DATA 442: Neural Networks & Deep Learning

Dan Runfola – danr@wm.edu

icss.wm.edu/data442/

https://www.sandvine.com/hubfs/Sandvine_Redesign_2019/Downloads/Internet%20Phenomena/Internet%20Phenomena%20Report%20Q32019%2020190910.pdf

GLOBAL APPLICATION CATEGORY TRAFFIC SHARE

1 VIDEO STREA
60.6%(+2.9)

2 WEB
13.1%(-3.8)

3 GAMING
8.0%(0.2)

4 SOCIAL
6.1%(+1.1)

5 FILE SHARIN
4.2%(+1.4)

6 MARKETPLA
2.6%(-1.9)

7 SECURITY AN
1.6%(+0.2)

8 MESSAGING
1.6%(-0.1)

9 CLOUD
1.4%(+0.01)

10 AUDIO STREAMING
0.4%(-0.5) 0.3%(-0.1)

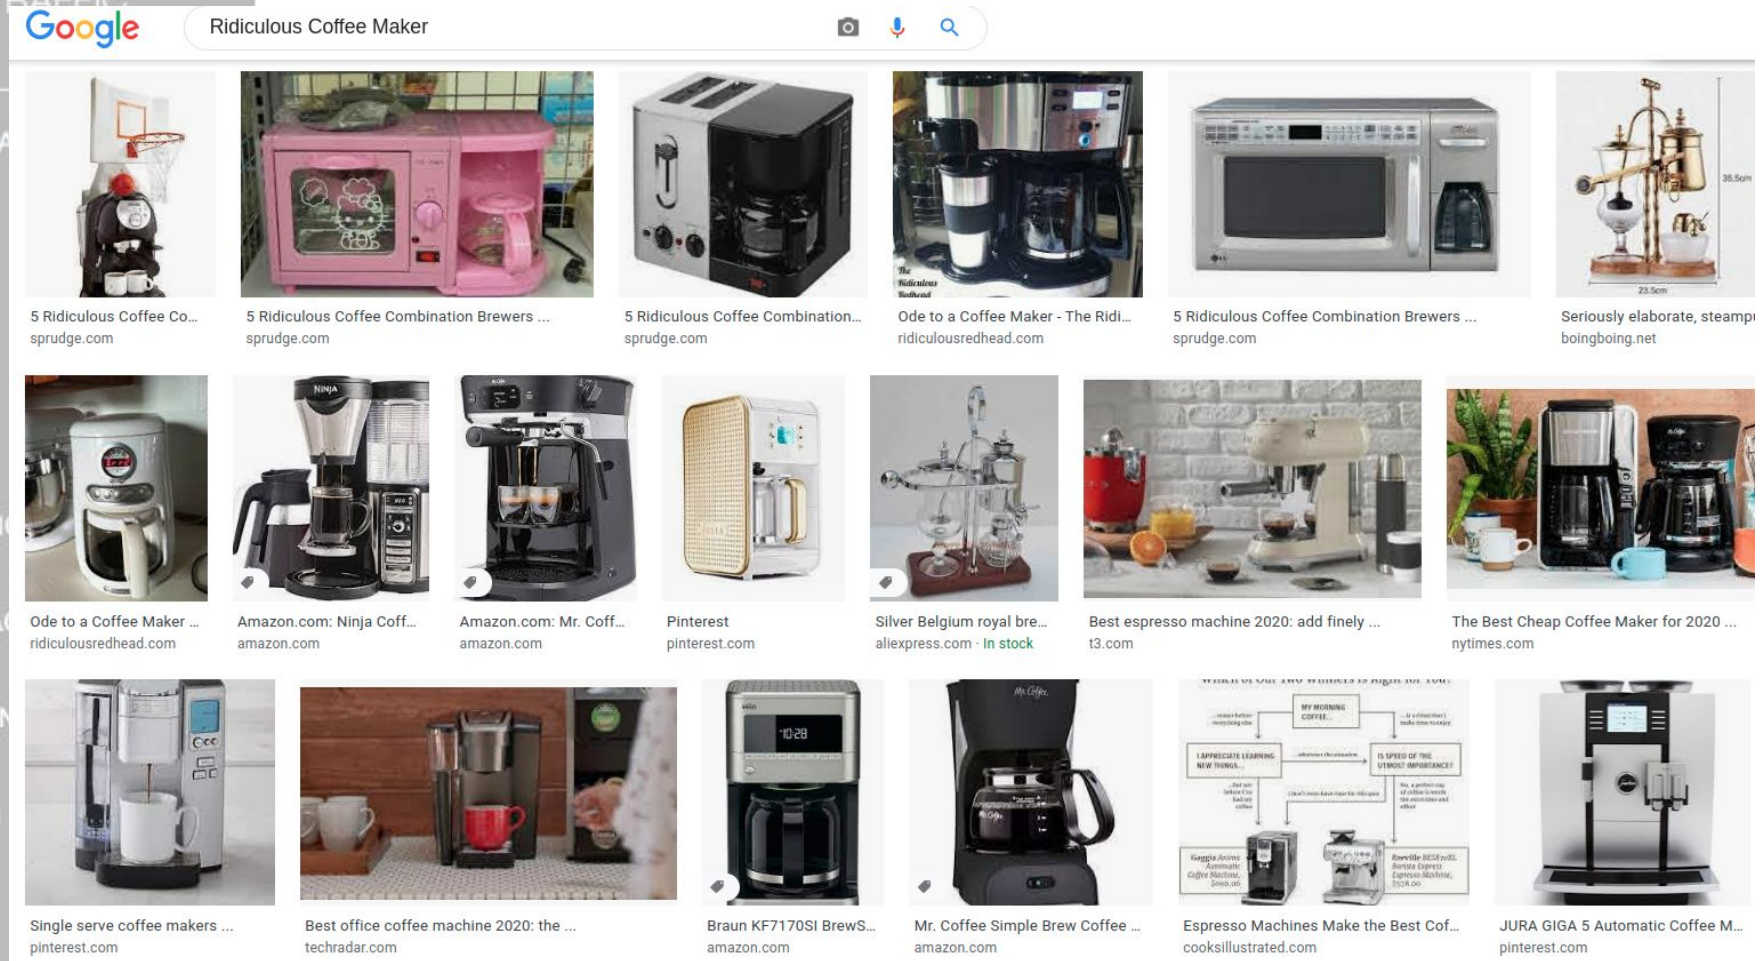

https://www.sandvine.com/hubfs/Sandvine_Redesign_2019/Downloads/Internet%20Phenomena/Internet%20Phenomena%20Report%20Q32019%2020190910.pdf

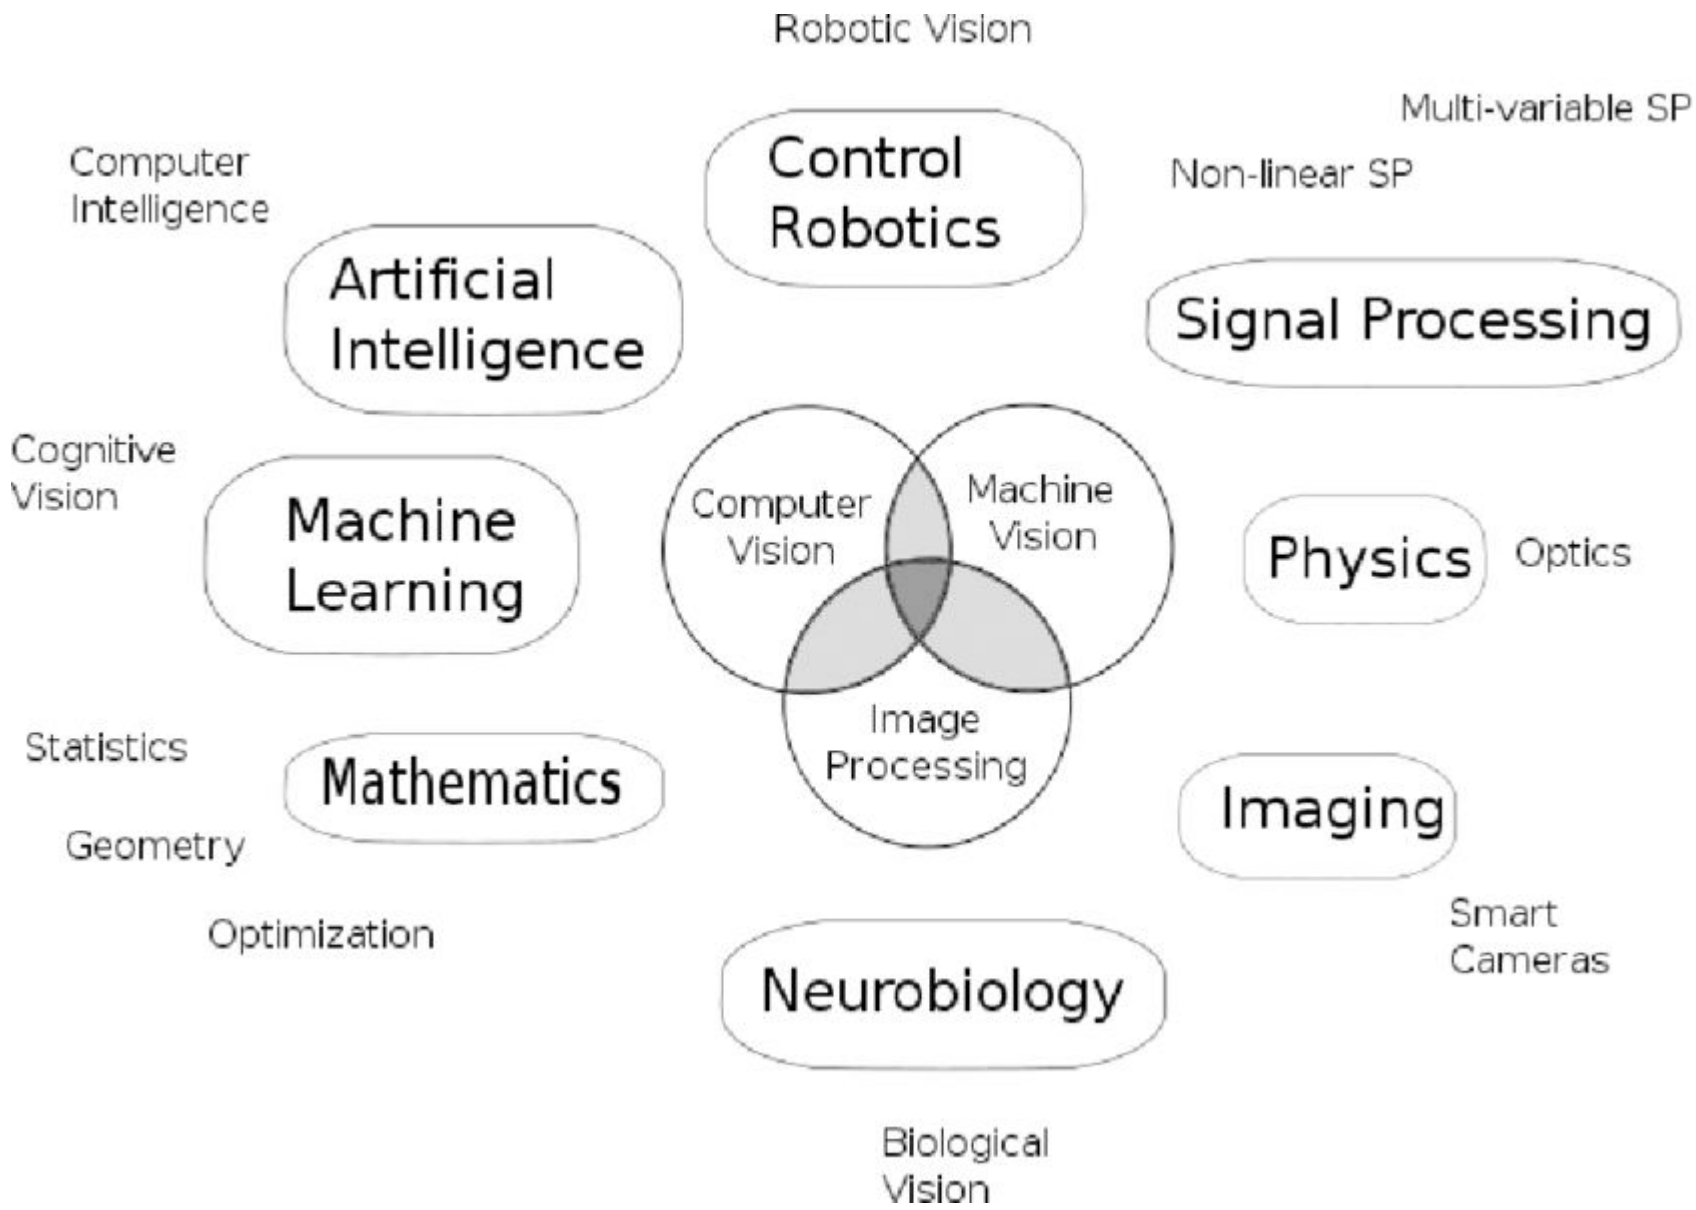

https://www.researchgate.net/figure/Computer-Vision-related-fields-22_fig5_318234623

