# lecture 8: optimization and deeper architectures deep learning for vision 

Yannis Avrithis

Inria Rennes-Bretagne Atlantique

Rennes, Nov. 2018 - Jan. 2019


## outline

optimizers
initialization
normalization deeper architectures

## optimizers

## gradient descent

- update rule

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\epsilon \mathbf{g}^{(\tau)}
$$

where

$$
\mathbf{g}^{(\tau)}:=\nabla f\left(\mathbf{x}^{(\tau)}\right)
$$

in a (continuous-time) physical analogy, if $\mathrm{x}^{(\tau)}$ represents the position of a particle at time $\tau$, then $-\mathbf{g}^{(\tau)}$ represents its velocity

$$
\frac{d \mathrm{x}}{d \tau}=-\mathrm{g}=-\nabla f(\mathrm{x})
$$

(where $\frac{d \mathrm{x}}{d \tau} \approx \frac{\mathrm{x}^{(\tau+1)}-\mathrm{x}^{(\tau)}}{\epsilon}$ )

- in the following, we examine a batch and a stochastic version: in the latter, each update is split into 10 smaller steps, with stochastic noise added to each step (assuming a batch update consists of 10 terms)


## gradient descent

- update rule

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\epsilon \mathbf{g}^{(\tau)}
$$

where

$$
\mathbf{g}^{(\tau)}:=\nabla f\left(\mathbf{x}^{(\tau)}\right)
$$

- in a (continuous-time) physical analogy, if $\mathbf{x}^{(\tau)}$ represents the position of a particle at time $\tau$, then $-\mathbf{g}^{(\tau)}$ represents its velocity

$$
\frac{d \mathbf{x}}{d \tau}=-\mathbf{g}=-\nabla f(\mathbf{x})
$$

(where $\frac{d \mathbf{x}}{d \tau} \approx \frac{\mathbf{x}^{(\tau+1)}-\mathbf{x}^{(\tau)}}{\epsilon}$ )
in the following, we examine a batch and a stochastic version: in the
latter, each update is split into 10 smaller steps, with stochastic noise added to each step (assuming a batch update consists of 10 terms)

## gradient descent

- update rule

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\epsilon \mathbf{g}^{(\tau)}
$$

where

$$
\mathbf{g}^{(\tau)}:=\nabla f\left(\mathbf{x}^{(\tau)}\right)
$$

- in a (continuous-time) physical analogy, if $\mathbf{x}^{(\tau)}$ represents the position of a particle at time $\tau$, then $-\mathbf{g}^{(\tau)}$ represents its velocity

$$
\frac{d \mathbf{x}}{d \tau}=-\mathbf{g}=-\nabla f(\mathbf{x})
$$

(where $\frac{d \mathbf{x}}{d \tau} \approx \frac{\mathbf{x}^{(\tau+1)}-\mathbf{x}^{(\tau)}}{\epsilon}$ )

- in the following, we examine a batch and a stochastic version: in the latter, each update is split into 10 smaller steps, with stochastic noise added to each step (assuming a batch update consists of 10 terms)
(batch) gradient descent

(batch) gradient descent

(batch) gradient descent

$\epsilon=0.14$, iteration 2
(batch) gradient descent

$\epsilon=0.14$, iteration 3
(batch) gradient descent

$\epsilon=0.14$, iteration 4
(batch) gradient descent

(batch) gradient descent

(batch) gradient descent

$\epsilon=0.14$, iteration 7
(batch) gradient descent

(batch) gradient descent



## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 0
$$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 1
$$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 2
$$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 3
$$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 4
$$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 5
$$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 6
$$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 7
$$

## (stochastic) gradient descent


$\epsilon=0.07$, iteration $10 \times 8$

## (stochastic) gradient descent



$$
\epsilon=0.07, \text { iteration } 10 \times 9
$$

## (stochastic) gradient descent


$\epsilon=0.07$, iteration $10 \times 10$

## (stochastic) gradient descent


$\epsilon=0.07$, iteration $10 \times 11$

## (stochastic) gradient descent


$\epsilon=0.07$, iteration $10 \times 12$

## (stochastic) gradient descent


$\epsilon=0.07$, iteration $10 \times 13$

## (stochastic) gradient descent


$\epsilon=0.07$, iteration $10 \times 14$

## (stochastic) gradient descent


$\epsilon=0.07$, iteration $10 \times 15$

## problems

- high condition number: oscillations, divergence
- plateaus, saddle points: no progress
- sensitive to stochastic noise


## gradient descent



- inspector needs to walk down the hill
- it is better to go skiing!


## gradient descent with momentum



- inspector needs to walk down the hill
- it is better to go skiing!


## gradient descent with momentum

[Rumelhart et al. 1986]

- in the same analogy, if the particle is of mass $m$ and moving in a medium with viscosity $\mu$, now $-\mathbf{g}$ represents a (gravitational) force and $f$ the potential energy, proportional to altitude

$$
m \frac{d^{2} \mathbf{x}}{d \tau^{2}}+\mu \frac{d \mathbf{x}}{d \tau}=-\mathbf{g}=-\nabla f(\mathbf{x})
$$

this formulation yields the update rule

where $\mathbf{v}:=\frac{d \mathbf{x}}{d \tau} \approx \mathbf{x}^{(\tau+1)}-\mathbf{x}^{(\tau)}$ represents the velocity, initialized to zero, $\frac{d^{2} \mathbf{x}}{d \tau^{2}} \approx \frac{\mathbf{v}^{(\tau+1)}-\mathbf{v}^{(\tau)}}{\delta}, \alpha:=\frac{m-\mu \delta}{m}$, and $\epsilon:=\frac{\delta}{m}$

## gradient descent with momentum

[Rumelhart et al. 1986]

- in the same analogy, if the particle is of mass $m$ and moving in a medium with viscosity $\mu$, now $-\mathbf{g}$ represents a (gravitational) force and $f$ the potential energy, proportional to altitude

$$
m \frac{d^{2} \mathbf{x}}{d \tau^{2}}+\mu \frac{d \mathbf{x}}{d \tau}=-\mathbf{g}=-\nabla f(\mathbf{x})
$$

- this formulation yields the update rule

$$
\begin{aligned}
& \mathbf{v}^{(\tau+1)}=\alpha \mathbf{v}^{(\tau)}-\epsilon \mathbf{g}^{(\tau)} \\
& \mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}+\mathbf{v}^{(\tau+1)}
\end{aligned}
$$

where $\mathbf{v}:=\frac{d \mathbf{x}}{d \tau} \approx \mathbf{x}^{(\tau+1)}-\mathbf{x}^{(\tau)}$ represents the velocity, initialized to zero, $\frac{d^{2} \mathbf{x}}{d \tau^{2}} \approx \frac{\mathbf{v}^{(\tau+1)}-\mathbf{v}^{(\tau)}}{\delta}, \alpha:=\frac{m-\mu \delta}{m}$, and $\epsilon:=\frac{\delta}{m}$

## gradient descent with momentum

[Rumelhart et al. 1986]

- when $\mathbf{g}$ is constant, $\mathbf{v}$ reaches terminal velocity

$$
\mathbf{v}^{(\infty)}=-\epsilon \mathbf{g} \sum_{\tau=0}^{\infty} \alpha^{\tau}=-\frac{\epsilon}{1-\alpha} \mathbf{g}
$$

e.g. if $\alpha=0.99$, this is 100 times faster than gradient descent
$\alpha \in[0,1)$ is another hyperparameter with $1-\alpha$ representing viscosity; usually $\alpha=0.9$

## gradient descent with momentum

[Rumelhart et al. 1986]

- when $\mathbf{g}$ is constant, $\mathbf{v}$ reaches terminal velocity

$$
\mathbf{v}^{(\infty)}=-\epsilon \mathbf{g} \sum_{\tau=0}^{\infty} \alpha^{\tau}=-\frac{\epsilon}{1-\alpha} \mathbf{g}
$$

e.g. if $\alpha=0.99$, this is 100 times faster than gradient descent

- $\alpha \in[0,1)$ is another hyperparameter with $1-\alpha$ representing viscosity; usually $\alpha=0.9$


## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum


$\epsilon=0.14$, iteration 2

Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (batch) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum



$$
\epsilon=0.07, \text { iteration } 10 \times 10
$$

## (stochastic) momentum


$\epsilon=0.07$, iteration $10 \times 11$

Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum


$\epsilon=0.07$, iteration $10 \times 12$

Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum


$\epsilon=0.07$, iteration $10 \times 13$

Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum


$\epsilon=0.07$, iteration $10 \times 14$

Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## (stochastic) momentum


$\epsilon=0.07$, iteration $10 \times 15$

Rumelhart, Hinton and Williams. N 1986. Learning Representations By Back-Propagating Errors.

## gradient descent with momentum

- good for high condition number: damps oscillations by its viscosity
- good for plateaus/saddle points: accelerates in directions with consistent gradient signs
- insensitive to stochastic noise, due to averaging


## adaptive learning rates

- the partial derivative with respect to each parameter may be very different, especially e.g. for units with different fan-in or for different layers
- we need separate, adaptive learning rate per parameter
- for batch learning, we can
- just use the the gradient sign on the agreement of gradient signs between iterations


## adaptive learning rates

- the partial derivative with respect to each parameter may be very different, especially e.g. for units with different fan-in or for different layers
- we need separate, adaptive learning rate per parameter
- for batch learning, we can
- just use the the gradient sign
- Rprop: also adjust the learning rate of each parameter depending on the agreement of gradient signs between iterations


## RMSprop

## [Tieleman and Hinton 2012]

- for mini-batch or online methods, we need to average over iterations
- $\operatorname{sgn} \mathbf{g}$ can be written as $\mathbf{g} /|\mathbf{g}|$ (element-wise) and we can replace $|\mathbf{g}|$ by an average
- maintain a moving average b of the squared gradient $\mathrm{g}^{2}$, then divide g by $\sqrt{\mathrm{b}}$

where all operations are taken element-wise
- e.g. $\beta=0.9, \delta=10^{-8}$


## RMSprop

## [Tieleman and Hinton 2012]

- for mini-batch or online methods, we need to average over iterations
- $\operatorname{sgn} \mathbf{g}$ can be written as $\mathbf{g} /|\mathbf{g}|$ (element-wise) and we can replace $|\mathbf{g}|$ by an average
- maintain a moving average $\mathbf{b}$ of the squared gradient $\mathbf{g}^{2}$, then divide g by $\sqrt{\mathbf{b}}$

$$
\begin{aligned}
& \mathbf{b}^{(\tau+1)}=\beta \mathbf{b}^{(\tau)}+(1-\beta)\left(\mathbf{g}^{(\tau)}\right)^{2} \\
& \mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\frac{\epsilon}{\delta+\sqrt{\mathbf{b}^{(\tau+1)}}} \mathbf{g}^{(\tau)}
\end{aligned}
$$

where all operations are taken element-wise

- e.g. $\beta=0.9, \delta=10^{-8}$


## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop


$\epsilon=0.14$, iteration 2

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop


$\epsilon=0.14$, iteration 10
Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop


$\epsilon=0.14$, iteration 11
Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop


$\epsilon=0.14$, iteration 12
Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



$$
\epsilon=0.14, \text { iteration } 13
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop


$\epsilon=0.14$, iteration 14
Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (batch) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



$$
\epsilon=0.07, \text { iteration } 10 \times 10
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



$$
\epsilon=0.07, \text { iteration } 10 \times 12
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



$$
\epsilon=0.07, \text { iteration } 10 \times 13
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



$$
\epsilon=0.07, \text { iteration } 10 \times 14
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



$$
\epsilon=0.07, \text { iteration } 10 \times 15
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



$$
\epsilon=0.07, \text { iteration } 10 \times 18
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## (stochastic) RMSprop



$$
\epsilon=0.07, \text { iteration } 10 \times 20
$$

Tieleman and Hinton 2012. Divide the gradient by a running average of its recent magnitude. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

## RMSprop

- good for high condition number plateaus/saddle points: gradient is amplified (attenuated) in directions of low (high) curvature
- still, sensitive to stochastic noise


## RMSprop

- good for high condition number plateaus/saddle points: gradient is amplified (attenuated) in directions of low (high) curvature
- still, sensitive to stochastic noise
- momentum is averaging the gradient:
- RMSprop is averaging the squared gradient: 2nd order moment
- combine both: maintain moving average a (b) of gradient g (squared gradient $\mathrm{g}^{2}$ ), then update by $\mathbf{a} / \sqrt{\mathbf{b}}$

$$
\begin{aligned}
& \mathbf{b}^{(\tau+1)}=\beta \mathbf{b}^{(\tau)}+(1-\beta)\left(\mathbf{g}^{(\tau)}\right)^{2} \\
& \mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\frac{\epsilon}{\delta+\sqrt{\mathbf{b}^{(\tau+1)}}} \mathbf{g}^{(\tau)}
\end{aligned}
$$

where all operations are taken element-wise


- bias correction for small $\tau$ not shown here
- momentum is averaging the gradient: 1st order moment
- RMSprop is averaging the squared gradient: 2nd order moment
- combine both: maintain moving average a (b) of gradient g (squared gradient $\mathrm{g}^{2}$ ), then update by $\mathrm{a} / \sqrt{\mathbf{b}}$

$$
\begin{aligned}
& \mathbf{a}^{(\tau+1)}=\alpha \mathbf{a}^{(\tau)}+\left(1-\alpha \mathbf{g}^{(\tau)}\right. \\
& \mathrm{b}^{(\tau+1)}=\beta \mathrm{b}^{(\tau)}+(1-\beta)\left(\mathrm{g}^{(\tau)}\right)^{2} \\
& \mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\frac{\epsilon}{\delta+\sqrt{\mathrm{b}(\tau+1)}} \mathbf{a}^{(\tau+1)}
\end{aligned}
$$

where all operations are taken element-wise
e.g.
$\alpha=0.9, \beta=0.999, \delta=10^{-8}$

- bias correction for small $\tau$ not shown here


## Adam

## [Kingma and Ba 2015]

- momentum is averaging the gradient: 1st order moment
- RMSprop is averaging the squared gradient: 2nd order moment
- combine both: maintain moving average $\mathbf{a}(\mathbf{b})$ of gradient $g$ (squared gradient $\mathbf{g}^{2}$ ), then update by $\mathbf{a} / \sqrt{\mathbf{b}}$

$$
\begin{aligned}
& \mathbf{a}^{(\tau+1)}=\alpha \mathbf{a}^{(\tau)}+(1-\alpha) \mathbf{g}^{(\tau)} \\
& \mathbf{b}^{(\tau+1)}=\beta \mathbf{b}^{(\tau)}+(1-\beta)\left(\mathbf{g}^{(\tau)}\right)^{2} \\
& \mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\frac{\epsilon}{\delta+\sqrt{\mathbf{b}^{(\tau+1)}}} \mathbf{a}^{(\tau+1)}
\end{aligned}
$$

where all operations are taken element-wise

- e.g. $\alpha=0.9, \beta=0.999, \delta=10^{-8}$
- bias correction for small $\tau$ not shown here


## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam


$\epsilon=0.14$, iteration 13

Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam



Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.

## (batch) Adam


$\epsilon=0.14$, iteration 20
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


$$
\epsilon=0.07, \text { iteration } 10 \times 10
$$

(stochastic) Adam

$\epsilon=0.07$, iteration $10 \times 11$
(stochastic) Adam

$\epsilon=0.07$, iteration $10 \times 12$

Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


$$
\epsilon=0.07, \text { iteration } 10 \times 13
$$

(stochastic) Adam

$\epsilon=0.07$, iteration $10 \times 14$
(stochastic) Adam

$\epsilon=0.07$, iteration $10 \times 15$

Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


$$
\epsilon=0.07, \text { iteration } 10 \times 16
$$

(stochastic) Adam


$$
\epsilon=0.07, \text { iteration } 10 \times 17
$$

(stochastic) Adam

$\epsilon=0.07$, iteration $10 \times 18$

Kingma and Ba. ICLR 2015. Adam: A Method for Stochastic Optimization.
(stochastic) Adam


$$
\epsilon=0.07, \text { iteration } 10 \times 19
$$

(stochastic) Adam


$$
\epsilon=0.07, \text { iteration } 10 \times 20
$$

## learning rate

- remember
- all these methods need to determine the learning rate
- to converge, the learning rate needs to be reduced during learning
- set a fixed learning rate schedule, e.g.
or, halve the learning rate every 10 epochs
- adjust to the current behavior, manually or automatically
- if the error is decreasing slowly and consistently, try increasing $\epsilon$
- if it is increasing, fluctuating, or stabilizing, try decreasing $\epsilon$


## learning rate

- remember
- all these methods need to determine the learning rate
- to converge, the learning rate needs to be reduced during learning
- set a fixed learning rate schedule, e.g.

$$
\epsilon_{\tau}=\epsilon_{0} e^{-\gamma \tau}
$$

or, halve the learning rate every 10 epochs

- adjust to the current behavior, manually or automatically
- if the error is decreasing slowly and consistently, try
- if it is increasing, fluctuating, or stabilizing, try decreasing $\epsilon$


## learning rate

- remember
- all these methods need to determine the learning rate
- to converge, the learning rate needs to be reduced during learning
- set a fixed learning rate schedule, e.g.

$$
\epsilon_{\tau}=\epsilon_{0} e^{-\gamma \tau}
$$

or, halve the learning rate every 10 epochs

- adjust to the current behavior, manually or automatically
- if the error is decreasing slowly and consistently, try increasing $\epsilon$
- if it is increasing, fluctuating, or stabilizing, try decreasing $\epsilon$


## second order optimization*

- remember, the gradient descent update rule

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\epsilon \mathbf{g}^{(\tau)}
$$

comes from assuming a second-order Taylor approximation of $f$ around $\mathbf{x}^{(\tau)}$ with an fixed, isotropic Hessian $H f(\mathbf{x})=\frac{1}{\epsilon} I$ everywhere, and making its gradient vanish

## if we knew the true Hessian matrix at $\mathrm{x}^{(\tau)}$, we would get the update rule instead


where

## second order optimization*

- remember, the gradient descent update rule

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\epsilon \mathbf{g}^{(\tau)}
$$

comes from assuming a second-order Taylor approximation of $f$ around $\mathbf{x}^{(\tau)}$ with an fixed, isotropic Hessian $H f(\mathbf{x})=\frac{1}{\epsilon} I$ everywhere, and making its gradient vanish

- if we knew the true Hessian matrix at $\mathbf{x}^{(\tau)}$, we would get the Newton update rule instead

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\left[H^{(\tau)}\right]^{-1} \mathbf{g}^{(\tau)}
$$

where

$$
H^{(\tau)}:=H f\left(\mathbf{x}^{(\tau)}\right)
$$

## second order optimization*

- remember, the gradient descent update rule

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\epsilon \mathbf{g}^{(\tau)}
$$

comes from assuming a second-order Taylor approximation of $f$ around $\mathbf{x}^{(\tau)}$ with an fixed, isotropic Hessian $H f(\mathbf{x})=\frac{1}{\epsilon} I$ everywhere, and making its gradient vanish

- if we knew the true Hessian matrix at $\mathbf{x}^{(\tau)}$, we would get the Newton update rule instead

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\left[H^{(\tau)}\right]^{-1} \mathbf{g}^{(\tau)}
$$

where

$$
H^{(\tau)}:=H f\left(\mathbf{x}^{(\tau)}\right)
$$

- unfortunately, computing and inverting $H^{(\tau)}$ is not an option


## Hessian-free optimization*

[Martens ICML 2010]

- Newton's method can solve all curvature-related problems

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\left[H^{(\tau)}\right]^{-1} \mathbf{g}^{(\tau)}
$$



- in practice, solve linear system

$$
H^{(\tau)} \mathbf{d}=\mathbf{g}^{(\tau)}
$$

by conjugate gradient (CG) method, where matrix-vector products of the form $H^{(\tau)} \mathbf{v}$ are computed by back-propagation

## Hessian-free optimization*

[Martens ICML 2010]

- Newton's method can solve all curvature-related problems

$$
\mathbf{x}^{(\tau+1)}=\mathbf{x}^{(\tau)}-\left[H^{(\tau)}\right]^{-1} \mathbf{g}^{(\tau)}
$$



- in practice, solve linear system

$$
H^{(\tau)} \mathbf{d}=\mathbf{g}^{(\tau)}
$$

by conjugate gradient (CG) method, where matrix-vector products of the form $H^{(\tau)} \mathbf{v}$ are computed by back-propagation
"well begun is half done"

## initialization

## remember CIFAR10 experiment?

## prepare

- vectorize $32 \times 32 \times 3$ images into $3072 \times 1$
- split training set e.g. into $n_{\text {train }}=45000$ training samples and $n_{\text {val }}=5000$ samples to be used for validation
- center vectors by subtracting mean over the training samples
- initialize network weights as Gaussian with standard deviation $10^{-4}$
- train for a few iterations and evaluate accuracy on the validation set for a number of learning rates $\epsilon$ and regularization strengths $\lambda$
- train for 10 epochs on the full training set for the chosen hyperparameters; mini-batch $m=200$
- evaluate accuracy on the test set


## remember CIFAR10 experiment?

## prepare

- vectorize $32 \times 32 \times 3$ images into $3072 \times 1$
- split training set e.g. into $n_{\text {train }}=45000$ training samples and $n_{\text {val }}=5000$ samples to be used for validation
- center vectors by subtracting mean over the training samples
- initialize network weights as Gaussian with standard deviation $10^{-4}$


