Foundations of Deep Learning: Mathematical Details

1 What Is the Main Objective of Deep Learning

Traditional neural networks were invented when computers were reasonably slow, which prevented them from solving important practical problems. For these computers, computation speed was the main objective. As we have shown, this need led to what we know as traditional neural networks.

Nowadays, computers are much faster. In most practical applications, speed is no longer the main problem. But now a new problem emerged: that the traditional neural networks, while fast, have limited accuracy of their predictions. It is therefore desirable to come up with devices that would achieve better prediction *accuracy* – even if to achieve this accuracy, they will be slower than the traditional neural networks. This is, in a nutshell, the main motivation behind what is now known as deep learning.

Possible question. *What is the main objective is deep learning?*

2 The More Models We Have, the More Accurately We Can Approximate

As a result of training a neural network – or any other machine learning tool – we get the values of some parameters for which the corresponding models provides the best approximation to the actual data. There are usually finitely many parameters, each of which – when represented in a computer – consists of finitely many bits. We can write down all these bits into a single binary sequence.

Let us denote the length of this sequence by N. If we have a parameters each of which consists of b bits, then $N = a \cdot b$.

Possible question. *If after training, we get the values of 3 parameters each of which is described by a 8-bit number, how many bits do we need to store the information about all possible combinations of these parameters?*

Answer. *In this case,* $a = 3$ *and* $b = 8$*, so we need* $3 \cdot 8 = 24$ *bits.* □

In these terms, different models that can be obtained from training can be described by different N-bit sequences.

There are two possible values of each bit. For each of these values, we can have two different 2-but extensions depending on whether we extend the original bit with 0 or 1, so we have $2 \cdot 2 = 2^2$ possible sequences of length 2. For each of these 2-bit sequences, we have two possible 3-bits extensions, so the overall number of 3-bit sequences is $2 \cdot 2^2 = 2^3$. In general, for b bits, we have 2^b possible b-bit sequences.

Possible question. *How many possible 10-bit binary sequences can we have?*

Answer. *We can have* $2^{10} = 1024$ *different 10-bit binary sequences.* □

In particular, if we use N bits to store all the information about all the tuned parameters, then we can have 2^N possible binary sequences and thus, 2^N possible models. In these terms, training simply means selecting one of these 2^N possible models.

Possible question. *In general, how many different models can you have if you use 2 variables with 8 bits each?*

Answer. *In this case, we have* $N = 2 \cdot 8 = 16$ *bits, so, in general, by using all these bits, we can represent* $2^{16} \approx 65,000$ *different models.* □.

If we have only one model to represent the actual dependence, this model will be a very lousy description. If we can have two models, we can have more accurate approximations. If we can have 100 different models, we can have an even more accurate one. In general, the more models we have, the more accurate representation we can have.

We can illustrate this idea on the example of approximating real numbers from the interval $[0, 1]$. If we have only one model – e.g., the value $x = 0.5$, then we approximate every other number with accuracy 0.5. If we can have 10 models, then we can take 10 values 0.05, 0.15, ..., 0.95. The first value approximates all the numbers from the interval $[0, 0.1]$ with accuracy 0.05. The second value approximates all the numbers from the interval $[0, 1, 0.2]$ with the same accuracy, etc. As a result, by selecting one these values, we can approximate any number from the interval $[0, 1]$ with accuracy 0.05.

If we can have 100 models, then can similarly take the values 0.005, 0.015, ..., all the way to 0.995 and this approximate any number from the interval $[0, 1]$ with accuracy 0.005.

3 How Many Models Can We Represent with a Traditional Neural Network

Let us consider a traditional neural network with K neurons. Each neuron k is characterized by several weights W_k and w_{ki}. Let b denote the number of bits needed to describe all the weights corresponding to a single neuron. Then, overall, to describe all possible bit sequences resulted from training, we need $N = K \cdot b$ bits.

As we mentioned, we can have 2^N different binary sequences of length N. So, at first glance, one may think that, in line with what we have discussed earlier, we can thus represent 2^N different models. However, the actual number of models is much smaller. The reason for this is that if we swap two neurons – i.e., swap all the bits corresponding to these two neurons – the resulting functions

$$f(x_1, \ldots, x_n) = \sum_{k=1}^{K} W_k \cdot s\left(\sum_{i=1}^{n} w_{ki} \cdot x_i - w_{k0}\right) - W_0$$

will not change – since the sum does not change if we swap two of added numbers.

