# Neural networks 

## Victor Kitov

v.v.kitov@yandex.ru

## Table of Contents

(1) Introduction
(2) Output generation
(3) Weight space symmetries

4 Neural network optimization
(5) Backpropagation algorithm
(6) Invariances
(7) Case study: ZIP codes recognition

## History

- Neural networks originally appeared as an attempt to model human brain

- Human brain consists of multiple interconnected neuron cells
- cerebral cortex (the largest part) is estimated to contain 15-33 billion neurons
- communication is performed by sending electrical and electro-chemical signals
- signals are transmitted through axons - long thin parts of neurons.


## Simple model of a neuron



- Neuron get's activated in the half-space, defined by $w_{0}+w_{1} x^{1}+w_{2} x^{2}+\ldots+w_{D} x^{D} \geq 0$.
- Each node is called a neuron
- Each edge is associated a weight
- Constant feature 1 stands for bias


## Multilayer perceptron architecture ${ }^{1}$

- Hierarchically nested set of neurons.
- Each node has its own weights.


This is structure of multilayer perceptron - acyclic directed graph.

[^0]
## Layers



- Structure of neural network:
- 1-input layer
- 2-hidden layers
- 3-output layer


## Continious activations

- Pitfall of $\mathbb{I}[]$ : it causes stepwise constant outputs, weight optimization methods become inapliccable.
- We can replace $\mathbb{I}\left[w^{T} x+w_{0} \geq 0\right]$ with smooth activation $\varphi\left(w^{\top} x+w_{0}\right)$



## Typical activation functions

- sigmoidal: $\sigma(x)=\frac{1}{1+e^{-x}}$
- 1-layer neural network with sigmoidal activation is equivalent to logistic regression
- hyperbolic tangent: $\operatorname{tangh}(x)=\frac{e^{x}-e^{-x}}{e^{x}+e^{-x}}$

- ReLu: $\varphi(x)=[x]_{+}$.


## Activation functions

Activation functions are smooth approximations of step functions:


$\operatorname{tangh}(a x)$ limits to $-1 / 1$-step function as $a \rightarrow \infty$

## Definition details

- Label each neuron with integer $j$.
- Denote: $l_{j}$ - input to neuron $j, O_{j}$ - output of neuron $j$
- Output of neuron $j: O_{j}=\varphi\left(I_{j}\right)$.
- Input to neuron $j: I_{j}=\sum_{k \in \operatorname{inc}(j)} w_{k j} O_{k}+w_{0 j}$,
- $w_{0 j}$ is the bias term
- inc $(j)$ is a set of neurons with outging edges incoming to neuron $j$.
- further we will assume that at each layer there is a vertex with constant output $O_{\text {const }} \equiv 1$, so we can simplify notation

$$
I_{j}=\sum_{k \in \operatorname{inc}(j)} w_{k j} O_{k}
$$

## Table of Contents

## (1) Introduction

(2) Output generation
(3) Weight space symmetries

4 Neural network optimization
(5) Backpropagation algorithm
6) Invariances
(7) Case study: ZIP codes recognition

## Output generation

- Forward propagation is a process of successive calculations of neuron outputs for given features.



## Activations at output layer

- Regression: $\varphi(I)=$ I
- Classification:
- binary: $y \in\{+1,-1\}$

$$
\varphi(I)=p(y=+1 \mid x)=\frac{1}{1+e^{-l}}
$$

- multiclass: $y \in 1,2, \ldots C$

$$
\varphi\left(O_{1}, \ldots O_{C}\right)=p(y=j \mid x)=\frac{e^{O_{j}}}{\sum_{k=1}^{C} e^{O_{k}}}, j=1,2, \ldots C
$$

where $O_{1}, \ldots O_{C}$ are outputs of output layer.

## Generalizations

- each neuron $j$ may have custom non-linear transformation $\varphi_{j}$
- weights may be constrained:
- non-negative
- equal weights
- etc.
- layer skips are possible

- Not considered here: RBF-networks, recurrent networks.


