Leveraging Big Data: Photo Completion

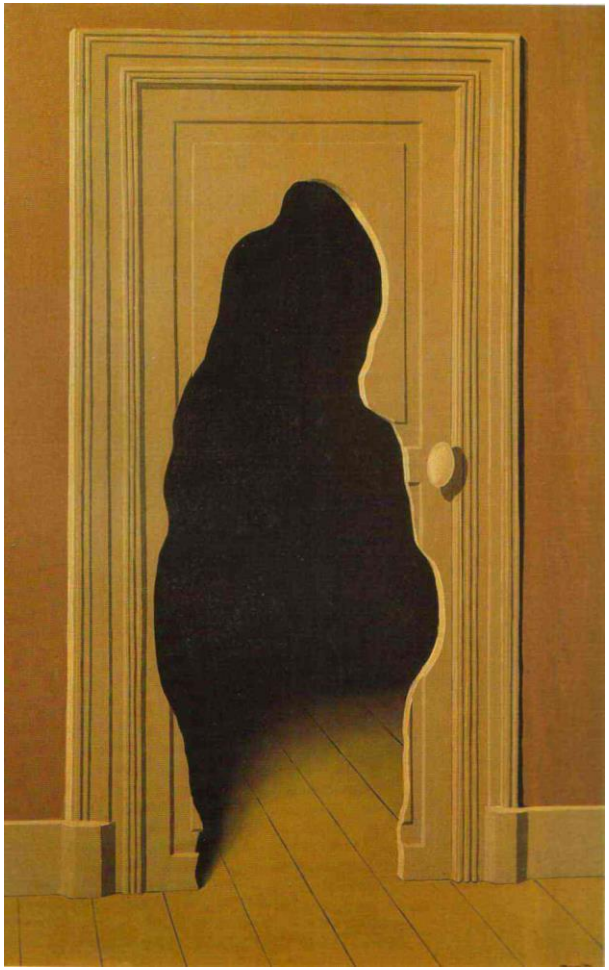

René Magritte, “The Unexpected Answer”

CSC320: Introduction to Visual Computing

Slides from Alyosha Efros, Derek Hoiem

Michael Guerzhoy

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 <u>in the data</u>

Google Translate

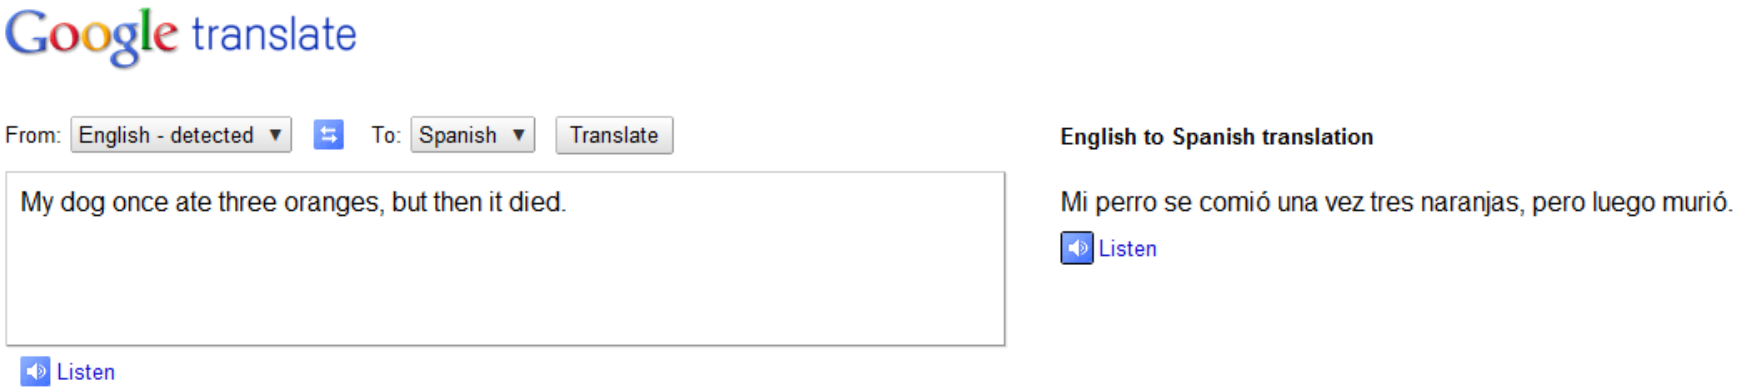

Chinese Room

- John Searle (1980)

SIGGRAPH2007 Computer Graphics Proceedings, Annual Conference Series, 2007

Scene Completion Using Millions of Photographs

James Hays Alexei A. Efros
Carnegie Mellon University

Original Image Input Scene Matches Output

What should the missing region contain?

