Data-driven Methods: Texture

© Darren Green (www.darrensworld.com)

15-463: Computational Photography

Alexei Efros, CMU, Fall 2007

Texture

- Texture depicts spatially repeating patterns
- Many natural phenomena are textures

radishes

rocks

yogurt

Texture Synthesis

- Goal of Texture Synthesis: create new samples of a given texture
- Many applications: virtual environments, hole-filling, texturing surfaces

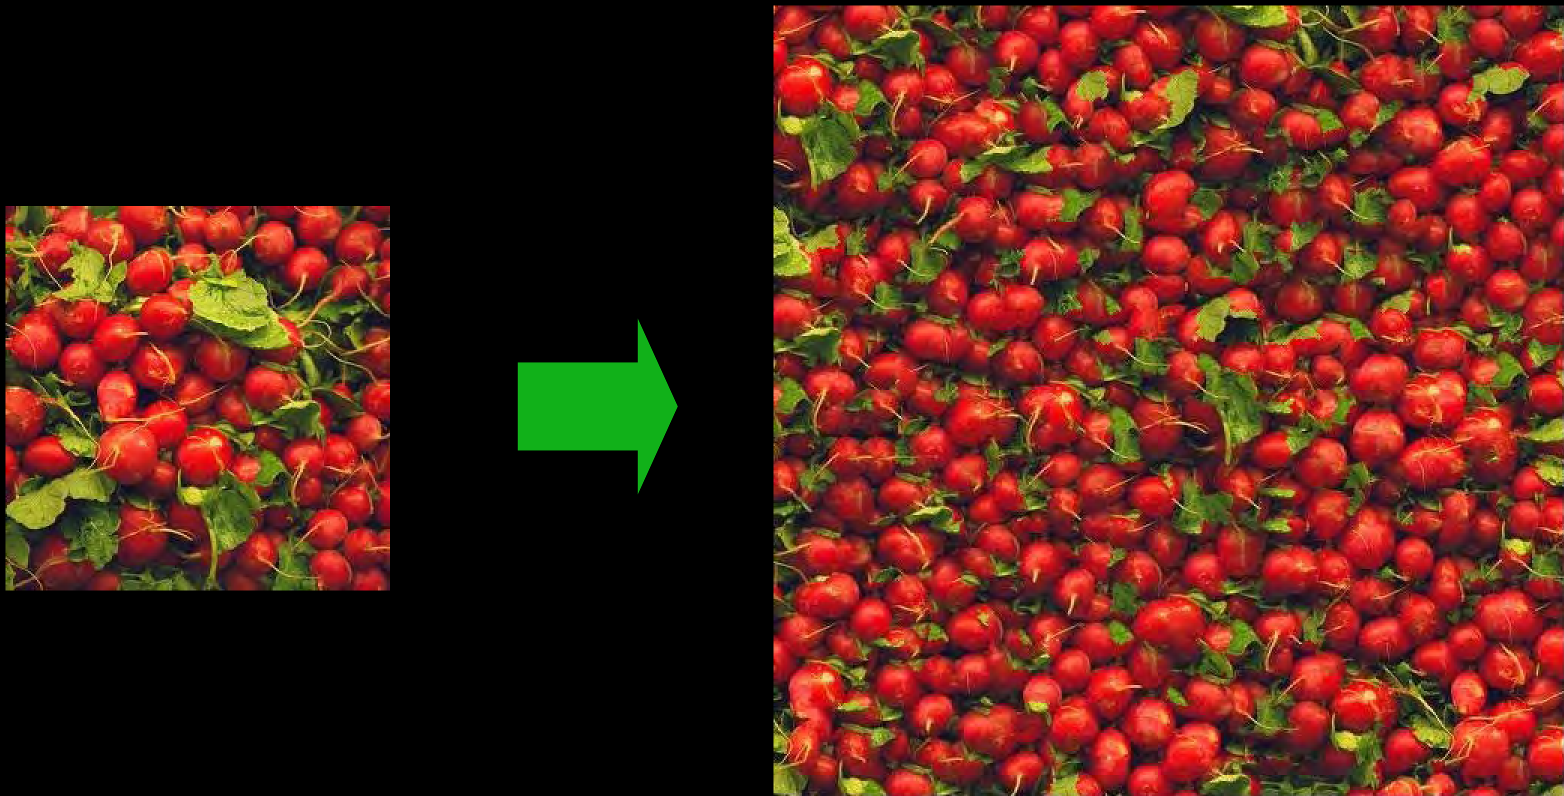

The Challenge

- Need to model the whole spectrum: from repeated to stochastic texture

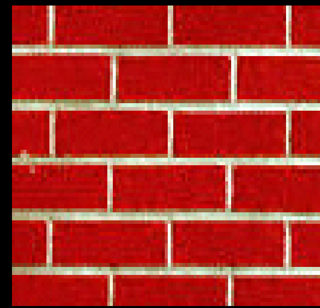

repeated

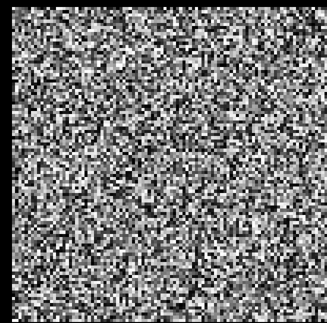

stochastic

Both?

Efros & Leung Algorithm

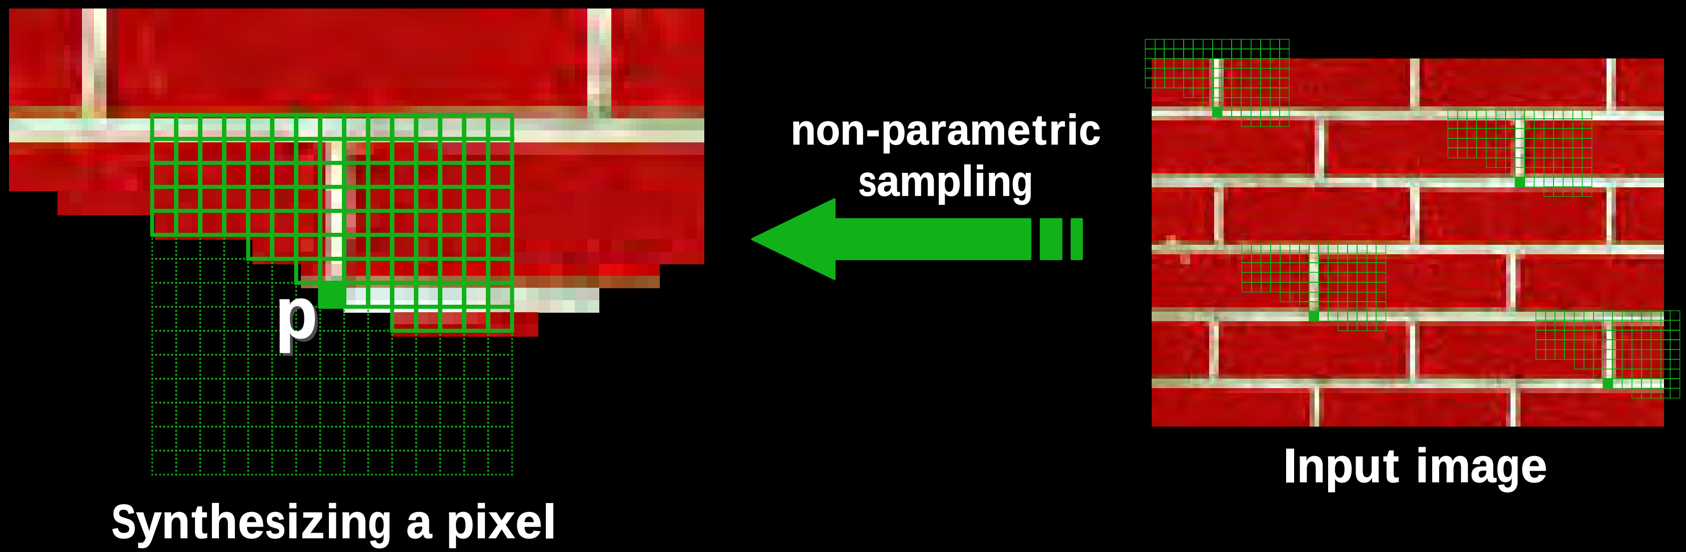

- Assuming Markov property, compute P(**p**|N(**p**))
  - Building explicit probability tables infeasible
  - Instead, we *search the input image* for all similar neighborhoods — that's our pdf for **p**
  - To sample from this pdf, just pick one match at random

Some Details

- Growing is in "onion skin" order
  - Within each "layer", pixels with most neighbors are synthesized first
  - If no close match can be found, the pixel is not synthesized until the end
- Using *Gaussian-weighted* SSD is very important
  - to make sure the new pixel agrees with its closest neighbors
  - Approximates reduction to a smaller neighborhood window if data is too sparse

