# Deep Learning 

Pontus Giselsson

## Outline

## - Deep learning

- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Deep learning

- Can be used both for classification and regression
- Deep learning training problem is of the form

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)
$$

where typically

- $L(u, y)=\frac{1}{2}\|u-y\|_{2}^{2}$ is used for regression
- $L(u, y)=\log \left(\sum_{j=1}^{K} e^{u_{j}}\right)-y^{T} u$ is used for $K$-class classification
- Difference to previous convex methods: Nonlinear model $m(x ; \theta)$
- Deep learning regression generalizes least squares
- DL classification generalizes multiclass logistic regression
- Nonlinear model makes training problem nonconvex


## Deep learning - Model

- Nonlinear model of the following form is often used:
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$ where $\theta$ contains all $W_{i}$ and $b_{i}$
- Each activation $\sigma_{j}$ constitutes a hidden layer in the model network
- We have no final layer activation (is instead part of loss)
- Graphical representation with three hidden layers

- Some reasons for using this structure:
- (Assumed) universal function approximators
- Efficient gradient computation using backpropagation


## No final layer activation in classification

- In classification, it is common to use
- Softmax final layer activation
- Cross entropy loss function
- Equivalent to
- no (identity) final layer activation
- multiclass logistic loss
which is what we use


## Activation functions

- Activation function $\sigma_{j}$ takes as input the output of $W_{j}(\cdot)+b_{j}$
- Often a function $\bar{\sigma}_{j}: \mathbb{R} \rightarrow \mathbb{R}$ is applied to each element
- Example: $\sigma_{j}: \mathbb{R}^{3} \rightarrow \mathbb{R}^{3}$ is $\sigma_{j}(u)=\left[\begin{array}{l}\bar{\sigma}_{j}\left(u_{1}\right) \\ \bar{\sigma}_{j}\left(u_{2}\right) \\ \bar{\sigma}_{j}\left(u_{3}\right)\end{array}\right]$
- We will use notation over-loading and call both functions $\sigma_{j}$


## Examples of activation functions

Graph


Tanh
$\frac{e^{u}-e^{-u}}{e^{-u}+e^{u}}$

ReLU
$\max (u, 0)$

LeakyReLU $\max (u, \alpha u)$

ELU

$$
\begin{cases}u & \text { if } u \geq 0 \\ \alpha\left(e^{u}-1\right) & \text { else }\end{cases}
$$



## Examples of affine transformations

- Dense (fully connected): Dense $W_{j}$
- Sparse: Sparse $W_{j}$
- Convolutional layer (convolution with small pictures)
- Fixed (random) sparsity pattern
- Subsampling: reduce size, $W_{j}$ fat (smaller output than input)
- average pooling


## Prediction

- Prediction as in least squares and multiclass logistic regression
- Assume model $m(x ; \theta)$ trained and "optimal" $\theta^{\star}$ found
- Regression:
- Predict response for new data $x$ using $\hat{y}=m\left(x ; \theta^{\star}\right)$
- Classification (with no final layer activation):
- We have one model $m_{j}\left(x ; \theta^{\star}\right)$ output for each class
- Predict class belonging for new data $x$ according to

$$
\underset{j \in\{1, \ldots, K\}}{\operatorname{argmax}} m_{j}\left(x ; \theta^{\star}\right)
$$

i.e., class with largest model value (since loss designed this way)

## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Learning features

- Convex methods use prespecified feature maps (or kernels)
- Deep learning instead learns feature map during training
- Define parameter dependent feature vector:

$$
\phi(x ; \theta):=\sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)
$$

- Model becomes $m(x ; \theta)=W_{n} \phi(x ; \theta)+b_{n}$
- Inserted into training problem:

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(W_{n} \phi\left(x_{i} ; \theta\right)+b_{n}, y_{i}\right)
$$

same as before, but with learned (parameter-dependent) features

- Learning features at training makes training nonconvex


## Learning features - Graphical representation

- Fixed features gives convex training problems

- Learning features gives nonconvex training problems

- Output of last activation function is feature vector


## Optimizing only final layer

- Assume:
- that parameters $\bar{\theta}_{f}$ in the layers in the square are fixed
- that we optimize only the final layer parameters
- that the loss is a (binary) logistic loss

- What can you say about the training problem?


## Optimizing only final layer

- Assume:
- that parameters $\bar{\theta}_{f}$ in the layers in the square are fixed
- that we optimize only the final layer parameters
- that the loss is a (binary) logistic loss

- What can you say about the training problem?
- It reduces to logistic regression with fixed features $\phi\left(x_{i} ; \bar{\theta}_{f}\right)$

$$
\operatorname{minimize}_{\theta=\left(W_{n}, b_{n}\right)} \sum_{i=1}^{N} L\left(W_{n} \phi\left(x_{i} ; \bar{\theta}_{f}\right)+b_{n}, y_{i}\right)
$$

- The training problem is convex


## Design choices

Many design choices in building model to create good features

- Number of layers
- Width of layers
- Types of layers
- Types of activation functions
- Different model structures (e.g., residual network)


## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Model properties - ReLU networks

- Recall model
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$ where $\theta$ contains all $W_{i}$ and $b_{i}$
- Assume that all activation functions are (Leaky)ReLU
- What can you say about the properties of $m(\cdot ; \theta)$ for fixed $\theta$ ?


