Advanced Deep Learning for Computer Vision

The Team

Lecturers

Prof. Dr. Matthias Niessner

Website

https://niessnerlab.org/

Tutors

Bartłomiej Baranowski

Ziya Erkoç

Nicolas von Lützow

Jonathan Schmidt

Visual Computing & AI Group at TUM

Photorealistic AI Avatars

NeRFs / 3D Gaussians

3D Semantics / Reconstruction

https://niessnerlab.org/publications.html
https://twitter.com/MattNiessner

Visual Computing & AI Group at TUM

MeshGPT: 3D Mesh Generation

SceneTex: 3D Scene Texturing

https://niessnerlab.org/publications.html
https://twitter.com/MattNiessner

Basics of DL

AI vs ML vs DL

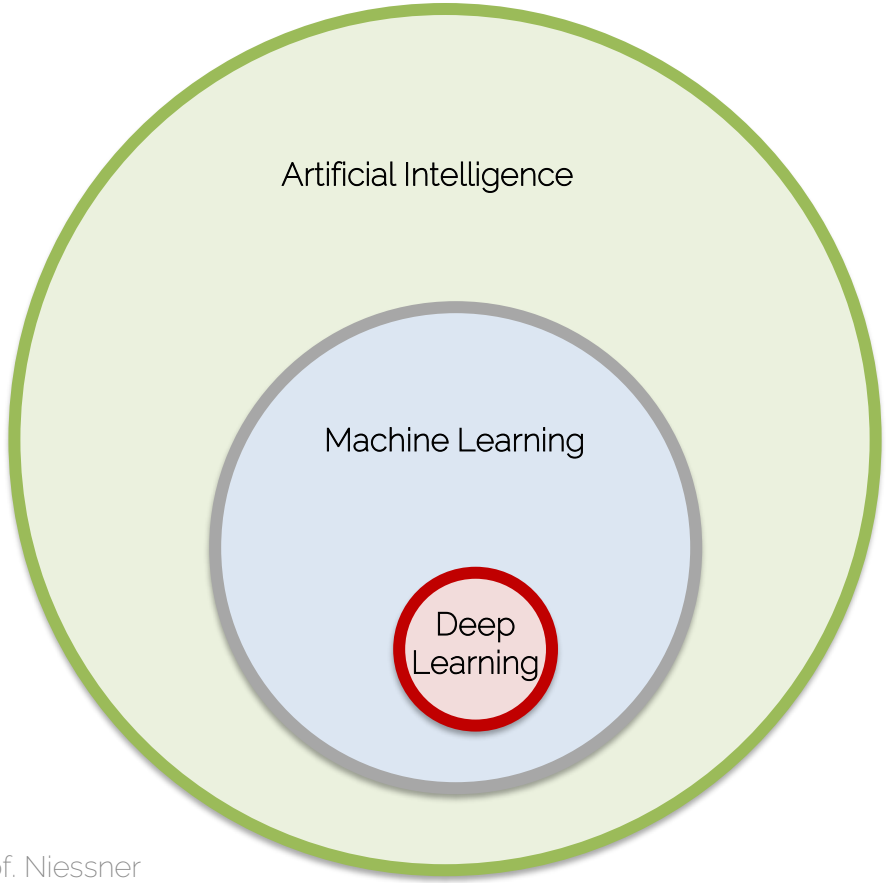

- Deep Learning
  - ML-methods leveraging neural networks
    - Fit non-linear function to training set through optimization
    - "Hope" that we generalize to unseen training samples

What we assume you know

- Linear Algebra & Programming!
- Basics from the Introduction to Deep Learning lecture
  - https://niessner.github.io/I2DL/
- PyTorch (can use TensorFlow…)
- You already trained several models + you know how to debug problems, observe training curves, prepare training/validation/test data

What is a neural network?

What is a neural network?

- Architecture:
  - Layers (MLP vs ConvNet vs Transformer)
  - Activations
- Loss Function
- Optimization
- Debugging ☺

Neural network

- Linear score function $f = Wx$

On CIFAR-10

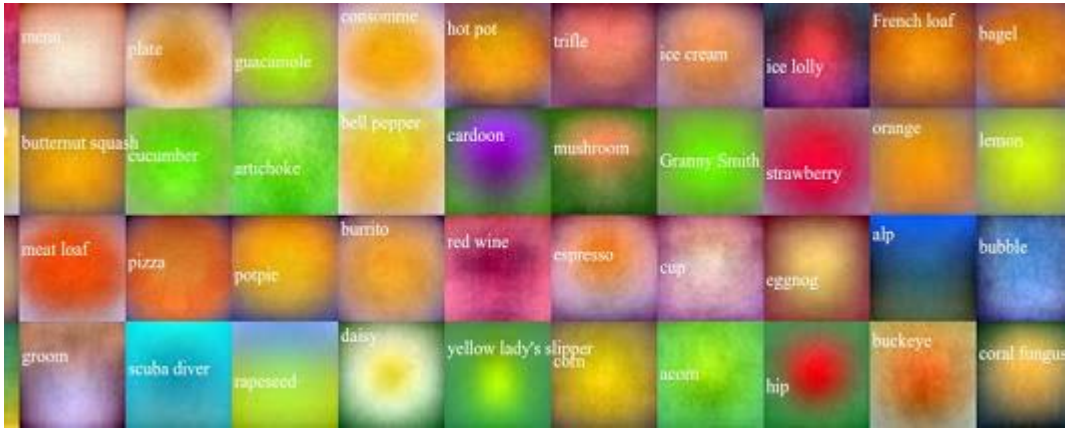

On ImageNet

Credit: Li/Karpathy/Johnson

Neural network

- Linear score function $f = Wx$

- Neural network is a nesting of 'functions'
  - 2-layers: $f = W_2 \max(0, W_1 x)$
  - 3-layers: $f = W_3 \max(0, W_2 \max(0, W_1 x))$
  - 4-layers: $f = W_4 \tanh(\mathrm{W}_3, \max(0, W_2 \max(0, W_1 x)))$
  - 5-layers: $f = W_5 \sigma(W_4 \tanh(\mathrm{W}_3, \max(0, W_2 \max(0, W_1 x))))$
  - … up to hundreds of layers

Neural network

1-layer network: $f = Wx$

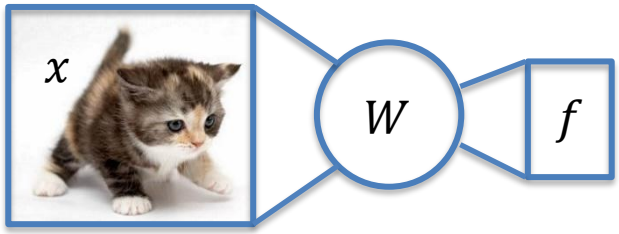

$128 \times 128 = 16384$ $\quad$ 10

2-layer network: $f = W_2 \max(0, W_1 x)$

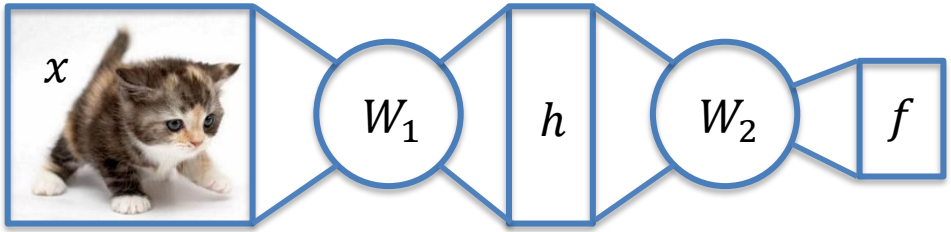

$128 \times 128 = 16384$ $\quad$ 1000 $\quad$ 10

Neural network

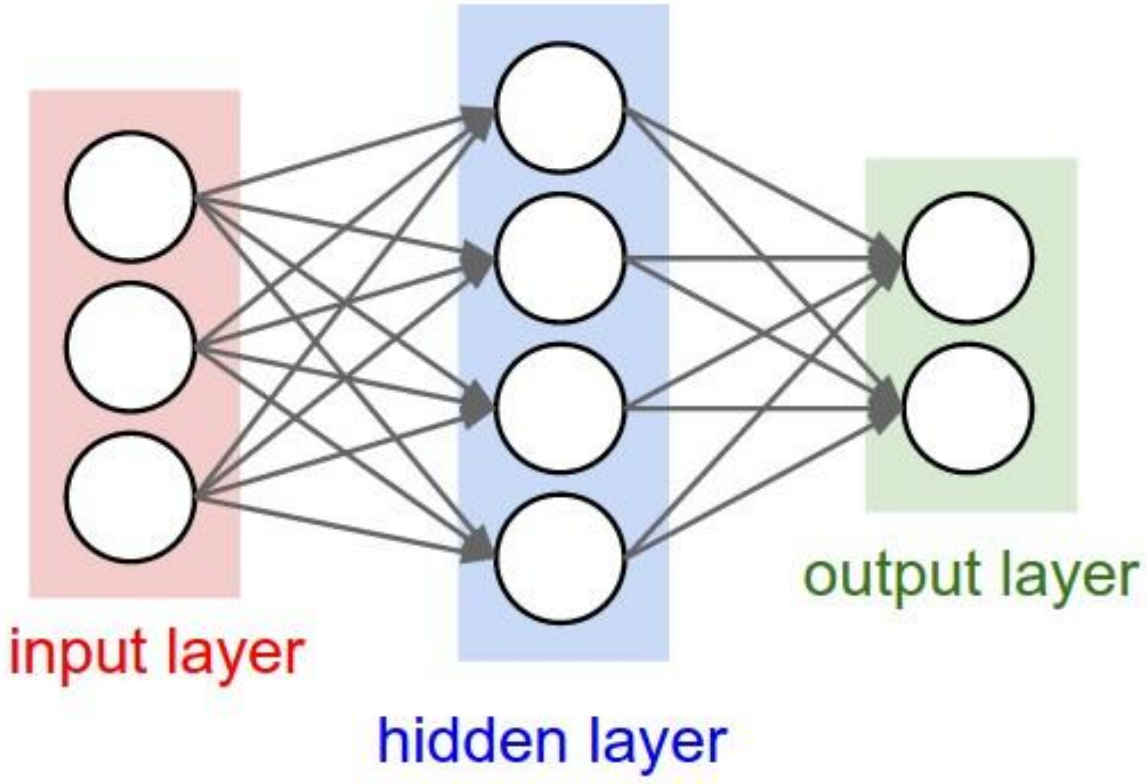

Credit: Li/Karpathy/Johnson

Loss functions

Neural networks

- What is the shape of this function?

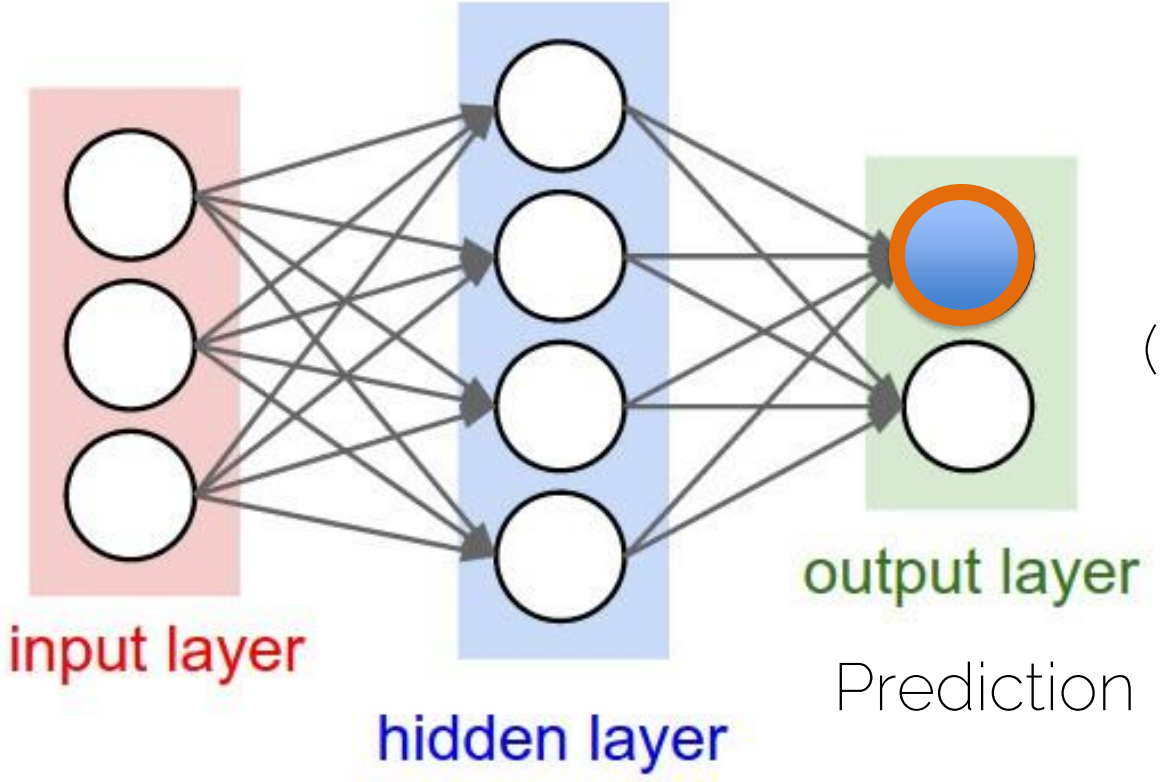

Loss functions

- Softmax loss function

$$L_i = -\log\left(\frac{e^{s_{y_i}}}{\sum_k e^{s_k}}\right)$$

Evaluate the ground truth score for the image

- Hinge Loss (derived from the Multiclass SVM loss)

$$L_i = \sum_{k \neq y_i} \max(0, s_k - s_{y_i} + 1)$$

Loss functions

- Softmax loss function
  - Optimizes until the loss is zero
- Hinge Loss (derived from the Multiclass SVM loss)
  - Saturates whenever it has learned a class "well enough"

Activation functions

Sigmoid

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

Forward

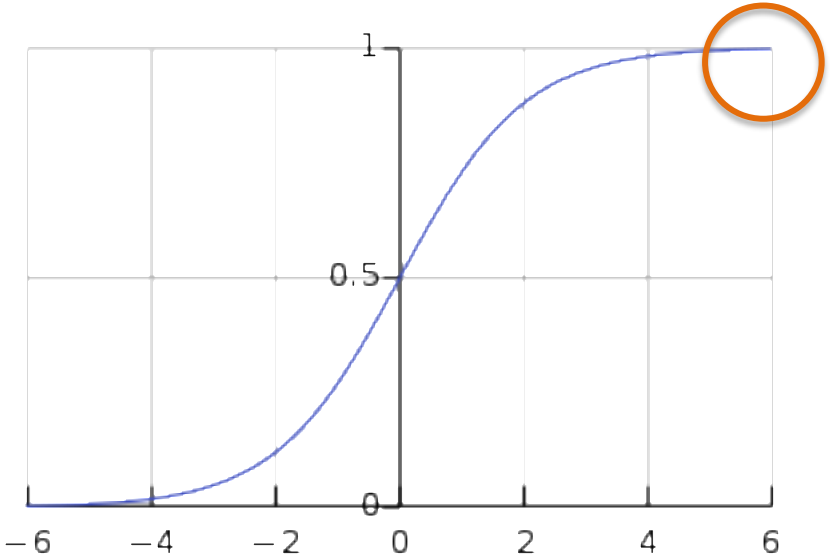

$x = 6$

✘ Saturated neurons kill the gradient flow

$$\frac{\partial L}{\partial x} = \frac{\partial \sigma}{\partial x}\frac{\partial L}{\partial \sigma}$$

$$\frac{\partial \sigma}{\partial x}$$

$$\frac{\partial L}{\partial \sigma}$$

tanh

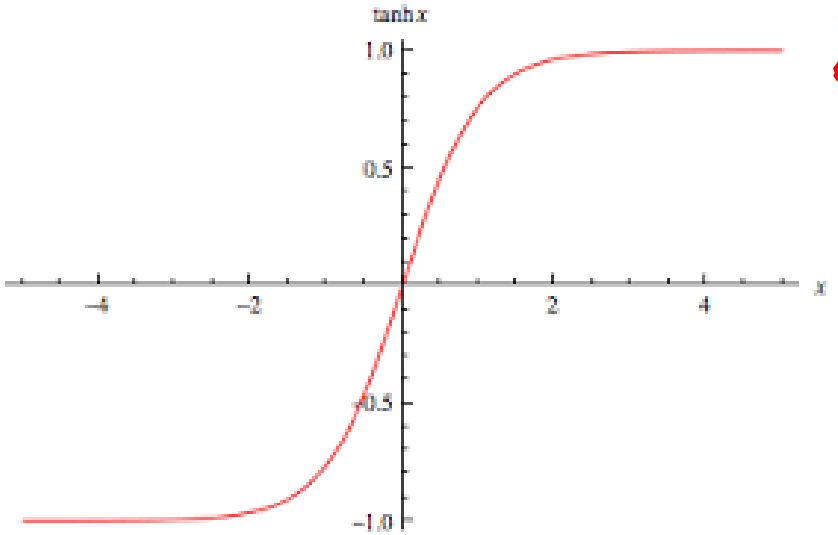

✘ Still saturates

✘ Still saturates

✓ Zero-centered

LeCun 1991

Rectified Linear Units (ReLU)

✗ Dead ReLU

What happens if a ReLU outputs zero?

