Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Machine Vision

Slides available at `jonkrohn.com/talks`

March 2nd, 2022

Jon Krohn

Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Outline

1. Convolutional Layers
2. Convolutional Neural Networks
3. Much Deeper CNNs
4. Transfer Learning
5. Object Detection
6. Image Segmentation
7. Capsule Networks

Outline

1. Convolutional Layers
2. Convolutional Neural Networks
3. Much Deeper CNNs
4. Transfer Learning
5. Object Detection
6. Image Segmentation
7. Capsule Networks

Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Outline

1. Convolutional Layers
2. Convolutional Neural Networks
3. Much Deeper CNNs
4. Transfer Learning
5. Object Detection
6. Image Segmentation
7. Capsule Networks

Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Outline

1. Convolutional Layers
2. Convolutional Neural Networks
3. Much Deeper CNNs
4. Transfer Learning
5. Object Detection
6. Image Segmentation
7. Capsule Networks

Outline

1. Convolutional Layers
2. Convolutional Neural Networks
3. Much Deeper CNNs
4. Transfer Learning
5. Object Detection
6. Image Segmentation
7. Capsule Networks

Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Outline

1. Convolutional Layers
2. Convolutional Neural Networks
3. Much Deeper CNNs
4. Transfer Learning
5. Object Detection
6. Image Segmentation
7. Capsule Networks

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

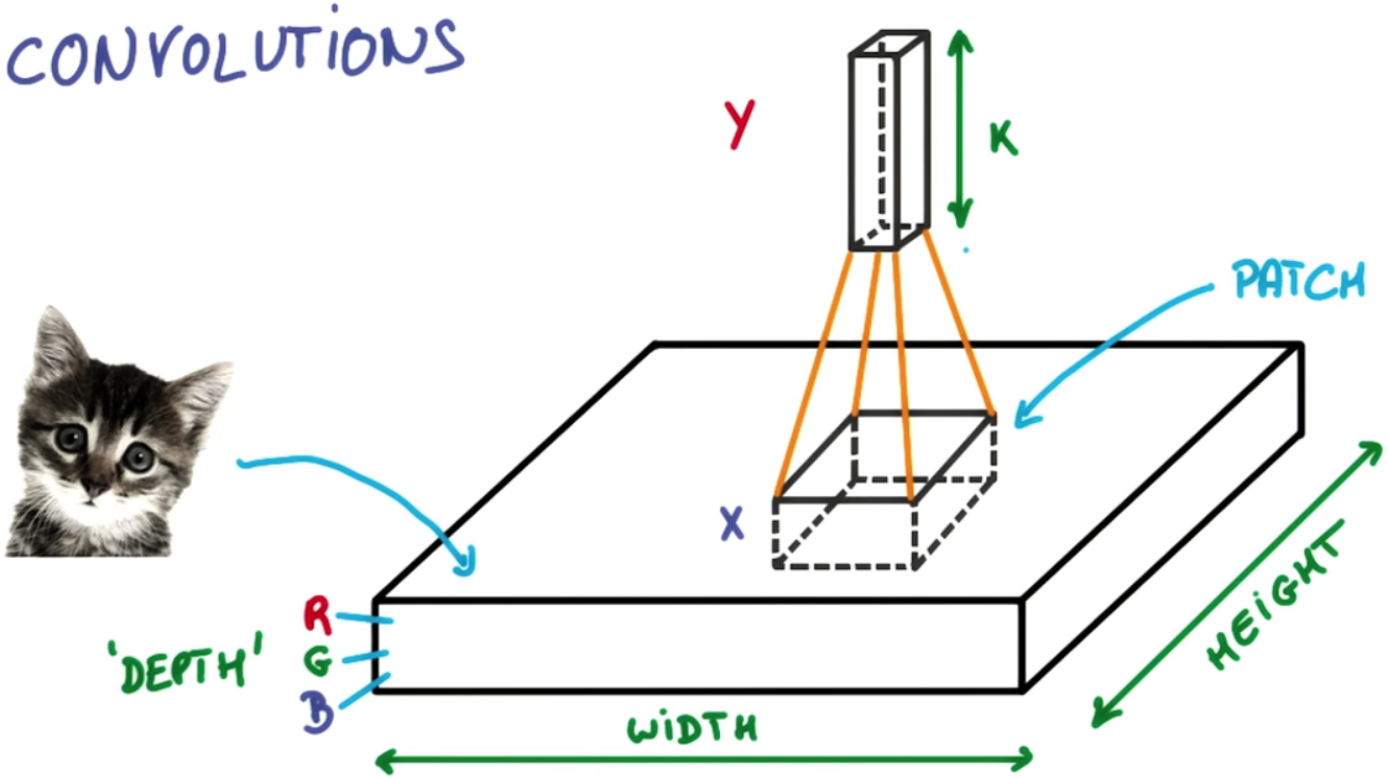
CONVOLUTIONS
Y
K
PATCH
X
'DEPTH'
R
G
B
WIDTH
HEIGHT

DeepVis

[deepvis]

Convolution Demo

from the illustrious [Andrej Karpathy]

$$\text{Activation map} = \frac{D - F + 2P}{S} + 1$$

...where:

- D is the size of the image (either width or height, depending on whether you're calculating the width or height of the activation map)
- F is the size of the filter
- P is the amount of padding, and
- and S is the stride length.

$$\text{Activation map} = \frac{D - F + 2P}{S} + 1$$

...where:

- D is the size of the image (either width or height, depending on whether you're calculating the width or height of the activation map)
- F is the size of the filter
- P is the amount of padding, and
- and S is the stride length.

$$\text{Activation map} = \frac{D - F + 2P}{S} + 1$$

...where:

- D is the size of the image (either width or height, depending on whether you're calculating the width or height of the activation map)
- F is the size of the filter
- P is the amount of padding, and
- and S is the stride length.

$$\text{Activation map} = \frac{D - F + 2P}{S} + 1$$

...where:

- D is the size of the image (either width or height, depending on whether you're calculating the width or height of the activation map)
- F is the size of the filter
- P is the amount of padding, and
- and S is the stride length.

