

Lecture 2: Backpropagation

Peter Bloem
Deep Learning

dlvu.github.io



Today's lecture will be entirely devoted to the backpropagation algorithm. The heart of all deep learning.

STOCHASTIC GRADIENT DESCENT

pick some initial weights θ (for the whole model)

loop: ← epochs

for x, t in Data:

$$\theta \leftarrow \theta - \alpha \cdot \nabla_{\theta} \text{loss}_{x,t}(\theta)$$

↑ learning rate

stochastic gradient descent: loss over one example per step (loop over data)

minibatch gradient descent: loss over a few examples per step (loop over data)

full-batch gradient descent: loss over the whole data

2



This is where we left things last lecture. We introduced **gradient descent** as the training algorithm for neural networks. The gradient collects all derivatives of the loss with respect to every weight in the neural network.

For simple models, like a one-layer linear or logistic regression, we can work out the gradient in *closed form*. That means we get a function of what the gradient is at every point in our parameter space.

For complex, multilayer functions, working out the gradient is not so easy. We need an algorithm to do it for us. That algorithm is called **backpropagation**, and it's what powers deep learning.

THE PLAN

part 1: scalar backpropagation

part 2: tensor backpropagation

part 3: automatic differentiation

3



Because backpropagation is so important, we are going to spend an entire lecture on it. Looking at it from different perspectives, and slowly building up to the completely automated, highly parallelized version of the algorithm that we use today.

In **the first part** we describe backpropagation in a *scalar* setting. That is, we will treat each individual element of the neural network as a single number, and simply loop over all these numbers to do backpropagation over the whole network. This simplifies the derivation, but it is ultimately a slow algorithm with a complicated notation.


In **the second part**, we translate neural networks to operations on vectors, matrices and tensors. This allows us to simplify our notation, and more importantly, massively speed up the computation of neural networks. Backpropagation on tensors is a little more difficult to do than backpropagation on scalars, but it's well worth the effort.

In **the third part**, we will make the final leap from manually worked out and implemented backpropagation system to full-fledged *automatic differentiation*: we will show you how to build a system that takes care of the gradient computation entirely by itself. This is the technology behind software like pytorch and tensorflow.

Lecture 2: Backpropagation

Peter Bloem
Deep Learning


dlvu.github.io



In the previous video, we looked at what neural networks are, and we saw that to train them, we need to work out the derivatives of the loss with respect to the parameters of the neural network: collectively these derivatives are known as *the gradient*.


PART ONE: SCALAR BACKPROPAGATION

How do we work out the gradient for a neural network?



|section|Scalar backpropagation|

|video|https://www.youtube.com/embed/idO5r5eWlrw?si=NnUUgtAroD3_Rich|



Here is a diagram of the sort of network we'll be encountering (this one is called the GoogLeNet). We can't work out a complete gradient for this kind of architecture by hand. We need help. What we want is some sort of algorithm that lets the computer work out the gradient for us.

GRADIENT COMPUTATION: THE SYMBOLIC APPROACH

WolframAlpha computational intelligence.

derivative of $l(w) = (d + z * 1/(1+e^{-(c + v*(b + 1/(1+e^{-(x^w))}))}))$

Extended Keyboard Upload Examples Random

Derivative: Approximate form Step-by-step solution

$$\frac{\partial}{\partial w} \left(l(w) = d + \frac{z}{1 + e^{-(c + v(b + 1/(1 + e^{-(x^w))}))}} \right) = \frac{\partial \text{Hold} \left[\frac{z}{1 + e^{-(c + v(b + 1/(1 + e^{-(x^w))}))}} + d \right]}{\partial w}$$

Alternate form:

$$l'(w) = \frac{v x z e^{w x}}{(e^{w x} + 1)^2 (e^{v(b + 1/(e^{-w x} + 1)) + c} + 1)} - \frac{v x z e^{w x}}{(e^{w x} + 1)^2 (e^{v(b + 1/(e^{-w x} + 1)) + c} + 1)^2}$$

7

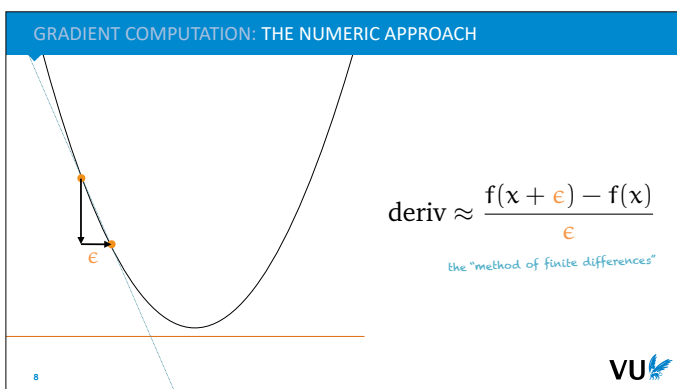
Of course, working out derivatives is a pretty mechanical process. We could easily take all the rules we know, and put them into some algorithm. This is called **symbolic differentiation**, and it's what systems like Mathematica and Wolfram Alpha do for us.

Unfortunately, as you can see here, the derivatives it returns get pretty horrendous the deeper the neural network gets. This approach becomes impractical very quickly.

As the depth of the network grows, the symbolic expression of its gradient (usually) grows exponentially in size.

Note that in symbolic differentiation we get a description of the derivative that is **independent of**

the input. We get a function that we can then feed any input to.



Another approach is to compute the gradient *numerically*. For instance by the method of finite differences: we take a small step ϵ and, see how much the function changes. The amount of change divided by the step size is a good estimate for the gradient if ϵ is small enough.

Numeric approaches are sometimes used in deep learning, but it's very expensive to make them accurate enough if you have a large number of parameters.

Note that in the numeric approach, you only get an answer **for a particular input**. If you want to compute the gradient at some other point in space, you have to compute another numeric approximation. Compare this to the symbolic approach (either with pen and paper or through wolfram alpha) where once the differentiation is done, all we have to compute is the derivative that we've worked out.

BACKPROPAGATION: THE MIDDLE GROUND

Work out parts of the derivative **symbolically**
chain these together in a **numeric computation**.

secret ingredient: **the chain rule**.

VU


Backpropagation is a kind of middle ground between symbolic and numeric approaches to working out the gradient. We break the computation into parts: we work out the derivatives of the parts *symbolically*, and then chain these together *numerically*.

The secret ingredient that allows us to make this work is the **chain rule** of differentiation.

THE CHAIN RULE

$$\frac{\partial f(g(x))}{\partial x} = \frac{\partial f(g(x))}{\partial g(x)} \frac{\partial g(x)}{\partial x}$$

shorthand:



$$\frac{\partial f}{\partial x} = \frac{\partial f}{\partial g} \frac{\partial g}{\partial x}$$

10

VU

Here is the chain rule: if we want the derivative of a function which is the composition of two other functions, in this case f and g , we can take the derivative of f with respect to the output of g and multiply it by the derivative of g with respect to the input x .

Since we'll be using the chain rule *a lot*, we'll introduce a simple shorthand to make it a little easier to parse. We draw a little diagram of which function feeds into which. This means we know what the argument of each function is, so we can remove the arguments from our notation.

We call this diagram a **computation graph**. We'll stick with simple diagrams like this for now. At the end of the lecture, we will expand our notation a little bit to

INTUITION FOR THE CHAIN RULE

slope of f over g

$$f(g) = s_f g + b_f$$

$$g(x) = s_g x + b_g$$

with $s_f = \frac{\partial f}{\partial g}$

$$s_g = \frac{\partial g}{\partial x}$$



$$f(g(x)) = s_f (s_g x + b_g) + b_f$$

$$= s_f s_g x + s_f b_g + b_f$$

slope of f over x
constant

$$\frac{\partial f}{\partial g} \frac{\partial g}{\partial x}$$

11

VU

Since the chain rule is the heart of backpropagation, and backpropagation is the heart of deep learning, we should probably take some time to see why the chain rule is true at all.


If we imagine that f and g are linear functions, it's pretty straightforward to show that this is true. They may not be, of course, but the nice thing about calculus is that locally, we can *treat* them as linear functions (if they are differentiable). In an infinitesimally small neighbourhood f and g are exactly linear.

If f and g are locally linear, we can describe their behavior with a slope s and an additive constant b . The slopes, s_f and s_g , are simply the derivatives. The additive constants we will show can be ignored.

In this linear view, what the chain rule says is this: if we approximate $f(x)$ as a linear function, then its slope is the slope of f as a function of g , times the slope of g as a function of x . To prove that this is true, we just write down $f(g(x))$ as linear functions, and multiply out the brackets.

Note that this doesn't quite count as a rigorous proof, but it's hopefully enough to give you some intuition for why the chain rule holds.

MULTIVARIATE CHAIN RULE



$$\frac{\partial f}{\partial x} = \frac{\partial f}{\partial g} \frac{\partial g}{\partial x} + \frac{\partial f}{\partial h} \frac{\partial h}{\partial x}$$

12

VU

Since we'll be looking at some pretty elaborate computation graphs, we'll need to be able to deal with this situation as well: we have a computation graph, as before, but f depends on x through *two different* operations. How do we take the derivative of f over x ?

The multivariate chain rule tells us that we can simply apply the chain rule along g , taking h as a constant, and sum it with the chain rule along h taking g as a constant.

INTUITION FOR THE MULTIVARIATE CHAIN RULE

$$f(g, h) = s_1 g + s_2 h + b_f$$

$$g(x) = s_g x + b_g \quad \text{with } s_1 = \frac{\partial f}{\partial g} \quad s_2 = \frac{\partial f}{\partial h} \quad s_g = \frac{\partial g}{\partial x} \quad s_h = \frac{\partial h}{\partial x}$$

$$h(x) = s_h x + b_h$$

$$f(g(x), h(x)) = s_1(s_g x + b_g) + s_2(s_h x + b_h) + b_f$$

$$= s_1 s_g x + s_1 b_g + s_2 s_h x + s_2 b_h + b_f$$


$$= \underbrace{(s_1 s_g + s_2 s_h)}_{\text{slope of } f \text{ over } x} x + \underbrace{s_1 b_g + s_2 b_h + b_f}_{\text{constant}}$$

$$\frac{\partial f}{\partial g} \frac{\partial g}{\partial x} + \frac{\partial f}{\partial h} \frac{\partial h}{\partial x}$$



We can see why this holds in the same way as before. The short story: since all functions can be taken to be linear, their slopes distribute out into a sum

MULTIVARIATE CHAIN RULE



$$\frac{\partial f}{\partial x} = \sum_i \frac{\partial f}{\partial g_i} \frac{\partial g_i}{\partial x}$$



If we have more than two paths from the input to the output, we simply sum over all of them.