## Model properties - ReLU networks

- Recall model
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$ where $\theta$ contains all $W_{i}$ and $b_{i}$
- Assume that all activation functions are (Leaky)ReLU
- What can you say about the properties of $m(\cdot ; \theta)$ for fixed $\theta$ ?
- It is continuous piece-wise affine


## 1D Regression - Model properties

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyReLU



## 1D Regression - Model properties

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyReLU

- Vertical lines show kinks


## 1D Regression - Model properties

- Fully connected, layers widths: 5,5,5,1,1 (78 params), Tanh

- No kinks for Tanh


## Identity activation

- Do we need nonlinear activation functions?
- What can you say about model if all $\sigma_{j}=\mathrm{Id}$ in
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$
where $\theta$ contains all $W_{j}$ and $b_{j}$


## Identity activation

- Do we need nonlinear activation functions?
- What can you say about model if all $\sigma_{j}=$ Id in
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$ where $\theta$ contains all $W_{j}$ and $b_{j}$
- We then get

$$
\begin{aligned}
m(x ; \theta) & :=W_{n}\left(W_{n-1}\left(\cdots\left(W_{2}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n} \\
& =\underbrace{W_{n} W_{n-1} \cdots W_{2} W_{1}}_{W} x+\underbrace{b_{n}+\sum_{l=2}^{n-1} W_{n} \cdots W_{l} b_{l-1}}_{b} \\
& =W x+b
\end{aligned}
$$

which is linear in $x$ (but training problem nonconvex)

## Network with identity activations - Example

- Fully connected, layers widths: 5,5,5,1,1 (78 params), Identity



## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Training problem properties

- Recall model
$m(x ; \theta):=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$
where $\theta$ includes all $W_{j}$ and $b_{j}$ and training problem

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)
$$

- If all $\sigma_{j}$ LeakyReLU and $L(u, y)=\frac{1}{2}\|u-y\|_{2}^{2}$, then for fixed $x, y$
- $m(x ; \cdot)$ is continuous piece-wise polynomial (cpp) of degree $n$ in $\theta$
- $L(m(x ; \theta), y)$ is cpp of degree $2 n$ in $\theta$
where both model output and loss can grow fast
- If $\sigma_{j}$ is instead Tanh
- model no longer piece-wise polynomial (but "more" nonlinear)
- model output grows slower since $\sigma_{j}: \mathbb{R} \rightarrow(-1,1)$


## Loss landscape - Leaky ReLU

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- First choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Leaky ReLU

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- Second choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Leaky ReLU

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- Third choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Tanh

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- First choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Tanh

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where
- $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- Second choice of $\theta_{1}$ and $\theta_{2}$ :



## Loss landscape - Tanh

- Fully connected, layers widths: 5,5,5,1,1 (78 params), LeakyRelu
- Regression problem, least squares loss
- Plot: $\sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta^{\star}+t_{1} \theta_{1}+t_{2} \theta_{2}\right), y_{i}\right)$ vs scalars $t_{1}, t_{2}$, where - $\theta^{\star}$ is numerically found solution to training problem
- $\theta_{1}$ and $\theta_{2}$ are random directions in parameter space
- Third choice of $\theta_{1}$ and $\theta_{2}$ :



## ReLU vs Tanh

Previous figures suggest:

- ReLU: more regular and similar loss landscape?
- Tanh: less steep (on macro scale)?
- Tanh: Minima extend over larger regions?


## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Performance with increasing depth

- Increasing depth can deteriorate performance
- Deep networks may even have worse training errors than shallow
- Intuition: deeper layers bad at approximating identity mapping


## Residual networks

- Add skip connections between layers
- Instead of network architecture with $z_{1}=x_{i}$ (see figure):

$$
z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right) \text { for } j \in\{1, \ldots, n-1\}
$$

use residual architecture

$$
z_{j+1}=z_{j}+\sigma_{j}\left(W_{j} z_{j}+b_{j}\right) \text { for } j \in\{1, \ldots, n-1\}
$$

- Assume $\sigma(0)=0, W_{j}=0, b_{j}=0$ for $j=1, \ldots, m(m<n-1)$ $\Rightarrow$ deeper part of network is identity mapping and does no harm
- Learns variation from identity mapping (residual)