Neighborhood Window

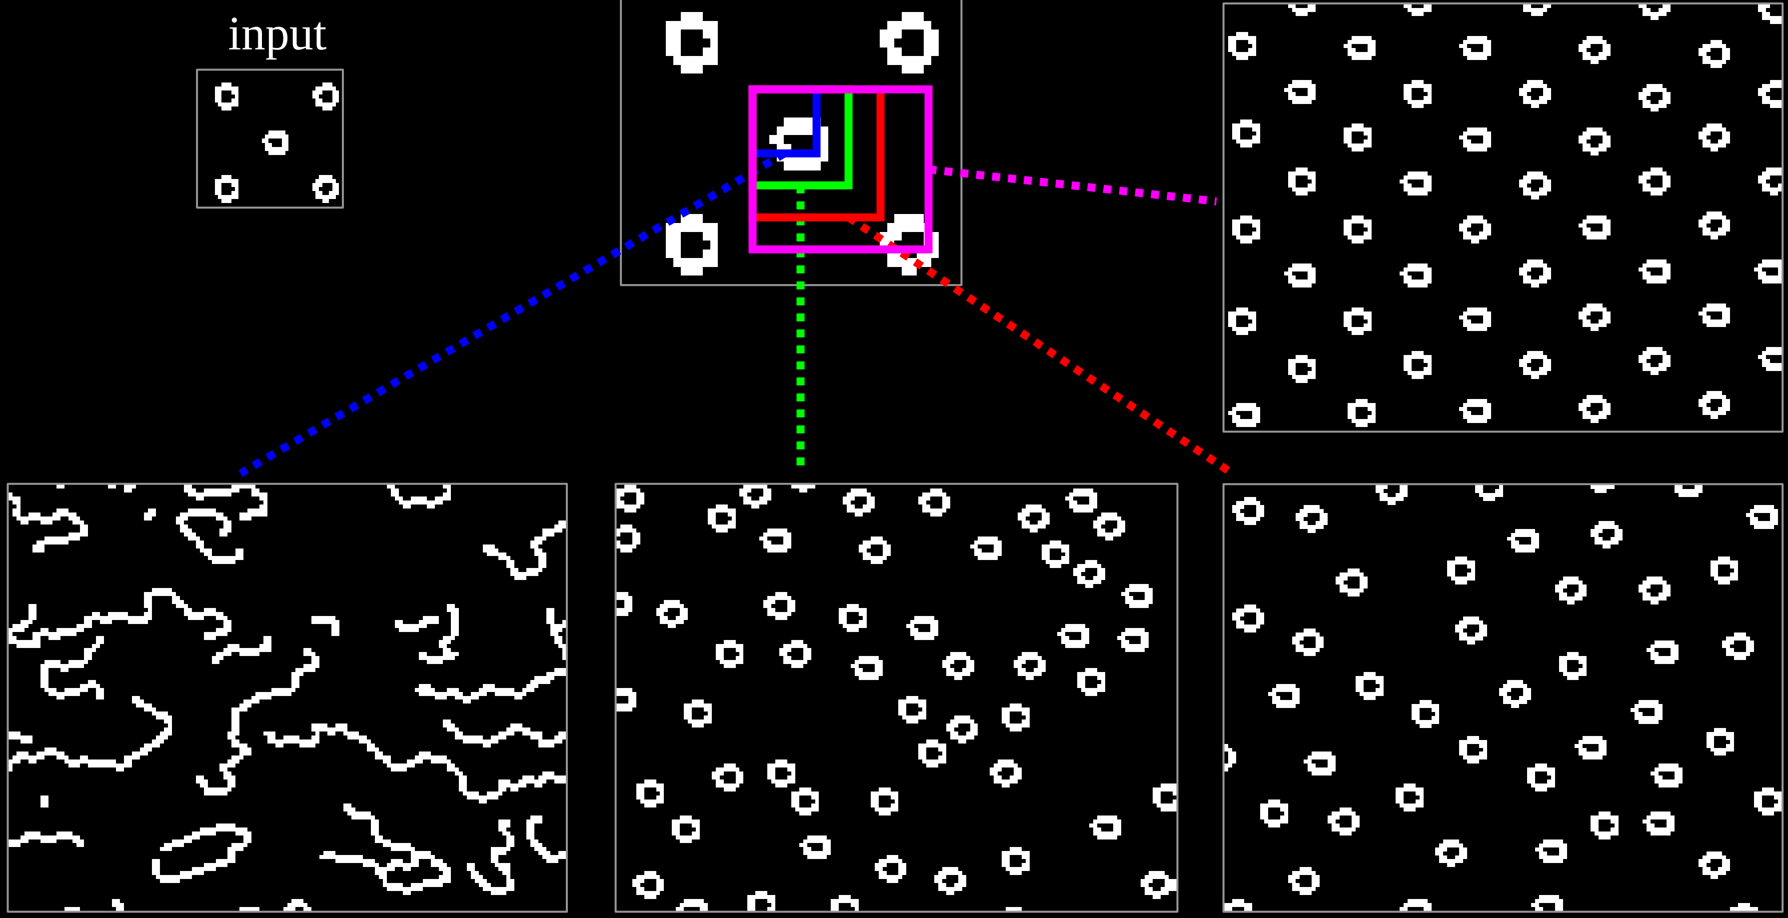

Varying Window Size

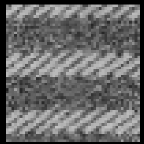

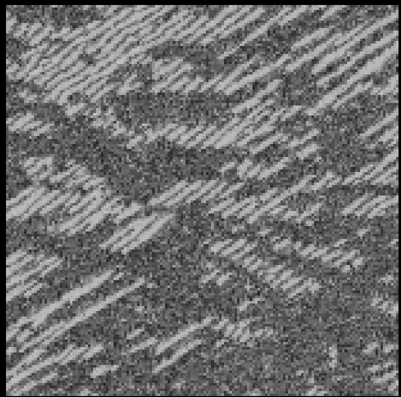

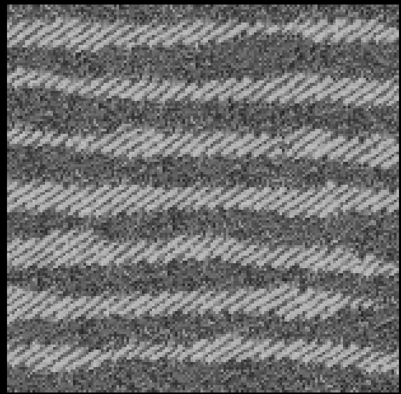

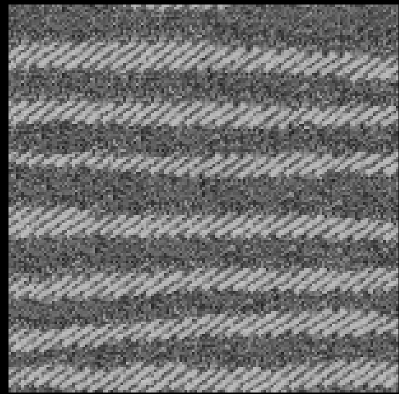

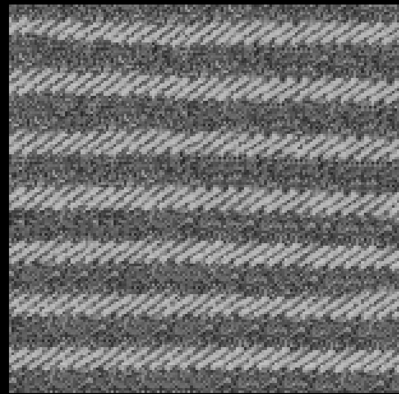

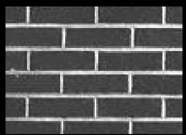

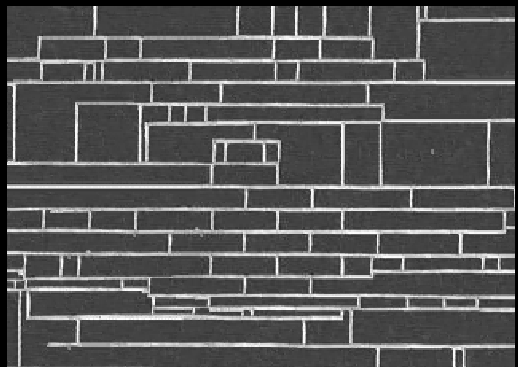

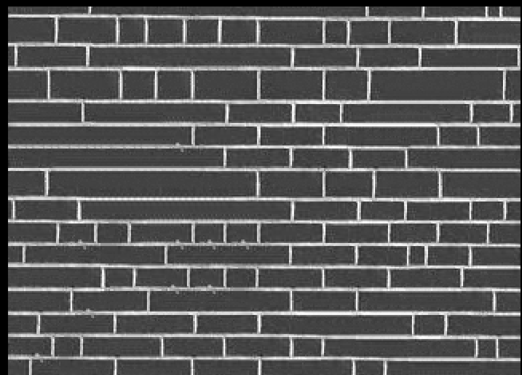

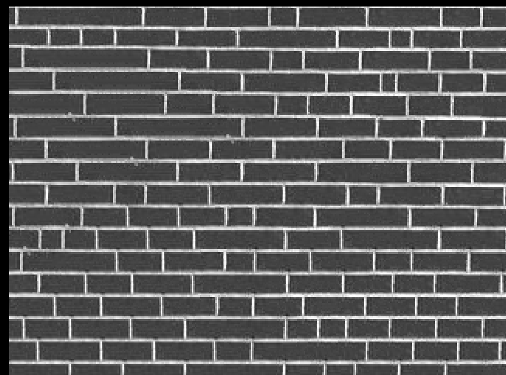

Increasing window size

Synthesis Results

french canvas

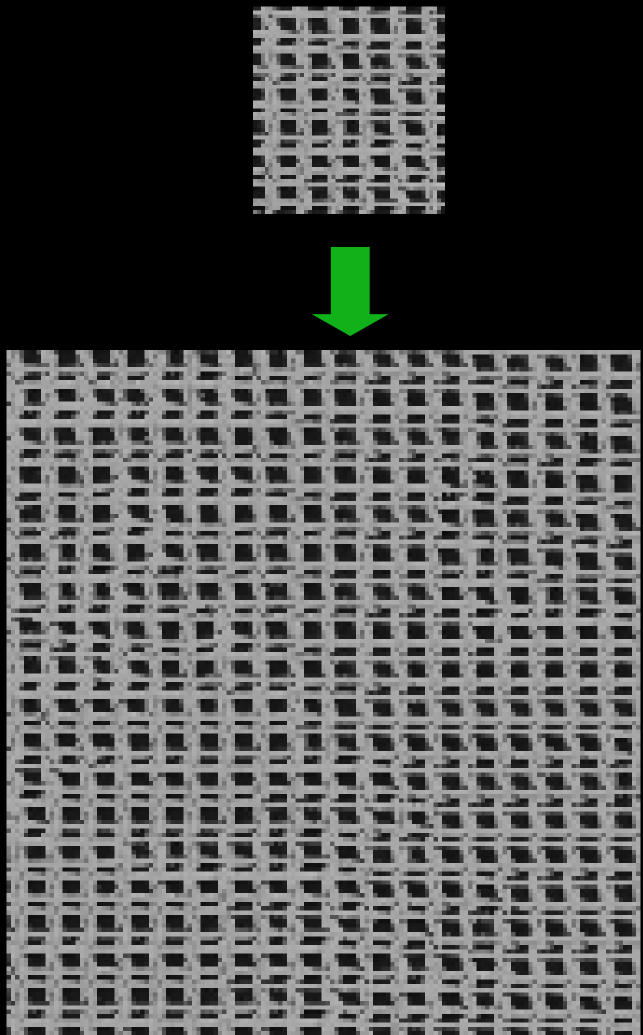

rafia weave

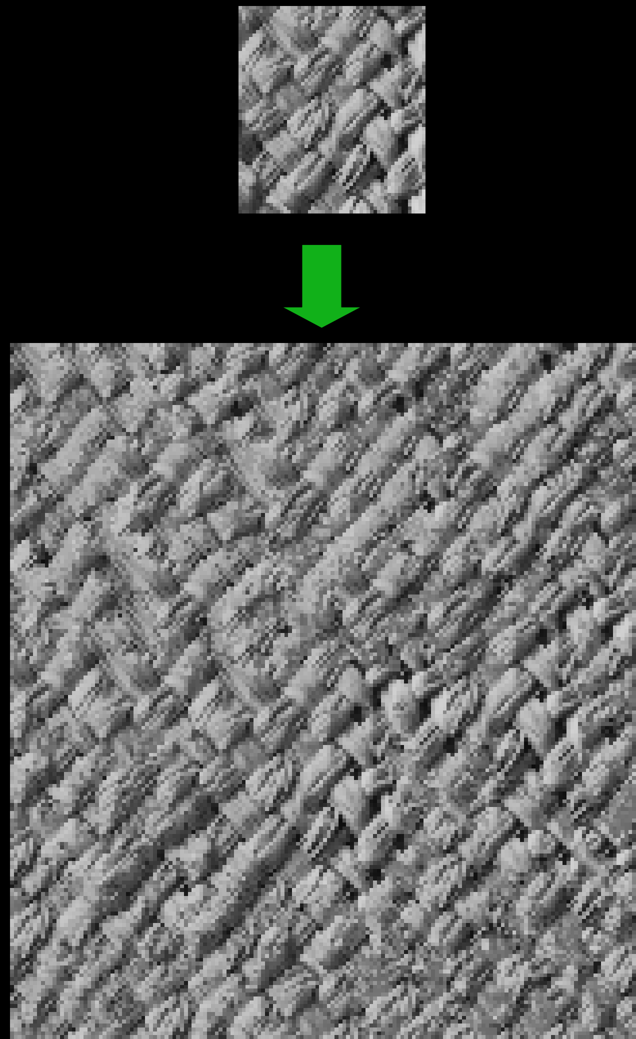

More Results

white bread

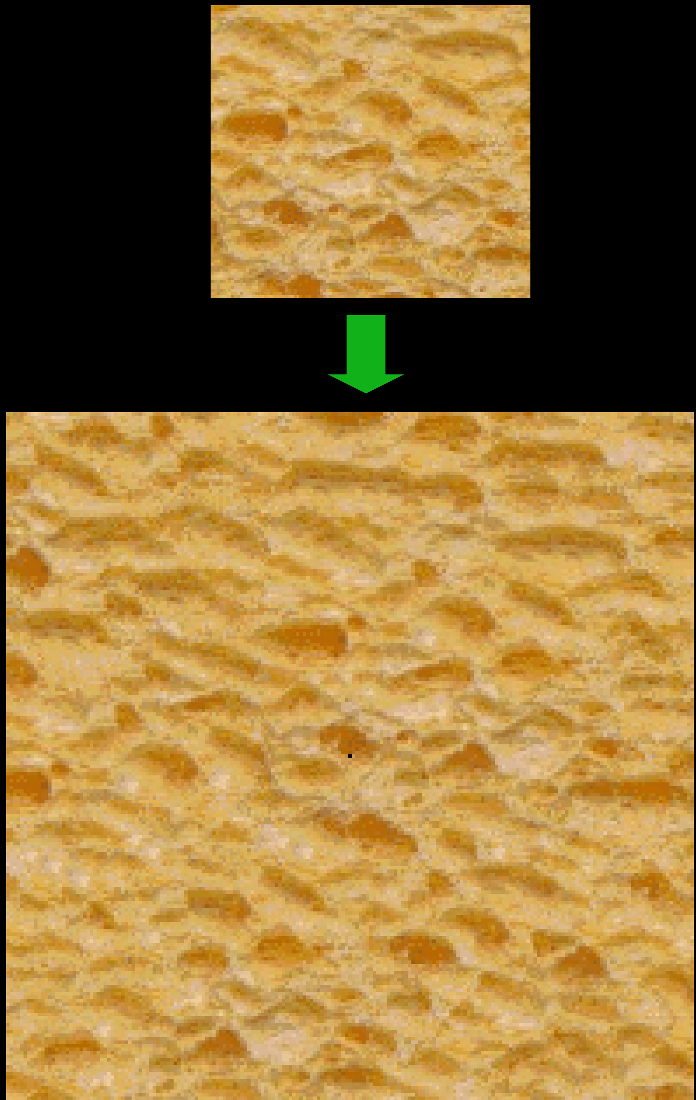

brick wall

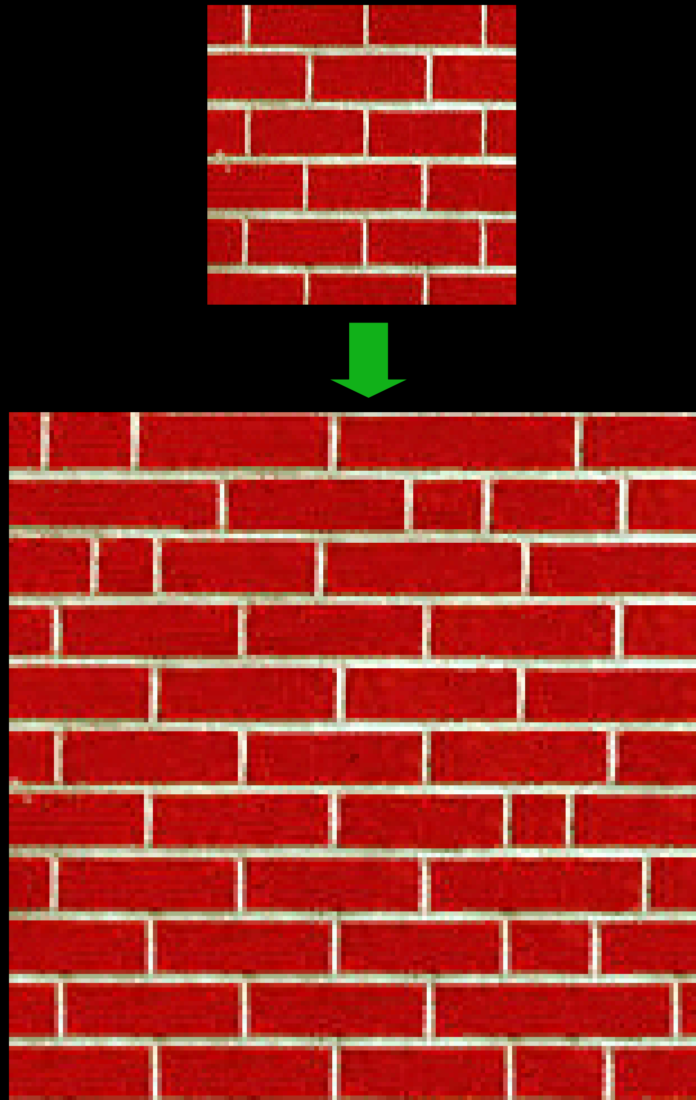

Homage to Shannon

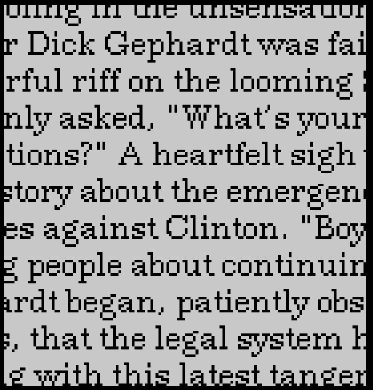

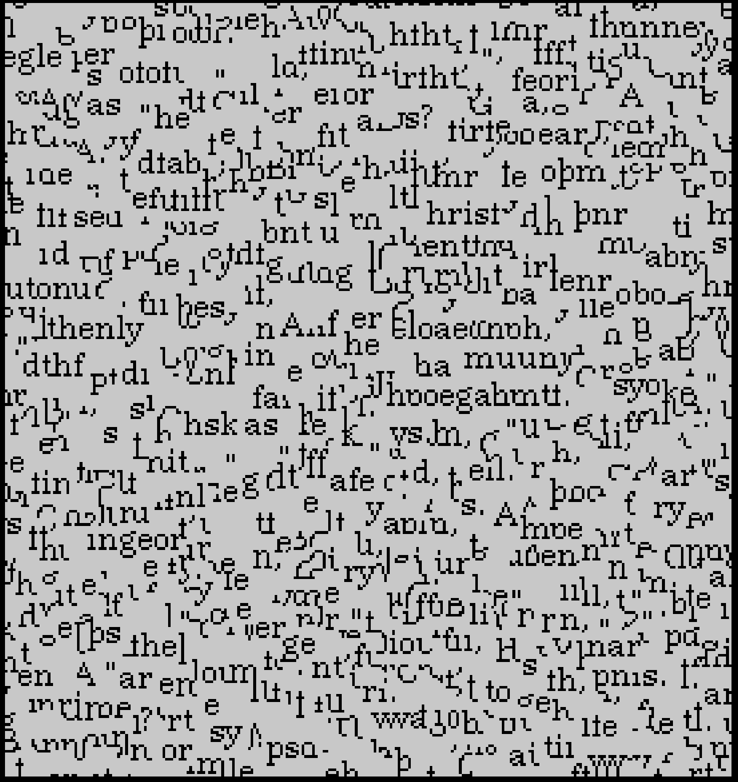

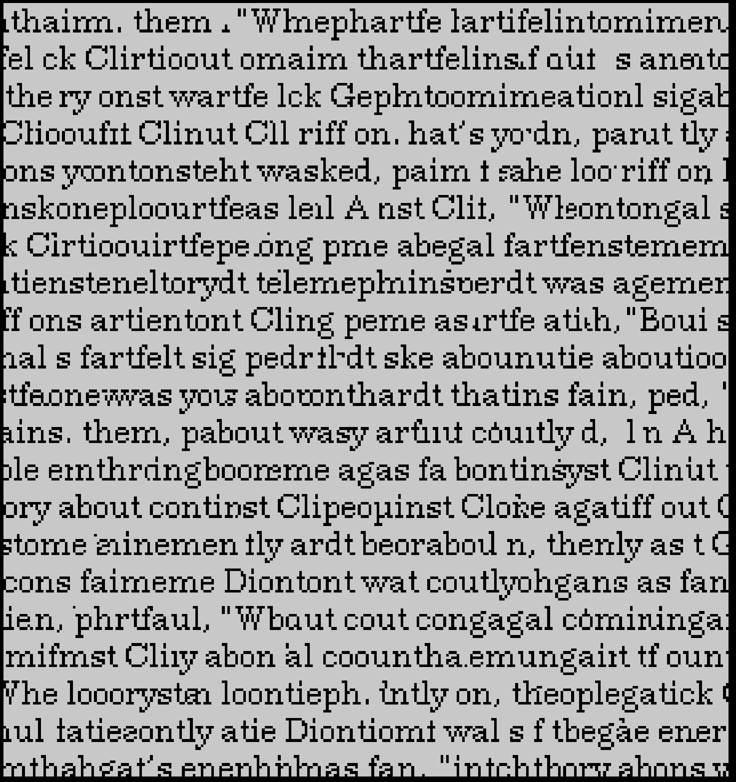

Hole Filling

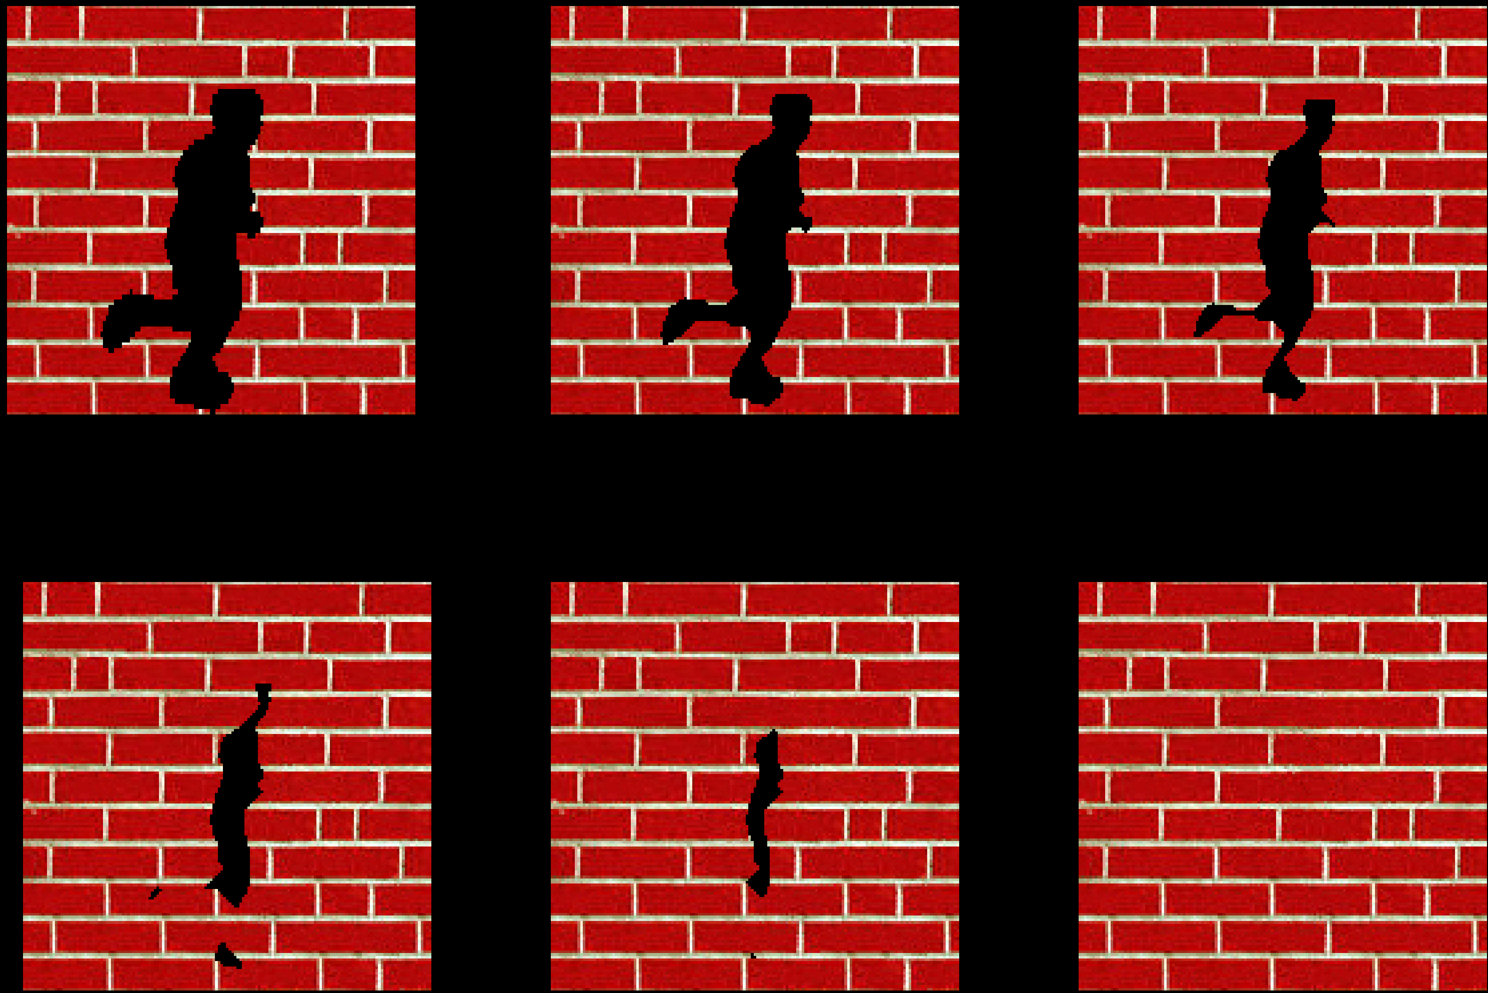

Extrapolation

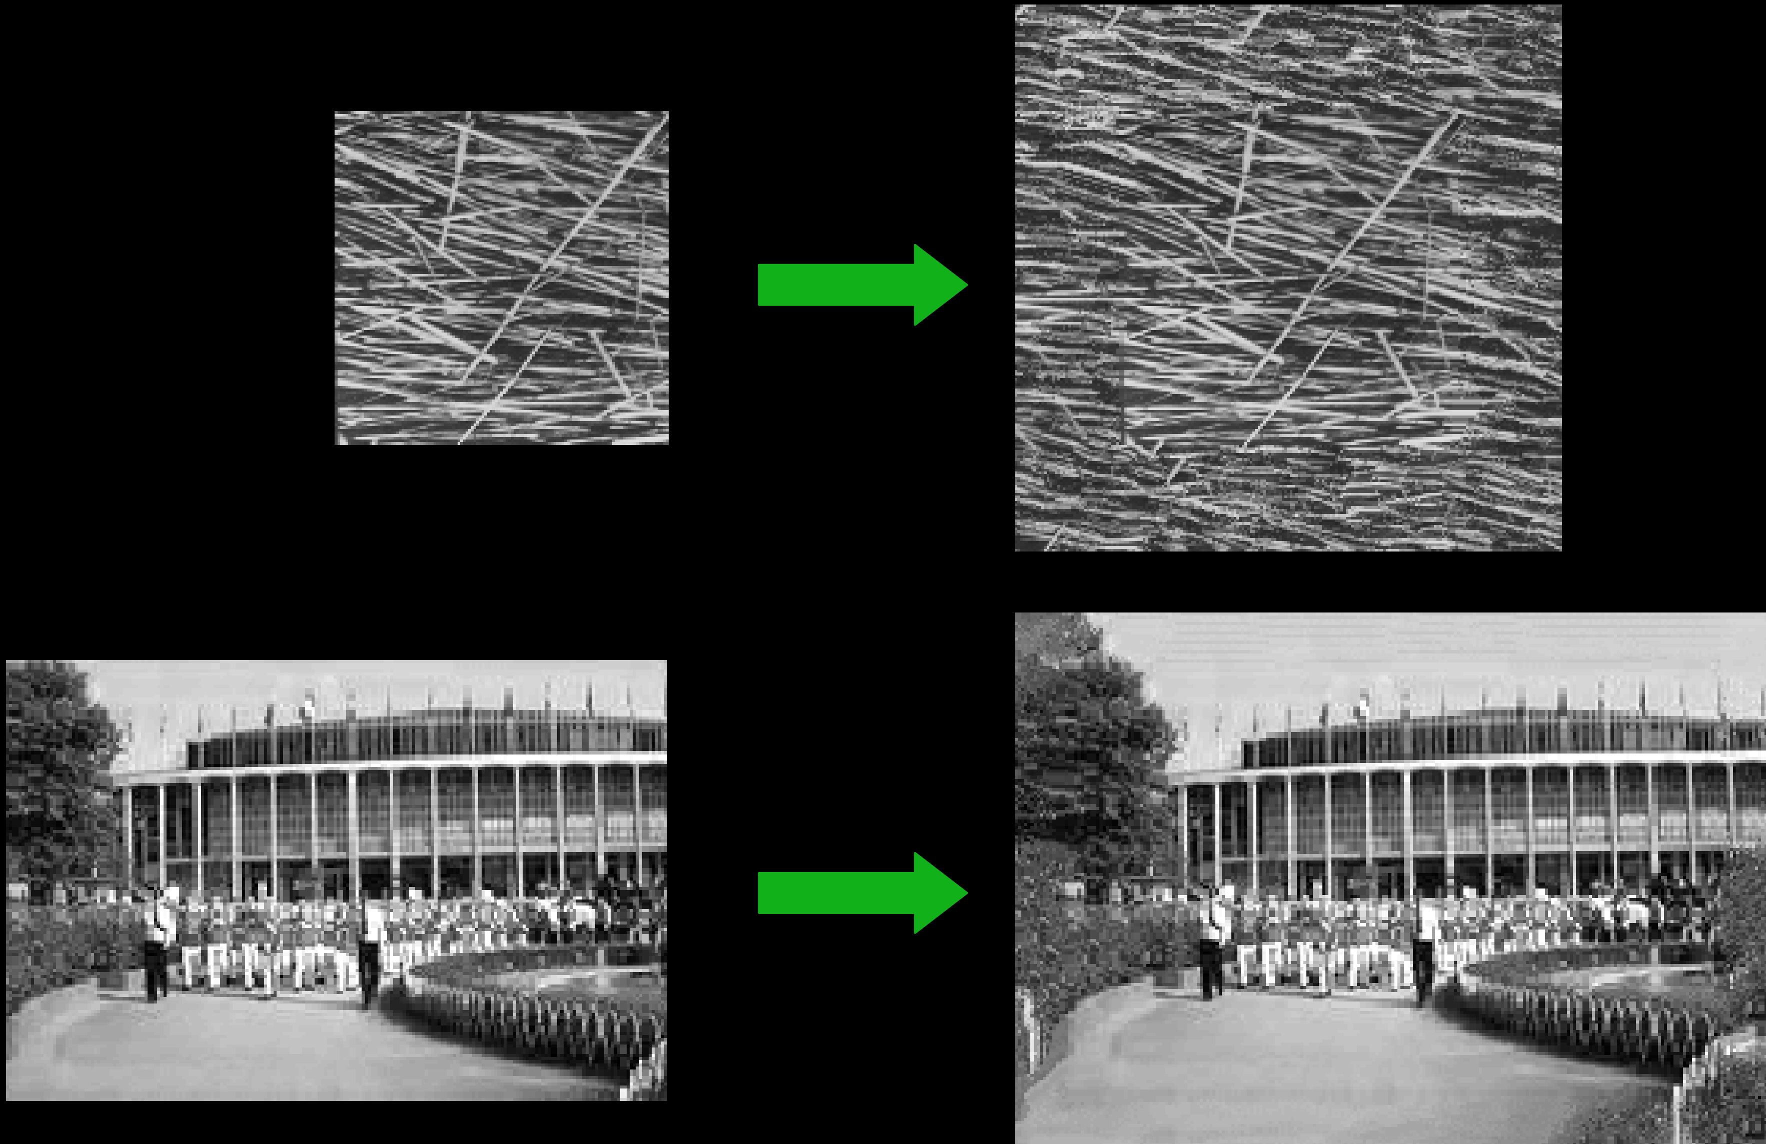

Summary

- The Efros & Leung algorithm
  - Very simple
  - Surprisingly good results
  - Synthesis is easier than analysis!
  - …but very slow

