Object Recognition

Danna Gurari
The University of Texas at Austin
Fall 2019

Review

- Last week:
  - Computer vision: past, present, & future
  - Computer vision: what makes it hard?
  - Introduction to crowdsourcing for computer vision
- Assignments (Canvas)
  - Reading assignment due yesterday
  - New reading assignment out due next week
  - Lab assignment out due in two weeks
- Questions?

Today’s Topics

- Object recognition applications
- Object recognition datasets: key steps in creating them
- Object recognition datasets: scaling up their size with *crowdsourcing*
- Scaling up community working on object recognition with *workshop challenges*
- Class Discussion
- Lab: cascading stylesheets and web page layout

Today’s Topics

- Object recognition applications
- Object recognition datasets: key steps in creating them
- Object recognition datasets: scaling up their size with *crowdsourcing*
- Scaling up community working on object recognition with *workshop challenges*
- Class Discussion
- Lab: cascading stylesheets and web page layout

Object Recognition Applications: Shopping

Take a picture of an object and find where to buy it

Object Recognition Applications: Vision Assistance for People Who Are Blind

Orcam Demo: https://www.youtube.com/watch?v=_3XVsCsscyw
(start video at 3:16)

Object Recognition Applications: Photo Organization

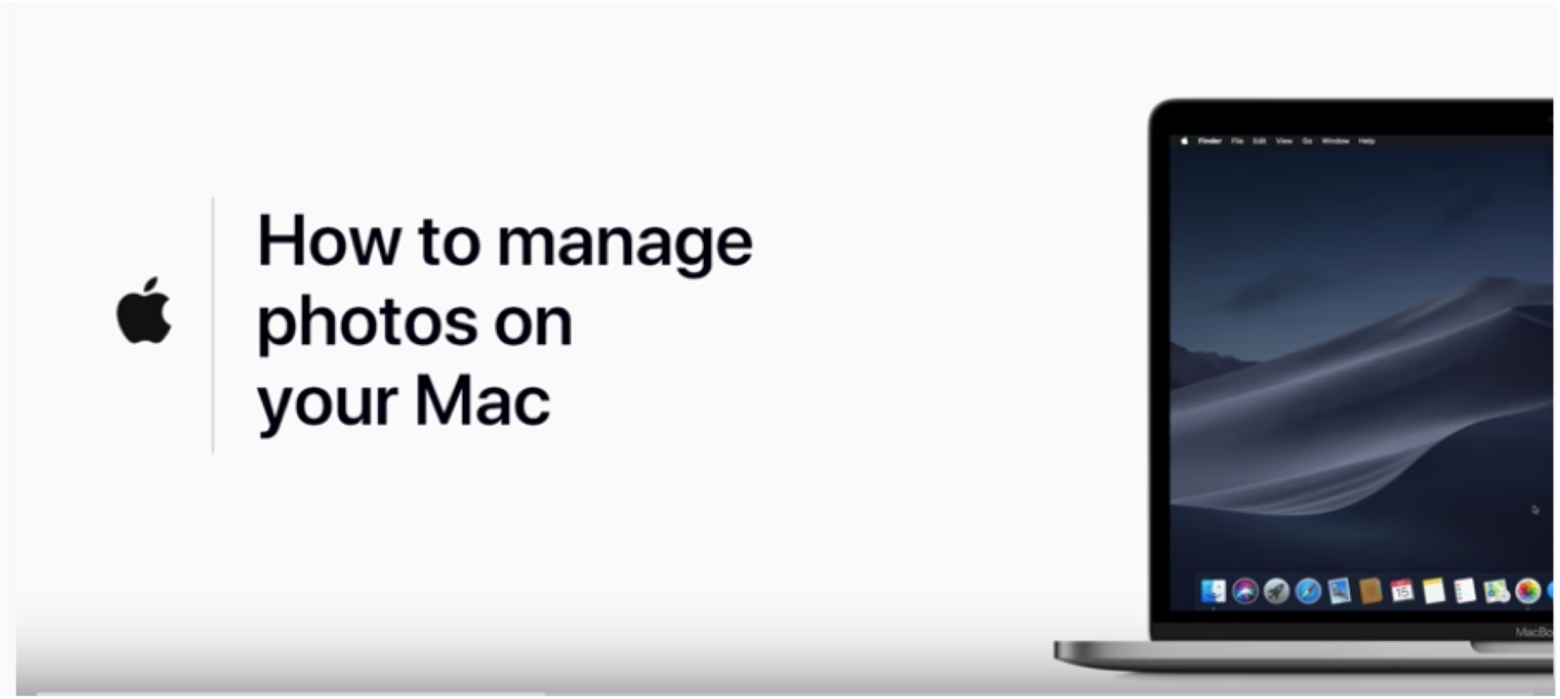

Apple Demo: https://www.youtube.com/watch?v=R3JTaxhpYzc
(start video at 2:36)

Object Recognition Applications: Image Search with Automated Keywording

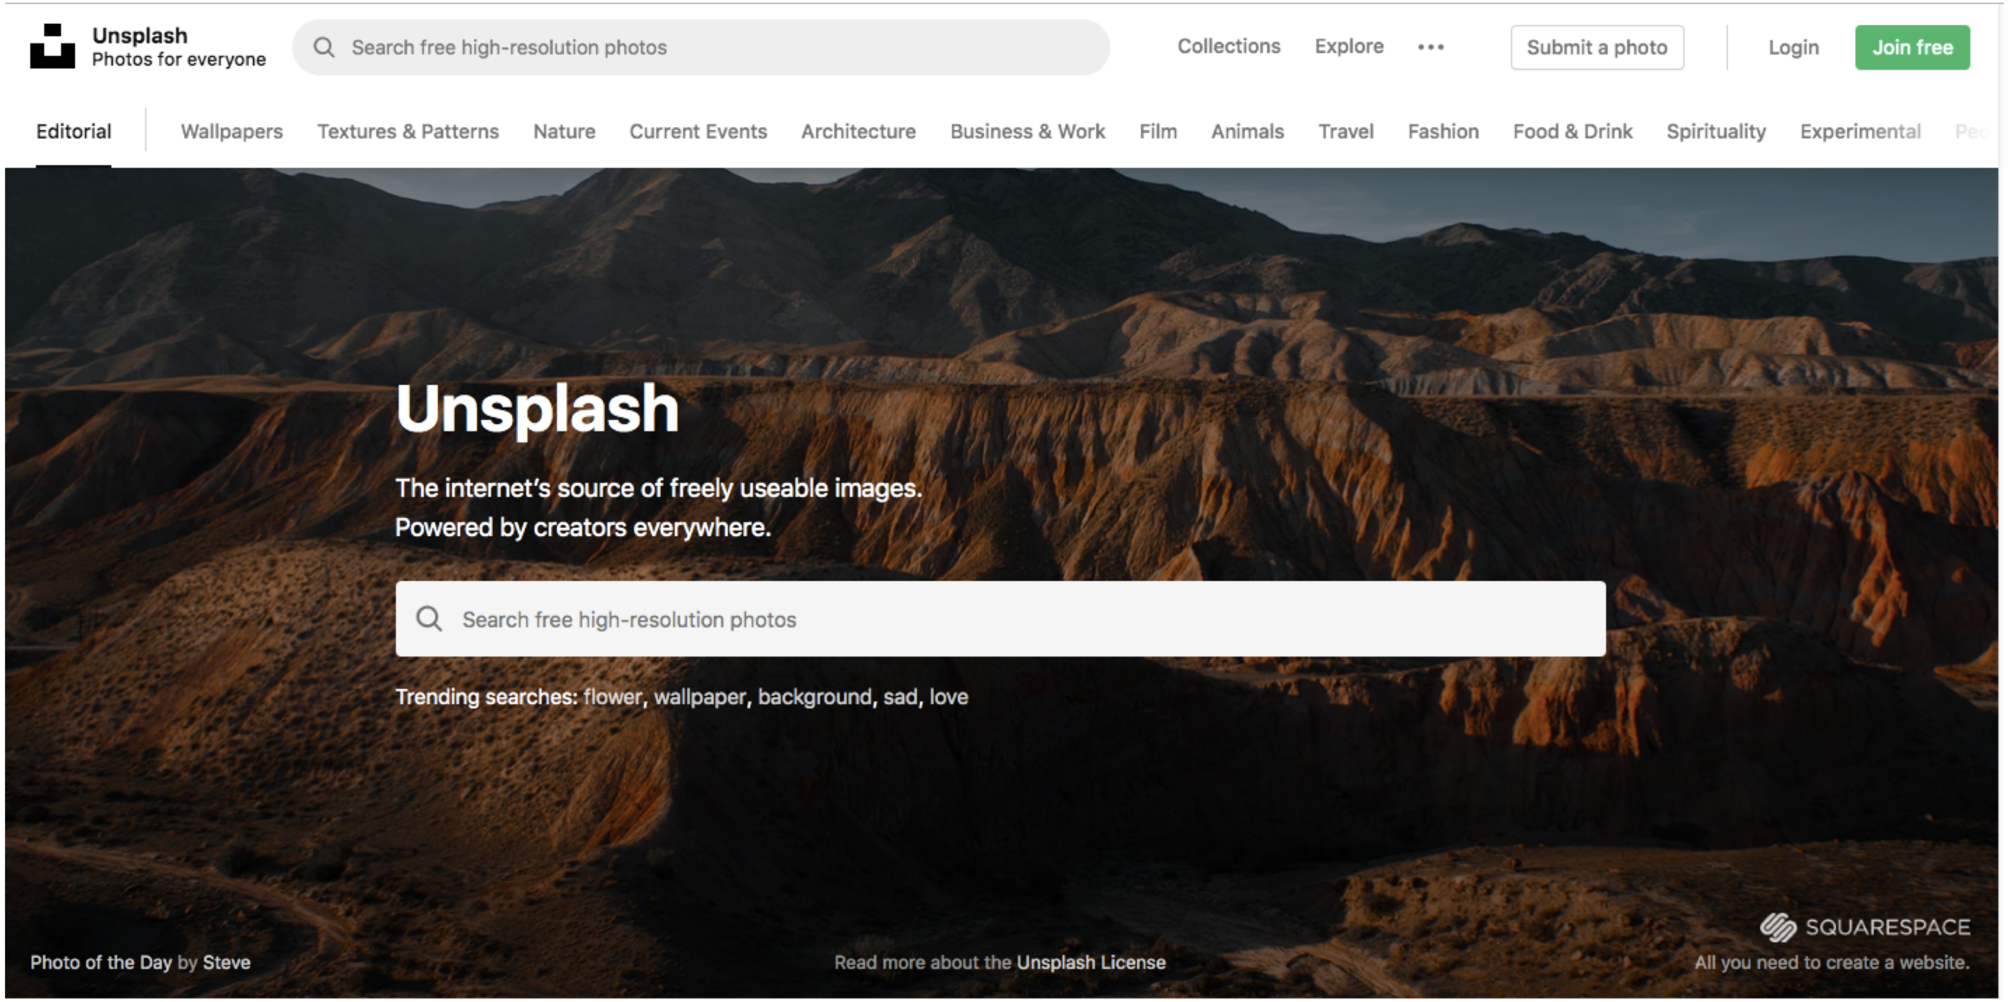

Object Recognition Applications: And Many More…

e.g., search on Google for “image recognition applications”

Object Recognition Applications Gone Wrong

- Ethical Mistake: Photo Tagging
  - http://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/

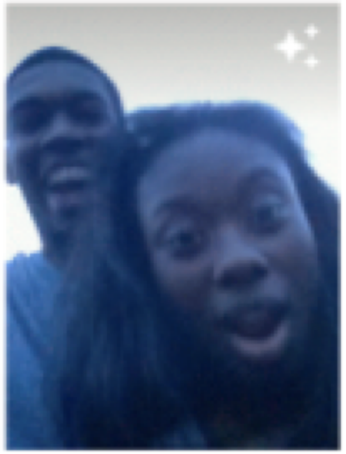

- Security Mistake: Person Recognition
  - https://www.cs.cmu.edu/~sbhagava/papers/face-rec-ccs16.pdf

Object Recognition Applications Gone Wrong

- Ethical Mistake: Photo Tagging
  - http://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/

Why do you think these systems make such mistakes?

- Security Mistake: Person Recognition
  - https://www.cs.cmu.edu/~sbhagava/papers/face-rec-ccs16.pdf

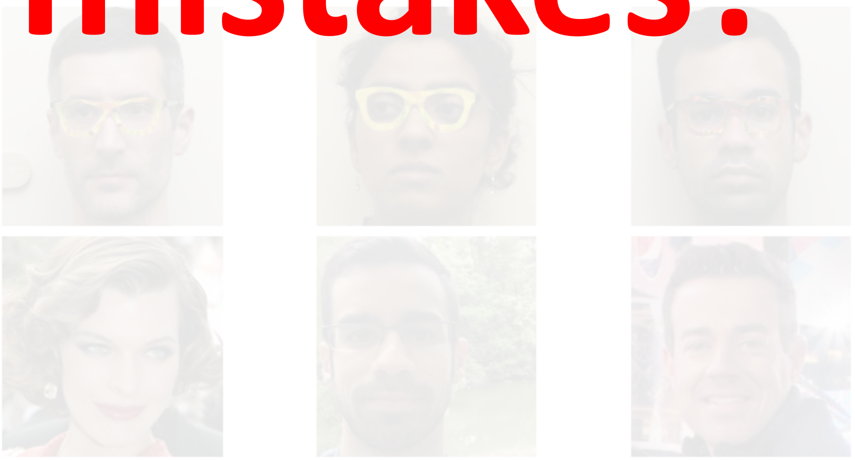

Object Recognition Applications Gone Wrong

- Ethical Mistake: Photo Tagging
  - http://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/

If you were the CEO, how would you change your product in response?

- Security Mistake: Person Recognition
  - https://www.cs.cmu.edu/~sbhagava/papers/face-rec-ccs16.pdf

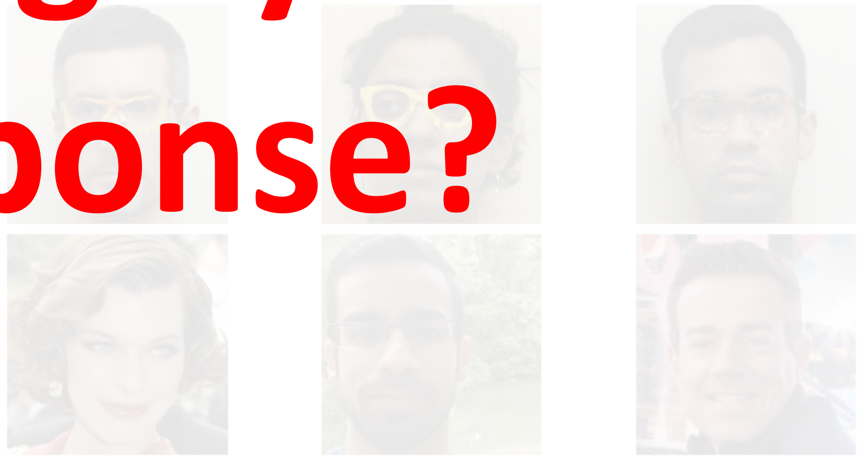

Today’s Topics

- Object recognition applications
- Object recognition datasets: key steps in creating them
- Object recognition datasets: scaling up their size with *crowdsourcing*
- Scaling up community working on object recognition with *workshop challenges*
- Class Discussion
- Lab: cascading stylesheets and web page layout