## Graphical representation

For graphical representation, first collapse nodes into single node


## Graphical representation

- Collapsed network representation

- Residual network

- If some $h_{j}=0$ gives same performance as shallower network


## Residual network - Example

- Fully connected - no residual layers, LeakyReLU activation
- Layers widths: $3 \times 5,1,1$ (depth: 5, 78 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - no residual layers, LeakyReLU activation
- Layers widths: $5 \times 5,1,1$ (depth: 7, 138 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - no residual layers, LeakyReLU activation
- Layers widths: $10 \times 5,1,1$ (depth: 12, 288 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - no residual layers, LeakyReLU activation
- Layers widths: $15 \times 5,1,1$ (depth: 17, 438 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - no residual layers, LeakyReLU activation
- Layers widths: $45 \times 5,1,1$ (depth: 47, 1,338 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - residual layers, LeakyReLU activation
- Layers widths: 3x5,1,1 (depth: 5, 78 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - residual layers, LeakyReLU activation
- Layers widths: $5 \times 5,1,1$ (depth: 7, 138 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - residual layers, LeakyReLU activation
- Layers widths: $10 \times 5,1,1$ (depth: 12, 288 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - residual layers, LeakyReLU activation
- Layers widths: $15 \times 5,1,1$ (depth: 17, 438 params)
- Trained for 5000 epochs



## Residual network - Example

- Fully connected - residual layers, LeakyReLU activation
- Layers widths: $45 \times 5,1,1$ (depth: 47, 1,338 params)
- Trained for 5000 epochs



## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Why overparameterization?

- Neural networks are often overparameterized in practice
- Why? They often perform better than underparameterized


## What is overparameterization?

- We mean that many solutions exist that can:
- fit all data points (0 training loss) in regression
- correctly classify all training examples in classification
- This requires (many) more parameters than training examples
- Need wide and deep enough networks
- Can result in overfitting
- Questions:
- Which of all solutions give best generalization?
- (How) can network design affect generalization?



## Overparameterization - An example

- Assume fully connected network with
- input data $x_{i} \in \mathbb{R}^{p}$
- $n$ layers and $N \approx p^{2}$ samples
- same width throughout (except last layer, which can be neglected)
- What is the relation between number of weights and samples?



## Overparameterization - An example

- Assume fully connected network with
- input data $x_{i} \in \mathbb{R}^{p}$
- $n$ layers and $N \approx p^{2}$ samples
- same width throughout (except last layer, which can be neglected)
- What is the relation between number of weights and samples?

- We have:
- Number of parameters approximately: $\left(W_{j}\right)_{l k}: p^{2} n$ and $\left(b_{j}\right)_{l}: p n$
- Then $\frac{\text { \#weights }}{\# \text { samples }} \approx \frac{p^{2} n}{p^{2}}=n$ more weights than samples


## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Generalization

- Most important for model to generalize well to unseen data
- General approach in training
- Train a model that is too expressive for the underlying data
- Overparameterization in deep learning
- Use regularization to
- find model of appropriate (lower) complexity
- favor models with desired properties


## Regularization

What regularization techniques in DL are you familiar with?

## Regularization techniques

- Reduce number of parameters
- Sparse weight tensors (e.g., convolutional layers)
- Subsampling (gives fewer parameters deeper in network)
- Explicit regularization term in cost function, e.g., Tikhonov
- Data augmentation - more samples, artificial often OK
- Early stopping - stop algorithm before convergence
- Dropouts
- ...


## Implicit vs explicit regularization

- Regularization can be explicit or implicit
- Explicit - Introduce something with intent to regularize:
- Add cost function to favor desirable properties
- Design (adapt) network to have regularizing properties
- Implicit - Use something with regularization as byproduct:
- Use algorithm that finds favorable solution among many
- Will look at implicit regularization via SGD


## Generalization - Our focus

Will here discuss generalization via:

- Norm of parameters - leads to implicit regularization via SGD
- Flatness of minima - leads to implicit regularization via SGD


## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Lipschitz continuity of ReLU networks

- Assume that all activation functions 1-Lipschitz continuous
- The neural network model $m(\cdot ; \theta)$ is Lipschitz continuous in $x$,

$$
\left\|m\left(x_{1} ; \theta\right)-m\left(x_{2} ; \theta\right)\right\|_{2} \leq L\left\|x_{1}-x_{2}\right\|_{2}
$$

for fixed $\theta$, e.g., the $\theta$ obtained after training

- This means output differerences are bounded by input differences
- A Lipschitz constant $L$ is given by

$$
L=\left\|W_{n}\right\|_{2} \cdot\left\|W_{n-1}\right\|_{2} \cdots\left\|W_{1}\right\|_{2}
$$

since activation functions are 1-Lipschitz continuous

- For residual layers each $\left\|W_{j}\right\|_{2}$ replaced by $\left(1+\left\|W_{j}\right\|_{2}\right)$


## Desired Lipschitz constant

- Overparameterization gives many solutions that perfectly fit data
- Would you favor one with high or low Lipschitz constant $L$ ?