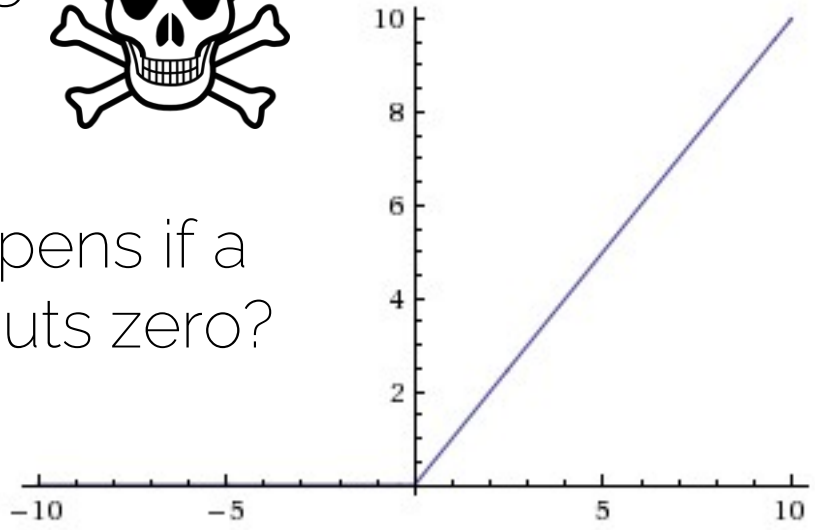

Large and consistent gradients ✓

✓ Fast convergence

✓ Does not saturate

Maxout units

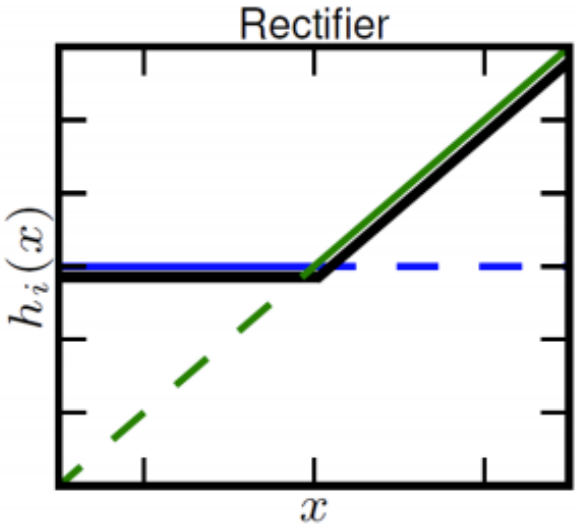

k=2

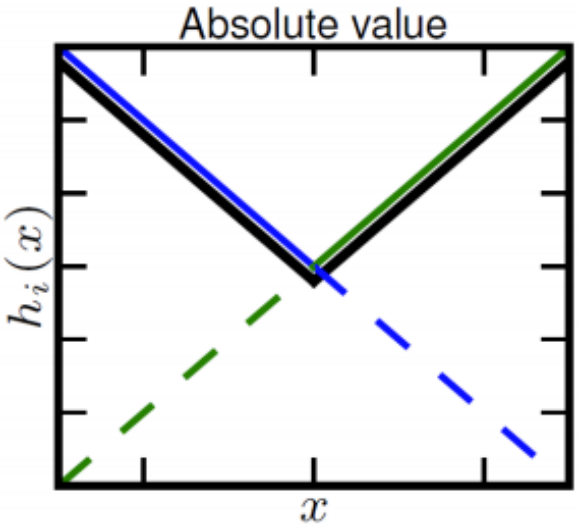

k=2

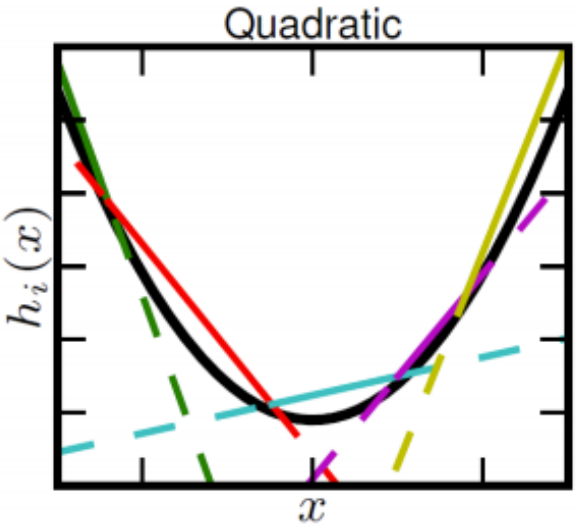

k=5

✓ Generalization of ReLUs

✓ Linear regimes

✓ Does not die

✓ Does not saturate

✗ Increase of the number of parameters

Optimization

Gradient Descent for Neural Networks

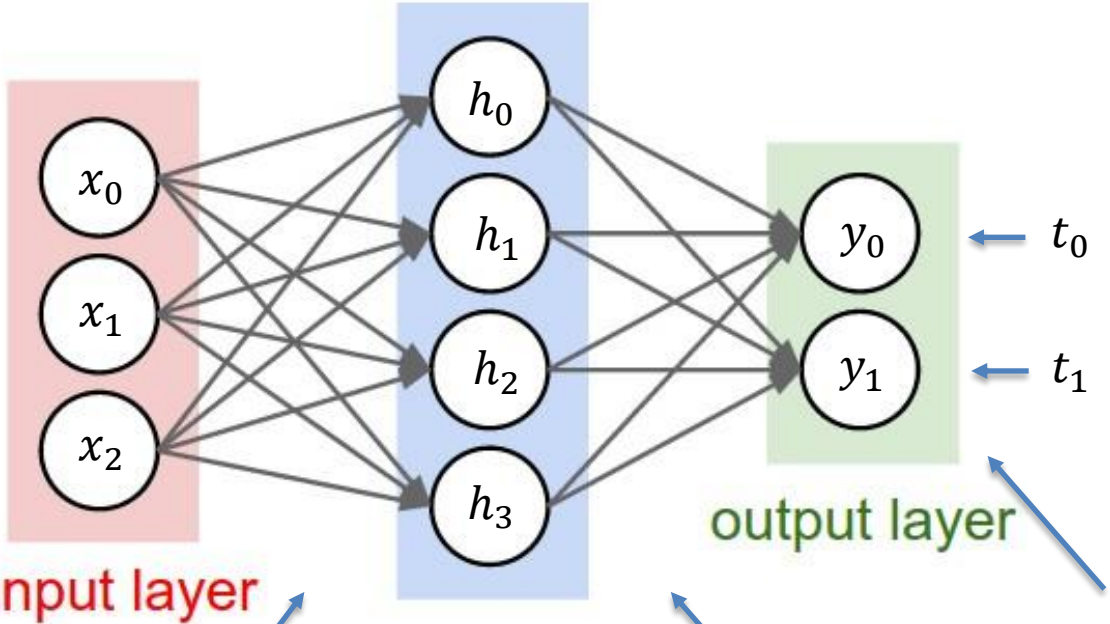

$$\nabla_{w,b} f_{\{x,t\}}(w) = \begin{bmatrix} \frac{\partial f}{\partial w_{0,0,0}} \\ \dots \\ \frac{\partial f}{\partial w_{l,m,n}} \\ \dots \\ \frac{\partial f}{\partial b_{l,m}} \end{bmatrix}$$

$$L_i = (y_i - t_i)^2$$

$$h_j = A\left(b_{0,j} + \sum_k x_k w_{0,j,k}\right)$$

$$y_i = A\left(b_{1,i} + \sum_j h_j w_{1,i,j}\right)$$

Just simple: $A(x) = \max(0, x)$

Stochastic Gradient Descent (SGD)

$$\theta^{k+1} = \theta^k - \alpha \nabla_\theta L(\theta^k, x_{\{1..m\}}, y_{\{1..m\}})$$

$$\nabla_\theta L = \frac{1}{m} \sum_{i=1}^{m} \nabla_\theta L_i$$

k now refers to k-th iteration

m training samples in the current batch

Gradient for the k-th batch

Note the terminology: iteration vs epoch

Gradient Descent with Momentum

$$v^{k+1} = \beta \cdot v^k + \nabla_\theta L(\theta^k)$$

accumulation rate ('friction', momentum)

velocity

Gradient of current minibatch

$$\theta^{k+1} = \theta^k - \alpha \cdot v^{k+1}$$

model

learning rate

velocity

Exponentially-weighted average of gradient

Important: velocity v^k is vector-valued!

Gradient Descent with Momentum

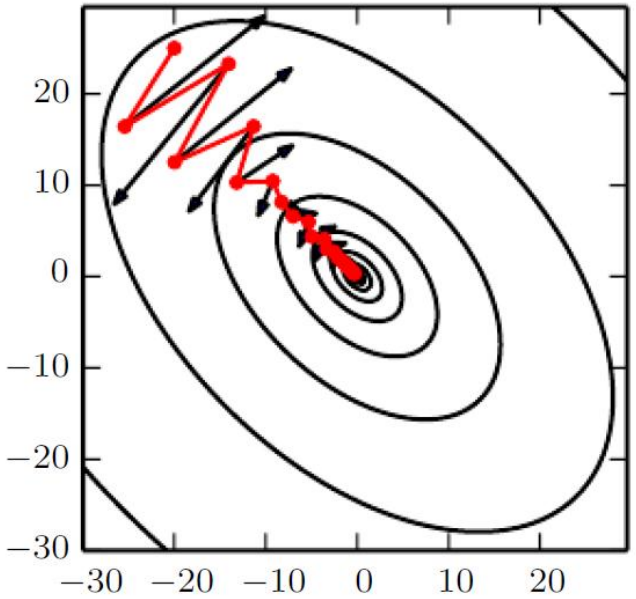

Step will be largest when a sequence of gradients all point to the same direction

Hyperparameters are α, β
β is often set to 0.9

$$\theta^{k+1} = \theta^{k} - \alpha \cdot v^{k+1}$$

Fig. credit: I. Goodfellow

RMSProp

$$s^{k+1} = \beta \cdot s^k + (1-\beta)[\nabla_\theta L \circ \nabla_\theta L]$$

Element-wise multiplication

$$\theta^{k+1} = \theta^k - \alpha \cdot \frac{\nabla_\theta L}{\sqrt{s^{k+1}} + \epsilon}$$

Hyperparameters: α, β, ϵ

Needs tuning!

Often 0.9

Typically 10^{-8}

RMSProp

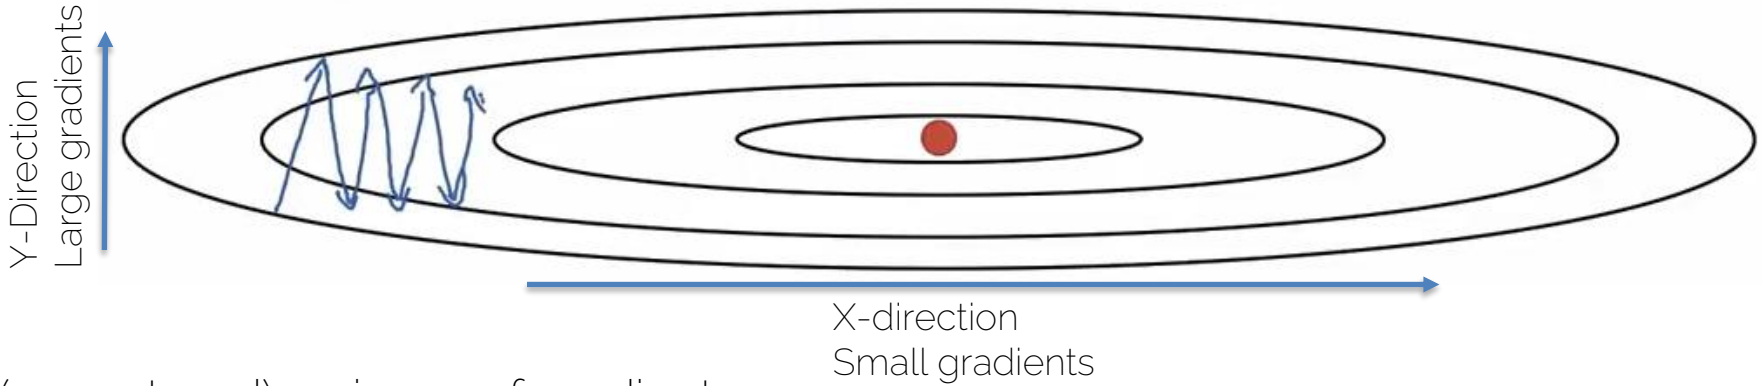

(uncentered) variance of gradients
-> second momentum

$$s^{k+1} = \beta \cdot s^k + (1-\beta)[\nabla_\theta L \circ \nabla_\theta L]$$

$$\theta^{k+1} = \theta^k - \alpha \cdot \frac{\nabla_\theta L}{\sqrt{s^{k+1}} + \epsilon}$$

We're dividing by square gradients:
- Division in Y-Direction will be large
- Division in X-Direction will be small

Can increase learning rate!