Initiative for Computational Societal and Security Research

BASED OUT OF THE OLDEST PUBLIC RESEARCH UNIVERSITY IN THE UNITED STATES, THE ICSS IS DESIGNING THE FUTURE OF LARGE SCALE SOCIETAL MODELING.

RESEARCH THEMES →

PROJECTS & PEOPLE →

Featured Research

Location-based AI

BASIC RESEARCH ON THE USE OF MACHINE LEARNING WITH SATELLITE AND OTHER LOCATION DATA.

Relationships with Other Courses @ W&M

DATA 141 / CSCI 140 - Programming for Data Science

DATA 310 / CSCI 416 - Core machine learning class, teaches intermediate concepts of machine learning (traditional neural networks, SVM, gradient descent, cost functions).

DATA 442 (This Class) - Heavy focus on computer vision and deep learning neural network approaches.

A brief history of Computer Vision

https://www.nature.com/articles/eye2017226

A brief history of Computer Vision

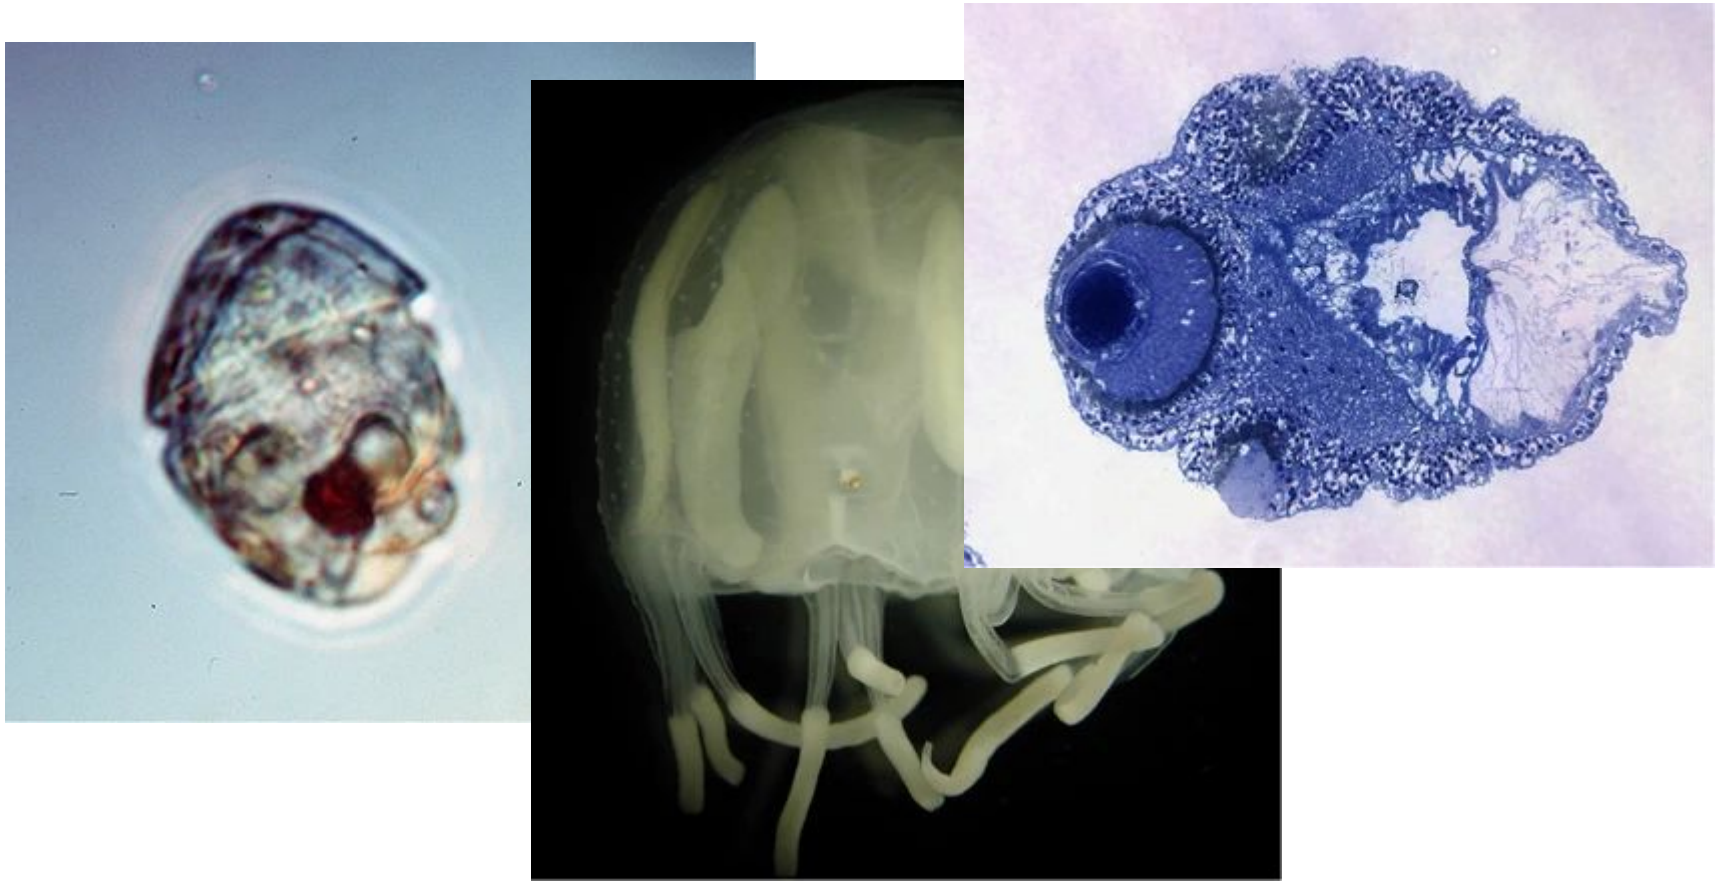

https://www.nature.com/articles/eye2017226

A brief history of Computer Vision

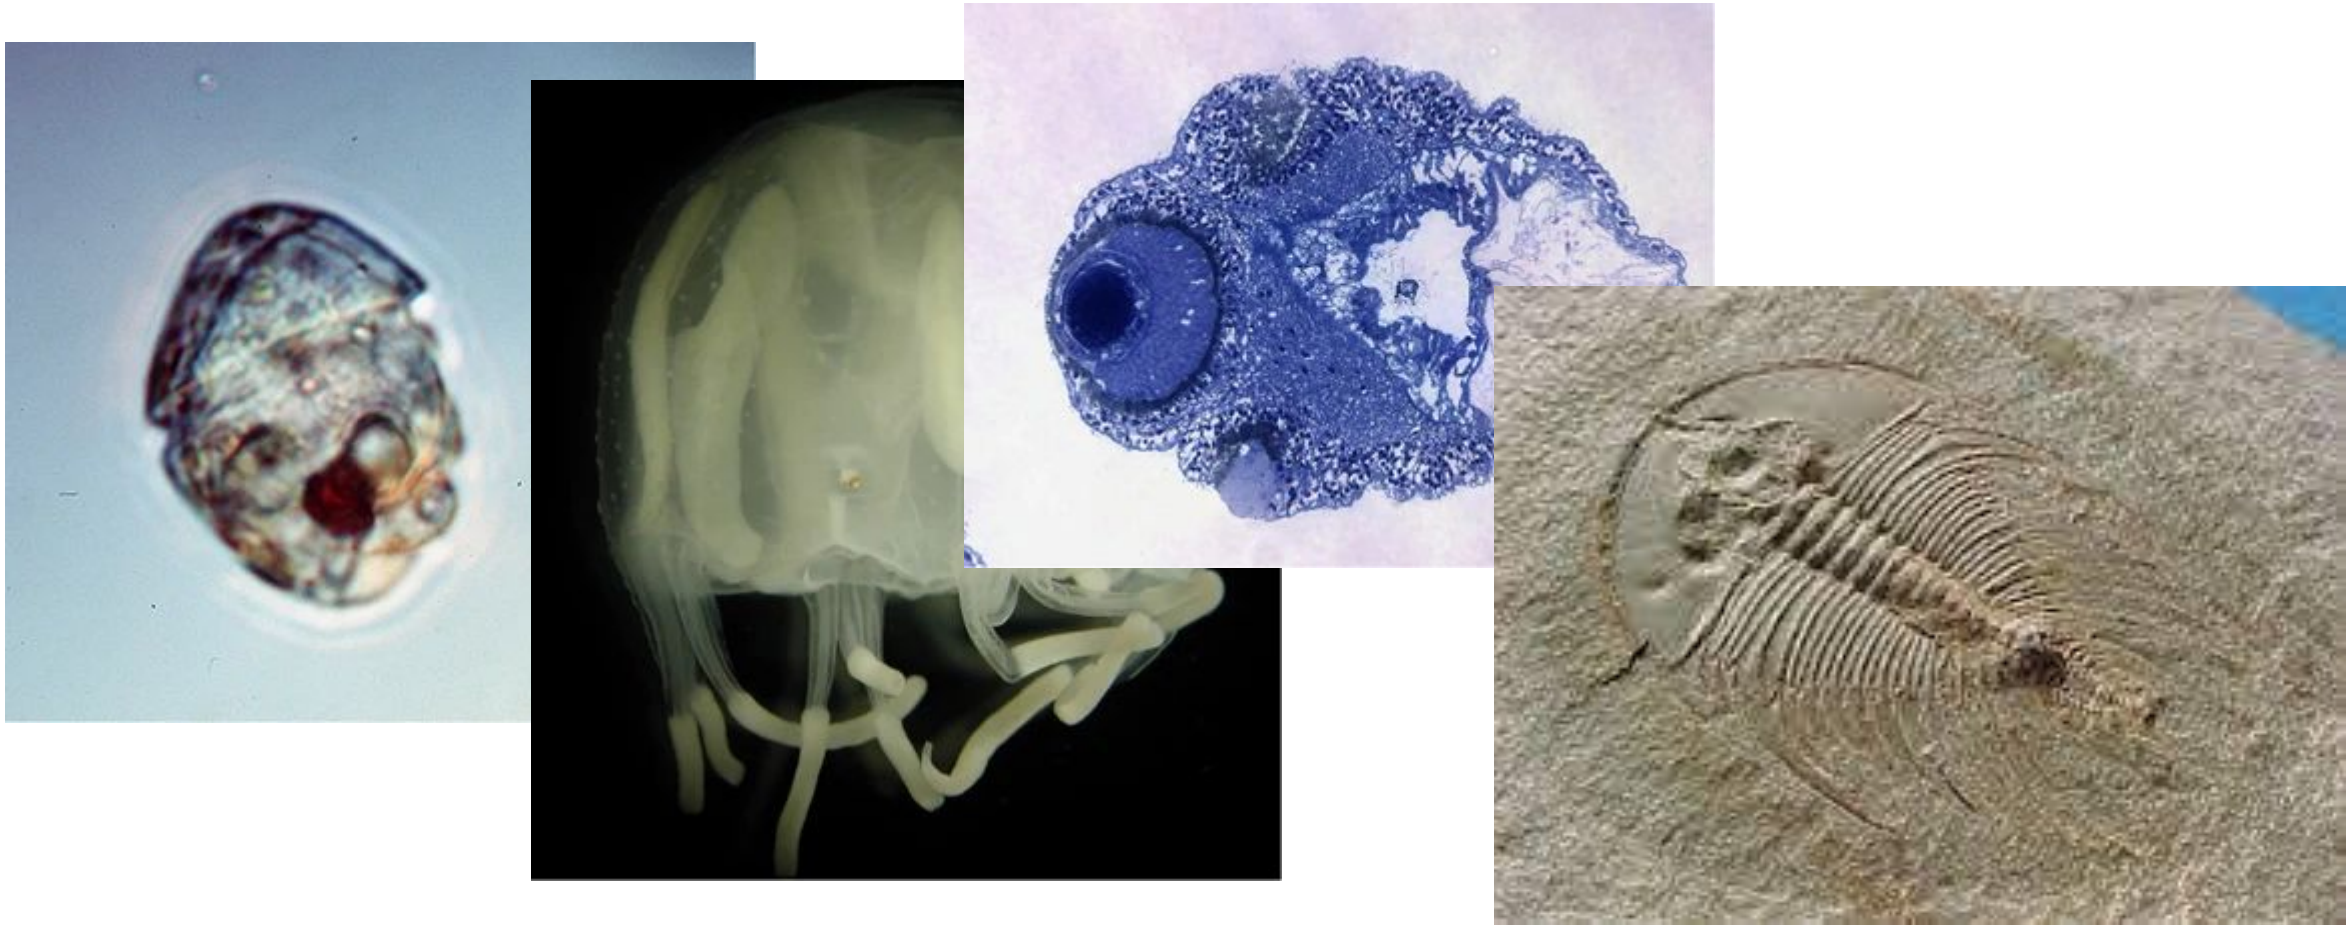

https://www.nature.com/articles/eye2017226

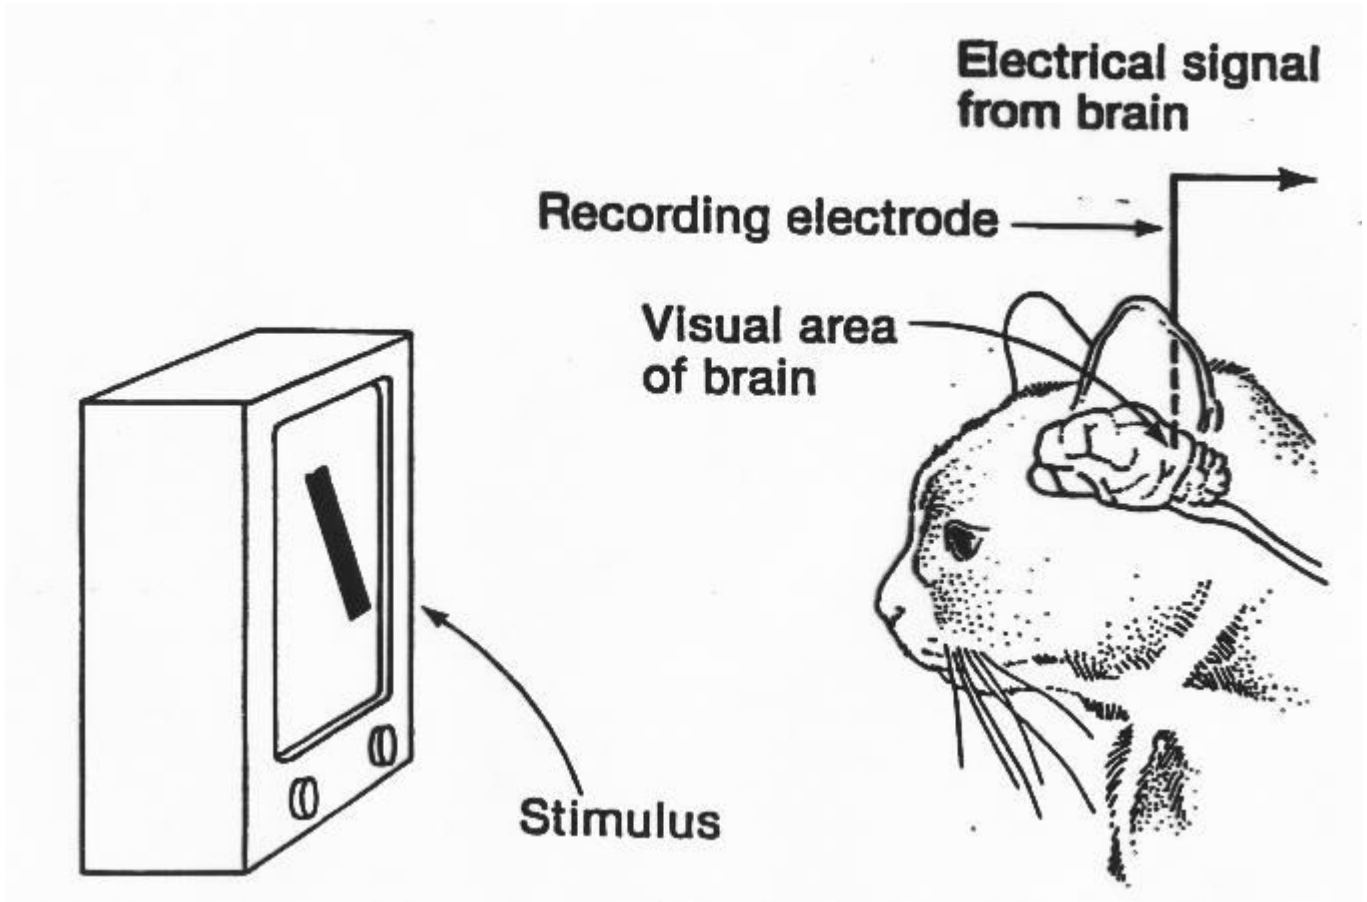
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus

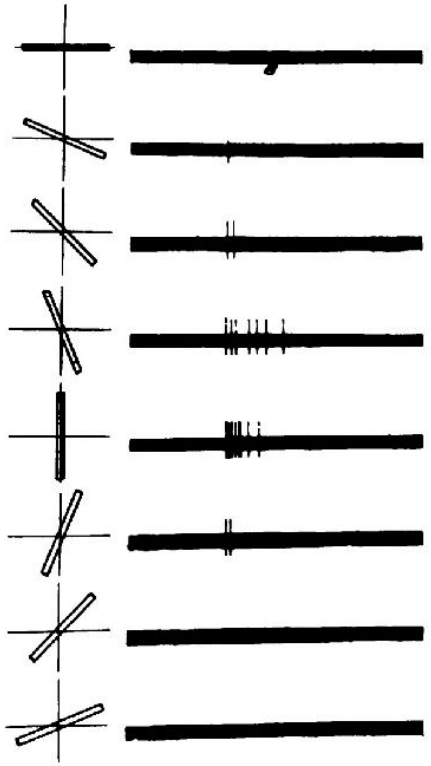

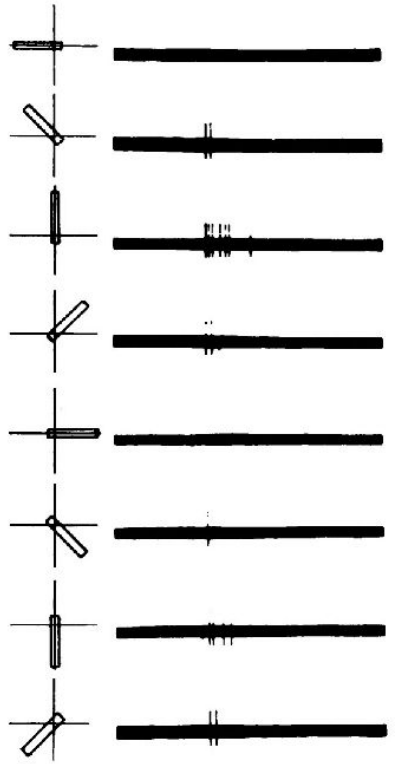

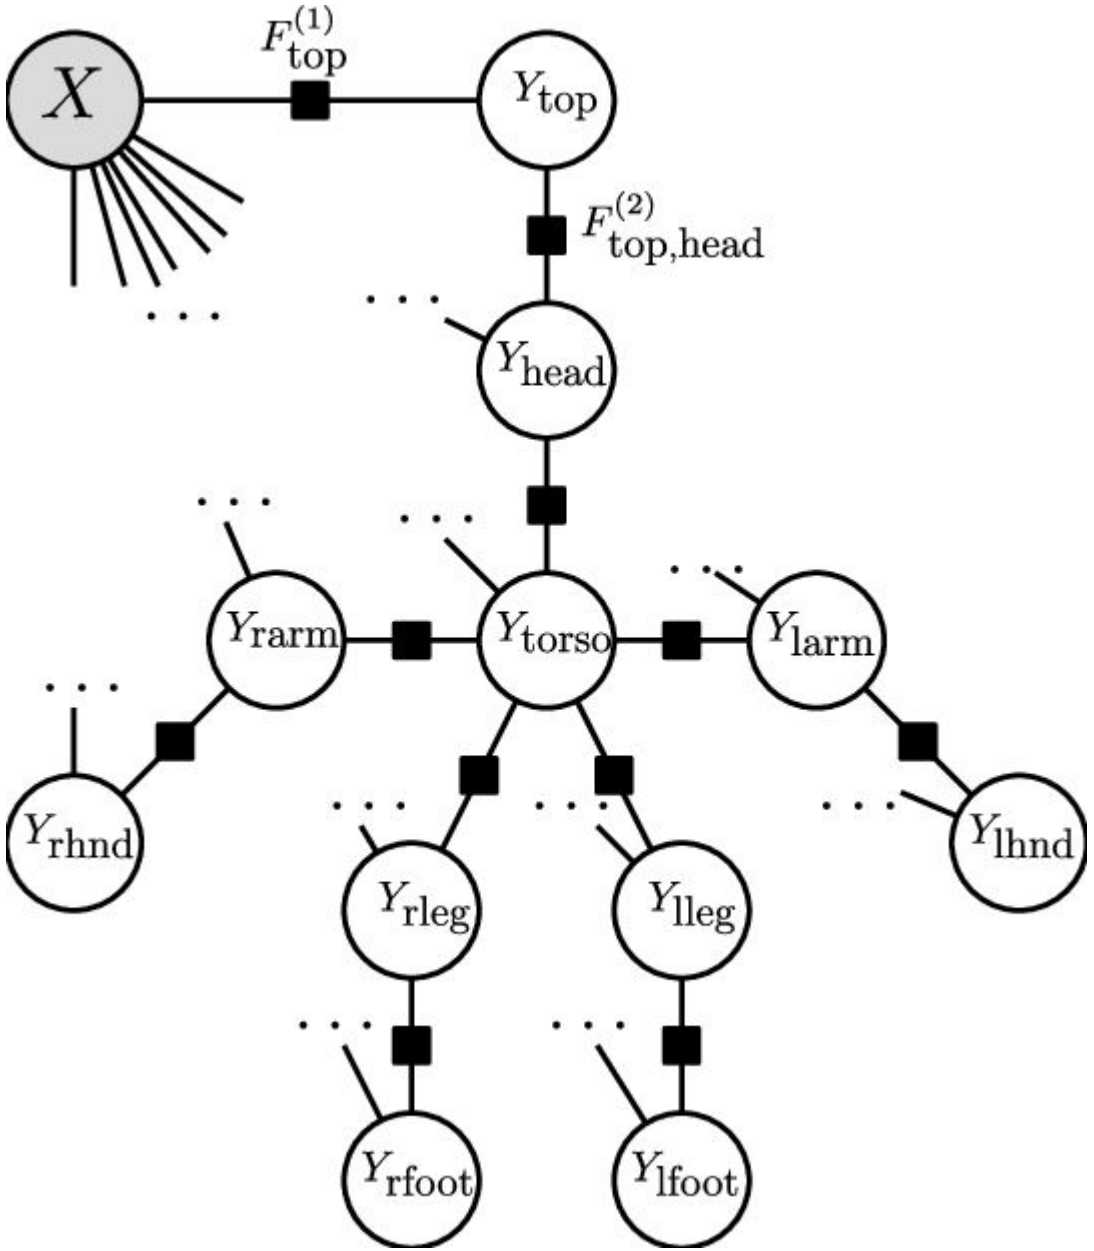

Figure 14: Successful matches between sets of image segments and particular viewpoints of the model.

Figure 15: Successfully matched image segments superimposed upon the original image.