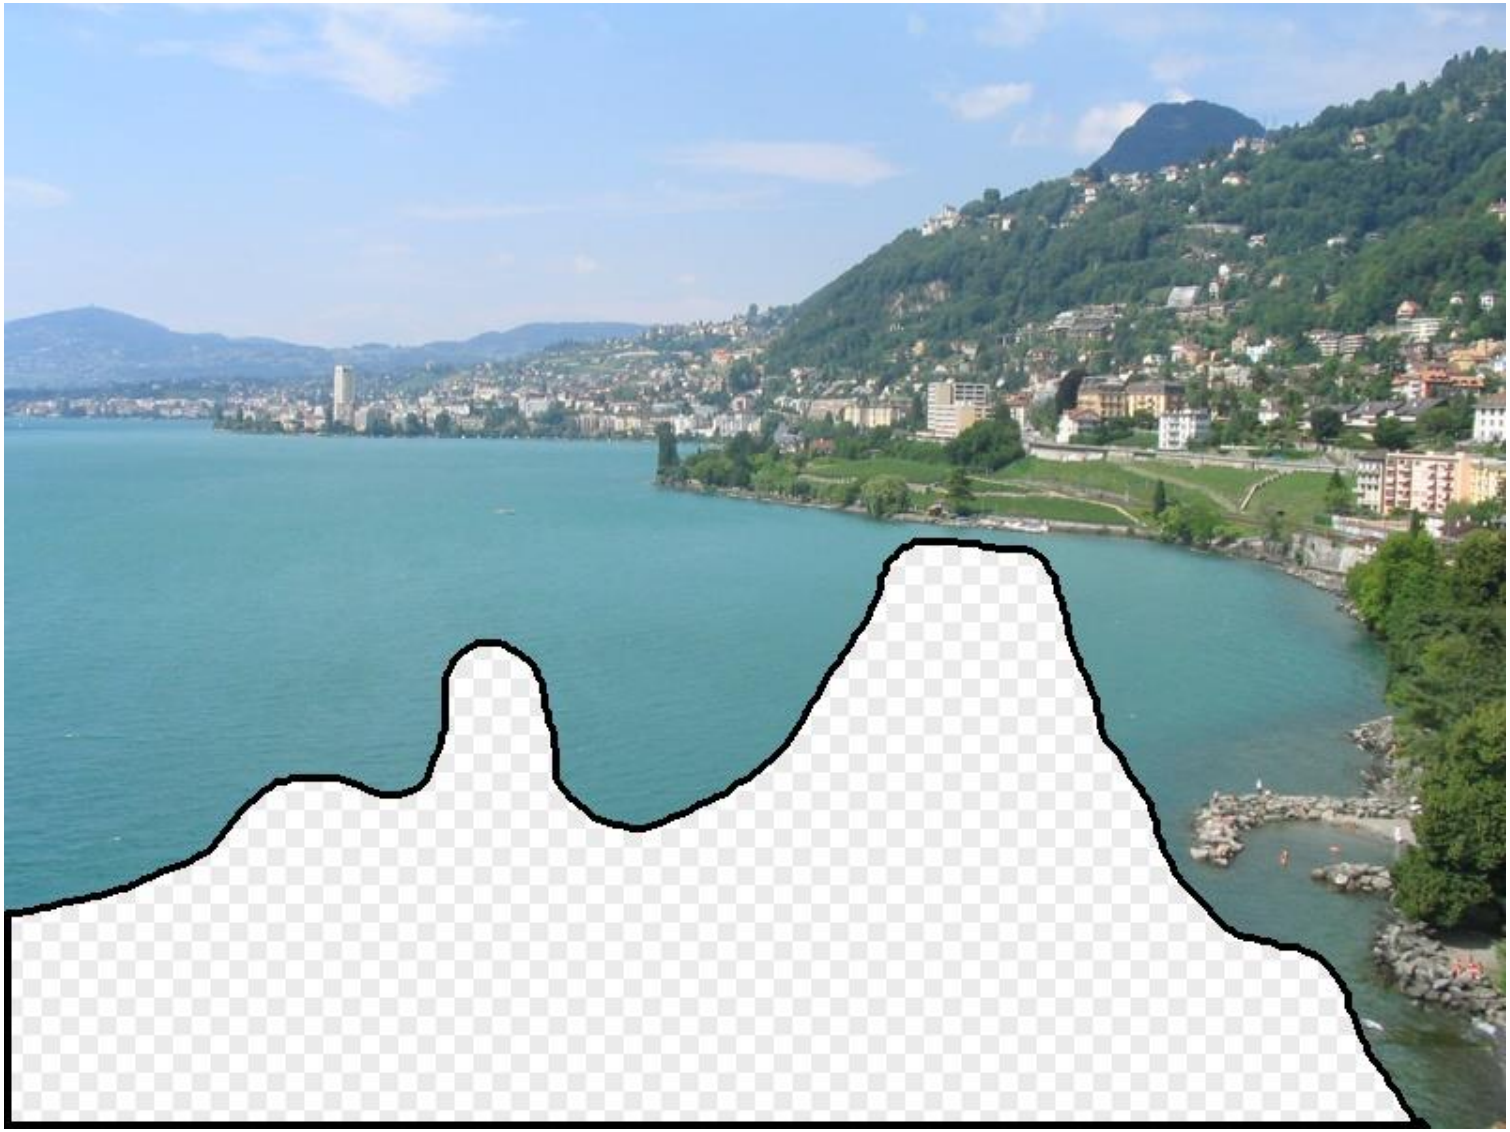

How it works

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

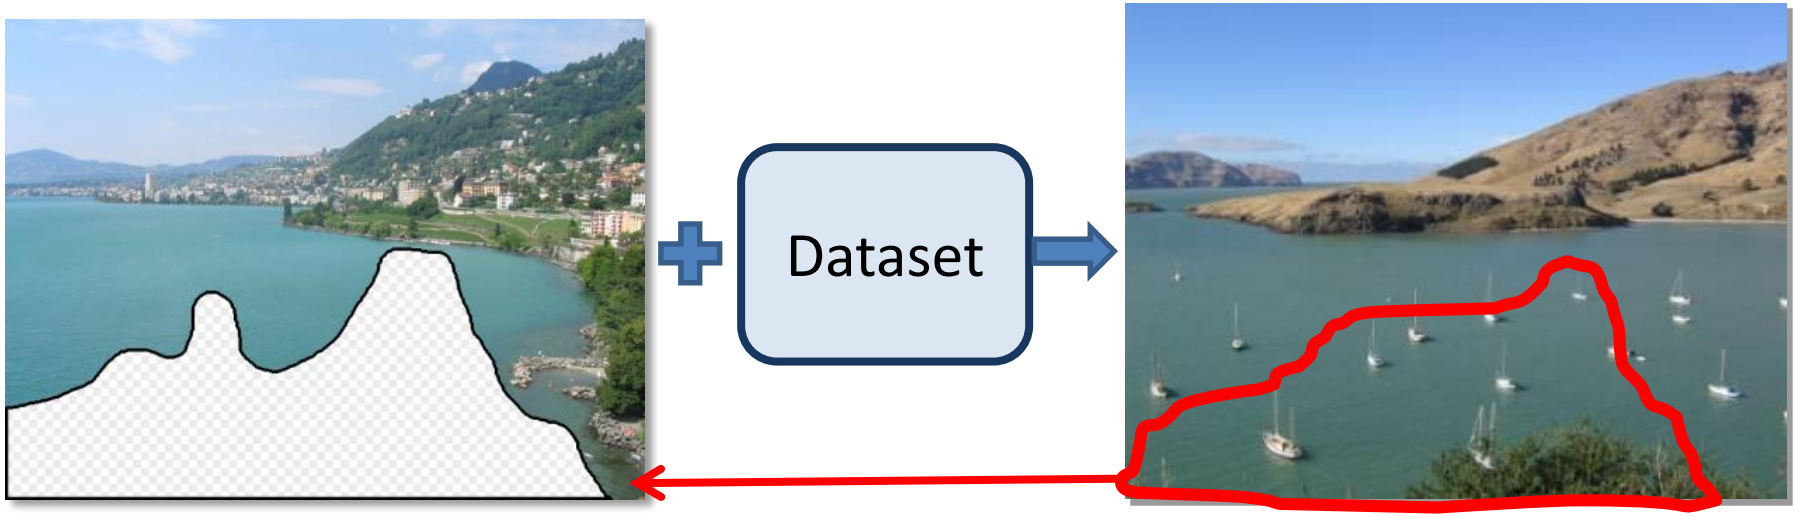

General Principal

