



# *Neural Networks*

*(P-ITEEA-0011)*

## **Gradient based optimization methods**

Akos Zarandy  
Lecture 4  
October 1, 2018




# Contents

- Recall
  - Single- and multilayer perceptron and its learning method
- Mathematical background
- Simple gradient based optimizers
  - 1<sup>st</sup> and 2<sup>nd</sup> order optimizers
- Advanced optimizers
  - Momentum
  - ADAM

# We have discussed...

- Architecture of the single- and multilayer artificial neural networks
- Operation of these networks
- Derivation of the parameters
  - Direct calculations
  - Training (most commonly used method)
- Various training methods
  - Single-layer perceptrons
  - Fully connected multilayer perceptrons



$$Net(\mathbf{W}, \mathbf{x}) = y = \phi \left( \sum_{i=1}^{n^L} w_i^{(L)} \cdot \phi \left( \sum_{j=1}^{n^{L-1}} w_{ij}^{(L-1)} \cdot \dots \cdot \phi \left( \sum_{m=1}^{n^1} w_{km}^{(1)} x_m \right) \cdot \dots \right) \right)$$

***What does training really mean mathematically?***



# What is training?

**Stochastic** process is a process, where we cannot observe the exact values. In these processes, our observations are always corrupted with some random noise.

- Given:
  - Definition of the network architecture
    - Topology
    - Initial weights
    - Activation functions (nonlinearities)
    - Training set  $(\mathbf{x}_i \rightarrow \mathbf{y}_i)$
- Goal:
  - Calculation of the optimal weight composition:  $\mathbf{W}_{\text{opt}}$ 
    1. *Having a known function to approximate*

$$\mathbf{w}_{\text{opt}} : \min_{\mathbf{w}} \|\mathbf{F}(\mathbf{x}) - \text{Net}(\mathbf{x}, \mathbf{w})\|^2 = \min_{\mathbf{w}} \int \dots \int (\mathbf{F}(\mathbf{x}) - \text{Net}(\mathbf{x}, \mathbf{w}))^2 dx_1 \dots dx_N$$

2. *Having a set of observations from a stochastic process*

$$\mathbf{w}_{\text{opt}}^{(K)} : \min_{\mathbf{w}} \frac{1}{K} \sum_{k=1}^K (d_k - \text{Net}(\mathbf{x}_k, \mathbf{w}))^2$$

**OPTIMIZATION!!!**



# Optimization

- Given an **Objective function** to optimize
  - Also called: Error function, Cost function, Loss function, Criterion
- Function types:
  - Quadratic, in case of regression (stochastic process)

$$R_{emp}(\mathbf{w}) = \frac{1}{K} \sum_{k=1}^K (d_k - \text{Net}(\mathbf{x}_k, \mathbf{w}))^2$$

- Conditional log-likelihood, in case of classification (classification process)
  - The sum of the negative logarithmic likelihood is minimized

$$\theta_{\text{ML}} = \arg \min_{\theta} \sum_{i=1}^m -\log P(\mathbf{y}^{(i)} | \mathbf{x}^{(i)}; \theta)$$



# Optimizations

- *Here we always minimize the objective function*
- *Optimal value is indicated with*  
$$x^* = \operatorname{argmin} f(x).$$



# Mathematics behind: Function analysis

- Assumptions

- Poor conditioning
- Conditioning number  
(Ratio of Eigen values):  $\max_{i,j} \left| \frac{\lambda_i}{\lambda_j} \right|$

$$f(x) = A^{-1}x \quad A \in \mathbb{R}^{n \times n}$$

- Applied functions should be Lipschitz continuous or have Lipschitz continuous derivative

$$\forall x, \forall y, \quad |f(x) - f(y)| \leq L \|x - y\|_2$$

**Conditioning** refers to how rapidly a function changes with respect to small changes in its inputs.

Functions that change rapidly when their inputs are perturbed slightly can be problematic for scientific computation because rounding errors in the inputs can result in large changes in the output.  
(e.g. Matrix inversion)

(where:

L is the Lipschitz constant)



# Basic idea of Gradient Descent


- Derivative means for small  $\varepsilon$

$$f(x + \varepsilon) \approx f(x) + \varepsilon f'(x)$$

- therefore*

$$f(x - \varepsilon \operatorname{sign}(f'(x))) \leq f(x)$$

- This technique is called **Gradient Descent** (Cauchy, 1847).




