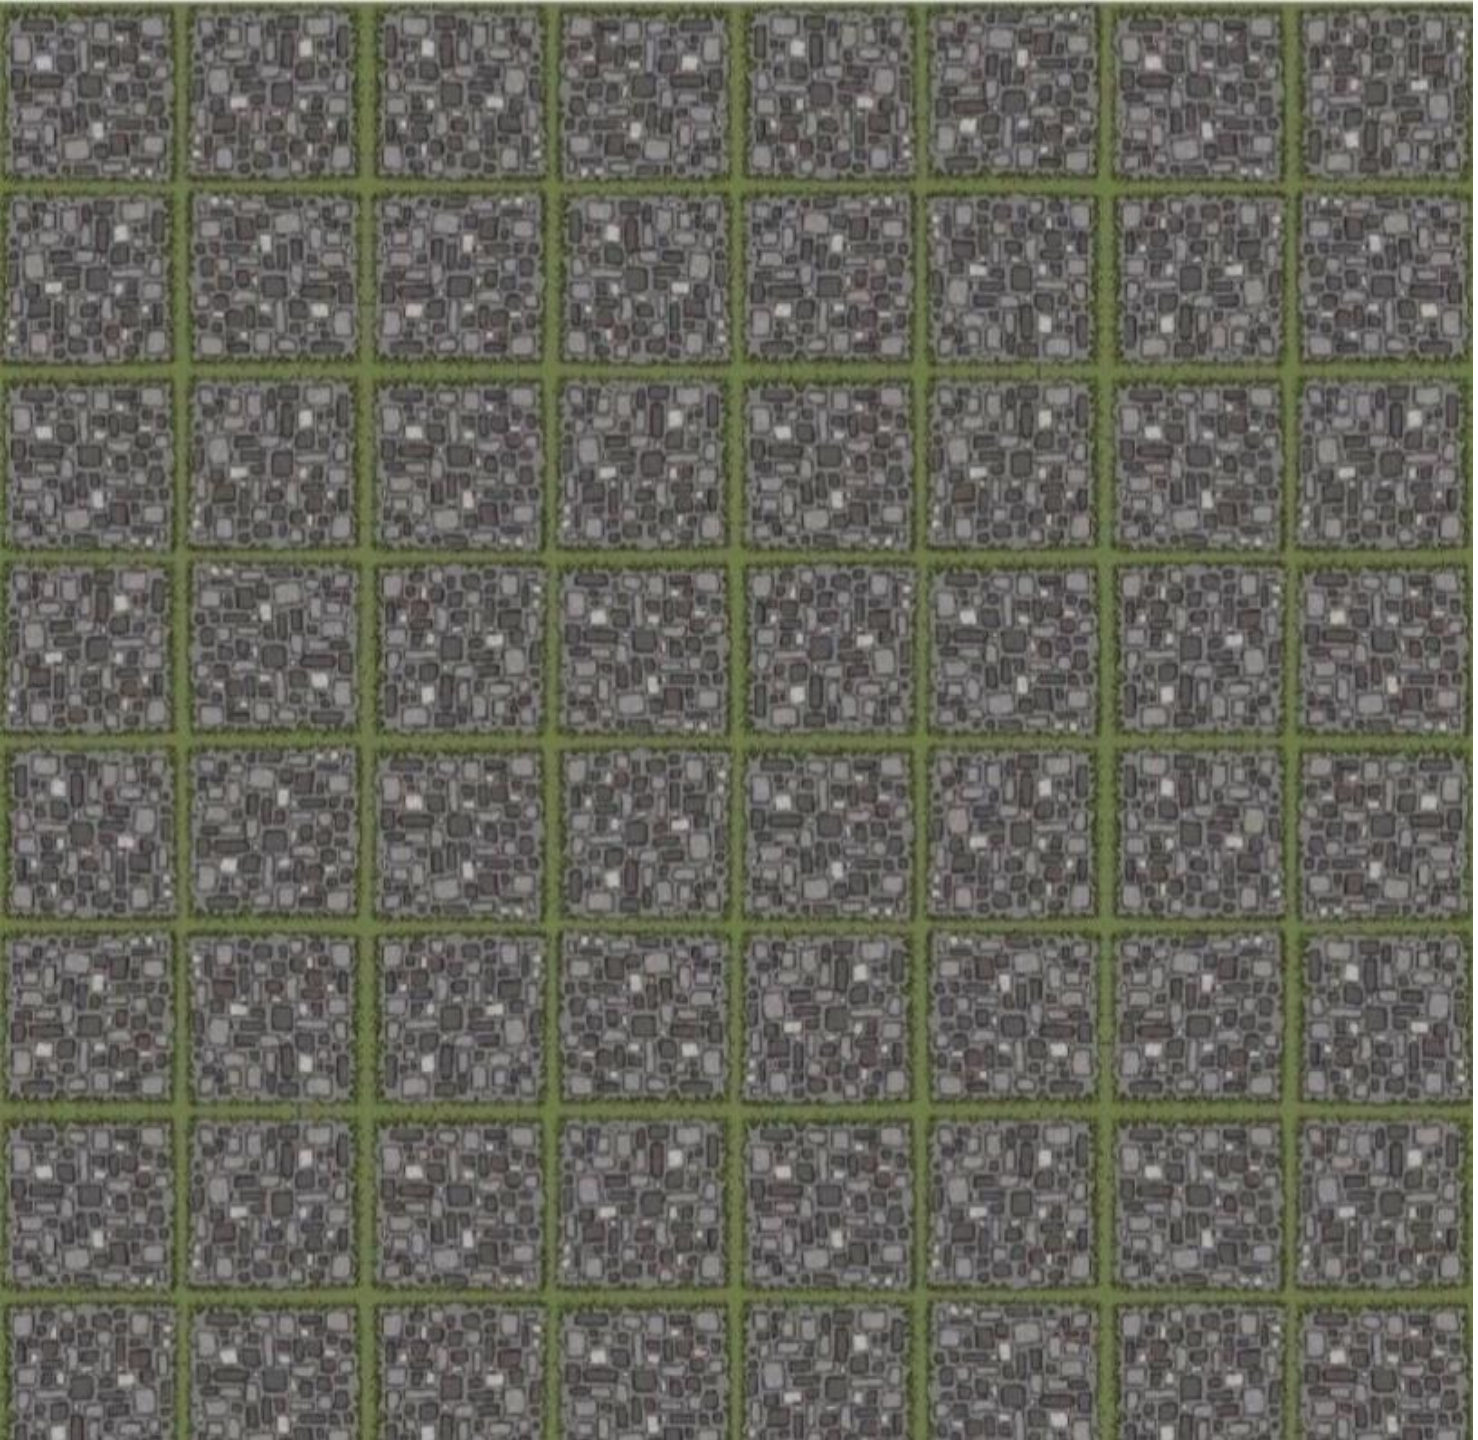

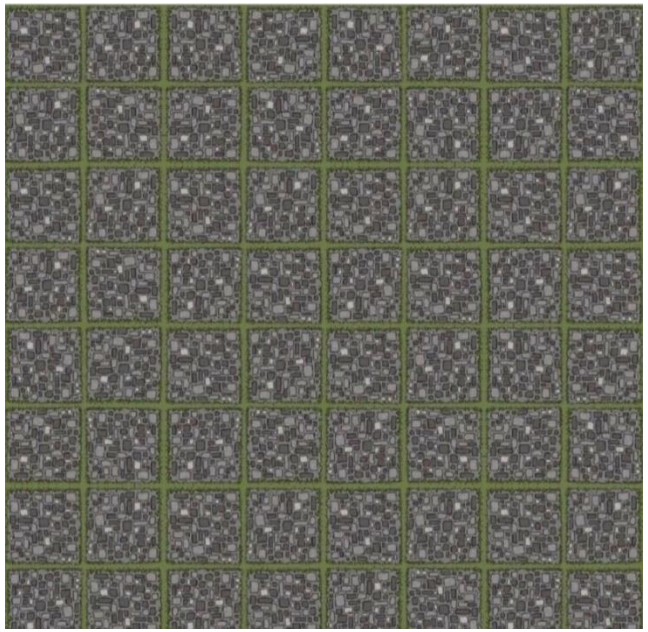

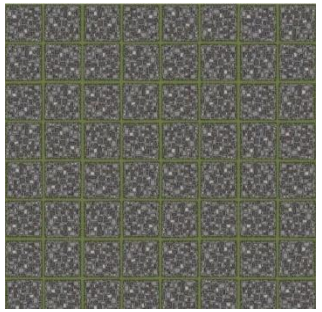

Big Data: Opportunities of Scale

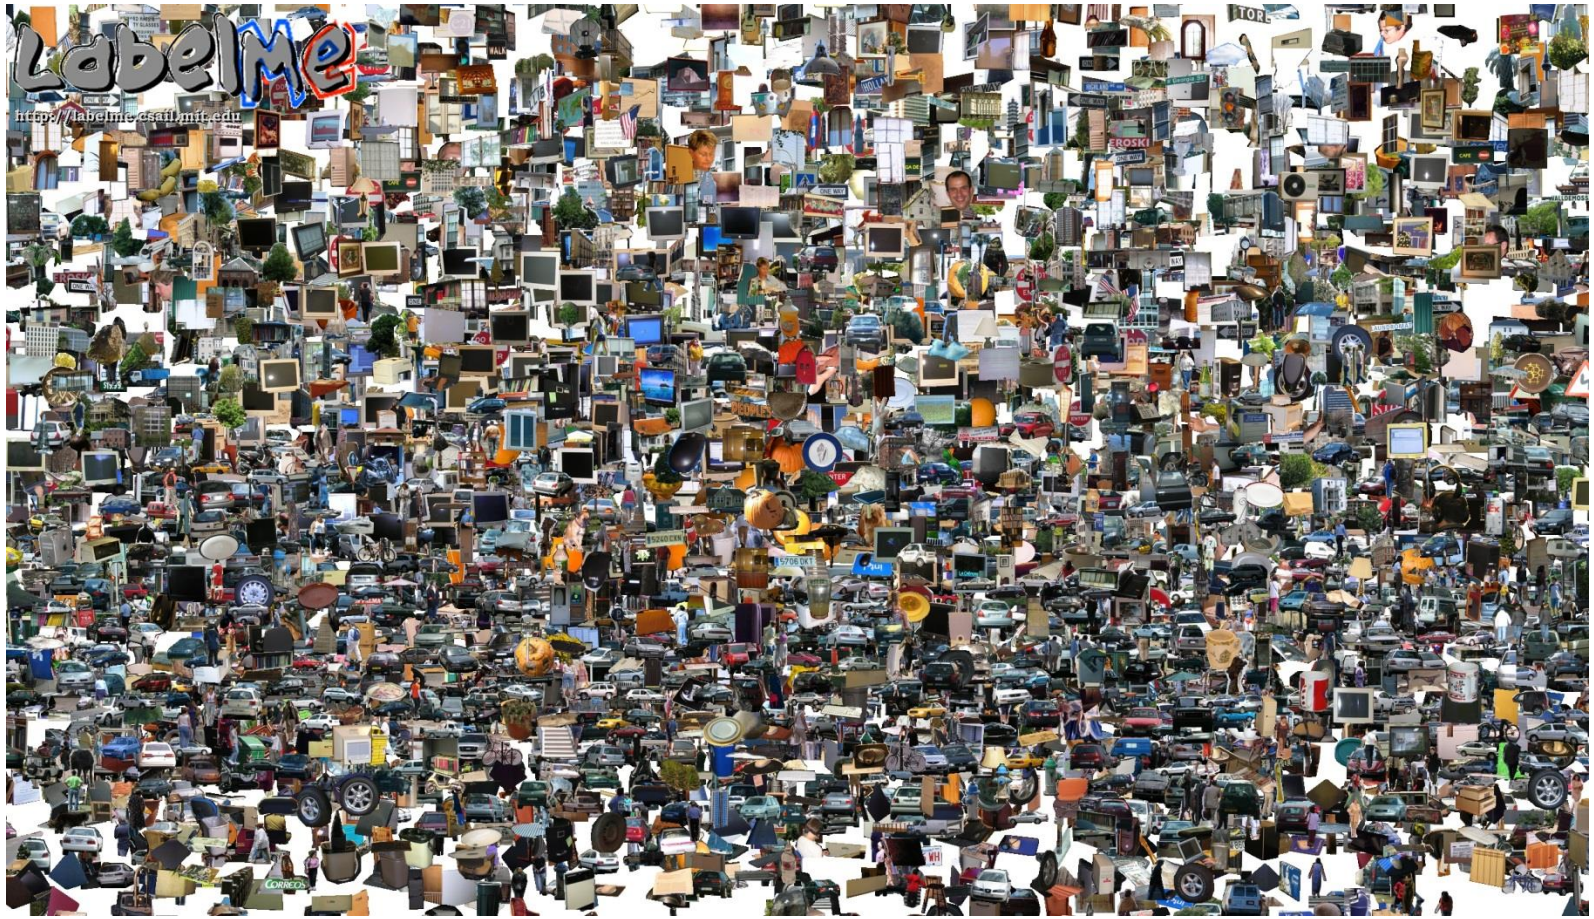

Computer Vision

James Hays

Many slides from James Hays, Alyosha Efros, and Derek Hoiem

Graphic from Antonio Torralba

Outline

Opportunities of Scale: Data-driven methods

- The Unreasonable Effectiveness of Data
- Scene Completion
- Image Geolocation

Computer Vision Class so far

- The geometry of image formation
  - Ancient / Renaissance
- Signal processing / Convolution
  - 1800s, but really the 50's and 60's
- Hand-designed Features for recognition, either instance-level or categorical
  - 1999 (SIFT), 2003 (Video Google), 2005 (Dalal-Triggs), 2006 (spatial pyramid bag of words)
- Learning from Data
  - 1991 (EigenFaces) but late 90's to now especially

What has changed in the last 20 years?

- The Internet
- Crowdsourcing
- Learning representations from the data these sources provide (deep learning)
- The inevitable Moore’s-law-esque increase in compute that allows large scale deep learning

Google and massive data-driven algorithms

A.I. for the postmodern world:

- all questions have already been answered…many times, in many ways
- Google is dumb, the “intelligence” is in the data