## Small norm likely to generalize better

- Smaller Lipschitz constant probably generalizes better if perfect fit
- "Similar inputs give similar outputs", recall

$$
\left\|m\left(x_{1} ; \theta\right)-m\left(x_{2} ; \theta\right)\right\|_{2} \leq L\left\|x_{1}-x_{2}\right\|_{2}
$$

with a Lipschitz constant is given by

$$
L=\left\|W_{n}\right\|_{2} \cdot\left\|W_{n-1}\right\|_{2} \cdots\left\|W_{1}\right\|_{2}
$$

or with $\left\|W_{j}\right\|_{2}$ replaced by $\left(1+\left\|W_{j}\right\|_{2}\right)$ for residual layers

- Smaller weight norms give better generalization if perfect fit


## Generalization - Norm of weights

- Fully connected - residual layers, LeakyReLU
- Layers widths: $30 \times 5,1,1$ ( 888 params)
- Norm of weights (with perfect fit): 72



## Generalization - Norm of weights

- Fully connected - residual layers, LeakyReLU
- Layers widths: $30 \times 5,1,1$ ( 888 params)
- Norm of weights (with perfect fit): 540



## Generalization - Norm of weights

- Fully connected - residual layers, LeakyReLU
- Layers widths: $30 \times 5,1,1$ (888 params)
- Norm of weights (with perfect fit): 540

- Same as previous, new scaling


## Generalization - Norm of weights

- Fully connected - residual layers, LeakyReLU
- Layers widths: $30 \times 5,1,1$ ( 888 params)
- Norm of weights (with perfect fit): 595

- Large norm, but seemingly fair generalization


## Generalization - Norm of weights

- Fully connected - residual layers, LeakyReLU
- Layers widths: $30 \times 5,1,1$ (888 params)
- Norm of weights (with perfect fit): 595

variable $x$
- Same as previous, new scaling


## Generalization - Norm of weights

- Fully connected - residual layers, LeakyReLU
- Layers widths: $30 \times 5,1,1$ (888 params)
- Norm of weights (with perfect fit): 72

- Same as first, new scaling - overfits less than large norm solutions


## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Flatness of minima

- Consider the following illustration of average loss:

- Depicts test loss as shifted training loss
- Motivation to that flat minima generalize better than sharp


## Flatness of minima

- Consider the following illustration of average loss:

- Depicts test loss as shifted training loss
- Motivation to that flat minima generalize better than sharp
- Is there a limitation in considering the average loss only?


## Generalization from loss landscape

- Training set $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{N}$ and training problem:

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)
$$

- Test set $\left\{\left(\hat{x}_{i}, \hat{y}_{i}\right)\right\}_{i=1}^{\hat{N}}, \theta$ generalizes well if test loss small

$$
\sum_{i=1}^{\hat{N}} L\left(m\left(\hat{x}_{i} ; \theta\right), \hat{y}_{i}\right)
$$

- By overparameterization, we can for each $\left(\hat{x}_{i}, \hat{y}_{i}\right)$ find $\hat{\theta}_{i}$ so that

$$
L\left(m\left(\hat{x}_{i} ; \theta\right), \hat{y}_{i}\right)=L\left(m\left(x_{j_{i}} ; \theta+\hat{\theta}_{i}\right), y_{j_{i}}\right)
$$

for all $\theta$ given a (similar) $\left(x_{j_{i}}, y_{j_{i}}\right)$ pair in training set

- Evaluate test loss by training loss at shifted points $\theta+\hat{\theta}_{\hat{\theta}}{ }^{1)}$
- Test loss small if original individual loss small at all $\theta+\hat{\theta}_{i}$
- Previous figure used same $\hat{\theta}_{i}=\hat{\theta}$ for all $i$

[^0]
## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?



## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses
- Let us evaluate test loss by shifting individual training losses


## Example

- Can flat (local) minima be different?
- Does one of the following minima generalize better?

- It depends on individual losses
- Let us evaluate test loss by shifting individual training losses
- Do not only want flat minima, want individual losses flat at minima


## Individually flat minima

- Both flat minima have $\nabla f(\theta)=0$, but
- One minima has large individual gradients $\left\|\nabla f_{i}(\theta)\right\|$
- Other minima has small individual gradients $\left\|\nabla f_{i}(\theta)\right\|$
- The latter (individually flat minima) seems to generalize better
- Want individually flat minima (with small $\left\|\nabla f_{i}(\theta)\right\|$ )
- This implies average flat minima
- The reverse implication may not hold
- Overparameterized networks:
- The reverse implication may often hold at global minima
- Why? $f(\theta)=0$ and $\nabla f(\theta)=0$ implies $f_{i}(\theta)=0$ and $\nabla f_{i}(\theta)=0$


## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Training algorithm

- Neural networks often trained using stochastic gradient descent
- DNN weights are updated via gradients in training
- Gradient of cost is sum of gradients of summands (samples)
- Gradient of each summand computed using backpropagation


## Backpropagation

- Backpropagation is reverse mode automatic differentiation
- Based on chain-rule in differentiation
- Backpropagation must be performed per sample
- Our derivation assumes:
- Fully connected layers ( $W$ full, if not, set elements in $W$ to 0 )
- Activation functions $\sigma_{j}(v)=\left(\sigma_{j}\left(v_{1}\right), \ldots, \sigma_{j}\left(v_{p}\right)\right)$ element-wise (overloading of $\sigma_{j}$ notation)
- Weights $W_{j}$ are matrices, samples $x_{i}$ and responses $y_{i}$ are vectors
- No residual connections