Optimization goal is to find the  $f'(x) = 0$  position.  
(Critical or stationary points)




# Stationary points

- Local minimum, where  $f'(x)=0$ , and  $f(x)$  is smaller than all neighboring points
- Local maximum, where  $f'(x)=0$ , and  $f(x)$  is larger than all neighboring points
- Saddle points, where  $f'(x)=0$ , and neither minimum nor maximum


Minimum




Maximum



Saddle point






# Local and global minimum



In neural network parameter optimization we usually settle for finding a value of  $f$  that is very low, but not necessarily minimal in any formal sense.

# Multidimensional input functions I

- In case of a vector scalar function
- In 2D, directional derivatives (slope towards  $x$  and  $y$ ):




# Multidimensional input functions II

- In case of a vector scalar function
- Gradient definition in 2D

$$f : \mathbb{R}^2 \rightarrow \mathbb{R}$$

$$\nabla f(x, y) := \begin{pmatrix} \frac{\partial f}{\partial x} & \frac{\partial f}{\partial y} \end{pmatrix}$$




A vector in the in the  $x$ - $y$  plane

# Multidimensional input functions III

- The gradient defines (hyper) plane approximating the function infinitesimally

$$\Delta z = \frac{\partial f(x, y)}{\partial x} \cdot \Delta x + \frac{\partial f(x, y)}{\partial y} \cdot \Delta y$$



# Multidimensional input functions IV



- Directional derivative to an arbitrary direction  $\mathbf{u}$  ( $\mathbf{u}$  is unit vector) is the slope of  $f$  in that direction:

$$\mathbf{u}^T \nabla_x f(\mathbf{x})$$

- $f$  decreases the fastest:

$$\min_{\mathbf{u}, \mathbf{u}^T \mathbf{u} = 1} \mathbf{u}^T \nabla_x f(\mathbf{x}) = \min_{\mathbf{u}, \mathbf{u}^T \mathbf{u} = 1} \|\mathbf{u}\|_2 \|\nabla_x f(\mathbf{x})\|_2 \cos \theta$$


Not changing with  $\mathbf{u}$  (pointing to  $\|\nabla_x f(\mathbf{x})\|_2$ )

minimum at 180 (pointing to  $\cos \theta$ )

- $\mathbf{u}$  is opposite to the gradient!!!

New points towards steepest descent:

$$\mathbf{x}' = \mathbf{x} - \varepsilon \nabla_x f(\mathbf{x})$$






# Gradient Descent in multidimensional input case



- Steepest gradient descent iteration

$$\mathbf{x}(n+1) = \mathbf{x}(n) - \varepsilon \nabla_x f(\mathbf{x}(n))$$

- $\varepsilon$  is the learning rate
- Choosing  $\varepsilon$ :
  - Small constant
  - Decreases as the iteration goes ahead
  - **Line search**: checked with several values, and the one selected, where  $\mathbf{x}'$  is the smallest
- Stopping condition of the gradient descent iteration
  - When the gradient is zero or close to zero





# Jacobian Matrix

- Partial derivative of a vector  $\rightarrow$  vector function
- Specifically, if we have a function  $\mathbf{f} : \mathbb{R}^m \rightarrow \mathbb{R}^n$  then the Jacobian matrix  $\mathbf{J} \in \mathbb{R}^{n \times m}$  of  $\mathbf{f}$  is defined such that:  $\mathbf{J}_{i,j} = \frac{\partial}{\partial x_j} f(x_i)$

$$\mathbf{J} = \begin{bmatrix} \frac{\partial \mathbf{f}}{\partial x_1} & \cdots & \frac{\partial \mathbf{f}}{\partial x_n} \end{bmatrix} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_n} \end{bmatrix}$$