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

LeNet-5

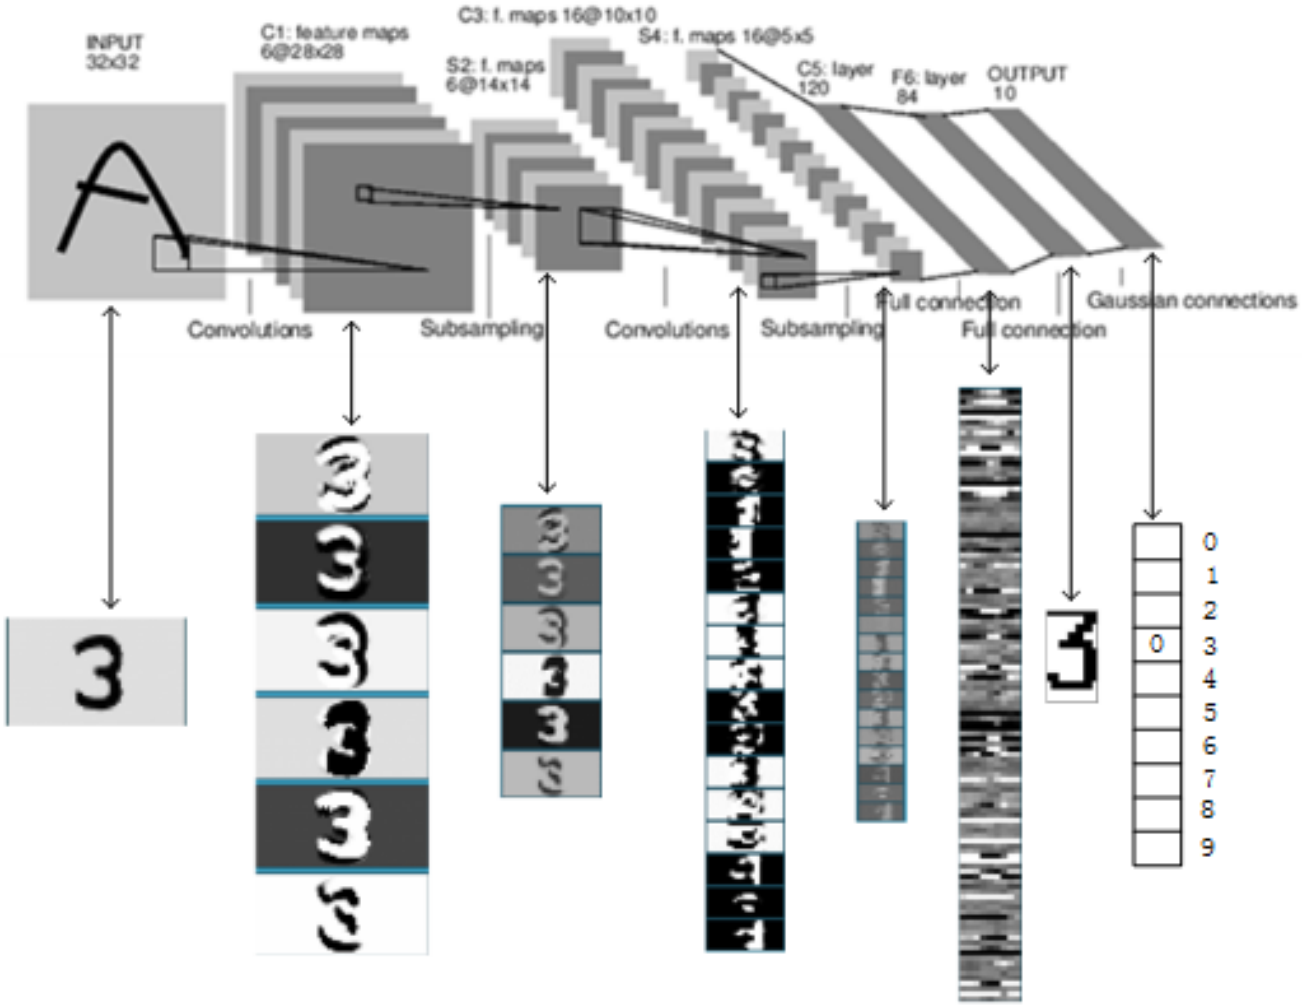

Jon Krohn

let's make our [deep net] *convolutional*!

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

AlexNet

Conv-Pool Blocks

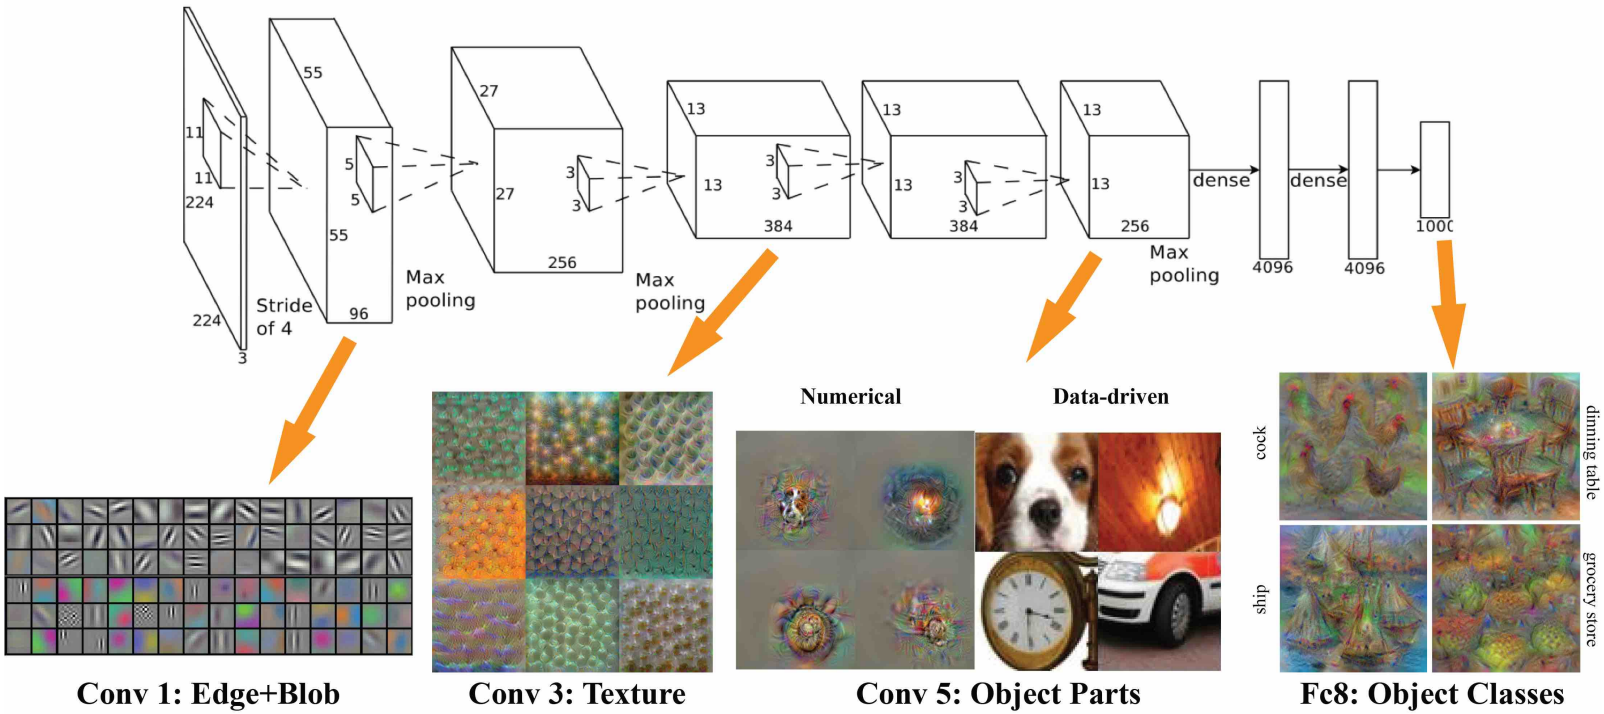

[AlexNet] from scratch

Jon Krohn

VGGNet (Simonyan & Zisserman, 2014)

Exercise III

- build VGGNet from AlexNet notebook
- be able to verbalize all Arsenal (Theory I-IV) items