For example, if we have two neurons, the first one is described by a binary sequence 010 and the second one by 110, then this whole configuration is described by the sequence 010110. If we swap the neurons, we get a different binary sequence 110010 which, however, represents the same dependence – i.e., the same model.

Similarly, if instead of swapping two neurons, we apply any permutation, we get the exact same model.

For K neurons, there are $K!$ possible permutations.

Possible question. *If we have 5 neurons, how many possible permutations can we have that do not change the model?*

Answer. *We can have* $5! = 1 \cdot 2 \cdot 3 \cdot 4 \cdot 5 = 120$ *different permutations.* □.

Thus, $K!$ different binary sequences represent the same model. So, by using N neurons, instead of 2^N possible models, we can only have $\dfrac{2^N}{K!}$ possible models.

Possible question. *How many models can be represent by a traditional neural network of K neuron that uses N bits to store all the weights? Explain your answer.*

Possible question. *If we use 5 neurons each of which uses 3 bits, how many possible models can we represent? How many models could we represent in general if we use all these bits?*

Answer. *We use* $N = 5 \cdot 3 = 15$ *bits. So, in general, we could represent* $2^{15} = 32768$ *different models.*

However, because of the permutations, we can use only $2^{15}/5! = 32768/120 \approx 273$ *models – which is much much less than the original 32 thousand.* □

4 How Can We Achieve Better Accuracy: The Main Idea Behind Deep Learning

As we have mentioned, the more models we can represent, the more accurate will be the resulting approximation. When the overall number of bits is fixed –

e.g., by the ability of our computers – the only way to increase the number of models is to decrease $K!$, i.e., to decrease K.

In the traditional neural networks, all the neurons are, in effect, in one layer – known as the hidden layer. The only way to decrease K is to make the number of neurons in each layer much smaller. This means that instead of placing the neurons into a single layer, we place then in many layers. Each of these layers is *narrow* – in the sense that the number of neurons in each later is small. We now have several layers – the construction is *deep* – so that the inputs go to the first layer, then the outputs of the first layer serve as inputs of the second layer, etc. – until we reach the final layer, which produces the resulting value y.

Possible question. *Explain why, to get a better accuracy, we need to go from the traditional neural network with many neurons in a single layer to a multi-layer structure with a small number of neurons in each layer.*

5 Which Activation Function Should We Use for Deep Learning?

To answer this question, we need to recall that usually, we process the values of physical quantities, but the numerical values of physical quantities depend on what measuring unit we use and – for some quantities like temperature or time – what starting point we select for the measurement.

If we change a measuring unit to a one which is λ times smaller – e.g., replace feet with inches which are 12 times smaller – then all numerical values get multiplied by λ – e.g., 2.5 feet becomes 30 inches. So, instead of the original numerical value x, we get a new numerical value $x' = \lambda \cdot x$.

Similarly, if we replace the original starting point with the new point which is x_0 units before, then each numerical value x is replaced by a new numerical value $x' = x + x_0$. For example, if instead of measuring years from our current Year 0 (supposedly the year when Jesus was born, although many historians disagree), we use Year -3761 (supposedly the year when the world was created, although here no historians agree), then current year 2019 becomes $2019 + 3761 = 5780$, which is our year in the traditional Jewish calendar.

We want to select an activation function $s(x)$ that would not depend on the choice of a measuring unit. In other words, we want to make sure that if $y = s(x)$ and we select a new measuring unit, i.e., switch to new numerical values $x' = \lambda \cdot x$ and $y' = \lambda \cdot y$, then for these new values x' and y', we will have the exact same dependence: $y' = s(x')$.

Substituting the expressions $x' = \lambda \cdot x$ and $y' = \lambda \cdot y$ into this formula, we conclude that $\lambda \cdot y = s(\lambda \cdot x)$. Here, $y = s(x)$, so we conclude that

$$s(\lambda \cdot x) = \lambda \cdot s(x)$$

for all possible x and $\lambda > 0$.

For $x = 1$, we conclude that $s(\lambda) = \lambda \cdot s(1)$. If we denote $s(1)$ by c_+, and rename λ into z, we conclude that for all $z > 0$, we get $s(z) = c_+ \cdot z$.