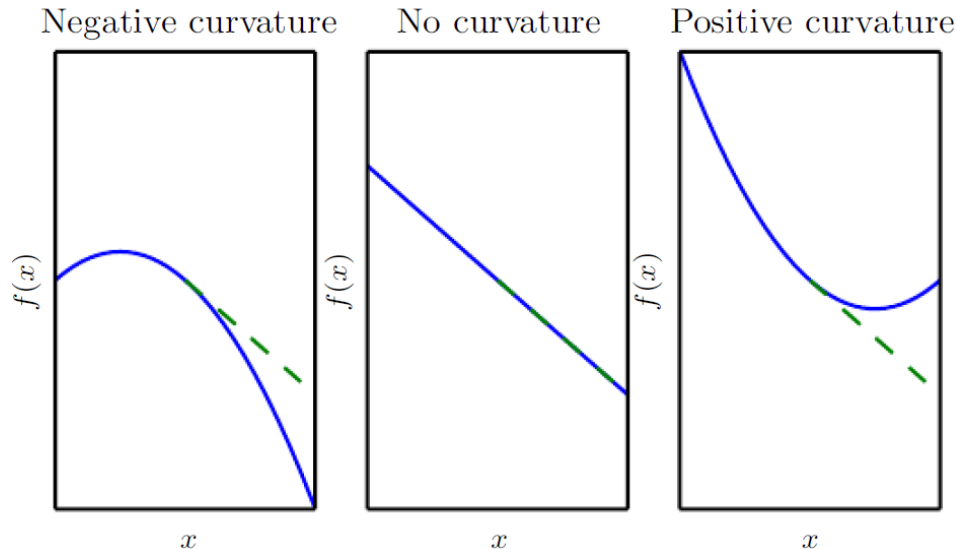


# 2<sup>nd</sup> derivatives

- 2<sup>nd</sup> derivative determines the curvature of a line in 1D
- In nD, it is described by the Hessian Matrix

$$H(f(x_{i,j})) = \frac{\partial^2}{\partial x_i \partial x_j} f(x) = \frac{\partial^2}{\partial x_j \partial x_i} f(x)$$

- The Hessian is the Jacobian of the gradient.





# 2<sup>nd</sup> order gradient descent method I

- 2<sup>nd</sup> derivative in a specific direction:  $\mathbf{u}^T \mathbf{H} \mathbf{u}$
- Second-order Taylor series approximation to the function  $f(\mathbf{x})$  *around the current point*  $\mathbf{x}_0$

$$f(\mathbf{x}) \approx f(\mathbf{x}_0) + (\mathbf{x} - \mathbf{x}_0)^T \mathbf{g} + \frac{1}{2} (\mathbf{x} - \mathbf{x}_0)^T \mathbf{H} (\mathbf{x} - \mathbf{x}_0)$$

where:

$\mathbf{g}$ : gradient at  $\mathbf{x}_0$

$\mathbf{H}$ : Hessian at  $\mathbf{x}_0$

- if we substitute  $\mathbf{x}_0 - \varepsilon \mathbf{g} = \mathbf{x} \rightarrow \mathbf{x} - \mathbf{x}_0 = -\varepsilon \mathbf{g}$

$$f(\mathbf{x}_0 - \varepsilon \mathbf{g}) \approx f(\mathbf{x}_0) - \varepsilon \mathbf{g}^T \mathbf{g} + \frac{1}{2} \varepsilon^2 \mathbf{g}^T \mathbf{H} \mathbf{g}$$

## 2<sup>nd</sup> order gradient descent method II



- Analyzing: 
$$f(\mathbf{x}_0 - \varepsilon \mathbf{g}) \approx f(\mathbf{x}_0) - \varepsilon \mathbf{g}^T \mathbf{g} + \frac{1}{2} \varepsilon^2 \mathbf{g}^T \mathbf{H} \mathbf{g}$$

Original value      Expected improvement      Correction due to curvature

- When the third term is too large, the gradient descent step can actually move uphill.
- When it *is zero or negative*, the Taylor series approximation predicts that increasing  $\varepsilon$  *forever will decrease*  $f$  forever.
- In practice, the Taylor series is unlikely to remain accurate for large  $\varepsilon$ , so one must resort to more heuristic choices of  $\varepsilon$  *in this case*.
- When it is positive*, solving for the optimal step

$$\varepsilon^* = \frac{\mathbf{g}^T \mathbf{g}}{\mathbf{g}^T \mathbf{H} \mathbf{g}}$$

