1950

Future Vision

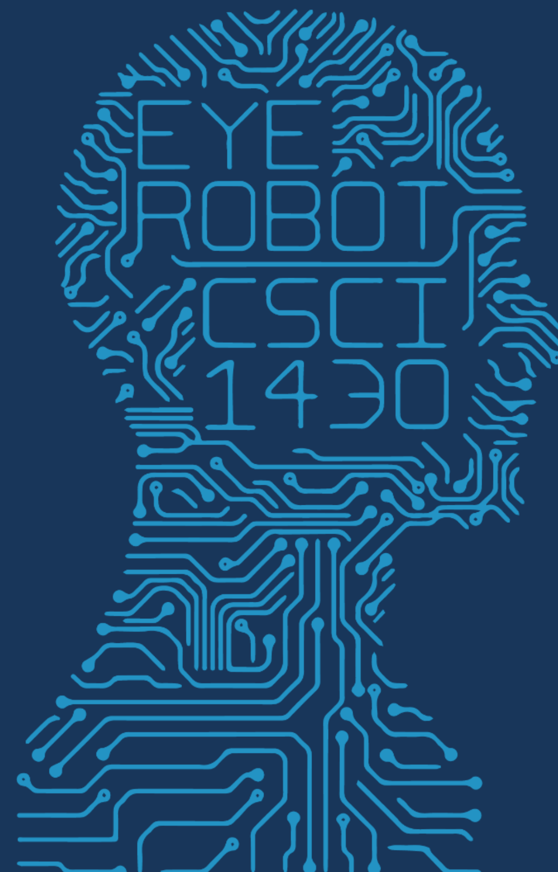

2017 MWF 1pm 368

Computer Vision

General Principal

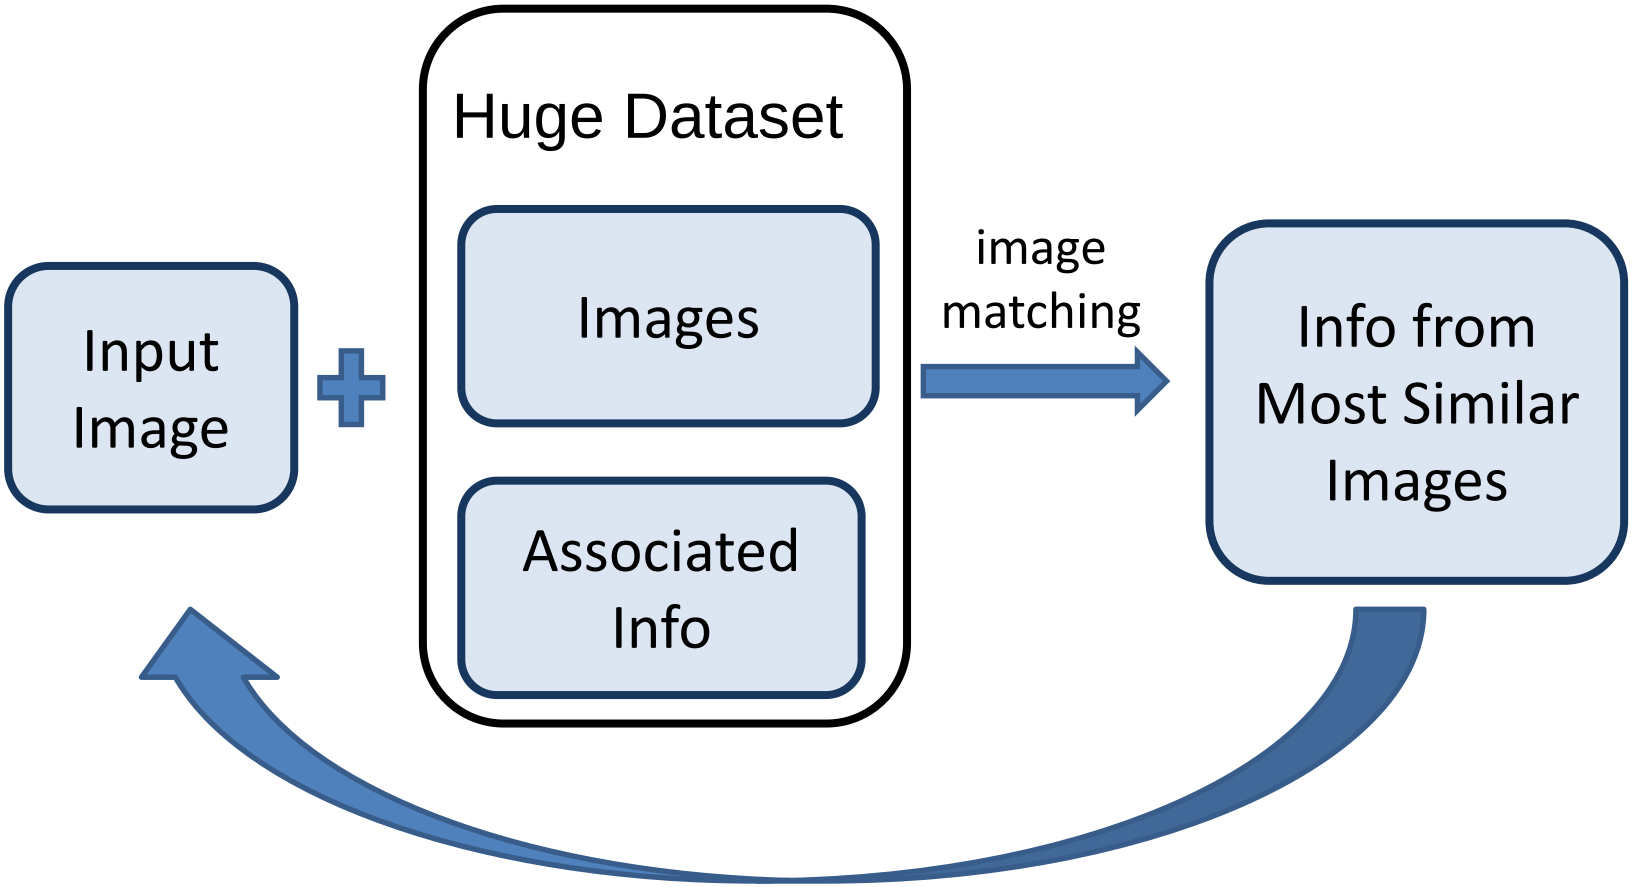

Hopefully, If you have enough images, the dataset will contain very similar images that you can find with simple matching methods.

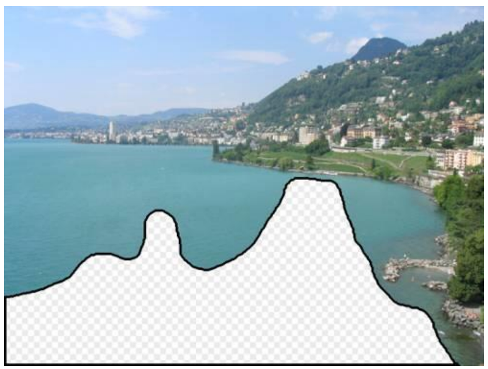

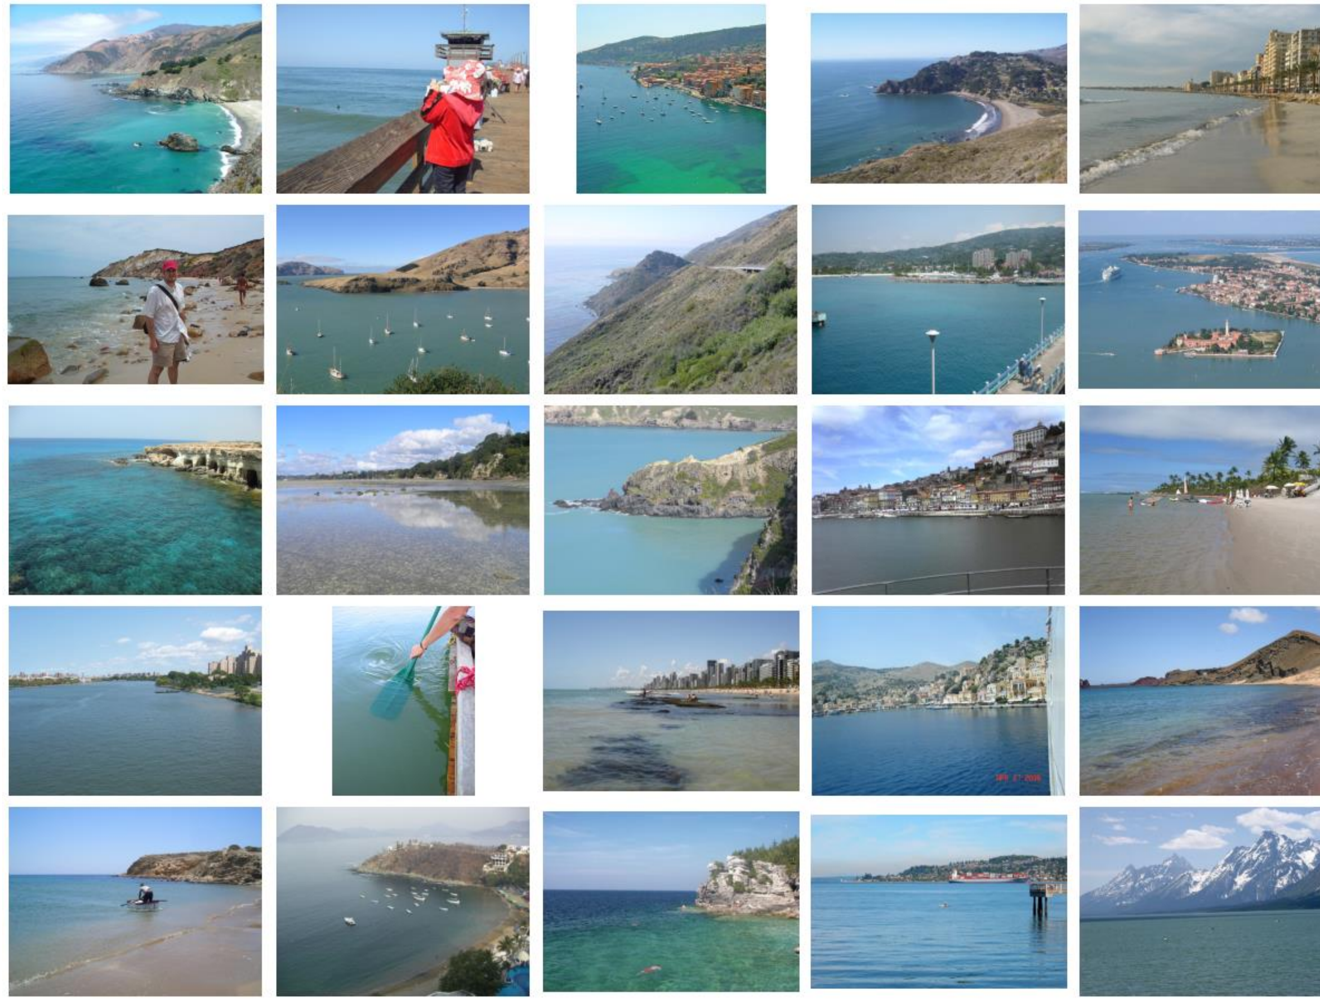

... 200 total

Graph cut + Poisson blending

How much can an image tell about its geographic location?

How much can an image tell about its geographic location?

6 million geo-tagged Flickr images

http://graphics.cs.cmu.edu/projects/im2gps/

im2gps (Hays & Efros, CVPR 2008)

Nearest Neighbors according to gist + bag of SIFT + color histogram + a few others