## Number of layers selection

- Number of layers usually denotes all layers except input layer (hidden layers+output layer)
- We will consider only continuous activation functions.
- Classification:
- single layer network selects arbitrary half-spaces
- 2-layer network selects arbitrary convex polyhedron (by intersection of 1-layer outputs)
- therefore it can approximate arbitrary convex sets
- 3-layer network selects (by union of 2-layer outputs) arbitrary finite sets of polyhedra
- therefore it can approximate almost all sets with well defined volume (Borel measurable)


## Number of layers selection

- Regression
- single layer can approximate arbitrary linear function
- 2-layer network can model indicator function of arbitrary convex polyhedron
- 3-layer network can uniformly approximate arbitrary continuous function (as sum weighted sum of indicators convex polyhedra)


## Sufficient amount of layers

Any continuous function on a compact space can be uniformly approximated by 2 -layer neural network with linear output and wide range of activation functions (excluding polynomial).

- In practice often it is more convenient to use more layers with less total amount of neurons
- model becomes more interpretable and easy to fit.


## Neural network architecture selection

- Network architecture selection:
- increasing complexity (control by validation error)
- decresing complexity ("optimal brain damage")
- may be used for feature selection


## Table of Contents

## (1) Introduction

(2) Output generation
(3) Weight space symmetries

4 Neural network optimization
(5) Backpropagation algorithm
6) Invariances
(7) Case study: ZIP codes recognition

## Weight space symmetries

- Consider a neural network with 1 hidden layer
- with tangh(x) activation functions
- consisting of $M$ neurons

Input
layer

Hidden
layer

Output
layer

## Weight space symmetries

- The following transformations in weight space lead to neural networks with equivalent outputs:
- for any neuron in hidden layer: simultaneous change of sign of input and output weights
- $2^{M}$ possible equivalent transformations of such kind
- for any pair of neurons in the hidden layer: interchange of input weights between the neurons and simultaneous interchange of output weights
- this is equivalent to reordering of neurons in the hidden layer, so there are $M$ ! such orderings
- $2^{M} M$ ! equivalent transformations exist in total.
- For neural network with $K$ hidden layers, consisting of $M_{k}, k=1,2, \ldots K$ neurons each, we obtain $\prod_{k=1}^{K} 2^{M_{k}} M_{k}$ ! equivalent neural networks.
- In general case these are the only symmetries existing in the weights space.


## Table of Contents

## (1) Introduction

(2) Output generation

## (3) Weight space symmetries

4 Neural network optimization
(5) Backpropagation algorithm
(6) Invariances
(7) Case study: ZIP codes recognition

# Network optimization: regression 

- Single output:

$$
\frac{1}{N} \sum_{n=1}^{N}\left(\widehat{y}_{n}\left(x_{n}\right)-y_{n}\right)^{2} \rightarrow \min _{w}
$$

Network optimization: regression

- Single output:

$$
\frac{1}{N} \sum_{n=1}^{N}\left(\widehat{y}_{n}\left(x_{n}\right)-y_{n}\right)^{2} \rightarrow \min _{w}
$$

- K outputs

$$
\frac{1}{N K} \sum_{n=1}^{N} \sum_{k=1}^{K}\left(\widehat{y}_{n k}\left(x_{n}\right)-y_{n k}\right)^{2} \rightarrow \min _{w}
$$

Network optimization: classification

- Two classes $(y \in\{0,1\}, p=P(y=1))$ :

$$
\prod_{n=1}^{N} p\left(y_{n}=1 \mid x_{n}\right)^{y_{n}}\left[1-p\left(y_{n}=1 \mid x_{n}\right)\right]^{1-y_{n}} \rightarrow \max _{w}
$$

Network optimization: classification

- Two classes $(y \in\{0,1\}, p=P(y=1))$ :

$$
\prod_{n=1}^{N} p\left(y_{n}=1 \mid x_{n}\right)^{y_{n}}\left[1-p\left(y_{n}=1 \mid x_{n}\right)\right]^{1-y_{n}} \rightarrow \max _{w}
$$

- $C$ classes $\left(y_{n c}=\mathbb{I}\left\{y_{n}=c\right\}\right)$ :

$$
\prod_{n=1}^{N} \prod_{c=1}^{C} p\left(y_{n}=c \mid x_{n}\right)^{y_{n c}} \rightarrow \max _{w}
$$