The Unreasonable Effectiveness of Data

Peter Norvig
Google

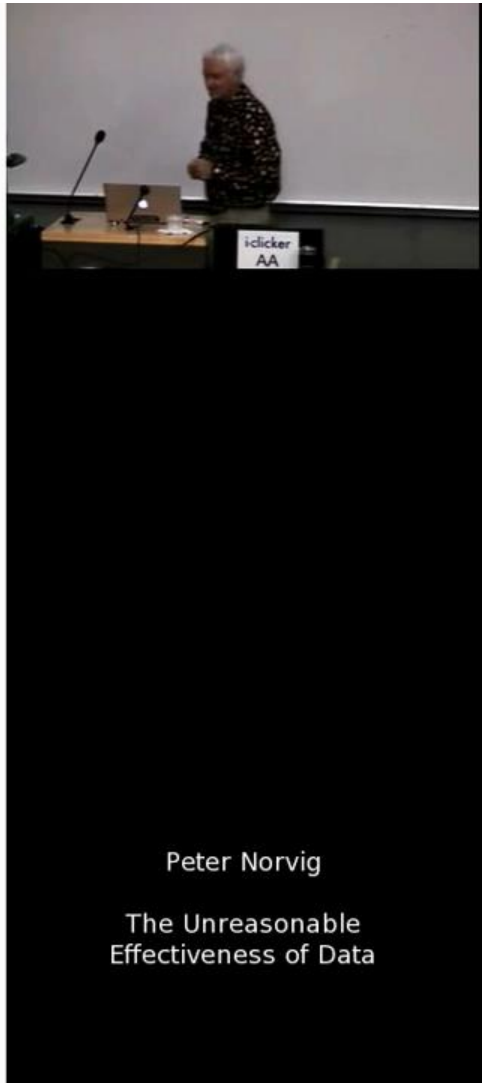

https://youtu.be/yvDCzhbjYWs?t=24
Watch until 9:42

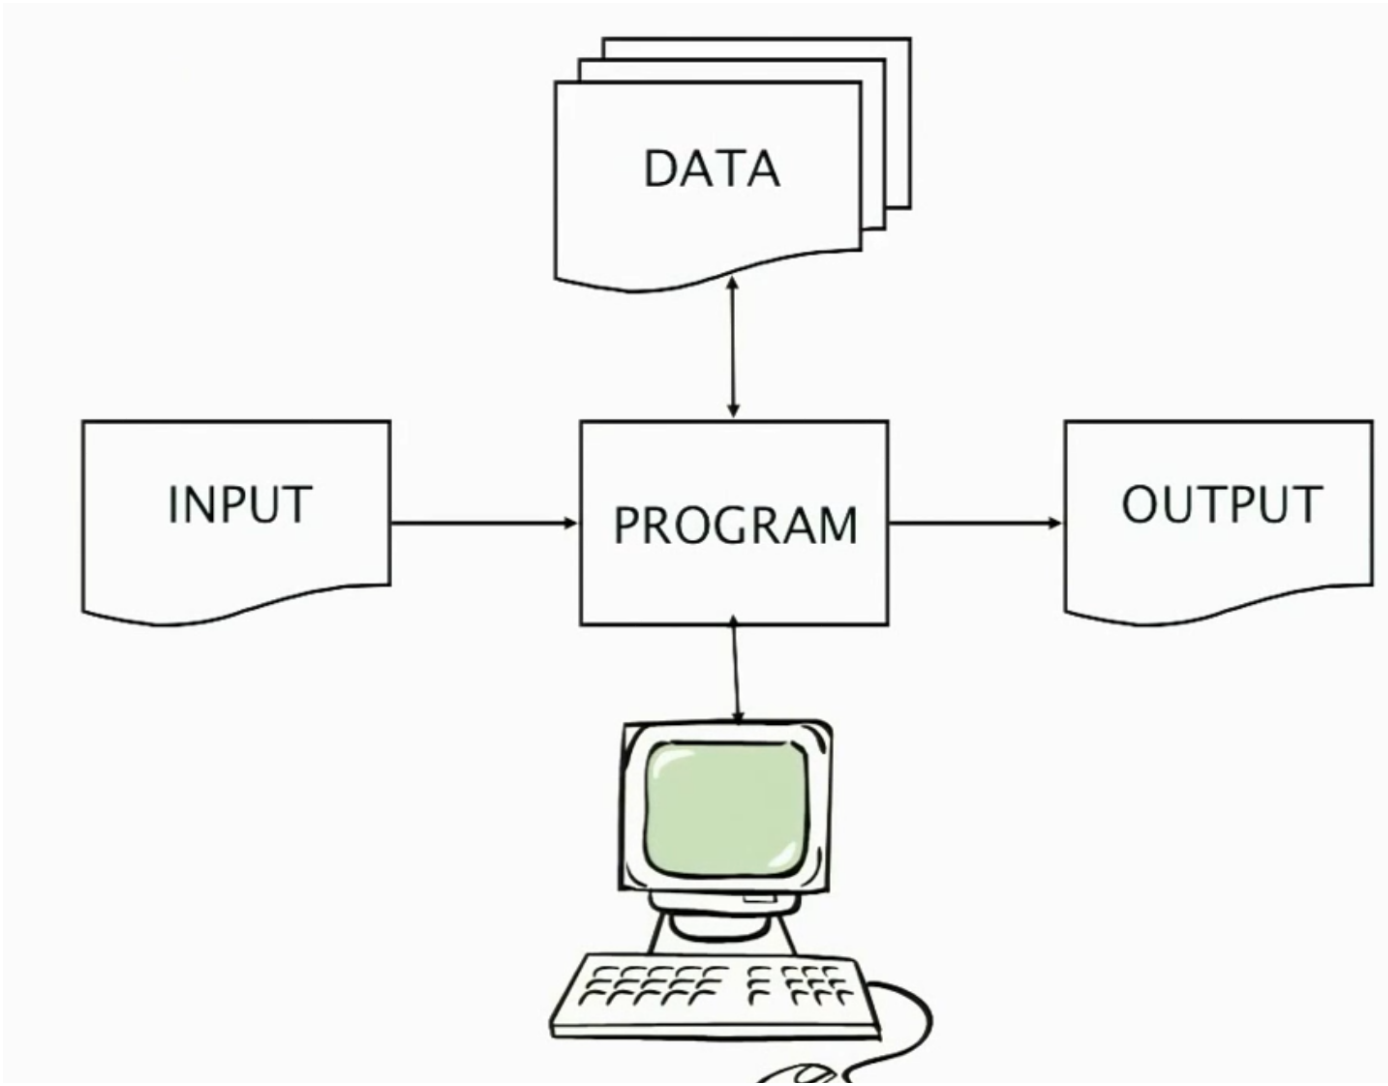
DATA
INPUT
PROGRAM
OUTPUT

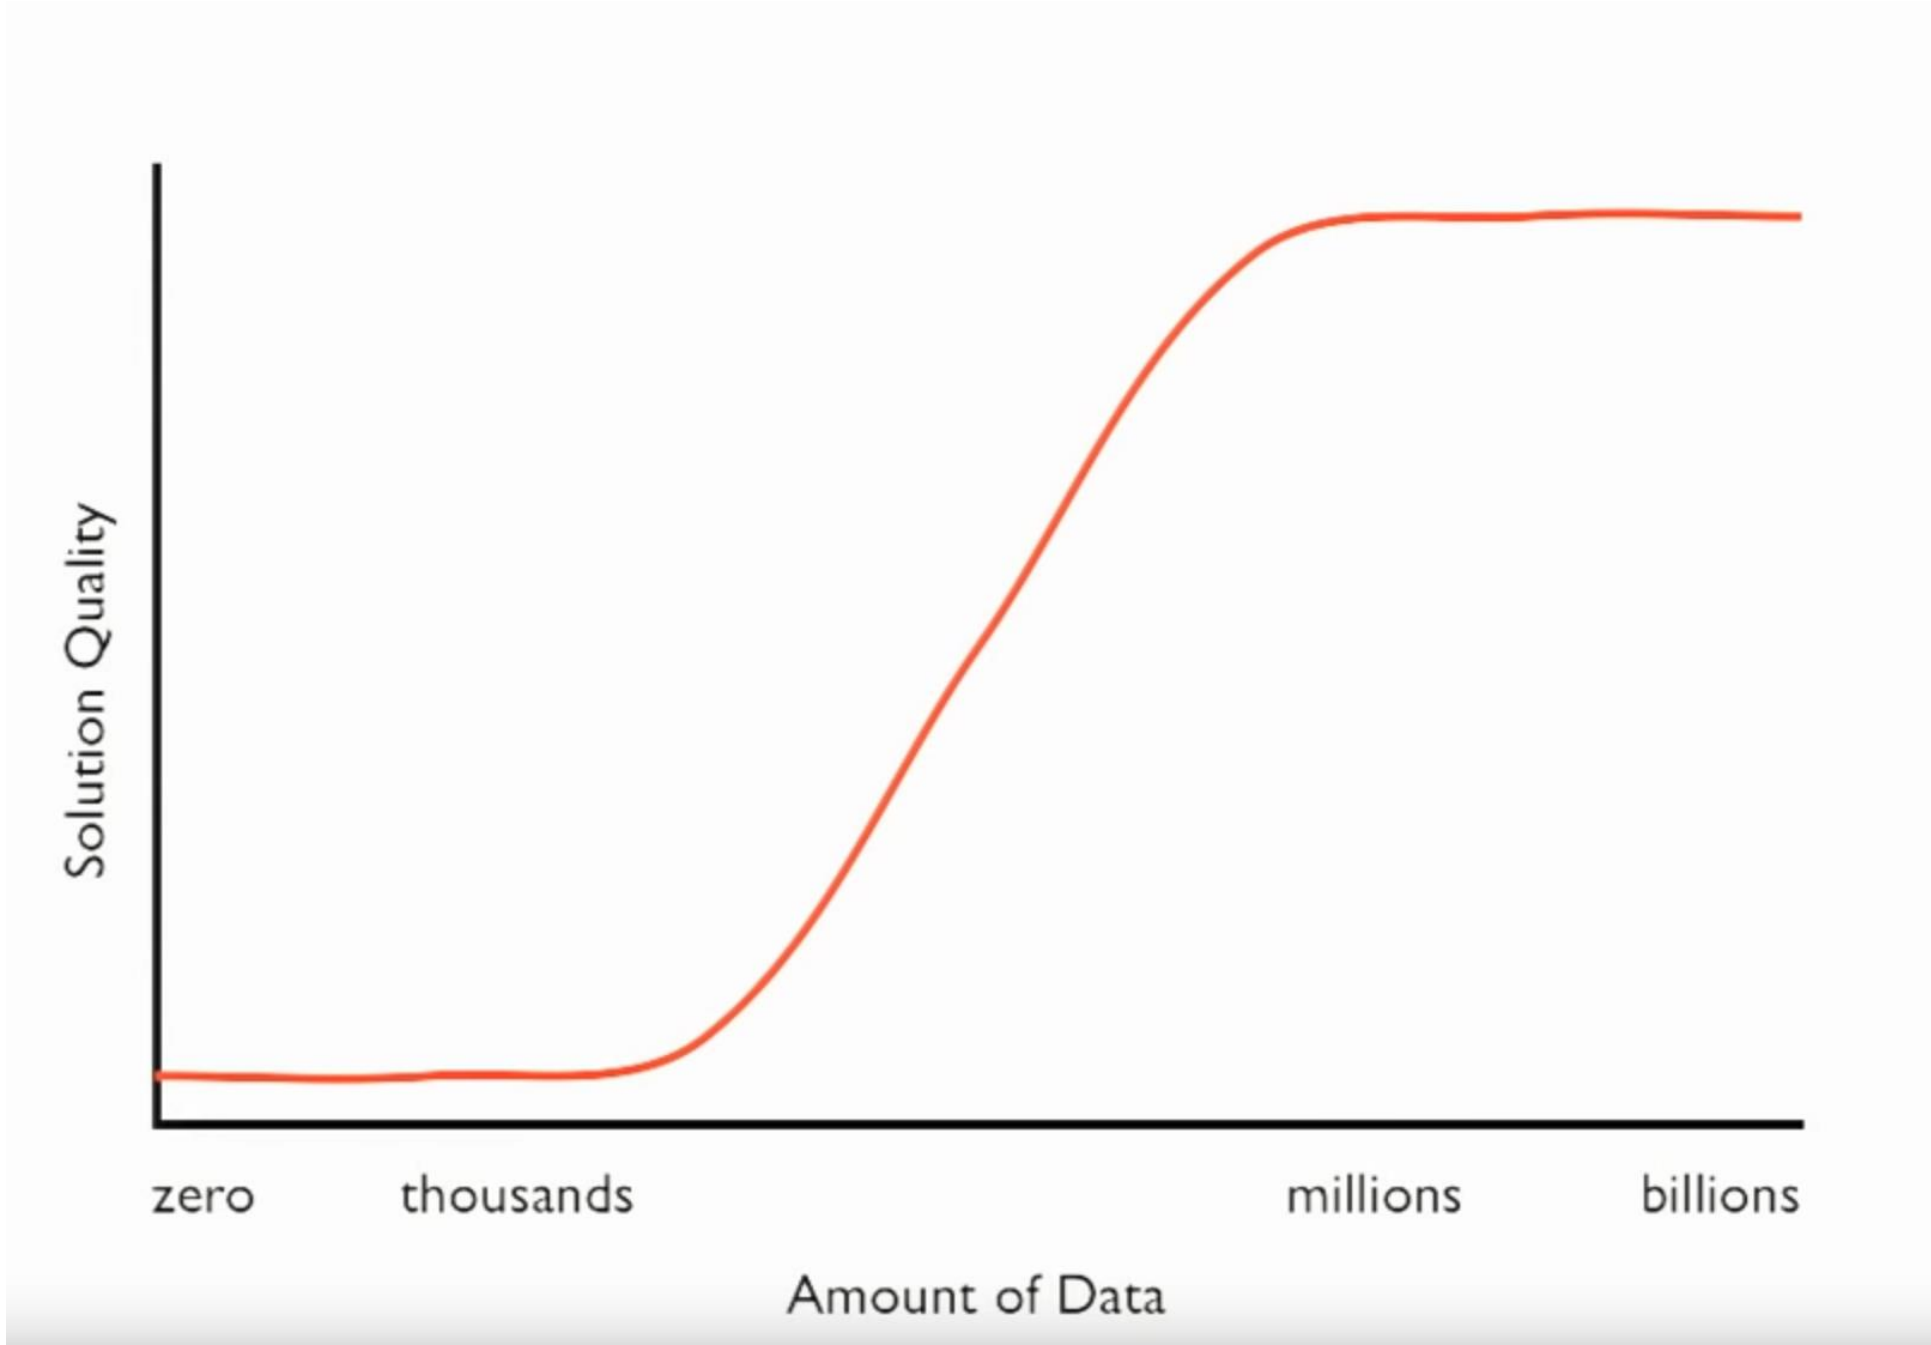
Solution Quality
zero
thousands
millions
billions
Amount of Data

The Unreasonable Effectiveness of Math

- **“**The miracle of the appropriateness of the language of mathematics…” **Eugene Wigner**

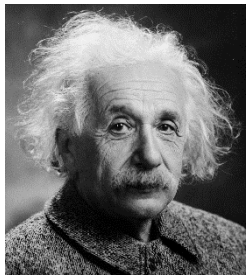

- “The most incomprehensible thing about the universe is that it is comprehensible.” **Albert Einstein**

- “There is only one thing which is more unreasonable than the unreasonable effectiveness of mathematics in physics, and this is the unreasonable ineffectiveness of mathematics in biology.” **Israel Gelfand**

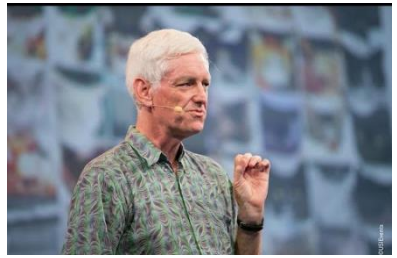

- “We should stop acting as if our goal is to author extremely elegant theories, and instead embrace complexity and make use of the best ally we have: the unreasonable effectiveness of data.” **Peter Norvig**

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

Big Idea

- Do we need computer vision systems to have strong AI-like reasoning about our world?
- What if invariance / generalization isn't actually the core difficulty of computer vision?
- What if we can perform high level reasoning with brute-force, data-driven algorithms?

Scene Completion

[Hays and Efros. Scene Completion Using Millions of Photographs. SIGGRAPH 2007 and CACM October 2008.]

Selected as one of SIGGRAPH's "Seminal papers" in 2023.

http://graphics.cs.cmu.edu/projects/scene-completion/

How it works

- Find a similar image from a large dataset
- Blend a region from that image into the hole

