.pdf) (http://scs.ryerson.ca/~aharley/vis/harley_vis_isvc15.pdf))

□

47 thoughts on “A Quick Introduction to Neural Networks”

1. ***snendroid*** says: [August 13, 2016 at 6:02 am](#)

That was an amazing article!

[Reply](#)

2. ***Scott*** says: [August 25, 2016 at 5:47 am](#)

This is a great post. The main question I have is what is gained by adding additional hidden layers? Does it reduce error? Is there diminishing returns by adding additional hidden layers into the network? How should someone new to Neural Networks think about the benefits of additional hidden layers vs. the additional CPU cycles and Memory resource costs of a bigger Neural Network?

[Reply](#)

3. ***Sivathanu K (@sivathanu_k)*** says: [August 31, 2016 at 4:16 pm](#)

Thanks for the article...!!

[Reply](#)

4. ***Minwei Shen*** says: [October 26, 2016 at 1:42 am](#)

This is exactly what ‘introduction to neural network’ should be! Thanks for this great article!

[Reply](#)

5. ***Muneera Hmdi*** says: [November 29, 2016 at 2:30 am](#)

Thanks a lot for this amazing article. Honestly, I learned many things from it. Kindly, can u provide like this artificial about the CNN? Thank again.

[Reply](#)

1. ***jxf0623*** says: [June 29, 2018 at 6:07 pm](#)

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

Here is it.

[Reply](#)

6. ***Ke Xu*** says: [January 8, 2017 at 9:09 am](#)

Thanks very much ! After reading several posts and watching couple of videos, I eventually found the most gentle introduction to CNN !

[Reply](#)

7. ***German*** says: [February 24, 2017 at 10:43 pm](#)

Thank you, awesome post!

[Reply](#)

8. ***Barry*** says: [March 3, 2017 at 7:37 am](#)

Good job! Great explanation for beginners!

[Reply](#)

9. ***nnitsew*** says: [April 1, 2017 at 11:53 pm](#)

Thank you for this good article.

Reply

10. **jupi** says: May 12, 2017 at 9:44 am
very helpful, finally I understood, thank you!

Reply

11. **Jasaz** says: May 15, 2017 at 12:09 am
Explained in very simple way. Loved it. Do keep publishing more plz.

Reply

12. **Nobiel Ho** says: July 2, 2017 at 8:51 pm
Great article and very helpful. Thank you!

Reply

13. **Ralph M.** says: July 18, 2017 at 10:19 pm
Great intro on neural networks. I just finished a course on experfy on machine learning – definitely recommend it to anyone who wants to learn more!
<https://www.experfy.com/training/courses/machine-learning-foundations-supervised-learning>

Reply

14. **AA** says: August 2, 2017 at 6:57 am
I think there is a mistake in figure 5, the error of the second output node should be -0.6, shouldn't it?

Reply

ujjwalkarn says: August 3, 2017 at 3:10 pm
Yes, will fix it, thanks for flagging!

Reply

15. **Andy** says: September 14, 2017 at 9:11 am
Fantastic article – this one explained from the ground up. Thank you!

Reply

16. **johnburner91** says: October 11, 2017 at 7:58 am
Thumbs up man. You'er a blessing

Reply

17. **Sam** says: October 17, 2017 at 8:28 pm
Great article, very helpful to me. This is a real "introduction" compares to other information on internet.

Reply

18. **mjbolanos9** says: November 19, 2017 at 11:40 pm
Easy to follow explanation.. Thanks!

Reply

19. **MJ** says: November 20, 2017 at 12:16 am
Detailed article and easy to follow.. Thank you!

Reply

20. **Amy** says: November 25, 2017 at 12:59 pm

This is a FANTASTIC article on ANN. My one question after reading this was “why multiple neurons in the hidden layer” and “why multiple hidden layers.” This began to answer my question: <https://datascience.stackexchange.com/questions/14028/what-is-the-purpose-of-multiple-neurons-in-a-hidden-layer/14030> but I have a lot more learning to do!

Reply

21. **Abu Noman** says: January 3, 2018 at 5:48 pm
Great article! It's easy to catch-up. Keep writing...

