

Module 11: NEURAL NETWORKS

TMA4268 Statistical Learning V2019

Mette Langaas

week 13 2019

Contents

Introduction	2
Learning material for this module	2
What will you learn?	2
Why a module on neural networks?	2
AI, machine learning and statistics	4
From statistics to artificial neural networks	4
Multiple linear regression	4
Logistic regression	8
Multiclass regression	14
Summing up	17
Feedforward networks	18
The single hidden layer feedforward network	18
Universal approximation property	20
Examples	22
Boston house prices	22
Using <code>nnet</code>	28
Neural network parts	28
Model	29
Method	32
Optimizers	33
Regularization	35
Metrics	38
Deep learning	38
Timeline	38
Deep?	38
The R keras package	39
MNIST dataset	40
Kaggle	44
Other analyses	44
Summing up	45
Not covered	45
Recommended exercises	45
Theoretical	45
Practical	48
Exam problems	52
V2018 Problem 4 Classification of diabetes cases	52

R packages	52
References and further reading	53
Acknowledgements	53

Last changes: (02.04: fixed scaling issue in cv-loop for finding number of hidden nodes, the scaling of the training and validation can only be made from the training part of the data, 28.03: material for L2, 24.03: first version)

Introduction

Learning material for this module

(on the reading list)

- Classnotes 25.03.2019
- Classnotes 28.03.2019

See also under References, for further reading material.

What will you learn?

- Why a module on neural networks?
 - Translating from statistical to neural networks language
 - linear regression
 - logistic regression
 - multiclass (multinomial) regression
 - Feedforward networks
 - one hidden layer
 - universal approximation theorem
-

- Neural network parts
 - model: architecture, activation functions
 - method: loss function
 - algorithm: how to estimate parameters, gradient descent and back-propagation
 - recent developments
 - Deep learning
 - the timeline
 - Keras
 - Summing up
 - Recommended exercises
 - References
-

Why a module on neural networks?

- Every day you read about the success of AI, machine learning — and in particular deep learning.

- In the last five years the field of deep learning has gone from low level performance to excellent performance — particularly in image recognition and speech transcription.

- Deep learning is based on a layered artificial neural network structure.

But, what is a *neural network*?






Image title: Neuron and myelinated axon, with signal flow from inputs at dendrites to outputs at axon terminals. Image credits: By Egm4313.s12 (Prof. Loc Vu-Quoc) - Own work, CC BY-SA 4.0, <https://commons.wikimedia.org/w/index.php?curid=72816083>

- Understanding the basic building blocks of artificial neural network, and how they are connected to concepts in mathematical statistics might be useful!
- The focus here will mainly be to relate what we have learned from regression and classification in this course to the basic framework of artificial neural network,
- and we will touch upon the new and hot stuff.

- There are several (self-study) learning resources (some listed under References) that the student may turn to for further knowledge into deep learning, but this presentation is heavily based on Chollet and Allaire (2018), with added formulas and theory.
- Or, you may plan to take the new IT3030 Deep learning course (Langseth and Downing), that will start up in the spring of 2020, where the main reading will be Goodfellow, Bengio, and Courville (2016).
- According to Chollet and Allaire (2018) (page xiv) *there are no difficult ideas in deep learning*, and this module will try to help the *democratization of deep learning*.



AI, machine learning and statistics

Artificial intelligence dates back to the 1950s, and can be seen as *the effort to automate intellectual tasks normally performed by humans* (page 4, Chollet and Allaire (2018)). AI was first based on hardcoded rules, but with the field of machine learning the shift is that a system is *trained* rather than explicitly programmed. See figure 1.1 of Chollet and Allaire (2018), but where is statistics?

Machine learning is related to mathematical statistics, but differ in many ways (as we will soon see). In short, machine learning attaches much larger, and more complex data sets than what is usually done in statistics. In addition, focus in machine learning is not on mathematical theory, but is more oriented towards engineering, and ideas are proven *empirically* rather than theoretically (which is the case in mathematical statistics).

According to Chollet and Allaire (2018) (page 19): *Machine learning isn't mathematics or physics, where major advancements can be done with a pen and a piece of paper. It's an engineering science.*

From statistics to artificial neural networks

Multiple linear regression

Problem

Recapitulate from Module 3 with Munich rent index, 1999: observations on the following variables.

- **rsqm**: the net rent per month per square meter (in Euro).
- **area**: Living area in square meters.
- **yearc**: year of construction.
- **loc**: quality of location: a factor indicating whether the location is average location, 1, good location, 2, and top location, 3.
- **bath**: quality of bathroom: a factor indicating whether the bath facilities are standard, 0, or premium, 1.
- **kitch**: Quality of kitchen: 0 standard 1 premium.
- **cheat**: central heating: a factor 0 without central heating, 1 with central heating.

We will now look at modelling the **rsqm** as respons and using **area**, **yearc**, **loc** (dummy variable coding, average is reference), **bath**, **kitch**, **cheat** as covariates - this will give us

- one numerical output (response), and
- seven covariates
- one intercept

Let n be the number of observations in the training set, here $n = 3082$.

Statistical model

(from Module 3)

For $i = 1, \dots, n$ we have:

$$Y_i = \beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip} + \varepsilon_i = \mathbf{x}_i^T \boldsymbol{\beta} + \varepsilon_i$$

which we wrote with vectors (for all observations $i = 1, \dots, n$ together):

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

Assumptions:

1. $E(\boldsymbol{\varepsilon}) = \mathbf{0}$.
2. $\text{Cov}(\boldsymbol{\varepsilon}) = E(\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^T) = \sigma^2\mathbf{I}$.
3. The design matrix has full rank, $\text{rank}(\mathbf{X}) = p + 1$. (We assume $n \gg (p + 1)$.)

The classical *normal* linear regression model is obtained if additionally

4. $\boldsymbol{\varepsilon} \sim N_n(\mathbf{0}, \sigma^2\mathbf{I})$ holds. Here N_n denotes the n -dimensional multivariate normal distribution.
-

From statistical model to network architecture

How can our statistical model be represented as a network?

New concepts:

- Our covariates are now presented as *input nodes* in an *input layer*.
 - The intercept is presented as a node, and is called a *bias* node. (New meaning to us!)
 - The response is presented as one *output node* in an *output layer*.
 - The regression coefficients are called *weights* and are often written on the lines (arrows) from the inputs to the output node.
 - All lines going into the output node signifies that we multiply the covariate values in the input nodes with the weights (regression coefficients), and then sum. This sum can be sent through a so-called *activation function*. The activation function for linear regression is just the identity function.
-

The regression function is then rewritten from the joint form

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

or the form for each observation i

$$Y_i = \mathbf{x}_i^T \boldsymbol{\beta} + \varepsilon_i$$

to the neural network model

$$y_1(\mathbf{x}_i) = \phi_o(w_0 + w_1 x_{i1} + \dots + w_p x_{ip})$$


where $\phi_o(x) = x$.

Observe: we do not say anything of what is random and fixed, and do not make any assumption distribution of a random variable.

From a statistical point of view we would instead have written $\hat{y}_1(\mathbf{x}_i)$ to specify that we are looking for using this formula to predict a value of the response for the given covariate value. To be able to distinguish this predicted response from the observed response we use the notation:

$$\hat{y}_1(\mathbf{x}_i) = \phi_o(w_0 + w_1x_{i1} + \dots + w_px_{ip})$$

The only difference to our MLR model is then that we would have called the w s $\hat{\beta}$ s instead.



Statistical parameter estimation

In multiple linear regression parameters in β are estimated with maximum likelihood and least squares. These two methods give the same estimator when we assume the normal linear regression model.

The estimator $\hat{\beta}$ is found by minimizing the RSS for a multiple linear regression model:

$$\begin{aligned} \text{RSS} &= \sum_{i=1}^n (y_i - \hat{y}_i)^2 = \sum_{i=1}^n (y_i - \hat{\beta}_0 - \hat{\beta}_1x_{i1} - \hat{\beta}_2x_{i2} - \dots - \hat{\beta}_px_{ip})^2 \\ &= \sum_{i=1}^n (y_i - \mathbf{x}_i^T \hat{\beta})^2 = (\mathbf{Y} - \mathbf{X}\hat{\beta})^T (\mathbf{Y} - \mathbf{X}\hat{\beta}) \end{aligned}$$

This problem has a solution given on closed form:

$$\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$$

Neural networks: loss function and gradient descent

We now translate what we did for the MLR into what is done for neural networks:

1. To estimate our parameters (now called our *network weights*) we define a *loss function*. We used for MLR the RSS, which is the sum over *all observations* in our *training data set*. We will still use that, but not this is called the *mean squared error* (and we scale with the number of observations in the training set)

$$J(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_1(\mathbf{x}_i))^2$$

here $J(\mathbf{w})$ gives focus to that the unknown parameters are the weights.

2. For MLR we found the parameter estimates by *minimizing* this loss function
 - we found the derivative of the loss function with respect to each of our parameters and
 - ended up with $(p + 1)$ linear equations, which we solved.
3. For this simple model the solution was given on closed form.

What if the solution was not available on closed form? What would we have done then?

Now we replace the words parameters with *weights* (and also think of this bias term as a weight).

- Assume we have initialized our weights by just some random numerical values, and then calculate the predicted values $\hat{y}_1(\mathbf{x}_i)$ for each observation, and then calculated the loss function. We want to find the weights that minimize this loss function.
 - We know from Mathematics 2 (Calculus) that the *gradient of a function* gives the direction of the steepest ascent for this function. Our function is a function of the $p + 1$ weights, so we are in $(p + 1)$ -dimensional space of real numbers. The gradient is vector in this $(p + 1)$ -dimensional space.
 - If we go in the direction of the *negative of the gradient* this is the direction that *decreases the function* most quickly. We want to minimize the loss (not to maximize it).
-

- How far should we go in that direction? This is the *step length*, and can be chosen in many ways (when we do numerical optimization). In neural networks this is referred to as the *learning rate* and is often set by the user.
 - We now update our weight to this new point in our $p + 1$ -dimensional weight space.
 - At this new point, we calculate new predicted values $\hat{y}_1(\mathbf{x}_i)$ for each observation, and then calculated the new value of the loss function, and then the new value of the gradient, and move to the new point, and repeat this until we are at the global (or local) optimum of our loss function.
-

Figures that give good illustration of the optimization problem.

Chollet and Allaire (2018):

- 2.11: SGD down a 1D loess curve
- 2.12: Gradient descent down a 2D loss surface
- 2.13: local and global minimum

Goodfellow, Bengio, and Courville (2016): see Figure 4.1 and 4.3.

Algorithm

1. Let $t = 0$ and denote the given initial values for the weights $\mathbf{w}^{(t)}$,
2. *propagate* the observations through the network and calculate the predictions $\hat{y}_1(\mathbf{x}_i)$
3. calculate the loss function $J(\mathbf{w}^{(0)})$,

4. find the gradient (direction) in the $(p + 1)$ -dimensional space of the weights, and evaluate this at the current weight values $\nabla J(\mathbf{w}^{(0)}) = \frac{\partial J}{\partial \mathbf{w}}(\mathbf{w}^{(0)})$
5. go with a given step length (learning rate) λ in the direction of the negative of the gradient of the loss function to get updated values for the weights

$$\mathbf{w}^{(t+1)} = \mathbf{w}^{(t)} - \lambda \nabla J(\mathbf{w}^{(t)})$$

6. Set $t = t + 1$, go to 2. and continue to 6. several times until you arrive at a (local) optimum
7. The final values of the weights in that $(p + 1)$ dimensional space are our parameter estimates and your network is *trained* and can be used for prediction on a test set.