Recall: Need Datasets to Train & Evaluate Algorithms

1. Create Training Data

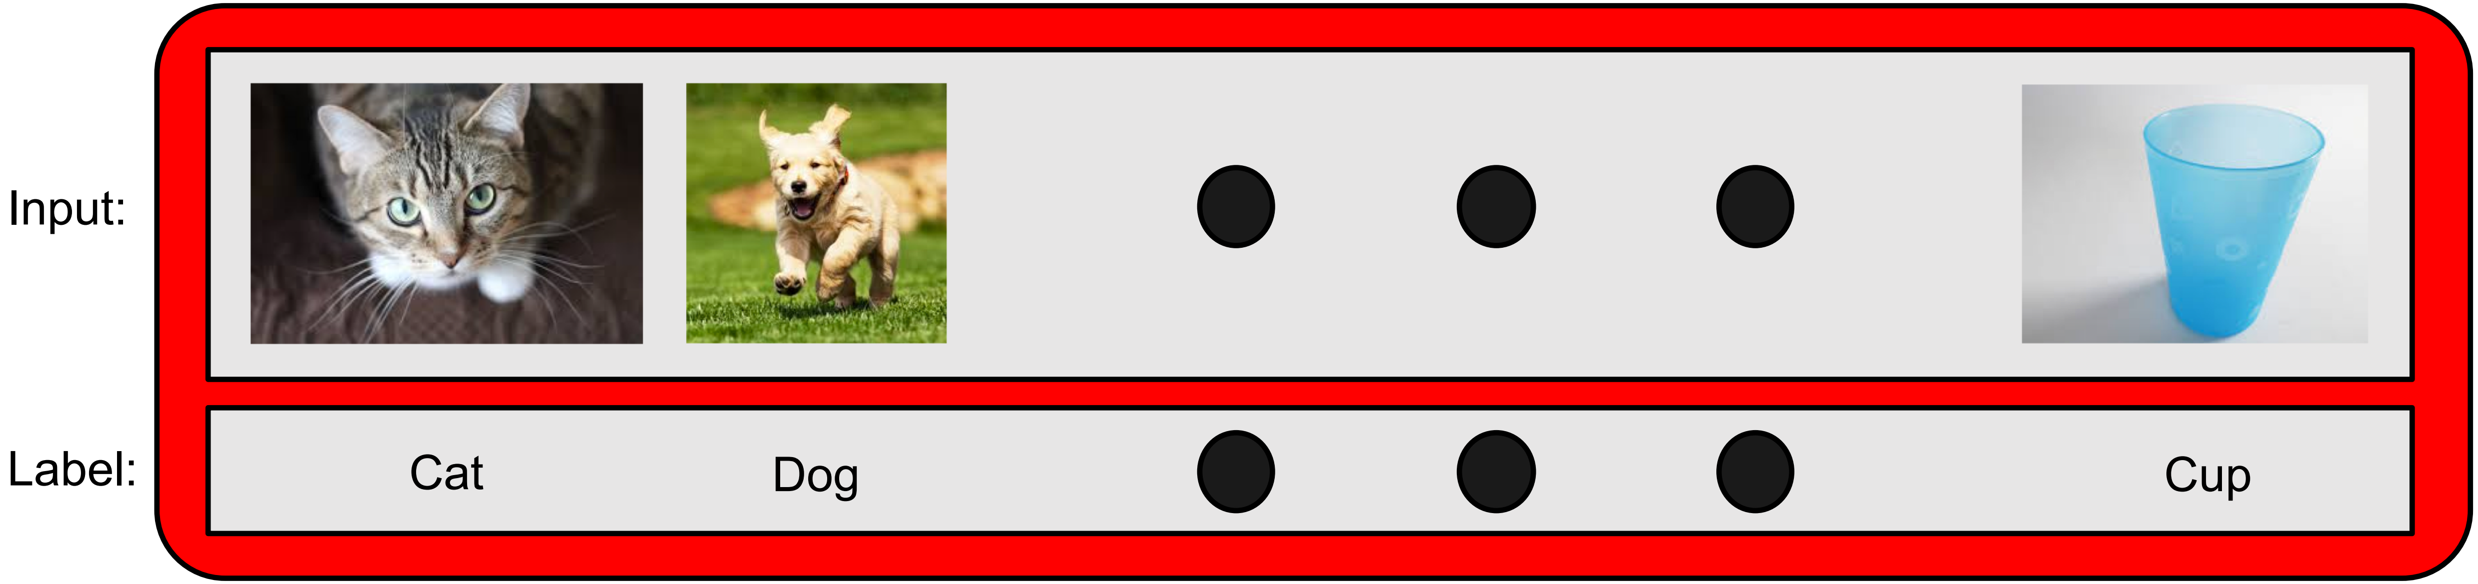

Recall: Need Datasets to Train & Evaluate Algorithms

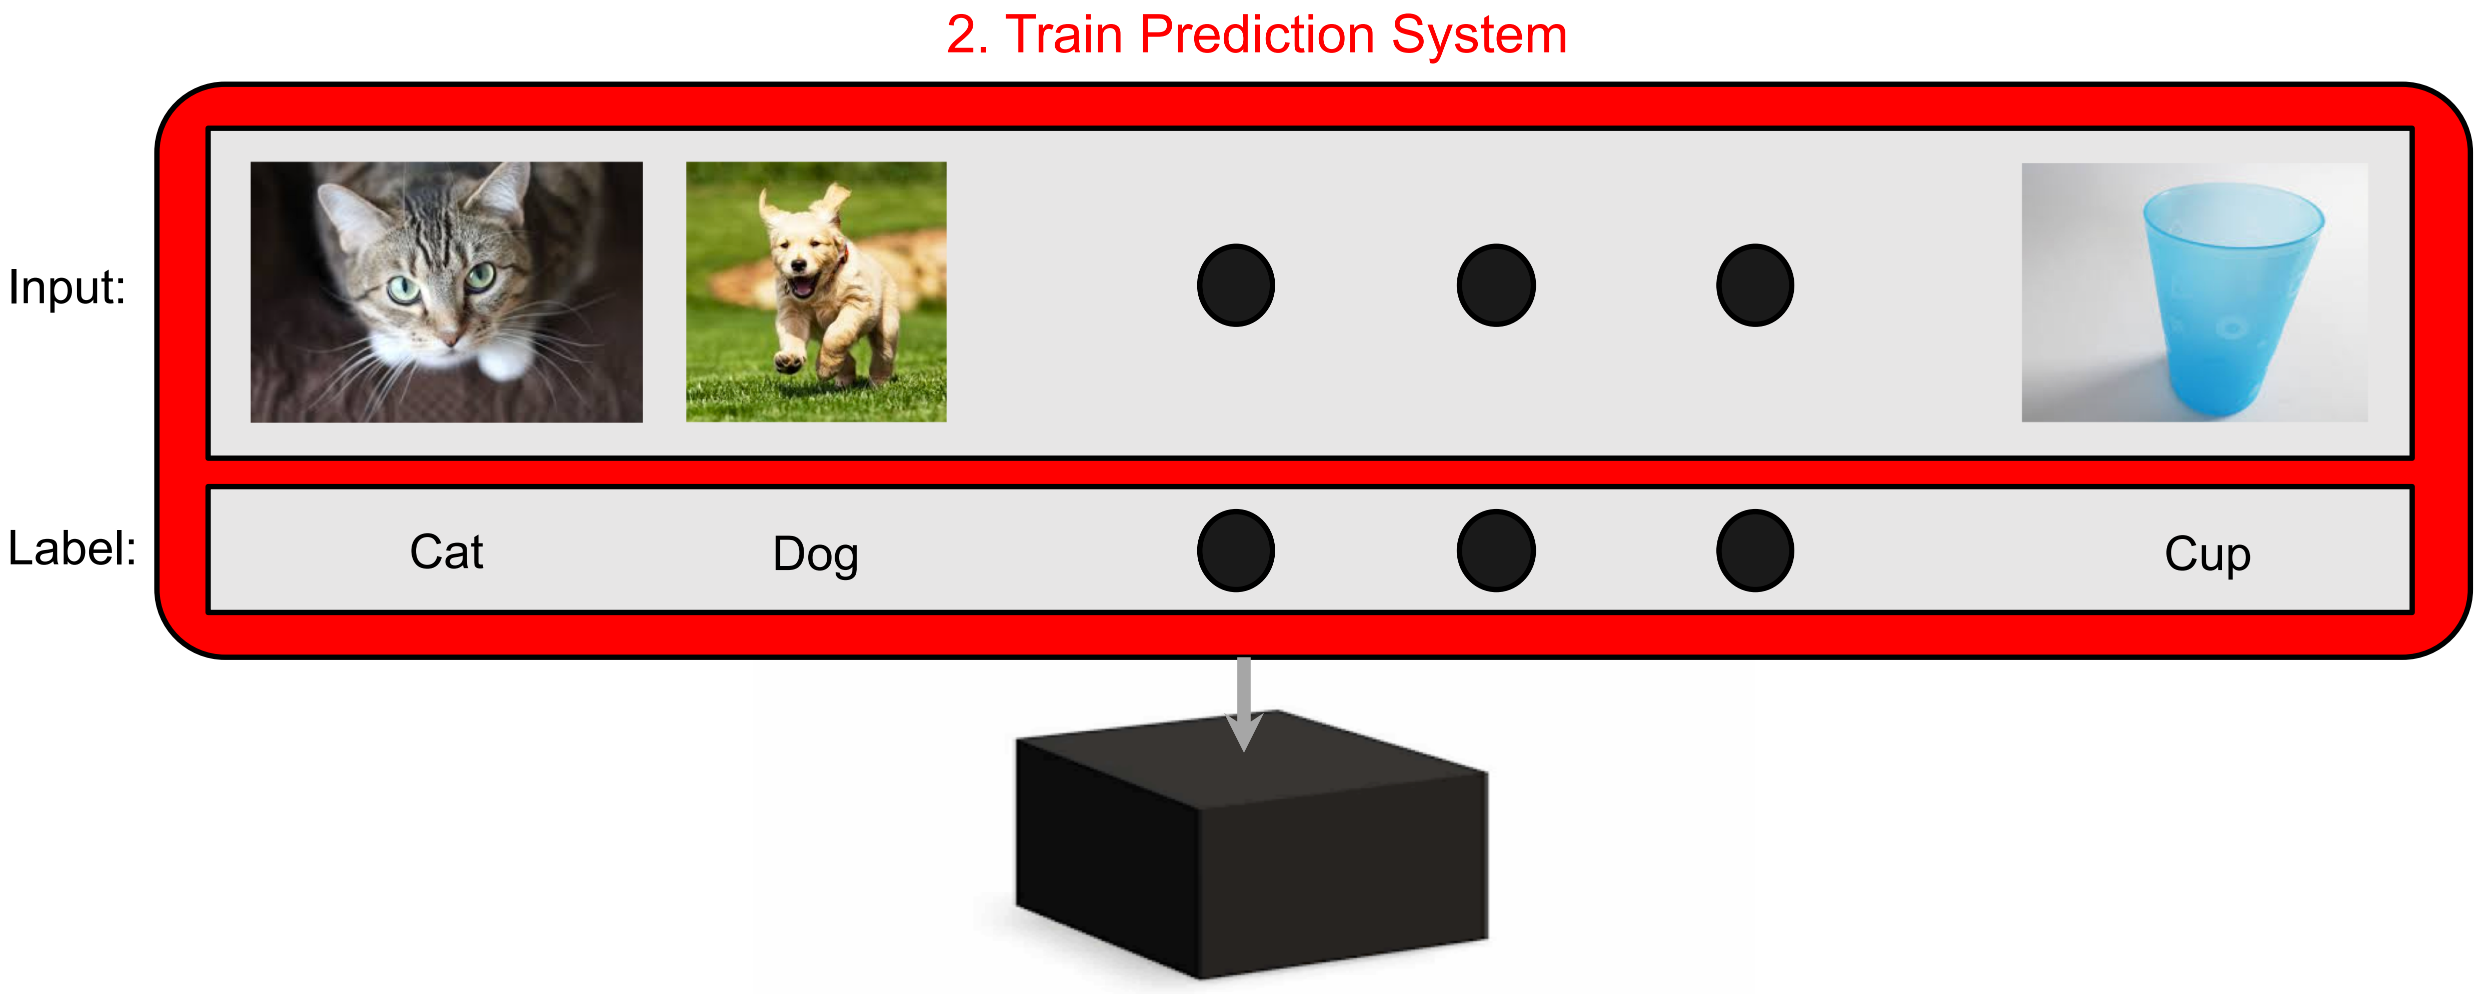

Recall: Need Datasets to Train & Evaluate Algorithms

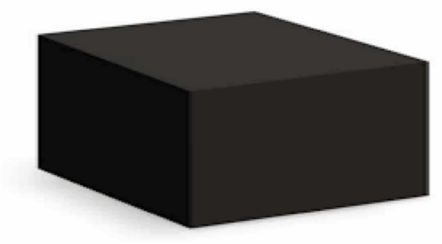

Prediction Model

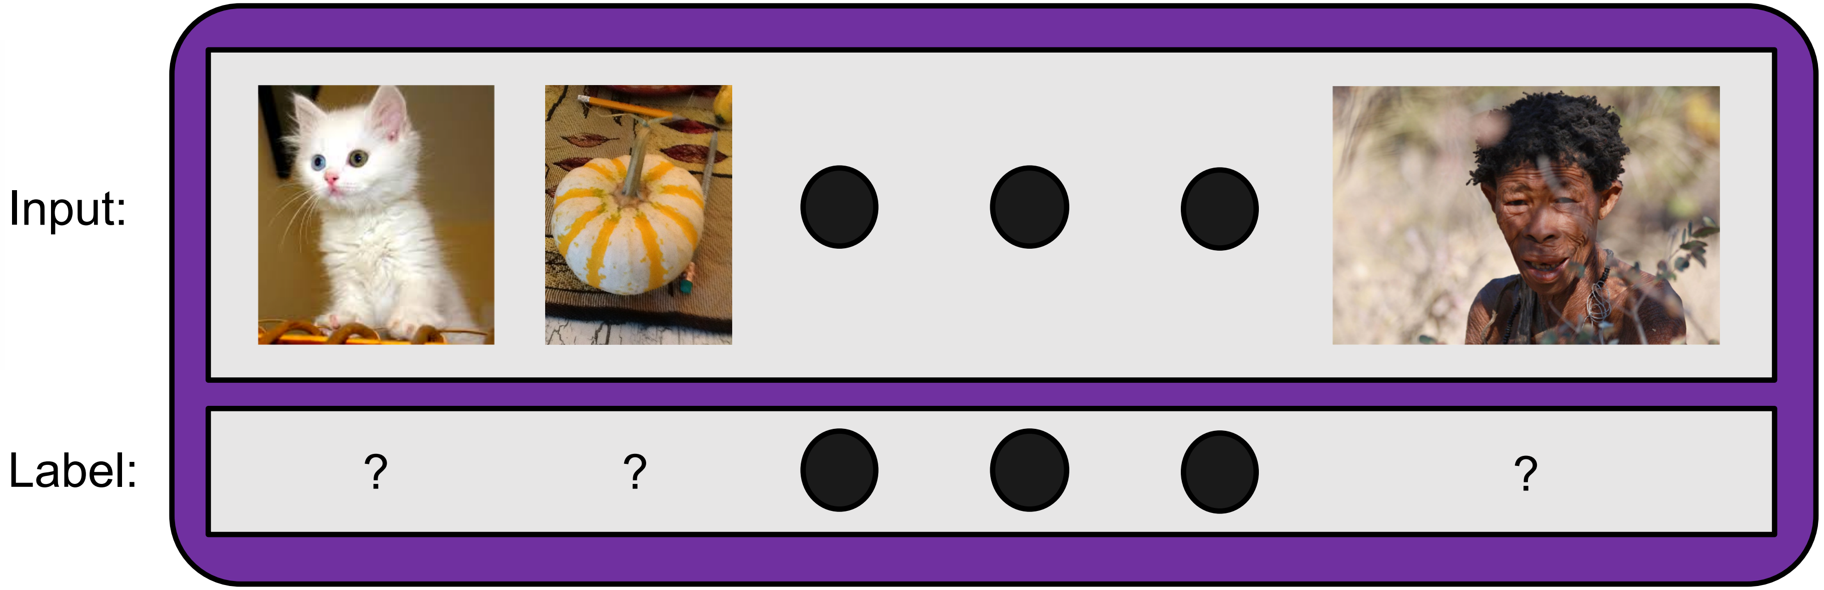

Recall: Need Datasets to Train & Evaluate Algorithms

1. Apply Prediction System

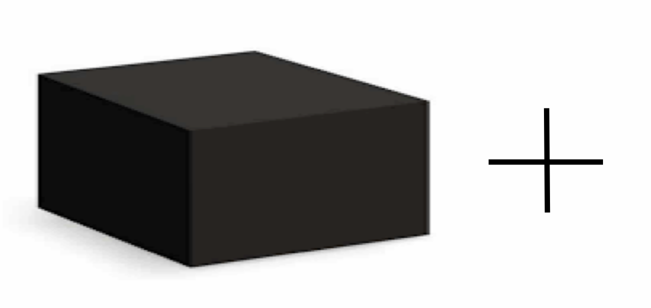

Prediction Model

+

Input:

Predicted Label:

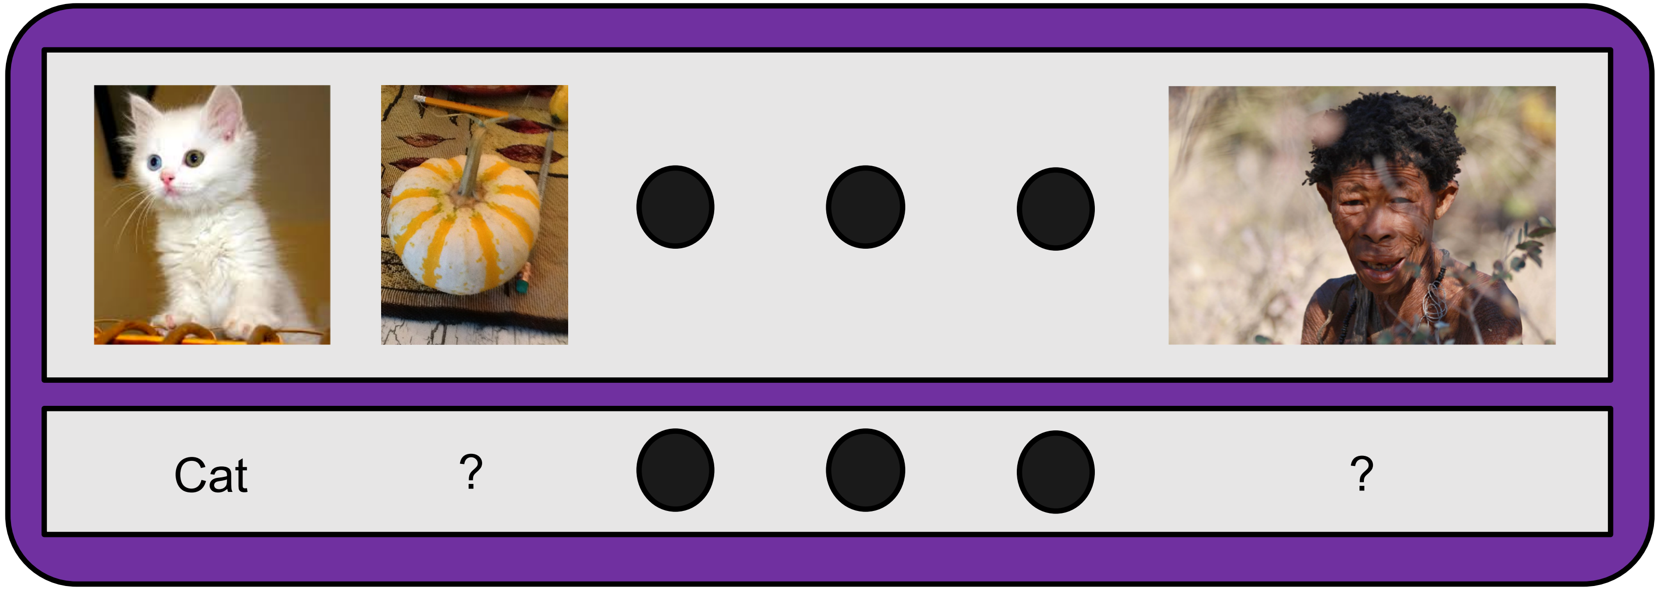

Recall: Need Datasets to Train & Evaluate Algorithms

1. Apply Prediction System

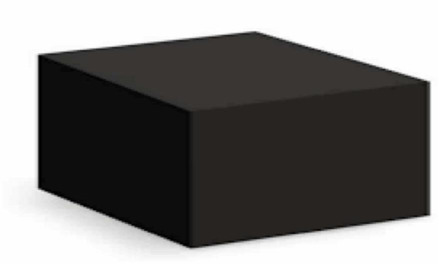

Prediction Model

+ Input:

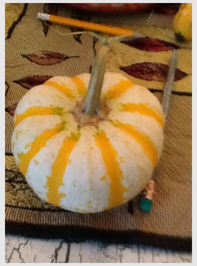

Predicted Label: Cat Giraffe ?

Recall: Need Datasets to Train & Evaluate Algorithms

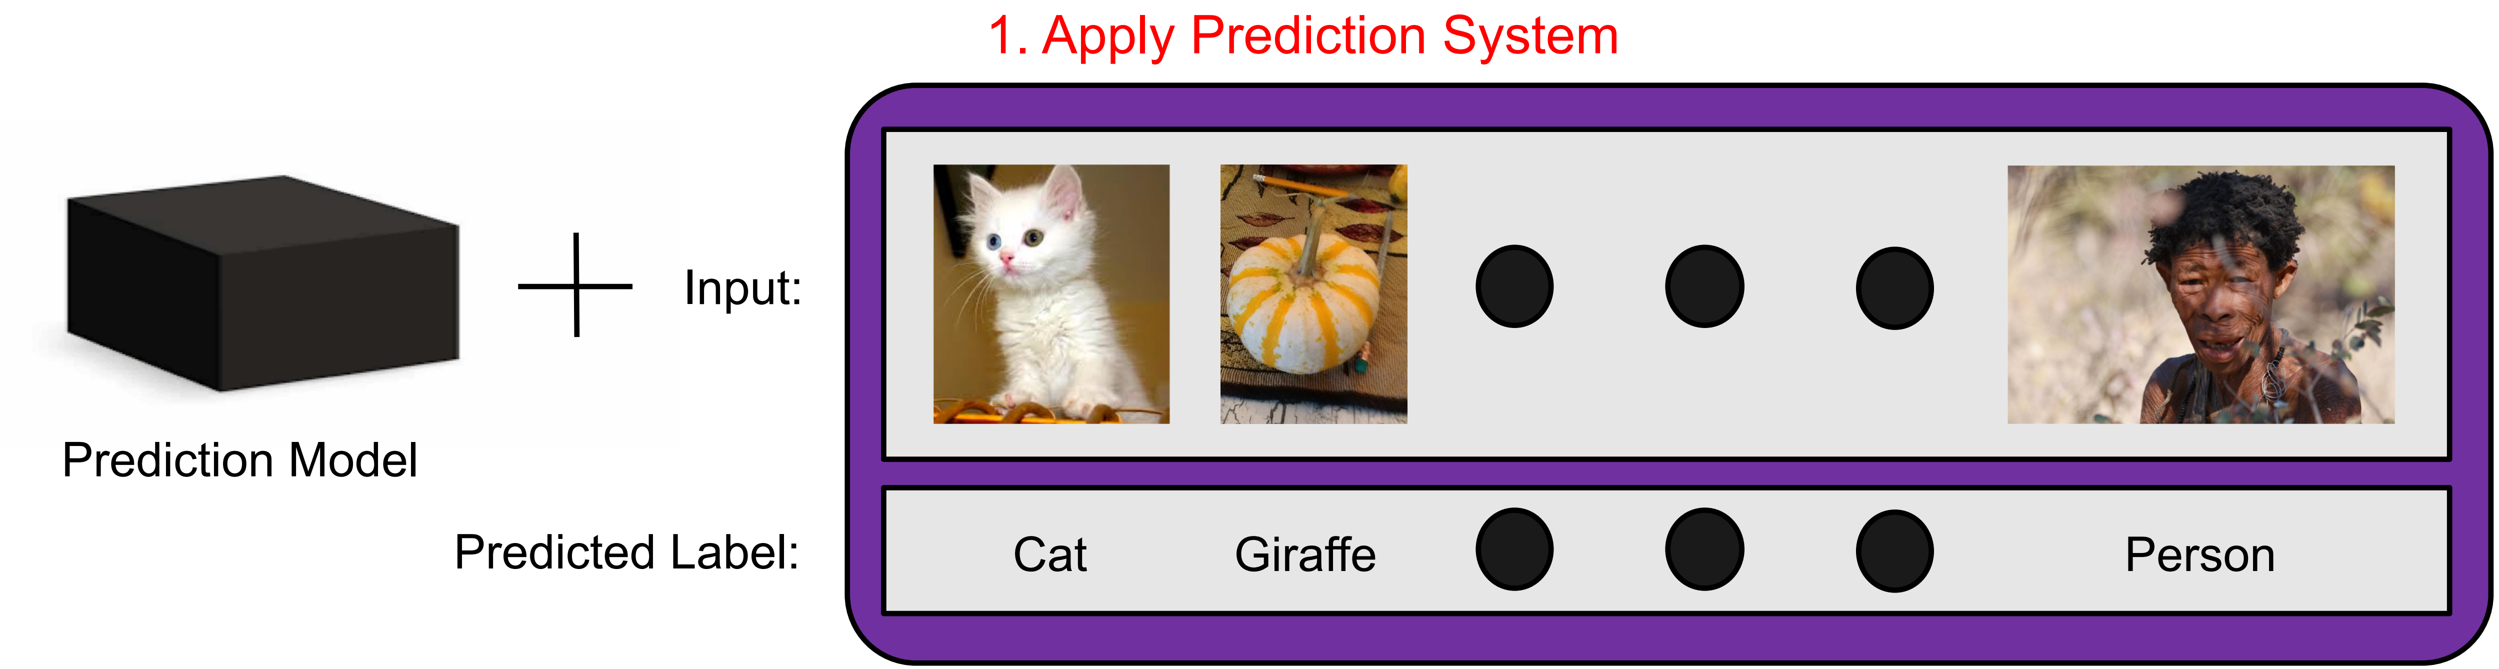

Recall: Need Datasets to Train & Evaluate Algorithms

2. Tally Percentage of Correct Results

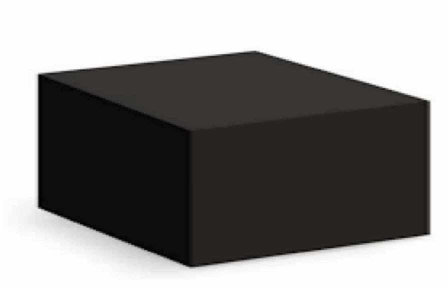

Prediction Model

+ Input:

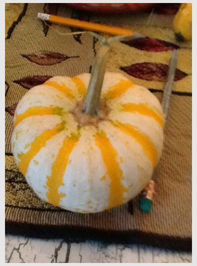

Predicted Label: Cat Giraffe Person

Recall: Need Datasets to Train & Evaluate Algorithms

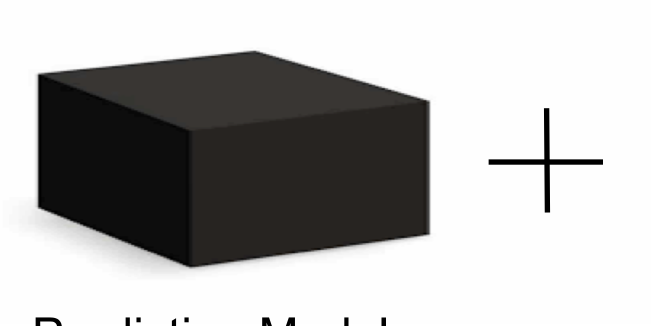

Prediction Model

+

2. Tally Percentage of Correct Results

Input:		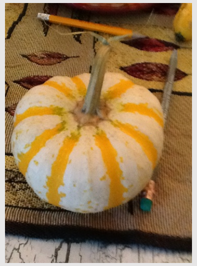	●	●	●	
Predicted Label:	Cat	Giraffe	●	●	●	Person
Human Annotated Label:	Cat	Pumpkin	●	●	●	Person