EXAMPLE

$$f(x) = \frac{2}{\sin(e^{-x})}$$



With that, we are ready to show how backpropagation works. We'll start with a fairly arbitrary function to show the principle before we move on to more realistic neural networks.

ITERATING THE CHAIN RULE

$$f(x) = \frac{2}{\sin(e^{-x})} \quad f(x) = d(c(b(a(x))))$$

operations:


$$d(c) = \frac{2}{c}$$

$$c(b) = \sin b$$

$$b(a) = e^a$$

$$a(x) = -x$$

computation graph:




The first thing we do is to break up its functional form into a series of smaller operations. The entire function f is then just a chain of these small operations chained together. We can draw this in a computation graph as we did before.

Normally, we wouldn't break a function up in such small operations. This is just a simple example to illustrate the principle.

ITERATING THE CHAIN RULE

$$\begin{aligned}\frac{\partial f}{\partial x} &= \frac{\partial d}{\partial x} \\ &= \frac{\partial d}{\partial c} \frac{\partial c}{\partial x} \\ &= \frac{\partial d}{\partial c} \frac{\partial c}{\partial b} \frac{\partial b}{\partial x} \\ &= \frac{\partial d}{\partial c} \frac{\partial c}{\partial b} \frac{\partial b}{\partial a} \frac{\partial a}{\partial x}\end{aligned}$$



Now, to work out the derivative of f , we can *iterate the chain rule*. We apply it again and again, until the derivative of f over x is expressed as a long product of derivatives of operation outputs over their inputs.

17



GLOBAL AND LOCAL DERIVATIVES

$$\frac{\partial f}{\partial x} = \frac{\partial d}{\partial c} \frac{\partial c}{\partial b} \frac{\partial b}{\partial a} \frac{\partial a}{\partial x}$$

global derivative
local derivatives

We call the larger derivative of f over x the **global derivative**. And we call the individual factors, the derivatives of f over x is expressed as a long product of derivatives of operation outputs over their inputs, the **local derivatives**.

18



BACKPROPAGATION

The **BACKPROPAGATION** algorithm:

- break your computation up into a sequence of operations
what counts as an operation is up to you
- work out the **local derivatives** symbolically.
- compute the **global derivative** numerically
by computing the local derivatives and multiplying them

This is how the backpropagation algorithm combines symbolic and numeric computation. We work out the local derivatives symbolically, and then work out the global derivative numerically.

19



WORK OUT THE LOCAL DERIVATIVES SYMBOLICALLY

$$f(x) = \frac{2}{\sin(e^{-x})}$$

operations:

$$d(c) = \frac{2}{c}$$

$$c(b) = \sin b$$

$$b(a) = e^a$$

$$a(x) = -x$$

$$\begin{aligned}\frac{\partial f}{\partial x} &= \frac{\partial d}{\partial c} \frac{\partial c}{\partial b} \frac{\partial b}{\partial a} \frac{\partial a}{\partial x} \\ &= -\frac{2}{c^2} \cdot \cos b \cdot e^a \cdot -1\end{aligned}$$

For each local derivative, we work out the symbolic derivative with pen and paper.

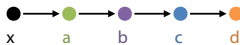
Note that we could fill in the a , b and c in the result, but we don't. We simply leave them as is. For the symbolic part, we are only interested in the derivative of the output of each sub-operation with respect to its immediate input.

The rest of the algorithm is performed numerically.

20



COMPUTE A FORWARD PASS (X = -4.499)

$$f(-4.499) = 2$$


$$d = \frac{2}{c} = 2$$

$$c = \sin b = 1$$

$$b = e^a = 90$$

$$a = -x = 4.499$$

21



This we are now computing things numerically, we need a specific input, in this case $x = -4.499$. We start by feeding this through the computation graph. For each sub-operation, we store the output value. At the end, we get the output of the function f . This is called a **forward pass**: a fancy term for computing the output of f for a given input.

Note that at this point, we are no longer computing solutions in general. We are computing our function for a specific input. We will be computing the gradient for this specific input as well.

COMPUTE A FORWARD PASS (X = -4.499)

$$f(-4.499) = 2 \quad \frac{\partial f}{\partial x} = -\frac{2}{c^2} \cdot \cos b \cdot e^a \cdot -1$$

$$d = \frac{2}{c} = 2 \quad = -\frac{2}{1^2} \cdot \cos 90 \cdot e^{4.499} \cdot -1$$

$$c = \sin b = 1 \quad = -2 \cdot 0 \cdot 90 \cdot -1 = 0$$

$$b = e^a = 90$$

$$a = -x = 4.499$$

22



Keeping all intermediate values from the forward pass in memory, we go back to our symbolic expression of the derivative. Here, we fill in the intermediate values a , b and c . After we do this, we can finish the multiplication numerically, giving us a numeric value of the gradient of f at $x = -4.499$. In this case, the gradient happens to be 0.

BACKPROPAGATION

The **BACKPROPAGATION** algorithm:

- break your computation up into a sequence of operations
what counts as an operation is up to you
- work out the **local derivatives** symbolically.
- compute the **global derivative** numerically
by computing the local derivatives and multiplying them
- Much more accurate than finite differences
only source of inaccuracy is the numeric computation of the operations.
- Much faster than symbolic differentiation
The backward pass has (broadly) the same complexity as the forward.


23



Before we try this on a neural network, here are the main things to remember about the backpropagation algorithm.

Note that backpropagation by itself does not train a neural net. It just provides a gradient. When people say that they trained a network by backpropagation, that's actually shorthand for training the network by gradient descent, with the gradients worked out by backpropagation.

BACKPROPAGATION IN A FEEDFORWARD NETWORK



24



To explain how backpropagation works in a neural network, we extend our neural network diagram a little bit, to make it closer to the actual computation graph we'll be using.

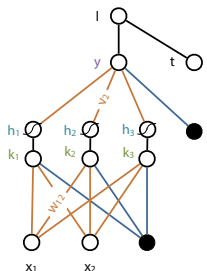
First, we separate the hidden node into the result of the linear operation k_i and the application of the nonlinearity h_i . Second, since we're interested in the derivative of the loss rather than the output of the network, we extend the network with one more step: **the computation of the loss** (over one example to keep things simple). In this final step, the output y of the network is compared to the target value t from the data, producing a loss value.

The loss is the function for which we want to work out the gradient, so the computation graph is the one that

computes first the model output, and then the loss based on this output (and the target).

Note that the model is now just a *subgraph* of the computation graph. You can think of t as another input node, like x_1 and x_2 , (but one to which the *model* doesn't have access).

BACKPROPAGATION IN A FEEDFORWARD NETWORK



$$\frac{\partial l}{\partial v_2}, \frac{\partial l}{\partial w_{12}}$$

$$l = (y - t)^2$$

$$y = v_1 h_1 + v_2 h_2 + v_3 h_3 + b_h$$

$$h_2 = \frac{1}{1 + \exp(-k_2)}$$

$$k_2 = w_{12} x_1 + w_{22} x_2 + b_x$$

25 VU

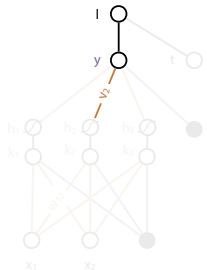
We want to work out the gradient of the loss. This is simply the collection of the derivative of the loss over each parameter.

We'll pick two parameters, v_2 in the second layer, and w_{12} in the first, and see how backpropagation operates. The rest of the parameters can be worked out in the same way to give us the rest of the gradient.

First, we have to break the computation of the loss into operations. If we take the graph on the left to be our computation graph, then we end up with the operations of the right.

To simplify things, we'll compute the loss over only one instance. We'll use a simple squared error loss.

BACKPROPAGATION IN A FEEDFORWARD NETWORK



$$l = (y - t)^2$$

$$y = v_1 h_1 + v_2 h_2 + v_3 h_3 + b_h$$

$$\frac{\partial l}{\partial v_2} = \frac{\partial l}{\partial y} \frac{\partial y}{\partial v_2}$$

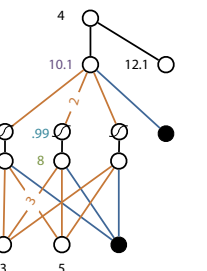
$$= 2(y - t) \cdot h_2$$

26 VU

For the derivative with respect to v_2 , we'll only need these two operations. Anything below doesn't affect the result.

To work out the derivative we apply the chain rule, and work out the local derivatives symbolically.

BACKPROPAGATION IN A FEEDFORWARD NETWORK



$$\frac{\partial l}{\partial v_2} = \frac{\partial l}{\partial y} \frac{\partial y}{\partial v_2}$$

$$= 2(y - t) \cdot h_2$$

$$= 2(10.1 - 12.1) \cdot .99 = -3.96$$

$$v_2 \leftarrow v_2 - \alpha \cdot -3.96$$


27 VU

We then do a forward pass with some values. We get an output of 10.1, which should have been 12.1, so our loss is 4. We keep all intermediate values in memory.

We then take our product of local derivatives, fill in the numeric values from the forward pass, and compute the derivative over v_2 .

When we apply this derivative in a gradient descent update, v_2 changes as shown below.

BACKPROPAGATION IN A FEEDFORWARD NETWORK



$$l = (y - t)^2$$

$$y = v_1 h_1 + v_2 h_2 + v_3 h_3 + b_h$$

$$h_2 = \frac{1}{1 + \exp(-k_2)}$$

$$k_2 = w_{12} x_1 + w_{22} x_2 + b_x$$


$$\frac{\partial l}{\partial w_{12}} = \frac{\partial l}{\partial y} \frac{\partial y}{\partial h_2} \frac{\partial h_2}{\partial k_2} \frac{\partial k_2}{\partial w_{12}}$$

$$= 2(y - t) \cdot v_2 \cdot h_2(1 - h_2) \cdot x_1$$

derivative of sigmoid: $\sigma'(x) = \sigma(x)(1 - \sigma(x))$

Let's try something a bit earlier in the network: the weight w_{12} . We add two operations, apply the chain rule and work out the local derivatives.