https://www.cs.ubc.ca/~lowe/papers/aij87.pdf

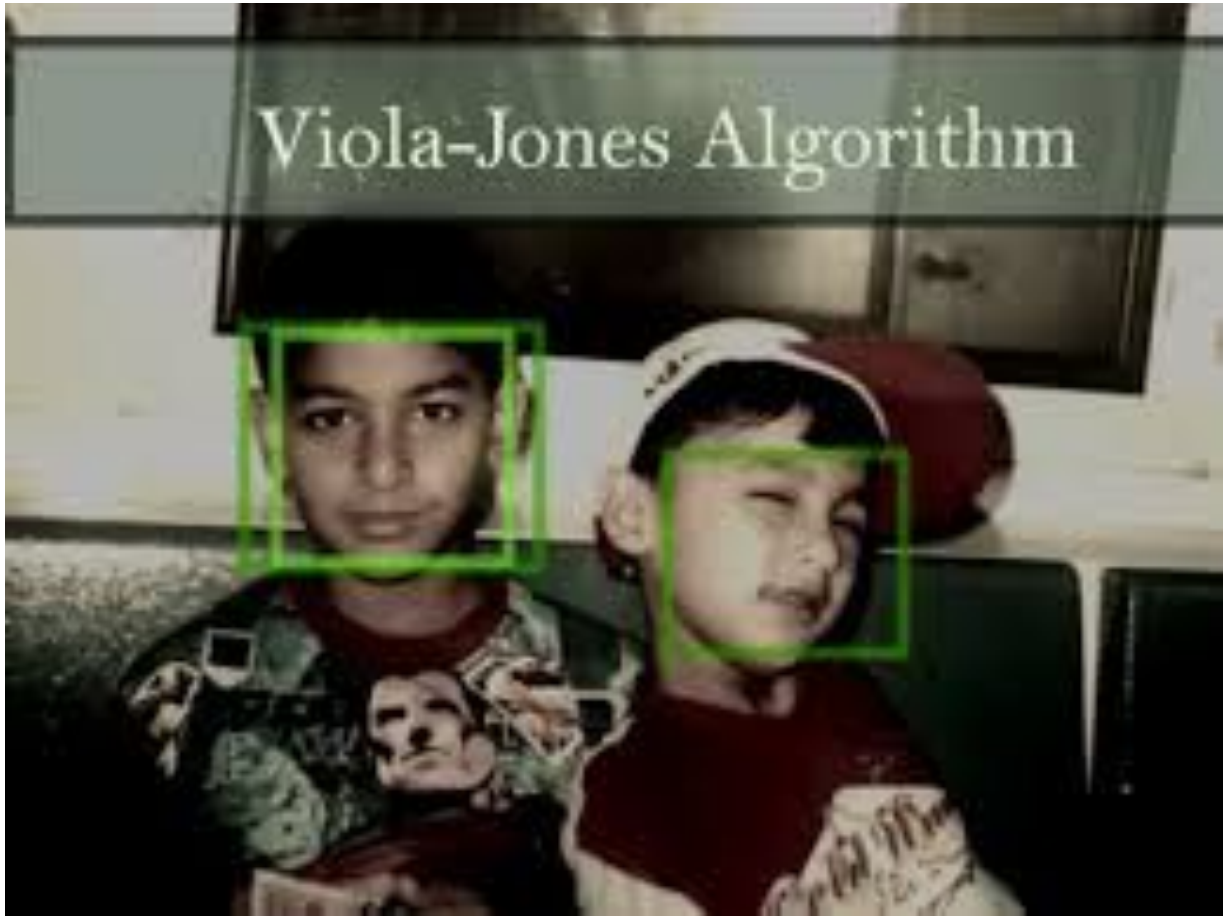
Viola-Jones Algorithm

Scale Invariant Feature Transformation (SIFT) Lowe 1999

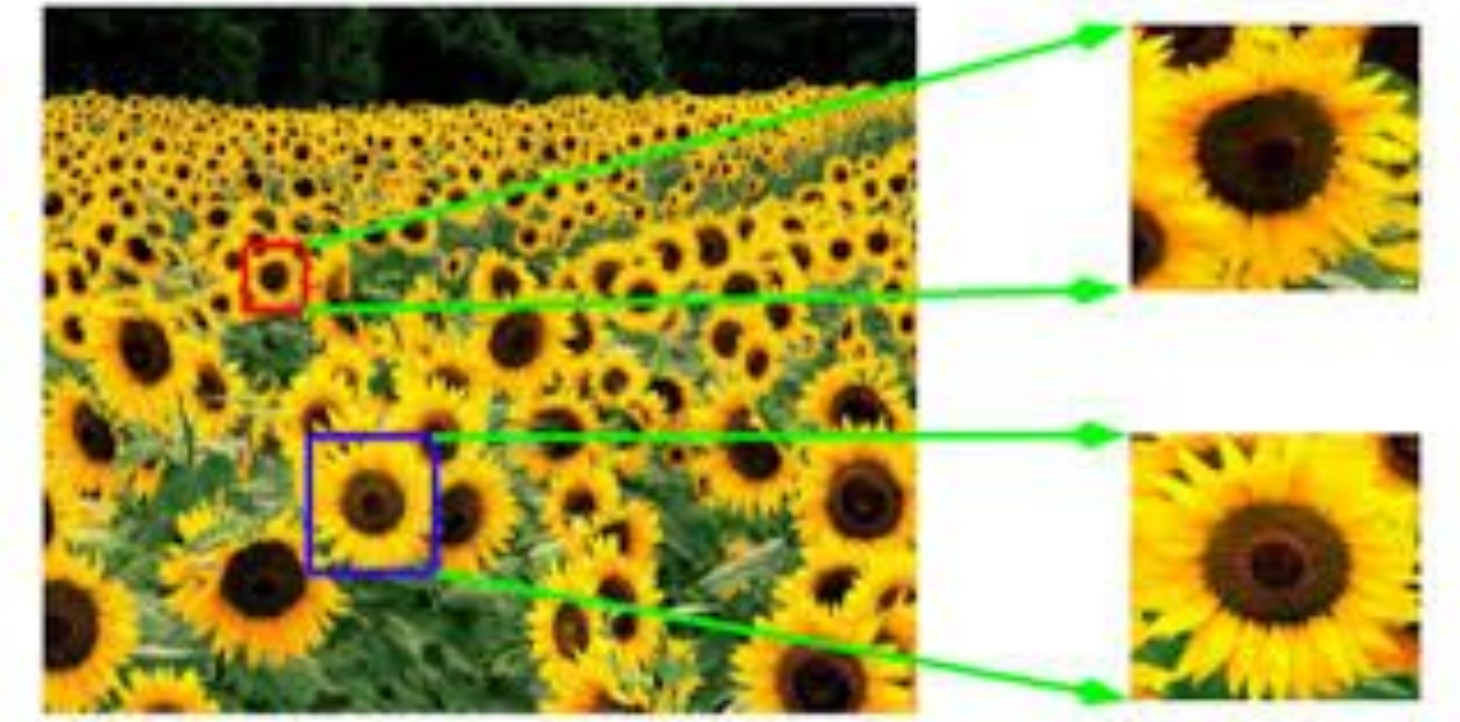

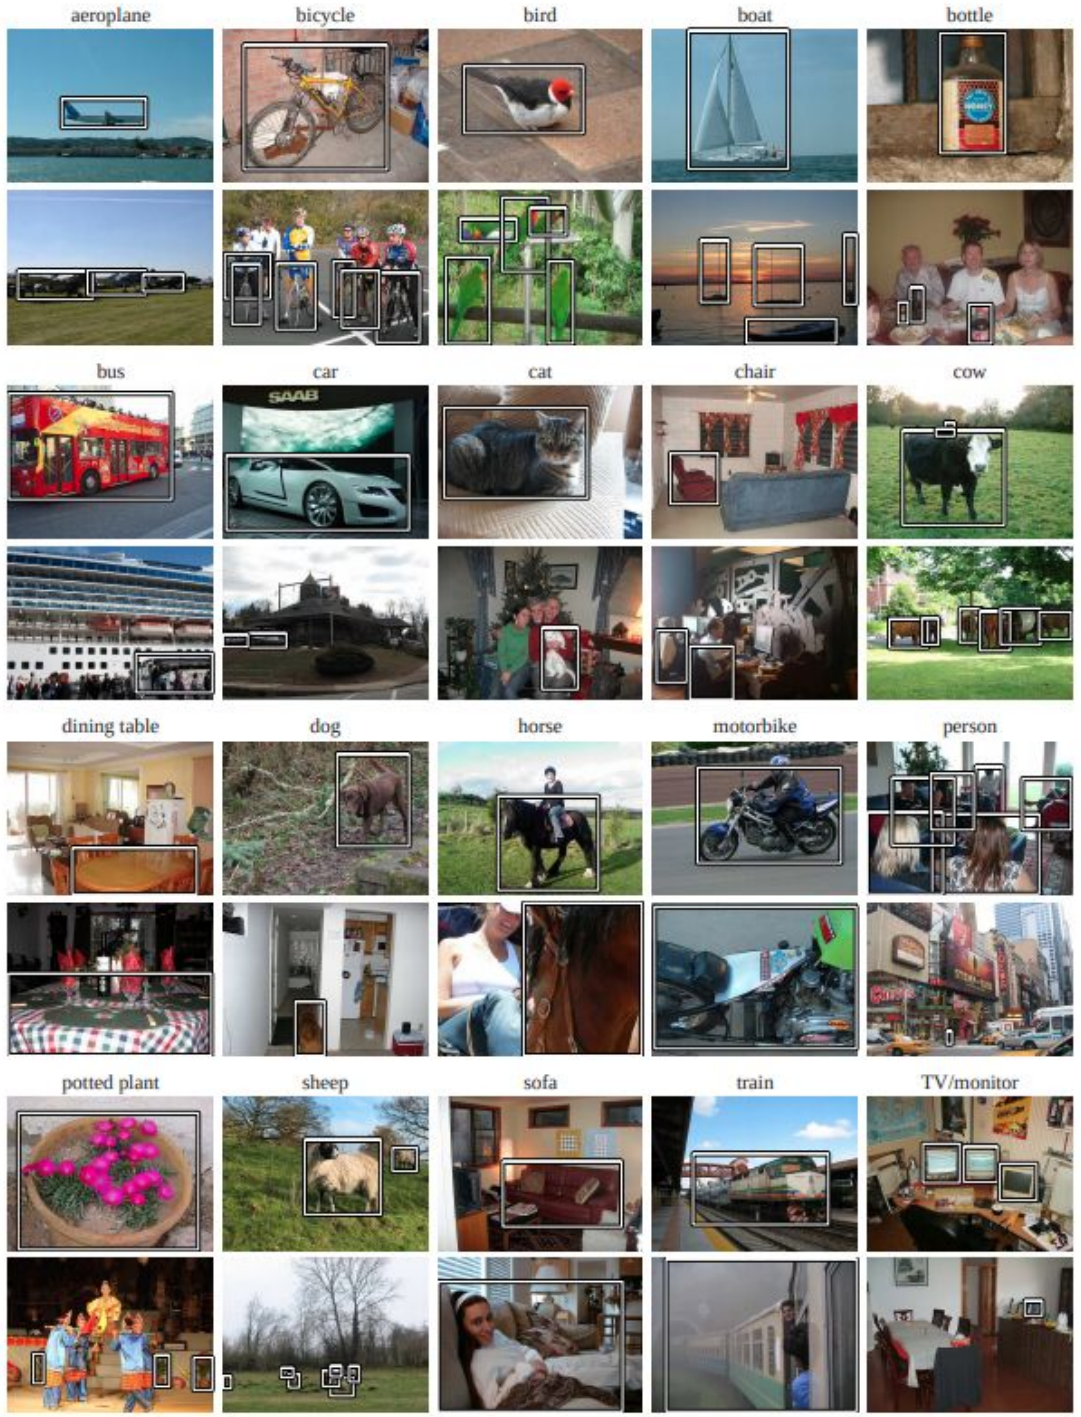
aeroplane
bicycle
bird
boat
bottle
bus
car
cat
chair
cow
dining table
dog
horse
motorbike
person
potted plant
sheep
sofa
train
TV/monitor

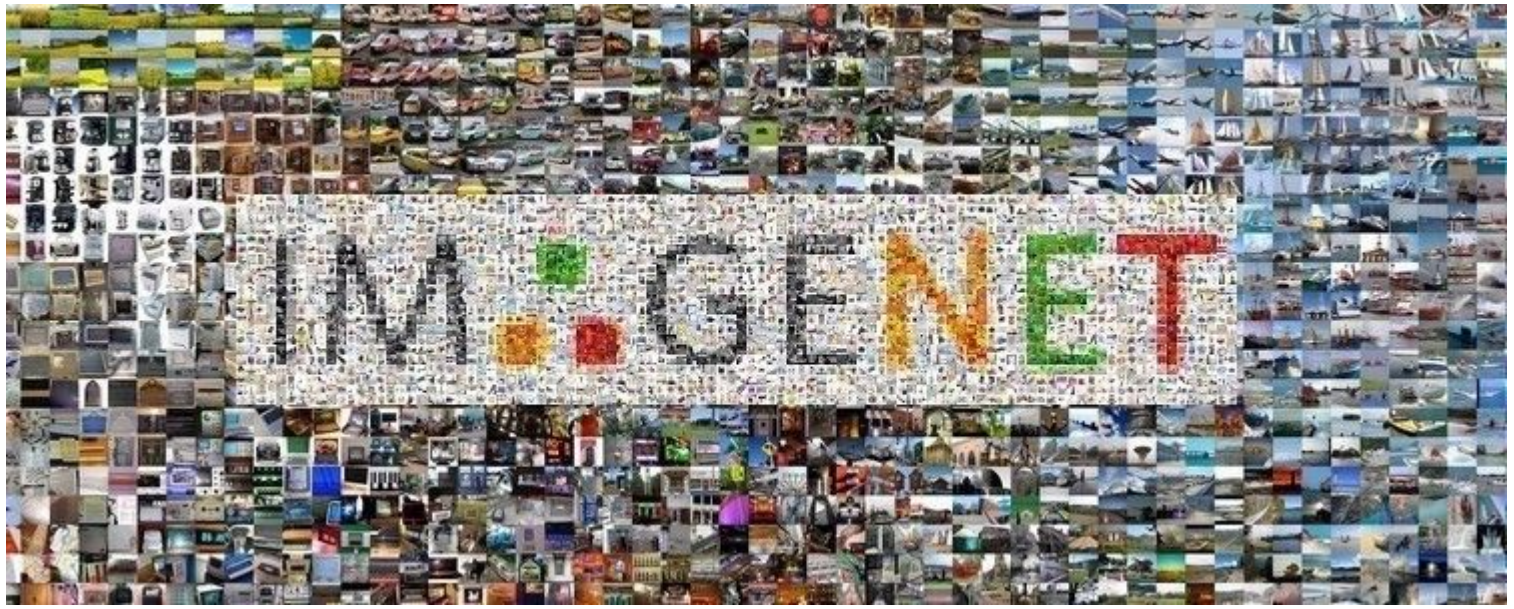

www.image-net.org

Image classification

Year	Codename	Error (percent)	99.9% Conf Int
2014	**GoogLeNet**	**6.66**	**6.40 - 6.92**
2014	VGG	7.32	7.05 - 7.60
2014	MSRA	8.06	7.78 - 8.34
2014	AHoward	8.11	7.83 - 8.39
2014	DeeperVision	9.51	9.21 - 9.82
2013	Clarifai†	11.20	10.87 - 11.53
2014	CASIAWS†	11.36	11.03 - 11.69
2014	Trimps†	11.46	11.13 - 11.80
2014	Adobe†	11.58	11.25 - 11.91
2013	**Clarifai**	**11.74**	**11.41 - 12.08**
2013	NUS	12.95	12.60 - 13.30
2013	ZF	13.51	13.14 - 13.87
2013	AHoward	13.55	13.20 - 13.91