## learn

- train for a few iterations and evaluate accuracy on the validation set for a number of learning rates $\epsilon$ and regularization strengths $\lambda$
- train for 10 epochs on the full training set for the chosen hyperparameters; mini-batch $m=200$
- evaluate accuracy on the test set


## result

- linear classifier: test accuracy $38 \%$
- two-layer classifier, 200 hidden units, relu: test accuracy $51 \%$ - eight-layer classifier, 100 hidden units per layer, relu:


## result

- linear classifier: test accuracy $38 \%$
- two-layer classifier, 200 hidden units, relu: test accuracy $51 \%$
- eight-layer classifier, 100 hidden units per layer, relu: nothing works


## CIFAR10 experiment, again

## prepare

- vectorize $32 \times 32 \times 3$ images into $3072 \times 1$
- split training set e.g. into $n_{\text {train }}=45000$ training samples and $n_{\text {val }}=5000$ samples to be used for validation
- center vectors by subtracting mean over the training samples
- initialize network weights as Gaussian with standard deviation $10^{-4}$


## learn

- train for a few iterations and evaluate accuracy on the validation set for a number of learning rates $\epsilon$ and regularization strengths $\lambda$
- train for 10 epochs on the full training set for the chosen hyperparameters; mini-batch $m=200$
- evaluate accuracy on the test set


## CIFAR10 experiment, again

## prepare

- vectorize $32 \times 32 \times 3$ images into $3072 \times 1$
- split training set e.g. into $n_{\text {train }}=45000$ training samples and $n_{\text {val }}=5000$ samples to be used for validation
- center vectors by subtracting mean over the training samples
- initialize network weights as Gaussian with standard deviation $10^{-4}$


## learn

- train for a few iterations and evaluate accuracy on the validation set for a number of learning rates $\epsilon$ and regularization strengths $\lambda$
- train for 10 epochs on the full training set for the chosen hyperparameters; mini-batch $m=200$
- evaluate accuracy on the test set


## affine layer initialization

- $k \times k^{\prime}$ weight matrix $W, k^{\prime} \times 1$ bias vector $\mathbf{b}$

$$
\mathbf{a}=W^{\top} \mathbf{x}+\mathbf{b}, \quad \mathbf{x}^{\prime}=h(\mathbf{a})=h\left(W^{\top} \mathbf{x}+\mathbf{b}\right)
$$

- each element $w$ of $W$ can be drawn at random, e.g.

$$
\begin{aligned}
& \text { - Gaussian } w \sim \mathcal{N}\left(0, \sigma^{2}\right) \text {, with } \operatorname{Var}(w)=\sigma^{2} \\
& w \sim U(-a, a) \text {, with } \operatorname{Var}(w)=\sigma^{2}=\frac{a^{2}}{3}
\end{aligned}
$$

- in any case, it is important to determine the standard deviation $\sigma$, which we call weight scale
- can be again Gaussian or uniform
- more commonly, constant e.g. zero
- the constant depends on the activation function $h$ and should be chosen such that $h$ does not saturate or 'die'


## affine layer initialization

- $k \times k^{\prime}$ weight matrix $W, k^{\prime} \times 1$ bias vector $\mathbf{b}$

$$
\mathbf{a}=W^{\top} \mathbf{x}+\mathbf{b}, \quad \mathbf{x}^{\prime}=h(\mathbf{a})=h\left(W^{\top} \mathbf{x}+\mathbf{b}\right)
$$

## weights

- each element $w$ of $W$ can be drawn at random, e.g.
- Gaussian $w \sim \mathcal{N}\left(0, \sigma^{2}\right)$, with $\operatorname{Var}(w)=\sigma^{2}$
- uniform $w \sim U(-a, a)$, with $\operatorname{Var}(w)=\sigma^{2}=\frac{a^{2}}{3}$
- in any case, it is important to determine the standard deviation $\sigma$, which we call weight scale
- can be again Gaussian or uniform
- more commonly, constant e.g. zero
- the constant depends on the activation function $h$ and should be chosen such that $h$ does not saturate or 'die'


## affine layer initialization

- $k \times k^{\prime}$ weight matrix $W, k^{\prime} \times 1$ bias vector $\mathbf{b}$

$$
\mathbf{a}=W^{\top} \mathbf{x}+\mathbf{b}, \quad \mathbf{x}^{\prime}=h(\mathbf{a})=h\left(W^{\top} \mathbf{x}+\mathbf{b}\right)
$$

## weights

- each element $w$ of $W$ can be drawn at random, e.g.
- Gaussian $w \sim \mathcal{N}\left(0, \sigma^{2}\right)$, with $\operatorname{Var}(w)=\sigma^{2}$
- uniform $w \sim U(-a, a)$, with $\operatorname{Var}(w)=\sigma^{2}=\frac{a^{2}}{3}$
- in any case, it is important to determine the standard deviation $\sigma$, which we call weight scale


## biases

- can be again Gaussian or uniform
- more commonly, constant e.g. zero
- the constant depends on the activation function $h$ and should be chosen such that $h$ does not saturate or 'die'


## weight scale sensitivity



- using $\mathcal{N}\left(0, \sigma^{2}\right)$, training on a small subset of the training set and cross-validating $\sigma$ reveals a narrow peak in validation accuracy


## weight scale sensitivity

- to understand why, we measure the distribution of features $\mathbf{x}$ in all layers, starting with Gaussian input $\sim \mathcal{N}(0,1)$
- we repeat with and without relu nonlinearity
- in each case, we try three different values of quantity $k \sigma$
linear units, input
$k \sigma^{2}=2 / 3$

$k \sigma^{2}=3 / 2$


$$
k \sigma^{2}=1
$$


linear units, input-layer 1
$k \sigma^{2}=2 / 3$

$k \sigma^{2}=3 / 2$


$$
k \sigma^{2}=1
$$


linear units, input-layer 2
$k \sigma^{2}=2 / 3$

$k \sigma^{2}=3 / 2$


$$
k \sigma^{2}=1
$$


linear units, input-layer 3
$k \sigma^{2}=2 / 3$

$k \sigma^{2}=3 / 2$


linear units, input-layer 4
$k \sigma^{2}=2 / 3$

$k \sigma^{2}=3 / 2$


$$
k \sigma^{2}=1
$$


linear units, input-layer 5
$k \sigma^{2}=2 / 3$

$k \sigma^{2}=3 / 2$


$$
k \sigma^{2}=1
$$


linear units, input-layer 6
$k \sigma^{2}=2 / 3$

$k \sigma^{2}=3 / 2$


$$
k \sigma^{2}=1
$$


linear units, input-layer 7
$k \sigma^{2}=2 / 3$
$k \sigma^{2}=3 / 2$



$$
k \sigma^{2}=1
$$


linear units, input-layer 8
$k \sigma^{2}=2 / 3$
$k \sigma^{2}=3 / 2$

$k \sigma^{2}=1$


## linear approximation

- assuming we are in a linear regime of the activation function, forward-backward relations are, recalling $W$ is $k \times k^{\prime}$

$$
\mathbf{x}^{\prime}=W^{\top} \mathbf{x}+\mathbf{b}, \quad d \mathbf{x}=W d \mathbf{x}^{\prime}, \quad d W=\mathbf{x}\left(d \mathbf{x}^{\prime}\right)^{\top}
$$

- forward: assuming $w_{i j}$ are i.i.d, $\operatorname{Var}\left(x_{i}\right)$ are the same, $w_{i j}$ and $x_{i}$ are independent, and $w_{i j}, x_{i}$ are centered, i.e. $\mathbb{E}\left(w_{i j}\right)=\mathbb{E}\left(x_{i}\right)=0$, $\operatorname{Var}\left(x_{j}^{\prime}\right)=\operatorname{Var}\left(\left(T^{\top} \mathrm{x}\right)_{j}\right)=k \operatorname{Var}(w) \operatorname{Var}(x)=k 0^{2} \operatorname{Var}(x)$
- backward, activation: under the same assumptions,

$$
\left.\operatorname{Var}\left(d x_{i}\right)=\operatorname{Var}\left(\left(W T \mathbf{N}^{\prime}\right)\right)_{i}\right)=k^{\prime} \operatorname{Var}(w) \operatorname{Var}\left(d x^{\prime}\right)=k^{\prime} \operatorname{Var}^{2}\left(d x^{\prime}\right)
$$

- backward, weights: also assuming that $x_{i}, d x_{j}^{\prime}$ are independent,



## linear approximation

- assuming we are in a linear regime of the activation function, forward-backward relations are, recalling $W$ is $k \times k^{\prime}$

$$
\mathbf{x}^{\prime}=W^{\top} \mathbf{x}+\mathbf{b}, \quad d \mathbf{x}=W d \mathbf{x}^{\prime}, \quad d W=\mathbf{x}\left(d \mathbf{x}^{\prime}\right)^{\top}
$$

- forward: assuming $w_{i j}$ are i.i.d, $\operatorname{Var}\left(x_{i}\right)$ are the same, $w_{i j}$ and $x_{i}$ are independent, and $w_{i j}, x_{i}$ are centered, i.e. $\mathbb{E}\left(w_{i j}\right)=\mathbb{E}\left(x_{i}\right)=0$,

$$
\operatorname{Var}\left(x_{j}^{\prime}\right)=\operatorname{Var}\left(\left(W^{\top} \mathbf{x}\right)_{j}\right)=k \operatorname{Var}(w) \operatorname{Var}(x)=k \sigma^{2} \operatorname{Var}(x)
$$

- backward, activation: under the same assumptions,

$$
\operatorname{Var}\left(d x_{i}\right)=\operatorname{Var}\left(\left(W d \mathbf{x}^{\prime}\right)_{i}\right)=k^{\prime} \operatorname{Var}(w) \operatorname{Var}\left(d x^{\prime}\right)=k^{\prime} \sigma^{2} \operatorname{Var}\left(d x^{\prime}\right)
$$

also assuming that $x_{i}, d x_{j}^{\prime}$ are independent,
$\operatorname{Var}\left(d w_{i j}\right)=\operatorname{Var}\left(x_{i}\right) \operatorname{Var}\left(d x_{j}^{\prime}\right)$

## linear approximation

- assuming we are in a linear regime of the activation function, forward-backward relations are, recalling $W$ is $k \times k^{\prime}$

$$
\mathbf{x}^{\prime}=W^{\top} \mathbf{x}+\mathbf{b}, \quad d \mathbf{x}=W d \mathbf{x}^{\prime}, \quad d W=\mathbf{x}\left(d \mathbf{x}^{\prime}\right)^{\top}
$$