Adaptive Moment Estimation (Adam)

Combines Momentum and RMSProp

$$m^{k+1} = \beta_1 \cdot m^k + (1 - \beta_1)\nabla_\theta L(\theta^k)$$

First momentum: mean of gradients

$$v^{k+1} = \beta_2 \cdot v^k + (1 - \beta_2)[\nabla_\theta L(\theta^k) \circ \nabla_\theta L(\theta^k)]$$

Second momentum: variance of gradients

$$\theta^{k+1} = \theta^k - \alpha \cdot \frac{m^{k+1}}{\sqrt{v^{k+1}} + \epsilon}$$

Adam

Combines Momentum and RMSProp

$$m^{k+1} = \beta_1 \cdot m^k + (1-\beta_1)\nabla_\theta L(\theta^k)$$

$$v^{k+1} = \beta_2 \cdot v^k + (1-\beta_2)[\nabla_\theta L(\theta^k) \circ \nabla_\theta L(\theta^k)]$$

m^{k+1} and v^{k+1} are initialized with zero
-> bias towards zero

Typically, bias-corrected moment updates

$$\hat{m}^{k+1} = \frac{m^k}{1-\beta_1}$$

$$\hat{v}^{k+1} = \frac{v^k}{1-\beta_2}$$

$$\theta^{k+1} = \theta^k - \alpha \cdot \frac{\hat{m}^{k+1}}{\sqrt{\hat{v}^{k+1}}+\epsilon}$$

Convergence

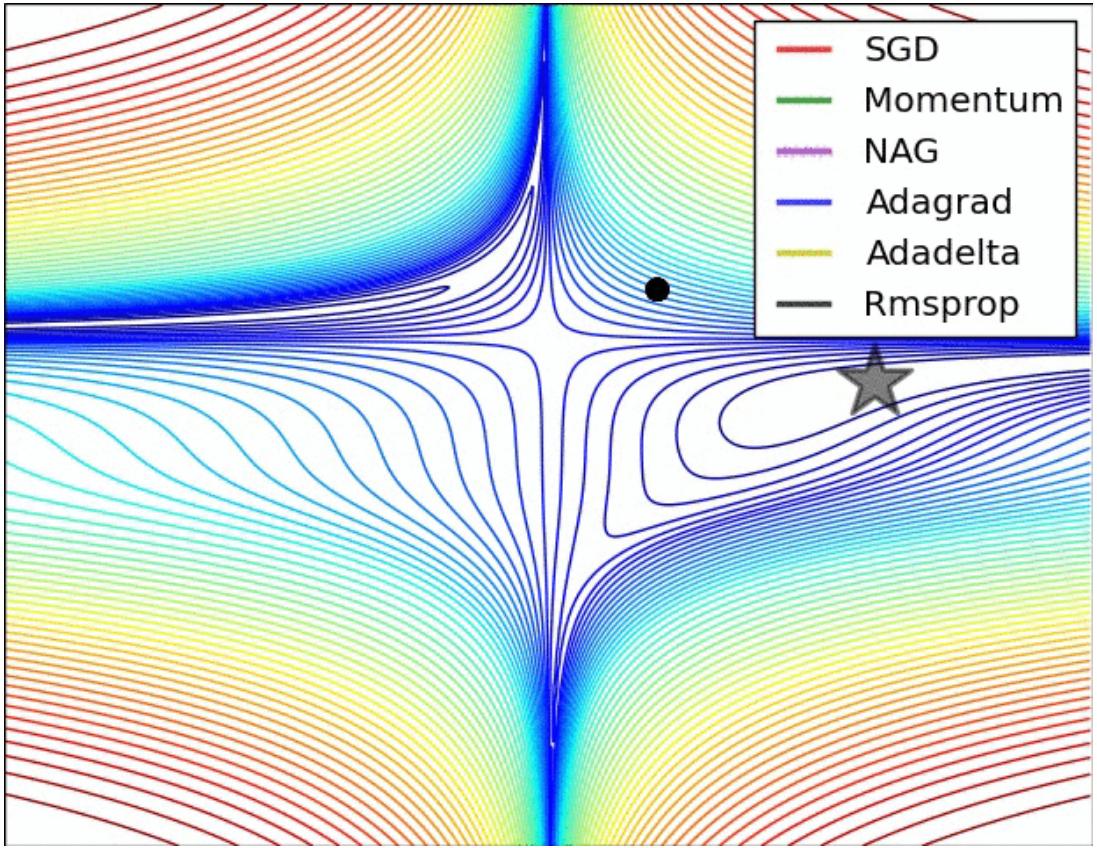

Training NNs

Importance of Learning Rate

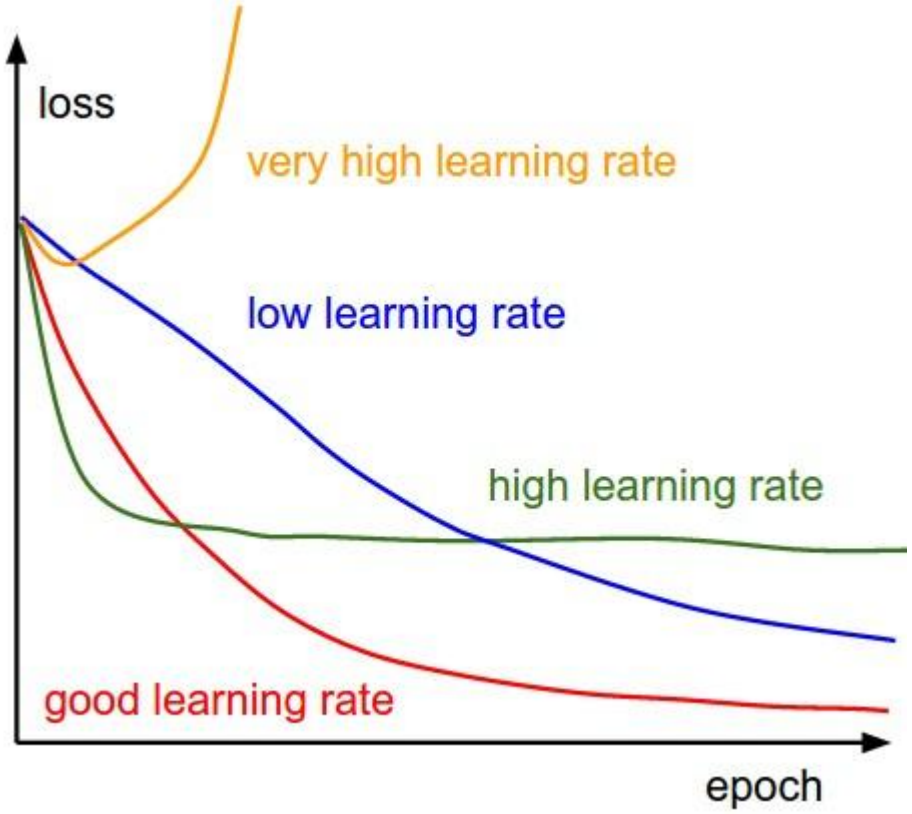

Over- and Underfitting

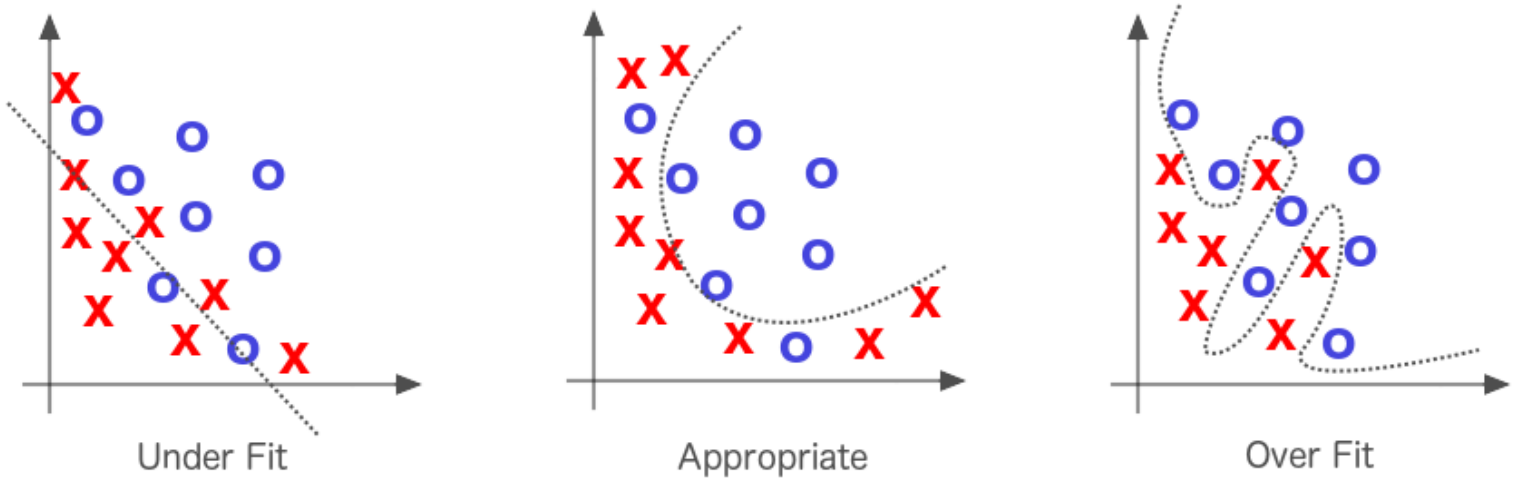

Figure extracted from Deep Learning by Adam Gibson, Josh Patterson, O'Reily Media Inc., 2017

