Machine Learning and
Computer Vision Group

IST AUSTRIA
Institute of Science and Technology

Deep Learning with TensorFlow

http://cvml.ist.ac.at/courses/DLWT_W18

Lecture 6:
AlexNet

Deep Learning *with* Tensorflow

AlexNet *with* Vyacheslav Li

December 17, 2018

Outline

- Historical introduction
- Architecture of AlexNet
- Key characteristics of AlexNet
- After AlexNet breakthrough

Outline

- **Historical introduction**
- Architecture of AlexNet
- Key characteristics of AlexNet
- After AlexNet breakthrough

ImageNet Classification with Deep Convolutional Neural Networks (2012)

Alex Krizhevsky

Ilya Sutskever

Geoffry Hinton

According to Google Scholar it has **32,680** citations

~ 15 million labeled high resolution images

~ 22,000 categories

ILSVRC

- 1000 images in each of 1000 categories
- **1.2 million** training images
- 50,000 validation images
- 150,000 testing images

Revolution of Depth

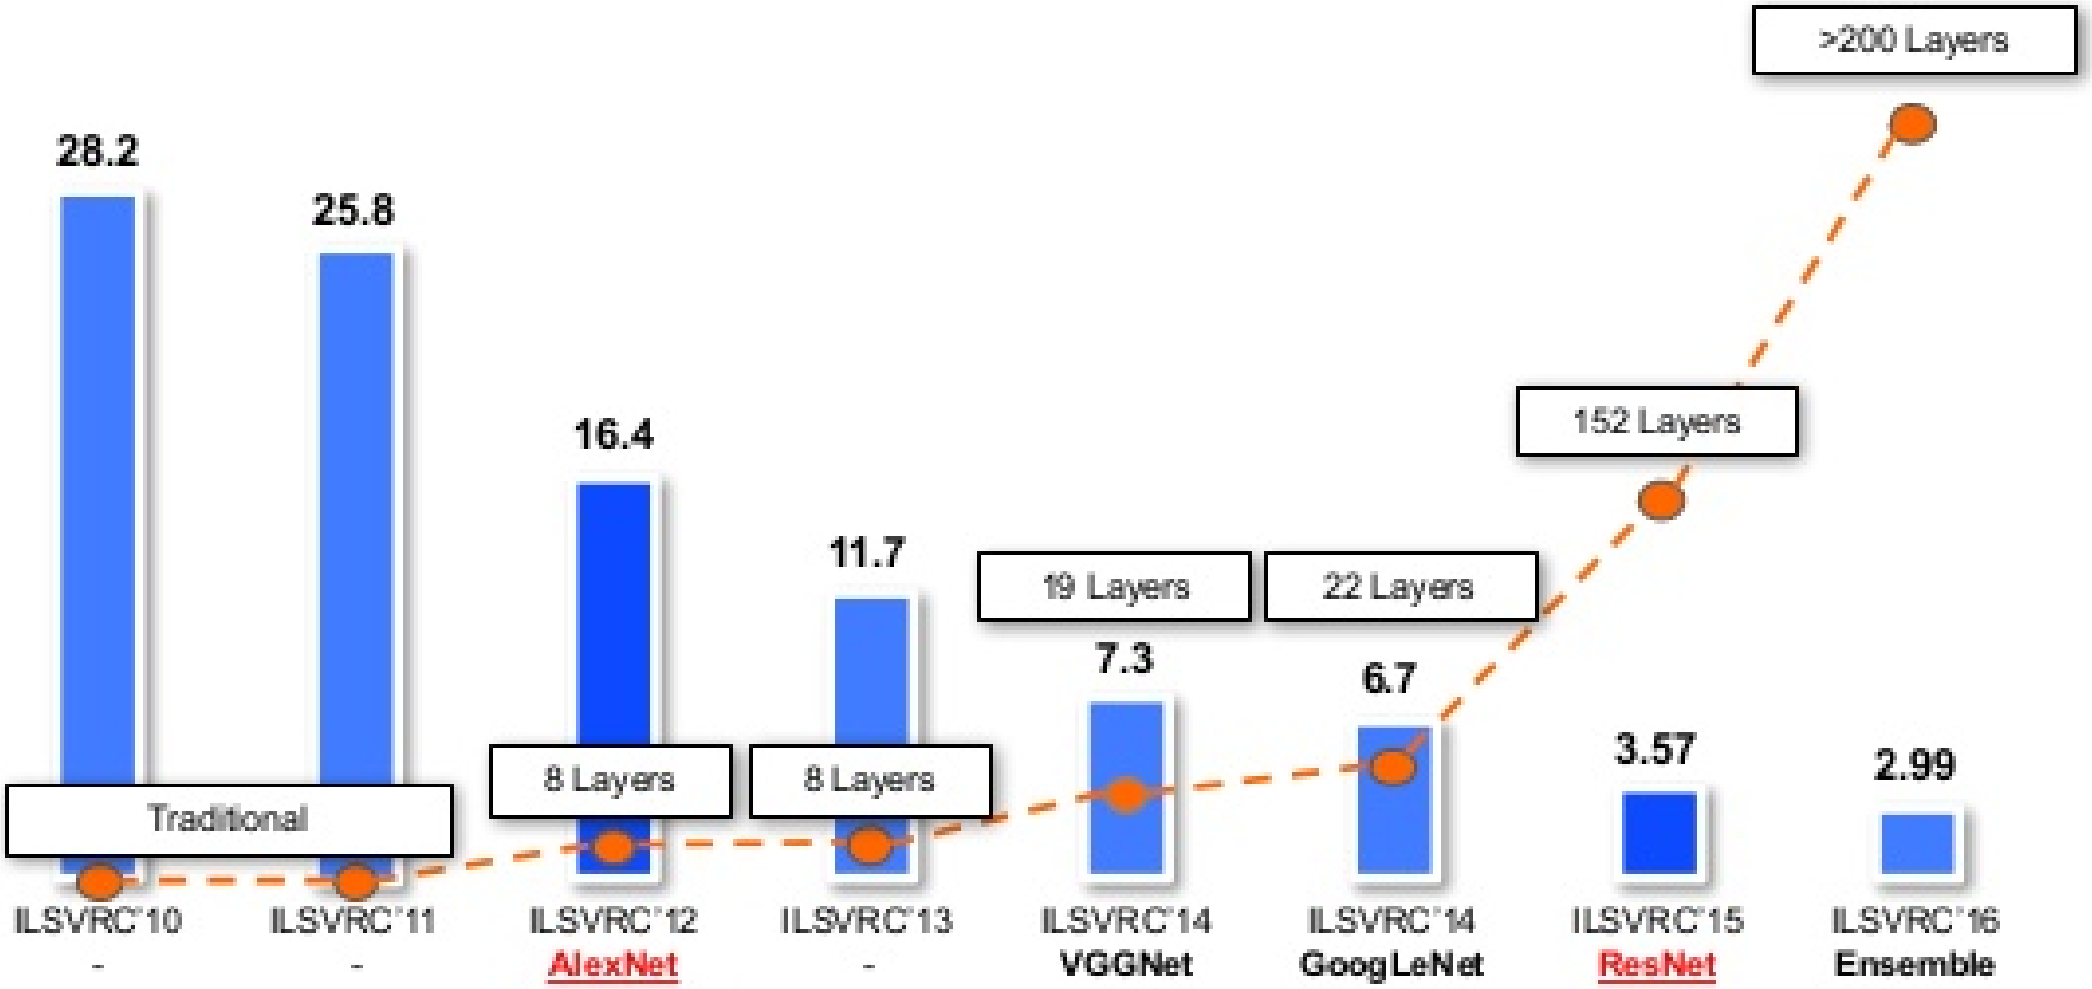

A bit of a history

- 1943: McCulloch & Pitts show that neurons can be combined to construct a Turing machine (using ANDs, ORs, & NOTs).
- 1985: The backpropagation algorithm by Geoffrey Hinton et al revitalizes the field
- 1998: CNNs with Backpropagation for document analysis by Yan LeCun
- 2012 : **AlexNet** by Alex Krizhevesky in 2012

Outline

- Historical introduction
- **Architecture of AlexNet**
- Key characteristics of AlexNet
- After AlexNet breakthrough