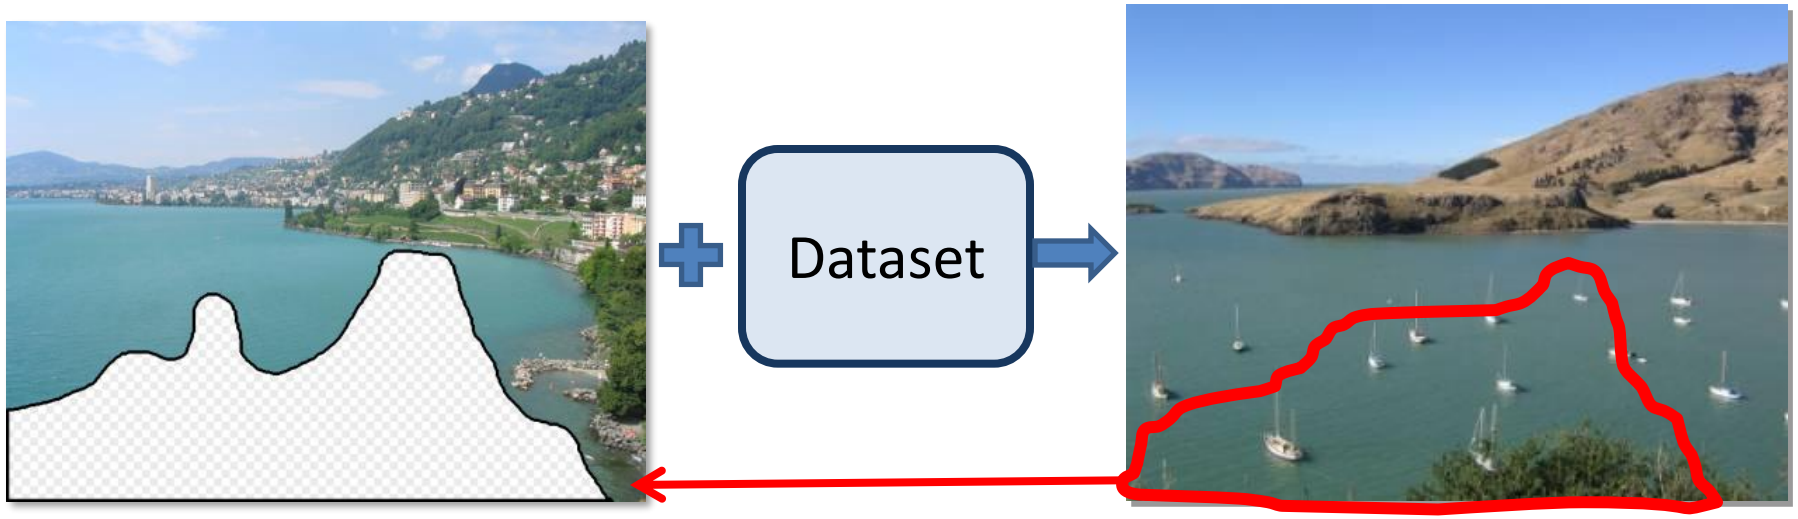

General Principal

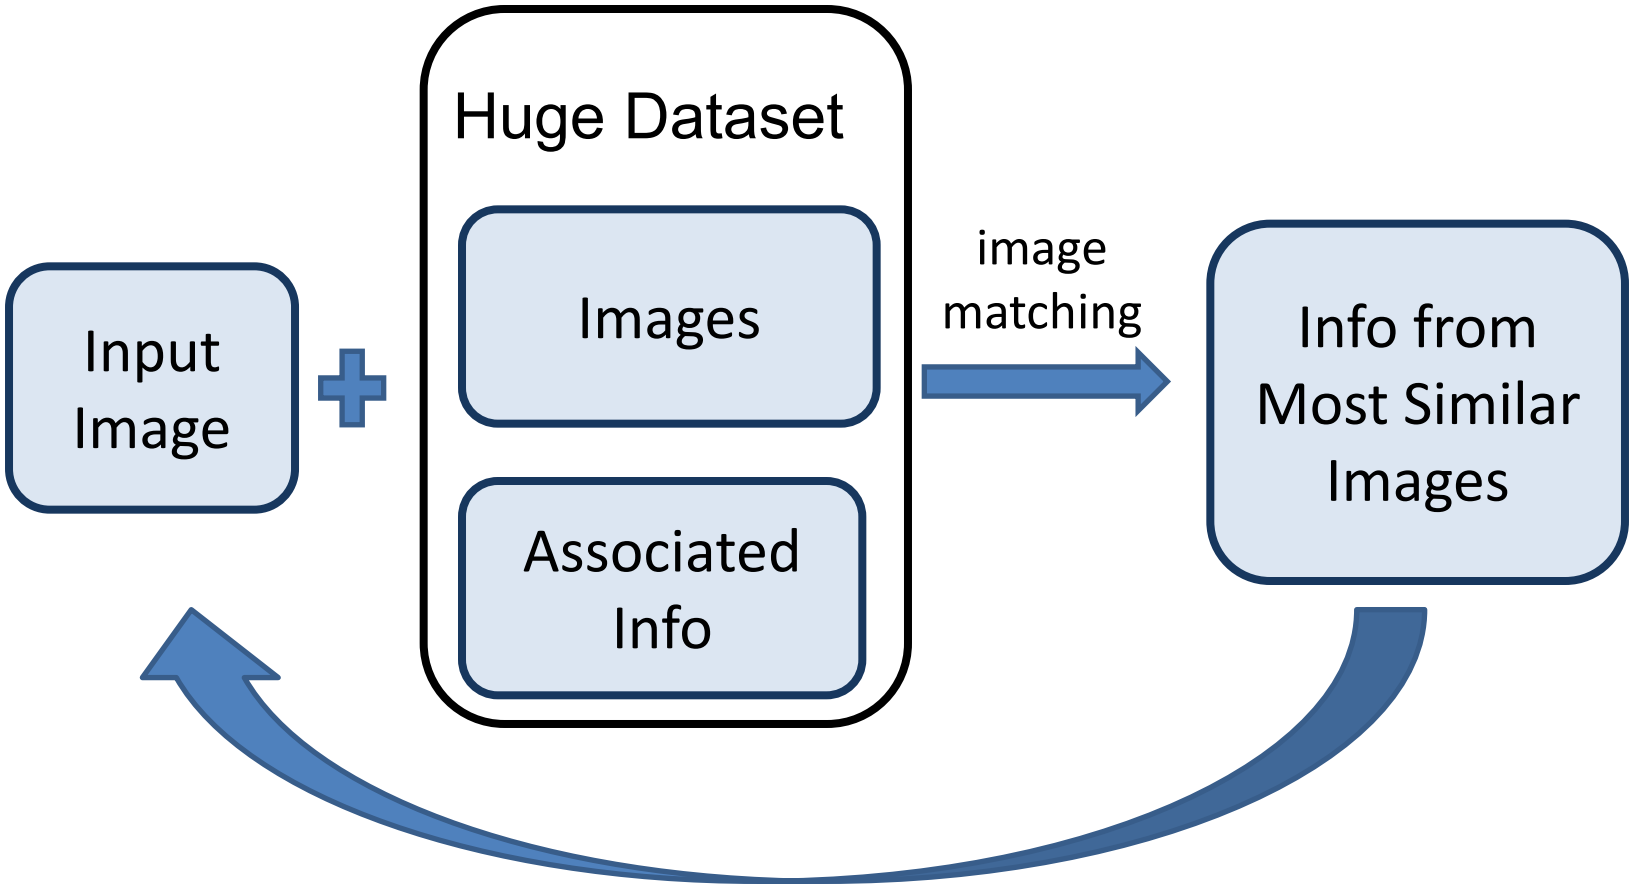

Hopefully, If you have enough images, the dataset will contain very similar images that you can find with simple matching methods.

How many images is enough?

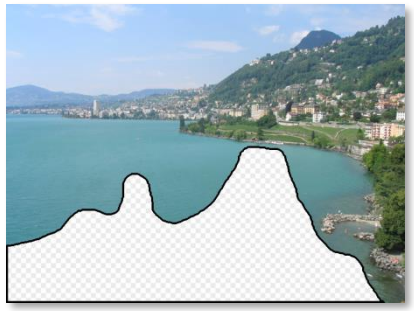

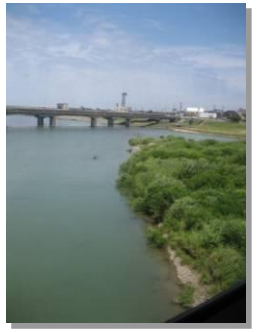




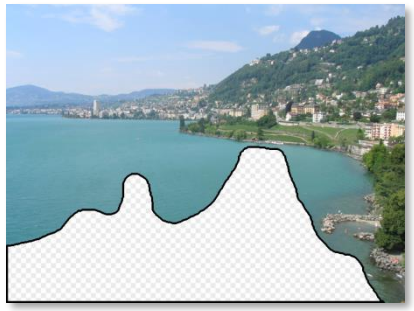




Nearest neighbors from a
collection of 20 thousand images

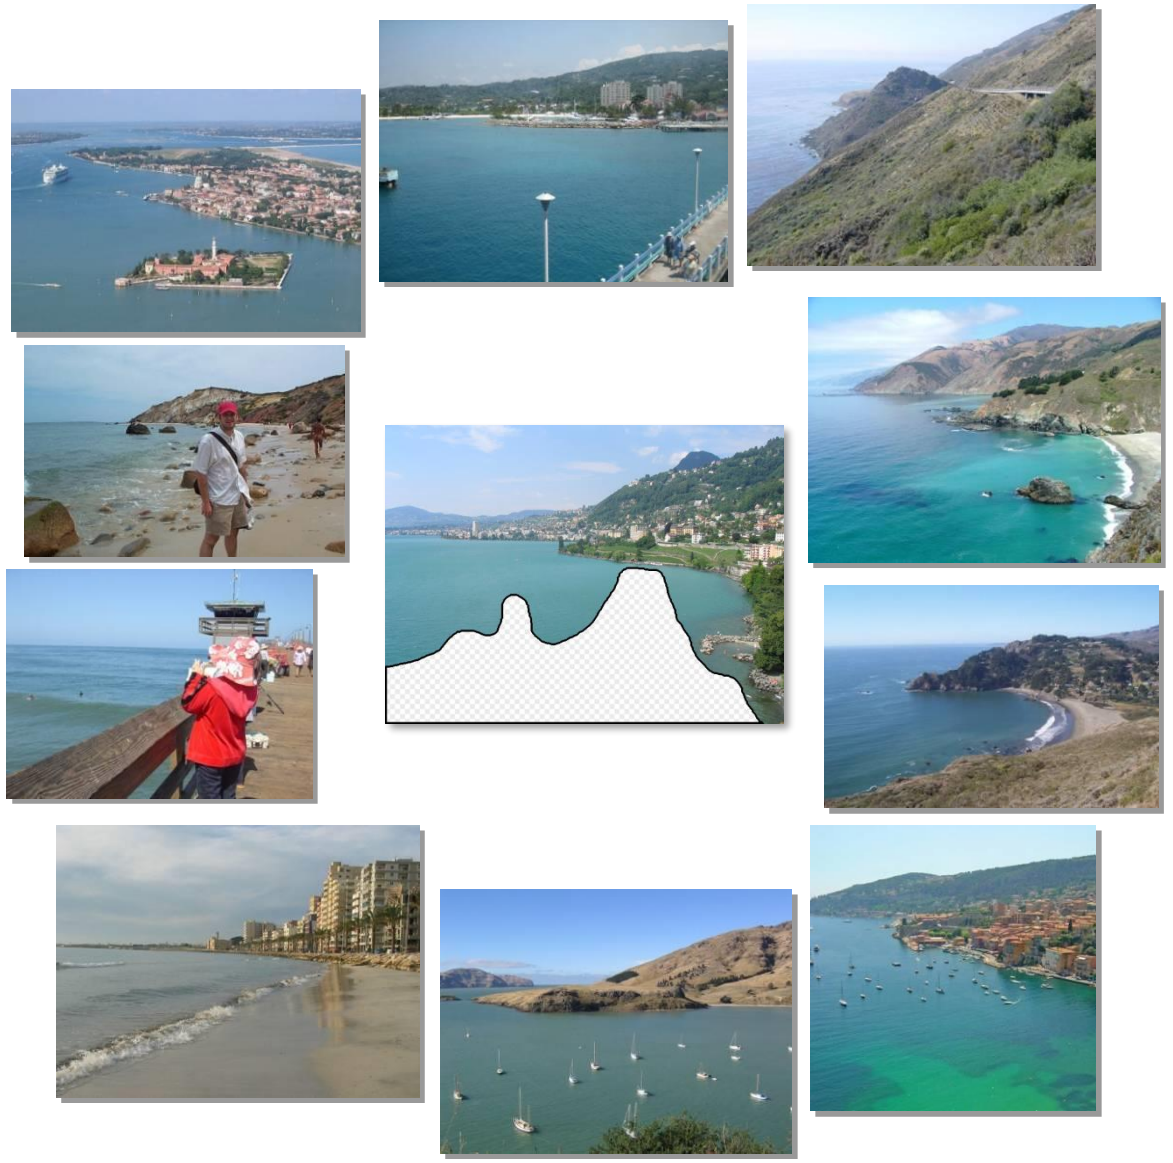

Nearest neighbors from a
collection of 2 million images

Image Data on the Internet

- Flickr (as of Sept. 19th, 2010)
  - 5 billion photographs
  - 100+ million geotagged images
- Facebook (as of 2009)
  - 15 billion

http://royal.pingdom.com/2010/01/22/internet-2009-in-numbers/

Image Data on the Internet

- Flickr (as of Nov 2013)
  - 10 billion photographs
  - 100+ million geotagged images
  - 3.5 million a day
- Facebook (as of Sept 2013)
  - 250 billion+
  - 300 million a day
- Instagram
  - 55 million a day

Scene Completion: how it works

[Hays and Efros. Scene Completion Using Millions of Photographs. SIGGRAPH 2007 and CACM October 2008.]

