

BACK -PROPAGATION ALGORITHM MECHANISM

1. Prof Sangamesh S K, 2. Prof Pavitra M G

Department of Computer Science and Engineering,
RTE Society's Rural Engineering College Hulkoti

Abstract—: In neural networks learning, the Back - Propagation is an area of extreme importance as an mechanism to reduce errors and get the correct output. That when a prediction is made by the network contains error or not can be identified by Back Propagation. The main aim of propagation is to send some information as motion or information as sound or light and other input in a predefined path or direction through a chosen medium. The primary motive is to delve into deep learning through Back Propagation is about communicating information and transmission that in turn gives a fair estimation of errors commented by the neural network while also suggests the type of data retrieved.

Keywords : - Model initialization, Forward propagate, Loss function, - Differentiation, Back-propagation, Weight update, Iterate until convergence.

1. INTRODUCTION

The back-propagation algorithm was originally introduced in the 1970s, but its importance wasn't fully appreciated until a famous 1986 paper by David Rumelhart, Geoffrey Hinton, and Ronald Williams.

Back-propagation is just a way of propagating the total loss back into the neural network to know how much of the loss every node is responsible for, and subsequently updating the weights in such a way that minimizes the loss by giving the nodes with higher error rates lower weights and vice versa.

Back-propagation explained in a simple way. Any complex system can be abstracted in a simple way, or at least dissected to its basic abstract components. Complexity arises by the accumulation of several simple layers. We will try to reduce the machine learning mechanism in NN to its basic abstract components.

A supervised neural network, at the highest and simplest representation, can be presented as a black box with 2 methods learn and predict as following:

The learning process takes the inputs and the desired outputs and updates its internal state accordingly, so the calculated output get as close as possible to the desired output. The predict process takes an input and generate, using the internal state, the most likely output according to its past "training experience". That's why machine learning is called sometimes model fitting

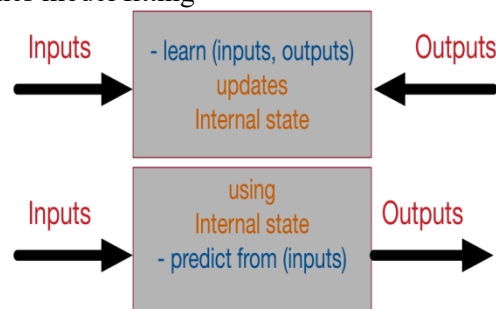


Fig 1: Learning Process

1. **Implementation of BACK-PROPAGATION ALGORITHM**

A simple numerical example

The easiest example to start with neural network and supervised learning, is to start simply with an input and an output and a linear relation between them. The goal of the supervised neural network is to try to search over all possible linear functions which one fits the best the data. Take for instance the following dataset:

Sl.no	Input	Desired output
1	0	0
2	1	2
3	2	3
4	3	4
5	4	6
6	5	8

Table 2.1: Sample Dataset

For this example, it might seem very obvious to you that the output = 2 x input, however it is not the case for most of the real datasets (where the relationship between the input and output is highly non-linear and not that obvious).

Step 1- Model initialization

The first step of the learning is to start from somewhere: the initial hypothesis. Like in genetic algorithms and evolution theory, neural networks can start from anywhere. Thus a random initialization of the model is a common practice. The rationale behind is that from wherever we start, if we are perseverant enough and through an iterative learning process, we can reach the pseudo-ideal model.

In order to give an analogy, take for instance a person who has never played football in his life. The very first time he tries to shoot the ball, he can just shoot it randomly. Similarly, for our numerical case study, let's consider the following random initialization: (Model 1): $y=3.x$. The number 3 here is generated at random. Another random initialization can be: (Model 2): $y=5.x$, or (Model3): $y=0,5.x$. We will explore later, how, through the learning process, all these models can converge to the ideal solution ($y=2.x$) (which we are trying to find).

In this example, we are exploring which model of the generic form $y=W.x$ can fit the best the current dataset. Where W is called the weights of the network and can be initialized randomly. These types of models are simply called linear layers.

Step 2- Forward propagate

The next natural step to do after initializing the model at random is to check its performance. We start from the input we have, we pass them through the network layer and calculate the actual output of the model straightforwardly.

Sl.no	Input	Actual output of model 1 ($y= 3.x$)
1	0	0
2	1	3
3	2	6
4	3	9
5	4	12

Table 2.2: Random Initialization of Model 1



Volume 9, Issue 8 - September 2021-2022 - Pages 1-11

This step is called forward-propagation, because the calculation flow is going in the natural forward direction from the input -> through the neural network -> to the output.