Recall: Need Datasets to Train & Evaluate Algorithms

2. Tally Percentage of Correct Results

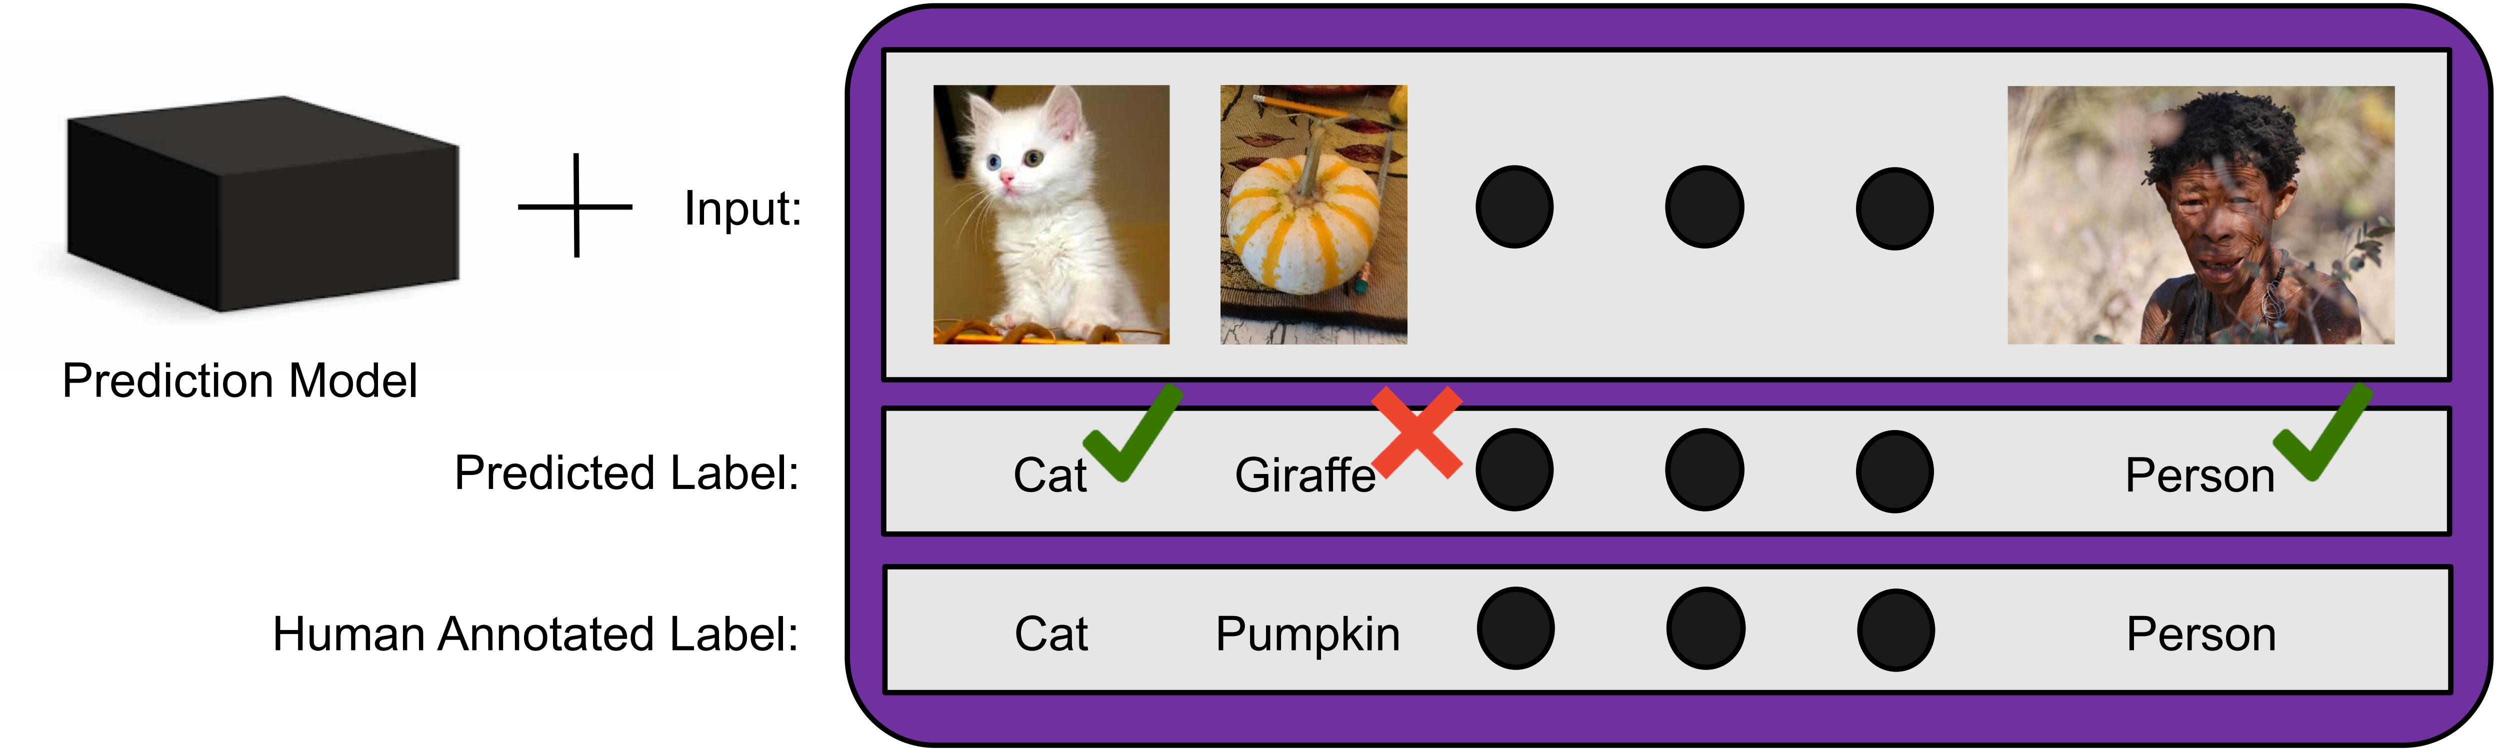

Object Recognition Datasets

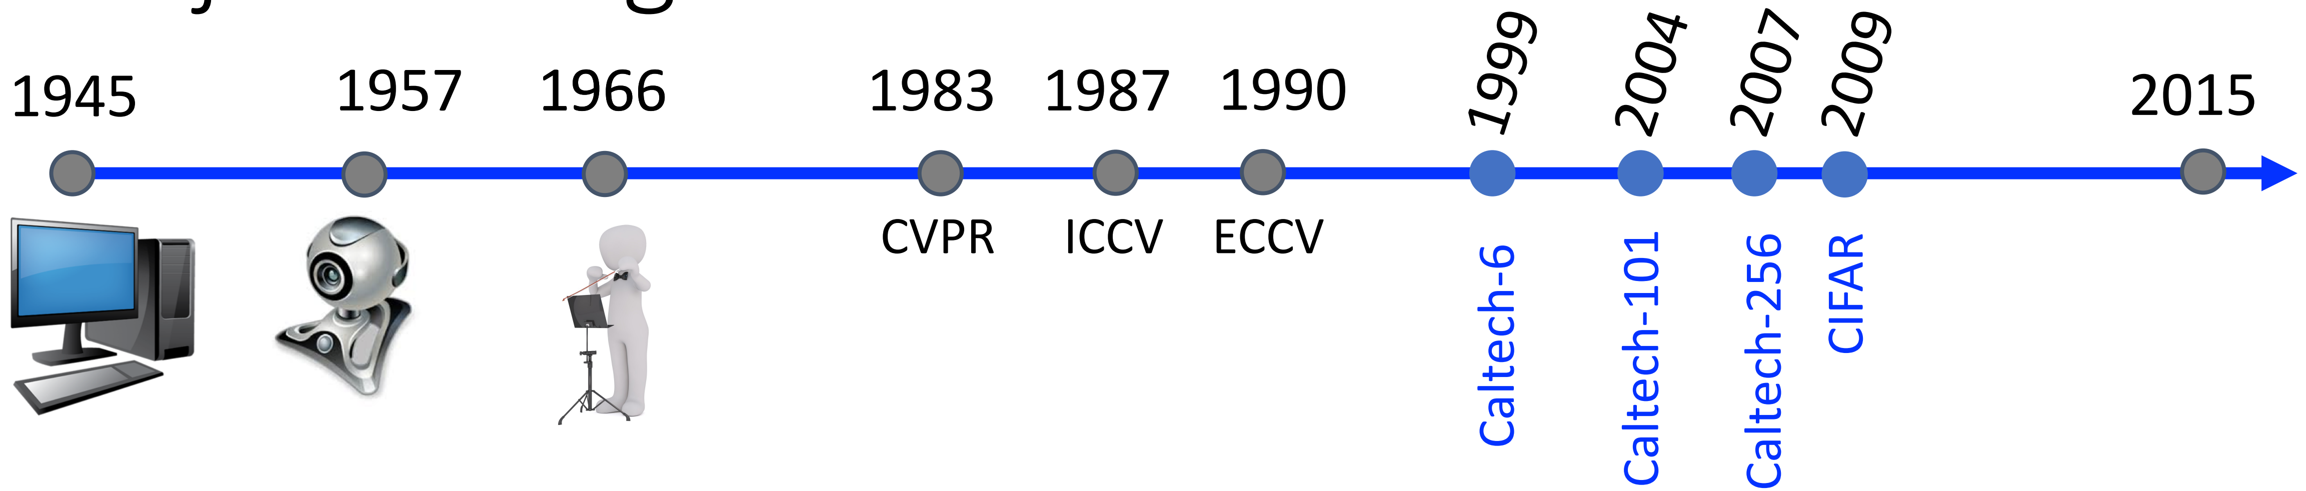

Object Recognition Datasets

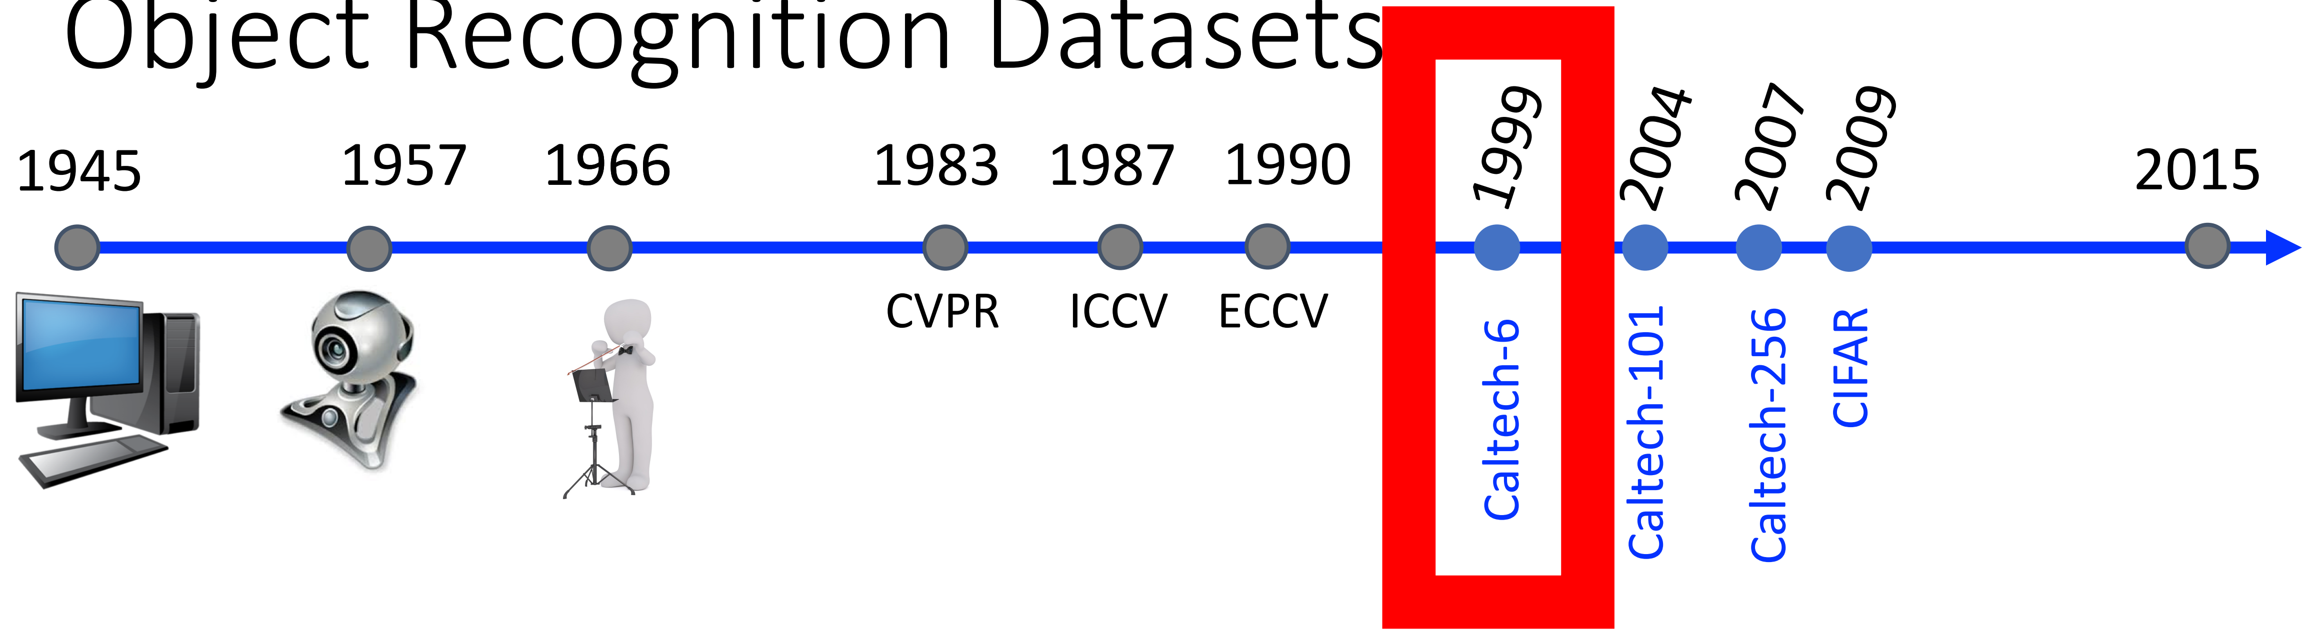

Object Recognition Datasets: Caltech-6

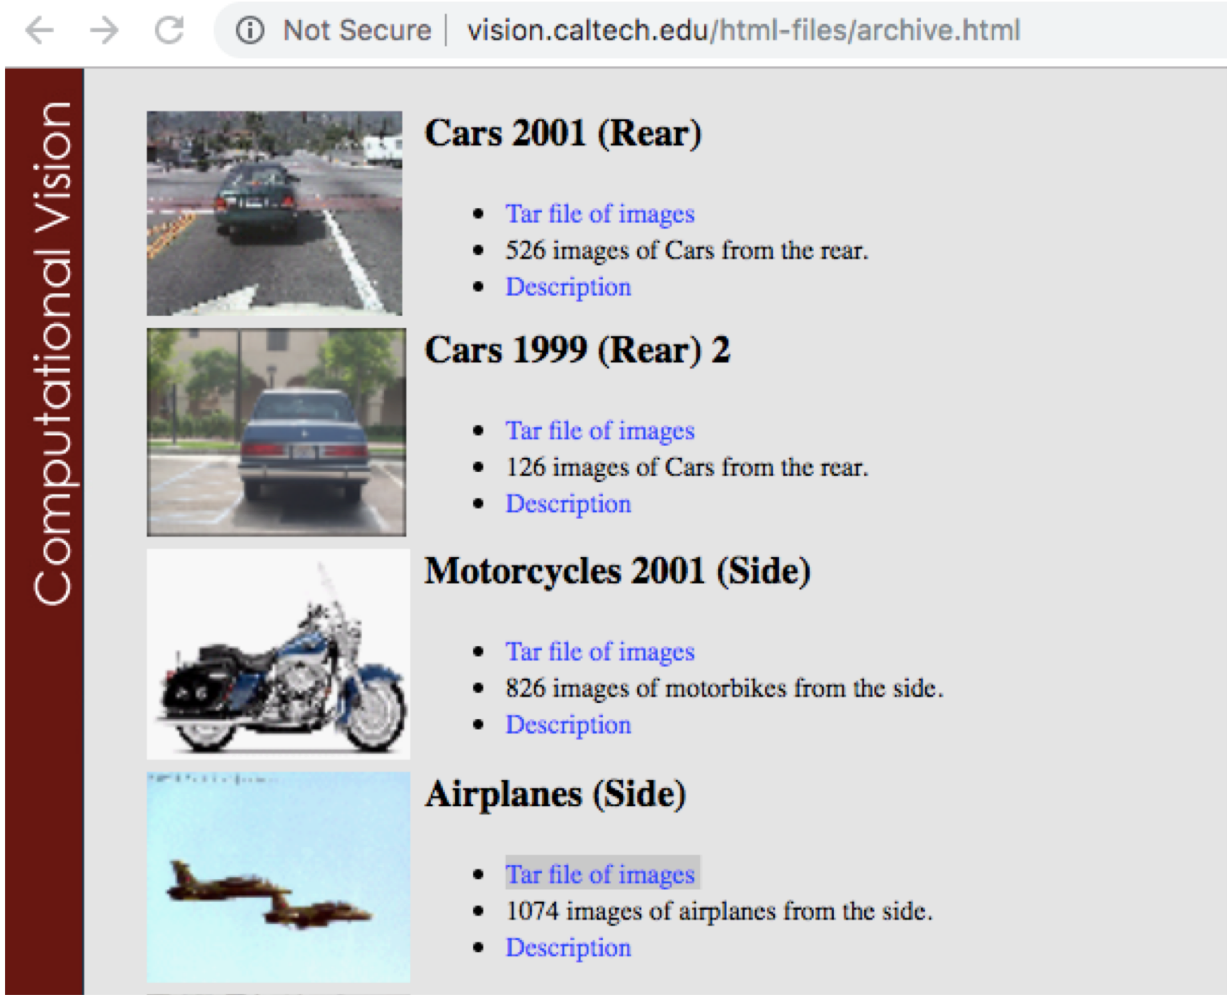
Not Secure | vision.caltech.edu/html-files/archive.html

Computational Vision

Cars 2001 (Rear)

- Tar file of images
- 526 images of Cars from the rear.
- Description

Cars 1999 (Rear) 2

- Tar file of images
- 126 images of Cars from the rear.
- Description

Motorcycles 2001 (Side)

- Tar file of images
- 826 images of motorbikes from the side.
- Description

Airplanes (Side)

- Tar file of images
- 1074 images of airplanes from the side.
- Description

Given an object category, students (1) took pictures or (2) collected images from the web of it

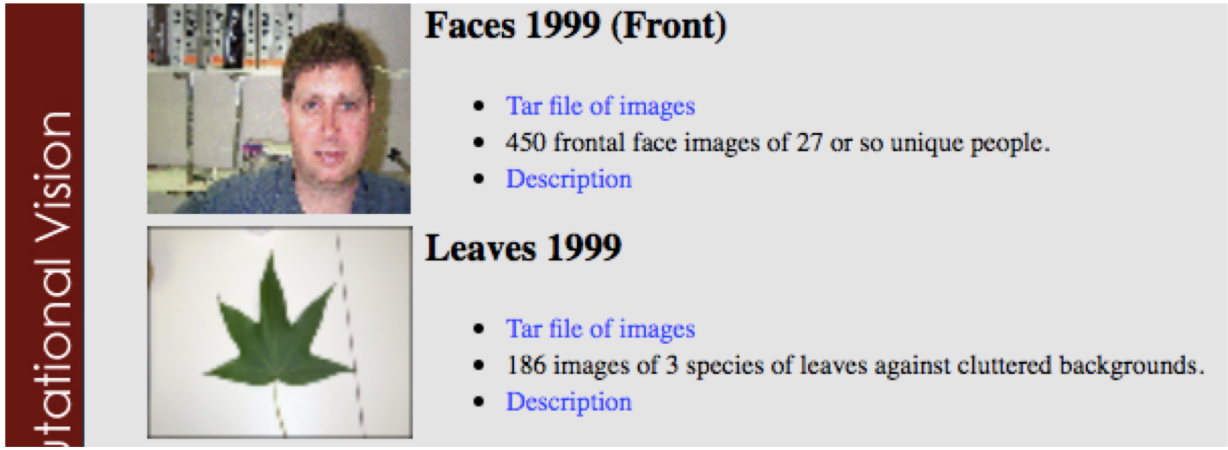
utational Vision

Faces 1999 (Front)

- Tar file of images
- 450 frontal face images of 27 or so unique people.
- Description

Leaves 1999

- Tar file of images
- 186 images of 3 species of leaves against cluttered backgrounds.
- Description

http://www.vision.caltech.edu/html-files/archive.html

Object Recognition Datasets

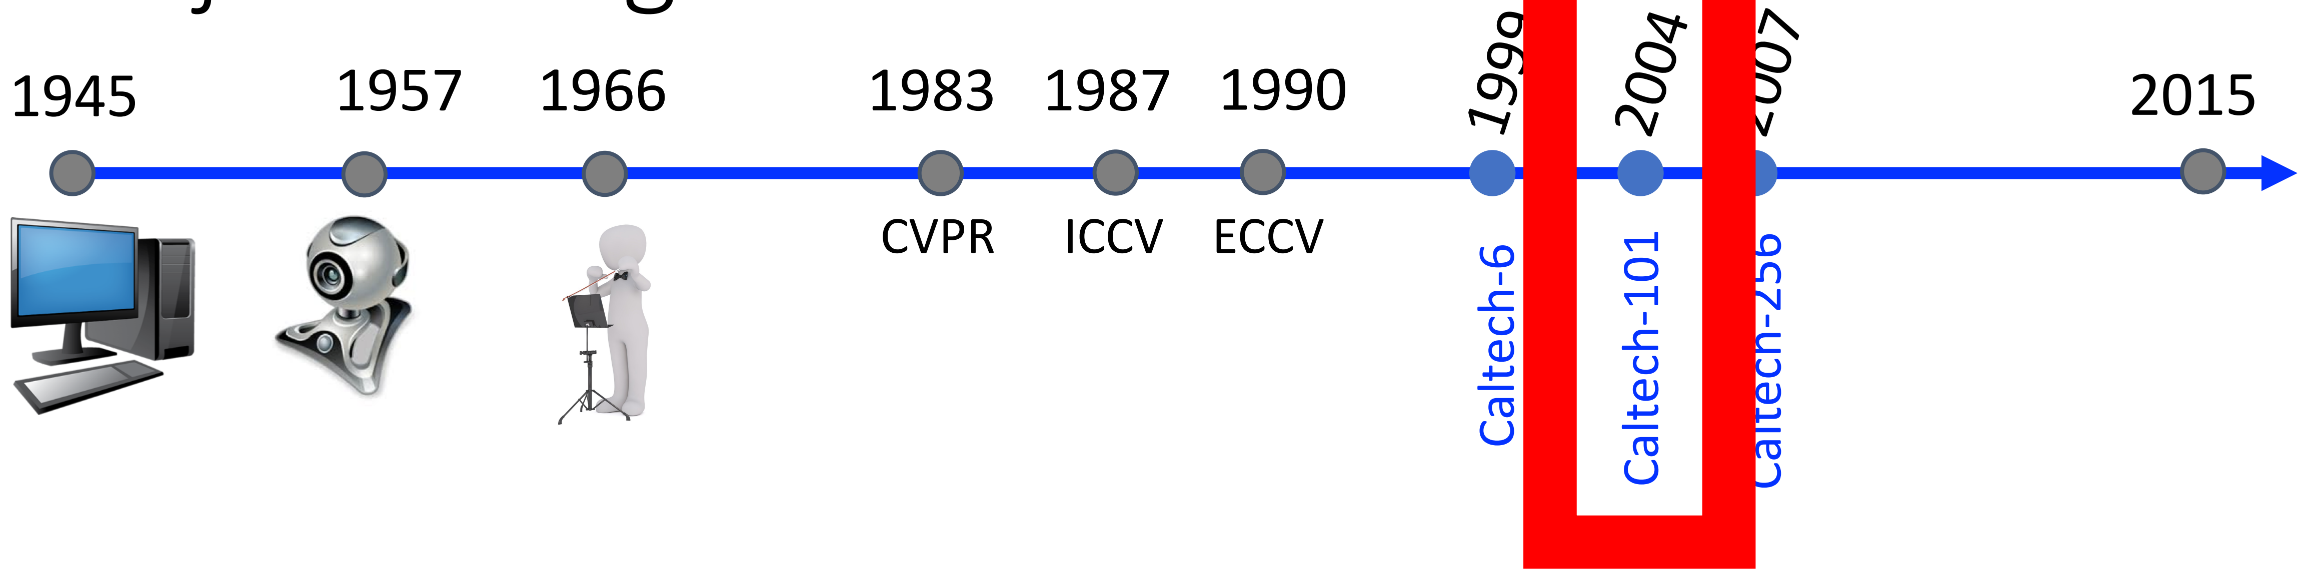

Object Recognition Datasets: Caltech-101

1. Category Selection

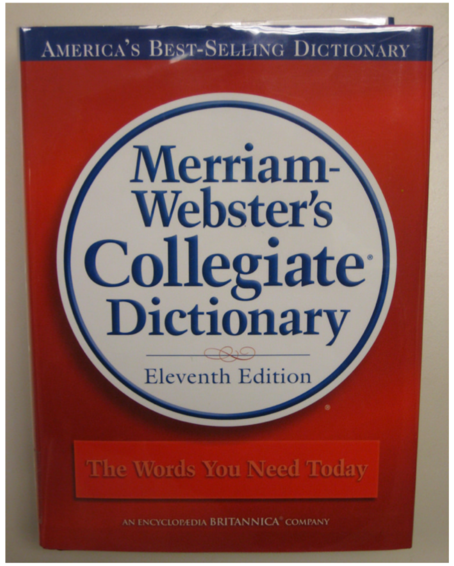

Flipped through a dictionary and chose 101 categories associated with a drawing

Li Fei Fei, Rob Fergus, & Pietro Perona. Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. CVPR 2004.