Reply

22. **Kelly** says: February 17, 2018 at 10:33 pm
“While a single layer perceptron can only learn linear functions” – Can't there be an activation function such as tanh, therefore it's learning a non-linear function? The Wikipedia article on perceptrons says, “Single layer perceptrons are only capable of learning *linearly separable patterns*” (emphasis added). I'm new to this, but “can only learn linear functions” seems inaccurate – what do you think?

Reply

23. **Asha** says: March 7, 2018 at 2:13 am
Well explained. Thanks!

Reply

24. **Babita Koundal** says: March 31, 2018 at 3:16 pm
A nice explained article...

Reply

25. **rishi** says: May 6, 2018 at 8:06 pm
fantastic article – really helped me understand ANN's for my computer science exam tomorrow!

Reply

26. **Charles Fu** says: May 7, 2018 at 2:42 pm
A nice writing. Figure 5 has a typo: Error = $(0-0.6) = (-0.4)$ should be (-0.6)

Reply

27. **Andrew Bannerman** says: May 16, 2018 at 9:08 am
Great post 😊

Reply

28. **Cedric** says: June 20, 2018 at 1:24 am
This was a great post to explain the very basics to those that are new to Neural Networks. Thanks.

Reply

29. **Than K Ful** says: June 20, 2018 at 3:42 am
Thank you! This is a great intro

Reply

30. **Abhishek Subramanian** says: September 7, 2018 at 6:05 pm
That is a very nice introduction to Neural networks. Great Job

Reply

31. **EricEhlert** says: January 22, 2019 at 7:12 am

Every neuron takes a single number and performs a fixed activation function on it, this I understand. My question is how can this hold true when the input to a neuron on the input layer representing a nominal or categorical feature is not in the form of a single number, but instead a vector?

Reply

1. *anonymous* says: March 24, 2019 at 2:45 pm

ConvNets deals with Vector representation of data like images, audio.

Reply

32. *Chudi Lin* says: April 6, 2019 at 3:54 pm

it's really a amazing article that gives me a lot inspiration,thank you very much!

Reply

33. *刘彦辉* says: June 2, 2019 at 3:18 pm

I'm a chinese.I have searched many papers about ANN. To be honest , this is the best neuron I have ever read.

Reply

34. *qiufengyiyeyu* says: June 2, 2019 at 3:19 pm

I'm a chinese.I have searched many papers about ANN. To be honest , this is the best neuron I have ever read.

Reply

1. *qiufengyiyeyu* says: June 2, 2019 at 3:21 pm
article

Reply

35. *WheresMyButt* says: June 10, 2019 at 2:00 pm

It's amazing that I had to read about 40 articles, tutorials, guides, and documentation pages before I found one that actually started at the beginning. I learned more about the fundamentals from this guide than all the others combined. Thank you, great job

Reply

36. *Curious* says: July 29, 2019 at 2:11 am

Thank you for posting this article. It is very helpful.

All the articles I read on the Web say that "weight" is a property of a connection between two neurons.

But a Wikipedia article says: "The connections between artificial neurons are called 'edges'. Artificial neurons and edges typically have a weight that adjusts as learning proceeds."

https://en.wikipedia.org/wiki/Artificial_neural_network

Does this sentence imply that a neuron also has a weight?

Does a neuron really have a weight?

Reply

37. *nori* says: August 18, 2019 at 9:44 am

thank you for this great article!good job! do you do any publications, so i can make your publication as one of my reference 😊

Reply

38. *Unnikrishnan* says: January 12, 2020 at 1:11 pm

Great article. Thank You.

Reply

39. **Vetrivel PS** says: March 30, 2020 at 5:31 pm

Is the Figure 3. Representing Feedback Neural Network wrong ?

Reply

40. **Vetrivel PS** says: March 30, 2020 at 8:33 pm

3D Visualization of a Multi Layer Perceptron – LINK is not working – Update please

Reply

41. **Himanshu** says: April 24, 2020 at 4:21 pm

this is the great article for beginners but the only question i have is why the data in datasheet is so weird.

Reply

42. **SW** says: May 20, 2020 at 6:55 pm

Excellent examples!

Reply

43. **zhichuv** says: March 4, 2022 at 5:52 pm

thank you for your excellent essay

Reply