Step 3- Loss function

At this stage, in one hand, we have the actual output of our randomly initialized neural network. On the other hand, we have the desired output we would like the network to learn. Let's put them both in the same table.

Sl.no	Input	Actual output of model	Desired output
1	0	0	0
2	1	3	2
3	2	6	3
4	3	9	4
5	4	12	6
			8

Table 2.3: Output of Randomly Initialized Model

If we compare this to our football player shooting for the first time, the actual output will be the final position of the ball; the desired output would be that the ball goes inside the goal. In the beginning, our player is just shooting randomly. Let's say the ball went most of the time, to the right side of the goal. What he can learn from this is that he needs to shoot a bit more to the left next time he trains.

In order to be able to generalize to any problem, we define what we call: loss function. Basically it is a performance metric on how well the NN manages to reach its goal of generating outputs as close as possible to the desired values.

The most intuitive loss function is simply loss = (Desired output — actual output). However this loss function returns positive values when the network undershoot (prediction < desired output), and negative values when the network overshoot (prediction > desired output). If we want the loss function to reflect an absolute error on the performance regardless if it's overshooting or undershooting we can define it as :loss = Absolute value of (desired — actual).

If we go back to our football player example, if our newbie guy shoots the ball 10m to the right or 10m to the left of the goal, we consider, in both cases, that he missed its target by 10m regardless the direction (right or left). In this case we will add a new column to the table -> the absolute error.

However, several situations can lead to the same total sum of errors: for instance, lot of small errors or few big errors can sum up exactly to the same total amount of error. Since we would like the prediction to work under any situation, it is more preferable to have a distribution of lot of small errors, rather than a few big ones. In order to encourage the NN to converge to such situation, we can define the loss function to be the sum of squares of the absolute errors (which is the most famous loss function in NN). This way, small errors are counted much less than large errors! (the square of 2 is 4, but the square of 10 is 100! So an error of 10, is penalized 25 times more than an error of 2 — not only 5 times!). Our table becomes the following:



Input Error	Actual	Desired	Absolute Error	Square Error
0	0	0	0	0
1	3	2	1	1
2	6	4	2	4
3	9	6	3	9
4	12	8	4	16
Total			10	30

Table 2.4: Final Result

Notice how, if we consider only the first input 0, we can say that the network predicted correctly the result! However, this is just the beginner’s luck in our football player analogy who can manage to score from the first shoot as well. What we care about, is to minimize the overall error over the whole dataset (total of the sum of the squares of the errors!).

As a summary, the loss function is an error metric, that gives an indicator on how much precision we lose, if we replace the real output by the actual output generated by our trained neural network model. That’s why it’s called loss!

Simply speaking, the machine learning goal becomes then to minimize the loss function (to reach as close as to 0). We can just transform our machine learning problem now to an optimization process that aims to minimize this loss function.

Step 4- Differentiation

Obviously we can use any optimization technique that modifies the internal weights of neural networks in order to minimize the total loss function that we previously defined. These techniques can include genetic algorithms or greedy search or even a simple brute-force search: In our simple numerical example, with only one parameter of weight to optimize W, we can search from -1000.0 to +1000.0 step 0.001, which W has the smallest sum of squares of errors over the dataset.

This might work if the model has only very few parameters and we don’t care much about precision. However, if we are training the NN over an array of 600x400 inputs (like in image processing), we can reach very easily models with millions of weights to optimize and brute force can’t be even be imaginable, since it’s a pure waste of computational resources!

Luckily for us, there is a powerful concept in mathematics that can guide us how to optimize the weights called differentiation. Basically it deals with the derivative of the loss function. In mathematics, the derivative of a function at a certain point, gives the rate or the speed of which this function is changing its values at this point. In order to see the effect of the derivative, we can ask ourselves the following question: how much the total error will change if we change the internal weight of the neural network with a certain small value δW . For the sake of simplicity will consider $\delta W=0.0001$. in reality it should be even smaller!.

Volume 9, Issue 8 - September 2021-2022 - Pages 1-11

Let's recalculate the sum of the squares of errors when the weight W changes very slightly:

Now as we can see from this table, if we increase W from 3 to 3.0001, the sum of squares of error will increase from 30 to 30.006. Since we know that the best function that fits this model is $y=2.x$, increasing the weights from 3 to 3.0001 should obviously create a little bit more error (because we are going further from the intuitive correct weight of 2. We have: $3.0001 > 3 > 2$ thus the error is higher).