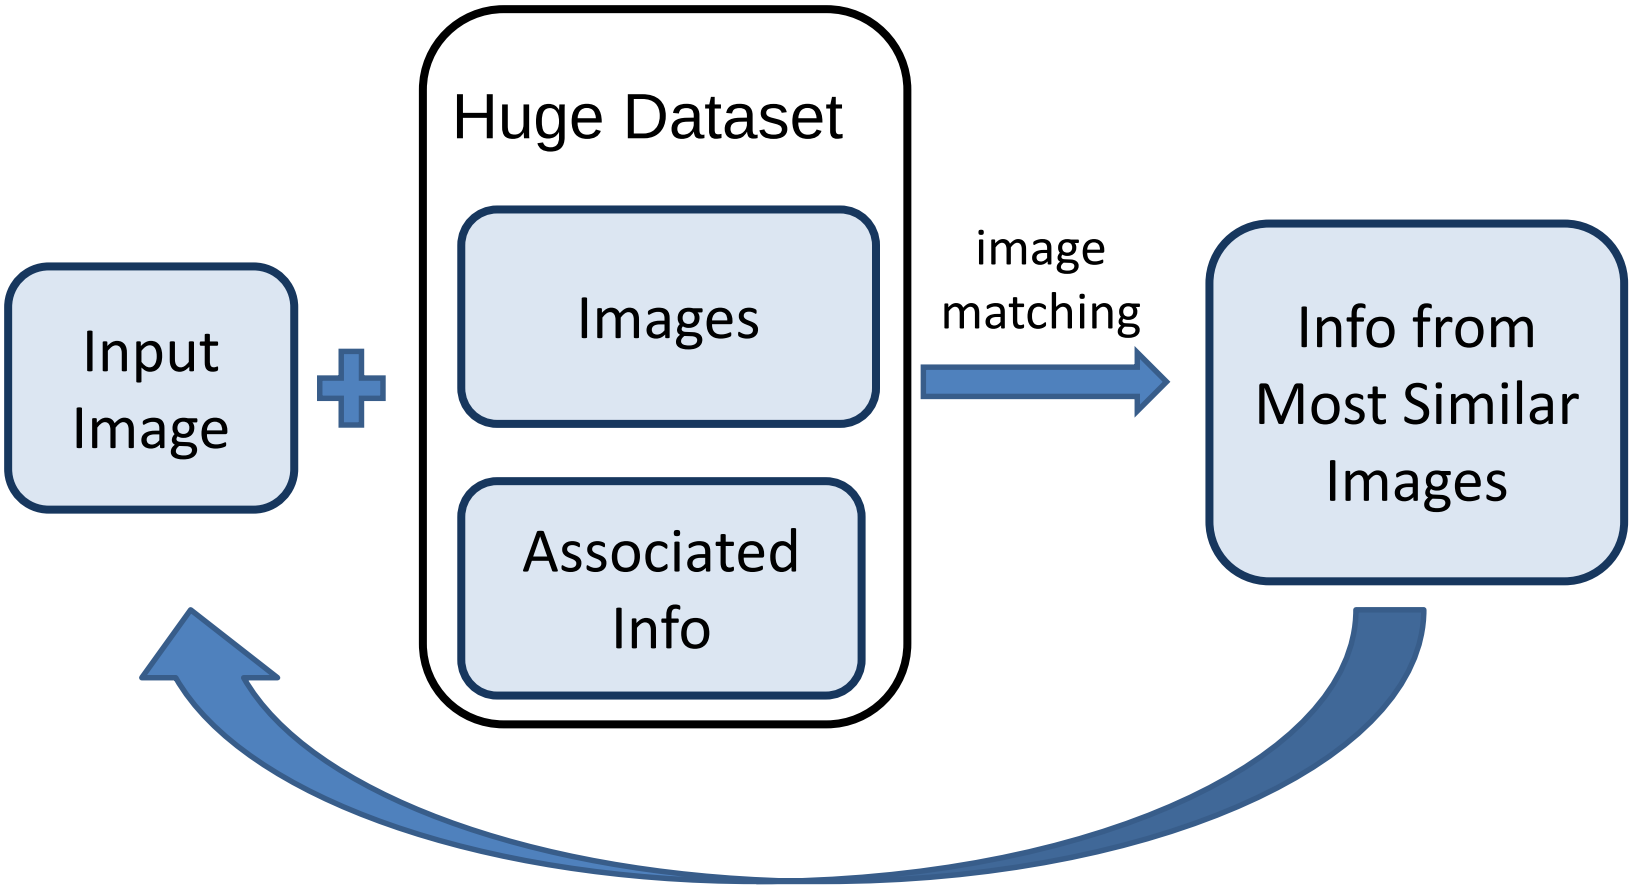

Trick: If you have enough images, the dataset will contain very similar images that you can find with simple matching methods.

How many images is enough?