In neural networks the gradient part of the gradient descent algorithm is implemented efficiently in an algorithm called *backpropagation* - which we will look into a bit later.

Here we compare the MLR solution with `lm` to the neural network solution with `nnet` (which in fact improves upon the gradient descent with Hessian information and the BFGS-algorithm).

Parameter estimation vs. network weights

```
fit = lm(rsqm ~ area + yearc + loc + bath + kitch + cheat, data = rent99)
fitnnet = nnet(rsqm ~ area + yearc + loc + bath + kitch + cheat, data = rent99,
  linout = TRUE, size = 0, skip = TRUE, maxit = 1000, entropy = FALSE)
```

```
## # weights: 8
## initial value 1679110405.489187
## iter 10 value 4547174.822698
## final value 12679.131795
## converged
```

```
cbind(fitnnet$wts, fit$coefficients)
```

```
##           [,1]      [,2]
## (Intercept) -45.47573121 -45.47548356
## area        -0.03233024 -0.03233033
## yearc        0.02695870  0.02695857
## loc2         0.77713364  0.77713297
## loc3         1.72506019  1.72506792
## bath1        0.76278335  0.76280784
## kitch1       1.13691082  1.13690814
## cheat1       1.76526280  1.76526110
```

Logistic regression

In the exam in 2018 (and on the compulsory exercise 2 in 2019), we have looked at predicting if a person has diabetes. The data is from a population of women of Pima Indian heritage in the US, available in the R `MASS` package. The following information is available for each woman:

- diabetes: 0= not present, 1= present
- npreg: number of pregnancies
- glu: plasma glucose concentration in an oral glucose tolerance test

- bp: diastolic blood pressure (mmHg)
- skin: triceps skin fold thickness (mm)
- bmi: body mass index (weight in kg/(height in m)²)
- ped: diabetes pedigree function.
- age: age in years

According to Chollet and Allaire (2018) (page 14) logistic regression (referred to as logreg) is the “hello world” of machine learning - that is, the first method a data scientist will try on a two-class data set (to get a feel of the problem).

The statistical model

We have $i = 1, \dots, n$ observations in the training set. We will use r (instead of p) to be the number of covariates, to avoid confusion with the probability p .

Assume that the response Y_i is coded ($\mathcal{C} = \{1, 0\}$ or {success, failure}), and we focus on success. We may assume that Y_i follows a Bernoulli distribution with probability of success p_i .


$$Y_i = \begin{cases} 1 & \text{with probability } p_i, \\ 0 & \text{with probability } 1 - p_i. \end{cases}$$

In logistic regression we *link* together our covariates \mathbf{x}_i with this probability p_i using a *logistic function*.

$$p_i = \frac{\exp(\beta_0 + \beta_1 x_{i1} + \dots + \beta_r x_{ir})}{1 + \exp(\beta_0 + \beta_1 x_{i1} + \dots + \beta_r x_{ir})} \tag{1}$$

$$= \frac{1}{1 + \exp(-\beta_0 - \beta_1 x_{i1} - \dots - \beta_r x_{ir})} \tag{2}$$

This function is S-shaped, and ranges between 0 and 1 (so the p_i is between 0 and 1).



Neural network architecture and activation function

Compared to the neural network version of the MLR, we only need to change one thing, the activation function of the output node. We had:

$$y_1(\mathbf{x}_i) = \phi_o(w_0 + w_1x_{i1} + \dots + w_rx_{ir})$$

where $\phi_o(x) = x$, for the MLR, and now we instead have the logistic function as the function ϕ_o , that is

$$\phi_o(x) = \frac{1}{1 + \exp(-x)}$$

This is now referred to as the *sigmoid* activation function and is often denoted $\sigma(x)$, and is said to *squash the values to be between 0 and 1*. Again, we prefer to use $\hat{y}_1(\mathbf{x}_i)$ and get:

$$\hat{y}_1(\mathbf{x}_i) = \frac{1}{1 + \exp(-(w_0 + w_1x_{i1} + \dots + w_rx_{ir}))}$$

Parameter estimation with maximum likelihood for logistic regression

We assume that pairs of covariates and responses $\{x_i, y_i\}$ are measured independently of each other. Given n such observation pairs, the likelihood function of a logistic regression model can be written as:

$$L(\boldsymbol{\beta}) = \prod_{i=1}^n L_i(\boldsymbol{\beta}) = \prod_{i=1}^n f(y_i; \boldsymbol{\beta}) = \prod_{i=1}^n (p_i)^{y_i} (1 - p_i)^{1-y_i},$$

where $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \dots, \beta_r)^T$ enters into p_i .

The maximum likelihood estimates are found by maximizing the log-likelihood (because this makes the maths easier).

$$\ln(L(\boldsymbol{\beta})) = l(\boldsymbol{\beta}) = \sum_{i=1}^n (y_i \ln p_i + (1 - y_i) \ln(1 - p_i))$$

- To maximize the log-likelihood function we find the $r + 1$ partial derivatives (to form the gradient), and set equal til 0.
- This gives us a set of $r + 1$ non-linear equations in the β s.
- This set of equations does not have a closed form solution.
- These equations are therefore solved numerically, using the *Newton-Raphson algorithm* (or Fisher Scoring):

$$\boldsymbol{\beta}^{(t+1)} = \boldsymbol{\beta}^{(t)} + \mathbf{F}(\boldsymbol{\beta}^{(t)})^{-1} \mathbf{s}(\boldsymbol{\beta}^{(t)})$$

where the gradient of the log-likelihood $\mathbf{s}(\boldsymbol{\beta}) = \frac{\partial l}{\partial \boldsymbol{\beta}}$ is called the score vector, and here the new quantity $\mathbf{F}(\boldsymbol{\beta}^{(t)})^{-1}$ is called the inverse *expected Fisher information matrix* and is the expected value of the negative of the gradient of the score vector (the negative of the Hessian matrix of the loglikelihood), and also the covariance matrix of the score vector. (Observe that we here are maximizing so we are going in the direction of the gradient, not the negative of the direction which is needed when we minimize.)

Neural networks: loss function and gradient descent

For parameter estimation we looked at maximizing the log-likelihood of the statistical model. For neural networks the negative of the binomial loglikelihood is a scaled version of the *binomial cross-entropy loss*.

$$J(\mathbf{w}) = -\frac{1}{n} \sum_{i=1}^n (y_i \ln(\hat{y}_1(\mathbf{x}_i)) + (1 - y_i) \ln(1 - \hat{y}_1(\mathbf{x}_i)))$$

The optimization is also now done using gradient descent, but observe that due to the activation function we can use the chain rule when calculating the partial derivatives to get the gradient direction. The Algorithm given for MLR is also applicable now, with the modification to the activation and loss function given here.

Parameter estimation vs. network weights

```
fitlogist = glm(diabetes ~ npreg + glu + bp + skin + bmi + ped + age,  
  data = train, family = binomial(link = "logit"))  
summary(fitlogist)
```

```
##  
## Call:  
## glm(formula = diabetes ~ npreg + glu + bp + skin + bmi + ped +  
##   age, family = binomial(link = "logit"), data = train)  
##  
## Deviance Residuals:  
##   Min       1Q   Median       3Q      Max  
## -1.9830  -0.6773  -0.3681   0.6439   2.3154  
##  
## Coefficients:  
##              Estimate Std. Error z value Pr(>|z|)  
## (Intercept) -9.773062   1.770386  -5.520 3.38e-08 ***  
## npreg        0.103183   0.064694   1.595 0.11073  
## glu          0.032117   0.006787   4.732 2.22e-06 ***  
## bp          -0.004768   0.018541  -0.257 0.79707  
## skin        -0.001917   0.022500  -0.085 0.93211  
## bmi         0.083624   0.042827   1.953 0.05087 .  
## ped         1.820410   0.665514   2.735 0.00623 **  
## age         0.041184   0.022091   1.864 0.06228 .  
## ---  
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1  
##  
## (Dispersion parameter for binomial family taken to be 1)  
##  
##   Null deviance: 256.41  on 199  degrees of freedom  
## Residual deviance: 178.39  on 192  degrees of freedom  
## AIC: 194.39  
##  
## Number of Fisher Scoring iterations: 5
```

```
library(nnet)  
library(NeuralNetTools)  
fitnnet = nnet(diabetes ~ npreg + glu + bp + skin + bmi + ped + age,  
  data = train, linout = FALSE, size = 0, skip = TRUE, maxit = 1000,  
  entropy = TRUE, Wts = fitlogist$coefficients + rnorm(8, 0, 0.1))
```

```
## # weights:  8  
## initial value 10987.990346  
## iter 10 value 117.355781  
## iter 20 value 89.212650  
## final value 89.195333  
## converged
```


```
# entropy=TRUE because default is least squares  
cbind(fitnnet$wts, fitlogist$coefficients)
```

```
##              [,1]      [,2]  
## (Intercept) -9.773059448 -9.773061533
```

```
## npreg      0.103183429  0.103183427
## glu        0.032116820  0.032116823
## bp         -0.004767551 -0.004767542
## skin       -0.001916624 -0.001916632
## bmi        0.083623875  0.083623912
## ped        1.820410616  1.820410367
## age        0.041183526  0.041183529
```

In the `nnet` R package a slightly different version is used, `entropy=maximum conditional likelihood` which is half the deviance for the logistic regression model.

```
plotnet(fitnnet)
```



But, there may also exist local minima.

```
set.seed(123)
fitnnet = nnet(diabetes ~ npreg + glu + bp + skin + bmi + ped + age,
  data = train, linout = FALSE, size = 0, skip = TRUE, maxit = 10000,
  entropy = TRUE, Wts = fitlogist$coefficients + rnorm(8, 0, 1))
```

```
## # weights:  8
## initial value 24315.298582
## final value 12526.062906
## converged
```

```
cbind(fitnnet$wts, fitlogist$coefficients)
```

```
##           [,1]      [,2]
## (Intercept) -36.733537 -9.773061533
## npreg        -77.126994  0.103183427
## glu         -2984.409175  0.032116823
## bp          -1835.934259 -0.004767542
## skin        -718.072629 -0.001916632
## bmi         -818.561311  0.083623912
## ped         -8.687473   1.820410367
## age         -773.023878  0.041183529
```

Multiclass regression

Which type of iris species?

The iris flower data set was introduced by the British statistician and biologist Ronald Fisher in 1936, and we studied that in module 4 on classification.

- Three plant species {setosa, virginica, versicolor} (50 observation of each), and
- four features: Sepal.Length, Sepal.Width, Petal.Length and Petal.Width.

The aim is to predict the species of an iris plant.