Image Quilting [Efros & Freeman]

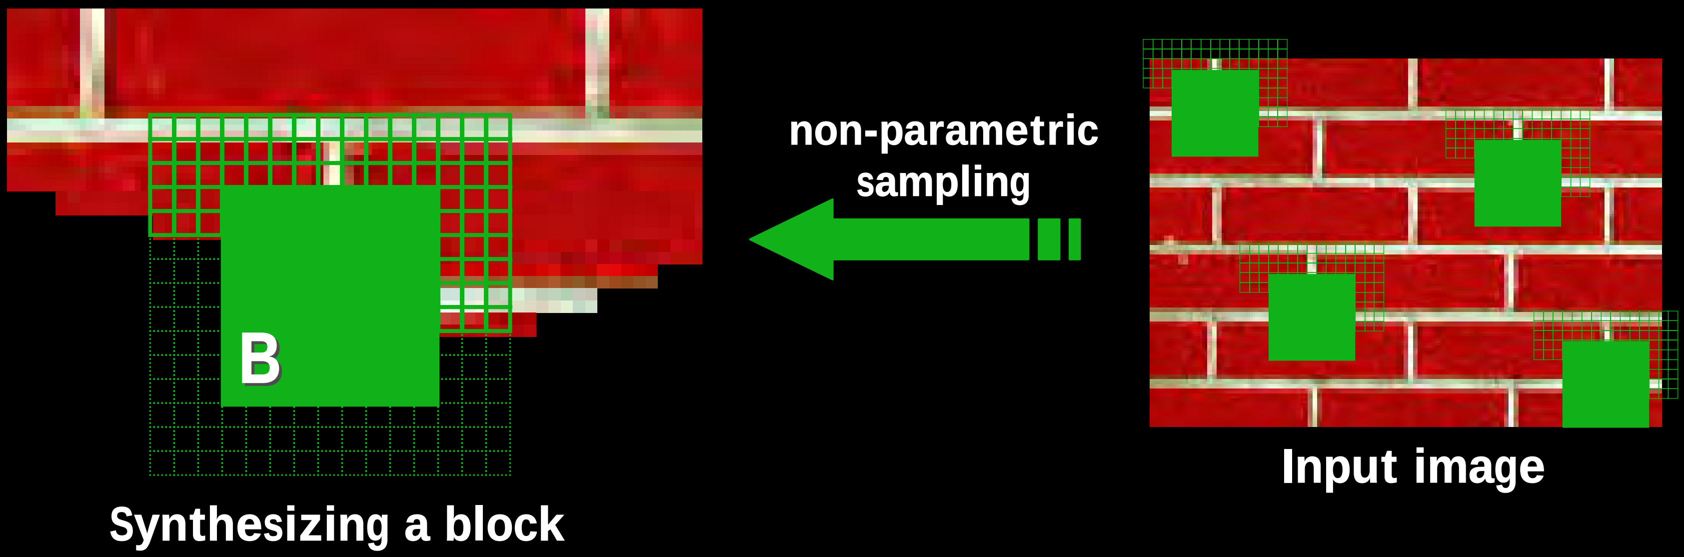

- Observation: neighbor pixels are highly correlated

Idea: unit of synthesis = block

- Exactly the same but now we want P(**B**|N(**B**))
- Much faster: synthesize all pixels in a block at once
- Not the same as multi-scale!

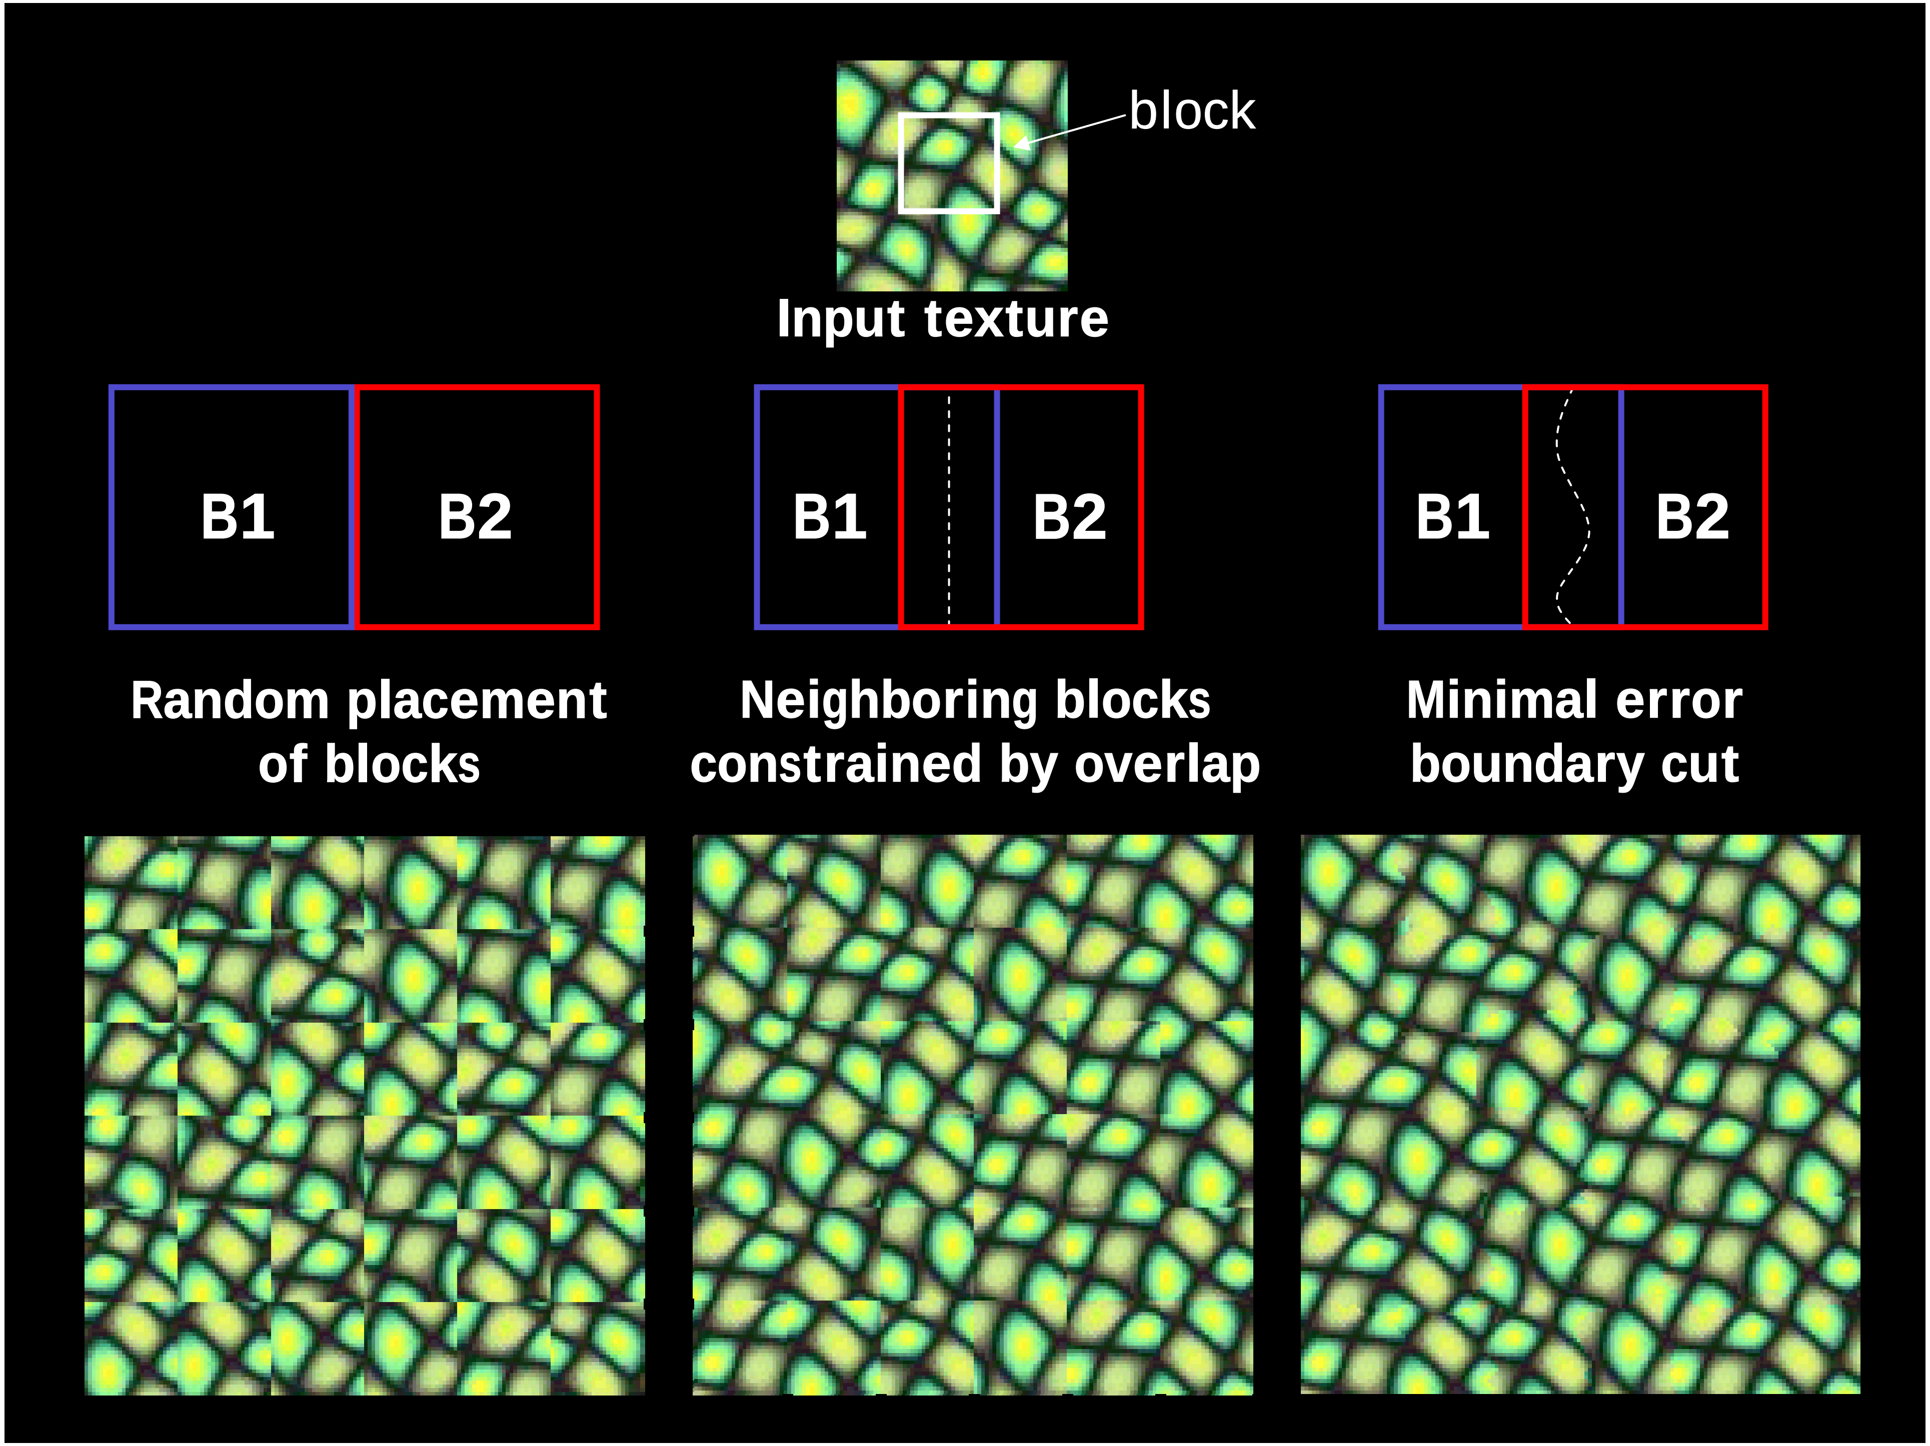
block
Input texture
B1
B2
B1
B2
B1
B2
Random placement of blocks
Neighboring blocks constrained by overlap
Minimal error boundary cut

Minimal error boundary

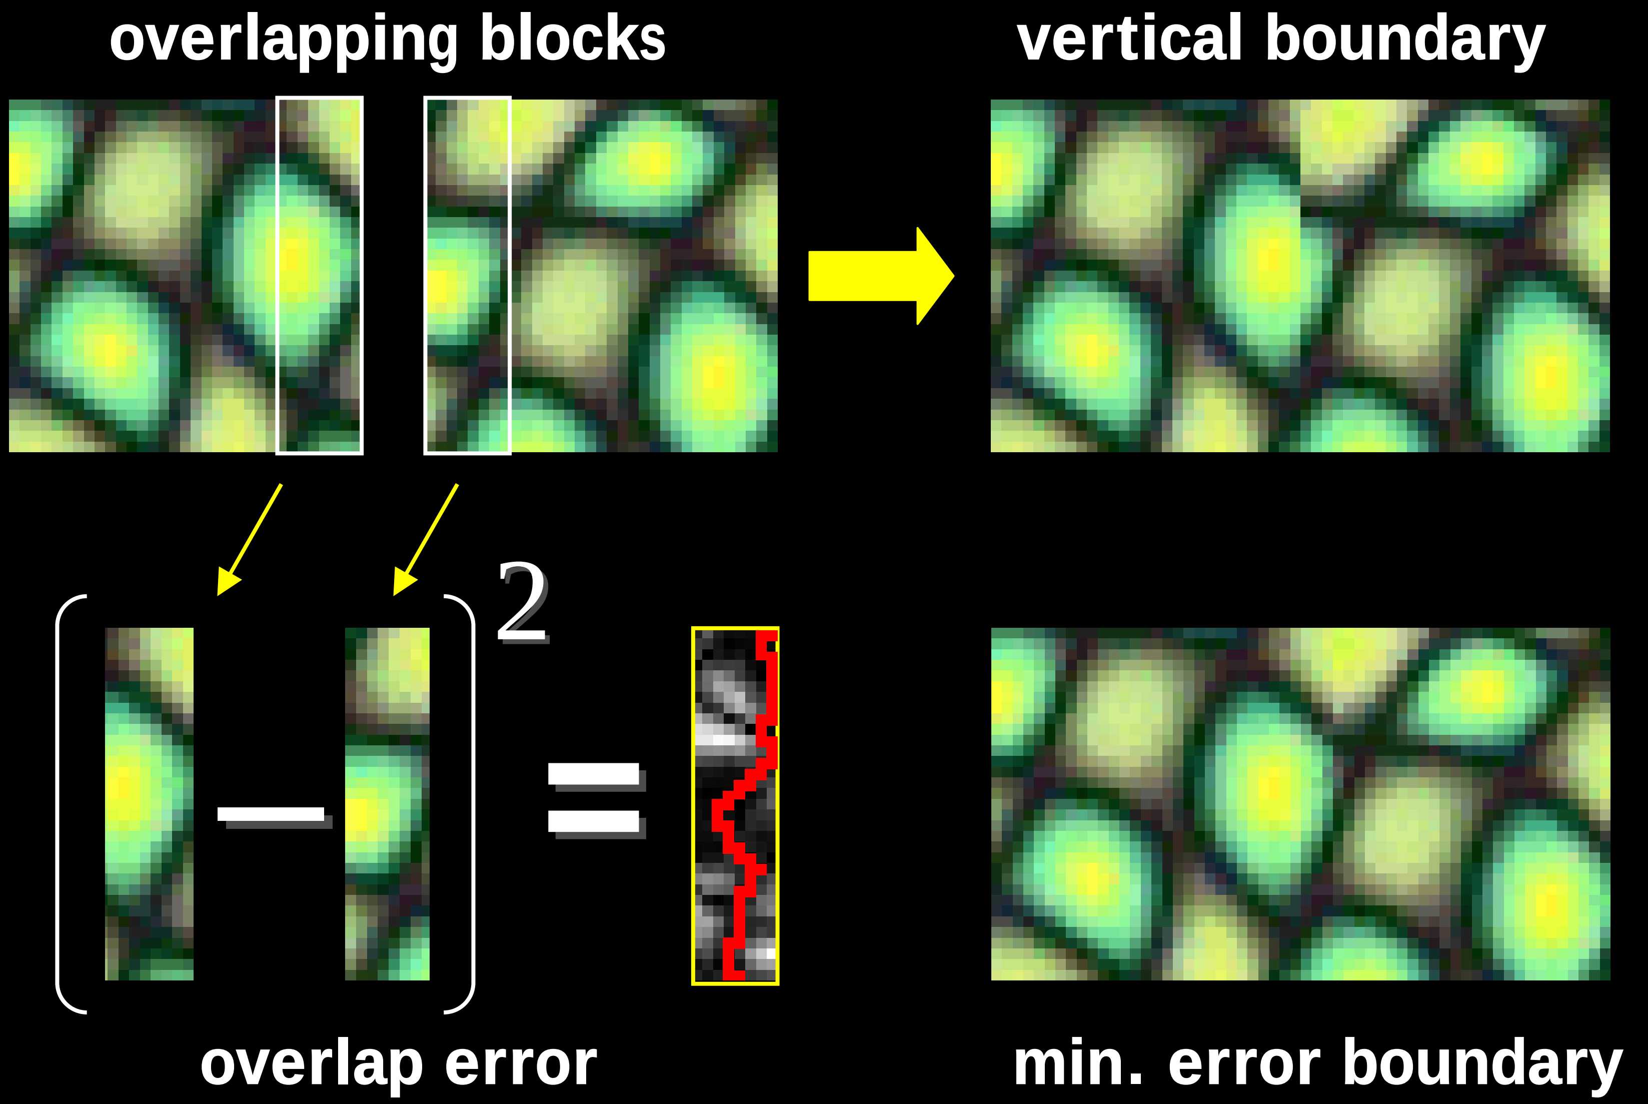

Our Philosophy

- The “Corrupt Professor’s Algorithm”:
  - Plagiarize as much of the source image as you can
  - Then try to cover up the evidence
- Rationale:
  - Texture blocks are by definition correct samples of texture so problem only connecting them together