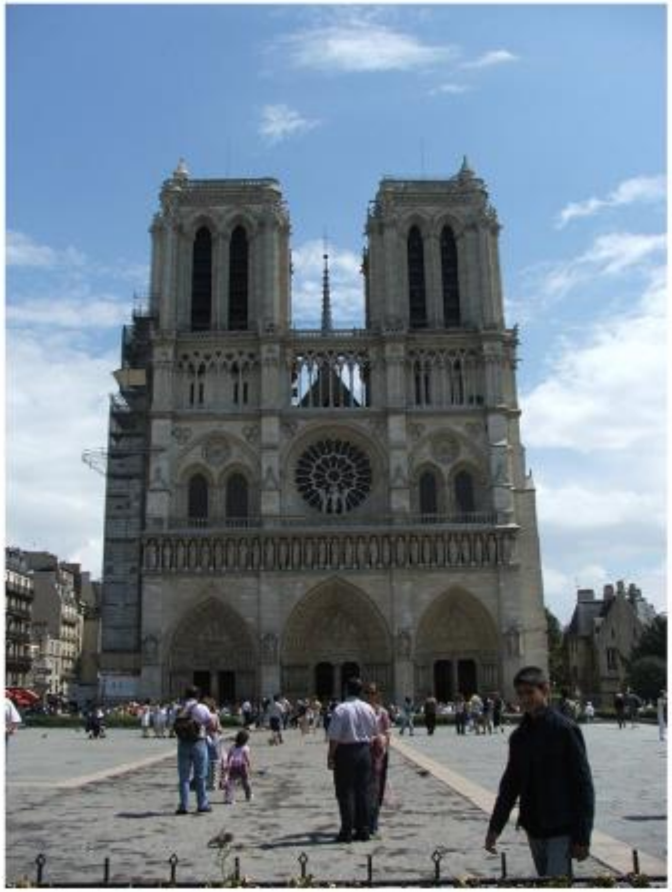

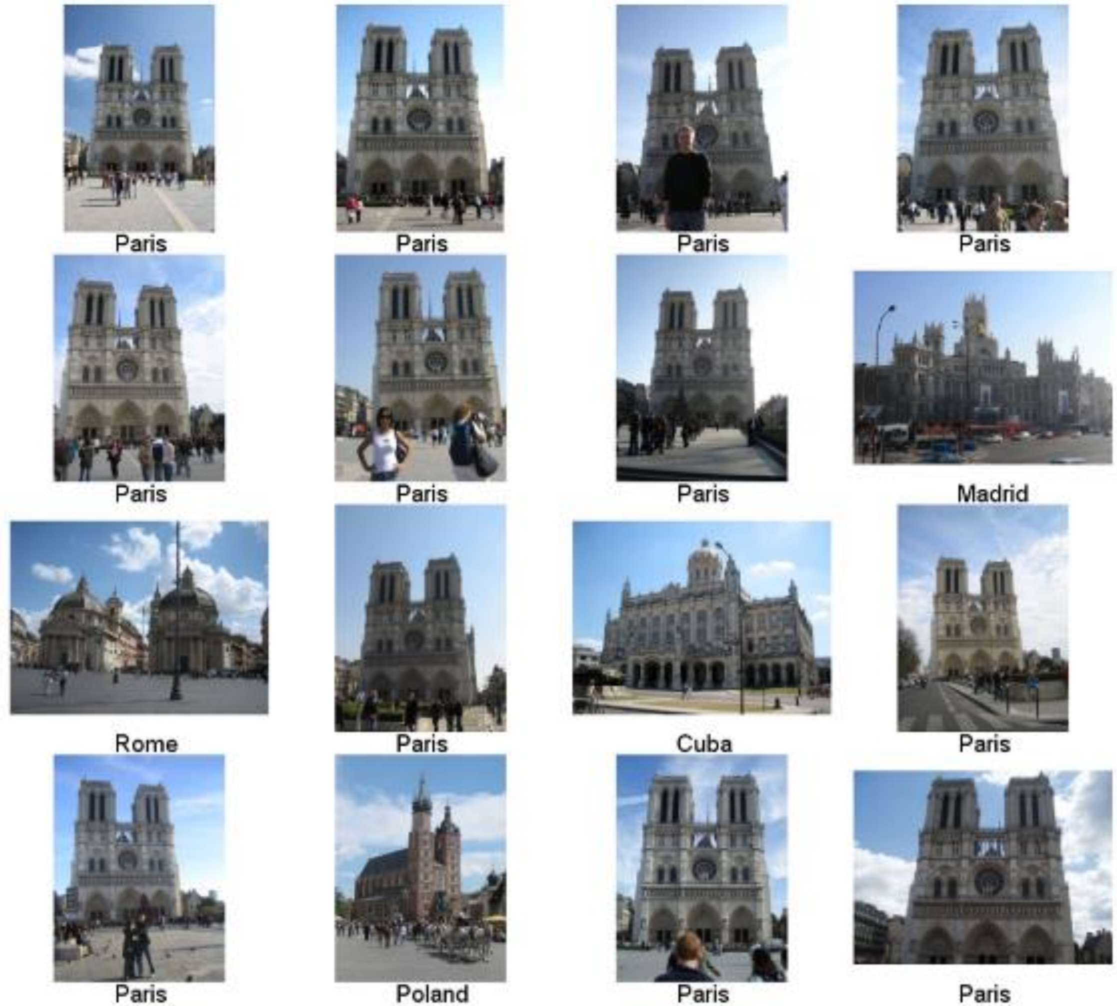


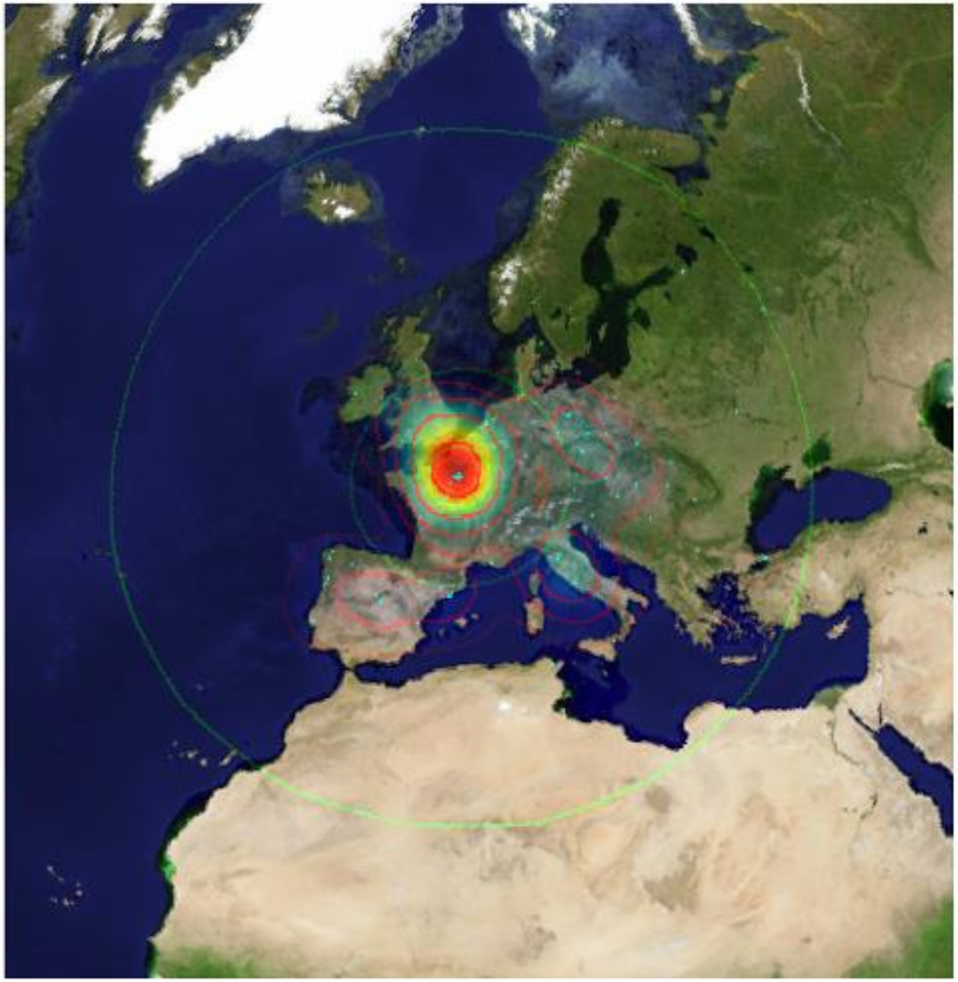

Im2gps

Example Scene Matches

Voting Scheme

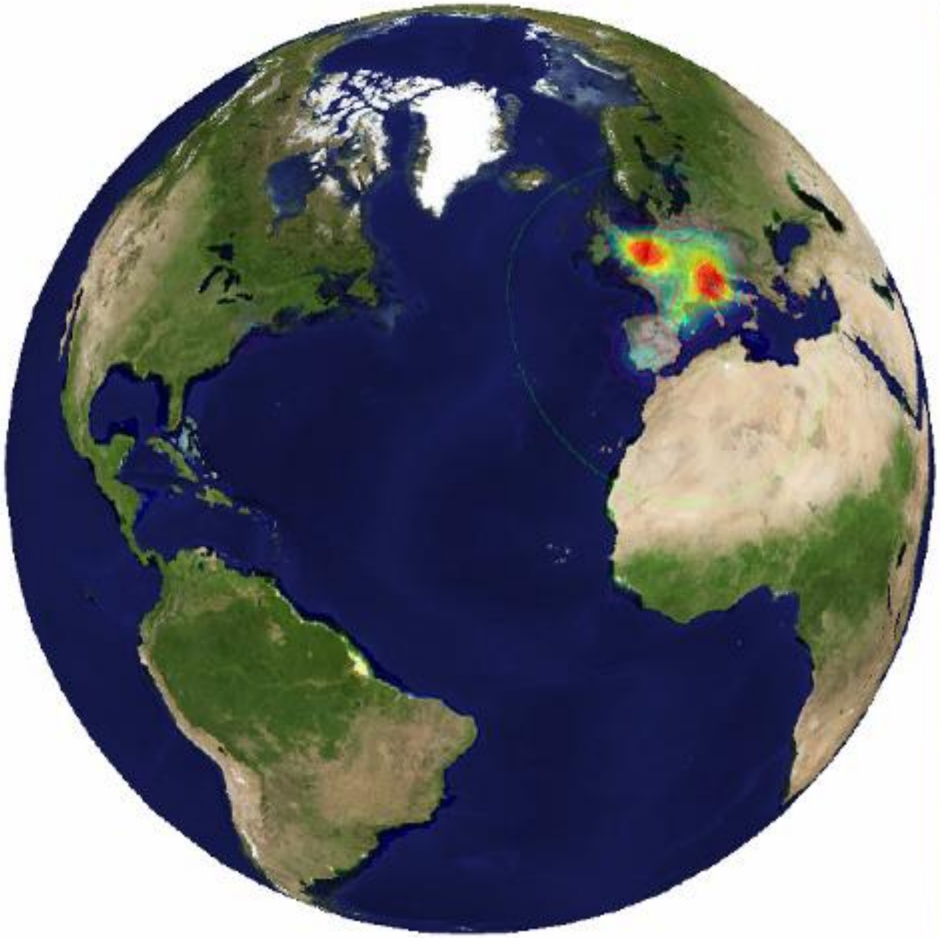

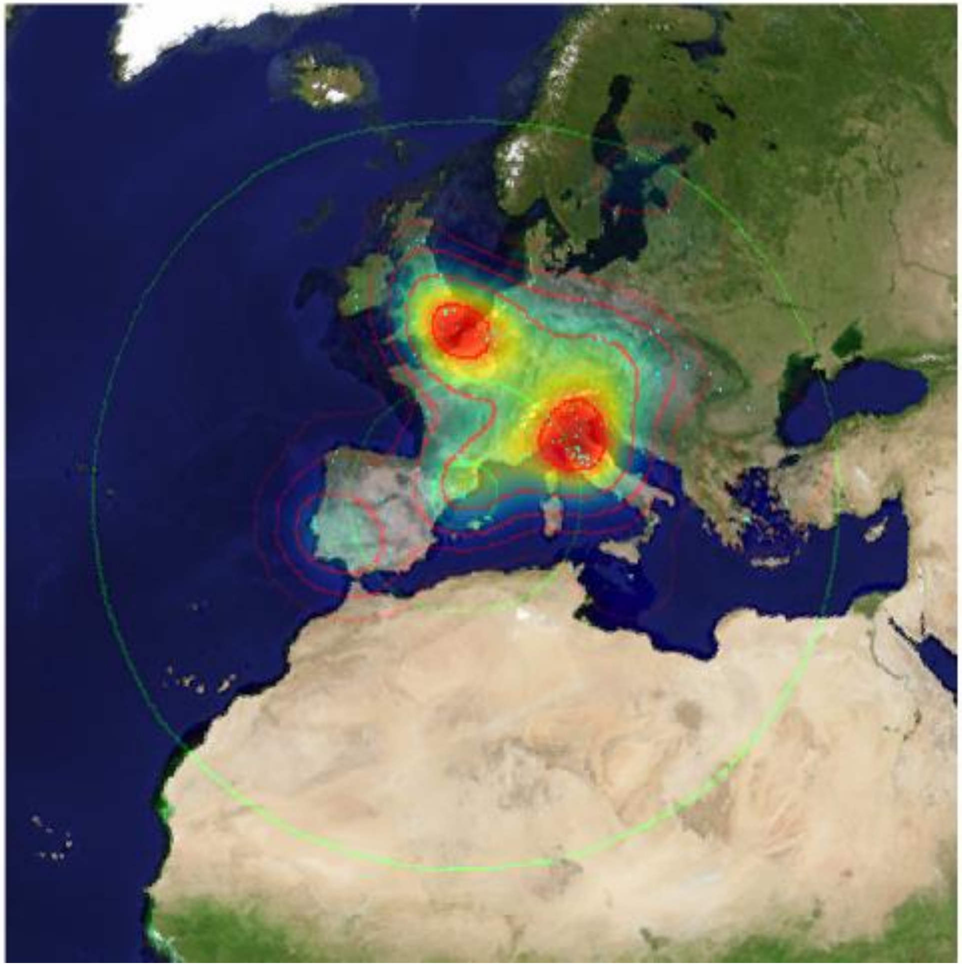

im2gps

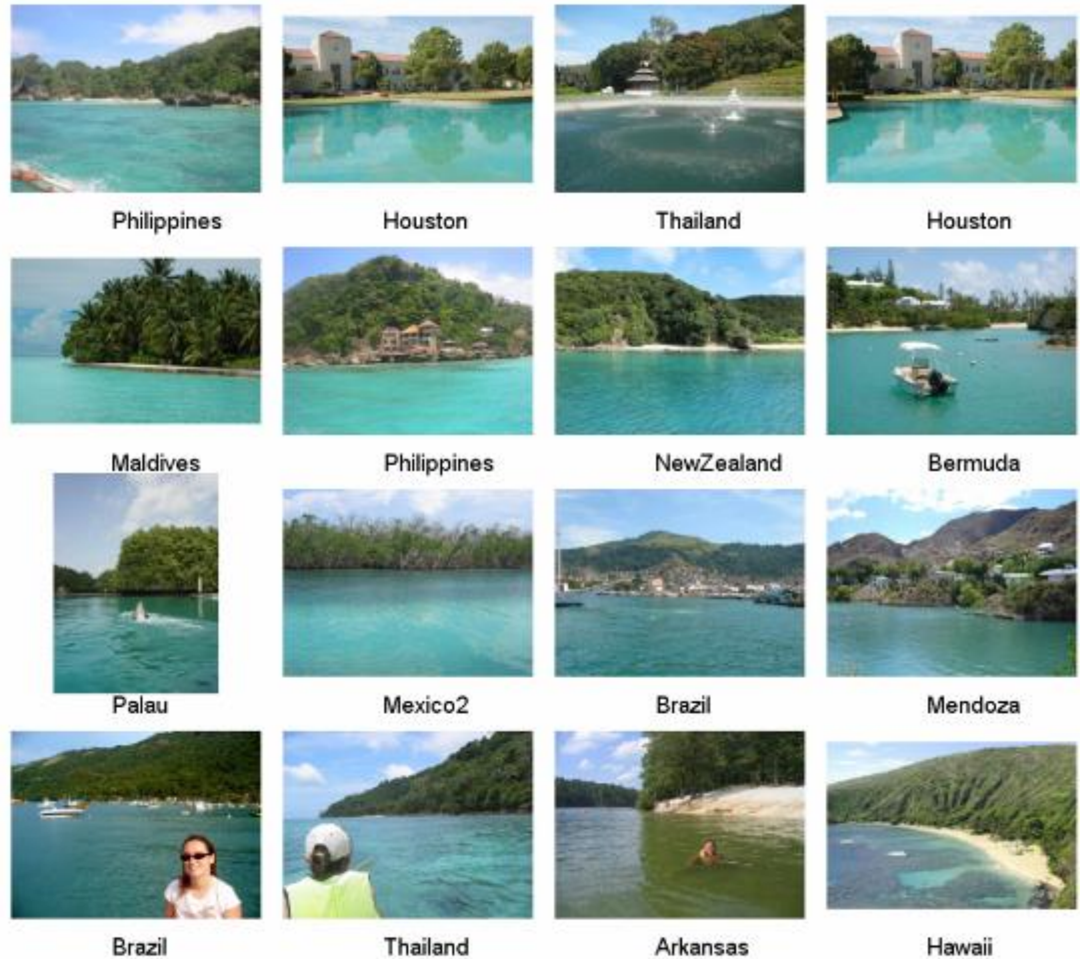
Philippines
Houston
Thailand
Houston
Maldives
Philippines
NewZealand
Bermuda
Palau
Mexico2
Brazil
Mendoza
Brazil
Thailand
Arkansas
Hawaii