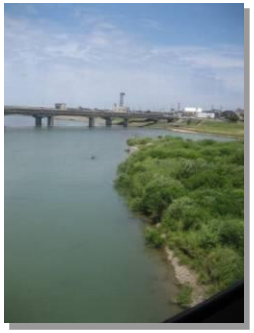




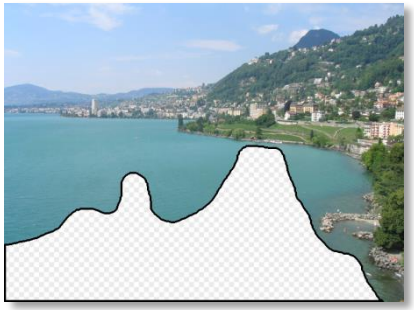

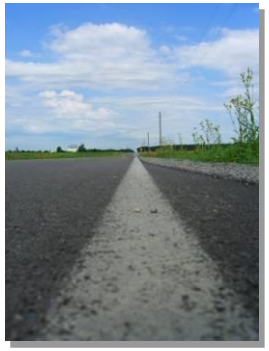


Nearest neighbors from a
collection of 20 thousand images

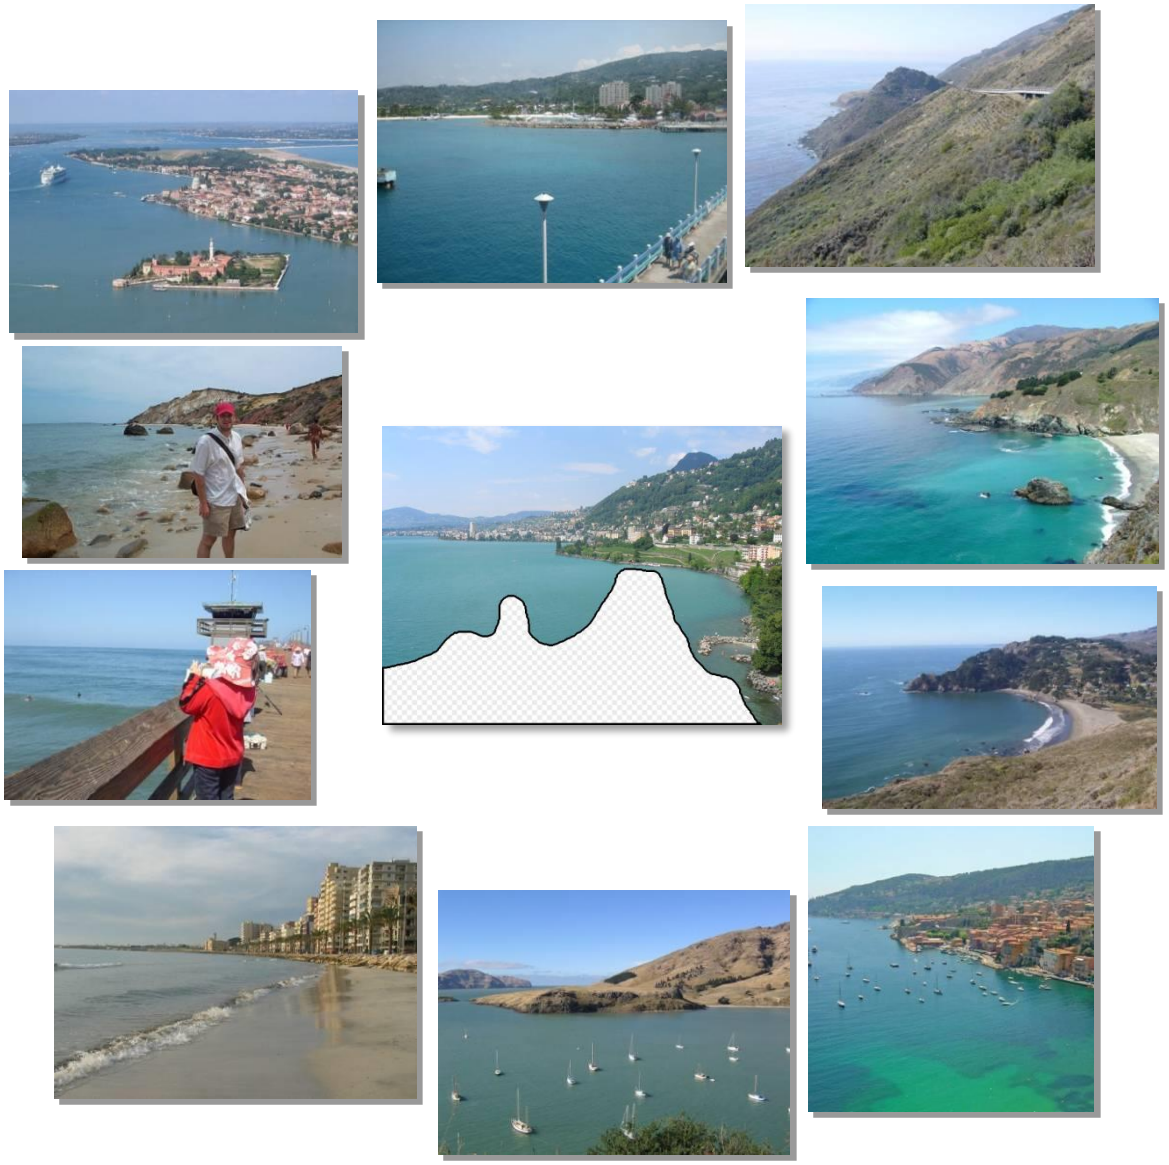

Nearest neighbors from a
collection of 2 million images

Image Data on the Internet

- Flickr (now)
  - 6 billion images per month
  - More than 100 petabytes of images/video
- Flickr (as of Sept. 19th, 2010)
  - 5 billion photographs
  - 100+ million geotagged images
- Imageshack (as of 2009)
  - 20 billion
- Facebook (as of 2009)
  - 15 billion

