

Training Neural Networks

Milan Straka

 March 2, 2020



Charles University in Prague
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics



unless otherwise stated

An *estimator* is a rule for computing an estimate of a given value, often an expectation of some random value(s).

The *bias* of an estimator is the difference of the expected value of the estimator and the true value being estimated.

If the bias is zero, we call the estimator *unbiased*, otherwise we call it *biased*.

If we have a sequence of estimates, it also might happen that the bias converges to zero. Consider the well known sample estimate of variance. Given x_1, \dots, x_n independent and identically distributed random variables, we might estimate mean and variance as

$$\hat{\mu} = \frac{1}{n} \sum_i x_i, \quad \hat{\sigma}^2 = \frac{1}{n} \sum_i (x_i - \hat{\mu})^2.$$

Such an estimate is biased, because $\mathbb{E}[\hat{\sigma}^2] = (1 - \frac{1}{n})\sigma^2$, but the bias converges to zero with increasing n .

Also, an unbiased estimator does not necessarily have small variance – in some cases it can have large variance, so a biased estimator with smaller variance might be preferred.

We usually have a **training set**, which is assumed to consist of examples generated independently from a **data generating distribution**.

The goal of *optimization* is to match the training set as well as possible.

However, the main goal of *machine learning* is to perform well on *previously unseen* data, so called **generalization error** or **test error**. We typically estimate the generalization error using a **test set** of examples independent of the training set, but generated by the same data generating distribution.

Challenges in machine learning:

- *underfitting*
- *overfitting*




Figure 5.2, page 113 of Deep Learning Book, <http://deeplearningbook.org>

We can control whether a model underfits or overfits by modifying its *capacity*.

- representational capacity
- effective capacity




Figure 5.3, page 115 of Deep Learning Book, <http://deeplearningbook.org>

The **No free lunch theorem** (Wolpert, 1996) states that averaging over *all possible* data distributions, every classification algorithm achieves the same *overall* error when processing unseen examples. In a sense, no machine learning algorithm is *universally* better than others.

Any change in a machine learning algorithm that is designed to *reduce generalization error* but not necessarily its training error is called **regularization**.

L_2 regularization (also called weighted decay) penalizes models with large weights (i.e., penalty of $\|\boldsymbol{\theta}\|^2$).




Figure 5.5, page 119 of Deep Learning Book, <http://deeplearningbook.org>

Hyperparameters are not adapted by the learning algorithm itself.

Usually a **validation set** or **development set** is used to estimate the generalization error, allowing to update hyperparameters accordingly.

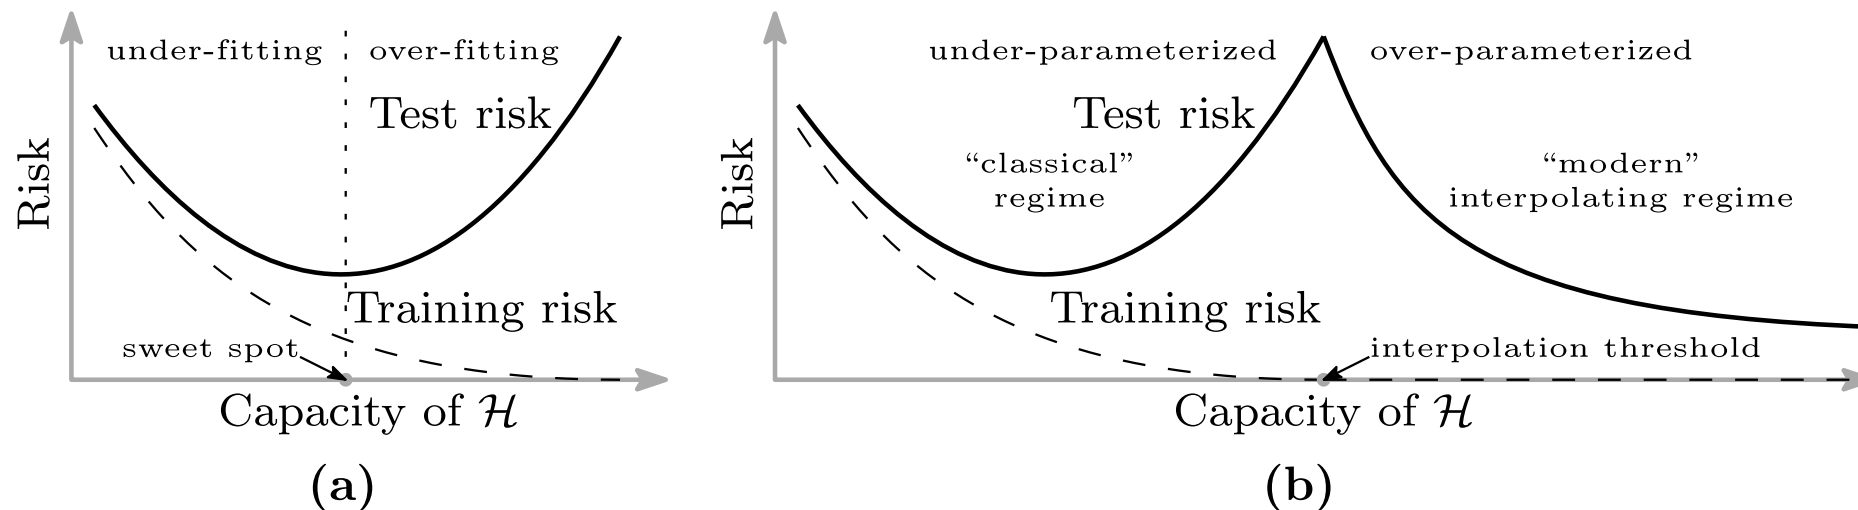


Figure 1: **Curves for training risk (dashed line) and test risk (solid line).** (a) The classical *U-shaped risk curve* arising from the bias-variance trade-off. (b) The *double descent risk curve*, which incorporates the U-shaped risk curve (i.e., the “classical” regime) together with the observed behavior from using high capacity function classes (i.e., the “modern” interpolating regime), separated by the interpolation threshold. The predictors to the right of the interpolation threshold have zero training risk.

Figure 1 of paper "Reconciling modern machine learning practice and the bias-variance trade-off", <https://arxiv.org/abs/1812.11118>.