Statistical model

We only briefly mentioned this method in module 4 (there our focus was on LDA and QDA when we had more than two classes), and the full theoretical treatment is given in the statistics course TMA4315 Generalized linear models, Module 6 on categorical regression (nominal response).

Assumptions:

We have independent observation pairs (Y_i, \mathbf{x}_i) , where the covariate vector \mathbf{x}_i consists of the same measurements for each response category (that is, not different covariate types that are measured for each response category).

Each observation can only belong to one response class, $Y_i = c$ for $c = 1, \dots, C$, and we define a dummy variable coding of the response in a C -dimensional vector: $\mathbf{y}_i = (0, 0, \dots, 0, 1, 0, \dots, 0)$ with a value of 1 in the c th element of \mathbf{y}_i if the class is c .

The probabilities are $p_{ic} = P(Y_i = c)$ that the response is category c for subject i , where $\sum_{c=1}^C p_{ic} = 1$ (which means that $p_{iC} = 1 - \sum_{c=1}^{C-1} p_{ic}$ and in statistics we don't include p_{iC} in the modelling).

To analyse this type of data a generalization of the model for the logistic regression (for two classes) can be used, and the model we fit is:

$$\begin{aligned} p_{ic} &= P(Y_i = c) = \frac{\exp(\mathbf{x}_i^T \boldsymbol{\beta}_c)}{1 + \sum_{s=1}^{C-1} \exp(\mathbf{x}_i^T \boldsymbol{\beta}_s)} \\ &= \frac{\exp(\beta_{0c} + \beta_{1c}x_{i1} + \dots + \beta_{rc}x_{ir})}{1 + \sum_{s=1}^{C-1} \exp(\beta_{0s} + \beta_{1s}x_{i1} + \dots + \beta_{rs}x_{ir})} \end{aligned}$$

See the difference to $p_i = P(Y_i = 1) = \frac{\exp(\mathbf{x}_i^T \boldsymbol{\beta})}{1 + \exp(\mathbf{x}_i^T \boldsymbol{\beta})}$ for two classes.

An observation is classified to the class with the highest probability, $\operatorname{argmax} p_{ic}$.

Q: How many parameters are we estimating?

Neural network architecture and activation function

The neural network uses the dummy variable coding of the responses, but call this *one-hot* coding, and builds an output layer with C nodes — and corresponding 0/1 targets (responses).

The activation function for the output layer is called *softmax* and is given as

$$\hat{y}_1(\mathbf{x}_i) = \frac{\exp(\mathbf{x}_i^T \mathbf{w}_1)}{\sum_{s=1}^C \exp(\mathbf{x}_i^T \mathbf{w}_s)} \quad (3)$$

$$\hat{y}_2(\mathbf{x}_i) = \frac{\exp(\mathbf{x}_i^T \mathbf{w}_2)}{\sum_{s=1}^C \exp(\mathbf{x}_i^T \mathbf{w}_s)} \quad (4)$$

$$\vdots = \vdots \quad (5)$$


$$\hat{y}_C(\mathbf{x}_i) = \frac{\exp(\mathbf{x}_i^T \mathbf{w}_C)}{\sum_{s=1}^C \exp(\mathbf{x}_i^T \mathbf{w}_s)} \quad (6)$$

Where each \mathbf{w}_s is a $r + 1$ dimensional vector of weights.

Observe that there is some redundancy here, since $\sum_{c=1}^C \hat{y}_{ci}(\mathbf{x}_i) = 1$, so we could have had $C - 1$ output nodes, but this is not done.

The focus of neural networks is not to interpret the weights, and there is no need to assume full rank of a matrix with output nodes.

Q: How many parameters are we estimating?



Parameter estimation

The likelihood of the multinomial regression model can be written as

$$\ln(L(\beta)) \propto \sum_{i=1}^n \sum_{c=1}^C y_{ic} \ln(p_{ic})$$

where $p_{iC} = 1 - p_{i1} - p_{i2} - \dots - p_{i,C-1}$, and the regression parameters enters via the p_{ic} .

Parameter estimation is done in the same way as for the logistic regression, with the Fisher scoring algorithm (with score vector and Fisher information matrix).

Neural networks: loss function and gradient descent

For parameter estimation we looked at maximizing the log-likelihood of the statistical model. For neural networks the negative of the multinomial loglikelihood is a scaled version of the *categorical cross-entropy loss*.

$$J(\mathbf{w}) = -\frac{1}{n} \sum_{i=1}^n \frac{1}{C} \sum_{c=1}^C (y_{ic} \ln(\hat{y}_c(\mathbf{x}_i)))$$

The optimization is done using gradient descent, with minor changes from what was done for the logistic regression due to the added sum and the small change in the activation function.

Parameter estimation vs. network weights

```
library(caret)
# multiclass model via neural network
fit=multinom(species~.,family=multinomial, data=iris_train)
```

```
## # weights: 18 (10 variable)
## initial value 54.930614
## iter 10 value 4.730606
## iter 20 value 3.290523
## iter 30 value 3.207769
## iter 40 value 3.189667
## iter 50 value 3.189261
## iter 60 value 3.189216
## final value 3.189211
## converged
```

```
summary(fit)
```

```
## Call:
## multinom(formula = species ~ ., data = iris_train, family = multinomial)
##
## Coefficients:
## (Intercept) Sepal.L. Sepal.W. Petal.L. Petal.W.
## s 1.452577 1.0478218 6.417625 -9.330466 -0.363278
## v -18.236797 -0.2376874 -7.611047 4.889676 10.349686
##
## Std. Errors:
## (Intercept) Sepal.L. Sepal.W. Petal.L. Petal.W.
## s 30.82585 95.516371 161.43601 65.427607 43.25393
## v 13.80006 4.789165 7.50082 6.173063 10.15204
##
## Residual Deviance: 6.378423
## AIC: 26.37842
```

```
testclass=predict(fit,new=iris_test)
confusionMatrix(data=testclass,reference=iris_test$species)$table
```

```
##           Reference
## Prediction c s v
##           c 26 0 1
```



```
##          s 0 33 0
##          v 2 0 38
```

```
library(NeuralNetTools)
#plotnet(iris.nnet)
summary(iris.nnet)
```

```
## a 4-0-3 network with 15 weights
## options were - skip-layer connections  softmax modelling  decay=5e-04
## b->o1 i1->o1 i2->o1 i3->o1 i4->o1
##  7.51 -0.90  1.20  0.90 -3.16
## b->o2 i1->o2 i2->o2 i3->o2 i4->o2
##  1.38  2.24  4.68 -6.59 -3.12
## b->o3 i1->o3 i2->o3 i3->o3 i4->o3
## -8.90 -1.34 -5.89  5.68  6.28
```

```
library(caret)
table(predict(iris.nnet, iris_test, type = "class"),iris_test$species)
```

```
##
##      c  s  v
## c 26  0  1
## s  0 33  0
## v  2  0 38
```

Summing up

- Fitting a multiple linear regression
 - can be achieved by building a neural network with one input layer and one node in the output layer,
 - linear activation function,
 - and mean squared loss,
 - Classification with two classes can be performed using logistic regression, and this
 - by building a neural network with one input layer and one node in the output layer,
 - sigmoid activation function,
 - and binary cross-entropy loss.
 - Remember: this will only give linear boundaries between the classes (in the output space).
-
- Classification with C classes, $C > 2$, can be performed using multinomial regression,
 - by building a neural network with one input layer and C nodes in the output layer,
 - softmax activation function,
 - and categorical cross-entropy loss.
 - Also here: this will only give linear boundaries between the classes (in the output space).
-
- We have seen that parameters (weights) can be found using gradient descent algorithms:
 - activation function “must” be differentiable
 - step length (learning rate) must be set
 - But, we have now only looked at linear models — which give linear boundaries in the classification case — now we need to move on to allow for non-linearities!

Feedforward networks

In feedforward networks we have only connections (weights) forward in the network, and no feedback connections that sends the output of the model back into the network. The word *multi-layer perceptron* (MLP) and *sequentially layered* networks are also used

The examples of MLR, logistic regression and multiclass regression are examples of feedforward networks without any so-called *hidden layers* (between the input and output layers).

We may have no hidden layer, one (to be studied next), or many.

The number of hidden layers is called the *depth* of the network, and the number of nodes in a layer is called the *width* of the layer.

The idea of using many layers of many nodes is inspired from neuroscience, but today we don't have the goal to model the brain — but instead to approximate function to perform statistical generalizations and maybe also insight into the problem at hand.

Now we will see how adding *hidden layers* with *non-linear activation functions* between the input and output layer will make nonlinear statistical models.

The single hidden layer feedforward network

The word neuron and node can be used interchangeably. We stick with greek letters α and β for parameters, but call them weights.

We use the following notation:

1. Inputs (input layer nodes), $j = 1, \dots, p$: x_1, x_2, \dots, x_p , or as a vector \mathbf{x} .
2. The nodes in the hidden layer, $m = 1, \dots, M$: z_1, z_2, \dots, z_m , or as vector \mathbf{z} , and the hidden layer activation function ϕ_h .

$$z_m(\mathbf{x}) = \phi_h(\alpha_{0m} + \sum_{j=1}^p \alpha_{jm}x_j)$$

where α_{jm} is the weight from input j to hidden node m , and α_{0m} is the bias term for the m th hidden node. The hidden nodes can be thought of as *latent variables*.

3. The node(s) in the output layer, $c = 1, \dots, C$: y_1, y_2, \dots, y_C , or as vector \mathbf{y} , and output layer activation function ϕ_o .

$$y_c(\mathbf{x}) = \phi_o(\beta_{0c} + \sum_{m=1}^M \beta_{mc}z_m(\mathbf{x}))$$


where β_{mc} is from hidden neuron m to output node c , and β_{0c} is the bias term for the c th output node.

To sum up: $c = 1, \dots, C$: y_1, y_2, \dots, y_C , where y_c is given as

$$y_c(\mathbf{x}) = \phi_o(\beta_{0c} + \sum_{m=1}^M \beta_{mc}z_m) = \phi_o(\beta_{0c} + \sum_{m=1}^M \beta_{mc}\phi_h(\alpha_{0m} + \sum_{j=1}^p \alpha_{jm}x_j))$$

Hands on:

- Identify p , M , C in the network figure below, and relate that to the $y_c(\mathbf{x})$ equation.
- How many parameters need to be estimated for this network?
- What decides the value of p and C ?
- What is the connection between problem, ϕ_o and C ?



Special case with linear activation function for the hidden layer

$$y_c(\mathbf{x}) = \phi_o\left(\beta_{0c} + \sum_{m=1}^M \beta_{mc} z_m\right) = \phi_o\left(\beta_{0c} + \sum_{m=1}^M \beta_{mc} \phi_h\left(\alpha_{0m} + \sum_{j=1}^p \alpha_{jm} x_j\right)\right)$$

Here we assume that $\phi_h(z) = z$, called linear or identity activation:

$$y_c(\mathbf{x}) = \phi_o\left(\beta_{0c} + \sum_{m=1}^M \beta_{mc} \left(\alpha_{0m} + \sum_{j=1}^p \alpha_{jm} x_j\right)\right)$$

Q: Does this look like something you have seen before?