The Algorithm

Scene Matching

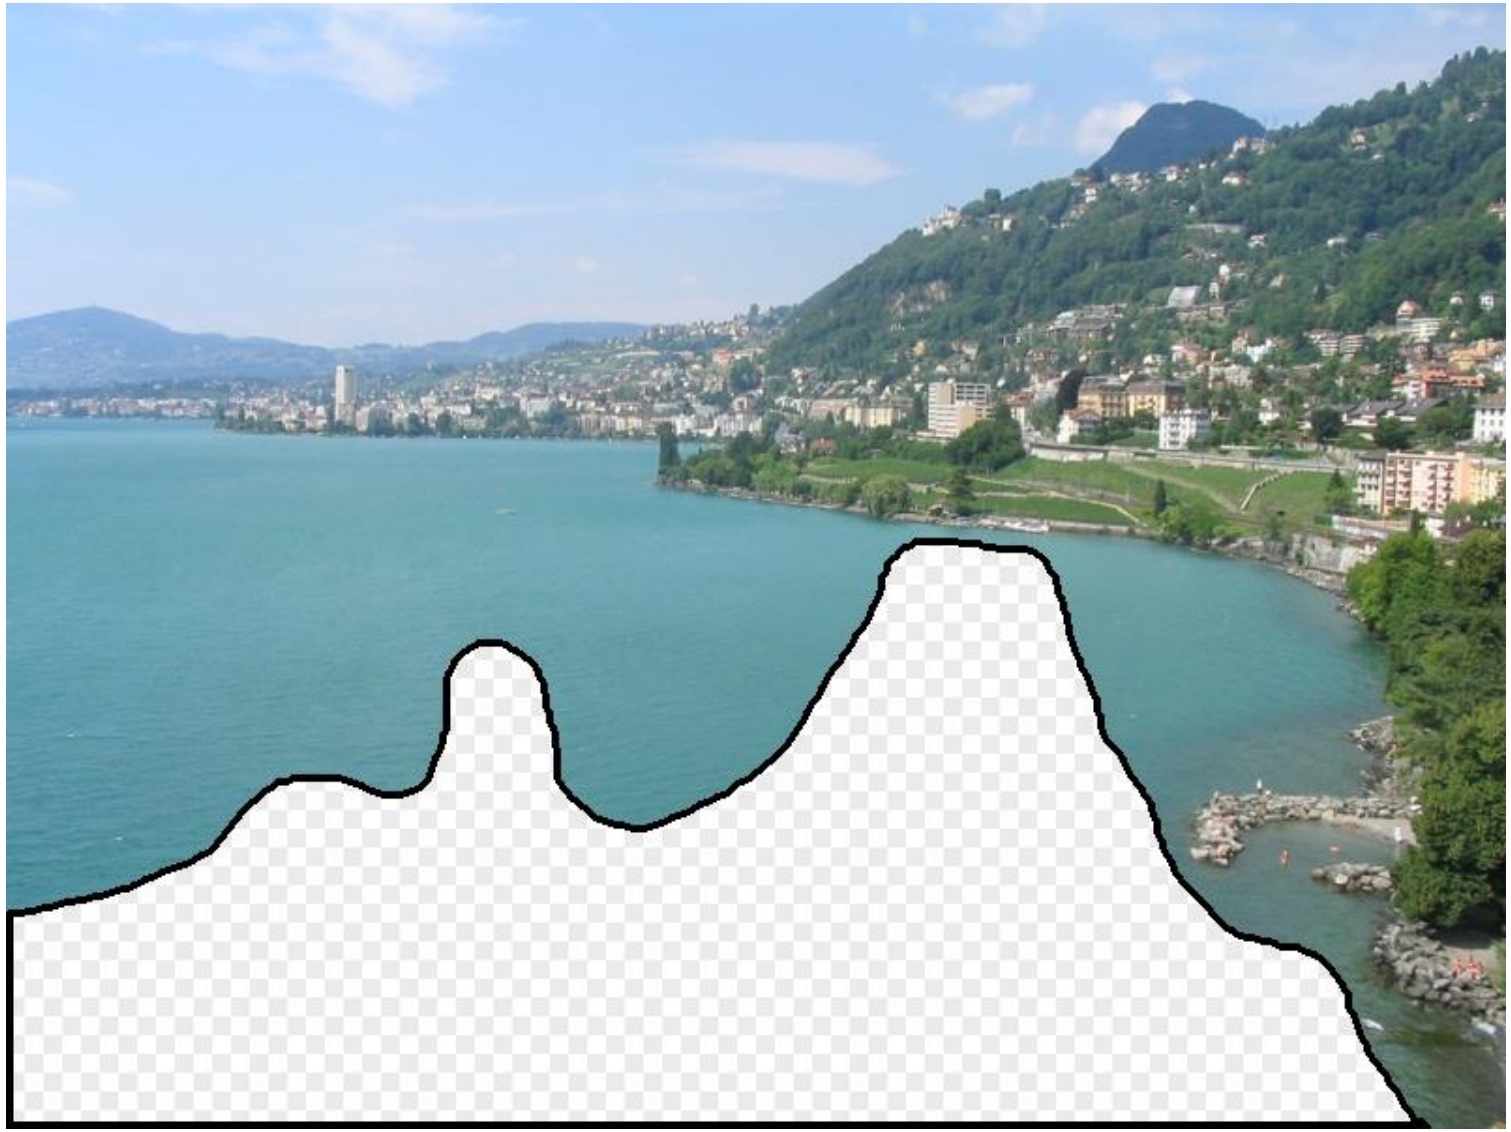

Scene Descriptor

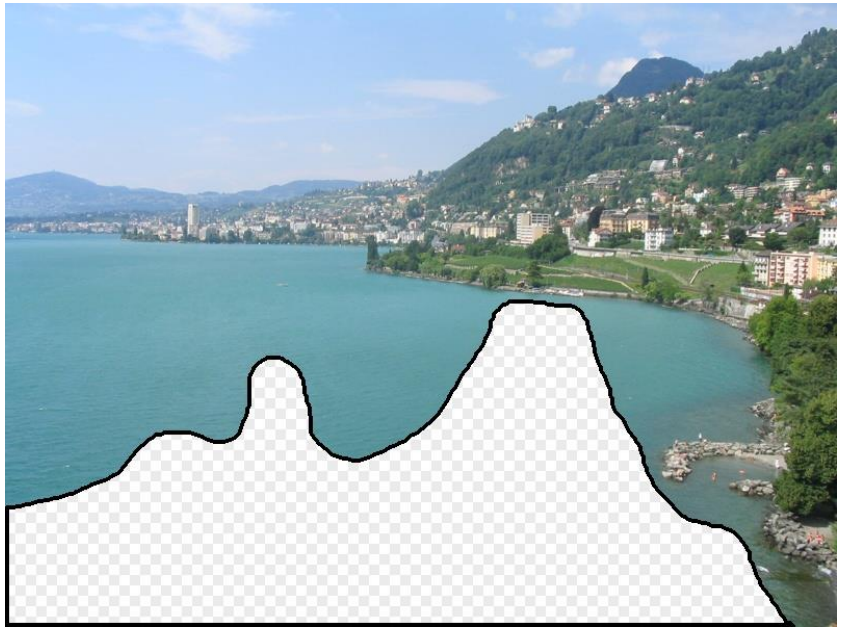

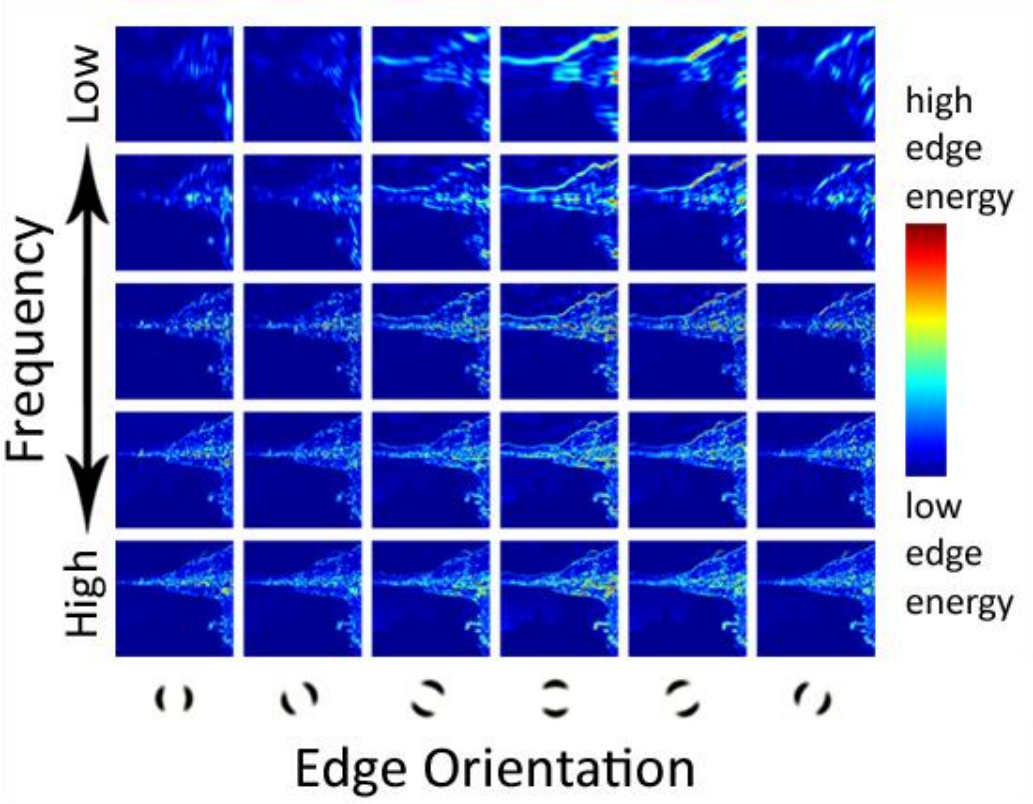

Scene Descriptor

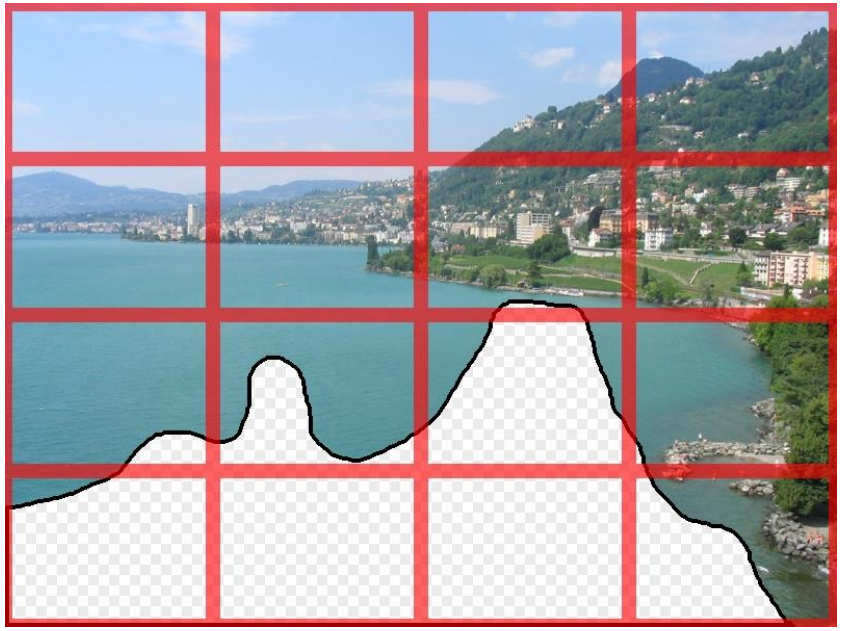

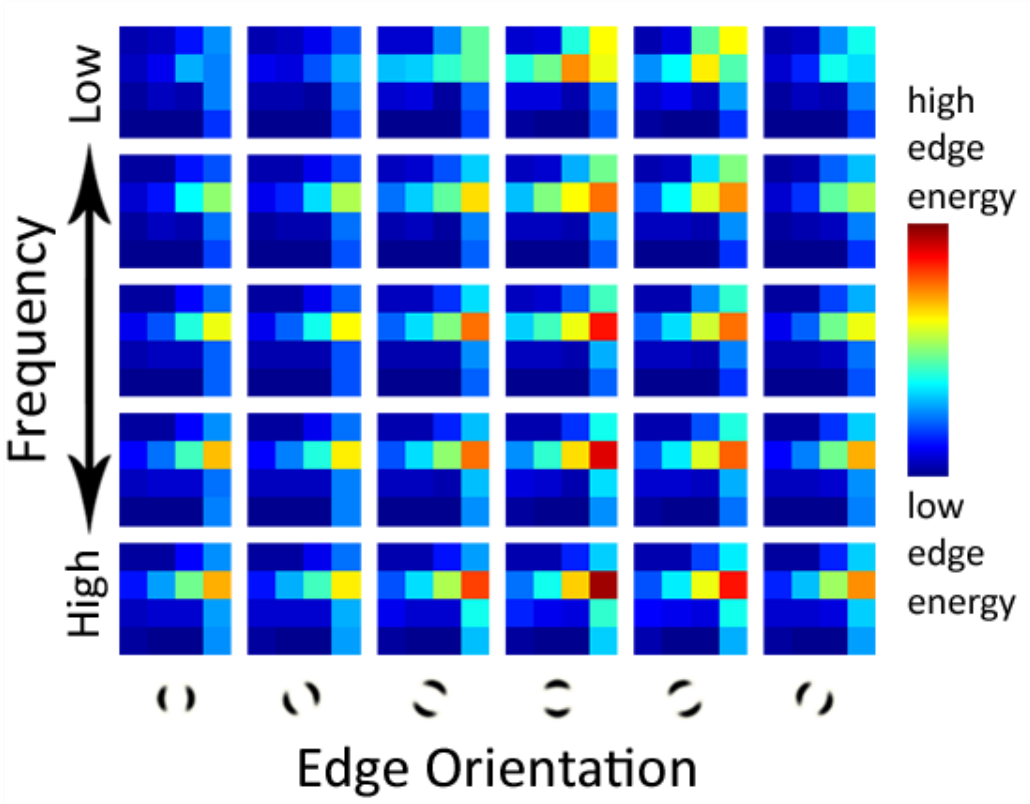

Scene Gist Descriptor
(Oliva and Torralba 2001)