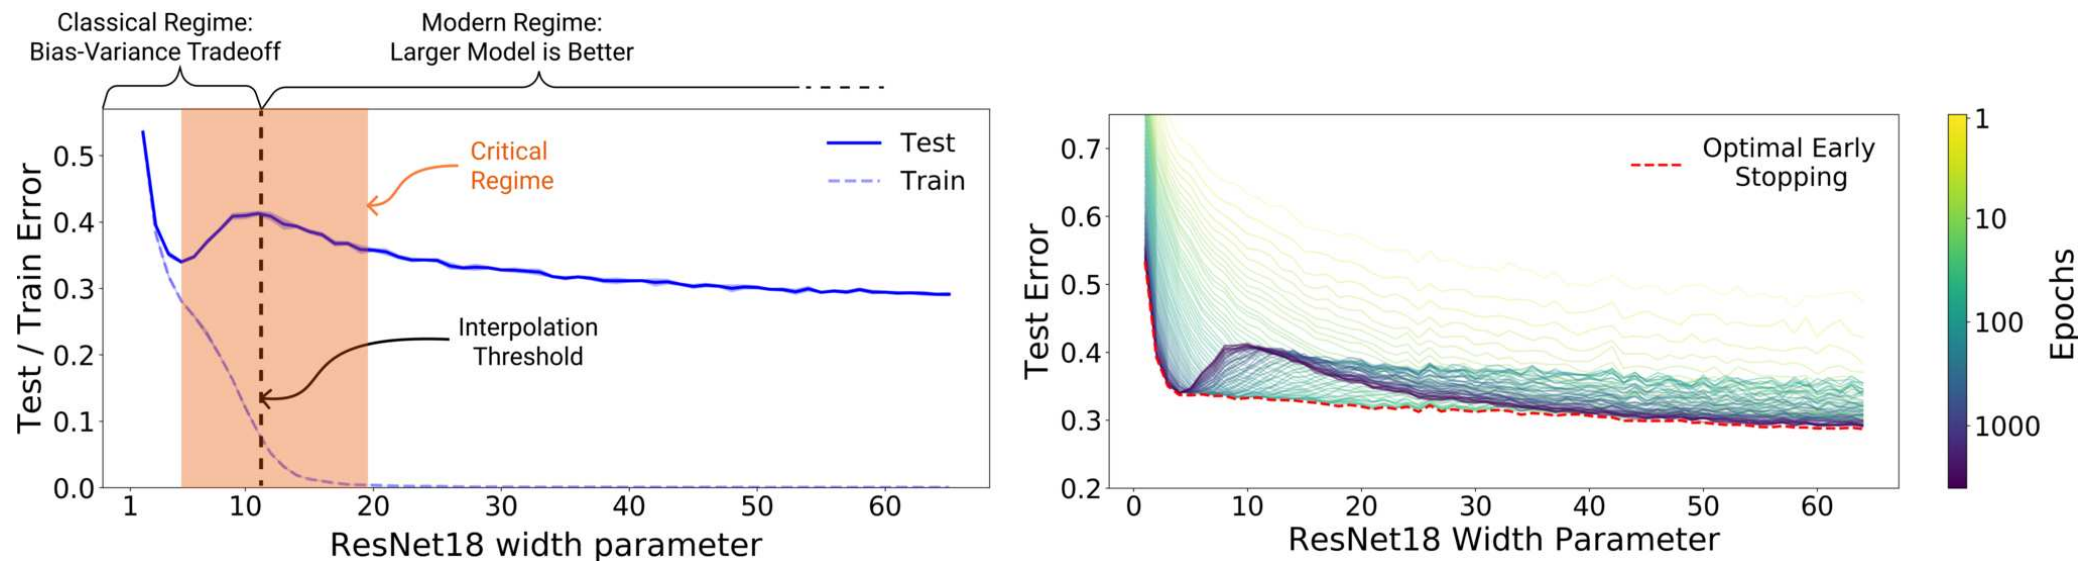






Figure 1: **Left:** Train and test error as a function of model size, for ResNet18s of varying width on CIFAR-10 with 15% label noise. **Right:** Test error, shown for varying train epochs. All models trained using Adam for 4K epochs. The largest model (width 64) corresponds to standard ResNet18.

Figure 1 of the paper "Deep Double Descent: Where Bigger Models and More Data Hurt", <https://arxiv.org/abs/1912.02292>.

Why do Neural Networks Generalize so Well



(a) **CIFAR-100.** There is a peak in test error even with no label noise.



(b) **CIFAR-10.** There is a “plateau” in test error around the interpolation point with no label noise, which develops into a peak for added label noise.

Figure 4 of the paper "Deep Double Descent: Where Bigger Models and More Data Hurt", <https://arxiv.org/abs/1912.02292>.