For regression (linear output activation also) we have looked at a similar model in Module 6: principal component regression. Then the weights α_{jm} were not estimated, but were defined to be the eigenvectors corresponding to the m th largest eigenvalue of the estimated covariance matrix of the covariates \mathbf{x}_i . Therefore we may think of the hidden nodes as *latent variables*. In PCR only the β s were estimated using the responses (and the latent variables).

In Module 6 we also touched upon Partial least squares, where the latent variables were found with the help of the response y . Here both the α s and β s were estimated.

In Module 4 we only considered “ordinary logistic regression”, but we could have used the principal components in place of the original covariates — which then could have been depicted as a feedforward network with one hidden layer, but where only the β s were estimated.

Universal approximation property

We may think of the goal of a feedforward network to approximate some function f , mapping our input vector \mathbf{x} to an output value \mathbf{y} (as we saw with regression and classification).

What type of mathematical function can a feedforward neural network with one hidden layer and linear output activation represent?


According to Goodfellow, Bengio, and Courville (2016), page 198: The *universal approximation theorem* says that a feedforward network with

- a *linear output layer*
- at least one hidden layer with a “squashing” activation function and “enough” hidden units

can approximate any (Borel measurable) function from one finite-dimensional space (our input layer) to another (our output layer) with any desired non-zero amount of error.

The theorem has been shown both for the sigmoid (logistic) and ReLU activation functions in the hidden layer.

The rectified linear unit ReLU $\phi_h(a) = \max(0, a)$



Earlier, the default activation function for the hidden layer was the sigmoid, but now the ReLU is default activation function to be used in the hidden layer(s) of a feedforward network — in particular when more than one hidden layer is used. The ReLU function is piecewise linear, but in total non-linear.

Even though a large feedforward network with one hidden layer may be able to represent a desired function, we may not be able to estimate the parameters of the function,

- we may choose a too many or too few nodes in the hidden layer
- our optimization routine may fail
- we may overfit/underfit the training data

Therefore, sometimes networks with more than one hidden layer is used — with fewer total number of nodes but more layers. A network with many hidden layers is called a *deep network*.

In module 7 we looked at additive models of “complex” functions (splines) of one covariate each.

Now we look at many rather simple non-linear function of linear combinations of covariates, and non-linear functions of non-linear functions of linear combinations of covariates.

Q: Is one better than the other when it comes to interpretation, and to prediction?

The `nnet` R package

In order to not make this too complex (since we only have one week to work with this module), we will use the `nnet` R package by Brian Ripley (instead of the currently very popular `keras` package for deep learning — however, we will also present the `keras` package for completeness).

`nnet` fits one hidden layer with sigmoid activation function. The implementation is not gradient descent, but instead BFGS using `optim`.

Description: Fit single-hidden-layer neural network, possibly with skip-layer connections.

Usage: (for formula class): `nnet(formula, data, weights, ..., subset, na.action, contrasts = NULL)`

`nnet(x, y, ...)`

Some arguments to `nnet()`

- formula, or x and y
- x= input variables, matrix or data frame
- y= response (target) values, if factor=classification
- size: number of units in the hidden layer. Can be zero if there are skip-layer units.
- data: the dataset to be used

- subset: index of entries in data that is the training sample

If the response in formula is a factor, an appropriate classification network is constructed; this has one output and entropy fit if the number of levels is two, and a number of outputs equal to the number of classes and a softmax output stage for more levels.

- linout: switch for linear output units. Default logistic output units.
 - entropy: switch for entropy (= maximum conditional likelihood) fitting. Default by least-squares.
 - softmax switch for softmax (log-linear model) and maximum conditional likelihood fitting. linout, entropy, softmax and censored are mutually exclusive.
 - censored: A variant on softmax, in which non-zero targets mean possible classes. Thus for softmax a row of (0, 1, 1) means one example each of classes 2 and 3, but for censored it means one example whose class is only known to be 2 or 3.
 - skip: switch to add skip-layer connections from input to output.
 - decay: parameter for weight decay. Default 0.
 - maxit: maximum number of iterations. Default 100.
 - MaxNWts: The maximum allowable number of weights. There is no intrinsic limit in the code, but increasing MaxNWts will probably allow fits that are very slow and time-consuming.
-

Examples

Boston house prices

The objective here is to predict the median price of homes in a given Boston suburb in the mid-1970s, and 10 input variables are given.

This data set is both available in the MASS and `keras` R package.

Preparing the data

- Few data points: only 506, split between 404 training samples and 102 test samples (this split already done in the `keras` library)
- Each feature in the input data (for example, the crime rate) has a different scale, some values are proportions, which take values between 0 and 1; others take values between 1 and 12, others between 0 and 100, and so on.
- The targets are the median values of owner-occupied homes, in thousands of dollars.

```
library(keras)
dataset <- dataset_boston_housing()
c(c(train_data, train_targets), c(test_data, test_targets)) %<-% dataset
str(train_targets)

## num [1:404(1d)] 15.2 42.3 50 21.1 17.7 18.5 11.3 15.6 15.6 14.4 ...

head(train_data)

##      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9] [,10] [,11]
## [1,] 1.23247 0.0 8.14 0 0.538 6.142 91.7 3.9769 4 307 21.0
## [2,] 0.02177 82.5 2.03 0 0.415 7.610 15.7 6.2700 2 348 14.7
## [3,] 4.89822 0.0 18.10 0 0.631 4.970 100.0 1.3325 24 666 20.2
## [4,] 0.03961 0.0 5.19 0 0.515 6.037 34.5 5.9853 5 224 20.2
## [5,] 3.69311 0.0 18.10 0 0.713 6.376 88.4 2.5671 24 666 20.2
## [6,] 0.28392 0.0 7.38 0 0.493 5.708 74.3 4.7211 5 287 19.6
##      [,12] [,13]
## [1,] 396.90 18.72
## [2,] 395.38 3.11
## [3,] 375.52 3.26
## [4,] 396.90 8.01
## [5,] 391.43 14.65
## [6,] 391.13 11.74

library(MASS) #no names for columns - take that from other source
colnames(Boston)

## [1] "crim" "zn" "indus" "chas" "nox" "rm" "age"
## [8] "dis" "rad" "tax" "ptratio" "black" "lstat" "medv"

colnames(train_data) = colnames(Boston)[-13]
colnames(test_data) = colnames(Boston)[-13]
```

A widespread best practice to deal with such data is to do feature-wise normalization. This is to make the optimization easier with gradient based methods.

```
org_train = train_data
mean <- apply(train_data, 2, mean)
std <- apply(train_data, 2, sd)
train_data <- scale(train_data, center = mean, scale = std)
test_data <- scale(test_data, center = mean, scale = std)
```

Note that the quantities used for normalizing the test data are computed using the training data. You should never use in your workflow any quantity computed on the test data, even for something as simple as data normalization. This also means that we need to do this standardization again (from scratch) if we need to do cross-validation.

Just checking out one hidden layer with 5 units to get going.

```
library(nnet)
fit5 <- nnet(train_targets ~ ., data = train_data, size = 5, linout = TRUE,
             maxit = 1000)
```

```
## # weights: 76
## initial value 228708.652654
## iter 10 value 17128.620279
## iter 20 value 10935.487830
## iter 30 value 7709.690111
## iter 40 value 6934.725601
## iter 50 value 6420.119056
## iter 60 value 6027.466099
## iter 70 value 5573.989985
## iter 80 value 5469.766138
## iter 90 value 5448.703370
## iter 100 value 5429.437453
## iter 110 value 5411.646330
## iter 120 value 5393.592332
## iter 130 value 5383.646278
## iter 140 value 5375.047585
## iter 150 value 5264.737607
## iter 160 value 5204.251899
## iter 170 value 5197.162280
## iter 180 value 5187.868231
## iter 190 value 5157.187655
## iter 200 value 5147.647563
## iter 210 value 5136.934079
## iter 220 value 5128.790793
## iter 230 value 5119.240823
## iter 240 value 5117.469949
## iter 250 value 5116.563716
## iter 260 value 5115.679294
## iter 270 value 5115.067277
## iter 280 value 5114.462243
## iter 290 value 5114.354583
## iter 300 value 5114.197427
## iter 310 value 5114.184209
## iter 320 value 5114.177852
## iter 330 value 5114.163112
## iter 340 value 5114.125749
## iter 350 value 5114.076195
## iter 360 value 5113.479009
## iter 370 value 5112.200202
## iter 380 value 5112.021734
## iter 390 value 5111.919923
## iter 400 value 5111.828983
## iter 410 value 5111.797230
## iter 420 value 5111.692630
## iter 430 value 5111.629206
## iter 440 value 5111.582660
## final value 5111.574978
```



```
## converged
```

```
summary(fit5)
```

```
## a 13-5-1 network with 76 weights
## options were - linear output units
## b->h1 i1->h1 i2->h1 i3->h1 i4->h1 i5->h1 i6->h1 i7->h1 i8->h1
## -299.43 54.08 -117.70 -73.56 -33.00 -207.84 221.65 17.81 -5.66
## i9->h1 i10->h1 i11->h1 i12->h1 i13->h1
## 348.83 51.33 -8.03 -180.14 -217.82
## b->h2 i1->h2 i2->h2 i3->h2 i4->h2 i5->h2 i6->h2 i7->h2 i8->h2
## 62.61 -58.90 142.44 43.09 75.26 158.14 -36.89 -112.93 57.77
## i9->h2 i10->h2 i11->h2 i12->h2 i13->h2
## 244.93 -252.41 -318.19 69.34 -53.10
## b->h3 i1->h3 i2->h3 i3->h3 i4->h3 i5->h3 i6->h3 i7->h3 i8->h3
## 96.00 -130.08 -12.30 -25.58 15.00 -17.83 -35.98 -33.04 -62.99
## i9->h3 i10->h3 i11->h3 i12->h3 i13->h3
## 2.58 -19.39 26.04 24.18 -71.27
## b->h4 i1->h4 i2->h4 i3->h4 i4->h4 i5->h4 i6->h4 i7->h4 i8->h4
## -145.39 -81.88 -5.39 -23.19 -7.03 -34.95 159.94 -5.57 -67.23
## i9->h4 i10->h4 i11->h4 i12->h4 i13->h4
## 29.45 -38.53 8.29 0.12 -64.11
## b->h5 i1->h5 i2->h5 i3->h5 i4->h5 i5->h5 i6->h5 i7->h5 i8->h5
## -231.75 -78.49 3.34 88.52 37.54 -51.99 150.31 10.32 57.97
## i9->h5 i10->h5 i11->h5 i12->h5 i13->h5
## 14.38 -5.61 16.41 36.89 10.03
## b->o h1->o h2->o h3->o h4->o h5->o
## 9.49 5.54 4.03 8.94 7.21 12.04
```

```
pred = predict(fit5, newdata = test_data, type = "raw")
sqrt(mean((pred[, 1] - test_targets)^2))
```

```
## [1] 6.082564
```

```
mean(abs(pred[, 1] - test_targets))
```

```
## [1] 4.162192
```