BACKPROPAGATION



$$\frac{\partial l}{\partial y}$$

$$\frac{\partial l}{\partial y} \frac{\partial y}{\partial h_2} = \frac{\partial l}{\partial h_2}$$

$$\frac{\partial l}{\partial y} \frac{\partial y}{\partial h_2} \frac{\partial h_2}{\partial k_2} = \frac{\partial l}{\partial k_2}$$


$$\frac{\partial l}{\partial y} \frac{\partial y}{\partial h_2} \frac{\partial h_2}{\partial k_2} \frac{\partial k_2}{\partial w_{12}} = \frac{\partial l}{\partial w_{12}}$$

Note that when we're computing the derivative for w_{12} , we are also, along the way computing the derivatives for y , h_2 and k_2 .

This is useful when it comes to implementing backpropagation. We can walk backward down the computation graph and compute the derivative of the loss for every node. For the nodes below, we just multiply the local gradient. This means we can very efficiently compute any derivatives we need.

In fact, this is *where the name backpropagation comes from*: the derivative of the loss propagates down the network in the opposite direction to the forward pass. We will show this more precisely in the last part of this lecture.

WALKING BACK DOWN THE NETWORK



- Break your function up into modules ◆
Call the last node the *loss*.
- Do a forward pass.
- Start at the loss node and work your way down.
- At each **input** to a **module**, multiplying
 - deriv. of loss over output ← already computed
 - deriv. of output over **input** ← local deriv.
 produces the deriv. of loss over **input**.



Here is a more abstract view of what is happening in backpropagation, which should apply to any kind of computation. We can think of the computations we do as modules with **inputs** and **outputs**, chained together to make a computation *graph*. The **output** of each module contributes ultimately to the result of the computation, which is the *loss*.

We want to know the derivative of the loss with respect to each of our **input** nodes. By the chain rule, this is the derivative of the loss with respect to the **output** times the derivative of the **output** with respect to the **input**.

If we take care to walk back down our computation graph in the correct order, then we can be sure that for every module we encounter, we will have already computed the first factor. We only need to compute the second, and multiply by the value we already have.

We'll extend this abstract view of backpropagation in the last part of the lecture.

$$v_2 \leftarrow v_2 - \alpha \cdot 2(y - t)h_2$$

$$v_1 \leftarrow v_1 - \alpha \cdot 20 \cdot 0.1$$

$$v_2 \leftarrow v_2 - \alpha \cdot 20 \cdot 0.5$$

$$v_3 \leftarrow v_3 - \alpha \cdot 20 \cdot 0.9$$

VU

Of course, we cannot change y directly. Instead, we have to change the values that influenced y .

Here we see what that looks like for the **weights** of the second layer. First note that the output y in this example was **too high**. Since all the hidden nodes have positive values (because of the sigmoid), we end up subtracting some positive value from all the weights. This will lower the output, as expected.

Second, note that the change is proportional to the input. The first hidden node h_1 only contributes a factor of 0.1 (times its weight) to the value of y , so it isn't changed as much as h_3 , which contributes much more to the erroneous value.

Note also that the current value of the weight doesn't factor into the update: whether v_1 is 1, 10 or 100, we subtract the same amount. Only how much influence the weight had on the value of y in the forward pass is taken into account. The higher the activation of the **source node**, the more the weight gets adjusted.

Finally, note how the sign of the the derivative wrt to y is taken into account. Here, the model output was too high, so the more a weight contributed to the output, the more it gets "punished" by being lowered. If the output had been too low, the opposite would be true, and we would be adding something to the value of each weight.

NEGATIVE ACTIVATION

$$v_1 \leftarrow v_1 - \alpha \cdot 20 \cdot 0.1$$

$$v_2 \leftarrow v_2 + \alpha \cdot 20 \cdot 0.5$$

$$v_3 \leftarrow v_3 - \alpha \cdot 20 \cdot 0.9$$

VU

The sigmoid activation we've used so far allows only positive values to emerge from the hidden layer. If we switch to an activation that also allows negative activations (like a linear activation or a **tanh** activation), we see that **backpropagation very naturally takes the sign into account**.

Note the negative activation on h_2 .

In this case, we want to update in such a way that y decreases, but we note that the weight v_2 is multiplied by a *negative* value. This means that (for this instance) v_2 contributes *negatively* to the model output, and its value should be increased if we want to decrease the output.

Note that the sign of v_2 itself doesn't matter. Whether

$$h_2 \leftarrow h_2 - \alpha \cdot 2(y - t)v_2$$

$$= h_2 + \alpha \cdot 20 \cdot 3$$

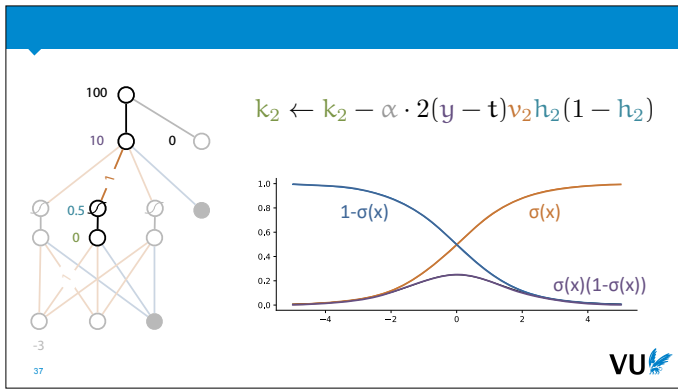
VU

We use the same principle to work our way back down the network. If we could change the output of the second node h_2 directly, this is how we'd do it.

Note that we now take the value of v_2 to be a constant. We are working out *partial derivatives*, so when we are focusing on one parameters, all the others are taken as constant.

Remember, that we want to decrease the output of the network. Since v_2 makes a *negative* contribution to the loss, we can achieve this by *increasing* the activation of the source node v_2 is multiplied by.

Note also that we've switched back to sigmoid activations.

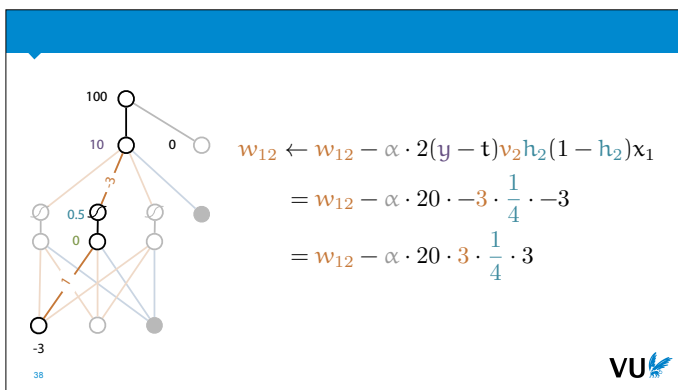


Moving down to k_2 , remember that the derivative of the sigmoid is the output of the sigmoid times 1 minus that output.

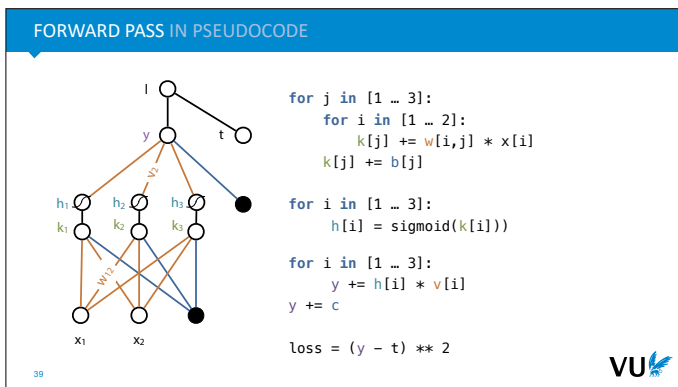
We see here, that in the extreme regimes, the sigmoid is *resistant to change*. The closer to 1 or 0 we get the smaller the weight update becomes.

This is actually a great *downside* of the sigmoid activation, and one of the big reasons it was eventually replaced by the ReLU as the default choice for hidden units. We'll come back to this in a later lecture.

Nevertheless, this update rule tells us what the change is to k_2 that we *want to achieve* by changing the gradients we can actually change (the *weights* of layer 1).

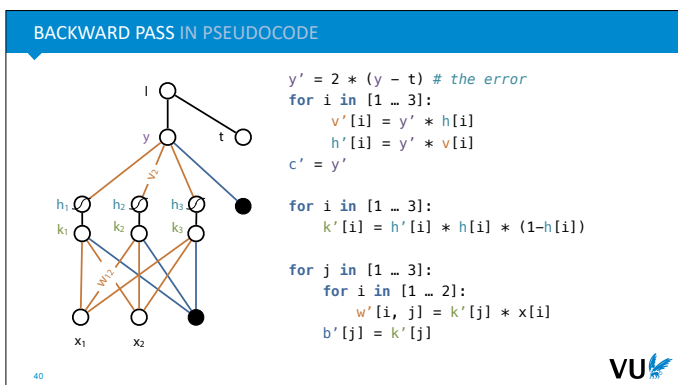


Finally, we come to the weights of the first layer. As before, we want the output of the network to *decrease*. To achieve this, we want h_2 to *increase* (because v_2 is negative). However, the input x_1 is negative, so we should decrease w_{12} to increase h_2 . This is all beautifully captured by the chain rule: the two negatives of x_1 and v_2 cancel out and we get a positive value which we subtract from w_{12} .



To finish up let's look at how you would implement this in code. Here is the forward pass: computing the model output and the loss, given the inputs and the target value.

Assume that k and y are initialized with 0s or random values. We'll talk about initialization strategies in the 4th lecture.



And here is the backward pass. We compute gradients for every node in the network, regardless of whether the node represents a parameter. When we do the gradient descent update, we'll use the gradients of the parameters, and ignore the rest.