## Network optimization: classification

- Two classes $(y \in\{0,1\}, p=P(y=1))$ :

$$
\prod_{n=1}^{N} p\left(y_{n}=1 \mid x_{n}\right)^{y_{n}}\left[1-p\left(y_{n}=1 \mid x_{n}\right)\right]^{1-y_{n}} \rightarrow \max _{w}
$$

- $C$ classes $\left(y_{n c}=\mathbb{I}\left\{y_{n}=c\right\}\right)$ :

$$
\prod_{n=1}^{N} \prod_{c=1}^{C} p\left(y_{n}=c \mid x_{n}\right)^{y_{n c}} \rightarrow \max _{w}
$$

- In practice log-likelihood is maximized.


## Neural network optimization

- Let $W$ denote the total dimensionality of weights space
- Let $E(\hat{y}, y)$ denote the loss function of output
- We may optimize neural network using gradient descent:

```
k=0
initialize randomly wo # small values for sigmoid and tangh
while (stop criteria not met):
    w}\mp@subsup{w}{}{k+1}:=\mp@subsup{w}{}{k}-\eta\nablaE(\mp@subsup{w}{}{k}
    k:=k+1
```

- Standardization of features makes gradient descend converge faster
- Other optimization methods are more efficient (such as conjugate gradients)
- Denote $W$ - total number of edges (and weights) in the neural net.


## Gradient calculation

- Direct $\nabla E(w)$ calculation, using

$$
\frac{\partial E}{\partial w_{i}}=\frac{E\left(w+\varepsilon_{i}\right)-E(w)}{\varepsilon}+O(\varepsilon)
$$

or better

$$
\frac{\partial E}{\partial w_{i}}=\frac{E\left(w+\varepsilon_{i}\right)-E\left(w-\varepsilon_{i}\right)}{2 \varepsilon}+O\left(\varepsilon^{2}\right)
$$

has complexity $O\left(W^{2}\right)$ [W forward propagations to evaluate W derivatives]

Backpropagation algorithm needs only $O(W)$ to evaluate all derivatives.

## Multiple local optima problem

- Optimization problem for neural nets is non-convex.
- Different optima will correspond to:
- different starting parameter values
- different training samples
- So we may solve task many times for different conditions and then
- select best model
- alternatively: average different obtained models to get ensemble


## Table of Contents

## (1) Introduction

(2) Output generation
(3) Weight space symmetries

4 Neural network optimization
(5) Backpropagation algorithm
(6) Invariances
(7) Case study: ZIP codes recognition

## Definitions

- Denote $w_{i j}$ be the weight of edge, connecting $i$-th and $j$-th neuron.
- Define $\delta_{j}=\frac{\partial E}{\partial I_{j}}=\frac{\partial E}{\partial O_{j}} \frac{\partial O_{j}}{\partial I_{j}}$
- Since $E$ depends on $w_{i j}$ through the following functional relationship $E\left(w_{i j}\right) \equiv E\left(O_{j}\left(I_{j}\left(w_{i j}\right)\right)\right)$, using the chain rule we obtain:

$$
\frac{\partial E}{\partial w_{i j}}=\frac{\partial E}{\partial I_{j}} \frac{\partial I_{j}}{\partial w_{i j}}=\delta_{j} O_{i}
$$

because $\frac{\partial I_{j}}{\partial w_{i j}}=\frac{\partial}{\partial w_{i j}}\left(\sum_{k \in \operatorname{inc}(j)} w_{k j} O_{k}\right)=O_{i}$, where $\operatorname{inc}(j)$ is a set of all neurons with outgoing edges to neuron $j$.

- $\frac{\partial E}{\partial I_{j}}=\frac{\partial E}{\partial O_{j}} \frac{\partial O_{j}}{\partial I_{j}}=\frac{\partial E}{\partial O_{j}} \varphi^{\prime}\left(I_{j}\right)$, where $\varphi$ is the activation function.