- forward: assuming $w_{i j}$ are i.i.d, $\operatorname{Var}\left(x_{i}\right)$ are the same, $w_{i j}$ and $x_{i}$ are independent, and $w_{i j}, x_{i}$ are centered, i.e. $\mathbb{E}\left(w_{i j}\right)=\mathbb{E}\left(x_{i}\right)=0$,

$$
\operatorname{Var}\left(x_{j}^{\prime}\right)=\operatorname{Var}\left(\left(W^{\top} \mathbf{x}\right)_{j}\right)=k \operatorname{Var}(w) \operatorname{Var}(x)=k \sigma^{2} \operatorname{Var}(x)
$$

- backward, activation: under the same assumptions,

$$
\operatorname{Var}\left(d x_{i}\right)=\operatorname{Var}\left(\left(W d \mathbf{x}^{\prime}\right)_{i}\right)=k^{\prime} \operatorname{Var}(w) \operatorname{Var}\left(d x^{\prime}\right)=k^{\prime} \sigma^{2} \operatorname{Var}\left(d x^{\prime}\right)
$$

- backward, weights: also assuming that $x_{i}, d x_{j}^{\prime}$ are independent,

$$
\operatorname{Var}\left(d w_{i j}\right)=\operatorname{Var}\left(x_{i}\right) \operatorname{Var}\left(d x_{j}^{\prime}\right)
$$

## linear approximation

- if $k \sigma^{2}<1$, activations vanish forward; if $k \sigma^{2}>1$ they explode, possibly driving nonlinearities to saturation
- if $k^{\prime} \sigma^{2}<1$, activation gradients vanish backward; if $k^{\prime} \sigma^{2}>1$ they explode, and everything is linear backwards
- interestingly, weight gradients are stable (why?), but only at initialization


## "Xavier" initialization

## [Glorot and Bengio 2010]

- forward requirement is $\sigma^{2}=1 / k$
- backward requirement is $\sigma^{2}=1 / k^{\prime}$
- as a compromise, initialize according to

$$
\sigma^{2}=\frac{2}{k+k^{\prime}}
$$

## a simpler alternative

## [LeCun et al. 1998]

- however, any of these alternatives would do

$$
\sigma^{2}=\frac{1}{k}, \quad \text { or } \quad \sigma^{2}=\frac{1}{k^{\prime}}
$$

in the sense that if the forward signal is properly initialized, then so is the backward signal, and vice versa (why?)

- so, initialize according to

$$
\sigma^{2}=\frac{1}{k}
$$

## relu units, input

$k \sigma^{2}=2 \times 2 / 3$

$k \sigma^{2}=2 \times 3 / 2$


relu units, input-layer 1
$k \sigma^{2}=2 \times 2 / 3$

$k \sigma^{2}=2 \times 3 / 2$


relu units, input-layer 2
$k \sigma^{2}=2 \times 2 / 3$
$k \sigma^{2}=2 \times 3 / 2$



relu units, input-layer 3
$k \sigma^{2}=2 \times 2 / 3$
$k \sigma^{2}=2 \times 3 / 2$




## relu units, input-layer 4

$k \sigma^{2}=2 \times 2 / 3$
$k \sigma^{2}=2 \times 3 / 2$



relu units, input-layer 5
$k \sigma^{2}=2 \times 2 / 3$
$k \sigma^{2}=2 \times 3 / 2$



relu units, input-layer 6
$k \sigma^{2}=2 \times 2 / 3$
$k \sigma^{2}=2 \times 3 / 2$




## relu units, input-layer 7

$k \sigma^{2}=2 \times 2 / 3$

$k \sigma^{2}=2 \times 3 / 2$


relu units, input-layer 8
$k \sigma^{2}=2 \times 2 / 3$
$k \sigma^{2}=2 \times 3 / 2$




## relu ("Kaiming/MSRA") initialization [He et al. 2015]

- because relu squeezes half of the volume, a corrective factor of 2 appears in the expectations of both forward and backward
- so any of the following will do

$$
\sigma^{2}=\frac{2}{k}, \quad \text { or } \quad \sigma^{2}=\frac{2}{k^{\prime}}
$$

## relu ("Kaiming/MSRA") initialization



- Xavier converges more slowly or not at all
- 30-layer network trained from scratch for the first time, but has worse performance than a 14-layer network


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## convolutional layer initialization



- a convolutional layer is just an affine layer with a special matrix structure
- it is actually represented by a 4d tensor $\mathbf{w}$ of size $r^{2} k k^{\prime}$, where $r$ is the kernel size and $k, k^{\prime}$ the input/output features
- initialization is the same, but with
- fan-in $k$ replaced by $r^{2} k$
- fan-out $k^{\prime}$ replaced by $r^{2} k^{\prime}$


## beyond Gaussian matrices*

- for linear and relu units, we can now keep the signal variance constant across layers, both forward and backward
- but this just holds on average
- how exactly are signals amplified or attenuated in each dimension?
- how does that affect the learning speed?
- we return to the linear case and examine the singular values of a product $W_{8} \cdots W_{1}$ of Gaussian matrices


## beyond Gaussian matrices*

- for linear and relu units, we can now keep the signal variance constant across layers, both forward and backward
- but this just holds on average
- how exactly are signals amplified or attenuated in each dimension?
- how does that affect the learning speed?
- we return to the linear case and examine the singular values of a product $W_{8} \cdots W_{1}$ of Gaussian matrices


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$
- a product $W_{2} \cdots W_{1}$ of $\ell=2$ such matrices


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$
- a product $W_{3} \cdots W_{1}$ of $\ell=3$ such matrices


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$
- a product $W_{4} \cdots W_{1}$ of $\ell=4$ such matrices


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$
- a product $W_{5} \cdots W_{1}$ of $\ell=5$ such matrices


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$
- a product $W_{6} \cdots W_{1}$ of $\ell=6$ such matrices


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$
- a product $W_{7} \cdots W_{1}$ of $\ell=7$ such matrices


## matrices as numbers*



- singular values of $k \times k$ Gaussian matrix $W$ with elements $\sim \mathcal{N}\left(0, \sigma^{2}\right)$, for $k=100$ and for different values of $k \sigma^{2}$
- a product $W_{8} \cdots W_{1}$ of $\ell=8$ such matrices has the same behavior as raising a scalar $w^{\ell}$ : vanishing for $w<1$, exploding for $w>1$


## orthogonal initialization*

[Saxe et al. 2014]

- choose $k \times k^{\prime}$ matrix $W$ to be a random (semi-)orthogonal matrix, i.e. $W^{\top} W=I$ if $k \geq k^{\prime}$ and $W W^{\top}=I$ if $k<k^{\prime}$
- for instance, with a random Gaussian matrix followed by QR or SVD decomposition
a scaled Gaussian matrix has singular values around 1 and preserves norm
a random orthogonal matrix has singular values exactly 1 and
preserves norm
a product of orthogonal matrices remains orthogonal, while a product of scaled Gaussian matrices becomes strongly non-isotropic


## orthogonal initialization*

[Saxe et al. 2014]

- choose $k \times k^{\prime}$ matrix $W$ to be a random (semi-)orthogonal matrix, i.e. $W^{\top} W=I$ if $k \geq k^{\prime}$ and $W W^{\top}=I$ if $k<k^{\prime}$
- for instance, with a random Gaussian matrix followed by QR or SVD decomposition
- a scaled Gaussian matrix has singular values around 1 and preserves norm on average

$$
\mathbb{E}_{w \sim \mathcal{N}(0,1 / k)}\left(\mathbf{x}^{\top} W^{\top} W \mathbf{x}\right)=\mathbf{x}^{\top} \mathbf{x}
$$

- a random orthogonal matrix has singular values exactly 1 and preserves norm exactly

$$
\mathbf{x}^{\top} W^{\top} W \mathbf{x}=\mathbf{x}^{\top} \mathbf{x}
$$

## orthogonal initialization*

## [Saxe et al. 2014]

- choose $k \times k^{\prime}$ matrix $W$ to be a random (semi-)orthogonal matrix, i.e. $W^{\top} W=I$ if $k \geq k^{\prime}$ and $W W^{\top}=I$ if $k<k^{\prime}$
- for instance, with a random Gaussian matrix followed by QR or SVD decomposition
- a scaled Gaussian matrix has singular values around 1 and preserves norm on average

$$
\mathbb{E}_{w \sim \mathcal{N}(0,1 / k)}\left(\mathbf{x}^{\top} W^{\top} W \mathbf{x}\right)=\mathbf{x}^{\top} \mathbf{x}
$$

- a random orthogonal matrix has singular values exactly 1 and preserves norm exactly

$$
\mathbf{x}^{\top} W^{\top} W \mathbf{x}=\mathbf{x}^{\top} \mathbf{x}
$$

- a product of orthogonal matrices remains orthogonal, while a product of scaled Gaussian matrices becomes strongly non-isotropic


## data-dependent initialization

- orthogonal initialization only applies to linear layers
- relu requires analyzing input-output variances to find the corrective factor of 2
- it is not possible to do this theoretical derivation for any kind of nonlinearity, e.g. maxout, max-pooling, normalization etc.
- a practical solution is to use actual data at the input of the network and compute weights according to output statistics


## layer-sequential unit-variance (LSUV) initialization*

[Mishkin and Matas 2016]

- begin by random orthogonal initialization
- then, for each affine layer ( $W, \mathbf{b}$ ), measure output variance over a mini-batch (not per feature) and iteratively normalize it to one

$$
\begin{aligned}
& \text { def lsuv(batch, }(W, \mathbf{b}), \tau=0.1) \text { : } \\
& \begin{array}{l}
\sigma=0 \\
\text { while }|\sigma-1| \geq \tau \text { : } \\
\quad X=\operatorname{batch}() \\
Y=\operatorname{dot}(X, W)+\mathbf{b} \\
\sigma=\operatorname{std}(Y) \\
W=W / \sigma \\
\text { return }(W, \mathbf{b})
\end{array}
\end{aligned}
$$

- as given by batch(), we use a new mini-batch per iteration and feed it forward through the network until we reach the input $X$ of that layer
- $X$ is $m \times k, W$ is $k \times k^{\prime}, Y$ is $m \times k^{\prime}$, where $m$ is the mini-batch size


## within-layer initialization*

[Krähenbühl et al. 2016]

- computed on a single mini-batch, non-iterative
- measure both mean and variance, initialize both bias and weights
- measurements are per feature

$$
\begin{aligned}
& \text { def } \operatorname{within}(X,(W, \mathbf{b})) \text { : } \\
& \quad Y=\operatorname{dot}(X, W)+\mathbf{b} \\
& \boldsymbol{\mu}, \boldsymbol{\sigma}=\operatorname{mean}_{0}(Y), \operatorname{std}_{0}(Y) \\
& W, \mathbf{b}=W / \boldsymbol{\sigma},-\boldsymbol{\mu} / \boldsymbol{\sigma} \\
& \text { return }(W, \mathbf{b})
\end{aligned}
$$

- vector operations are element-wise
- matrix-vector operations are broadcasted


## data-dependent initialization*

- weights initialized by PCA or (spherical) $k$-means on mini-batch samples
- within-layer initialization normalizes affine layer outputs to zero mean, unit variance
- between-layer initialization iteratively normalizes weights and biases of different layers
- as a result, all parameters are learned at the same "rate"


## data-dependent initialization: CaffeNet



- data-dependent initialization is better at first 100k iterations
- but random initialization catches up after the second learning rate drop


## data-dependent initialization: CaffeNet




- data-dependent initialization is better at first 100k iterations
- but random initialization catches up after the second learning rate drop


## data-dependent initialization: CaffeNet



- data-dependent initialization is better at first 100k iterations
- but random initialization catches up after the second learning rate drop


## data-dependent initialization: CaffeNet


nearest neighbors of given input image in feature space

Krähenbühl, Doersch, Donahue and Darell. ICLR 2016. Data-Dependent Initializations of Convolutional Neural Networks.