Architecture of AlexNet

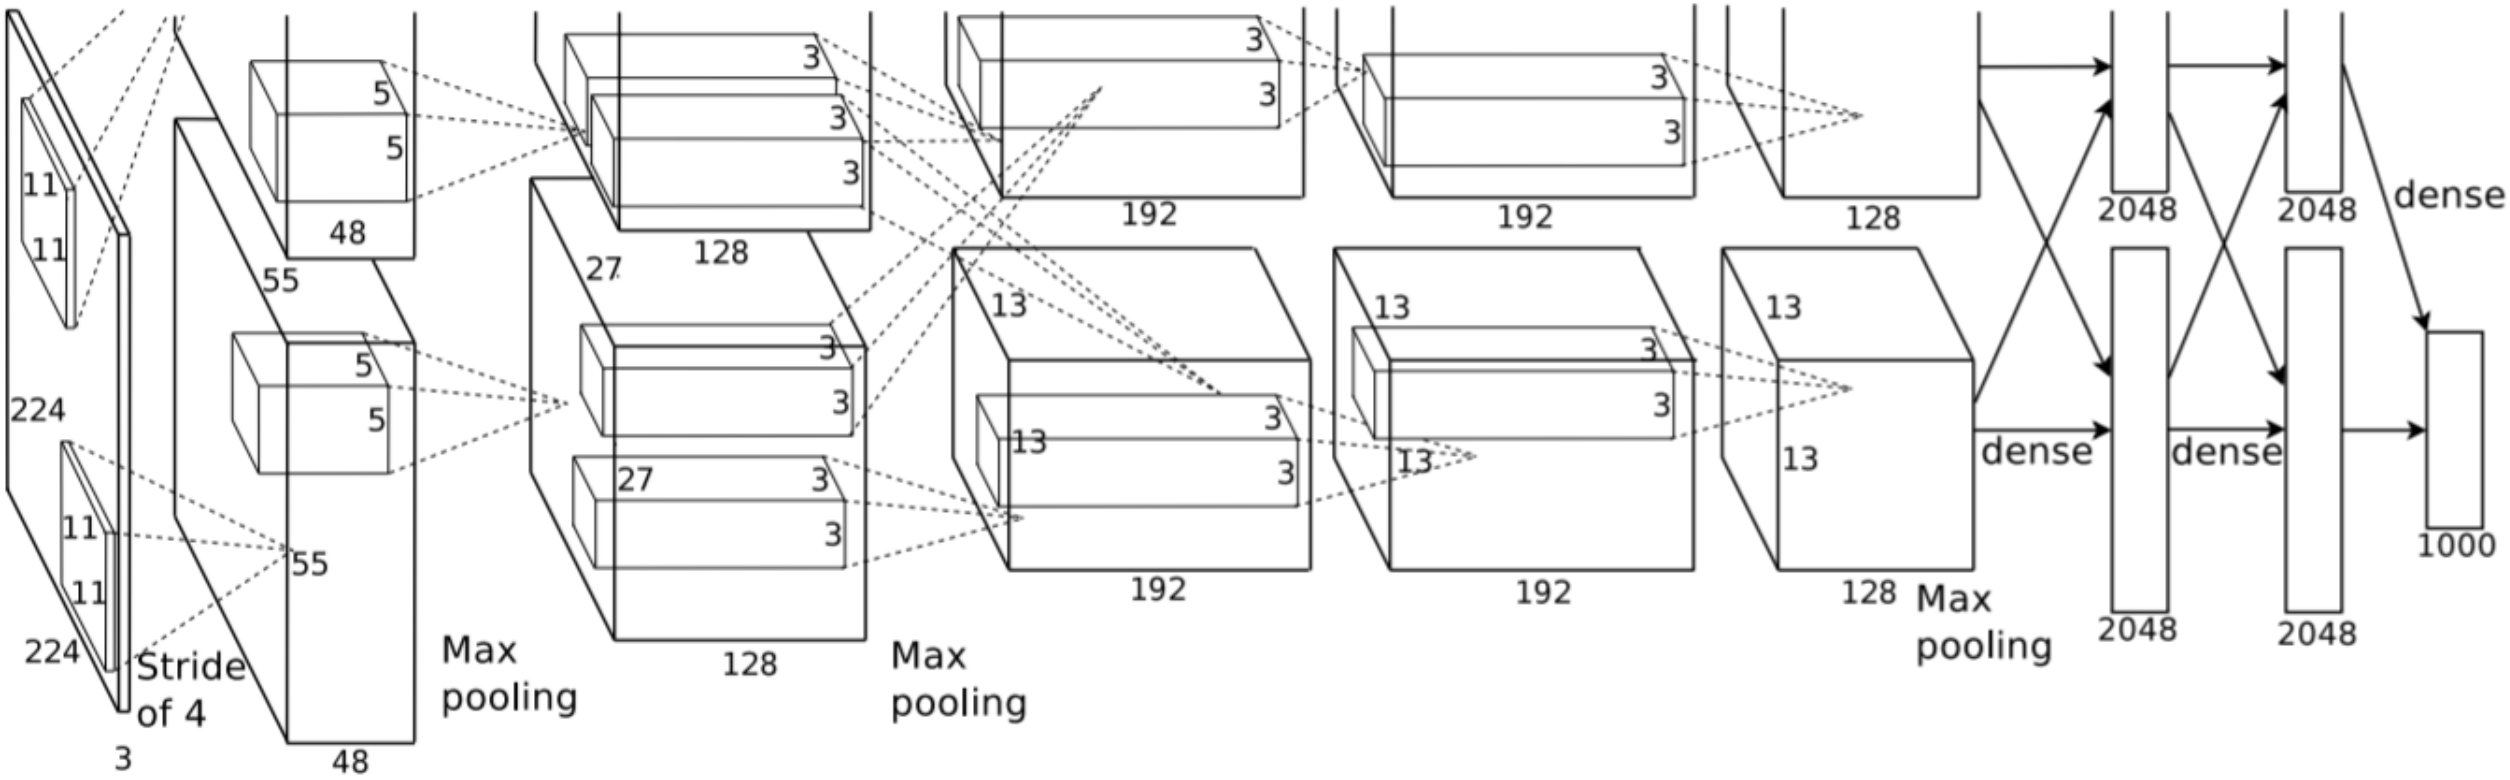

62.3 million parameters and needs **1.1 billion** computation units in a forward pass

Architecture of AlexNet

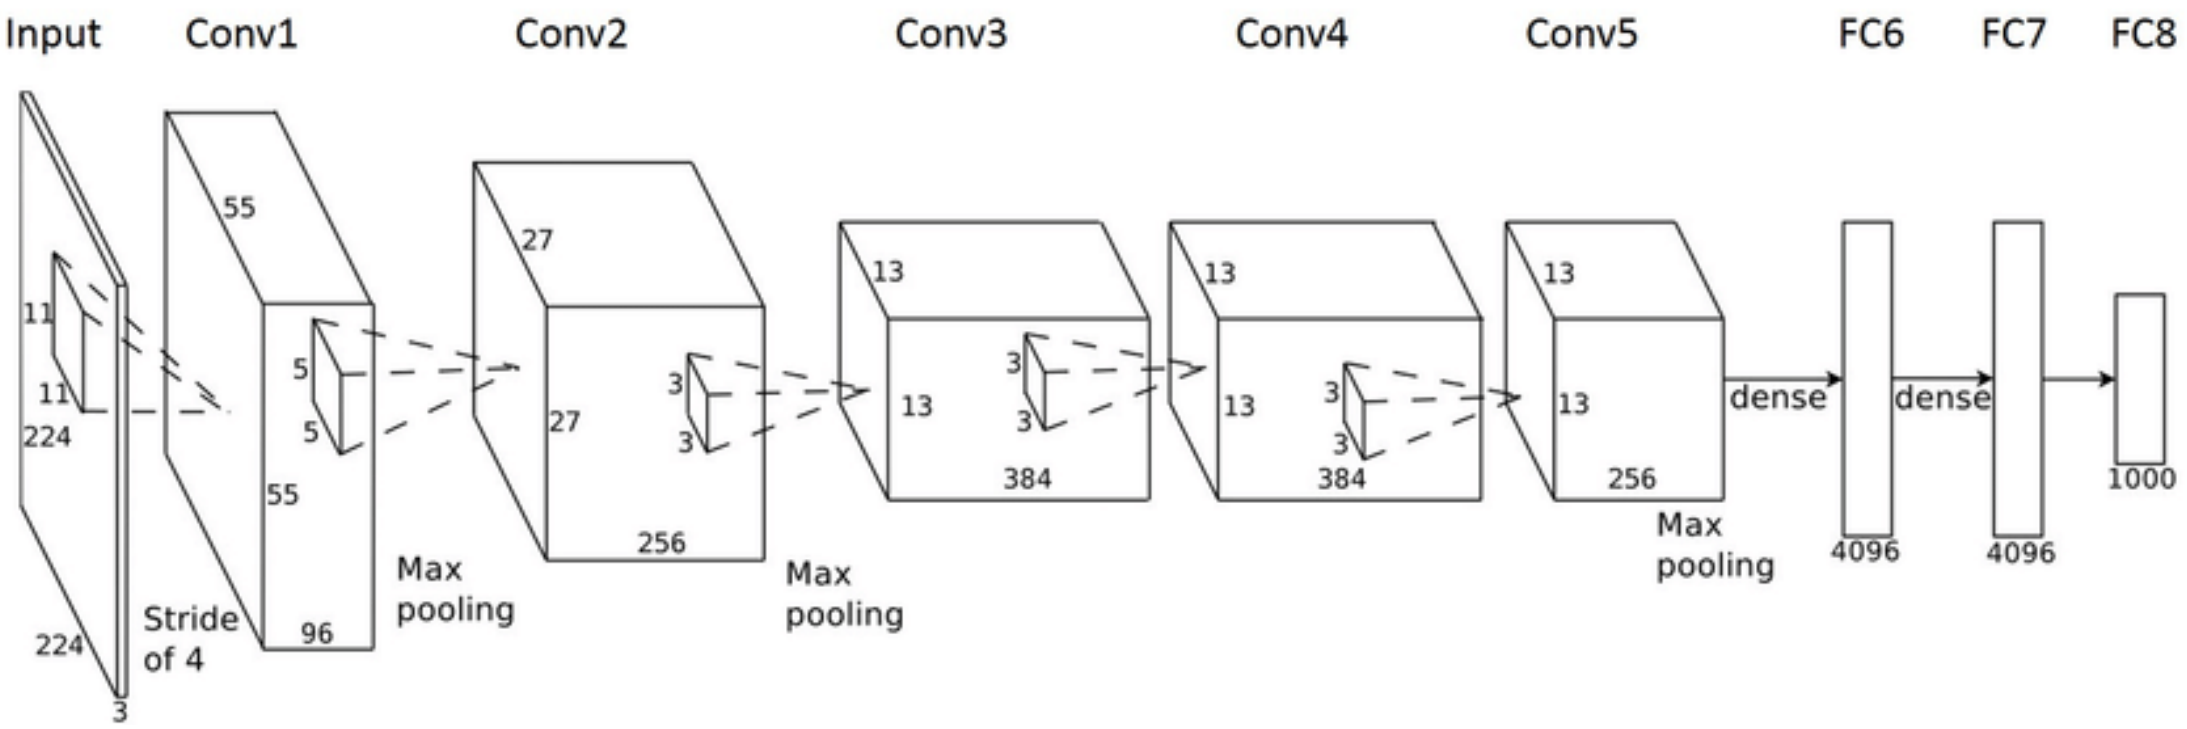

Outline

- Historical introduction
- Architecture of AlexNet
- **Key characteristics of AlexNet**
- After AlexNet breakthrough