VGGNet (Simonyan & Zisserman, 2014)

Exercise III

- build VGGNet from AlexNet notebook
- be able to verbalize all Arsenal (Theory I-IV) items

Residual Networks

Hardt & Ma, 2018

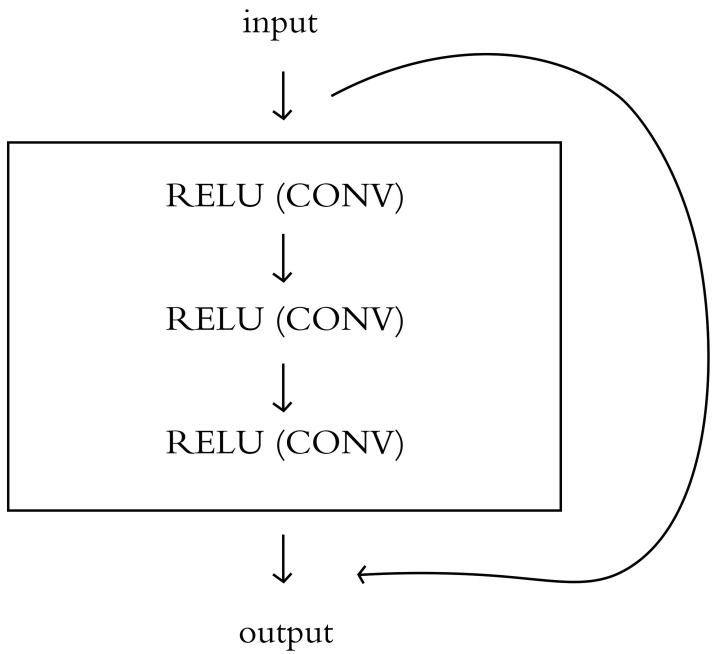

Residual Networks

Hardt & Ma, 2018

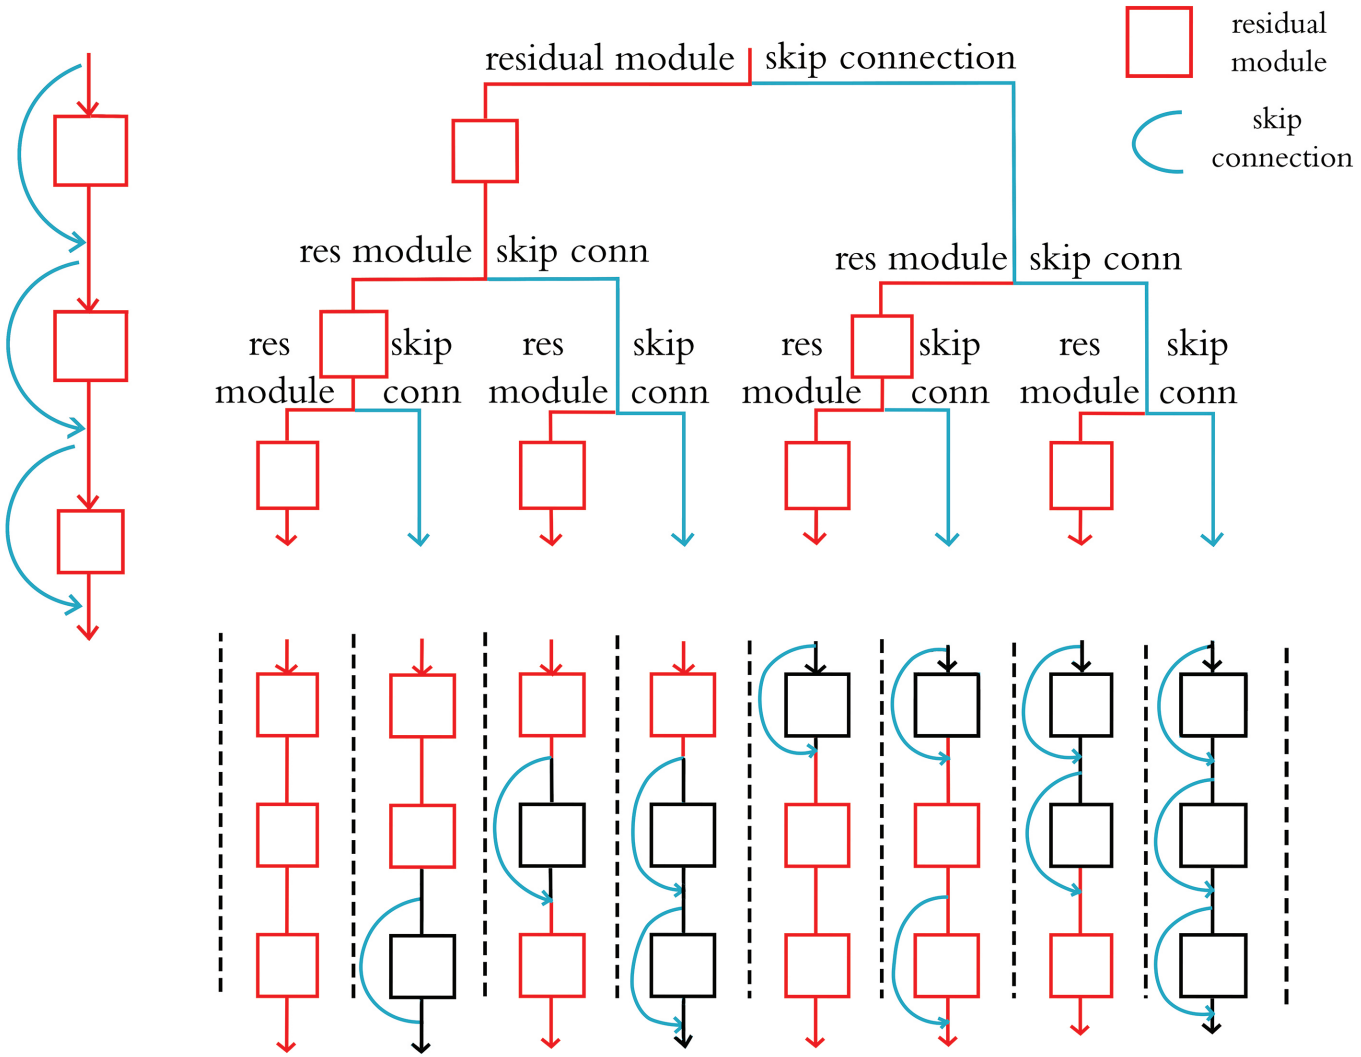

Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

[transfer learning Jupyter notebook]

Other examples:

- [toy-sized]
- [pre-trained model weights in Keras]
- [beefy bottleneck features example]

[transfer learning Jupyter notebook]

Other examples:

- [toy-sized]
- [pre-trained model weights in Keras]
- [beefy bottleneck features example]

[transfer learning Jupyter notebook]

Other examples:

- [toy-sized]
- [pre-trained model weights in Keras]
- [beefy bottleneck features example]