# *Simplest 2<sup>nd</sup> order Gradient descent method:*

## Newton Method



$$f(\mathbf{x}) \approx f(\mathbf{x}_0) + (\mathbf{x} - \mathbf{x}_0)^T \nabla_x f(\mathbf{x}_0) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_0)^T \mathbf{H}(f(\mathbf{x}_0)) (\mathbf{x} - \mathbf{x}_0)$$

- Replacing  $(\mathbf{x} - \mathbf{x}_0) \rightarrow \Delta \mathbf{x}$  and differentiating it with  $\Delta \mathbf{x}$ , assuming that we are in a critical point  $\nabla_x f(\mathbf{x}_0) \approx 0$

$$0 = \frac{\partial}{\partial \Delta \mathbf{x}} \left( f(\mathbf{x}_0) + \Delta \mathbf{x}^T \nabla_x f(\mathbf{x}_0) + \frac{1}{2} \Delta \mathbf{x}^T \mathbf{H}(f(\mathbf{x}_0)) \Delta \mathbf{x} \right) = \nabla_x f(\mathbf{x}_0) + \mathbf{H}(f(\mathbf{x}_0)) \Delta \mathbf{x}$$

$$\Delta \mathbf{x} = -\nabla_x f(\mathbf{x}_0) \mathbf{H}(f(\mathbf{x}_0))^{-1} \quad \mathbf{x}(n+1) = -\nabla_x f(\mathbf{x}(n)) \mathbf{H}(f(\mathbf{x}(n)))^{-1}$$

Newton optimization:




# Properties of Newton optimization method

- When  $f$  is a *positive definite quadratic function*, Newton's method jumps in a single step to the minimum of the function directly.
- *Newton's method* can reach the critical point much faster than 1<sup>st</sup> order gradient descent.


$$\Delta \mathbf{x} = -\nabla_x f(\mathbf{x}_0) \mathbf{H}(f(\mathbf{x}_0))^{-1} \quad \mathbf{x}(n+1) = -\nabla_x f(\mathbf{x}(n)) \mathbf{H}(f(\mathbf{x}(n)))^{-1}$$

Newton optimization:

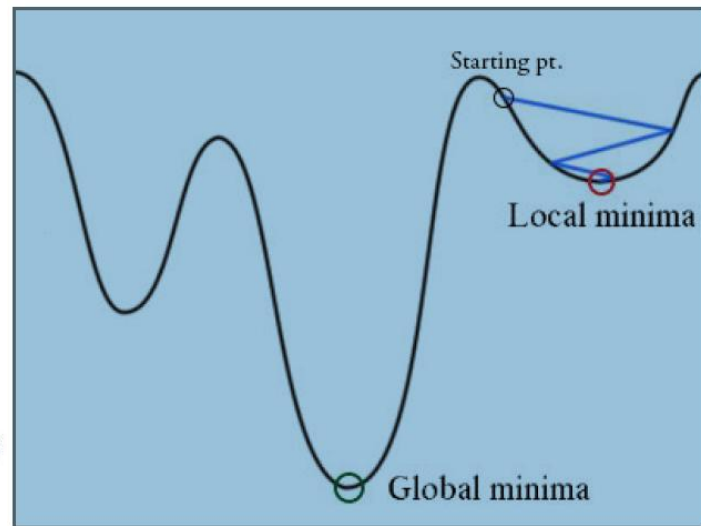
# Convex and non-convex functions



Strongly convex  
function:  
1 local minimum




Non-Strongly convex  
function: infinity local  
touching minima with  
the same values



Non-convex function:  
multiple non-touching  
local minima with  
different values

# Local optimization in non-convex case

- Optimization is done locally in a certain domain, where the function is assumed to be convex
- Multiple local optimization is used to find global minimum



# Most commonly applied gradient descent methods



- Algorithms with changing but not adaptive learning rate
  - Stochastic Gradient Descent algorithm
  - Momentum algorithm
- Algorithms with adaptive learning rate
  - AdaGrad algorithm
  - RMSProp algorithm
  - ADAM algorithm
- 2<sup>nd</sup> order algorithm
  - Newton algorithm



# Most commonly applied gradient descent methods



- Algorithms with changing but not adaptive learning rate
  - Stochastic Gradient Descent algorithm
  - Momentum algorithm
- Algorithms with adaptive learning rate
  - AdaGrad algorithm
  - RMSProp algorithm
  - ADAM algorithm
- 2<sup>nd</sup> order algorithm
  - Newton algorithm