## data-dependent initialization

- PCA is orthogonal but data-dependent rather than random
- $k$-means is non-orthogonal, but centroids are still only weakly correlated
we cannot fail to notice that
- codebooks are now the initial weights, computed layer-wise
- bag-of-words representations are now the initial features
- compared to the conventional approach, now the entire pipeline is optimized end-to-end


## data-dependent initialization

- PCA is orthogonal but data-dependent rather than random
- $k$-means is non-orthogonal, but centroids are still only weakly correlated
- we cannot fail to notice that
- codebooks are now the initial weights, computed layer-wise
- bag-of-words representations are now the initial features
- compared to the conventional approach, now the entire pipeline is optimized end-to-end


## normalization

## input normalization: zero mean, unit variance

- input $X$ is an $n \times d$ matrix, where $n$ is the number of samples and $d$ is the dimension of a vectorized image
- measure empirical mean and variance and normalize per dimension
def $\operatorname{norm}(X)$ :

$$
\boldsymbol{\mu}, \boldsymbol{\sigma}=\operatorname{mean}_{0}(X), \operatorname{std}_{0}(X)
$$

$$
\text { return }(X-\boldsymbol{\mu}) / \boldsymbol{\sigma}
$$

## input normalization: zero mean, unit variance




## input normalization: zero mean, unit variance





## input normalization: zero mean, unit variance





## input normalization: PCA and whitening

- center data to zero mean as before
- using SVD, measure the eigenvalues $\sigma$ and eigenvectors $V$ of the covariance matrix $\frac{1}{n} X^{\top} X$
- PCA-rotate by $V^{-1}=V^{\top}$ to decorrelate the data
- whiten by $1 / \sigma$ to unit variance def whiten $(X)$ :
$n=X$.shape $[0]$
$X-=\operatorname{mean}_{0}(X)$
$U, \boldsymbol{\sigma}, V=\operatorname{svd}(X / \operatorname{sqrt}(n))$
return $\operatorname{dot}\left(X, V^{\top}\right) / \boldsymbol{\sigma}$
input normalization: PCA and whitening



input normalization: PCA and whitening





## input normalization: PCA and whitening






## input normalization: PCA and whitening






## in practice: only centering

- the network is expected to discover nonlinear manifold structure, so in principle it should have no difficulty discovering the linear PCA + whitening structure
- in practice, only centering is enough:
- subtract the mean value per pixel (mean image)
- subtract the mean value per color channel (mean color or intensity, just one or three scalars)


## in practice: only centering

- the network is expected to discover nonlinear manifold structure, so in principle it should have no difficulty discovering the linear PCA + whitening structure
- in practice, only centering is enough:
- subtract the mean value per pixel (mean image)
- subtract the mean value per color channel (mean color or intensity, just one or three scalars)


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## why is centering important?

- each weight derivative $d w_{i}$ of layer 1 is $(d a) x_{i}$ where $d a$ is the derivative of the activation and $x_{i}$ is the corresponding input
- if all inputs are positive, then updates on weights $w_{i}$ are either all positive (if $d a<0$, quadrant 1 ) or all negative (if $d a<0$, quadrant 3)

- weights can only all increase or all decrease together for a given sample
- to follow the direction of w, we can only do so by zig-zagging


## activation normalization

- if normalization is important at the input, why not at every layer activation?
- this is even more important in the presence of saturating nonlinearities: given a wrong offset or scale, activation functions can 'die'
- and even more important in the presence of stochastic updates, where statistics change at every mini-batch and at every update (internal covariate shift)


## activation functions



$\tanh (x)=\frac{e^{x}-e^{-x}}{e^{x}+e^{-x}}=2 \sigma(x)-1$
hyperbolic tangent


$$
\operatorname{relu}(x)=[x]_{+}=\max (0, x)
$$

rectified linear unit (ReLU)


$$
\begin{gathered}
\zeta(x)=\log \left(1+e^{x}\right) \\
\text { softplus }
\end{gathered}
$$

## activation functions: non-localized



$\tanh (x)=\frac{e^{x}-e^{-x}}{e^{x}+e^{-x}}=2 \sigma(x)-1$
hyperbolic tangent


$$
\operatorname{relu}(x)=[x]_{+}=\max (0, x)
$$

rectified linear unit (ReLU)


$$
\begin{gathered}
\zeta(x)=\log \left(1+e^{x}\right) \\
\text { softplus }
\end{gathered}
$$

## activation functions: centering



$\tanh (x)=\frac{e^{x}-e^{-x}}{e^{x}+e^{-x}}=2 \sigma(x)-1$
hyperbolic tangent


$$
\operatorname{relu}(x)=[x]_{+}=\max (0, x)
$$

rectified linear unit (ReLU)


$$
\begin{gathered}
\zeta(x)=\log \left(1+e^{x}\right) \\
\text { softplus }
\end{gathered}
$$

## activation functions: centering



## activation functions: centering



## activation functions: centering



## activation functions: centering




## activation functions: centering



$$
f(x)=\max (\alpha x, x)
$$

parametric ReLU: $\alpha$ is learned

$f(x)= \begin{cases}x, & \text { if } x>0 \\ \alpha\left(e^{x}-1\right), & \text { if } x \leq 0\end{cases}$
exponential linear unit (ELU)

## activation functions: self-normalizing!


parametric ReLU: $\alpha$ is learned

$f(x)= \begin{cases}x, & \text { if } x>0 \\ \alpha\left(e^{x}-1\right), & \text { if } x \leq 0\end{cases}$ exponential linear unit (ELU)

$f(\mathbf{x})=\max _{j}\left(\mathbf{w}_{j}^{\top} \mathbf{x}+b_{j}\right)$
maxout


$$
f(x)=\lambda \begin{cases}x, & \text { if } x>0 \\ \alpha\left(e^{x}-1\right), & \text { if } x \leq 0\end{cases}
$$

$$
\text { scaled ELU }(\lambda>1)
$$

## batch normalization (BN)

## [loffe and Szegedy 2015]



- if $\mathbf{x}=\left(x_{1}, \ldots, x_{k}\right)$ is the activation or feature at any layer, normalize it element-wise

$$
\hat{x}_{j}=\frac{x_{j}-\mathbb{E}\left(x_{j}\right)}{\sqrt{\operatorname{Var}\left(x_{j}\right)}}
$$

to have zero-mean, unit-variance, where $\mathbb{E}$ and Var are empirical over the training set
insert this layer after convolutional or fully-connected layers and before nonlinear activation functions

## batch normalization (BN)

## [loffe and Szegedy 2015]



- if $\mathbf{x}=\left(x_{1}, \ldots, x_{k}\right)$ is the activation or feature at any layer, normalize it element-wise

$$
\hat{x}_{j}=\frac{x_{j}-\mathbb{E}\left(x_{j}\right)}{\sqrt{\operatorname{Var}\left(x_{j}\right)}}
$$

to have zero-mean, unit-variance, where $\mathbb{E}$ and Var are empirical over the training set

- insert this layer after convolutional or fully-connected layers and before nonlinear activation functions


## batch normalization (BN)

## [loffe and Szegedy 2015]



- if $\mathbf{x}=\left(x_{1}, \ldots, x_{k}\right)$ is the activation or feature at any layer, normalize it element-wise

$$
\hat{x}_{j}=\frac{x_{j}-\mathbb{E}\left(x_{j}\right)}{\sqrt{\operatorname{Var}\left(x_{j}\right)}}
$$

to have zero-mean, unit-variance, where $\mathbb{E}$ and Var are empirical over the training set

- insert this layer after convolutional or fully-connected layers and before nonlinear activation functions (although this is not clear)


## batch normalization: parameters

- normalized features may remain in the linear regime of the following nonlinearity, limiting the representational power of the network
- introduce parameters $\beta=\left(\beta_{1}, \ldots, \beta_{k}\right), \gamma=\left(\gamma_{1}, \ldots, \gamma_{k}\right)$ and let the output of the BN layer be $\mathbf{y}=\left(y_{1}, \ldots, y_{k}\right)$ with

or, element-wise,

- then, with

we can recover the identity mapping if needed


## batch normalization: parameters

- normalized features may remain in the linear regime of the following nonlinearity, limiting the representational power of the network
- introduce parameters $\boldsymbol{\beta}=\left(\beta_{1}, \ldots, \beta_{k}\right), \boldsymbol{\gamma}=\left(\gamma_{1}, \ldots, \gamma_{k}\right)$ and let the output of the BN layer be $\mathbf{y}=\left(y_{1}, \ldots, y_{k}\right)$ with

$$
y_{j}=\gamma_{j} \hat{x}_{j}+\beta_{j}
$$

or, element-wise,

$$
\mathbf{y}=\gamma \hat{\mathbf{x}}+\beta
$$

- then, with

we can recover the identity mapping if needed


## batch normalization: parameters

- normalized features may remain in the linear regime of the following nonlinearity, limiting the representational power of the network
- introduce parameters $\boldsymbol{\beta}=\left(\beta_{1}, \ldots, \beta_{k}\right), \boldsymbol{\gamma}=\left(\gamma_{1}, \ldots, \gamma_{k}\right)$ and let the output of the BN layer be $\mathbf{y}=\left(y_{1}, \ldots, y_{k}\right)$ with

$$
y_{j}=\gamma_{j} \hat{x}_{j}+\beta_{j}
$$

or, element-wise,

$$
\mathbf{y}=\gamma \hat{\mathbf{x}}+\beta
$$

- then, with

$$
\beta_{j}=\mathbb{E}\left(x_{j}\right), \quad \gamma_{j}=\sqrt{\operatorname{Var}\left(x_{j}\right)}
$$

we can recover the identity mapping if needed

## batch normalization: training

- as the name suggests, BN learns using the mini-batch statistics
- given an index set $I$ of mini-batch samples with $|I|=m$, the BN layer with parameters $\boldsymbol{\beta}, \gamma$ yields, for each sample feature $\mathbf{x}_{i}$ with $i \in I$,

$$
\mathbf{y}_{i}=\operatorname{BN}_{\boldsymbol{\beta}, \boldsymbol{\gamma}}\left(\mathbf{x}_{i}\right):=\gamma \frac{\mathbf{x}_{i}-\boldsymbol{\mu}_{I}}{\sqrt{\mathbf{v}_{I}+\delta}}+\boldsymbol{\beta}
$$

