## Advanced Machine Learning Lecture 10

Backpropagation
05.12.2016

Bastian Leibe
RWTH Aachen
http://www.vision.rwth-aachen.de/
leibe@vision.rwth-aachen.de

This Lecture: Advanced Machine Learning

- Regression Approaches
, Linear Regression
- Regularization (Ridge, Lasso)
, Kernels (Kernel Ridge Regression)
, Gaussian Processes

- Approximate Inference
, Sampling Approaches
- MCMC
- Deep Learning
, Linear Discriminants
, Neural Networks
, Backpropagation
- CNNs, RNNs, ResNets, etc.
- One output node per class


Output layer Weights
Input layer

- Outputs
, Linear outputs
With output nonlinearity
$y_{k}(\mathbf{x})=\sum_{i=0}^{d} W_{k i} x_{i}$
$y_{k}(\mathbf{x})=g\left(\sum_{i=0}^{d} W_{k i} x_{i}\right)$
$\Rightarrow$ Can be used to do multidimensional linear regression or multiclass classification.

Recap: Non-Linear Basis Functions

- Straightforward generalization


Output layer Weights
Feature layer Mapping (fixed) Input layer

- Remarks
, Perceptrons are generalized linear discriminants!
, Everything we know about the latter can also be applied here.
, Note: feature functions $\phi(\mathbf{x})$ are kept fixed, not learned!


## Recap: Perceptron Learning

- Process the training cases in some permutation
- If the output unit is correct, leave the weights alone.
, If the output unit incorrectly outputs a zero, add the input vector to the weight vector.

Output layer Weights
Feature layer Mapping (fixed)
Input layer

$$
\begin{array}{lr}
\text { Linear outputs } & \text { with output nonlinearity } \\
y_{k}(\mathbf{x})=\sum_{i=0}^{d} W_{k i} \phi\left(x_{i}\right) & y_{k}(\mathbf{x})=g\left(\sum_{i=0}^{d} W_{k i} \phi\left(x_{i}\right)\right)
\end{array}
$$

- Outputs


## Recap: Non-Linear Basis Functions

- Straightforward generalization

, If the output unit incorrectly outputs a one, subtract the input vector from the weight vector.
- Translation

$$
w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\eta\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right)
$$

, This is the Delta rule a.k.a. LMS rule!
$\Rightarrow$ Perceptron Learning corresponds to $1^{\text {stt-order (stochastic) }}$ Gradient Descent of a quadratic error function!

## Recap: Loss Functions

- We can now also apply other loss functions



## Topics of This Lecture <br> - Learning with Hidden Units

- Obtaining the Gradients
- Naive analytical differentiation
, Numeric differentiation
, Backpropagation
, Computational graphs
- Automatic differentiation
- Practical Issues
, Nonlinearities
, Sigmoid outputs and the $L_{2}$ loss
, Implementing Softmax correctly

Gradient Descent

- Two main steps

1. Computing the gradients for each weight
today
2. Adjusting the weights in the direction of

Thursday the gradient

## Topics of This Lecture

- Learning with Hidden Units
- Obtaining the Gradients
, Naive analytical differentiation
, Numeric differentiation
- Backpropagation
, Computational graphs
- Automatic differentiation
- Practical Issues

Nonlinearities
Sigmoid outputs and the $L_{2}$ loss
Implementing Softmax correctly


Output layer

Hidden layer

Input layer

- Output
$y_{k}(\mathbf{x})=g^{(2)}\left(\sum_{i=0}^{h} W_{k i}^{(2)} g^{(1)}\left(\sum_{j=0}^{d} W_{i j}^{(1)} x_{j}\right)\right)$


## Learning with Hidden Units

- How can we train multi-layer networks efficiently?
- Need an efficient way of adapting all weights, not just the last layer.
- Idea: Gradient Descent
, Set up an error function

$$
E(\mathbf{W})=\sum_{n} L\left(t_{n}, y\left(\mathbf{x}_{n} ; \mathbf{W}\right)\right)+\lambda \Omega(\mathbf{W})
$$

with a loss $L(\cdot)$ and a regularizer $\Omega(\cdot)$.
, E.g., $L(t, y(\mathbf{x} ; \mathbf{W}))=\sum_{n}\left(y\left(\mathbf{x}_{n} ; \mathbf{W}\right)-t_{n}\right)^{2} \quad \mathbf{L}_{2}$ loss

$$
\begin{array}{lr}
\Omega(\mathbf{W})=\|\mathbf{W}\|_{F^{\prime}}^{2} & \begin{array}{c}
\mathrm{L}_{2} \text { regularizer } \\
\text { ("weight decay") }
\end{array} \\
\Rightarrow \text { Update each weight } W_{i j}^{(k)} \text { in the direction of the gradient } \frac{\partial E(\mathbf{W})}{\partial W_{i, 10}^{(k)}} \\
\text { B. Leibe } &
\end{array}
$$