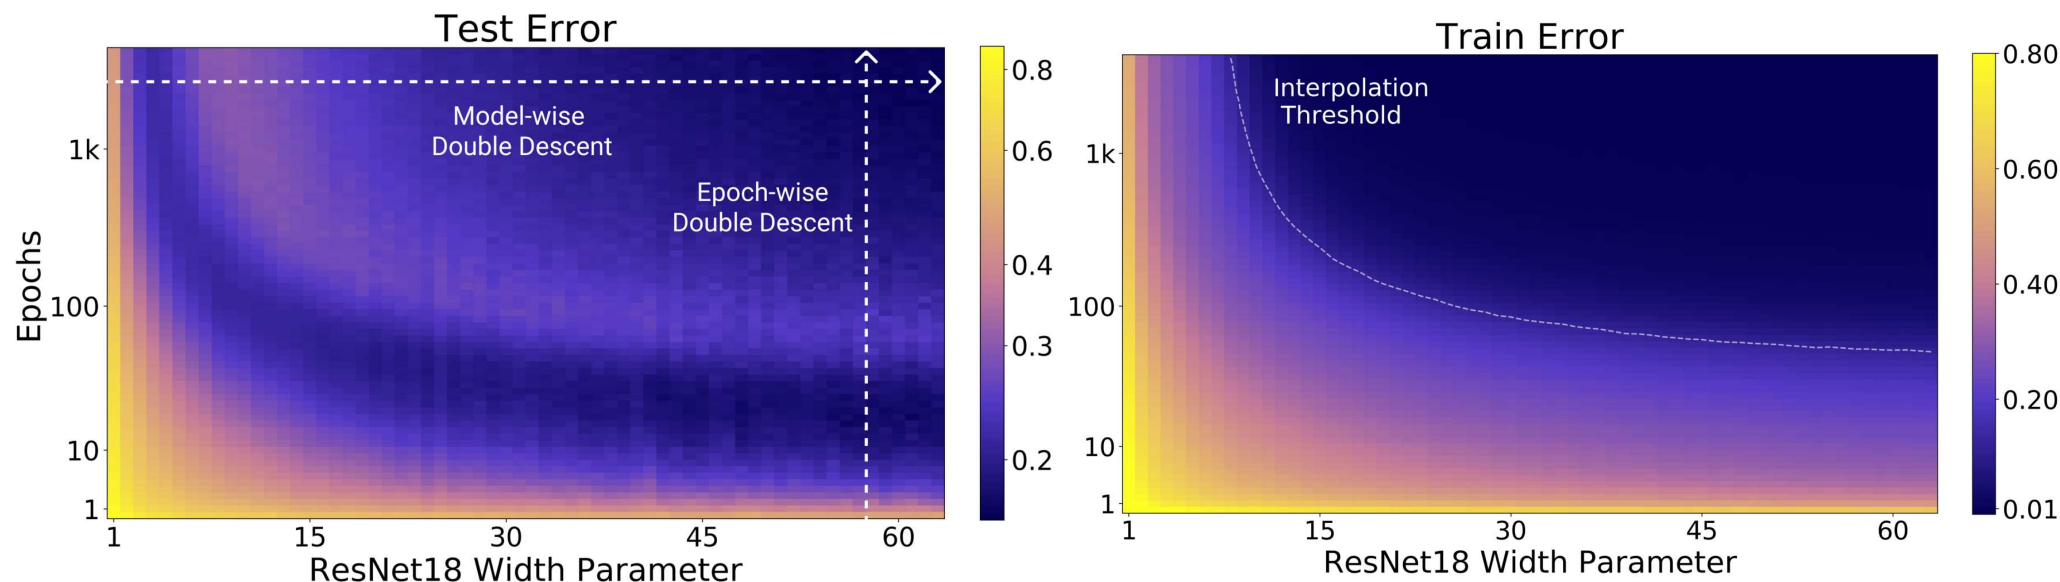


Figure 2: **Left:** Test error as a function of model size and train epochs. The horizontal line corresponds to model-wise double descent—varying model size while training for as long as possible. The vertical line corresponds to epoch-wise double descent, with test error undergoing double-descent as train time increases. **Right** Train error of the corresponding models. All models are Resnet18s trained on CIFAR-10 with 15% label noise, data-augmentation, and Adam for up to 4K epochs.

Figure 2 of the paper "Deep Double Descent: Where Bigger Models and More Data Hurt", <https://arxiv.org/abs/1912.02292>.

A model is usually trained in order to minimize the *loss* on the training data.

Assuming that a model computes $f(\mathbf{x}; \boldsymbol{\theta})$ using parameters $\boldsymbol{\theta}$, the *mean square error* is computed as

$$\frac{1}{m} \sum_{i=1}^m \left(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}) - y^{(i)} \right)^2.$$

A common principle used to design loss functions is the *maximum likelihood principle*.

Maximum Likelihood Estimation

Let $\mathbb{X} = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(m)}\}$ be training data drawn independently from the data-generating distribution p_{data} . We denote the empirical data distribution as \hat{p}_{data} .

Let $p_{\text{model}}(\mathbf{x}; \boldsymbol{\theta})$ be a family of distributions. The *maximum likelihood estimation* of $\boldsymbol{\theta}$ is:

$$\begin{aligned}
 \boldsymbol{\theta}_{\text{ML}} &= \arg \max_{\boldsymbol{\theta}} p_{\text{model}}(\mathbb{X}; \boldsymbol{\theta}) \\
 &= \arg \max_{\boldsymbol{\theta}} \prod_{i=1}^m p_{\text{model}}(\mathbf{x}^{(i)}; \boldsymbol{\theta}) \\
 &= \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^m -\log p_{\text{model}}(\mathbf{x}^{(i)}; \boldsymbol{\theta}) \\
 &= \arg \min_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{x} \sim \hat{p}_{\text{data}}} [-\log p_{\text{model}}(\mathbf{x}; \boldsymbol{\theta})] \\
 &= \arg \min_{\boldsymbol{\theta}} H(\hat{p}_{\text{data}}, p_{\text{model}}(\mathbf{x}; \boldsymbol{\theta})) \\
 &= \arg \min_{\boldsymbol{\theta}} D_{\text{KL}}(\hat{p}_{\text{data}} || p_{\text{model}}(\mathbf{x}; \boldsymbol{\theta})) + H(\hat{p}_{\text{data}})
 \end{aligned}$$

MLE can be easily generalized to a conditional case, where our goal is to predict y given \mathbf{x} :

$$\begin{aligned}\boldsymbol{\theta}_{\text{ML}} &= \arg \max_{\boldsymbol{\theta}} p_{\text{model}}(\mathbb{Y}|\mathbb{X}; \boldsymbol{\theta}) \\ &= \arg \max_{\boldsymbol{\theta}} \prod_{i=1}^m p_{\text{model}}(y^{(i)}|\mathbf{x}^{(i)}; \boldsymbol{\theta}) \\ &= \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^m -\log p_{\text{model}}(y^{(i)}|\mathbf{x}^{(i)}; \boldsymbol{\theta})\end{aligned}$$

The resulting *loss function* is called *negative log likelihood*, or *cross-entropy* or *Kullback-Leibler divergence*.

Assume that the true data generating distribution p_{data} lies within the model family $p_{\text{model}}(\cdot; \boldsymbol{\theta})$, and assume there exists a unique $\boldsymbol{\theta}_{p_{\text{data}}}$ such that $p_{\text{data}} = p_{\text{model}}(\cdot; \boldsymbol{\theta}_{p_{\text{data}}})$.

- MLE is a *consistent* estimator. If we denote $\boldsymbol{\theta}_m$ to be the parameters found by MLE for a training set with m examples generated by the data generating distribution, then $\boldsymbol{\theta}_m$ converges in probability to $\boldsymbol{\theta}_{p_{\text{data}}}$.

Formally, for any $\varepsilon > 0$, $P(\|\boldsymbol{\theta}_m - \boldsymbol{\theta}_{p_{\text{data}}}\| > \varepsilon) \rightarrow 0$ as $m \rightarrow \infty$.

- MLE is in a sense the most *statistically efficient*. For any consistent estimator, we might consider the average distance of $\boldsymbol{\theta}_m$ and $\boldsymbol{\theta}_{p_{\text{data}}}$, formally $\mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_m \sim p_{\text{data}}} [\|\boldsymbol{\theta}_m - \boldsymbol{\theta}_{p_{\text{data}}}\|_2^2]$. It can be shown (Rao 1945, Cramér 1946) that no consistent estimator has lower mean squared error than the maximum likelihood estimator.