Over- and Underfitting

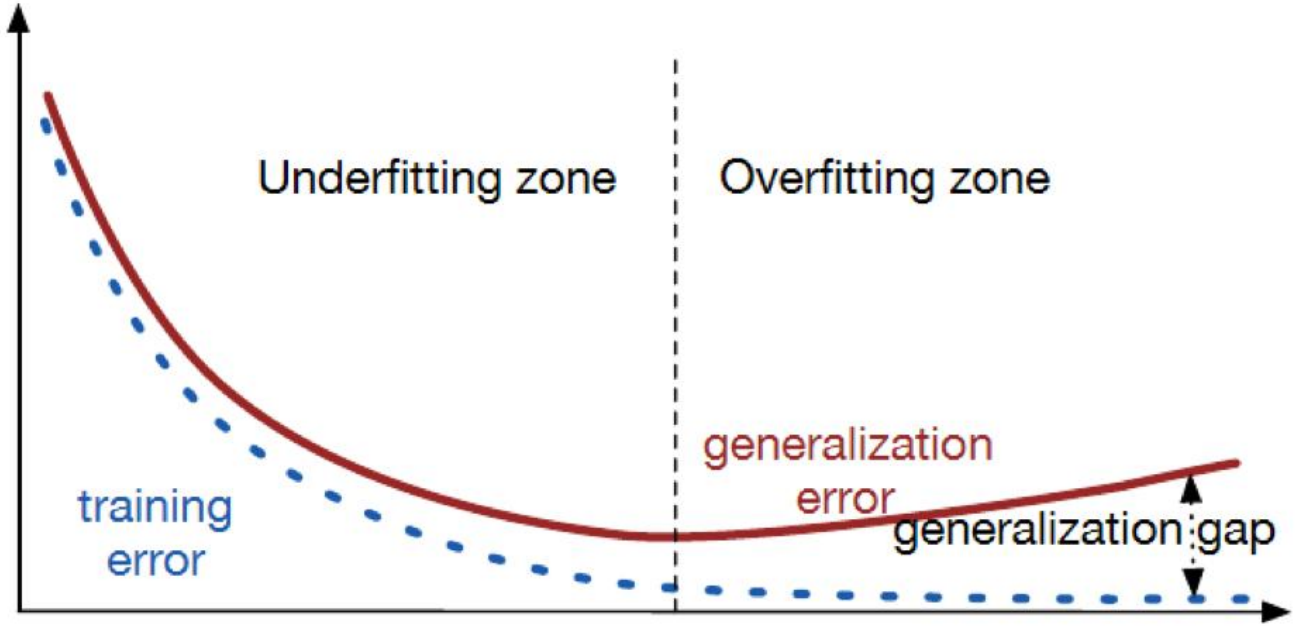

Source: http://srdas.github.io/DLBook/ImprovingModelGeneralization.html

Basic recipe for machine learning

- Split your data

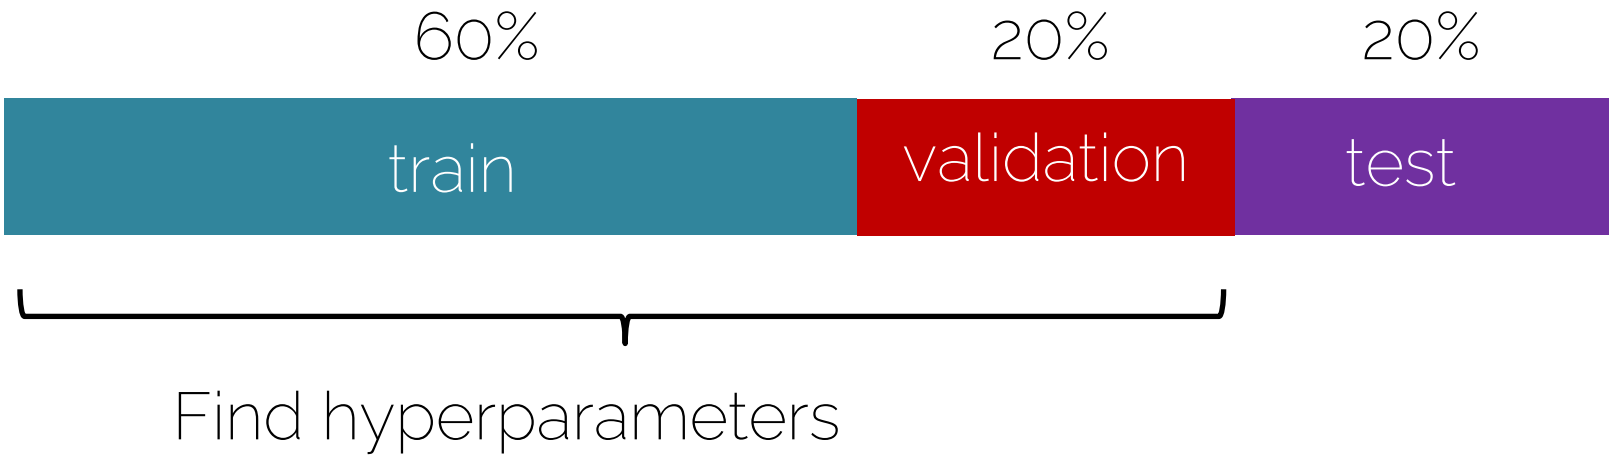

Basic recipe for machine learning

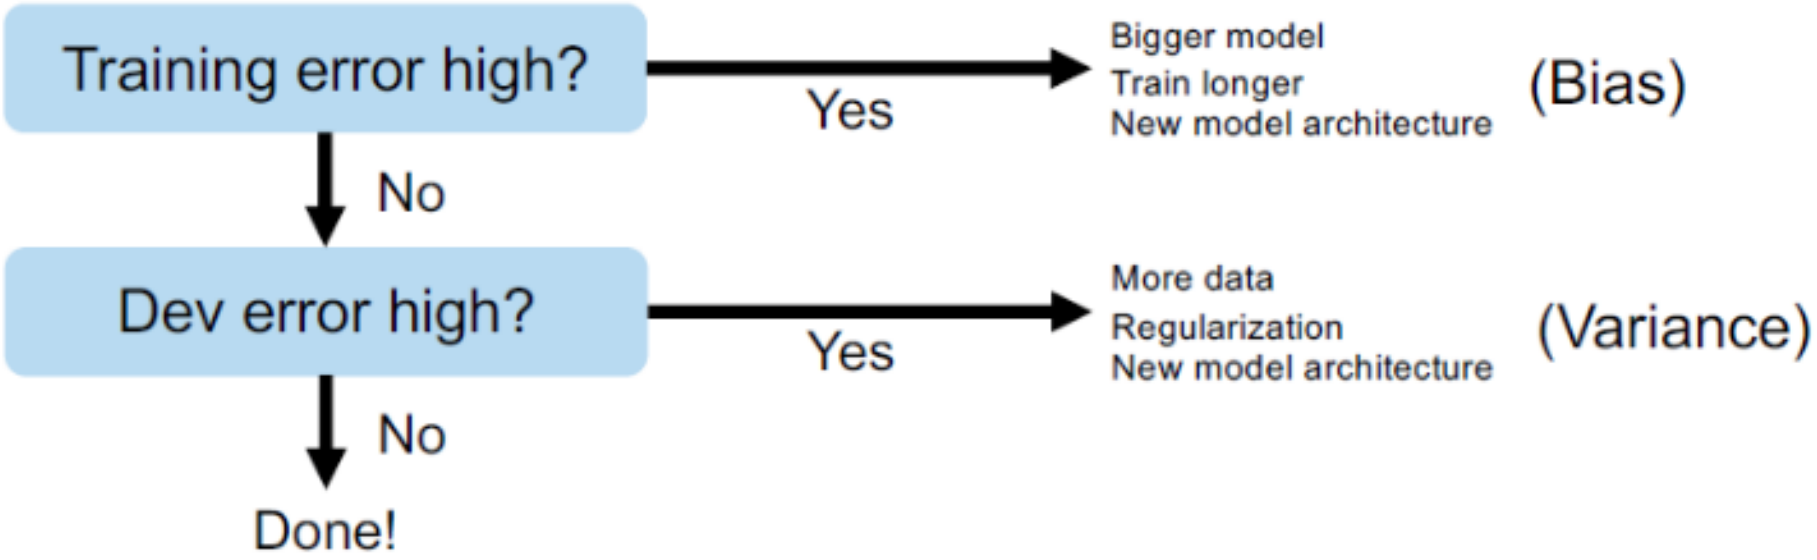

How to deal with images?

Using CNNs in Computer Vision

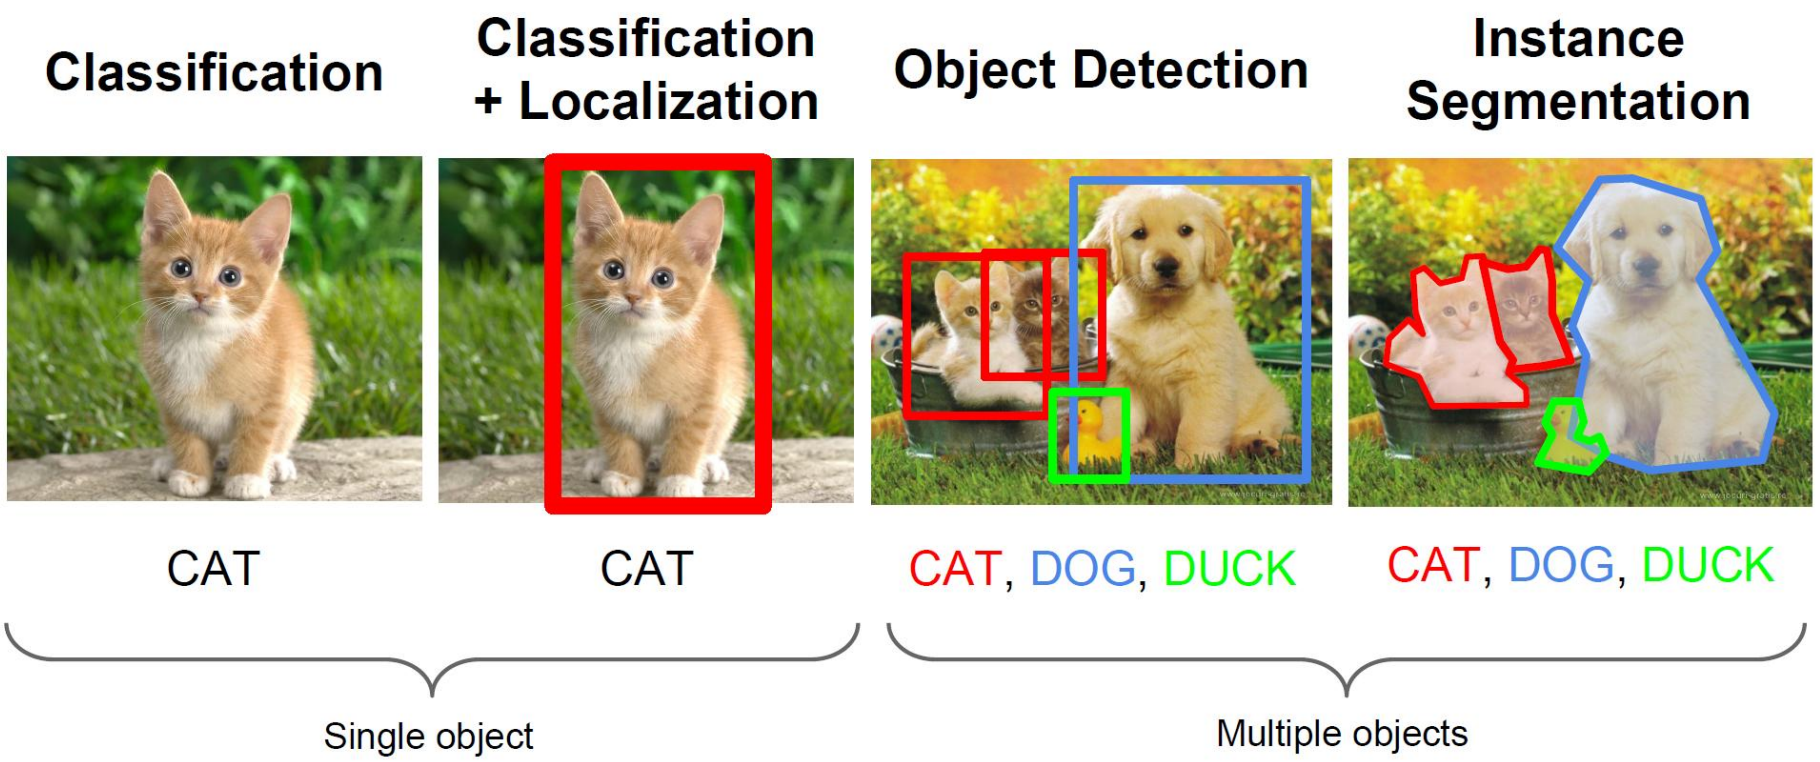

Credit: Li/Karpathy/Johnson

Image filters

- Each kernel gives us a different image filter

Input

Edge detection

$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Box mean

$$\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

LET'S LEARN THESE FILTERS!