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

Scene Matching

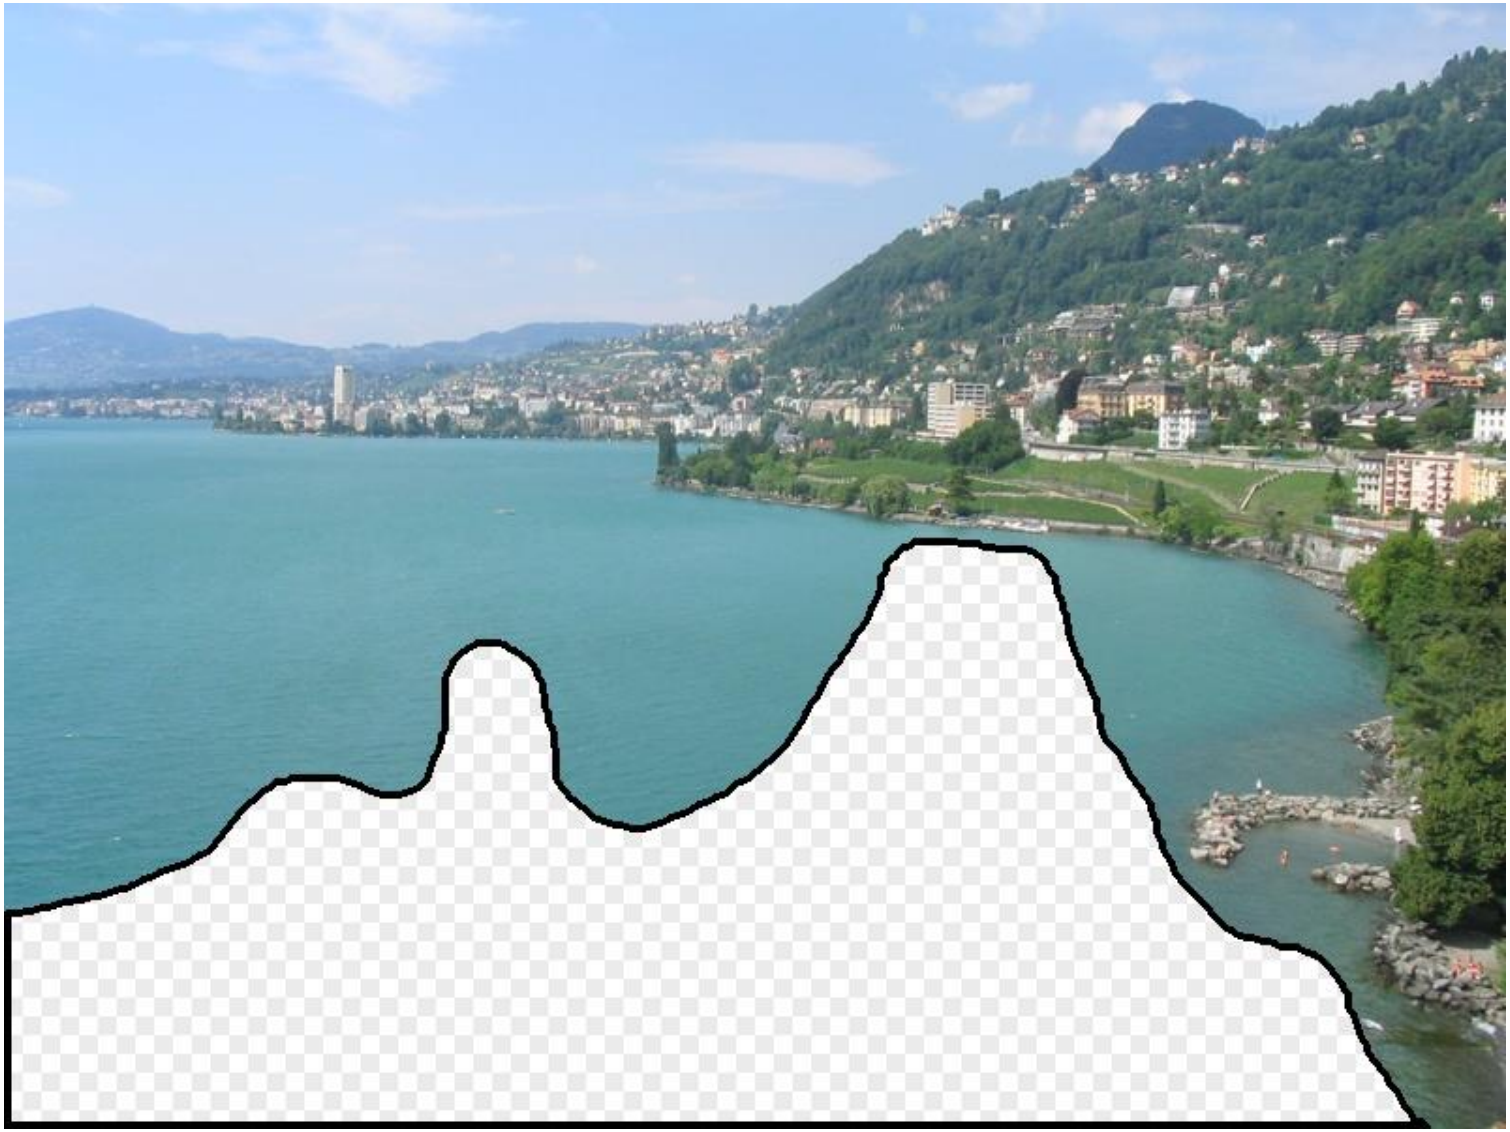

Scene Descriptor

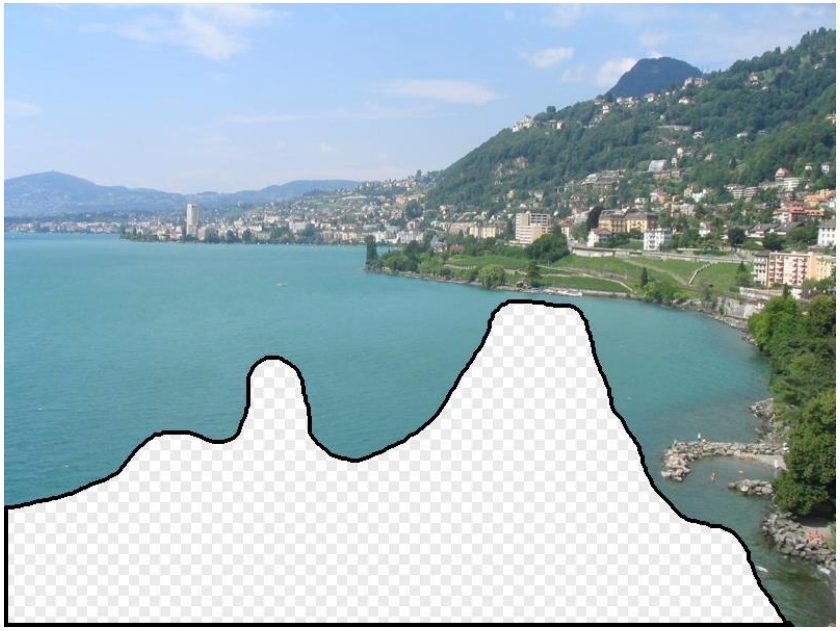

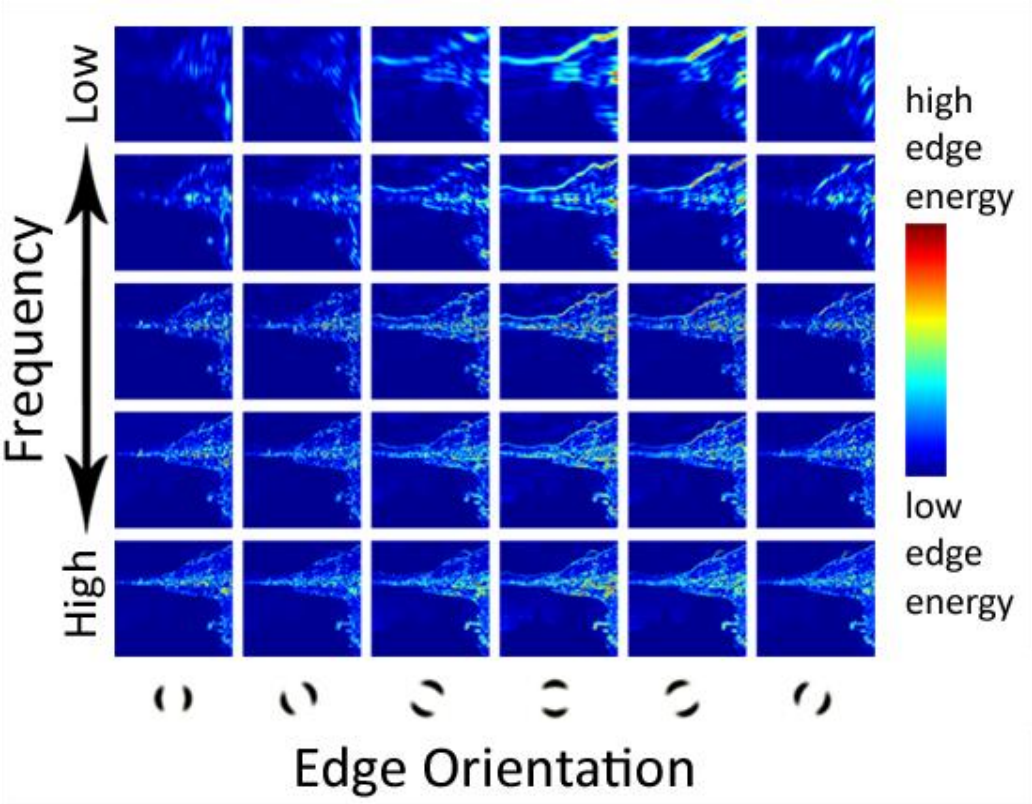

Scene Descriptor

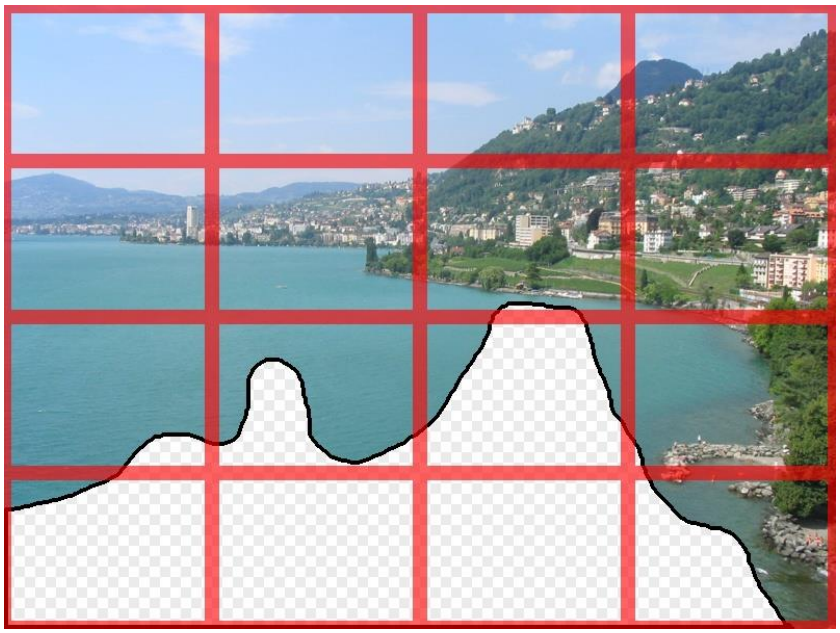