RWIHAACHE
UNIVERSIT
Recap: Multi-Layer Perceptrons

- Adding more layers

Input layer

- One-dimensional case: Scalar functions


## Excursion: Chain Rule of Differentiation



$$
\begin{aligned}
& \Delta z=\frac{\mathrm{d} z}{\mathrm{~d} y} \Delta y \\
& \Delta y=\frac{\mathrm{d} y}{\mathrm{~d} x} \Delta x \\
& \Delta z=\frac{\mathrm{d} z}{\mathrm{~d} y} \frac{\mathrm{~d} y}{\mathrm{~d} x} \Delta x \\
& \frac{\mathrm{~d} z}{\mathrm{~d} x}=\frac{\mathrm{d} z}{\mathrm{~d} y} \frac{\mathrm{~d} y}{\mathrm{~d} x}
\end{aligned}
$$

## Excursion: Chain Rule of Differentiation

- Multi-dimensional case: Total derivative

$$
\begin{array}{ll}
\frac{\partial z}{\partial y_{1}} y_{2} \ldots \int_{2 z}^{\partial y_{k}} y_{k} & \frac{\partial z}{\partial x}
\end{array}=\frac{\partial z}{\partial y_{1}} \frac{\partial y_{1}}{\partial x}+\frac{\partial z}{\partial y_{2}} \frac{\partial y_{2}}{\partial x}+\ldots
$$

$\Rightarrow$ Need to sum over all paths that lead to the target variable $x$.

## Obtaining the Gradients

- Approach 1: Naive Analytical Differentiation


$$
\begin{aligned}
& \frac{\partial E(\mathbf{W})}{\partial W_{10}^{(2)}} \cdots \frac{\partial E(\mathbf{W})}{\partial W_{k h}^{(2)}} \\
& \frac{\partial E(\mathbf{W})}{\partial W_{10}^{(1)}} \ldots \frac{\partial E(\mathbf{W})}{\partial W_{h d}^{(1)}}
\end{aligned}
$$

. Compute the gradients for each variable analytically.
, What is the problem when doing this?
$\Rightarrow$ With increasing depth, there will be exponentially many paths!
$\Rightarrow$ Infeasible to compute this way.
B. Leibe

## Topics of This Lecture

- Learning with Hidden Units
- Obtaining the Gradients
, Naive analytical differentiation
, Numerical differentiation
, Backpropagation
, Computational graphs
, Automatic differentiation
- Practical Issues

Nonlinearities

- Sigmoid outputs and the $L_{2}$ loss

Implementing Softmax correctly

## Obtaining the Gradients

- Approach 2: Numerical Differentiation

, Given the current state $\mathbf{W}^{(\tau)}$, we can evaluate $E\left(\mathbf{W}^{(\tau)}\right)$.
, Idea: Make small changes to $\mathbf{W}^{(\tau)}$ and accept those that improve $E\left(\mathbf{W}^{(\tau)}\right)$.
$\Rightarrow$ Horribly inefficient! Need several forward passes for each weight. Each forward pass is one run over the entire dataset!

| Topics of This Lecture <br> - Learning with Hidden Units <br> - Obtaining the Gradients <br> , Naive analytical differentiation <br> , Numerical differentiation <br> , Backpropagation <br> - Computational graphs <br> - Automatic differentiation <br> - Practical Issues <br> , Nonlinearities <br> , Sigmoid outputs and the $L_{2}$ loss <br> = Implementing Softmax correctly | RWIHMAELHET |
| :---: | :---: |
|  | 19 |



## Backpropagation Algorithm



- Efficient propagation scheme
, $y_{i}$ is already known from forward pass! (Dynamic Programming)
$\Rightarrow$ Propagate back the gradient from layer $j$ and multiply with $y_{i}$.