Now, use cross-validation to find the best number of hidden nodes — but that took some time, so only results shown below. Remember the scaling need to be done within the loop.

```
grid = c(5, 10, 15, 20, 25, 30, 50)
```

```
k <- 4
```

```
set.seed(123)
```

```
indices <- sample(1:nrow(train_data))
```

```
folds <- cut(indices, breaks = k, labels = FALSE)
```

```
# have now assigned the training data to 4 different folds all of the  
# same size (101)
```

```
resmat = matrix(NA, ncol = k, nrow = length(grid))
```

```
for (j in 1:k) {
```

```
  thistrain = (1:dim(train_data)[1])[folds != j]
```

```
  thisvalid = (1:dim(train_data)[1])[folds == j]
```

```
  mean <- apply(org_train[thistrain, ], 2, mean)
```

```
  std <- apply(org_train[thistrain, ], 2, sd)
```

```

new <- scale(org_train, center = mean, scale = std)
for (i in 1:length(grid)) {
  thissize = grid[i]

  fit = nnet(train_targets ~ ., data = new, size = thissize, linout = TRUE,
            maxit = 5000)
  pred = predict(fit, newdata = new[thisvalid, ], type = "raw")
  resmat[i, j] = sum((pred[, 1] - train_targets[thisvalid])^2)
}
}
mse = apply(resmat, 1, sum)/nrow(train_data)
plot(grid, mse, type = "l")
mse
# 7.727735e+00 2.550359e+00 9.582672e-01 3.458643e-01 4.290271e-01
# 1.337877e-03 1.786437e-07

```

The best model here was the model with 50 nodes, the largest model we tried. Fitting that model on the full training set and testing on the test set

```

library(nnet)
fit50 <- nnet(train_targets ~ ., data = train_data, size = 50, linout = TRUE,
             maxit = 5000)

```

```

## # weights: 751
## initial value 279094.836571
## iter 10 value 8761.532788
## iter 20 value 2473.681039
## iter 30 value 1406.953259
## iter 40 value 894.700879
## iter 50 value 588.256747
## iter 60 value 447.801098
## iter 70 value 309.160656
## iter 80 value 180.667669
## iter 90 value 118.253736
## iter 100 value 78.970658
## iter 110 value 55.352917
## iter 120 value 40.223424
## iter 130 value 26.990290
## iter 140 value 19.095481
## iter 150 value 13.942489
## iter 160 value 10.127897
## iter 170 value 7.467880
## iter 180 value 5.843444
## iter 190 value 4.457298
## iter 200 value 2.776633
## iter 210 value 1.848395
## iter 220 value 1.162447
## iter 230 value 0.700183
## iter 240 value 0.445388
## iter 250 value 0.301094
## iter 260 value 0.221092
## iter 270 value 0.163954
## iter 280 value 0.131316

```

```

## iter 290 value 0.101229
## iter 300 value 0.070024
## iter 310 value 0.048856
## iter 320 value 0.034301
## iter 330 value 0.024344
## iter 340 value 0.016799
## iter 350 value 0.011626
## iter 360 value 0.007468
## iter 370 value 0.005091
## iter 380 value 0.003482
## iter 390 value 0.002534
## iter 400 value 0.001707
## iter 410 value 0.001054
## iter 420 value 0.000539
## iter 430 value 0.000213
## final value 0.000098
## converged

```

```
head(summary(fit50))
```

```

## $n
## [1] 13 50 1
##
## $nunits
## [1] 65
##
## $nconn
## [1] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 14 28
## [18] 42 56 70 84 98 112 126 140 154 168 182 196 210 224 238 252 266
## [35] 280 294 308 322 336 350 364 378 392 406 420 434 448 462 476 490 504
## [52] 518 532 546 560 574 588 602 616 630 644 658 672 686 700 751
##
## $conn
## [1] 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8
## [24] 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3
## [47] 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12
## [70] 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7
## [93] 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2
## [116] 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11
## [139] 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6
## [162] 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1
## [185] 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10
## [208] 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5
## [231] 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0
## [254] 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9
## [277] 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4
## [300] 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13
## [323] 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8
## [346] 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3
## [369] 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12
## [392] 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7
## [415] 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2
## [438] 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11
## [461] 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1 2 3 4 5 6
## [484] 7 8 9 10 11 12 13 0 1 2 3 4 5 6 7 8 9 10 11 12 13 0 1

```

```

## [507]  2  3  4  5  6  7  8  9 10 11 12 13  0  1  2  3  4  5  6  7  8  9 10
## [530] 11 12 13  0  1  2  3  4  5  6  7  8  9 10 11 12 13  0  1  2  3  4  5
## [553]  6  7  8  9 10 11 12 13  0  1  2  3  4  5  6  7  8  9 10 11 12 13  0
## [576]  1  2  3  4  5  6  7  8  9 10 11 12 13  0  1  2  3  4  5  6  7  8  9
## [599] 10 11 12 13  0  1  2  3  4  5  6  7  8  9 10 11 12 13  0  1  2  3  4
## [622]  5  6  7  8  9 10 11 12 13  0  1  2  3  4  5  6  7  8  9 10 11 12 13
## [645]  0  1  2  3  4  5  6  7  8  9 10 11 12 13  0  1  2  3  4  5  6  7  8
## [668]  9 10 11 12 13  0  1  2  3  4  5  6  7  8  9 10 11 12 13  0  1  2  3
## [691]  4  5  6  7  8  9 10 11 12 13  0 14 15 16 17 18 19 20 21 22 23 24 25
## [714] 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
## [737] 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63
##
## $nsunits
## [1] 64
##
## $decay
## [1] 0

```

```

pred = predict(fit50, newdata = test_data, type = "raw")
sqrt(mean((pred[, 1] - test_targets)^2))

```

```
## [1] 5.063217
```

```

mae = mean(abs(pred[, 1] - test_targets))
mae

```

```
## [1] 3.76992
```

See here: https://www.math.ntnu.no/emner/TMA4268/2018v/11NN/11-neural_networks_boston_housing.html for running two hidden layers of 64 nodes each with `keras`. (This gave a mean absolute error on the test set of 2.682524 after early stopping.)

Using `nnet`

- only one hidden layer available
- only sigmoid hidden layer activation function
- optimization done with an optimization method where the learning rate is automatically calculated
- weight decay (L_2 regularization as for the ridge regression) is available (see below)
- optimization is supposed to be done until convergence
- the possible choice of hyper parameter is the number of nodes in the hidden layer

We suggest you use `nnet` in **Compulsory exercise 2, Problem 3**.

Neural network parts

We now focus on the different elements of neural networks.

Illustration drawn in class, motivated from Figure 1.9/3.1 from Chollet and Allaire (2018).

Model

Output layer activation

These choices have been guided by solutions in statistics (multiple linear regression, logistic regression, multiclass regression)


- Linear: for regression problems
- Sigmoid: for two-class classification problems
- Softmax: for more than two classes classification problems

Remark: it is important that the output activation is matched with an appropriate loss function.

Hidden layer activation

Q: Why can we not just use linear activation function in all hidden layers?

A: Then each layer would only be able to do linear transformations of the input data and a deep stack of linear layers would still implement a linear operation. The activation functions sigmoid and relu add non-linearity to the model. And, the universal approximation property is dependent on a squashing type activation function.



- Sigmoid: $g(a) = 1/(1 + \exp(-a))$. The previous standard choice. These units were found to saturate at a high value when the input was very positive and at a low value when the input was very negative, and that they were only strongly sensitive when the input was close to 0. This saturation lead to problems in the gradient descent routines.

However, for our Compulsory exercise 2, Problem 3, there should be quite ok to use sigmoid hidden layer activation in the `nnet` R package (default hidden layer activation is sigmoid, and that can not be changed).

-
- ReLU: $\phi_h(a) = \max(0, a)$. The standard choice for deep networks (many hidden layers and many nodes) today.

The function is piecewise linear, but in total non-linear. It has shown to be easy to use with gradient descent — even though the function is not differentiable at 0. As we will touch upon later, we don't expect to train a network until the gradient is 0. The derivative from the left at 0 is 0, and the derivative from the right is 1. See Goodfellow, Bengio, and Courville (2016) page 192 for a discussion on this topic.

When initializing the parameter going into a ReLU node, the advice is to set weight of the bias node to be small and positive, e.g. 0.1, which makes it likely that node will be active (larger than 0) for most inputs in the training set.

Goodfellow, Bengio, and Courville (2016), page 226, reports this (replacing sigmoid with ReLU) to be one of the major changes that have improved the performance of the feedforward networks.

ReLU can also be motivated from biology.


- For some inputs a biological neuron can be completely inactive
- For some inputs a biological neuron output can be proportional to the input
- But, most of the time a biological neuron is inactive.

According to Goodfellow, Bengio, and Courville (2016), page 197, hidden unit design is an *active area of research*.

Other possible choices for hidden layer activation functions are:

- Radial basis functions: as we looked at in Module 9.
 - Softplus: $\phi_h(a) = \ln(1 + \exp(a))$
 - Hard tanh: $\phi_h(a) = \max(-1, \min(1, a))$
-

Neural motivation for squashing functions



https://commons.wikimedia.org/wiki/File:Action_potential.svg

Network architecture

How many nodes in the network and how are the nodes connected to each other? This depends on the problem, and here experience is important.

We will only consider feedforward networks, where all nodes in one layer is connected to all the nodes in the

next layer. The layers are then *fully connected* and *dense*.

Choosing architecture to us is to chose the *depth of the network* (*how many hidden layers*) and the *width of each layer*.

However, the recent practice, see e.g. Chollet and Allaire (2018), Section 4.5.6 and Goodfellow, Bengio, and Courville (2016), page 229, is to

- choose a too large network (too many nodes and/or too many layers) so that if trained until convergence (optimum) then the this would result in overfitting, and
- then use other means to avoid this (various variants of regularization and hyperparameter optimization).

This makes the choice of network architecture to be to *choose a large enough network*. See also Hyperparameter optimization below.

Method

The method part is to choose the loss function for the output layer.

The choice of loss function is closely related to the output layer activation function. To sum up, the popular problem types, output activation and loss functions are:

Problem	Output activation	Loss function
Regression	linear	mse
Classification (C=2)	sigmoid	binary_crossentropy
Classification (C>2)	softmax	categorical_crossentropy

where the mathematical formulas are given for n training samples.

Let all network parameters (weights) be denoted by θ .

Mean squared error

$$J(\theta) = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_1(\mathbf{x}_i))^2$$

$\hat{y}_1(\mathbf{x}_i)$ is the output from the linear output node, and y_i is the response.

Binary cross-entropy

$$J(\theta) = -\frac{1}{n} \sum_{i=1}^n (y_i \ln(\hat{y}_1(\mathbf{x}_i)) + (1 - y_i) \ln(1 - \hat{y}_1(\mathbf{x}_i)))$$

where $\hat{y}_1(\mathbf{x}_i)$ is the output from the sigmoid output node, and y_i is the 0/1 observed class.

Categorical cross-entropy

$$J(\boldsymbol{\theta}) = -\frac{1}{n} \sum_{i=1}^n \frac{1}{C} \sum_{c=1}^C (y_{ic} \ln(\hat{y}_c(\mathbf{x}_i)))$$

where $\hat{y}_c(\mathbf{x}_i)$ is the output from the softmax output node c and y_{ic} is the observed one-hot coding (0/1) for class c .