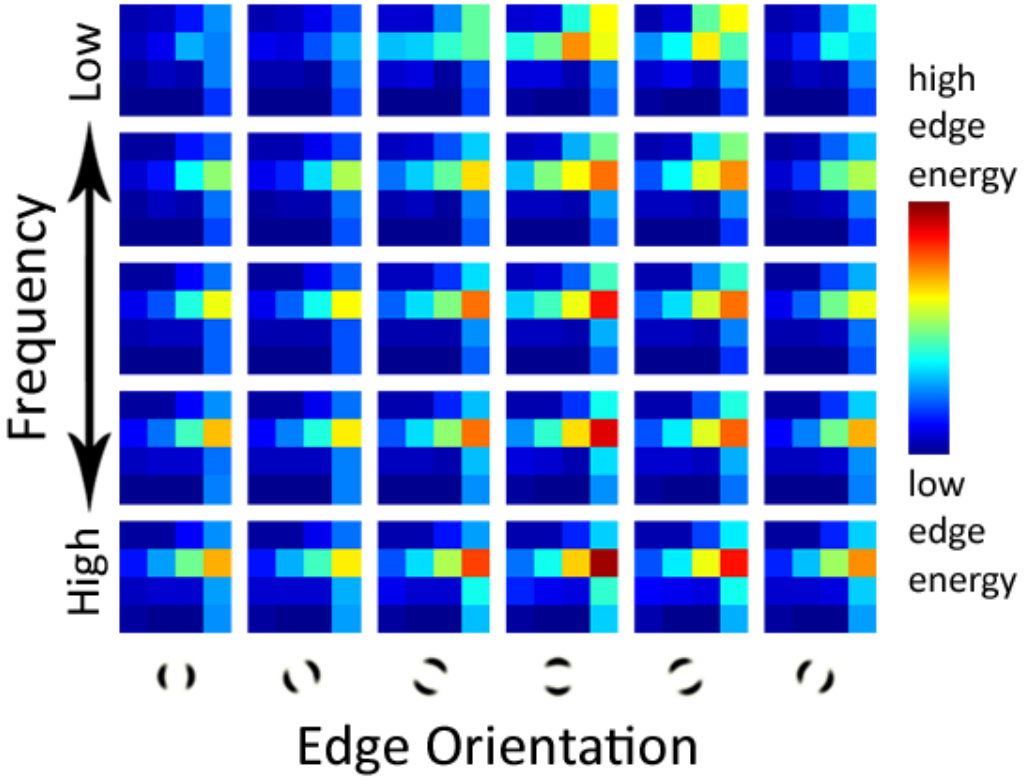

Scene Gist Descriptor
(Oliva and Torralba 2001)

Scene Descriptor

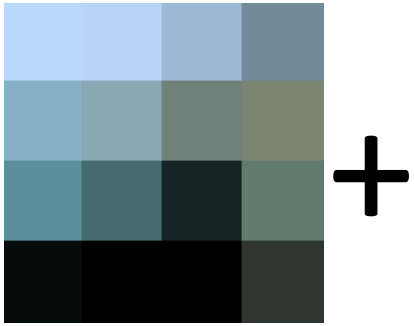

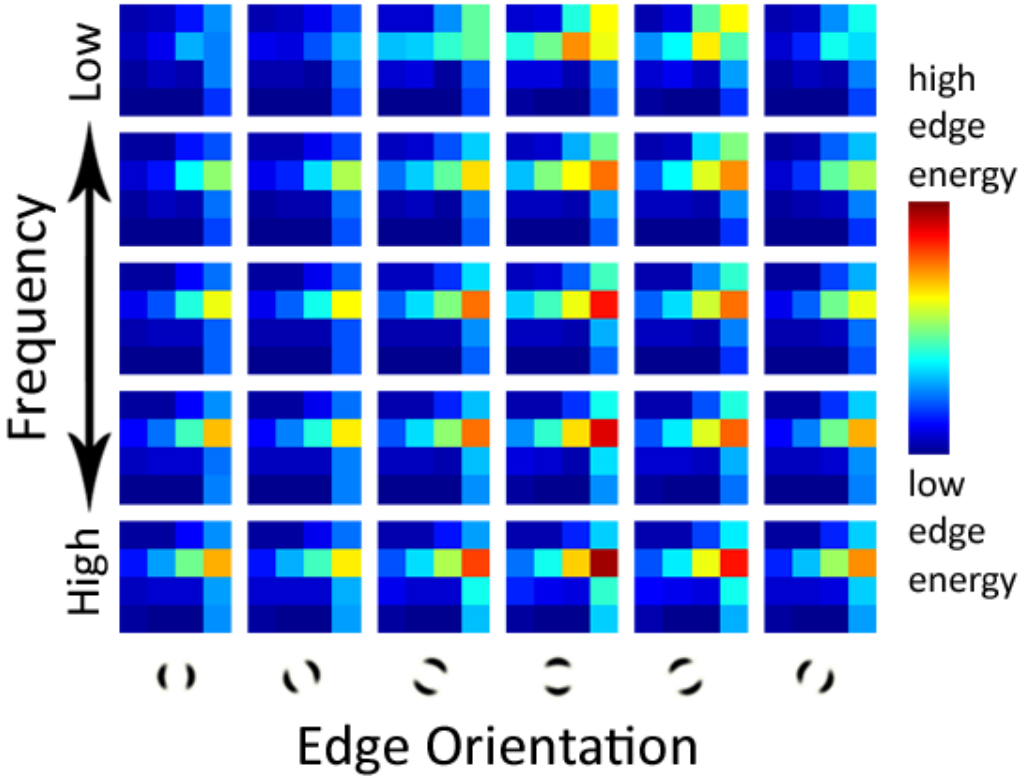

Scene Gist Descriptor
(Oliva and Torralba 2001)


