Introduction to Deep Learning with Python

Samar Haider

University of Engineering and Technology, Lahore

PyCon Pakistan

16/12/2017

PyCon Pakistan 2017

Artificial intelligence, machine learning & deep learning

The AI universe

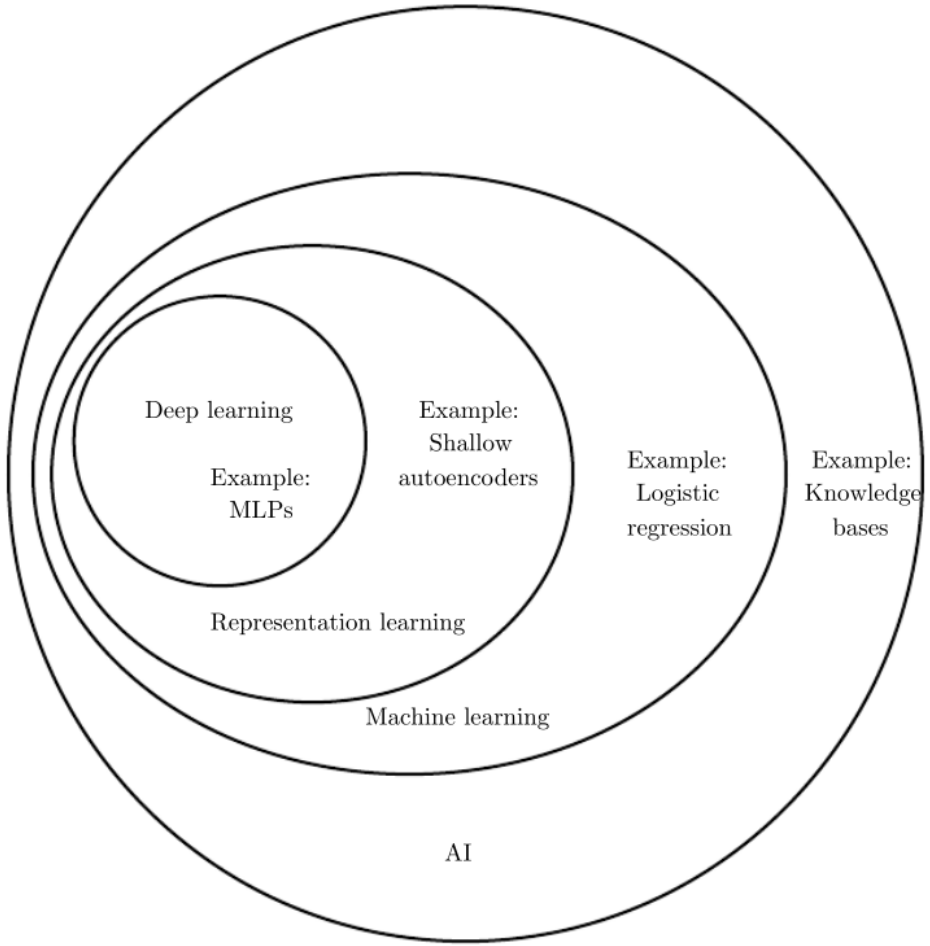

Goodfellow et. al., “Deep Learning.” *MIT Press (2016)*

The AI universe

“What's the difference between AI and ML?”

“It's AI when you're raising money, it's ML when you're trying to hire people.”

Machine learning

Data

Machine learning

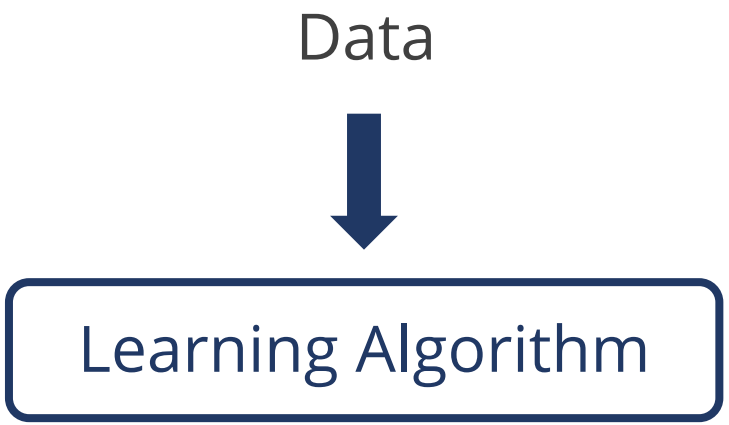

Machine learning

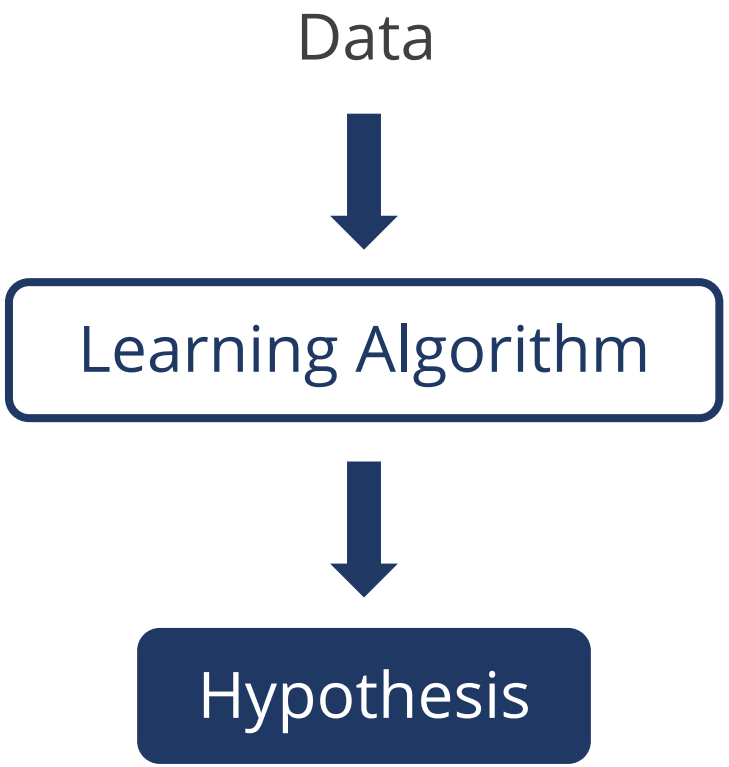

Machine learning

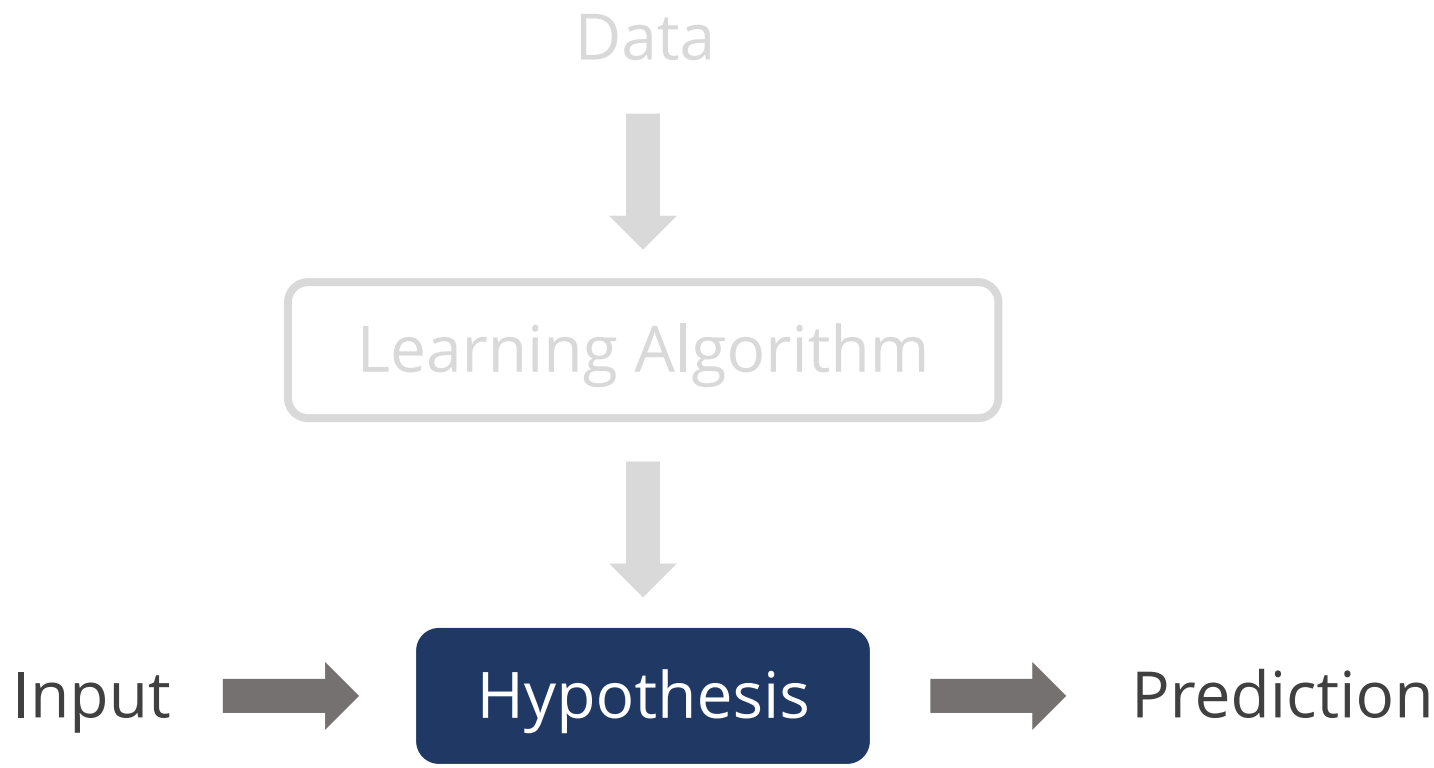

Machine learning

Deep learning

The rise of deep learning

The rise of deep learning

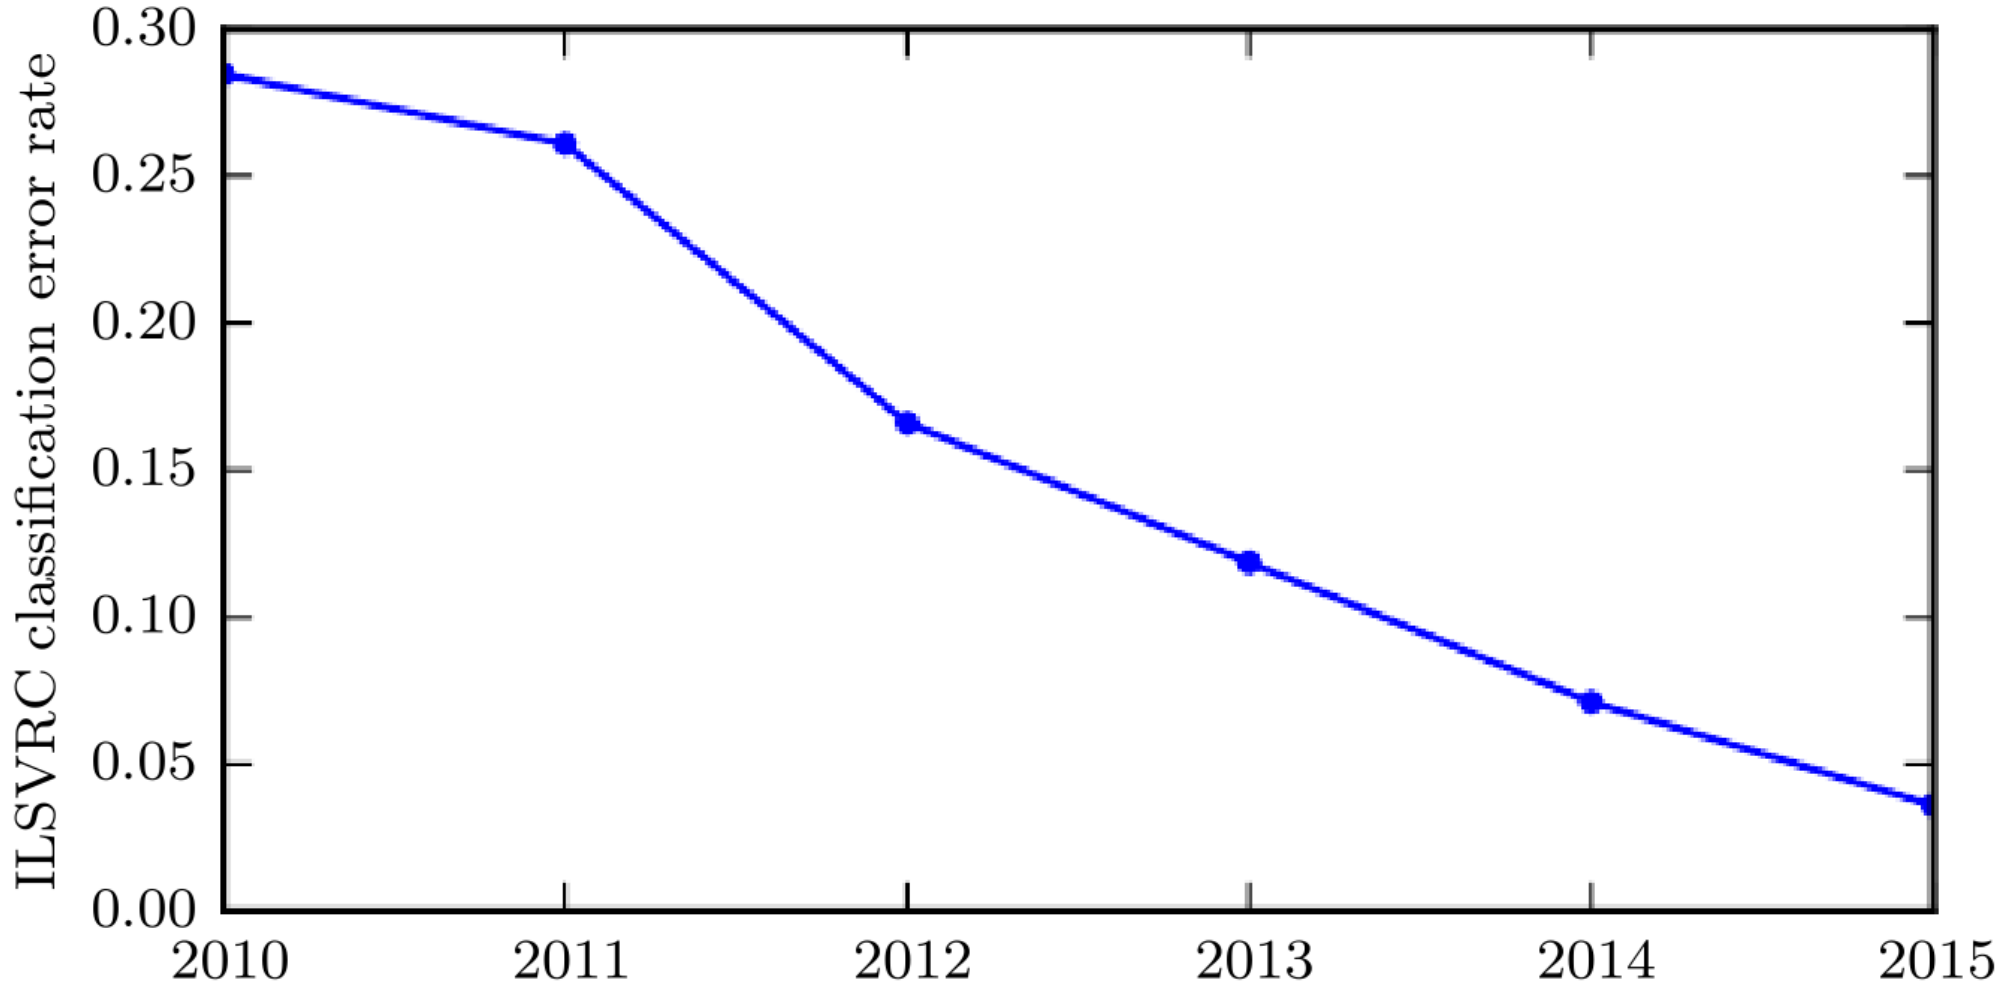

Goodfellow et. al., “Deep Learning.” *MIT Press (2016)*

The rise of deep learning

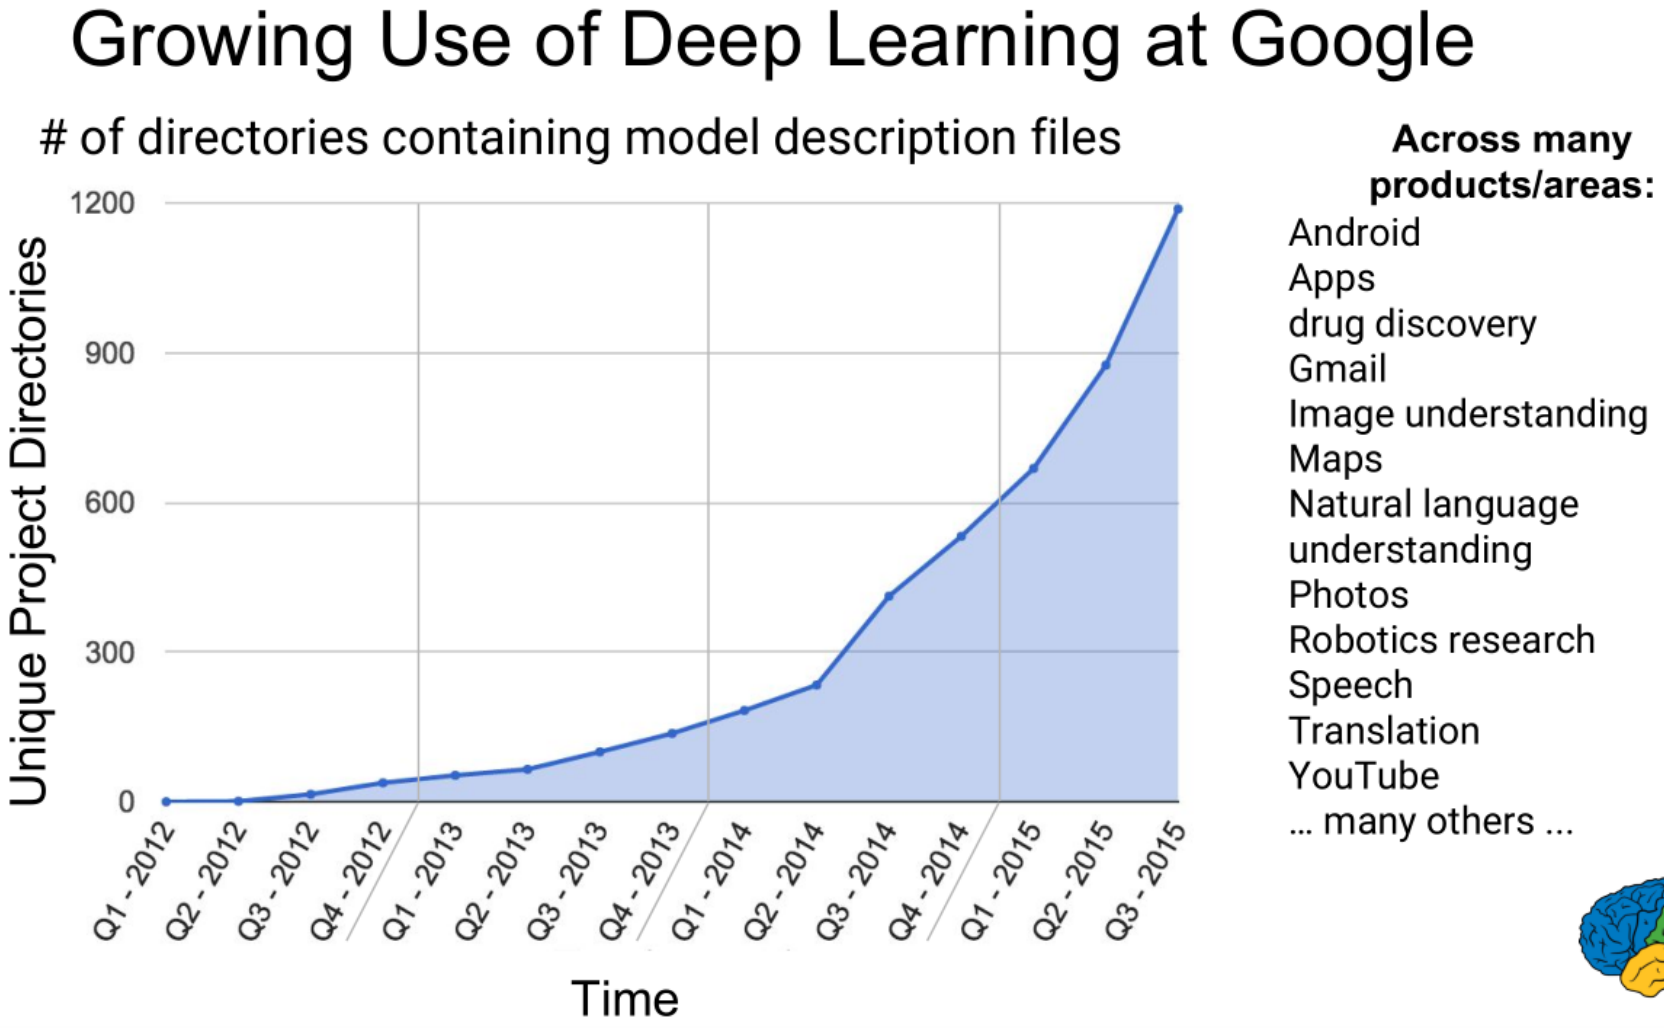

Across many products/areas:

Android
Apps
drug discovery
Gmail
Image understanding
Maps
Natural language understanding
Photos
Robotics research
Speech
Translation
YouTube
... many others ...

Dean, "Large-Scale Deep Learning for Intelligent Computer Systems." *WSDM (2016)*

The rise of deep learning

1. Bigger datasets
2. More computational power
3. Improvements in algorithms (due to 1 and 2)

The rise of deep learning

"If big data is the new oil, deep learning is the new internal combustion engine."

– Yann LeCun

(Director, Facebook AI Research)

The rise of deep learning

"AI is the new electricity: Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years."

– Andrew Ng

(Founder, deeplearning.ai)