(element-wise), where $\boldsymbol{\mu}_{I}, \mathbf{v}_{I}$ are the mini-batch mean and variance

$$
\begin{aligned}
\boldsymbol{\mu}_{I} & :=\frac{1}{m} \sum_{i \in I} \mathbf{x}_{i} \\
\mathbf{v}_{I} & :=\frac{1}{m} \sum_{i \in I}\left(\mathbf{x}_{i}-\boldsymbol{\mu}_{I}\right)^{2}
\end{aligned}
$$

## batch normalization: inference

- at inference, BN operates with global statistics
- given a test sample feature $\mathbf{x}$, the BN layer with parameters $\boldsymbol{\beta}, \gamma$ yields (element-wise)

$$
\mathbf{y}=\mathrm{BN}_{\boldsymbol{\beta}, \boldsymbol{\gamma}}^{\mathrm{inf}}(\mathbf{x}):=\boldsymbol{\gamma} \frac{\mathbf{x}-\boldsymbol{\mu}}{\sqrt{\mathbf{v}+\delta}}+\boldsymbol{\beta}
$$

where $\boldsymbol{\mu}, \mathbf{v}$ are moving averages of the training set mean and variance, updated at every mini-batch $I$ during training as

$$
\begin{aligned}
\boldsymbol{\mu}^{(\tau+1)} & :=\alpha \boldsymbol{\mu}^{(\tau)}+(1-\alpha) \boldsymbol{\mu}_{I} \\
\mathbf{v}^{(\tau+1)} & :=\alpha \mathbf{v}^{(\tau)}+(1-\alpha) \mathbf{v}_{I}
\end{aligned}
$$

so they track the accuracy of the model as it trains

## batch normalization: derivatives*

- input mini-batch $m \times k$ matrix $X$, output $m \times k$ matrix $Y$
- forward

$$
Y=\mathrm{BN}(X,(\boldsymbol{\beta}, \boldsymbol{\gamma}))
$$

- backward: exercise


## batch normalization: derivatives*

- input mini-batch $m \times k$ matrix $X$, output $m \times k$ matrix $Y$
- forward

$$
Y=\mathrm{BN}(X,(\boldsymbol{\beta}, \boldsymbol{\gamma}))
$$

- backward: exercise

$$
\begin{aligned}
d X & =\ldots d Y \ldots \\
d \boldsymbol{\beta} & =\ldots d Y \ldots \\
d \boldsymbol{\gamma} & =\ldots d Y \ldots
\end{aligned}
$$

## batch normalization: convolution

- same as fully-connected, only now mean and variance are computed per feature map rather than per feature
- i.e. we average over mini-batch samples and spatial positions
- if feature map volumes are $w \times h \times k$, the effective mini-batch size at training becomes $m^{\prime}=m w h$, and

$$
\begin{aligned}
\boldsymbol{\mu}_{I} & :=\frac{1}{m^{\prime}} \sum_{i \in I} \sum_{\mathbf{n}} \mathbf{x}_{i}[\mathbf{n}] \\
\mathbf{v}_{I} & :=\frac{1}{m^{\prime}} \sum_{i \in I} \sum_{\mathbf{n}}\left(\mathbf{x}_{i}[\mathbf{n}]-\boldsymbol{\mu}_{I}\right)^{2}
\end{aligned}
$$

## remember weight scale sensitivity?



- using $\mathcal{N}\left(0, \sigma^{2}\right)$, training on a small subset of the training set and cross-validating $\sigma$ reveals a narrow peak in validation accuracy
- BN allows convergence over a much wider range of weight scales


## remember weight scale sensitivity?



- using $\mathcal{N}\left(0, \sigma^{2}\right)$, training on a small subset of the training set and cross-validating $\sigma$ reveals a narrow peak in validation accuracy
- BN allows convergence over a much wider range of weight scales


## batch normalization: weight scale*

- if BN is connected at the output activation of an affine layer

$$
\mathbf{a}=W^{\top} \mathbf{x}+\mathbf{b}, \quad \mathbf{x}^{\prime}=h(\mathbf{a})=h\left(W^{\top} \mathbf{x}+\mathbf{b}\right)
$$

the bias $\mathbf{b}$ is absorbed into $\boldsymbol{\beta}$ and the layer is replaced by

$$
\mathbf{x}^{\prime}=h\left(\mathrm{BN}\left(W^{\top} \mathbf{x}\right)\right)
$$

- the layer and its Jacobian are then unaffected by weight scale

- moreover, larger weights yield smaller gradients, stabilizing growth



## batch normalization: weight scale*

- if BN is connected at the output activation of an affine layer

$$
\mathbf{a}=W^{\top} \mathbf{x}+\mathbf{b}, \quad \mathbf{x}^{\prime}=h(\mathbf{a})=h\left(W^{\top} \mathbf{x}+\mathbf{b}\right)
$$

the bias $\mathbf{b}$ is absorbed into $\boldsymbol{\beta}$ and the layer is replaced by

$$
\mathbf{x}^{\prime}=h\left(\mathrm{BN}\left(W^{\top} \mathbf{x}\right)\right)
$$

- the layer and its Jacobian are then unaffected by weight scale

$$
\begin{aligned}
\operatorname{BN}\left(a W^{\top} \mathbf{x}\right) & =\mathrm{BN}\left(W^{\top} \mathbf{x}\right) \\
\frac{\partial \mathrm{BN}\left(a W^{\top} \mathbf{x}\right)}{\partial \mathbf{x}} & =\frac{\partial \mathrm{BN}\left(W^{\top} \mathbf{x}\right)}{\partial \mathbf{x}}
\end{aligned}
$$

- moreover, larger weights yield smaller gradients, stabilizing growth


## batch normalization: weight scale*

- if BN is connected at the output activation of an affine layer

$$
\mathbf{a}=W^{\top} \mathbf{x}+\mathbf{b}, \quad \mathbf{x}^{\prime}=h(\mathbf{a})=h\left(W^{\top} \mathbf{x}+\mathbf{b}\right)
$$

the bias $\mathbf{b}$ is absorbed into $\boldsymbol{\beta}$ and the layer is replaced by

$$
\mathbf{x}^{\prime}=h\left(\operatorname{BN}\left(W^{\top} \mathbf{x}\right)\right)
$$

- the layer and its Jacobian are then unaffected by weight scale

$$
\begin{aligned}
\mathrm{BN}\left(a W^{\top} \mathbf{x}\right) & =\mathrm{BN}\left(W^{\top} \mathbf{x}\right) \\
\frac{\partial \mathrm{BN}\left(a W^{\top} \mathbf{x}\right)}{\partial \mathbf{x}} & =\frac{\partial \mathrm{BN}\left(W^{\top} \mathbf{x}\right)}{\partial \mathbf{x}}
\end{aligned}
$$

- moreover, larger weights yield smaller gradients, stabilizing growth

$$
\frac{\partial \mathrm{BN}\left(a W^{\top} \mathbf{x}\right)}{\partial(a W)}=\frac{1}{a} \frac{\partial \mathrm{BN}\left(W^{\top} \mathbf{x}\right)}{\partial W}
$$

## batch normalization: modified GoogLeNet



- allows to
- increase learning rate, accelerate learning rate decay
- reduce weight decay, reduce or remove dropout
- remove data augmentation such as photometric distortions
- remove local response normalization


## layer normalization* <br> [Ba et al. 2016]

- the LN layer with parameters $\boldsymbol{\beta}, \boldsymbol{\gamma}$ yields, for each sample feature $\mathbf{x}=\left(x_{1}, \ldots, x_{k}\right)$,

$$
\mathbf{y}=\operatorname{LN}_{\boldsymbol{\beta}, \boldsymbol{\gamma}}(\mathbf{x}):=\gamma \frac{\mathbf{x}-\mu}{\sqrt{v+\delta}}+\boldsymbol{\beta}
$$

(element-wise), where $\mu, v$ are the sample mean and variance

$$
\begin{aligned}
\mu & :=\frac{1}{k} \sum_{j=1}^{k} x_{j} \\
v & :=\frac{1}{k} \sum_{j=1}^{k}\left(x_{j}-\mu\right)^{2}
\end{aligned}
$$

- training and inference are now identical and independent of mini-batch


## group normalization*

[Wu and He 2018]


- training and inference are identical and independent of mini-batch like layer normalization.
- statistics are measured over groups of channels


## group normalization*




- ResNet50 validation error on ImageNet
- batch norm is sensitive to mini-batch size, group norm is not


## weight normalization*

## [Salimans and Kingma 2016]

- considering a single affine unit $\mathbf{y}=h\left(\mathbf{w}^{\top} \mathbf{x}+b\right)$, weights $\mathbf{w}$ are re-parametrized

$$
\mathbf{w}=g \frac{\mathbf{v}}{\|\mathbf{v}\|}
$$

- its derivatives are given by

$$
d g=d \mathbf{w}^{\top} \frac{\mathbf{v}}{\|\mathbf{v}\|}, \quad d \mathbf{v}^{\top}=\frac{g}{\|\mathbf{v}\|} d \mathbf{w}^{\top}\left(I-\frac{\mathbf{v} \mathbf{v}^{\top}}{\|\mathbf{v}\|^{2}}\right)
$$

- $d \mathbf{w}$ is scaled by $\frac{g}{\|\mathbf{v}\|}$ and projected in a direction normal to $\mathbf{v}$ (and $\mathbf{w}$ )
- during learning, $\|\mathbf{v}\|$ increases monotonically: $\left\|\mathbf{v}^{(\tau+1)}\right\| \geq\left\|\mathbf{v}^{(\tau)}\right\|$
- if $\|d \mathbf{v}\|$ is large, the scaling factor $\frac{g}{\|\mathbf{v}\|}$ decreases; and if it is small, $\|\mathbf{v}\|$ stops increasing: the effect is similar to RMSprop


## summary (so far)

- the deeper the network, the more we need to learn all parameters at the same rate
- in the absence of second order derivatives, optimizers attempt to do so by moving averages and normalization over the training iterations and parameter derivatives are initially well balanced
- it is more effective to modify the objective function itself such that these properties are maintained during optimization


## summary (so far)

- the deeper the network, the more we need to learn all parameters at the same rate
- in the absence of second order derivatives, optimizers attempt to do so by moving averages and normalization over the training iterations
- initialization should be designed such that activations, their derivatives and parameter derivatives are initially well balanced
- it is more effective to modify the objective function itself such that these properties are maintained during optimization