# Stochastic Gradient Descent (SGD) algorithm



- Introduced in 1945
- Gradient Descent method, plus:
  - Applying mini batches
  - Changing the learning rate during the iteration

# Batched approach at SGD



- Single vector update approach
  - Weights are updated after each input vector
- Batched update approach
  - All the input vectors are applied
  - Updates ( $\Delta w_{ij}$ ) are calculated for each vector, and averaged
  - Update is done with the averaged values ( $\Delta w_{ij}$ ) after the entire batch is calculated
- Mini batch approach
  - When the number of inputs are very high (billions), batch would be ineffective
  - Random selection of  $m$  input vectors ( $m$  is a few hundred)
  - Updates ( $\Delta w_{ij}$ ) are calculated for each vector, and averaged
  - Update is done with the averaged values ( $\Delta w_{ij}$ ) after the mini batch is calculated
  - Works efficiently when far away from minimum, but inaccurate close to minimum
  - Requires reducing learning rate



# Learning rate at SGD

- Sufficient conditions to guarantee convergence of SGD:

$$\sum_{k=1}^{\infty} \epsilon_k = \infty, \quad \text{and} \quad \sum_{k=1}^{\infty} \epsilon_k^2 < \infty.$$

- In practice:

$$\epsilon_k = (1 - \alpha)\epsilon_0 + \alpha\epsilon_{\tau} \quad \alpha = \frac{k}{\tau}$$

- After iteration  $\tau$ , it is common to leave  $\epsilon$  constant



# Stochastic Gradient Descent algorithm

---

**Algorithm** Stochastic gradient descent (SGD) update at training iteration  $k$

---

**Require:** Learning rate  $\epsilon_k$ .

**Require:** Initial parameter  $\theta$

**while** stopping criterion not met **do**

Sample a minibatch of  $m$  examples from the training set  $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$  with corresponding targets  $\mathbf{y}^{(i)}$ .

Compute gradient estimate:  $\hat{\mathbf{g}} \leftarrow +\frac{1}{m} \nabla_{\theta} \sum_i L(f(\mathbf{x}^{(i)}; \theta), \mathbf{y}^{(i)})$

Apply update:  $\theta \leftarrow \theta - \epsilon \hat{\mathbf{g}}$


**end while**

---

where:  $L$  is the cost function


# Stochastic Gradient Descent algorithm

- This very elongated quadratic function resembles a long canyon.
- Gradient descent wastes time repeatedly descending canyon walls, because they are the steepest feature.
- Because the step size is somewhat too large, it has a tendency to overshoot the bottom of the function and thus needs to descend the opposite canyon wall on the next iteration.



# Momentum I

- Introduced in 1964
- Physical analogy
- The idea is to simulate a unity weight mass
- It flows through on the surface of the error function
- Follows Newton's laws of dynamics
- Having  $v$  velocity
- Momentum correctly traverses the canyon lengthwise, while gradient steps waste time moving back and forth across the narrow axis of the canyon.





# Momentum II: velocity considerations

The update rule is given by:

$$\begin{aligned}\mathbf{v} &\leftarrow \alpha \mathbf{v} - \epsilon \nabla_{\boldsymbol{\theta}} \left( \frac{1}{m} \sum_{i=1}^m L(\mathbf{f}(\mathbf{x}^{(i)}; \boldsymbol{\theta}), \mathbf{y}^{(i)}) \right), \\ \boldsymbol{\theta} &\leftarrow \boldsymbol{\theta} + \mathbf{v}.\end{aligned}$$

The velocity  $\mathbf{v}$  accumulates the gradient elements  $\nabla_{\boldsymbol{\theta}} \left( \frac{1}{m} \sum_{i=1}^m L(\mathbf{f}(\mathbf{x}^{(i)}; \boldsymbol{\theta}), \mathbf{y}^{(i)}) \right)$ . The larger  $\alpha$  is relative to  $\epsilon$ , the more previous gradients affect the current direction.