ReLU

The ReLU activation was first used in AlexNet

$$y = \max(0, x)$$

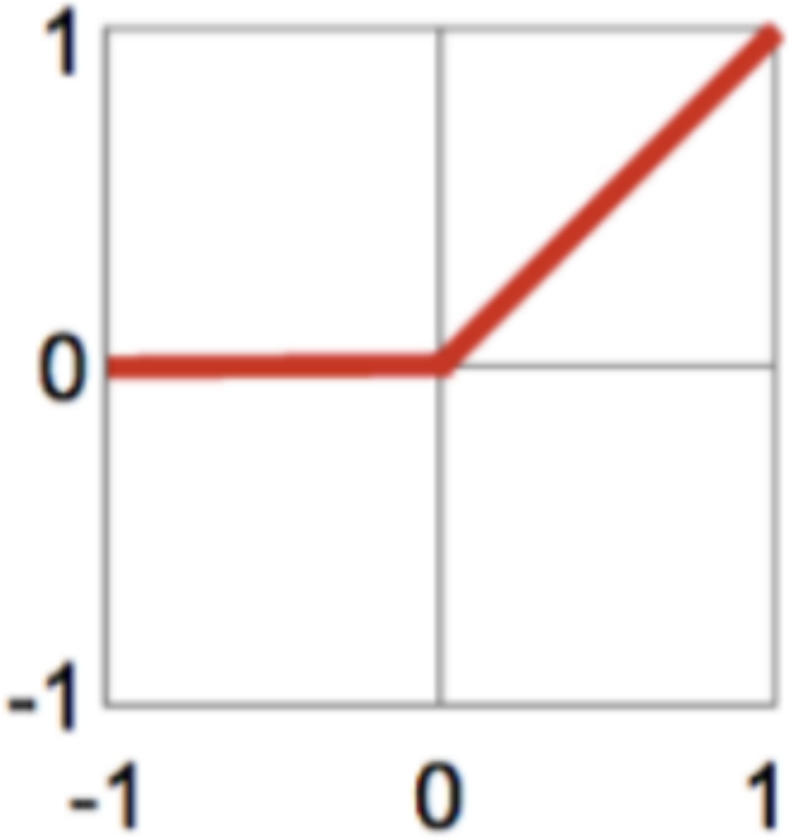

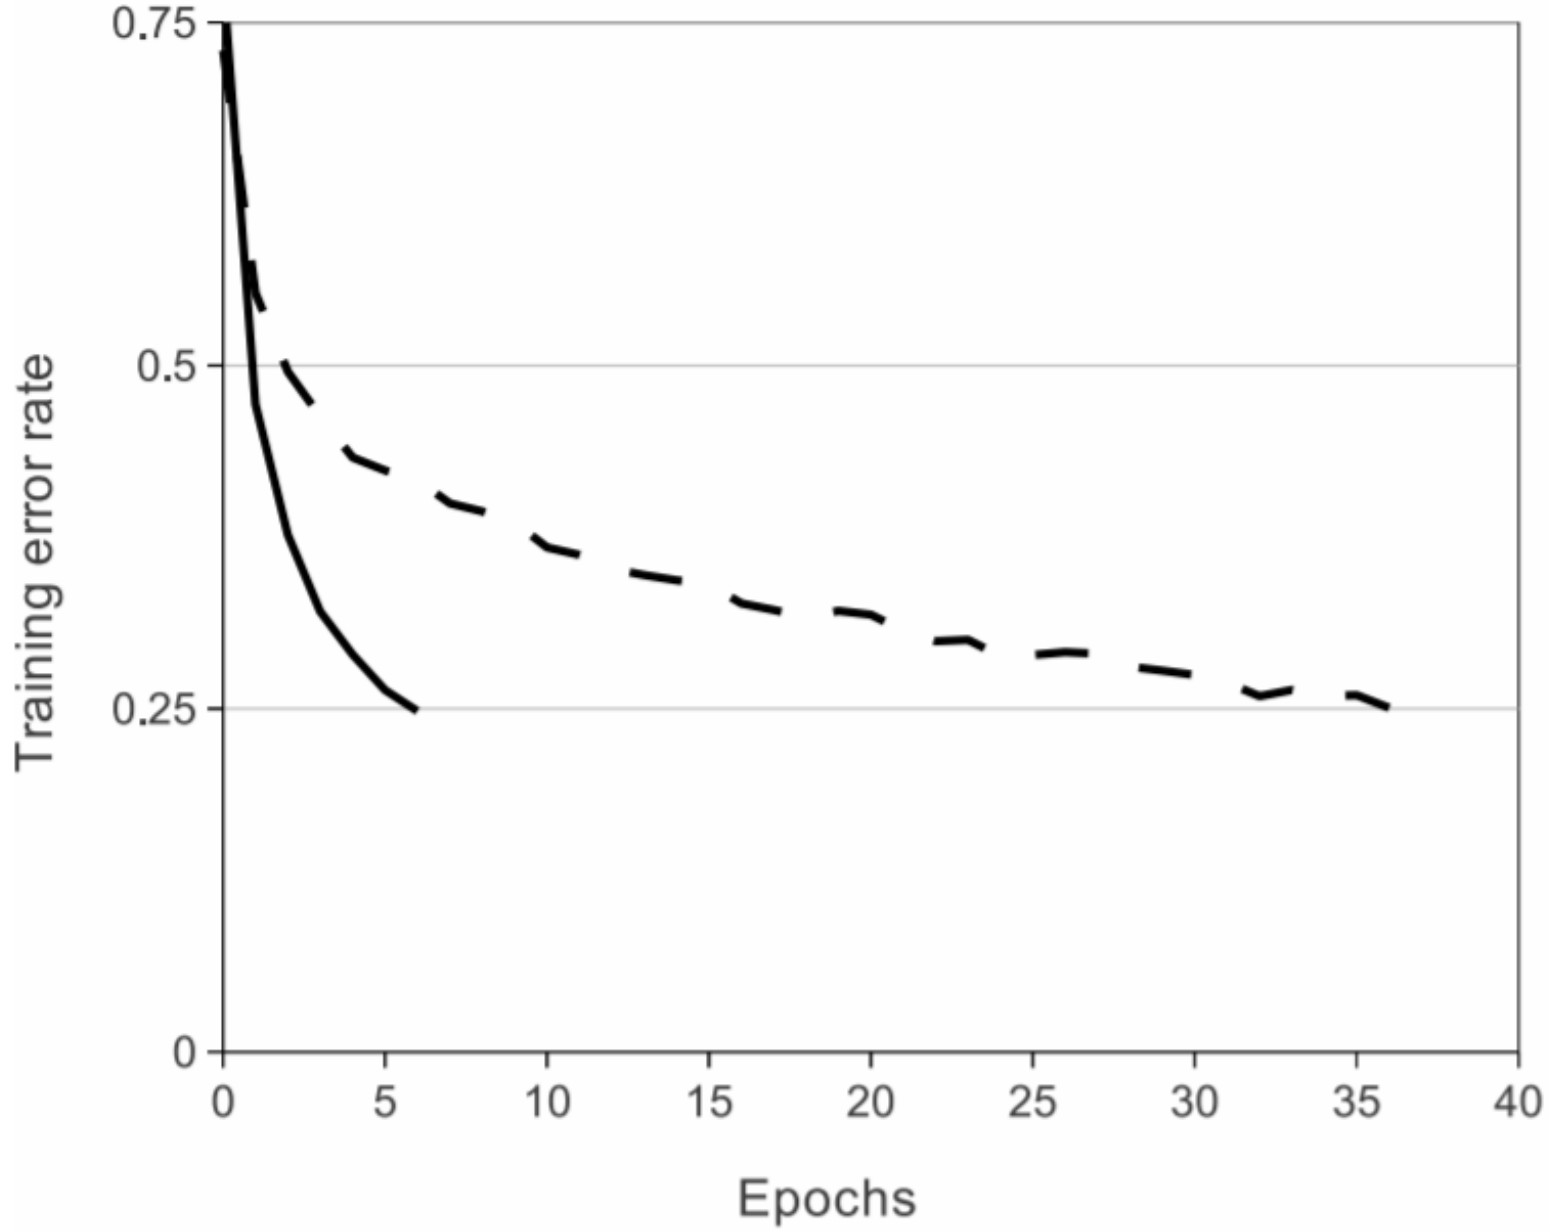

Trained on CIFAR – 10 dataset

Why ReLU faster?