For $x = -1$, we conclude that $s(-\lambda) = \lambda \cdot s(-1)$. If we denote $-s(-1)$ by c_- (so that $s(-1) = -c_-$) and denote $-\lambda$ by z (so that $\lambda = -z$), we conclude that for all negative values z, we have $s(z) = (-c_-) \cdot (-z) = c_- \cdot z$.

Thus, we conclude that the activation function $s(z)$ should have the following *piecewise linear* form:

- for $z > 0$, we have $s(z) = c_+ \cdot z$;
- for $z < 0$, we have $s(z) = c_- \cdot z$.

Possible question. *Explain why for deep neural networks, the activation function should be piece-wise linear.*

Comment. We must have $c_+ \neq c_-$, since otherwise, the function $s(z)$ would be linear, and we know that with linear functions, we can only describe linear dependencies.

Possible question. *Why should we have $c_+ \neq c_-$ in the activation function?*

6 What Activation Function Is Actually Used in Deep Learning? Why

To uniquely determine a piecewise linear function, we need to select two real numbers: c_+ and c_-. The simplest possible real numbers is 0 and 1. Thus, the simplest possible piecewise linear function has the form:

- for $z > 0$, we have $s(z) = z$;
- for $z < 0$, we have $s(z) = 0$.

In other words, $s(z) = \max(z, 0)$. This function is known as *rectified linear function*, it is actually used in deep learning.

Possible question. *What activation function is actually used in deep learning? Why?*

7 It Does Not Matter Which Piecewise Linear Activation Function to Use

Why not use some other piecewise linear function? Because it does not matter much. Indeed, for all the layers – expect for the last one – the output of a neuron is linearly combined with other signals anyway. And any piecewise linear function can be represented as a linear combination of the rectified linear function $\max(z, 0)$ and the identity function z:

$$s(z) = c_- \cdot z + (c_+ - c_-) \cdot \max(z, 0).$$

Indeed:

- for $z > 0$, the right-hand side is equal to $c_- \cdot z + c_+ \cdot 0 = c_- \cdot z$, while
- for $z < 0$, the right-hand side is equal to

$$c_- \cdot z + (c_+ - c_-) \cdot z = (c_- + (c_+ - c_-)) \cdot z = c_+ \cdot z.$$

Possible question. *Prove that every piece-wise linear function can be represented as a linear combination of the rectified linear function and the identity function.*

Possible question. *Represent the following function as a linear combination of the rectified linear function and the identity function:*

- *for $z > 0$, we have $s(z) = 3z$;*
- *for $z < 0$, we have $s(z) = -2z$.*

8 Why Cannot We Require Shift-Invariance Instead of Scale-Invariance?

We mentioned that the numerical value of a physical quantity changes when we change the measuring unit *and* when we change the starting point. However, when we looked for the activation function, we only considered invariance with respect to changing the unit (it is known as *scale-invariance*). Why not consider invariance with respect to changing the starting point (which is called *shift-invariance*)?

Let us give it a try. Similarly to the case of scale-invariance, we want to make sure that when $y = s(x)$ then for $x' = x + x_0$ and $y' = y + x_0$, we will have $y' = s(x')$. Substituting the expressions for x' and y' into the formula $y' = s(x')$, we get $y + x_0 = s(x + x_0)$. Here, $s(x) = y$, so we have

$$s(x + x_0) = s(x) + x_0$$

for all possible values x and x_0.

In particular, for $x = 0$, we get $s(x_0) = s(0) + x_0$. Renaming $s(0)$ as a and x_0 as z, we conclude that $s(z) = z + a$. This is a linear function – thus, such neurons cannot describe any non-linear process.

Possible question. *Why cannot we require shift-invariance instead of scale-invariance?*

9 Need for Pooling

Often, we have a lot of data points to process. For example, an image is usually represented in a computer by the intensities at each pixel. Even for a not very good 1000 by 1000 picture, we have 1,000,000 pixels – so to process such an image, we need to process 1,000,000 numbers.

In a traditional neural network, we could use as many neurons as needed, but in a deep neural network, there are only a few neurons in the first layer. Thus, before we start processing, we need to combined several input values into one (a similar procedure can also be applied at a later stage). This operation of combining several values into one is known as *pooling*.