Sharpen

$$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$$

Gaussian blur

$$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

Convolutions on RGB Images

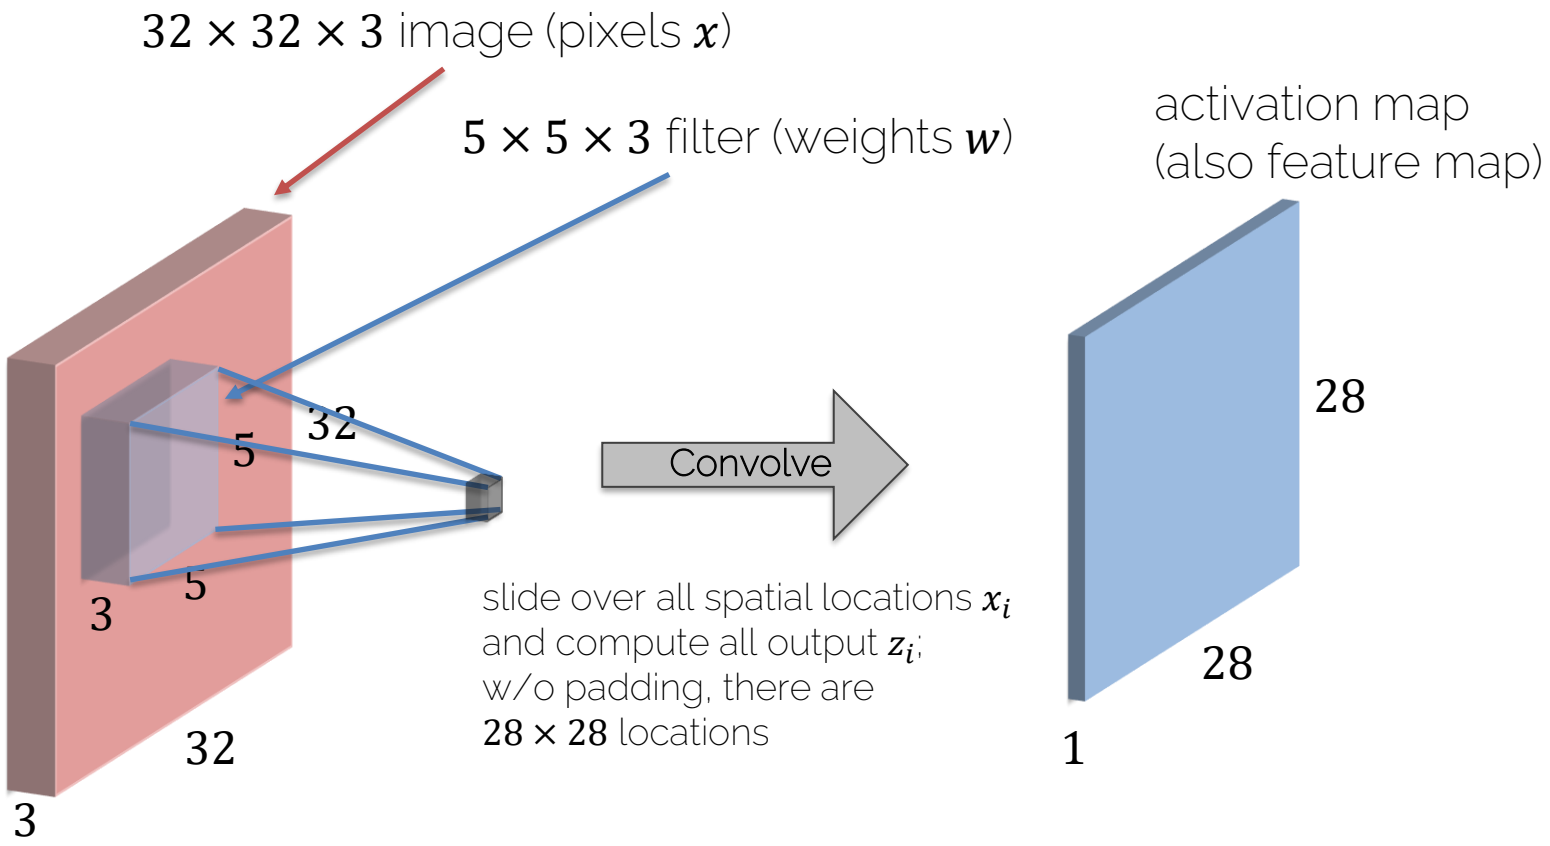

Convolution Layer

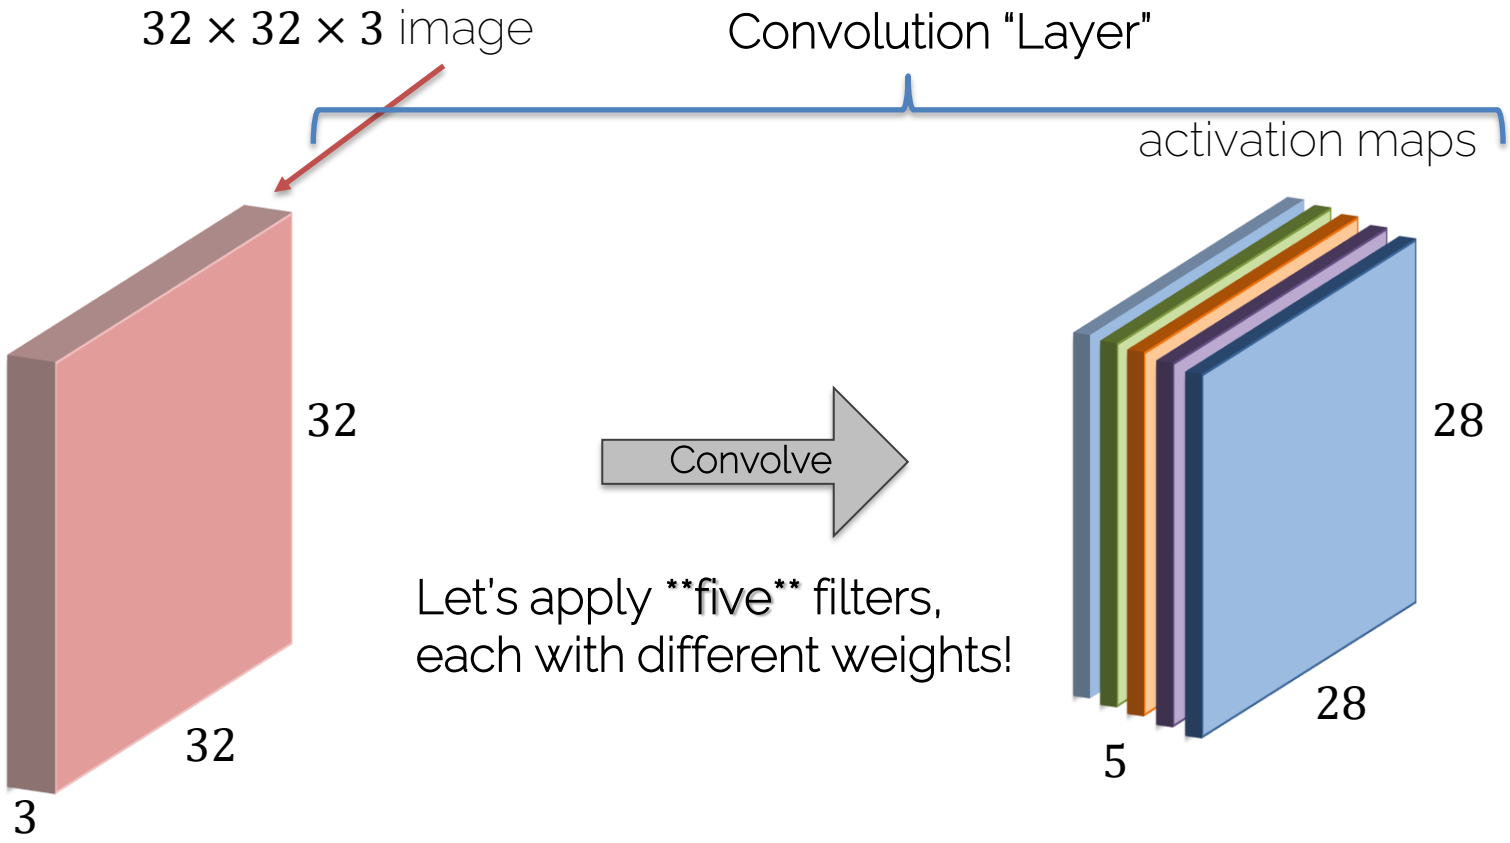

CNN Prototype

ConvNet is concatenation of Conv Layers and activations

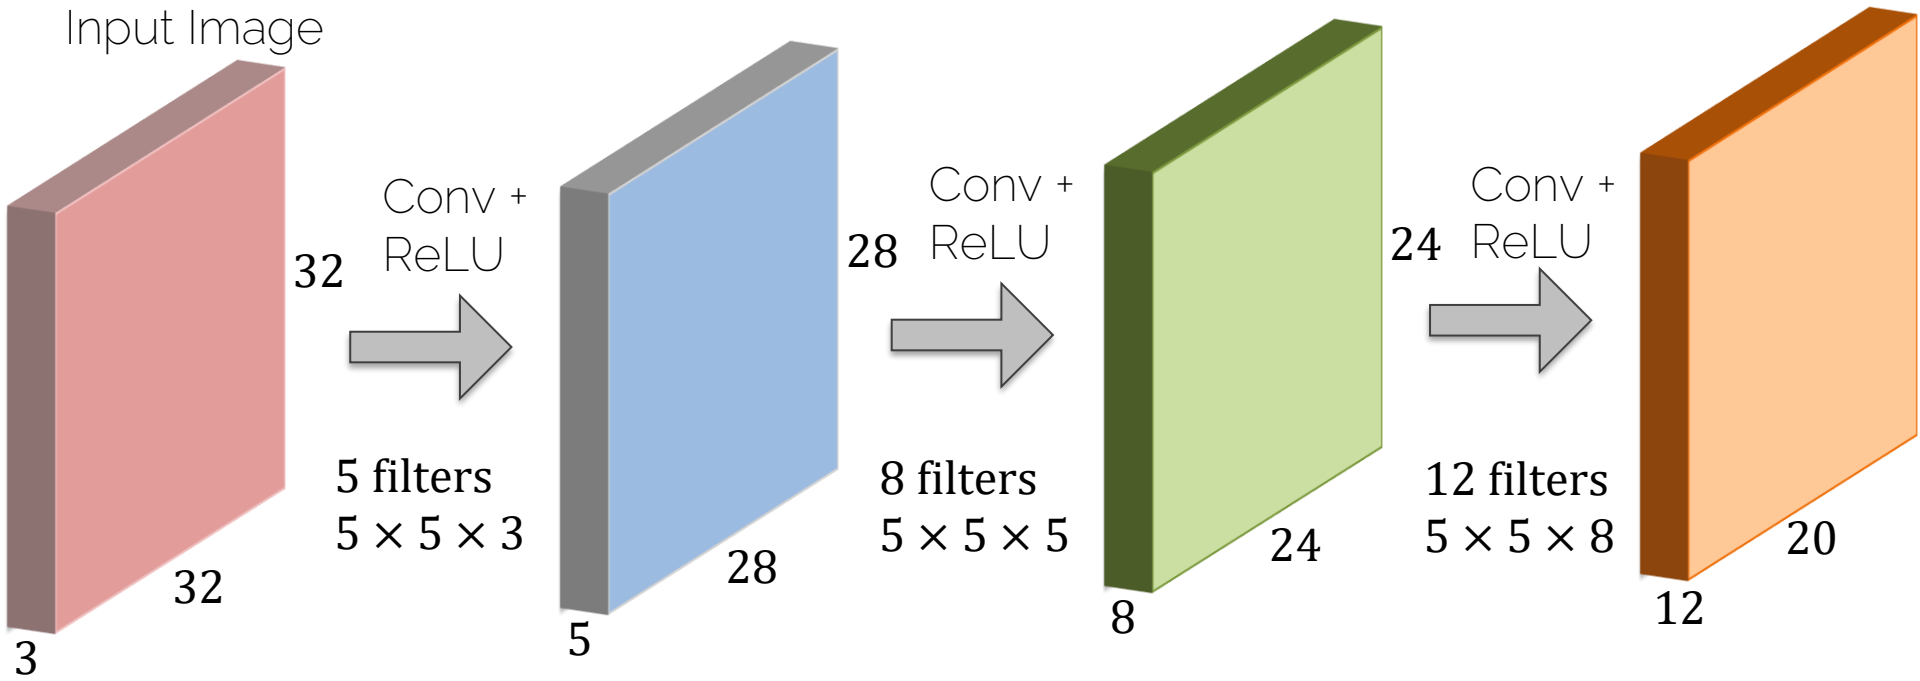

CNN learned filters

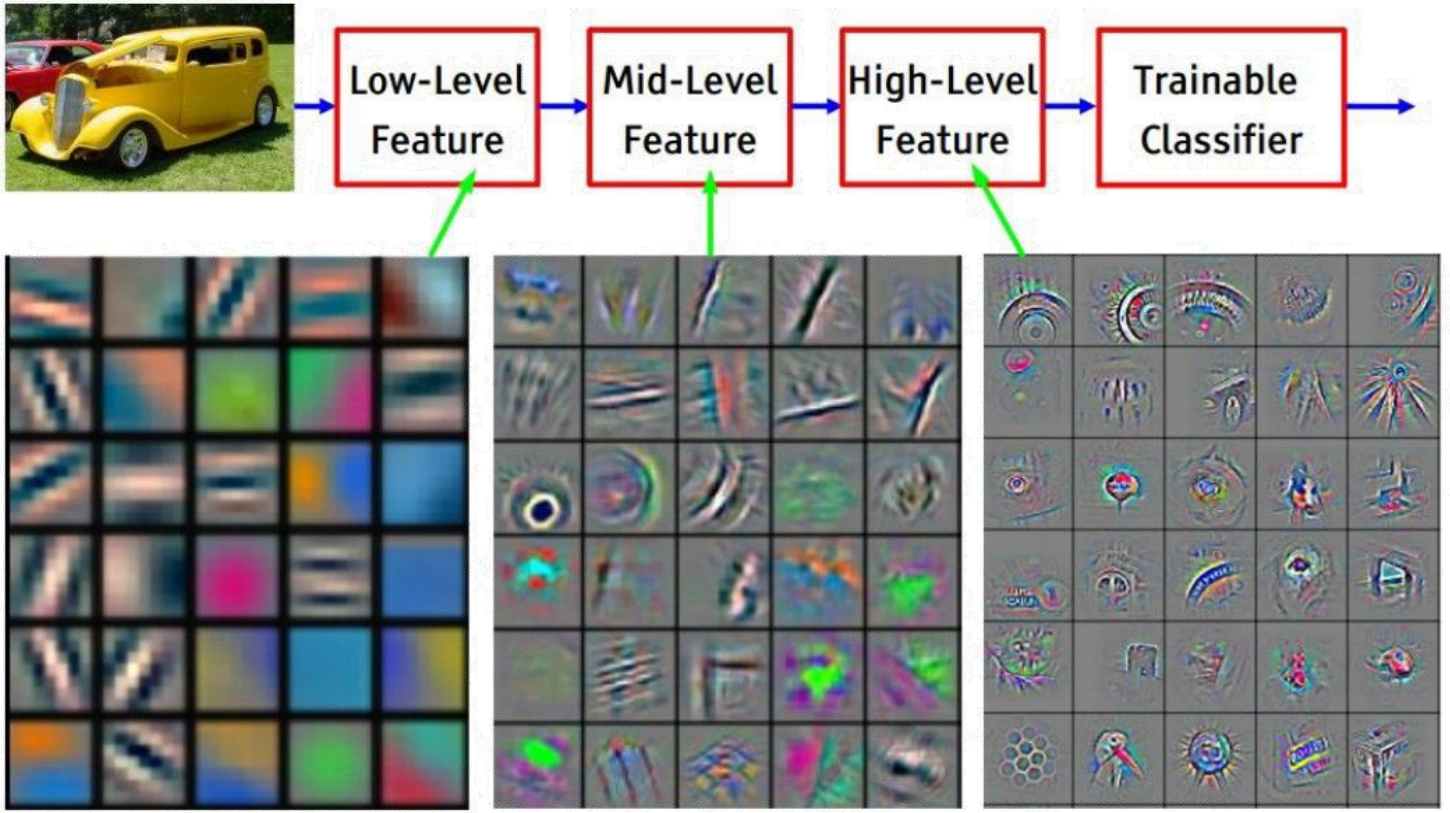

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Pooling Layer: Max Pooling

Single depth slice of input

3	1	3	5
6	0	7	9
3	2	1	4
0	2	4	3

Max pool with
2 × 2 filters and stride 2

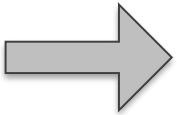

'Pooled' output

6	9
3	4

CNN Architectures

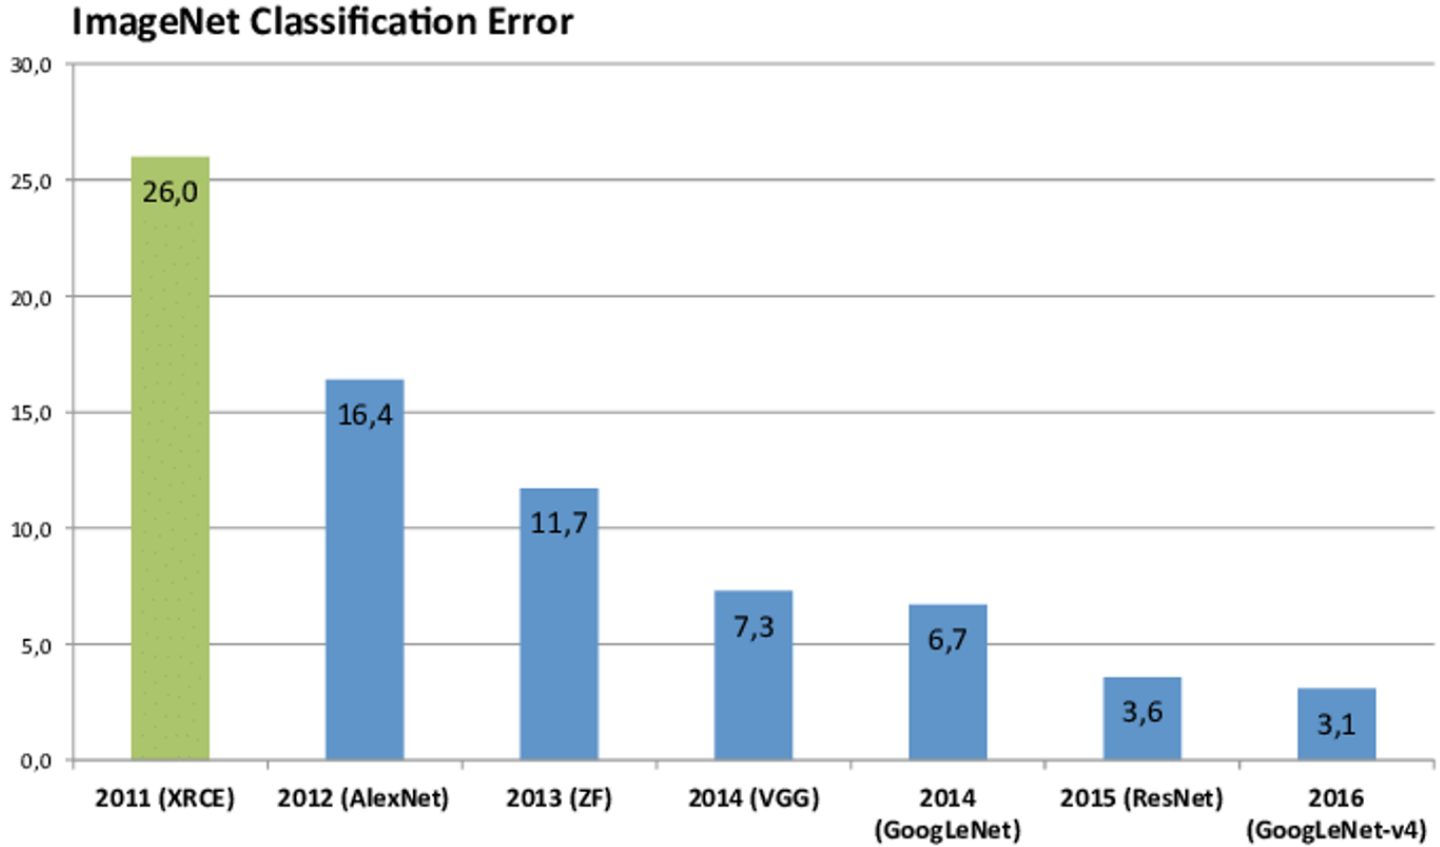

Transformers

Transformers

Fully connected layer

Multi-Head Attention on the "encoder"

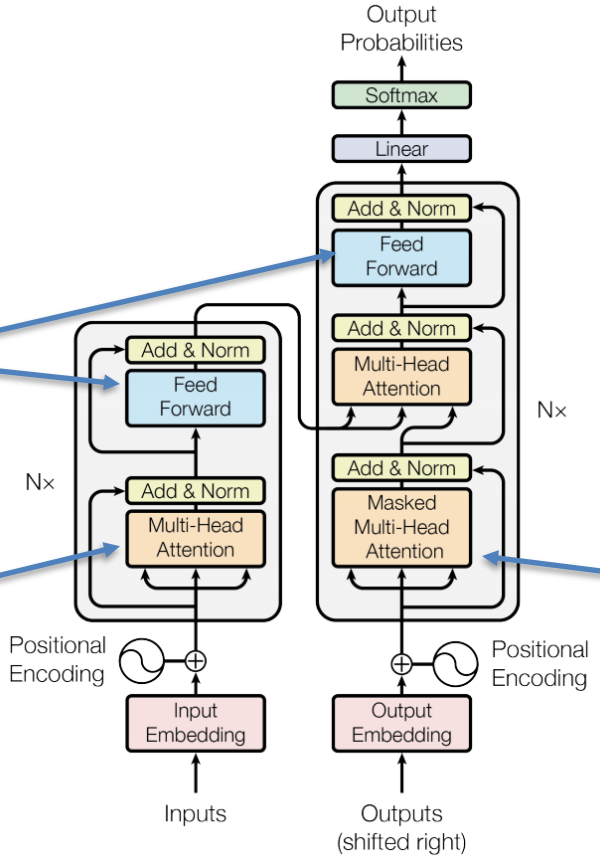

Masked Multi-Head Attention on the "decoder"

Multi-Head Attention

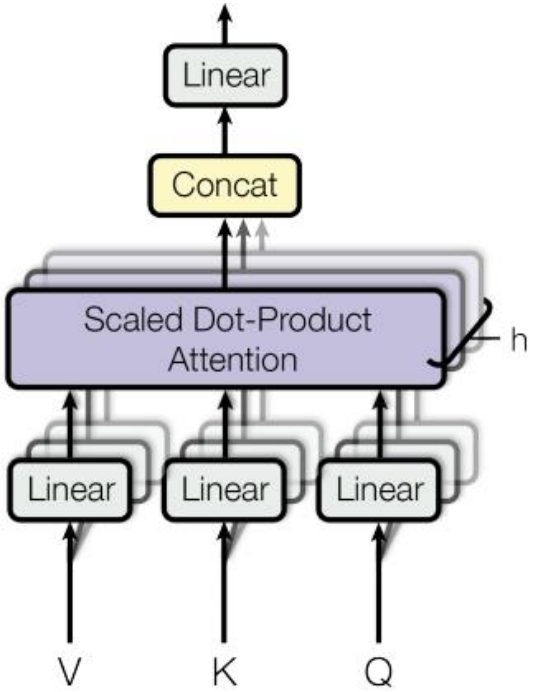

Intuition: Take the query Q, find the most similar key K, and then find the value V that corresponds to the key.