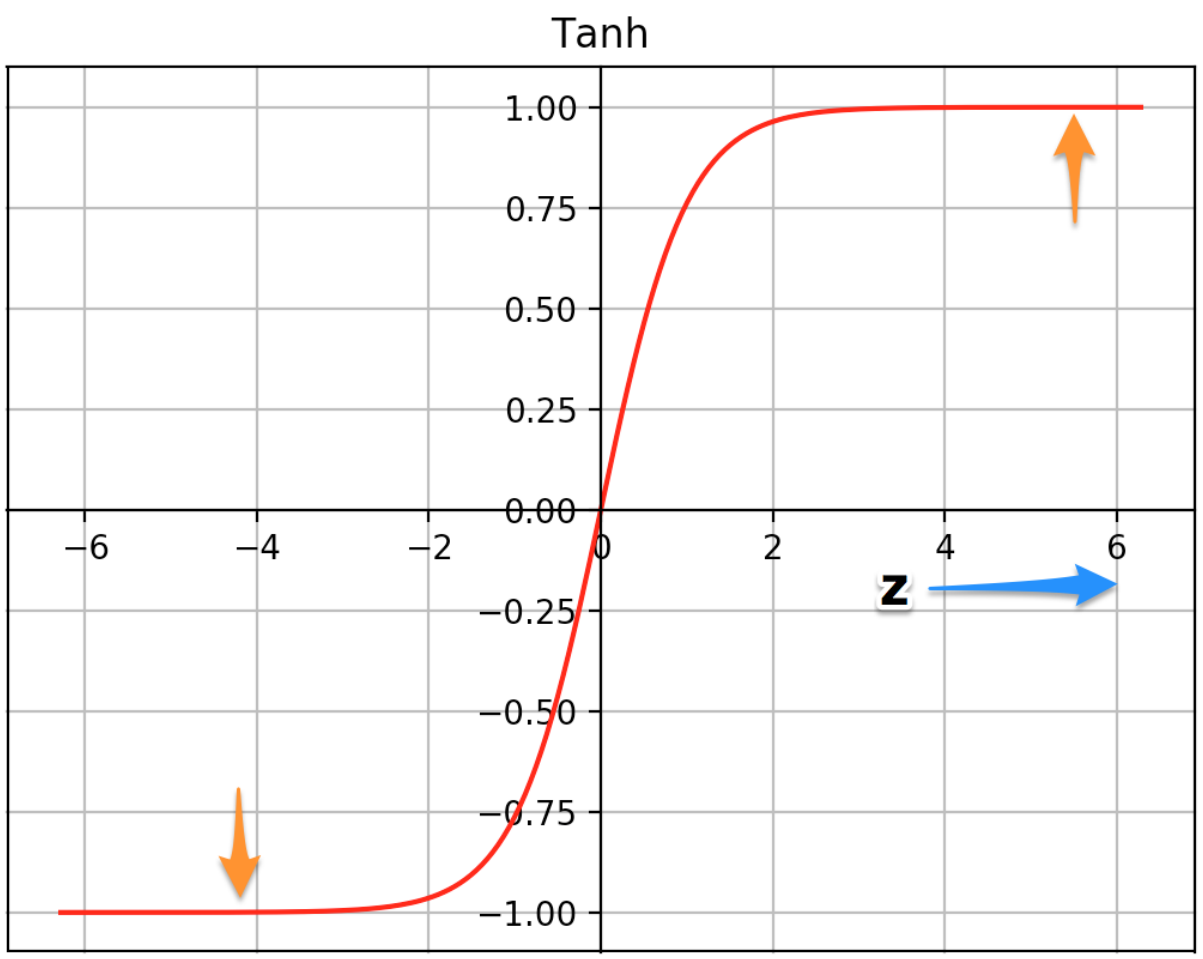

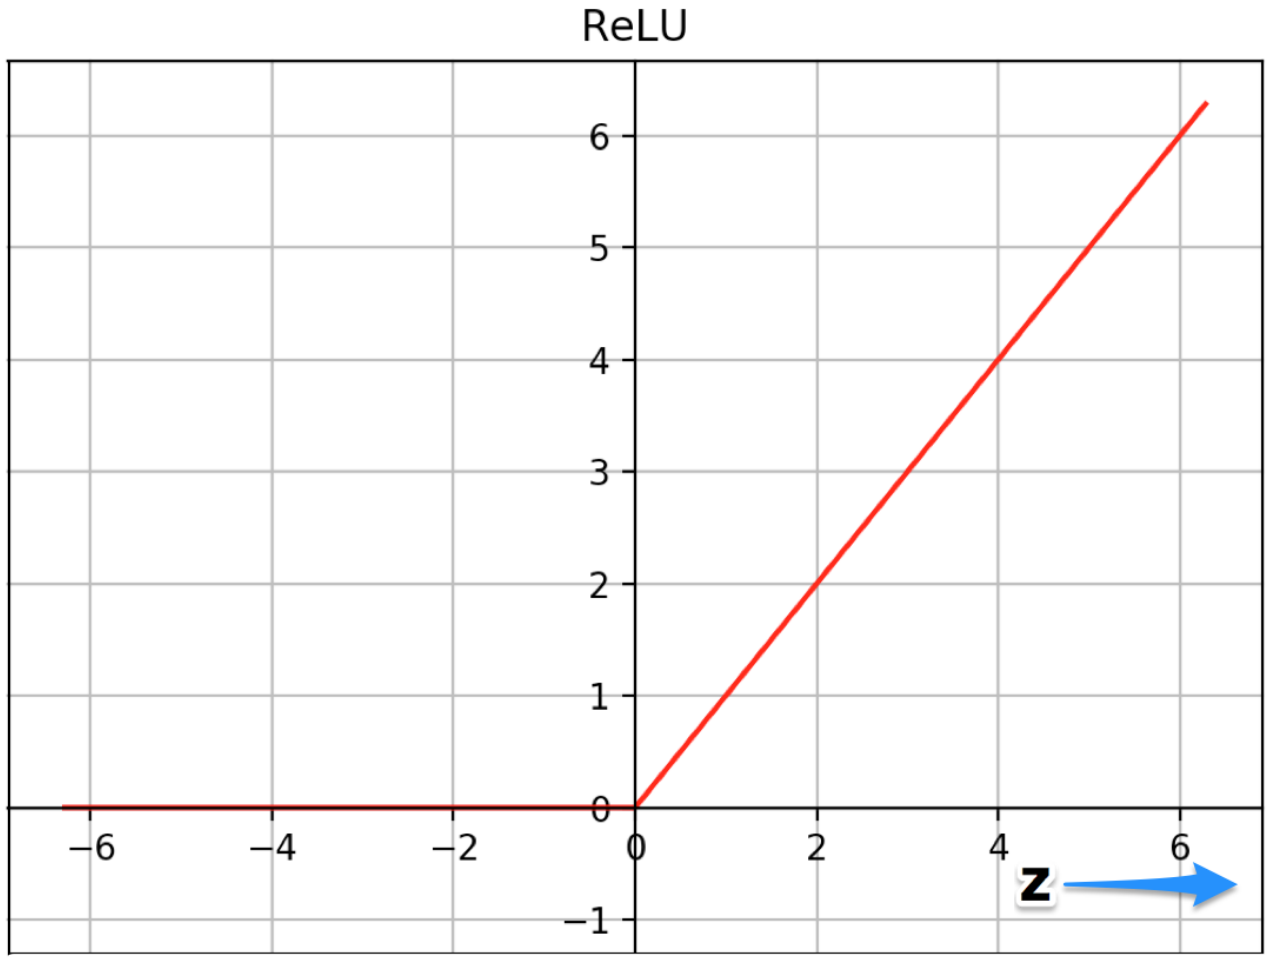

Local Response Normalization

$$b^i_{x,y} = a^i_{x,y} / \left(k + \alpha \sum_{j=\max(0,i-n/2)}^{\min(N-1,i+n/2)} (a^j_{x,y})^2 \right)^{\beta}$$

Response normalization reduced top-1 and top-5 error rates by **1.4%** and **1.2%**, respectively.