## Output layer

- If neuron $j$ belongs to the output node, then error $\frac{\partial E}{\partial O_{j}}$ is calculated directly.
- For output layer deltas are calculated directly:

$$
\begin{equation*}
\delta_{j}=\frac{\partial E}{\partial O_{j}} \frac{\partial O_{j}}{\partial I_{j}}=\frac{\partial E}{\partial O_{j}} \varphi^{\prime}\left(I_{j}\right) \tag{1}
\end{equation*}
$$

- example for training set $=\{$ single point $x$ and true vector of outputs $\left.\left(y_{1}, \ldots y_{\mid O L}\right)\right\}$ :
- for $E=\frac{1}{2} \sum_{j \in O L}\left(O_{j}-y_{j}\right)^{2}$ :

$$
\frac{\partial E}{\partial O_{j}}=O_{j}-y_{j}
$$

- for $\varphi(I)=\operatorname{sigm}(I)$ :

$$
\varphi^{\prime}\left(I_{j}\right)=\sigma\left(I_{j}\right)\left(1-\sigma\left(I_{j}\right)\right)=O_{j}\left(1-O_{j}\right)
$$

- finally

$$
\delta_{j}=\left(O_{j_{4}-6}-y_{j}\right) O_{j}\left(1-O_{j}\right)
$$

## Inner layer

- If neuron $j$ belongs some hidden layer, denote out $(j)=\left\{k_{1}, k_{2}, \ldots k_{m}\right\}$ the set of all neurons, receiving output from neuron $j$.
- The effect of $O_{j}$ on $E$ is fully absorbed by $I_{k_{1}}, I_{k_{2}}, \ldots I_{k_{m}}$, so

$$
\frac{\partial E\left(O_{j}\right)}{\partial O_{j}}=\frac{\partial E\left(I_{k_{1}}, I_{k_{2}}, \ldots I_{k_{m}}\right)}{\partial O_{j}}=\sum_{k \in \text { out }(j)}\left(\frac{\partial E}{\partial I_{k}} \frac{\partial I_{k}}{\partial O_{j}}\right)=\sum_{k \in \text { out }(j)}\left(\delta_{k} w_{j k}\right)
$$

- So for layers other than output layer we have:

$$
\begin{equation*}
\delta_{j}=\frac{\partial E}{\partial I_{j}}=\frac{\partial E}{\partial O_{j}} \frac{\partial O_{j}}{\partial I_{j}}=\sum_{k \in \text { out }(j)}\left(\delta_{k} w_{j k}\right) \varphi^{\prime}\left(I_{j}\right) \tag{2}
\end{equation*}
$$

- Weight derivatives are calculated using errors and outputs:

$$
\begin{equation*}
\frac{\partial E}{\partial w_{i j}}=\frac{\partial E}{\partial I_{j}} \frac{\partial I_{j}}{\partial w_{i j}}=\delta_{j} O_{i} \tag{3}
\end{equation*}
$$

## Backpropagation

- Backpropagation algorithm:
(1) Forward propagate $x_{n}$ to the neural network, store all inputs $I_{i}$ and outputs $O_{i}$ for each neuron.
(2) Calculate $\delta_{i}$ for all $i \in$ output layer using (1).
(3) Backpropagate $\delta_{i}$ from final layer backwards layer by layer using (2).
(1) Using calculated deltas and outputs calculate $\frac{\partial E}{\partial w_{i j}}$ with (3).
- Algorithm complexity: $O(W)$, where $W$ is total number of edges.
- Updates:
- batch
- stochastic
- using minibatches of objects


## Regularization

- Constrain model complexity directly
- constrain number of neurons
- constrain number of layers
- impose constraints on weights
- Take a flexible model
- use early stopping during iterative evaluation (by controlling validation error)
- quadratic regularization

$$
\tilde{E}(w)=E(w)+\lambda \sum_{i} w_{i}^{2}
$$

## Table of Contents

## (1) Introduction

(2) Output generation
(3) Weight space symmetries

4 Neural network optimization
(5) Backpropagation algorithm
(6) Invariances
(7) Case study: ZIP codes recognition

## Invariances

- It may happen that solution should not depend on certain kinds of transformations in the input space.
- Example: character recognition task
- translation invariance
- scale invariance
- invariance to small rotations
- invariance to small uniform noise



## Invariances

- Approaches to build an invariant model:
- augment training objects with their transformed copies according to given invariances
- amount of possible transformations grows exponentially with the number of invariances
- add regularization term to the target cost function, which penalizes changes in output after invariant transformations
- see tangent propagation
- extract features that are invariant to transformations
- build the invariance properties into the structure of neural network
- see convolutional neural networks