## deeper architectures

## going even deeper



- when initialization, normalization and optimization are appropriately addressed, we can train networks with 50 layers "from scratch"
- a degradation of test error is now exposed with increasing depth, which looks like overfitting (CIFAR10 shown here)
- however, the same degradation appears also at training error


## going even deeper




- when initialization, normalization and optimization are appropriately addressed, we can train networks with 50 layers "from scratch"
- a degradation of test error is now exposed with increasing depth, which looks like overfitting (CIFAR10 shown here)
- however, the same degradation appears also at training error


## remember GoogLeNet auxiliary classifiers?



- GoogLeNet has two auxiliary classifiers that are discarded at inference
- these classifiers inject gradient signal deeper backwards
- we now transform the network in ways that are not necessarily equivalent, but maintain this backward flow pattern
- the result is two skip


## remember GoogLeNet auxiliary classifiers?



- GoogLeNet has two auxiliary classifiers that are discarded at inference
- these classifiers inject gradient signal deeper backwards
we now transform the network in
ways that are not necessarily equivalent, but maintain this backward flow pattern
- the result is two


## remember GoogLeNet auxiliary classifiers?



- GoogLeNet has two auxiliary classifiers that are discarded at inference
- these classifiers inject gradient signal deeper backwards
- we now transform the network in ways that are not necessarily equivalent, but maintain this backward flow pattern
- the result is two


## remember GoogLeNet auxiliary classifiers?



- GoogLeNet has two auxiliary classifiers that are discarded at inference
- these classifiers inject gradient signal deeper backwards
- we now transform the network in ways that are not necessarily equivalent, but maintain this backward flow pattern



## remember GoogLeNet auxiliary classifiers?



- GoogLeNet has two auxiliary classifiers that are discarded at inference
- these classifiers inject gradient signal deeper backwards
- we now transform the network in ways that are not necessarily equivalent, but maintain this backward flow pattern



## remember GoogLeNet auxiliary classifiers?



- GoogLeNet has two auxiliary classifiers that are discarded at inference
- these classifiers inject gradient signal deeper backwards
- we now transform the network in ways that are not necessarily equivalent, but maintain this backward flow pattern
- the result is two skip connections that can be maintained at inference


## skip connections are not new

the network diagram:

represents a four-layer series-coupled system, whereas the diagram

represents a three-layer cross coupled system, since all A-units are at least the same logical distance from the sensory units (see Definition 18,

## ImageNet classification performance



# residual networks 

[He et al. 2016]



- $3.57 \%$ top-5 error on ILSVRC'15
- won first place on several ILSVRC and COCO 2015 tasks
- depth increased to 152 layers, kernel size mostly $3 \times 3$
- residual unit repeated up to 50 times
- $1 \times 1$ kernels used as "bottleneck" layers
- up to $10 \times$ more operations but same parameters as AlexNet


## skip connections and residual



- by copying the features of a shallow model and setting the new mapping to the identity, a deeper model performs at least as well as the shallow one
- "if an identity mapping were optimal, it would be easier to push a residual to zero than to fit an identity mapping by a stack of nonlinear layers"


## skip connections and residual



- "plain" unit: $f$ is the mapping

$$
\mathbf{y}=f(\mathbf{x})
$$

- residual unit: $f$ is the residual

$$
\mathbf{y}=\mathbf{x}+f(\mathbf{x})
$$

- by copying the features of a shallow model and setting the new mapping to the identity, a deeper model performs at least as well as the shallow one
- "if an identity mapping were optimal, it would be easier to push a residual to zero than to fit an identity mapping by a stack of nonlinear layers"


## skip connections and residual



- "plain" unit: $f$ is the mapping

$$
\mathbf{y}=f(\mathbf{x})
$$

- residual unit: $f$ is the residual

$$
\mathbf{y}=\mathbf{x}+f(\mathbf{x})
$$

- by copying the features of a shallow model and setting the new mapping to the identity, a deeper model performs at least as well as the shallow one
> "if an identity mapping were optimal, it would be easier to push a residual to zero than to fit an identity mapping by a stack of nonlinear layers"


## skip connections and residual



- "plain" unit: $f$ is the mapping

$$
\mathbf{y}=f(\mathbf{x})
$$

- residual unit: $f$ is the residual

$$
\mathbf{y}=\mathbf{x}+f(\mathbf{x})
$$

- by copying the features of a shallow model and setting the new mapping to the identity, a deeper model performs at least as well as the shallow one
- "if an identity mapping were optimal, it would be easier to push a residual to zero than to fit an identity mapping by a stack of nonlinear layers"


## residual unit



- "plain" unit, with nonlinearities shown separately, and batch normalization included in each convolutional layers


## residual unit



$$
\operatorname{res}(3,64)
$$

- residual unit, with a skip connection over the two convolutional layers and the relu between them


## residual unit



$$
\operatorname{res}(3,64, s 2)
$$

- stride 2 in the first convolutional layer, along with downsampling on the skip connection


## residual unit



$$
\operatorname{res}(3,128, s 2)
$$

- increasing the number of features, along with a $1 \times 1$ convolution on the skip connection to project to the new feature space


## residual unit



$$
\operatorname{res}(3,128, s 2)
$$

- which is the same as a single $1 \times 1$ convolution with stride 2 , both downsampling and projecting


## residual bottleneck unit



- "plain" bottleneck unit, with $1 \times 1$ convolutions

He, Zhang, Ren, Sun. CVPR 2016. Deep Residual Learning for Image Recognition.

## residual bottleneck unit



$$
\operatorname{res}(3,(64,256))
$$

- residual bottleneck unit with a skip connection, always projecting


## residual bottleneck unit



$$
\operatorname{res}(3,(64,256), s 2)
$$

- stride 2 in the first convolutional and the skip layer


## ResNet-34

|  |  | parameters | operations | volume |
| :---: | :---: | :---: | :---: | :---: |
|  | input(224, 3) | 0 | 0 | $224 \times 224 \times 3$ |
|  | $\operatorname{conv}(7,64, p 3, s 2)$ | 9,472 | 118, 816, 768 | $112 \times 112 \times 64$ |
|  | pool $(3,2, p 1)$ | 0 | 802, 816 | $56 \times 56 \times 64$ |
| $3 \times$ | res ( 3,64 ) | 221,568 | 694, 837, 248 | $56 \times 56 \times 64$ |
|  | res(3, 128, $s 2$ ) | 229, 760 | 180, 182, 016 | $28 \times 28 \times 128$ |
| $3 \times$ | res $(3,128)$ | 885,504 | 694, 235, 136 | $28 \times 28 \times 128$ |
|  | res(3, 256, $s 2$ ) | 918,272 | 180, 006, 400 | $14 \times 14 \times 256$ |
| $5 \times$ | res $(3,256)$ | 5, 900, 800 | 1,156, 556, 800 | $14 \times 14 \times 256$ |
|  | res(3, 512, $s 2)$ | 3, 671,552 | 179, 918, 592 | $7 \times 7 \times 512$ |
| $2 \times$ | res $(3,512)$ | 9, 439, 232 | 462, 522, 368 | $7 \times 7 \times 512$ |
|  | $\operatorname{avg}(7)$ | 0 | 25, 088 | 512 |
|  | $\mathrm{fc}(1000)$ | 513, 000 | 513, 000 | 1000 |
|  | softmax | 0 | 1,000 | 1000 |

- $3 \times$ more operations but $3 \times$ less parameters comparing to AlexNet


## ResNet-101

|  |  | parameters | operations | volume |
| :---: | :---: | :---: | :---: | :---: |
|  | input(224, 3) | 0 | 0 | $224 \times 224 \times 3$ |
|  | $\operatorname{conv}(7,64, p 3, s 2)$ | 9,472 | 118, 816, 768 | $112 \times 112 \times 64$ |
|  | pool $(3,2, p 1)$ | 0 | 802, 816 | $56 \times 56 \times 64$ |
| $3 \times$ | $\operatorname{res}(3,(64,256))$ | 214,400 | 672, 358, 400 | $56 \times 56 \times 256$ |
|  | $\operatorname{res}(3,(128,512), s 2)$ | 378,112 | 296,640, 512 | $28 \times 28 \times 512$ |
| $3 \times$ | res(3, $(128,512)$ ) | 837, 888 | 656, 904, 192 | $28 \times 28 \times 512$ |
|  | $\operatorname{res}(3,(256,1024), s 2)$ | 1,509, 888 | 296, 038, 400 | $14 \times 14 \times 1024$ |
| $22 \times$ | res(3, $(256,1024)$ ) | 24, 544, 256 | 4, 810, 674,176 | $14 \times 14 \times 1024$ |
|  | $\operatorname{res}(3,(512,2048), s 2)$ | 6,034,432 | 295, 737, 344 | $7 \times 7 \times 2048$ |
| $2 \times$ | res(3, (512, 2048)) | 8,919, 040 | 437, 032, 960 | $7 \times 7 \times 2048$ |
|  | $\operatorname{avg}(7)$ | 0 | 100, 352 | 2048 |
|  | fc(1000) | 2, 049, 000 | 2, 049, 000 | 1000 |
|  | softmax | 0 | 1,000 | 1000 |

- $7 \times$ more operations but $1.5 \times$ less parameters comparing to AlexNet


## ResNet-34: ImageNet



- a plain network exhibits degradation with increasing depth
- while a residual network gains from increasing depth