Object Recognition Datasets: Caltech-101

1. Category Selection

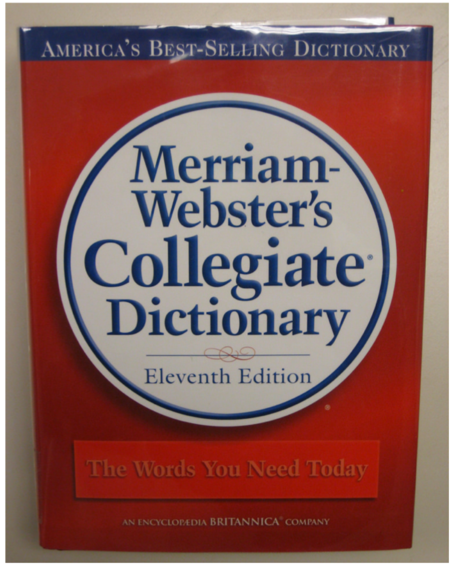

Flipped through a dictionary and chose 101 categories associated with a drawing

2. Image Collection

Search for each category

Li Fei Fei, Rob Fergus, & Pietro Perona. Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. CVPR 2004.

Object Recognition Datasets: Caltech-101

1. Category Selection

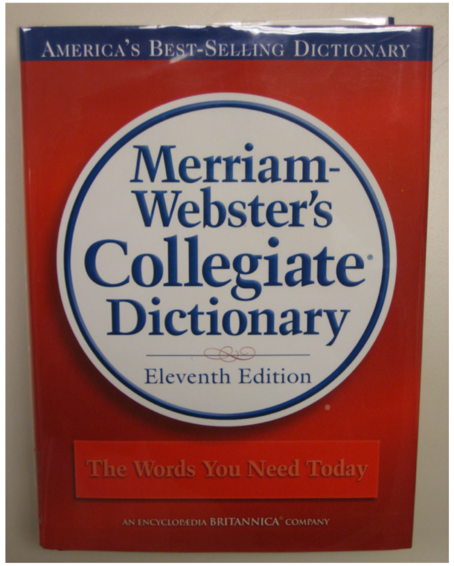

Flipped through a dictionary and chose 101 categories associated with a drawing

2. Image Collection

Search for each category

3. Human Verification

- 2 graduate students reviewed & discarded irrelevant images

- Result is 9,144 grayscale 300x200 pixel images with 45-400 images per category

Li Fei Fei, Rob Fergus, & Pietro Perona. Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. CVPR 2004.

Object Recognition Datasets: Caltech-101

Two random samples per category

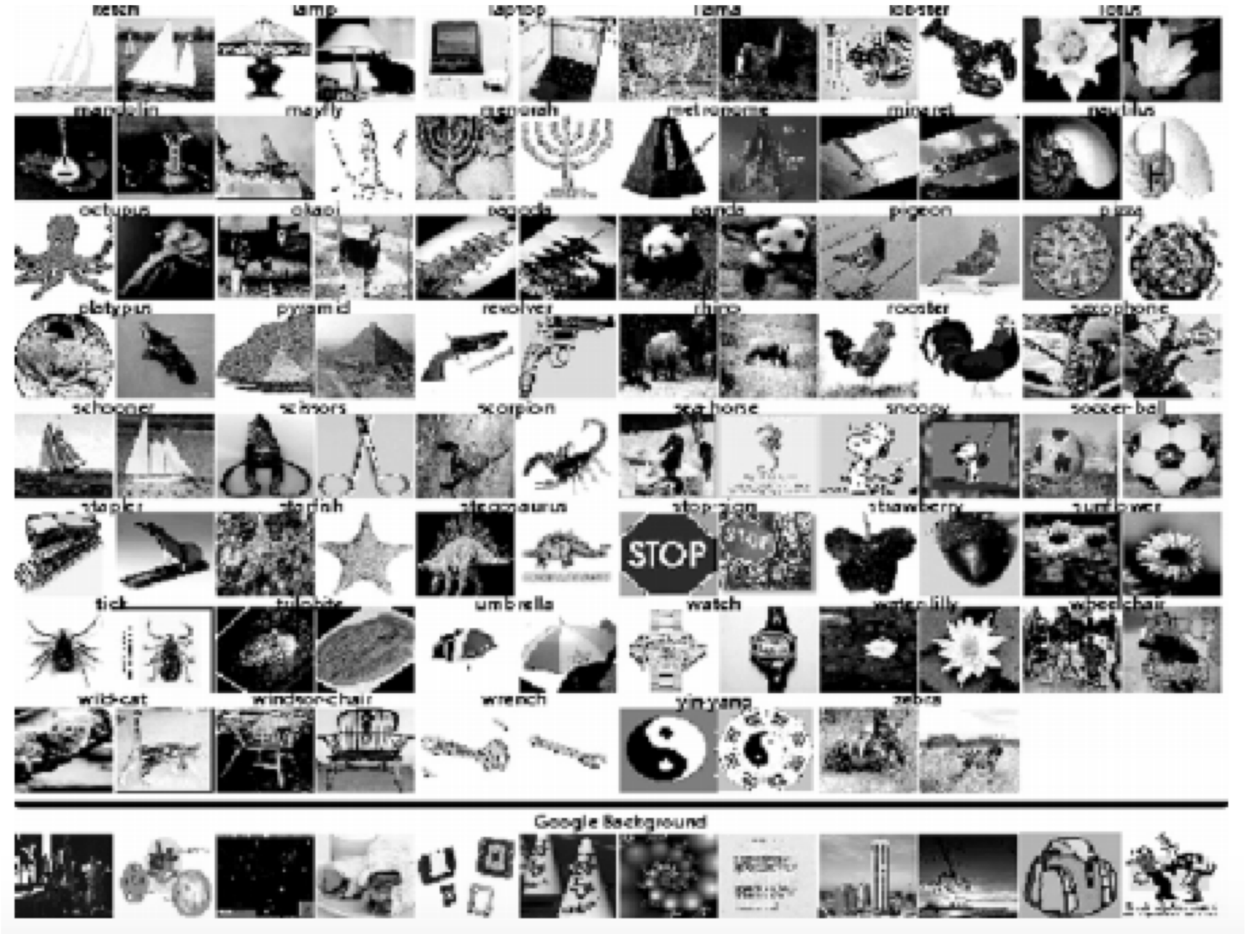

Dataset location: http://vision.caltech.edu

Object Recognition Datasets: Caltech-101

Two random samples per category

Was 15x Larger Than Existing Datasets

Dataset location: http://vision.caltech.edu

Object Recognition Datasets: Caltech-101

Charting progress of algorithms

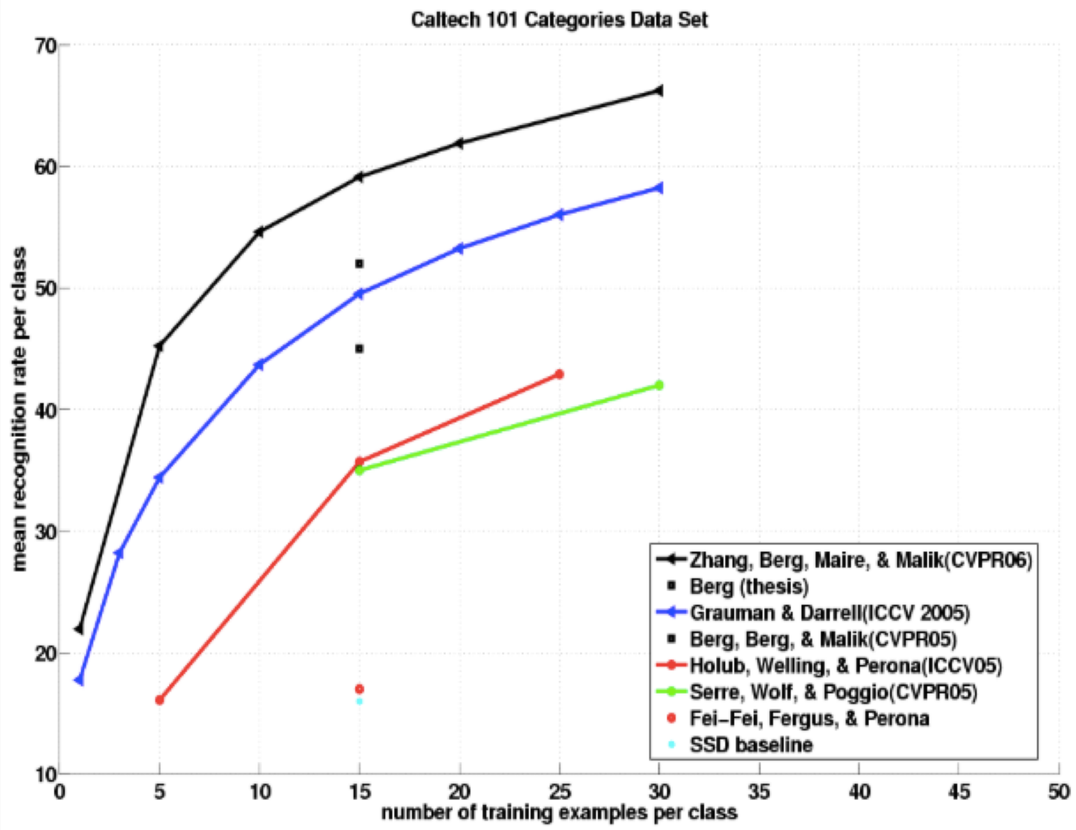

Latest results (March 2006) on the Caltech 101 from a variety of groups. (published results only).

If you would like to include your algorithm's performance please email us at holub@caltech.edu or greg@vision.caltech.edu with a citation and your results. Thanks!

We are also interested in the time it takes to run your algorithm. Both during the training and during the classification stage

Plot courtesy of Hao Zhang.

Update by holub, April 2006.

Dataset location: http://vision.caltech.edu

Object Recognition Datasets

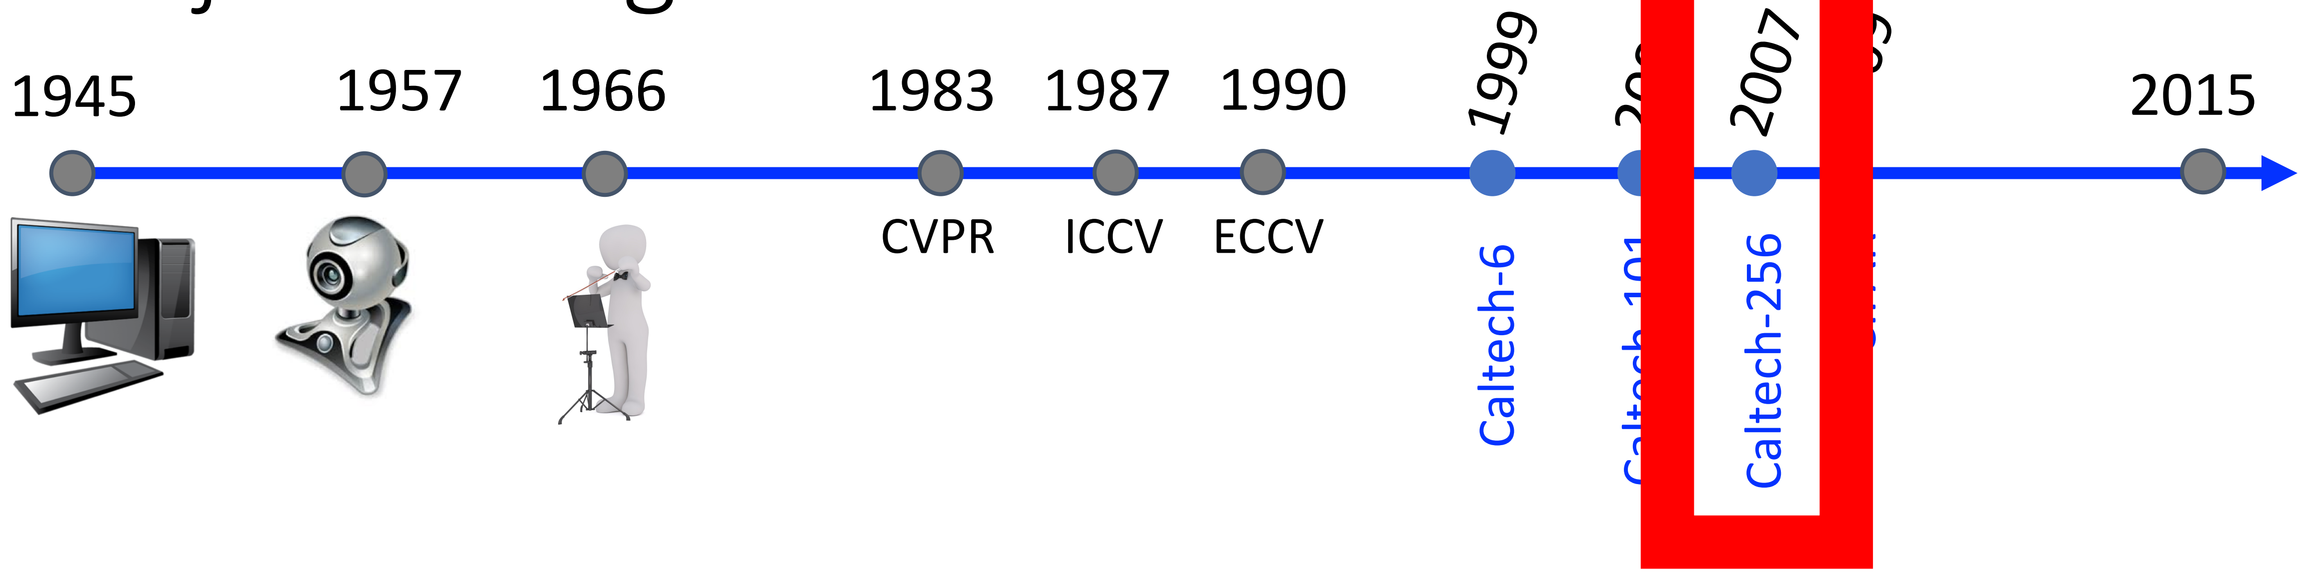

Caltech-256

1. Category Selection

- Several individuals chose ~300 object categories

- Taxonomy of categories grouped around (in)animate

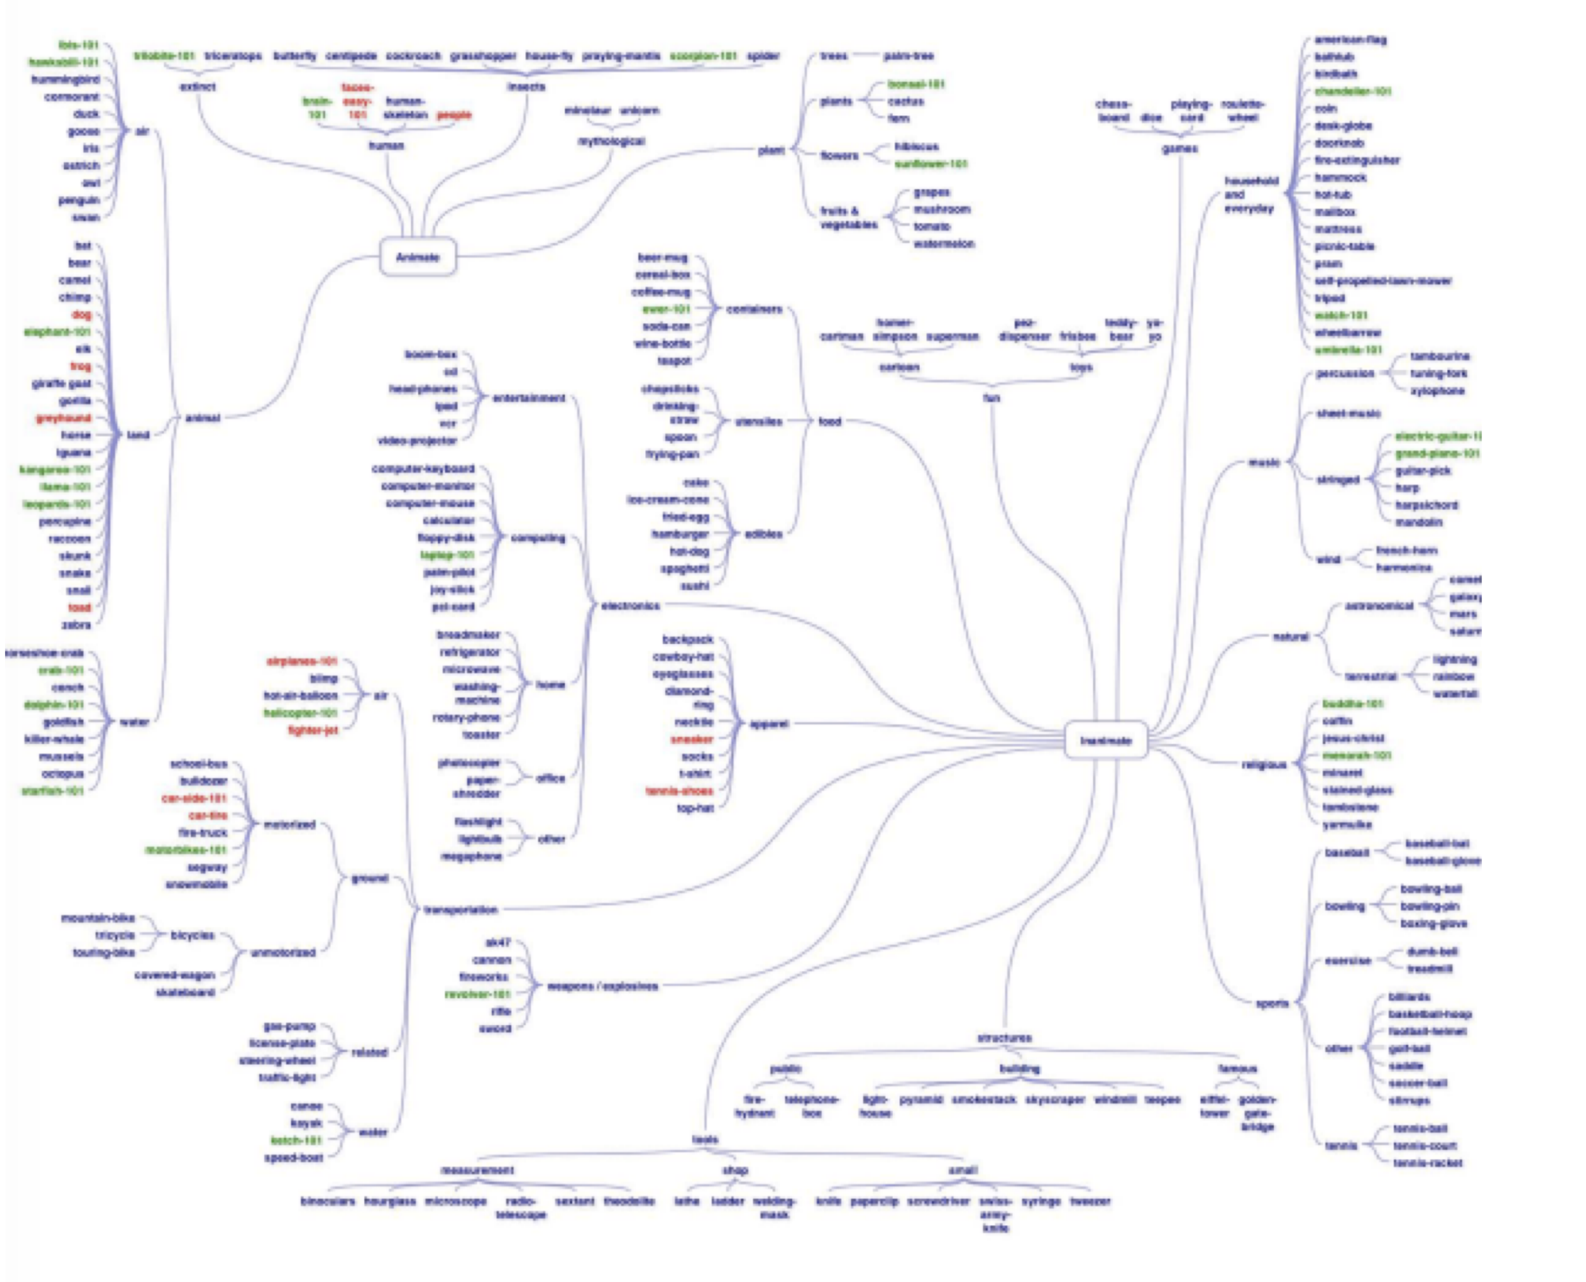

Greg Griffin, Alex Holub, & Pietro Perona. Caltech-256 Object Category Dataset. Technical Report 2007.