Neural networks

Biological neuron

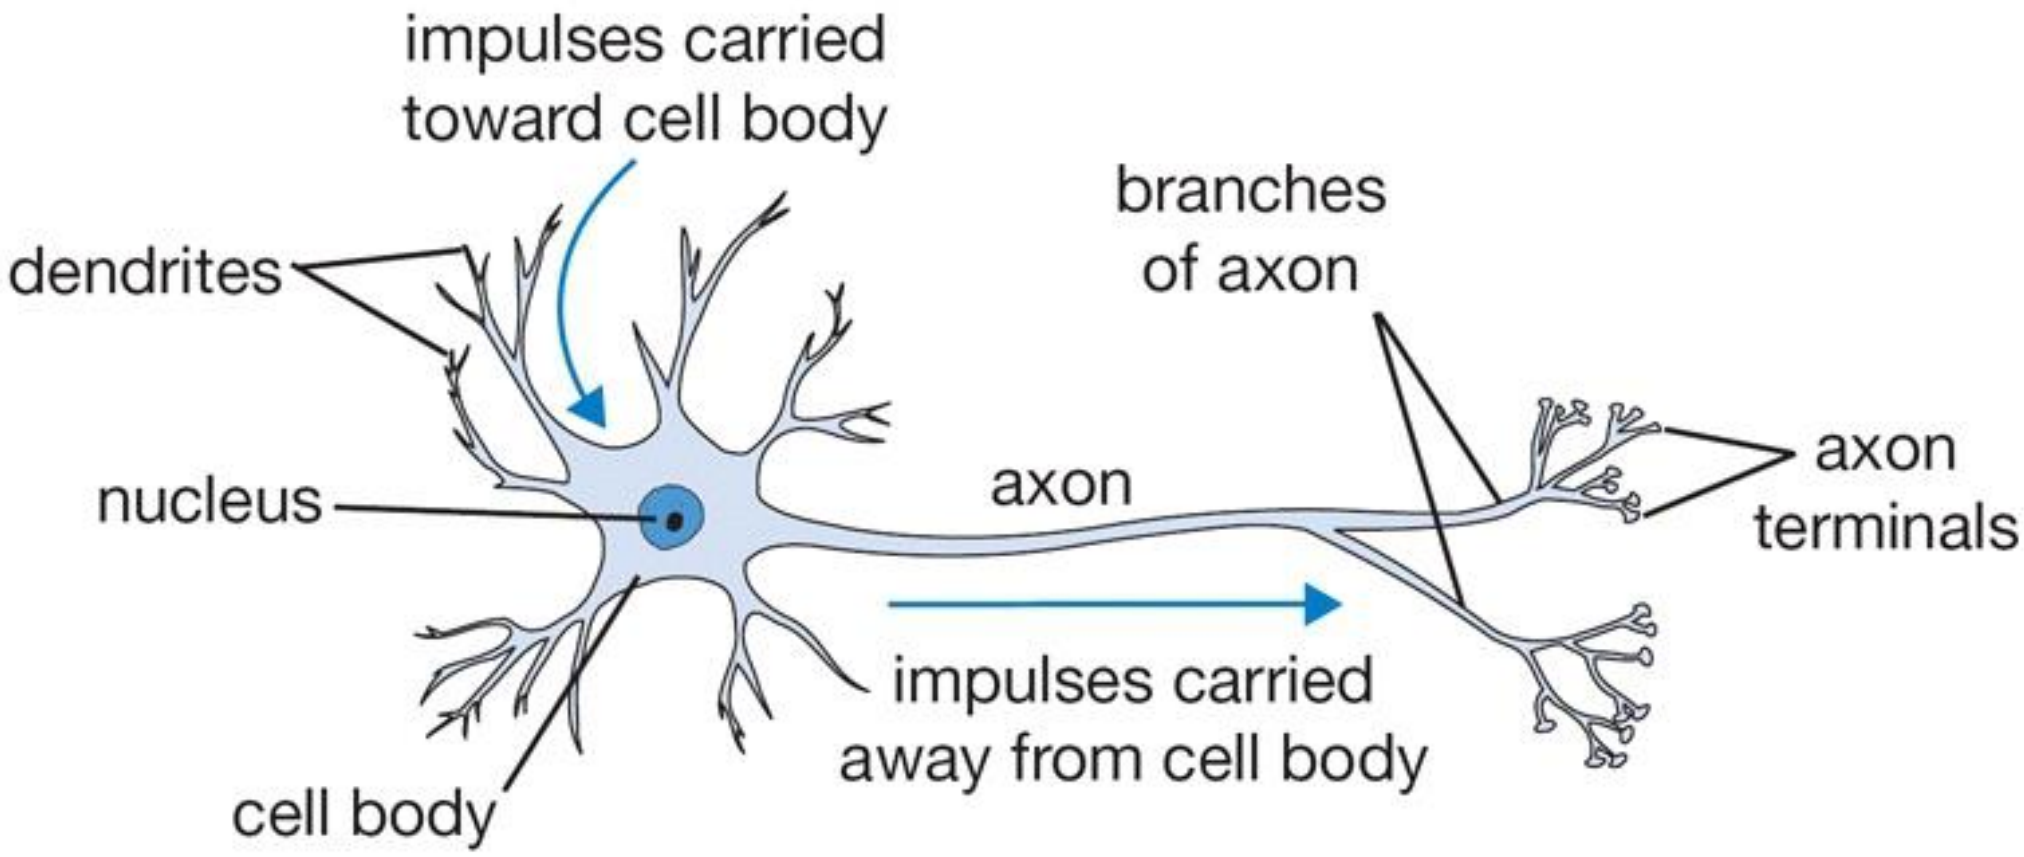

Artificial neuron

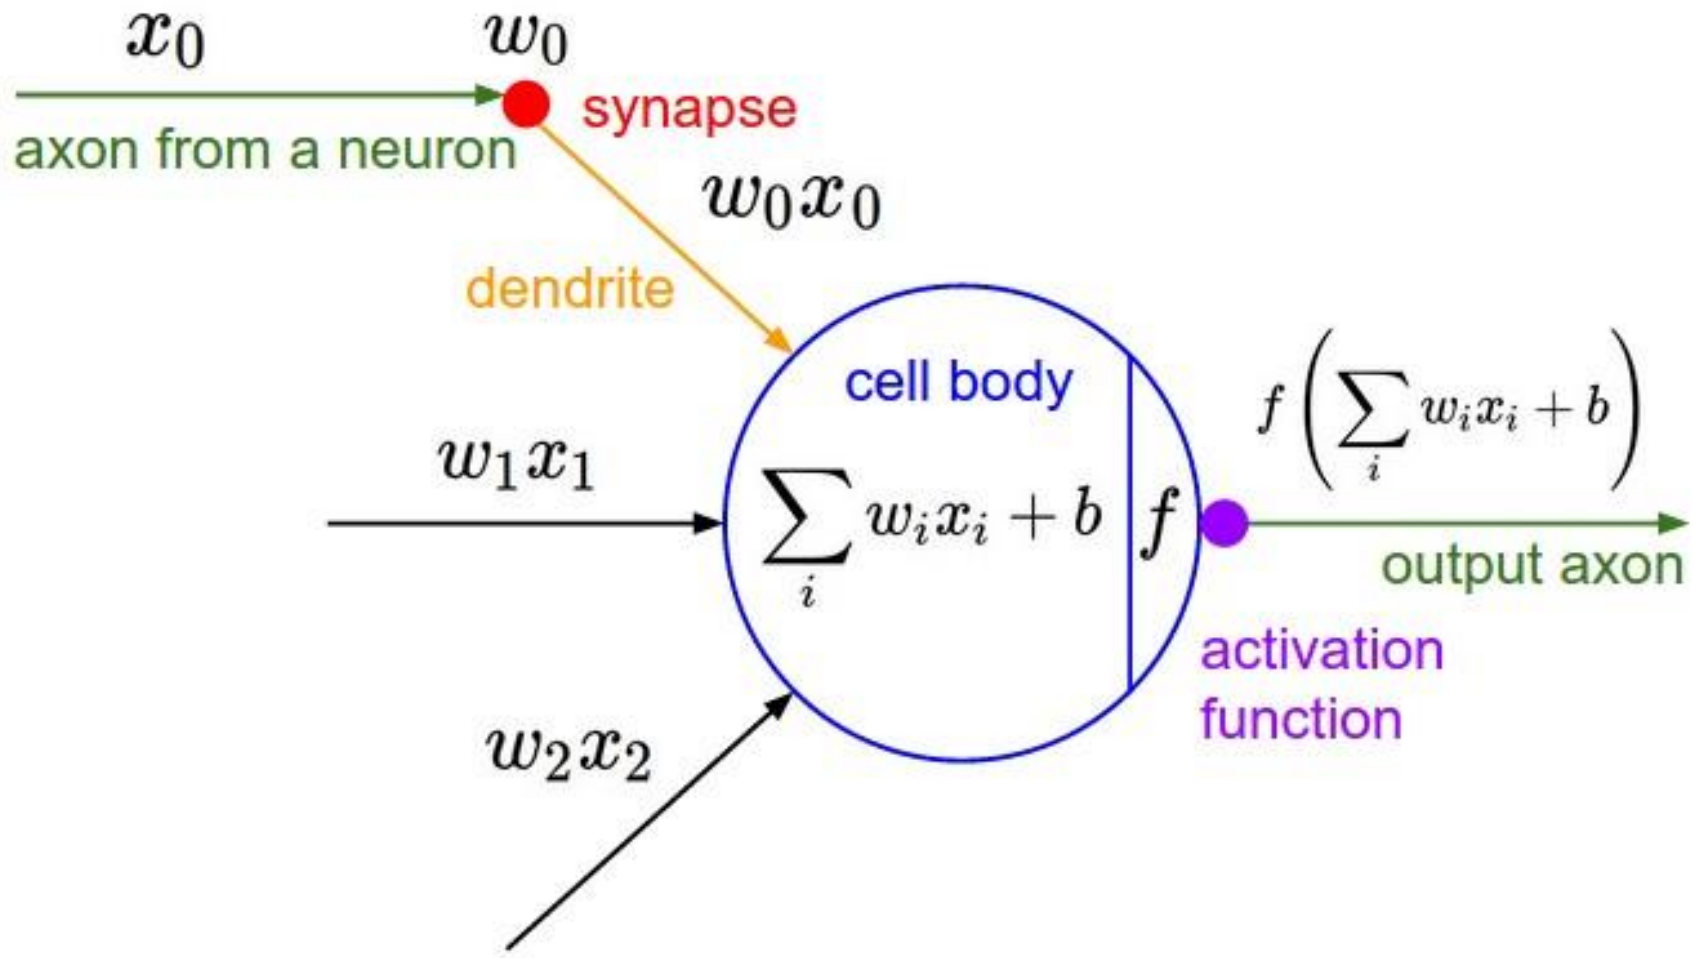

Biological vs artificial neuron

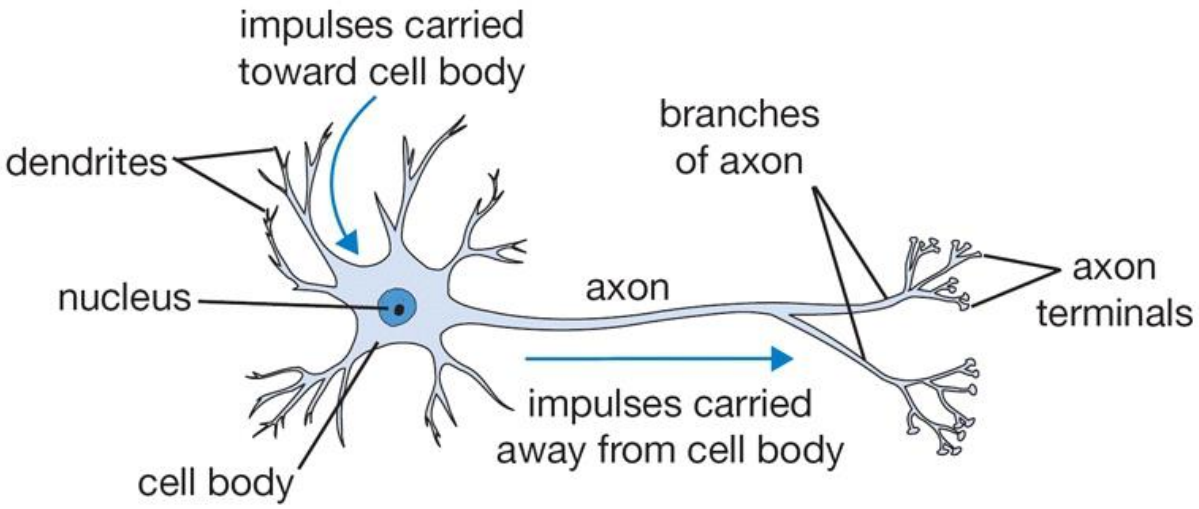

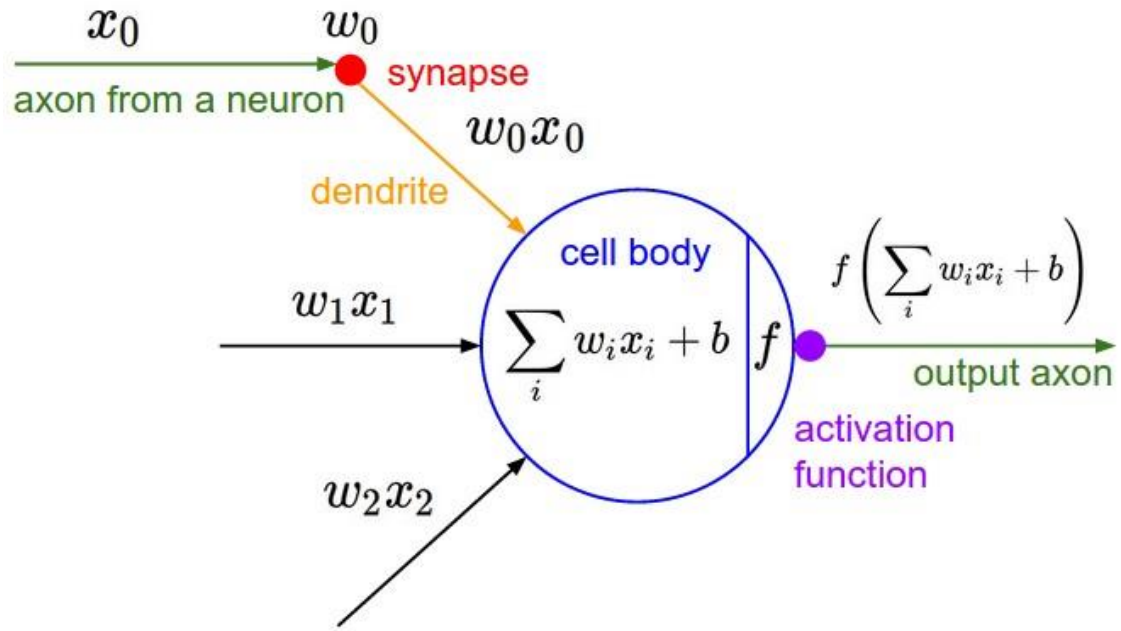

Activation functions

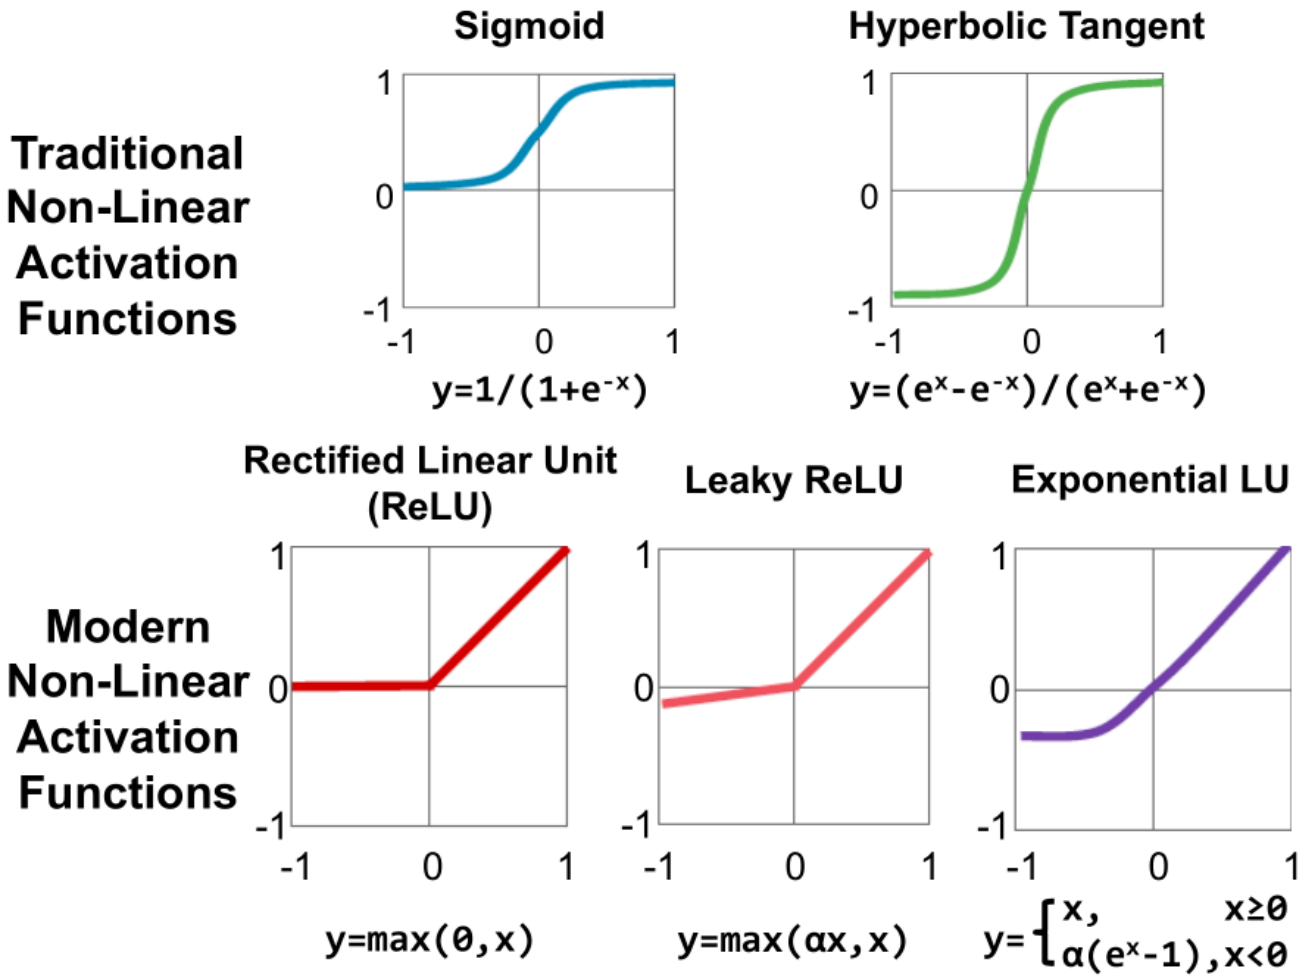

Sze et. al., "Efficient Processing of Deep Neural Networks: A Tutorial and Survey." *arXiv (2017)*