2013	OverFeat	14.18	13.83 - 14.54
2014	Orange†	14.80	14.43 - 15.17
2012	SuperVision†	15.32	14.94 - 15.69
2012	**SuperVision**	**16.42**	**16.04 - 16.80**
2012	ISI	26.17	25.71 - 26.65
2012	VGG	26.98	26.53 - 27.43
2012	XRCE	27.06	26.60 - 27.52
2012	UvA	29.58	29.09 - 30.04

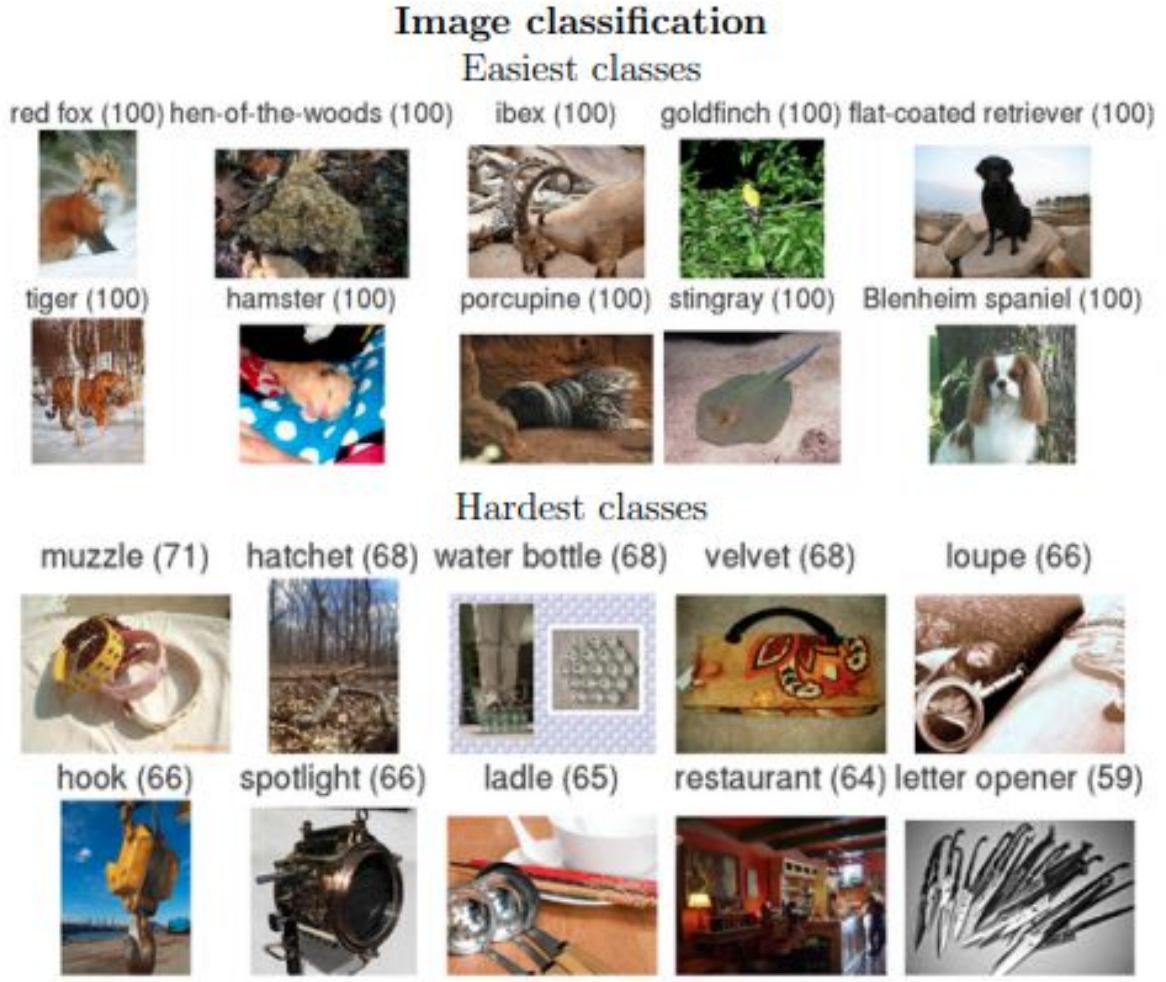

https://arxiv.org/pdf/1409.0575.pdf

Single-object localization			
Year	Codename	Error (percent)	99.9% Conf Int
2014	**VGG**	**25.32**	**24.87 - 25.78**
2014	GoogLeNet	26.44	25.98 - 26.92
2013	**OverFeat**	**29.88**	**29.38 - 30.35**
2014	Adobe†	30.10	29.61 - 30.58
2014	SYSU	31.90	31.40 - 32.40
2012	SuperVision†	33.55	33.05 - 34.04
2014	MIL	33.74	33.24 - 34.25
2012	**SuperVision**	**34.19**	**33.67 - 34.69**
2014	MSRA	35.48	34.97 - 35.99
2014	Trimps†	42.22	41.69 - 42.75
2014	Orange†	42.70	42.18 - 43.24
2013	VGG	46.42	45.90 - 46.95
2012	VGG	50.03	49.50 - 50.57
2012	ISI	53.65	53.10 - 54.17
2014	CASIAWS†	61.96	61.44 - 62.48