Max pooling

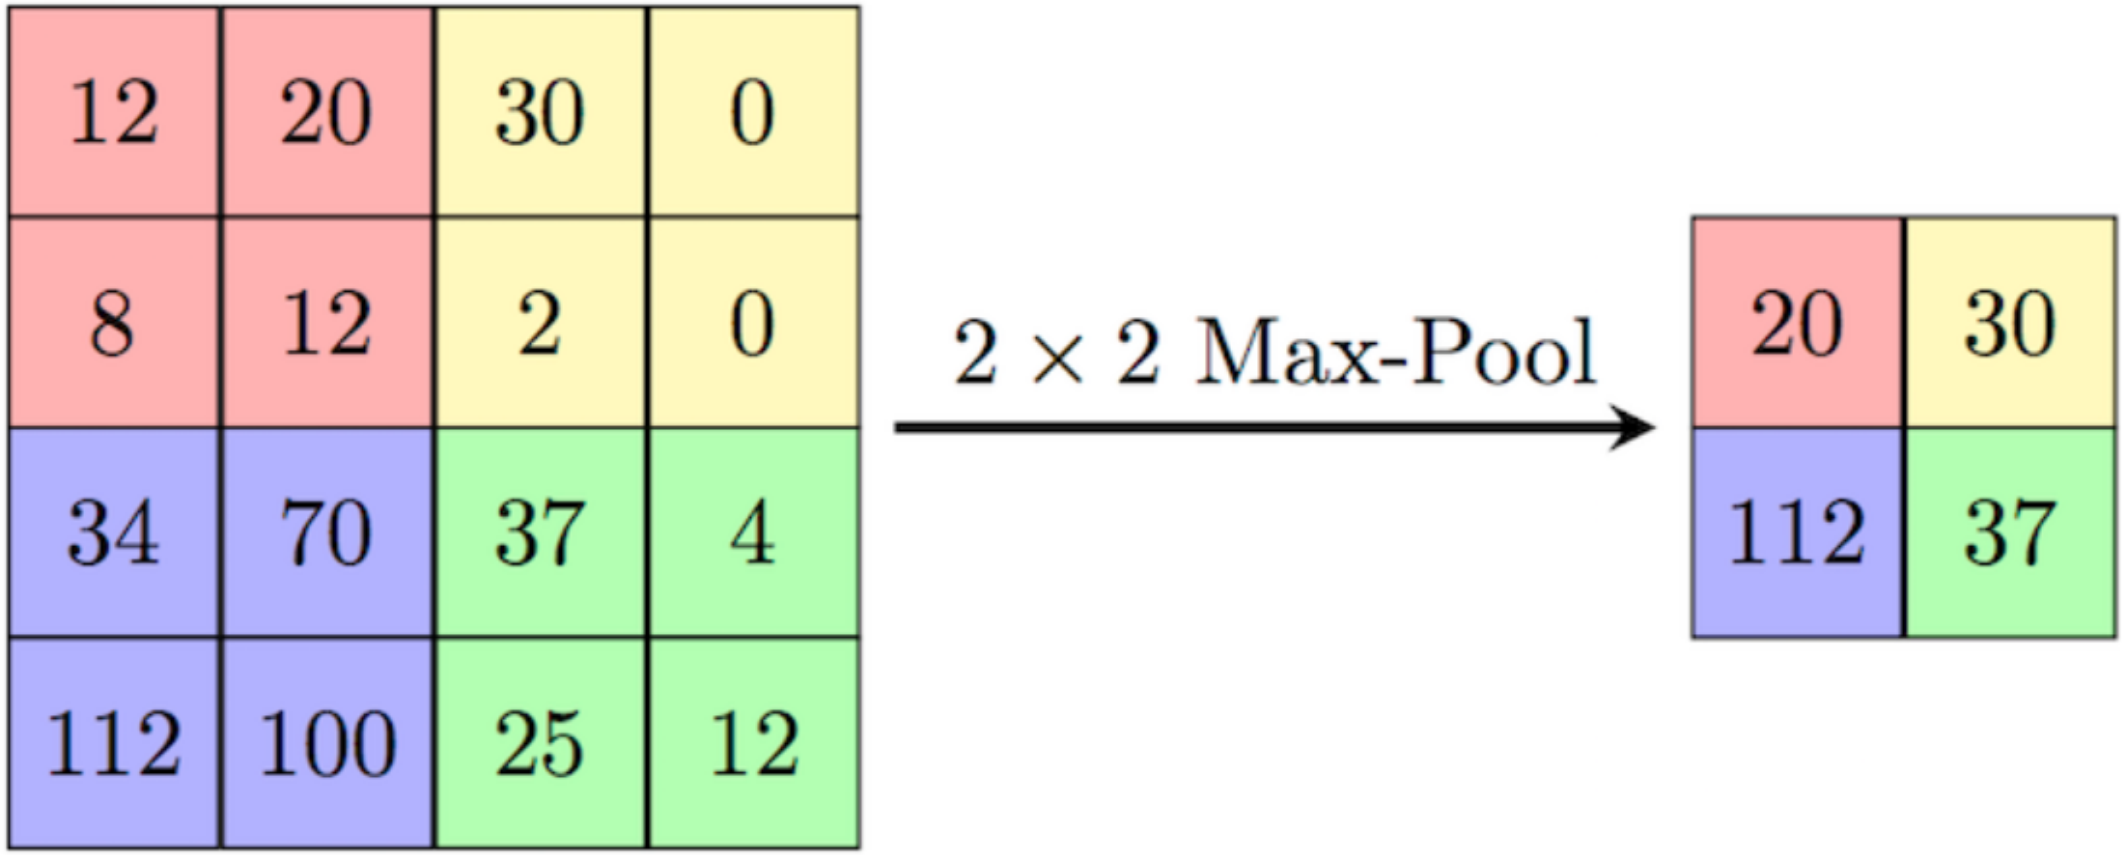

Overlapping Max pooling

S < Z instead of traditional S=Z (S: stride, Z: size of filter matrix)
This scheme reduces the top-1 and top-5 error rates by 0.4% and 0.3%, respectively

