

# Neural networks

Adam Boulton ([www.bou.lt](http://www.bou.lt))

February 10, 2023

# Contents

Preface	2
<b>I Supervised neural networks</b>	<b>3</b>
1 Multi-layer perceptrons with sigmoids and backpropagation	4
2 Regularising neural networks	8
3 Normalising neural networks	10
4 Addressing gradient problems with the Rectified Linear Unit (ReLU)	11
<b>II Supervised neural networks cont.</b>	<b>14</b>
5 Alternatives to backpropagation	15
6 Multi-class neural networks	16
7 Probabilistic neural networks	17
8 Pre-training neural networks	18
9 Convolutional layers for neural networks	19
10 Attention in neural networks	22
<b>III Generative neural networks</b>	<b>23</b>
11 Autoencoders and Variational Autoencoders (VAE)	24
12 Restricted Boltzmann Machines (RBMs)	25

<i>CONTENTS</i>	2
<b>13 Self-organising maps</b>	<b>26</b>
<b>14 Generative neural networks</b>	<b>27</b>
<b>IV Estimating time series models</b>	<b>28</b>
<b>15 Recurrent neural networks</b>	<b>29</b>
<b>16 Recurrent Neural Network (RNN) encoders and decoders</b>	<b>31</b>

# Preface

This is a live document, and is full of gaps, mistakes, typos etc.

## Part I

# Supervised neural networks

# Chapter 1

## Multi-layer perceptrons with sigmoids and backpropagation

### 1.1 Forward pass through a Multi-Layer Perceptron (MLP)

#### 1.1.1 Hidden layers

In the perceptron we have input vector  $x$ , and output:

$$a = a(wx)$$

We can augment the perceptron by adding a hidden layer.

Now the output on the activation function is an input to a second layer. By using different weights, we can create a second vector of inputs to the second layer.

#### 1.1.2 The parameters of a feed forward model

$\Theta^j$  is a matrix of weights for mapping layer  $j$  to  $j + 1$ . So we have  $\Theta^1$  and  $\Theta^2$ .

If we have  $s$  units in the hidden layer,  $n$  features and  $k$  classes:

- The dimension of  $\Theta^1$  is  $(n + 1) \times s$
- The dimension of  $\Theta^2$  is  $(s + 1) \times k$

These include the offsets for each layer.

### 1.1.3 The activation function of a multi-layer perceptron

For a perceptron we had  $a = f(wx)$ . Now we have:

$$a_i^j = f(a_{j-1}\Theta_{j-1})$$

We refer to the value of a node as  $a_i^j$ , the activation of unit  $i$  in layer  $j$ .

### 1.1.4 Dummies in neural networks

## 1.2 Back propagation

### 1.2.1 Adapting the delta rule

To arrive at the delta rule we considered the cost function:

$$E = \sum_j \frac{1}{2} (y_j - a_j)^2$$

And used the chain rule:

$$\frac{\delta E}{\delta \theta_i} = \frac{\delta E}{\delta a_j} \frac{\delta a_j}{\delta z_j} \frac{\delta z_j}{\delta \theta_i}$$

This gave us:

$$\Delta \theta_i = \alpha \sum_j (y_j - a_j) a'(z_j) x_{ij}$$

$$\text{Or, setting } \delta_i = -\frac{\delta E}{\delta z_j} = \sum_j (y_j - a_j) a'(z_j)$$

$$\Delta \theta_i = \alpha \delta_j x_{ij}$$

Let's update the rule for multiple layers:

$$\frac{\delta E}{\delta \theta_{li}} = \frac{\delta E}{\delta a_{lj}} \frac{\delta a_{lj}}{\delta z_{lj}} \frac{\delta z_{lj}}{\delta \theta_{li}}$$

Previously  $\frac{\delta z_{lj}}{\delta \theta_{li}} = x_i$ . We now use the more general  $a_{li}$ . For the first layer, these will be the same.

We can then instead write:

$$\Delta \theta_i = \alpha \delta_{lj} a_{li}$$

### 1.2.2 Calculating delta values

Now we need a way of calculating the value of  $\delta_{lj}$  for all neurons.

$$\delta_i = -\frac{\delta E}{\delta z_{lj}}$$

If this is an output node, then this is simply  $\sum_j (y_j - a_j) a'(z_j)$

If this is not an output node, then the impact of change in the parameter will affect the results through all intermediate neurons.