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

Machine Vision Applications

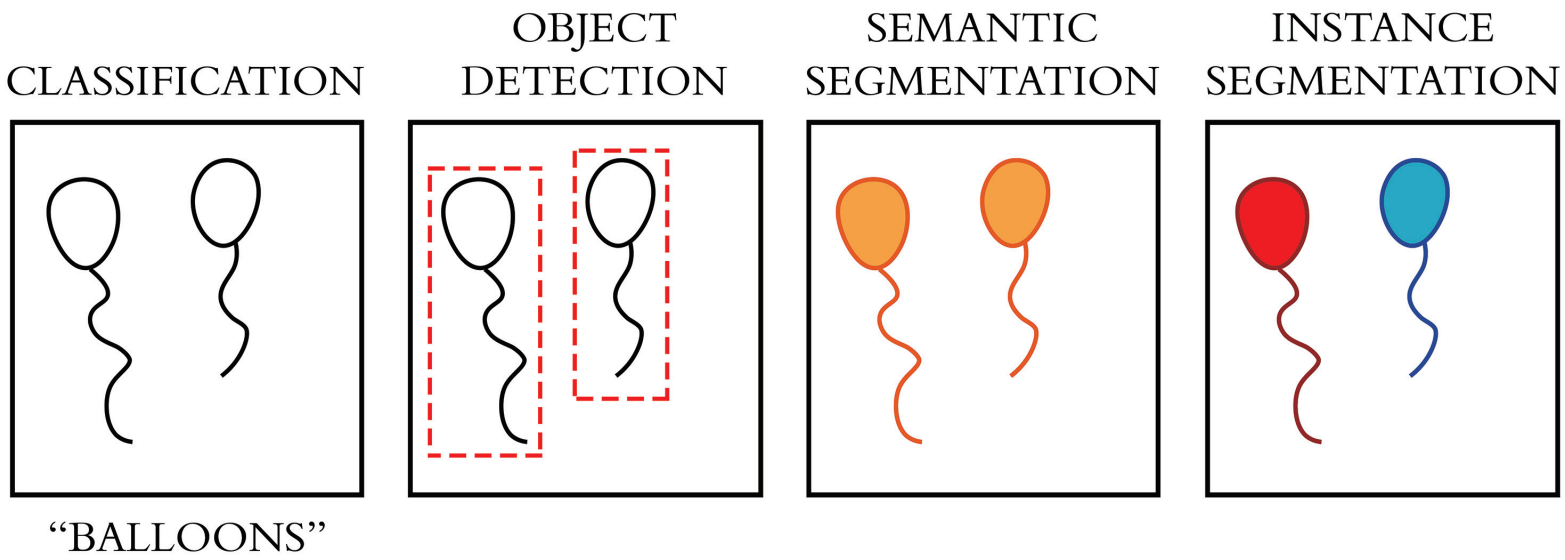

Battle algorithm

Artificial intelligence is changing every aspect of war

A new type of arms race could be on the cards

Print edition | Science and technology ›

Sep 7th 2019

Krohn

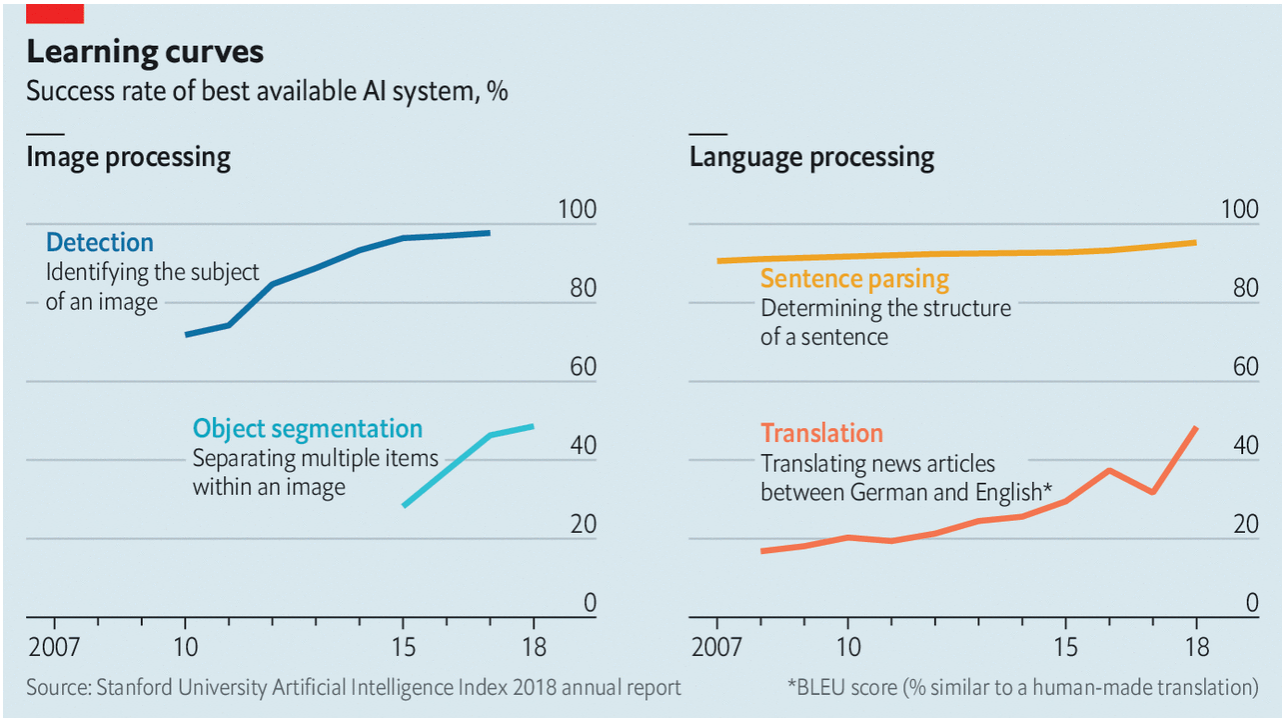

The Economist

Object Detection

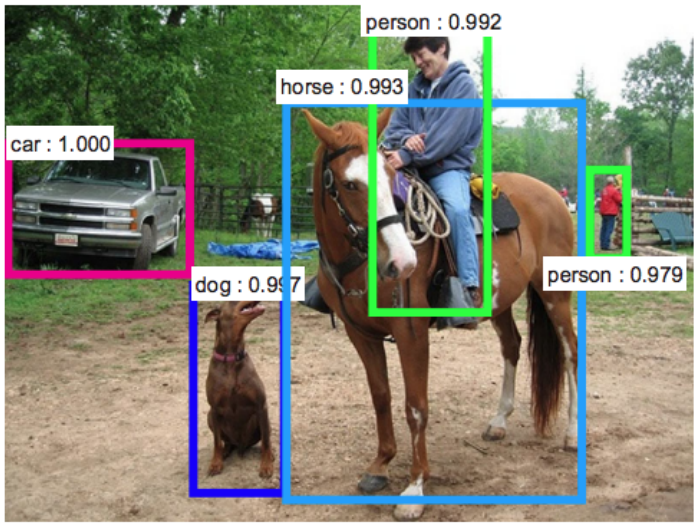

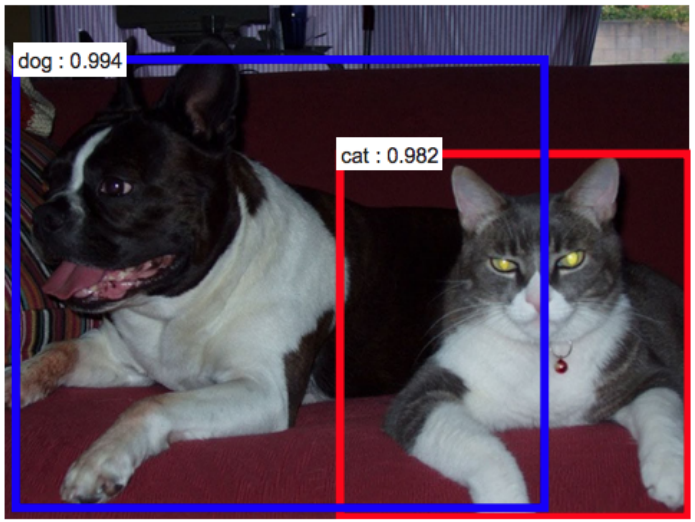

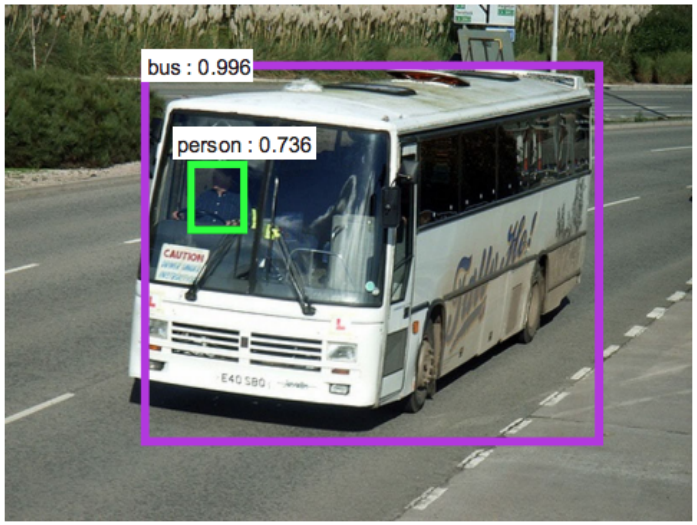

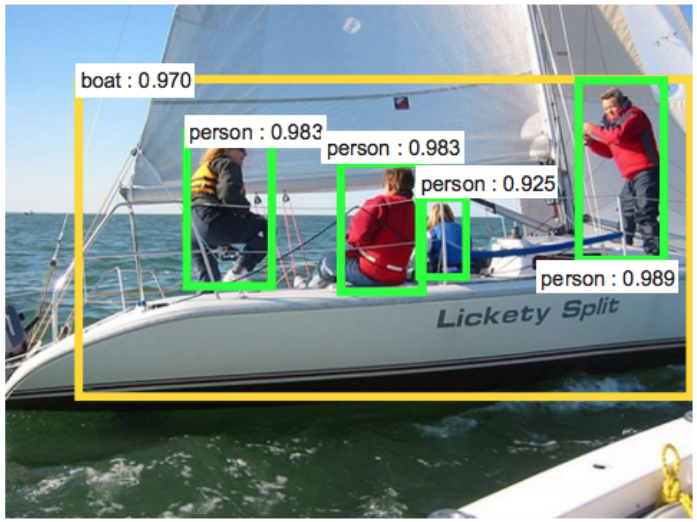

Typical Process

1. Identify region of interest (ROI)
2. Perform automatic feature extraction on ROI
3. Classify ROI

Typical Process

1. Identify region of interest (ROI)
2. Perform automatic feature extraction on ROI
3. Classify ROI

Typical Process

1. Identify region of interest (ROI)
2. Perform automatic feature extraction on ROI
3. Classify ROI