Overall

Due to how estimation is done (see below), the loss functions chosen “need” to be:

- differentiable
- possible to compute for each single training data point (or a mini-batch — to be explained soon)

In the 1980-1990s the mean squared error was the prominent loss function also for classification problems, but this has subsequently changed — motivated by the spread of maximum likelihood from statistics to machine learning.

Observe that we have only given the formula for the loss function, and not explicitly assumed anything about any probability distribution of the responses (not even assumed that the responses are random variables). However, we know which statistical model assumptions would give the loss functions as related to the negative of the loglikelihood.

Optimizers

Let the unknown parameters be denoted $\boldsymbol{\theta}$ (what we have previously denotes as α s and β s), and the loss function to be minimized $J(\boldsymbol{\theta})$.

Illustration: Figure 1.9 from Chollet and Allaire (2018) (again).

Gradient descent

Let $\nabla J(\boldsymbol{\theta}^{(t)})$ be the gradient of the loss function evaluated at the current estimate $\boldsymbol{\theta}^{(t)}$, then a step $t + 1$ of the algorithm the new values (estimates) of of the parameters are given as:

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \lambda \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}^{(t)})$$

The *learning rate* λ is often set to some small value. In `keras` the default learning rate is 0.01.

Remember that the gradient is the vector of partial derivative of the loss function with respect to each of the parameter in the network.

Q: Why are we moving in the direction of the negative of the gradient? Why not the positive?

Mini-batch stochastic gradient descent (SGD)

The loss function is computed as a mean over all training examples.

$$J(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n J(\mathbf{x}_i, y_i)$$

This means that the gradient will also be a mean over the gradient contribution from each training example.

$$\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \nabla_{\boldsymbol{\theta}} J(\mathbf{x}_i, y_i)$$

To build a network that generalizes well, it is important to have many training examples, but that would make us spend a lot of time and computer resources at calculating each gradient descent step.

We observe that we may see the gradient as an average over many individual gradients, and think of this as an estimator for an expectation. This expectation can we (also) approximate by the average gradient over just a *mini-batch* (random sample) of the observations.

The idea here is that the optimizer will converge much faster if they can rapidly compute approximate estimates of the gradient, instead of slowly computing the exact gradient (using all training data).

In addition with multicore systems, mini-batches may be processed in parallel and the batch size is often a power of 2 (32 or 256).

It also turns out that small batches also serves as a regularization effect maybe due to the variability they bring to the optimization process.

In the 4th video (on backpropagation) from 3Blue1Brown there is nice example of one trajectory from gradient decent and one from SGD (13:50 minutes into the video): <https://www.youtube.com/watch?v=tIeHLnjs5U8>

This means that for (mini-batch) stochastic gradient descent we do as follows:

1. Divide all the training samples randomly into mini-batches.
2. For each mini-batch: Make predictions of the reponses in the mini-batch in a *forward pass*.
3. Compute the loss for the training data in this batch.
4. Compute the gradient of the loss with regard to the model's parameters (*backward pass*) based on the training data in the batch. $\nabla_{\boldsymbol{\theta}}^* J(\boldsymbol{\theta}^{(t)})$
5. Update all weighs, but just using the average gradient from the mini-batch $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \lambda \nabla_{\boldsymbol{\theta}}^* J(\boldsymbol{\theta}^{(t)})$
6. Repeat 2-5 until convergence. (Could have gone back to 1, but often not done.)

-
- The algorithm defined above is called mini-batch SGD. The Stochastic part comes from the fact that we are randomly sampling batches from the training data.
 - Stochastic gradient descent (SGD) for size equals to 1.
 - Mini-batch SGD is a compromise between SGD (one sample per iteration) and full gradient descent (full dataset per iteration)
-

Backpropagation algorithm

Computing the analytical expression for the gradient ∇J is not difficult, but the numerical evaluation may be expensive. The Backpropagation algorithm is a simple and inexpensive way to calculate the gradient.

The chain rule is used to compute derivatives of functions of other functions where the derivatives are known, this is done efficiently with backpropagation.

Backpropagation starts with the value of the loss function and works backward from the top layers to the bottom layers, applying the chain rule to compute the contribution that each parameter has in the loss value.

More information:

- Mathematical details in Goodfellow, Bengio, and Courville (2016) Section 6.5 (pages 204-238),
 - 3Blue1Brown: video overview: <https://www.youtube.com/watch?v=Ilg3gGewQ5U> and chain rule maths <https://www.youtube.com/watch?v=tIeHLnjs5U8>
-

Variations of SDG — with adaptive learning rates

General concept:

- momentum term: previous gradients are allowed to contribute.

Named variants: In `keras` the “default optimizer” is RMSprop.

- AdaGrad: individually adapt the learning rates of all model parameters by scaling them inversely proportional to the square root of the sum of all their historical squared values. (Nice properties for convex problems, but with non-linear hidden activation function we do not have a convex problem.)
 - RMSprop: modification to AdaGrad in non-convex setting. Scales with exponentially weighted moving average instead of all historical squared gradient values.
-

Further reading on optimizers

- Keras documentation for Optimizers
 - An overview of gradient descent optimization algorithms
 - Overview of mini-batch gradient descent
 - Andrew Ng explains about RMSprop in Coursera course: Improving Deep Neural Networks: Hyperparameter tuning, Regularization and Optimization
-

Regularization

Goodfellow, Bengio, and Courville (2016), Chapter 7, define regularization as *any modification we make to a learning algorithm that is intended to reduce its generalization error but not its training error*.

We looked at regularization in Module 6, where our aim was to trade *increased bias* for *reduced variance*. Another way of looking at this is we need not focus entirely on finding the model of “correct size”, but instead find a large model that has been regularized properly.

In Module 6 we looked in particular at adding a penalty to the loss function. The penalties we looked at were of type absolute value of parameter (L_1 , lasso, where we looked at this as model selection) and square value of parameter (L_2 , ridge regression). This can also be done for neural networks.

In neural networks, *weight decay* is the expression for adding a L_2 -penalty to the loss function, and is available in the `nnet` R package.

Other versions of regularization are *dataset augmentation* and *label smoothing*:

- Dataset augmentation means adding fake data to the dataset, in order that the trained model will generalize better. For some learning task it is straightforward to create the new fake data — for image data this can be done by rotating and scaling the images.
- Label smoothing is motivated by the fact that the training data may contain errors in the responses recorded, and replaced the one-hot coding for C classes with $\epsilon/(C - 1)$ and $1 - \epsilon$ for some small ϵ .

Early stopping

(Based on Goodfellow, Bengio, and Courville (2016), Section 7.8)

The most commonly used form of regularization is *early stopping*.

If we have chosen a sufficiently large model with the capacity to overfit the training data, we would observe that the training error decreases steadily during training, but the error on the validation set at some point begins to increase.

If we stop the learning early and return the parameters giving the test performance on the validation set, this model would hopefully be a better model than if we trained the model until convergence — and this model will then also give a better test set error.

It is possible to think of the number of *training steps* as a hyperparameter. This hyperparameter can easily be set, and the cost is just that we need to monitor the performance on the validation set during training. Alternatively, cross-validation can be used.

One strategy is to first find the optimal stopping time for the training based on the validation set (or with cross-validation with small data sets), and then retrain the full training set (including the validation part) and stop at the selected stopping time.

Why is early stopping a regularization technique? By early stopping the optimization procedure has only seen a relatively small part of the parameter space in the neighbourhood of the initial parameter value. See Goodfellow, Bengio, and Courville (2016), page 250 for a relationship with L_2 regularization.

Hyperparameter optimization

How to avoid overfitting:

- reduce network size,
- collect more observations,
- regularization (including early stopping and drop-out)

The network architecture, the number of batches to run before terminating the optimization, the drop-out rate, are all examples of hyperparameters that need to be chosen in a sensible way before fitting the final model.

It is important that the hyperparameters are chosen on a validation set or by cross-validation.

However, a “popular” term is *validation-set overfitting* and refers to using the validation set to decide many hyperparameters, so many that you may effectively overfit the validation set.

In statistics we use design of experiments to explore these hyperparameters, and just using marginal grids (one hyperparameter at a time) is common in machine learning.

Example on DOE:hyperparameter optimization with boosting (which of course also can be used for neural networks). Article: Design of experiments and response surface methodology to tune machine learning hyperparameters, with a random forest case-study (2018), Gustavo A. Lujan-Moreno, Phillip R. Howard, Omar G. Rojas, Douglas Montgomery, Expert Systems with Applications, Volume 109, <https://doi.org/10.1016/j.eswa.2018.05.024>

Dropout

(Based on Goodfellow, Bengio, and Courville (2016), Section 7.12, and Chollet and Allaire (2018) 4.4.3)

Dropout was developed by Geoff Hinton and his students.

- During training: randomly *dropout* (set to zero) outputs in a layer. Drop-out rates may be chosen between 0.2 and 0.5.
- During test: not dropout, but scale down the layer output values by a factor equal to the drop-out rate (since now more units are active than we had during training)

Alternatively, the drop-out and scaling (now upscaling) can be done during training.

One way to look at dropout is on the lines of what we did in Module 8 when we used bootstrapping to produce many data sets and then fitted a model to each of them and then took the average (bagging). But randomly dropping out outputs in a layer, this can be looked at as mimicking bagging — in an efficient way.

See Goodfellow, Bengio, and Courville (2016), Section 7.12 for more insight into the mathematics behind drop-out.

https://www.reddit.com/r/MachineLearning/comments/4w6tstv/ama_we_are_the_google_brain_team_wed_love_to/

The following is a direct quotation.

figplucker: **How was ‘Dropout’ conceived? Was there an ‘aha’ moment?**

geoffhinton (2 years ago)

There were actually three aha moments. One was in about 2004 when Radford Neal suggested to me that the brain might be big because it was learning a large ensemble of models. I thought this would be a very inefficient use of hardware since the same features would need to be invented separately by different models. Then I realized that the “models” could just be the subset of active neurons. This would allow combinatorially many models and might explain why randomness in spiking was helpful.

Soon after that I went to my bank. The tellers kept changing and I asked one of them why. He said he didn’t know but they got moved around a lot. I figured it must be because it would require cooperation between employees to successfully defraud the bank. **This made me realize that randomly removing a different subset of neurons on each example would prevent conspiracies and thus reduce overfitting.**

I tried this out rather sloppily (I didn't have an adviser) in 2004 and it didn't seem to work any better than keeping the squared weights small so I forgot about it.

Then in 2011, Christos Papadimitriou gave a talk at Toronto in which he said that the whole point of sexual reproduction was to break up complex co-adaptations. He may not have said it quite like that, but that's what I heard. It was clearly the same abstract idea as randomly removing subsets of the neurons. So I went back and tried harder and in collaboration with my grad students we showed that it worked really well.

Metrics

Provides different forms to measure how well the predictions are compared with the true values.

- **accuracy**: Percentage of correct classifications
- **mae**: Mean absolute error.

These metrics can be monitored during training (on validation set) or in the end on the test set, and can be the basis for the choice when to top training.