## Jacobians

- The Jacobian of a function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ is given by

$$
\frac{\partial f}{\partial x}=\left[\begin{array}{ccc}
\frac{\partial f_{1}}{\partial x_{1}} & \cdots & \frac{\partial f_{1}}{\partial x_{n}} \\
\vdots & \vdots & \vdots \\
\frac{\partial f_{m}}{\partial x_{1}} & \cdots & \frac{\partial f_{m}}{\partial x_{n}}
\end{array}\right] \in \mathbb{R}^{m \times n}
$$

- The Jacobian of a function $f: \mathbb{R}^{p \times n} \rightarrow \mathbb{R}$ is given by

$$
\frac{\partial f}{\partial x}=\left[\begin{array}{ccc}
\frac{\partial f}{\partial x_{11}} & \cdots & \frac{\partial f}{\partial x_{1 n}} \\
\vdots & \vdots & \vdots \\
\frac{\partial f}{\partial x_{p 1}} & \cdots & \frac{\partial f}{\partial x_{p n}}
\end{array}\right] \in \mathbb{R}^{p \times n}
$$

- The Jacobian of a function $f: \mathbb{R}^{p \times n} \rightarrow \mathbb{R}^{m}$ is at layer $j$ given by

$$
\left[\frac{\partial f}{\partial x}\right]_{:, j,:}=\left[\begin{array}{ccc}
\frac{\partial f_{1}}{\partial x_{j 1}} & \cdots & \frac{\partial f_{1}}{\partial x_{j n}} \\
\vdots & \vdots & \vdots \\
\frac{\partial f_{m}}{\partial x_{j 1}} & \cdots & \frac{\partial f_{m}}{\partial x_{j n}}
\end{array}\right] \in \mathbb{R}^{m \times n}
$$

the full Jacobian is a 3D tensor in $\mathbb{R}^{m \times p \times n}$

## Jacobian vs gradient

- The Jacobian of a function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is given by

$$
\frac{\partial f}{\partial x}=\left[\begin{array}{lll}
\frac{\partial f}{\partial x_{1}} & \cdots & \frac{\partial f}{\partial x_{n}}
\end{array}\right]
$$

- The gradient of a function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is given by

$$
\nabla f=\left[\begin{array}{c}
\frac{\partial f}{\partial x_{1}} \\
\vdots \\
\frac{\partial f}{\partial x_{n}}
\end{array}\right]
$$

i.e., transpose of Jacobian for $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$

- Chain rule holds for Jacobians:

$$
\frac{\partial f}{\partial x}=\frac{\partial f}{\partial z} \frac{\partial z}{\partial x}
$$

## Jacobian vs gradient - Example

- Consider differentiable $f: \mathbb{R}^{m} \rightarrow \mathbb{R}$ and $M \in \mathbb{R}^{m \times n}$
- Compute Jacobian of $g=(f \circ M)$ using chain rule:
- Rewrite as $g(x)=f(z)$ where $z=M x$
- Compute Jacobian by partial Jacobians $\frac{\partial f}{\partial z}$ and $\frac{\partial z}{\partial x}$ :

$$
\frac{\partial g}{\partial x}=\frac{\partial g}{\partial z} \frac{\partial z}{\partial x}=\frac{\partial f}{\partial z} \frac{\partial z}{\partial x}=\nabla f(z)^{T} M=\nabla f(M x)^{T} M \in \mathbb{R}^{1 \times n}
$$

- Know gradient of $(f \circ M)(x)$ satisfies

$$
\nabla(f \circ M)(x)=M^{T} \nabla f(M x) \in \mathbb{R}^{n}
$$

which is transpose of Jacobian

## Backpropagation - Introduce states

- Compute gradient/Jacobian of

$$
L\left(m\left(x_{i} ; \theta\right), y_{i}\right)
$$

w.r.t. $\theta=\left\{\left(W_{j}, b_{j}\right)\right\}_{j=1}^{n}$, where
$m\left(x_{i} ; \theta\right)=W_{n} \sigma_{n-1}\left(W_{n-1} \sigma_{n-2}\left(\cdots\left(W_{2} \sigma_{1}\left(W_{1} x_{i}+b_{1}\right)+b_{2}\right) \cdots\right)+b_{n-1}\right)+b_{n}$

- Rewrite as function with states $z_{j}$

$$
\begin{aligned}
& L\left(z_{n+1}, y_{i}\right) \\
\text { where } & z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right) \text { for } j \in\{1, \ldots, n\} \\
\text { and } & z_{1}=x_{i}
\end{aligned}
$$

where $\sigma_{n}(u) \equiv u$

## Graphical representation

- Per sample loss function

$$
\begin{aligned}
& L\left(z_{n+1}, y_{i}\right) \\
\text { where } & z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right) \text { for } j \in\{1, \ldots, n\} \\
\text { and } & z_{1}=x_{i}
\end{aligned}
$$

where $\sigma_{n}(u) \equiv u$

- Graphical representation



## Backpropagation - Chain rule

- Jacobian of $L$ w.r.t. $W_{j}$ and $b_{j}$ can be computed as

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}} \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial W_{j}} \\
\frac{\partial L}{\partial b_{j}} & =\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}} \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial b_{j}}
\end{aligned}
$$

where we mean derivative w.r.t. first argument in $L$

- Backpropagation evaluates partial Jacobians as follows

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\left(\left(\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}}\right) \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}}\right) \frac{\partial z_{j+1}}{\partial W_{j}} \\
\frac{\partial L}{\partial b_{j}} & =\left(\left(\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}}\right) \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}}\right) \frac{\partial z_{j+1}}{\partial b_{j}}
\end{aligned}
$$