Activation functions

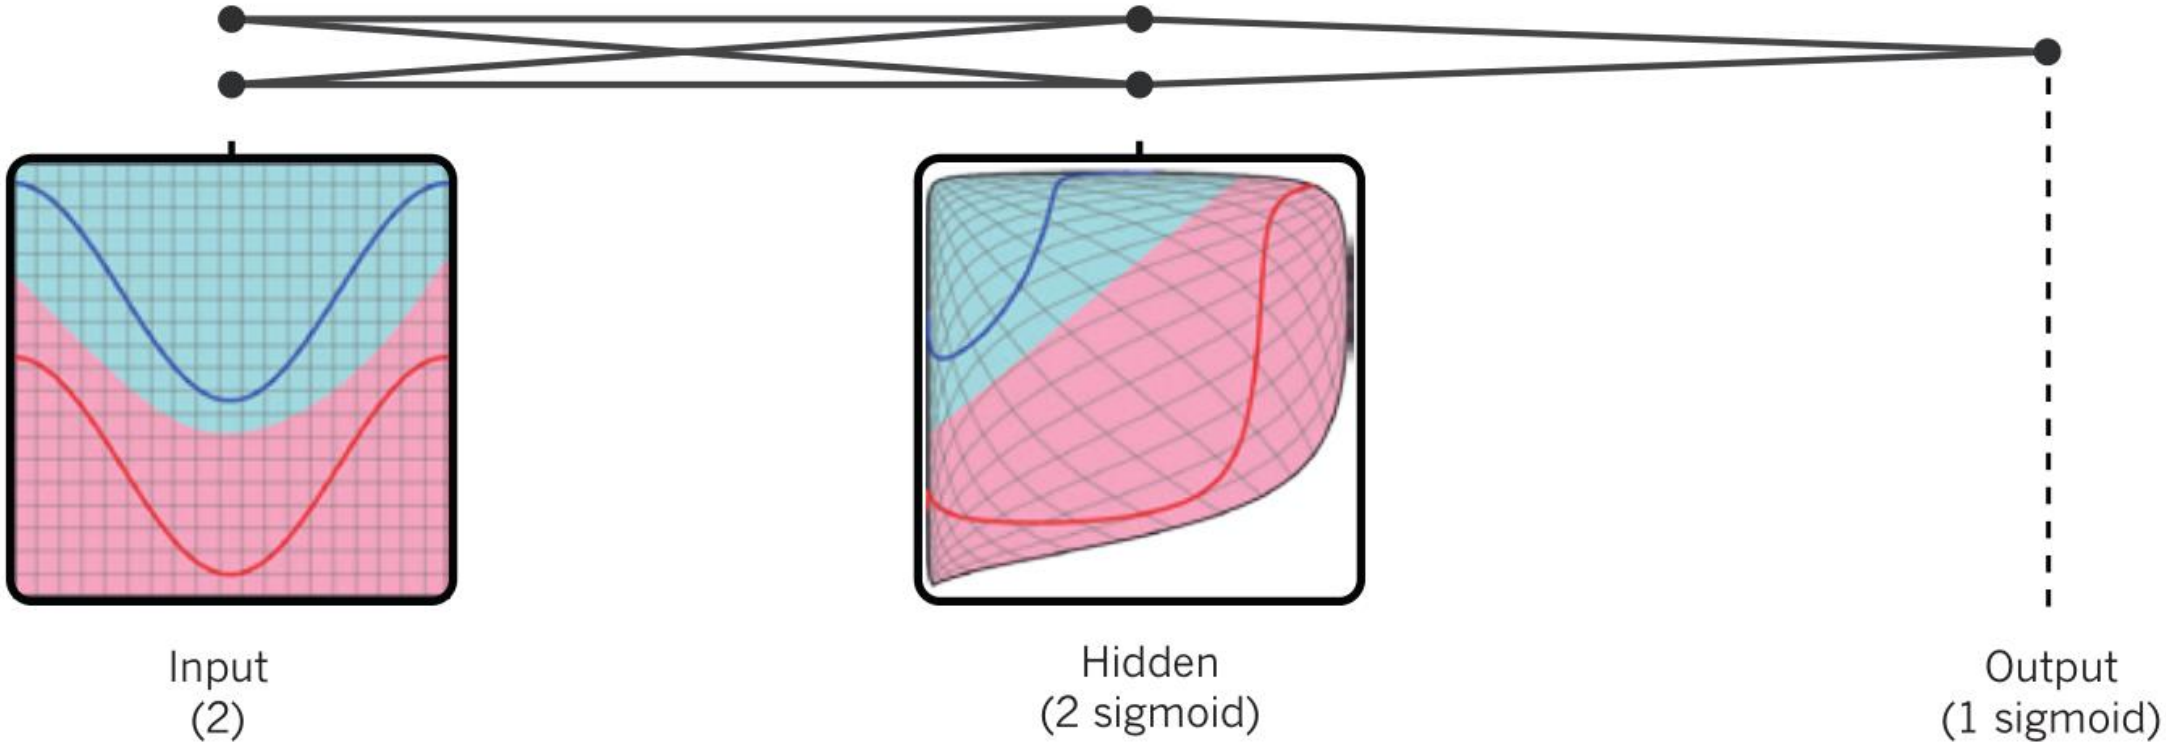

LeCun et. al., “Deep Learning.” *Nature (2015)*

A shallow neural network

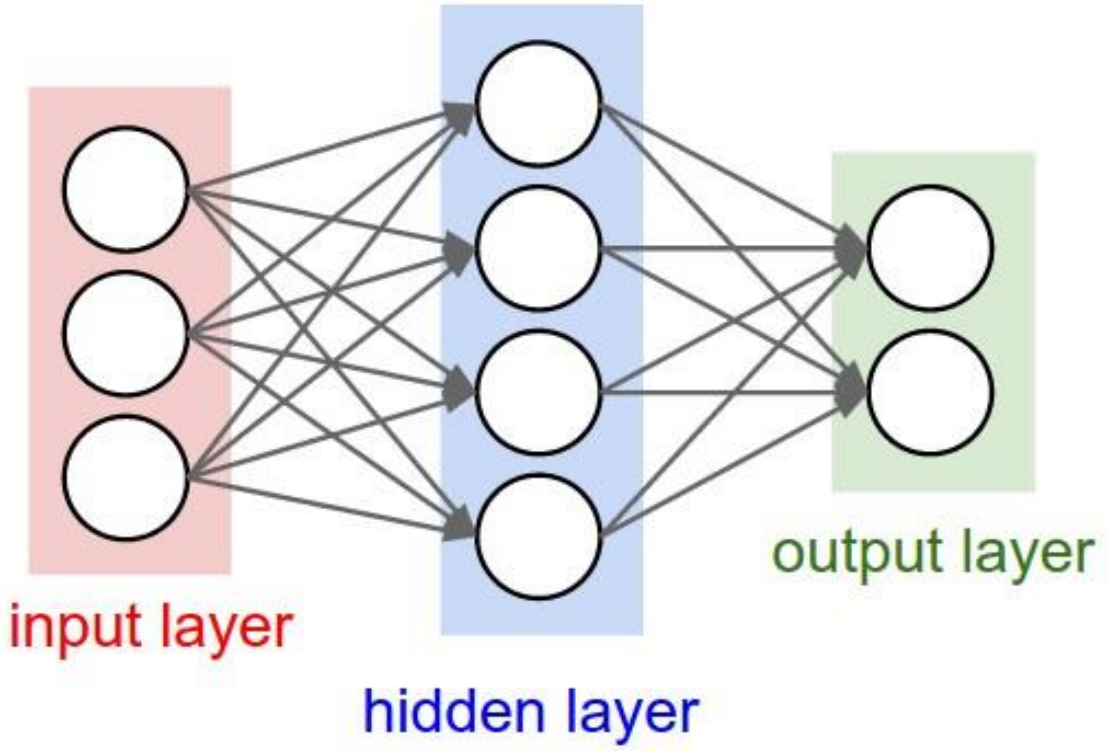

A deep neural network

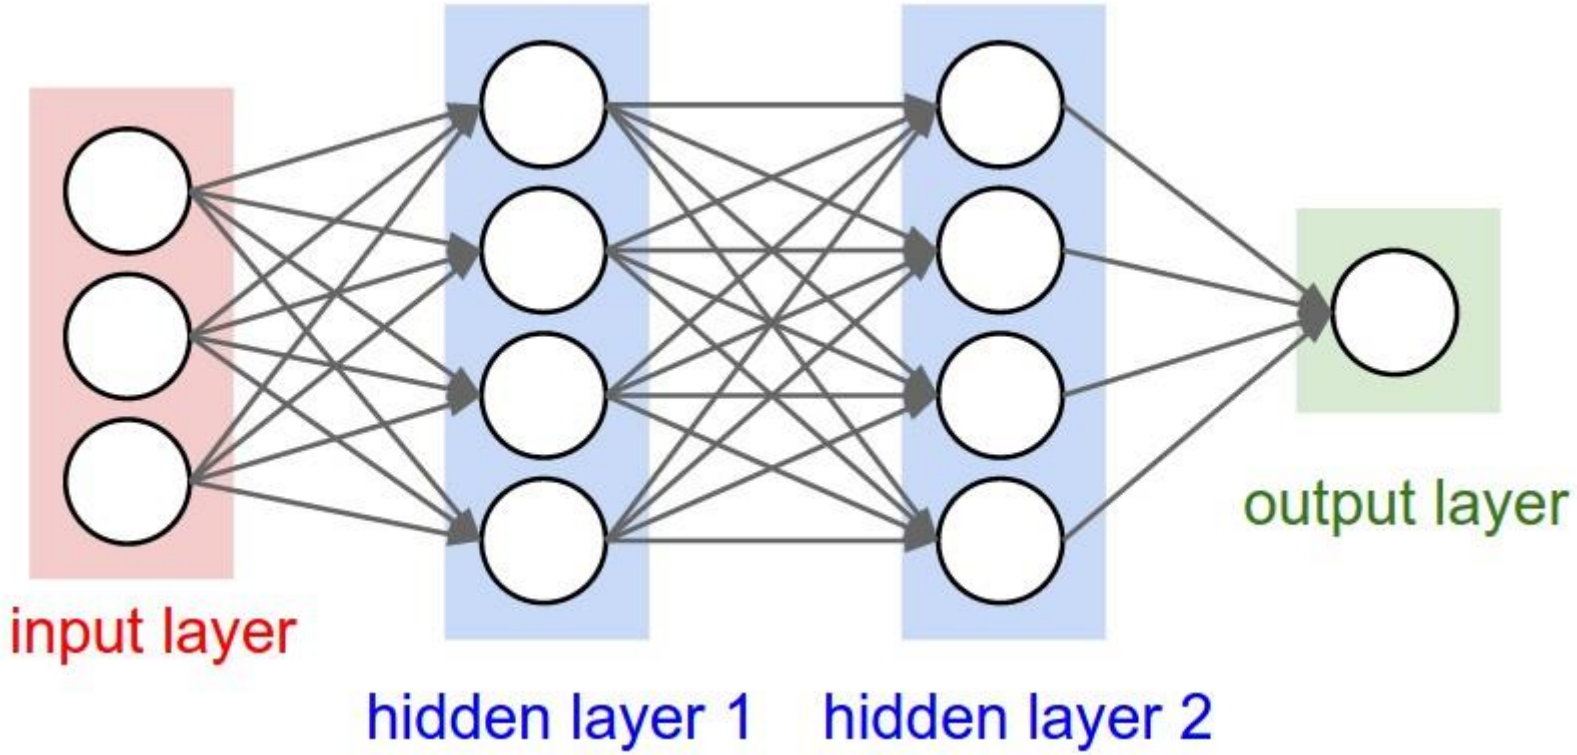

What we want

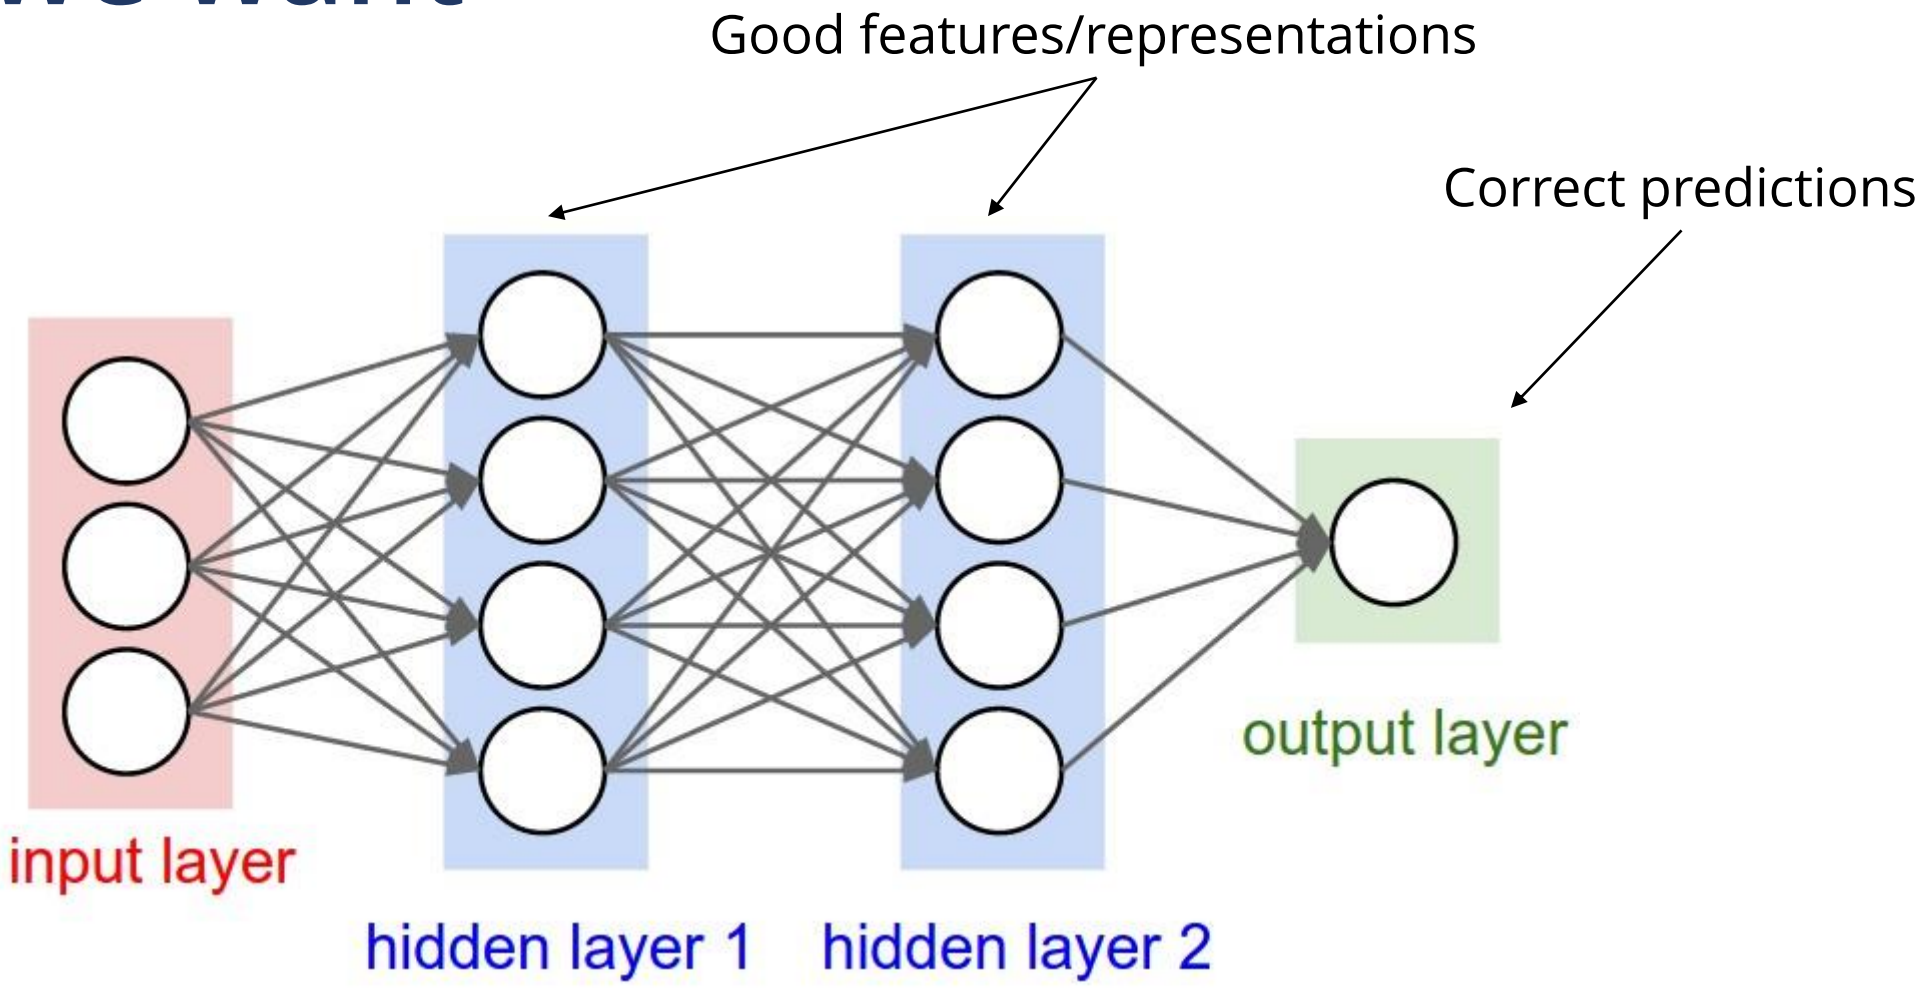
Good features/representations
Correct predictions
input layer
hidden layer 1
hidden layer 2
output layer