In other words, learn V, K, Q where:
V – here is a bunch of interesting things.
K – here is how we can index some things.
Q – I would like to know this interesting thing.

Loosely connected to Neural Turing Machines (Graves et al.).

Multi-Head Attention

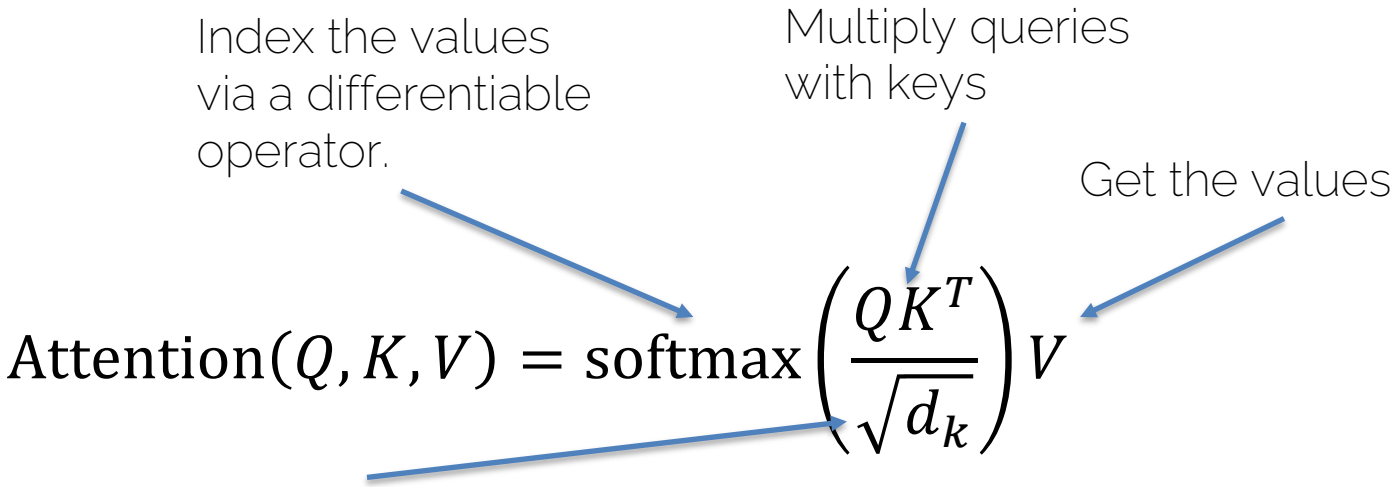

To train them well, divide by $\sqrt{d_k}$, "probably" because for large values of the key's dimension, the dot product grows large in magnitude, pushing the softmax function into regions where it has extremely small gradients.

How to train your neural network?

Setup Visualizations

- Always visualize train and validation loss curves.

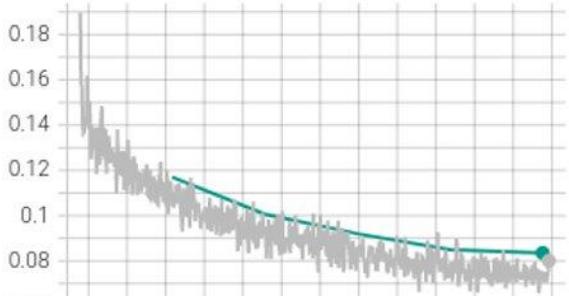

- Check data loading and augmentation by visualizing samples.

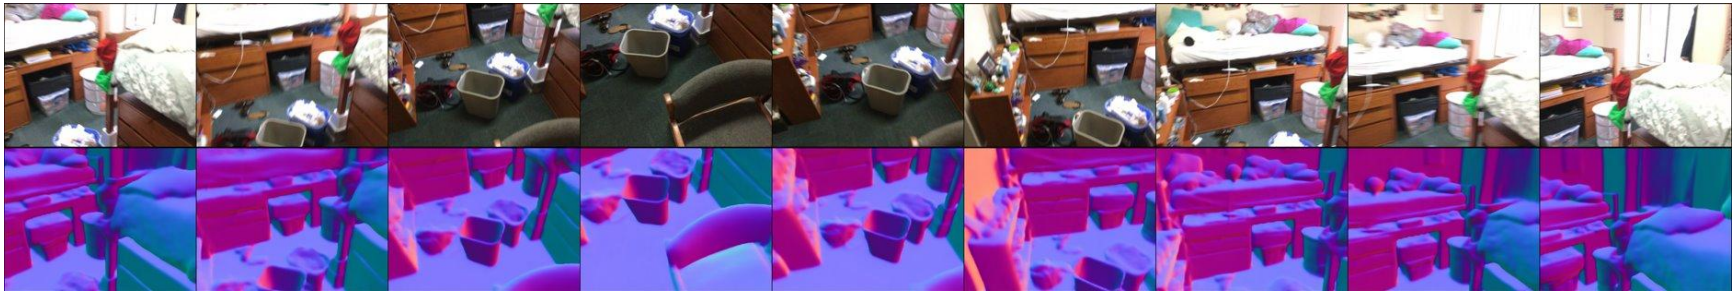

Setup Visualizations

- TensorBoard is easy to setup

  https://pytorch.org/tutorials/recipes/recipes/tensorboard_with_pytorch.html
  https://www.tensorflow.org/tensorboard/

- And provides an easy-to-use interface for visualizing image batches, metrics, histograms, videos …

  https://pytorch.org/docs/stable/tensorboard.html?highlight=summarywriter#

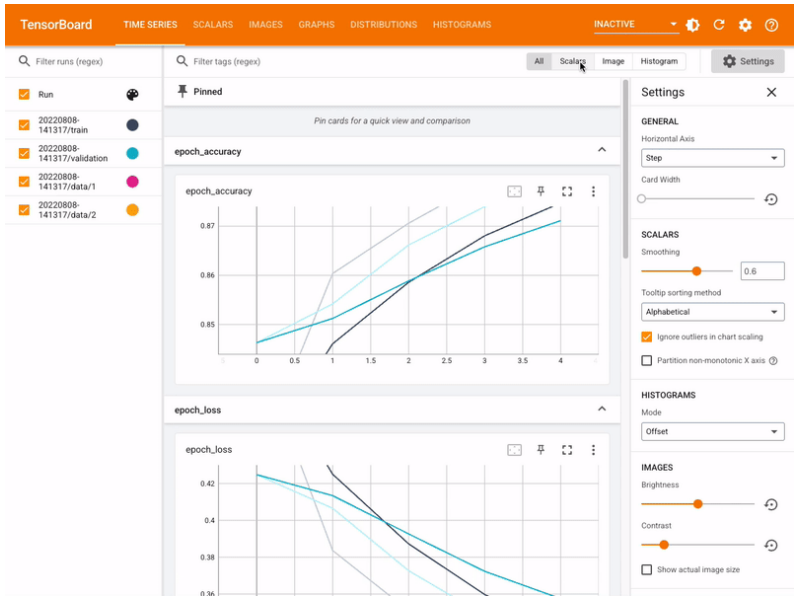

Is data loading correct?

- Data output (target): overfit to single training sample (needs to have 100% because it just memorizes input)
  - It's irrespective of input !!!

- Data input: overfit to a handful (e.g., 4) training samples
  - It's now conditioned on input data

Overfitting curves

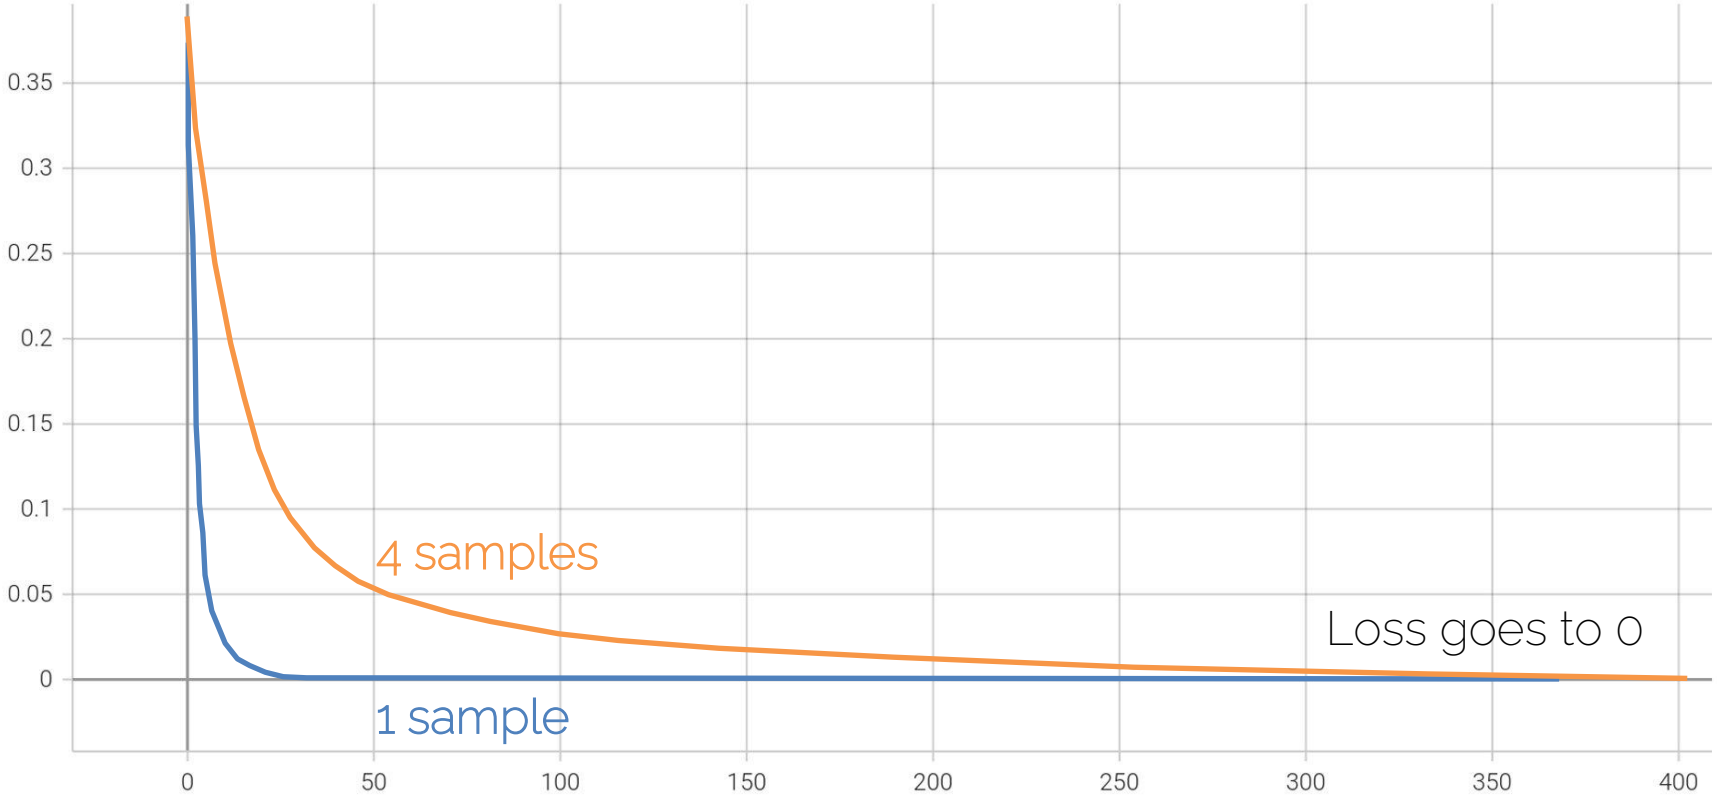

Debugging: overfitting -> generalization

- Move from overfitting to a hand-full of samples
  - 5, 10, 100, 1000…
  - At some point, we should see generalization

- Apply common sense: can we overfit to the current number of samples?

- Always be aware of network parameter count!

Check timings

- How long does each iteration take?
  - Get precise timings!!!
  - If an iteration takes > 500ms, things get dicey…

- Where is the bottleneck: data loading vs backprop?
  - Speed up data loading: smaller resolutions, compression, train from SSD – e.g., network training is good idea
  - Speed up backprop

- Estimate total timings: how long until you see some pattern? How long till convergence?

Network architecture

- 100% mistake so far: “let’s use super big network and train for two weeks and we see where we stand.” [because we desperately need those 2%…]

- Start with simplest network possible: rule of thumb divide #layers you started with by 5.

- Get debug cycles down – ideally, minutes!!!

Debugging

- Need train/val/test curves
  - Evaluation needs to be consistent!
  - Numbers need to be comparable

- Only make one change at a time
  - "I've added 5 more layers and double the training size, and now I also trained 5 days longer" – it's better, but WHY?

Overfitting

- ONLY THINK ABOUT THIS ONCE YOUR TRAINING LOSS GOES DOWN AND YOU CAN OVERFIT!

- Typically try this order:
- Network too big – makes things also faster ☺
- More regularization; e.g., weight decay
- Not enough data - makes things slower!
- Dropout - makes things slower!
- Guideline: make training harder -> generalize better

Severe overfitting!