Therefore, for reasons of consistency and efficiency, maximum likelihood is often considered the preferred estimator for machine learning.

Mean Square Error as MLE

Assume our goal is to perform regression, i.e., to predict $p(y|\mathbf{x})$ for $y \in \mathbb{R}$.

Let $\hat{y}(\mathbf{x}; \boldsymbol{\theta})$ give a prediction of the mean of y .

We define $p(y|\mathbf{x})$ as $\mathbb{N}(y; \hat{y}(\mathbf{x}; \boldsymbol{\theta}), \sigma^2)$ for a given fixed σ^2 . Then:

$$\begin{aligned} \arg \max_{\boldsymbol{\theta}} p(y|\mathbf{x}; \boldsymbol{\theta}) &= \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^m -\log p(y^{(i)}|\mathbf{x}^{(i)}; \boldsymbol{\theta}) \\ &= \arg \min_{\boldsymbol{\theta}} -\sum_{i=1}^m \log \sqrt{\frac{1}{2\pi\sigma^2}} e^{-\frac{(y^{(i)} - \hat{y}(\mathbf{x}^{(i)}; \boldsymbol{\theta}))^2}{2\sigma^2}} \\ &= \arg \min_{\boldsymbol{\theta}} -m \log(2\pi\sigma^2)^{-1/2} - \sum_{i=1}^m -\frac{(y^{(i)} - \hat{y}(\mathbf{x}^{(i)}; \boldsymbol{\theta}))^2}{2\sigma^2} \\ &= \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^m \frac{(y^{(i)} - \hat{y}(\mathbf{x}^{(i)}; \boldsymbol{\theta}))^2}{2\sigma^2} = \arg \min_{\boldsymbol{\theta}} \frac{1}{m} \sum_{i=1}^m (y^{(i)} - \hat{y}(\mathbf{x}^{(i)}; \boldsymbol{\theta}))^2. \end{aligned}$$

Let a model compute $f(\mathbf{x}; \boldsymbol{\theta})$ using parameters $\boldsymbol{\theta}$, and for a given loss function L denote

$$J(\boldsymbol{\theta}) = \mathbb{E}_{(\mathbf{x}, y) \sim \hat{p}_{\text{data}}} L(f(\mathbf{x}; \boldsymbol{\theta}), y).$$

Assuming we are minimizing an error function

$$\arg \min_{\boldsymbol{\theta}} J(\boldsymbol{\theta})$$

we may use *gradient descent*:

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta})$$

The constant α is called a *learning rate* and specifies the “length” of a step we perform in every iteration of the gradient descent.




Figure 4.1, page 83 of Deep Learning Book, <http://deeplearningbook.org>

(Regular) Gradient Descent

We use all training data to compute $J(\boldsymbol{\theta})$.

Online (or Stochastic) Gradient Descent

We estimate the expectation in $J(\boldsymbol{\theta})$ using a single randomly sampled example from the training data. Such an estimate is unbiased, but very noisy.

$$J(\boldsymbol{\theta}) \approx L(f(\mathbf{x}; \boldsymbol{\theta}), y) \text{ for randomly chosen } (\mathbf{x}, y) \text{ from } \hat{p}_{\text{data}}.$$

Minibatch SGD

The minibatch SGD is a trade-off between gradient descent and SGD – the expectation in $J(\boldsymbol{\theta})$ is estimated using m random independent examples from the training data.

$$J(\boldsymbol{\theta}) \approx \frac{1}{m} \sum_{i=1}^m L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)}) \text{ for randomly chosen } (\mathbf{x}^{(i)}, y^{(i)}) \text{ from } \hat{p}_{\text{data}}.$$

It can be proven (under some reasonable conditions; see Robbins and Monro algorithm, 1951) that if the loss function is convex and continuous, then SGD converges to the unique optimum almost surely if the sequence of learning rates α_i fulfills the following conditions:


$$\sum_i \alpha_i = \infty, \quad \sum_i \alpha_i^2 < \infty.$$

For non-convex loss functions, we can get guarantees of converging to a *local* optimum only. Note that finding a global minimum of an arbitrary function is *at least NP-hard*.


Stochastic Gradient Descent Convergence

Convex functions mentioned on the previous slide are such that for x_1, x_2 and real $0 \leq t \leq 1$,

$$f(tx_1 + (1 - t)x_2) \leq tf(x_1) + (1 - t)f(x_2).$$



<https://upload.wikimedia.org/wikipedia/commons/c/c7/ConvexFunction.svg>



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

A twice-differentiable function is convex iff its second derivative is always non-negative.

A local minimum of a convex function is always the unique global minimum.

Well-known examples of convex functions are x^2 , e^x and $-\log x$.

In 2018, there have been several improvements:

- Under some models with high capacity, it can be proven that SGD will reach global optimum by showing it will reach zero training error.
- Neural networks can be easily modified so that the augmented version has no local minimums. Therefore, if such a network converges, it converged to a global minimum. However, the training process can still fail to converge by increasing the size of the parameters $||\boldsymbol{\theta}||$ beyond any limit.

Loss Function Visualization

Visualization of loss function of ResNet-56 (0.85 million parameters) with/without skip connections:




Figure 1 of paper "Visualizing the Loss Landscape of Neural Nets", <https://arxiv.org/abs/1712.09913>.

Loss Function Visualization

Visualization of loss function of ResNet-110 without skip connections and DenseNet-121.