12	20	30	0
8	12	2	0
34	70	37	4
112	100	25	12

20	30	30
70	70	37
112	100	37

Reducing Overfitting

Data Augmentation

1024

500

Resize & Crop

256

256

Reducing overfitting

Data Augmentation

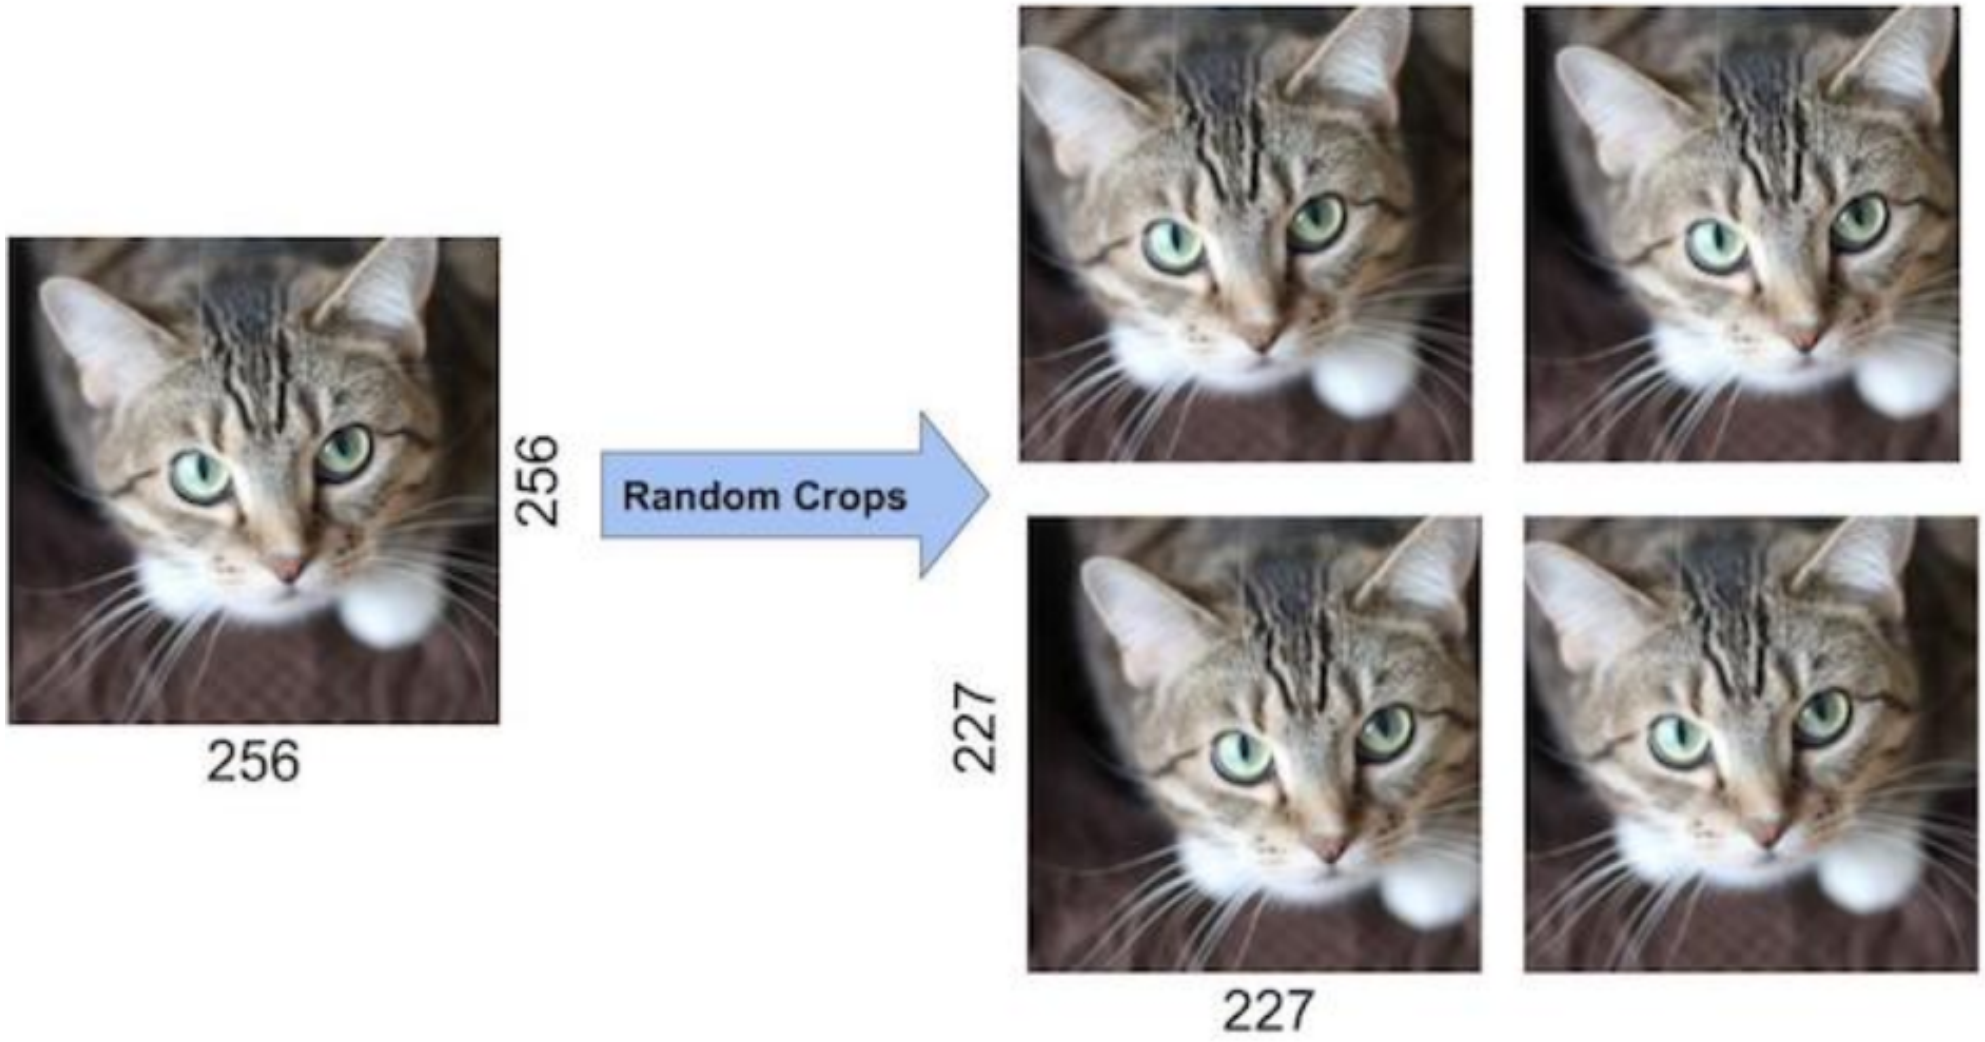

Reducing overfitting

Data Augmentation

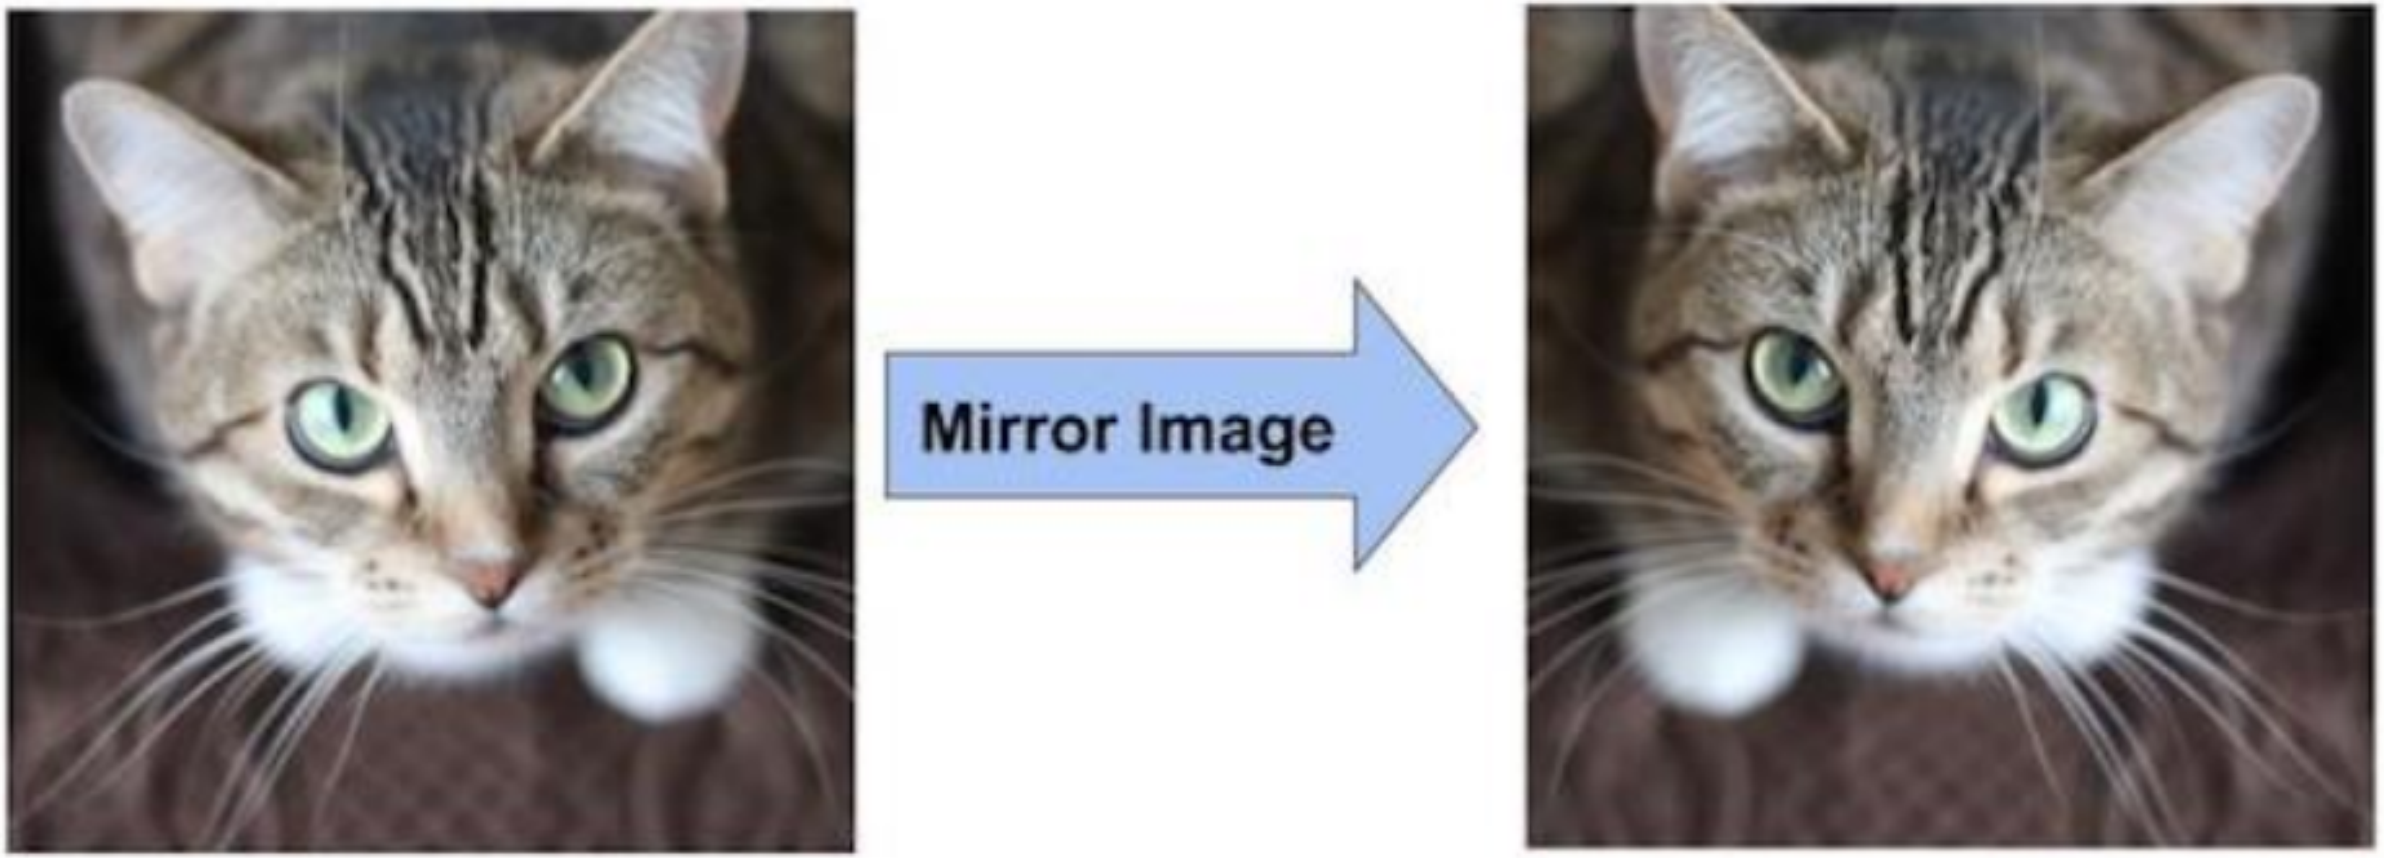

Increased the number of training data by factor of **2048!**

Reducing overfitting

Data Augmentation

Reducing overfitting

Data Augmentation

Principal Components Analysis was performed on the set of RGB pixel values. Varying intensities helped to reduce error by over **1 %**

$$I_{xy} = [I^R_{xy}, I^G_{xy}, I^B_{xy}]^T$$

$$[\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3][\alpha_1\lambda_1, \alpha_2\lambda_2, \alpha_3\lambda_3]^T$$