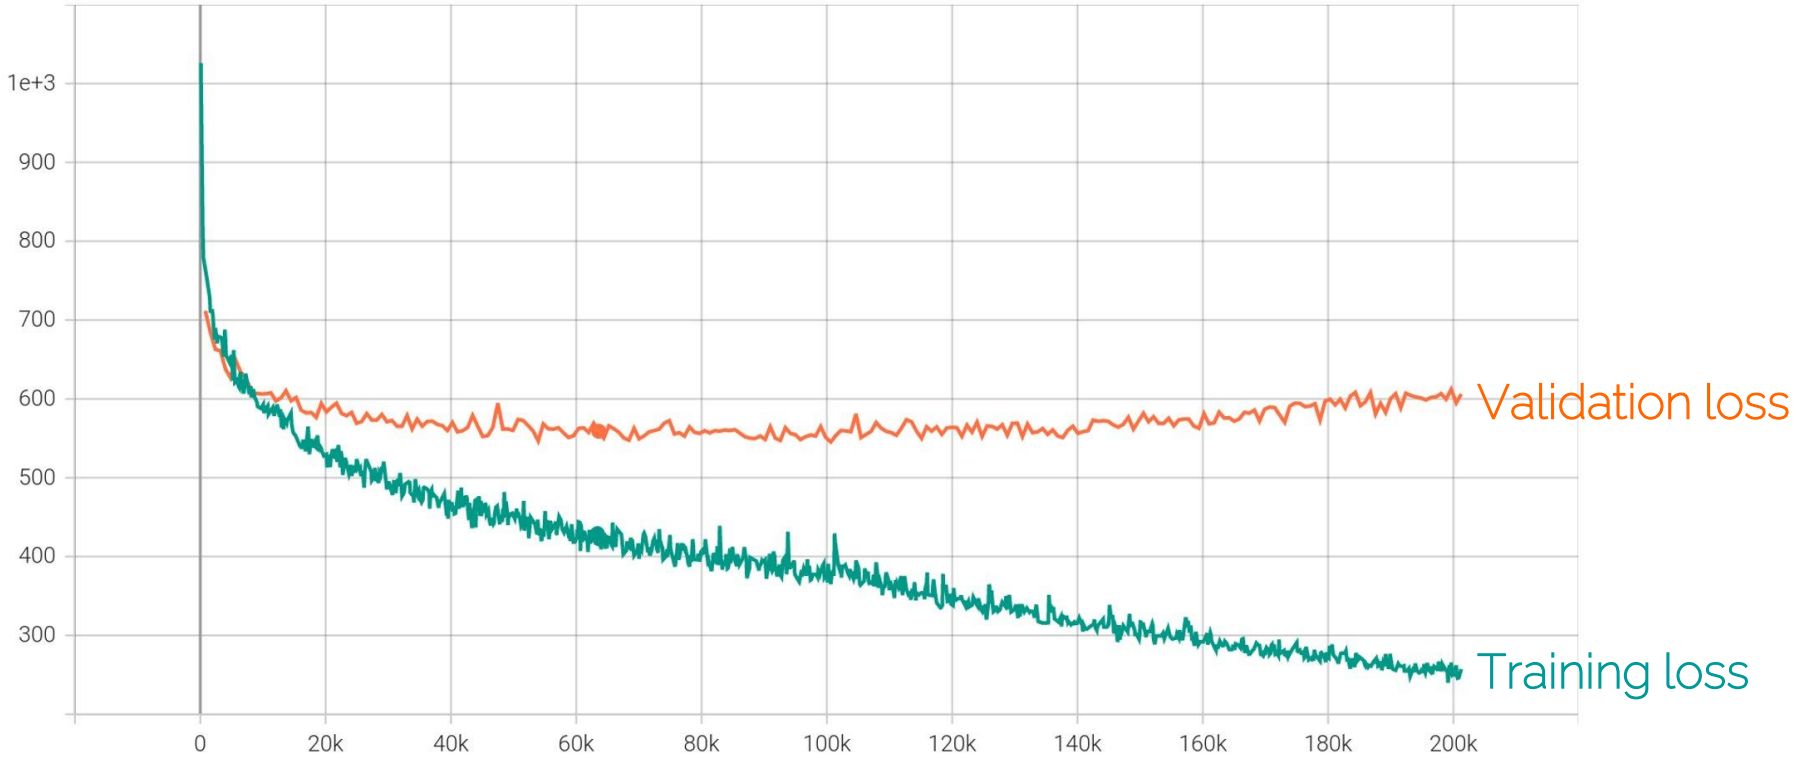

→ Try smaller network size, data augmentations, regularizations.

Moderate overfitting

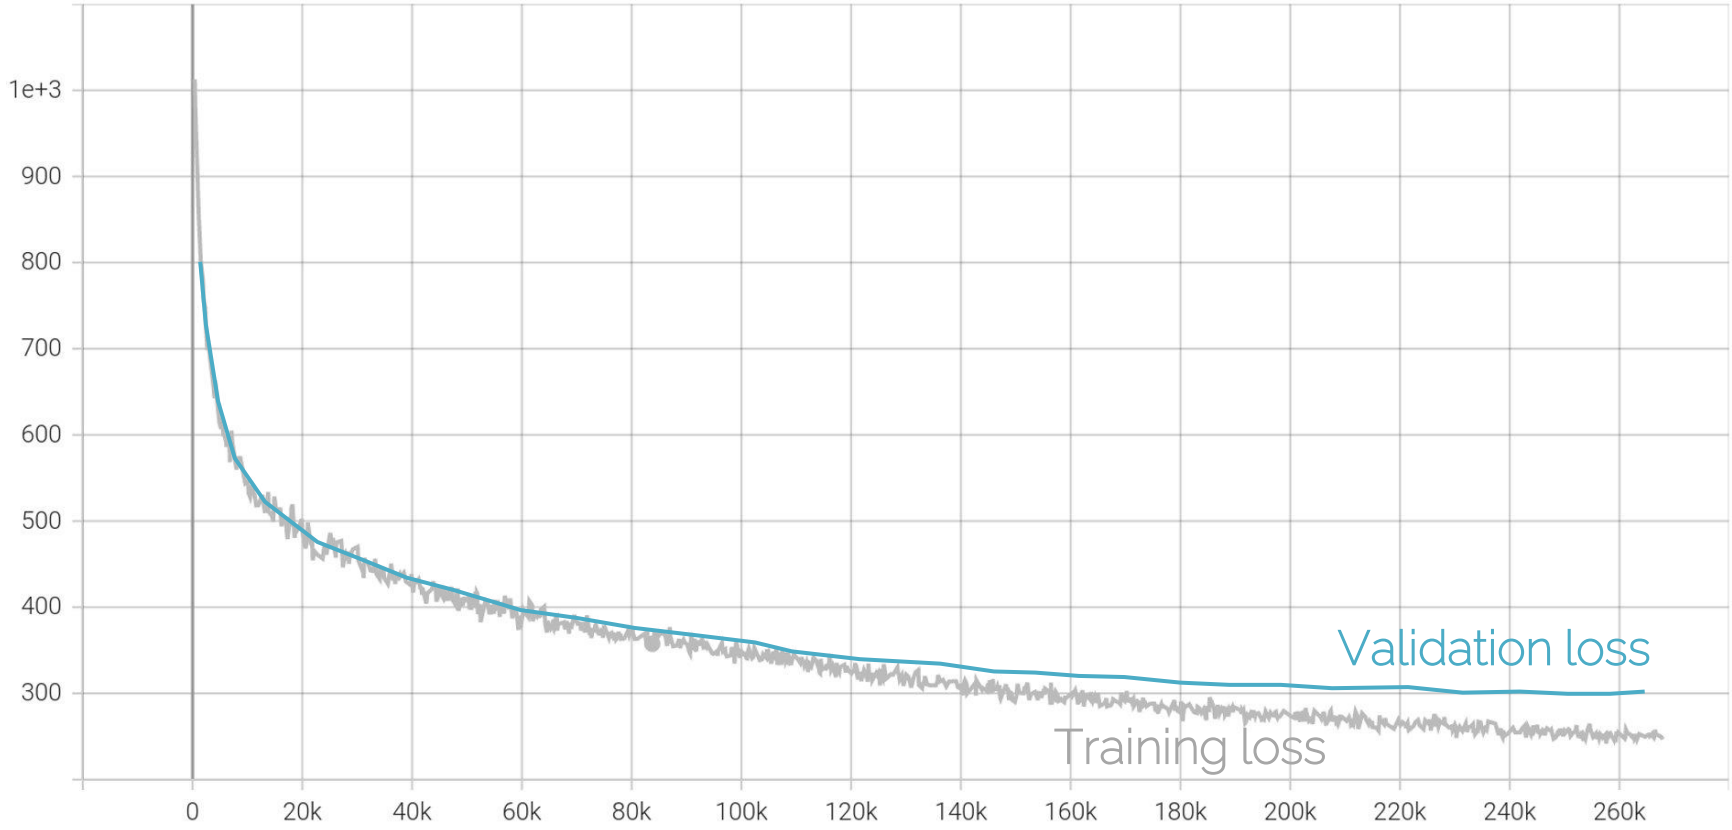

No overfitting

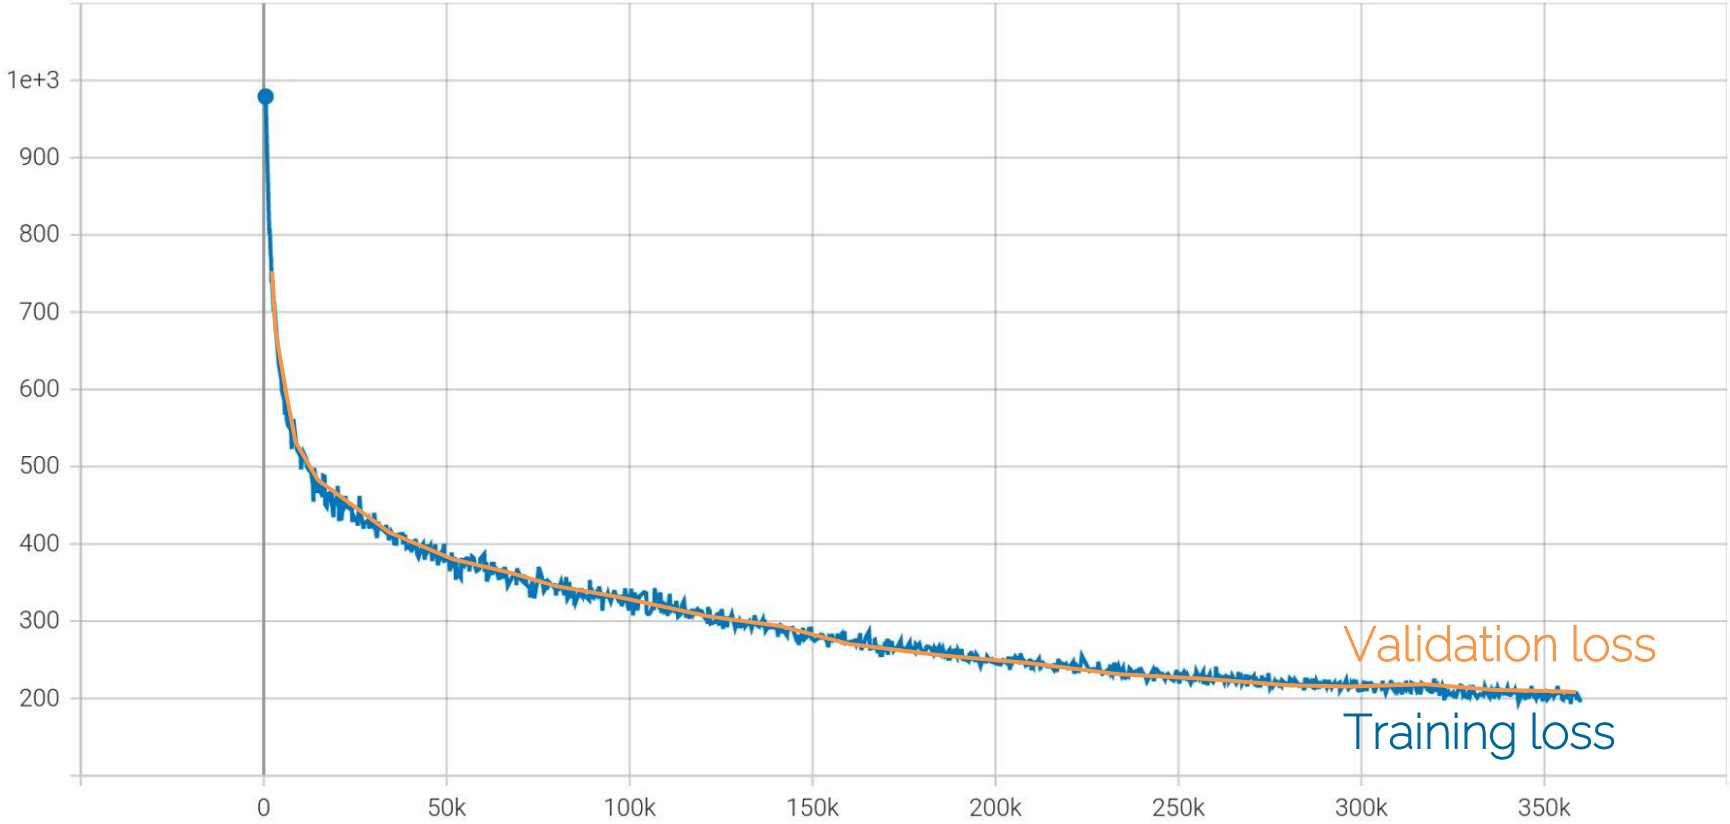

Pushing the limits!

- PROCEED ONLY IF YOU GENERALIZE AND YOU ADDRESSED OVERFITTING ISSUES!

- Bigger network -> more capacity, more power - needs also more data!
- Better architecture -> ResNet is typically standard, but InceptionNet architectures perform often better (e.g., InceptionNet v4, XceptionNet, etc.)
- Schedules for learning rate decay
- Class-based re-weighting (e.g., give under-represented classes higher weight)
- Hyperparameter tuning: e.g., grid search; apply common sense!

Bad signs…

- Train error doesn't go down…
- Validation error doesn't go down… (ahhh we don't learn)
- Validation performs better than train… (trust me, this scenario is very unlikely – unless you have a bug ☺)
- Test on train set is different error than train… (forgot dropout?)
- Often people mess up the last batch in an epoch…

- You are training set contains test data…
- You debug your algorithm on test data…

Visualizing importance

The occlusion experiment

- Block different parts of the image and see how the classification score changes

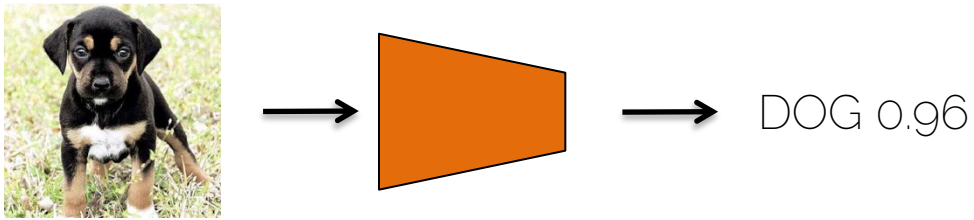

Zeiler and Fergus. „Visualizing and understanding convolutional neural networks". ECCV 2014

The occlusion experiment

- Block different parts of the image and see how the classification score changes

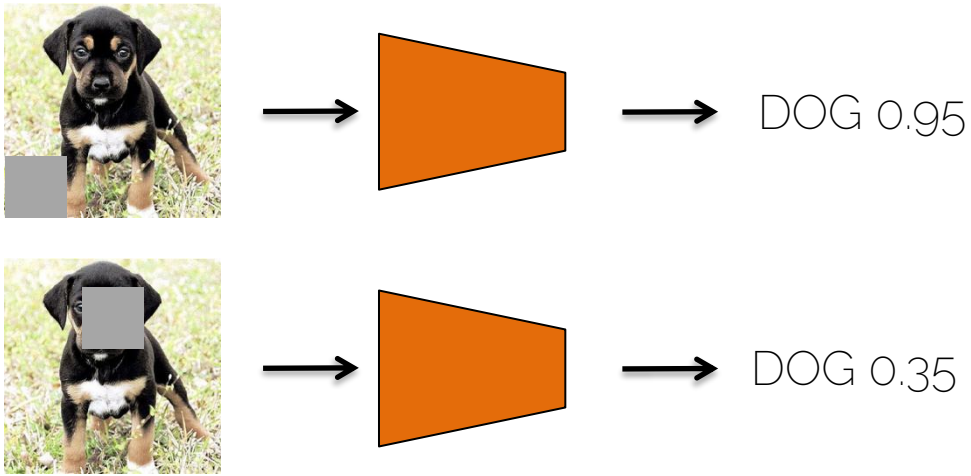

Zeiler and Fergus. „Visualizing and understanding convolutional neural networks". ECCV 2014