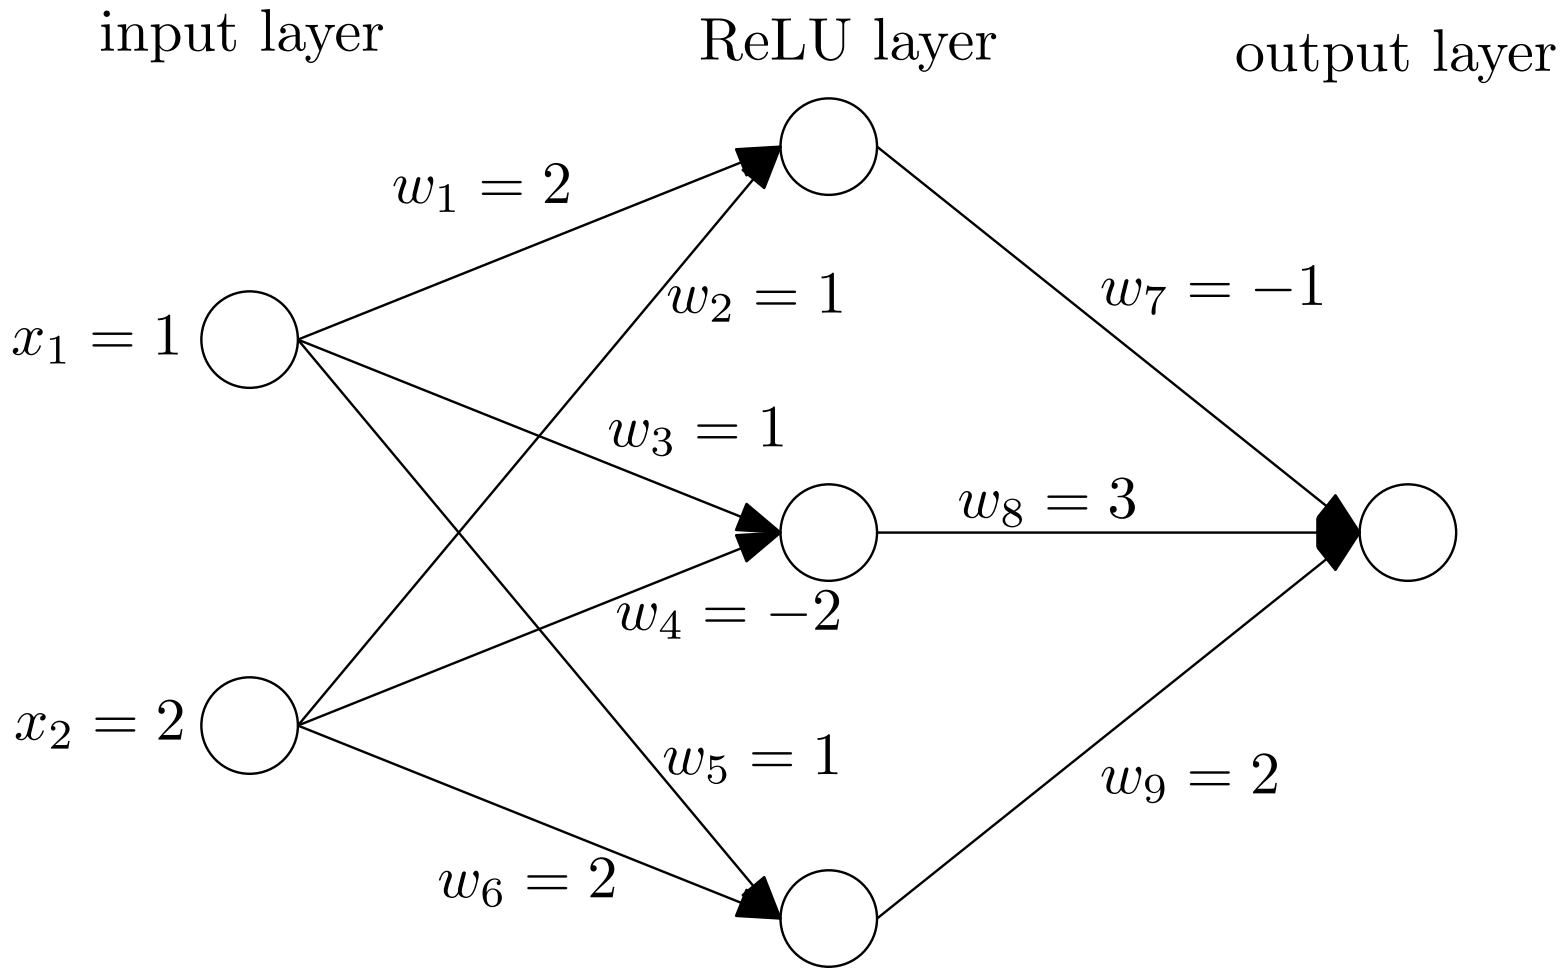


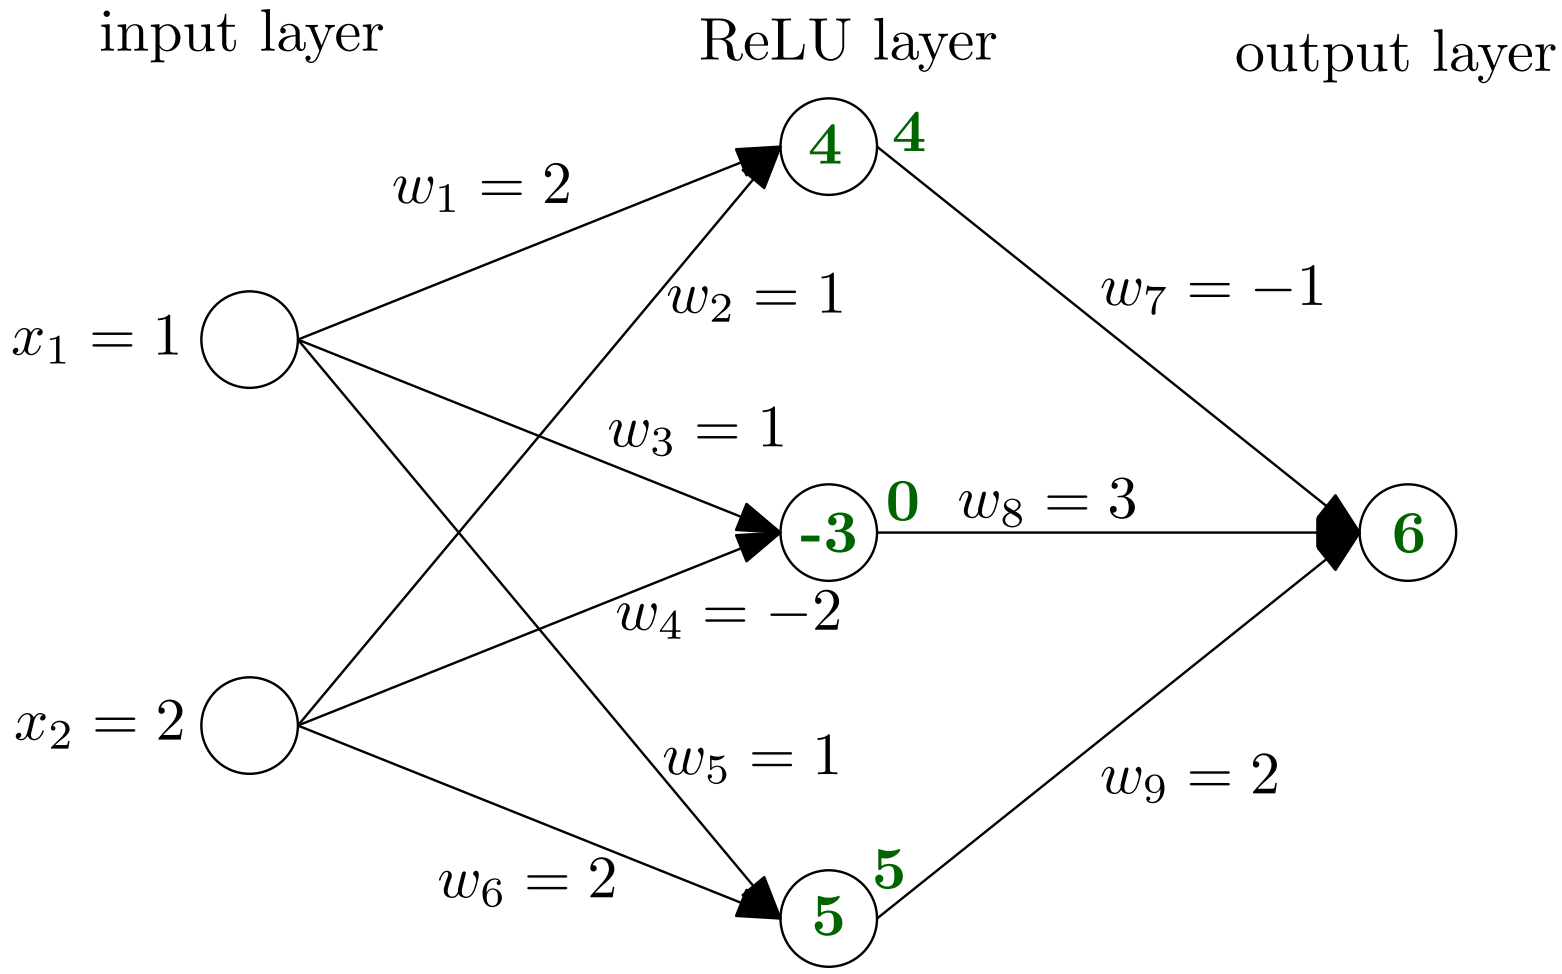
Figure 4 of paper "Visualizing the Loss Landscape of Neural Nets", <https://arxiv.org/abs/1712.09913>.