What we want

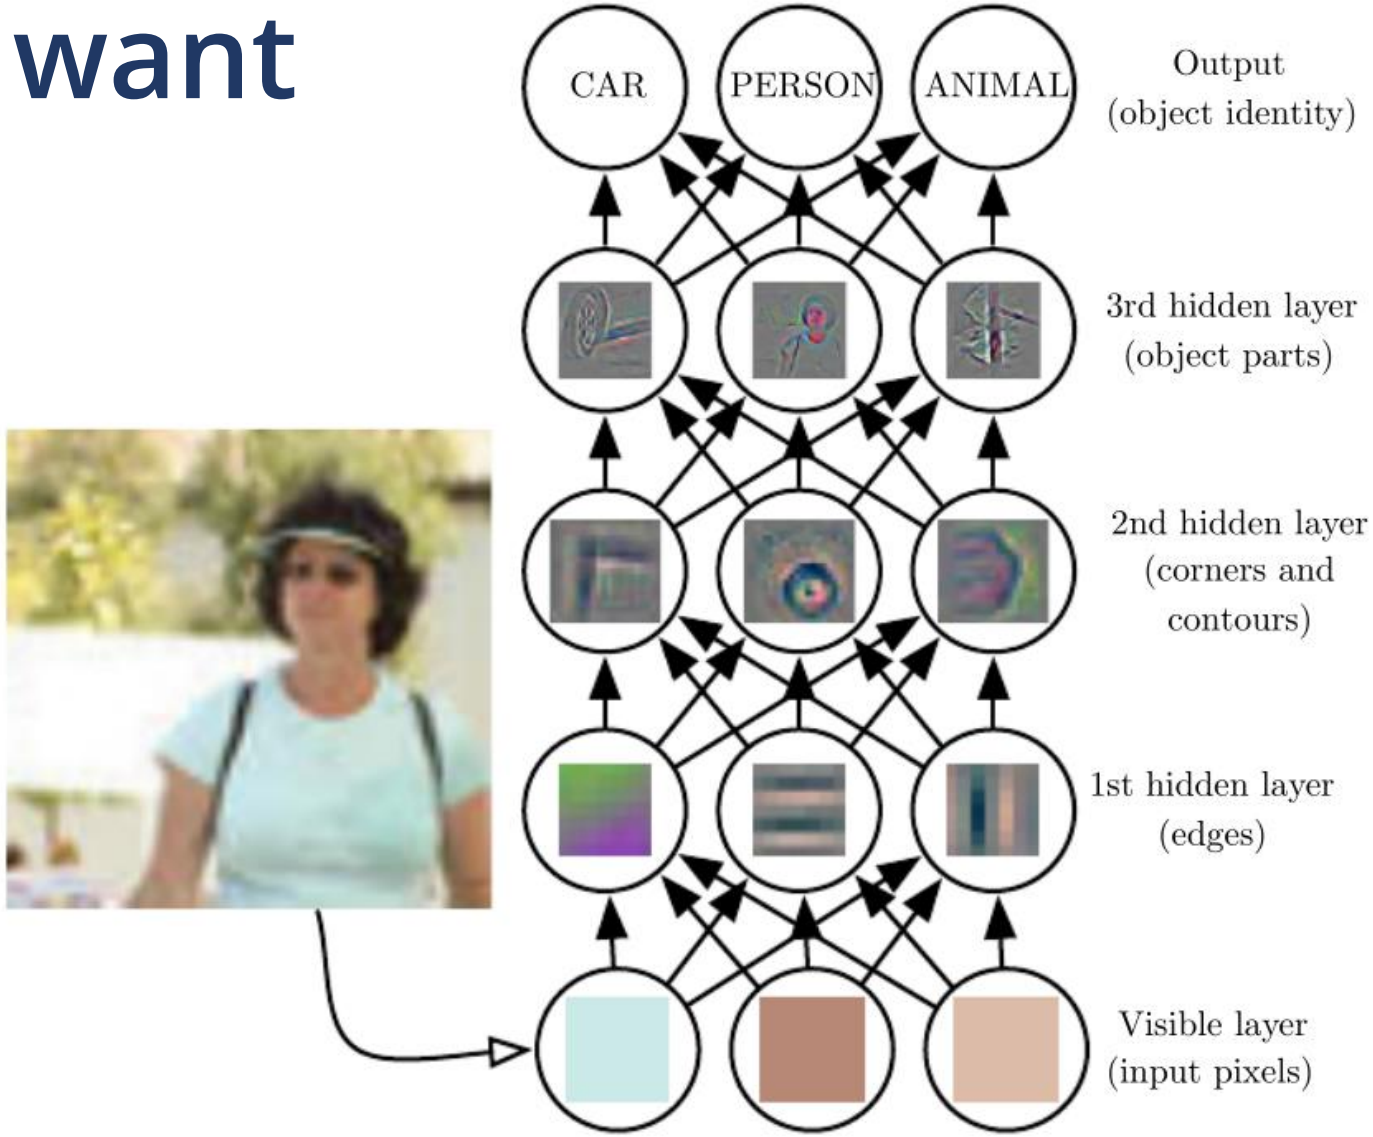

Goodfellow et. al., “Deep Learning.” *MIT Press (2016)*

The need for depth

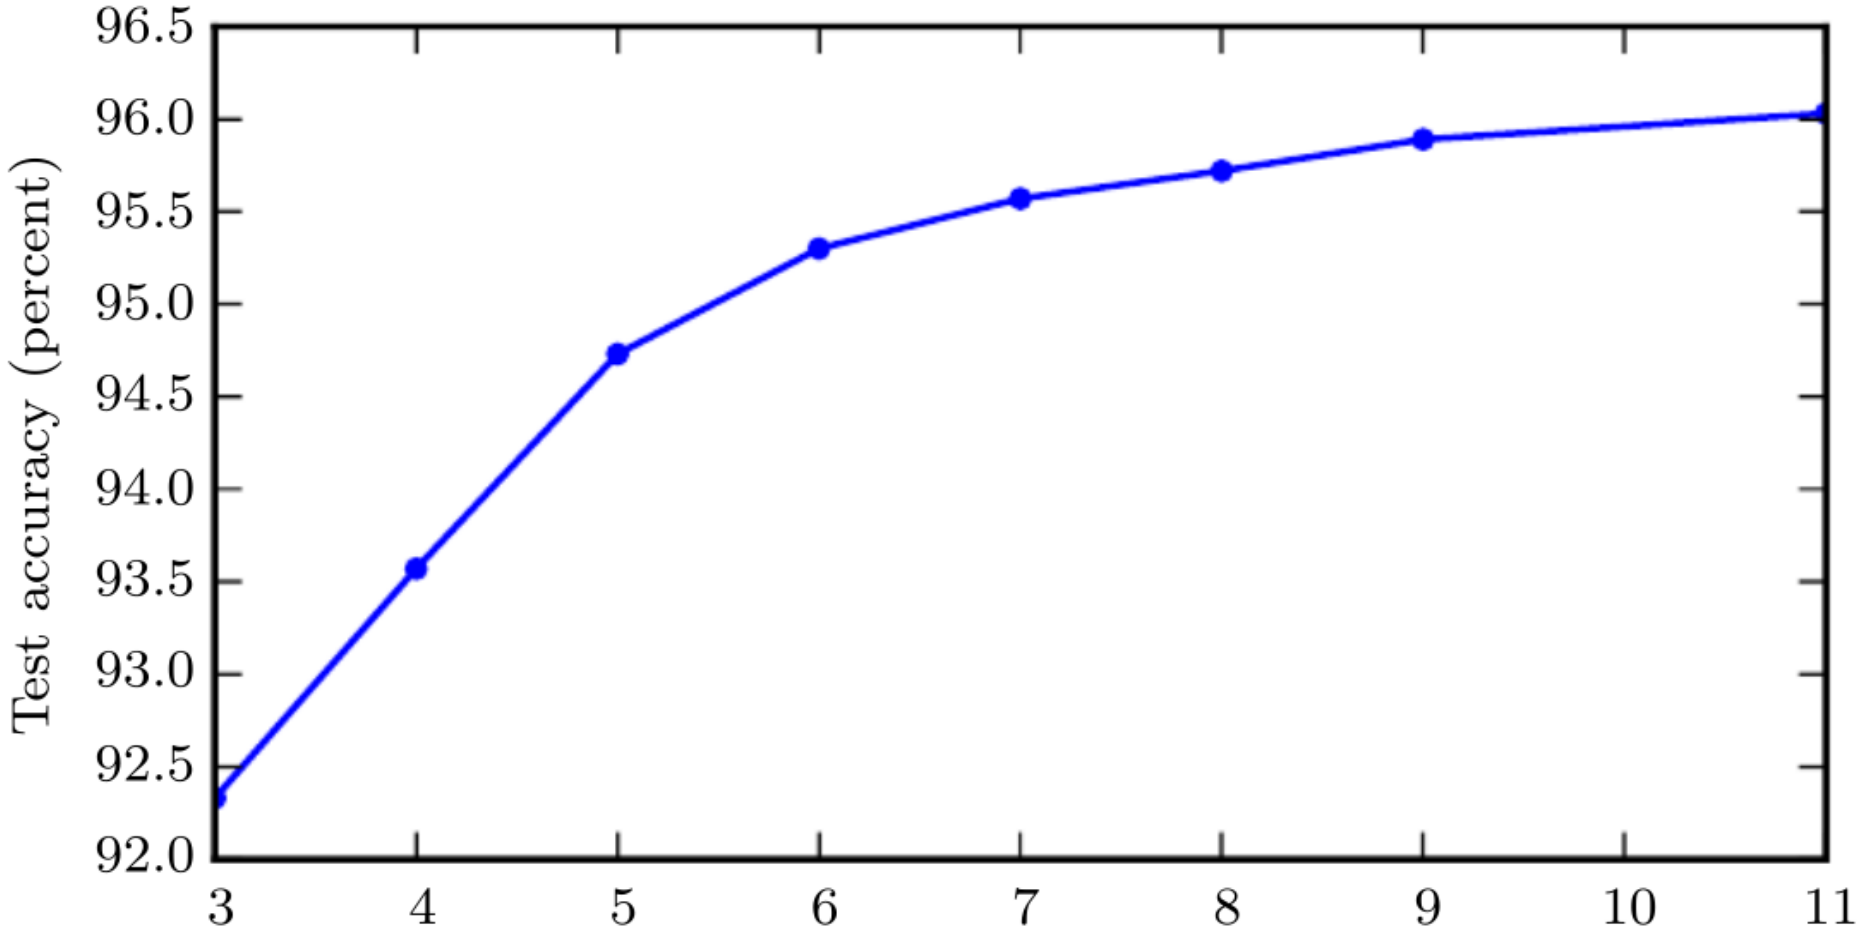

Goodfellow et. al., “Deep Learning.” *MIT Press (2016)*

… and even more depth

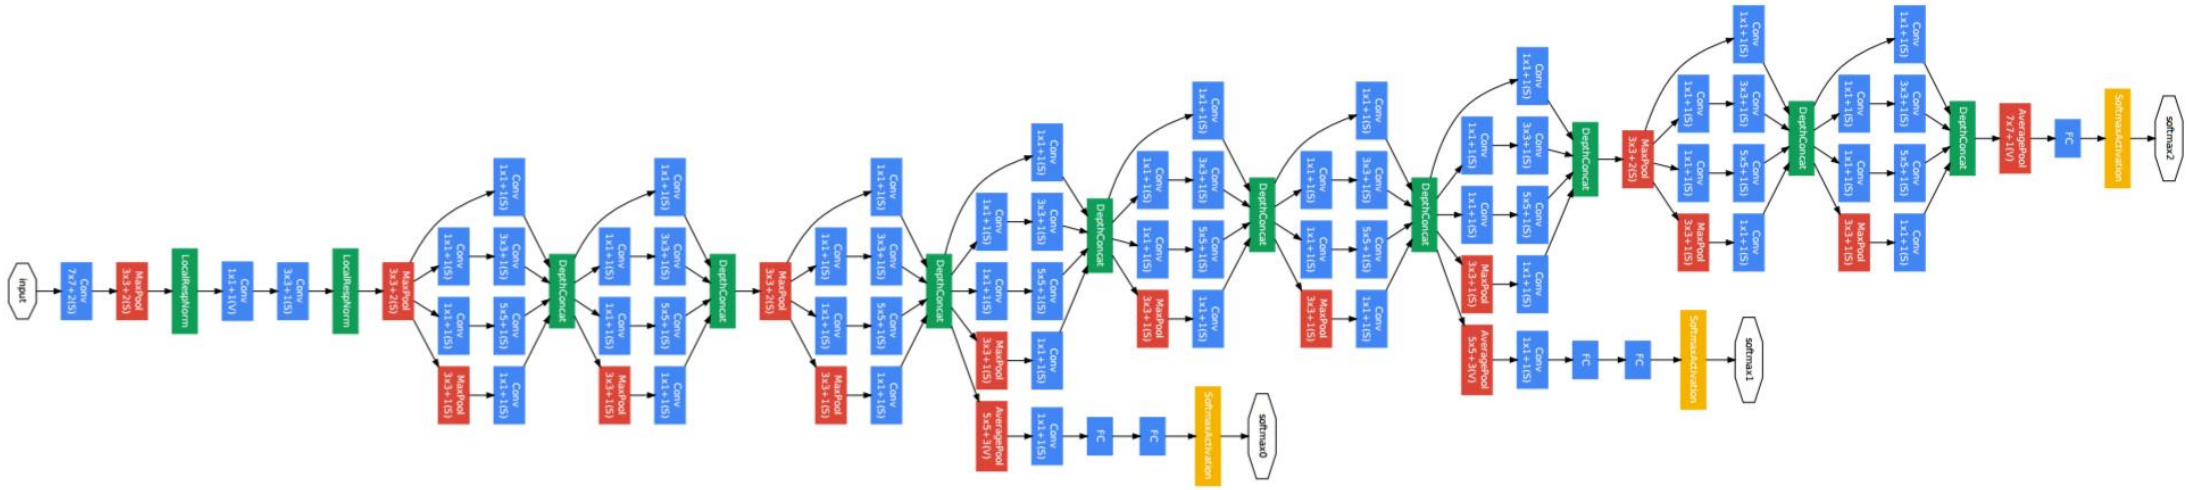

Szegedy et. al., “Going Deeper with Convolutions.” *CVPR (2015)*

Specialized architectures

Vision: Convolutional Neural Networks

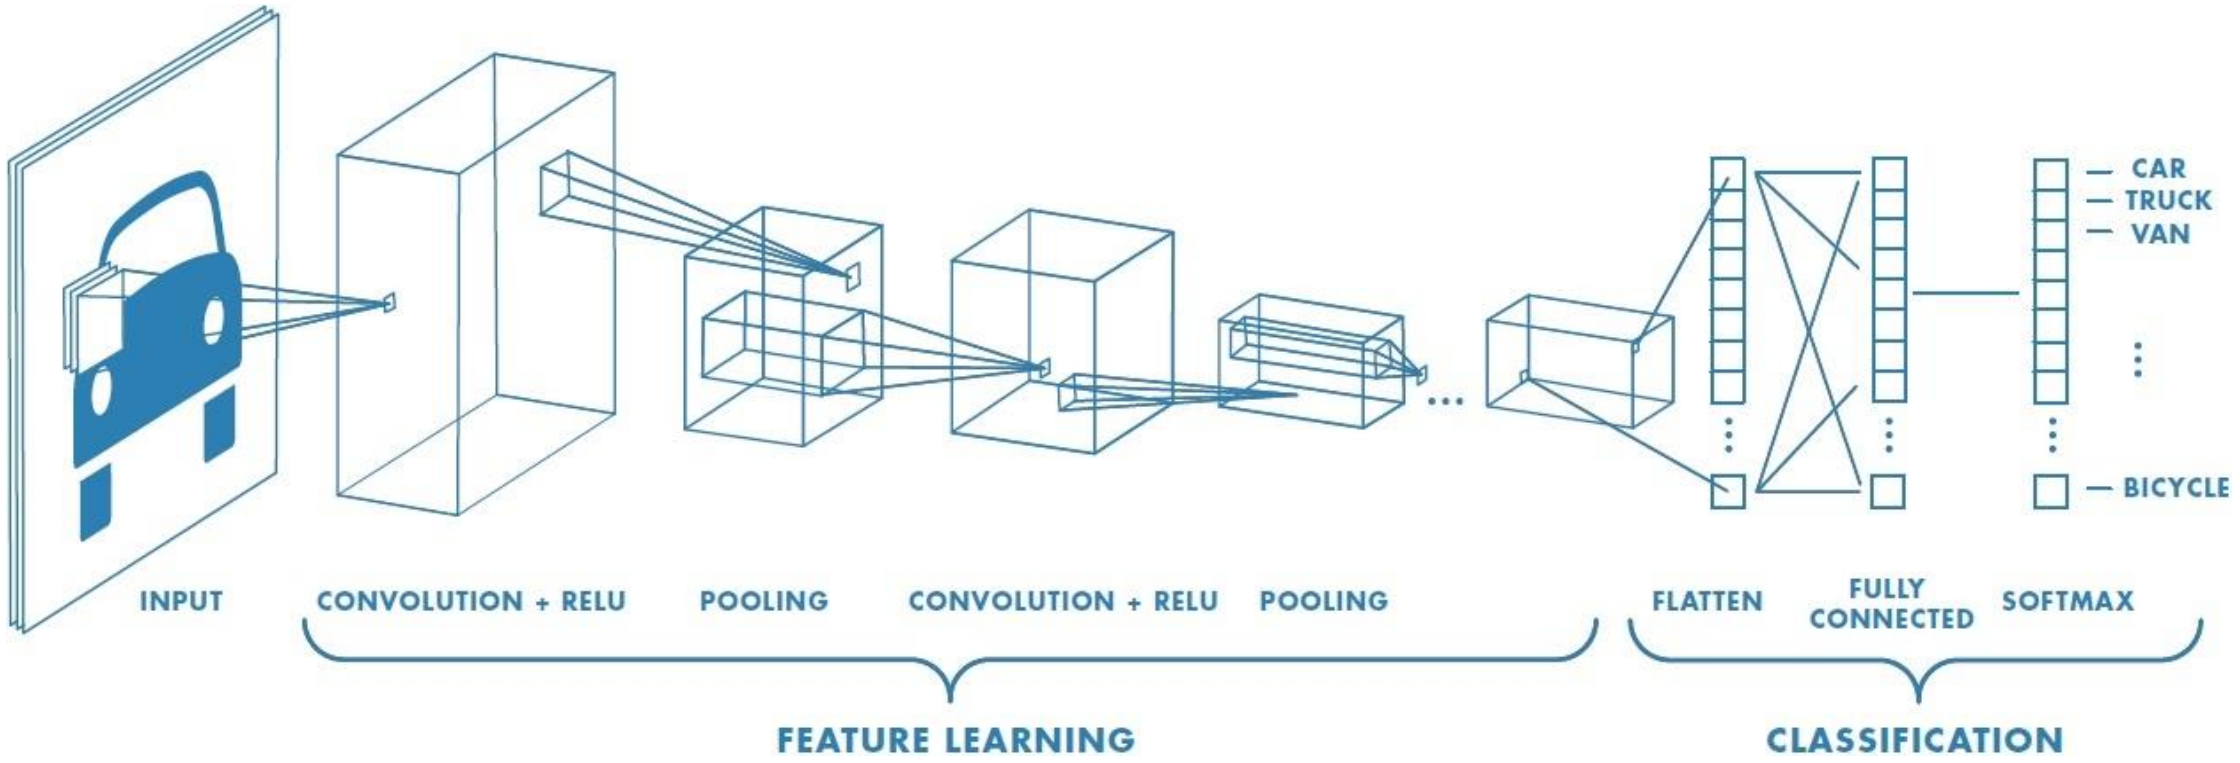

Specialized architectures

Language: Recurrent Neural Networks

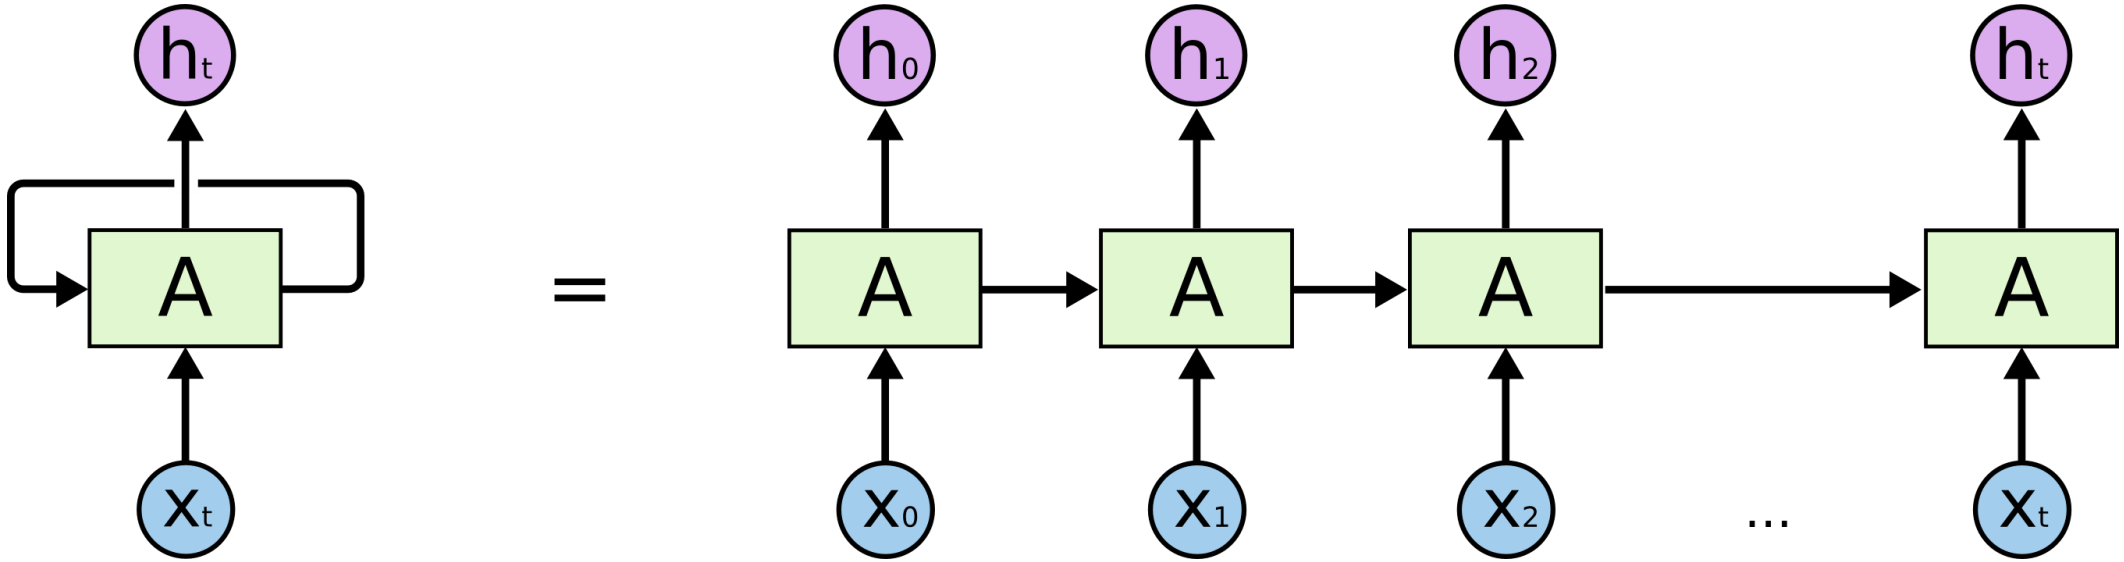

Deep learning

The learning process

1. Pick a training example
2. Make a prediction for it
3. Compare your prediction with the truth (= error)
4. Modify your weights in order to minimize this error
5. Repeat until convergence

Minimizing the error

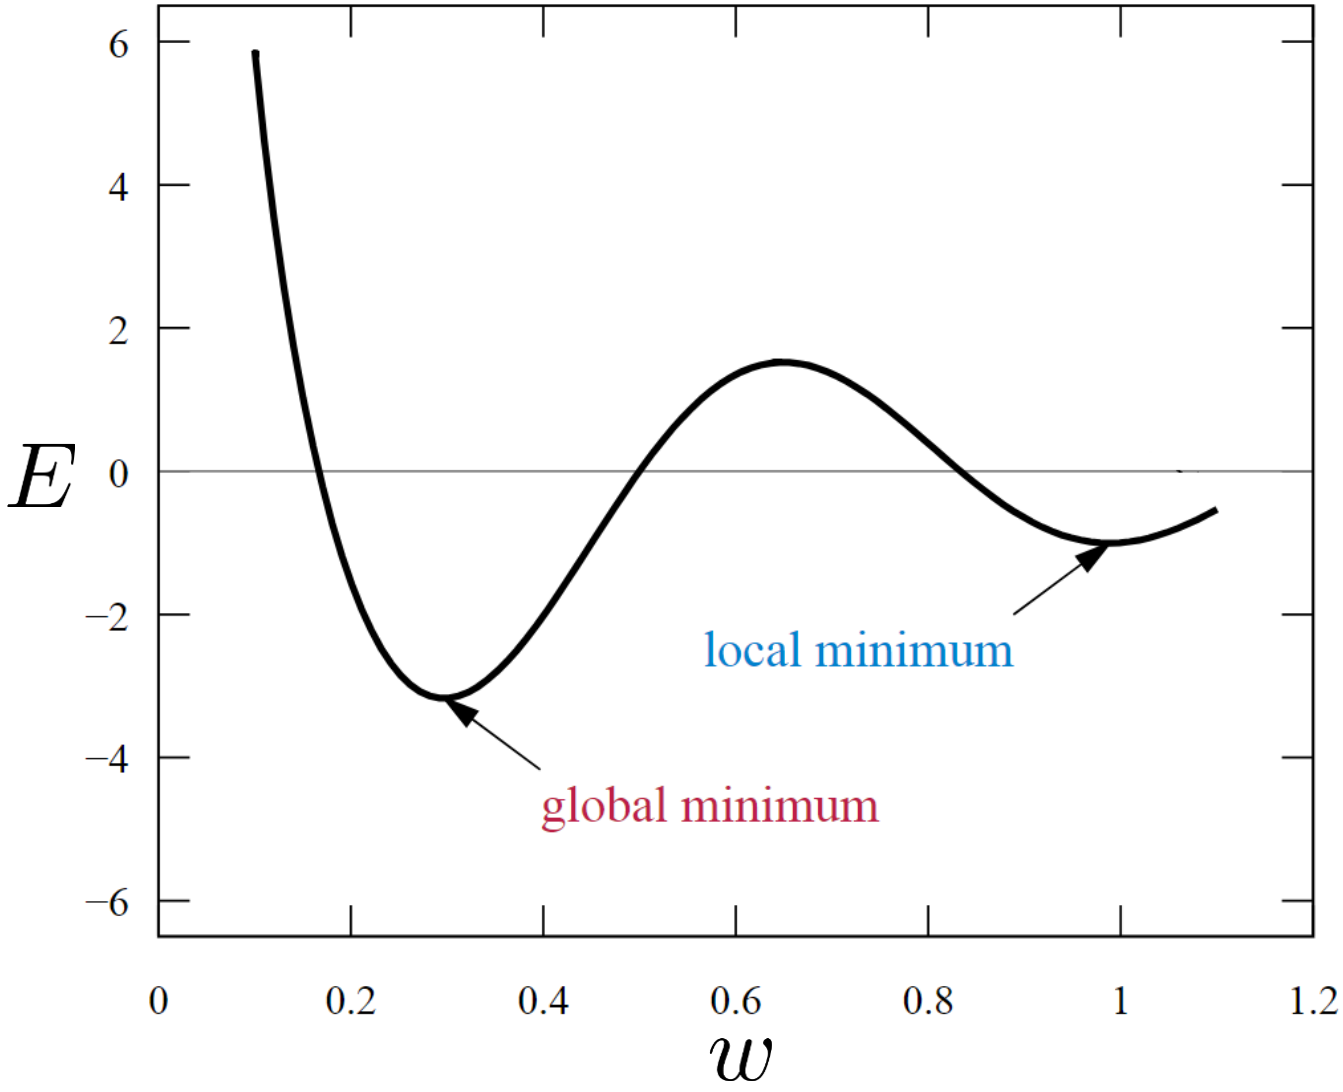

Gradient descent

$$w' = w - \alpha \frac{\partial E}{\partial w}$$

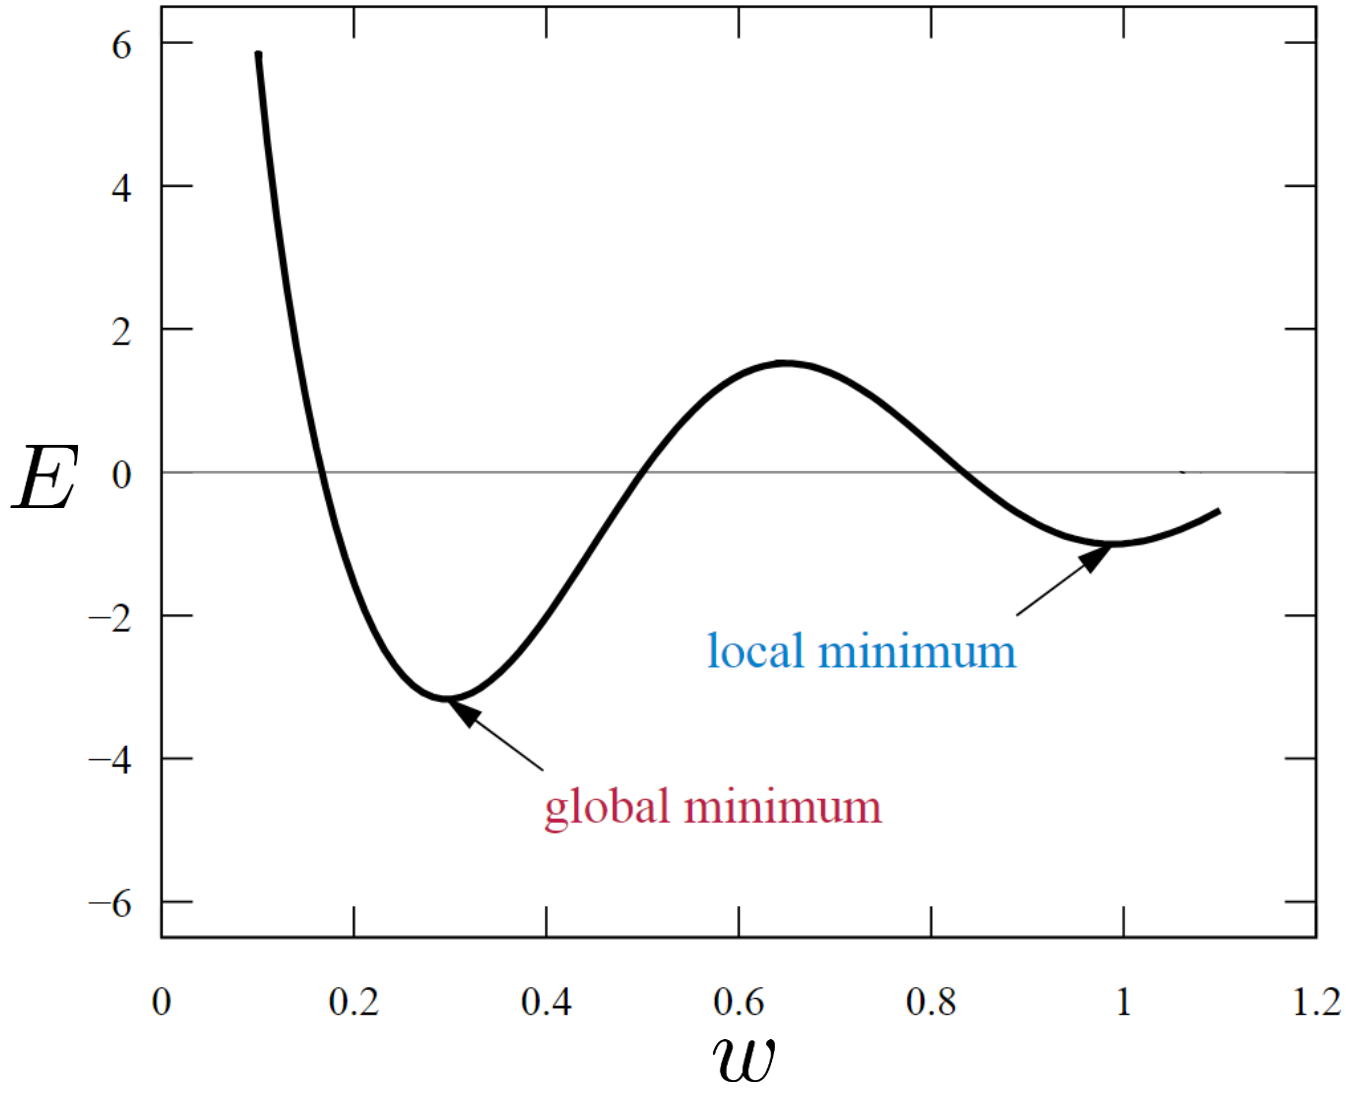

Gradient descent in higher dimensions

$$w_i' = w_i - \alpha \frac{\partial E}{\partial w_i}$$

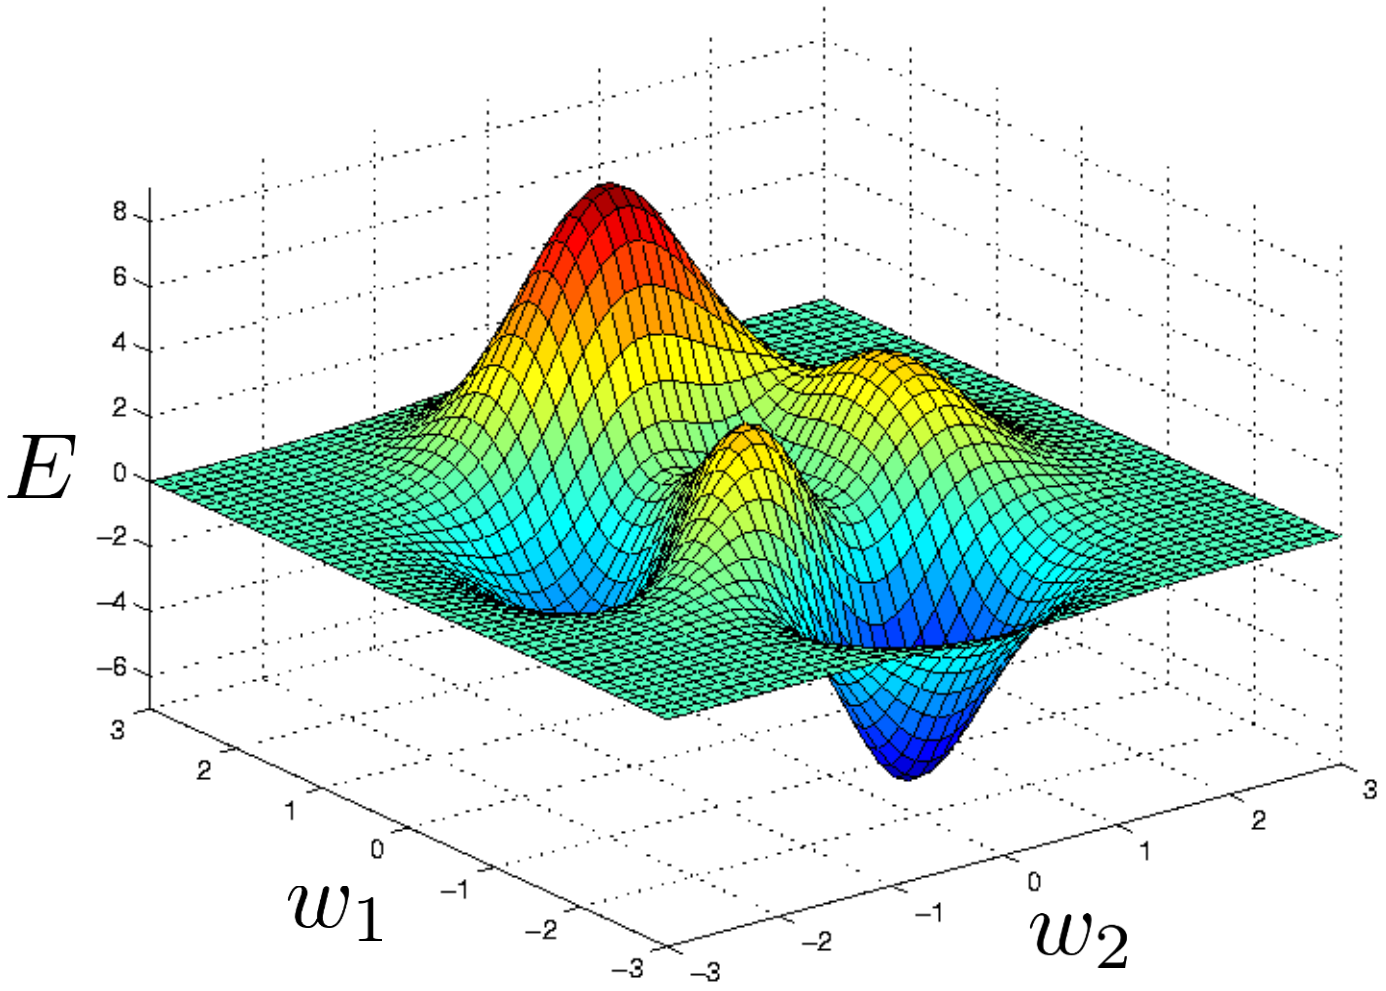

Learning multiple layers

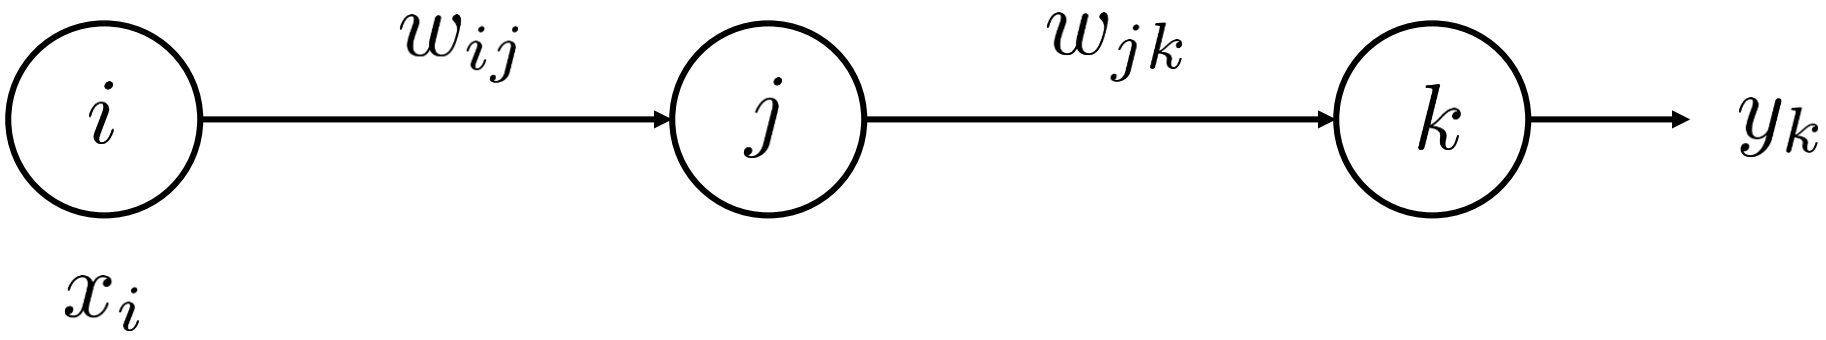

Forward pass

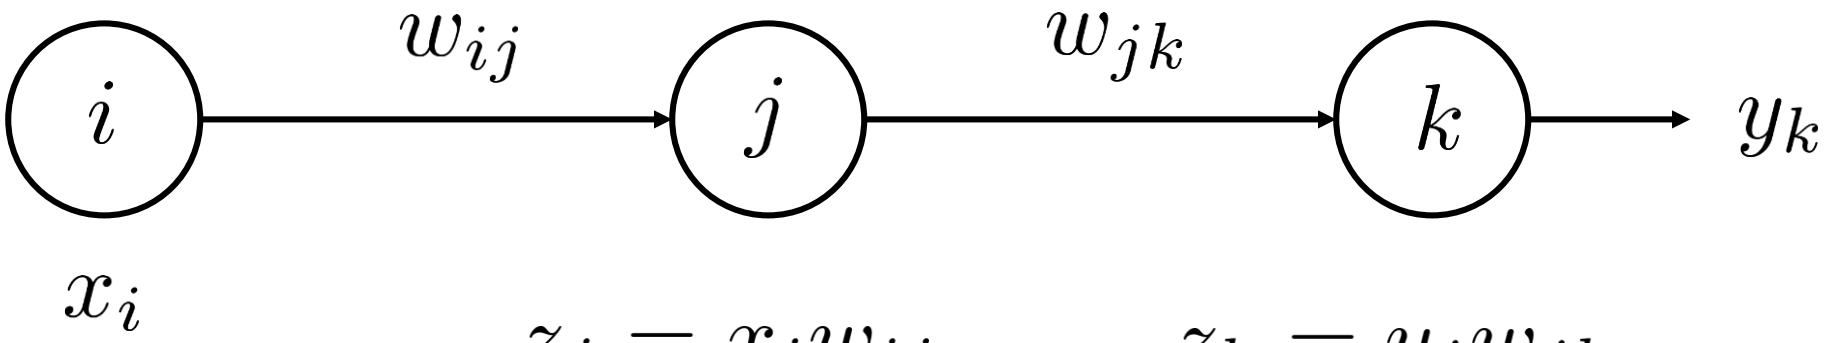

$$z_j = x_i w_{ij} \qquad z_k = y_j w_{jk}$$

$$y_j = f(z_j) \qquad y_k = f(z_k)$$

Backward propagation

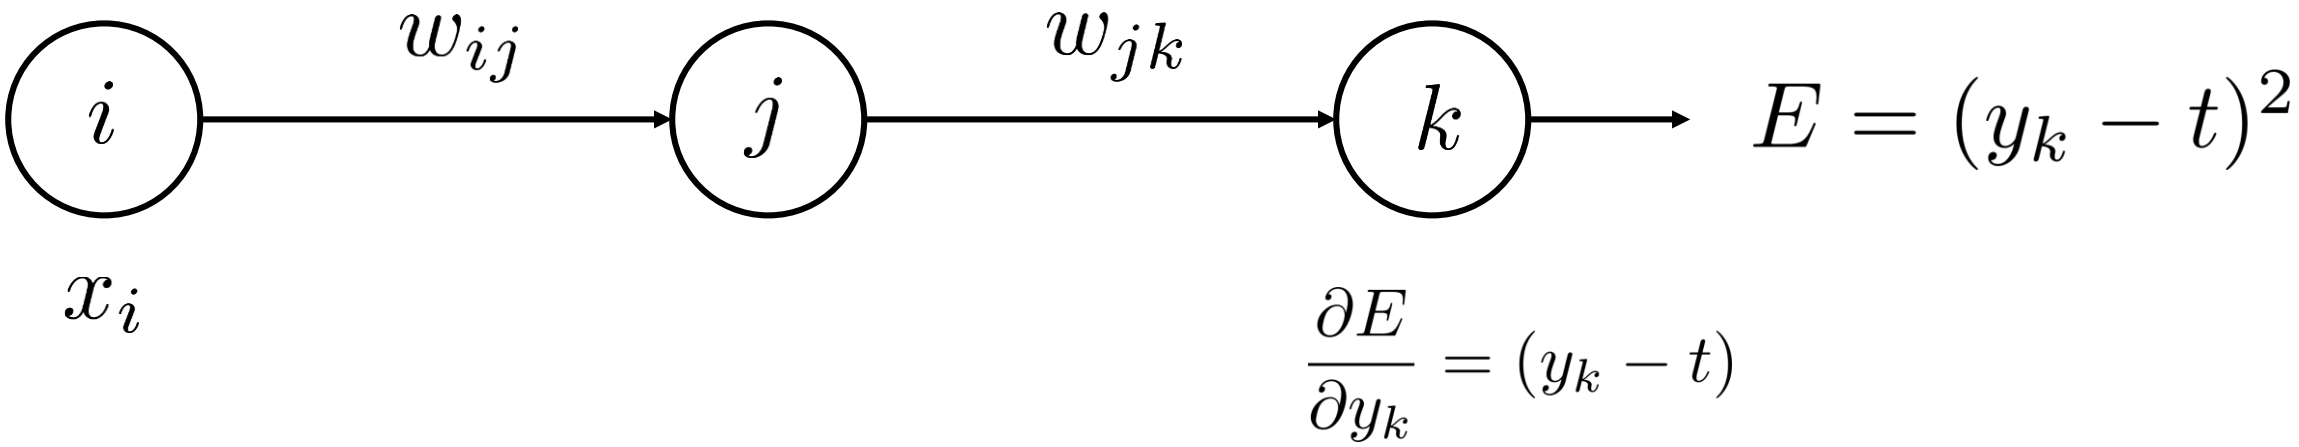

Backward propagation

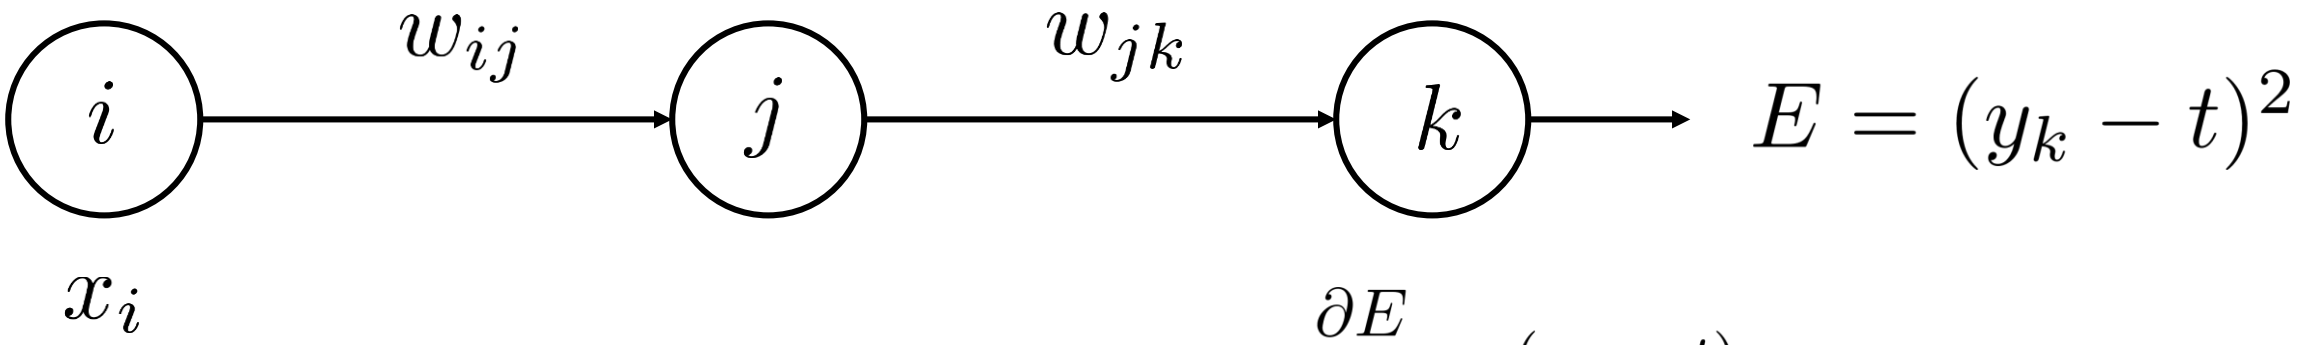

$$\frac{\partial E}{\partial y_k} = (y_k - t)$$

$$\frac{\partial E}{\partial z_k} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}$$