## Analysis: Backpropagation

- Backpropagation is the key to make deep NNs tractable However...
- The Backprop algorithm given here is specific to MLPs
- It does not work with more complex architectures e.g. skip connections or recurrent networks!
- Whenever a new connection function induces a different functional form of the chain rule, you have to derive a new Backprop algorithm for it.
$\Rightarrow$ Tedious...
- Let's analyze Backprop in more detail
, This will lead us to a more flexible algorithm formulation
B. Leibe


## Summary: MLP Backpropagation

- Forward Pass
$\mathbf{y}^{(0)}=\mathbf{x}$
for $k=1, \ldots, l$ do
$\mathbf{z}^{(k)}=\mathbf{W}^{(k)} \mathbf{y}^{(k-1)}$
$\mathbf{y}^{(k)}=g_{k}\left(\mathbf{z}^{(k)}\right)$
endfor
$\mathbf{y}=\mathbf{y}^{(l)}$
$E=L(\mathbf{t}, \mathbf{y})+\lambda \Omega(\mathbf{W})$
- Backward Pass
$\mathbf{h} \leftarrow \frac{\partial E}{\partial \mathbf{y}}=\frac{\partial}{\partial \mathbf{y}} L(\mathbf{t}, \mathbf{y})+\lambda \frac{\partial}{\partial \mathbf{y}} \Omega$
for $k=l, l-1, \ldots, 1$ do
$\mathbf{h} \leftarrow \frac{\partial E}{\partial \mathbf{z}^{(k)}}=\mathbf{h} \odot g^{\prime}\left(\mathbf{y}^{(k)}\right)$ $\frac{\partial E}{\partial \mathbf{W}^{(k)}}=\mathbf{h} \mathbf{y}^{(k-1) \top}+\lambda \frac{\partial \Omega}{\partial \mathbf{W}^{(k)}}$
$\mathbf{h} \leftarrow \frac{\partial E}{\partial \mathbf{y}^{(k-1)}}=\mathbf{W}^{(k) \top} \mathbf{h}$
endfor
- Notes

For efficiency, an entire batch of data $\mathbf{X}$ is processed at once.

- $\odot$ denotes the element-wise product


## Computational Graphs

- We can think of mathematical expressions as graphs
E.g., consider the expression $e=(a+b) *(b+1)$
- We can decompose this into the operations
$c=a+b$
$d=b+1$
$e=c * d$

and visualize this as a computational graph.
- Evaluating partial derivatives $\frac{\partial Y}{\partial X}$ in such a graph

General rule: sum over all possible paths from $Y$ to $X$ and multiply the derivatives on each edge of the path together.

Slide inssired by Christonher Olah B. Leibe

## Factoring Paths

- Problem: Combinatorial explosion
, Example:

- There are 3 paths from $X$ to $Y$ and 3 more from $Y$ to $Z$.
- If we want to compute $\frac{\partial Z}{\partial X}$, we need to sum over $3 \times 3$ paths:

$$
\frac{\partial Z}{\partial X}=\alpha \delta+\alpha \epsilon+\alpha \zeta+\beta \delta+\beta \epsilon+\beta \zeta+\gamma \delta+\gamma \epsilon+\gamma \zeta
$$

- Instead of naively summing over paths, it's better to factor them $\frac{\partial Z}{\partial X}=(\alpha+\beta+\gamma) *(\delta+\epsilon+\zeta)$


## Efficient Factored Algorithms



Apply operator $\frac{\partial}{\partial X}$
to every node.


## - Efficient algorithms for computing the sum

Instead of summing over all of the paths explicitly, compute the sum more efficiently by merging paths back together at every node.

## Why Do We Care?

- Let's consider the example again
, Using reverse-mode differentiation from $e$ down...
, Runtime: $\mathcal{O}$ (\#edges)
, Result: derivative of $e$ with respect to every node.

$\Rightarrow$ This is what we want to compute in Backpropagation!
- Forward differentiation needs one pass per node. With backward differentiation can compute all derivatives in one single pass.
$\Rightarrow$ Speed-up in $\mathcal{O}$ (\#inputs) compared to forward differentiation!
Slide inspired by Christopher Olah B. Leibe Imaee source: Christopher Olah, colah, pithub id


## Topics of This Lecture

- Learning with Hidden Units
- Obtaining the Gradients
. Naive analytical differentiation
, Numerical differentiation
, Backpropagation
- Computational graphs
- Automatic differentiation
- Practical Issues

Nonlinearities
Sigmoid outputs and the $L_{2}$ loss
Implementing Softmax correctly

## Obtaining the Gradients

- Approach 4: Automatic Differentiation

, Convert the network into a computational graph.
, Each new layer/module just needs to specify how it affects the forward and backward passes.
- Apply reverse-mode differentiation.
$\Rightarrow$ Very general algorithm, used in today's Deep Learning packages B. Leibe
- Solution in many current Deep Learning libraries
, Provide a limited form of automatic differentiation
, Restricted to "programs" composed of "modules" with a predefined set of operations.
- Each module is defined by two main functions