## Backpropagation - Forward and backward pass

- Jacobian of $L\left(z_{n+1}, y_{i}\right)$ w.r.t. $z_{n+1}$ (transpose of gradient)
- Computing Jacobian of $L\left(z_{n+1}, y_{i}\right)$ requires $z_{n+1}$ $\Rightarrow$ forward pass: $z_{1}=x_{i}, z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)$
- Backward pass, store $\delta_{j}$ :

$$
\frac{\partial L}{\partial z_{j+1}}=\underbrace{\underbrace{(\underbrace{\left.\left.\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}}\right) \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}}\right)}_{\delta_{n+1}^{T}}}_{\delta_{n}^{T}})}_{\delta_{j+1}^{T}}
$$

- Compute

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial W_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial W_{j}} \\
\frac{\partial L}{\partial b_{j}} & =\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial b_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial b_{j}}
\end{aligned}
$$

## Dimensions

- Let $z_{j} \in \mathbb{R}^{n_{j}}$, consequently $W_{j} \in \mathbb{R}^{n_{j+1} \times n_{j}}, b_{j} \in \mathbb{R}^{n_{j+1}}$
- Dimensions

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\underbrace{\frac{\partial L}{\partial b_{j}}}_{\underbrace{(\underbrace{\frac{\partial L}{\partial z_{n+1}}}_{1 \times n_{n+1}} \underbrace{\frac{\partial z_{n+1}}{\partial z_{n}}}_{n_{n+1} \times n_{n}}) \cdots \underbrace{\frac{\partial z_{j+2}}{\partial z_{j+1}}}_{n_{j+2} \times n_{j+1}})}_{1 \times n_{j+1}} \underbrace{\frac{\partial z_{j+1}}{\partial W_{j}}}_{n_{j+1} \times n_{j+1} \times n_{j}}}=\underbrace{\left(\left(\frac{\partial L}{\partial z_{n+1}} \frac{\partial z_{n+1}}{\partial z_{n}}\right) \cdots \frac{\partial z_{j+2}}{\partial z_{j+1}}\right)}_{1 \times n_{j+1}} \underbrace{\frac{\partial z_{j+1}}{\partial b_{j}}}_{n_{j+1} \times n_{j+1}}
\end{aligned}
$$

- Vector matrix multiplies except for in last step
- Multiplication with tensor $\frac{\partial z_{j+1}}{\partial W_{j}}$ can be simplified
- Backpropagation variables $\delta_{j} \in \mathbb{R}^{n_{j}}$ are vectors (not matrices)


## Partial Jacobian $\frac{\partial z_{j+1}}{\partial z_{j}}$

- Recall relation $z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)$ and let $v_{j}=W_{j} z_{j}+b_{j}$
- Chain rule gives

$$
\begin{aligned}
\frac{\partial z_{j+1}}{\partial z_{j}} & =\frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial z_{j}}=\operatorname{diag}\left(\sigma_{j}^{\prime}\left(v_{j}\right)\right) \frac{\partial v_{j}}{\partial z_{j}} \\
& =\operatorname{diag}\left(\sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) W_{j}
\end{aligned}
$$

where, with abuse of notation (notation overloading)

$$
\sigma_{j}^{\prime}(u)=\left[\begin{array}{c}
\sigma_{j}^{\prime}\left(u_{1}\right) \\
\vdots \\
\sigma_{j}^{\prime}\left(u_{n_{j+1}}\right)
\end{array}\right]
$$

- Reason: $\sigma_{j}(u)=\left[\sigma_{j}\left(u_{1}\right), \ldots, \sigma_{j}\left(u_{n_{j+1}}\right)\right]^{T}$ with $\sigma_{j}: \mathbb{R}^{n_{j+1}} \rightarrow \mathbb{R}^{n_{j+1}}$, gives

$$
\frac{d \sigma_{j}}{d u}=\left[\begin{array}{ccc}
\sigma_{j}^{\prime}\left(u_{1}\right) & & \\
& \ddots & \\
& & \sigma_{j}^{\prime}\left(u_{n_{j+1}}\right)
\end{array}\right]=\operatorname{diag}\left(\sigma_{j}^{\prime}(u)\right)
$$