Assume we want to compute partial derivatives of a given loss function J and let $\frac{\partial J}{\partial z}$ be known.

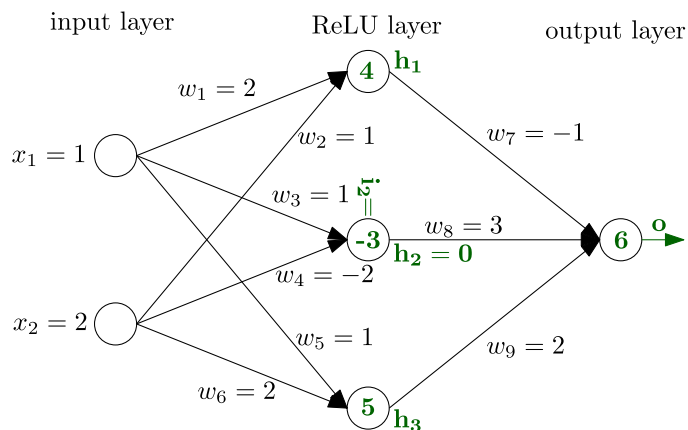


$$\frac{\partial J}{\partial y_i} = \frac{\partial J}{\partial z} \frac{\partial z}{\partial y_i} = \frac{\partial J}{\partial z} \frac{\partial g(\mathbf{y})}{\partial y_i}$$
$$\frac{\partial J}{\partial x_i} = \frac{\partial J}{\partial z} \frac{\partial z}{\partial y_i} \frac{\partial y_i}{\partial x_i} = \frac{\partial J}{\partial z} \frac{\partial g(\mathbf{y})}{\partial y_i} \frac{\partial f(\mathbf{x}_i)}{\partial x_i}$$





Backpropagation Example



$$\frac{\partial L}{\partial o} = 2(\text{output} - \text{gold}) = 6$$

$$\frac{\partial L}{\partial w_7} = \frac{\partial L}{\partial o} \frac{\partial o}{\partial w_7} = \frac{\partial L}{\partial o} h_1 = 24$$

$$\frac{\partial L}{\partial w_8} = \frac{\partial L}{\partial o} \frac{\partial o}{\partial w_8} = \frac{\partial L}{\partial o} h_2 = 0$$

$$\frac{\partial L}{\partial w_9} = \frac{\partial L}{\partial o} \frac{\partial o}{\partial w_9} = \frac{\partial L}{\partial o} h_3 = 30$$

$$\frac{\partial L}{\partial w_1} = \frac{\partial L}{\partial i_1} \frac{\partial i_1}{\partial w_1} = \frac{\partial L}{\partial i_1} x_1 = -6$$

$$\frac{\partial L}{\partial w_2} = \frac{\partial L}{\partial i_1} \frac{\partial i_1}{\partial w_2} = \frac{\partial L}{\partial i_1} x_2 = -12$$

$$\frac{\partial L}{\partial w_3} = \frac{\partial L}{\partial i_2} \frac{\partial i_2}{\partial w_3} = \frac{\partial L}{\partial i_2} x_1 = 0$$

$$\frac{\partial L}{\partial w_4} = \frac{\partial L}{\partial i_2} \frac{\partial i_2}{\partial w_4} = \frac{\partial L}{\partial i_2} x_2 = 0$$

$$\frac{\partial L}{\partial w_5} = \frac{\partial L}{\partial i_3} \frac{\partial i_3}{\partial w_5} = \frac{\partial L}{\partial i_3} x_1 = 12$$

$$\frac{\partial L}{\partial w_6} = \frac{\partial L}{\partial i_3} \frac{\partial i_3}{\partial w_6} = \frac{\partial L}{\partial i_3} x_2 = 24$$

$$\frac{\partial L}{\partial h_1} = \frac{\partial L}{\partial o} \frac{\partial o}{\partial h_1} = \frac{\partial L}{\partial o} w_7 = -6$$

$$\frac{\partial L}{\partial h_2} = \frac{\partial L}{\partial o} \frac{\partial o}{\partial h_2} = \frac{\partial L}{\partial o} w_8 = 18$$

$$\frac{\partial L}{\partial h_3} = \frac{\partial L}{\partial o} \frac{\partial o}{\partial h_3} = \frac{\partial L}{\partial o} w_9 = 12$$

$$\frac{\partial L}{\partial i_1} = \frac{\partial L}{\partial h_1} \frac{\partial h_1}{\partial i_1} = \frac{\partial L}{\partial h_1} 1 = -6$$

$$\frac{\partial L}{\partial i_2} = \frac{\partial L}{\partial h_2} \frac{\partial h_2}{\partial i_2} = \frac{\partial L}{\partial h_2} 0 = 0$$

$$\frac{\partial L}{\partial i_3} = \frac{\partial L}{\partial h_3} \frac{\partial h_3}{\partial i_3} = \frac{\partial L}{\partial h_3} 1 = 12$$

$$\frac{\partial L}{\partial x_1} = \sum_j \frac{\partial L}{\partial i_j} \frac{\partial i_j}{\partial x_1} = 0$$

$$\frac{\partial L}{\partial x_2} = \sum_j \frac{\partial L}{\partial i_j} \frac{\partial i_j}{\partial x_2} = 18$$

This is meant to be frightening – you do **not** do this manually when training.