Reducing overfitting

Drop out 0.5 of neurons during training

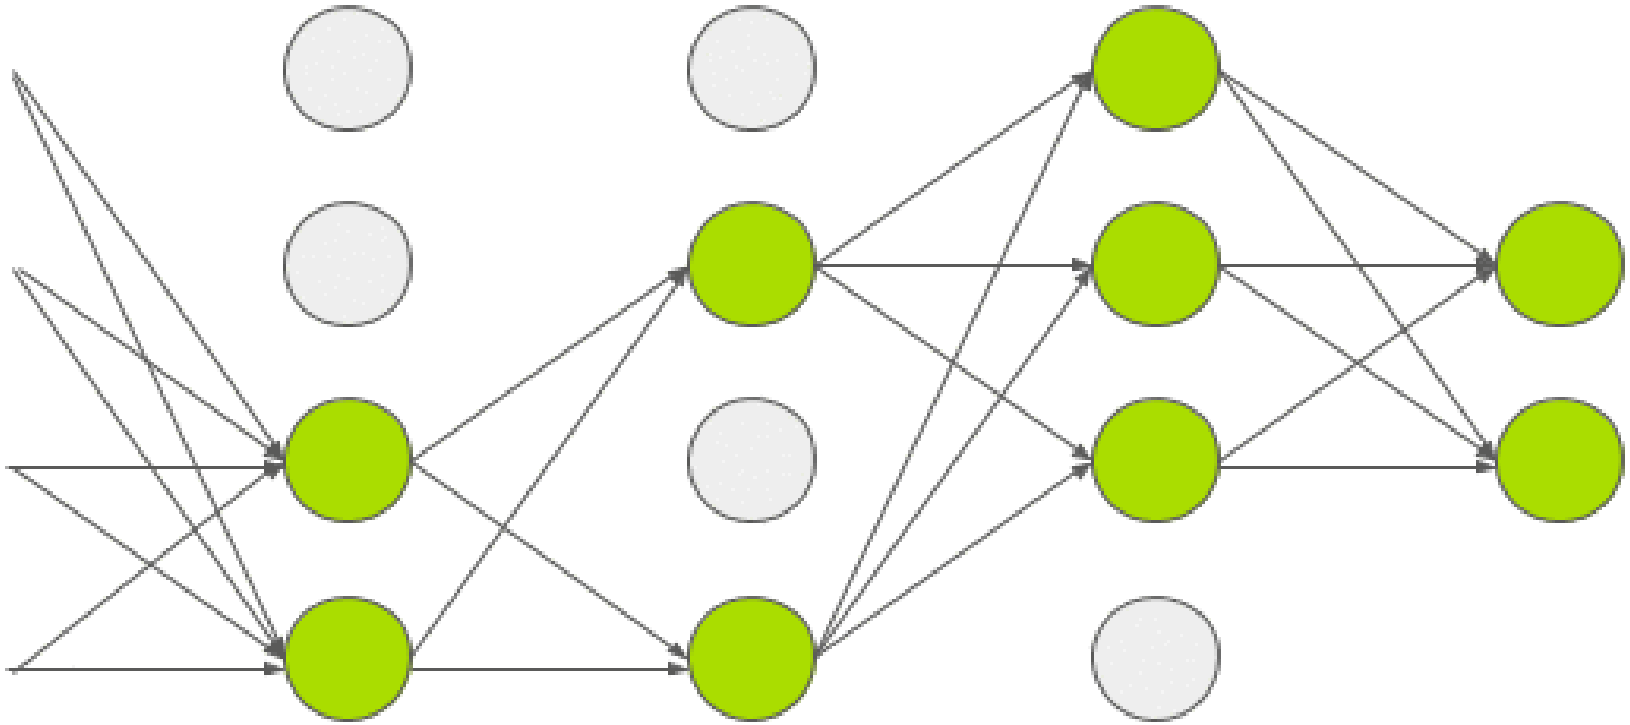

Used GPU to increase training speed

Double GPU reduced top-1 and top-5 error rates by **1.7%** and **1.2%**, respectively

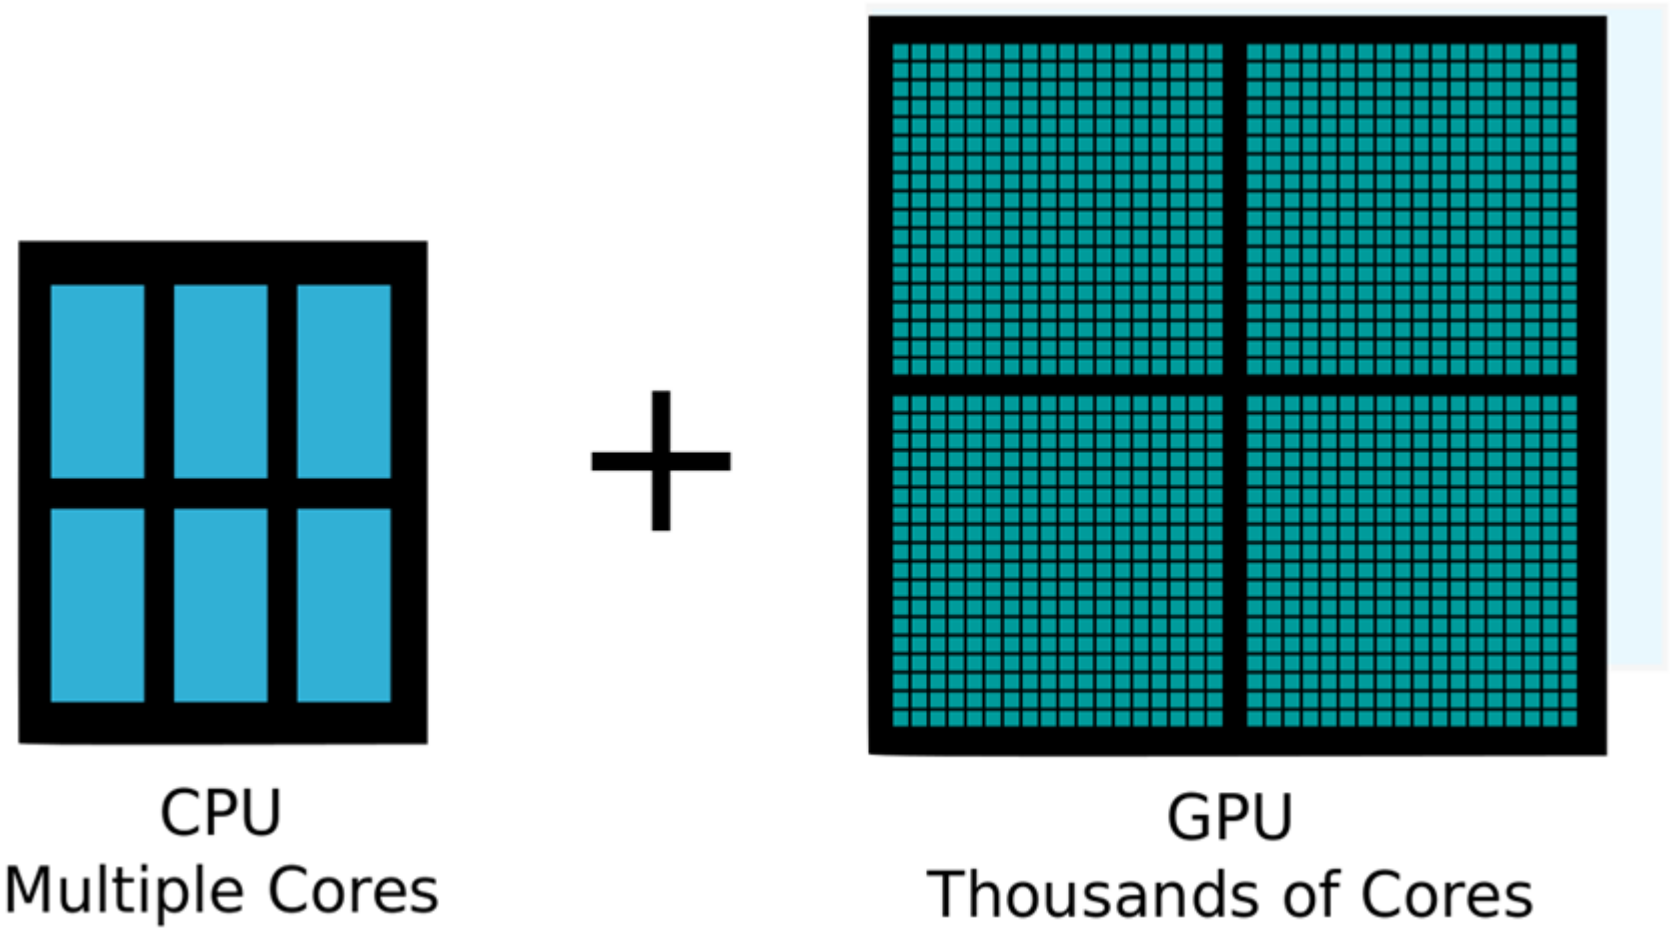

Outline

- Historical introduction
- Architecture of AlexNet
- Key characteristics of AlexNet
- **After AlexNet breakthrough**