Scene Descriptor

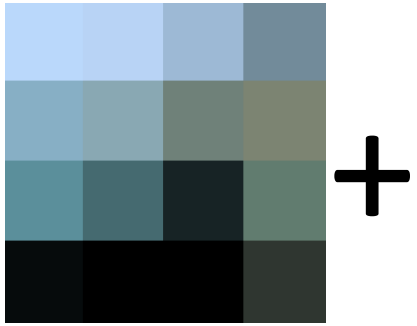

+

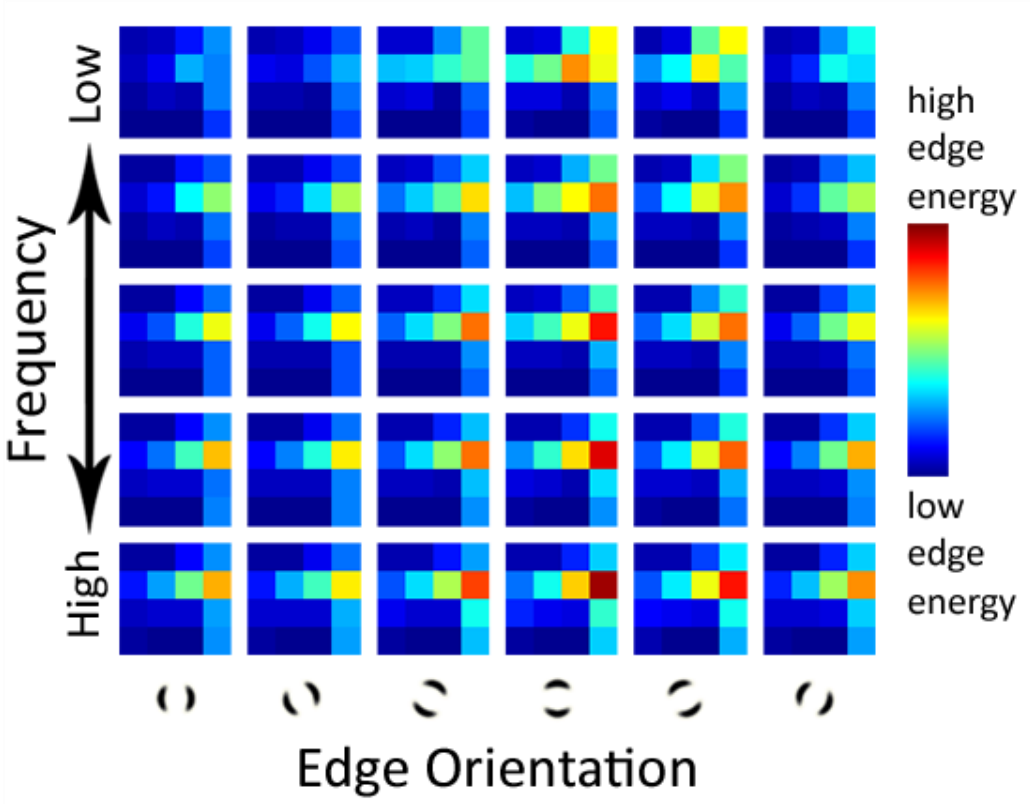

Scene Gist Descriptor
(Oliva and Torralba 2001)


2 Million Flickr Images
22,500 thumbnails

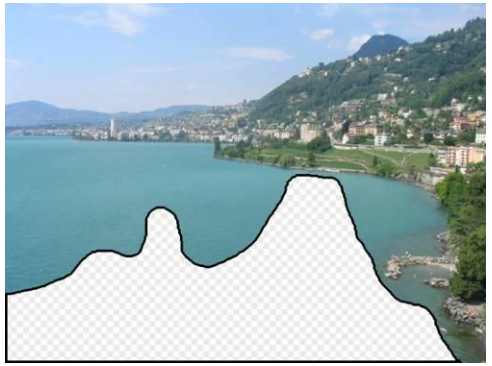

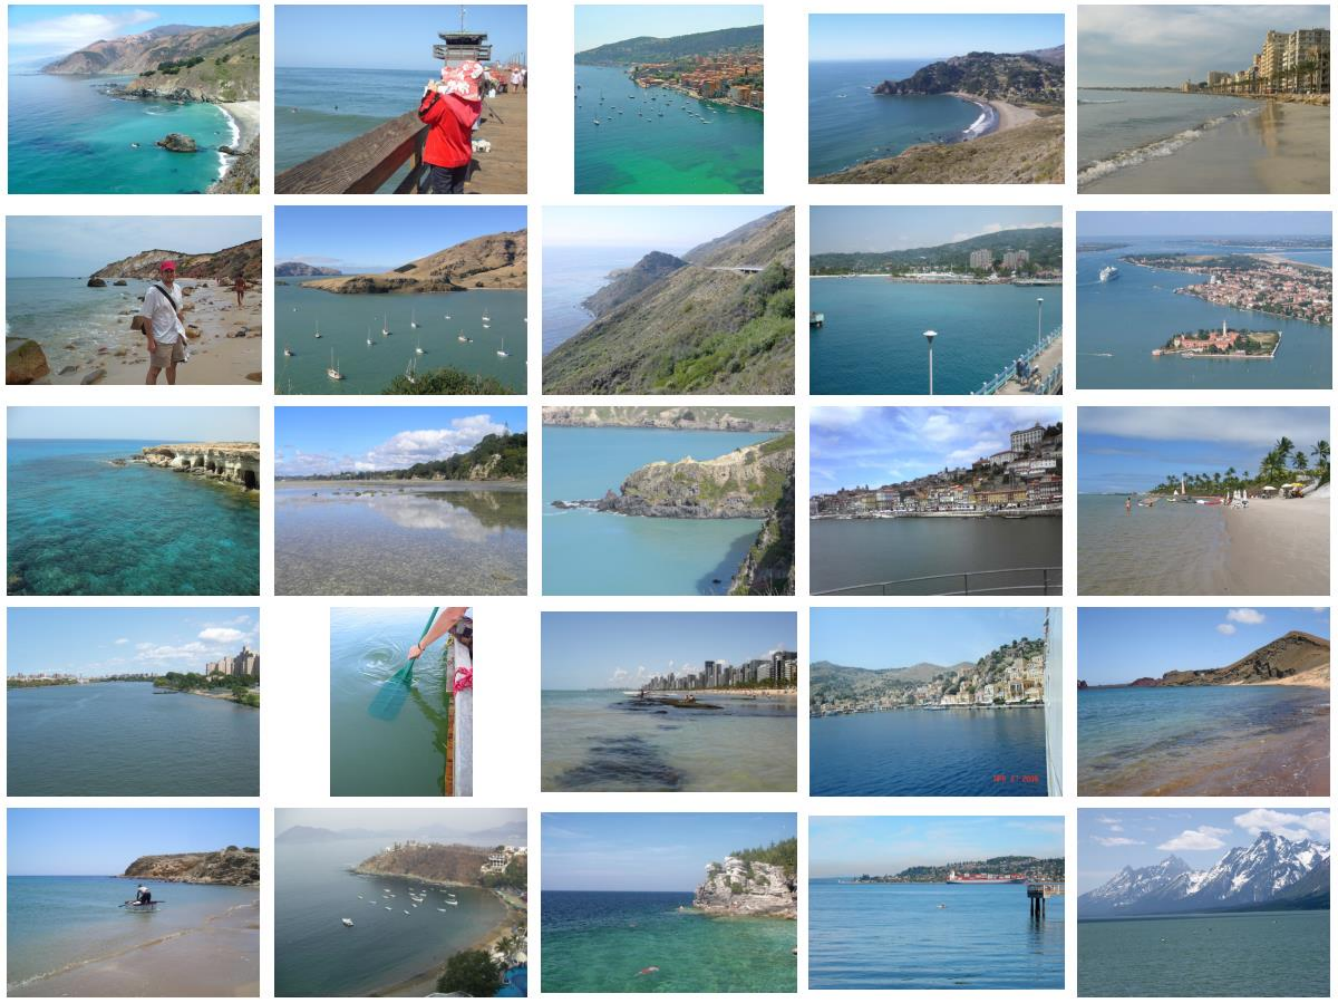

… 200 total

Context Matching

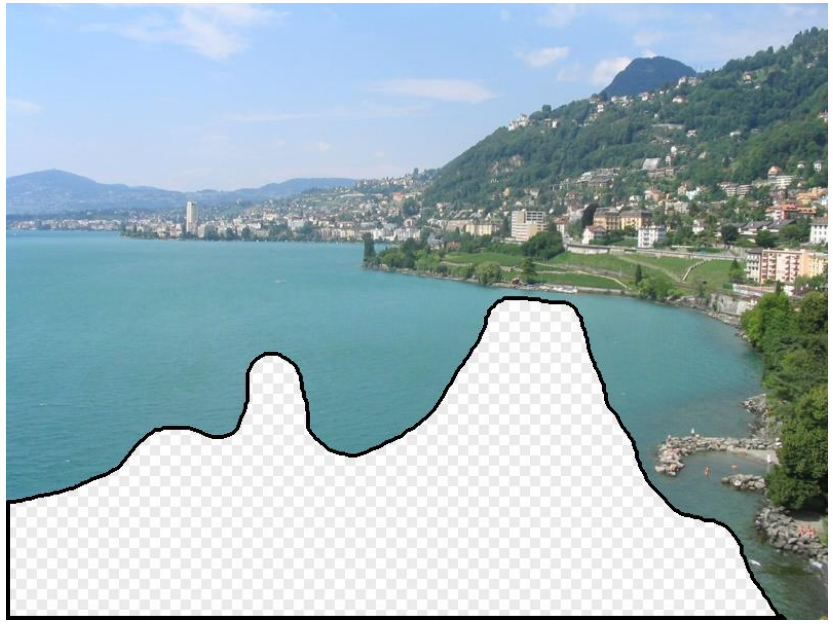


Graph cut + Poisson blending

Result Ranking

We assign each of the 200 results a score which is the sum of:

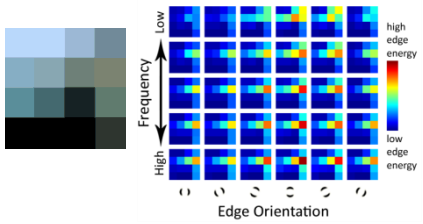

The scene matching distance

The context matching distance (color + texture)