Terminal velocity is applied when it finds descending gradient permanently:

$$\frac{\epsilon \|\mathbf{g}\|}{1 - \alpha}$$





# Momentum III

---

**Algorithm** Stochastic gradient descent (SGD) with momentum

---

**Require:** Learning rate  $\epsilon$ , momentum parameter  $\alpha$ .

**Require:** Initial parameter  $\theta$ , initial velocity  $v$ .

**while** stopping criterion not met **do**

Sample a minibatch of  $m$  examples from the training set  $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$  with corresponding targets  $\mathbf{y}^{(i)}$ .

Compute gradient estimate:  $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\theta} \sum_i L(f(\mathbf{x}^{(i)}; \theta), \mathbf{y}^{(i)})$

Compute velocity update:  $\mathbf{v} \leftarrow \alpha \mathbf{v} - \epsilon \mathbf{g}$

Apply update:  $\theta \leftarrow \theta + \mathbf{v}$

**end while**

---

# Most commonly applied gradient descent methods



- Algorithms with changing but not adaptive learning rate
  - Stochastic Gradient Descent algorithm
  - Momentum algorithm
- Algorithms with adaptive learning rate
  - AdaGrad algorithm
  - RMSProp algorithm
  - ADAM algorithm
- 2<sup>nd</sup> order algorithm
  - Newton algorithm



# AdaGrad algorithm

- The AdaGrad algorithm (2011) individually adapts the learning rates of all model parameters by scaling them inversely proportional to the square root of the sum of all of their historical squared values
- The parameters with the largest partial derivative of the loss have a correspondingly rapid decrease in their learning rate, while parameters with small partial derivatives have a relatively small decrease in their learning rate
- The net effect is greater progress in the more gently sloped directions of parameter space
- AdaGrad performs well for some but not all deep learning models

# AdaGrad algorithm



**Algorithm** The AdaGrad algorithm

*Remembers the  
entire history  
evenly*

**Require:** Global learning rate  $\epsilon$

**Require:** Initial parameter  $\theta$

**Require:** Small constant  $\delta$ , perhaps  $10^{-7}$ , for numerical stability

Initialize gradient accumulation variable  $\mathbf{r} = \mathbf{0}$

**while** stopping criterion not met **do**

Sample a minibatch of  $m$  examples from the training set  $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$  with corresponding targets  $\mathbf{y}^{(i)}$ .

Compute gradient:  $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\theta} \sum_i L(f(\mathbf{x}^{(i)}; \theta), \mathbf{y}^{(i)})$

Accumulate squared gradient:  $\mathbf{r} \leftarrow \mathbf{r} + \mathbf{g} \odot \mathbf{g}$

Compute update:  $\Delta \theta \leftarrow -\frac{\epsilon}{\delta + \sqrt{\mathbf{r}}} \odot \mathbf{g}$ . (Division and square root applied element-wise)

Apply update:  $\theta \leftarrow \theta + \Delta \theta$

**end while**

# RMSP algorithm



- The RMSProp algorithm (2012) modifies AdaGrad to perform better in the non-convex setting by changing the gradient accumulation into an exponentially weighted moving average
- AdaGrad shrinks the learning rate according to the entire history of the squared gradient and may have made the learning rate too small before arriving at such a convex structure
- RMSProp uses an exponentially decaying average to discard history from the extreme past so that it can converge rapidly after finding a convex bowl, as if it were an instance of the AdaGrad algorithm initialized within that bowl

# RMSP algorithm



**Algorithm** The RMSProp algorithm

*The closer parts of the history are counted more strongly.*

**Require:** Global learning rate  $\epsilon$ , decay rate  $\rho$ .

**Require:** Initial parameter  $\theta$

**Require:** Small constant  $\delta$ , usually  $10^{-6}$ , used to stabilize division by small numbers.

Initialize accumulation variables  $\mathbf{r} = 0$

**while** stopping criterion not met **do**

Sample a minibatch of  $m$  examples from the training set  $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$  with corresponding targets  $\mathbf{y}^{(i)}$ .