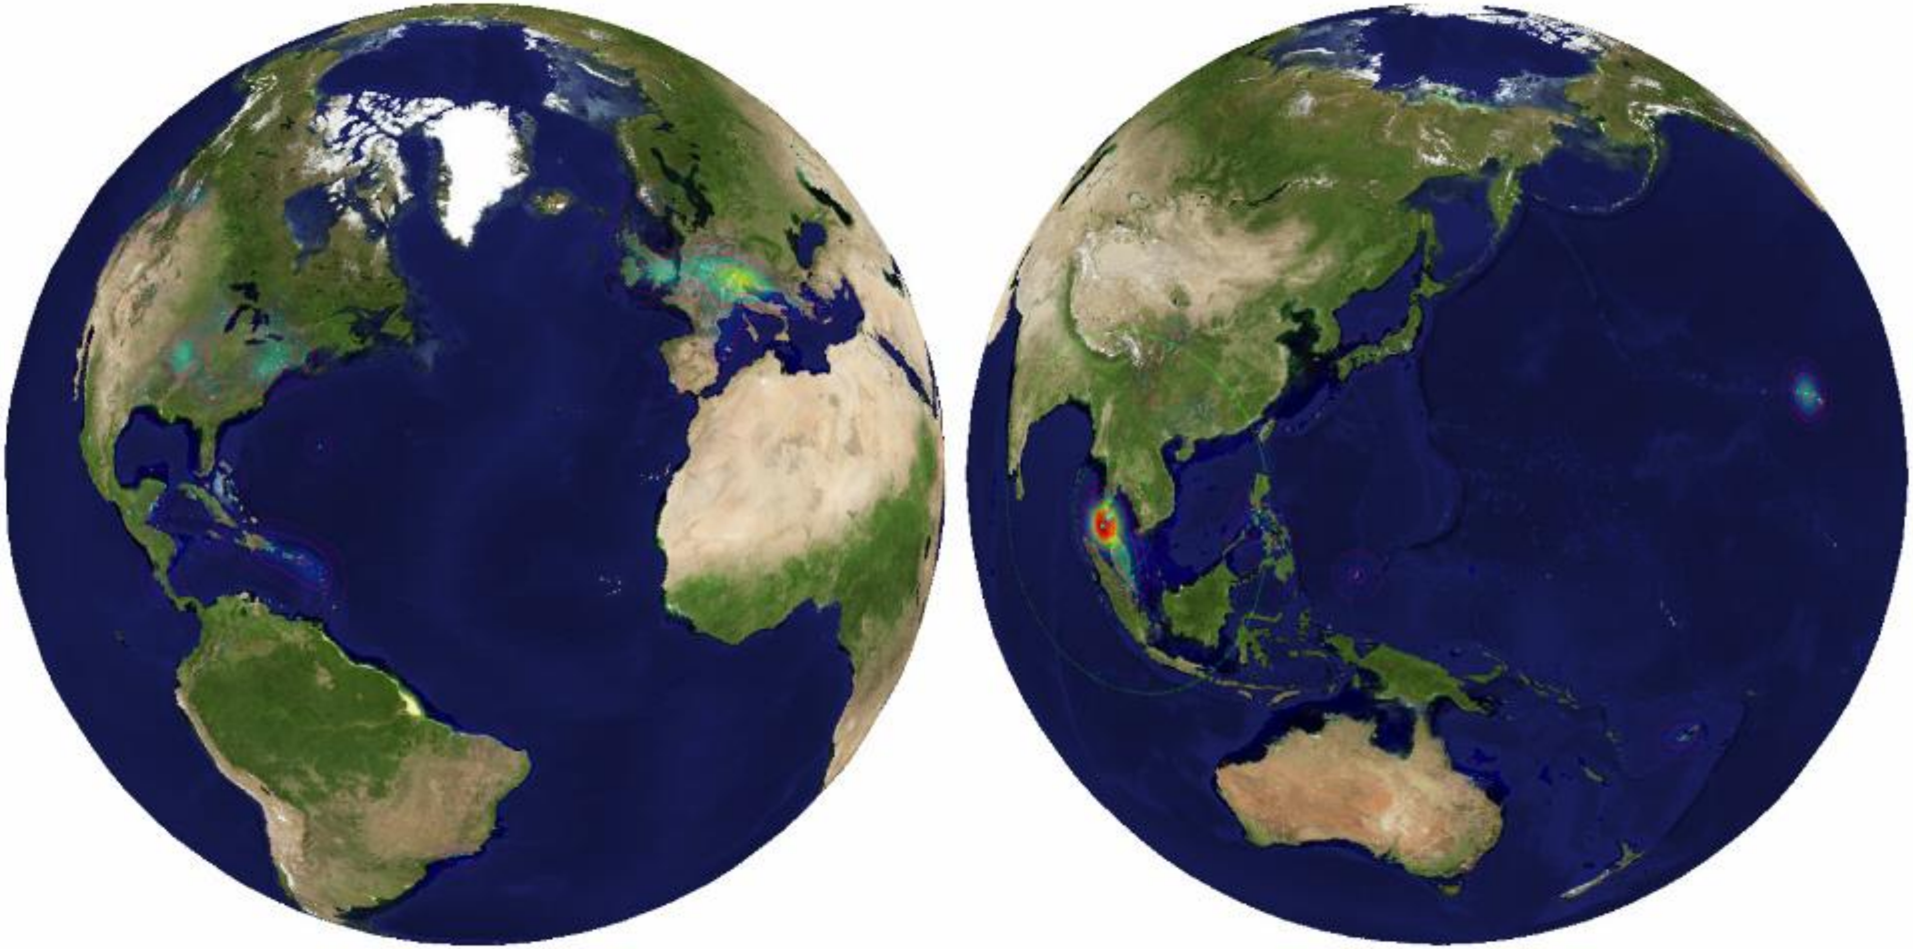

Effect of Dataset Size

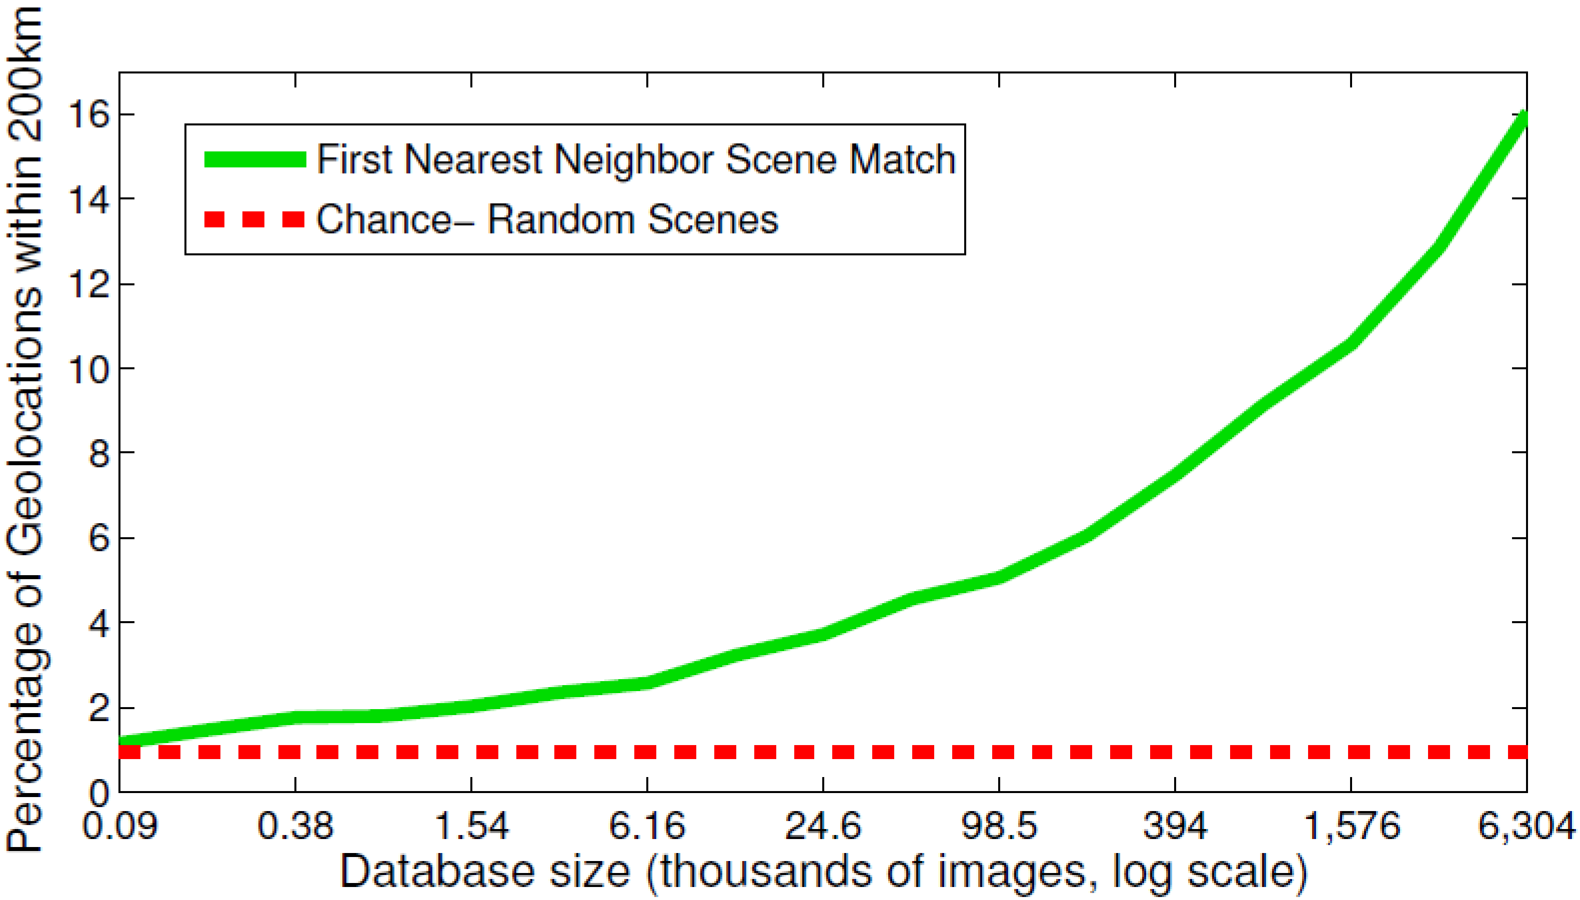

Where is This?

[Vesselova, Kalogerakis, Hertzmann, Hays, Efros. Image Sequence Geolocation. ICCV'09]

Where is This?

Where are These?

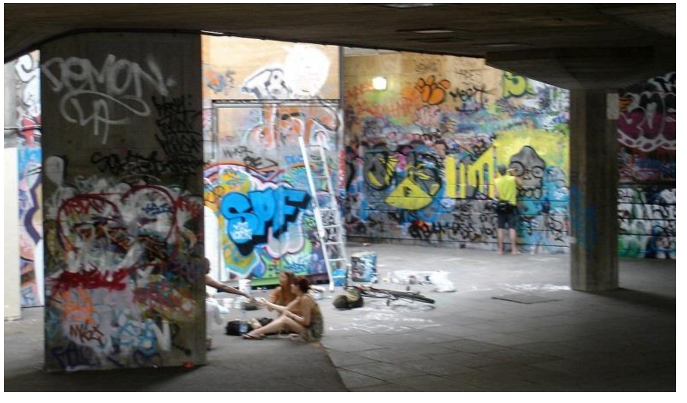

15:14,
June 18th, 2006

16:31,
June 18th, 2006

Where are These?

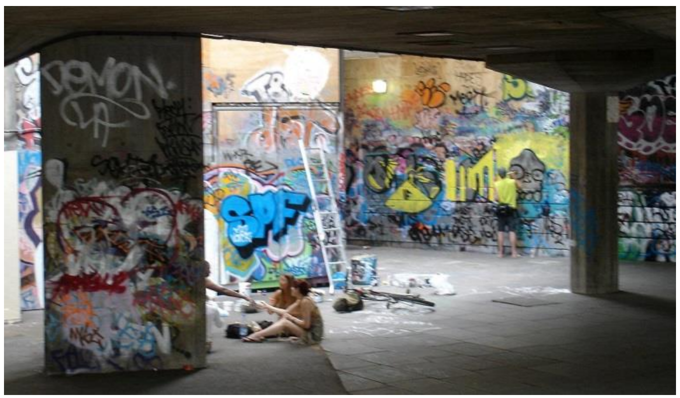

15:14,
June 18th, 2006

16:31,
June 18th, 2006

17:24,
June 19th, 2006

Results

- im2gps – 10% (geo-loc within 400 km)
- temporal im2gps – 56%

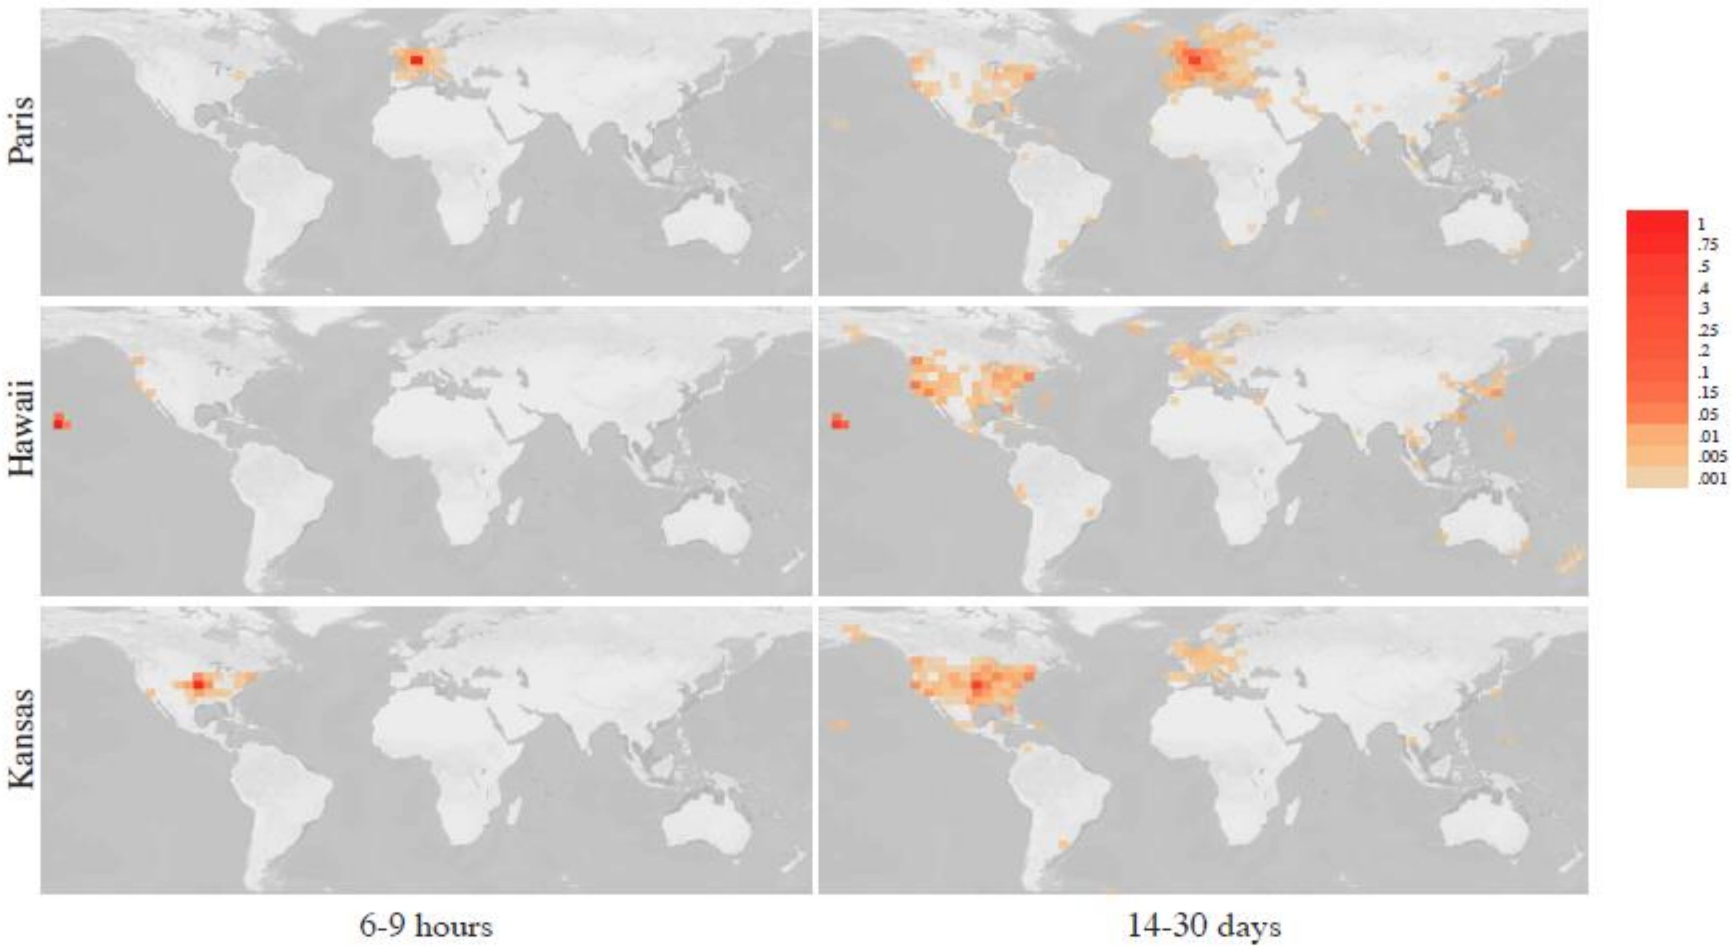

Tiny Images

80 million tiny images: a large dataset for non-parametric object and scene recognition

Antonio Torralba, Rob Fergus and William T. Freeman. PAMI 2008.
http://groups.csail.mit.edu/vision/TinyImages/

32x32


256x256
32x32


256x256
32x32

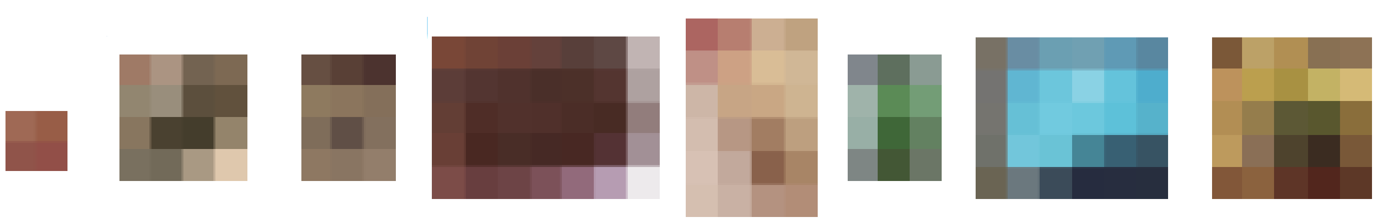

c) Segmentation of 32x32 images

Human Scene Recognition

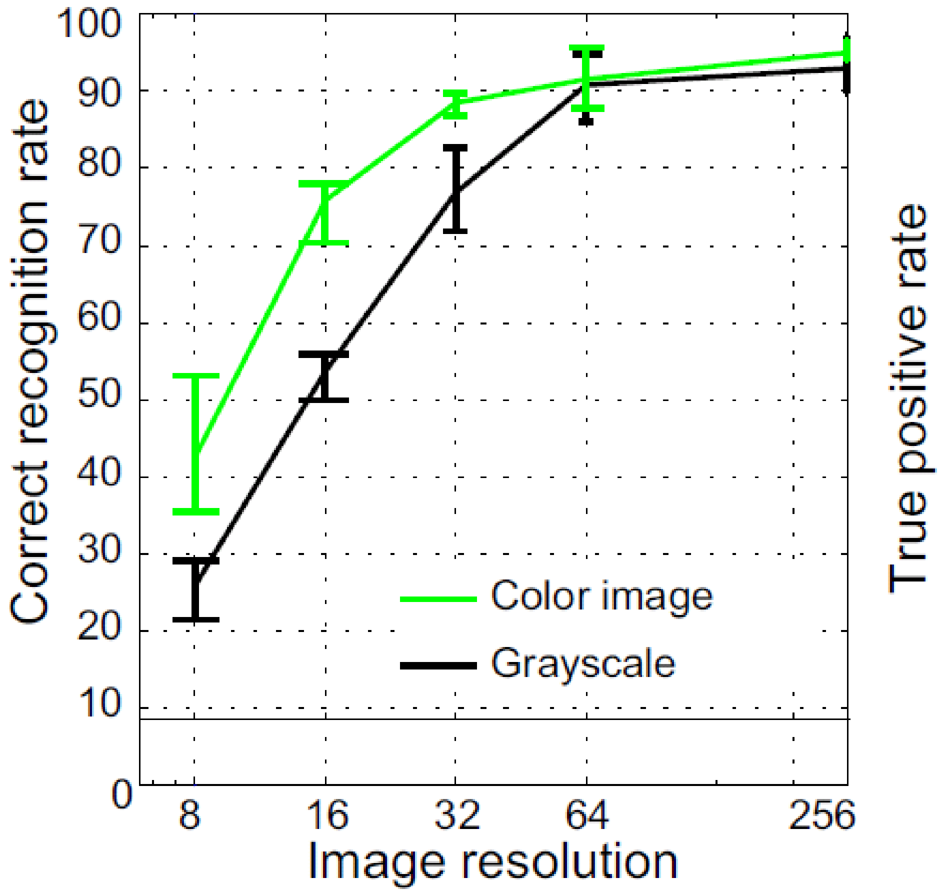

Humans vs. Computers: Car-Image Classification

Humans for 32 pixel tall images

Various computer vision algorithms for full resolution images

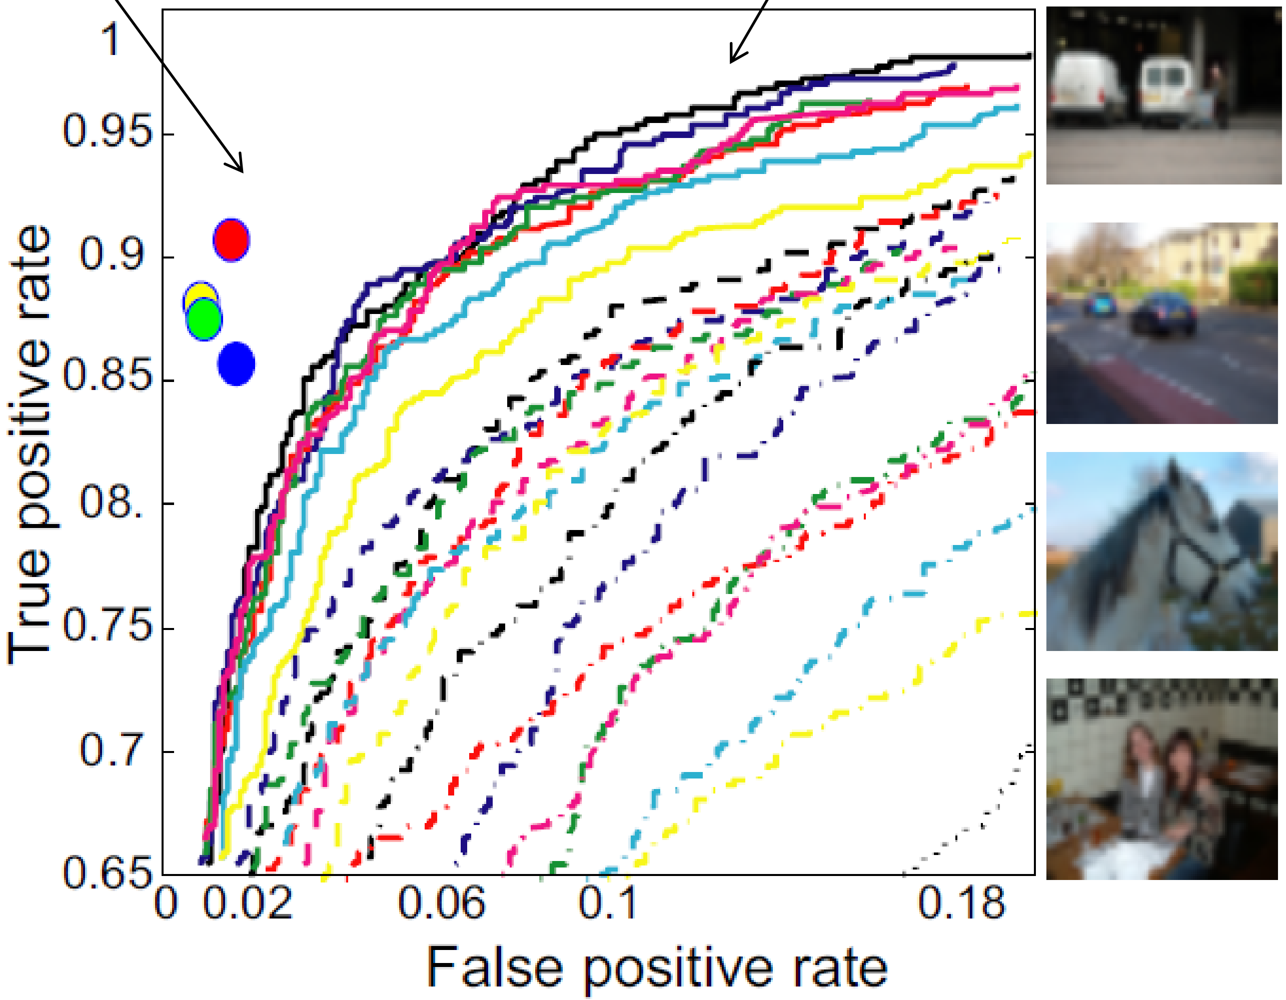

Powers of 10

Number of images on my hard drive: 10^{6}

Number of images seen during my first 10 years: 10^{8}

(3 images/second * 60 * 60 * 16 * 365 * 10 = 630,720,000)