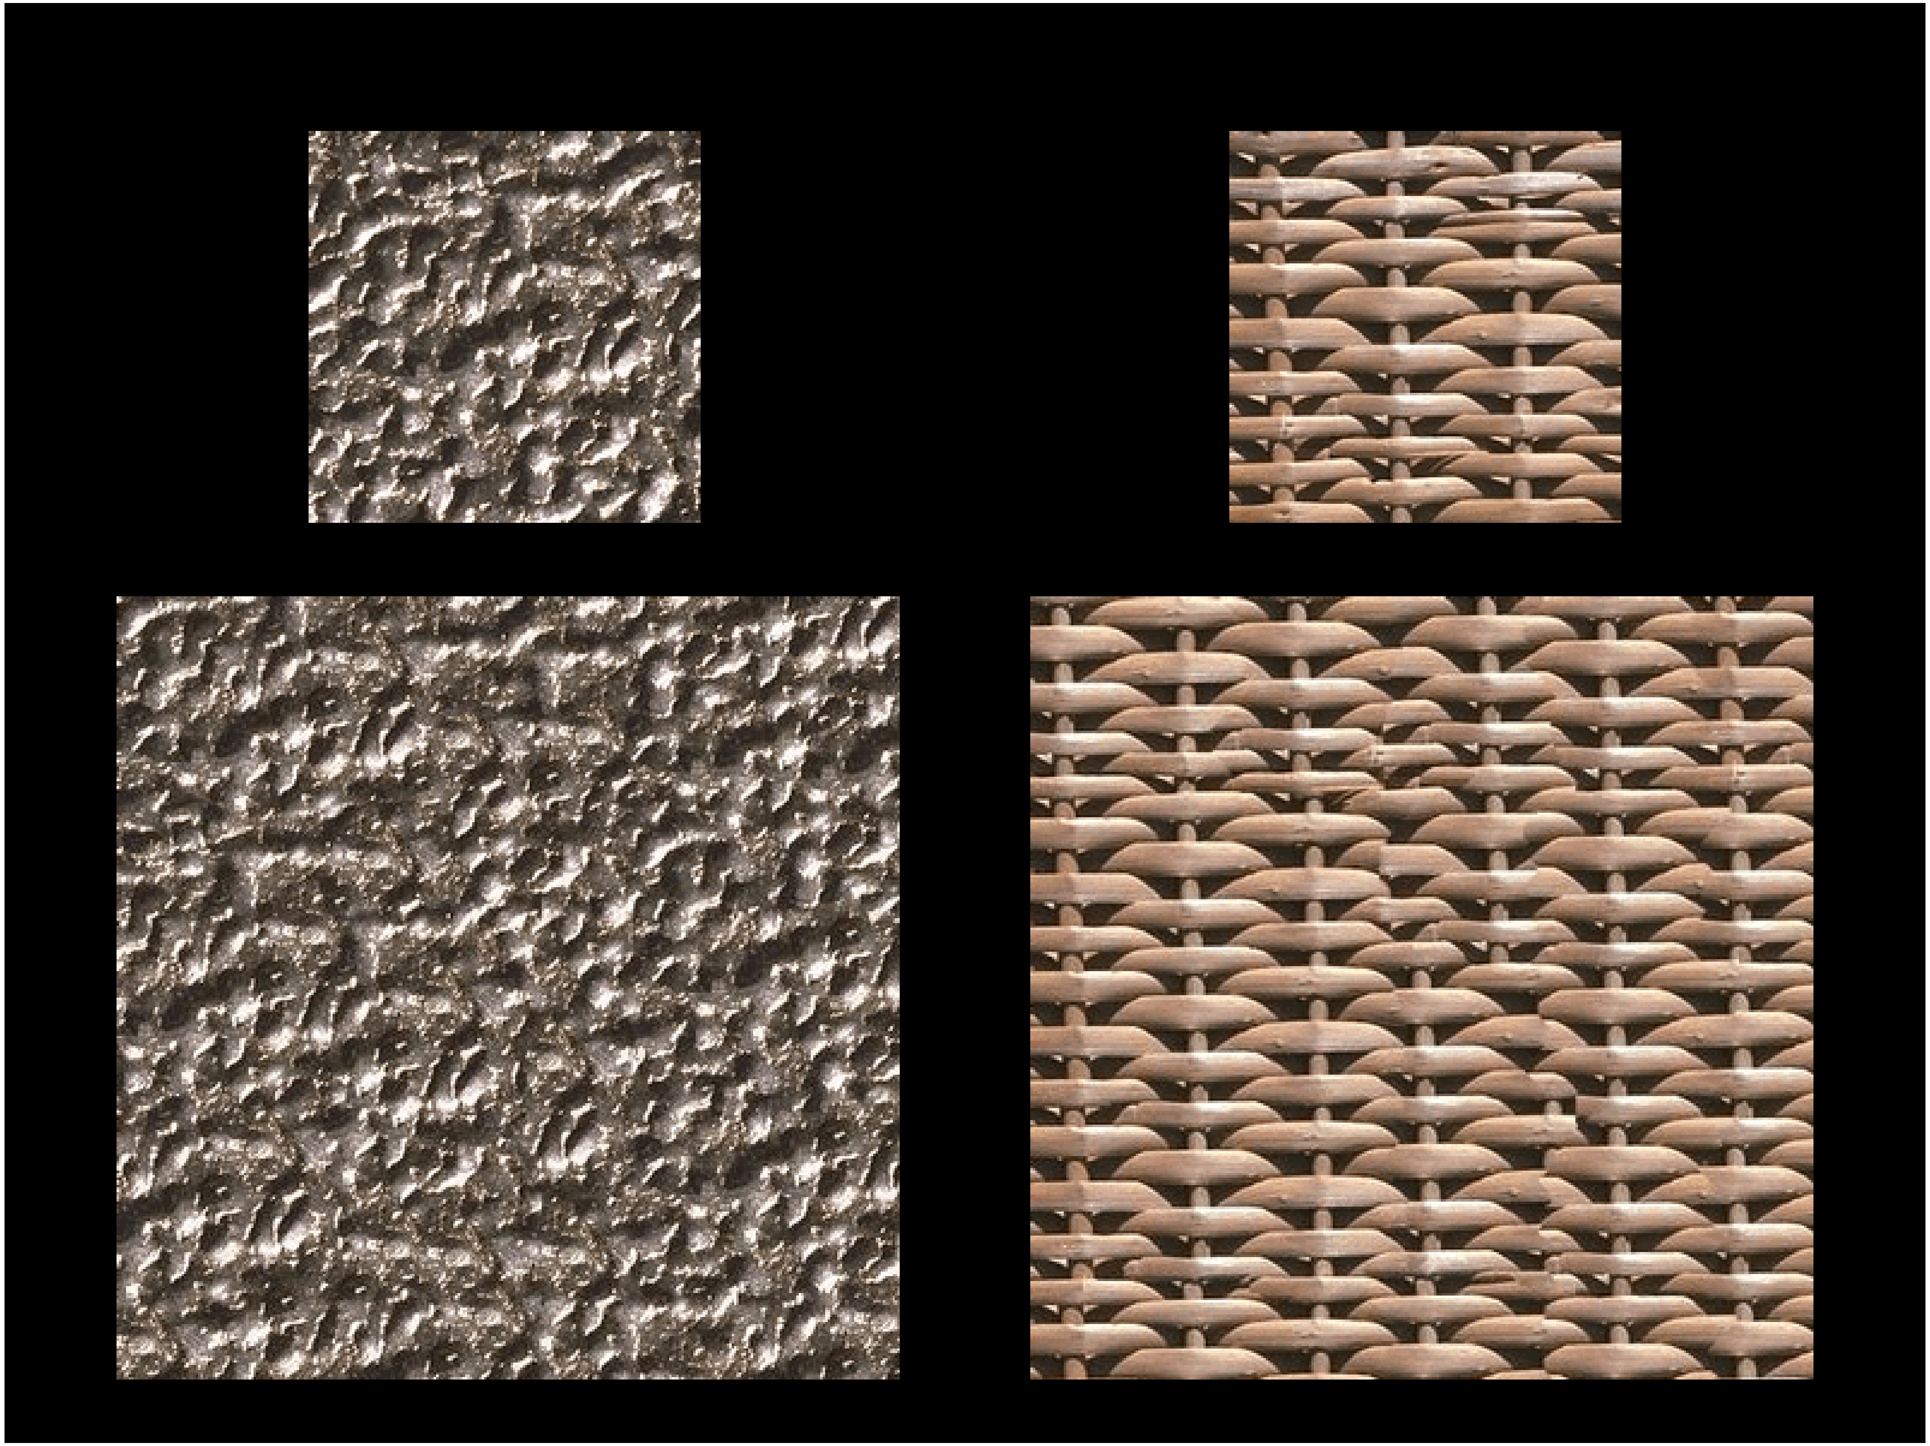

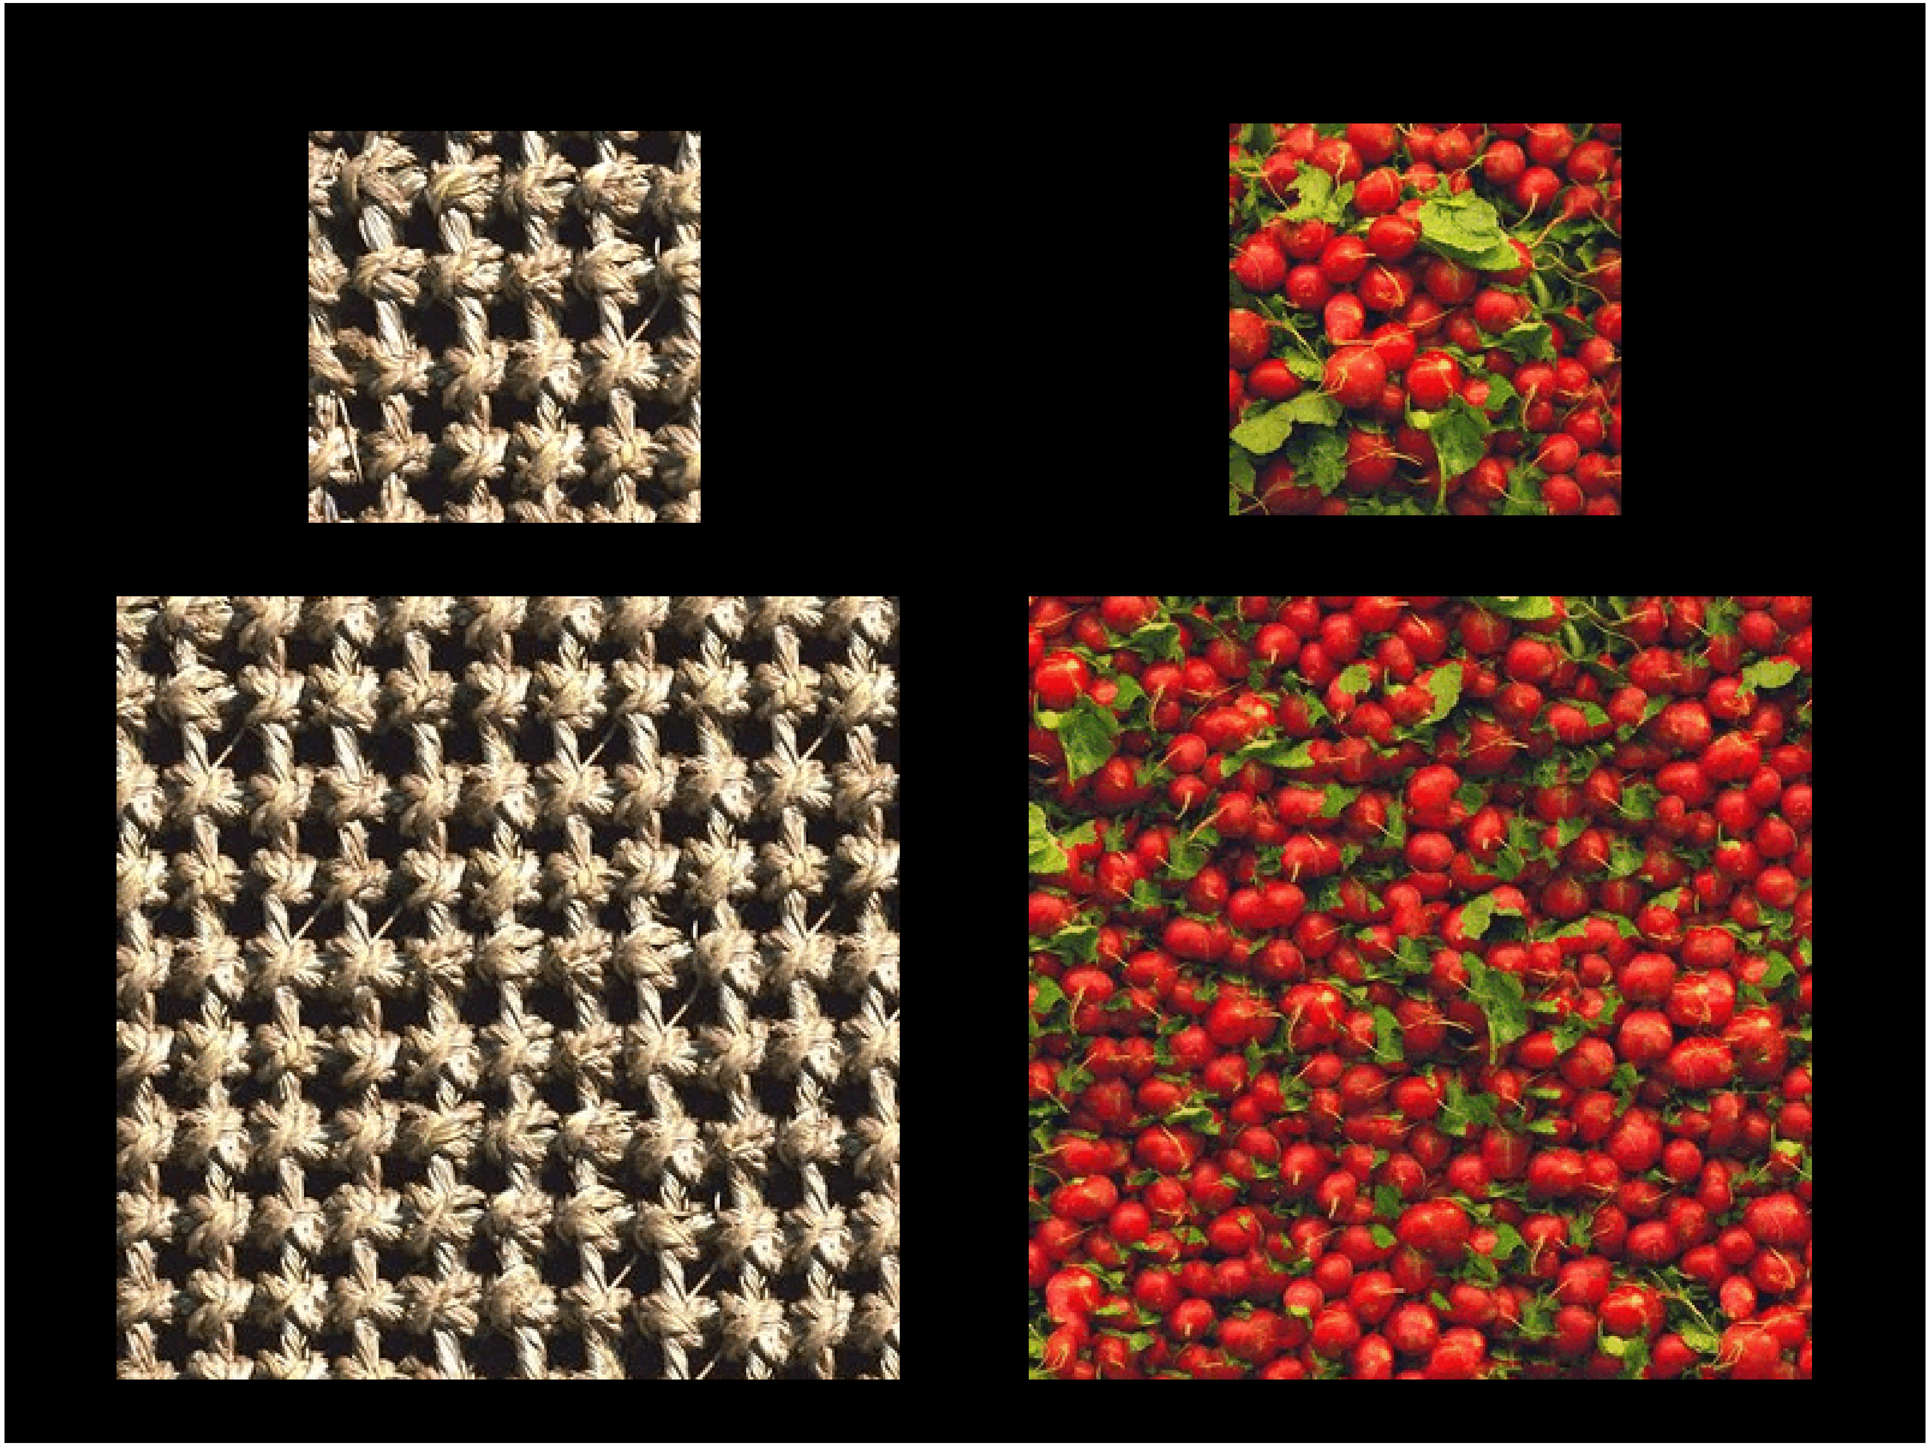




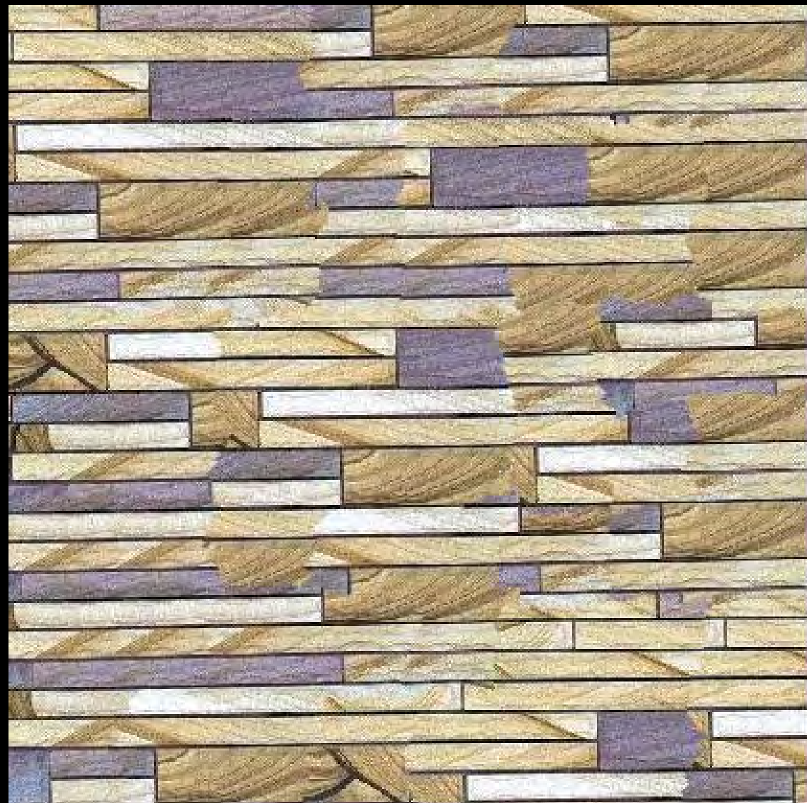


College Inn
Chicken Broth


College Inn
Chicken Broth


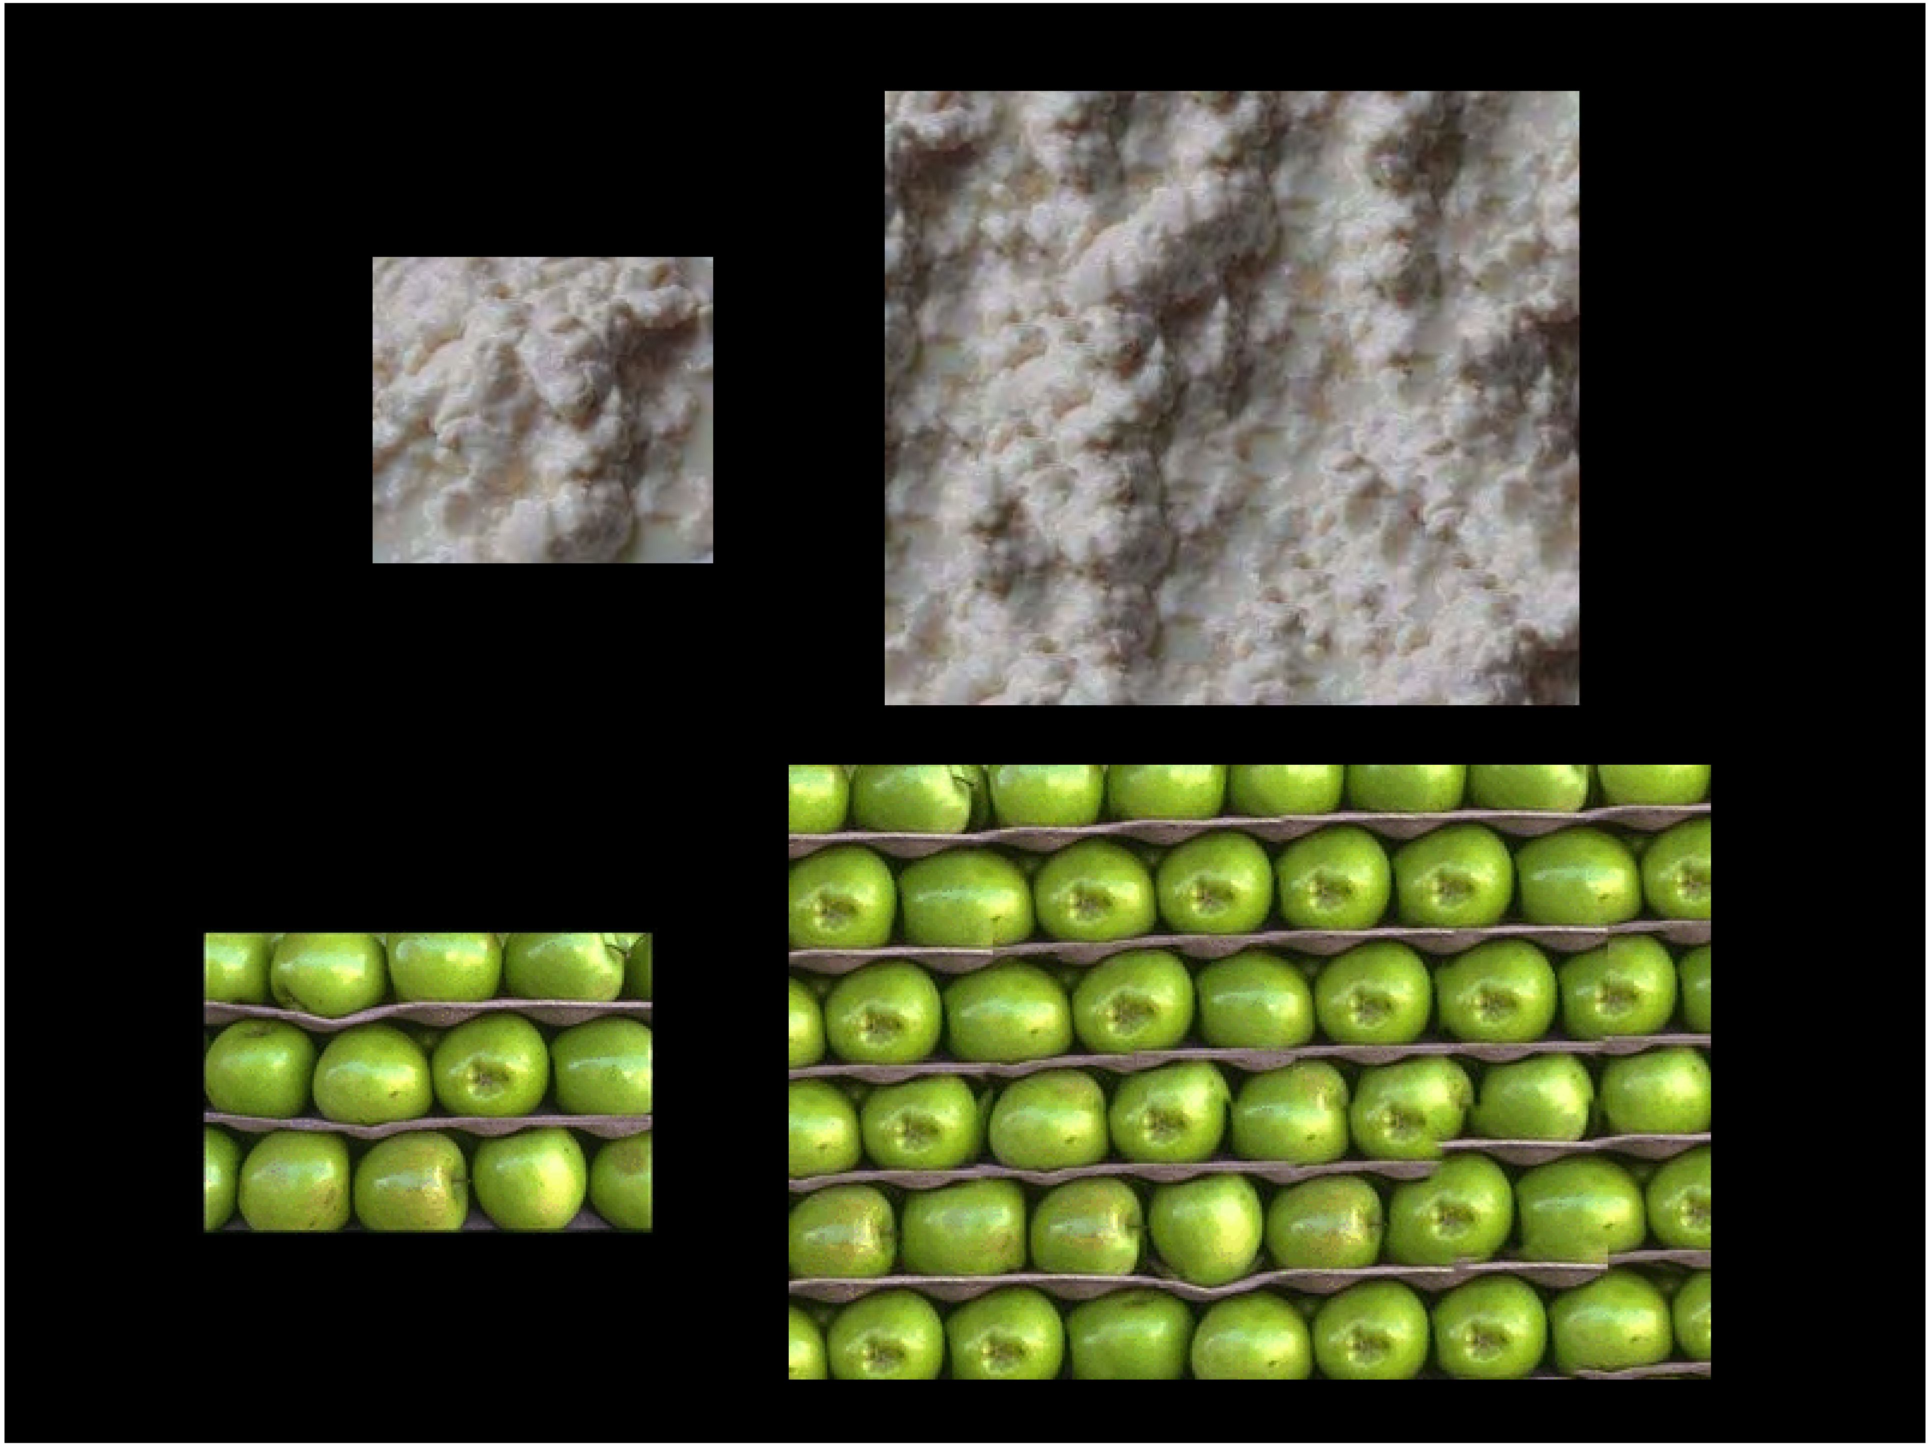