Note that we don't implement the derivations from slide 44 directly. Instead, we work backwards down the neural network: computing the derivative of each node as we go, by taking the derivative of the loss over the outputs and multiplying it by the local derivative.

RECAP

Backpropagation: method for computing derivatives.

Combined with **gradient descent** to train NNs.

Middle ground between symbolic and numeric computation.

- Break computation up into operations.
- Work out **local derivatives** symbolically.
- Work out **global derivative** numerically.

41



Lecture 2: Backpropagation

Peter Bloem
Deep Learning

dlvu.github.io



PART THREE: TENSOR BACKPROPAGATION

Expressing backpropagation in vector, matrix and tensor operations.



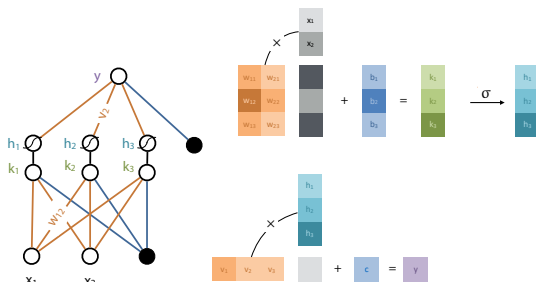
|section| Tensor backpropagation |

|video| <https://www.youtube.com/embed/O-xs8IyP4bQ?si=YWT0e5PGkU37kOVf> |

We've seen what neural networks are, how to train them by gradient descent and how to compute that gradient by backpropagation.

In order to scale up to larger and more complex structures, we need to make our computations as efficient as possible, and we'll also need to simplify our notation. There's one insight that we are going to get a lot of mileage out of.

IT'S ALL JUST LINEAR ALGEBRA



44



When we look at the computation of a neural network, we can see that most operations can be expressed very naturally in those of linear algebra.

The multiplication of weights by their inputs is a multiplication of a **matrix of weights** by a vector of inputs.

Note that the weight w_{12} (as we've called it so far, because it goes from node 1 to node 2) actually goes into element w_{21} of the matrix W . You can tell the difference by whether we're using a lower-case or capital w .

The addition of a bias is the addition of a **vector of bias parameters**.

The nonlinearities can be implemented as simple


element-wise operations.

This perspective buys us two things. First...

$$f(\mathbf{x}) = \mathbf{V} \sigma(\mathbf{W}\mathbf{x} + \mathbf{b}) + \mathbf{c}$$

Our notation becomes extremely simple. We can describe the whole operation of a neural network with one simple equation. This expressiveness will be sorely needed when we move to more complicated networks.

MATRIX MULTIPLICATION



```
# k, w, k, h, etc are zero'd arrays
for j in [1 ... 2]:
  for i in [1 ... 3]:
    k[i] += w[i,j] * x[i]
  k[i] += b[i]

for i in [1 ... 3]:
  h[i] = sigmoid(k[i])

for i in [1 ... 3]:
  y += h[i] * v[i]
y += c

loss = (y - c) ** 2
```



The second reason is that the biggest computational bottleneck in a neural network is the multiplication of the layer input by the layer weights. The matrix multiplication in the notation of the previous slide. This operation is more than quadratic while everything else is linear. We can see that in our pseudocode, because we have one loop, nested in another.

Matrix multiplication (and other tensor operations like it) can be parallelized and implemented efficiently but it's a lot of work. Ideally, we'd like to let somebody else do all that work (like the implementers of numpy) and then just call their code to do the matrix multiplications.

This is especially important if you have access to a GPU. A matrix multiplication can easily be offloaded to the GPU for much faster processing, but for a loop over an array, there's no benefit.

This is called **vectorizing**: taking a piece of code written in for loops, and getting rid of the loops by expressing the function as a sequence of linear algebra operations.

VECTORIZING

Express everything as operations on vectors, matrices and tensors.

Get rid of all the loops

Makes the notation *simpler*.


Makes the execution *faster*.

47



Without vectorized implementations, deep learning would be painfully slow, and GPUs would be useless. Turning a naive, loop-based implementation into a vectorized one is a key skill for DL researchers and programmers.

FORWARD



$$l = \sum (y - t)^2$$

$$k = w \cdot \text{dot}(x) + b$$

$$y = Vh + c$$

$$h = \text{sigmoid}(k)$$

$$h = \sigma(k)$$

$$y = v \cdot \text{dot}(h) + c$$

$$k = Wx + b$$

$$l = (y - t) ** 2$$

48




Here's what the vectorized forward pass looks like as a computation graph, in symbolic notation and in pseudocode.

When you implement this in numpy it'll look almost the same.

Note that this doesn't just represent the network in the previous part, it represents any such network, regardless of the sizes of the input and hidden layers. We've abstracted away some details about the specifics of the network. We've also made the output and the target label vectors for the sake of generality

So far so good. The forward pass is easy enough to vectorize.

BUT WHAT ABOUT THE BACKWARD PASS?



$$l = \sum (y - t)^2$$

Can we do something like this?

$$y = Vh + c$$

$$\frac{\partial l}{\partial W} \stackrel{?}{=} \frac{\partial l}{\partial y} \frac{\partial y}{\partial h} \frac{\partial h}{\partial k} \frac{\partial k}{\partial W}$$

$$h = \sigma(k)$$

$$k = Wx + b$$

49



Of course, we lose a lot of the benefit of vectorizing if the backward pass (the backpropagation algorithm) isn't also expressed in terms of matrix multiplications. **How do we vectorize the backward pass?** That's the question we'll answer in the rest of this video.

On the left, we see the forward pass of our loss computation as a set of operations on vectors and matrices.

To generalize backpropagation to this view, we might ask if something similar to the chain rule exists for vectors and matrices. Firstly, can we define something analogous to the derivative of one matrix over another, and secondly, can we break this apart into a product of local derivatives, possibly giving us a sequence of matrix multiplications?

The answer is yes, there are many ways of applying calculus to vectors and matrices, and there are many chain rules available in these domains. However, things can quickly get a little hairy, so we need to tread carefully.

GRADIENTS, JACOBIANS, ETC

$f(\mathbf{A}) = \mathbf{B}$, $\frac{\partial \mathbf{B}}{\partial \mathbf{A}}$: derivatives of every element of \mathbf{A} over every element of \mathbf{B}

for instance:

$$f \left(\begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \right) = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$$

$$J_f = \begin{pmatrix} \frac{\partial b_1}{\partial a_1} & \frac{\partial b_2}{\partial a_1} \\ \frac{\partial b_1}{\partial a_2} & \frac{\partial b_2}{\partial a_2} \\ \frac{\partial b_1}{\partial a_3} & \frac{\partial b_2}{\partial a_3} \end{pmatrix}$$

		function returns a		
		scalar	vector	matrix
arrange derivatives in a:				
scalar		scalar	vector	matrix
vector	input is a	vector	matrix	?
matrix		matrix	?	?

50

The derivatives of high-dimensional objects are easily defined. We simply take the derivative of every number in the input over every number in the output, and we arrange all possibilities into some logical shape. For instance, if we have a vector-to-vector function, the natural analog of the derivative is a matrix with all the partial derivatives in it.

However, once we get to matrix/vector operations or matrix/matrix operations, the only way to logically arrange every input with every output is a tensor of higher dimension than 2.

NB: We don't normally apply the differential symbol to non-scalars like this. We'll introduce better notation later.

$$\frac{\partial l}{\partial \mathbf{W}} \stackrel{?}{=} \frac{\partial l}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \mathbf{h}} \frac{\partial \mathbf{h}}{\partial \mathbf{k}} \frac{\partial \mathbf{k}}{\partial \mathbf{W}}$$

uh oh

51



As we see, even if we could come up with this kind of chain rule, one of the local derivatives is now a vector over a matrix. The result could only be represented in a 3-tensor. There are two problems with this:

- If the operation has n inputs and n outputs, we are computing n^3 derivatives, even though we were only ultimately interested in n^2 of them (the derivatives of \mathbf{W}). In the scalar algorithm we only ever had two nested loops (an n^2 complexity), and we only ever stored one gradient for one node in the computation graph. Now we suddenly have n^3 complexity and n^3 memory use. We were supposed to be making things faster.
- We can easily *represent* a 3-tensor, but there's no obvious, default way to multiply a 3-tensor with a matrix, or with a vector (in fact there are many different ways). The multiplication of the chain rule becomes very complicated this way.

In short, we need a more careful approach.

SIMPLIFYING ASSUMPTIONS

1) The computation graph *always* has a single, **scalar** output: l

2) We are only ever interested in the derivative of l .


$$\frac{\partial l}{\partial \mathbf{W}} \leftarrow \begin{matrix} \text{scalar} \\ \text{tensor of} \\ \text{any dimension} \end{matrix}$$

		scalar	vector	matrix
scalar		scalar	vector	matrix
vector		vector	matrix	?
matrix		matrix	?	?
3-tensor		3-tensor	?	?

52



To work out how to do this we make these following simplifying assumptions.



Note that this doesn't mean we can only ever train neural networks with a single scalar output. Our neural networks can have any number of outputs of any shape and size. We can make neural networks that generate images, natural language, raw audio, even video.

However, the *loss* we define over those outputs needs to map them all to a single scalar value. The computation graph is always the model, plus the computation of the loss.

THE GRADIENT

We will call $\frac{\partial l}{\partial \mathbf{W}}$ the gradient of l (with respect to \mathbf{W})

Commonly written as $\nabla_{\mathbf{W}} l$

Nonstandard assumption: $\nabla_{\mathbf{W}} l$ has the same *shape* as \mathbf{W}

$$[\nabla_{\mathbf{W}} l]_{ijk} = \frac{\partial l}{\partial W_{ijk}}$$



We call this derivative **the gradient of \mathbf{W}** . This is a common term, but we will deviate from the standard approach in one respect.

Normally, the gradient is defined as a row vector, so that it can operate on a space of column vectors by matrix multiplication.

In our case, we are never interested in multiplying the gradient by anything. We only ever want to *sum* the gradient of l wrt \mathbf{W} with the original matrix \mathbf{W} in a gradient update step. For this reason **we define the gradient as having the same shape** as the tensor \mathbf{W} with respect to which we are taking the derivative.