Number of images seen by all humanity: 10^{20}

106,456,367,669 humans[1] * 60 years * 3 images/second * 60 * 60 * 16 * 365 =

1 from http://www.prb.org/Articles/2002/HowManyPeopleHaveEverLivedonEarth.aspx

Number of photons in the universe: 10^{88}

Number of all 32x32 images: 10^{7373}

$256^{32*32*3} \sim 10^{7373}$

Understanding scenes encompasses all kinds of knowledge

But not all scenes are so original








Lots

Of

Images

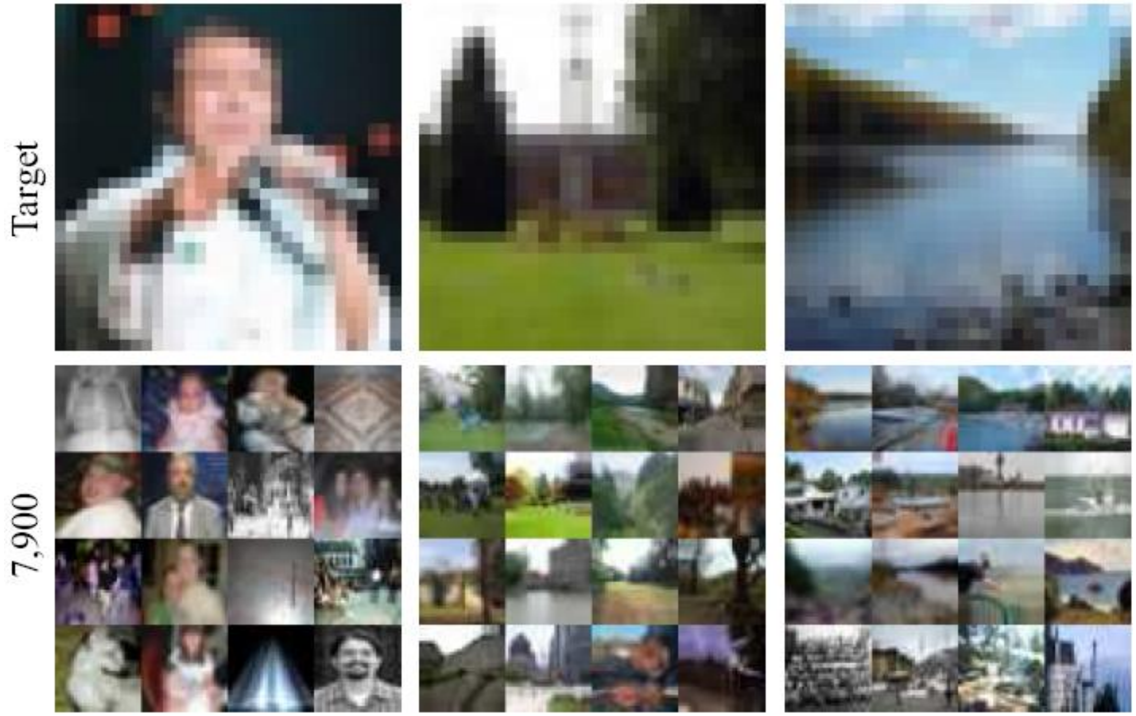

A. Torralba, R. Fergus, W.T.Freeman. PAMI 2008

Lots

Of

Images

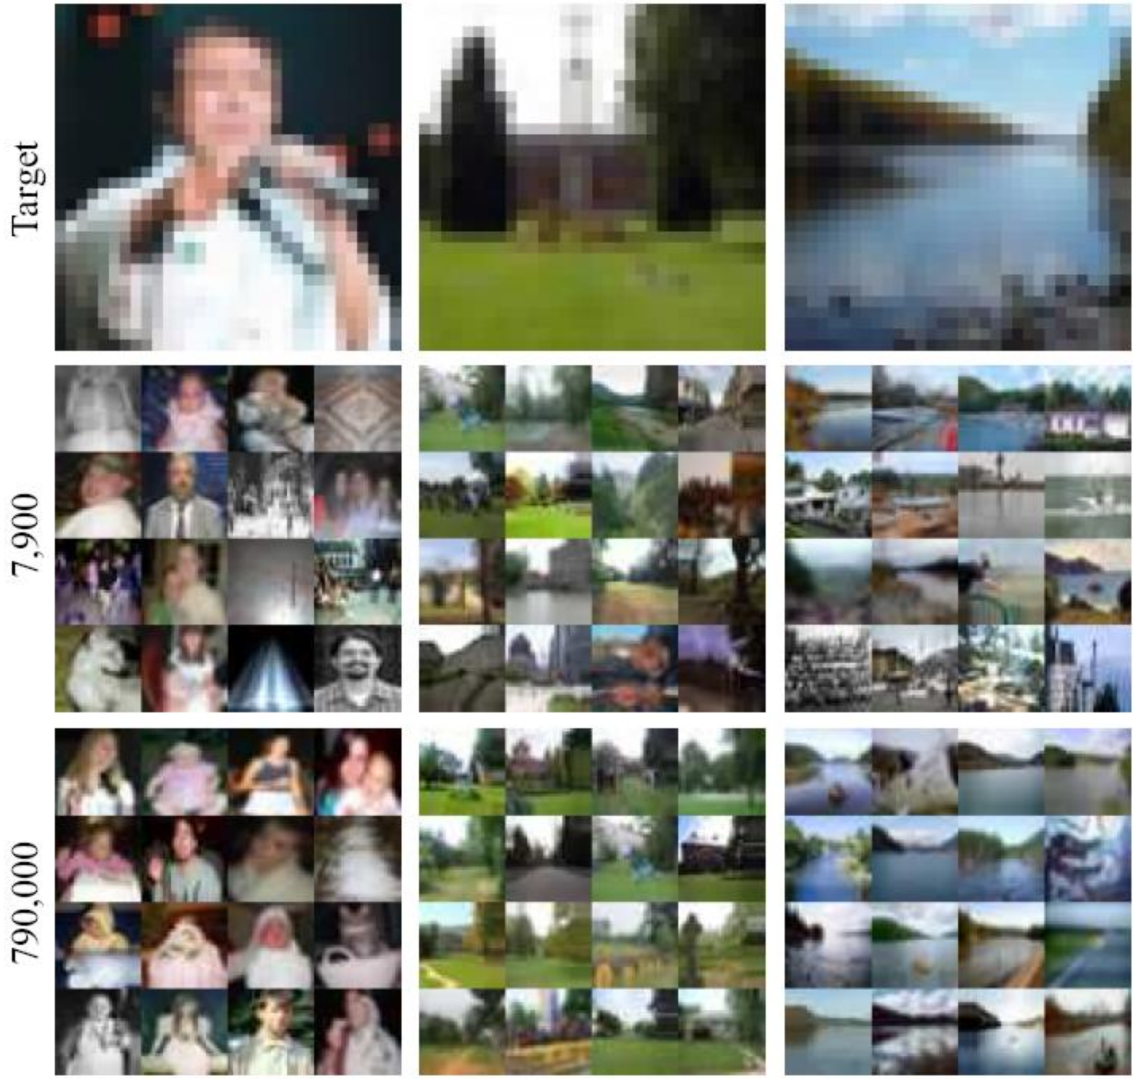

A. Torralba, R. Fergus, W.T.Freeman. PAMI 2008

Lots

Of

Images

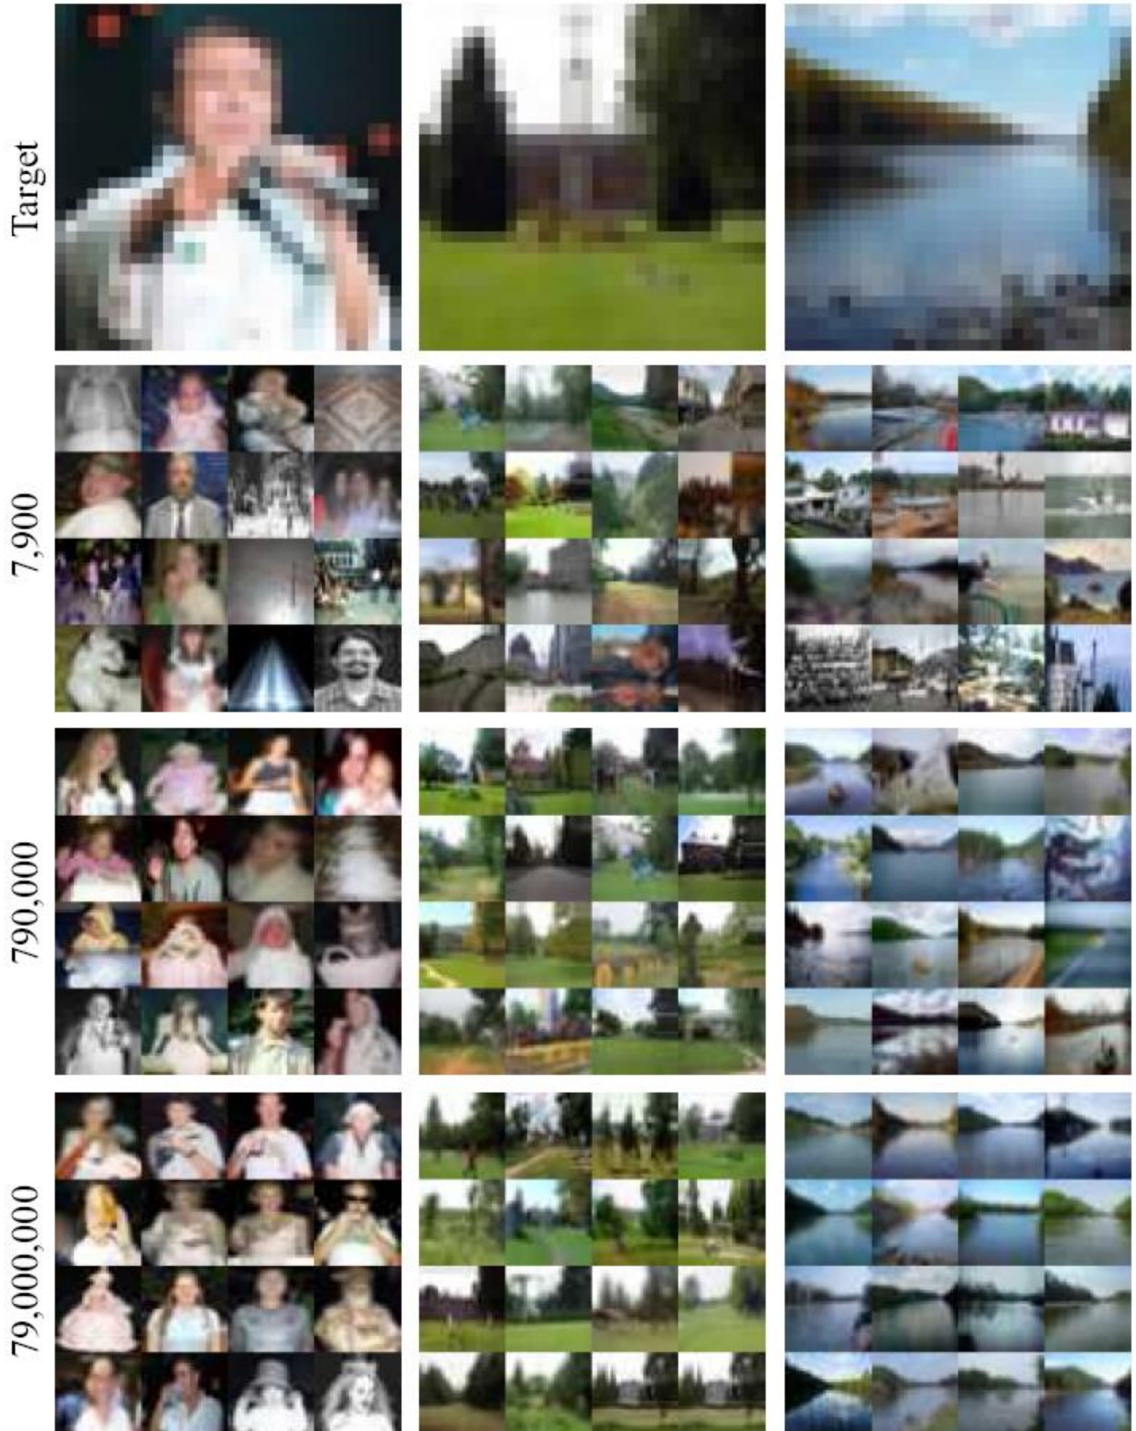

Application: Automatic Colorization

Input

Color Transfer

Color Transfer

Matches (gray)

Matches (w/ color)

Avg Color of Match

Application: Automatic Colorization

Input

Color Transfer

Color Transfer

Matches (gray)

Matches (w/ color)

Avg Color of Match

Short cuts to AI

- With billions of images on the web, it’s often possible to find a close nearest neighbor.

- In such cases, we can shortcut hard problems by “looking up” the answer, stealing the labels from our nearest neighbor.

So what is intelligence?

Weak AI: The simulation of a ‘mind’ is a model for the ‘mind’.

Strong AI: The simulation of a ‘mind’ is a ‘mind’.

Chinese Room, John Searle (1980)

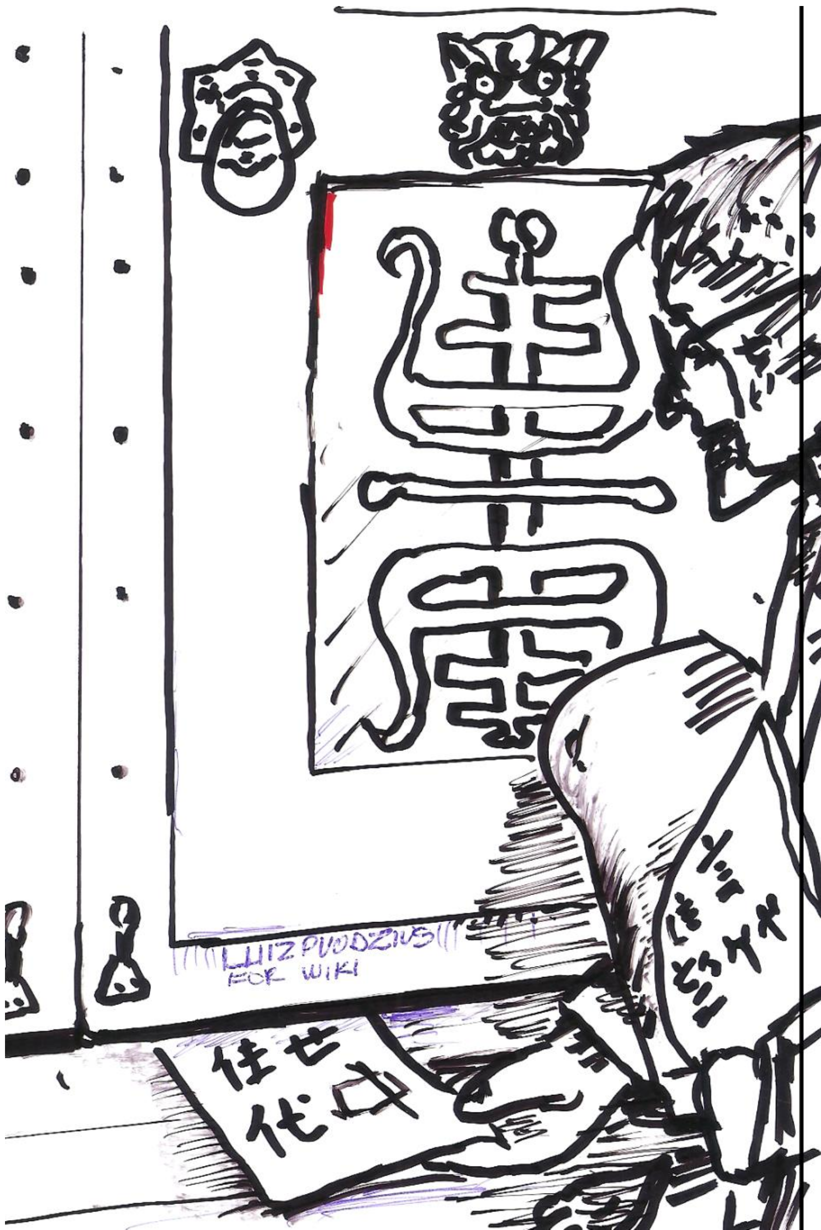

Chinese Room, John Searle (1980)

If a machine can convincingly simulate an intelligent conversation, does it necessarily understand? In the experiment, Searle imagines himself in a room, acting as a computer by manually executing a program that convincingly simulates the behavior of a native Chinese speaker.

Most of the discussion consists of attempts to refute it. "The overwhelming majority," notes *BBS* editor Stevan Harnad,“ still think that the Chinese Room Argument is dead wrong." The sheer volume of the literature that has grown up around it inspired Pat Hayes to quip that the field of cognitive science ought to be redefined as "the ongoing research program of showing Searle's Chinese Room Argument to be false.

Yann LeCun

October 23 at 9:58pm ·

Questions from the piece:

Q1. Does the Chinese Room argument prove the impossibility of machine consciousness?
A1: Hell no. ... See More

Can Machines Become Moral?

The question is heard more and more often, both from those who think that machines cannot become moral, and who think that to believe otherwise is a dangerous illusion, and from those who think that machines must become moral,...

BIGQUESTIONSONLINE.COM | BY DON HOWARD

You and 156 others

30 Comments 20 Shares

Like Comment Share

Mechanical Turk

- von Kempelen, 1770.
- Robotic chess player.
- Clockwork routines.
- Magnetic induction (not vision)

- Toured the world; played Napoleon Bonaparte and Benjamin Franklin.

Mechanical Turk

- It was all a ruse!
- Ho ho ho.