Forward Propagation

Input: Network with nodes $u^{(1)}, u^{(2)}, \dots, u^{(n)}$ numbered in topological order.

Each node's value is computed as $u^{(i)} = f^{(i)}(\mathbb{A}^{(i)})$ for $\mathbb{A}^{(i)}$ being a set of values of the predecessors $P(u^{(i)})$ of $u^{(i)}$.

Output: Value of $u^{(n)}$.

- For $i = 1, \dots, n$:
 - $\mathbb{A}^{(i)} \leftarrow \{u^{(j)} \mid j \in P(u^{(i)})\}$
 - $u^{(i)} \leftarrow f^{(i)}(\mathbb{A}^{(i)})$
- Return $u^{(n)}$


Simple Variant of Backpropagation

Input: The network as in the Forward propagation algorithm.

Output: Partial derivatives $g^{(i)} = \frac{\partial u^{(n)}}{\partial u^{(i)}}$ of $u^{(n)}$ with respect to all $u^{(i)}$.

- Run forward propagation to compute all $u^{(i)}$
- $g^{(n)} = 1$
- For $i = n - 1, \dots, 1$:
 - $g^{(i)} \leftarrow \sum_{j:i \in P(u^{(j)})} g^{(j)} \frac{\partial u^{(j)}}{\partial u^{(i)}}$
- Return g

In practice, we do not usually represent networks as collections of scalar nodes; instead we represent them as collections of tensor functions – most usually functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Then $\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}}$ is a Jacobian. However, the backpropagation algorithm is analogous.



There is a weight on each edge, and an activation function f is performed on the hidden layers, and optionally also on the output layer.

$$h_i = f \left(\sum_j w_{i,j} x_j + b_i \right)$$

If the network is composed of layers, we can use matrix notation and write:

$$\mathbf{h} = f(\mathbf{W}\mathbf{x} + \mathbf{b})$$




Figure 5.1 of Pattern Recognition and Machine Learning.

Hidden Layers Derivatives

- σ :

$$\frac{d\sigma(x)}{dx} = \sigma(x) \cdot (1 - \sigma(x))$$

- \tanh :

$$\frac{d \tanh(x)}{dx} = 1 - \tanh(x)^2$$

- ReLU:

$$\frac{d \text{ReLU}(x)}{dx} = \begin{cases} 1 & \text{if } x > 0 \\ \text{NaN} & \text{if } x = 0 \\ 0 & \text{if } x < 0 \end{cases}$$

Stochastic Gradient Descent (SGD) Algorithm

Input: NN computing function $f(\mathbf{x}; \boldsymbol{\theta})$ with initial value of parameters $\boldsymbol{\theta}$.

Input: Learning rate α .

Output: Updated parameters $\boldsymbol{\theta}$.

- Repeat until stopping criterion is met:
 - Sample a minibatch of m training examples $(\mathbf{x}^{(i)}, y^{(i)})$
 - $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \mathbf{g}$

SGD With Momentum

Input: NN computing function $f(\mathbf{x}; \boldsymbol{\theta})$ with initial value of parameters $\boldsymbol{\theta}$.

Input: Learning rate α , momentum β .

Output: Updated parameters $\boldsymbol{\theta}$.

- Repeat until stopping criterion is met:
 - Sample a minibatch of m training examples $(\mathbf{x}^{(i)}, y^{(i)})$
 - $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$
 - $\mathbf{v} \leftarrow \beta \mathbf{v} - \alpha \mathbf{g}$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \mathbf{v}$




Figure 8.5, page 297 of Deep Learning Book, <http://deeplearningbook.org>



SGD With Nesterov Momentum

Input: NN computing function $f(\mathbf{x}; \boldsymbol{\theta})$ with initial value of parameters $\boldsymbol{\theta}$.

Input: Learning rate α , momentum β .

Output: Updated parameters $\boldsymbol{\theta}$.

- Repeat until stopping criterion is met:
 - Sample a minibatch of m training examples $(\mathbf{x}^{(i)}, y^{(i)})$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \beta \mathbf{v}$
 - $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$
 - $\mathbf{v} \leftarrow \beta \mathbf{v} - \alpha \mathbf{g}$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \mathbf{g}$



<https://github.com/cs231n/cs231n.github.io/blob/master/assets/nn3/nesterov.jpeg>

AdaGrad (2011)

Input: NN computing function $f(\mathbf{x}; \boldsymbol{\theta})$ with initial value of parameters $\boldsymbol{\theta}$.

Input: Learning rate α , constant ε (usually 10^{-8}).

Output: Updated parameters $\boldsymbol{\theta}$.

- Repeat until stopping criterion is met:
 - Sample a minibatch of m training examples $(\mathbf{x}^{(i)}, y^{(i)})$
 - $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$
 - $\mathbf{r} \leftarrow \mathbf{r} + \mathbf{g}^2$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \frac{\alpha}{\sqrt{\mathbf{r} + \varepsilon}} \mathbf{g}$

AdaGrad has favourable convergence properties (being faster than regular SGD) for convex loss landscapes. In this settings, gradients converge to zero reasonably fast.

However, for non-convex losses, gradients can stay quite large for a long time. In that case, the algorithm behaves as if decreasing learning rate by a factor of $1/\sqrt{t}$, because if each

$$\mathbf{g} \approx \mathbf{g}_0,$$

then after t steps

$$\mathbf{r} \approx t \cdot \mathbf{g}_0^2$$

and therefore

$$\frac{\alpha}{\sqrt{\mathbf{r} + \varepsilon}} \approx \frac{\alpha/\sqrt{t}}{\sqrt{\mathbf{g}_0^2 + \varepsilon/t}}.$$

RMSProp (2012)

Input: NN computing function $f(\mathbf{x}; \boldsymbol{\theta})$ with initial value of parameters $\boldsymbol{\theta}$.