In this case:

$$\frac{\delta E}{\delta z_{lj}} = \sum_{k \in \text{succl}} \frac{\delta E}{\delta z_k} \frac{\delta z_k}{\delta z_{lj}}$$

$$\frac{\delta E}{\delta z_{lj}} = \sum_{k \in \text{succl}} -\delta_k \frac{\delta z_k}{\delta a_{kj}} \frac{\delta a_{kj}}{\delta z_{lj}}$$

$$\frac{\delta E}{\delta z_{lj}} = \sum_{k \in \text{succl}} -\delta_k \theta_{kj} a'_{kj}$$

$$\delta_i = a'_{kj} \sum_{k \in \text{succl}} \delta_k \theta_{kj}$$

For each layer there is a matrix, where the columns and rows represent the *theta* between the current layer and the next layer. We have a matrix for each layer in the network.

### 1.2.3 Initialising parameters

We start by randomly initialising the value of each  $\theta$ .

We do this to prevent each neuron from moving in sync.

## 1.3 Regressing on unbounded outputs with neural networks

### 1.3.1 Introduction

Can not have a sigmoid function at the last step.

Alternatively can apply a sigmoid function to the unbounded output to make it bounded.

## 1.4 Deep neural networks can represent complicated functions

### 1.4.1 More layers allow for more complex function

With additional hidden layers we can map more complex functions.

These allow the effective combination of logic gates.

### 1.4.2 2 hidden layers can map highly complex functions

With only two hidden layers we can map any function for classification, including discontinuous functions.



### 1.4.3 Increasing numbers of dimensions in a unit

Topology of layers. Increasing number of units in subsequent layers is like increasing dimension.

We are trying to make data linearly separable. it may be that we need additional dimensions to do this, rather than a series of transformations within the existing number of dimensions.

eg for a circle of data within a circle of data, there is no linear separable line, so no depth without increasing dimensions will split data.

## Chapter 2

# Regularising neural networks

### 2.1 Regularising neural networks

#### 2.1.1 Feature normalisation

#### 2.1.2 Dropout, and dropout layers

#### 2.1.3 $L_2$ regularisation (including how to change backprop algorithm)

#### 2.1.4 Sparse networks

Parameters are set to 0 and not trained.

#### 2.1.5 Parameter sharing

Parameters share the same value and are trained together.

#### 2.1.6 Weight decay

After each update, multiply the parameter by  $p < 1$ .

#### 2.1.7 The anomaly detection problem

Can change input to get any classification.

### 2.1.8 Early stopping

### 2.1.9 Residual blocks

In a node we have:

$$a_{ij} = \sigma_{ij}(W_{ij}a_{i-1})$$

That is, the value of a node, is the activation on the sum of the weights of the previous layer.

Residual block however look further back that one layer. They include the full data from an older layer (without weights)

$$a_{ij} = \sigma_{ij}(W_{ij}a_{i-1} + a_k)$$

## Chapter 3

# Normalising neural networks

### 3.1 Optimisation

#### 3.1.1 Input layer normalisation

We normalise the input layer.

This speeds up training, and makes regularisation saner.

#### 3.1.2 Batch normalisation

We can normalise other layers. We take each input, subtract the batch mean divided by the batch standard deviation.

#### **Batch normalisation and covariance shift**

Batch normalisation can make networks better adapted for related problems.

#### **Training with batch normalisation**

## Chapter 4

# Addressing gradient problems with the Rectified Linear Unit (ReLU)

### 4.1 Link/activation functions: Regression

#### 4.1.1 Absolute value rectification

$$a(x) = |x|$$

#### 4.1.2 Rectified Linear Unit (ReLU)

**The function**

$$a(z) = \max(0, z)$$

**The derivative**

Its differential is 1 for values of  $z$  above 0, and 0 for values of  $z$  below 0.

The differential is undefined at  $z = 0$ , however this is unlikely to occur in practice.

**Notes**

The ReLU activation function induces sparsity.

**4.1.3 Noisy ReLU****4.1.4 Leaky ReLU****4.1.5 Parametric ReLU****4.1.6 Softplus****The function**

$$a(z) = \ln(1 + e^z)$$

**The derivative**

Its derivative is the sigmoid function:

$$a'(z) = \frac{1}{1 + e^{-z}}$$

**Notes**

The softplus function is a smooth approximation of the ReLU function.

Unlike the ReLU function, Softplus does not induce sparsity.

**4.1.7 Exponential Linear Unit (ELU)****4.2 Link/activation functions: Regression****4.2.1 Convexity**

The error function for neural networks is nearly convex.

**4.2.2 Unstable gradient problem****Vanishing gradient problem**

Gradients can become small, and so propagation can be very slow.

**Exploding gradient problem**

Gradients can become too large and not converge.

**ReLU**

This addresses the unstable gradient problem.

**4.2.3 Curse of dimensionality**

As parameter space gets bigger, data requirements get bigger.