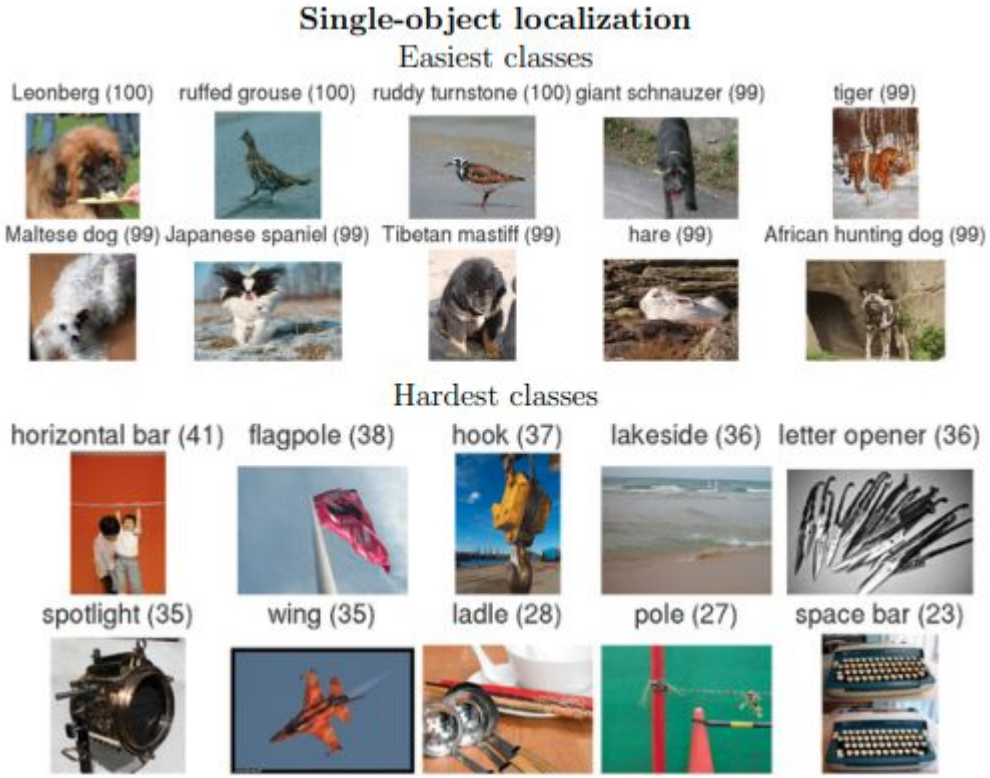

https://arxiv.org/pdf/1409.0575.pdf

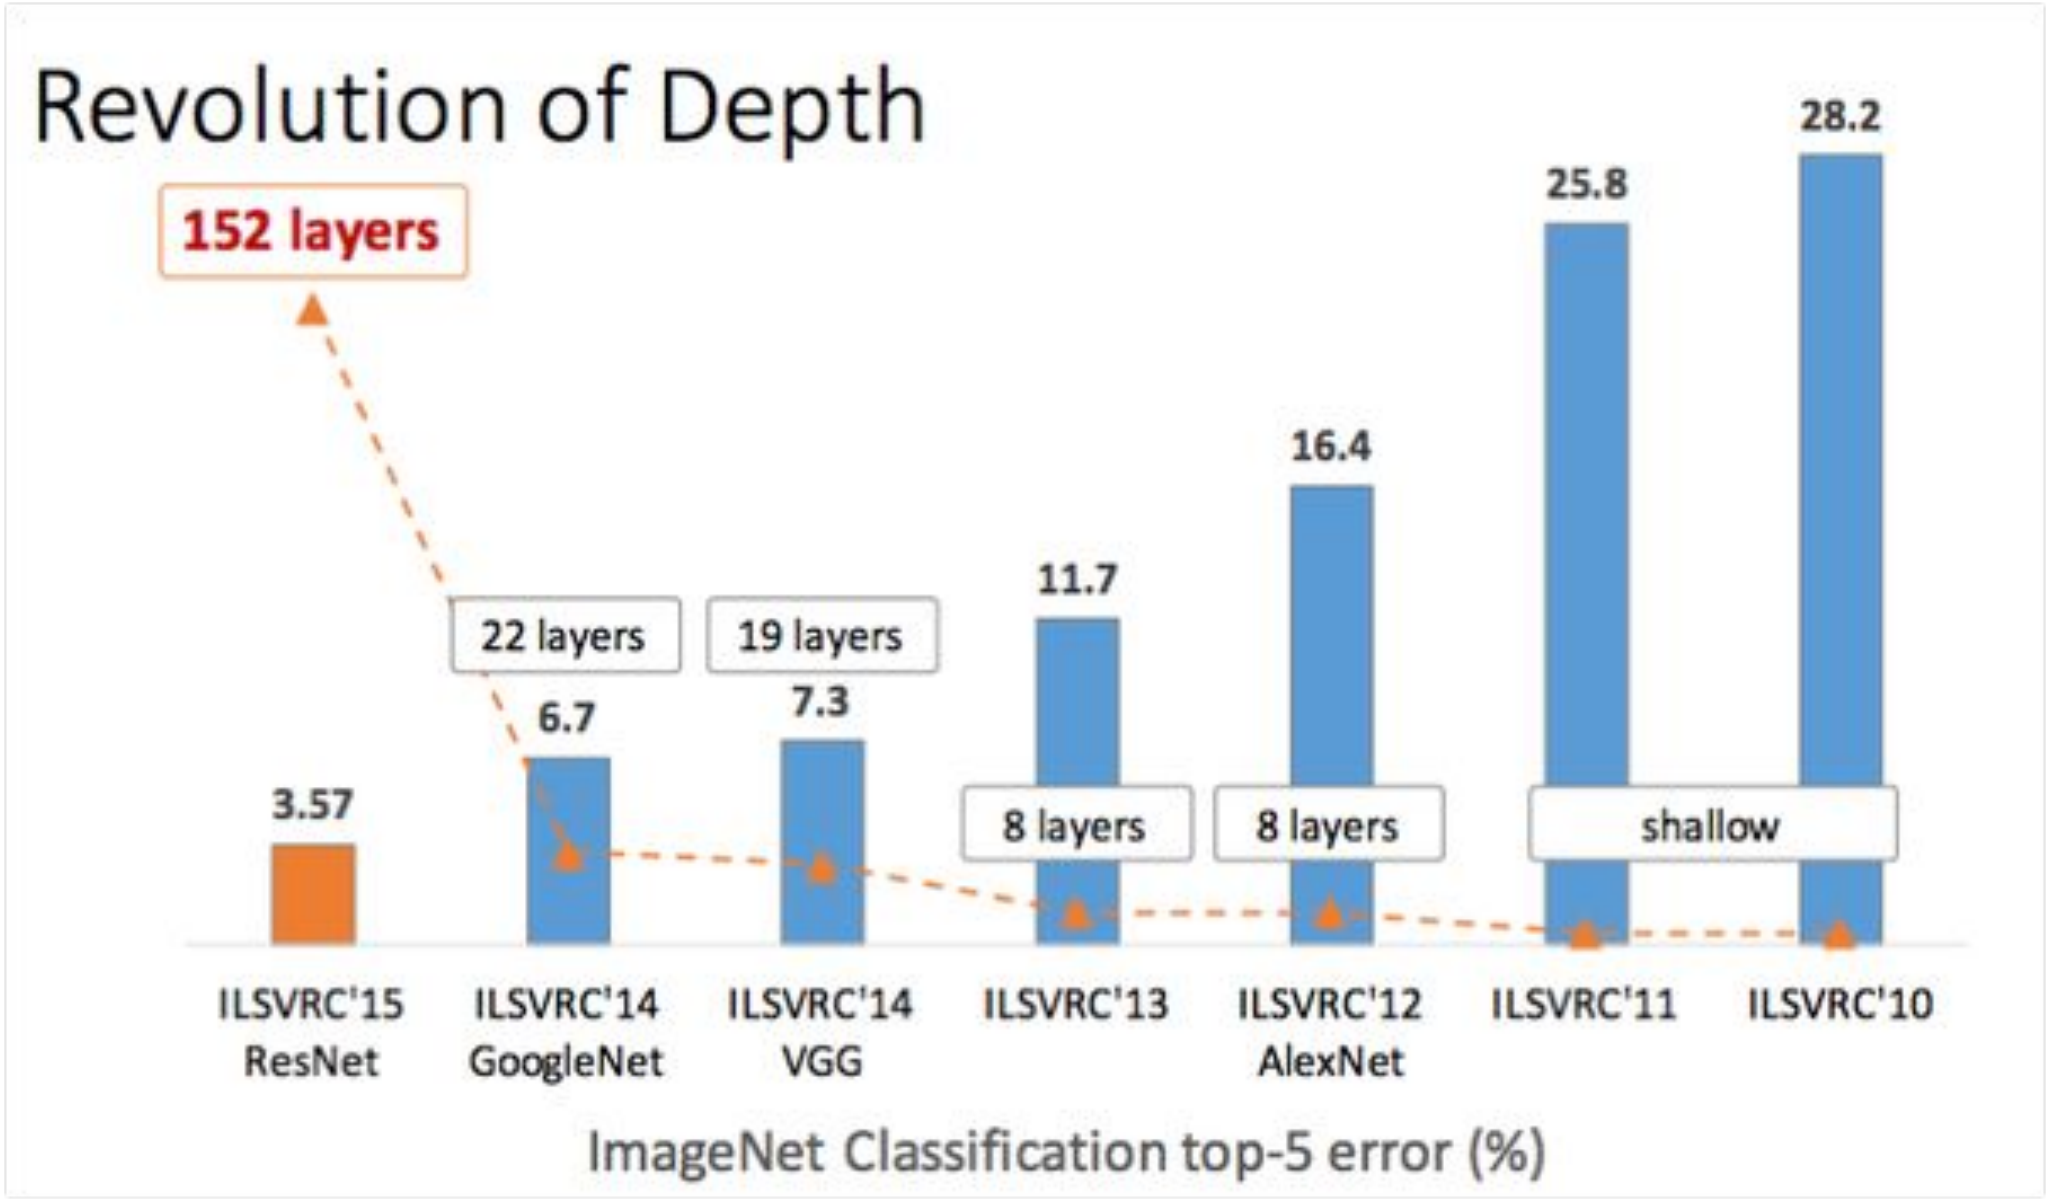
Revolution of Depth
152 layers
22 layers
19 layers
8 layers
8 layers
shallow
3.57
6.7
7.3
11.7
16.4
25.8
28.2
ILSVRC'15
ResNet
ILSVRC'14
GoogleNet
ILSVRC'14
VGG
ILSVRC'13
ILSVRC'12
AlexNet
ILSVRC'11
ILSVRC'10
ImageNet Classification top-5 error (%)

Back to DATA 442

- We will primarily be focusing on:
  - *Image Classification* - the process of reading in an image, and then selecting the most likely class that the image belongs to.
  - *Image Segmentation* - the process of identifying pixels in an image which belong to the same group.
  - *Object Detection* - The process of identifying an object exists, and where they are.
  - *Image Captioning* - The process of describing the relationships between objects in an image.