Possible question. *Why do we need pooling?*

10 Which Pooling Operation Shall We Use?

Let us consider the case when we pool two values a and b into a single value c. Let us denote the resulting value c by $p(a, b)$. Of course, the pooling should not depend on the order, i.e., we should have $p(a, b) = p(b, a)$. In other words, the pooling operation should be *commutative*.

Similarly to the above discussions, it is reasonable to require that the result of pooling will not change if we simply change the measuring unit or change the starting point for measurement. So, we make two requirements:

- if $c = p(a, b)$, then $c' = p(a', b')$, where $a' = \lambda \cdot a$, $b' = \lambda \cdot b$, and $c' = \lambda \cdot c$;
- if $c = p(a, b)$, then $c' = p(a', b')$, where $a' = a + a_0$, $b' = b + a_0$, and $c' = c + a_0$.

From the first requirement, substituting the expressions $a' = \lambda \cdot a$, $b' = \lambda \cdot b$, and $c' = \lambda \cdot c$ into the formula $c' = p(a', b')$, we conclude that $\lambda \cdot c = p(\lambda \cdot a, \lambda \cdot b)$. Here, $c = p(a, b)$, so we conclude that

$$p(\lambda \cdot a, \lambda \cdot b) = \lambda \cdot p(a, b).$$

From the second requirement, substituting the expressions $a' = a + a_0$, $b' = b + a_0$, and $c' = c + a_0$ into the formula $c' = p(a', b')$, we conclude that $c + a_0 = p(a + a_0, b + a_0)$. Here, $c = p(a, b)$, so we conclude that

$$p(a + a_0, b + a_0) = p(a, b) + a_0.$$

Let us use the resulting formulas to find the value $p(x, y)$ for all possible pairs (x, y). Without losing generality, we can assume that $x < y$. Then, substituting $a = 0$, $a_0 = x$, and $b = y - x$ into the formula $p(a + a_0, b + a_0) = p(a, b) + a_0$, we conclude that

$$p(x, y) = p(0, y - x) + x.$$

Substituting $\lambda = y - x$, $a = 0$, and $b = 1$ into the formula $p(\lambda \cdot a, \lambda \cdot b) = \lambda \cdot p(a, b)$, we conclude that $p(0, y - x) = (y - x) \cdot p(0, 1)$. Substituting this expression into the formula $p(x, y) = p(0, y - x) + x$ and denoting $p(0, 1)$ by α, we conclude that

$$p(x, y) = x + \alpha \cdot (y - x) = \alpha \cdot y + (1 - \alpha) \cdot x.$$

Once we learn how to pool two values, we can pool four values easily:

- divide the four values into two pairs,
- pool results within each pair, and then
- pool the two pooling results into a single value.

It is reasonable to require that the result should not depend on how exactly we divide the four original values into pairs. Let us consider the values 0, 1, 1, and 2.

First, we combine 0 with 1 and 1 with 2. Pooling 0 and 1 results in

$$\alpha \cdot 1 + (1 - \alpha) \cdot 0 = \alpha.$$

Pooling 1 and 2 results in

$$\alpha \cdot 2 + (1 - \alpha) \cdot 1 = 2\alpha + 1 - \alpha = 1 + \alpha.$$

Here always $1 + \alpha$ is larger than α, so combining the results α and $1 + \alpha$ leads to

$$\alpha \cdot (1 + \alpha) + (1 - \alpha) \cdot \alpha = \alpha + \alpha^2 + \alpha - \alpha^2 = 2\alpha.$$

What if we instead combine 1 with 1 and 0 with 2? Combining 1 with 1 results in

$$\alpha \cdot 1 + (1 - \alpha) \cdot 1 = 1.$$

Pooling 0 with 2 results in

$$\alpha \cdot 2 + (1 - \alpha) \cdot 0 = 2\alpha.$$

The resulting of pooling the resulting too values 1 and 2α depends on which of the two values is larger.

- If $2\alpha \geq 1$, i.e., if $\alpha \geq 0.5$, then we get

  $$\alpha \cdot (2\alpha) + (1 - \alpha) \cdot 1 = 2\alpha^2 + 1 - \alpha.$$

  In this case, the desired equality is

  $$2\alpha^2 + 1 - \alpha = 2\alpha,$$

  i.e.,

  $$2\alpha^2 - 3\alpha + 1 = 0.$$

  One can easily check that this quadratic equation has two solutions: $\alpha = 0.5$ and $\alpha = 1$.