Object Recognition Datasets: Caltech-256

1. Category Selection

- Several individuals chose ~300 object categories

- Taxonomy of categories grouped around (in)animate

2. Image Collection

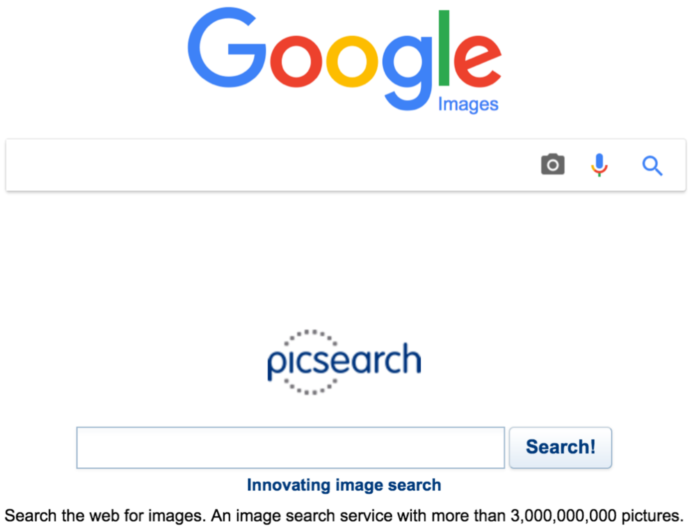

3. Human Verification

- One of 4 people rate each image for 92,652 images

- Keep only "good" images

- Result is 9,104 images spanning 256 categories that each have >80 good images

Greg Griffin, Alex Holub, & Pietro Perona. Caltech-256 Object Category Dataset. Technical Report 2007.

Object Recognition Datasets: Caltech-256

<u>Rating Instructions</u>

1. Good: A clear example of the visual category
2. Bad: A confusing, occluded, or artistic example
    1. Image is very cluttered
    2. Image is a line drawing
    3. Image is an abstract artistic representation
    4. Object only occupies small fraction of the image
3. Not applicable: Not an example of the category

3. Human Verification

- One of 4 people rate each image for 92,652 images

- Keep only “good” images

- Result is 9,104 images spanning 256 categories that each have >80 good images

Greg Griffin, Alex Holub, & Pietro Perona. Caltech-256 Object Category Dataset. Technical Report 2007.

Object Recognition Datasets: Caltech-256

e.g., dice

3. Human Verification

- One of 4 people rate each image for 92,652 images
- Keep only "good" images
- Result is 9,104 images spanning 256 categories that each have >80 good images

Greg Griffin, Alex Holub, & Pietro Perona. Caltech-256 Object Category Dataset. Technical Report 2007.

Object Recognition Datasets: Caltech-256

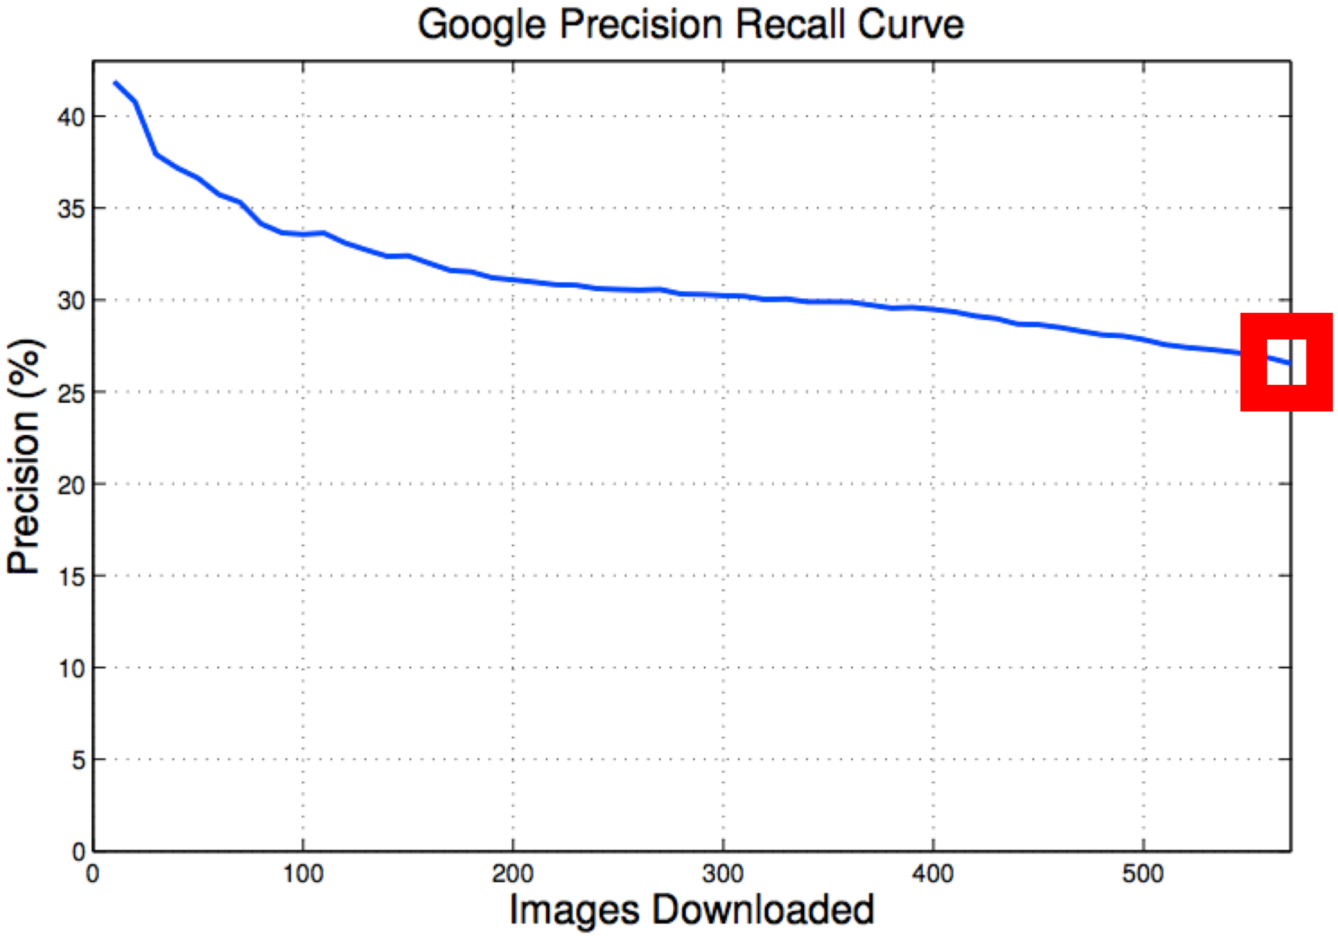

e.g., What percentage of Google images judged ``good"?

3. Human Verification

- One of 4 people rate each image for 92,652 images
- Keep only "good" images
- Result is 9,104 images spanning 256 categories that each have >80 good images

Search engine accuracy?

Greg Griffin, Alex Holub, & Pietro Perona. Caltech-256 Object Category Dataset. Technical Report 2007.

Object Recognition Datasets: Caltech-256

Charting progress of algorithms

If you would like to share performance results as well as your confusion matrix, please send them to `caltech256@vision.caltech.edu`. We will try to keep our comparison of performance as up-to-date as possible. For more details see

`http://www.vision.caltech.edu/Image_Datasets/Caltech256`

Greg Griffin, Alex Holub, & Pietro Perona. Caltech-256 Object Category Dataset. Technical Report 2007.

Object Recognition Datasets

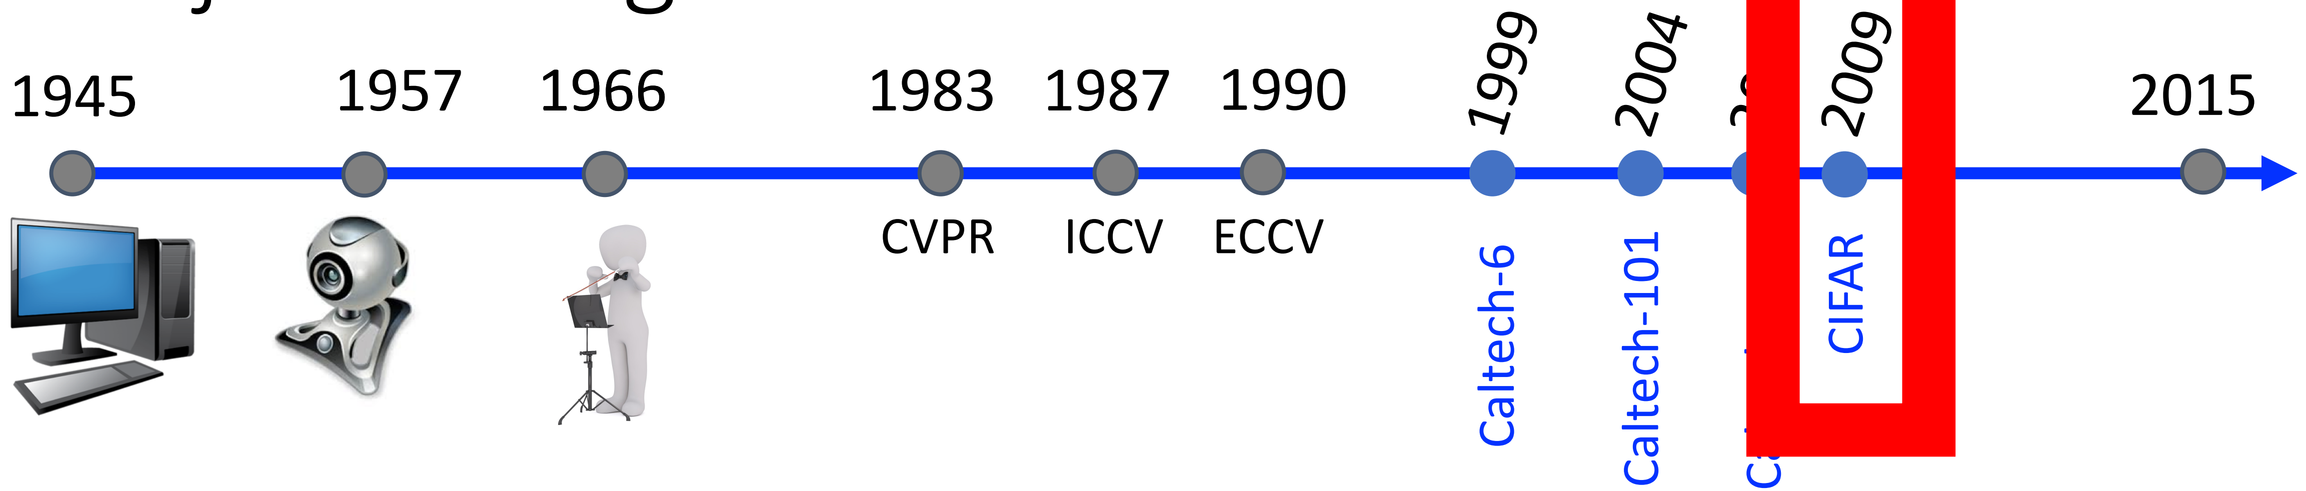

Object Recognition Datasets: CIFAR

1. Category Selection

100 categories taken from the “tiny images” dataset

→

2. Image Collection

Images taken from from the “tiny images” dataset

→

3. Human Verification

- Students paid to reject images not in category

- Authors verified labels

Alex Krizhevsky & Geoffrey Hinton. Learning Multiple Layers of Features from Tiny Images. Technical Report 2009.

CIFAR

Criteria for deciding whether to include an image

1. The main test is: Would you be quite likely to say the category name if asked to give a single basic category to describe the main object in the image?
2. It's worse to include one that shouldn't be included than to exclude one. False positives are worse than false negatives.
3. If there is more than one object that is roughly equally prominent, reject even if they are all of the right class.

4. If it is a line drawing or cartoon, reject. You can accept fairly photorealistic drawings that have internal texture.

5. Do not reject just because the viewpoint is unusual or the object is partially occluded (provided you think you might have assigned the right label without priming). We want ones with unusual viewpoints.

6. Do not reject just because the background is cluttered. We want some cluttered backgrounds. But also, do not reject just because the background is uniform.
7. Do not worry too much about accepting duplicates or near duplicates. If you are pretty sure it's a duplicate, reject it. But we will eliminate any remaining duplicates later, so including duplicates is not a bad error.
8. If a category has two meanings (like mouse), only include the main meaning. If there is doubt about what this is, then ask.

3. Human Verification

- Students paid to reject images not in category

- Authors verified labels

Alex Krizhevsky & Geoffrey Hinton. Learning Multiple Layers of Features from Tiny Images. Technical Report 2009.

Object Recognition Datasets

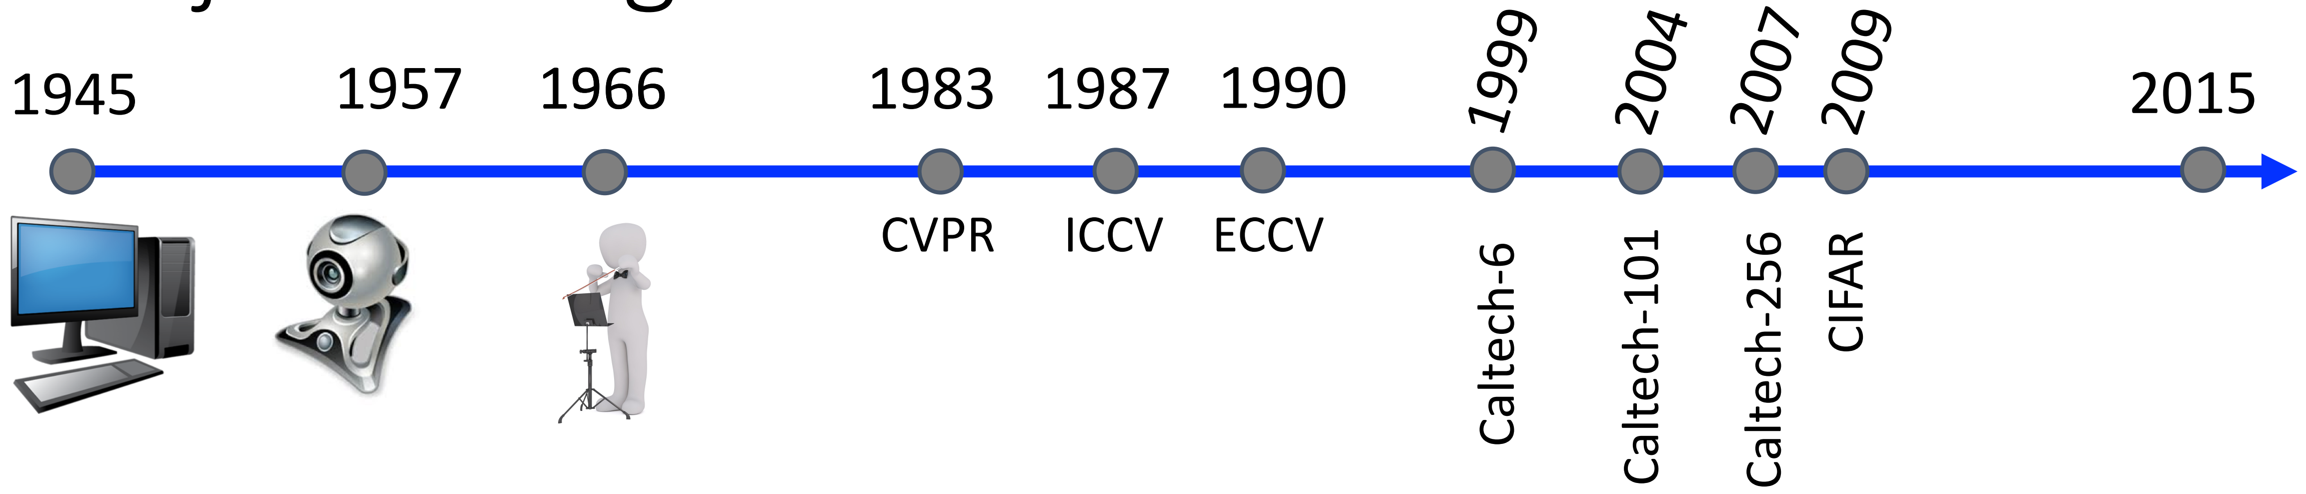

	Caltech-6	Caltech-101	Caltech-256	CIFAR
# Categories:	6	101	256	100
Minimum # Samples/Category:	126+	31+	80+	600
# Images:	3,738	9,144	30,607	114,000

Object Recognition Datasets: Summary

- Key steps in creating dataset:

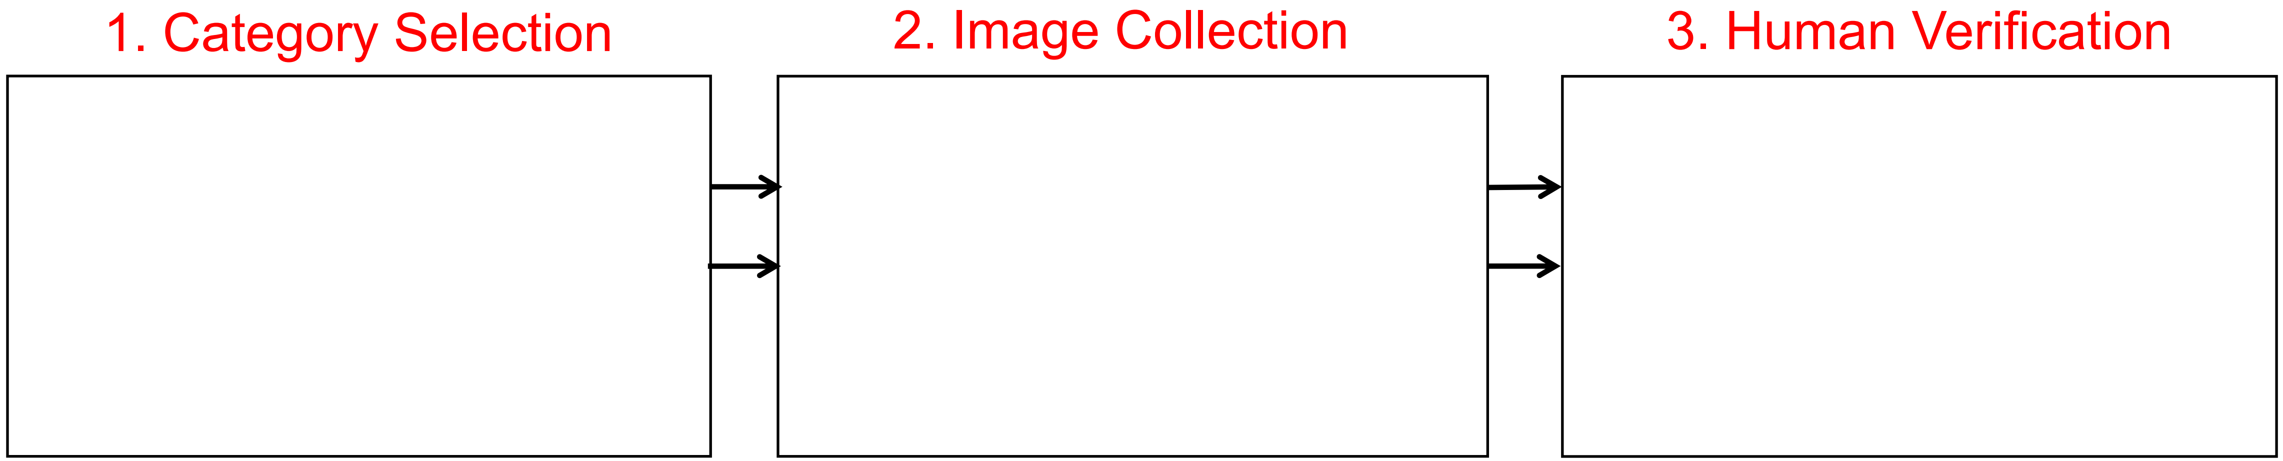

Object Recognition Datasets: Summary

- Key steps in creating dataset:

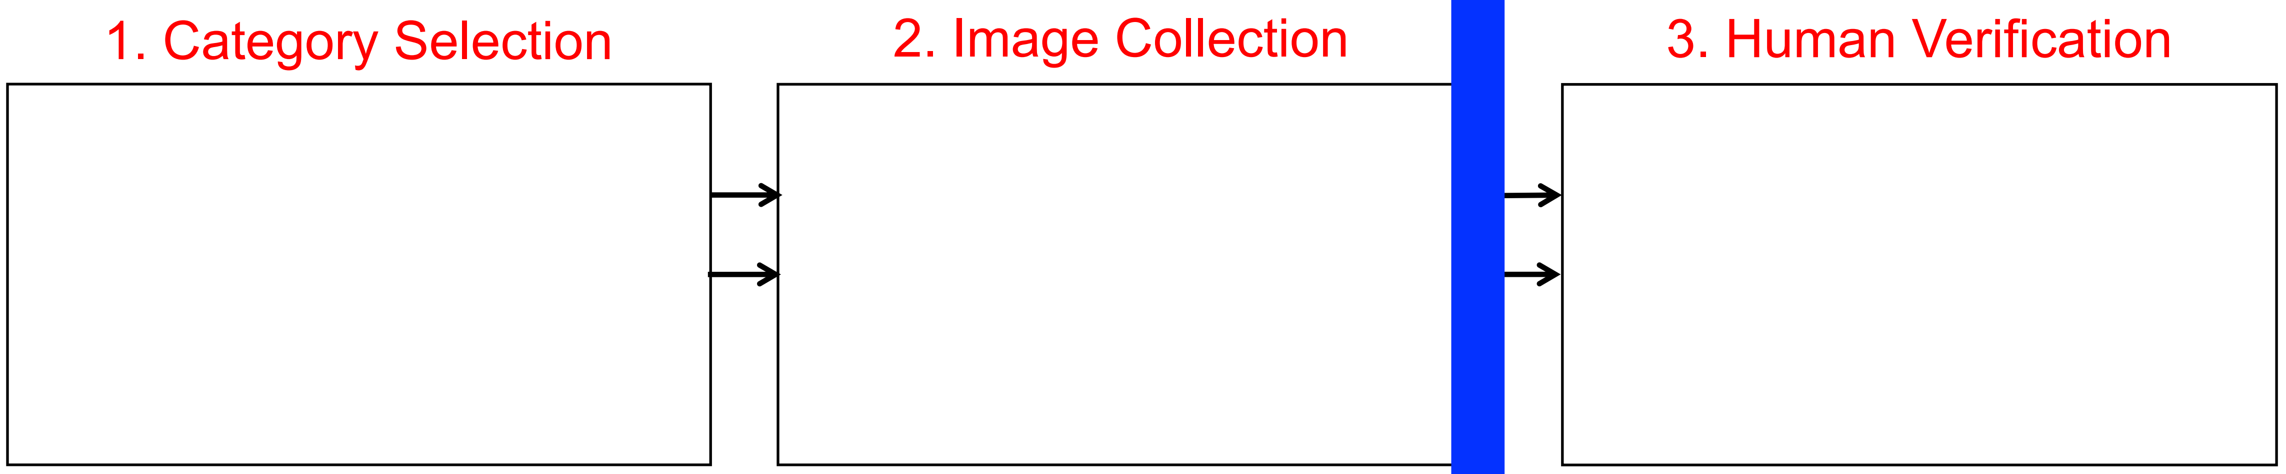

Limitation: relies on a few in-house people to verify images

Today’s Topics

- Object recognition applications
- Object recognition datasets: key steps in creating them
- Object recognition datasets: scaling up their size with *crowdsourcing*
- Scaling up community working on object recognition with *workshop challenges*
- Class Discussion
- Lab: cascading stylesheets and web page layout