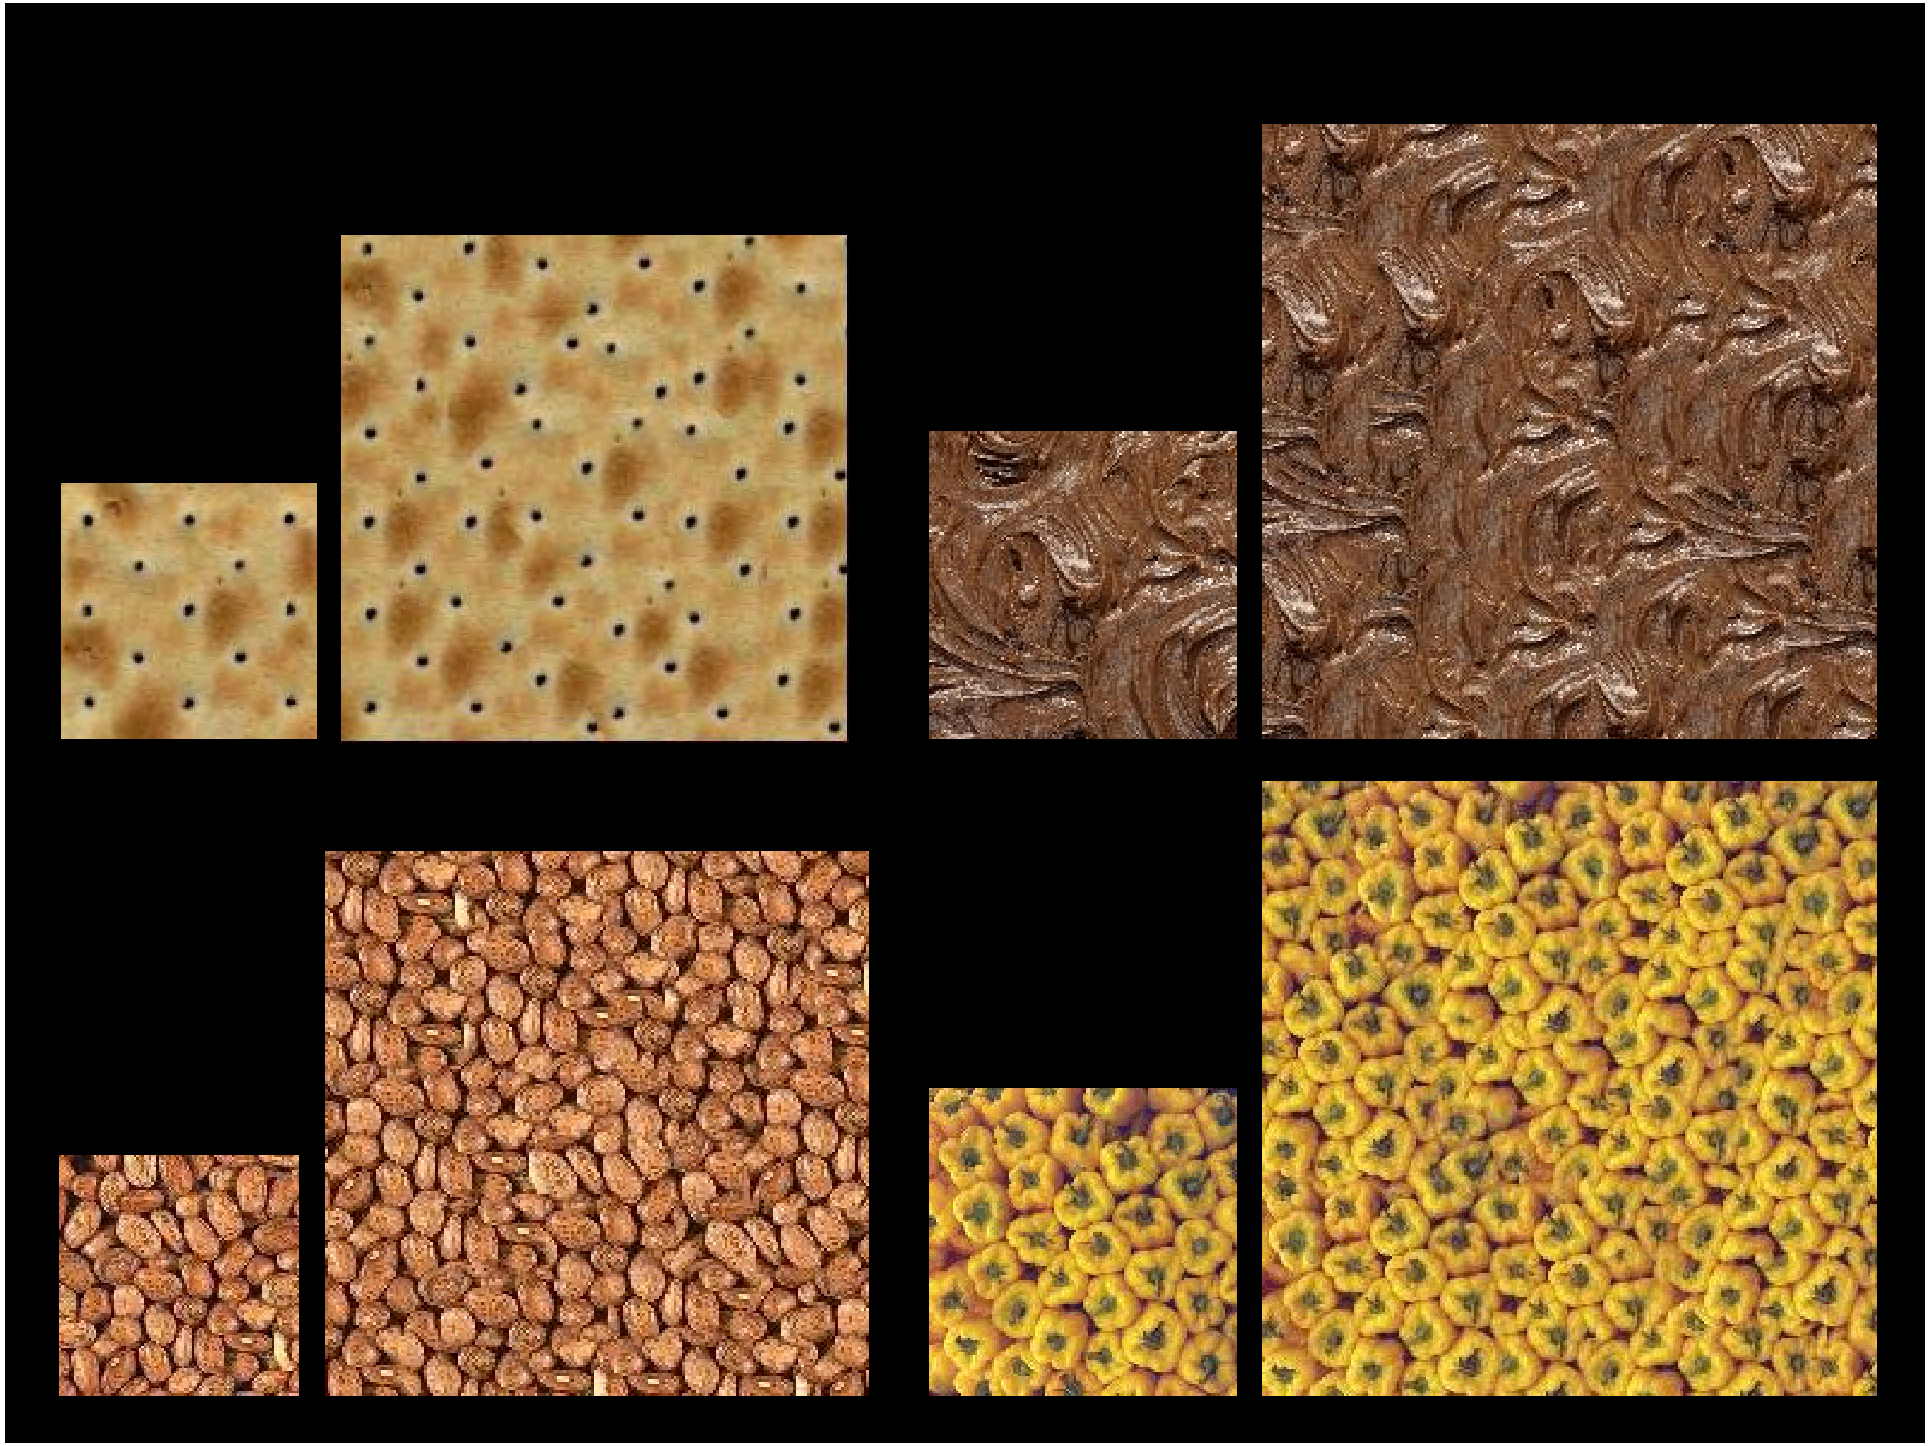

Failures
(Chernobyl Harvest)

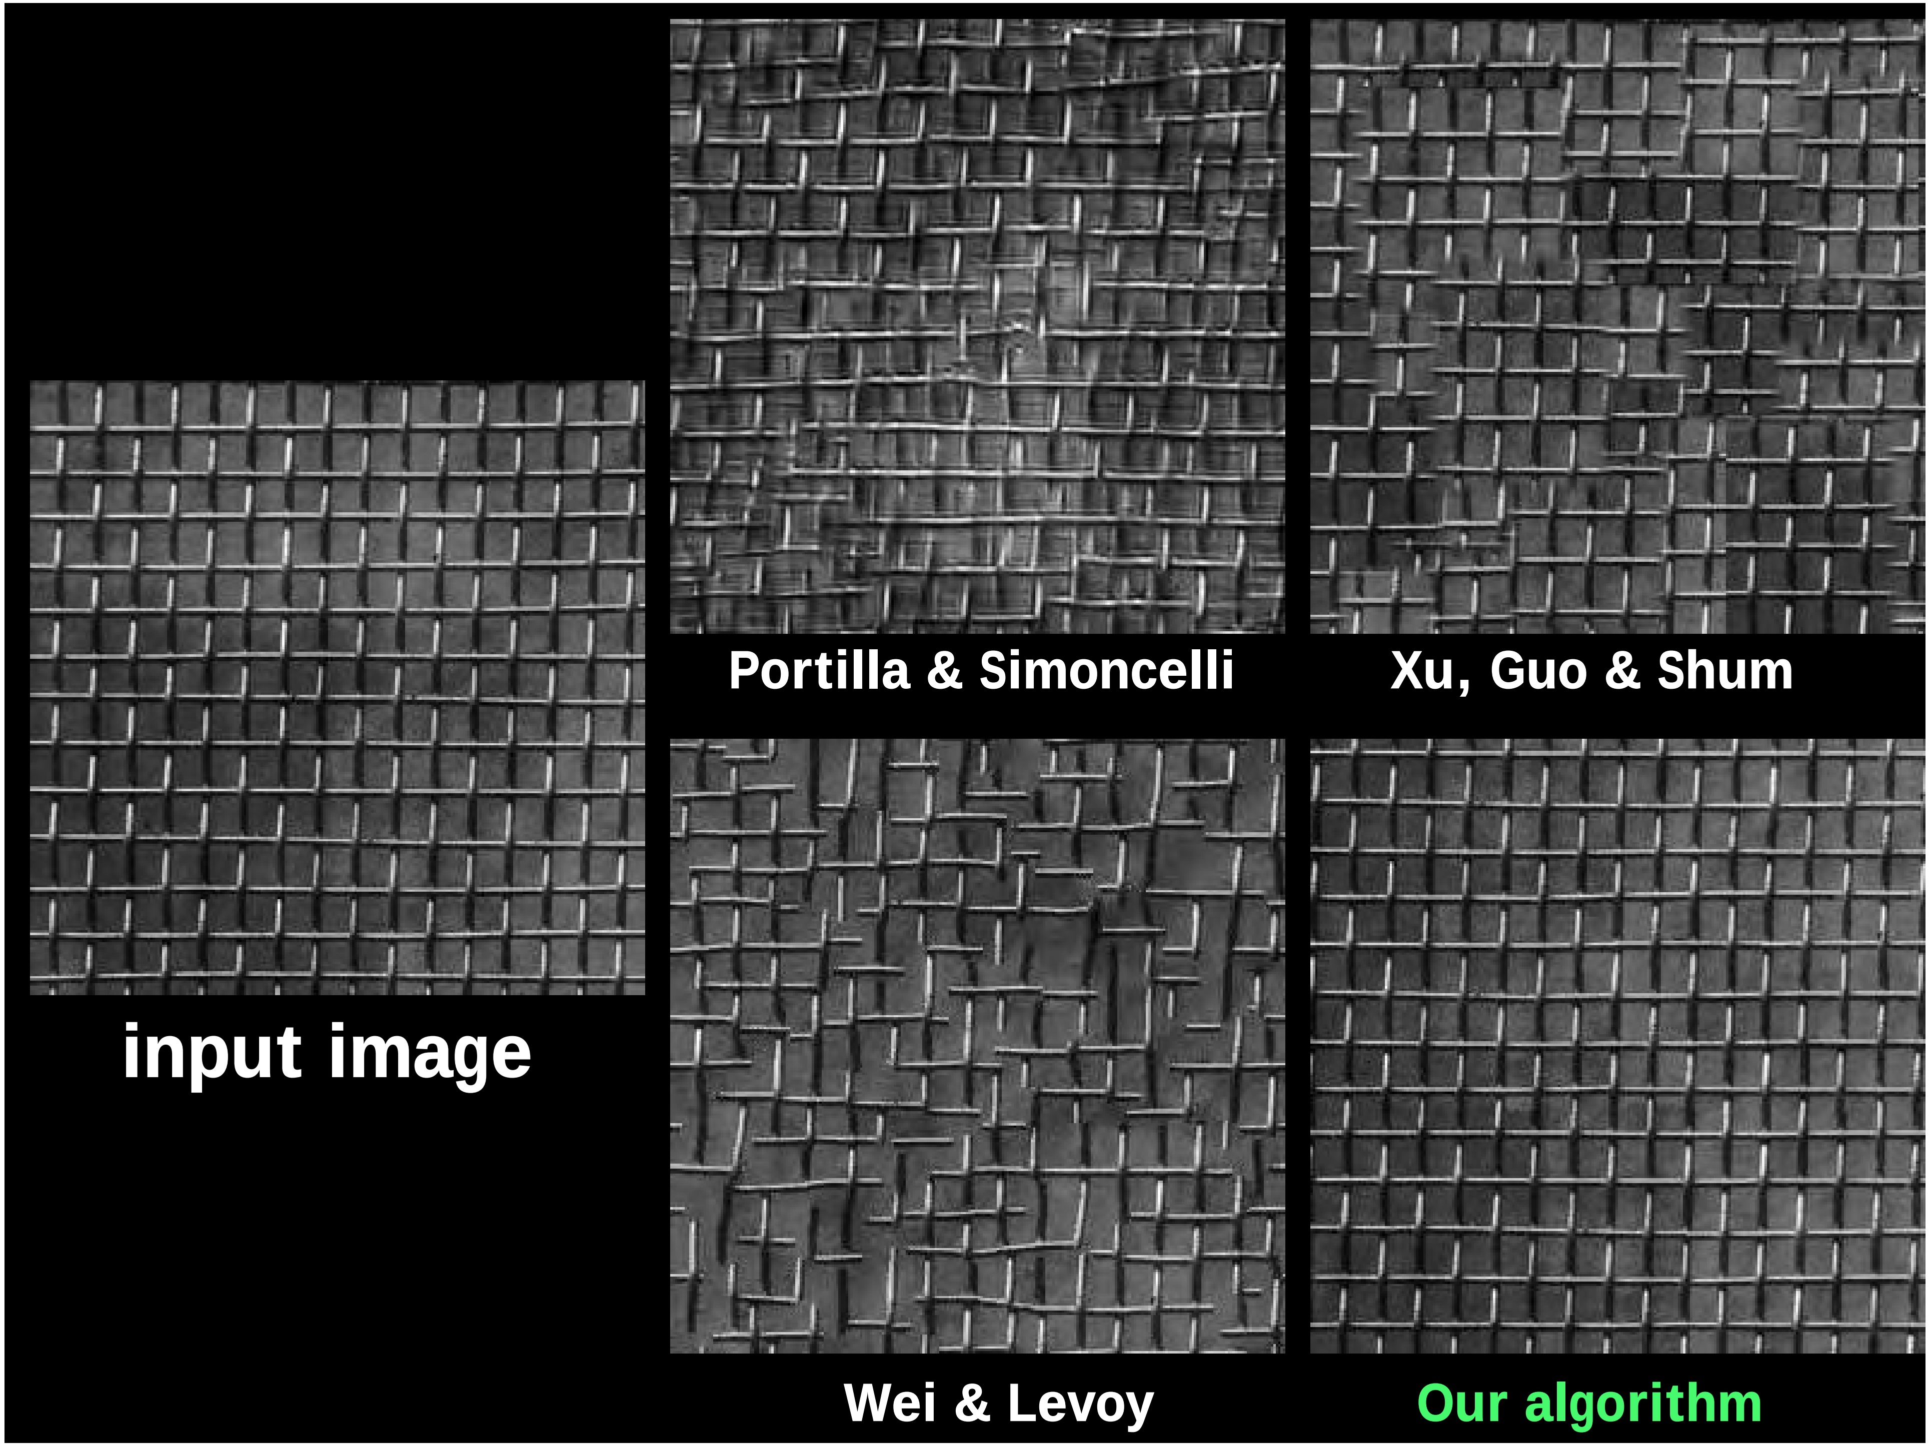
Portilla & Simoncelli
Xu, Guo & Shum
input image
Wei & Levoy
Our algorithm