## Augmentation of training samples

(1) generate a random set of invariant transformations
(2) apply these transformations to training objects
(3) obtain new training objects


## Tangent propagation

- Denote $s(x, \xi)$ be vector $x$ after invariant transformation parametrized by $\xi$.
- Denote

$$
\tau_{n}=\left.\frac{\partial s\left(x_{n}, \xi\right)}{\partial \xi}\right|_{\xi=0}, \quad J_{k i}=\frac{\partial y_{k}}{\partial x_{i}}
$$

- We want $\left.\frac{\partial y_{k}}{\partial \xi}\right|_{\xi=0}$ to be as small, as possible.
- Sensitivity of $y_{k}$ to small invariant transformation:

$$
\left.\frac{\partial y_{k}}{\partial \xi}\right|_{\xi=0}=\sum_{i=1}^{D} \frac{\partial y_{k}}{\partial x_{i}} \frac{\partial x_{i}}{\partial \xi}=\sum_{i=1}^{D} J_{k i} \tau_{i}
$$

- Tangent propagation - modify target cost function:

$$
\tilde{E}=E+\lambda \sum_{n} \sum_{k}\left(\sum_{i=1}^{D} J_{n k i} \tau_{n i}\right)^{2}
$$

## Convolutional neural networks

- Convolutional neural network:
- Used for image analysis
- Consists of a set of convolutional layer / sub-sampling layer pairs and aggregating layer



## Convolutional neural networks

- Convolutional layer
- Convolutional layer consists of a number of feature maps
- Feature map has the same dimensionality as input layer
- Locality: each neuron in the feature map takes output from small neigborhood of input layer neurons
- Equivalence: the same transformation is applied by each neuron in the feature map
- obtained by constraining sets of weights to each feature map layer neuron to be equal
- similar to convolution with moving adaptive kernel
- effectively it is feature extraction from a region


## Convolutional neural networks

- Sub-sampling layer
- Consists of a number of planes, each corresponding to respective feature map on the previous convolutional layer
- Locality: Sub-sampling layer neurons take output from small neigborhood of respective feature map neurons
- neigbourhoods are chosen to be contiguous and non-overlapping
- Aggregation: input of each neuron $i$ is: $w_{i 0}+w_{i 1} F$, where $w_{i 0}, w_{i 1}$ are adjustable weights and $F$ is aggregation function (sum or max of activations of respective feature map neurons)
- Implements small translational invariance
- There may be a sequence of convolutional and sub-sampling layers
- gradual dimensionality reduction


## Table of Contents

## (1) Introduction

(2) Output generation
(3) Weight space symmetries

4 Neural network optimization
(5) Backpropagation algorithm
6) Invariances
(7) Case study: ZIP codes recognition

## Case study (due to Hastie et al. The Elements of Statistical Learning)

ZIP code recognition task

$$
\begin{array}{llllllllll}
0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\
0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\
0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\
0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\
0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9
\end{array}
$$

## Neural network structures

Net1: no hidden layer
Net2: 1 hidden layer, 12 hidden units fully connected
Net3: 2 hidden layers, locally connected
Net4: 2 hidden layers, locally connected with weight sharing Net5: 2 hidden layers, locally connected, 2 levels of weight sharing


## Results



## Addition

- Deep learning
- Neural networks weights may be constrained to belong to mixture density
- $\tilde{E} \leftarrow E-\lambda P(w)$, where $P(w)$ is the mixture probability of weights
- soft forcing of weights to group into similar clusters
- Neural networks may model not only real value outputs, but densities
- each output - frequency of histogram bin
- each output - either prior or mean or variance of mixture of parametrized density (normal, beta, etc.)


## Conclusion

- Advantages of neural networks:
- can model accurately complex non-linear relationships
- easily parallelizable
- Disadvantages of neural networks:
- hardly interpretable ("black-box" algorithm)
- optimization requires skill
- too many parameters
- may converge slowly
- may converge to inefficient local minimum far from global one


[^0]:    ${ }^{1}$ Propose neural networks estimating OR,AND, XOR functions on boolean inputs.