## Partial Jacobian $\delta_{j}^{T}=\frac{\partial L}{\partial z_{j}}$

- For any vector $\delta_{j+1} \in \mathbb{R}^{n_{j+1} \times 1}$, we have

$$
\begin{aligned}
\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial z_{j}} & =\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) W_{j} \\
& =\left(W_{j}^{T}\left(\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right)^{T}\right)^{T} \\
& =\left(W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right)^{T}
\end{aligned}
$$

where $\odot$ is element-wise (Hadamard) product

- We have defined $\delta_{n+1}^{T}=\frac{\partial L}{\partial z_{n+1}}$, then

$$
\delta_{n}^{T}=\frac{\partial L}{\partial z_{n}}=\delta_{n+1}^{T} \frac{\partial z_{n+1}}{\partial z_{n}}=(\underbrace{W_{n}^{T}\left(\delta_{n+1} \odot \sigma_{n}^{\prime}\left(W_{n} z_{n}+b_{n}\right)\right)}_{\delta_{n}})^{T}
$$

- Consequently, using induction:

$$
\delta_{j}^{T}=\frac{\partial L}{\partial z_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial z_{j}}=(\underbrace{W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)}_{\delta_{j}})^{T}
$$

## Information needed to compute $\frac{\partial L}{\partial z_{j}}$

- To compute first Jacobian $\frac{\partial L}{\partial z_{n}}$, we need $z_{n} \Rightarrow$ forward pass
- Computing

$$
\frac{\partial L}{\partial z_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial z_{j}}=\left(W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right)^{T}=\delta_{j}^{T}
$$

is done using a backward pass

$$
\delta_{j}=W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)
$$

- All $z_{j}$ (or $v_{j}=W_{j} z_{j}+b_{j}$ ) need to be stored for backward pass



## Partial Jacobian $\frac{\partial L}{\partial W_{j}}$

- Computed by

$$
\frac{\partial L}{\partial W_{j}}=\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial W_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial W_{j}}
$$

where $z_{j+1}=\sigma_{j}\left(v_{j}\right)$ and $v_{j}=W_{j} z_{j}+b_{j}$

- Recall $\frac{\partial z_{j+1}}{\partial W_{l}}$ is 3D tensor, compute Jacobian w.r.t. row $l\left(W_{j}\right)_{l}$

$$
\begin{aligned}
\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial\left(W_{j}\right)_{l}} & =\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial\left(W_{j}\right)_{l}}=\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(v_{j}\right)\right)\left[\begin{array}{c}
0 \\
\vdots \\
z_{j}^{T} \\
\vdots \\
0
\end{array}\right] \\
& =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)^{T}\left[\begin{array}{c}
0 \\
\vdots \\
z_{j}^{T} \\
\vdots \\
0
\end{array}\right]=\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)_{l} z_{j}^{T}
\end{aligned}
$$

## Partial Jacobian $\frac{\partial L}{\partial W_{j}}$ cont'd

- Stack Jacobians w.r.t. rows to get full Jacobian:

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial W_{j}}=\left[\begin{array}{c}
\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial\left(W_{j}\right)_{1}} \\
\vdots \\
\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial\left(W_{j}\right)_{n_{j+1}}}
\end{array}\right]=\left[\begin{array}{c}
\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)_{1} z_{j}^{T} \\
\vdots \\
\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)_{n_{j+1}} z_{j}^{T}
\end{array}\right] \\
& =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) z_{j}^{T}
\end{aligned}
$$

for all $j \in\{1, \ldots, n-1\}$

- Dimension of result is $n_{j+1} \times n_{j}$, which matches $W_{j}$
- This is used to update $W_{j}$ weights in algorithm


## Partial Jacobian $\frac{\partial L}{\partial b_{j}}$

- Recall $z_{j+1}=\sigma_{j}\left(v_{j}\right)$ where $v_{j}=W_{j} z_{j}+b_{j}$
- Computed by

$$
\begin{aligned}
\frac{\partial L}{\partial b_{j}} & =\frac{\partial L}{\partial z_{j+1}} \frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial b_{j}}=\delta_{j+1}^{T} \frac{\partial z_{j+1}}{\partial v_{j}} \frac{\partial v_{j}}{\partial b_{j}}=\delta_{j+1}^{T} \operatorname{diag}\left(\sigma_{j}^{\prime}\left(v_{j}\right)\right) \\
& =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)^{T}
\end{aligned}
$$

## Backpropagation summarized

1. Forward pass: Compute and store $z_{j}$ (or $v_{j}=W_{j} z_{j}+b_{j}$ ):

$$
z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)
$$

where $z_{1}=x_{i}$ and $\sigma_{n}=\mathrm{Id}$
2. Backward pass:

$$
\delta_{j}=W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)
$$

with $\delta_{n+1}=\frac{\partial L}{\partial z_{n+1}}$
3. Weight update Jacobians (used in SGD)

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) z_{j}^{T} \\
\frac{\partial L}{\partial b_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} x_{j}+b_{j}\right)\right)^{T}
\end{aligned}
$$