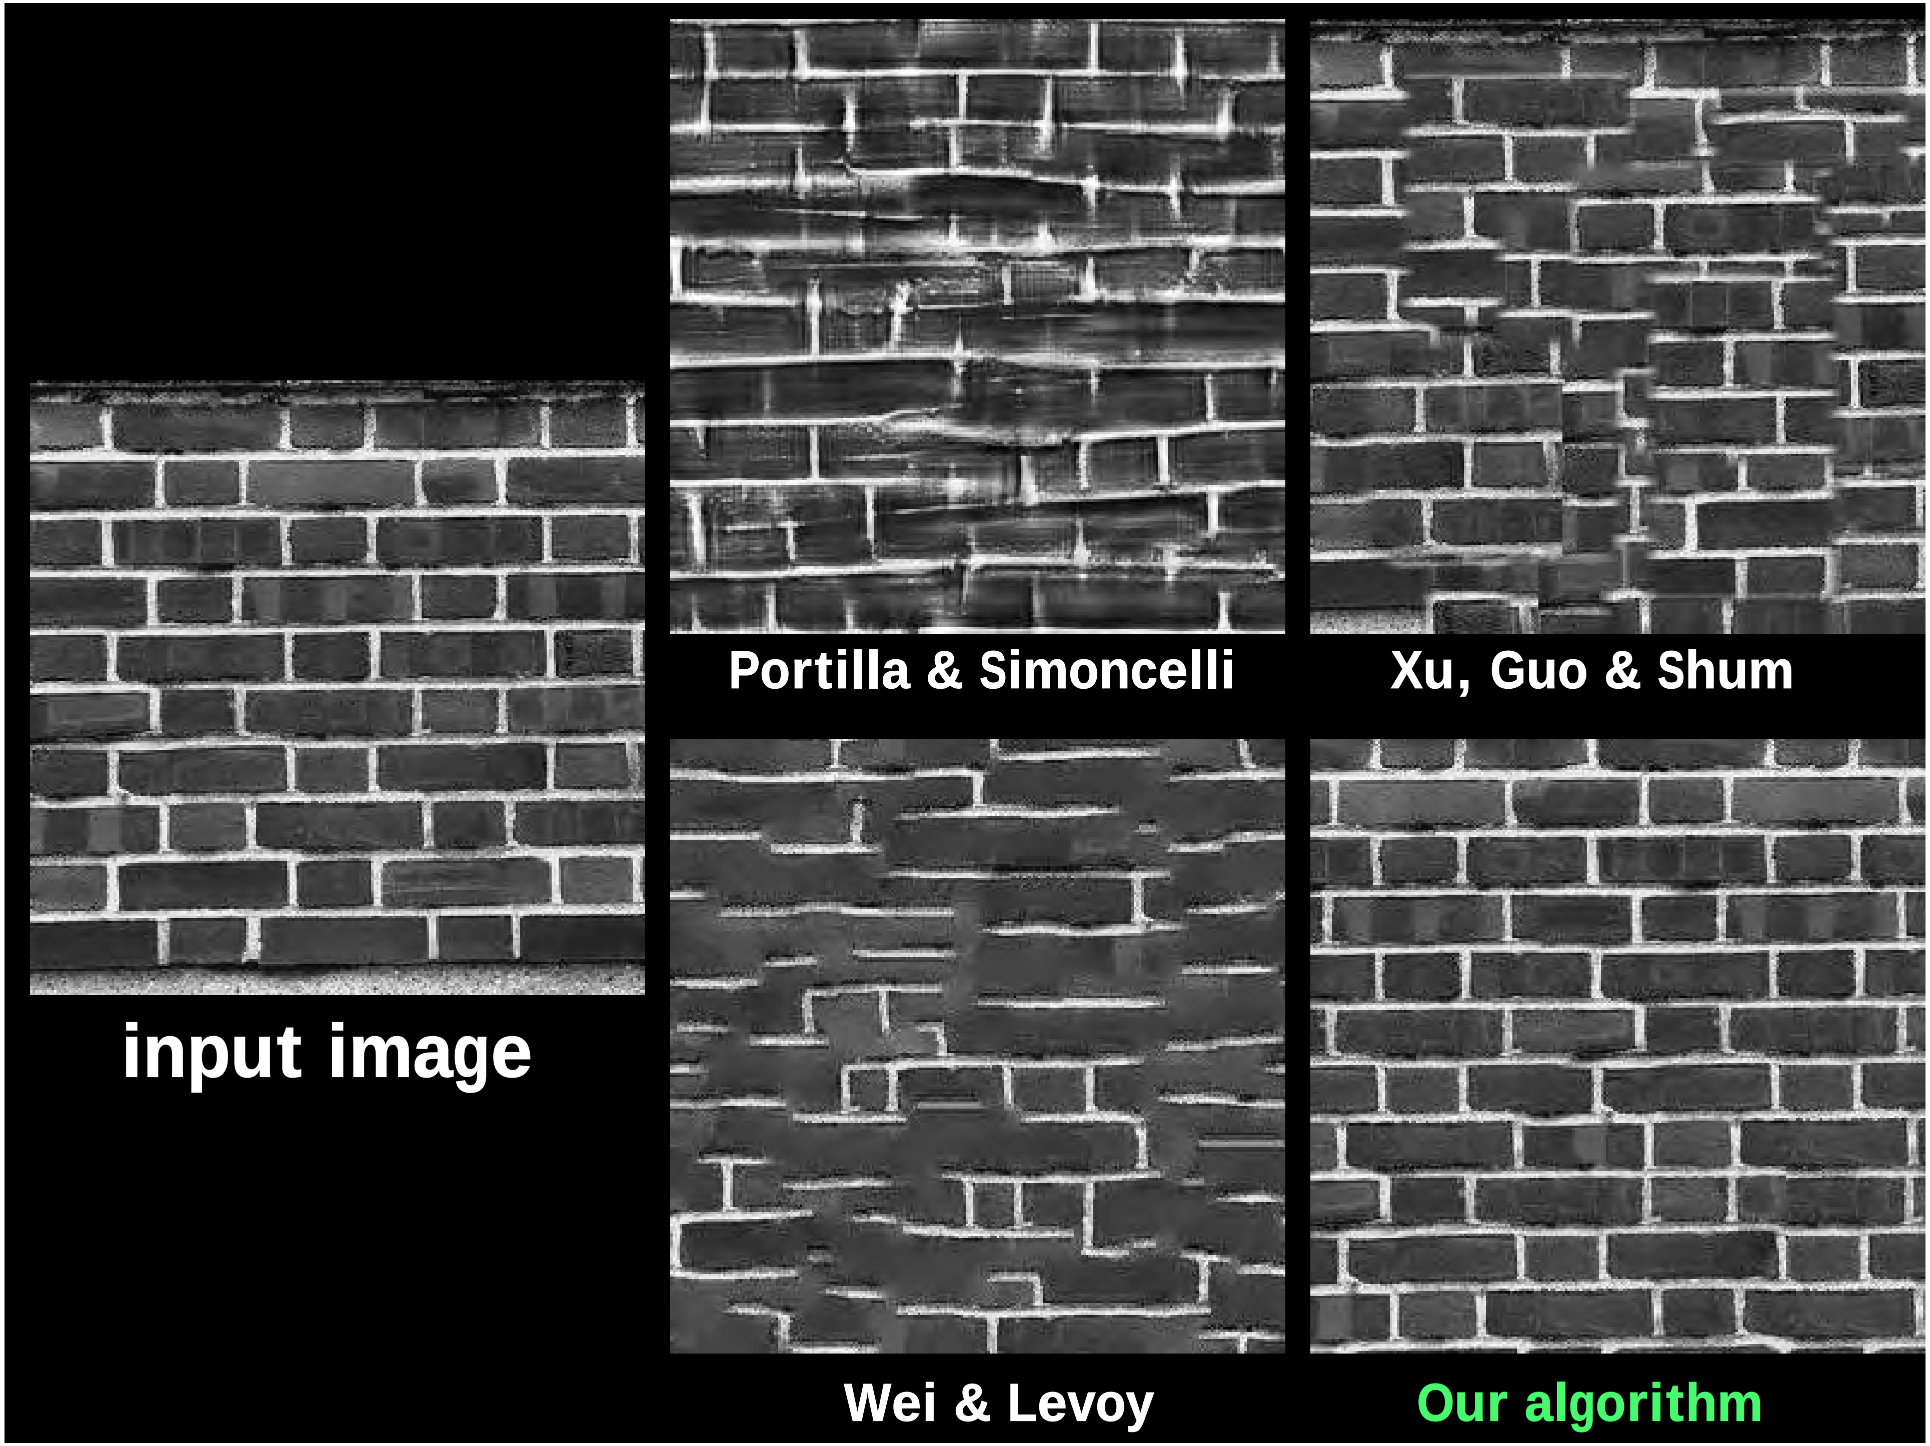
input image
Portilla & Simoncelli
Xu, Guo & Shum
Wei & Levoy
Our algorithm

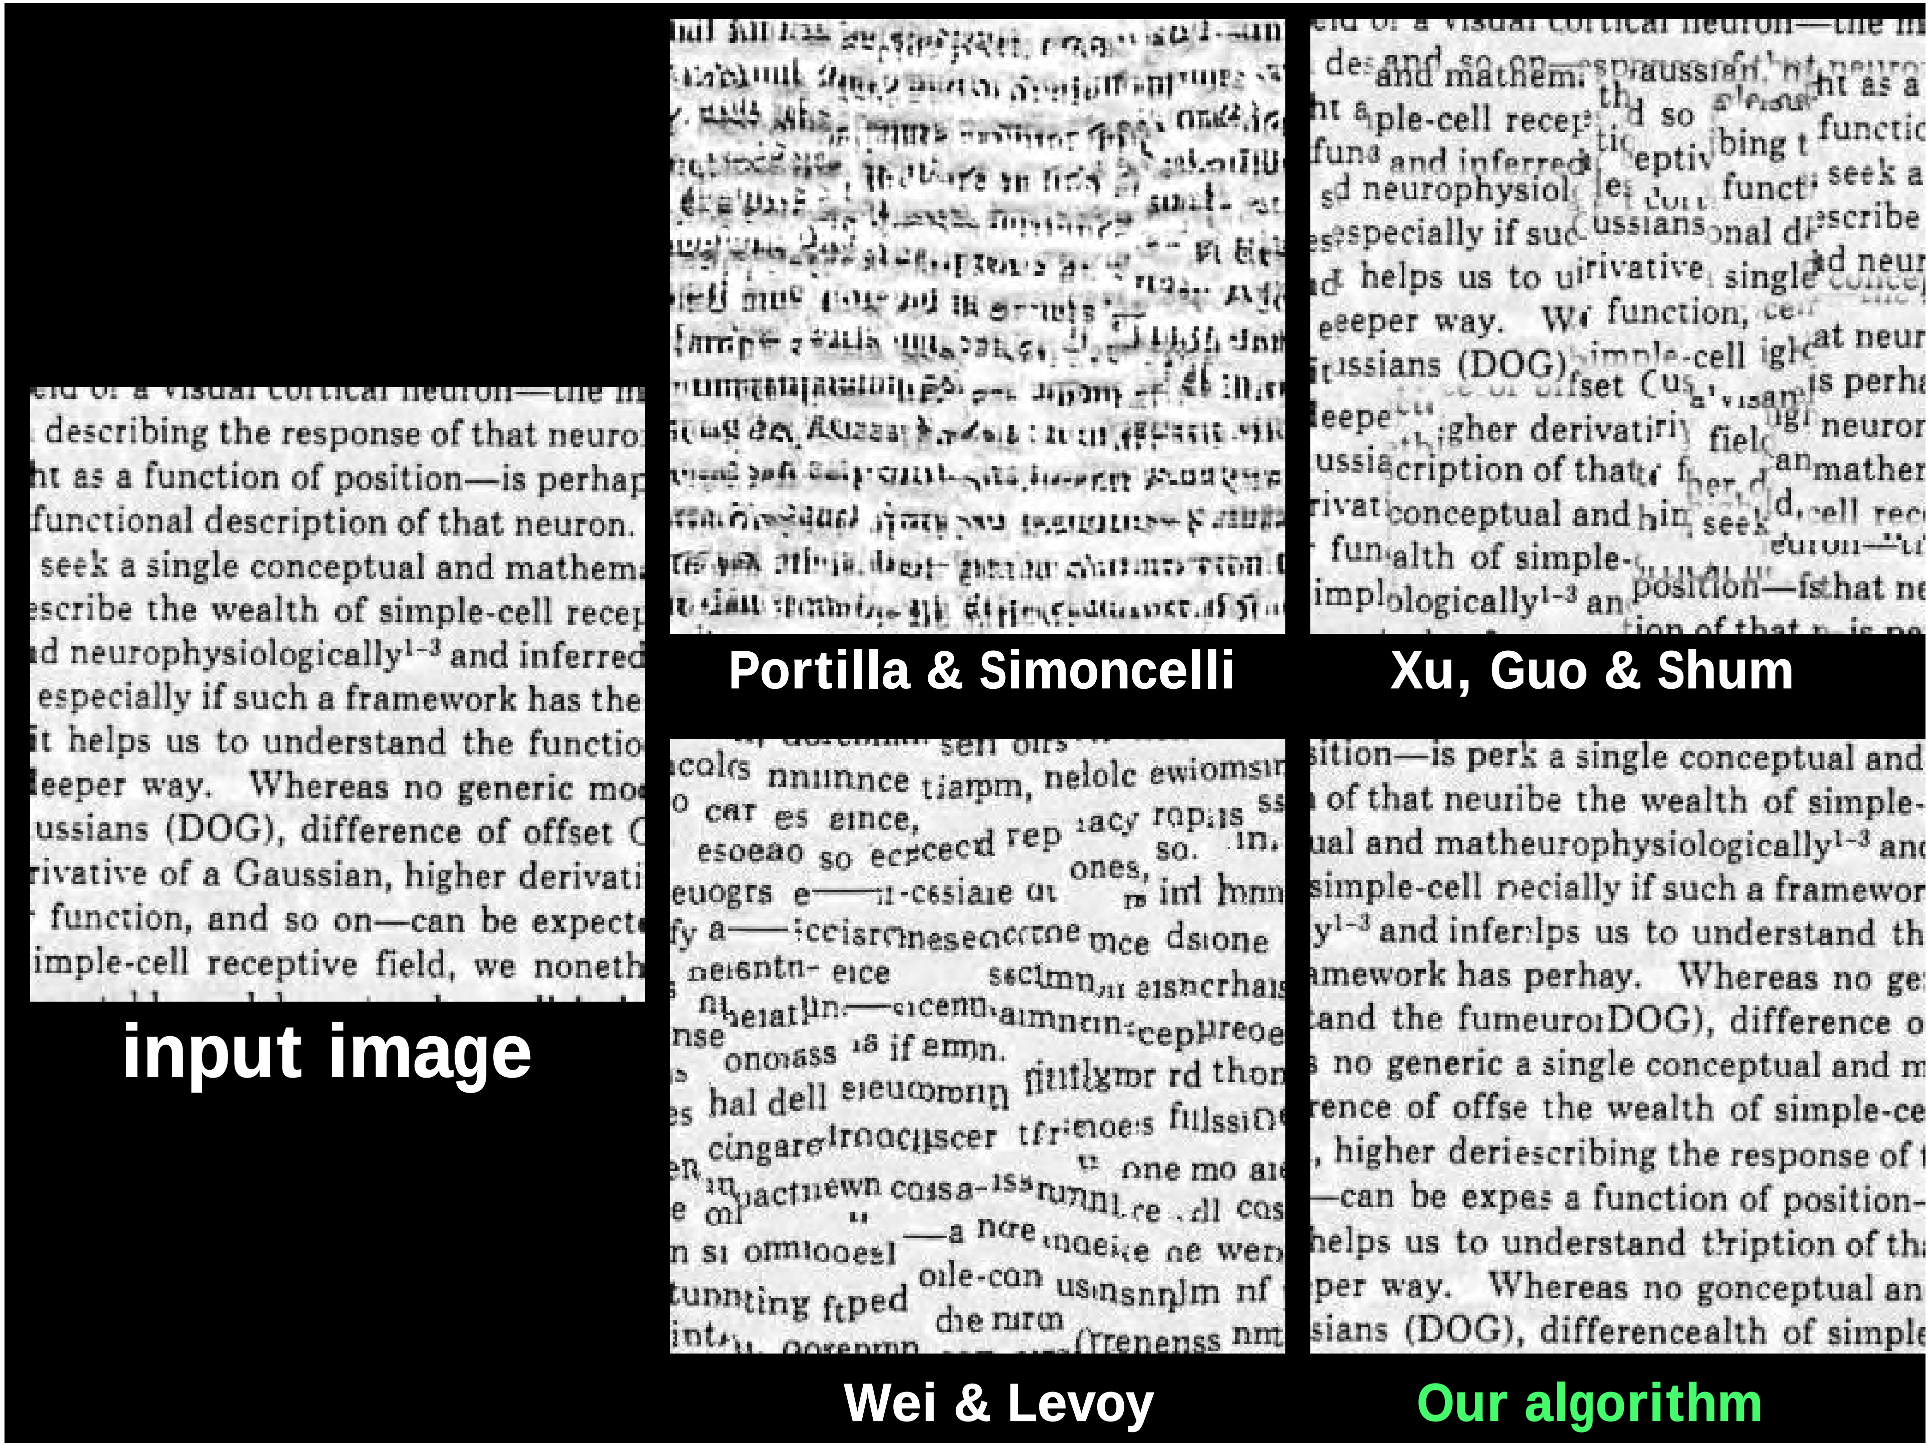
input image
Portilla & Simoncelli
Xu, Guo & Shum
Wei & Levoy
Our algorithm

Political Texture Synthesis!