In the example shown, \mathbf{W} is a 3-tensor (a kind of matrix with 3 dimensions instead of 2). The gradient of l wrt \mathbf{W} has the same shape as \mathbf{W} , and at element (i, j, k) it holds the scalar derivative of l wrt W_{ijk} .

With these rules, we can use tensors of any shape and dimension and always have a welldefined gradient.

NEW NOTATION: THE GRADIENT FOR \mathbf{W}

$$\nabla_{\mathbf{W}} l \leftarrow \text{always the same} = \mathbf{W}^{\nabla}$$

← actual object of interest

$$\mathbf{b}^{\nabla} = \nabla_{\mathbf{b}} l = \left[\frac{\partial l}{\partial b_1} \dots \frac{\partial l}{\partial b_n} \right]^T$$

$$y^{\nabla} = \nabla_y l = \frac{\partial l}{\partial y}$$



The standard gradient notation isn't very suitable for our purposes. It puts the loss front and center, but that will be the same in all cases. The object that we're actually interested in is relegated to a subscript. Also, it isn't very clear from the notation what the shape is of the tensor that we're looking at.


For this reason, we'll introduce a new notation. This isn't standard, so don't expect to see it anywhere else, but it will help to clarify our mathematics a lot as we go forward.

We've put the \mathbf{W} front and center, so it's clear that the result of taking the gradient is also a matrix, and we've removed the loss, since we've assume that we are always taking the gradient of the loss.

You can think of the nabla as an operator like a transposition or taking an inverse. It turns a matrix into another matrix, a vector into another vector and so on.

The notation works the same for vectors and even for scalars. This is the gradient of l with respect to \mathbf{W} . Since l never changes, we'll refer to this as "the gradient for \mathbf{W} ".

$$[\mathbf{W}^\nabla]_{ij} = [W_{ij}]^\nabla = \frac{\partial l}{\partial W_{ij}}$$

$$= W_{ij}^\nabla$$


When we refer to a single element of the gradient, we will follow our convention for matrices, and use the non-bold version of its letter.

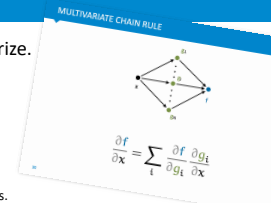
Element (i, j) for the gradient of \mathbf{W} is the same as the gradient for element (i, j) of \mathbf{W} .

To denote this element we follow the convention we have for elements of tensors, that we use the same letter as we use for the matrix, but non-bold.


TENSOR BACKPROPAGATION

Work out *scalar* derivatives first, *then* vectorize.
Use the **multivariate chain rule** to derive the scalar derivative.

Apply the chain rule step by step.
Start at the loss and work backward, *accumulating* the gradients.



MULTIVARIATE CHAIN RULE

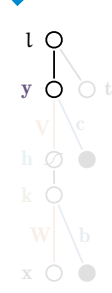
$$\frac{\partial f}{\partial \mathbf{x}} = \sum_i \frac{\partial f}{\partial g_i} \frac{\partial g_i}{\partial \mathbf{x}}$$


This gives us a good way of thinking about the gradients, but we still don't have a chain rule to base out backpropagation on.

The main trick we will use is to stick to **scalar derivatives** as much as possible.

Once we have worked out the derivative in purely scalar terms (on pen and paper), we will then find a way to vectorize their computation.


GRADIENT FOR \mathbf{y}



$$y_i^\nabla = \frac{\partial l}{\partial y_i}$$

$$= \frac{\partial \sum_k (y_k - t_k)^2}{\partial y_i} = \sum_k \frac{\partial (y_k - t_k)^2}{\partial y_i}$$

$$= \sum_k 2(y_k - t_k) \frac{\partial y_k - t_k}{\partial y_i} = 2(y_i - t_i)$$

$$\mathbf{y}^\nabla = \begin{pmatrix} y_1^\nabla \\ \vdots \\ y_N^\nabla \end{pmatrix} = \begin{pmatrix} 2(y_1 - t_1) \\ \vdots \\ 2(y_N - t_N) \end{pmatrix} = 2(\mathbf{y} - \mathbf{t})$$


We start simple: what is the gradient for \mathbf{y} ? This is a vector, because \mathbf{y} is a vector. Let's first work out the derivative of the i -th element of \mathbf{y} . This is purely a scalar derivative so we can simply use the rules we already know. We get $2(y_i - t_i)$ for that particular derivative.

Then, we just re-arrange all the derivatives for y_i into a vector, which gives us the complete gradient for \mathbf{y} .

The final step requires a little creativity: we need to figure out how to compute this vector using only basic linear algebra operations on the given vectors and matrices. In this case it's not so complicated: we get the gradient for \mathbf{y} by element-wise subtracting \mathbf{t} from \mathbf{y} and multiplying by 2.

We haven't needed any chain rule yet, because our computation graph for this part has only one edge.

GRADIENT FOR \mathbf{V}

$$\begin{aligned}
 v_{ij}^{\nabla} &= \frac{\partial l}{\partial v_{ij}} \\
 &= \sum_k \frac{\partial l}{\partial y_k} \frac{\partial y_k}{\partial v_{ij}} = \sum_k y_k^{\nabla} \frac{\partial y_k}{\partial v_{ij}} \\
 &= \sum_k y_k^{\nabla} \frac{\partial [V_{kj}h_j + c_k]}{\partial v_{ij}} = \sum_k y_k^{\nabla} \frac{\partial V_{kj}h_j + c_k}{\partial v_{ij}} \\
 &= \sum_k y_k^{\nabla} \frac{\partial \sum_l V_{kl}h_l + c_k}{\partial v_{ij}} \\
 &= \sum_{kl} y_k^{\nabla} \frac{\partial V_{kl}h_l}{\partial v_{ij}} = y_i^{\nabla} \frac{\partial V_{ij}h_j}{\partial v_{ij}} \\
 &= y_i^{\nabla} h_j
 \end{aligned}$$

$$\mathbf{v}^{\nabla} = \begin{pmatrix} \vdots \\ \dots & y_i^{\nabla} h_j & \dots \\ \vdots \end{pmatrix} = \mathbf{y}^{\nabla} \mathbf{h}^T$$

Let's move one step down and work out the gradient for \mathbf{V} . We start with the scalar derivative for an arbitrary element of \mathbf{V} : v_{ij} .

Now, we need a chain rule, to backpropagate through \mathbf{y} . However, since we are sticking to scalar derivatives, we can simply use the **scalar multivariate chain rule**. This tells us that however many intermediate values we have, we can work out the derivative for each, keeping the others constant, and sum the results.

This gives us the sum in the second equality. We've worked out the gradient \mathbf{y}^{∇} already, so we can fill that in and focus on the derivative of y_k over v_{ij} . the value of y_k is defined in terms of linear algebra operations like matrix multiplication, but with a little thinking we can always rewrite these as a simple scalar sums.

In the end we find that the derivative of y_k over v_{ij} reduces to the value of h_j .

This tells us the values of every element (i, j) of \mathbf{v}^{∇} . All that's left is to figure out how to compute this in a vectorized way. In this case, we can compute \mathbf{v}^{∇} as the outer product of the gradient for \mathbf{y} , which we've computed already, and the vector \mathbf{h} , which we can save during the forward pass.

RECIPE

To work out a gradient \mathbf{X}^{∇} for some \mathbf{X} :

- Write down a **scalar derivative** of the loss wrt *one element* of \mathbf{X} .
- Use the **multivariate chain** rule to sum over all outputs.
- Work out the scalar derivative.
- **Vectorize** the computation of \mathbf{X}^{∇} in terms of the original inputs.

VU

GRADIENT FOR h

$$\begin{aligned}
 h_i^\nabla &= \frac{\partial l}{\partial h_i} \\
 &= \sum_k \frac{\partial l}{\partial y_k} \frac{\partial y_k}{\partial h_i} = \sum_k y_k^\nabla \frac{\partial y_k}{\partial h_i} \\
 &= \sum_k y_k^\nabla \frac{\partial (\mathbf{Vh} + c_k)}{\partial h_i} = \sum_k y_k^\nabla \frac{\partial \sum_i V_{ki} h_i + c_k}{\partial h_i} \\
 &= \sum_{ki} y_k^\nabla \frac{\partial V_{ki} h_i + c_k}{\partial h_i} = \sum_k y_k^\nabla \frac{\partial V_{ki} h_i}{\partial h_i} \\
 &= \sum_k y_k^\nabla V_{ki} \\
 \mathbf{h}^\nabla &= \begin{pmatrix} \vdots \\ \vdots \\ \vdots \end{pmatrix} = (\mathbf{y}^\nabla \mathbf{V})^T = \mathbf{V}^T \mathbf{y}^\nabla
 \end{aligned}$$

Since this is an important principle to grasp, let's keep going until we get to the other set of parameters, **W**. We'll leave the biases as an exercise.

For the gradient for h, most of the derivation is the same as before, until we get to the point where the scalar derivative is reduced to a matrix multiplication. Unlike the previous derivation, the sum over k *doesn't disappear*. We need to take it into the description of the gradient for h.

To vectorize this, we note that each element of this vector is a dot product of **y** and a column of **V**. This means that if we premultiply **y** (transposed) by **V**, we get the required result as a row vector. Transposing the result (to make it a column vector) is equivalent to post multiplying **y** by the transpose of **V**.

Note the symmetry of the forward and the backward operation.

GRADIENT FOR k

$$\begin{aligned}
 k_i^\nabla &= \frac{\partial l}{\partial k_i} \\
 &= \sum_m h_m^\nabla \frac{\partial h_m}{\partial k_i} = \sum_m h_m^\nabla \frac{\partial \sigma(k_i)}{\partial k_i} \\
 &= h_i^\nabla \frac{\partial \sigma(k_i)}{\partial k_i} \\
 &= h_i^\nabla \sigma'(k_i) = h_i^\nabla \sigma(k_i)(1 - \sigma(k_i)) \\
 &= h_i^\nabla h_i(1 - h_i)
 \end{aligned}$$

derivative of sigmoid
 $\sigma'(x) = \sigma(x)(1 - \sigma(x))$

$$\mathbf{k}^\nabla = \mathbf{h}^\nabla \odot \mathbf{h} \odot (1 - \mathbf{h})$$

Now, at this point, when we analyze **k**, remember that we already have the gradient over **h**. This means that we no longer have to apply the chain rule to anything above **h**. We can draw this scalar computation graph, and work out the local gradient for **k** in terms for the gradient for **h**.