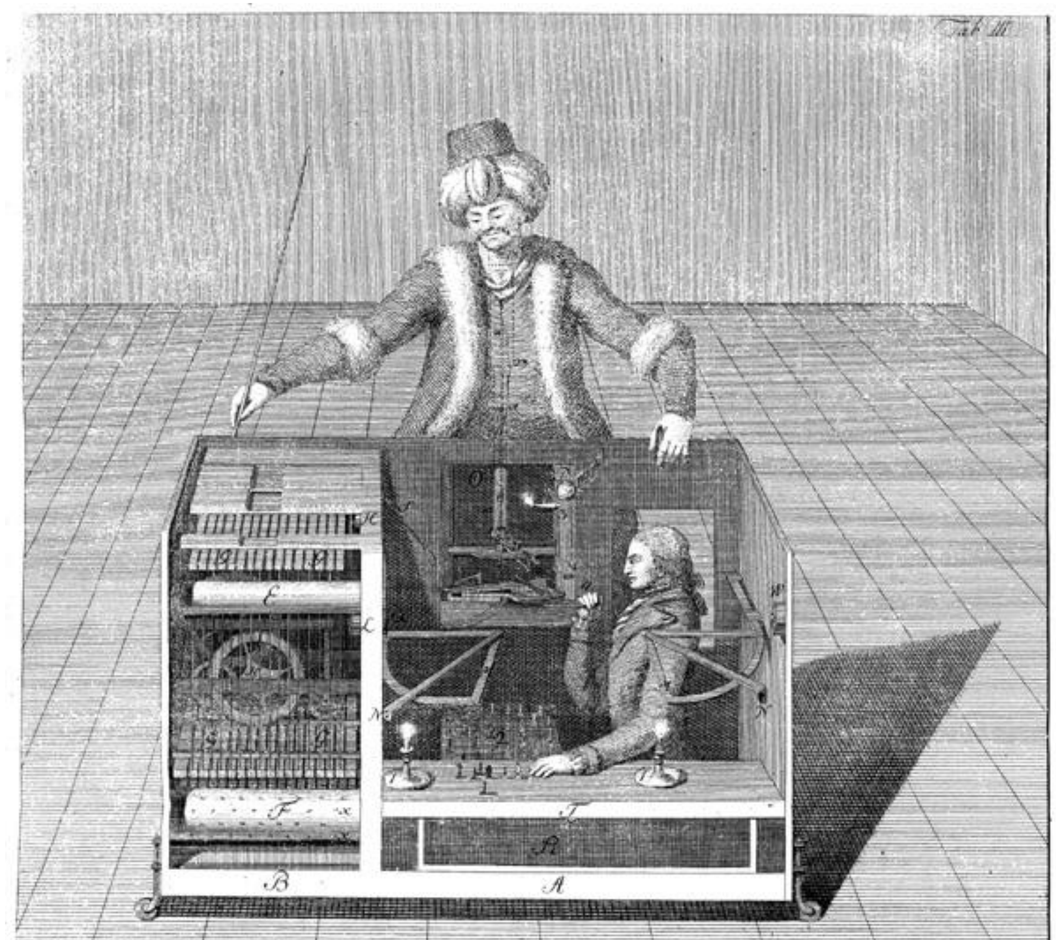

Amazon Mechanical Turk

Artificial artificial intelligence.

Launched 2005.
Small tasks, small pay.
Used extensively in data collection.

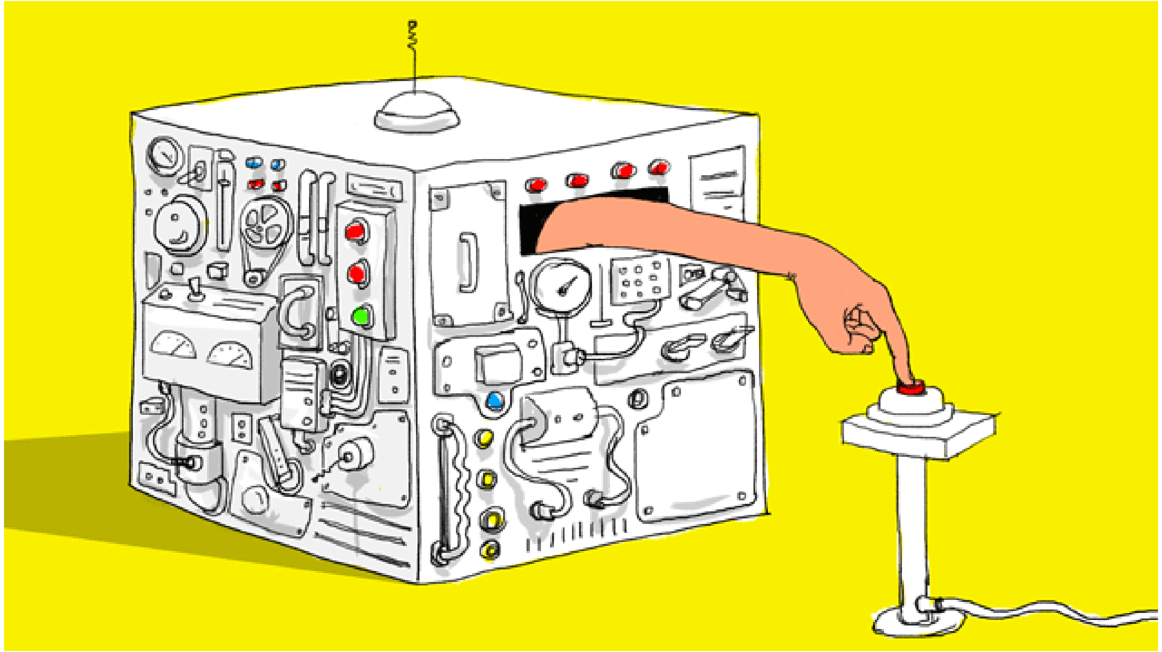

Amazon Mechanical Turk

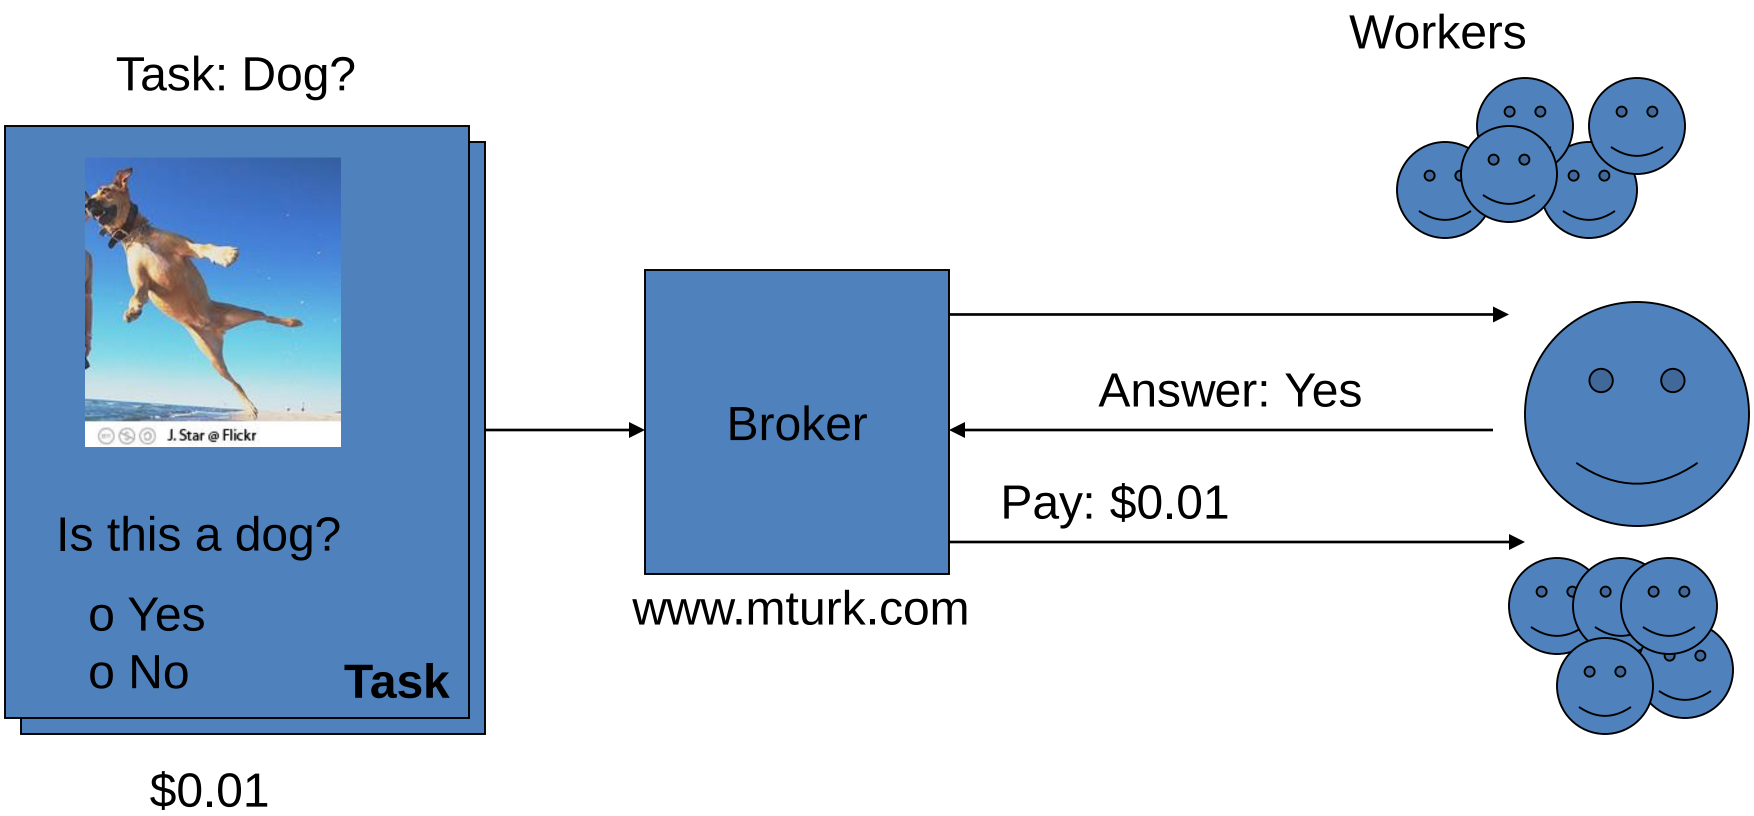

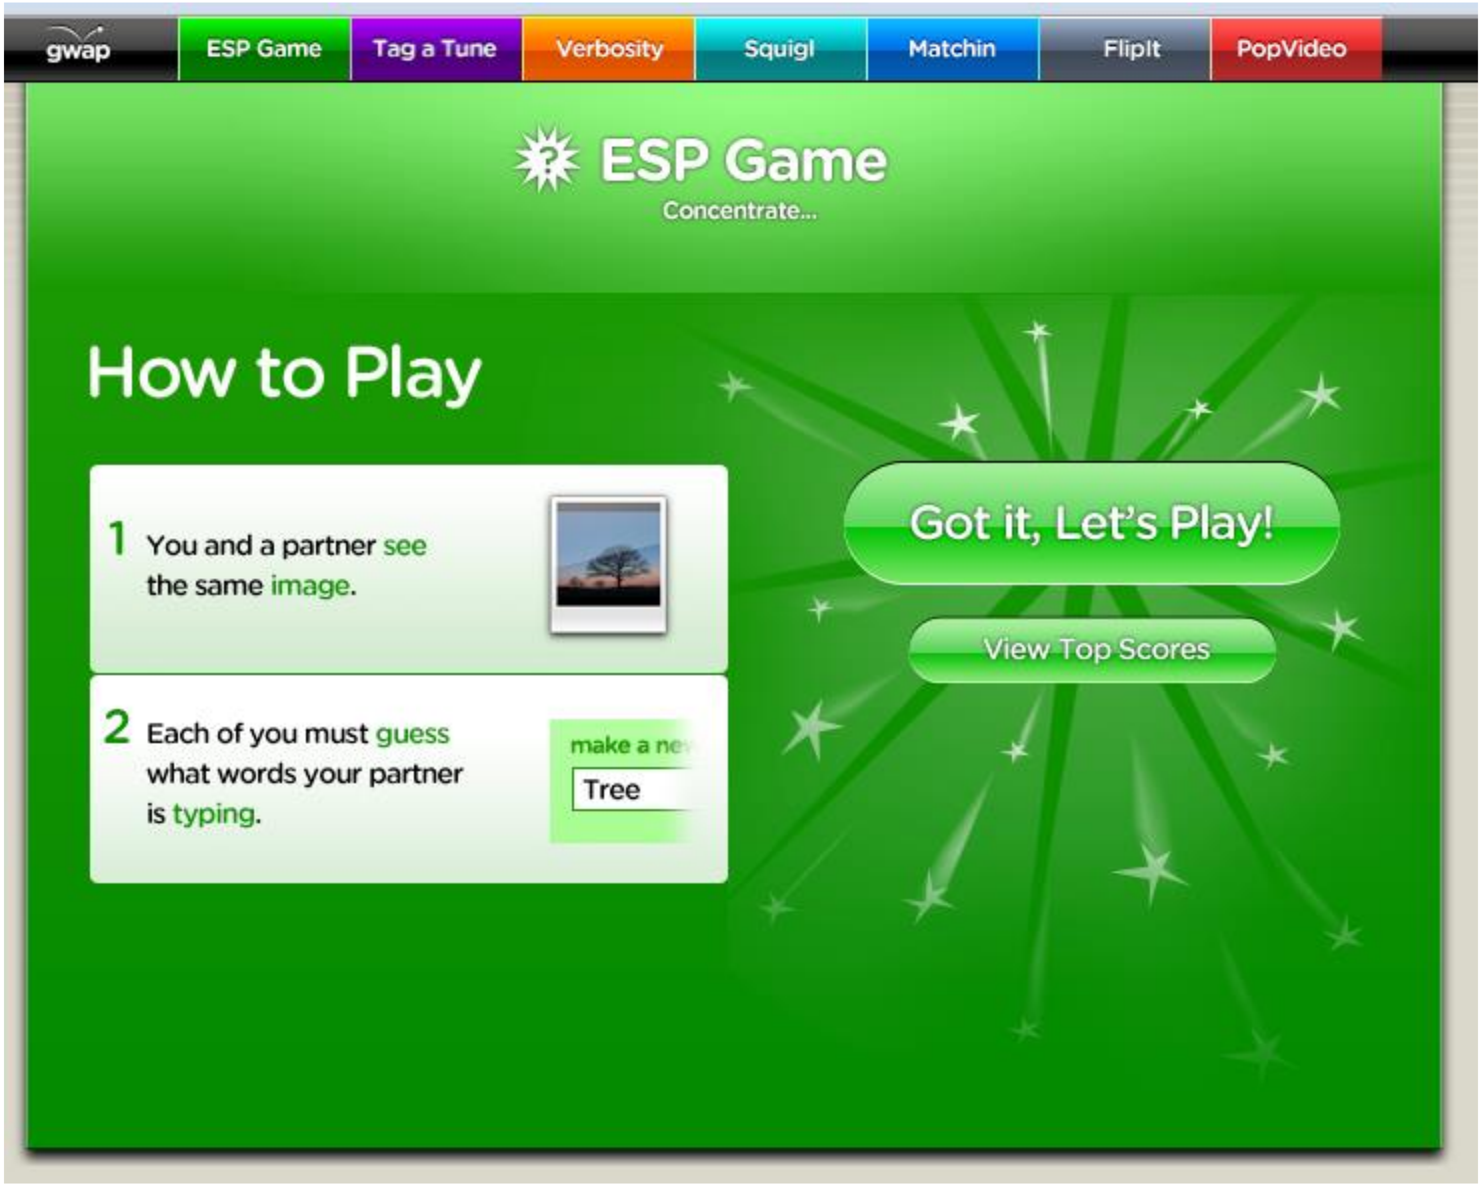

Luis von Ahn and Laura Dabbish. Labeling Images with a Computer Game. ACM Conf. on Human Factors in Computing Systems, CHI 2004

ESP Game
Tag a Tune
Verbosity
Squigl
Matchin
logged i
score
0
ESP Game
Concentrate...
time
2:56
What do you see?
taboo words
student
guesses
submit
pass
Play Anonymously
Screencast-O-Matic.com

Vision (Segmentation): LabelMe

http://labelme.csail.mit.edu

"Open world" database annotated by the community*

Notes on Image Annotation, Barriuso and Torralba 2012. http://arxiv.org/abs/1210.3448

Utility data annotation via Amazon Mechanical Turk

X 100 000 = $5000

Alexander Sorokin

David Forsyth

CVPR Workshops 2008

Slides by Alexander Sorokin

Building datasets

6000 images
from flickr.com

Annotators

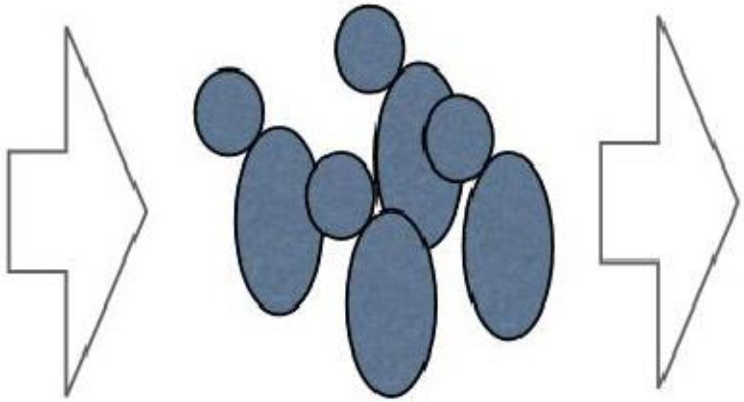

amazonmechanical turk
Artificial Artificial Intelligence

Is there an Indigo bunting in the image?