- If $2\alpha \leq 1$, i.e., if $\alpha \leq 0.5$, then we get

  $$\alpha \cdot 1 + (1 - \alpha) \cdot 2\alpha = \alpha + 2\alpha - 2\alpha^2 = 3\alpha - 2\alpha^2.$$

  In this case, the desired equality is

  $$3\alpha - 2\alpha^2 = 2\alpha,$$

i.e.,

$$2\alpha^2 - \alpha = 0.$$

One can easily check that this quadratic equation has two solutions: $\alpha = 0$ and $\alpha = 0.5$.

So, we have three options: $\alpha = 0$, $\alpha = 0.5$, and $\alpha = 1$.

- If $\alpha = 0$, then the pooling formula takes the form $p(x, y) = 0 \cdot y + 1 \cdot x = x$, i.e., since $x \le y$, the form

$$p(x, y) = \min(x, y).$$

- If $\alpha = 0.5$, then the pooling formula takes the form $p(x, y) = 0.5 \cdot y + 0.5 \cdot x$, i.e., to the arithmetic average

$$p(x, y) = \frac{x + y}{2}.$$

- If $\alpha = 1$, then the pooling formula takes the form $p(x, y) = 1 \cdot y + 0 \cdot x = y$, i.e., since $x \le y$, the form

$$p(x, y) = \max(x, y).$$

These three operations – minimum, maximum, and arithmetic averge – are indeed the ones which work most successfully in deep learning.

Possible question. *Which pooling operations work most successfully in deep learning? Provide a theoretical explanation for these operations.*

Possible question. *Suppose that instead of the three operations, we use the operation*

$$p(x, y) = \alpha \cdot \max(x, y) + (1 - \alpha) \cdot \min(x, y)$$

with $\alpha = 0.4$. Provide an example that for this operation, the result of pooling four numbers may depend on how we divide the original four numbers into pairs.

11 Sensitivity of Deep Learning: Phenomenon

One of the problems with deep learning is that its results are often too sensitive to minor changes in the inputs. For example, changing a few pixels in a picture of a cat may result in this picture being misclassified as a dog.

This is a very troublesome features, since in practice, signals often come with noise, and it is not good that a small noise can ruin the results.

Possible question. *Explain what researchers mean when they say that deep learning is too sensitive. Explain why this is a big problem in practical applications.*

12 Sensitivity of Deep Leaning: An Explanation

Each neuron is affected by the noise. It can take the original noise level δ and amplify it to a higher level $c \cdot \alpha$ for some $c > 1$.

In deep learning, we have several (L) layers. In the first layer, each neuron amplifies the noise level δ to $c \cdot \delta$. Neurons in the second layer amplify it even more, to $c \cdot (c \cdot \delta) = c^2 \cdot \delta$. After the third layer, we get $c^3 \cdot \delta$, etc., and after all L layers, we get $c^L \cdot \delta$. The exponential function c^L grows very fast with L, so, not surprisingly, we get a much higher noise level than for the traditional neural networks.

Possible question. *Explain why deep neural networks are very vulnerable to the noise in the original data, much more vulnerable than the traditional neural networks.*

Possible question. *If neurons on each layer amplify the noise by a factor of 2, how much will the noise be amplified after passing a 10-layer neural network?*

Answer. *Here, $c = 2$ and $L = 10$, so the noise will be amplified by a factor $c^L = 2^{10} = 1024$.* $\square$.

13 How to Deal with Sensitivity of Deep Learning

To train a traditional neural network, we feed it with actually observed patterns $\left(x_1^{(p)}, \ldots, x_n^{(p)}, y^{(p)}\right)$, and find the values of the corresponding weights that match all these patterns. As a result, the trained network usually works well not only for the original patterns, but also for modified versions of these patterns – e.g., it still works well when we add some noise to the original pattern.

For deep learning, we do not have automatic success on noised patterns. So, to achieve such success, it is reasonable to artificially add noise to the patterns and add such simulated-noise modification to the original patterns when training a network. This idea seems to work reasonably well.

14 Remaining Big Problem with Deep Learning

The remaining big problem with deep learning is that the network does not explain how it got its results. We know that sometimes it makes a mistake, but without an explanation, it is not clear whether a strange recommendation is a stroke of genius or a silly mistake (such as interpreting a clear cat picture as a dog).