Results

Growing Use of Deep Learning at Google

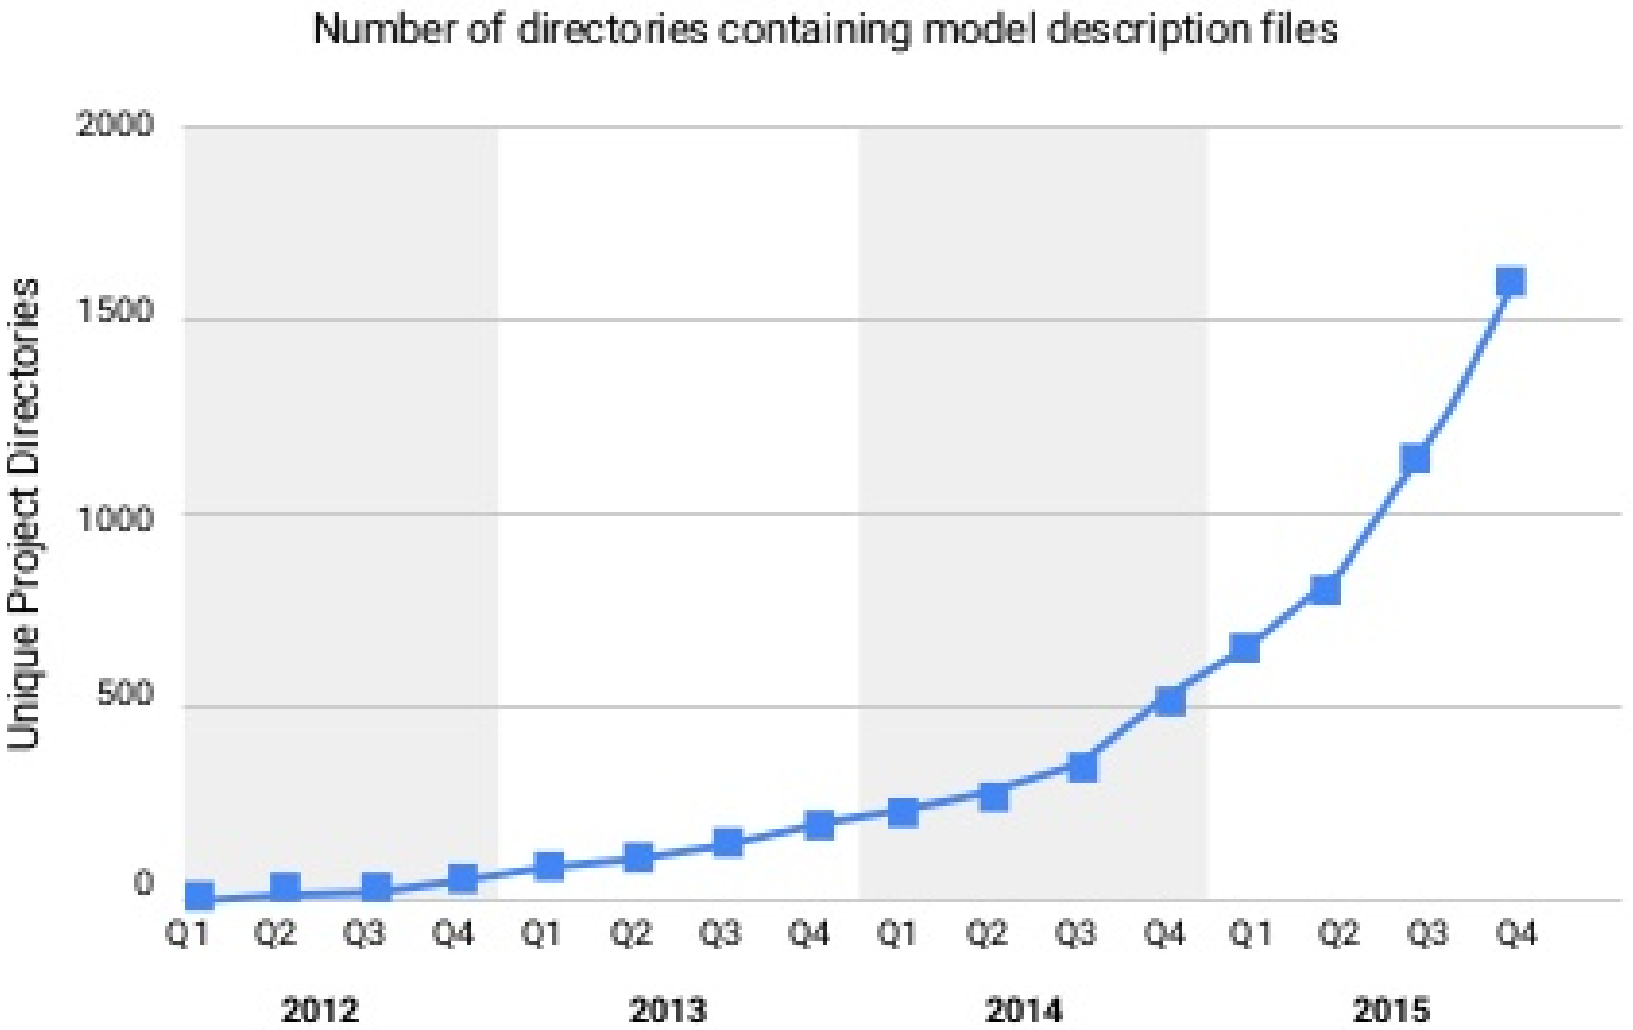

Across many products/areas

- Apps
- Maps
- Photos
- Gmail
- Speech
- Android
- YouTube
- Translation
- Robotics Research
- Image Understanding
- Natural Language Understanding
- Drug Discovery

Google

Results

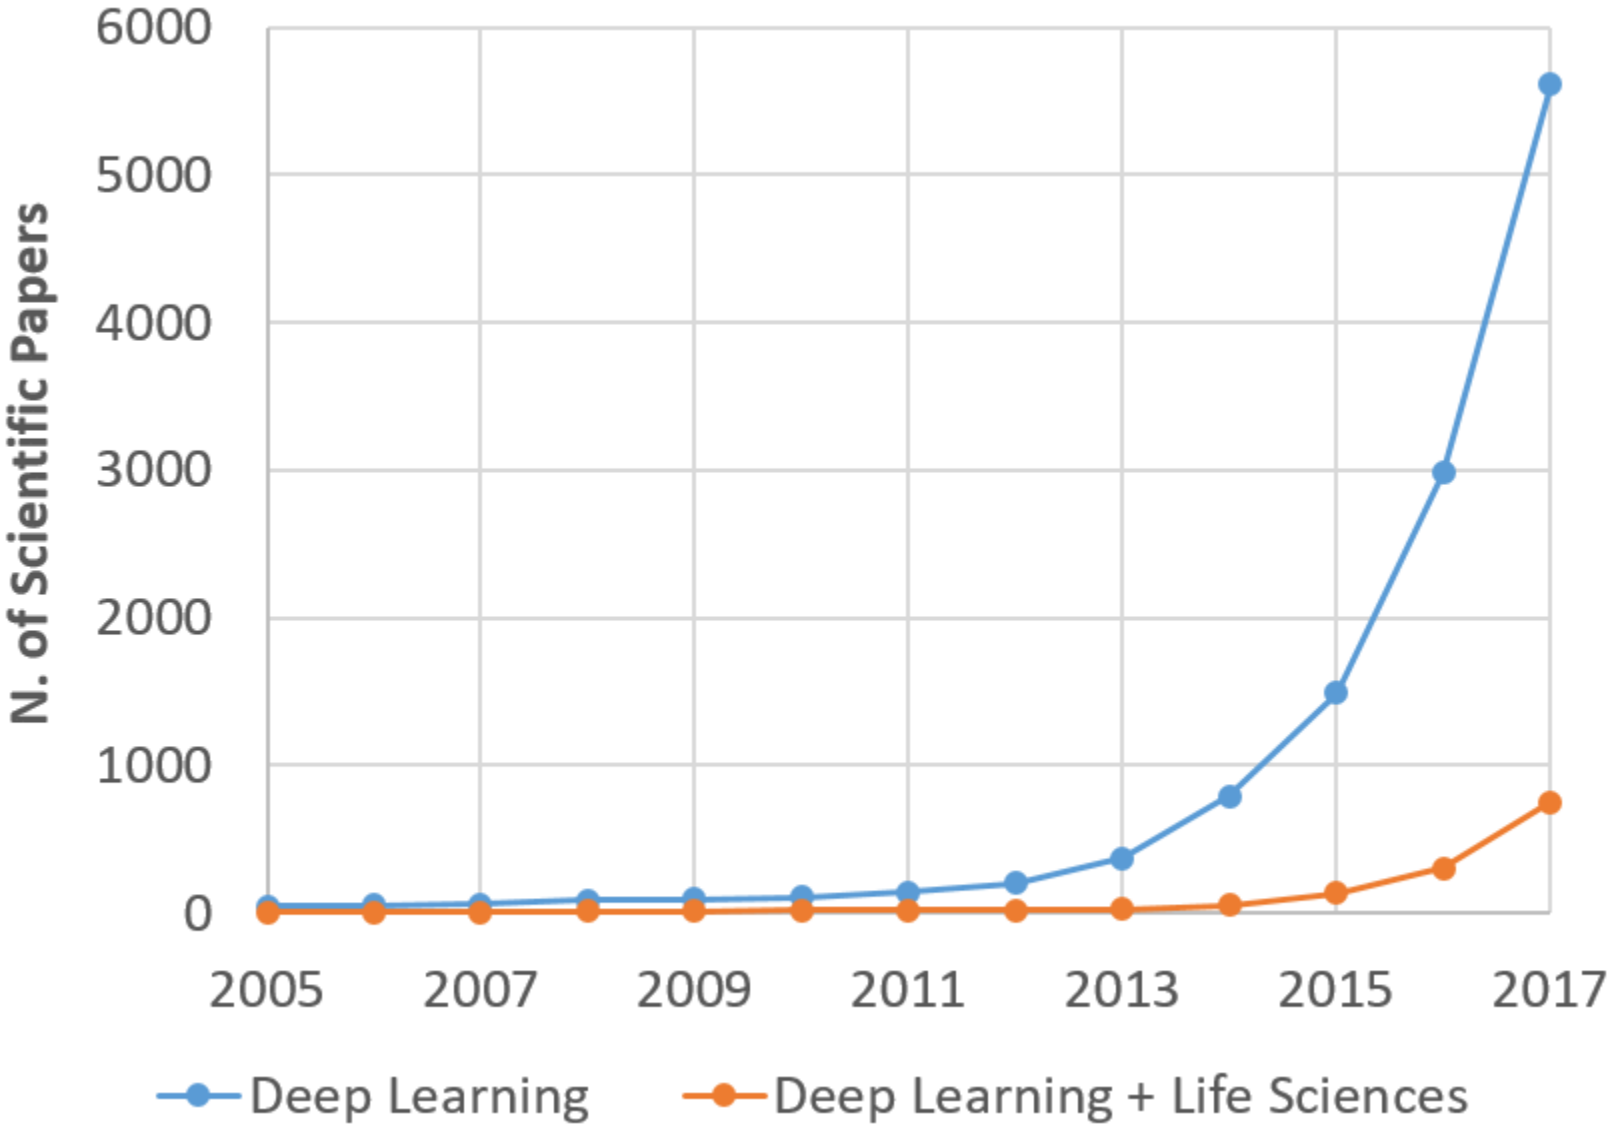
6000
5000
4000
3000
2000
1000
0
N. of Scientific Papers
2005
2007
2009
2011
2013
2015
2017
Deep Learning
Deep Learning + Life Sciences

References

- https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf
- https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html
- https://qz.com/1307091/the-inside-story-of-how-ai-got-good-enough-to-dominate-silicon-valley/
- https://medium.com/@smallfishbigsea/a-walk-through-of-alexnet-6cbd137a5637
- https://www.learnopencv.com/understanding-alexnet/