Backward propagation

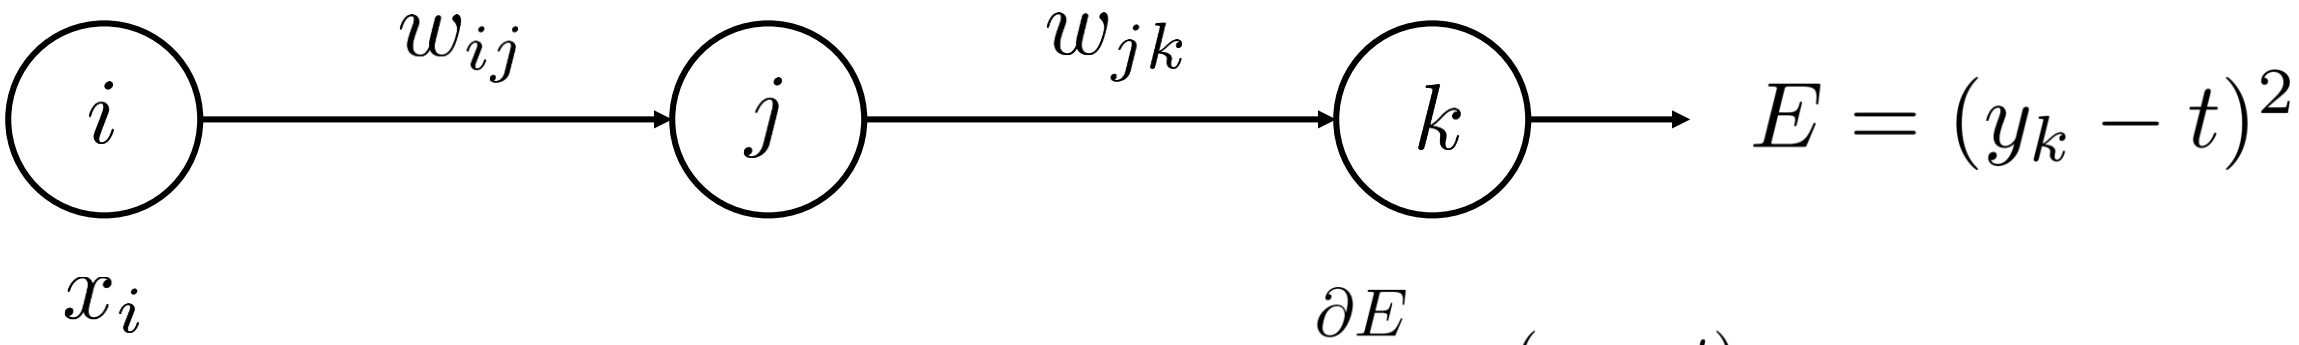

$$\frac{\partial E}{\partial y_k} = (y_k - t)$$

$$\frac{\partial E}{\partial z_k} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}$$

$$\frac{\partial E}{\partial w_{jk}} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}\frac{\partial z_k}{\partial w_{jk}}$$

Backward propagation

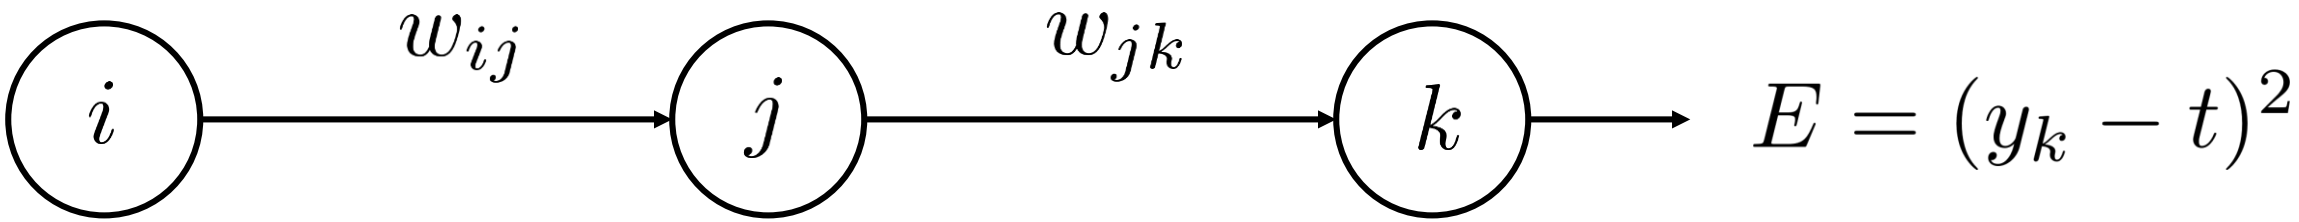

x_i

$$\frac{\partial E}{\partial y_j} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}\frac{\partial z_k}{\partial y_j}$$

$$\frac{\partial E}{\partial y_k} = (y_k - t)$$

$$\frac{\partial E}{\partial z_j} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}\frac{\partial z_k}{\partial y_j}\frac{\partial y_j}{\partial z_j}$$

$$\frac{\partial E}{\partial z_k} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}$$

$$\frac{\partial E}{\partial w_{ij}} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}\frac{\partial z_k}{\partial y_j}\frac{\partial y_j}{\partial z_j}\frac{\partial z_j}{\partial w_{ij}}$$

$$\frac{\partial E}{\partial w_{jk}} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}\frac{\partial z_k}{\partial w_{jk}}$$

Backpropagation algorithm in full

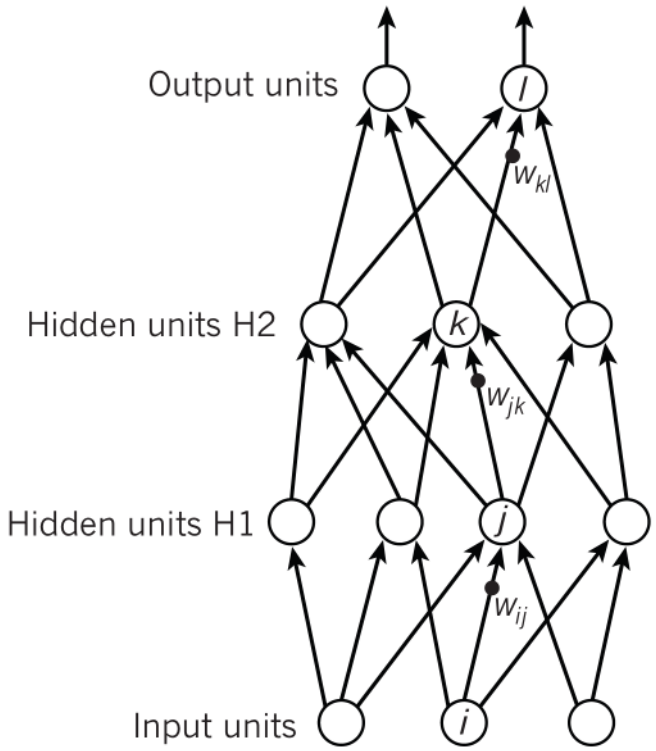

$$y_l = f(z_l)$$
$$z_l = \sum_{k \,\varepsilon\, \text{H2}} w_{kl}\, y_k$$

$$y_k = f(z_k)$$
$$z_k = \sum_{j \,\varepsilon\, \text{H1}} w_{jk}\, y_j$$

$$y_j = f(z_j)$$
$$z_j = \sum_{i \,\varepsilon\, \text{Input}} w_{ij}\, x_i$$

Compare outputs with correct answer to get error derivatives

$$\frac{\partial E}{\partial y_l} = y_l - t_l$$
$$\frac{\partial E}{\partial z_l} = \frac{\partial E}{\partial y_l}\frac{\partial y_l}{\partial z_l}$$

$$\frac{\partial E}{\partial y_k} = \sum_{l \,\varepsilon\, \text{out}} w_{kl}\frac{\partial E}{\partial z_l}$$

$$\frac{\partial E}{\partial z_k} = \frac{\partial E}{\partial y_k}\frac{\partial y_k}{\partial z_k}$$

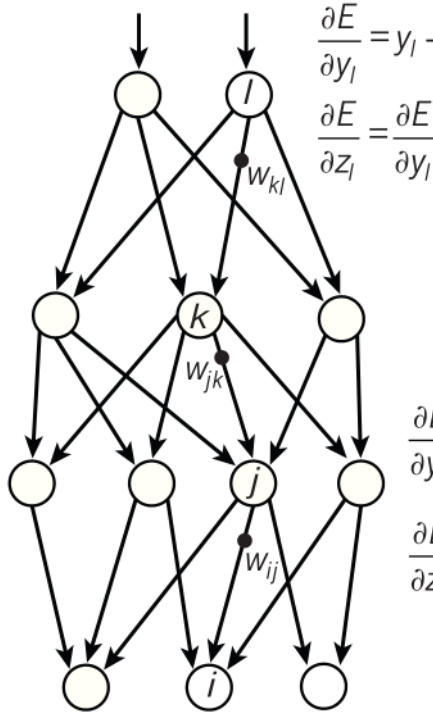

$$\frac{\partial E}{\partial y_j} = \sum_{k \,\varepsilon\, \text{H2}} w_{jk}\frac{\partial E}{\partial z_k}$$

$$\frac{\partial E}{\partial z_j} = \frac{\partial E}{\partial y_j}\frac{\partial y_j}{\partial z_j}$$

LeCun et. al., “Deep Learning.” *Nature (2015)*

Applications

Object Detection

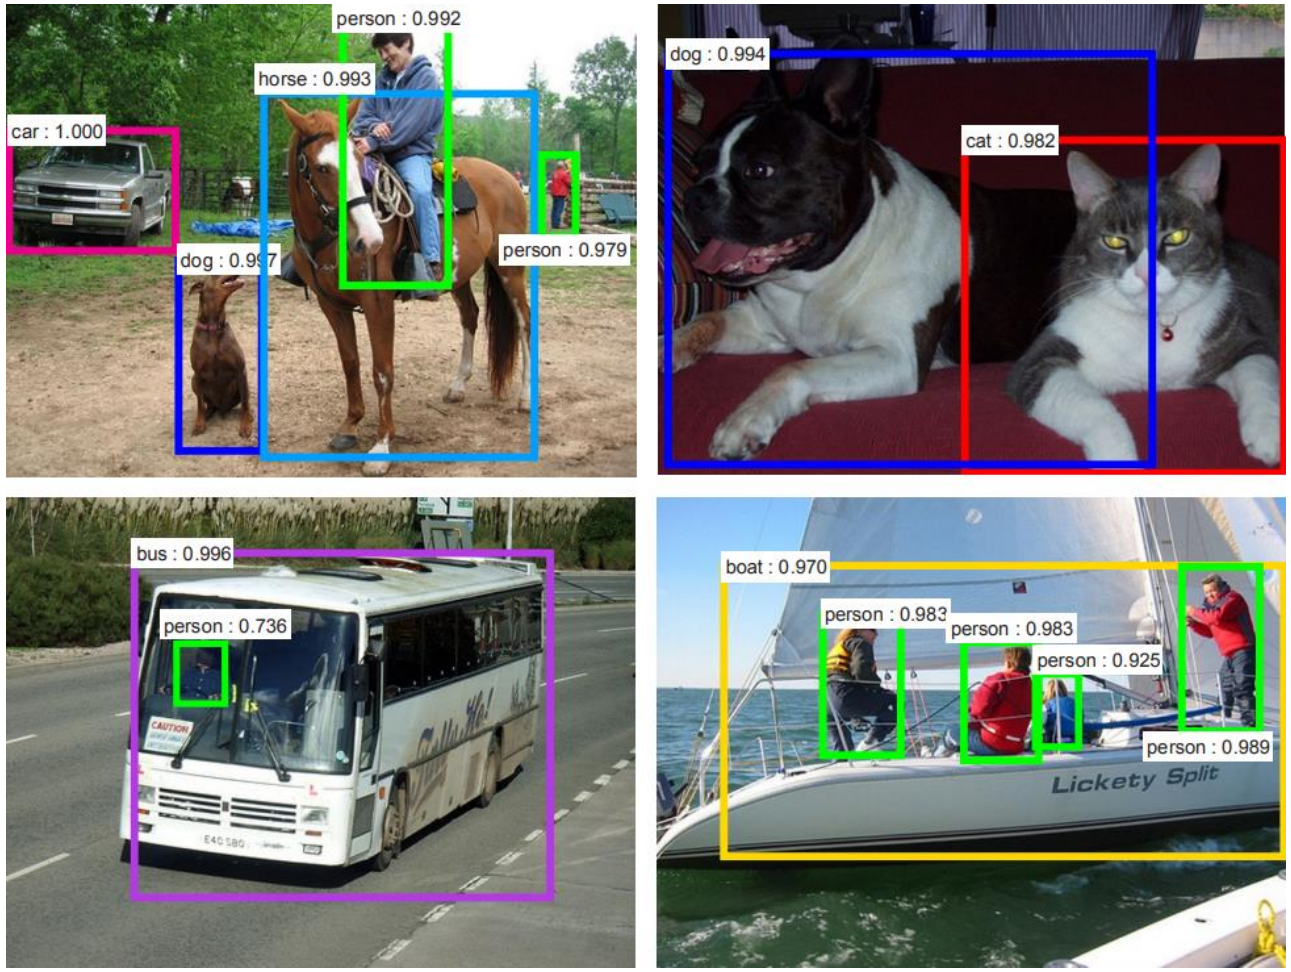

Ren et. al., “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks.” *NIPS (2015)*

Scene segmentation

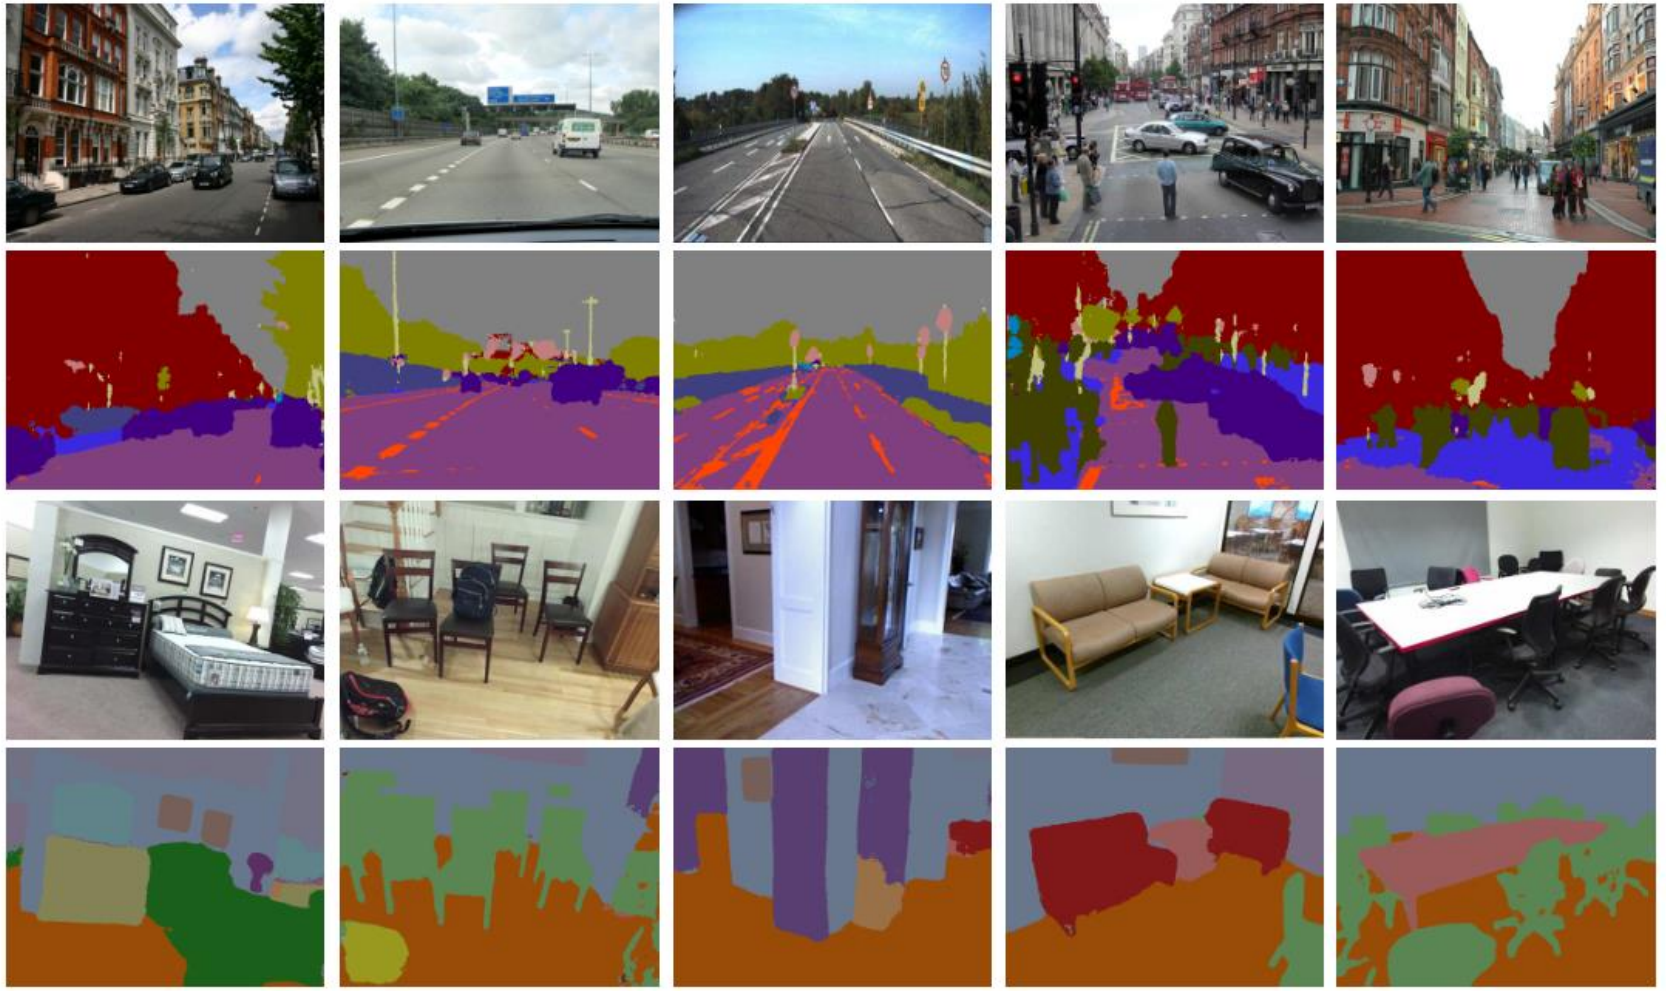

Badrinarayanan et. al., “SegNet: A Deep Convolutional Encoder-Decoder Architecture...” *PAMI (2016)*

Super resolution

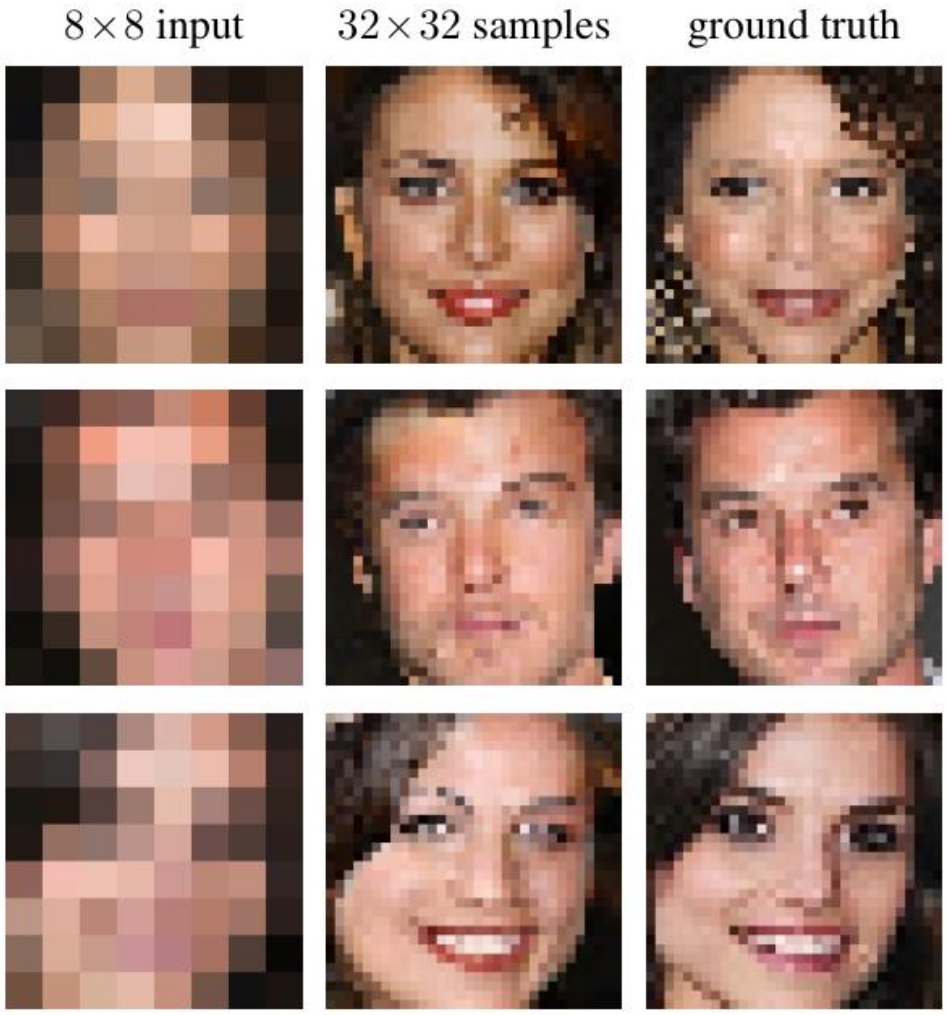

Dahl et. al., “Pixel Recursive Super Resolution.” *arXiv (2017)*

Style transfer

Gatys et. al., “A Neural Algorithm of Artistic Style.” *arXiv (2015)*

Image translation

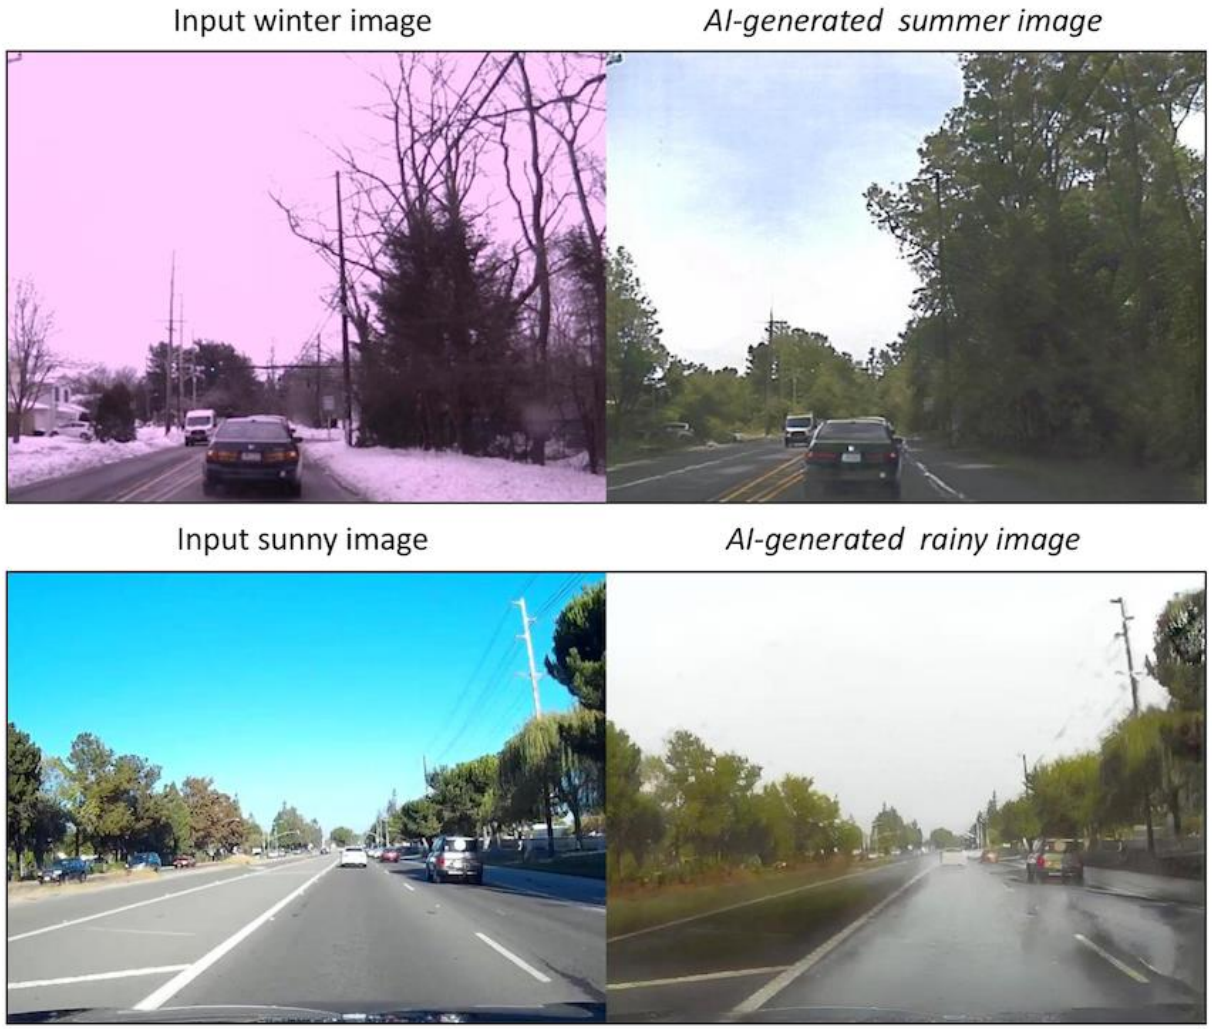

Liu et. al., “Unsupervised Image-to-Image Translation Networks.” *NIPS (2017)*

Image generation

Karras et. al., “Progressive Growing of GANs for Improved Quality, Stability, and Variation.” *arXiv (2017)*