Image Classification

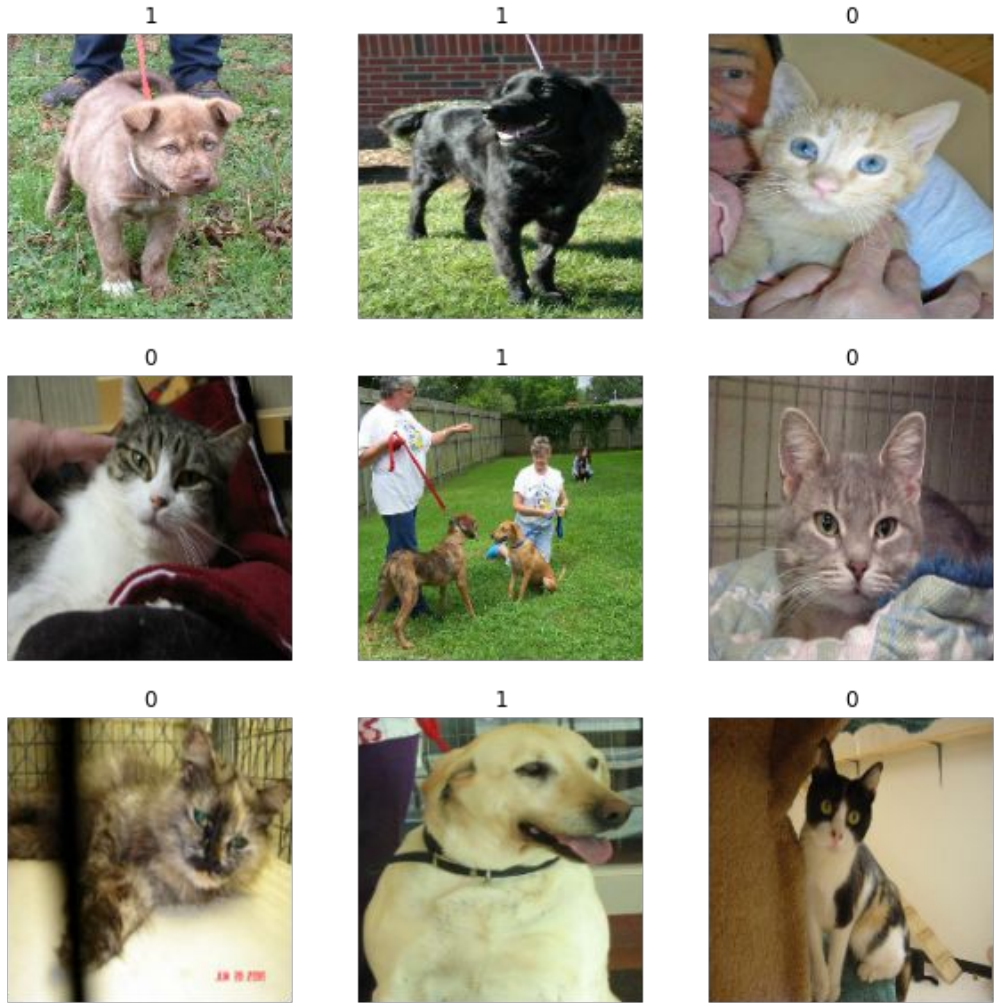

Image Segmentation

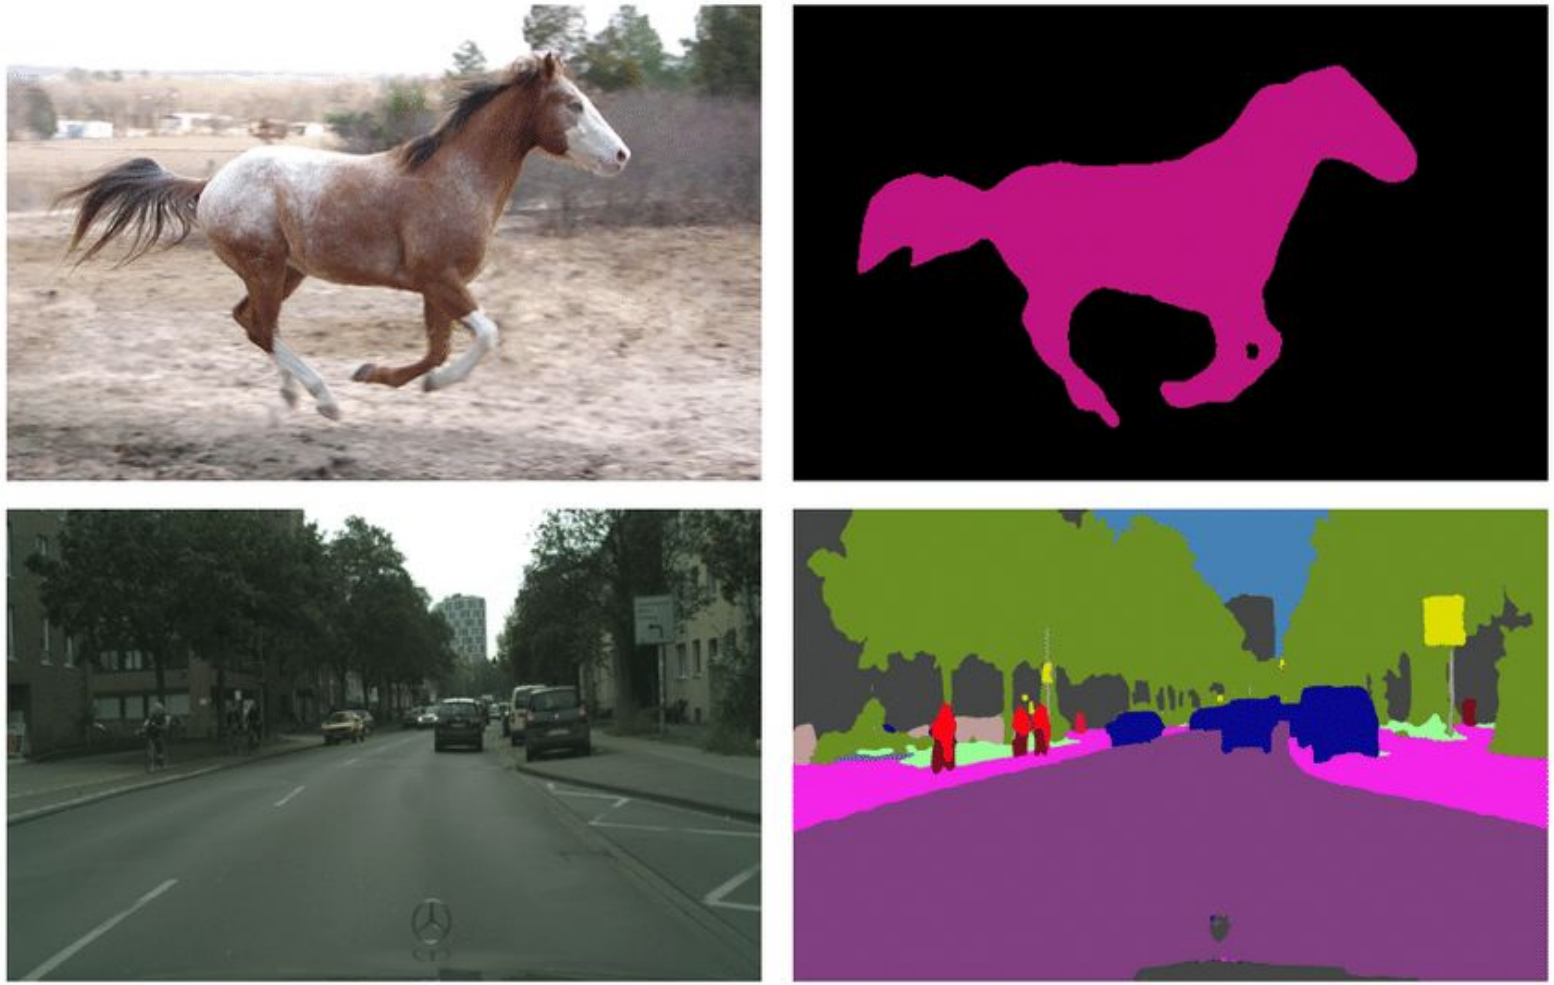

Object Detection

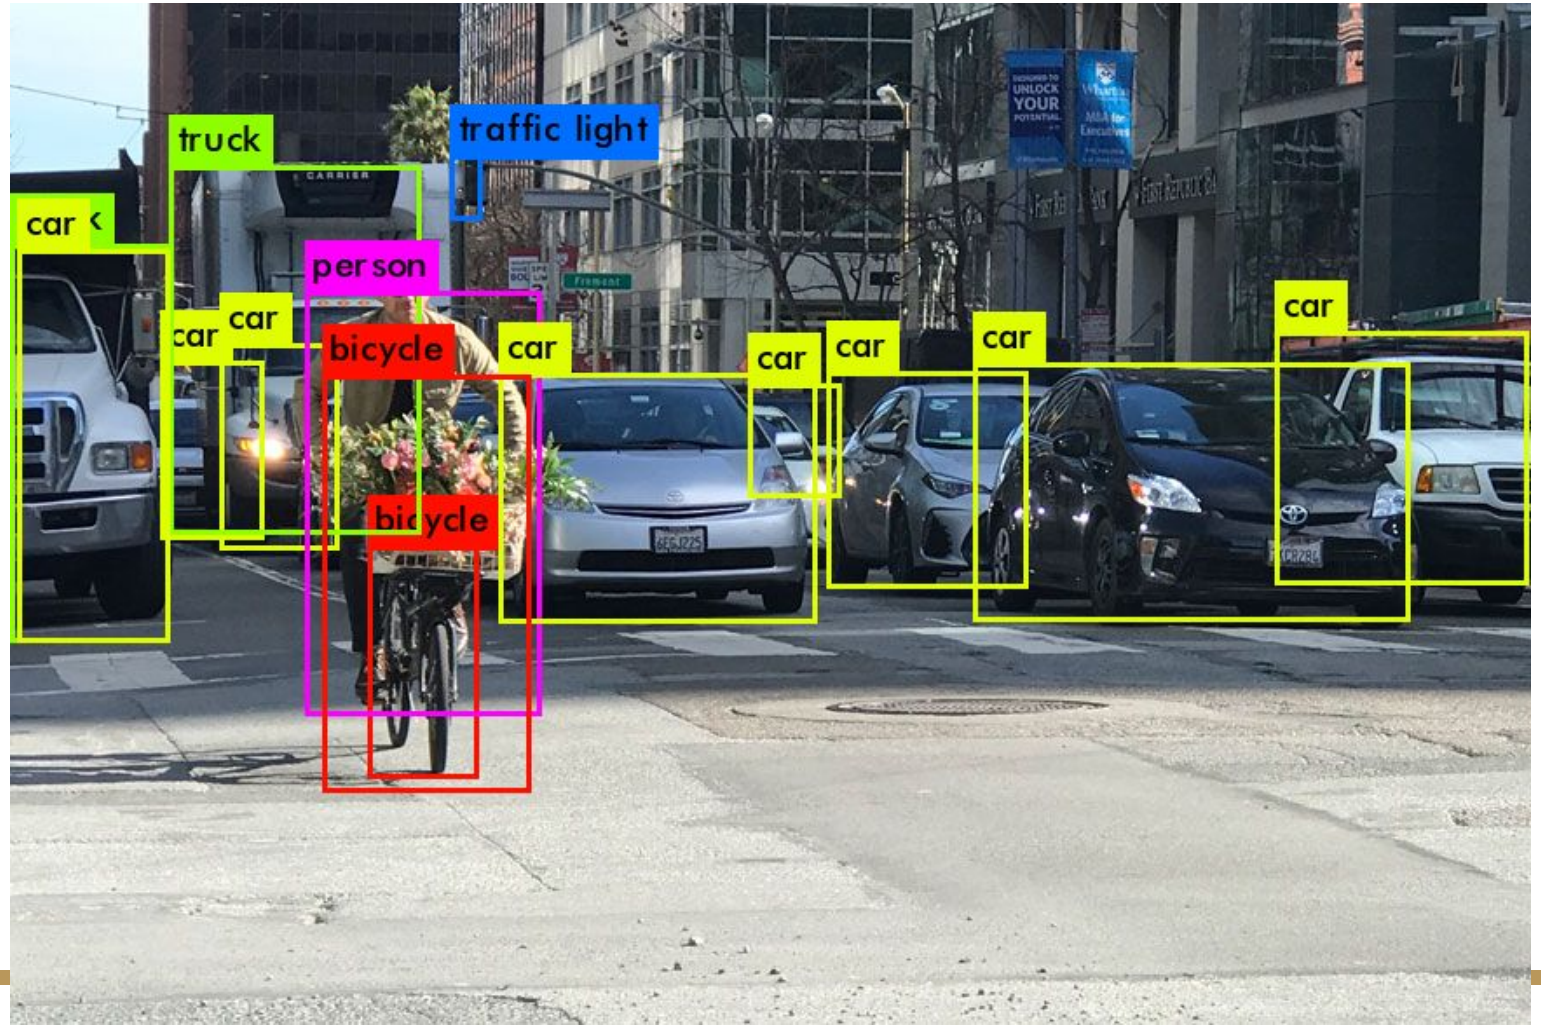

Image Captioning

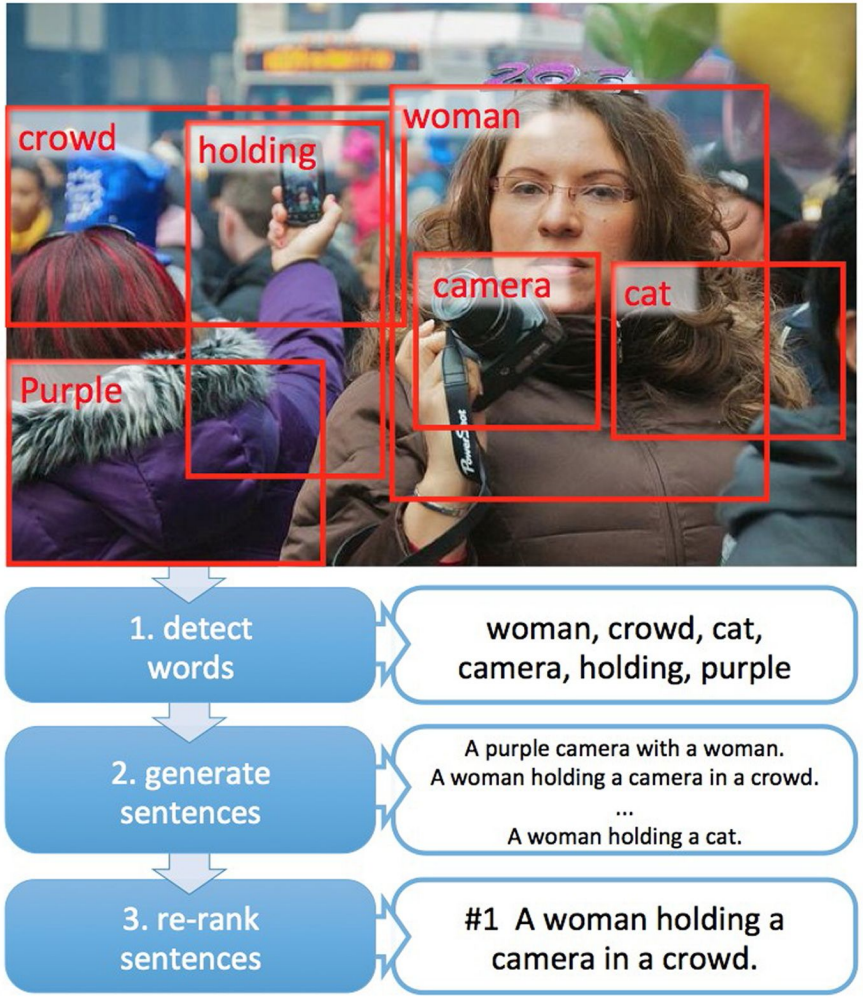

The Beating Heart of this Course: Convolutional Neural Networks

Since 2012, Neural Networks have been the king for image recognition. But, the algorithms themselves have been around since LeCun et al. in 1998 (Bell Labs).

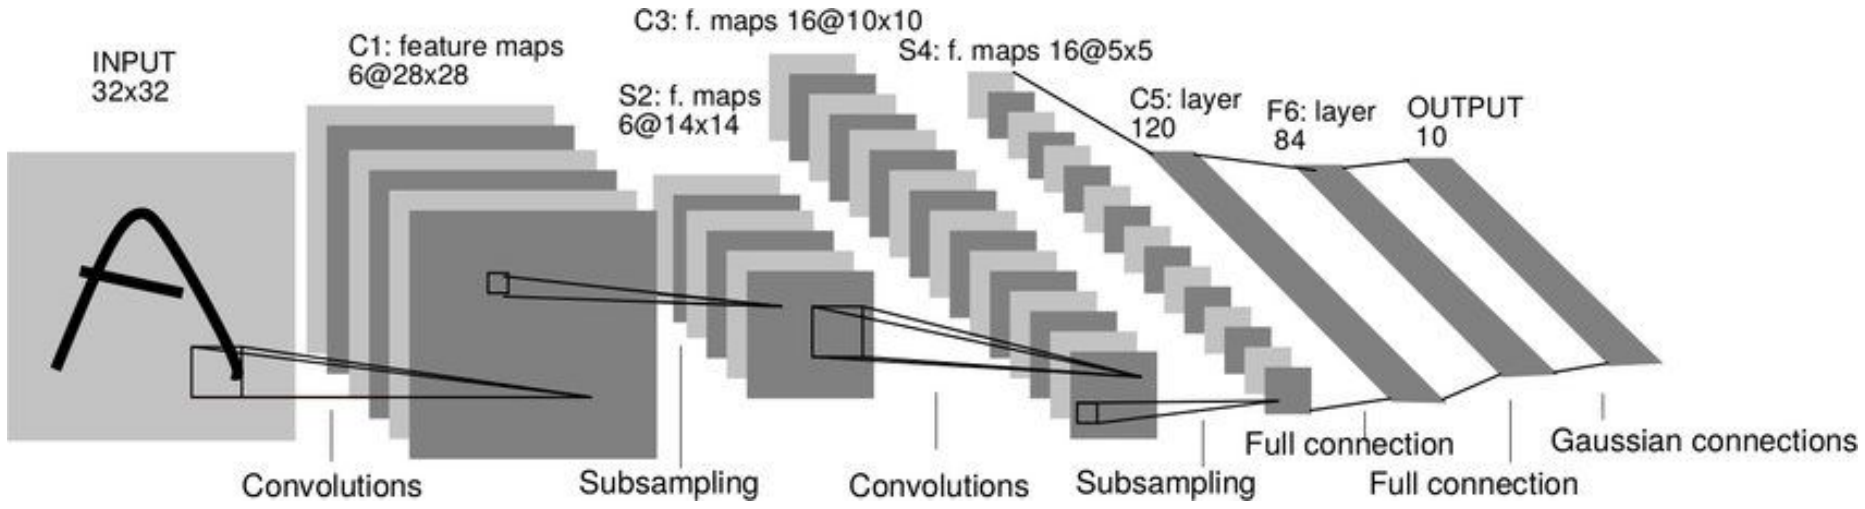

Leaps Forward in Consumer-Grade Hardware and Open Data

1998 Hardware

Intel Pentium II-450, Released August 1998

1 core, 7.5 million transistors, 0.45 GHz Clock Speed

Voodoo II 12MB GPU, Released CY 1998

4 million transistors

Training Data

MNIST 98

11.6 Megabytes, 70k Images

2020 Hardware

AMD Threadripper 3970X

32 core / 64 Thread, 23.94 ***billion*** transistors, 4.5 GHz Clock Speed

NVIDIA A100 40GB Tensor Core GPU

54 billion transistors

Training Data

Tencent ML Images

~35.4 Terabytes, 17.7m Images

Not Just Quantity

LSUN - A Scenes Database

Moving
MNIST

Crowd-sourced Emotional Mutimodal Actors Dataset (CREMA-D)

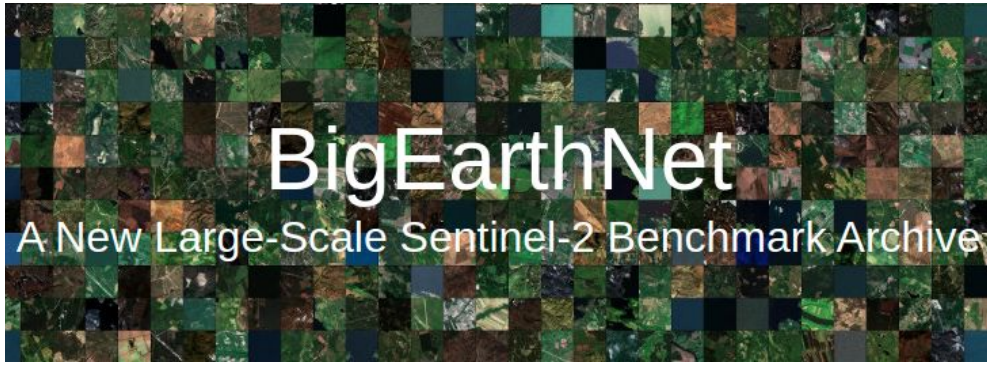

DATA 442 - Learning Paradigms

- Most Course Communication is done on Piazza.
  - Questions can be posted anonymously.
  - Professors and TAs will respond regularly.
  - We encourage you to share ideas or findings as you explore!
  - Up-to-date links for scheduling office hours, syllabus, submitting assignments, and more will be on Piazza.
- We will continually offer additional resources if you want to do a deeper dive on any topic.