100s of
training images

Slide credit: Welinder et al

Issues

- Quality?
  - How good is it?
  - How to be sure?
- Price?
  - Trade off between throughput and cost
    - *NOT* as much of a trade off with quality
  - Higher pay can actually attract scammers

Annotation quality

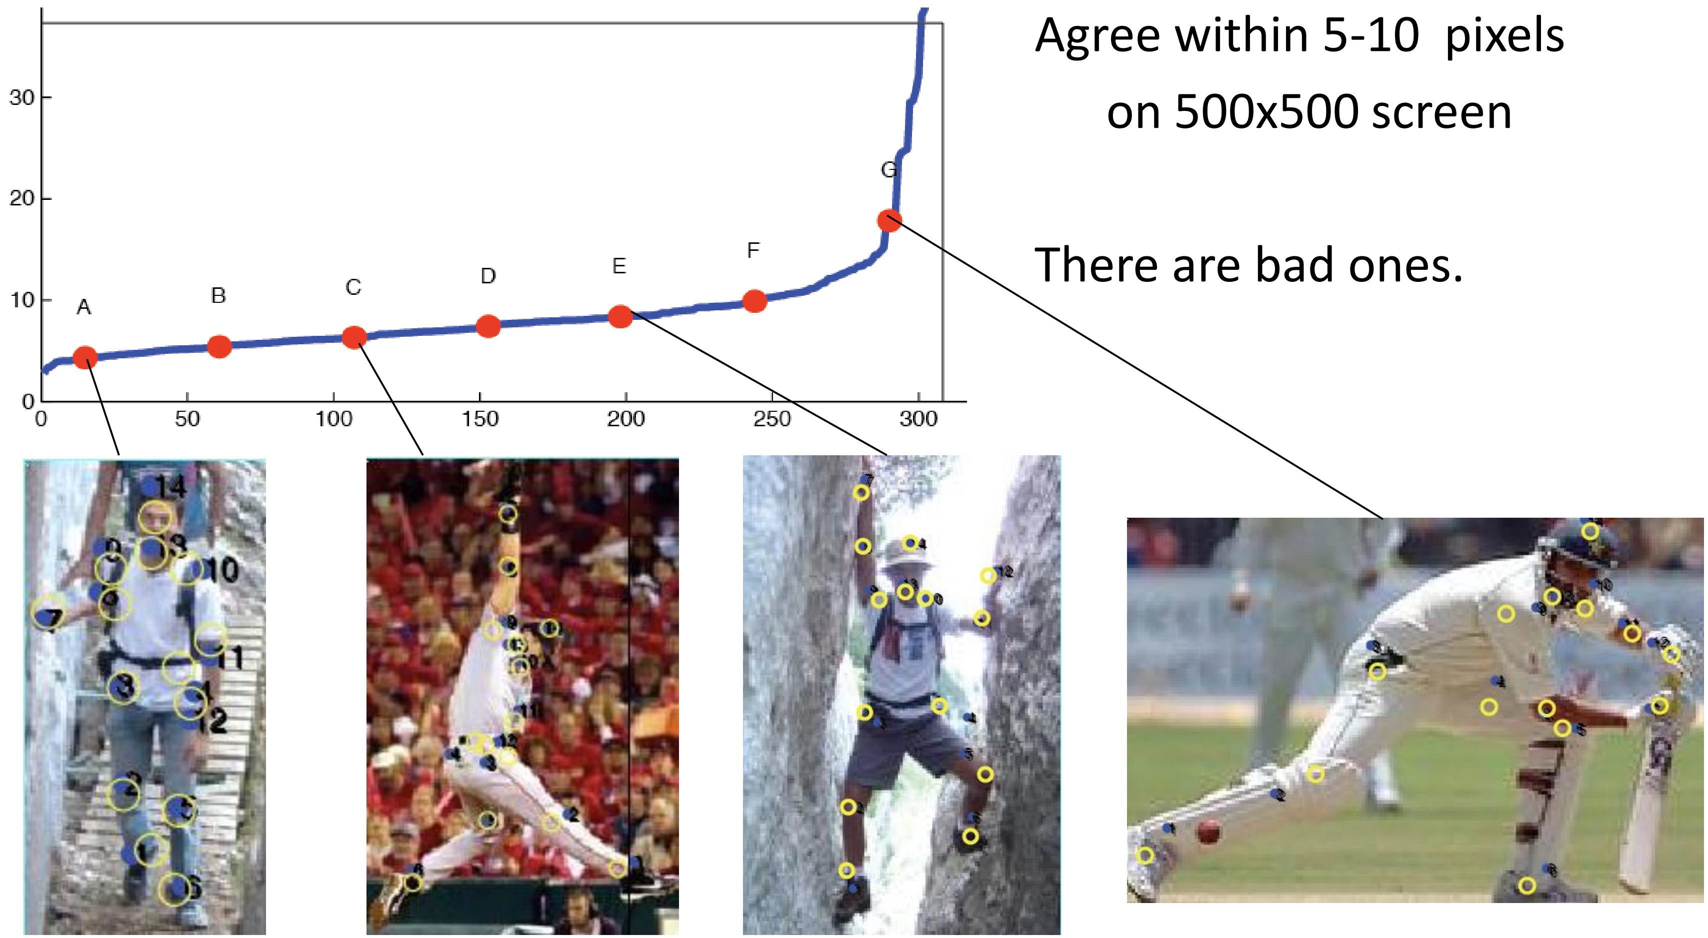

A C E G

Agree within 5-10 pixels on 500x500 screen

There are bad ones.

Ensuring Annotation Quality

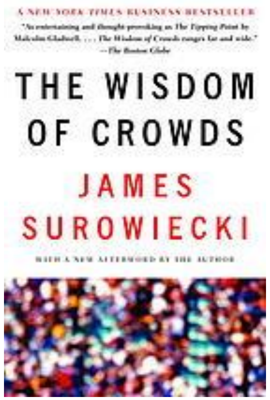

- Consensus / Multiple Annotation / "Wisdom of the Crowds"

  Not enough on its own, but widely used

- Gold Standard / Sentinel
  - Special case: qualification exam

  Widely used and most important. Find good annotators and keep them honest.