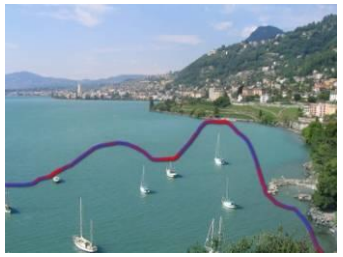

The graph cut cost

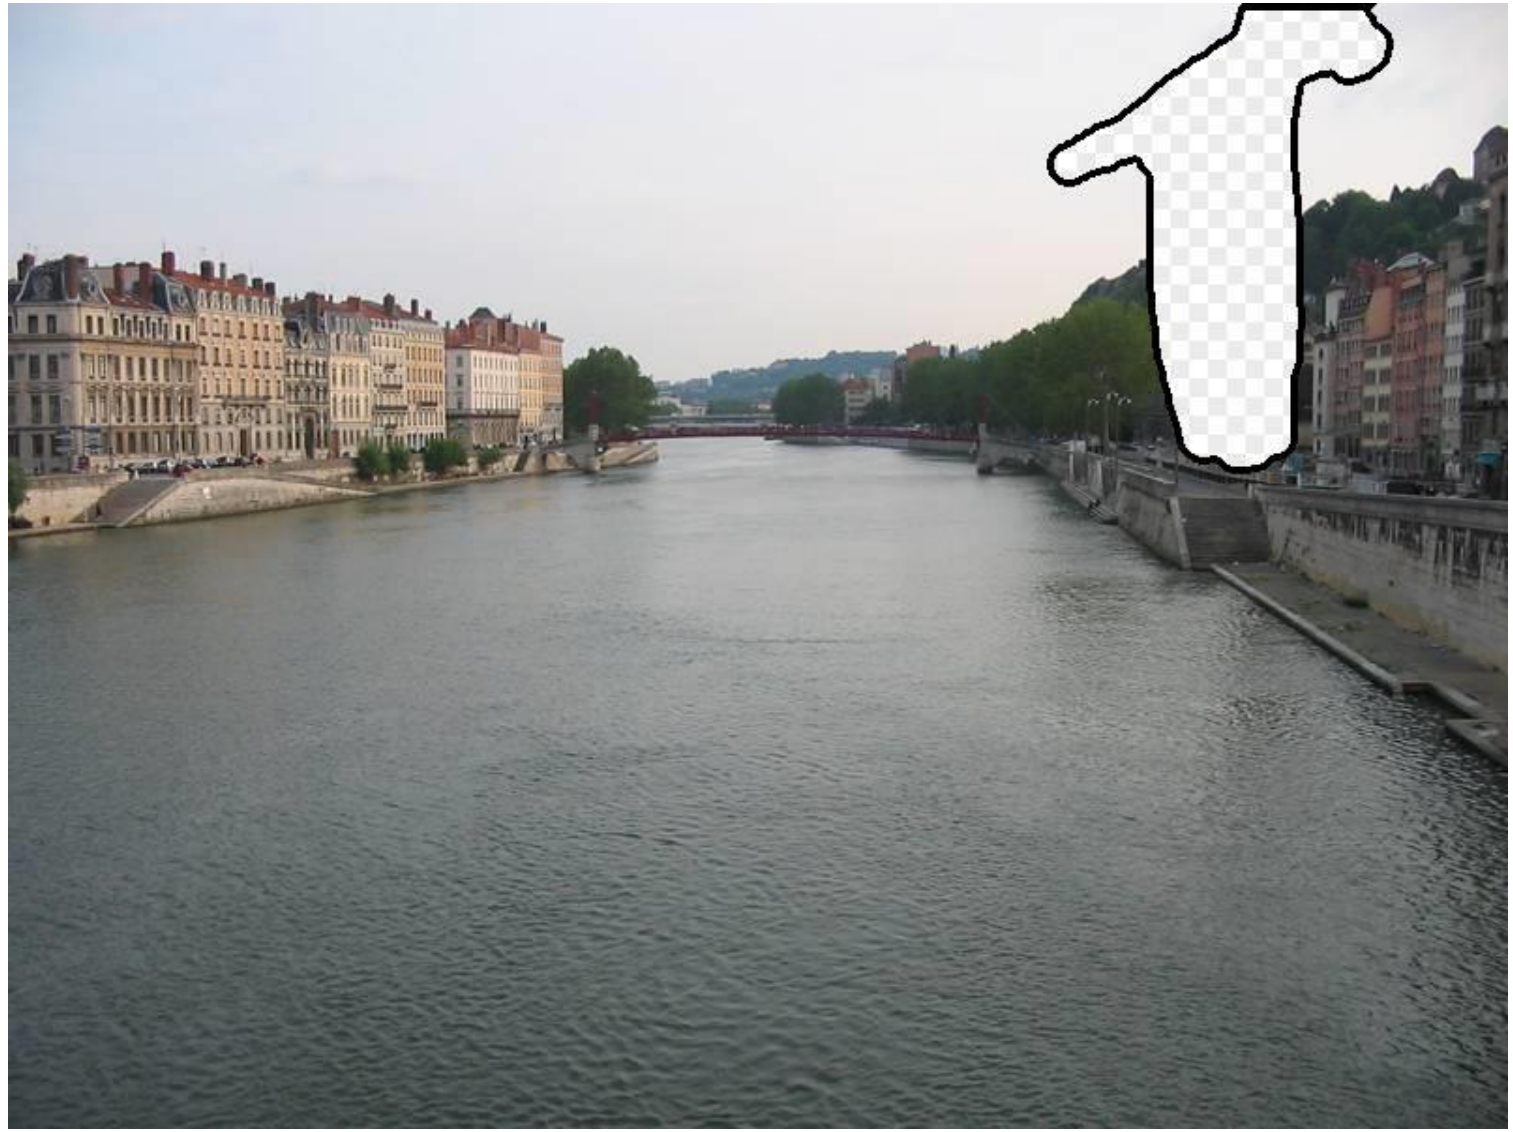

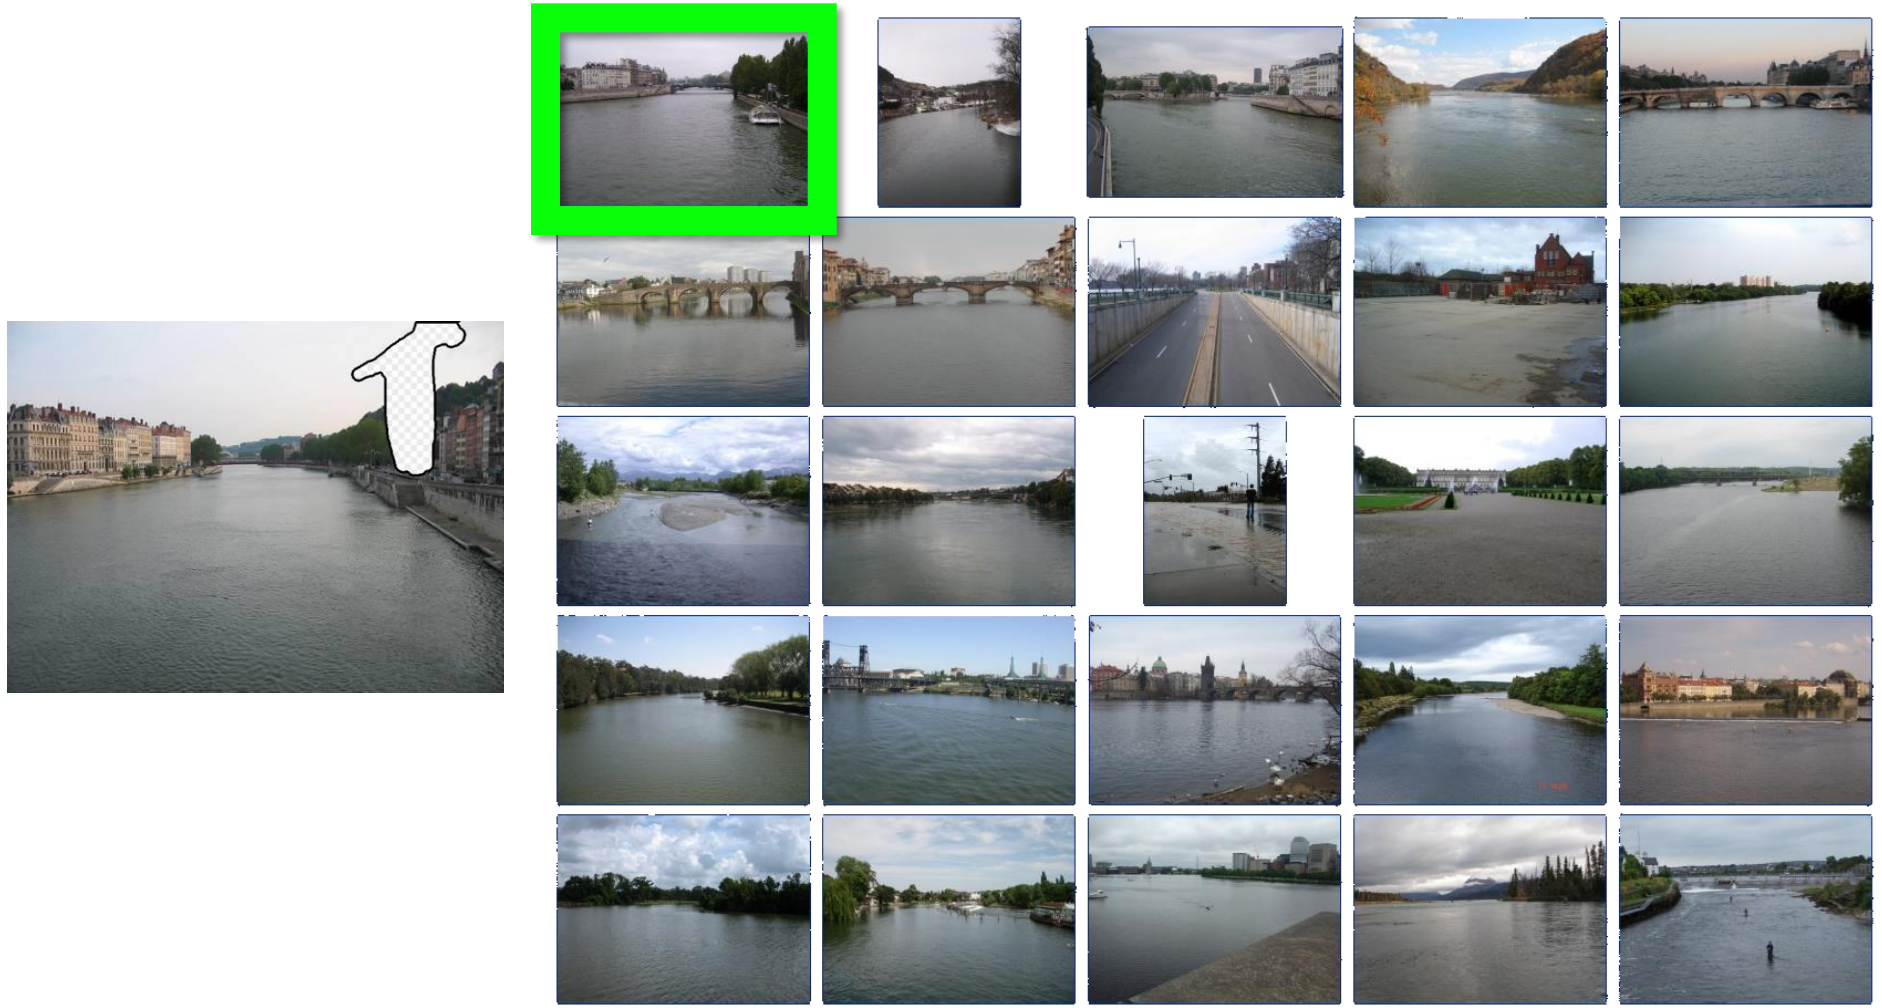

… 200 scene matches

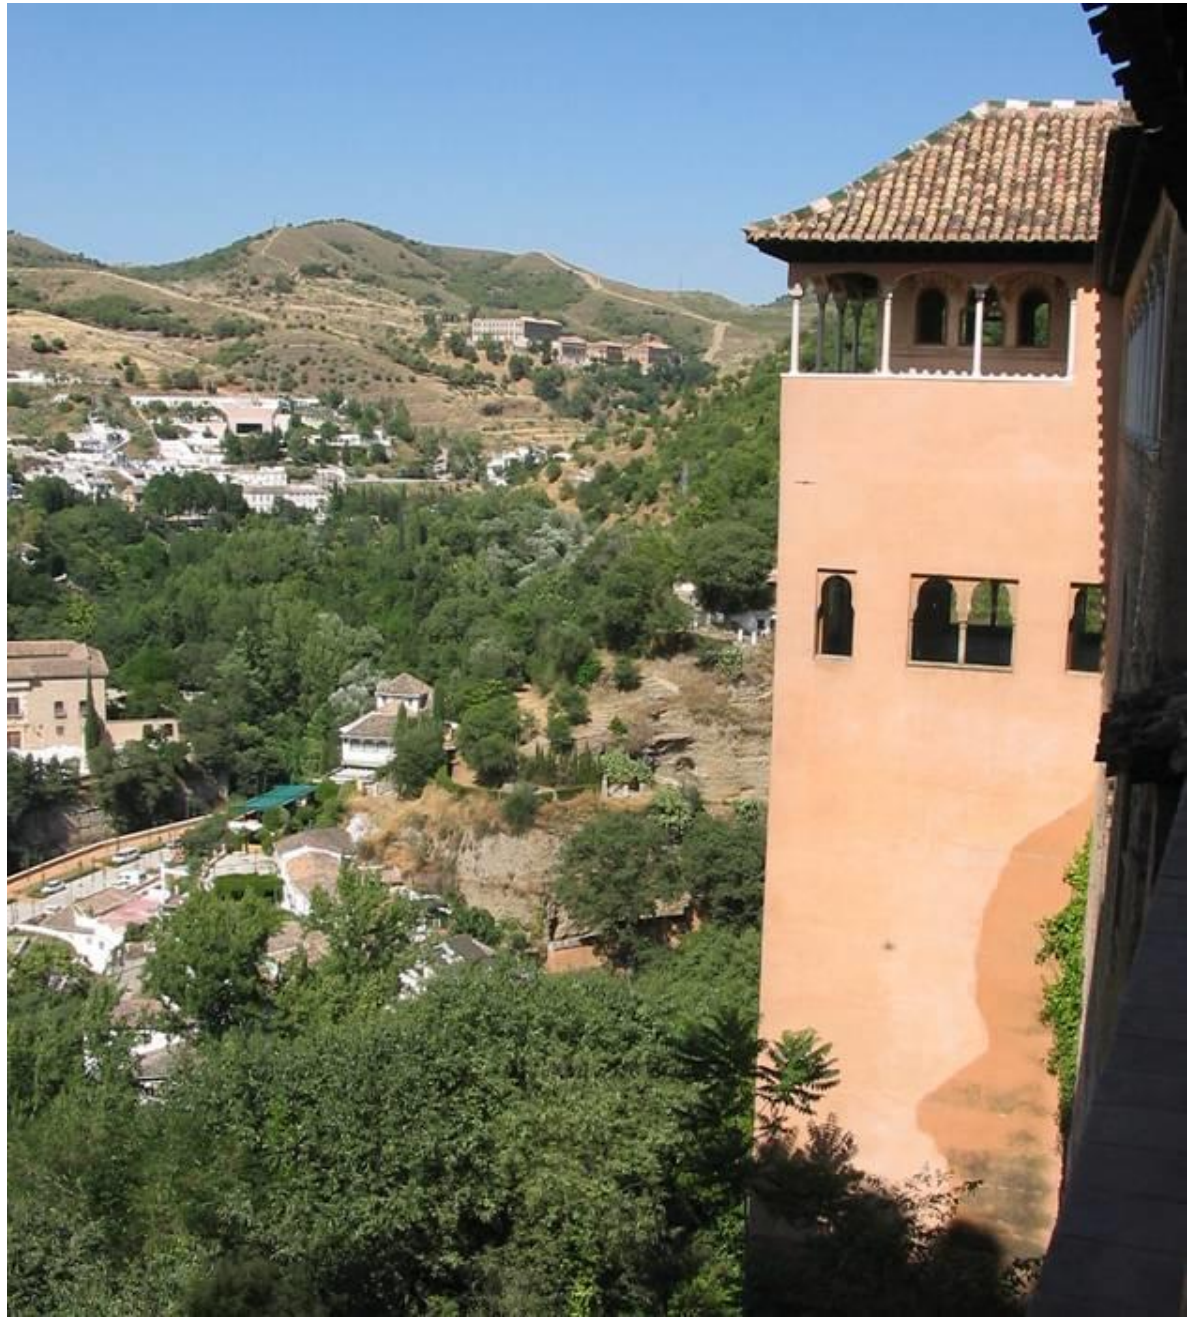

Outline

Opportunities of Scale: Data-driven methods

- The Unreasonable Effectiveness of Data
- Scene Completion