Given that, working out the gradient for **k** is relatively easy, since the operation from **k** to **h** is an element-wise one.

GRADIENT FOR W

$$\begin{aligned}
 W_{ij}^\nabla &= \frac{\partial l}{\partial W_{ij}} \\
 &= \sum_m \frac{\partial l}{\partial k_m} \frac{\partial k_m}{\partial W_{ij}} = \sum_m k_m^\nabla \frac{\partial k_m}{\partial W_{ij}} \\
 &= \sum_m k_m^\nabla \frac{\partial (\mathbf{Wx} + b)_m}{\partial W_{ij}} = \sum_m k_m^\nabla \frac{\partial \sum_j W_{mj} x_j + b_m}{\partial W_{ij}} \\
 &= \sum_m k_m^\nabla \frac{\partial \sum_j W_{mj} x_j + b_m}{\partial W_{ij}} \\
 &= \sum_{mj} k_m^\nabla \frac{\partial W_{mj} x_j}{\partial W_{ij}} = k_i^\nabla \frac{\partial W_{ij} x_j}{\partial W_{ij}} \\
 &= k_i^\nabla x_j \\
 \mathbf{W}^\nabla &= \begin{pmatrix} \vdots & \vdots & \vdots \\ \dots & k_i^\nabla x_j & \dots \\ \vdots & \vdots & \vdots \end{pmatrix} = \mathbf{k}^\nabla \mathbf{x}^T
 \end{aligned}$$

Finally, the gradient for **W**. The situation here is exactly the same as we saw earlier for **V** (matrix in, vector out, matrix multiplication), so we should expect the derivation to have the same form (and indeed it does).

PSEUDOCODE: BACKWARD



64

VU

Here's the backward pass that we just derived.

We've left the derivatives of the **bias parameters** out. You'll have to work these out to implement the first homework exercise.

RECAP

Vectorizing makes notation simpler, and computation faster.

Vectorizing the **forward pass** is usually easy.

Backward pass: work out the scalar derivatives, accumulate, then vectorize.

65

VU

Lecture 2: Backpropagation

Peter Bloem
Deep Learning

dlvu.github.io

VU
Vrije Universiteit
AMSTERDAM

PART FOUR: AUTOMATIC DIFFERENTIATION

Letting the computer do all the work.

VU

|section|Automatic differentiation|

|video|<https://www.youtube.com/embed/9H-o8wESCxI?si=NyQ6djR2xavcBsIL>|

We've simplified the computation of derivatives a lot. All we've had to do is work out local derivatives and chain them together. We can go one step further: if we let the computer keep track of our computation graph, and provide some backwards functions for basic operations, the computer can work out the whole backpropagation algorithm for us. This is called **automatic differentiation** (or sometimes **autograd**).

THE STORY SO FAR

$v' = y'.mm(h.T)$
 $h' = (y'[None, :] * v).sum(axis=1)$
 $k' = sigmoid(k) * sigmoid(1 - k)$
 $w' = k'.mm(x.T)$

pen and paper in the computer

VU

THE IDEAL

$k = w.dot(x) + b$
 $h = sigmoid(k)$
 $y = v.dot(h) + c$
 $l = (y - t) ** 2$
`l.backward() # start backprop`

pen and paper in the computer

VU

This is what we want to achieve. We work out on pen and paper the local derivatives of various *modules*, and then we chain these modules together in a computation graph *in our code*. The computer keeps the graph in memory, and can automatically work out the backpropagation.

THE KEY IDEA

$$W_{ij}^v = \frac{\partial l}{\partial W_{ij}}$$

$$= \sum_m \frac{\partial l}{\partial k_m} \frac{\partial k_m}{\partial W_{ij}} = \sum_m k_m^v \frac{\partial k_m}{\partial W_{ij}}$$

$$= \sum_m k_m^v \frac{\partial (W_{ij}x_j + b_m)}{\partial W_{ij}} = \sum_m k_m^v \frac{\partial (W_{ij}x_j + b_m)}{\partial W_{ij}}$$

$$= \sum_m k_m^v \frac{\partial W_{ij}x_j}{\partial W_{ij}} = k_i^v x_j$$

$$W^v = \begin{pmatrix} \dots & \dots & \dots \\ \dots & k_i^v x_j & \dots \\ \dots & \dots & \dots \end{pmatrix} \cdot k^v.T$$

Ignored once we have k^v

pen and paper in the computer

VU

Whenever we work out a gradient, as we did in the previous video, we always look at the node above the one for which we're computing the gradient, and use the multivariate chain rule to split the gradient in two. In this case, we are working out the gradient for **W** and we apply the multivariate chain to break that gradient into the gradient for **k** and the derivatives of **k** with respect to **W**.

The key idea, that powers automatic differentiation is that **once we have k^v , we no longer care about anything else that happens above **k** in our computation graph**. All we need is k^v and we can work out **W** nably.

Define your computation graph in terms of operations chained together.


For each operation $x \rightarrow y$, if you know the gradient y^v for the output, you can work out the gradient x^v for the input.

It doesn't matter what happens in the rest of the computation graph, before or after the operation.

With this knowledge, **you can start at the top of the graph, and work your way down**, computing all the gradients.

VU

AUTOMATIC DIFFERENTIATION



72

VU

This kind of algorithm is called automatic differentiation. What we've been doing so far is called **backward, or reverse mode automatic differentiation**. This is efficient if you have few output nodes. Note that if we had two output nodes, we'd need to do a separate backward pass for each.

If you have few inputs, it's more efficient to start with the inputs and apply the chain rule working forward. This is called forward mode automatic differentiation.

Since we assume we only have one output node (where the gradient computation is concerned), we will always use reverse mode in deep learning.

THE PLAN

Tensors

Scalars, vectors, matrices, etc.

Building computation graphs

Define operations with tensor inputs and outputs, chain them together.

Working out backward functions

For each operation, work out how to compute the gradient for the input given the gradient for the output


73

VU

To make this idea precise, and applicable to as broad a range of functions as possible, we'll need to set up a few ingredients. Our values will be tensors: scalars, vectors, matrices, or their higher-dimensional analogues.

We then build computation graphs out of operations

TENSORS (AKA MULTIDIMENSIONAL ARRAYS)



74

VU

The basic datastructure of our system will be the tensor. A tensor is a generic name for family of datastructures that includes a scalar, a vector, a matrix and so on.

There is no good way to visualize a 4-dimensional structure. We will occasionally use this form to indicate that there is a fourth dimension along which we can also index the tensor.

We will assume that whatever data we work with (images, text, sounds), will in some way be encoded into one or more tensors, before it is fed into our system. Let's look at some examples.

A CLASSIFICATION TASK AS TWO TENSORS

word count	nr. of recipients	class
30	3	ham
340	4	ham
121	2	spam
11	1	spam
23	1	spam
455	1	spam
512	2	spam
2	12	ham
32	1	ham
432	1	ham
23	2	spam

$$X = \begin{bmatrix} 30 & 3 \\ 340 & 4 \\ 121 & 2 \\ 11 & 1 \\ 23 & 1 \\ 455 & 1 \\ 512 & 2 \\ 2 & 12 \\ 32 & 1 \\ 432 & 1 \\ 23 & 2 \end{bmatrix} \quad y = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$


75

VU

A simple dataset, with numeric features can simply be represented as a matrix. For the labels, we usually create a separate corresponding vector for the labels.

Any categorical features or labels should be converted to numeric features (normally by one-hot coding).

AN IMAGE AS A 3-TENSOR



76


image source: <http://www.adell.com/scanning101.html>



Images can be represented as 3-tensors. In an RGB image, the color of a single pixel is represented using three values between 0 and 1 (how red it is, how green it is and how blue it is). This means that an RGB image can be thought of as a stack of three color channels, represented by matrices.

This stack forms a 3-tensor.

A DATASET OF IMAGES



```
In [5]: from keras.datasets import cifar10
(x_train, y_train), (x_test, y_test) = cifar10.load_data()
x_train.shape
Out[5]: (50000, 32, 32, 3)
```

77



If we have a dataset of images, we can represent this as a 4-tensor, with dimensions indexing the instances, their width, their height and their color channels respectively. Below is a snippet of code showing that when you load the CIFAR10 image data in Keras, you do indeed get a 4-tensor like this.

There is no agreed standard ordering for the dimensions in a batch of images. Tensorflow and Keras use (batch, height, width, channels), whereas Pytorch uses (batch, channels, height, width).

(You can remember the latter with the mnemonic “bachelor chow”.)

Tensors


Building computation graphs

Working out backward functions

78



COMPUTATION GRAPHS



bipartite: only op-to-tensor or tensor-to-op edges
tensor nodes: no input or one input, multiple outputs
operation nodes: multiple inputs and outputs

79

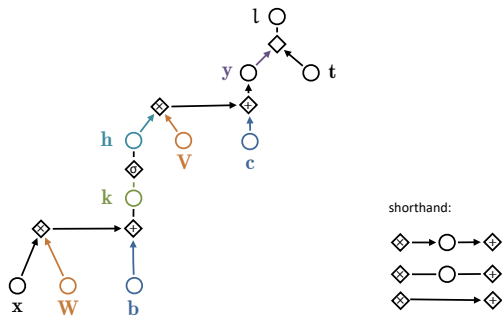


We'll need to be a little more precise in our notations. From now on we'll draw computation graphs with both the operations and the values as nodes. The graph is always bipartite and directed. A tensor node is connected to the ops for which it is an input, and an operation node is connected to the tensor nodes that it produces as outputs.

There doesn't seem to be a standard visual notation. This is what we'll use in most of this course.

In Tensorflow operations are called ops, and in Pytorch they're called functions.

FEEDFORWARD NETWORK



VU

As an example, here is our MLP expressed in our new notation.

Just as before, it's up to us how *granular* we make the computation. For instance, we wrap the whole computation of the loss in a single operation, but we could also separate out the subtraction and the squaring.