## ResNet models

| layer name | output size | 18-layer | 34-layer | 50-layer | 101-layer | 152-layer |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| conv1 | $112 \times 112$ | $7 \times 7,64$, stride 2 |  |  |  |  |
| conv2_x | $56 \times 56$ | $3 \times 3$ max pool, stride 2 |  |  |  |  |
|  |  | $\left[\begin{array}{l}3 \times 3,64 \\ 3 \times 3,64\end{array}\right] \times 2$ | $\left[\begin{array}{l}3 \times 3,64 \\ 3 \times 3,64\end{array}\right] \times 3$ | $\left[\begin{array}{c}1 \times 1,64 \\ 3 \times 3,64 \\ 1 \times 1,256\end{array}\right] \times 3$ | $\left[\begin{array}{c}1 \times 1,64 \\ 3 \times 3,64 \\ 1 \times 1,256\end{array}\right] \times 3$ | $\left[\begin{array}{c}1 \times 1,64 \\ 3 \times 3,64 \\ 1 \times 1,256\end{array}\right] \times 3$ |
| conv3_X | $28 \times 28$ | $\left[\begin{array}{l}3 \times 3,128 \\ 3 \times 3,128\end{array}\right] \times 2$ | $\left[\begin{array}{l}3 \times 3,128 \\ 3 \times 3,128\end{array}\right] \times 4$ | $\left[\begin{array}{l}1 \times 1,128 \\ 3 \times 3,128 \\ 1 \times 1,512\end{array}\right] \times 4$ | $\left[\begin{array}{l}1 \times 1,128 \\ 3 \times 3,128 \\ 1 \times 1,512\end{array}\right] \times 4$ | $\left[\begin{array}{l}1 \times 1,128 \\ 3 \times 3,128 \\ 1 \times 1,512\end{array}\right] \times 8$ |
| conv4_x | $14 \times 14$ | $\left[\begin{array}{l}3 \times 3,256 \\ 3 \times 3,256\end{array}\right] \times 2$ | $\left[\begin{array}{l}3 \times 3,256 \\ 3 \times 3,256\end{array}\right] \times 6$ | $\left[\begin{array}{c}1 \times 1,256 \\ 3 \times 3,256 \\ 1 \times 1,1024\end{array}\right] \times 6$ | $\left[\begin{array}{c}1 \times 1,256 \\ 3 \times 3,256 \\ 1 \times 1,1024\end{array}\right] \times 23$ | $\left[\begin{array}{c}1 \times 1,256 \\ 3 \times 3,256 \\ 1 \times 1,1024\end{array}\right] \times 36$ |
| conv5_x | $7 \times 7$ | $\left[\begin{array}{l}3 \times 3,512 \\ 3 \times 3,512\end{array}\right] \times 2$ | $\left[\begin{array}{l}3 \times 3,512 \\ 3 \times 3,512\end{array}\right] \times 3$ | $\left[\begin{array}{c}1 \times 1,512 \\ 3 \times 3,512 \\ 1 \times 1,2048\end{array}\right] \times 3$ | $\left[\begin{array}{c}1 \times 1,512 \\ 3 \times 3,512 \\ 1 \times 1,2048\end{array}\right] \times 3$ | $\left[\begin{array}{c}1 \times 1,512 \\ 3 \times 3,512 \\ 1 \times 1,2048\end{array}\right] \times 3$ |
|  | $1 \times 1$ |  |  | age pool, $1000-\mathrm{d} \mathrm{fc}$ | softmax |  |
|  | OPs | $1.8 \times 10^{9}$ | $3.6 \times 10^{9}$ | $3.8 \times 10^{9}$ | $7.6 \times 10^{9}$ | $11.3 \times 10^{9}$ |

- downsampling by 2 at layers conv3_1, conv4_1, conv5_1

He, Zhang, Ren, Sun. CVPR 2016. Deep Residual Learning for Image Recognition.

## network performance



## identity mappings*

[He et al. 2016]

- original residual unit, with relu and BN
 shown separately, where $h$ is relu

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- re-designed unit, with a more direct path through skip connections, and relu and BN acting as pre-activation
- recursively, there is a residual between


He, Zhang, Ren and Sun. ECCV 2016. Identity Mappings in Deep Residual Networks.

## identity mappings*

[He et al. 2016]

- original residual unit, with relu and BN
 shown separately, where $h$ is relu

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- re-designed unit, with a more direct path through skip connections, and relu and BN acting as pre-activation

$$
\mathbf{x}_{i+1}=\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)
$$

- recursively, there is a residual between


## identity mappings*

[He et al. 2016]

- original residual unit, with relu and BN
 shown separately, where $h$ is relu

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- re-designed unit, with a more direct path through skip connections, and relu and BN acting as pre-activation

$$
\mathbf{x}_{i+1}=\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)
$$

- recursively, there is a residual between any units $\ell_{1}, \ell_{2}$

$$
\mathbf{x}_{\ell_{2}}=\mathbf{x}_{\ell_{1}}+\sum_{i=\ell_{1}}^{\ell_{2}-1} f_{i}\left(\mathbf{x}_{i}\right)
$$

He, Zhang, Ren and Sun. ECCV 2016. Identity Mappings in Deep Residual Networks.

## residual networks as ensembles*

[Veit et al. 2016]


- residual network with identity mappings
- "unraveled" view where residual units are duplicated
- ensemble of networks of different lengths, with cardinality exponential in network depth
- dronning a laver is just zeroing half of the paths
- in a network of 110 layers, most gradient comes from paths that are 10-34 layers deep


## residual networks as ensembles*

[Veit et al. 2016]


- residual network with identity mappings
- "unraveled" view where residual units are duplicated
- ensemble of networks of different lengths, with cardinality exponential in network depth
- dropping a layer is just zeroing half of the paths
- in a network of 110 layers, most gradient comes from paths that are 10-34 layers deep


## residual networks as ensembles*

[Veit et al. 2016]


- residual network with identity mappings
- "unraveled" view where residual units are duplicated
- ensemble of networks of different lengths, with cardinality exponential in network depth
- dropping a layer is just zeroing half of the paths
- in a network of 110 layers, most gradient comes from paths that are 10-34 layers deep


## residual networks as ensembles*

## [Veit et al. 2016]



- residual network with identity mappings
- "unraveled" view where residual units are duplicated
- ensemble of networks of different lengths, with cardinality exponential in network depth
- dropping a layer is just zeroing half of the paths
- in a network of 110 layers, most
 10-34 layers deep


## residual networks as ensembles*

## [Veit et al. 2016]



- residual network with identity mappings
- "unraveled" view where residual units are duplicated
- ensemble of networks of different lengths, with cardinality exponential in network depth
- dropping a layer is just zeroing half of the paths
- in a network of 110 layers, most gradient comes from paths that are 10-34 layers deep


## networks with stochastic depth*

[Huang et al. 2016]


- (original) residual network
- at each training iteration, randomly drop a subset of layers
where $b_{i} \in\{0,1\}$ a Bernoulli random variable
- at inference, use all layers weighted by survival probabilities $p_{i}=\mathbb{E}\left(b_{i}\right)$

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+p_{i} f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

## - speeds up training, reduces test error

## networks with stochastic depth*

[Huang et al. 2016]


- (original) residual network
- at each training iteration, randomly drop a subset of layers

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+b_{i} f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

where $b_{i} \in\{0,1\}$ a Bernoulli random variable

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+p_{i} f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- speeds up training, reduces test error


## networks with stochastic depth*

## [Huang et al. 2016]



- (original) residual network
- at each training iteration, randomly drop a subset of layers

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+b_{i} f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

where $b_{i} \in\{0,1\}$ a Bernoulli random variable

- at inference, use all layers weighted by survival probabilities $p_{i}=\mathbb{E}\left(b_{i}\right)$

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+p_{i} f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- speeds up training, reduces test error


## networks with stochastic depth*

## [Huang et al. 2016]



- (original) residual network
- at each training iteration, randomly drop a subset of layers

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+b_{i} f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

where $b_{i} \in\{0,1\}$ a Bernoulli random variable

- at inference, use all layers weighted by survival probabilities $p_{i}=\mathbb{E}\left(b_{i}\right)$

$$
\mathbf{x}_{i+1}=h\left(\mathbf{x}_{i}+p_{i} f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- speeds up training, reduces test error


## densely connected networks

[Huang et al. 2017]


- residual unit with identity mapping: add

$$
\mathbf{x}_{i+1}=\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)
$$

- densely connected unit: concatenate
- feature map dimension increases by
growth rate $k$ at each unit
- a dense block is a chain of densely connected units
- a transition laver reduces feature map dimension by a factor $\theta=2$


## densely connected networks

[Huang et al. 2017]


- residual unit with identity mapping: add

$$
\mathbf{x}_{i+1}=\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)
$$

- densely connected unit: concatenate

$$
\mathbf{x}_{i+1}=\left(\mathbf{x}_{i}, f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- feature map dimension increases by growth rate $k$ at each unit
 connected units



## densely connected networks

[Huang et al. 2017]


- residual unit with identity mapping: add

$$
\mathbf{x}_{i+1}=\mathbf{x}_{i}+f_{i}\left(\mathbf{x}_{i}\right)
$$

- densely connected unit: concatenate

$$
\mathbf{x}_{i+1}=\left(\mathbf{x}_{i}, f_{i}\left(\mathbf{x}_{i}\right)\right)
$$

- feature map dimension increases by growth rate $k$ at each unit
- a dense block is a chain of densely connected units
- a transition layer reduces feature map dimension by a factor $\theta=2$


## densely connected networks



- dense block followed by transition layer

Huang, Liu, van der Maaten and Weinberger. CVPR 2017. Densely Connected Convolutional Networks.

## DenseNet models

| Layers | Output Size | DenseNet-121( $k=32$ ) | DenseNet-169 $(k=32)$ | DenseNet-201 $(k=32)$ | DenseNet-161 $(k=48)$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Convolution | $112 \times 112$ | $7 \times 7$ conv, stride 2 |  |  |  |
| Pooling | $56 \times 56$ | $3 \times 3$ max pool, stride 2 |  |  |  |
| Dense Block <br> (1) | $56 \times 56$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 6$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 6$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 6$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 6$ |
| Transition Layer | $56 \times 56$ | $1 \times 1$ conv |  |  |  |
| (1) | $28 \times 28$ | $2 \times 2$ average pool, stride 2 |  |  |  |
| Dense Block <br> (2) | $28 \times 28$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 12$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 12$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 12$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 12$ |
| Transition Layer | $28 \times 28$ | $1 \times 1$ conv |  |  |  |
| (2) | $14 \times 14$ | $2 \times 2$ average pool, stride 2 |  |  |  |
| Dense Block <br> (3) | $14 \times 14$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 24$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 32$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 48$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 36$ |
| Transition Layer | $14 \times 14$ | $1 \times 1$ conv |  |  |  |
| (3) | $7 \times 7$ | $2 \times 2$ average pool, stride 2 |  |  |  |
| Dense Block <br> (4) | $7 \times 7$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 16$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 32$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 32$ | $\left[\begin{array}{l}1 \times 1 \text { conv } \\ 3 \times 3 \text { conv }\end{array}\right] \times 24$ |
| Classification | $1 \times 1$ | $7 \times 7$ global average pool |  |  |  |
| Layer |  | 1000D fully-connected, softmax |  |  |  |

- input is $224 \times 224$; first convolutional layer produces $2 k$ features; transition layer reduces dimension and resolution by 2


## DenseNet vs. ResNet: ImageNet



- top-1 single-crop ImageNet validation error
- encourages feature re-use and reduces the number of parameters


## summary

- optimizers: gradient descent, momentum, RMSprop, Adam, Hessian-free*
- initialization: Gaussian matrices, unit variance, orthogonal*, data-dependent*
- normalization: input, batch, layer*, group*, weight*
- deeper architectures: residual networks, identity mappings*, networks with stochastic depth*, densely connected networks
- all parameters should be learned at the same rate, and all features computed by some layer should be re-used by the following layers
- initialization, normalization and architecture should be designed such that these properties hold initially and are maintained during training