## Backpropagation - Residual networks

1. Forward pass: Compute and store $z_{j}$ (or $v_{j}=W_{j} z_{j}+b_{j}$ ):

$$
z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)+z_{j}
$$

where $z_{1}=x_{i}$ and $\sigma_{n}=\mathrm{Id}$
2. Backward pass:

$$
\delta_{j}=W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)+\delta_{j+1}
$$

with $\delta_{n+1}=\frac{\partial L}{\partial z_{n+1}}$
3. Weight update Jacobians (used in SGD)

$$
\begin{aligned}
\frac{\partial L}{\partial W_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right) z_{j}^{T} \\
\frac{\partial L}{\partial b_{j}} & =\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} x_{j}+b_{j}\right)\right)^{T}
\end{aligned}
$$

## Outline

- Deep learning
- Learning features
- Model properties and activation functions
- Loss landscape
- Residual networks
- Overparameterized networks
- Generalization and regularization
- Generalization - Norm of weights
- Generalization - Flatness of minima
- Backpropagation
- Vanishing and exploding gradients


## Vanishing and exploding gradient problem

- For some activation functions, gradients can vanish
- For other activation functions, gradients can explode


## Vanishing gradient example: Sigmoid

- Assume $\left\|W_{j}\right\| \leq 1$ for all $j$ and $\left\|\delta_{n+1}\right\| \leq C$
- Maximal derivative of sigmoid $(\sigma)$ is 0.25
- Then

$$
\begin{aligned}
\left\|\frac{\partial L}{\partial z_{j}}\right\| & =\left\|\delta_{j}\right\|=\left\|W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right\| \leq 0.25\left\|\delta_{j+1}\right\| \\
& \leq 0.25^{n-j+1}\left\|\delta_{n+1}\right\| \leq 0.25^{n-j+1} C
\end{aligned}
$$

- Hence, as $n$ grows, gradients can become very small for small $i$
- In general, vanishing gradient if $\sigma^{\prime}<1$ everywhere
- Similar reasoning: exploding gradient if $\sigma^{\prime}>1$ everywhere
- Hence, need $\sigma^{\prime}=1$ in important regions


## Vanishing gradients - Residual networks

- Residual networks with forward pass

$$
z_{j+1}=\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)+z_{j}
$$

and backward pass

$$
\delta_{j}=W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)+\delta_{j+1}
$$

- Gradients do not vanish in passes despite small $\sigma$ gain


## Examples of activation functions

Activation functions that (partly) avoid vanishing gradients


## Exploding gradient - Example

- Assume $L$-Lipschitz activation (ReLU, Tanh etc have $L=1$ )
- Forward pass estimation:

$$
\begin{aligned}
\left\|z_{j+1}\right\|_{2} & =\left\|\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)\right\|_{2} \leq L\left\|W_{j} z_{j}+b_{j}\right\|_{2} \leq L\left(\left\|W_{j} z_{j}\right\|_{2}+\left\|b_{j}\right\|_{2}\right) \\
& \leq L\left\|W_{j}\right\|_{2}\left\|z_{j}\right\|_{2}+L\left\|b_{j}\right\|_{2}
\end{aligned}
$$

- Backward pass estimation:

$$
\begin{aligned}
\left\|\delta_{j}\right\|_{2} & =\left\|W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right\|_{2} \\
& \leq\left\|W_{j}^{T}\right\|_{2}\left\|\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right\|_{2} \\
& \leq L\left\|W_{j}\right\|_{2}\left\|\delta_{j+1}\right\|_{2}
\end{aligned}
$$

- If $L \leq 1,\left\|W_{j}\right\|_{2} \leq 1$ and $\left\|b_{j}\right\|_{2}$ small, gradients do not explode
- ReLU "average" $L=0.5$ reduces "average estimate"
- Tanh reduces "average estimates" more since
- $\sigma_{j}$-outputs are constrained to $(-1,1)$
- "average Lipschitz constant" is smaller


## Exploding gradient - Residual network

- Assume $L$-Lipschitz activation (ReLU, Tanh have $L=1$ )
- Forward pass estimation:

$$
\left\|z_{j+1}\right\|_{2}=\left\|\sigma_{j}\left(W_{j} z_{j}+b_{j}\right)\right\|_{2}+\left\|z_{j}\right\|_{2} \leq\left(1+L\left\|W_{j}\right\|_{2}\right)\left\|z_{j}\right\|_{2}+L\left\|b_{j}\right\|_{2}
$$

- Backward pass estimation:

$$
\begin{aligned}
\left\|\delta_{j}\right\|_{2} & =\left\|W_{j}^{T}\left(\delta_{j+1} \odot \sigma_{j}^{\prime}\left(W_{j} z_{j}+b_{j}\right)\right)\right\|_{2}+\delta_{j+1} \\
& \leq\left(1+L\left\|W_{j}\right\|_{2}\right)\left\|\delta_{j+1}\right\|_{2}
\end{aligned}
$$

- Larger estimates than for non-residual networks
- Activations with $L \leq 1$ to avoid exploding and vanishing gradients:
- $\alpha \times$ ReLU with $\alpha \in(0,1)$
- $\alpha \times$ Tanh with $\alpha \in(0,1)$


[^0]:    1) Don't compute in practice, just thought experiment to connect generalization to training loss