Compute gradient:  $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\theta} \sum_i L(f(\mathbf{x}^{(i)}; \theta), \mathbf{y}^{(i)})$

Accumulate squared gradient:  $\mathbf{r} \leftarrow \rho \mathbf{r} + (1 - \rho) \mathbf{g} \odot \mathbf{g}$

Compute parameter update:  $\Delta \theta = -\frac{\epsilon}{\sqrt{\delta + \mathbf{r}}} \odot \mathbf{g}$ . ( $\frac{1}{\sqrt{\delta + \mathbf{r}}}$  applied element-wise)

Apply update:  $\theta \leftarrow \theta + \Delta \theta$

**end while**

# ADAM algorithm (2014)



- The name “Adam” derives from the phrase “adaptive moments.”
- In the context of the earlier algorithms, it is perhaps best seen as a variant on the combination of RMSProp and momentum with a few important distinctions.
- in Adam, momentum is incorporated directly as an estimate of the first order moment (with exponential weighting) of the gradient.
- Adam includes bias corrections to the estimates of both the first-order moments (the momentum term) and the (uncentered) second-order moments to account for their initialization at the origin

# ADAM algorithm

$s$  estimates the  
gradient from the  
history (moment)

$r$  estimates the  
curvature of the  
gradient

Both of them are  
biased to reduce  
anomalies at the  
initialization

---

**Algorithm** The Adam algorithm

---

**Require:** Step size  $\epsilon$  (Suggested default: 0.001)

**Require:** Exponential decay rates for moment estimates,  $\rho_1$  and  $\rho_2$  in  $[0, 1)$ .  
(Suggested defaults: 0.9 and 0.999 respectively)

**Require:** Small constant  $\delta$  used for numerical stabilization. (Suggested default:  $10^{-8}$ )

**Require:** Initial parameters  $\theta$

Initialize 1st and 2nd moment variables  $s = \mathbf{0}$ ,  $r = \mathbf{0}$

Initialize time step  $t = 0$

**while** stopping criterion not met **do**

Sample a minibatch of  $m$  examples from the training set  $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$  with  
corresponding targets  $\mathbf{y}^{(i)}$ .

Compute gradient:  $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\theta} \sum_i L(f(\mathbf{x}^{(i)}; \theta), \mathbf{y}^{(i)})$

$t \leftarrow t + 1$

Update biased first moment estimate:  $\mathbf{s} \leftarrow \rho_1 \mathbf{s} + (1 - \rho_1) \mathbf{g}$

Update biased second moment estimate:  $\mathbf{r} \leftarrow \rho_2 \mathbf{r} + (1 - \rho_2) \mathbf{g} \odot \mathbf{g}$

Correct bias in first moment:  $\hat{\mathbf{s}} \leftarrow \frac{\mathbf{s}}{1 - \rho_1^t}$

Correct bias in second moment:  $\hat{\mathbf{r}} \leftarrow \frac{\mathbf{r}}{1 - \rho_2^t}$

Compute update:  $\Delta \theta = -\epsilon \frac{\hat{\mathbf{s}}}{\sqrt{\hat{\mathbf{r}} + \delta}}$  (operations applied element-wise)

Apply update:  $\theta \leftarrow \theta + \Delta \theta$

**end while**

---



# Most commonly applied gradient descent methods



- Algorithms with changing but not adaptive learning rate
  - Stochastic Gradient Descent algorithm
  - Momentum algorithm
- Algorithms with adaptive learning rate
  - AdaGrad algorithm
  - RMSProp algorithm
  - ADAM algorithm
- 2<sup>nd</sup> order algorithm
  - Newton algorithm



# Newton's algorithm

---

**Algorithm** Newton's method with objective  $J(\boldsymbol{\theta}) = \frac{1}{m} \sum_{i=1}^m L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$ .

---

**Require:** Initial parameter  $\boldsymbol{\theta}_0$

**Require:** Training set of  $m$  examples

**while** stopping criterion not met **do**

    Compute gradient:  $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), \mathbf{y}^{(i)})$

    Compute Hessian:  $\mathbf{H} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}}^2 \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), \mathbf{y}^{(i)})$

    Compute Hessian inverse:  $\mathbf{H}^{-1}$

    Compute update:  $\Delta\boldsymbol{\theta} = -\mathbf{H}^{-1} \mathbf{g}$

    Apply update:  $\boldsymbol{\theta} = \boldsymbol{\theta} + \Delta\boldsymbol{\theta}$

**end while**

---



# Back propagation

- We have seen last time how to calculate the gradient in a multilayer fully connected network using back propagation
  - The introduced method was based on gradient descent method
- However, being able to calculate gradient, we might select any of the above methods, which might lead to faster convergence