You may note that this graph is not directed or bipartite everywhere. This is just visual shorthand. Occasionally, we will omit the direction of an edge for clarity, or omit intermediate nodes when it's clear that a tensor node or op node should be there. The actual computation graph we are describing in such cases is always bipartite and entirely directed.

RESTATING OUR ASSUMPTIONS


One output node l with a scalar value.


All gradients are of l , with respect to the tensors on the tensor nodes.
We compute gradients for all tensor nodes.

VU

We hold on to the same assumptions we had before. They are required for automatic differentiation to work efficiently.

IMPLEMENTATIONS

TensorNode: 
value: <Tensor>
gradient: <Tensor>
source: <OpNode>

OpNode: 
inputs: List<TensorNode>
outputs: List<TensorNode>
op: <Op>

VU

To store a computation graph in memory, we need three classes of objects. The first two are shown here.

A **TensorNode** object holds the tensor value at that node. It holds the gradient of the tensor at that node (to be filled by the backpropagation algorithm) and it holds a pointer to the Operation Node that produced it (which can be null for leaf nodes).

An **OpNode** object represents an *instance* of a particular operation being applied in the computation. It stores the particular op being used (multiplication, additions, etc) and the inputs and the outputs. In the MLP example, we perform several matrix multiplications: each of these becomes a separate OpNode in the computation graph, all recording an instance of the matrix multiplication operation being used. Each of them refers to a single object representing this operation.

This is the third class: the Op.

DEFINING AN OPERATION

```

class Op:
    def forward(context, inputs):
        # given the inputs, compute the outputs
        ...

    def backward(context, outputs_gradient):
        # given the gradient of the loss wrt to the outputs
        # compute the gradient of the loss wrt to the inputs
        ...

```

stores anything we need for the backward pass

$f(A) = B$
 $\text{forward}_f : A \mapsto B$
 $\text{backward}_f : B^\nabla \mapsto A^\nabla$

83



If the difference between an Op and an OpNode is unclear consider that a single neural network may apply, say, a sigmoid operation many times. The Op object for the sigmoid defines how to compute the sigmoid operations (and its gradient). Then for each place in the computation graph where the sigmoid is applied, we create a separate OpNode object which references the inputs and outputs specific to that application of the sigmoid, together with a single pointer to the single sigmoid Op. This is where we actually define how to perform the computation.

An Op is defined by two functions.

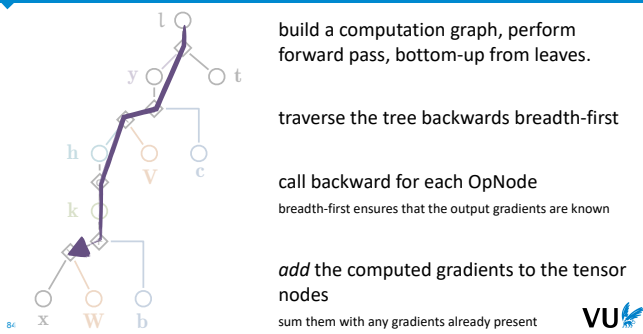
The function `forward` computes the outputs given the inputs (just like any function in any programming language).

The function `backward` takes the gradient for the outputs (the gradient of the loss wrt to the outputs) and produces the gradient for the inputs.

Both functions are also given a `context` object. This is a data structure (a dictionary or a list) to which the forward can add any value which it needs to save for the backward, we'll see an example of this later.

Note that the backward function does not compute the local derivative as we did in the scalar backpropagation: it computes the accumulated gradient of the loss over the inputs (given the accumulated gradient of the loss over the outputs).

BACKPROPAGATION



84



We'll see how to implement ops later. First, let's see how all this works put together. Using Ops, OpNodes and TensorNodes, we can build a computation graph (again, more on how we do that later). The OpNodes reference the TensorNodes that are their inputs and outputs and they reference the Ops that contain the actual code for their forward and backward.

Once we have this computation graph, computing our model output and loss and performing backpropagation becomes almost trivially easy. We simply traverse the tree forwards from the inputs to the loss calling all the forward functions, and then backwards from the loss to the inputs, calling all the backwards functions.

The only thing we need to keep track of is that when we call `forward()` on an OpNode's Op, all the ancestors of that OpNode have been called, so that we have all its inputs. And, when we traverse the tree backward, when we call `backward()` on an OpNode's Op that all the descendants have been called so that we have the gradients for all of its outputs.

Note the last point: some TensorNodes will be inputs to multiple operation nodes. By the multivariate chain rule, we should sum the gradients they get from all the OpNodes they feed into. We can achieve this easily by just initializing the gradients to zero and adding the gradients we compute to any that are already there.

BUILDING COMPUTATION GRAPHS IN CODE

Lazy execution: build your graph, compile it, feed data through.

Eager execution: perform forward pass, keep track of computation graph.

85



All this assumes that the computation graph is already there. But how do we tell the system what our computation graph should be? It turns out that the nicest way to do this is to describe the computation itself in code. This makes sense: code is a language explicitly for describing computations, so it would be nice to simply describe a computation as a piece of Python code, and have the system build a computation graph for us.

This can be done through the magic of *operator overloading*. We simply change what operators like `*` and `+` mean, when they are applied to `TensorNode` objects. There are two strategies common: lazy and eager execution.


EXAMPLE: LAZY EXECUTION

```
a = TensorNode()
b = TensorNode()
```

```
c = a * b
```

```
m = Model(
    in=(a, b),
    loss=c)
```

```
m.train((1, 2))
```



86



Here's one example of how a lazy execution API might look.

Note that when we're building the graph, we're not telling it which values to put at each node. We're just defining the *shape* of the computation, but not performing the computation itself.

When we create the model, we define which nodes are the input nodes, and which node is the loss node. We then provide the input values and perform the forward and backward passes.

BUILDING THE COMPUTATION GRAPH: LAZY EXECUTION

Tensorflow 1.x default, Keras default

Define the computation graph.

- Compile it.
- Iterate backward/forward over the data

Fast. Many possibilities for optimization. Easy to serialise models. Easy to make training parallel.

Difficult to debug. Model must remain static during training.

87



In lazy execution, which was the norm during the early years of deep learning, we build the computation graph, but we don't yet specify the values on the tensor nodes. When the computation graph is finished, we define the data, and we feed it through.


This is fast since the deep learning system can optimize the graph structure during compilation, but it makes models hard to debug: if something goes wrong during training, you probably made a mistake while defining your graph, but you will only get an error while passing data through it. The resulting stack trace will never point to the part of your code where you made the mistake.

EXAMPLE: EAGER EXECUTION

```
a = TensorNode(value=1)
b = TensorNode(value=2)
```

```
c = a * b
```

```
c.backward()
```



88



In eager mode deep learning systems, we create a node in our computation graph (a `TensorNode`) by specifying what data it should contain. The result is a tensor object that stores both the data, and the gradient over that data (which will be filled later).

Here we create the variables `a` and `b`. If we now apply an operation to these, for instance to multiply their values, the result is another variable `c`. Languages like python allow us to *overload* the `*` operator it looks like we're just computing multiplication, but behind the scenes, we are creating a computation graph that records all the computations we've done.

We compute the data stored in `c` by running the computation immediately, but we also store references to the variables that were used to create `c`,

and the operation that created it. **We create the computation graph on the fly, as we compute the forward pass.**

Using this graph, we can perform the backpropagation from a given node that we designate as **the loss node** (node **c** in this case). We work our way down the graph computing the derivative of each variable with respect to **c**. At the start the TensorNodes do not have their grad's filled in, but at the end of the backward, all gradients have been computed.

Once the gradients have been computed, and the gradient descent step has been performed, we clear the computation graph. **It's rebuilt from scratch for every forward pass.**

BUILDING THE COMPUTATION GRAPH: EAGER EXECUTION

PyTorch, Tensorflow 2.0 default, Keras option

- Build the computation graph on the fly during the forward pass.

Easy to debug, problems in the model occur as the module is executing. Flexible: the model can be entirely different from one forward to the next.

More difficult to optimize. A little more difficult to serialize.



In eager execution, we simply execute all operations immediately on the data, and collect the computation graph on the fly. We then execute the backward and ditch the graph we've collected.

This makes debugging much easier, and allows for more flexibility in what we can do in the forward pass. It can, however be a little difficult to wrap your head around. Since we'll be using pytorch later in the course, we'll show you how eager execution works step by step.

Tensors

Building computation graphs


Working out backward functions



WORKING OUT THE BACKWARD FUNCTION

```
class Plus(Op):
    def forward(context, a, b):
        # a, b are matrices of the same size
        return a + b

    def backward(context, goutput):
        return goutput, goutput
```



The final ingredient we need is a large collection of operations with backward functions worked out. We'll show how to do this for a few examples.

First, an operation that sums two matrices element-wise.

The recipe is the same as we saw in the last part:

- 1) Write out the scalar derivative for a single element.
- 2) Use the multivariate chain rule to sum over all outputs.
- 3) Vectorize the result.

Note that when we draw the computation graph, we can think of everything that happens between **S** and **l**

as a single module: we are already given the gradient of l over S , so it doesn't matter if it's one operation or a million.

Again, we can draw the computation graph at any granularity we like: very fine individual operations like summing and multiplying or very coarse-grained operations like entire NNs.

For this operation, the context object is not needed, we can perform the backward pass without remembering anything about the forward pass.

BACKWARD: A FEW MORE EXAMPLES

Sigmoid

Row-wise sum

Expand

92



To finish up, we'll show you the implementation of some more operations. You'll be asked to do a few of these in the second homework exercise.