- Image Geolocation

General Principal

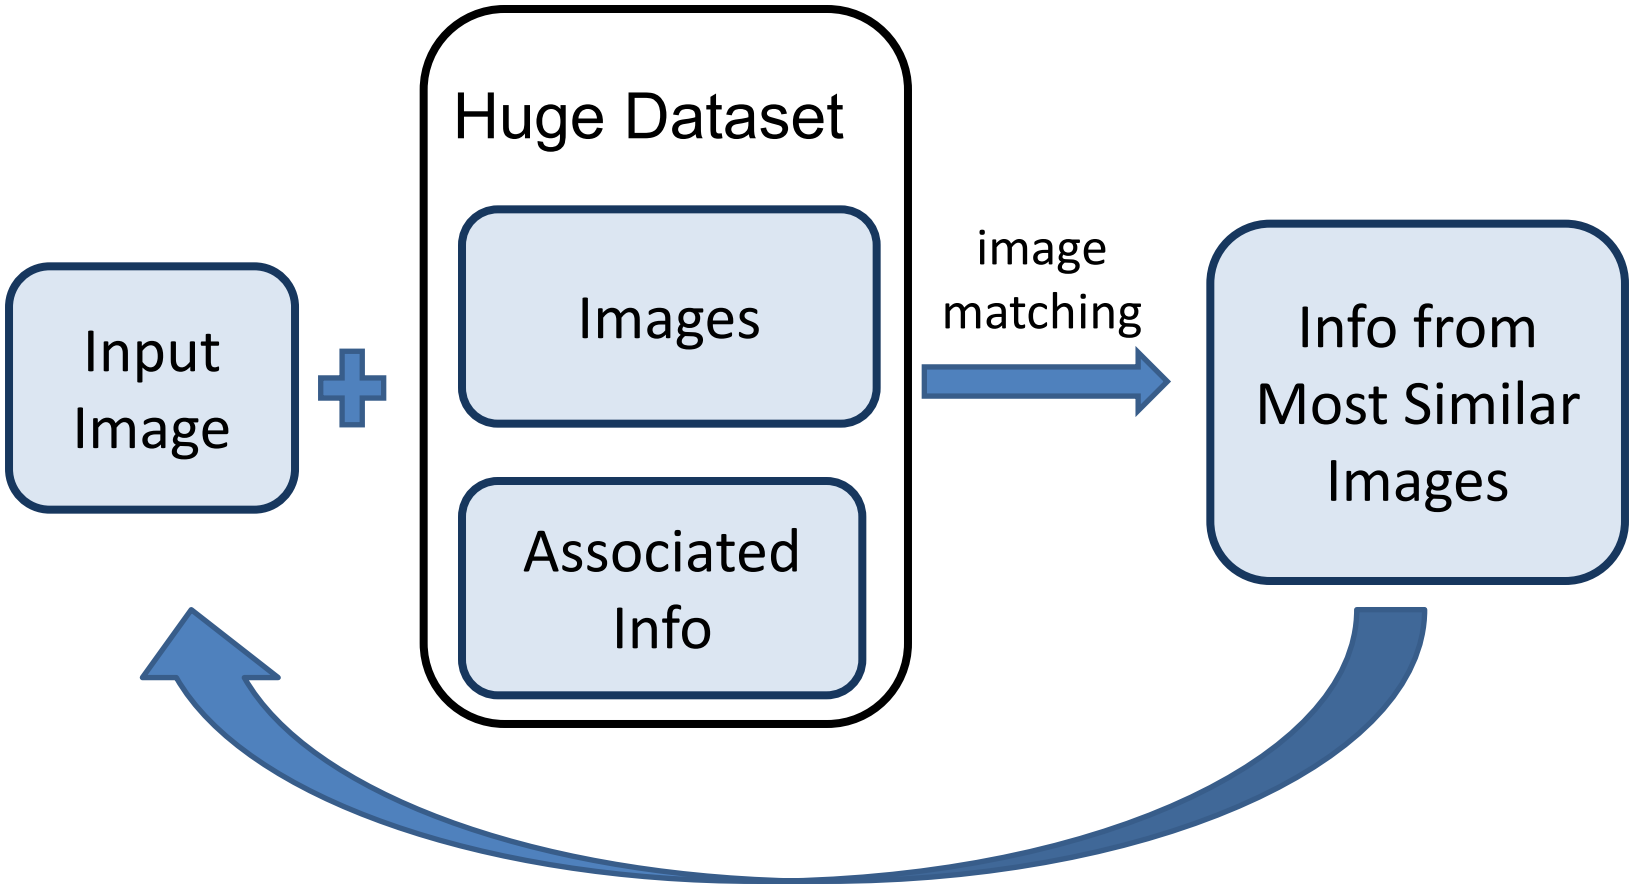

Hopefully, If you have enough images, the dataset will contain very similar images that you can find with simple matching methods.

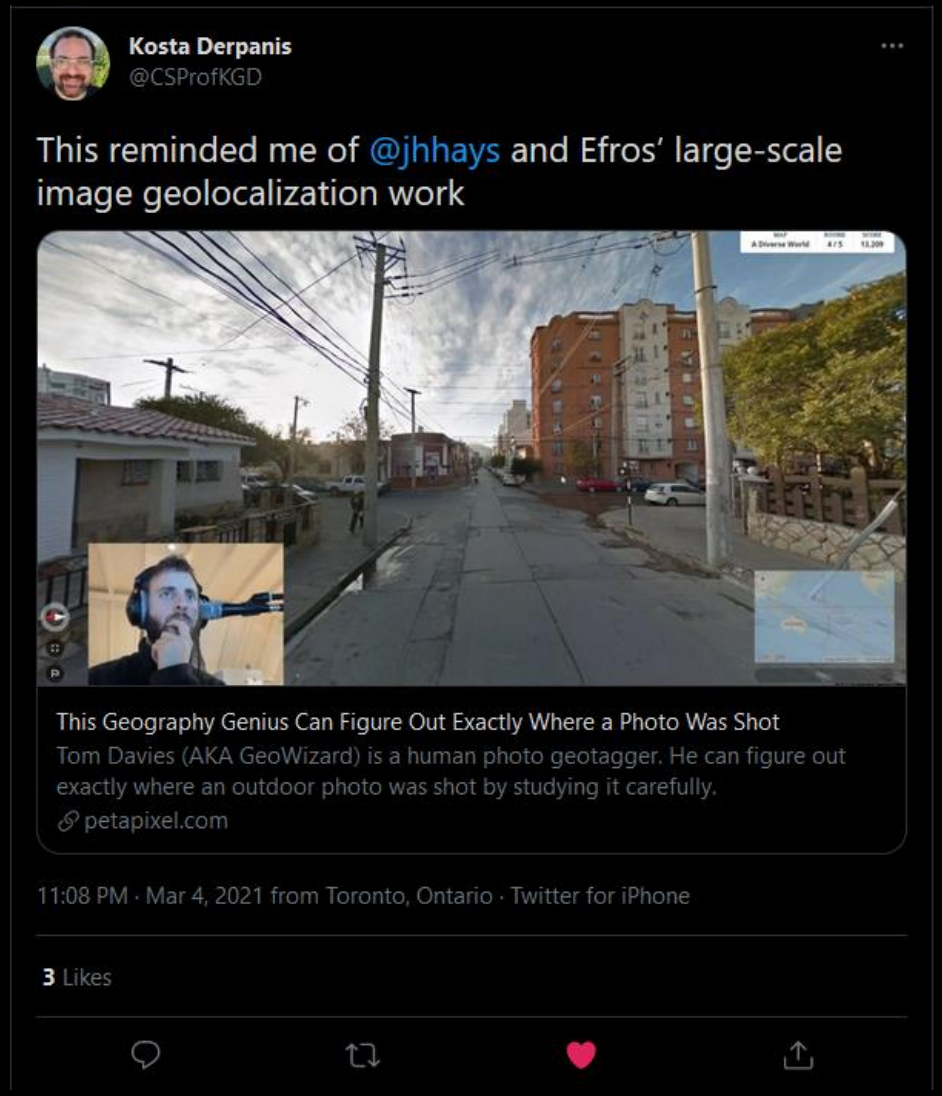

https://www.geoguessr.com/

https://www.youtube.com/c/GeoWizard/videos