Deep learning

Timeline

(based on Chollet and Allaire (2018))

Neural networks were investigated in “toy form” in the 1950s. The first big step was taken in the 1980s when the backpropagation algorithm were developed (rediscovered) to perform gradient descent in an efficient way.

In 1989 (Bell Labs, Yann LeCun) used convolutional neural networks to classifying handwritten digits, and *LeNet* was used in the US Postal Service for reading ZIP codes in the 1990s.

Not so much seen (?) activity in the neural network field in the 2000s.

In 2011 neural networks with many layers (and trained with GPUs) were performing well on image classification tasks. The *ImageNet* classification challenge (classify high resolution colour images into 1k different categories after training on 1.4M images) was won by solutions with deep convolutional neural networks (convnets). In 2011 the accuracy was 74.3%, in 2012 83.6% and in 2015 96.4%.

From 2012, convnets is the general solution for computer vision tasks. Other application areas are natural language processing.

Deep?

Deep learning does not mean a deeper understanding, but refers to successive layers of representations - where the number of layers gives the *depth* of the model. Often tens to hundreds of layers are used.

Deep neural networks are not seen as models of the brain, and are not related to neurobiology.

A deep network can be seen as many stages of *information-destillation* (Chollet and Allaire (2018), page 9), where each stage performs a simple data transformation. These transformations are not curated by the data analyst, but is estimated in the network.

In statistics we first select a set of inputs, then look at how these inputs should be transformed (projections in simple form or high-dimensional and nonlinear forms), before we apply some statistical methods. This transformation step can be called *feature engineering* and has been automated in deep learning.

Deep Learning is an algorithm which has no theoretical limitations of what it can learn; the more data you give and the more computational time you provide, the better it is.

Geoffrey Hinton (Google)

In addition, this built-in feature engineering of the deep network is not performed in a greedy fashion, but *jointly* with estimating/learning the full model.

The success of deep learning is dependent upon the breakthroughs in hardware development, especially with faster CPUs and massively parallel graphical processing units (GPUs). Tensor processing units (TPUs) is the next step.

Achievements of deep learning includes high quality (near-human to super human) image classification, speech recognition, handwriting transcription, machine translation, digital assistants, autonomous driving, advertise targeting, web searches, playing Go and chess.

The R keras package

Earlier good programming skills in C++ was essential to work in deep learning. In addition also skills on programming for GPUs were needed (e.g NVIDIA CUDA programming interface). With the launch of the Keras library now users may only need basic skills in Python or R.

Keras can be seen as a way to use LEGO bricks in deep learning. To quote the web-page:

Keras is a high-level neural networks API, written in Python and capable of running on top of TensorFlow, CNTK, or Theano. It was developed with a focus on enabling fast experimentation. Being able to go from idea to result with the least possible delay is key to doing good research.

Tensorflow is a *symbolic-tensor manipulation* framework that also can perform autodifferentiation.

A tensor is defined by its number of axis (below), shape (dimension) and data type (double, integer, character)

- 0D tensor is a scalar
- 1D tensor is a vector
- 2D tensor is a matrix - and the two *axis* of the tensor is the rows and columns
- 3D tensor generalization of a matrix, and may be used for time series data
- 4D tensors may be used for images (samples, height, width, channels), and
- 5D tensors may be used for video (as 4D, plus frames).

Tensor operations (reshaping, dot product) are performed in TensorFlow.

The use of Tensorflow in R Keras is referred to as “using Tensorflow as a backend engine”. Tensorflow is the default backend.

More information on the R solution: <https://keras.rstudio.com/>

Cheat-sheet for R Keras: <https://github.com/rstudio/cheatsheets/raw/master/keras.pdf>

Simple data analysis workflow

1. Select training and test data
 2. Define the model(s): layers, nodes in layers, activation function
 3. Configure the learning process by choosing a loss function, an optimizer, and some metrics to monitor.
If needed decide on an evaluation protocol (validation set, cross-validation, iterated cross-validation)
 4. Perform any needed preprocessing of data to fit the choices in 2 and 3
 5. Fit the model to the data
 6. Make predictions for test data or use the model
-

Chollet and Allaire (2018) Section 4.5 has the following recommendations for step 3+5:

- Develop a first model that performs better than a basic baseline (maybe the basic baseline could be standard statistics solutions)
- Develop a model that overfits the data
- Regularize and tune hyperparameters

Here is a tutorial: https://keras.rstudio.com/articles/tutorial_overfit_underfit.html



MNIST dataset

This is a larger image version of the handwritten digits data set (a different version, ZIP-codes is found under Recommended exercises).

This data analysis is based on https://www.math.ntnu.no/emner/TMA4268/2018v/11NN/8-neural_networks_mnist.html and the R `keras` cheat sheet. An advanced version using convolutional neural nets is found here: https://www.math.ntnu.no/emner/TMA4268/2018v/11NN/12-neural_networks_convolution_mnist.html

Objective: classify the digit contained in an image (128×128 greyscale).

Problem type: Multiclass classification based on image data.



Labels for the training data:

```
## train_labels
##  0  1  2  3  4  5  6  7  8  9
## 5923 6742 5958 6131 5842 5421 5918 6265 5851 5949
```

1. Training and test data

60 000 images for training and 10 000 images for testing.

```
# Training data
train_images <- mnist$train$x
train_labels <- mnist$train$y

# Test data
test_images <- mnist$test$x
test_labels <- mnist$test$y
org_testlabels <- test_labels
```

The `train_images` is a tensor (generalization of a matrix) with 3 axis, (`samples`, `height`, `width`).

2. Defining the model

In this case we are using a `layer_dense` (fully connected) which expects an input tensor of rank equal to two (`sample`, `features`) where each `sample` should contain $28 \times 28 = 784$ pixels. Adding a bias term (`intercept`) is default for `layer_dense`.

```
network <- keras_model_sequential() %>% layer_dense(units = 512, activation = "relu",
  input_shape = c(28 * 28)) %>% layer_dense(units = 10, activation = "softmax")
summary(network)
```

```
## -----
## Layer (type)                Output Shape          Param #
## =====
## dense_5 (Dense)             (None, 512)          401920
## -----
## dense_6 (Dense)             (None, 10)           5130
## =====
## Total params: 407,050
## Trainable params: 407,050
## Non-trainable params: 0
## -----
```

As activation function we use `relu` for the hidden layer, and `softmax` for the output layer - since we have 10 classes (where one is correct each time). We could have included more layers in our model - and then maybe used early stopping if our model was chosen too big.

3. Configure the learning process

We next need to choose the loss function we will use, which is cross-entropy, and then the version of the optimizer. Here it is RMSprop. Finally, which measure - metrics - do we want to monitor in our training phase? Here we choose accuracy (=percentage correctly classified).

```
network %>% compile(optimizer = "rmsprop", loss = "categorical_crossentropy",
  metrics = c("accuracy"))
```

4. Preprocessing to match with model inputs and outputs

The training data is scored in an array of dimension (60000, 28, 28) of type integer with values in the [0, 255] interval. The data must be reshaped into the shape the network expects (28*28). In addition the grey scale values are scales to be in the [0, 1] interval.


Also, the response must be transformed from 0-10 to a vector of 0s and 1s (dummy variable coding) aka one-hot-coding.

```
train_images <- array_reshape(train_images, c(60000, 28 * 28))
train_images <- train_images/255
train_labels <- to_categorical(train_labels)

test_images <- array_reshape(test_images, c(10000, 28 * 28))
test_images <- test_images/255
test_labels <- to_categorical(test_labels)
```

5. Fit the model

```
fitted<-network %>% fit(train_images, train_labels,  
  epochs = 30, batch_size = 128)  
library(ggplot2)  
plot(fitted)+ggtitle("Fitted model")
```



6. Evaluation and prediction

```
network %>% evaluate(test_images, test_labels)  
testres = network %>% predict_classes(test_images)  
# $loss [1] 0.1194063 $acc [1] 0.9827  
confusionMatrix(factor(testres), factor(org_testlabels))
```

Confusion Matrix and Statistics

	Reference									
Prediction	0	1	2	3	4	5	6	7	8	9
0	971	0	3	0	1	2	5	1	1	1
1	1	1128	2	0	0	0	2	2	2	3
2	1	1	1009	3	3	0	2	7	2	0
3	0	1	2	997	0	11	1	3	4	3
4	1	0	2	0	969	1	4	2	3	7
5	0	1	0	1	0	871	3	0	1	4
6	3	2	2	0	3	4	940	0	1	0
7	1	1	4	2	1	0	0	1002	2	3
8	2	1	7	0	0	2	1	4	953	1
9	0	0	1	7	5	1	0	7	5	987

Overall Statistics

Accuracy : 0.9827
 95% CI : (0.9799, 0.9852)
 No Information Rate : 0.1135
 P-Value [Acc > NIR] : < 2.2e-16

Kappa : 0.9808
 McNemar's Test P-Value : NA

Kaggle

Kaggle has hosted machine-learning competitions since 2010, and by looking at solutions to competitions it is possible to get an overview of what works. In 2016-2017 gradient boosting methods won the competitions with structured data (“shallow” learning problems), while deeplearning won perceptual problems (as image classification), Chollet and Allaire (2018) (page 18). Kaggle has helped (and is helping) the rise in deep learning.

Other analyses

Boston housing price

taken from Chollet and Allaire (2018): https://www.math.ntnu.no/emner/TMA4268/2018v/11NN/11-neural_networks_boston_housing.html

Movie data base

https://www.math.ntnu.no/emner/TMA4268/2018v/11NN/9-neural_networks_imdb.html

Reuters data

https://www.math.ntnu.no/emner/TMA4268/2018v/11NN/10-neural_networks_reuters.html

Summing up

- Feedforward network architecture: mathematical formula - layers of multivariate transformed (`relu`, `linear`, `sigmoid`) inner products - sequentially connected.
 - What is the number of parameters that need to be estimated? Intercept term (for each layer) is possible and is referred to as “bias term”.
 - Loss function to minimize (on output layer): regression (mean squared), classification binary (binary crossentropy), classification multiple classes (categorical crossentropy) — and remember to connect to the correct choice of output activation function: mean squared loss goes with linear activation, binary crossentropy with sigmoid, categorical crossentropy with softmax.
 - How to minimize the loss function: gradient based (chain rule) back-propagation - many variants.
 - Technicalities: `nnet` in R
 - Optional: `keras` in R. Use of tensors. Piping sequential layers, piping to estimation and then to evaluation (metrics).
-

Not covered

Sadly.

(The first two topics are covered in Chollet and Allaire (2018))

Recurrent networks: extending the feedforward network to also have feedback connections. This is a popular type of network to analyse time series data and natural language applications.

Convolutional networks: some layers in a sequential network contain operations specially suitable for grid-like topology (images). Convolution is used in place of general layer (where we do matrix multiplication) in at least one layer. A popular operation is *pooling*.


Explainable AI (XAI): how to use methods on the network or the predictions of the network to figure out the underlying reasoning of the network. Popular methods are called LIME (local linear regression), Shaply (concept from game theory). The DALEX R package contains different so-called *explainers* <https://arxiv.org/abs/1806.08915>.