2 Million Flickr Images

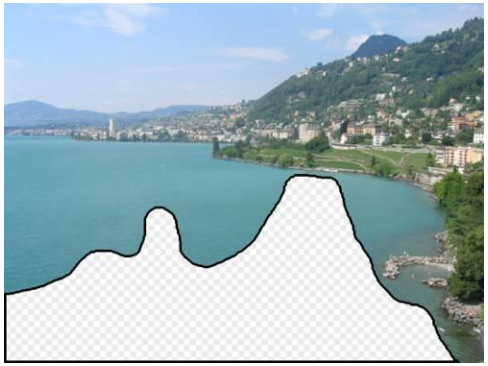

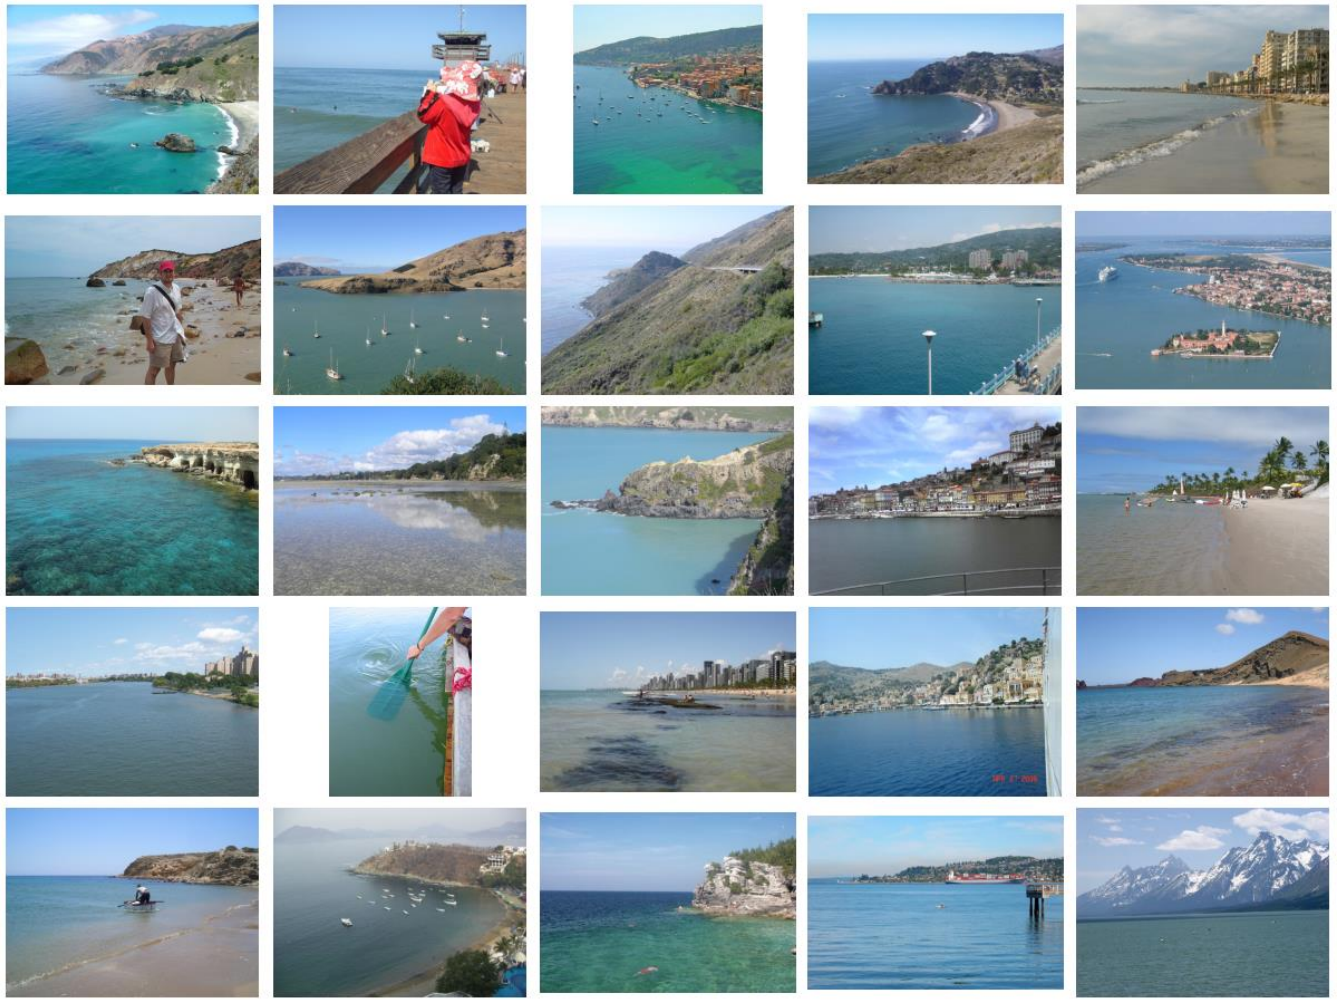

… 200 total

Context Matching

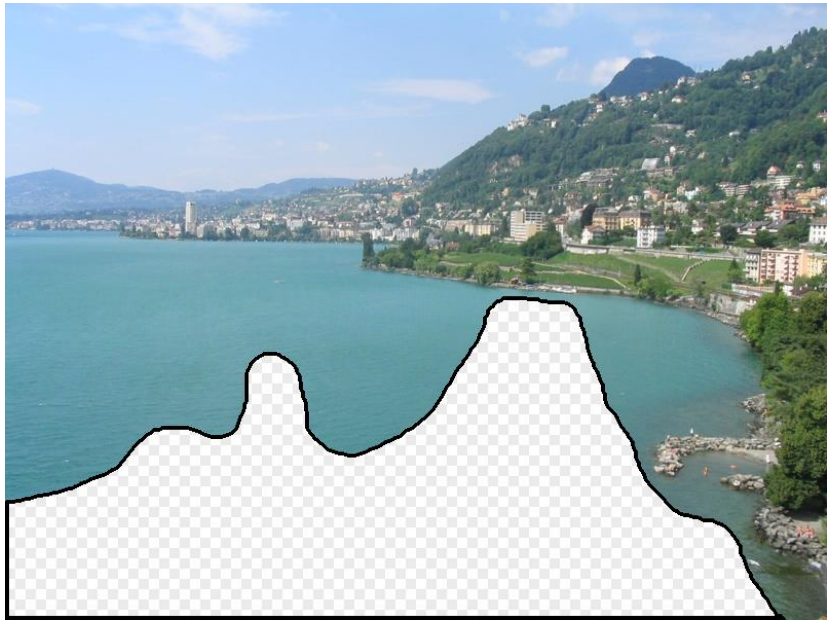