Image generation

Karras et. al., “Progressive Growing of GANs for Improved Quality, Stability, and Variation.” *arXiv (2017)*

Learning word representations

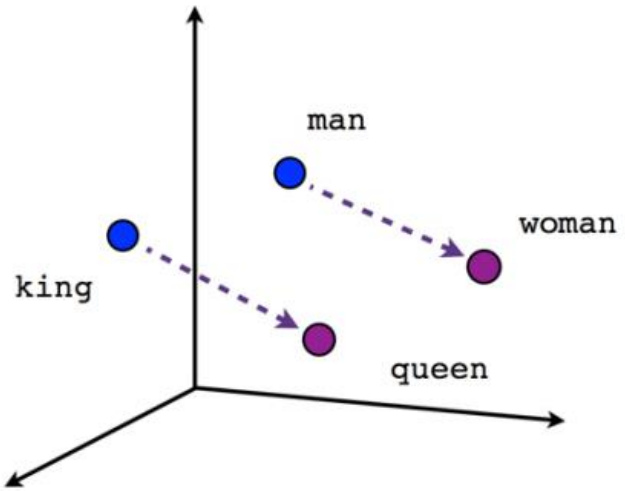

Male-Female

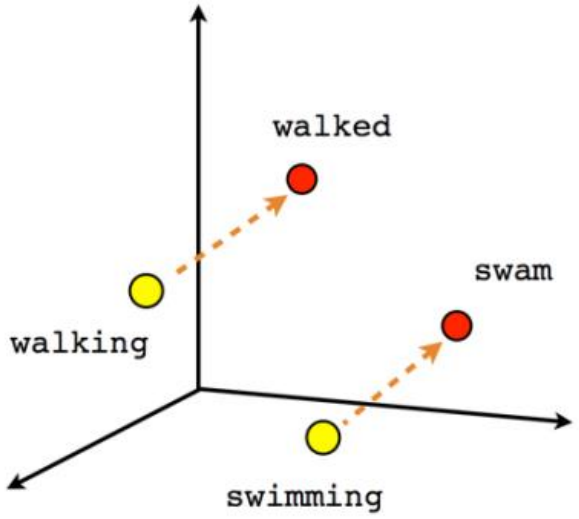

Verb tense

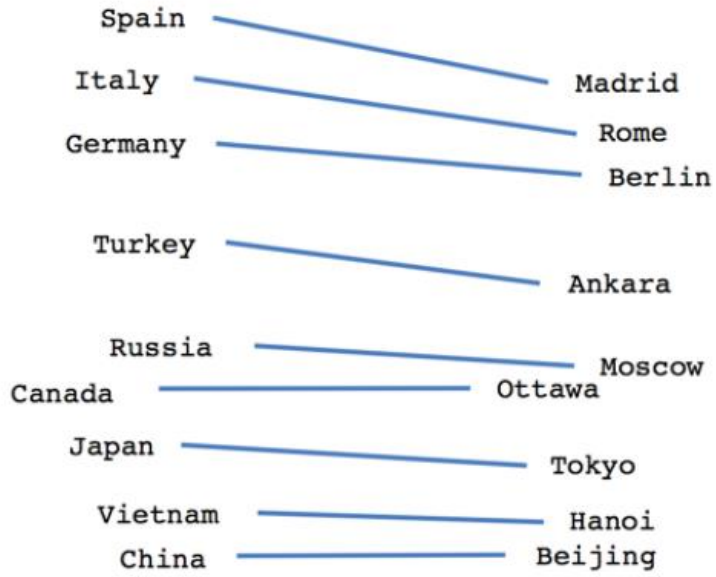

Country-Capital

Mikolov et. al., "Efficient Estimation of Word Representations in Vector Space." *arXiv (2013)*

Learning sentiment representations

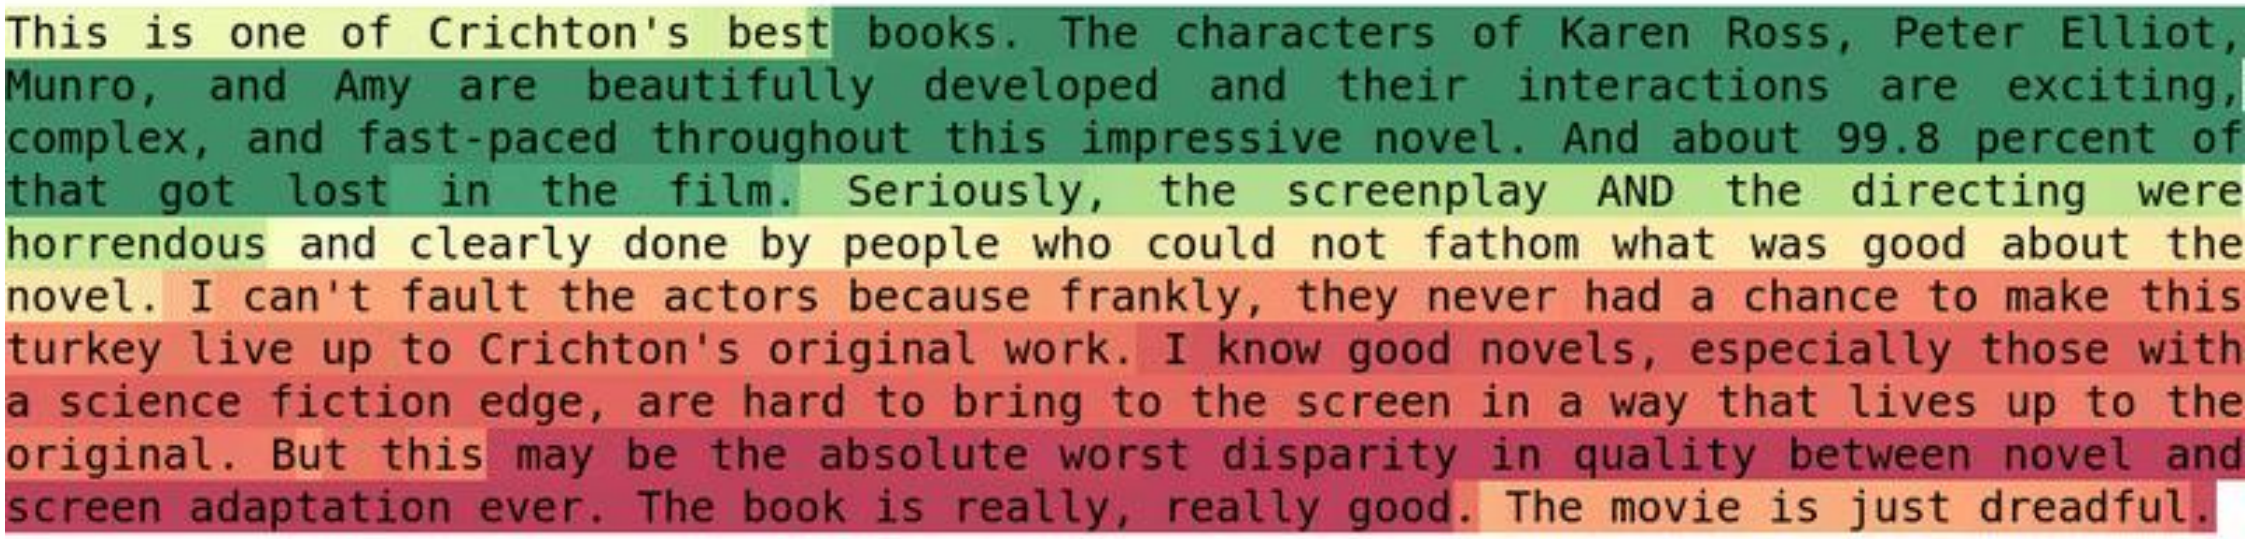

Radford et. al., “Learning to Generate Reviews and Discovering Sentiment.” *arXiv (2017)*

Image captioning

A woman is throwing a **frisbee** in a park.

A **dog** is standing on a hardwood floor.

A **stop** sign is on a road with a mountain in the background

A little **girl** sitting on a bed with a teddy bear.

A group of **people** sitting on a boat in the water.

A giraffe standing in a forest with **trees** in the background.

Vinyals et. al., "Show and Tell: A Neural Image Caption Generator." *CVPR (2015)*

Visual question answering

What are pulling a man on a wagon down on dirt road?
Answer: horses Prediction: horses

What is the color of the box ?
Answer: red Prediction: red

What next to the large umbrella attached to a table?
Answer: trees Prediction: tree

How many people are going up the mountain with walking sticks?
Answer: four Prediction: four

Yang et. al., "Stacked Attention Networks for Image Question Answering." *CVPR (2016)*

Playing games

Silver et. al., “Mastering the Game of Go with Deep Neural Networks and Tree Search.” *Nature (2016)*

Building better neural networks

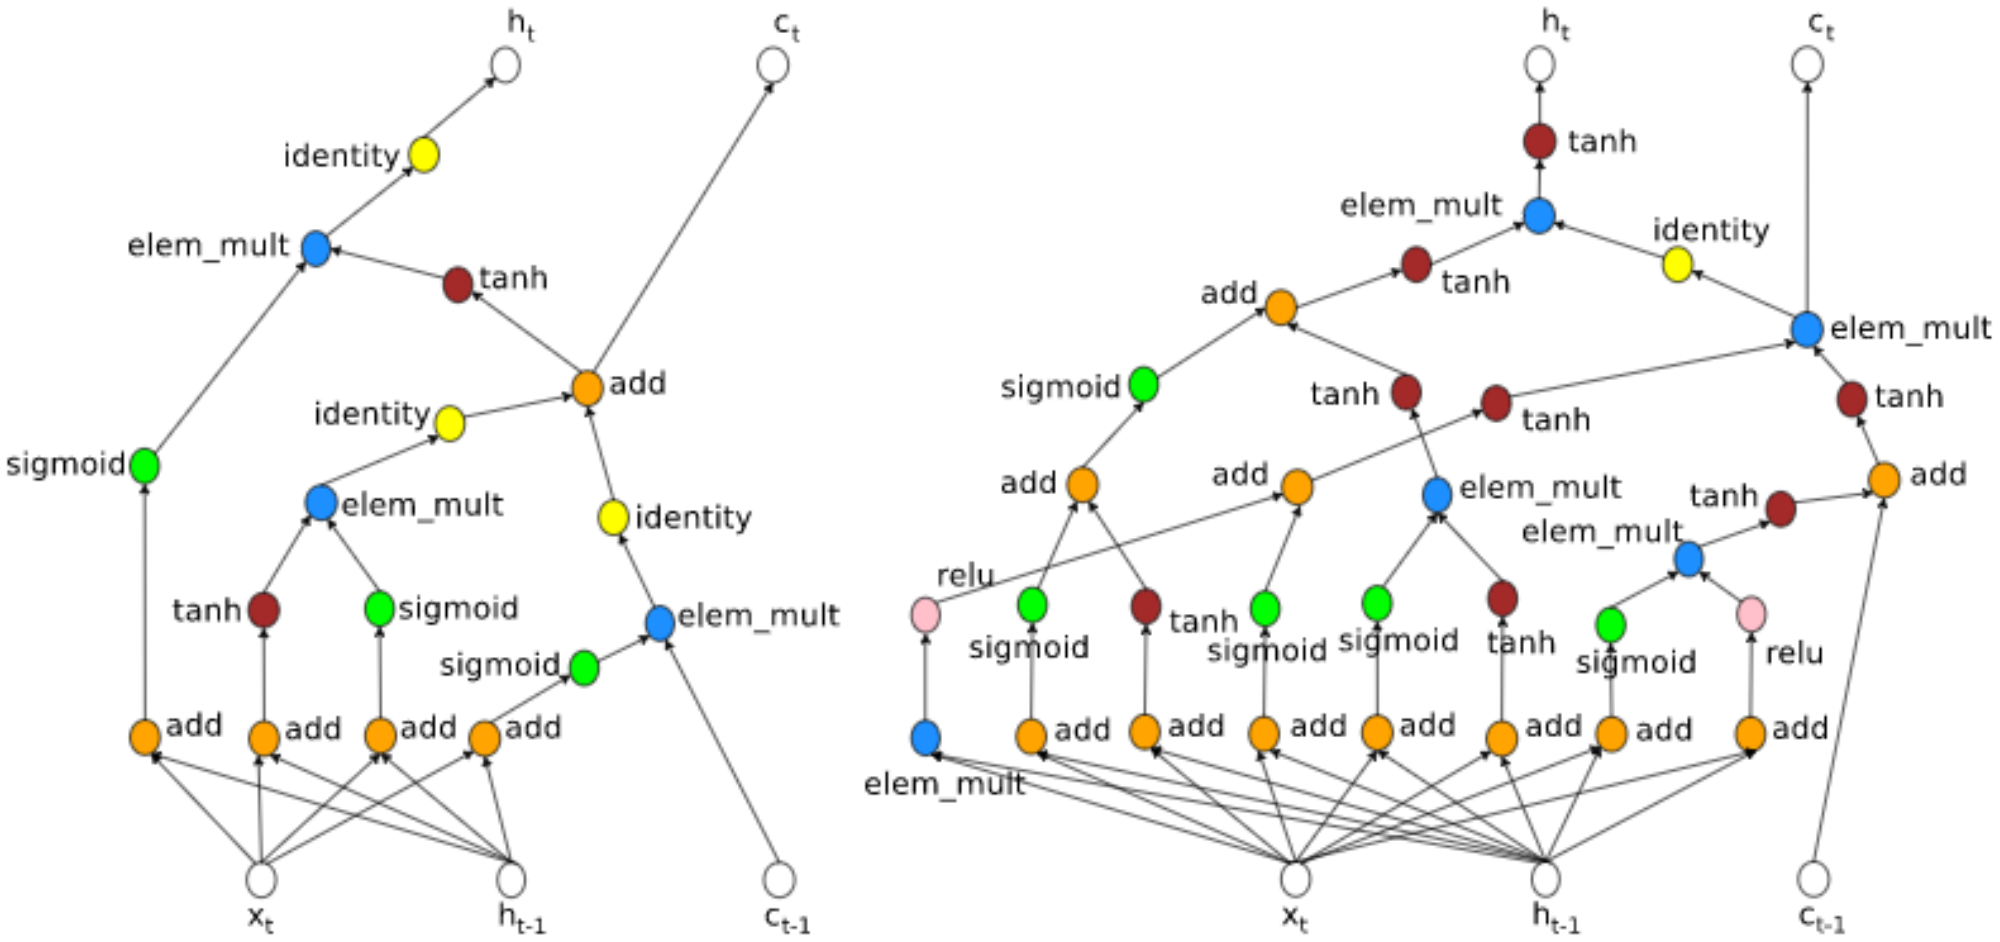

Zoph et. al., "Neural Architecture Search with Reinforcement Learning." *ICLR (2017)*

Building better software

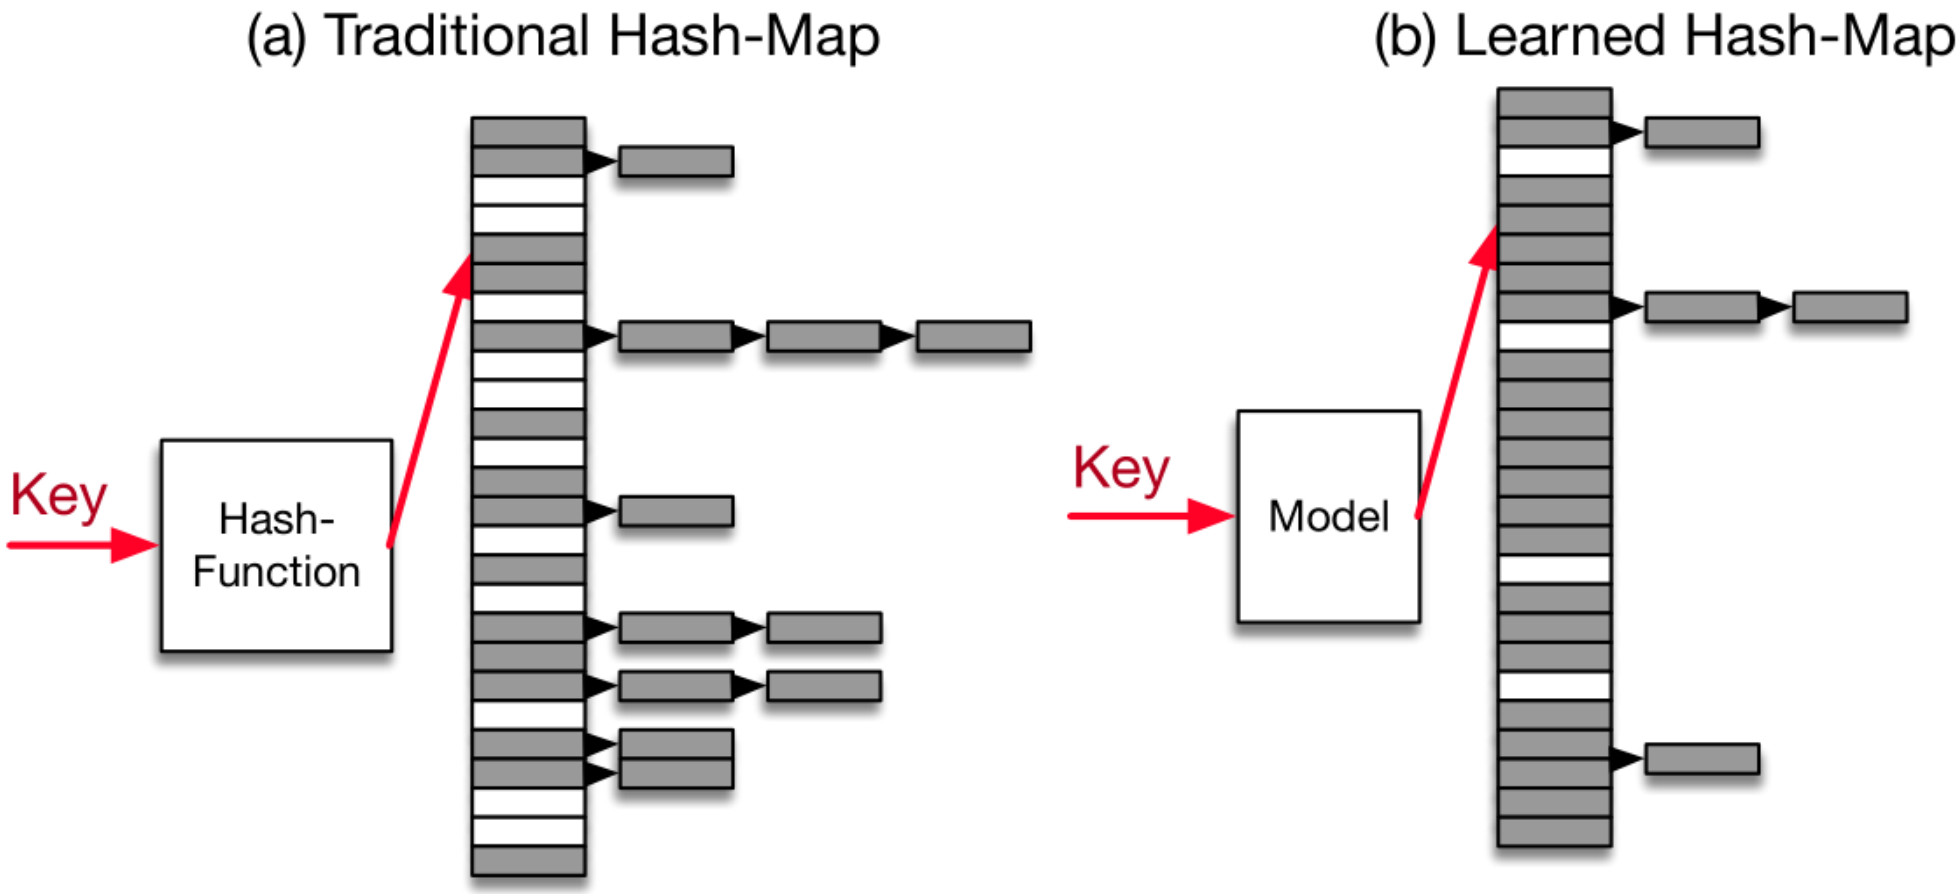

Kraska et. al., "The Case for Learned Index Structures." *arXiv (2017)*

What you need to get started with deep learning

These, pretty much

… plus a handful of other stuff

https://medium.com/towards-data-science/building-your-own-deep-learning-box-47b918aea1eb

Building a deep learning rig

Component		Selection	Base	Promo	Shipping	Tax	Price	Where	
CPU		Intel - Core i7-7700K 4.2GHz Quad-Core Processor	$318.69				$318.69	OutletPC	Buy
CPU Cooler		Cooler Master - Hyper 212 EVO 82.9 CFM Sleeve Bearing CPU Cooler	$29.99	-$10.00	FREE		$19.99	Newegg	Buy
Motherboard		MSI - Z270-A PRO ATX LGA1151 Motherboard	$114.88	-$10.00			$104.88	OutletPC	Buy
Memory		Corsair - Vengeance LPX 32GB (2 x 16GB) DDR4-3200 Memory	$366.59		FREE		$366.59	Newegg Marketplace	Buy
Storage		Samsung - 850 EVO-Series 500GB 2.5" Solid State Drive	$149.89				$149.89	OutletPC	Buy
Video Card		Zotac - GeForce GTX 1080 8GB AMP! Edition Video Card	$564.75				$564.75	OutletPC	Buy
Case		NZXT - S340 (White) ATX Mid Tower Case	$69.99	-$10.00	FREE		$59.99	Newegg	Buy
Power Supply		EVGA - SuperNOVA G2 750W 80+ Gold Certified Fully-Modular ATX Power Supply	$99.99	-$20.00	$5.99		$85.98	Newegg	Buy

Base Total:	$1714.77
Mail-in Rebates:	-$50.00
Shipping:	$5.99
Total:	**$1670.76**

https://pcpartpicker.com/list/FRp8XH

An alternative

Amazon Web Services

Amazon EC2 P2 Instances

Powerful, Scalable GPU instances for high-performance computing

Amazon EC2 P2 Instances are powerful, scalable instances that provide GPU-based parallel compute capabilities. For customers with graphics requirements, see G2 instances for more information.

P2 instances, designed for general-purpose GPU compute applications using CUDA and OpenCL, are ideally suited for machine learning, high performance databases, computational fluid dynamics, computational finance, seismic analysis, molecular modeling, genomics, rendering, and other server-side workloads requiring massive parallel floating point processing power.

Use the Amazon Linux AMI, pre-installed with popular deep learning frameworks such as Caffe and Mxnet, so you can get started quickly. You can also use the NVIDIA AMI with GPU driver and CUDA toolkit pre-installed for rapid onboarding.

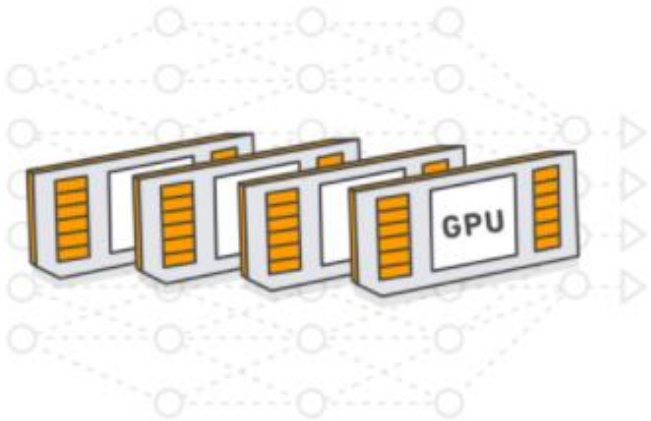

Get Started with T2 for Free

Create a Free Account

AWS Free Tier includes 750 hours of both Linux and Windows t2.micro instances each month for one year for new AWS customers. To stay within the Free Tier, use only t2.micro instances.