Recommended exercises

Theoretical

a)


Write down the equation describing this network. What would you call such a network?



b)

The following image is the illustration of an artificial neural network at Wikipedia.

- What can you say about this network architecture
- What do you think it can be used for (regression/classification)?



https://commons.wikimedia.org/wiki/File:Colored_neural_network.svg

c)

Given the following problems, what are sensible feedforward network architectures (depth, width, activation function) and methods (loss function, algorithms) that you would explore?

- Regression with one univariate response, 10 possible covariates, 500 observations.
- Classification with two classes, 100 possible covariates, 10000 observations.
- Classification with 10 classes, image data (like the MNIST), 50000 observations.

d)

What are the similarities and differences between a feedforward neural network with one hidden layer with **linear** activation and **sigmoid** output (one output) and logistic regression?

e)

In a feedforward neural network you may have 10000 weights to estimate but only 1000 observations. How is this possible?

f)

Which network architecture and activation functions does this formula give?

$$\hat{y}_1(\mathbf{x}) = \beta_{01} + \sum_{m=1}^5 \beta_{m1} \cdot \max(\alpha_{0m} + \sum_{j=1}^{10} \alpha_{jm} x_j, 0)$$

How many parameters are estimated in this network?

g)

Which network architecture and activation functions does this formula give?

$$\hat{y}_1(\mathbf{x}) = (1 + \exp(-\beta_{01} - \sum_{m=1}^5 \beta_{m1} \max(\gamma_{0m} + \sum_{l=1}^{10} \gamma_{lm} \max(\sum_{j=1}^4 \alpha_{jl} x_j, 0), 0)))^{-1}$$

How many parameters are estimated in this network?

h)

In a regression setting: Consider

- a sum of non-linear functions of each covariate in Module 7
- a sum of many non-linear functions of sums of covariates in feedforward neural networks (one hidden layer, non-linear activation in hidden layer) in Module 11.

Explain how these two ways of thinking differ? Pros and cons?

i)

What is the most interesting aspect of neural network (your opinion)? How would you compare how feedforward neural networks are fitted as compared to fitting multiple linear regression and logistic regression? Compare how model selection (which covariates) are performed.

Practical

You should be able to

- Fit regression
- Fit classification with two classes
- Fit classification with more than two classes

all with the `nnet` package. Below is an example for classification with more than two classes, and a small part with `keras` is added for completeness.

Handwritten digit recognition data

See Friedman, Hastie, and Tibshirani (2001) Section 11.7.

Description:

(direct quote from `?zip.train` in `ElemStatLearn` R package)

Normalized handwritten digits, automatically scanned from envelopes by the U.S. Postal Service. The original scanned digits are binary and of different sizes and orientations; the images here have been deslanted and size normalized, resulting in 16 x 16 grayscale images.


The data are in two gzipped files, and each line consists of the digit id (0-9) followed by the 256 grayscale values.

There are 7291 training observations and 2007 test observations.

The test set is notoriously “difficult”, and a 2.5 excellent. These data were kindly made available by the neural network group at AT&T research labs (thanks to Yann Le Cunn).

A selection of some images

```
library(ElemStatLearn)
# code from help(zip.train)
findRows <- function(zip, n) {
  # Find n (random) rows with zip representing 0,1,2,...,9
  res <- vector(length = 10, mode = "list")
  names(res) <- 0:9
  ind <- zip[, 1]
  for (j in 0:9) {
    res[[j + 1]] <- sample(which(ind == j), n)
  }
  return(res)
}
digits <- vector(length = 10, mode = "list")
names(digits) <- 0:9
rows <- findRows(zip.train, 6)
for (j in 0:9) {
  digits[[j + 1]] <- do.call("cbind", lapply(as.list(rows[[j + 1]]),
    function(x) zip2image(zip.train, x)))
}
im <- do.call("rbind", digits)
image(im, col = gray(256:0/256), zlim = c(0, 1), xlab = "", ylab = "")
```



Preprocessing data - scaling

Important to decide on the scale based on the training data, and then apply to both training and test data.

```
train_data = zip.train[, -1]
train_labels = factor(zip.train[, 1])
test_data = zip.test[, -1]
test_labels = factor(zip.test[, 1])
mean <- apply(train_data, 2, mean)
std <- apply(train_data, 2, sd)
train_data <- scale(train_data, center = mean, scale = std)
test_data <- scale(test_data, center = mean, scale = std)
```

Check out 5 hidden nodes

5 hidden nodes: $257 * 5 + 6 * 10 = 257 * 5 + 6 * 10$ parameters

```
library(nnet)
zipnnet5 <- nnet(train_labels ~ ., data = train_data, size = 5, MaxNWts = 3000,
  maxit = 5000)
summary(zipnnet5)
pred = predict(zipnnet5, newdata = test_data, type = "class")
library(caret)
confusionMatrix(factor(pred), test_labels)
```

The above took some time to run, the results were:

```
> zipnnet5<- nnet(train_labels~., data=train_data,size=5,MaxNWts=3000,maxit=5000)
iter2960 value 864.566658
final value 864.561810
converged
```

```

> summary(zipnet5)
a 256-5-10 network with 1345 weights
options were - softmax modelling
  b->h1  i1->h1  i2->h1  i3->h1  i4->h1  i5->h1  i6->h1  i7->h1  i8->h1  i9->h1  i10->h1  i11->h1
-49.27   9.15   1.24  21.03  -2.82  17.97   4.63  11.60  -4.31   2.28  -4.57  -0.00
  i19->h1 i20->h1 i21->h1 i22->h1 i23->h1 i24->h1 i25->h1 i26->h1 i27->h1 i28->h1 i29->h1 i30->h1
> confusionMatrix(factor(pred),test_labels)
Confusion Matrix and Statistics

          Reference
Prediction 0  1  2  3  4  5  6  7  8  9
0 324  0  6  5  3  9  7  0  3  0
1  1 245  7  0  1  0  0  0  7  5
2  4  7 148  8 12  5 11  0  3  0
3  2  0  6 128  4 10  0  4  5  3
4  4  1  4  0 152  1  1  9  2  4
5  1  0  2 18  1 117  5  0  7  1
6 21  3 11  0  6  1 146  0  3  0
7  0  1  6  4  5  1  0 122  9  5
8  2  2  7  3  6 15  0  1 113  3
9  0  5  1  0 10  1  0 11  14 156

Overall Statistics

          Accuracy : 0.8226
          95% CI   : (0.8052, 0.8391)

```

Slightly better and faster with keras, but we see that we probably need more hidden layers since this is a difficult problem! An introduction to keras is given later in this module.

```

library(keras)
network <- keras_model_sequential() %>% layer_dense(units = 5, activation = "sigmoid",
  input_shape = c(16 * 16)) %>% layer_dense(units = 10, activation = "softmax")
network %>% compile(optimizer = "rmsprop", loss = "categorical_crossentropy",
  metrics = c("accuracy"))
train_data <- array_reshape(train_data, c(7291, 16 * 16))
train_data <- train_data/255
train_labels <- to_categorical(train_labels)

test_data <- array_reshape(test_data, c(2007, 16 * 16))
test_data <- test_data/255
org_test_labels <- test_labels
test_labels <- to_categorical(test_labels)

fitted <- network %>% fit(train_data, train_labels, epochs = 500, batch_size = 128,
  validation_split = 0.2)
library(ggplot2)
plot(fitted) + ggtitle("Fitted model") #same as above

network %>% evaluate(test_data, test_labels)
# acc : num 0.873
library(caret)
res <- network %>% predict_classes(test_data)

```

```
confusionMatrix(factor(as.character(res)), org_test_labels)
```

10 hidden nodes

10 hidden nodes: $257 * 10 + 11 * 10 = 257 * 10 + 11 * 10$ parameters

```
library(ElemStatLearn)
train_data = zip.train[, -1]
train_labels = factor(zip.train[, 1])
test_data = zip.test[, -1]
test_labels = factor(zip.test[, 1])
mean <- apply(train_data, 2, mean)
std <- apply(train_data, 2, sd)
train_data <- scale(train_data, center = mean, scale = std)
test_data <- scale(test_data, center = mean, scale = std)
library(keras)
network <- keras_model_sequential() %>% layer_dense(units = 10, activation = "sigmoid",
  input_shape = c(16 * 16)) %>% layer_dense(units = 10, activation = "softmax")
network %>% compile(optimizer = "rmsprop", loss = "categorical_crossentropy",
  metrics = c("accuracy"))
train_data <- array_reshape(train_data, c(7291, 16 * 16))
train_data <- train_data/255
train_labels <- to_categorical(train_labels)

test_data <- array_reshape(test_data, c(2007, 16 * 16))
test_data <- test_data/255
org_test_labels <- test_labels
test_labels <- to_categorical(test_labels)

fitted <- network %>% fit(train_data, train_labels, epochs = 500, batch_size = 128,
  validation_split = 0.2)
allfitted <- network %>% fit(train_data, train_labels, epochs = 500,
  batch_size = 128)
network %>% evaluate(test_data, test_labels)
# acc : num 0.9133034
```

Exam problems

V2018 Problem 4 Classification of diabetes cases

<https://www.math.ntnu.no/emner/TMA4268/2018v/Compulsory3.html>

R packages

These packages needs to be install before knitting this R Markdown file.

```
install.packages("gamlss.data") #munich daata
install.packages("MASS") #Boston data
install.packages("nnet") # running feedforward nets with one hidden sigmoid layer
```

```
install.packages("NeuralNetTools") #plotting a net
install.packages("keras") #MNIST - and running deep nns
# for keras you also need to do install_keras() and dependencies to
# Tensorflow will be set - and you need then python installed
install.packages("knitr") # knitting the R Markdown file
install.packages("ggplot2") # plotting
install.packages("caret") #confusion matrices
```

References and further reading

- <https://youtu.be/aircAruvnKk> from 3BLUE1BROWN - 4 videos - using the MNIST-data set as the running example
- Look at how the hidden layer behave: <https://playground.tensorflow.org>
- Friedman, Hastie, and Tibshirani (2001), Chapter 11: Neural Networks
- Efron and Hastie (2016), Chapter 18: Neural Networks and Deep Learning
- Chollet and Allaire (2018)
- Goodfellow, Bengio, and Courville (2016) (to be used in IT3030)
- Explaining backpropagation <http://neuralnetworksanddeeplearning.com/chap2.html>
- Slides from MA8701 (Thiago Martins) <https://www.math.ntnu.no/emner/MA8701/2019v/DeepLearning/>

Acknowledgements

Chollet, François, and J. J. Allaire. 2018. *Deep Learning with R*. Manning Press.

Efron, Bradley, and Trevor Hastie. 2016. *Computer Age Statistical Inference - Algorithms, Evidence, and Data Science*. Cambridge University Press.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2001. *The Elements of Statistical Learning*. Vol. 1. Springer series in statistics New York.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. Vol. 112. Springer.

Ripley, Brian D. 1996. *Pattern Recognition and Neural Networks*. Cambridge University Press.

Venables, W. N., and B. D. Ripley. 2002. *Modern Applied Statistics with S*. Springer.