SIGMOID

```
class Sigmoid(Op):
    def forward(context, x):
        # x is a tensor of any shape
        sigx = 1 / (1 + (-x).exp())
        context['sigx'] = sigx
        return sigx

    def backward(context, goutput):
        sigx = context['sigx']
        return goutput * sigx * (1 - sigx)
```

$$\begin{aligned}
 X_{ijk}^{\nabla} &= \sum_{abc} Y_{abc}^{\nabla} \frac{\partial Y_{abc}}{\partial X_{ijk}} \\
 &= \sum_{abc} Y_{abc}^{\nabla} \frac{\partial \sigma(X_{abc})}{\partial X_{ijk}} \\
 &= Y_{ijk}^{\nabla} \frac{\partial \sigma(X_{ijk})}{\partial X_{ijk}} \\
 &= Y_{ijk}^{\nabla} \sigma(X_{ijk}) (1 - \sigma(X_{ijk})) \\
 &= Y_{ijk}^{\nabla} Y_{ijk} (1 - Y_{ijk}) \\
 X^{\nabla} &= Y^{\nabla} \otimes Y \otimes (1 - Y)
 \end{aligned}$$

93

This is a pretty simple derivation, but it shows two things:

- 1) We can easily do a backward over functions that output high dimensional tensors, but we should sum over all dimensions of the output when we apply the multivariate chain rule.
- 2) The backward function illustrates the utility of the **context object**. To work out the backward, we need to know the original input. We could just have stored that for every operation, but as we see here, we've done part of the computation already in the forward (and in other backwards, like the sum we saw earlier, the inputs *aren't* necessary at all). If we give the operation control over what information it wants to store for the backward pass, we are

ROW-WISE SUM

```
class RowSum(Op):
    def forward(context, x):
        # x is a matrix
        sumd = x.sum(axis=1)
        context['m'] = x.shape[1]
        return sumd

    def backward(context, gy):
        n, m = gy.shape[0], context['m']
        return gy[:, None].expand(n, m)
```

$$\begin{aligned}
 X_{ij}^{\nabla} &= \sum_k y_k^{\nabla} \frac{\partial y_k}{\partial X_{ij}} = \sum_k y_k^{\nabla} \frac{\partial \sum_l X_{kl}}{\partial X_{ij}} \\
 &= \sum_{kl} y_k^{\nabla} \frac{\partial X_{kl}}{\partial X_{ij}} = y_i^{\nabla} \frac{\partial X_{ij}}{\partial X_{ij}} \\
 &= y_i^{\nabla} \\
 X^{\nabla} &= y^{\nabla} \mathbf{1}^T
 \end{aligned}$$

94



Note that the output is a vector, because we've summed out one dimension. Therefore, when we apply the multivariate chain rule, we sum over only one dimension.


This is a sum like the earlier example, so we don't need to save any tensor values from the forward pass. However, we do need to remember the size of the dimension we summed out. Therefore, we use the context to store this information. The None keyword in the slice in the last line adds a singleton dimension: it turns the shape from (n) into $(n, 1)$.

The `expand` function we use here is not available in `numpy`, but it does exist in `pytorch`. In `numpy` we can use `repeat`.

EXPAND

```
class Expand(Op):
    def forward(context, x, size):
        # x is a scalar
        return np.full(x, size=size)

    def backward(context, gy):
        return gy.sum(), None
```



95



Our final example includes a constant argument. We take a scalar, and expand it to a matrix of the given size, filled with the value x .

In this case, we do not care about the gradient for size (and even if we did, integer values don't yield meaningful gradients without some extra work). We consider this a constant.

This can be done in different ways. In pytorch, you return `None` for the gradient for a constant (as we've shown here).

In our toy system (`vugrad`), we consider all *keyword arguments* constants.

Tensors

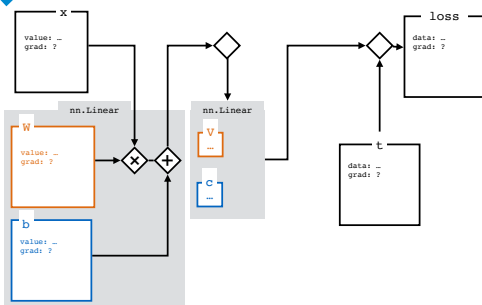
Building computation graphs

Working out backward functions

96



MODULES OR LAYERS



97

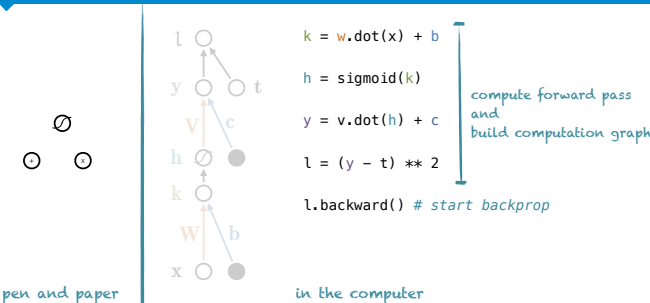


Most deep learning frameworks also have a way of combining model parameters and computation into a single unit, often called a **module** or a **layer**.

In this case a Linear module (as it is called in Pytorch) takes care of implementing the computation of a single layer of a neural network (sans activation) and of remembering the weights and the bias.

Modules have a **forward** function which performs the computation of the layer, but they *don't* have a **backward** function. They just chain together operations, and the backpropagation calls the backwards on the operations. In other words, it doesn't matter to the backpropagation whether we use a module or apply all the operations by hand.

THE IDEAL



pen and paper
work out forward and
backward functions

in the computer

98



Which completes the picture we wanted to create: a system where all we have to do is perform some computations on some tensors, just like we would do in numpy, and all the building of computation graphs and all the backpropagating is handled for us automatically.

So long as the forward and backward have been defined for all the operations we want to compute, either by us, or as part of a standard library, the system will do the rest.

LECTURE RECAP

backpropagation: work out **local gradients** symbolically, **global gradients** numerically. Work from back to front, applying the chain rule.

tensor backpropagation: assume scalar loss l , always compute the gradient of the loss.

automatic differentiation: Build computation graph on the fly. Backpropagation as a *breadth-first* walk down from the loss.

99



So

THANK YOU FOR YOUR ATTENTION

dlvu@peterbloem.nl

100



DELETED SCENES: TENSOR DETAILS



|section-nv| Tensor details* |


These slides were cut to lighten the amount of material covered in this lecture. They're not part of the exam material, but the information is still useful if you're working with DL systems. It's good to skim them quickly, so you know to come back here when you encounter any of the issues we describe here.

TENSORS IN MEMORY: ROW MAJOR ORDERING

A =

5	3	1
2	6	0

row-major ordering



102



It's important to realize that even though we think of tensors as multidimensional arrays, in memory, they are necessarily laid out as a single line of numbers. The Tensor object knows the shape that these numbers should take, and uses that to compute whatever we need it to compute.

We can do this in two ways: scanning along the rows first and then the columns is called **row-major ordering**. This is what numpy and pytorch both do. The other option is (unsurprisingly) called column major ordering. In higher dimensional tensors, the dimensions further to the right are always scanned before the dimensions to their left.

Imagine looping over all elements in a tensor with shape (a, b, c, d) in four nested loops. If you want to

loop over the elements in the order they are in memory, then the loop over d would be the innermost loop, then c, then b and then a.

This allows us to perform some operations very cheaply, **but we have to be careful**.

RESHAPE/VIEW

A =

5	3	1
2	6	0

A' =

5	3
1	2
6	0

Tensor A

data: [5, 3, 1, 2, 6, 0]

shape: (2, 3)

Tensor A'

data: [1, 2, 6, 0]

shape: (3, 2)

VU

Here is one such cheap operation: a reshape. This changes the shape of the tensor, but not the data.

To do this cheaply, we can create a new tensor object with the same data as the old one, but with a different shape.

RESHAPE/VIEW

A =

5	3	1
2	6	0

a' =

5	3	1	2	6	0
---	---	---	---	---	---

Tensor A

data: [5, 3, 1, 2, 6, 0]

shape: (3, 2)

Tensor a'

data: [5, 3, 1, 2, 6, 0]

shape: (6,)

VU

EXAMPLE: NORMALIZE COLORS

```

images = load_images(...)
n, h, w, c = images.size()

images = images.reshape(n*h*w, c)
images = images / images.max(dim=0)
images = images.reshape(n, h, w, c)

# this makes no sense, but it is allowed.
images = images.reshape(n*h, w*c)
images = images.reshape(h, n, c, w)

```

VU

Imagine that we wanted to scale a dataset of images in such way that over the whole dataset, each color channel has a maximal value of 1 (independent of the other channels).

To do this, we need a 3-vector of the maxima per color channel. Pytorch only allows us to take the maximum in one direction, but we can reshape the array so that all the directions we're not interested in are collapsed into one dimension.

Afterwards, we can reverse the reshape to get our image dataset back.

We have to be careful: the last three lines in this slide form a perfectly valid reshape, but the **c dimension** in the result does *not* contain our color values.

In general, you're fine if you **collapse dimensions that are next to each other**, and uncollapse them *in the same order*.

TRANSPOSE

A =

5	3	1
2	6	0

A^T =

5	2
3	6
1	0

Tensor A

data: [5, 3, 1, 2, 6, 0]

shape: (2, 3)

Tensor A^T

parent: ←

shape: reverse(parent.shape)

```
def get(i, j):
    return parent.get(j, i)
```

106

A transpose operation can also be achieved cheaply. In this case, we just keep a reference to the old tensor, and whenever a user requests the element (i, j) we return (j, i) from the original tensor.

NB: This isn't precisely how pytorch does this, but the effect is the same: we get a transposed tensor in constant time, by viewing the same data in memory in a different way.

SLICE

A =

5	3	1
2	6	0

A[:, :2] =

5	3
2	6

Tensor A

data: [5, 3, 1, 2, 6, 0]

shape: (3, 2)

Tensor A'

parent: ←

shape: (2, 2)

```
def get(i, j):
    return parent(i, j)
```

107

Even slicing can be accomplished by referring to the original data.

NON-CONTIGUOUS

Contiguous tensor: the data are directly laid out in row-major ordering for the shape with no gaps or shuffling required.

```
x = torch.randn(2, 3)
x = x[:, 1:]
x.view(6)
Traceback (most recent call last):
...
RuntimeError: view size is not compatible with input tensor's size and stride (at least one dimension spans across two contiguous subspaces). Use .reshape(...) instead.
```

x.contiguous(): make contiguous (by copying the data).
x.view(): reshape if possible without making contiguous.
x.reshape(): reshape, call contiguous() if necessary.

VU

108

For some operations, however, the data needs to be contiguous. That is, the tensor data in memory needs to be one uninterrupted string of data in row major ordering with no gaps. If this isn't the case, pytorch will throw an exception like this.

You can fix this by calling `.contiguous()` on the tensor. The price you pay is the linear time and space complexity of copying the data.