Object Recognition Datasets

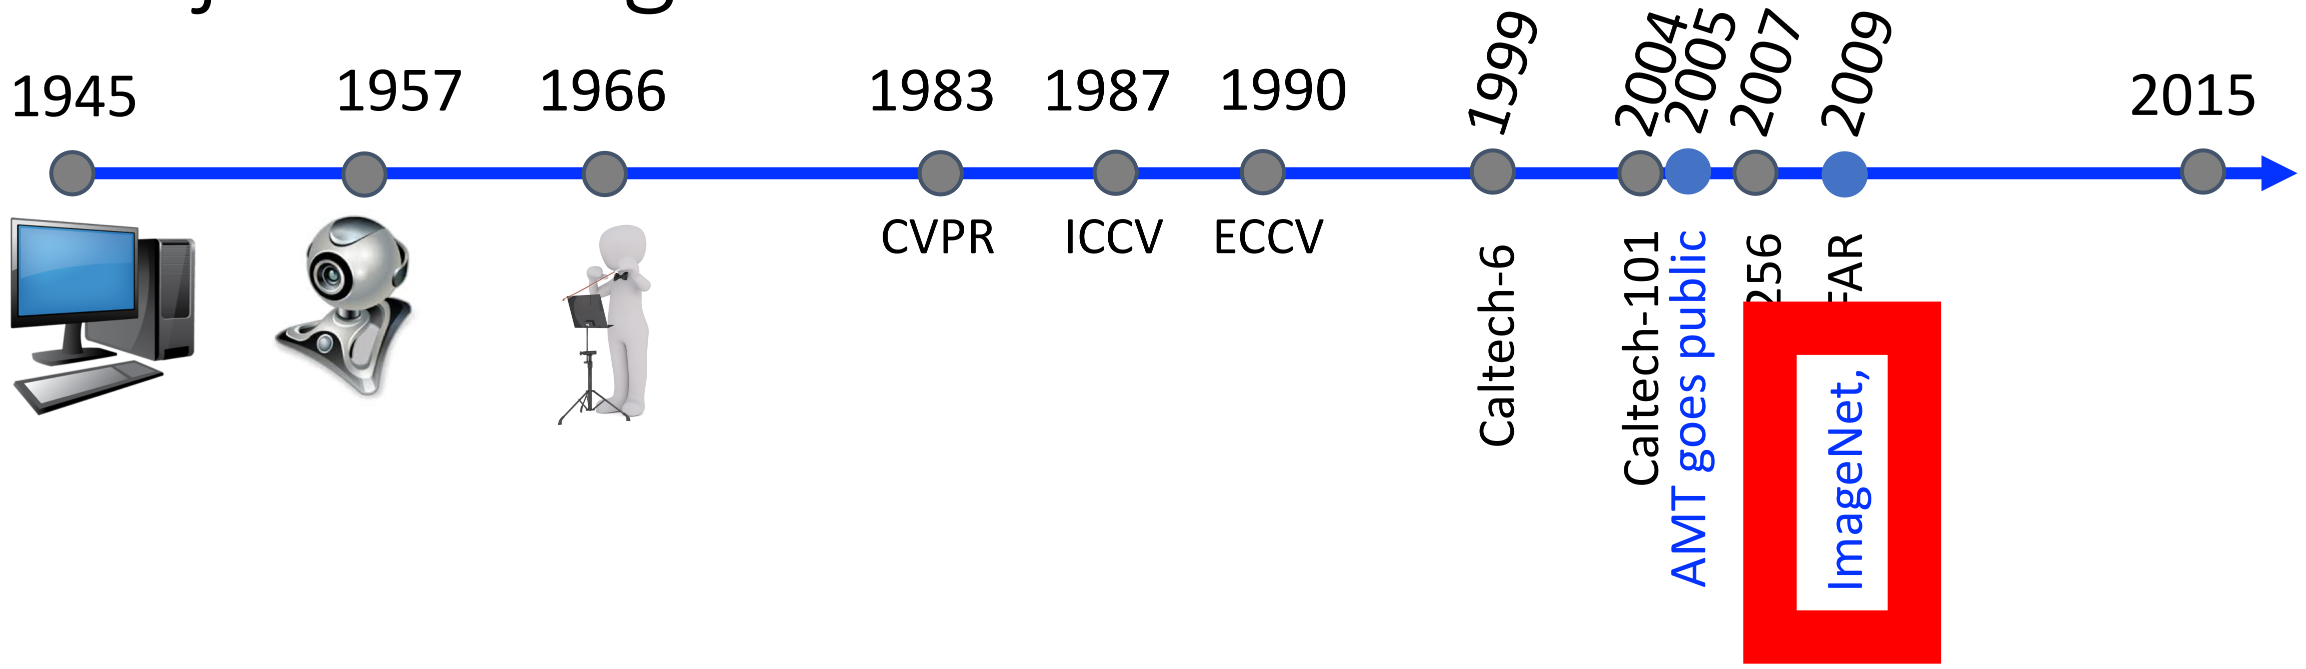

ImageNet’s founder, Fei-Fei Li, tells ImageNet’s story:

(Note: she previously created Caltech-101)

https://www.youtube.com/watch?v=40riCqvRoMs

(5:44 – 9:35)

Object Recognition Datasets: ImageNet

1. Category Selection

~10% of concepts (synonym sets) in WordNet taxonomy

e.g., two root-to-leaf branches of ImageNet with nine examples for each "synonym set"

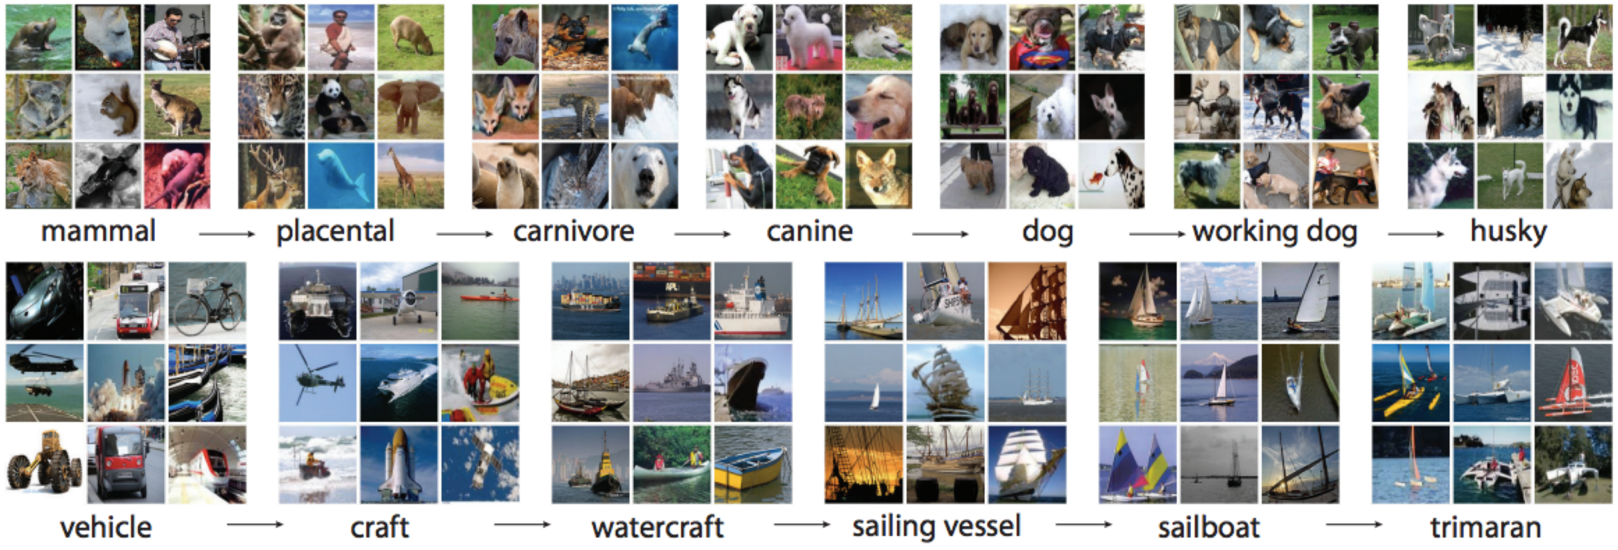

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet

1. Category Selection

~10% of concepts (synonym sets) in WordNet taxonomy

2. Image Collection

flickr

(& more search engines)

Query expansion:
- Augment queries
- Translate queries to different languages

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet

1. Category Selection

~10% of concepts (synonym sets) in WordNet taxonomy

2. Image Collection

flickr

(& more search engines)

Query expansion:
- Augment queries
- Translate queries to different languages

3. Human Verification

- Users verify if image contains queried object

- Use majority vote decision from multiple humans to support high quality results

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet Task

Definition of the target synonym set with link to Wikipedia.

3. Human Verification

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet Workers

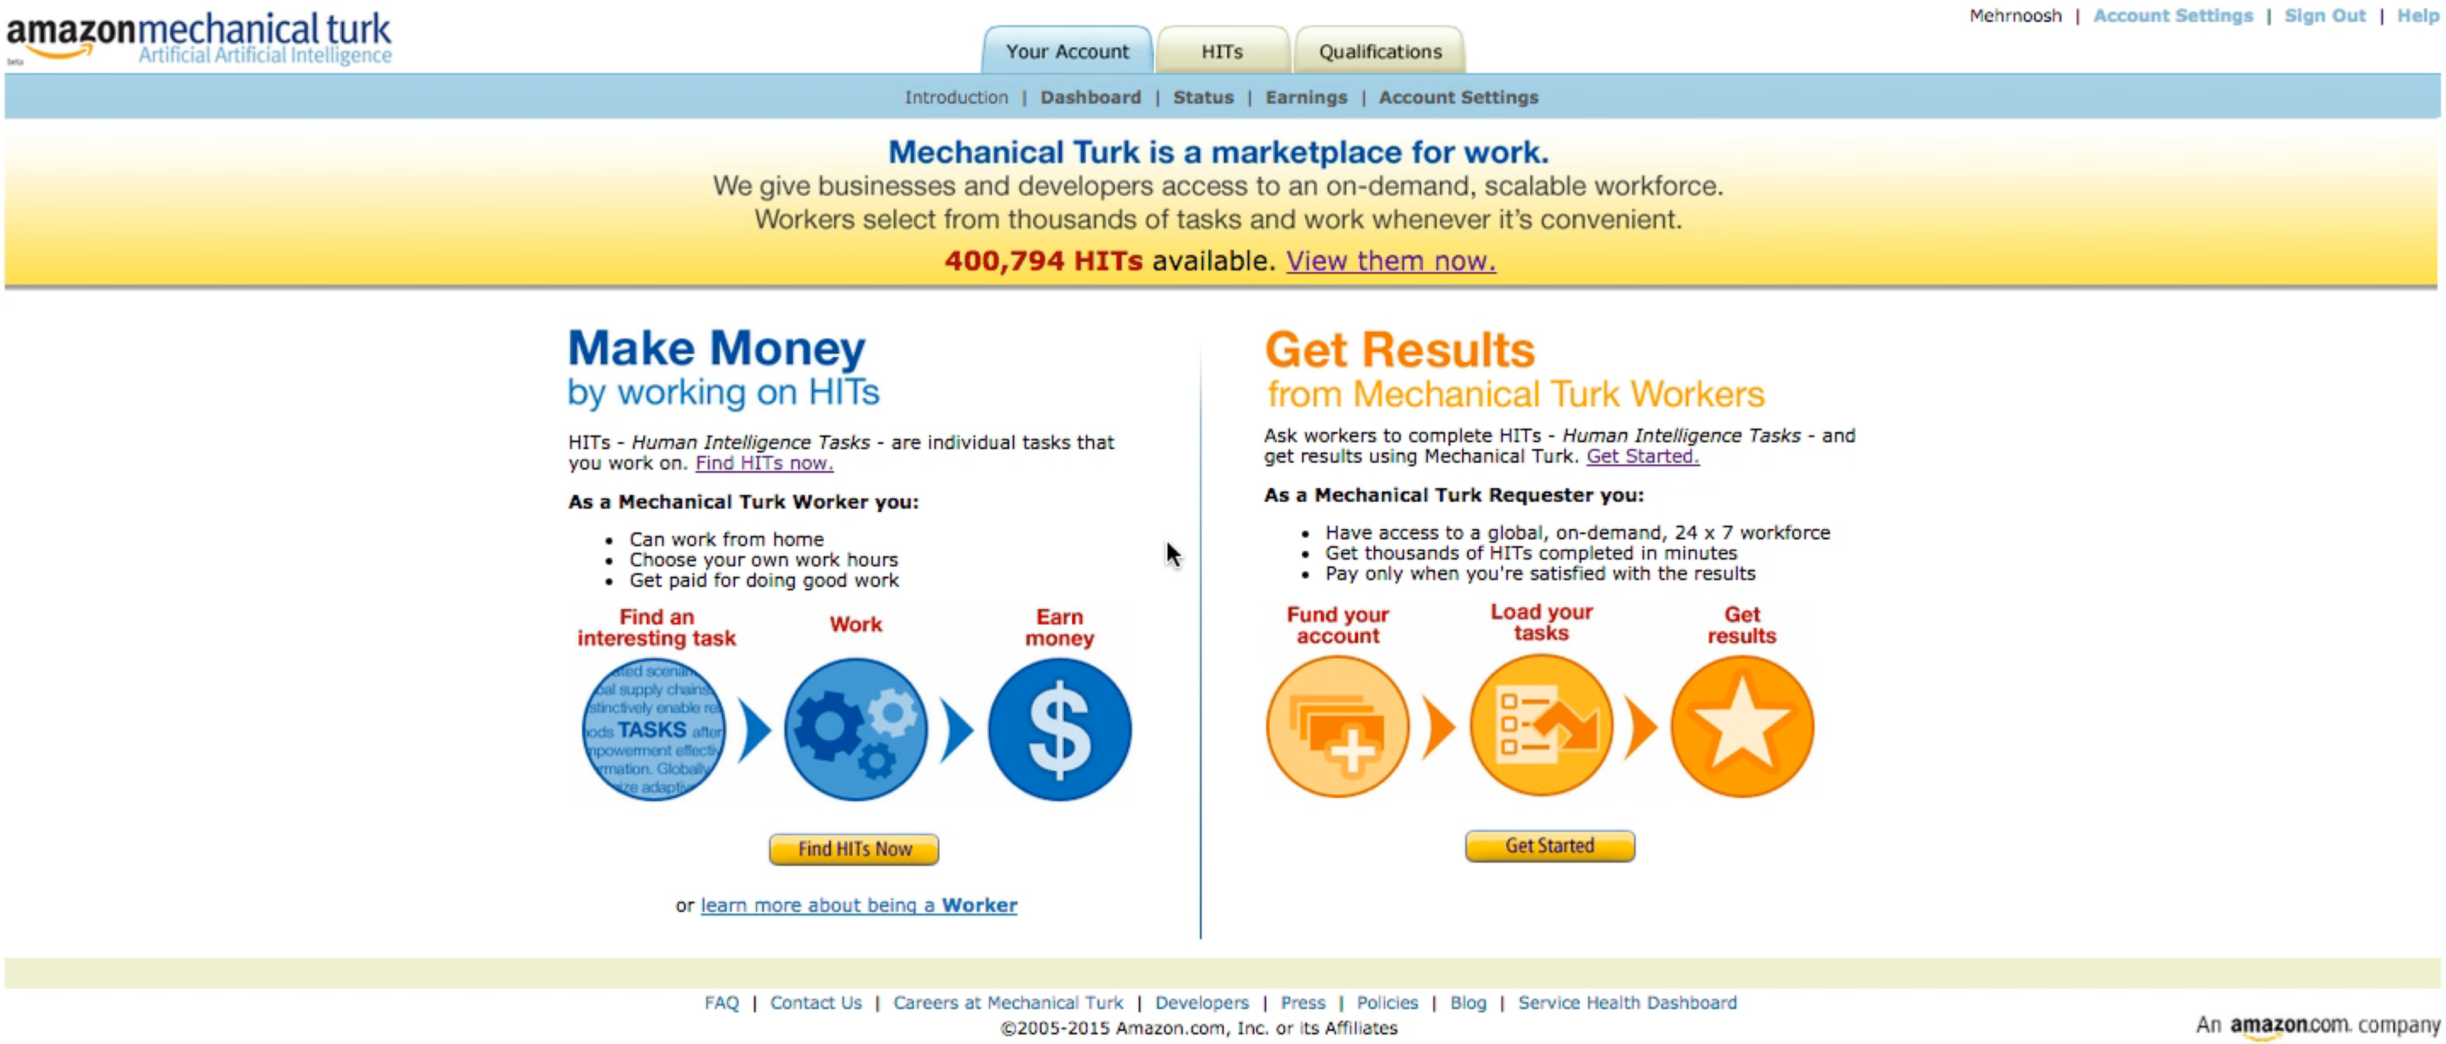

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet

Cost of naïve labeling approach?

1,500,000 images x 1,000 objects per image x 5 people per image x $0.01 per person per image x 1.2 (Amazon overhead fee) = $90,000,000

Inefficient (e.g., slow, expensive)!

3. Human Verification

- Users verify if image contains queried object

- Use majority vote decision from multiple humans to support high quality results

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet

Strategy 1: dynamically determine # of agreements needed per category

User 1	Y	Y	Y
User 2	N	Y	Y
User 3	N	Y	Y
User 4	Y	N	Y
User 5	Y	Y	Y
User 6	N	N	Y

#Y	# N	Conf Cat	Conf BCat
0	1	0.07	0.23
1	0	0.85	0.69
1	1	0.46	0.49
2	0	0.97	0.83
0	2	0.02	0.12
3	0	0.99	0.90
2	1	0.85	0.68

3. Human Verification

- Users verify if image contains queried object

- Use majority vote decision from multiple humans to support high quality results

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet

Strategy 2: embrace correlation, hierarchy, & sparsity to reduce human involvement

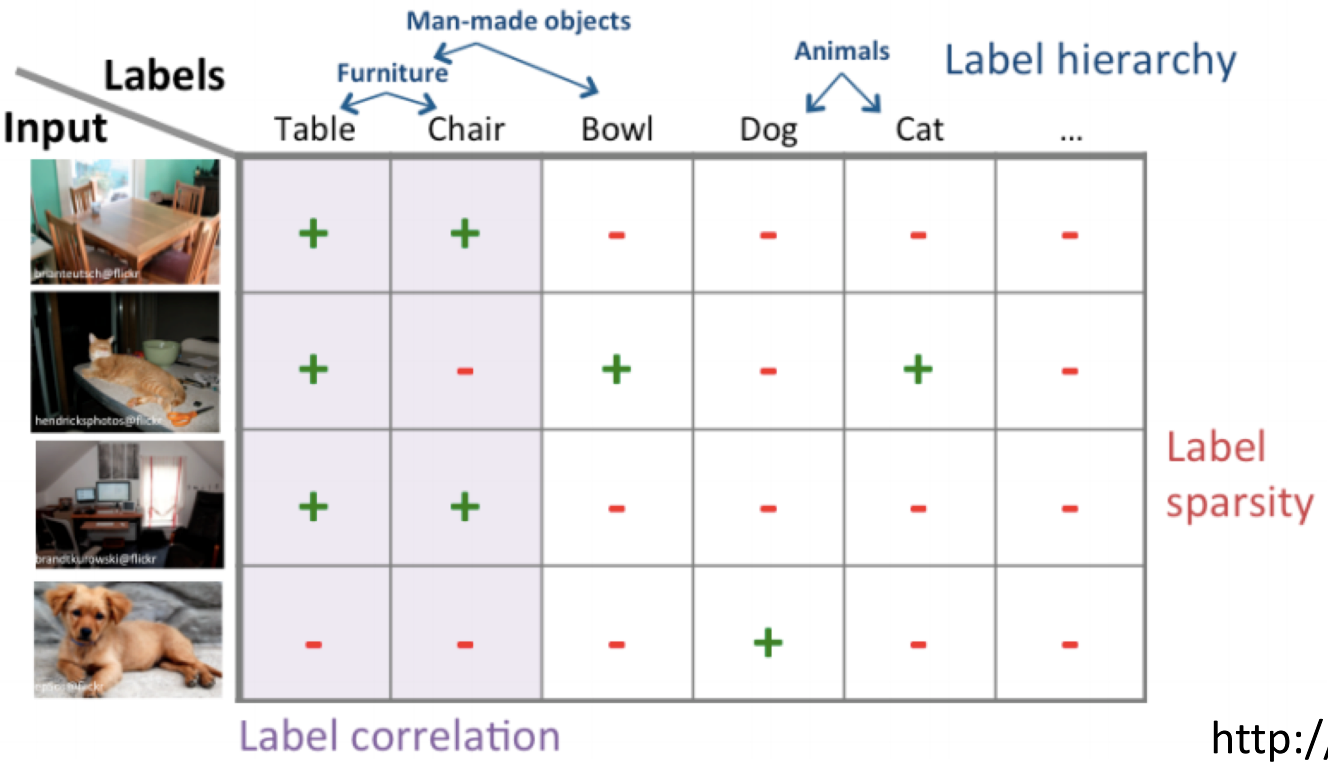

3. Human Verification

- Users verify if image contains queried object
- Use majority vote decision from multiple humans to support high quality results

http://ai.stanford.edu/~jkrause/papers/chi14_pres.pdf

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets: ImageNet

Strategy 2: embrace correlation, hierarchy, & sparsity to reduce human involvement

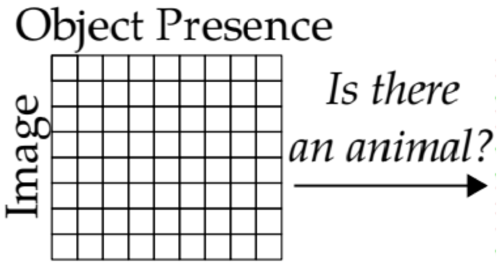

e.g., applying algorithm for one label (cat) on a set of images

3. Human Verification

- Users verify if image contains queried object

- Use majority vote decision from multiple humans to support high quality results

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

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, & Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. CVPR 2009.