View AWS Free Tier Details

Amazon Web Services

	vCPU	ECU	Memory (GiB)	Instance Storage (GB)	Linux/UNIX Usage
GPU Instances - Current Generation					
p2.xlarge	4	12	61	EBS Only	$0.9 per Hour
p2.8xlarge	32	94	488	EBS Only	$7.2 per Hour
p2.16xlarge	64	188	732	EBS Only	$14.4 per Hour
g3.4xlarge	16	47	122	EBS Only	$1.14 per Hour
g3.8xlarge	32	94	244	EBS Only	$2.28 per Hour
g3.16xlarge	64	188	488	EBS Only	$4.56 per Hour

Amazon Web Services

AWS AI Blog

AWS Deep Learning AMI Now Supports PyTorch, Keras 2 and Latest Deep Learning Frameworks

by Cynthya Peranandam | on 18 OCT 2017 | in Artificial Intelligence, AWS Deep Learning AMIs | Permalink | Share

Today, we're pleased to announce an update to the AWS Deep Learning AMI.

The AWS Deep Learning AMI, which lets you spin up a complete deep learning environment on AWS in a single click, now includes PyTorch, Keras 1.2 and 2.0 support, along with popular machine learning frameworks such as TensorFlow, Caffe2 and Apache MXNet.

Using PyTorch for fast prototyping

The AMI now includes PyTorch 0.2.0, allowing developers to create dynamic neural networks in Python, a good fit for dynamic inputs such as text and time series. Developers can get started quickly using these beginner and advanced tutorials, including setting up distributed training with PyTorch.

Improved Keras support

The AMI now supports the most recent version of Keras, v2.0.8. By default, your Keras code will run against TensorFlow as a backend; you can also swap to other supported backends such as Theano and CNTK. We've also included a modified version of Keras 1.2.2 which runs on the Apache MXNet backend with better training performance.

Deep learning software ecosystem

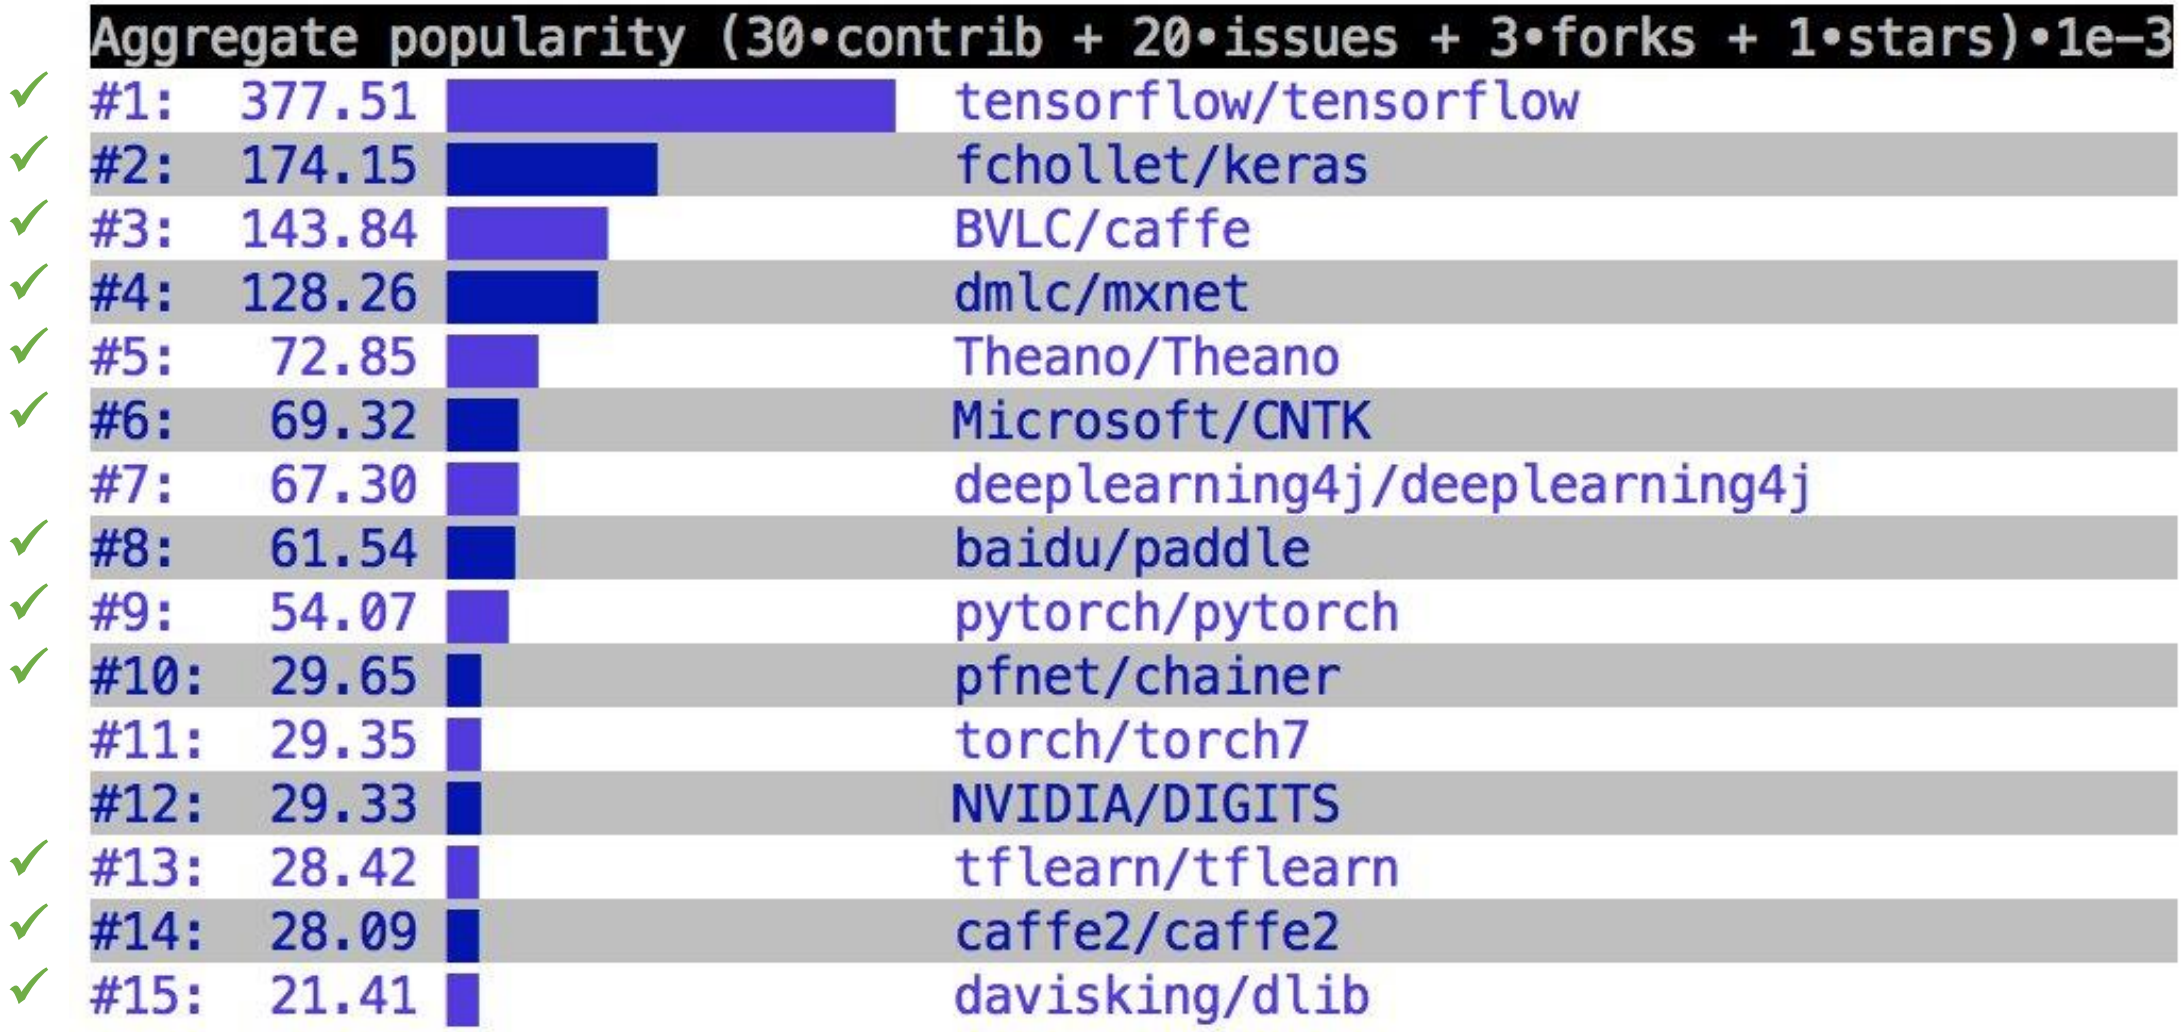

https://twitter.com/fchollet/status/915366704401719296

A typical beginner stack

A typical beginner stack

A typical beginner stack

A typical beginner stack

Deep learning in a day

1. Create an AWS account
2. Launch an EC2 instance
3. SSH into your instance
4. Launch a Jupyter Notebook
5. ???
6. Profit!

EC2 Dashboard
Events
Tags
Reports
Limits
INSTANCES
Instances
Launch Templates
Spot Requests
Reserved Instances
Dedicated Hosts
Scheduled Instances
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORK & SECURITY
Security Groups
Elastic IPs
Placement Groups
Key Pairs

Resources

You are using the following Amazon EC2 resources in the US West (Oregon) region:

- 0 Running Instances
- 0 Dedicated Hosts
- 0 Volumes
- 1 Key Pairs
- 0 Placement Groups
- 0 Elastic IPs
- 0 Snapshots
- 0 Load Balancers
- 3 Security Groups

Create Instance

To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.

Launch Instance

Note: Your instances will launch in the US West (Oregon) region

Service Health

Service Status:

US West (Oregon):
This service is operating normally

Availability Zone Status:

us-west-2a:
Availability zone is operating normally

us-west-2b:
Availability zone is operating normally

us-west-2c:
Availability zone is operating normally

Scheduled Events

US West (Oregon):

No events

Account Attributes

Supported Platforms
VPC

Default VPC
vpc-9e94a5fb

Resource ID length management

Additional Information

Getting Started Guide
Documentation
All EC2 Resources
Forums
Pricing
Contact Us

AWS Marketplace

Find free software trial products in the AWS Marketplace from the EC2 Launch Wizard. Or try these popular AMIs:

Barracuda NextGen Firewall F-Series - PAYG

Provided by Barracuda Networks, Inc.
Rating ★★★☆☆
Starting from $0.60/hr or from $4,599/yr (12% savings) for software + AWS usage fees

EC2 Dashboard
Events
Tags
Reports
Limits
INSTANCES
Instances
Launch Templates
Spot Requests
Reserved Instances
Dedicated Hosts
Scheduled Instances
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORK & SECURITY
Security Groups
Elastic IPs
Placement Groups
Key Pairs

Launch Instance | Connect | Actions

Filter by tags and attributes or search by keyword

None found

You do not have any running instances in this region.

First time using EC2? Check out the Getting Started Guide.

Click the Launch Instance button to start your own server.

Launch Instance

Select an instance above

Step 1: Choose an Amazon Machine Image (AMI)

Cancel and Exit

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.

Quick Start
My AMIs
AWS Marketplace
Community AMIs

Operating system

- Amazon Linux
- Cent OS
- Debian
- Fedora
- Gentoo
- OpenSUSE
- Other Linux
- Red Hat
- SUSE Linux
- Ubuntu
- Windows

Architecture

deep learning

1 to 50 of 66 AMIs

awsmarketplace — 40 results for "deep learning" on AWS Marketplace
Partner software pre-configured to run on AWS

Deep Learning AMI (Ubuntu) Version 2.0 - ami-3b6bce43
Deep Learning AMI with Conda-based virtual environments for Apache MXNet, TensorFlow, Caffe2, PyTorch, Theano, CNTK and Keras
Root device type: ebs Virtualization type: hvm ENA Enabled: Yes
Select
64-bit

Deep Learning AMI (Amazon Linux) Version 2.0 - ami-5c60c524
Deep Learning AMI with Conda-based virtual environments for Apache MXNet, TensorFlow, Caffe2, PyTorch, Theano, CNTK and Keras
Root device type: ebs Virtualization type: hvm ENA Enabled: Yes
Select
64-bit

Deep Learning Base AMI (Ubuntu) Version 2.0 - ami-041db87c
Deep Learning base AMI with NVidia drivers like CUDA 8 and 9, CuDNN 6 and 7, CuBLAS 8 and 9, NCCL and more
Root device type: ebs Virtualization type: hvm ENA Enabled: Yes
Select
64-bit

Deep Learning Base AMI (Amazon Linux) Version 2.0 - ami-8414b1fc
Select

Step 2: Choose an Instance Type

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instances are virtual servers that can run applications. They have varying combinations of CPU, memory, storage, and networking capacity, and give you the flexibility to choose the appropriate mix of resources for your applications. Learn more about instance types and how they can meet your computing needs.

Filter by: GPU compute | Current generation | Show/Hide Columns

Currently selected: p2.xlarge (11.75 ECUs, 4 vCPUs, 2.7 GHz, E5-2686v4, 61 GiB memory, EBS only)

	Family	Type	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance	IPv6 Support
☑	GPU compute	p2.xlarge	4	61	EBS only	Yes	High	Yes
☐	GPU compute	p2.8xlarge	32	488	EBS only	Yes	10 Gigabit	Yes
☐	GPU compute	p2.16xlarge	64	732	EBS only	Yes	25 Gigabit	Yes
☐	GPU compute	p3.2xlarge	8	61	EBS only	Yes	Up to 10 Gigabit	Yes
☐	GPU compute	p3.8xlarge	32	244	EBS only	Yes	10 Gigabit	Yes
☐	GPU compute	p3.16xlarge	64	488	EBS only	Yes	25 Gigabit	Yes

Cancel | Previous | Review and Launch | Next: Configure Instance Details

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

AMI Details

Edit AMI

Deep Learning AMI (Ubuntu) Version 2.0 - ami-3b6bce43

Deep Learning AMI with Conda-based virtual environments for Apache MXNet, TensorFlow, Caffe2, PyTorch, Theano, CNTK and Keras

Root Device Type: ebs Virtualization type: hvm

Instance Type

Edit instance type

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance
p2.xlarge	11.75	4	61	EBS only	Yes	High

Security Groups

Edit security groups

Security group name launch-wizard-3
Description launch-wizard-3 created 2017-12-15T23:16:49.904+05:00

Type	Protocol	Port Range	Source	Description

This security group has no rules

Instance Details

Edit instance details

Cancel Previous **Launch**

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

AMI Details

Edit AMI

Deep Learning AMI (Ubuntu) V

Deep Learning AMI with Conda-base

Root Device Type: ebs Virtualization type

Instance Type

Edit instance type

Instance Type	ECUs	vC	Network Performance
p2.xlarge	11.75	4	High

Security Groups

Edit security groups

Security group name launch-wizard

Description launch-wizard

Type	Pr	Description

This security group has no rules

Instance Details

Edit instance details

Cancel Previous Launch

Select an existing key pair or create a new key pair

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about removing existing key pairs from a public AMI.

Choose an existing key pair

Select a key pair

aws-key

☑ I acknowledge that I have access to the selected private key file (aws-key.pem), and that without this file, I won't be able to log into my instance.

Cancel Launch Instances

EC2 Dashboard
Events
Tags
Reports
Limits
INSTANCES
Instances
Launch Templates
Spot Requests
Reserved Instances
Dedicated Hosts
Scheduled Instances
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORK & SECURITY
Security Groups
Elastic IPs
Placement Groups
Key Pairs

Launch Instance | Connect | Actions

Filter by tags and attributes or search by keyword

1 to 1 of 1

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS (IPv4)	IPv4
	i-01ea3aedef789eca9	p2.xlarge	us-west-2b	running	2/2 checks ...	None	ec2-52-38-150-159.us-...	52.3

Instance: i-01ea3aedef789eca9 **Public DNS: ec2-52-38-150-159.us-west-2.compute.amazonaws.com**

Description | Status Checks | Monitoring | Tags

Instance ID	i-01ea3aedef789eca9	Public DNS (IPv4)	ec2-52-38-150-159.us-west-2.compute.amazonaws.com
Instance state	running	IPv4 Public IP	52.38.150.159
Instance type	p2.xlarge	IPv6 IPs	-
Elastic IPs		Private DNS	ip-172-31-37-122.us-west-2.compute.internal
Availability zone	us-west-2b	Private IPs	172.31.37.122
Security groups	launch-wizard-1. view inbound rules	Secondary private IPs	
Scheduled events	No scheduled events	VPC ID	vpc-9e94a5fb
AMI ID	Deep Learning AMI (Ubuntu) Version 2.0	Subnet ID	subnet-d3adf5a4

```
$ssh -i ./aws-key.pem ubuntu@ec2-34-211-139-121.us-west-2.compute.amazonaws.com

$jupyter notebook
```

```
$ ssh -i ./aws-key.pem ubuntu@ec2-34-211-139-121.us-west-2.compute.amazonaws.com
The authenticity of host 'ec2-34-211-139-121.us-west-2.compute.amazonaws.com (34
.211.139.121)' can't be established.
ECDSA key fingerprint is SHA256:kHrJKN1SjbiW2Igs1S8oPEyv6aFdpOAYd64vvV3fmpw.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'ec2-34-211-139-121.us-west-2.compute.amazonaws.com,3
4.211.139.121' (ECDSA) to the list of known hosts.
Welcome to Ubuntu 16.04.1 LTS (GNU/Linux 4.4.0-36-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:     https://landscape.canonical.com
 * Support:        https://ubuntu.com/advantage

  Get cloud support with Ubuntu Advantage Cloud Guest:
    http://www.ubuntu.com/business/services/cloud

469 packages can be updated.
192 updates are security updates.


Last login: Sun Oct 16 02:04:33 2016 from 124.110.152.24
ubuntu@ip-172-31-38-216:~$ nvidia-smi
Fri Dec 15 16:43:05 2017
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 367.48                 Driver Version: 367.48                    |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla K80           Off  | 0000:00:1E.0     Off |                    0 |
| N/A   48C    P0    56W / 149W |      0MiB / 11439MiB |     99%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID  Type  Process name                               Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
ubuntu@ip-172-31-38-216:~$ jupyter notebook
[I 16:43:48.924 NotebookApp] [nb_conda_kernels] enabled, 2 kernels found
[I 16:43:48.928 NotebookApp] Writing notebook server cookie secret to /run/user/1000/jupyter
/notebook_cookie_secret
[W 16:43:52.264 NotebookApp] WARNING: The notebook server is listening on all IP addresses a
nd not using encryption. This is not recommended.
[I 16:43:53.063 NotebookApp] ✓ nbpresent HTML export ENABLED
[W 16:43:53.063 NotebookApp] ✗ nbpresent PDF export DISABLED: No module named nbbrowserpdf.e
xporters.pdf
[I 16:43:53.069 NotebookApp] [nb_conda] enabled
[I 16:43:54.522 NotebookApp] [nb_anacondacloud] enabled
[I 16:43:54.527 NotebookApp] Serving notebooks from local directory: /home/ubuntu
[I 16:43:54.527 NotebookApp] 0 active kernels
[I 16:43:54.527 NotebookApp] The Jupyter Notebook is running at: http://[all ip addresses on
 your system]:8888/
[I 16:43:54.527 NotebookApp] Use Control-C to stop this server and shut down all kernels (tw
ice to skip confirmation).
```

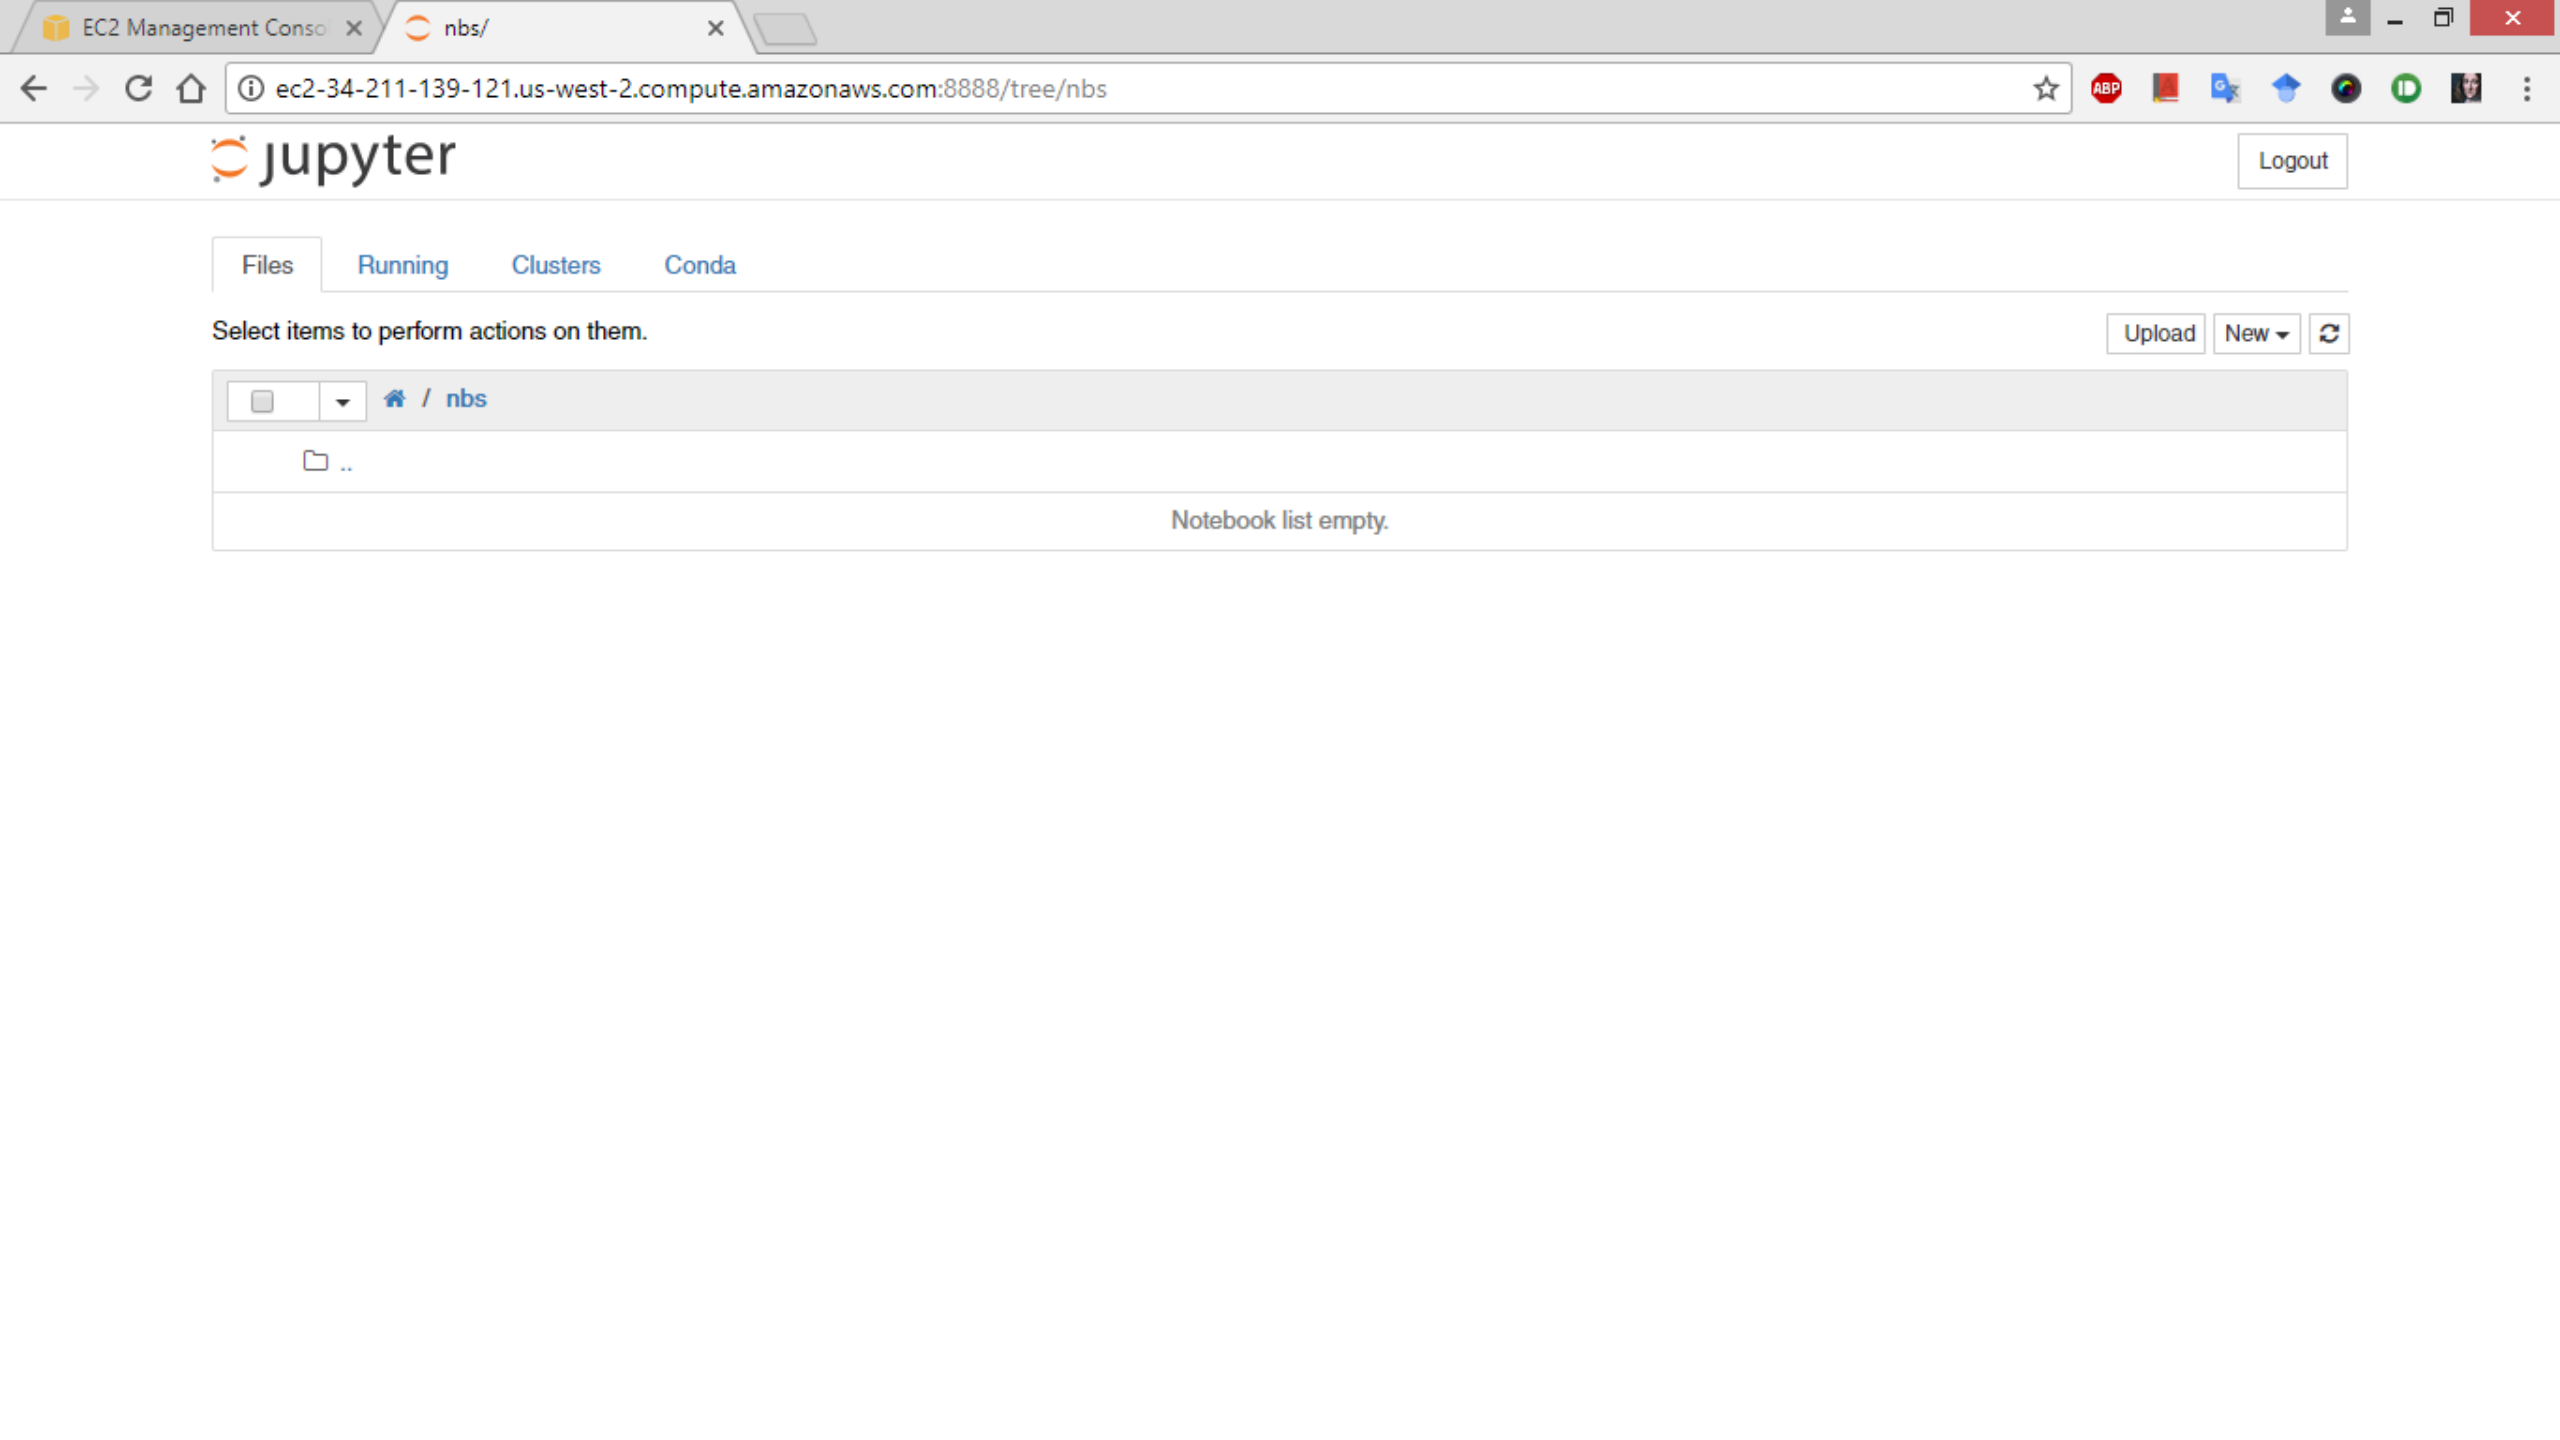
EC2 Management Conso
nbs/
ec2-34-211-139-121.us-west-2.compute.amazonaws.com:8888/tree/nbs
jupyter
Logout
Files
Running
Clusters
Conda
Select items to perform actions on them.
Upload
New
/ nbs
..
Notebook list empty.