Find the best seam + blend

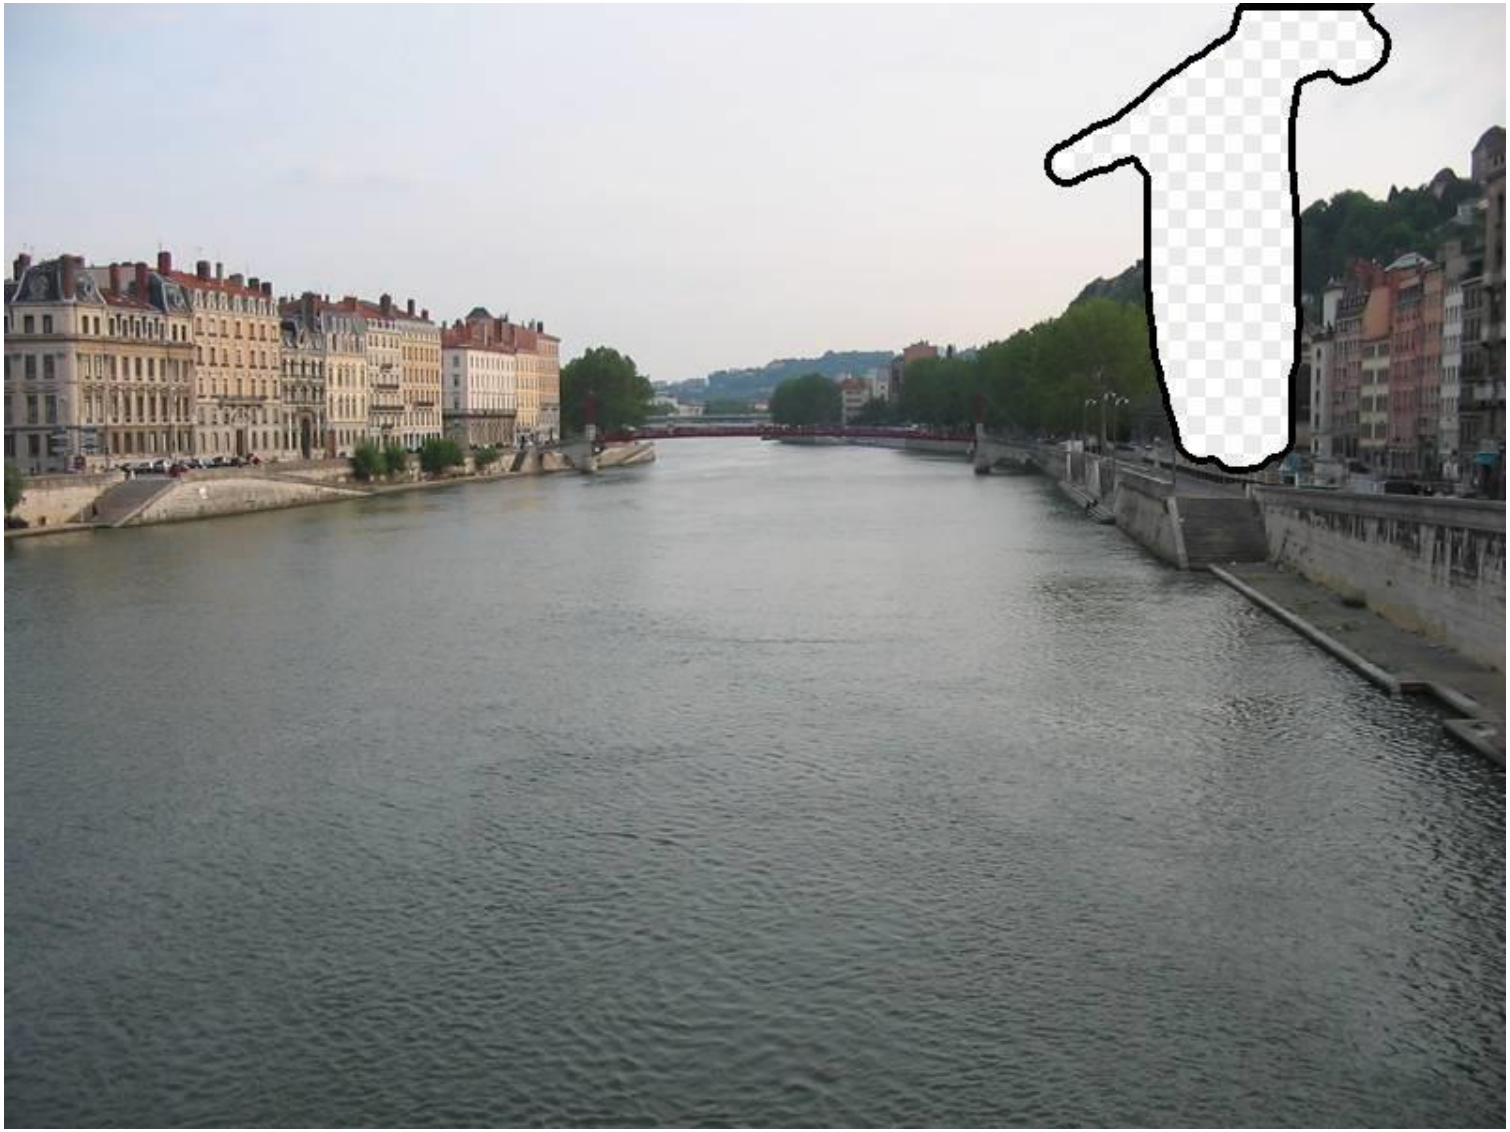

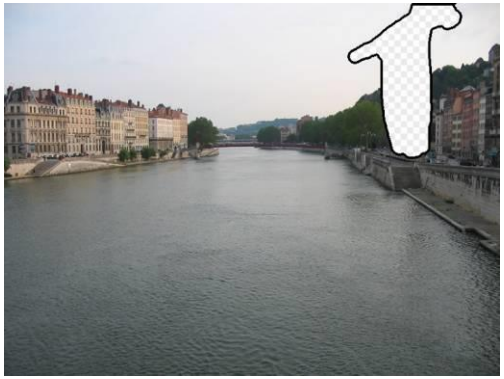

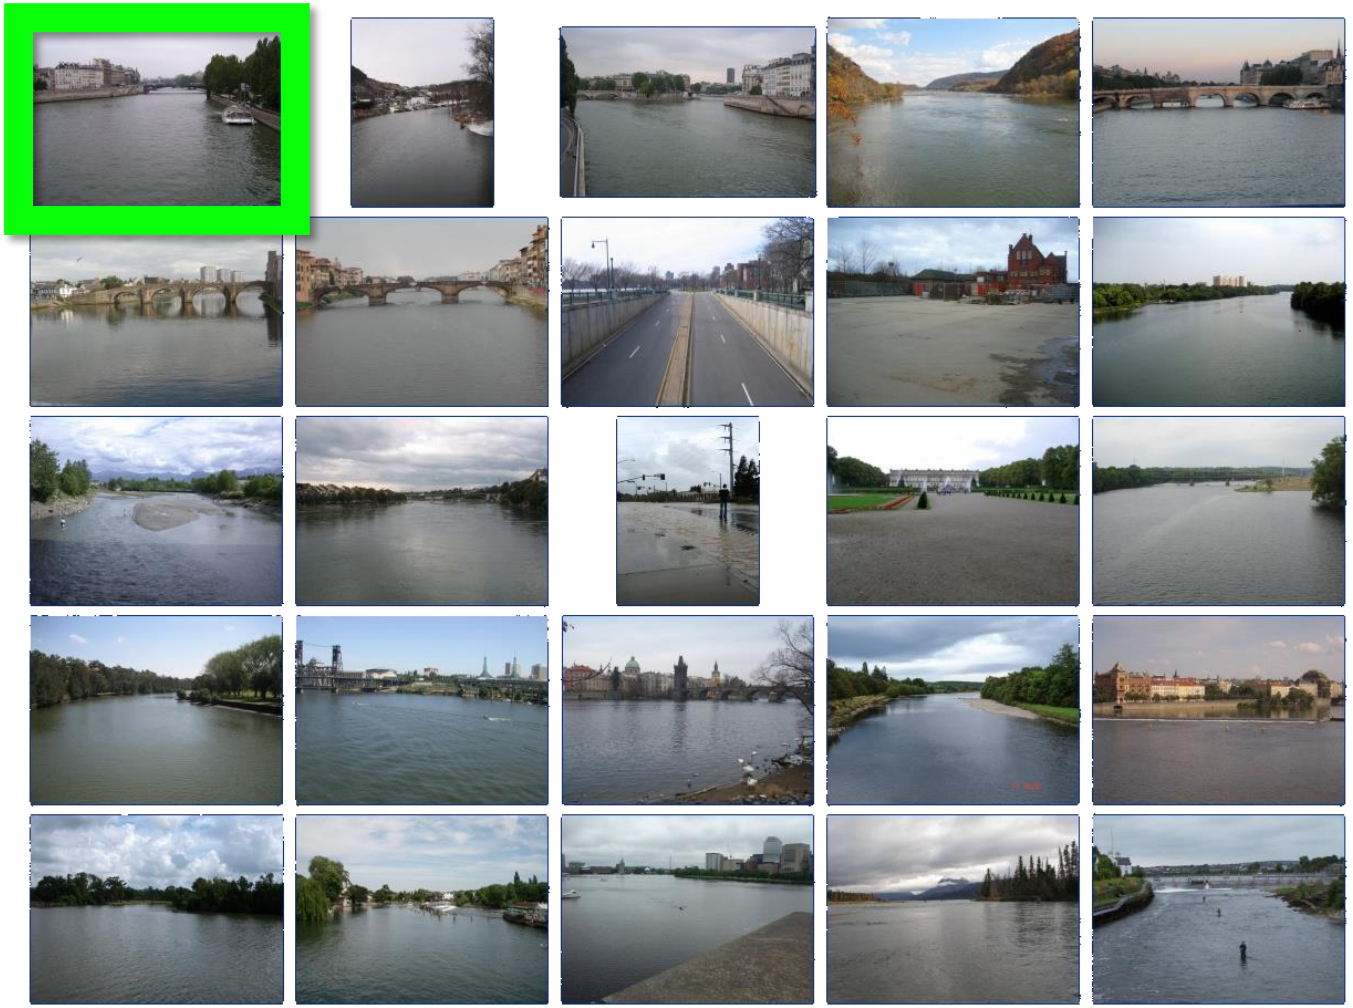

… 200 scene matches