The occlusion experiment

- Create a map, where each pixel represents the classification probability if an occlusion square is placed in that region

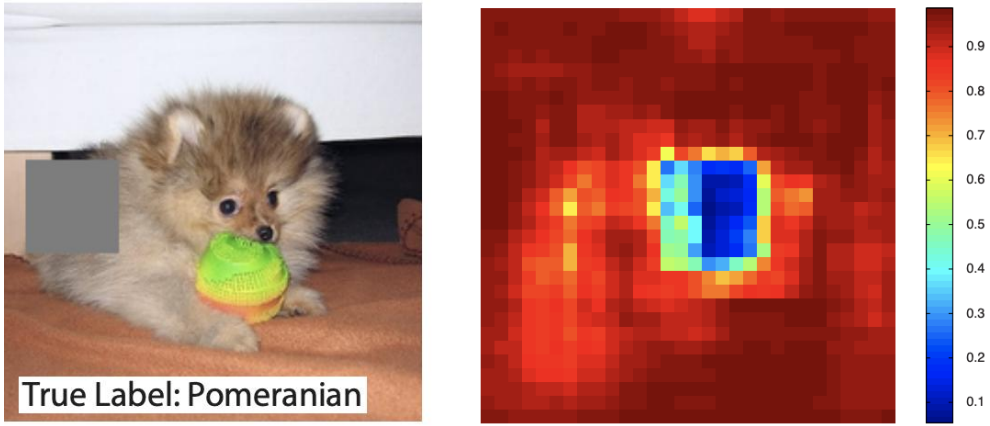

Zeiler and Fergus. „Visualizing and understanding convolutional neural networks". ECCV 2014

The occlusion experiment

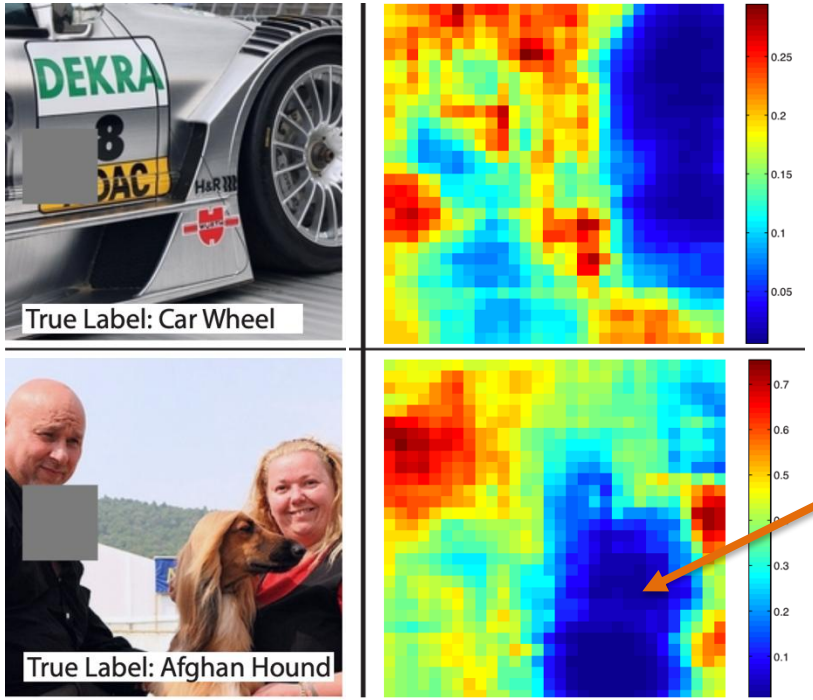

Most important pixels for classification

Zeiler and Fergus. „Visualizing and understanding convolutional neural networks". ECCV 2014

t-SNE

Intuition

- We want to visualize the last FC layer of AlexNet which dimension 4096

- We do a forward pass of all the images and get their 4096 representations

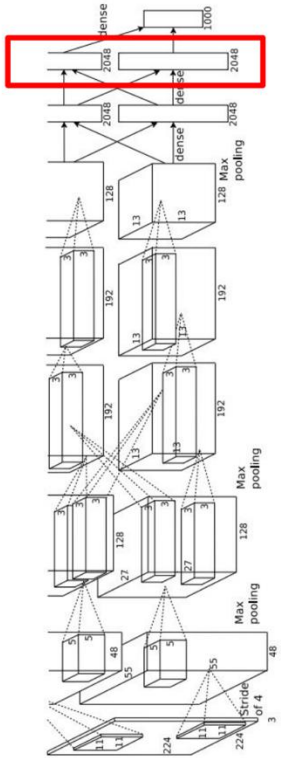

Intuition

- Nearest neighbor visualization

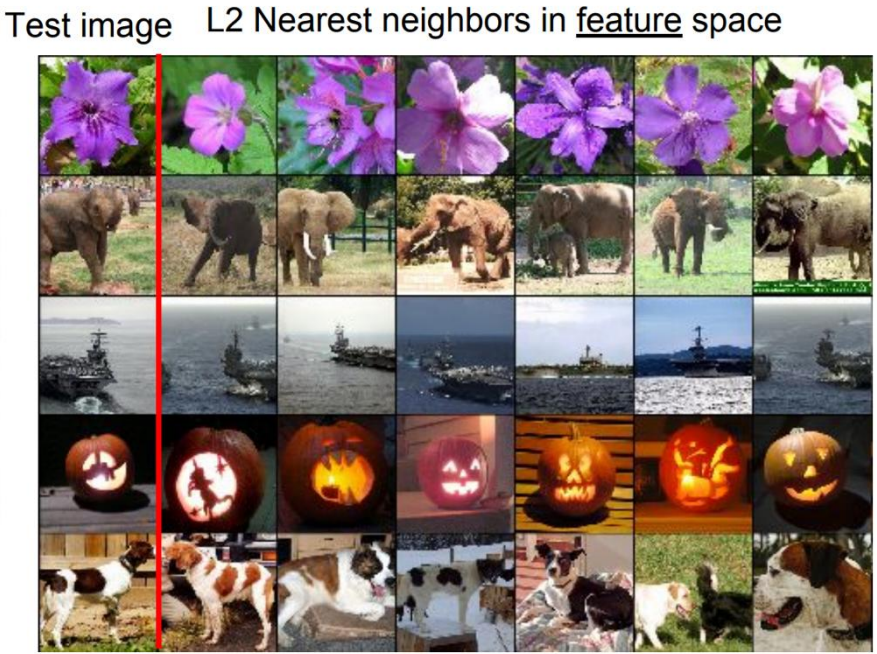

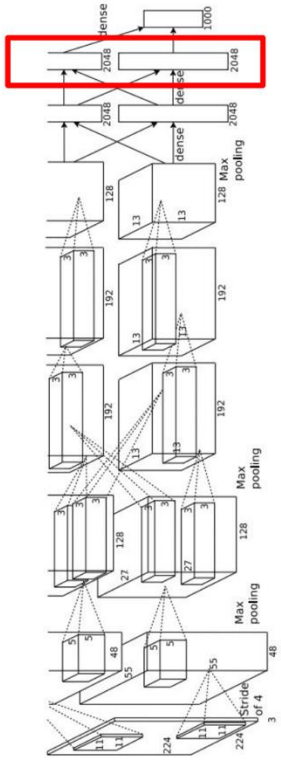

Image credit: Fei-Fei, Yeung, Johnson

Intuition

- How can I visualize these clusters in feature space?
- Map high-dimensional embedding to 2D map which preserves the pairwise distance of the points
- This mapping is done by t-SNE

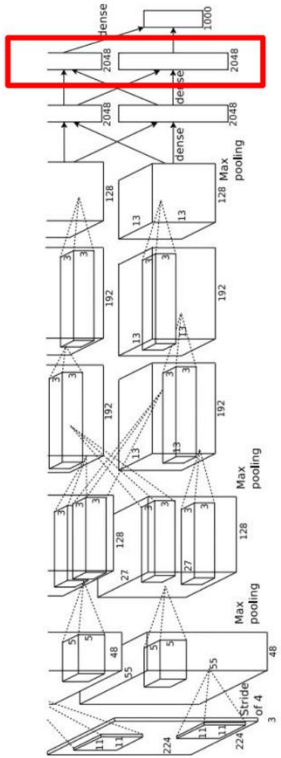

Image credit: Fei-Fei, Yeung, Johnson

t-SNE Visualization: MNIST

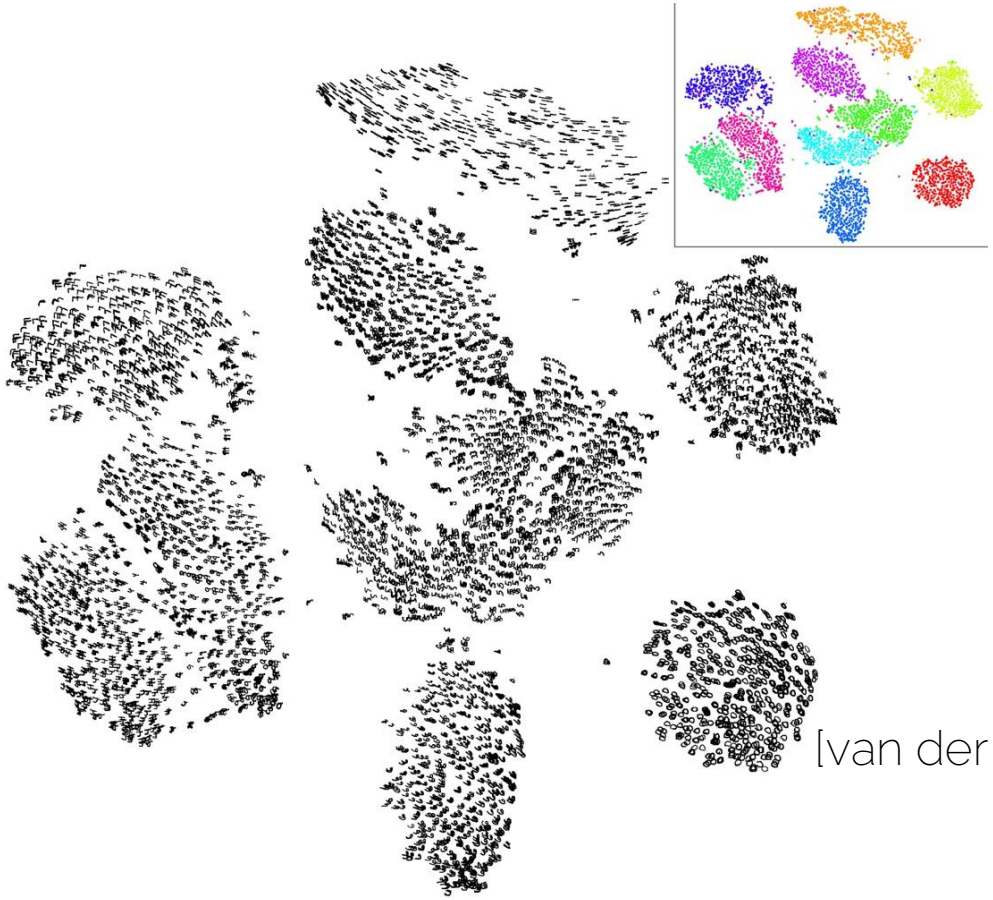

[van der Maaten et al.] t-SNE

t-SNE Visualization: ImageNet

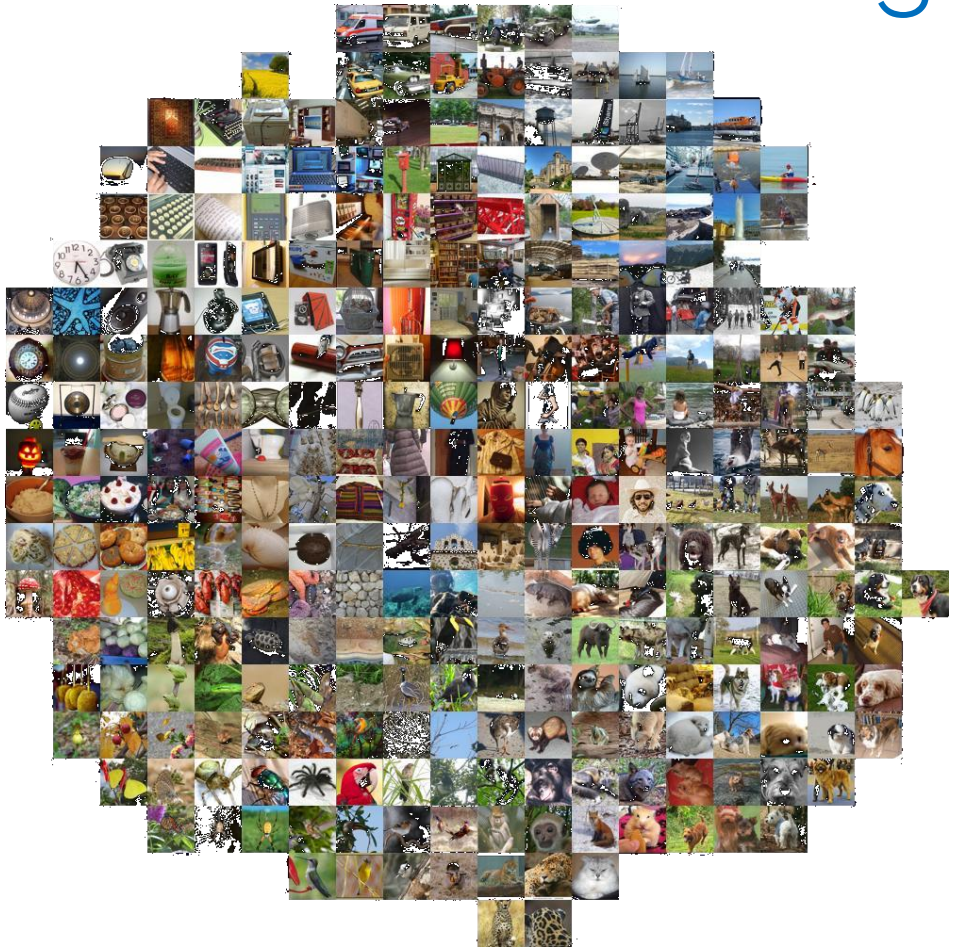

t-SNE Visualization: ShapeNet

When is t-SNE worth using?

- You can use it to debug your network
- Good for visualizing the clusters created by a Siamese network

Reading Homework

- [van der Maaten et al. 08] Visualizing Data using t-SNE
  - https://www.jmlr.org/papers/volume9/vandermaaten08a/vandermaaten08a.pdf

- TensorBoard Visualization
  - https://pytorch.org/tutorials/recipes/recipes/tensorboard_with_pytorch.html
  - https://www.tensorflow.org/tensorboard/

Literature

- I2DL Lecture
  - https://niessner.github.io/I2DL/
- Latest Research
  - https://niessnerlab.org/publications.html
- Social Media
  - How to start a research project: https://twitter.com/MattNiessner/status/1441027241870118913
  - Many good feeds for latest research papers

Thanks for watching!