Seminal Architectures

- R-CNN (Girshick et al., 2013)
- Fast R-CNN (Girshick et al., 2015)
- Faster R-CNN (Ren et al., 2015)
- YOLO, YOLO9000 & YOLOv3 (Redmon et al., 2015-8)

Seminal Architectures

- R-CNN (Girshick et al., 2013)
- Fast R-CNN (Girshick et al., 2015)
- Faster R-CNN (Ren et al., 2015)
- YOLO, YOLO9000 & YOLOv3 (Redmon et al., 2015-8)

Seminal Architectures

- R-CNN (Girshick et al., 2013)
- Fast R-CNN (Girshick et al., 2015)
- Faster R-CNN (Ren et al., 2015)
- YOLO, YOLO9000 & YOLOv3 (Redmon et al., 2015-8)

Seminal Architectures

- R-CNN (Girshick et al., 2013)
- Fast R-CNN (Girshick et al., 2015)
- Faster R-CNN (Ren et al., 2015)
- YOLO, YOLO9000 & YOLOv3 (Redmon et al., 2015-8)

[YOLOv3 Jupyter notebook]

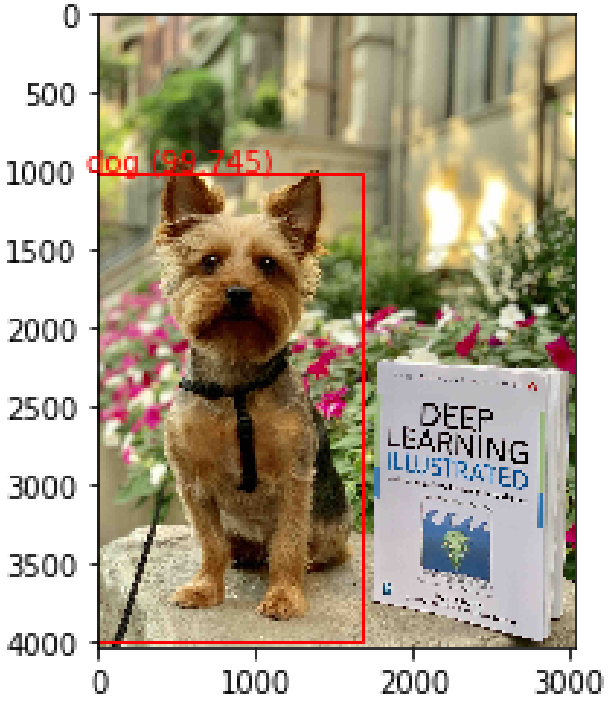

GitHub repositories:

- [Mask R-CNN]
- [RetinaNet]
- [YOLOv3]

[YOLOv3 Jupyter notebook]

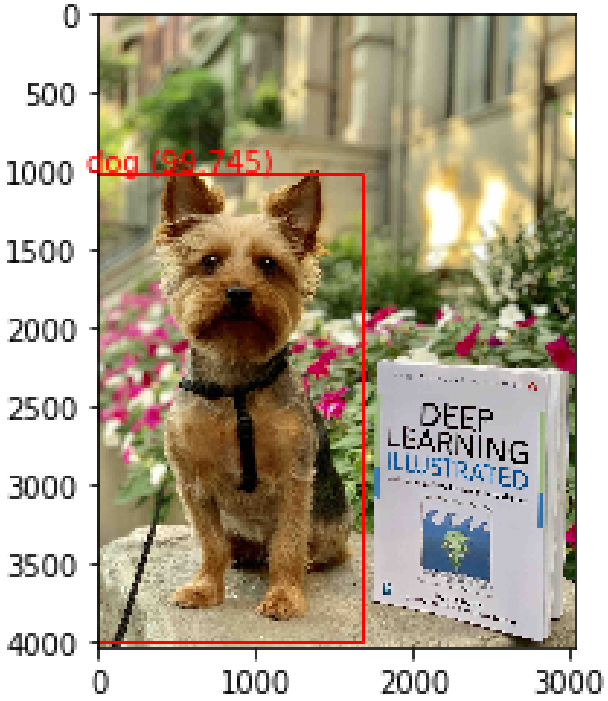

GitHub repositories:

- [Mask R-CNN]
- [RetinaNet]
- [YOLOv3]

[YOLOv3 Jupyter notebook]

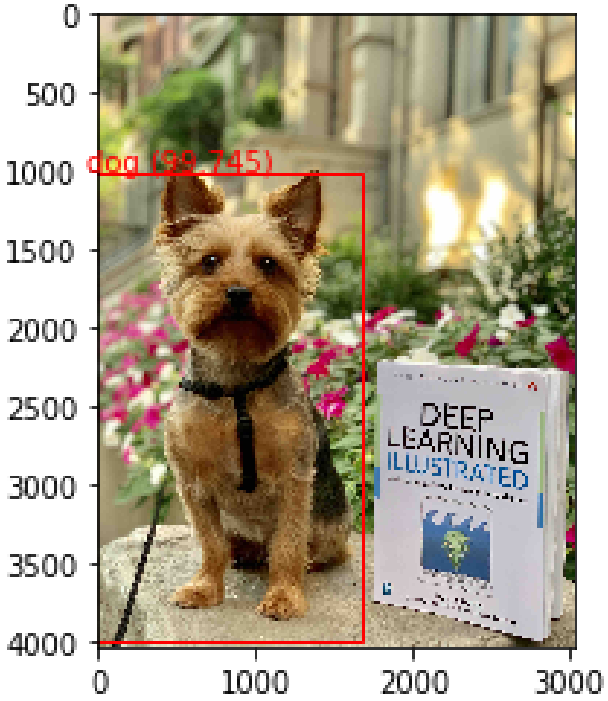

GitHub repositories:

- [Mask R-CNN]
- [RetinaNet]
- [YOLOv3]