But what we really care about is the rate of which the error changes relatively to the changes on the weight. Basically here this rate is the increase of 0.006 in the total error for each 0.0001 increasing weight \rightarrow that's a rate of $0.006/0.0001 = 60x!$ It works in both direction, so basically if we decrease the weights by 0.0001, we should be able to decrease the total error by 0.006 as well! Here is the proof, if we run again the calculation at $W=2.9999$, we get an error of 29.994. We managed to decrease the total error by 0.006, we learned something! We improved the model! We could have guessed this rate by calculating directly the derivative of the loss function. The advantage of using the mathematical derivative is that it is much faster and more precise to calculate (less floating point precision problems).

Here is what our loss function looks like:

- If $w=2$, we have a loss of 0, since the neural network actual output will fit perfectly the training set.
- If $w < 2$, we have a positive loss function, but the derivative is negative, meaning that an increase of weight will decrease the loss function.
- At $w=2$, the loss is 0 and the derivative is 0, we reached a perfect model, nothing is needed.
- If $w > 2$, the loss becomes positive again, but the derivative is as well positive, meaning that any more increase in the weight, will increase the losses even more!!

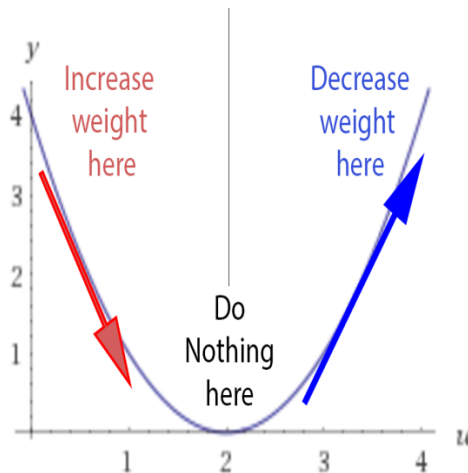


Fig 2: Differentiation curve


Arrows represent the derivative at the corresponding points

If we initialize randomly the network, we are putting any random point on this curve (let's say $w=3$). The learning process is actually saying this:

- Let's check the derivative.
- If it is positive, meaning the error increases if we increase the weights, then we should decrease the weight.
- If it's negative, meaning the error decreases if we increase the weights, then we should increase the weight.
- If it's 0, we do nothing, we reach our stable point.

Volume 9, Issue 8 - September 2021-2022 - Pages 1-11

In a simple matter, we are designing a process that acts like gravity. No matter where we randomly initialize the ball on this error function curve, there is a kind of force field that drives the ball back to the lowest energy level of ground 0.



Schematic of gradient descent.

Fig 3: Schematic of Gradient Descent

Step 5- Back-propagation

In this example, we used only one layer inside the neural network between the inputs and the outputs. In many cases, more layers are needed, in order to reach more variations in the functionality of the neural network.

For sure, we can always create one complicated function that represents the composition over the whole layers of the network. For instance, if layer 1 is doing: $3 \cdot x$ to generate a hidden output z , and layer 2 is doing: z^2 to generate the final output, the composed network will be doing $(3 \cdot x)^2 = 9 \cdot x^2$. However in most cases composing the functions is very hard. Plus for every composition one has to calculate the dedicated derivative of the composition (which is not at all scalable and very error prone).

In order to solve the problem, luckily for us, derivative is decomposable, thus can be back-propagated.

We have the starting point of errors, which is the loss function, and we know how to calculate its derivative, and if we know how to calculate the derivative of each function from the composition, we can propagate back the error from the end to the start.

Let's consider the simple linear example: where we multiply the input 3 times to get a hidden layer, then we multiply the hidden (middle layer) 2 times to get the output.

input $\rightarrow 3 \cdot x \rightarrow 2 \cdot x \rightarrow$ output.

A 0.001 delta change on the input will be translated to a 0.003 delta change after the first layer, then to 0.006 delta change on the output, which is the case if we compose both functions into one:

input $\rightarrow 6 \cdot x \rightarrow$ output.

Volume 9, Issue 8 - September 2021-2022 - Pages 1-11

Similarly an error on the output of 0.006 can be back-propagated to an error of 0.003 in the middle hidden stage, then to 0.001 on the input.

If we create a library of differentiable functions or layers where for each function we know how to forward-propagate (by directly applying the function) and how to back-propagate (by calculating the derivative of the function), we can compose any complex neural network.