Input: Learning rate α , momentum β (usually 0.9), constant ε (usually 10^{-8}).

Output: Updated parameters $\boldsymbol{\theta}$.

- Repeat until stopping criterion is met:
 - Sample a minibatch of m training examples $(\mathbf{x}^{(i)}, y^{(i)})$
 - $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$
 - $\mathbf{r} \leftarrow \beta \mathbf{r} + (1 - \beta) \mathbf{g}^2$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \frac{\alpha}{\sqrt{\mathbf{r} + \varepsilon}} \mathbf{g}$

However, after first step, $\mathbf{r} = (1 - \beta) \mathbf{g}^2$, which for default $\beta = 0.9$ is

$$\mathbf{r} = 0.1 \mathbf{g}^2,$$

a biased estimate (but the bias converges to zero exponentially fast).

Adam (2014)

Input: NN computing function $f(\mathbf{x}; \boldsymbol{\theta})$ with initial value of parameters $\boldsymbol{\theta}$.

Input: Learning rate α (default 0.001), constant ε (usually 10^{-8}).

Input: Momentum β_1 (default 0.9), momentum β_2 (default 0.999).

Output: Updated parameters $\boldsymbol{\theta}$.

- $\mathbf{s} \leftarrow 0, \mathbf{r} \leftarrow 0, t \leftarrow 0$
- Repeat until stopping criterion is met:
 - Sample a minibatch of m training examples $(\mathbf{x}^{(i)}, y^{(i)})$
 - $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$
 - $t \leftarrow t + 1$
 - $\mathbf{s} \leftarrow \beta_1 \mathbf{s} + (1 - \beta_1) \mathbf{g}$ (*biased first moment estimate*)
 - $\mathbf{r} \leftarrow \beta_2 \mathbf{r} + (1 - \beta_2) \mathbf{g}^2$ (*biased second moment estimate*)
 - $\hat{\mathbf{s}} \leftarrow \mathbf{s} / (1 - \beta_1^t), \hat{\mathbf{r}} \leftarrow \mathbf{r} / (1 - \beta_2^t)$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \frac{\alpha}{\sqrt{\hat{\mathbf{r}} + \varepsilon}} \hat{\mathbf{s}}$

Adam (2014)

Input: NN computing function $f(\mathbf{x}; \boldsymbol{\theta})$ with initial value of parameters $\boldsymbol{\theta}$.

Input: Learning rate α (default 0.001), constant ε (usually 10^{-8}).

Input: Momentum β_1 (default 0.9), momentum β_2 (default 0.999).

Output: Updated parameters $\boldsymbol{\theta}$.

- $\mathbf{s} \leftarrow 0, \mathbf{r} \leftarrow 0, t \leftarrow 0$
- Repeat until stopping criterion is met:
 - Sample a minibatch of m training examples $(\mathbf{x}^{(i)}, y^{(i)})$
 - $\mathbf{g} \leftarrow \frac{1}{m} \nabla_{\boldsymbol{\theta}} \sum_i L(f(\mathbf{x}^{(i)}; \boldsymbol{\theta}), y^{(i)})$
 - $t \leftarrow t + 1$
 - $\mathbf{s} \leftarrow \beta_1 \mathbf{s} + (1 - \beta_1) \mathbf{g}$ (*biased first moment estimate*)
 - $\mathbf{r} \leftarrow \beta_2 \mathbf{r} + (1 - \beta_2) \mathbf{g}^2$ (*biased second moment estimate*)
 - $\alpha_t \leftarrow \alpha \sqrt{1 - \beta_2^t} / (1 - \beta_1^t)$
 - $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \frac{\alpha_t}{\sqrt{\mathbf{r} + \varepsilon}} \mathbf{s}$


After t steps, we have

$$\mathbf{r}_t = (1 - \beta_2) \sum_{i=1}^t \beta_2^{t-i} \mathbf{g}_i^2.$$


Assuming that the second moment $\mathbb{E}[\mathbf{g}_i^2]$ is stationary, we have

$$\begin{aligned} \mathbb{E}[\mathbf{r}_t] &= \mathbb{E} \left[(1 - \beta_2) \sum_{i=1}^t \beta_2^{t-i} \mathbf{g}_i^2 \right] \\ &= \mathbb{E}[\mathbf{g}_t^2] \cdot (1 - \beta_2) \sum_{i=1}^t \beta_2^{t-i} \\ &= \mathbb{E}[\mathbf{g}_t^2] \cdot (1 - \beta_2^t) \end{aligned}$$


and analogously for correction of \mathbf{s} .




http://2.bp.blogspot.com/-q6l20Vs4P_w/VPmIC7sEhnl/AAAAAAAAACC4/g3UOUX2r_yA/s400/ found at <http://www.denizyuret.com/2015/03/alec-radfords-animations-for.html>



<http://2.bp.blogspot.com/-L98w-SBmF58/VPmICljKEKI/AAAAAAAAACCs/rrFz3VetYmM/s400/> found at <http://www.denizyuret.com/2015/03/alec-radfords-animations-for.html>



http://3.bp.blogspot.com/-nrtJPrdBWuE/VPmIB46F2aI/AAAAAAAAACCw/vaE_B0SVy5k/s400/ found at <http://www.denizyuret.com/2015/03/alec-radfords-animations-for.html>



http://1.bp.blogspot.com/-K_X-yud8nj8/VPmIBxwGsl/AAAAAAAAACC0/JS-h1fa09EQ/s400/ found at <http://www.denizyuret.com/2015/03/alec-radfords-animations-for.html>

Even if RMSProp and Adam are adaptive, they still usually require carefully tuned decreasing learning rate for top-notch performance.

- *Exponential decay*: learning rate is multiplied by a constant each batch/epoch/several epochs.
 - $\alpha = \alpha_{\text{initial}} \cdot c^t$
 - Often used for convolutional networks (image recognition etc.).
- *Polynomial decay*: learning rate is multiplied by some polynomial of t .
 - *Inverse time decay* uses $\alpha = \alpha_{\text{initial}} \cdot \frac{1}{t}$ and has theoretical guarantees of convergence, but is usually too fast for deep neural networks.
 - *Inverse-square decay* uses $\alpha = \alpha_{\text{initial}} \cdot \frac{1}{\sqrt{t}}$ and is currently used by best machine translation models.
- Cosine decay, restarts, warmup, ...

The `tf.optimizers.schedules` offers several such learning rate schedules, which can be passed to any Keras optimizer directly as a learning rate.