Sparse models can help. Eg ReLU.

#### 4.2.4 Representational sparsity

This is where the values in nodes are often 0, as opposed to just the parameters.

## Part II

# Supervised neural networks cont.



## Chapter 5

# Alternatives to backpropagation

### 5.1 Alternatives to backpropagation

#### 5.1.1 Greedy pretraining

#### 5.1.2 Cascade-correlation learning architecture

This is a method for both building and training.

We start with a bare bones network. We then add nodes one by one, training and then fixing their values.

#### 5.1.3 Extreme learning machines

This is an alternative to backpropagation for training a feedforward neural network.

We start with random parameters for each layer  $W_i$ .

We have:

$$\hat{y} = W_2 \sigma(W_1 x)$$

Etc.

We calculate:

$$W_2 = \sigma(W_1 x) + Y$$

So  $W_1$  is random and not updated.

$W_2$  is assigned to minimise loss, where  $W_2$  has no activation function.

## Chapter 6

# Multi-class neural networks

### 6.1 More than 2 classes

#### 6.1.1 Local Winner Takes All (LWTA) layer

The output for all nodes in a layer is 0 unless it is the greatest.

#### 6.1.2 Maxout layer

Single node, spits out the max of all inputs.

## Chapter 7

# Probabilistic neural networks

### 7.1 Probabilistic neural networks

#### 7.1.1 Introduction

#### 7.1.2 Input layer

The input is the feature vector.

#### 7.1.3 Pattern layer

The first layer has a node for each observation in the training set.

In each node, the value is the distance from the input to the comparator.

This can be calculated using Gaussian distribution, or another method.

#### 7.1.4 Summation layer

One neuron for each category.

We map from the pattern layer to the summation layer according to the actual label of each training item.

Ie, if a sample is red, it will be fed only to the red neuron.

The values are summed.

Largest value is selected.

## Chapter 8

# Pre-training neural networks

### 8.1 Pre-training

#### 8.1.1 Pre-training

Pre-training. eg train on general pictures before specific stuff. means you fit many parameters for detecting edges etc firsts

#### **Learning rate**

Reduce the learning rate.

#### 8.1.2 Resetting parameters

Replace last layer (softmax) for new problem.

#### 8.1.3 Freeze layers

Freeze feature learning of early layers.

#### 8.1.4 Catastrophic interference

When retraining model on new data, model may forget answers to old data.

## Chapter 9

# Convolutional layers for neural networks

### 9.1 Convolutional layers

#### 9.1.1 Convolutional layers

Can connect each node in first hidden layer to a subset of the input layer, eg one node for each 5x5 pixels

We also share weights for each of the first layer. Much fewer parameters, and can learn all good stuff

This also uses windows. Instead of max we multiply the window by a matrix elementwise and sum the values

Each matrix can represent some feature, like a curve.

We can use multiply convolution matrices to create multiple output matrices.

Matrices are called kernels. they are trained. start off random

#### Training convolutional layers

#### 9.1.2 Invariance of convolutional layers (rotation, translation)

#### 9.1.3 Flattening layers

We split the data up everytime we use convolutional layers

Flattening layers bring them all back together

Parameters are that for pooling layers (height, width, stride, padding, but also set of convolutions).

#### 9.1.4 Multi-scale convolutions

We use different window sizes in parallel.

#### 9.1.5 Inception modules

### 9.2 Pooling (max pooling, average pooling, sub-sampling)

#### 9.2.1 Pooling layers

The input is a matrix. We place a number of windows on the input matrix. The max of each window is an input to the next layer.

Means fewer parameters, easier to compute, less chance of overfitting

Parameters: height, width of window, stride (amount shifts by each window)

We can also add padding to the edge of the image so we don't lose data.

Same padding (use 0), valid padding (no padding)

Pooling layer compresses, takes 2x2. Max pooling returns highest activation

#### 9.2.2 Vector of Locally Aggregated Descriptors (VLAD)

### 9.3 Other window layers

#### 9.3.1 Capsules

##### Primary capsule layers

Outputs of convolutions are scalars. however we can also create vectors, if we associate some convolutions with each other

eg if we have 6 convolutions, the output of these can be used to create a 6 dimensional vector for each window.

##### Normalisation in primary capsule layers (vector squishing)

We can normalise the length of these vectors to between 0 and 1.

The output of this represents the chance of finding the feature they are looking for, and the orientation

If the vector length is low, feature not found. if high, feature found.

We have orientation from vector, and position from window

### **Routing capsule layers**

We now have a layer of position and orientation of basic shapes (triangles, rectangles etc)

We want to know which more complex thing they are part of.