We only need to keep a stack of the function calls during the forward pass and their parameters, in order to know the way back to back-propagate the errors using the derivatives of these functions. This can be done by de-stacking through the function calls. This technique is called auto-differentiation, and requires only that each function is provided with the implementation of its derivative. In a future blog post, we will explain how to accelerate auto-differentiation by implementing basic mathematical operations over vectors/matrices/and tensors.

In neural network, any layer can forward its results to many other layers, in this case, in order to do back-propagation, we sum the deltas coming from all the target layers. Thus our linear calculation stack can become a complex calculation graph.

This figure shows the process of back-propagating errors following this schemas:

Input -> Forward calls -> Loss function -> derivative -> back-propagation of errors. At each stage we get the deltas on the weights of this stage.

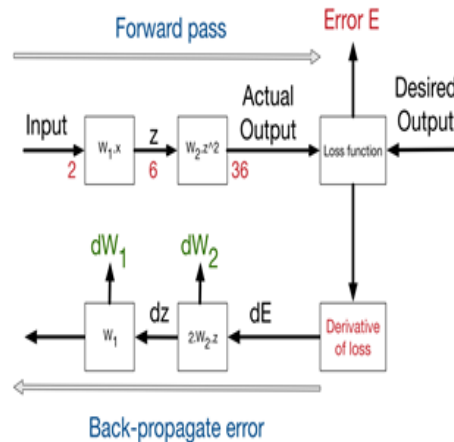


Fig 4: Diagram of Forward and Backward paths




Fig 5: Diagram of Forward and Backward propagation

Cool animation for the forward and backward paths

Step 6- Weight update

As we presented earlier, the derivative is just the rate of which the error changes relatively to the weight changes. In the numerical example presented earlier, this rate is 60x. Meaning that 1 unit of change in weights leads to 60 total units of change in the loss function.

And since we know that the loss error is currently at 30 units, by extrapolating the rate, in order to reduce the error to 0, we need to reduce the weights by 0.5 units.

However, for real-life problems we shouldn't update the weights with such big steps. Since there are lot of non-linearity, any big change in weights will lead to a chaotic behavior. We should not forget that the derivative is only local at the point where we are calculating the derivative.

Thus as a general rule of weight updates is the delta rule:

New weight = old weight — Derivative * learning rate

The learning rate is introduced as a constant (usually very small), in order to force the weight to get updated very smoothly and slowly (to avoid big steps and chaotic behavior). (To remember: Learn slow and steady!)

In order to validate this equation:

- If the derivative rate is positive, it means that an increase in weight will increase the error, thus the new weight should be smaller.
- If the derivative rate is negative, it means that an increase in weight will decrease the error, thus we need to increase the weights.
- If the derivative is 0, it means that we are in a stable minimum. Thus, no update on the weights is needed -> we reached a stable state.

Now several weight update methods exist. These methods are often called optimizers. The delta rule is the most simple and intuitive one, however it has several draw-backs. This excellent blog post presents the different methods available to update the weights.



Volume 9, Issue 8 - September 2021-2022 - Pages 1-11

In the numerical example we presented here, we only have 5 input/output training set. In reality, we might have millions of entries. Previously, we were talking about minimizing the error cost function (the loss function) over the whole dataset. This is called batch learning, and might be very slow for big data. What we can do instead, is to update the weights every $\text{BatchSize}=N$ of training, providing that the dataset is shuffled randomly. This is called mini-batch gradient descent. And if $N=1$, we call this case full online-learning or stochastic gradient descent, since we are updating the weights after each single input output observed!

Any optimizer can work with these 3 modes (full online/mini-batch/full-batch).

Step 7- Iterate until convergence

Since we update the weights with a small delta step at a time, it will take several iterations in order to learn.

This is very similar to genetic algorithms where after each generation we apply a small mutation rate and the fittest survives.

In neural network, after each iteration, the gradient descent force updates the weights towards less and less global loss function.

The similarity is that the delta rule acts as a mutation operator, and the loss function acts a fitness function to minimize.

The difference is that in genetic algorithms, the mutation is blind. Some mutations are bad, some are good, but the good ones have higher chance to survive. The weight update in NN are however smarter since they are guided by the decreasing gradient force over the error.

How many iterations are needed to converge?

•This depends on how strong the learning rate we are applying. High learning rate means faster learning, but with higher chance of instability and higher chance of sub-optimal results.