Object Recognition Datasets

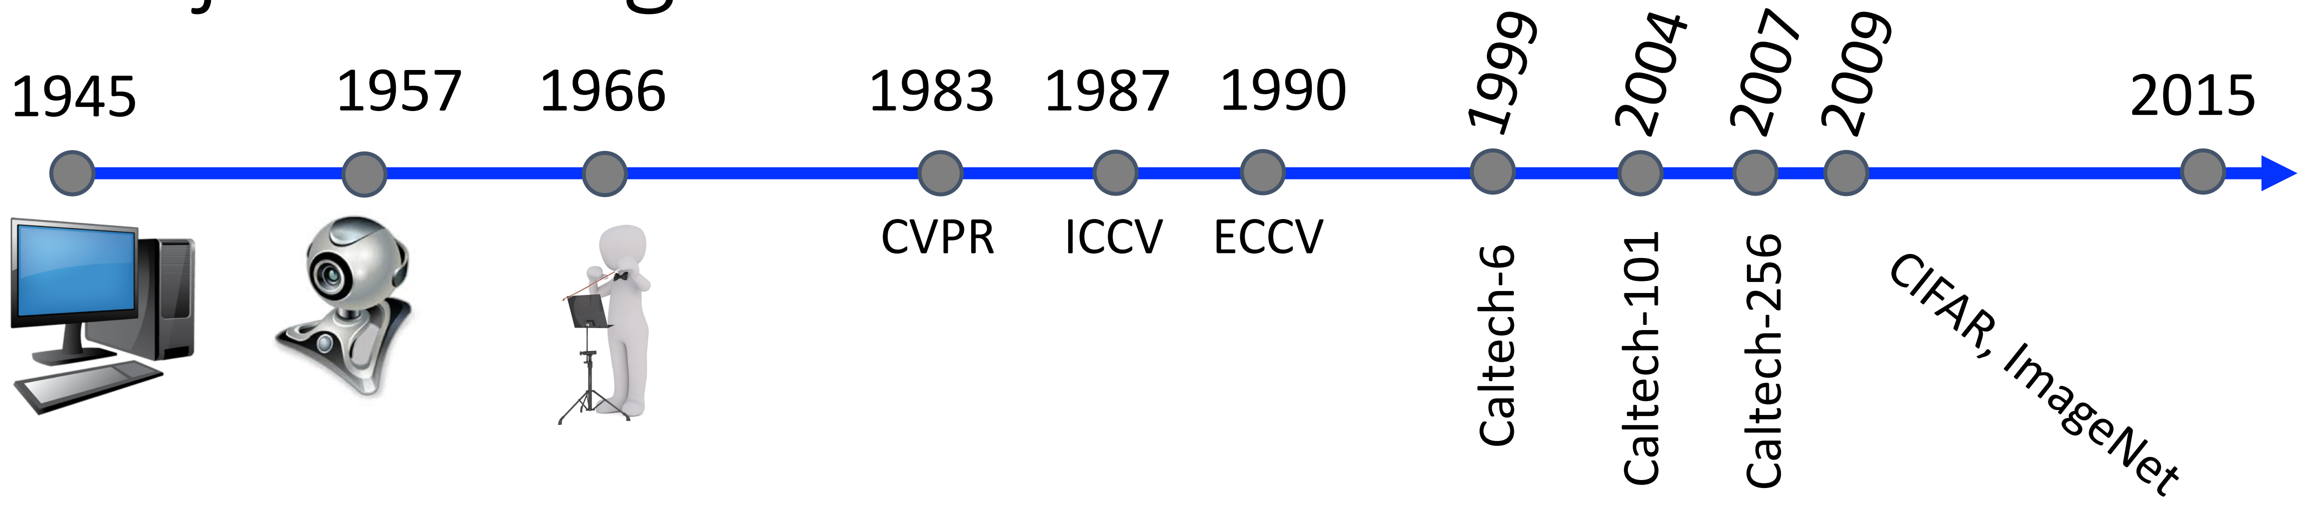

# Categories:	6	101	256	100	1000
Minimum # Samples/Category:	126+	31+	80+	600	668
# Images:	3,738	9,144	30,607	114,000	1,461,406

Today’s Topics

- Object recognition applications
- Object recognition datasets: key steps in creating them
- Object recognition datasets: scaling up their size with *crowdsourcing*
- Scaling up community working on object recognition with *workshop challenges*
- Class Discussion
- Lab: cascading stylesheets and web page layout

Engaging Larger Community: Typical Approach

1. Identify an AI problem

2. **Create infrastructure to work on the problem**: a big labelled dataset with a quantitative approach to evaluate algorithms

3. **Scale**: encourage community involvement in developing algorithms by publicly sharing the data with evaluation server and hosting a workshop to announce winners

Engaging Larger Community: ImageNet

1. Problem: Object Recognition

2. **Infrastructure:** ImageNet with evaluation metrics released

3. **Scale**: workshop challenge created

Engaging Larger Community: ImageNet Challenge

Engaging Larger Community: ImageNet Challenge

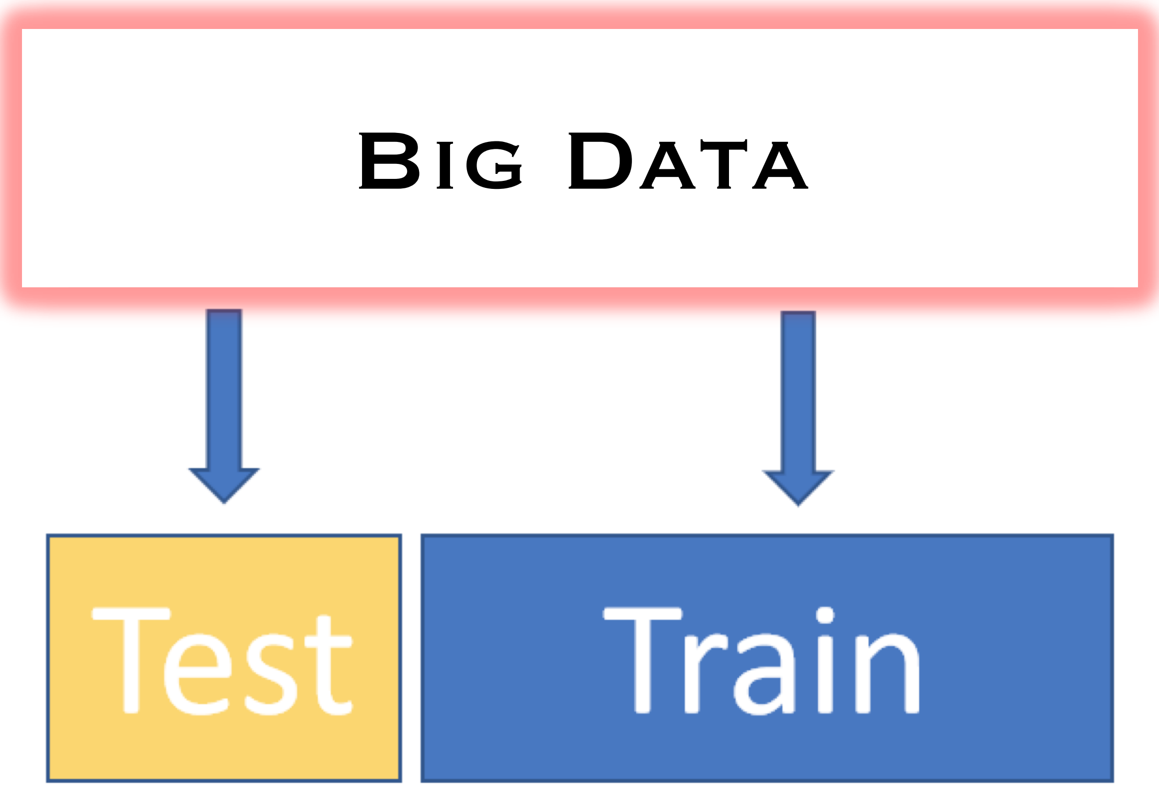

(Human labels hidden) (Human labels public)

Winner: highest scoring method on the hidden test set

Engaging Larger Community: ImageNet Challenge

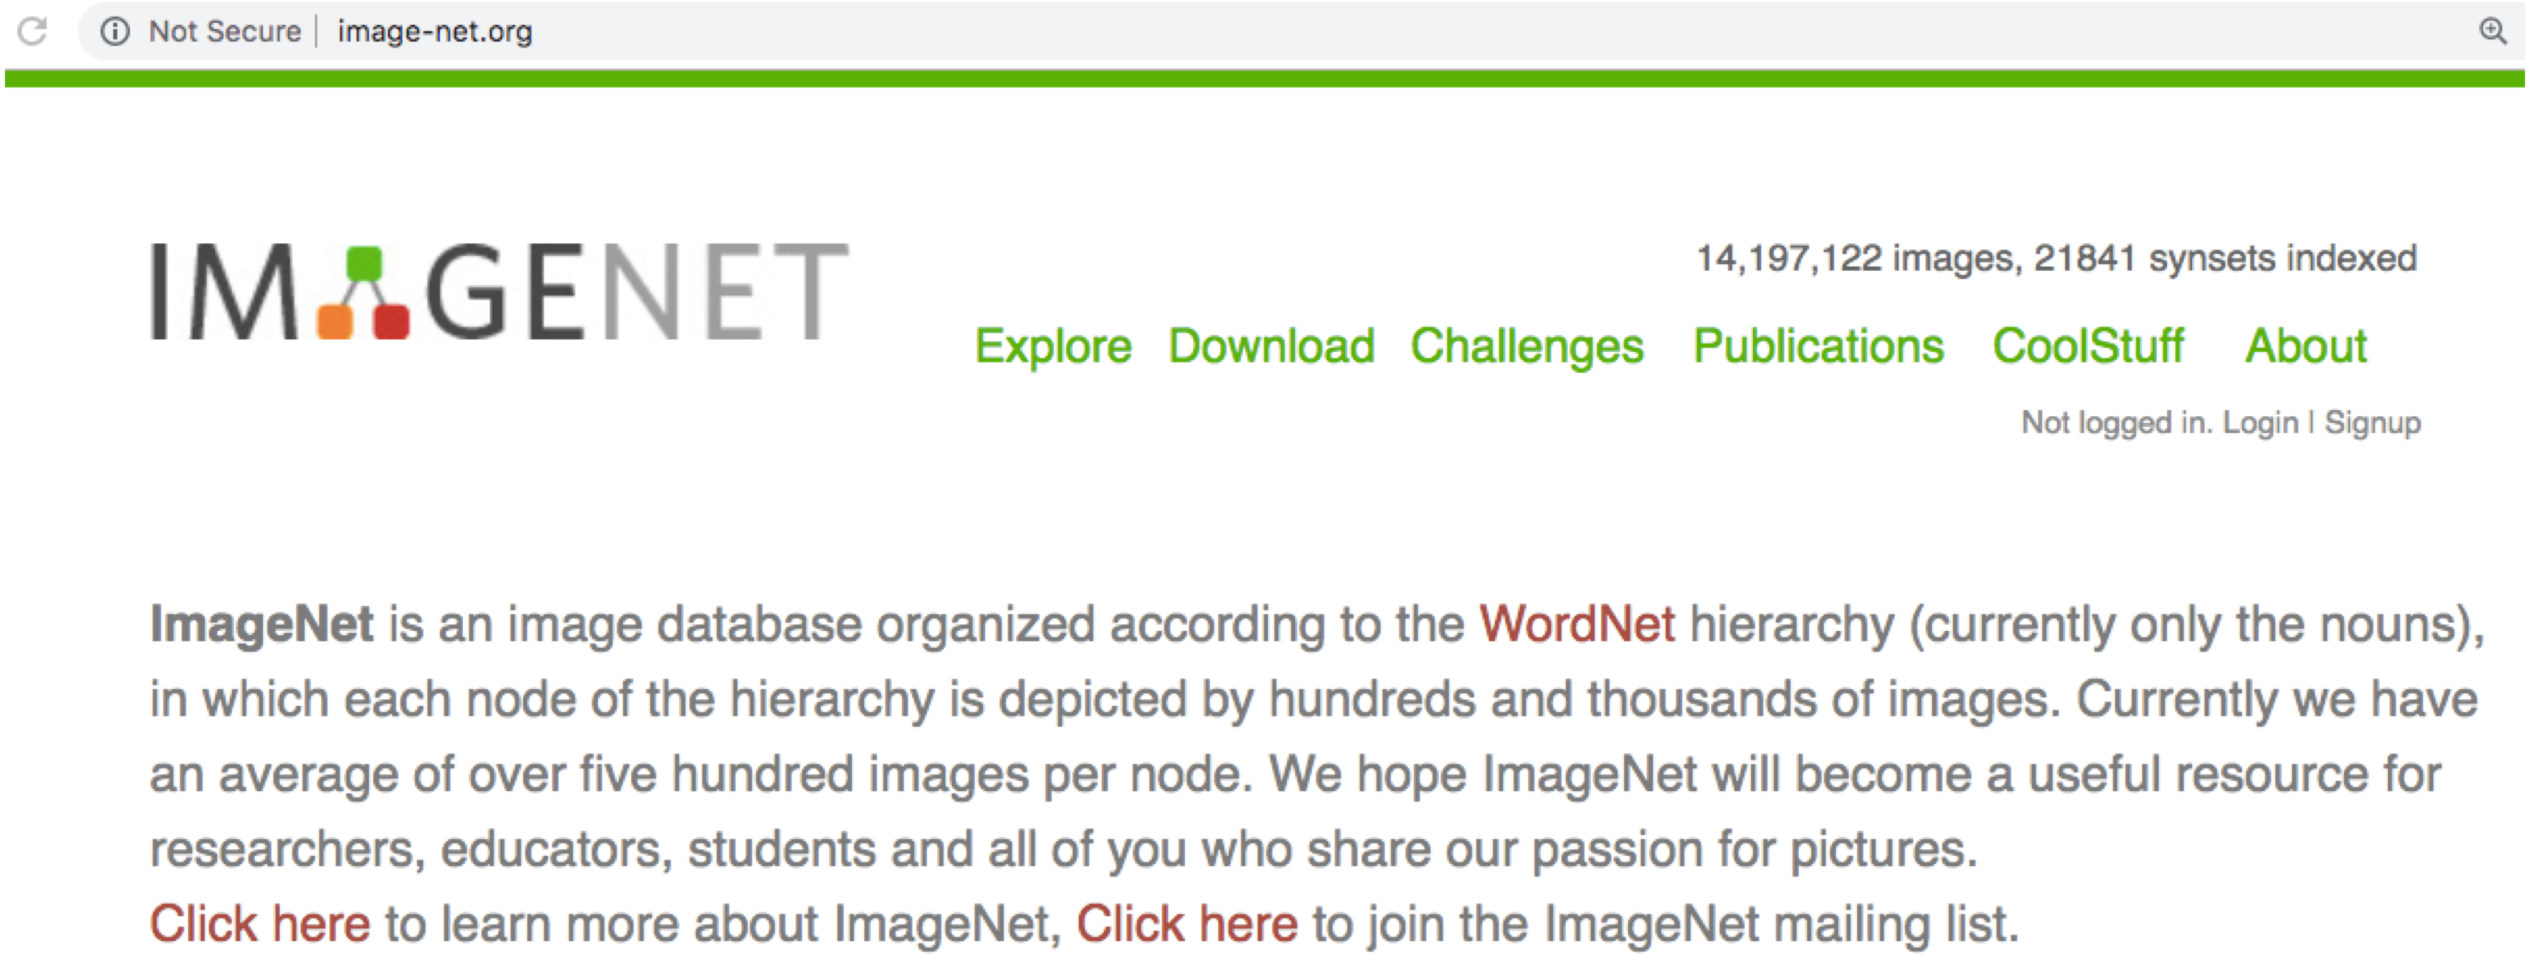

Demo: http://image-net.org/challenges/LSVRC/2010/index

Engaging Larger Community: ImageNet Challenge

Charting progress of algorithms

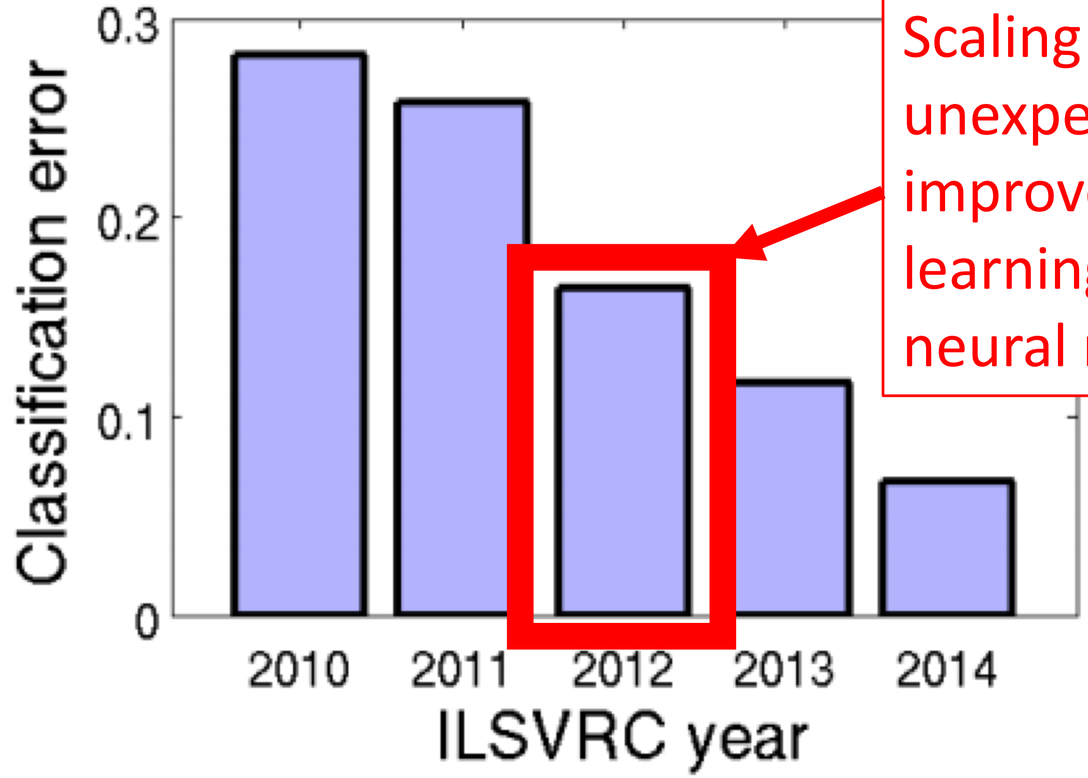

Scaling up the dataset led to unexpected, unprecedented improvements from a “deep learning” algorithm (i.e., deep neural network) called AlexNet

(15.3% error, 10.8 percentage points less than that of the runner up)

Olga Russakovsky et al. ImageNet Large Scale Visual Recognition Challenge. IJCV 2015.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet Classification with Deep Neural Networks. NIPS 2012.

Engaging Larger Community: ImageNet Challenge

"Suddenly people started to pay attention, not just within the AI community but across the technology industry as a whole."

- Economist

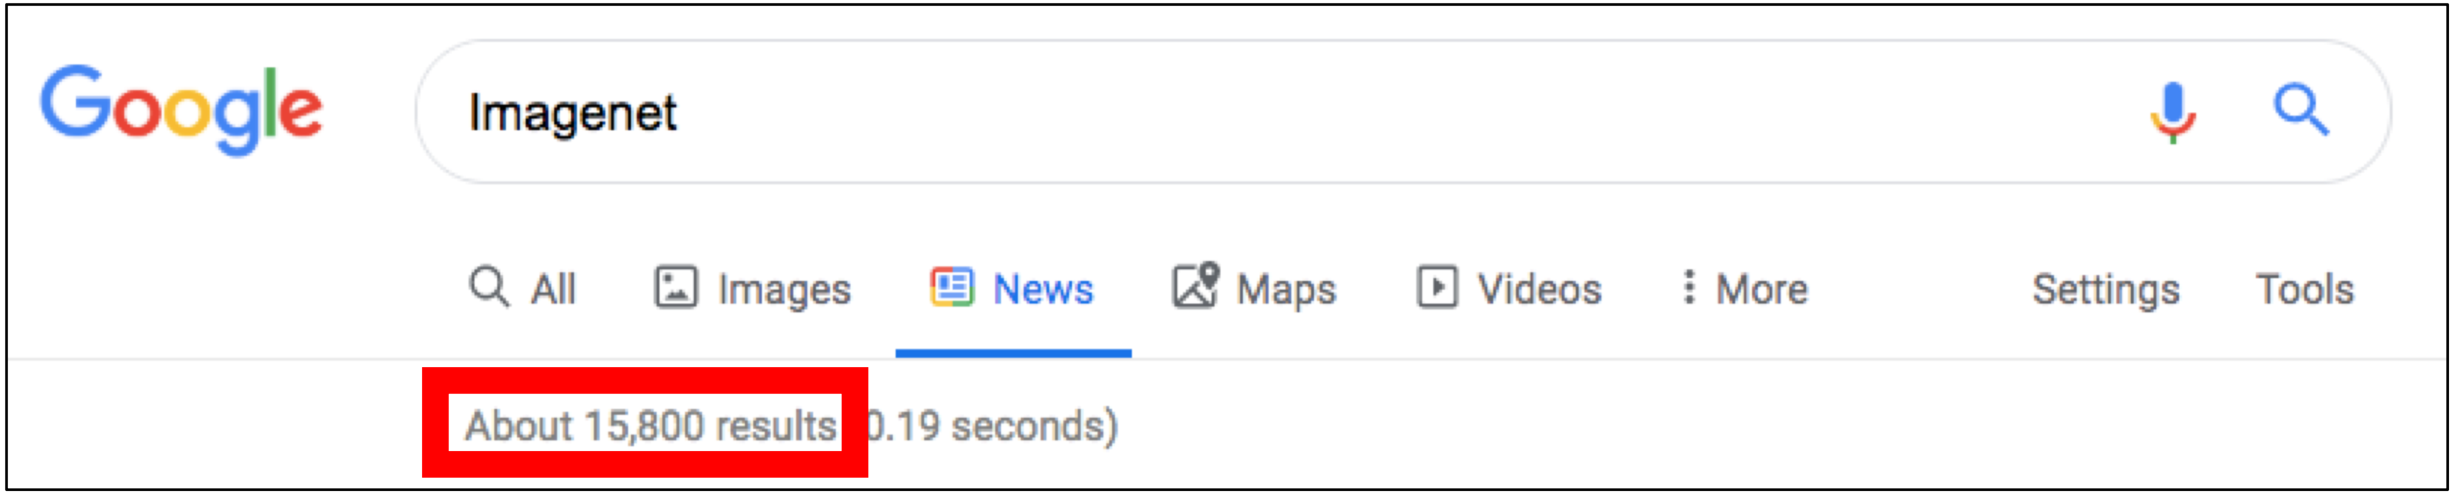

From not working to neural networking". *The Economist*. 25 June 2016. Retrieved 3 February 2018.