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

Machine Vision Applications

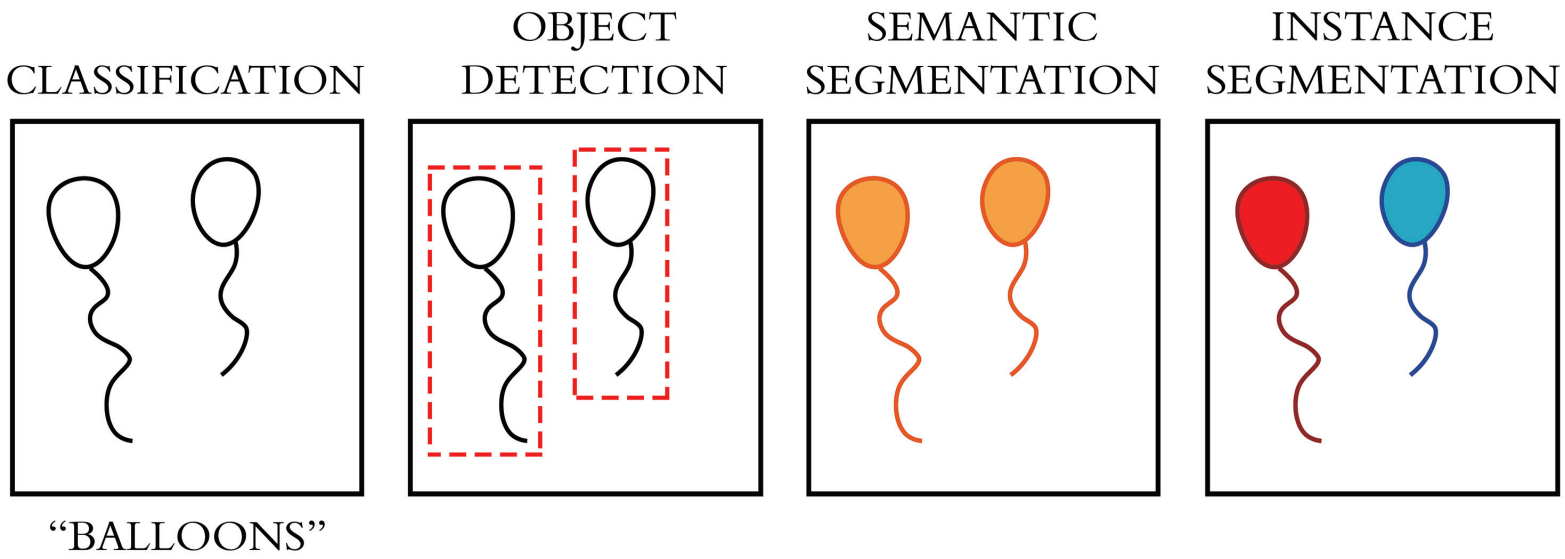

Image Segmentation

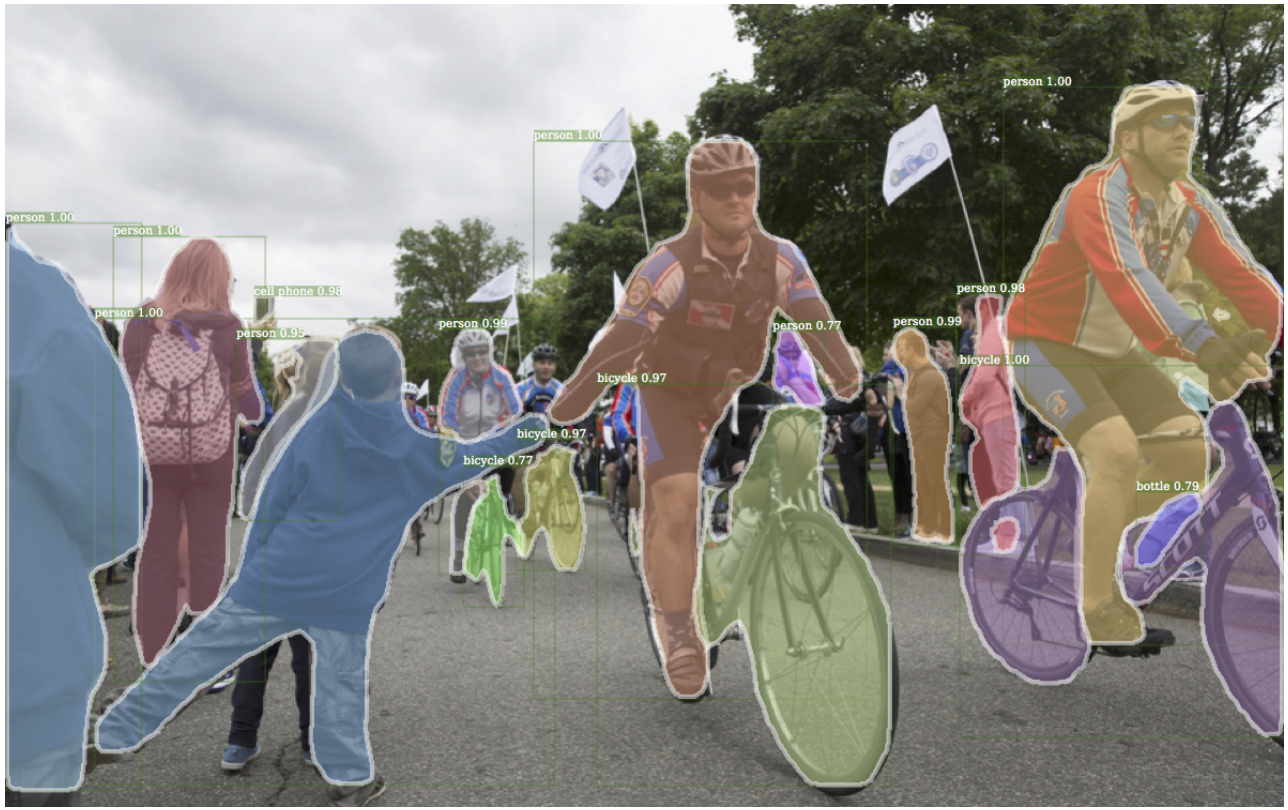

Seminal architectures:

- Mask R-CNN
- U-Net

Jon Krohn

Image Segmentation

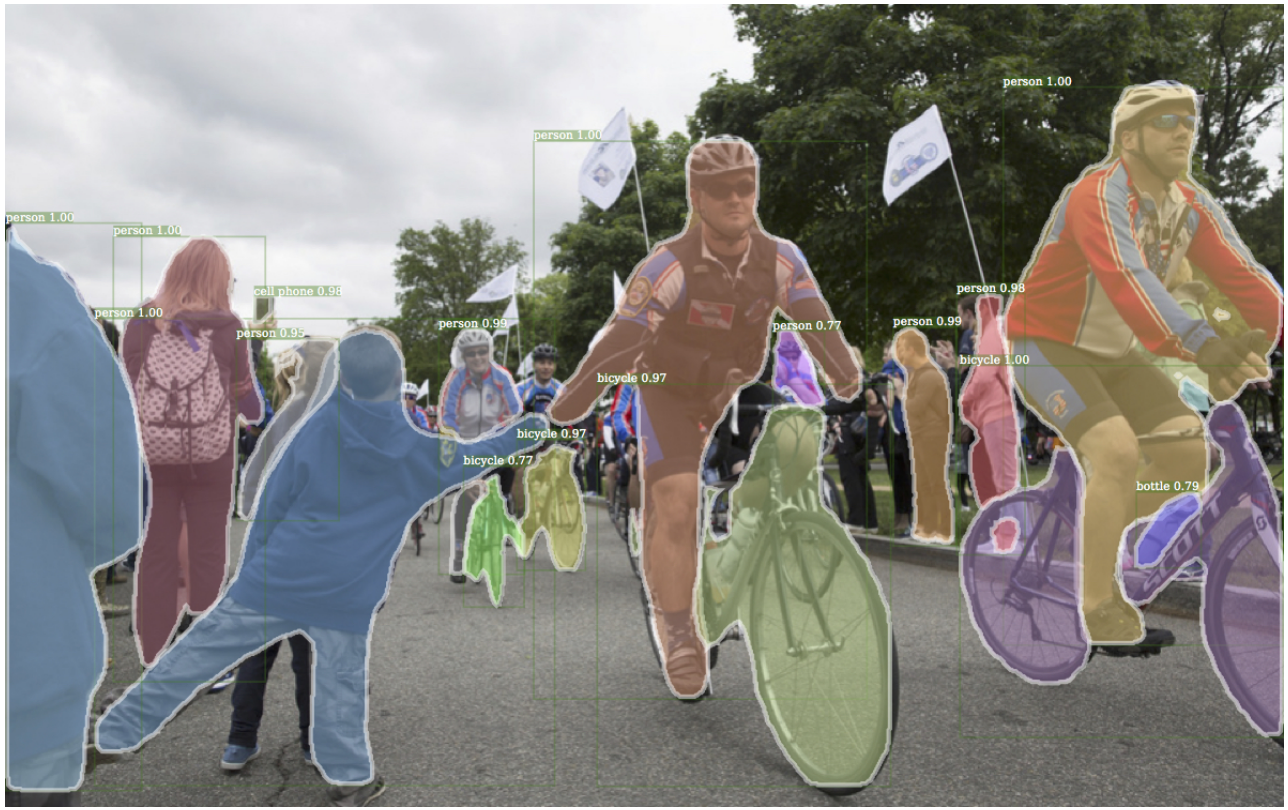

Seminal architectures:

- Mask R-CNN
- U-Net

Mask R-CNN

(He et FAIR, 2017)

1. Faster R-CNN model proposes ROIs
2. ROI classifier predicts class of object in ROI and refines bounding box
3. extract CNN's feature maps from within bounding box
4. fully CNN model outputs object-specific mask

Mask R-CNN

(He et FAIR, 2017)

1. Faster R-CNN model proposes ROIs
2. ROI classifier predicts class of object in ROI and refines bounding box
3. extract CNN's feature maps from within bounding box
4. fully CNN model outputs object-specific mask

Mask R-CNN

(He et FAIR, 2017)

1. Faster R-CNN model proposes ROIs
2. ROI classifier predicts class of object in ROI and refines bounding box
3. extract CNN's feature maps from within bounding box
4. fully CNN model outputs object-specific mask

Mask R-CNN

(He et FAIR, 2017)

1. Faster R-CNN model proposes ROIs
2. ROI classifier predicts class of object in ROI and refines bounding box
3. extract CNN's feature maps from within bounding box
4. fully CNN model outputs object-specific mask