So the output of this step is again a matrix with position and orientation, but of more complex features

To determine the activation from each basic shape to the next feature we use routing-by-agreement.

This takes each basic shape and works out what it would look like if the complex feature was present.

If a complex feature has two basic shapes, they will both have the same predicted complex shape. Otherwise the relationship is spurious and they will not

If they agree we have a high weight

This process is complex and computationally expensive.

However we don't need pooling layers now

### **caps net**

Does normal conv first, then primary, then secondary.

### **caps: reconstruction**

We have vector space of feature position and orientation. we can recreate output

## Chapter 10

# Attention in neural networks



## Part III

# Generative neural networks

## Chapter 11

# Autoencoders and Variational Autoencoders (VAE)

### 11.1 Autoencoders

#### 11.1.1 Autoencoders

Autoencoders are a type of neural network.

For autoencoders the goal for the output layer is the reconstructed input layer, rather than a classification.

By including sparsity in the neural network we can reduce the dimensions. This splits the network into an encoder and a decoder.

Middle vector is called latent variables.

### 11.2 Variational autoencoders

#### 11.2.1 Variational Autoencoders (VAE)

Like AE, but force middle vector to have unit gaussian by adding new loss function

Now we can generate new images by sampling for latent normal of unit 1.

## Chapter 12

# Restricted Boltzmann Machines (RBMs)

### 12.1 Restricted Boltzmann Machines (RBMs)

#### 12.1.1 Restricted Boltzmann Machines (RBMs)

#### 12.1.2 Deep Belief Networks (DBNs)

#### 12.1.3 Training with Contrastive Divergence (CD)

## Chapter 13

# Self-organising maps

### 13.1 Self-organising maps

# Chapter 14

## Generative neural networks

### 14.1 Generative adversarial networks

#### 14.1.1 Generative Adversarial Networks (GANs)

##### Introduction

##### The GAN generator

Generator: take latent multivar normal as input layer

Output of generator is of same dimension as input iamges

Generative NN: need probability distribution on input layer

Generator typically deconvolutional

Downside of gan. only discriminates between real and fake.

##### The GAN discriminator

Discriminator: use normal CNN for deep

## Part IV

# Estimating time series models

# Chapter 15

## Recurrent neural networks

### 15.1 Simple recurrent neural networks

#### 15.1.1 Simple Recurrent Neural Networks (RNNs)

##### Introduction

Recurrent Neural Networks (RNN) are an alternative to feedforward networks. These have loops.

##### Motivation

We have inputs which are not independent. For example speech input, where each input is a the recording for a length of time.

##### Unrolling RNNs

The activation unit takes the input, and an outcome from the previous activation unit. It then performs its activation function.

This allows information to be kept across time.

However this degrades, and relevant information was from much earlier, it will be lost.

#### 15.1.2 Backpropagation Through Time (BPTT)

We can do backpropagation on the unrolled network, backpropagating over time.

## 15.2 Long Short-Term Memory (LSTM)

### 15.2.1 Long Short-Term Memory (LSTM)

#### Introduction

These are a more complex RNN architecture.

#### Cell state

Each cell has as an input the cell state from the previous cell  $C_{t-1}$

The LSTM cell updates the cell state to  $C_t$  and pushes it to the next cell.

#### Other inputs to the cell

We have  $x_t$ , the input of the cell, and  $h_{t-1}$ , the output of the previous cell.

#### Cell output and the output gate

We run an activation function on the cell state  $C_t$  to get a candidate output.

We multiply this by the outcome of the output gate to get the actual result.

#### The input gate

We create a candidate change to the state.

We multiply this by the input gate value, and add it to the state.

#### The forget gate

This is a multiplication factor. What % of the state should be removed?

## 15.3 Variants

### 15.3.1 Peephole LSTM

### 15.3.2 Gated Recurrent Units (GRUs)

## 15.4 Forecasting with recurrent neural networks

### 15.4.1 Introduction

## 15.5 Other

### 15.5.1 Attention and Neural Turing Machines



## Chapter 16

# Recurrent Neural Network (RNN) encoders and decoders

### 16.1 Recurrent Neural Network (RNN) encoders and decoders

#### 16.1.1 Recurrent Neural Network (RNN) encoders

Final output is the encoding.

The end of the sequence is identified through an End-Of-Sequence token.

`seq2seq`

#### 16.1.2 Recurrent Neural Network (RNN) decoders

##### Introduction

We take the encoded vector and pass this through to the decoder. This spits out decoded output.

As we output a word, the word (and previous words) are sent as inputs to the following RNN cells.

##### Encoding the outputs

As we create outputs, we can pass this as an encoded vector in the target language.