EC2 Management Conso | nbs/ | Untitled

Not secure | ec2-34-211-139-121.us-west-2.compute.amazonaws.com:8888/notebooks/nbs/Untitled.ipynb

jupyter Untitled Last Checkpoint: 10 minutes ago (unsaved changes)

Logout

File Edit View Insert Cell Kernel Widgets Help

Python [default]

Code CellToolbar

In [1]:
```
"Hello World"
```

Out[1]: 'Hello World'

In [2]:
```
1 + 1
```

Out[2]: 2

```
import keras
from keras.models import Sequential

model = Sequential()
model.add(Dense(512, activation='relu', input_shape=(784,)))
model.add(Dense(512, activation='relu'))
model.add(Dense(10, activation='softmax'))

model.compile(loss='categorical_crossentropy', optimizer='sgd')
model.fit(x_train, y_train, epochs=5, batch_size=32)

classes = model.predict(x_test, batch_size=128)
score = model.evaluate(x_test, y_test)
```

Where to learn more

Courses

fast.ai

by Jeremy Howard

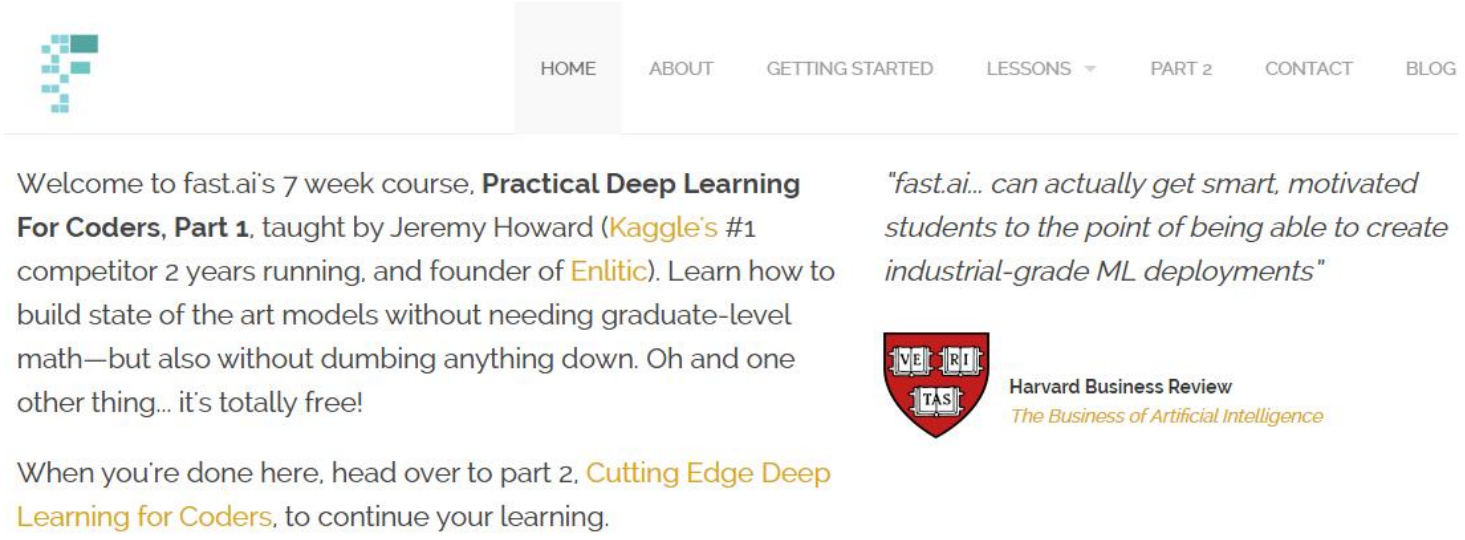

Courses

deeplearning.ai

by Andrew Ng

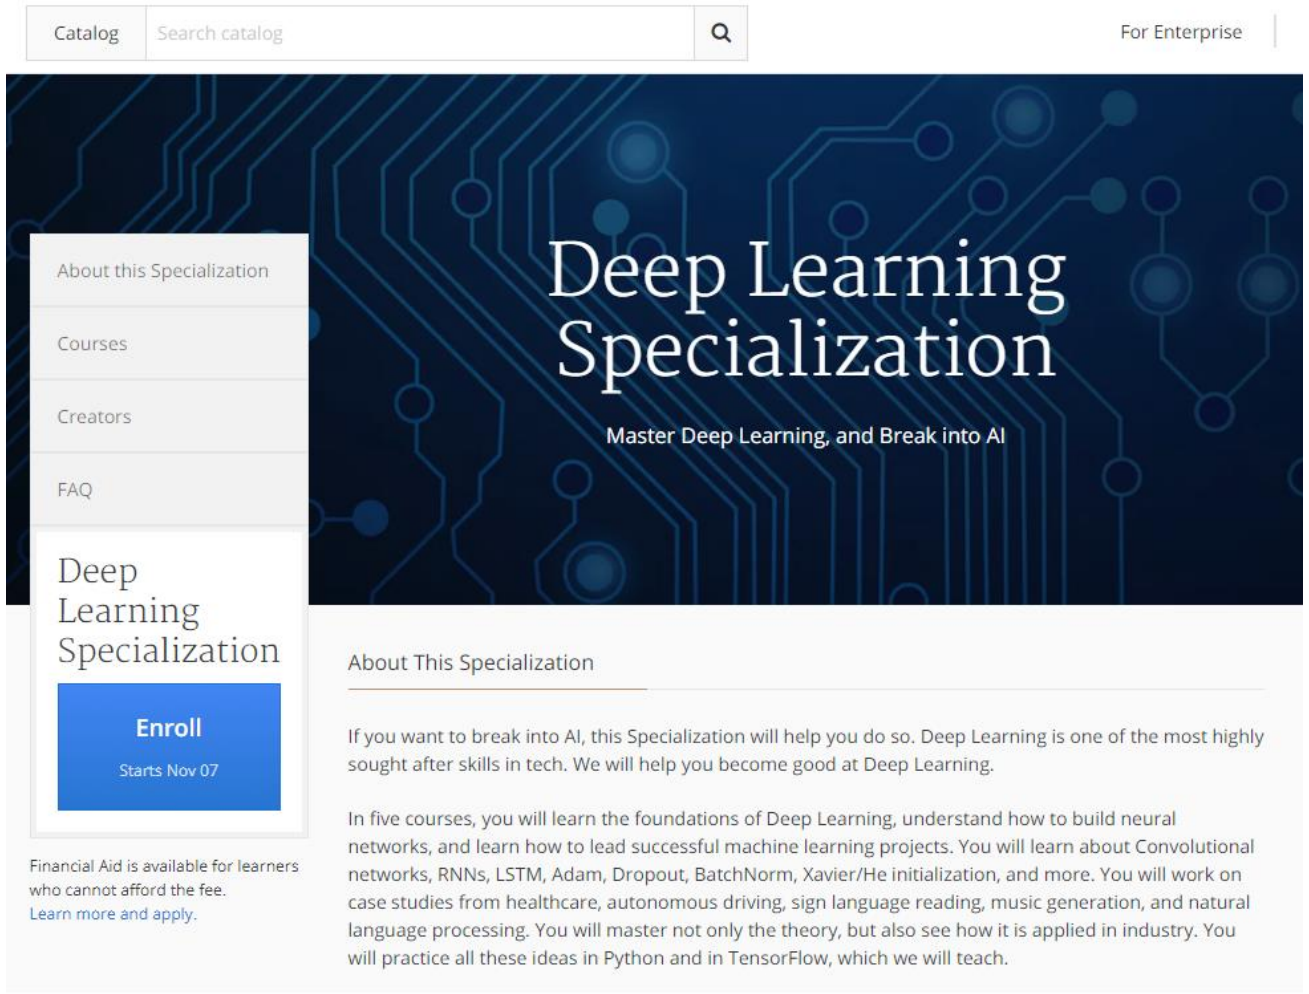

Books

Goodfellow
Bengio
& Courville

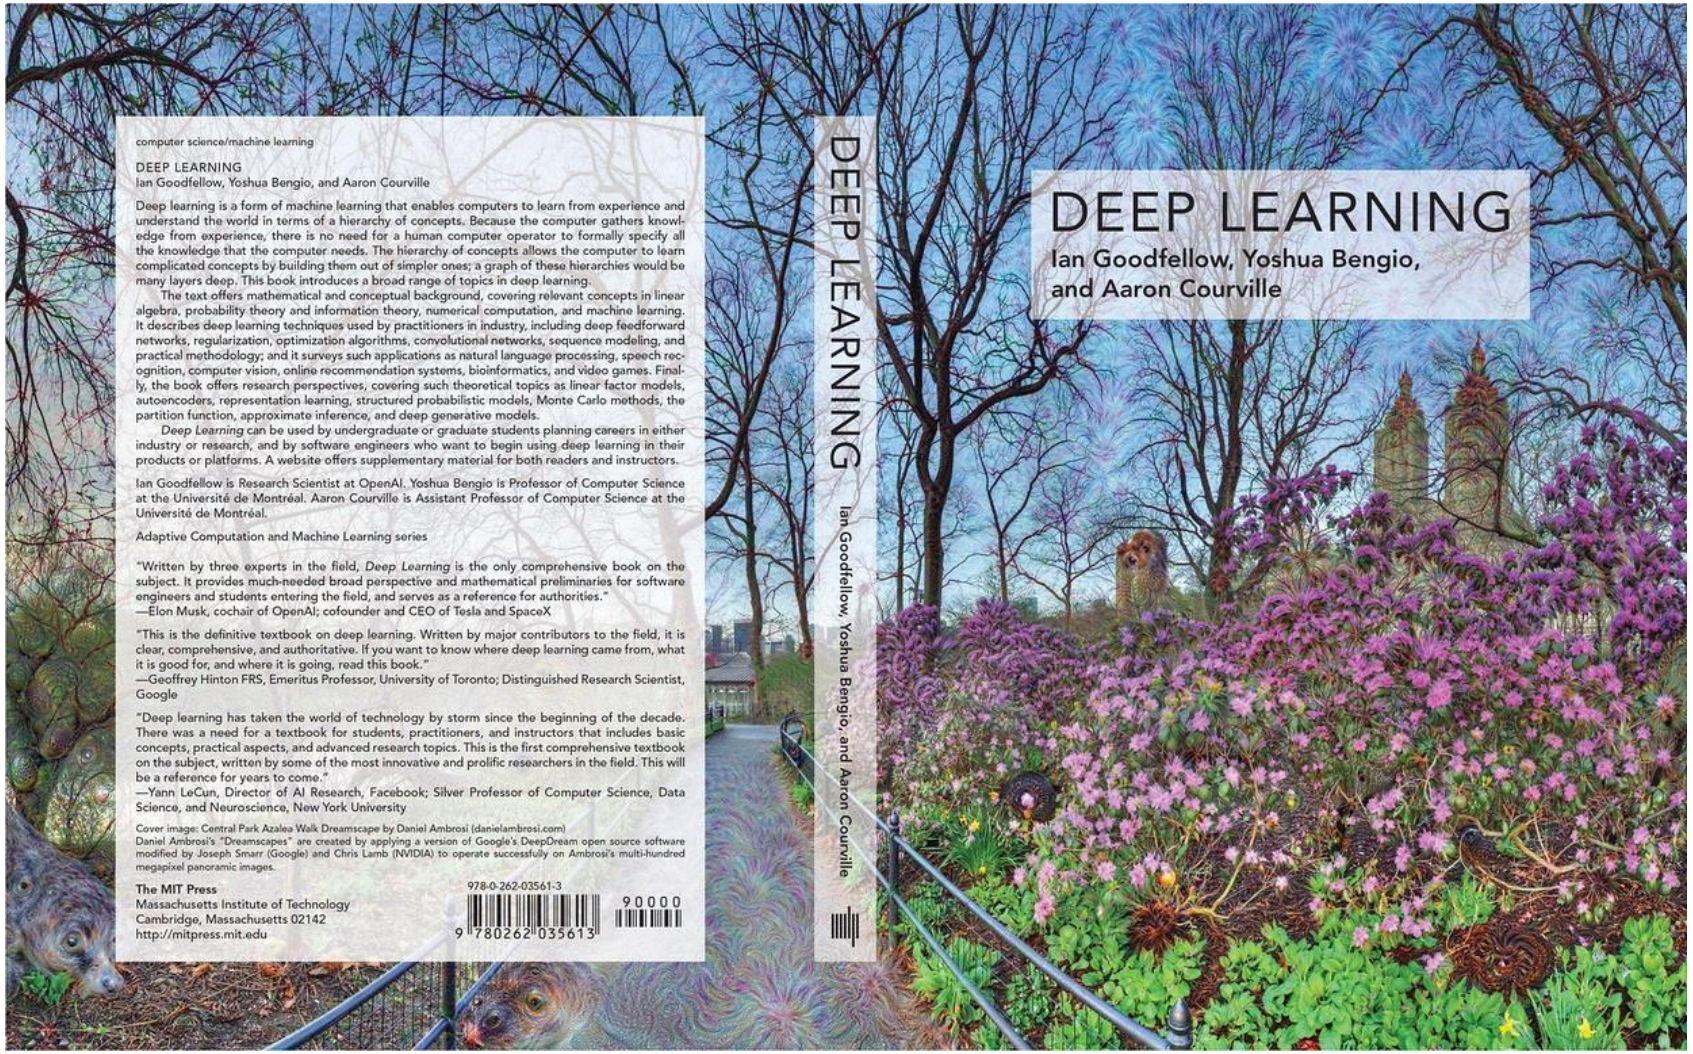

Books

Michael Nielsen

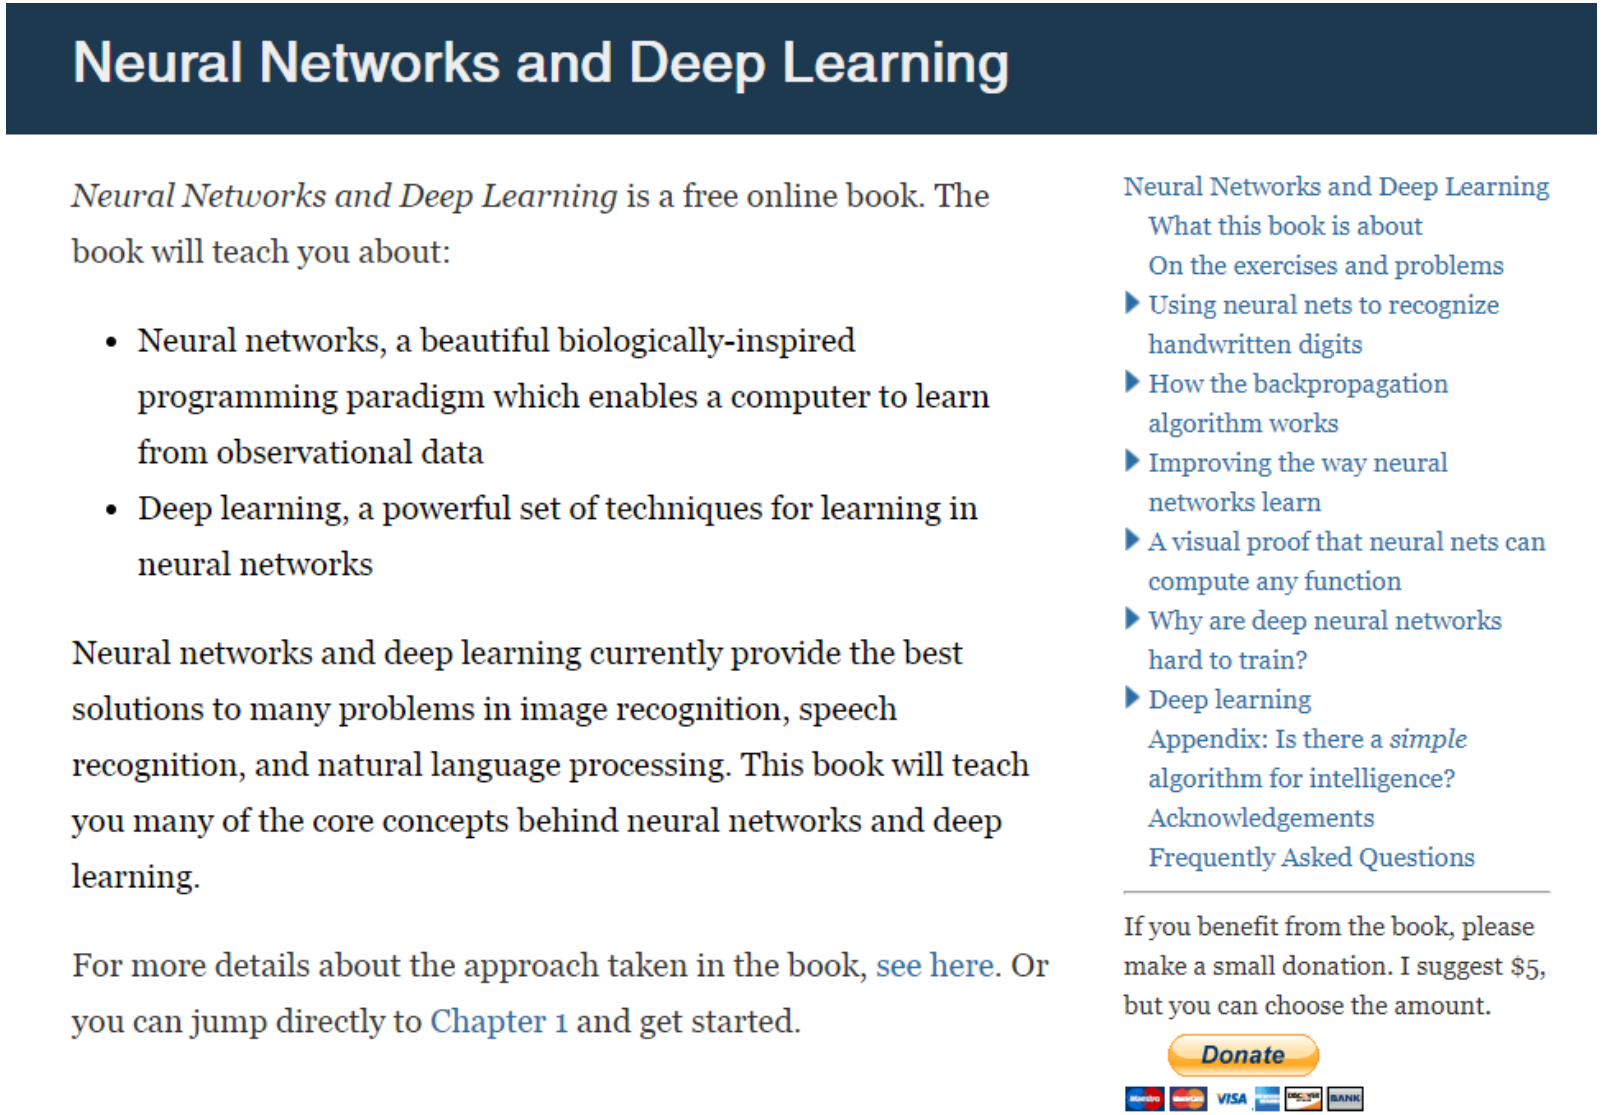

Neural Networks and Deep Learning

Neural Networks and Deep Learning is a free online book. The book will teach you about:

- Neural networks, a beautiful biologically-inspired programming paradigm which enables a computer to learn from observational data
- Deep learning, a powerful set of techniques for learning in neural networks

Neural networks and deep learning currently provide the best solutions to many problems in image recognition, speech recognition, and natural language processing. This book will teach you many of the core concepts behind neural networks and deep learning.

For more details about the approach taken in the book, see here. Or you can jump directly to Chapter 1 and get started.

If you benefit from the book, please make a small donation. I suggest $5, but you can choose the amount.

Donate

Papers

★ Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, “Deep Learning.” *Nature (2015)*

Most Cited Deep Learning Papers

https://github.com/terryum/awesome-deep-learning-papers

Deep Learning Papers Reading Roadmap

https://github.com/songrotek/Deep-Learning-Papers-Reading-Roadmap

Demos

TensorFlow Playground

https://playground.tensorflow.org/

ConvNetJS

https://cs.stanford.edu/people/karpathy/convnetjs/

Quick, Draw!

https://quickdraw.withgoogle.com/

“Software is eating the world, but AI is going to eat software.”

– Jensen Huang
(CEO, Nvidia)

Thank you

@samarhdr