U-Net

Ronneburger et al., 2015

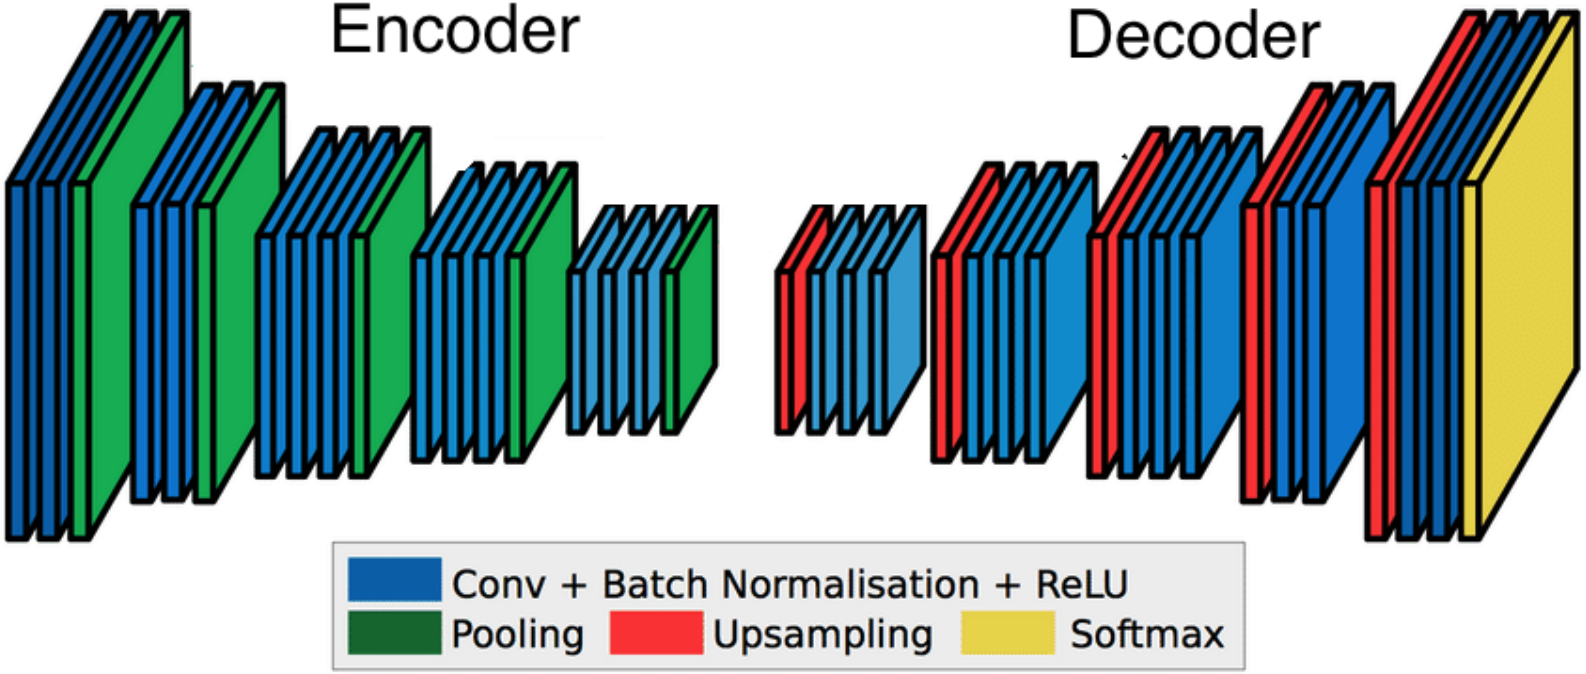

Contracting Path + Expanding Path

U-Net Transfer Learning

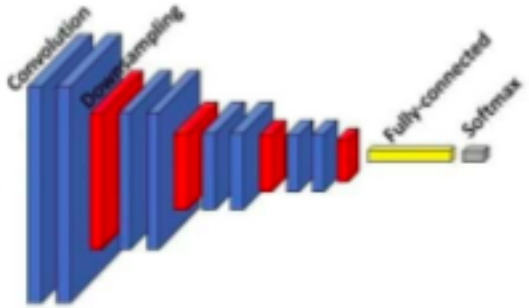

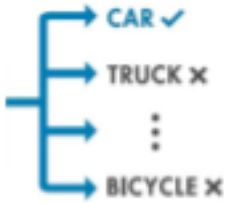

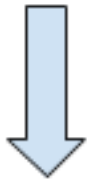

Transfer of weights

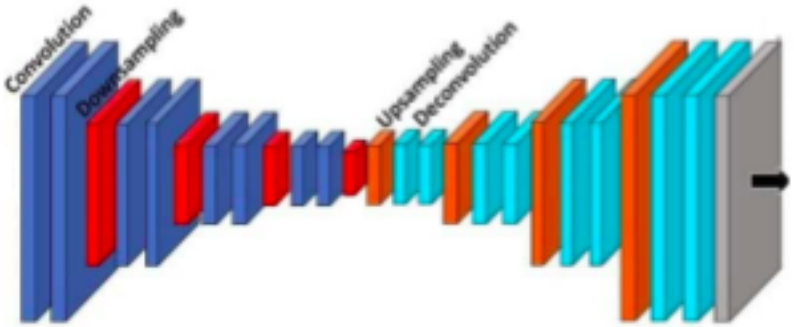

VGGNet U-Net

Validation output after 1 & 9 epochs

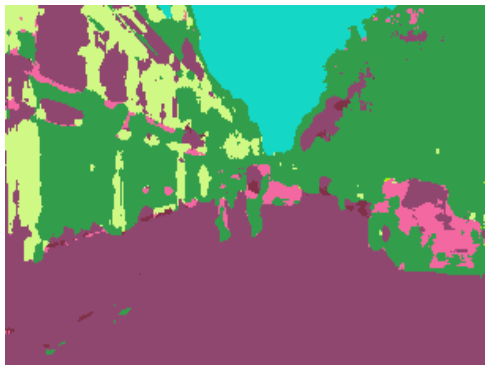

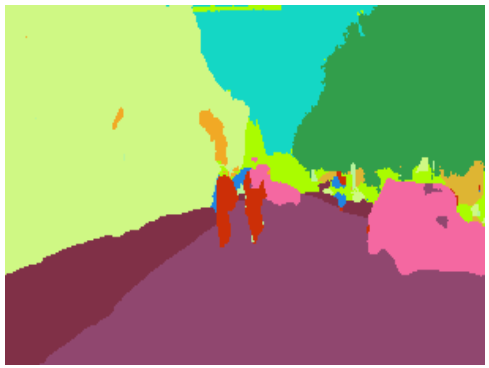

[Jupyter notebook]

ResNet U-Net

VGG after 9 epochs & ResNet after 4

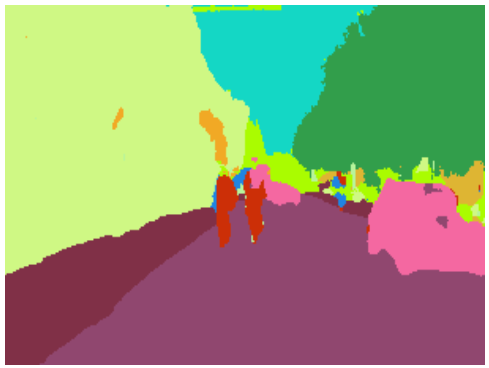

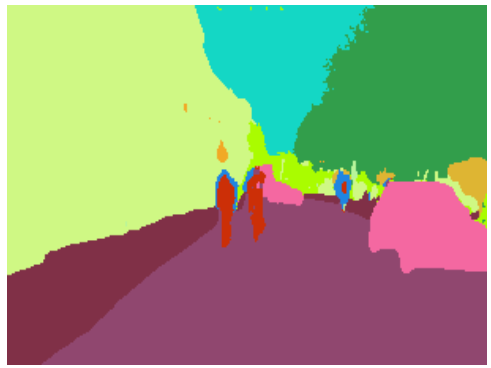

[Jupyter notebook]

Conv Layers
ConvNets
Deeper CNNs
Transfer Learning
Object Detection
Image Segmentation
Capsule Nets

Outline

1 Convolutional Layers
2 Convolutional Neural Networks
3 Much Deeper CNNs
4 Transfer Learning
5 Object Detection
6 Image Segmentation
7 Capsule Networks

Capsule Networks

Sabour & Hinton, 2017