1. Computing the outputs $y$ of the module given its inputs $x$

$$
\mathbf{y}=\operatorname{module} . \operatorname{fprop}(\mathbf{x})
$$

where $\mathbf{x}, \mathrm{y}$, and intermediate results are stored in the module.
2. Computing the gradient $\partial E / \partial \mathbf{x}$ of a scalar cost w.r.t. the inputs $\mathbf{x}$ given the gradient $\partial E / \partial \mathbf{y}$ w.r.t. the outputs $\mathbf{y}$

$$
\frac{\partial E}{\partial \mathbf{x}}=\text { module.bprop }\left(\frac{\partial E}{\partial \mathbf{y}}\right)
$$

## Topics of This Lecture

- Learning with Hidden Units
- Obtaining the Gradients

Naive analytical differentiation
Numerical differentiation
Backpropagation
Computational graphs
Automatic differentiation

- Practical Issues
, Nonlinearities
, Sigmoid outputs and the $L_{2}$ loss
, Implementing Softmax correctly
, Efficient batch processing


## Commonly Used Nonlinearities

- Sigmoid

$$
\begin{aligned}
g(a) & =\sigma(a) \\
& =\frac{1}{1+\exp \{-a\}}
\end{aligned}
$$

- Hyperbolic tangent

$$
\begin{aligned}
g(a) & =\tanh (a) \\
& =2 \sigma(2 a)-1
\end{aligned}
$$

- Softmax

$$
g(\mathbf{a})=\frac{\exp \left\{-a_{i}\right\}}{\sum_{j} \exp \left\{-a_{j}\right\}}
$$

## Commonly Used Nonlinearities (2)

- Hard tanh

$$
g(a)=\max \{-1, \min \{1, a\}\}
$$

- Rectified linear unit (ReLU)

$$
g(a)=\max \{0, a\}
$$

- Maxout

$$
g(\mathbf{a})=\max _{i}\left\{\mathbf{w}_{i}^{\top} \mathbf{a}+b_{i}\right\}
$$



Usage

- Output nodes
- Typically, a sigmoid or tanh function is used here.

Sigmoid for nice probabilistic interpretation (range [0,1]). tanh for regression tasks

- Internal nodes
- Historically, tanh was most often used.
, tanh is better than sigmoid for internal nodes, since it is already centered.
, Internally, tanh is often implemented as piecewise linear function (similar to hard tanh and maxout).
- More recently: ReLU often used for classification tasks.


## Topics of This Lecture

- Learning with Hidden Units
- Obtaining the Gradients

Naive analytical differentiation
Numeric differentiation
Backpropagation
Computational graphs
Automatic differentiation

- Practical Issues
, Nonlinearities
, Sigmoid outputs and the $L_{2}$ loss
- Implementing Softmax correctly
Another Note on Error Functions
Zero gradient!
- Squared error on sigmoid/tanh output function
, Avoids penalizing "too correct" data points.
, But: zero gradient for confidently incorrect classifications!
$\Rightarrow$ Do not use $\mathrm{L}_{2}$ loss with sigmoid outputs (instead: cross-entropy)!


## Topics of This Lecture

- Learning with Hidden Units
- Obtaining the Gradients

Naive analytical differentiation
Numerical differentiation
Backpropagation
Computational graphs
Automatic differentiation

- Practical Issues


## , Nonlinearities

, Sigmoid outputs and the $\mathrm{L}_{2}$ loss

- Implementing Softmax correctly
- Softmax output
- De-facto standard for multi-class outputs

$$
E(\mathbf{w})=-\sum_{n=1}^{N} \sum_{k=1}^{K}\left\{\mathbb{I}\left(t_{n}=k\right) \ln \frac{\exp \left(\mathbf{w}_{k}^{\top} \mathbf{x}\right)}{\sum_{j=1}^{K} \exp \left(\mathbf{w}_{j}^{\top} \mathbf{x}\right)}\right\}
$$

- Practical issue
, Exponentials get very big and can have vastly different magnitudes.
, Trick 1: Do not compute first softmax, then log,
but instead directly evaluate log-exp in the denominator.
Trick 2: Softmax has the property that for a fixed vector $\mathbf{b}$ $\operatorname{softmax}(\mathbf{a}+\mathbf{b})=\operatorname{softmax}(\mathbf{a})$

$$
\operatorname{softmax}(\mathbf{a}+\mathbf{b})=\operatorname{softmax}(\mathbf{a})
$$

$\Rightarrow$ Subtract the largest weight vector $\mathbf{w}_{j}$ from the others.

RNIHAACHE
References and Further Reading

- More information on Backpropagation can be found in Chapter 6 of the Goodfellow \& Bengio book

Ian Goodfellow, Aaron Courville, Yoshua Bengio lan Goodfellow
Deep Learning MIT Press, in preparation
https://goodfeli.github.io/dlbook/Chapter 6 of the Goodellow a Bengio book