- Grading Tasks
  - A second tier of workers who grade others

  Not widely used

Task: Find the Indigo Bunting

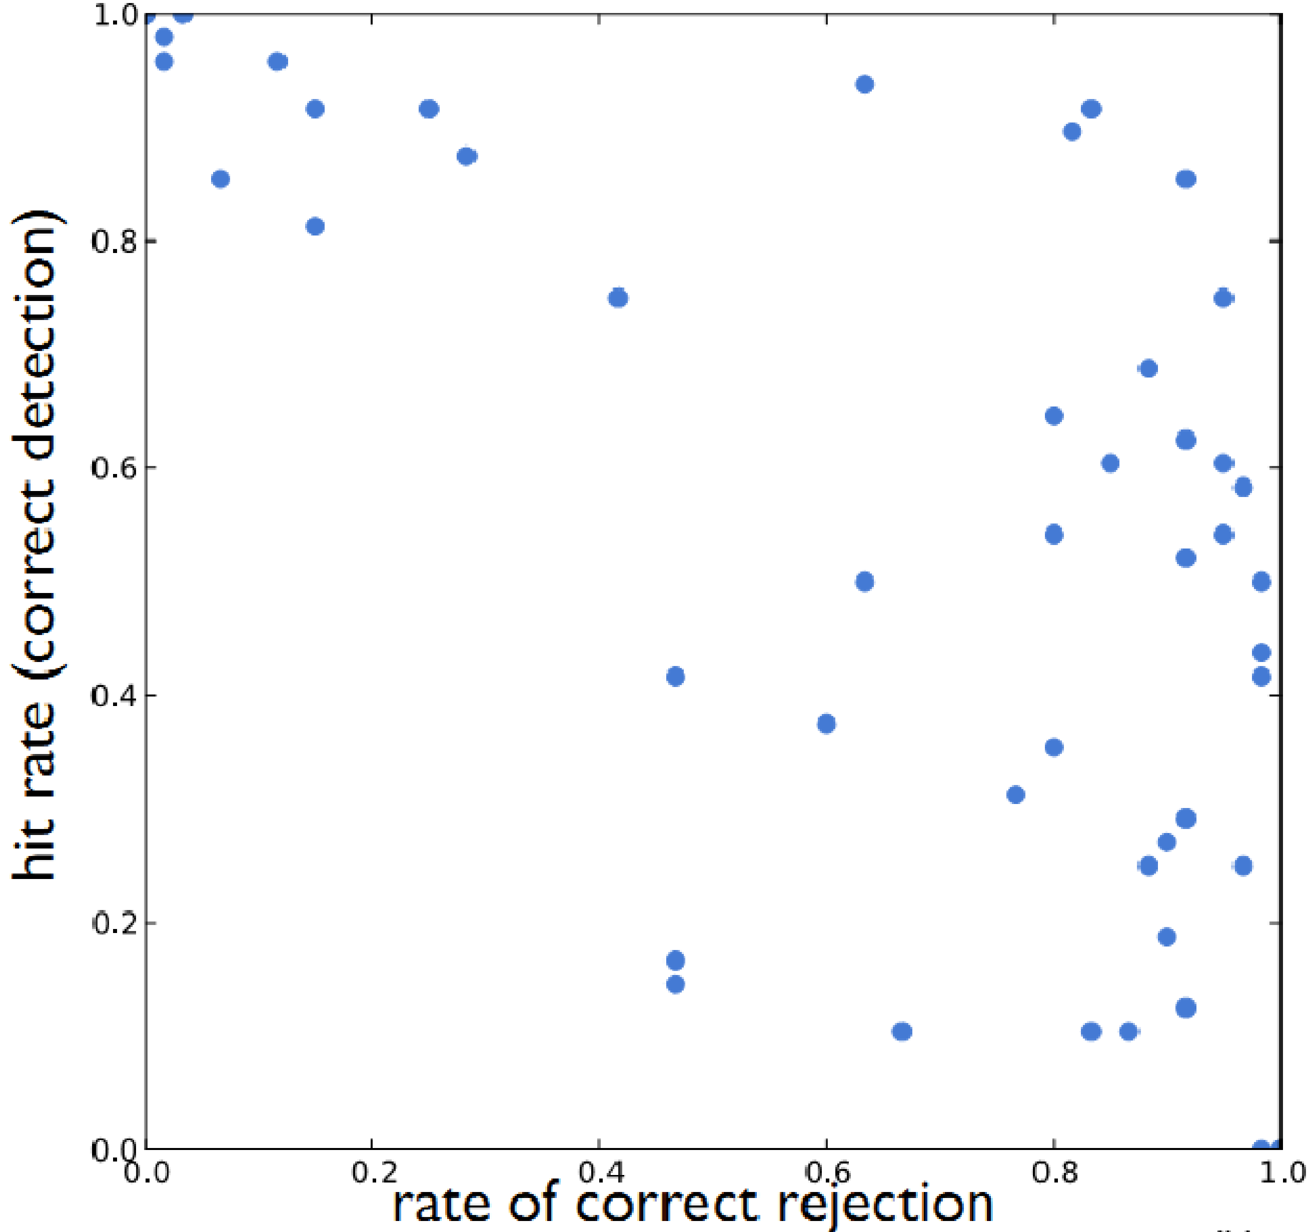

Slide credit: Welinder et al

Task: Find the Indigo Bunting

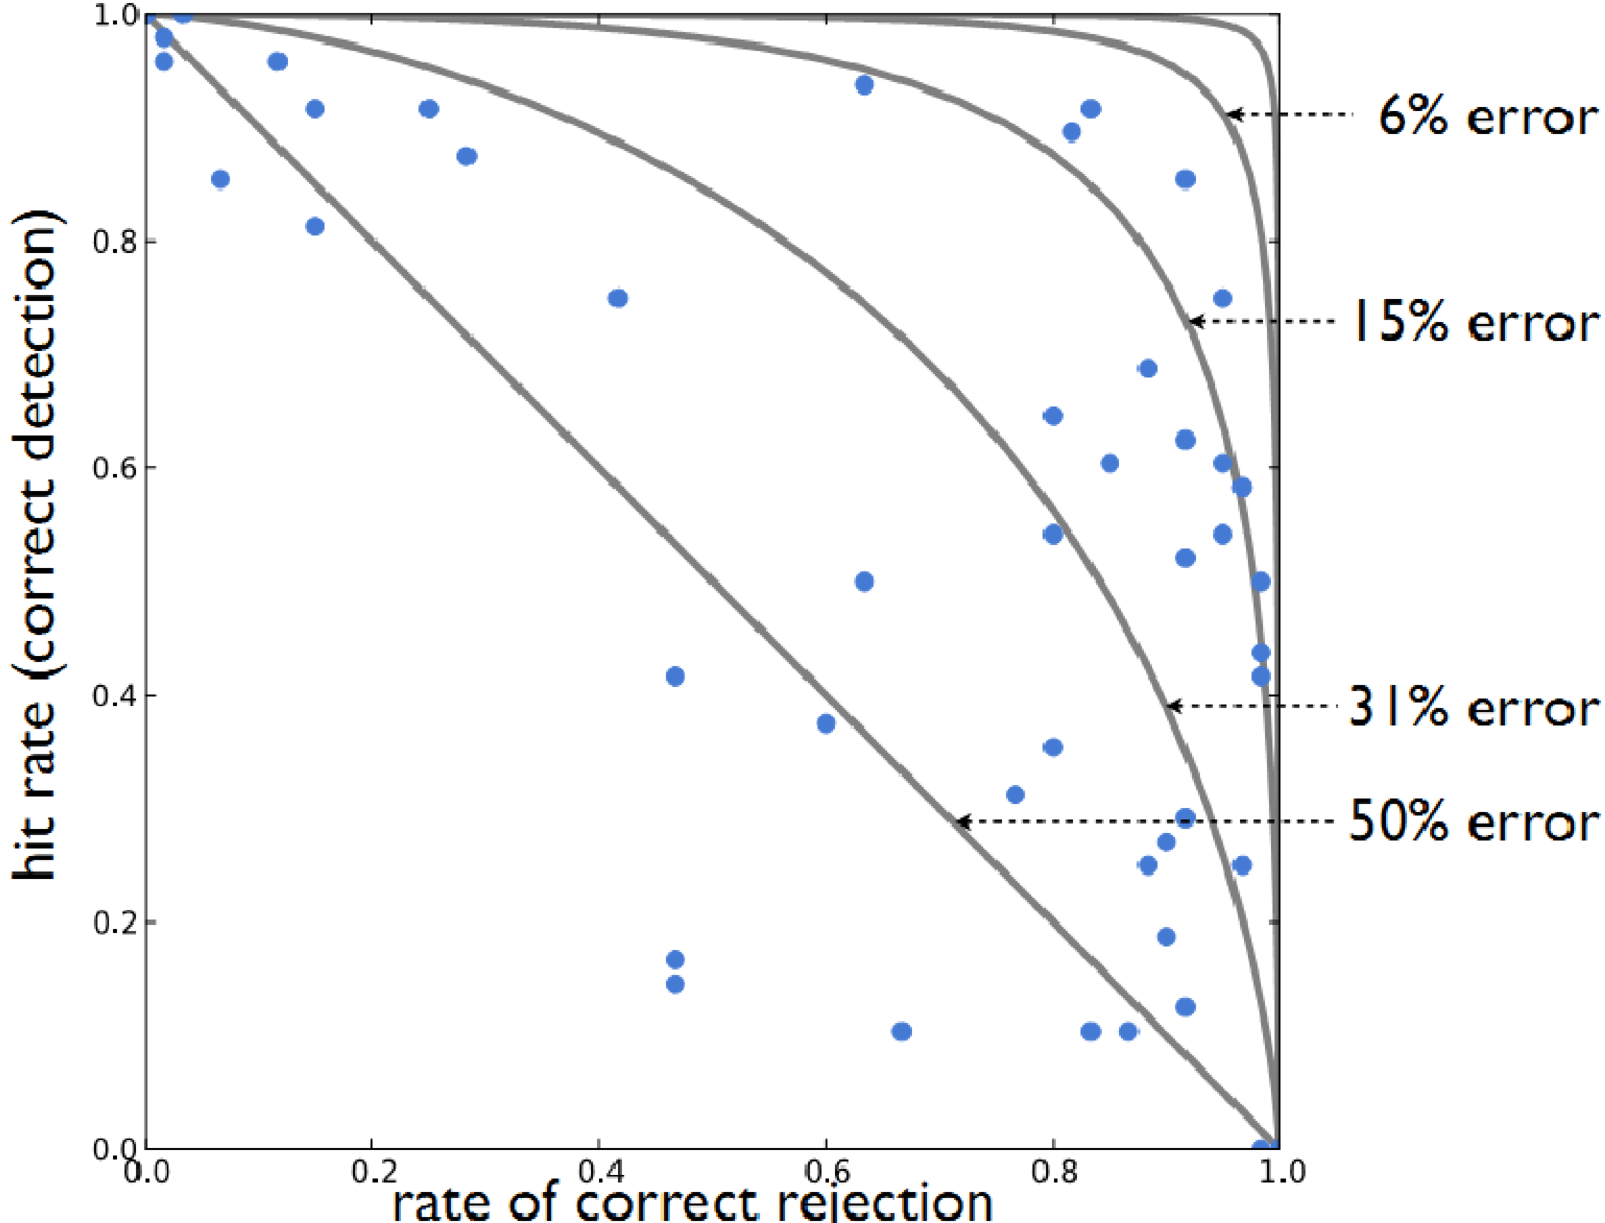

Slide credit: Welinder et al

Task: Find the Indigo Bunting

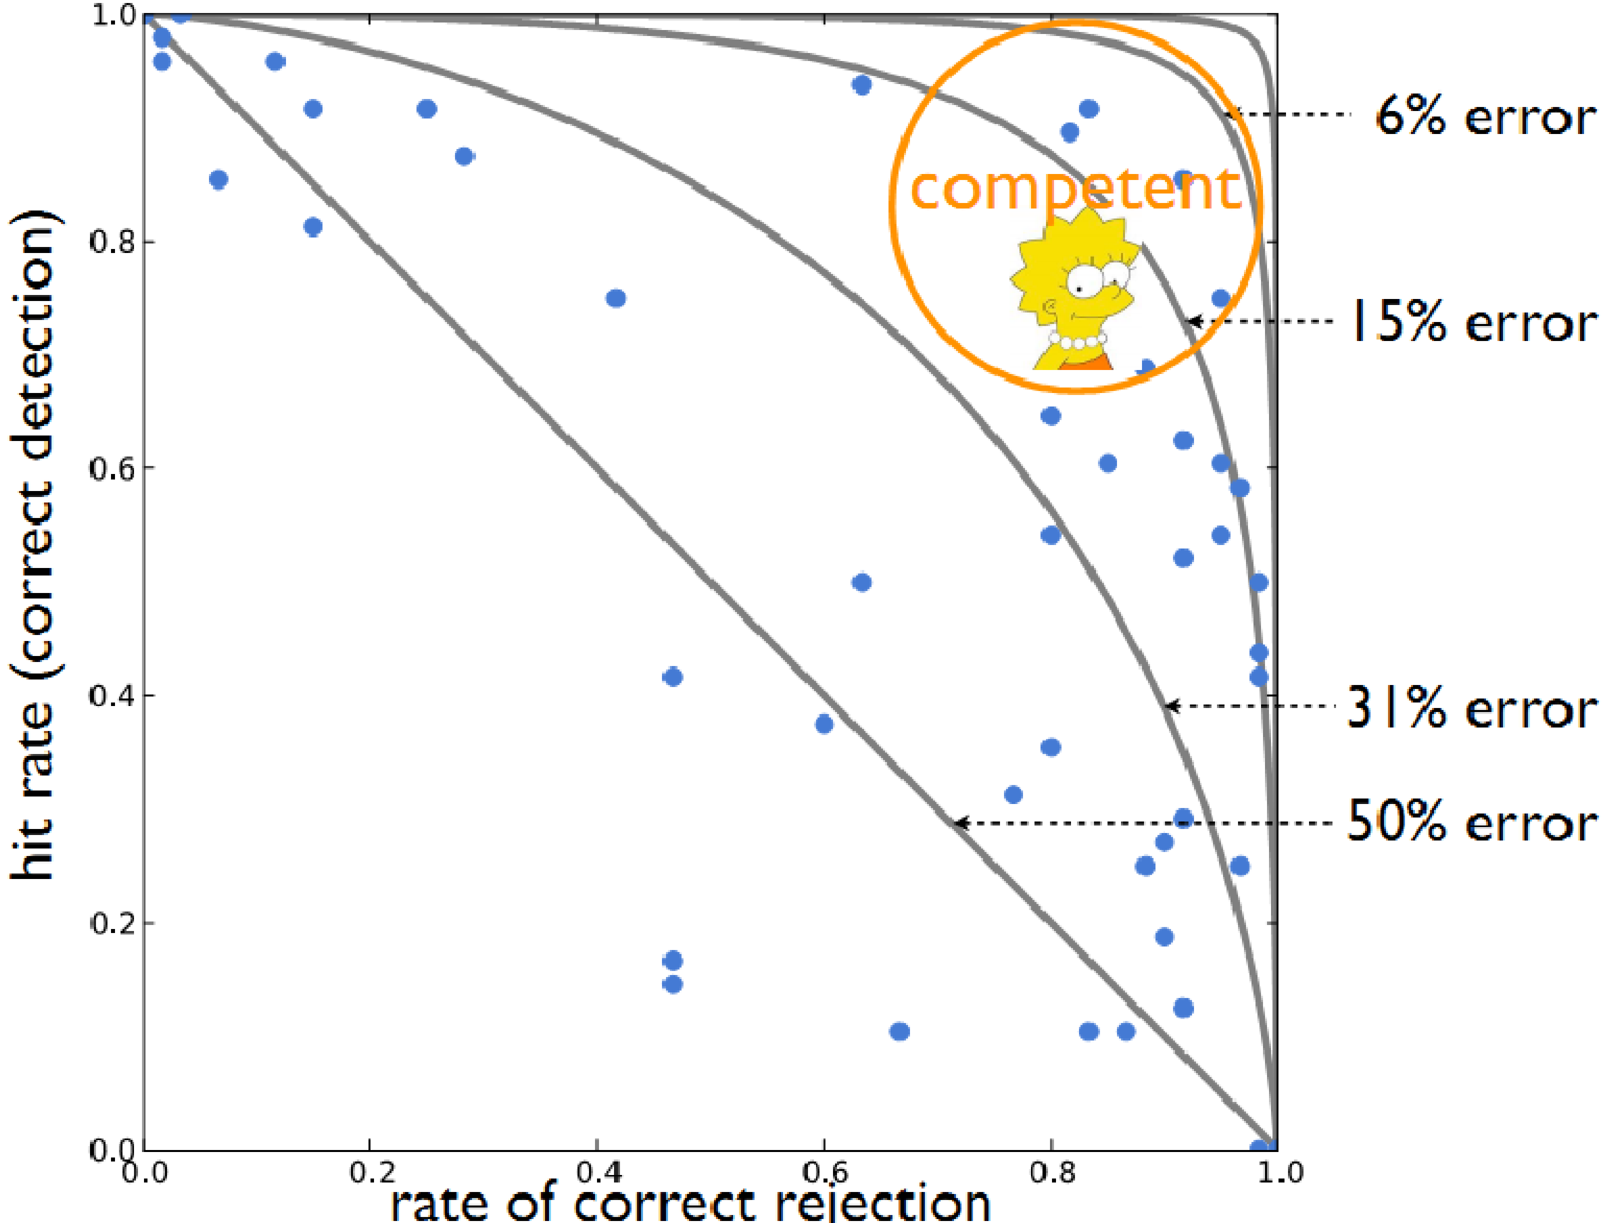

Slide credit: Welinder et al

Task: Find the Indigo Bunting