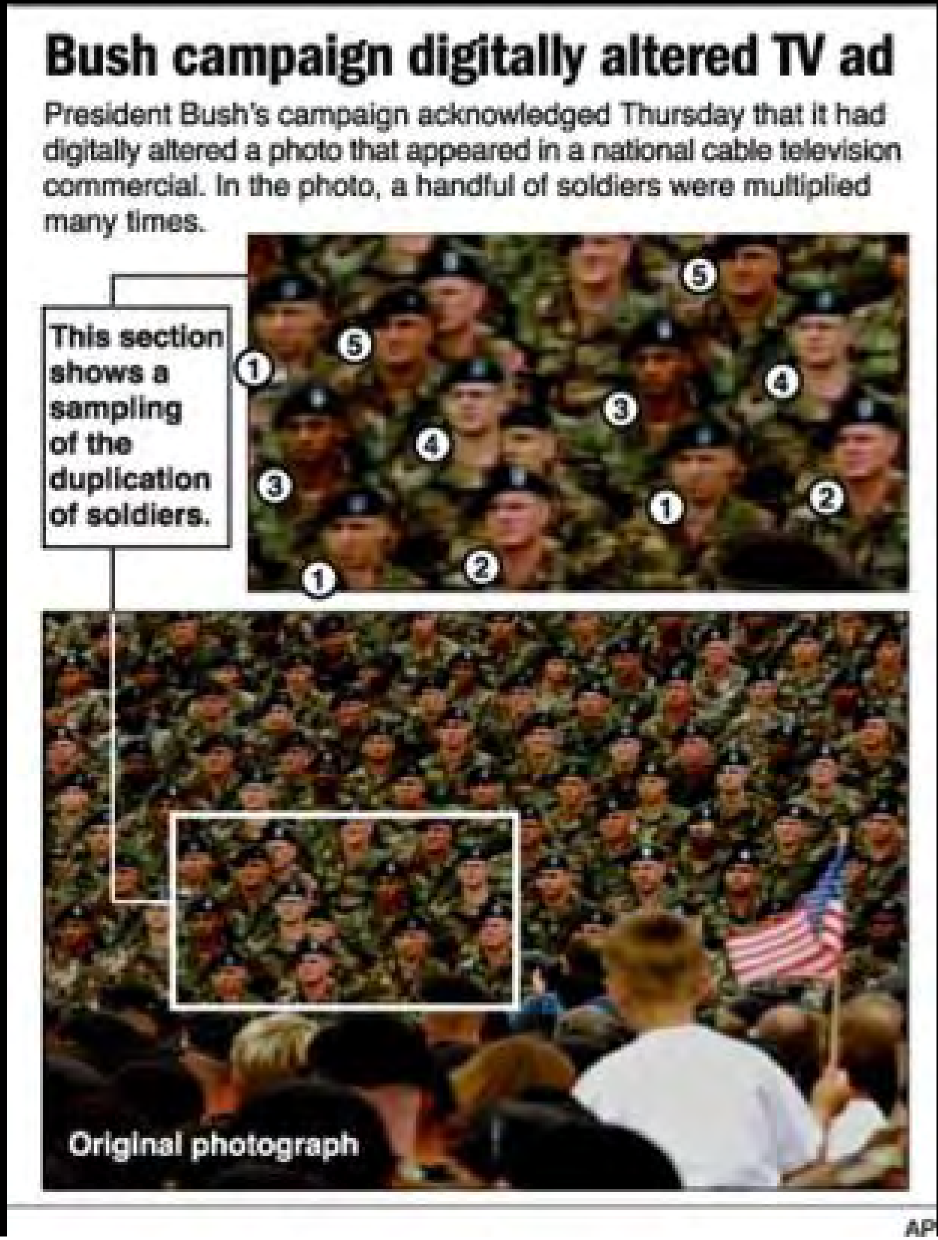

Fill Order

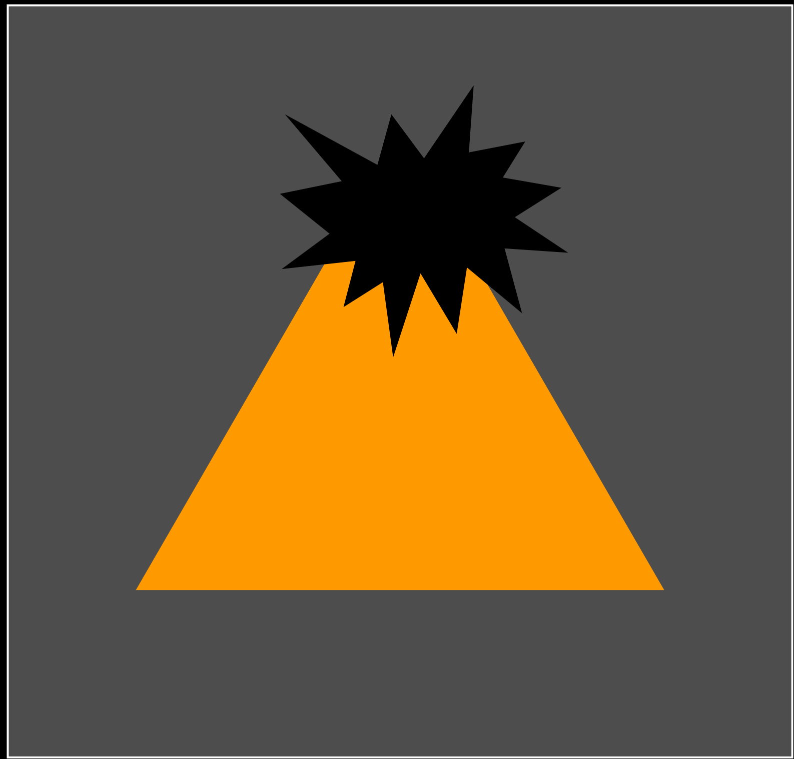

- In what order should we fill the pixels?

Fill Order

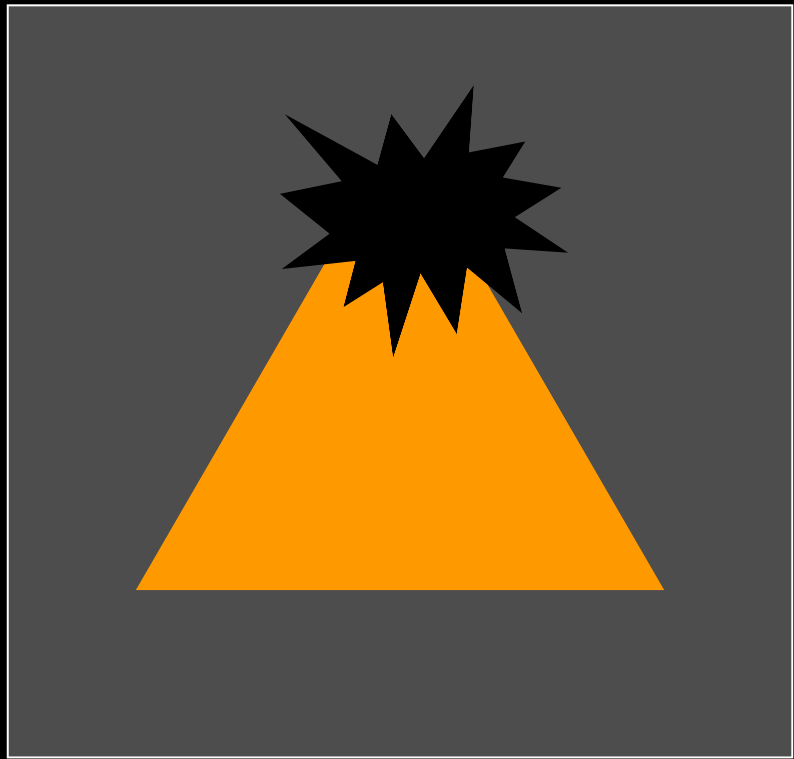

- In what order should we fill the pixels?
  - choose pixels that have more neighbors filled
  - choose pixels that are continuations of lines/curves/edges

Criminisi, Perez, and Toyama. "Object Removal by Exemplar-based Inpainting," Proc. CVPR, 2003.

Exemplar-based Inpainting demo

http://research.microsoft.com/vision/cambridge/i3l/patchworks.htm

Application: Texture Transfer

- Try to explain one object with bits and pieces of another object:

Texture Transfer

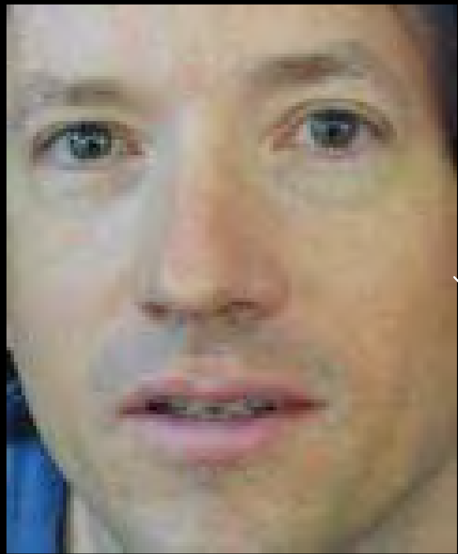

Constraint

Texture sample

Texture Transfer

- Take the texture from one image and “paint” it onto another object

Same as texture synthesis, except an additional constraint:

1. Consistency of texture
2. Similarity to the image being “explained”

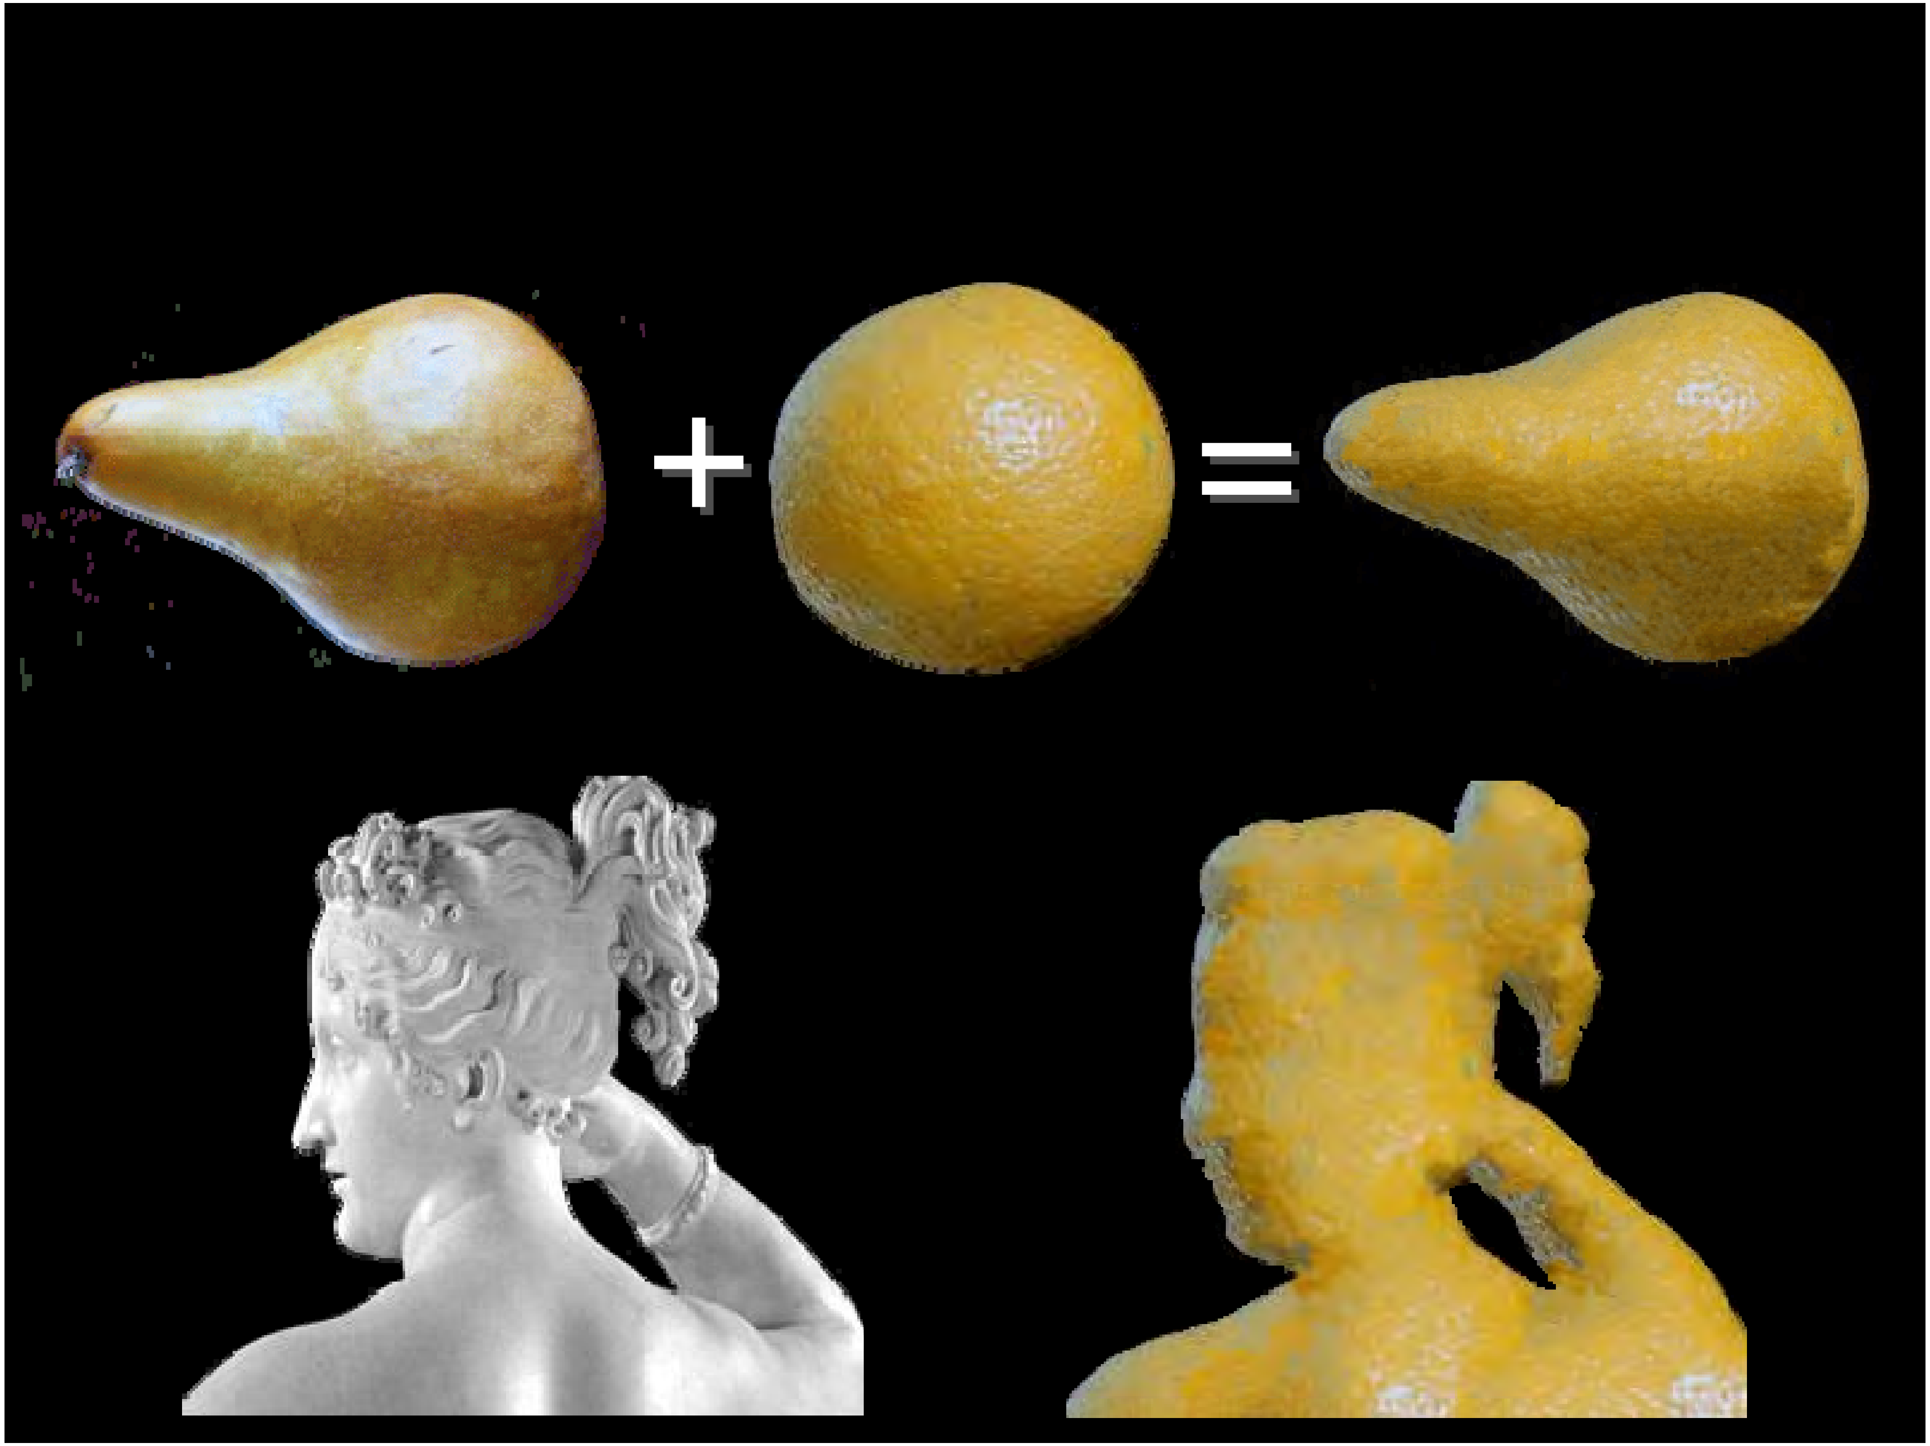



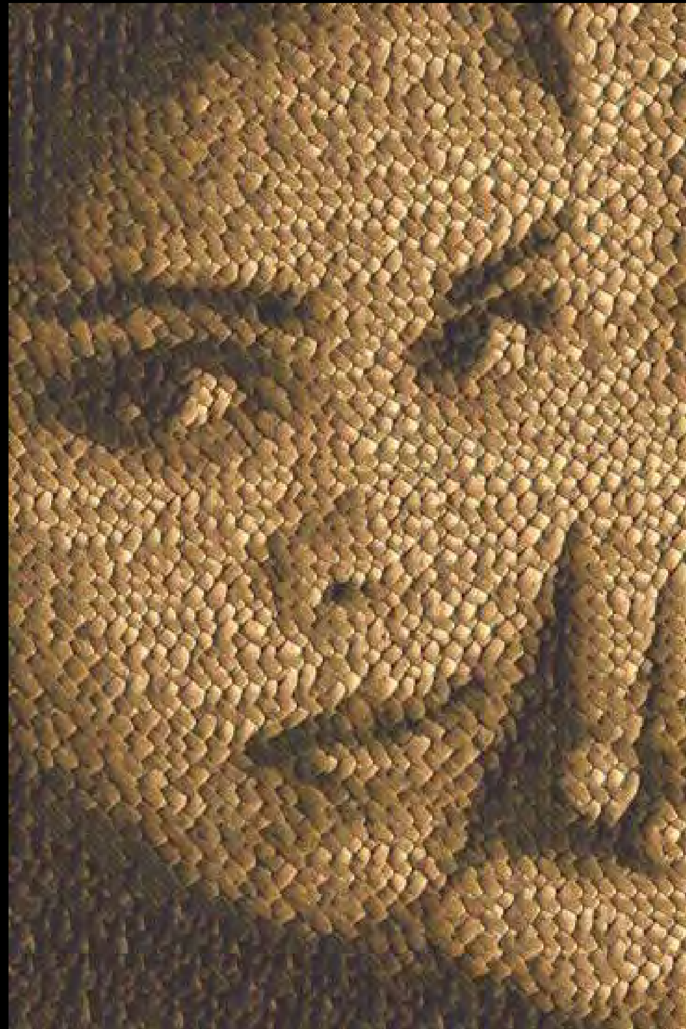

Image Analogies

Aaron Hertzmann[1,2]