  - Optional textbook: Deep Learning by Goodfellow, Bengio and Courville - free online from MIT Press!

The Goals

- Deep Understanding of “What’s going on with these networks”
  - You should be able to create, debug, train, test, and tweak convolutional neural networks.
- Applied
  - Teach you the advantages and disadvantages of different strategies for fitting networks, software packages, hardware architectures, and more knowledge required to successfully create your own nets.
- Beautiful and Scary
  - We’ll be showing some of the newer innovations in CNNs, and giving you the tools to experiment.

Grading

- "Fantastic Five"
  - 3 Assignments, 20% each (60% total)
  - Midterm - 20%
  - Final - 20%

- Late Submissions *are not accepted, excepting in documented circumstances (i.e., an illness with a doctor's note). We* ***highly recommend you submit assignments early!***

- Collaboration is highly encouraged, but the work you submit should be your own. Group submissions are not allowed.

Baseline Assumptions about You

Based on the prerequisites for this course, we are making a number of assumptions about you. If these assumptions are not true, you may struggle to keep up!

- Proficiency in Python (*all assignments in python!)*
- Very basic knowledge of how to take derivatives, matrix algebra (multiplication, etc.)
- Comfortable with concepts like cost functions, gradient descent optimization, and basic ML techniques like knn classifiers.

That's it for Today!

Head over to the course website to get all of the resources you'll need to succeed:

icss.wm.edu/data442/