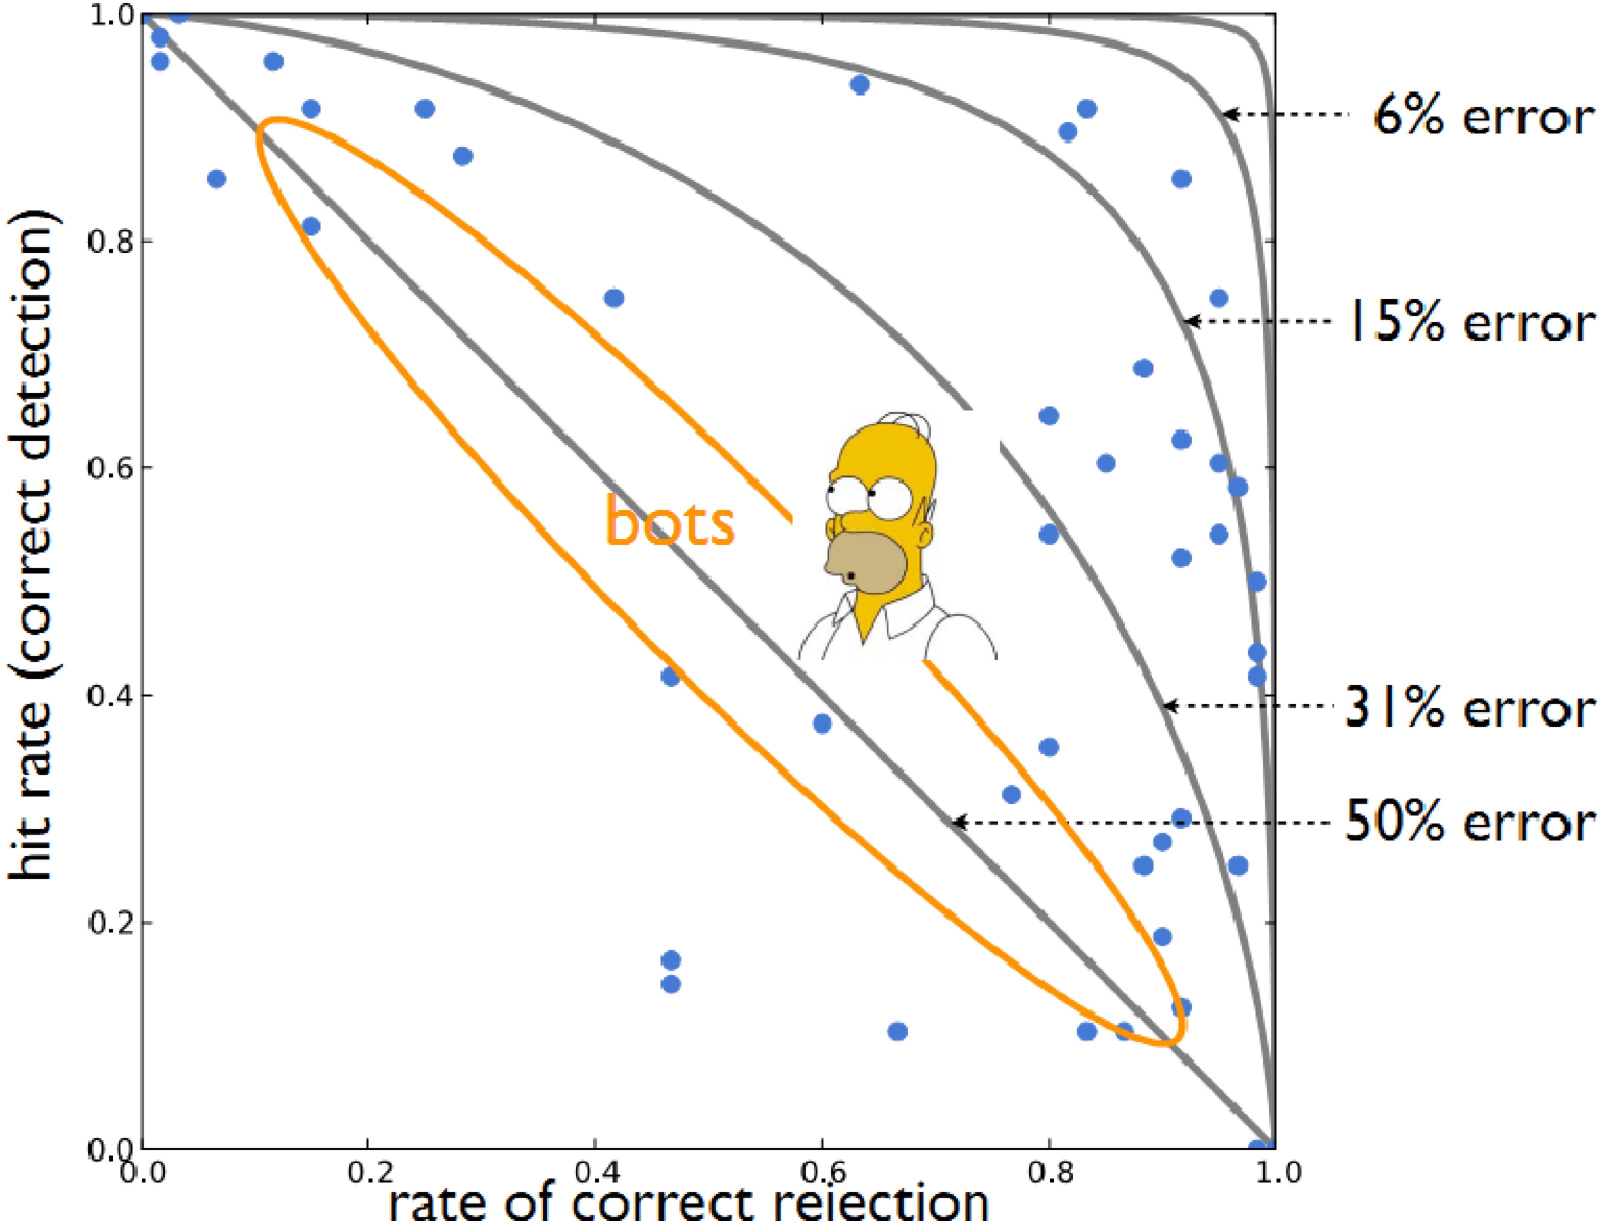

Slide credit: Welinder et al

Task: Find the Indigo Bunting

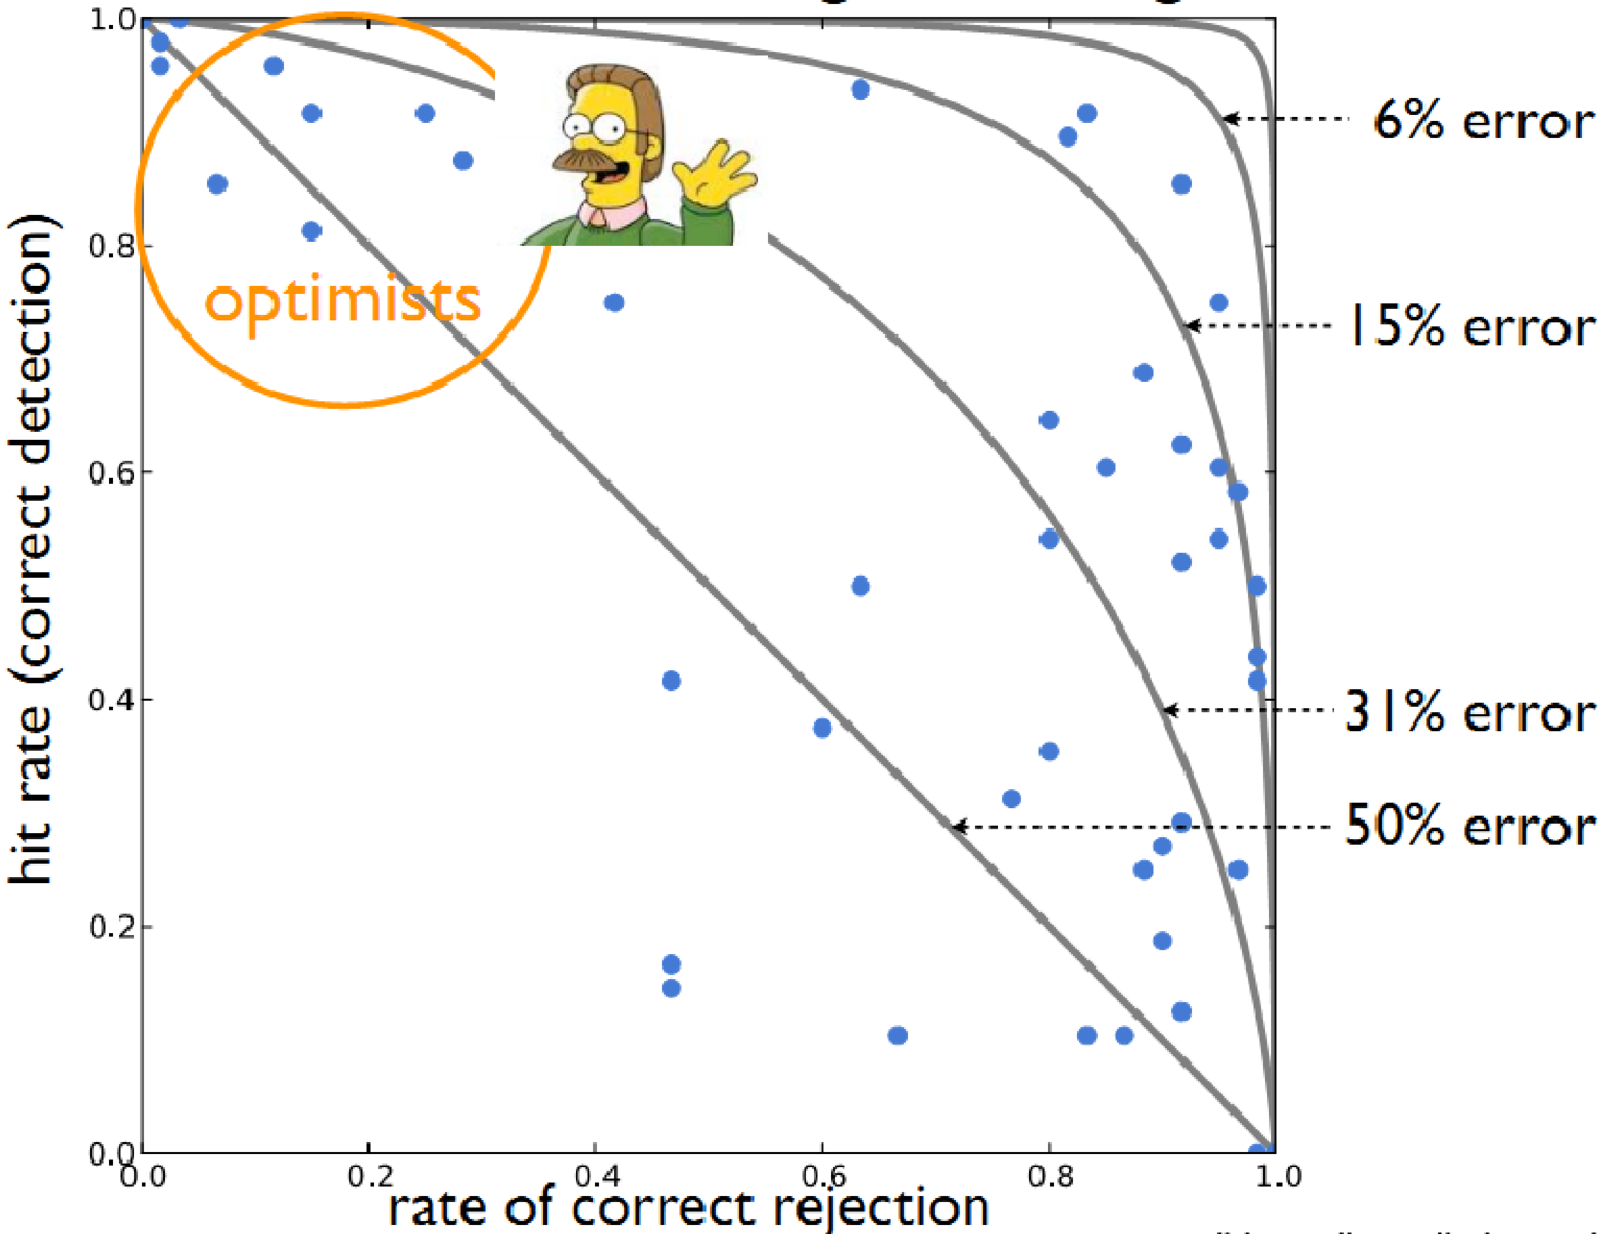

Slide credit: Welinder et al

Task: Find the Indigo Bunting

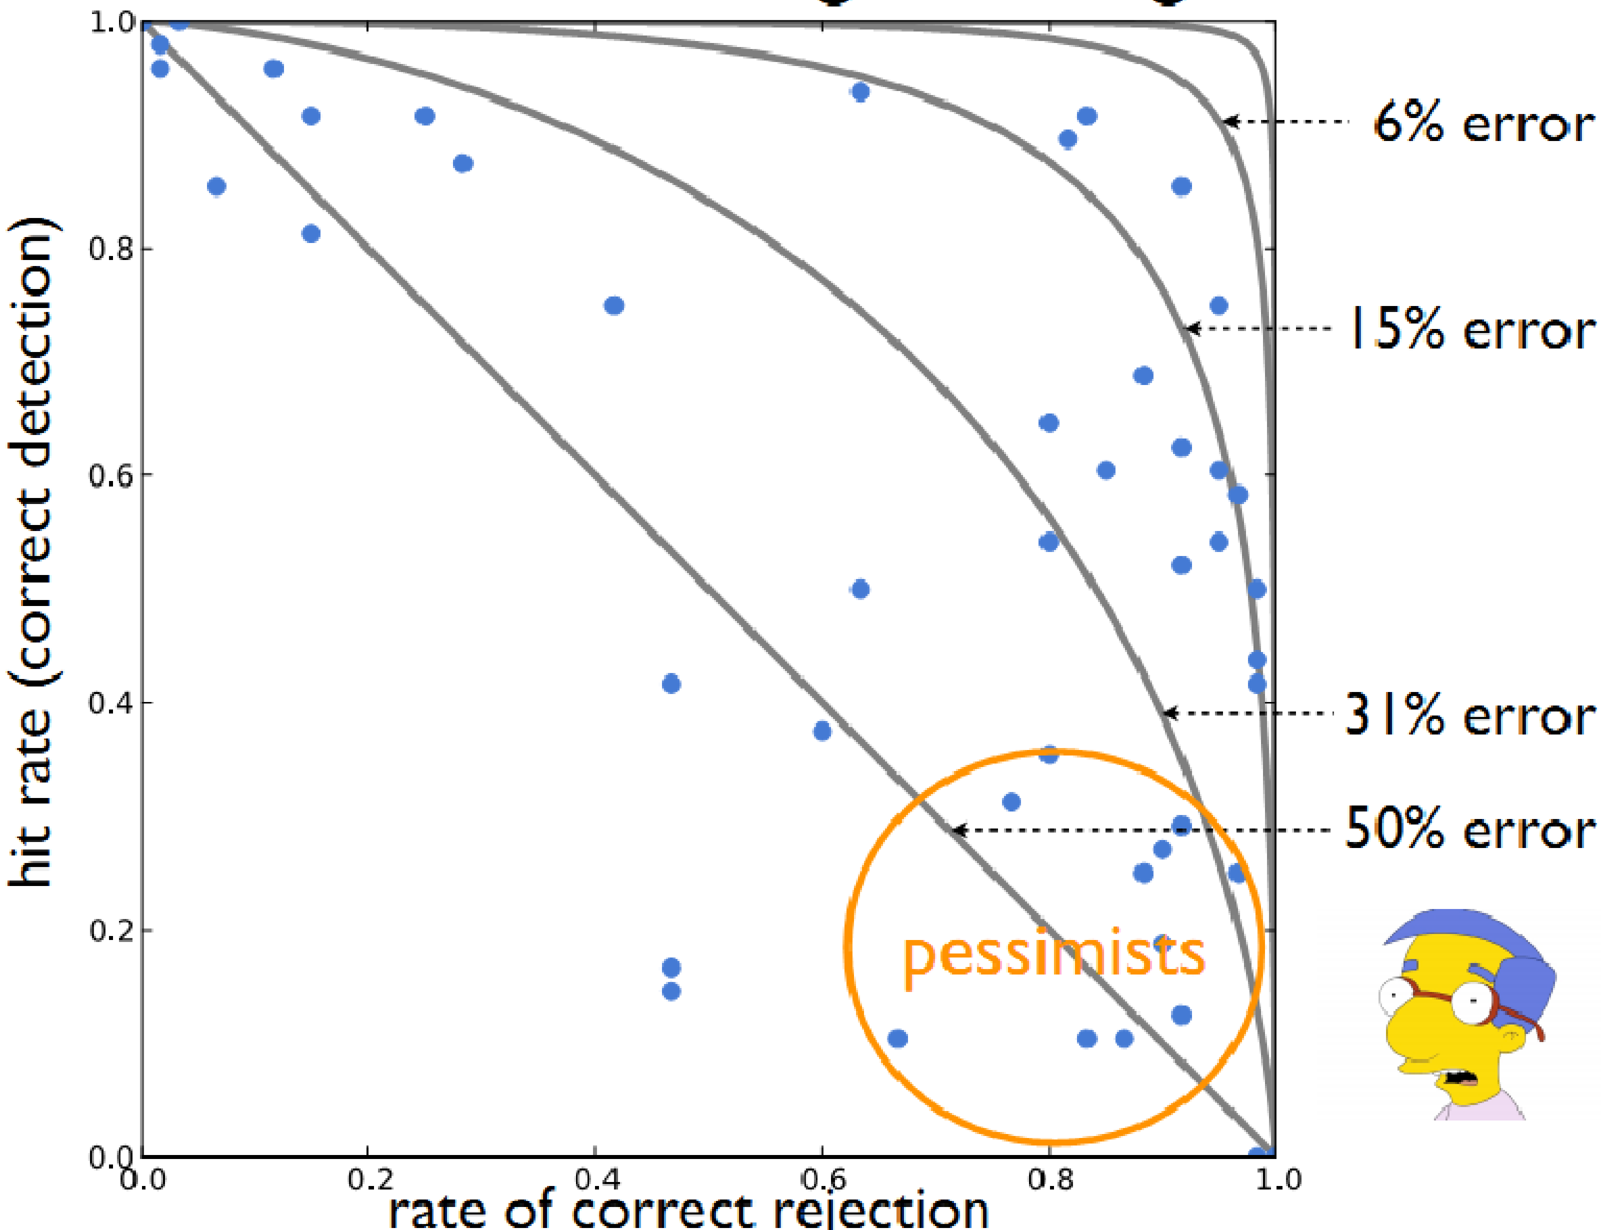

Slide credit: Welinder et al

Task: Find the Indigo Bunting

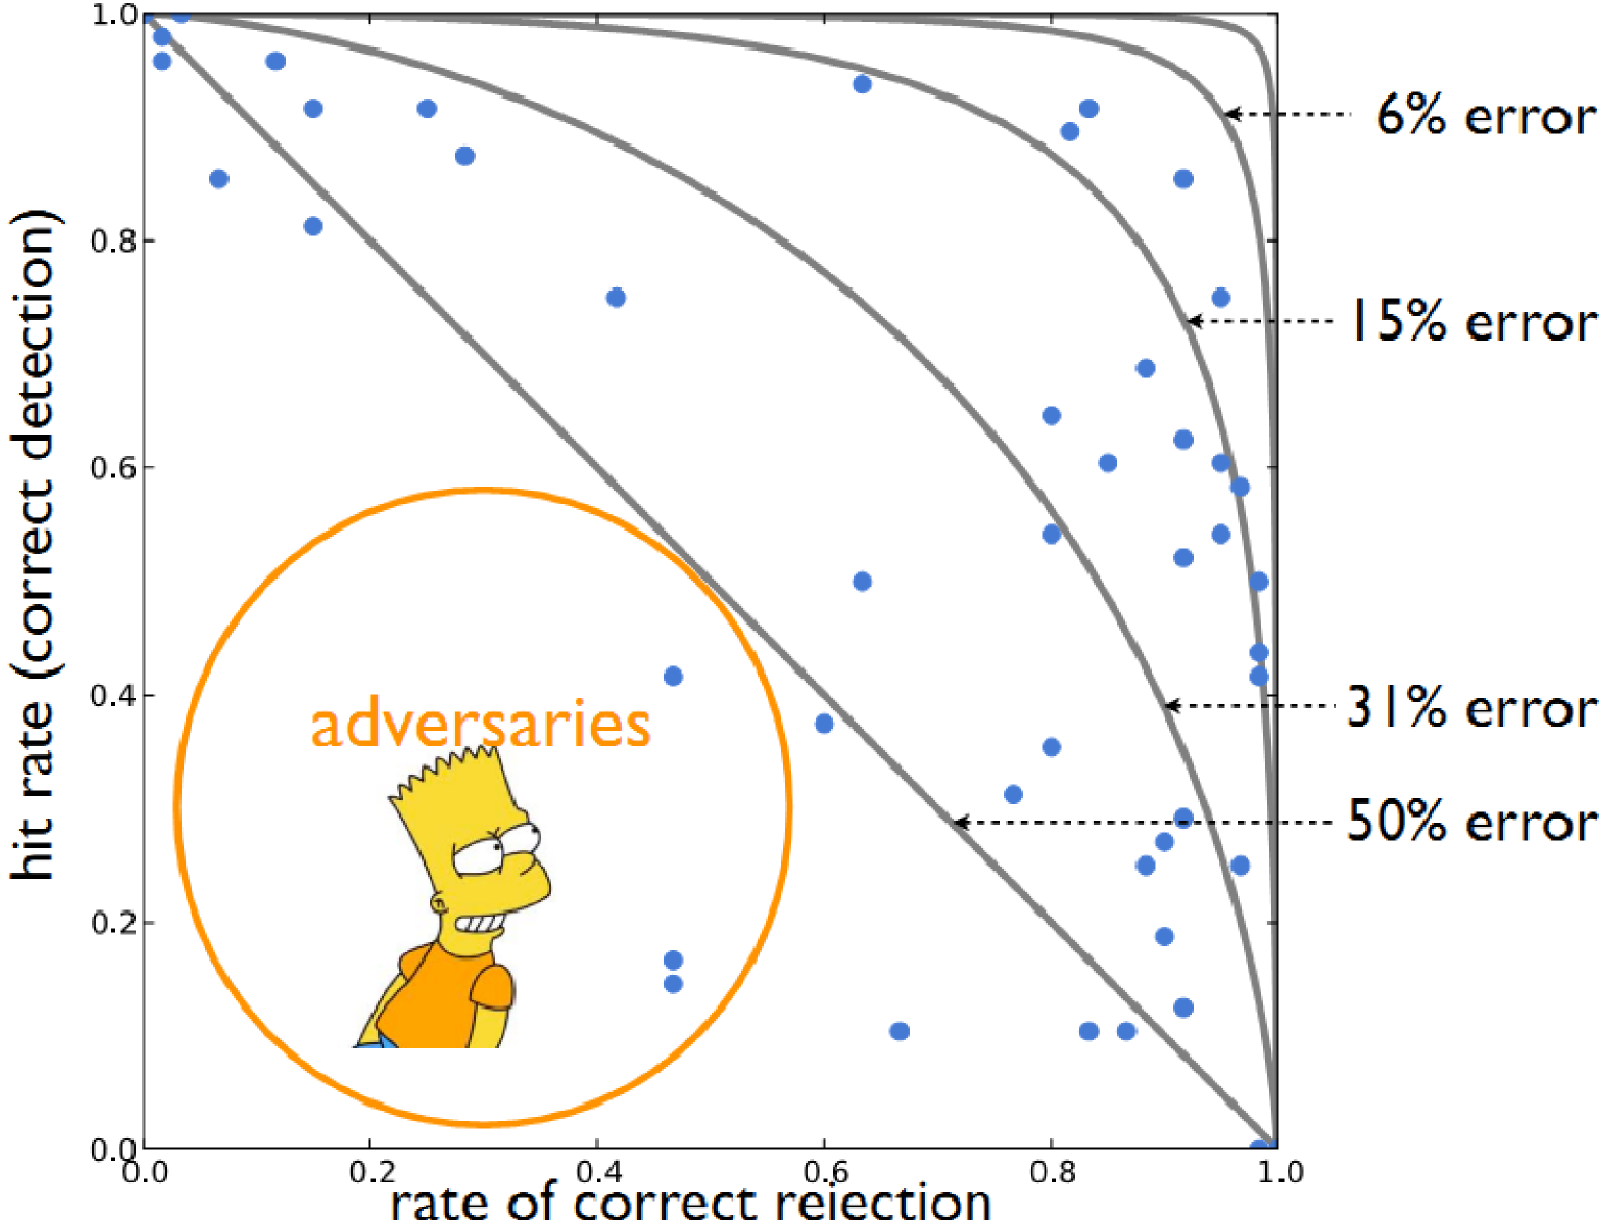

Slide credit: Welinder et al

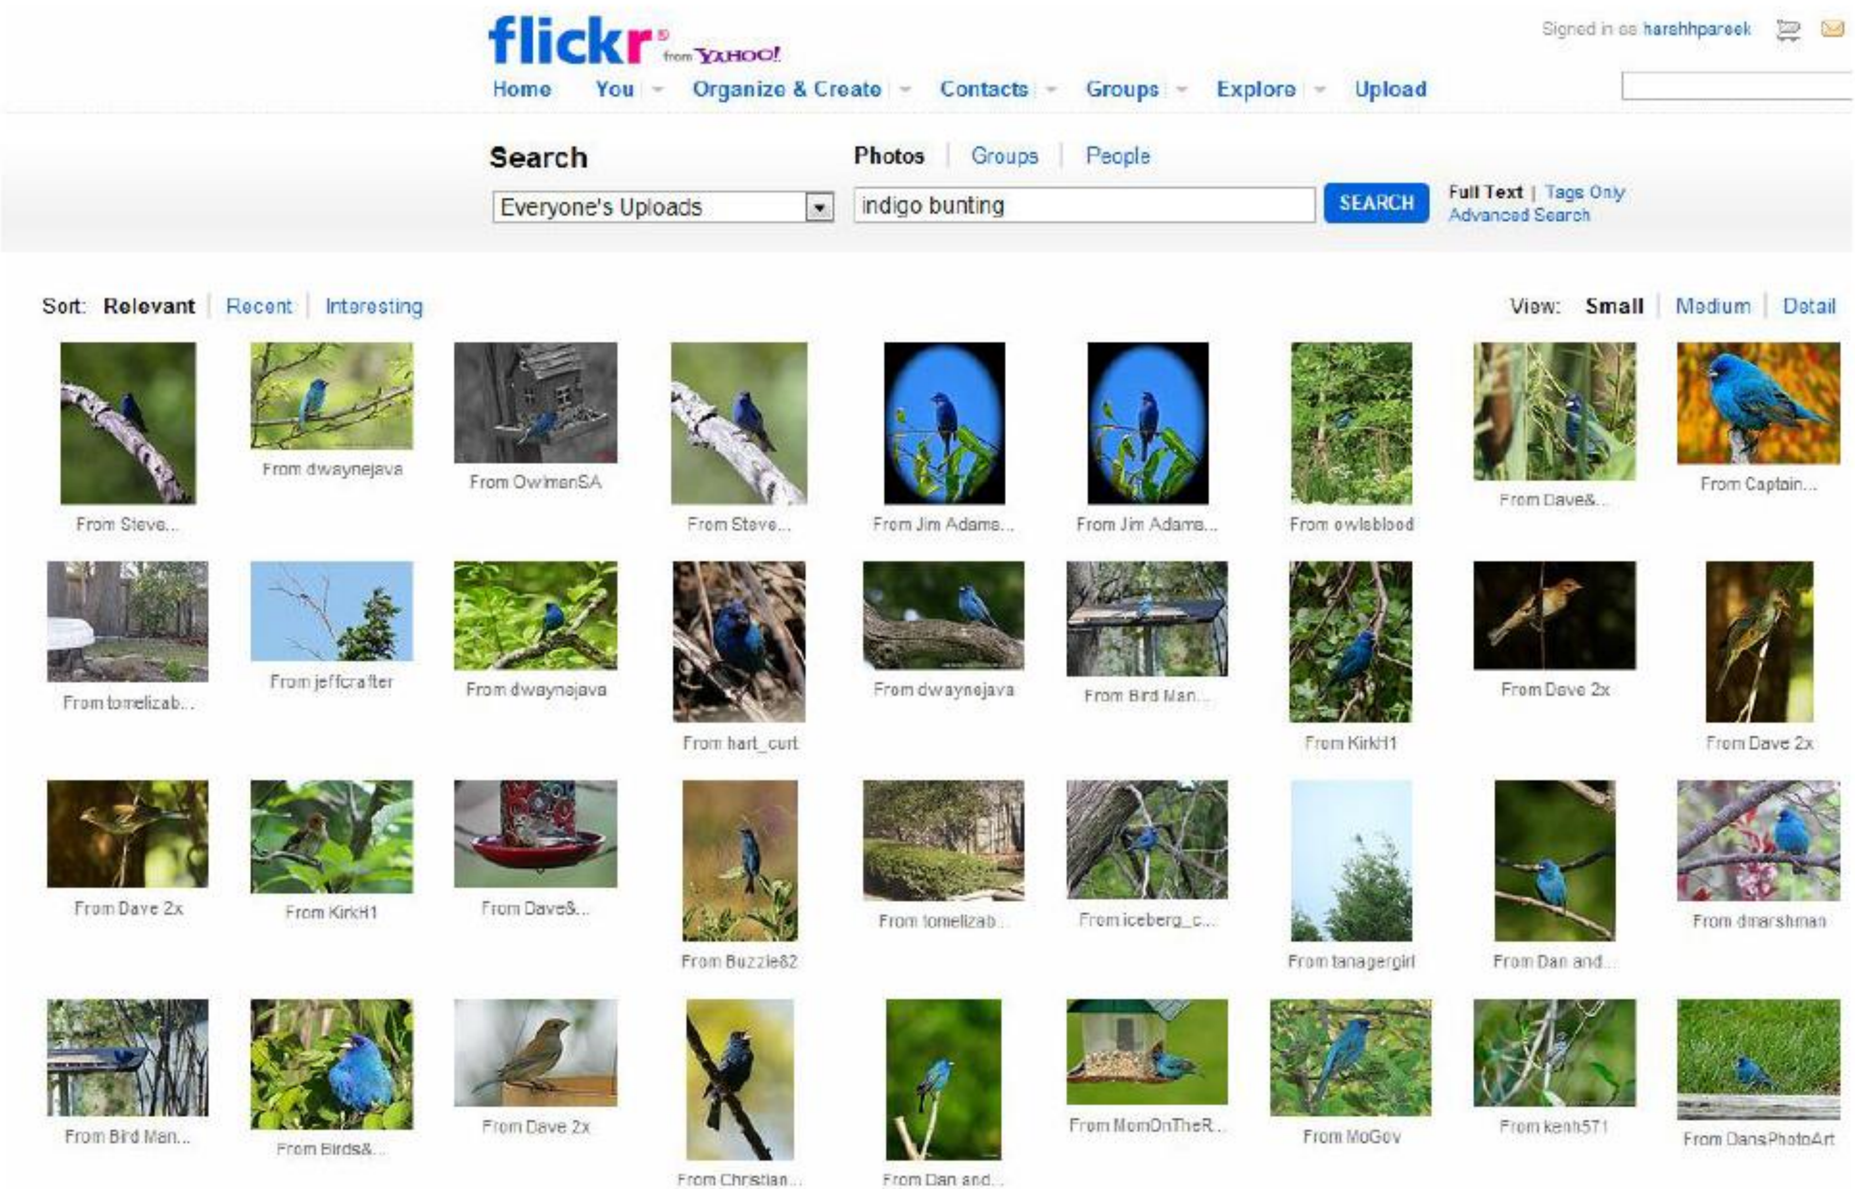

Image credit: Flickr.com