- It depends as well on the meta-parameters of the network (how many layers, how complex the non-linear functions are). The more it has variables the more it takes time to converge, but the higher precision it can reach.
- It depends on the optimization method used, some weight updates rule are proven to be faster than others.
- It depends on the random initialization of the network. Maybe with some luck you will initialize the network with ****W=1.99**** and you are only just 0.01 step away from the optimal solution.
- It depends on the quality of the training set. If the input and output has no correlation between each other, the neural network will not do magic and can't learn a random correlation.

In order to summarize, here is what the learning process on neural networks looks like (A full picture):

Volume 9, Issue 8 - September 2021-2022 - Pages 1-11

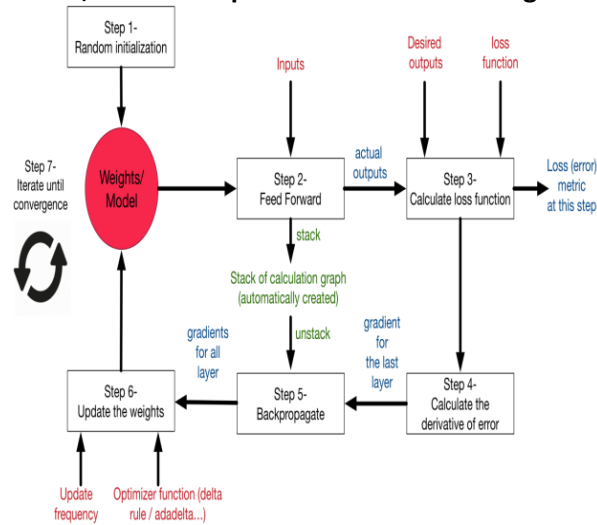


Fig 6: Neural networks step-by-step

3. ADVANTAGES OF BACK-PROPAGATION ALGORITHM

- Simplifies the network structure by elements weighted links that have the least effect on the trained network
- You need to study a group of input and activation values to develop the relationship between the input and hidden unit layers.
- It helps to assess the impact that a give n input variable has on a network output. The knowledge gained from this analysis should be represented in rules.
- Back-propagation is especially useful for deep neural networks working on error-prone projects, such as image or speech recognition.
- Back-propagation takes advantage of the chain and power rules allows back-propagation to function with any number of outputs.

References:

[1] Atkinson, P., and Tatnall, A. 1997. "Introduction Neural Networks in Remote Sensing." International Journal of Remote Sensing 18 (4): 699-709.

[2] Jindal, S., and Josan, G. 2007. "Neural Network and Fuzzy Logic Approach for Satellite Image Classification: A Review." In Proceedings of COLT, 1-4.

[3] Jiang, J., Zhang, J., Yang, G., Zhang, D., and Zhang, L. 2010. "Application of Back Propagation Neural Network in the Classification of High Resolution Remote Sensing Image: Take Remote Sensing Image of Beijing for Instance." In Proceedings of 18th International Conference on Geoinformatics, IEEE, 1-6.

[4] Nicoletti, G. 2000. "An Analysis of Neural Networks as Simulators and Emulators." Cybernetics and Systems 31 (3): 253-82.

[5] Alsmadi, M., Omar, K., and Noah, S. 2009. "Back Propagation Algorithm: The Best Algorithm among the Multi-layer Perceptron Algorithm," IJCSNS International Journal of Computer Science and Network Security 9 (4): 378-83.

[6] Abu-Mahfouz, I. 2005. "A Comparative Study of Three Artificial Neural Networks for the Detection and Classification of Gear Faults." International Journal of General Systems 34 (3): 261-77.



Volume 9, Issue 8 - September 2021-2022 - Pages 1-11

- [7] Graupe, D. 2007. Principles of Artificial Neural Networks, Singapore, Hackensack, N. J.: World Scientific.
- [8] Rojas, R. 1996. Neural Networks: A Systematic Introduction. Berlin: Springer.
- [9] Wang, T. S., Chen, L., Tan, C. H., Yeh, H. C., and Tsai, Y. C. 2009. "BPN for Land Cover Classification by Using Remotely Sensed Data." In Proceedings of Fifth International Conference on Natural Computation, IEEE, 535-9.
- [10] Richards, J. 2006, Remote Sensing Digital Image Analysis, Berlin: Springer-Verlag.
- [11] Mas, J. F., Mas, and Flores, J. J. 2008. "The Application of Artificial Neural Networks to the Analysis of Remotely Sensed Data." International Journal of Remote Sensing 39 (3): 617-63
- [12].<https://www.guru99.com/backpropogation-neural-network.html>
- [13]. Tom M. Mitchell, Machine Learning, India Edition 2013, McGraw Hill Education.