Engaging Larger Community

1945 — 1957 — 1966 — 1983 CVPR — 1987 ICCV — 1990 ECCV — 1999 Caltech-6 — 2004 Caltech-101 — 2007 Caltech-256 — 2009 ImageNet, CIFAR — 2015

Recall: rapid growth in the number of applications that use computer vision

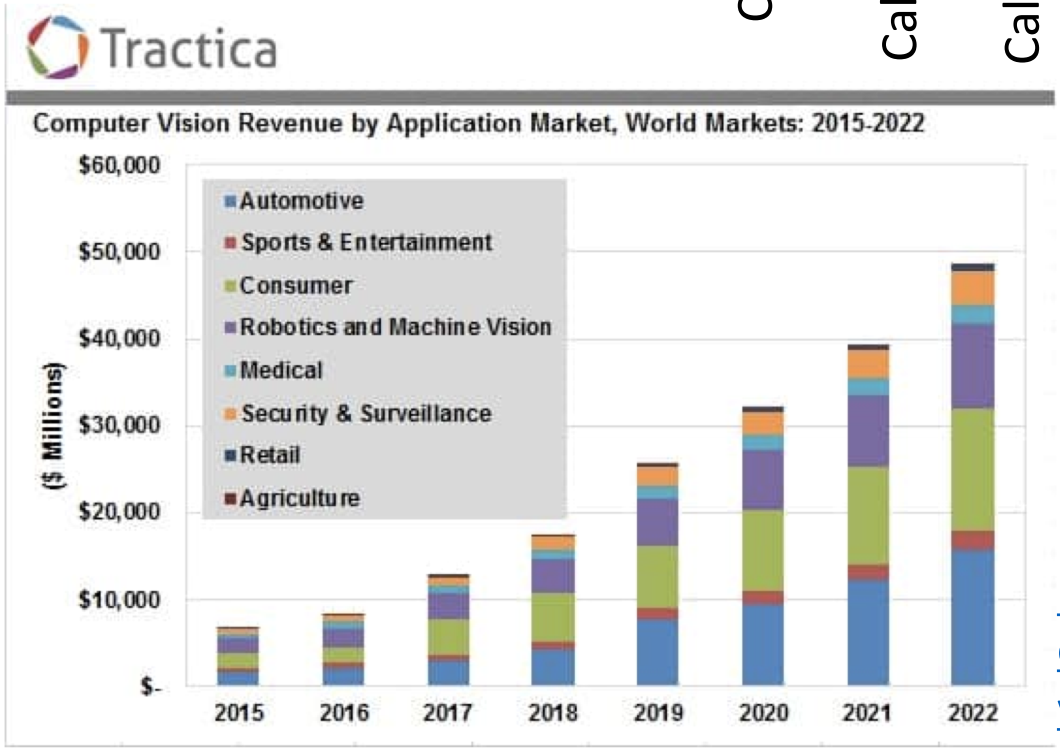

https://sevenshinestudios.wordpress.com/computer-vision-and-deep-learning/

2019 CVPR Community Recognition

1945 — 1957 — 1966 — 1983 CVPR — 1987 ICCV — 1990 ECCV — 1999 Caltech-6 — 2004 Caltech-101 — 2007 Caltech-256 — 2009 ImageNet — 2015

PAMI Longuet-Higgins Prize

Retrospective Most Impactful Paper from CVPR 2009

ImageNet: A large-scale hierarchical image database

Jia Deng, Wei Dong, Richard Socher,
Li-Jia Li, Kai Li, and Li Fei-Fei

https://syncedreview.com/2019/06/18/cvpr-2019-attracts-9k-attendees-best-papers-announced-imagenet-honoured-10-years-later/

2019 CVPR Community Recognition

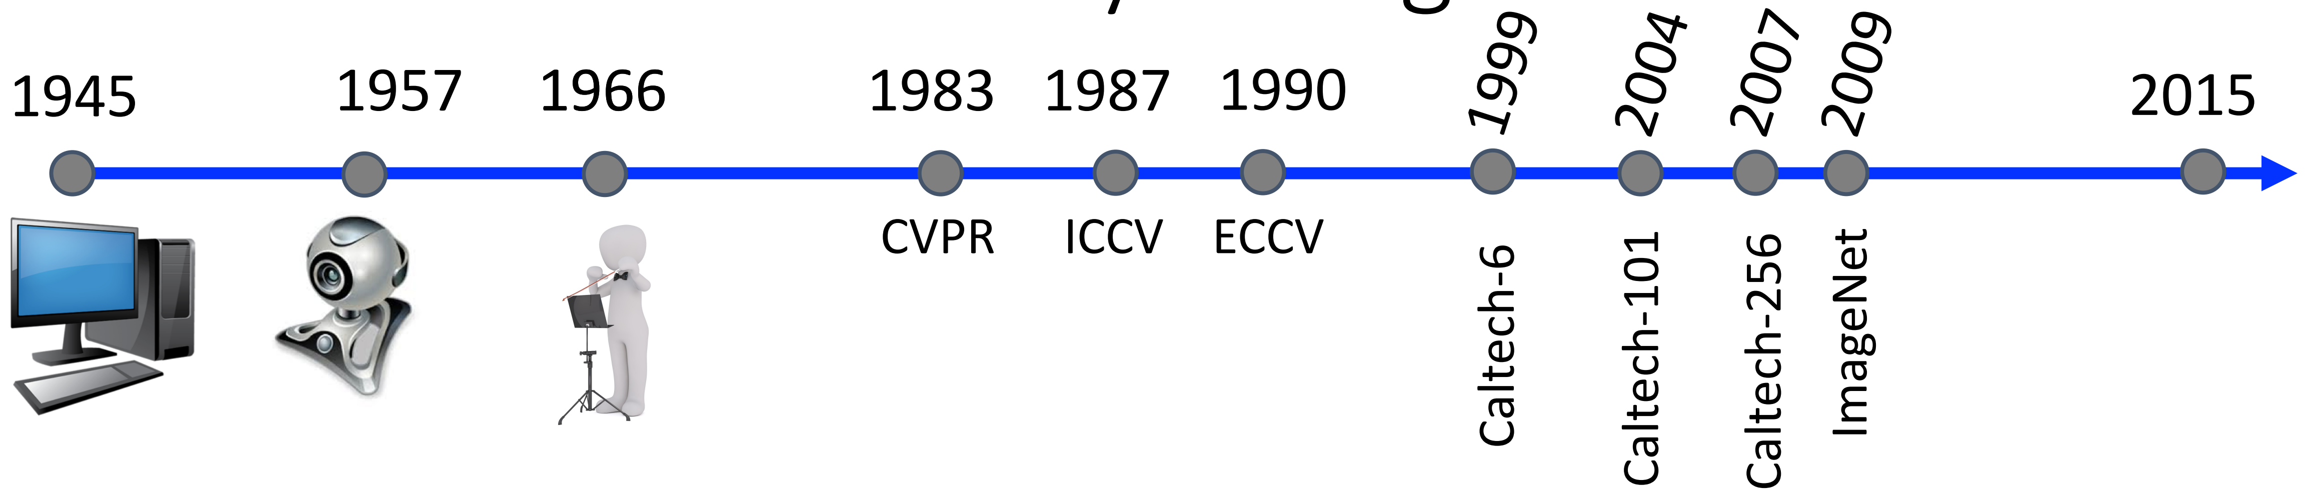

PAMI Longuet-Higgins Prize

Retrospective Most Impactful Paper from CVPR 2009

ImageNet: A large-scale hierarchical image database

Jia Deng, Wei Dong, Richard Socher,
Li-Jia Li, Kai Li, and Li Fei-Fei

"In 2009, ImageNet was not the most mainstream work, but all of us who did this project believed that it would have a big impact, so we put in a lot of efforts. One of the revelations it gives me is that you don't have to do the most popular things, but do what you believe will have an impact."

-First author, Jia Deng

https://syncedreview.com/2019/06/18/cvpr-2019-attracts-9k-attendees-best-papers-announced-imagenet-honoured-10-years-later/

ImageNet: Great Start…

Now what about being more inclusive? (e.g., households around world with low household incomes?)

Ground truth: Soap **Nepal, 288 $/month**

Azure: food, cheese, bread, cake, sandwich
Clarifai: food, wood, cooking, delicious, healthy
Google: food, dish, cuisine, comfort food, spam
Amazon: food, confectionary, sweets, burger
Watson: food, food product, turmeric, seasoning
Tencent: food, dish, matter, fast food, nutriment

Ground truth: Soap **UK, 1890 $/month**

Azure: toilet, design, art, sink
Clarifai: people, faucet, healthcare, lavatory, wash closet
Google: product, liquid, water, fluid, bathroom accessory
Amazon: sink, indoors, bottle, sink faucet
Watson: gas tank, storage tank, toiletry, dispenser, soap dispenser
Tencent: lotion, toiletry, soap dispenser, dispenser, after shave

Terrance DeVries, Ishan Misra, Changhan Wang, Laurens van der Maaten l. Does Object Recognition Work for Everyone? 2019.

ImageNet: Great Start…

Now what about being more inclusive? (e.g., households around world with low household incomes?)

Ground truth: Spices **Phillipines, 262 $/month**

Azure: bottle, beer, counter, drink, open
Clarifai: container, food, bottle, drink, stock
Google: product, yellow, drink, bottle, plastic bottle
Amazon: beverage, beer, alcohol, drink, bottle
Watson: food, larder food supply, pantry, condiment, food seasoning
Tencent: condiment, sauce, flavorer, catsup, hot sauce

Ground truth: Spices **USA, 4559 $/month**

Azure: bottle, wall, counter, food
Clarifai: container, food, can, medicine, stock
Google: seasoning, seasoned salt, ingredient, spice, spice rack
Amazon: shelf, tin, pantry, furniture, aluminium
Watson: tin, food, pantry, paint, can
Tencent: spice rack, chili sauce, condiment, canned food, rack

Terrance DeVries, Ishan Misra, Changhan Wang, Laurens van der Maaten l. Does Object Recognition Work for Everyone? 2019.

ImageNet: Great Start…

Why do you think the algorithms made these mistakes?

Ground truth: Spices **Phillipines, 262 $/month**

Azure: bottle, beer, counter, drink, open
Clarifai: container, food, bottle, drink, stock
Google: product, yellow, drink, bottle, plastic bottle
Amazon: beverage, beer, alcohol, drink, bottle
Watson: food, larder food supply, pantry, condiment, food seasoning
Tencent: condiment, sauce, flavorer, catsup, hot sauce

Ground truth: Spices **USA, 4559 $/month**

Azure: bottle, wall, counter, food
Clarifai: container, food, can, medicine, stock
Google: seasoning, seasoned salt, ingredient, spice, spice rack
Amazon: shelf, tin, pantry, furniture, aluminium
Watson: tin, food, pantry, paint, can
Tencent: spice rack, chili sauce, condiment, canned food, rack

Terrance DeVries, Ishan Misra, Changhan Wang, Laurens van der Maaten l. Does Object Recognition Work for Everyone? 2019.

ImageNet: Great Start…

Geographical distribution of images in the ImageNet using Flickr metadata:

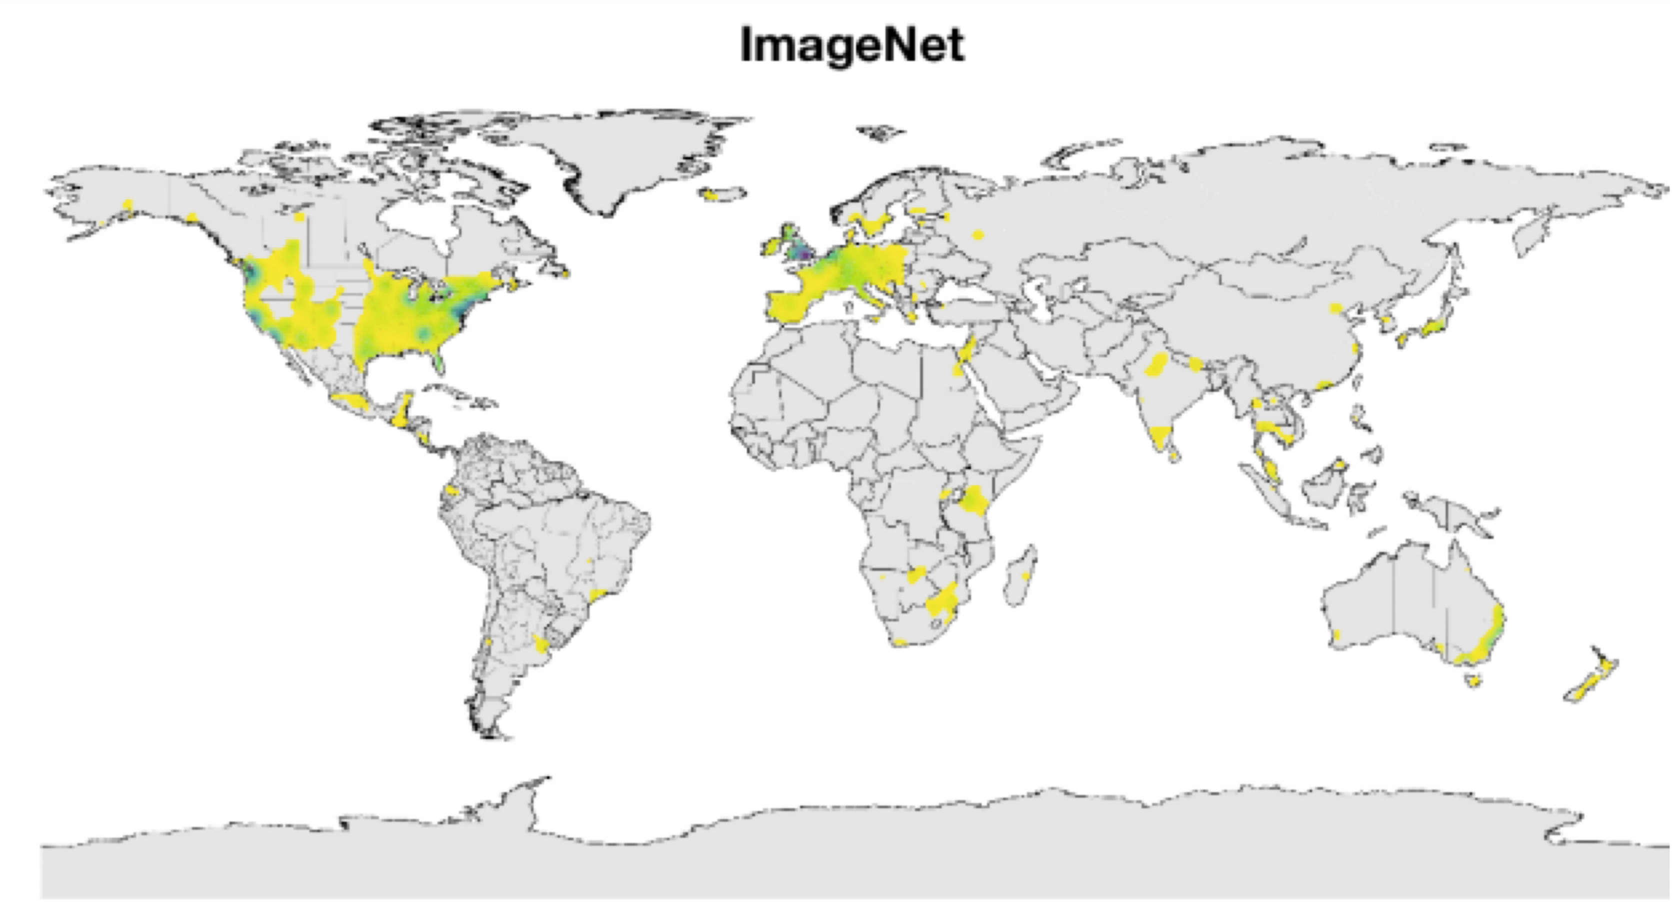

Jieyu Zhao et al. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. 2017.

ImageNet: Great Start…

Now what about not amplifying existing biases?

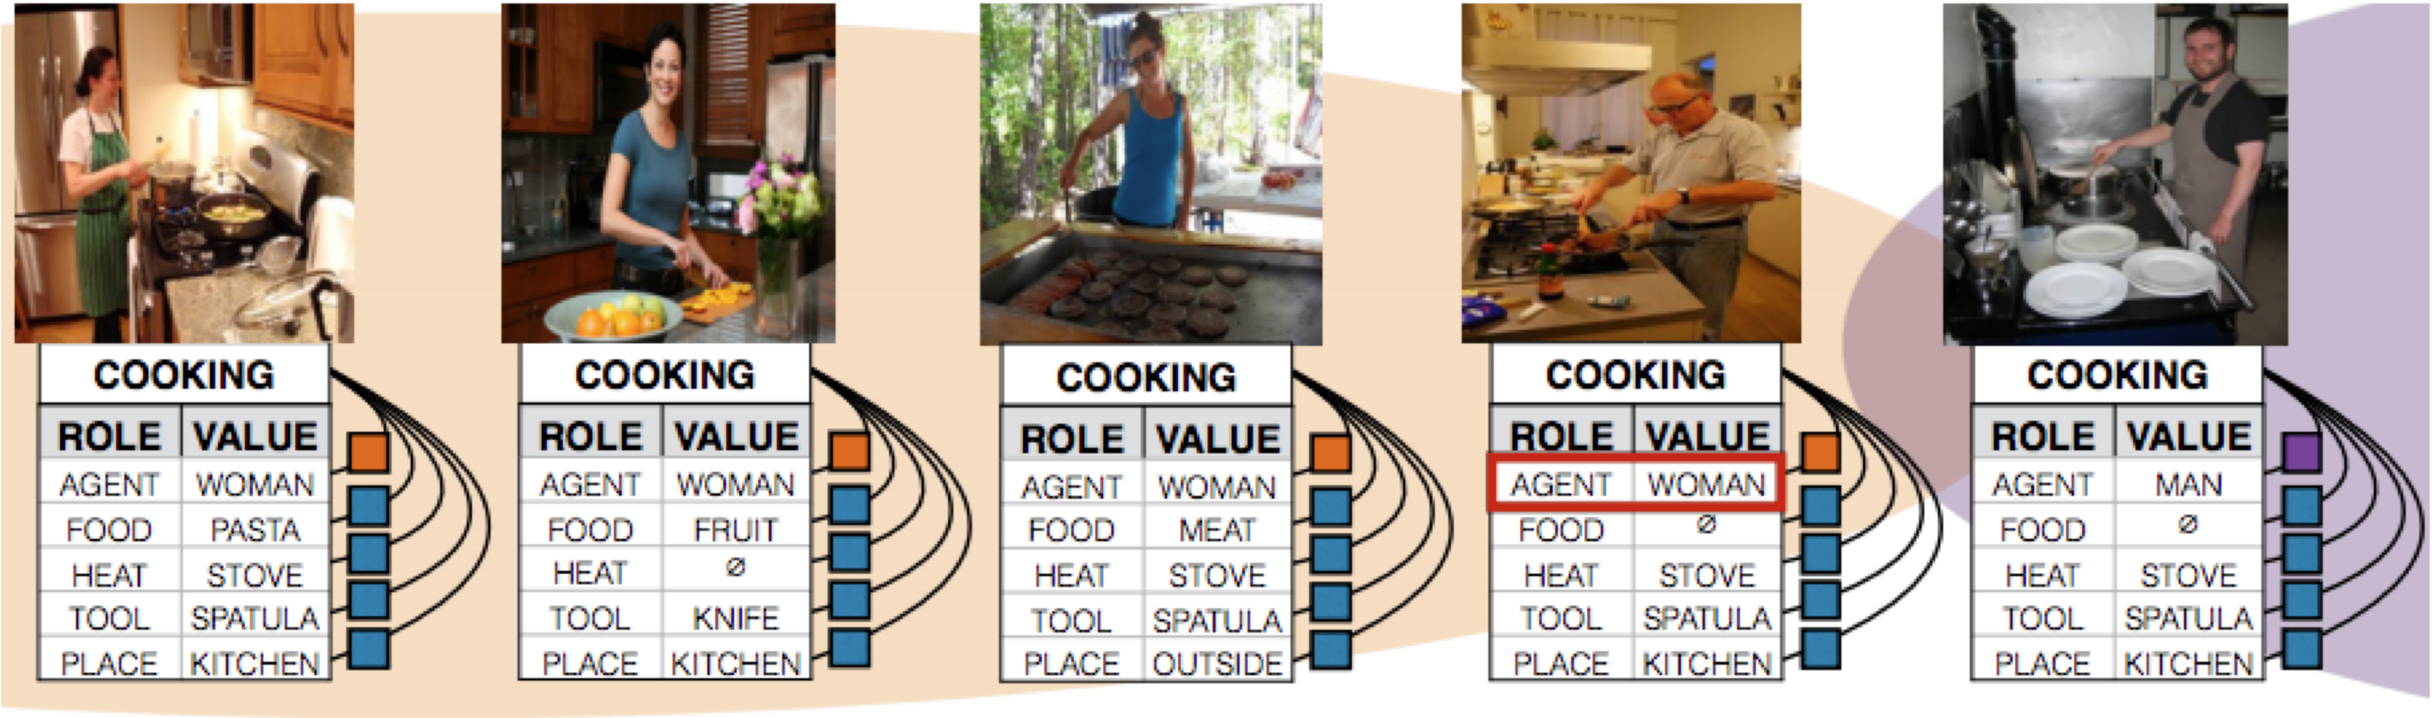

Algorithm identifies men in kitchens as women.

Jieyu Zhao et al. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. 2017.

ImageNet: Great Start…

Now what about not amplifying existing biases?

Algorithm identifies these pictures as men.

Yuchi Tian, Ziyuan Zhong, Vicente Ordonez, Baishakhi Ray. Testing Deep Neural Network based Image Classifiers. 2019.

ImageNet: Great Start…

Now what about not amplifying existing biases?

Why do you think the algorithm made these mistakes?

Algorithm identifies these pictures as men.

Yuchi Tian, Ziyuan Zhong, Vicente Ordonez, Baishakhi Ray. Testing Deep Neural Network based Image Classifiers. 2019.

Today’s Topics

- Object recognition applications
- Object recognition datasets: key steps in creating them
- Object recognition datasets: scaling up their size with *crowdsourcing*
- Scaling up community working on object recognition with *workshop challenges*
- Class Discussion
- Lab: cascading stylesheets and web page layout

Today’s Topics

- Object recognition applications
- Object recognition datasets: key steps in creating them
- Object recognition datasets: scaling up their size with *crowdsourcing*
- Scaling up community working on object recognition with *workshop challenges*
- Class Discussion
- Lab: cascading stylesheets and web page layout