Chuck Jacobs[2]

Nuria Oliver[2]

Brian Curless[3]

David Salesin[2,3]

[1]New York University

[2]Microsoft Research

[3]University of Washington

Image Analogies

A

A'

B

B'

Blur Filter

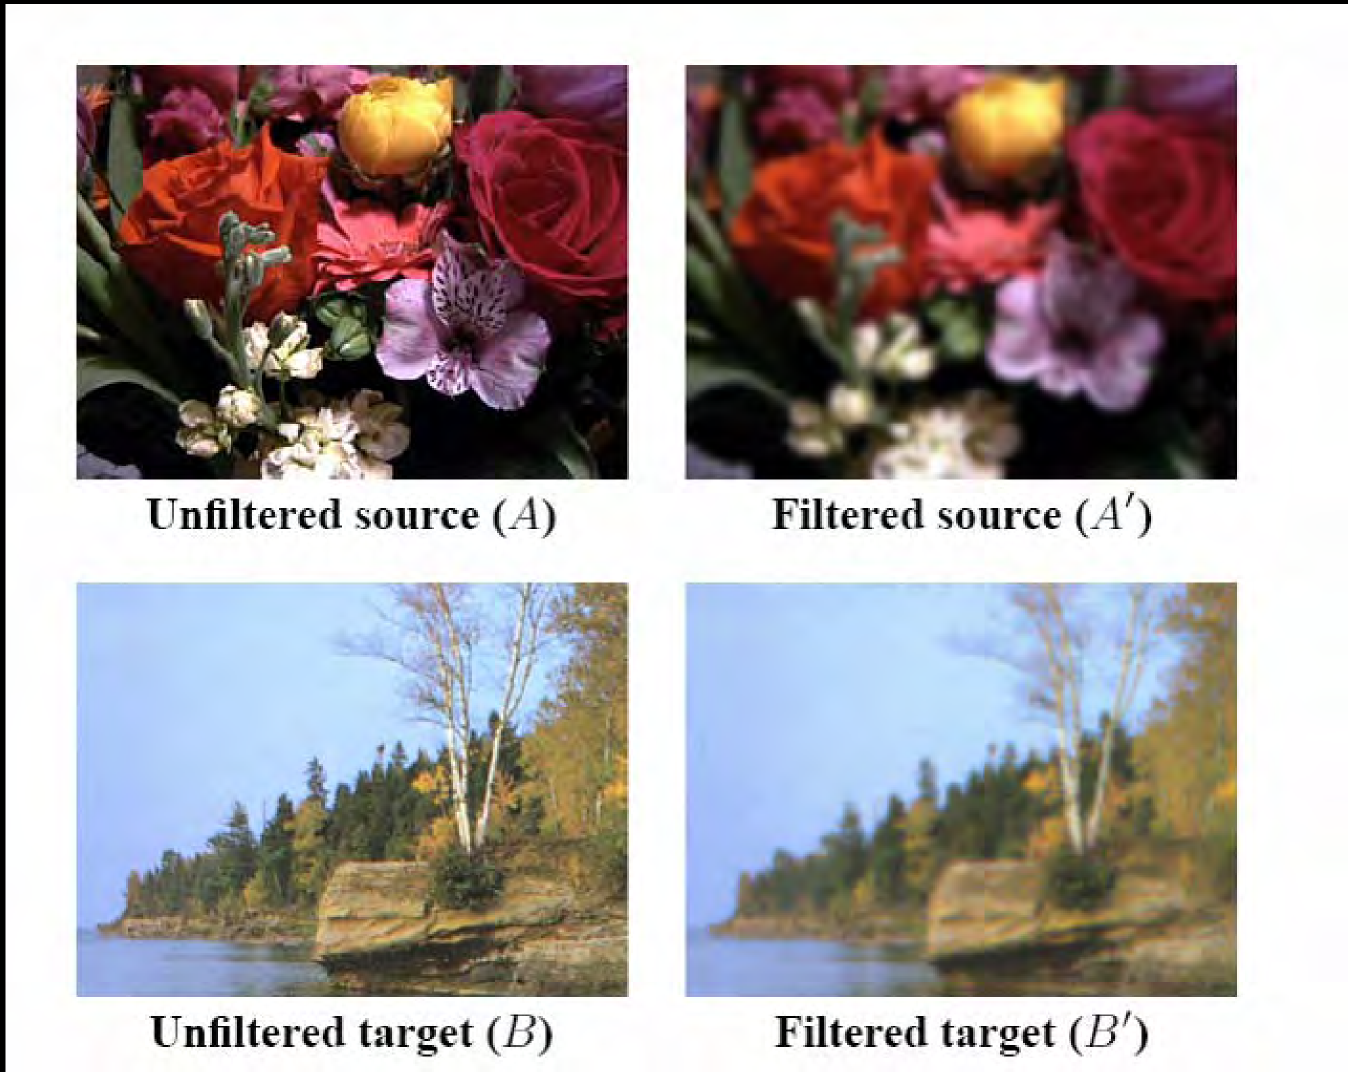
Unfiltered source (A)
Filtered source (A')
Unfiltered target (B)
Filtered target (B')

Edge Filter

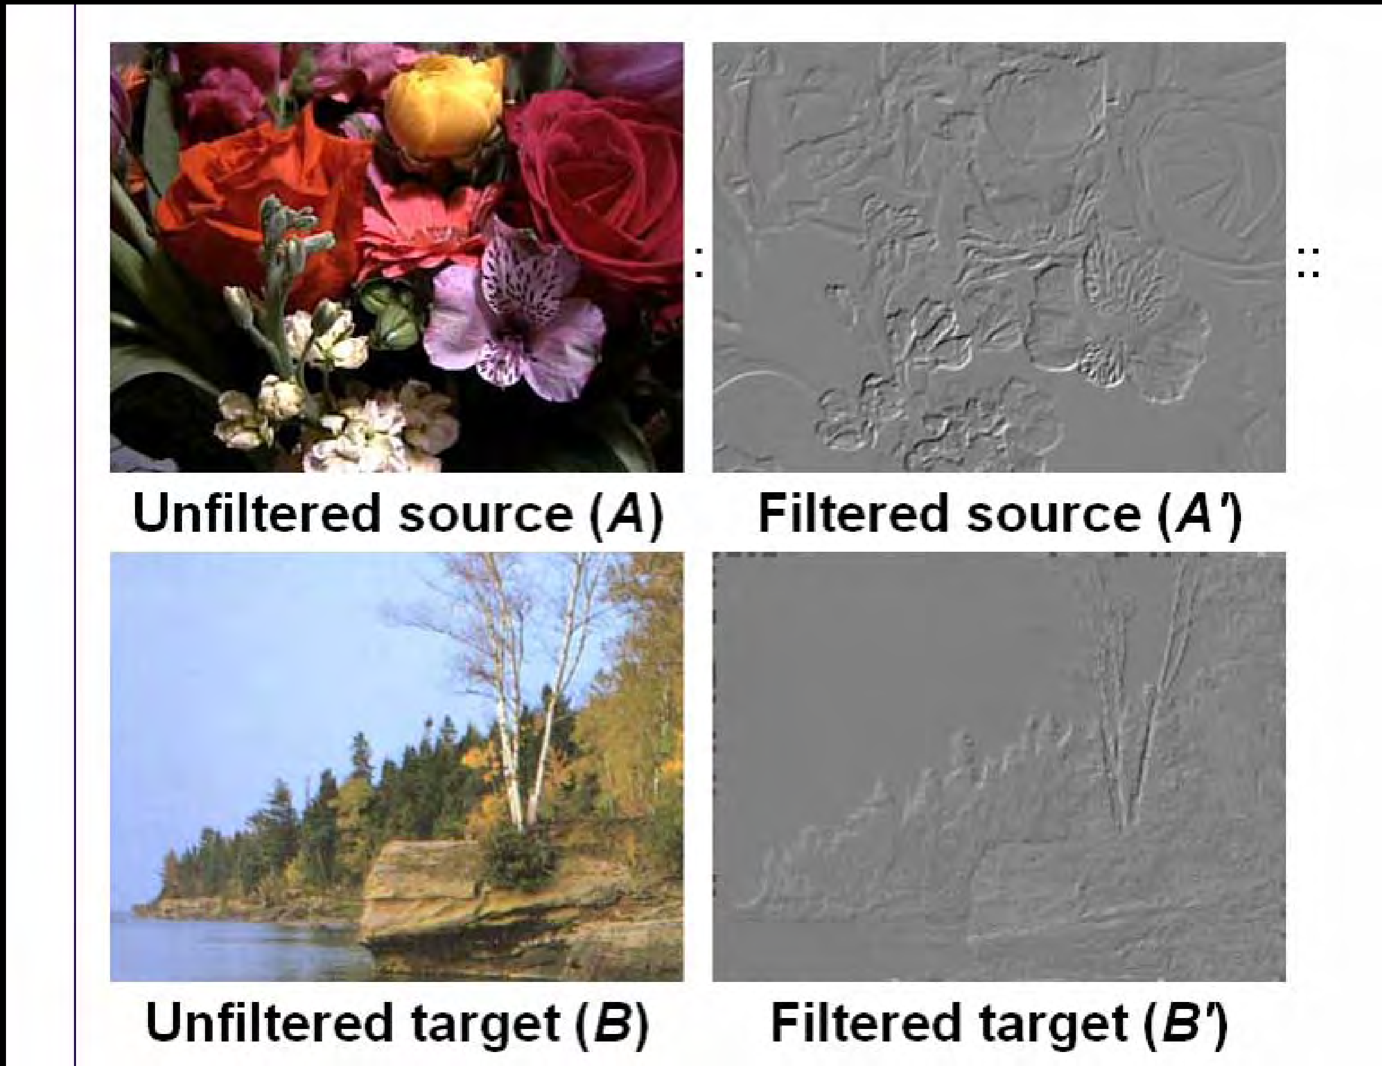
Unfiltered source (A)
:
Filtered source (A')
::
Unfiltered target (B)
Filtered target (B')

Artistic Filters

A

A’

B

B’

Colorization

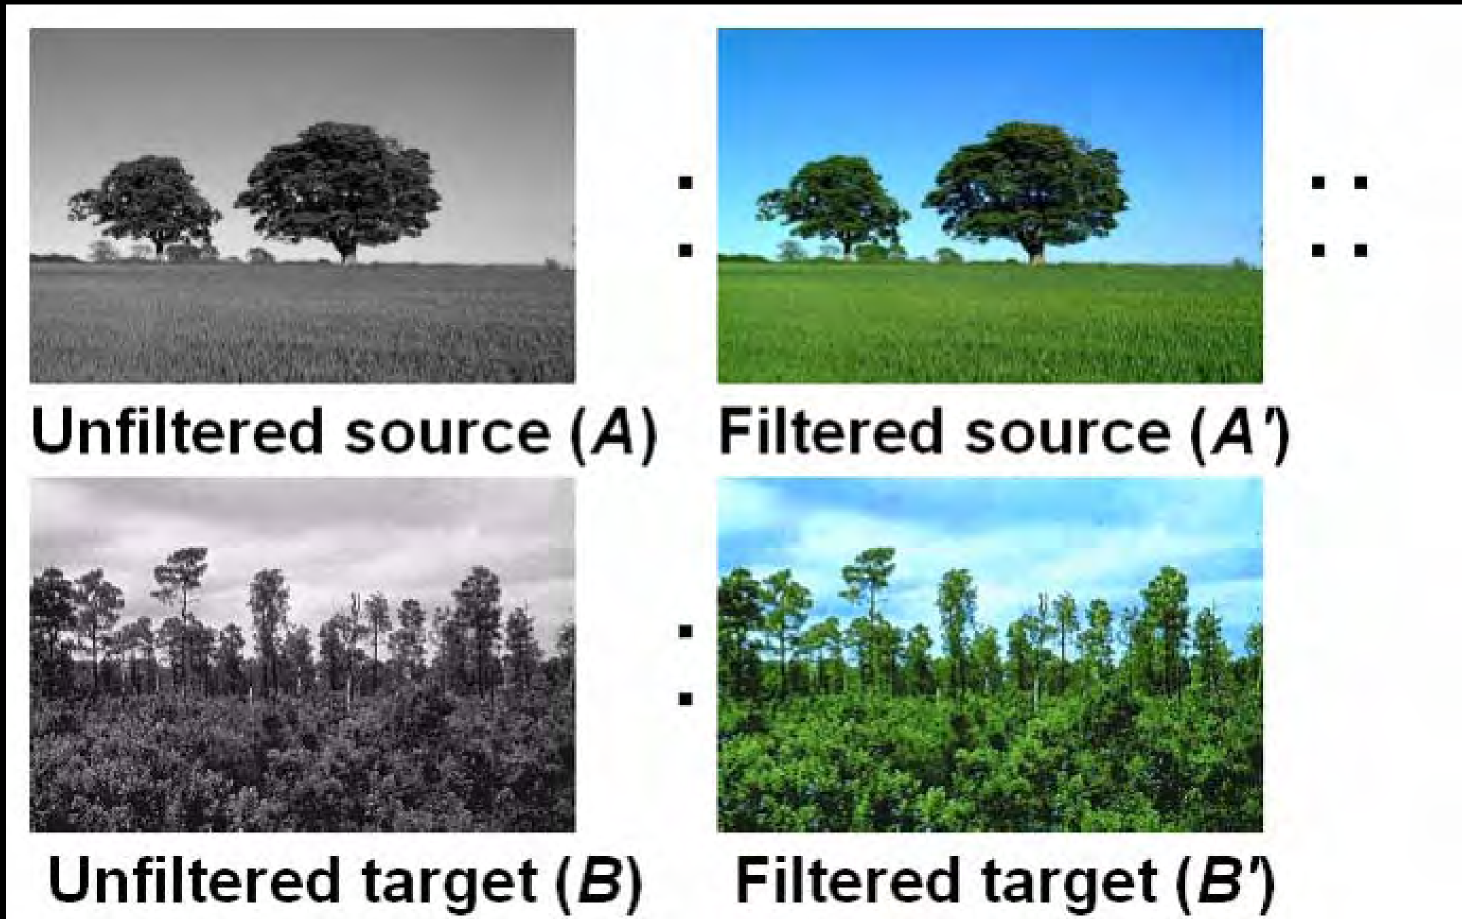

Texture-by-numbers

A

A'

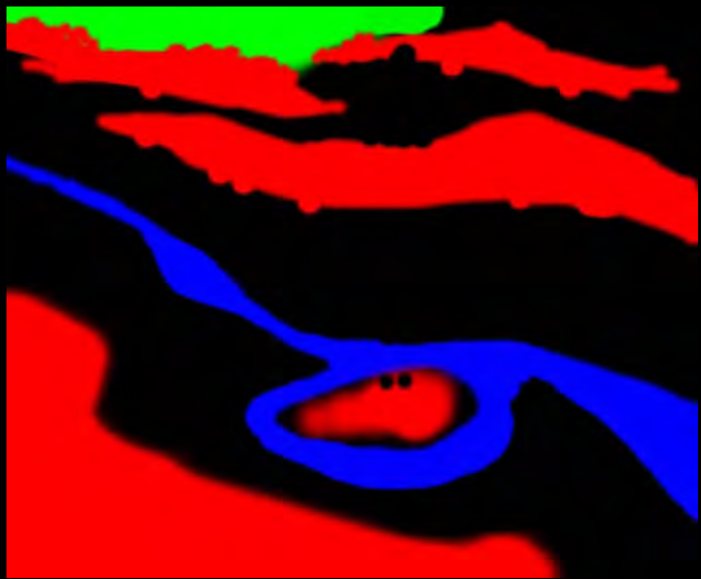

B

B'

Super-resolution

A

A’

Super-resolution (result!)

B

B’

Scene Completion Using Millions of Photographs

James Hays and Alexei A. Efros

Carnegie Mellon University


5407 BWL

Efros and Leung result

Criminisi et al. result

Criminisi et al. result

Scene Matching for Image Completion

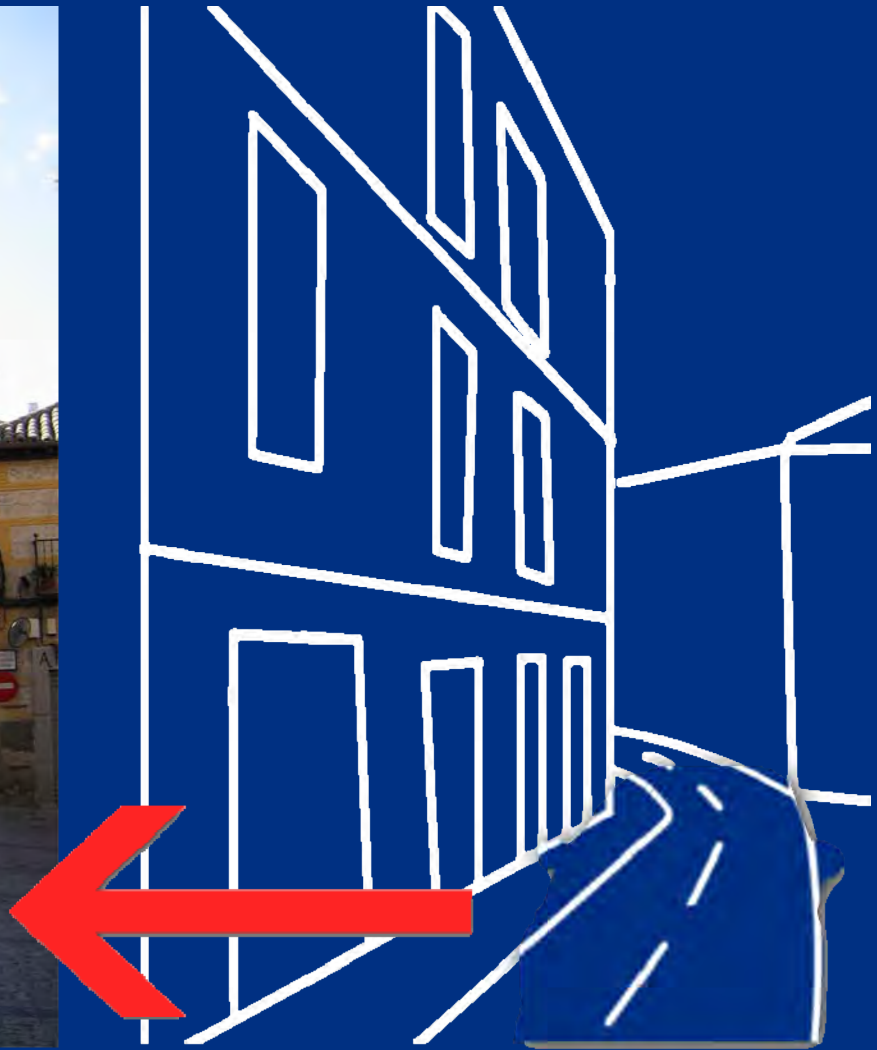

alley | Search Images | Search the Web | Advanced Image Search | Preferences

Strict SafeSearch is on

Images Showing: All image sizes

Results **1** - **20** of about **908,000** for **alley** [definition] with **Safesearch on**. (**0.07** seconds)

Change **Alley** Aerial Plaza with its ...
300 x 400 - 21k
en.wikipedia.org

The Printer's **Alley** sign looking **...**
679 x 450 - 469k - jpg
franklin.thefuntimesguide.com

Looking west past Printers **Alley**.
679 x 450 - 464k - jpg
franklin.thefuntimesguide.com

More Bubble Gum **Alley** photos can be **...**
764 x 591 - 33k - gif
www.locallinks.com

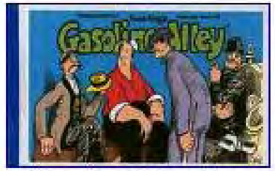

Gasoline **Alley** gang
692 x 430 - 177k - jpg
newcritics.com

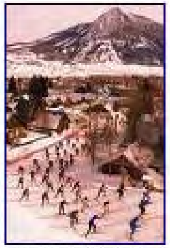

2007 **Alley** Loop Sponsors
300 x 453 - 51k - jpg
www.cbnordic.org

Change **Alley** : interior
550 x 413 - 98k
infopedia.nlb.gov.sg

Earl G. **Alley ...**
321 x 383 - 19k - jpg
www.msstate.edu

Gun **Alley** 8.5x11 Full Color Ink Wash **...**
390 x 301 - 14k - jpg
www.rorschachentertainment.com

Grace Court **Alley**
732 x 549 - 98k - jpg
www.bridgeandtunnelclub.com

Grace Court **Alley**
732 x 549 - 80k - jpg
www.bridgeandtunnelclub.com

panoramic photo of Alligator **Alley**
4902 x 460 - 1048k - jpg
sflwww.er.usgs.gov

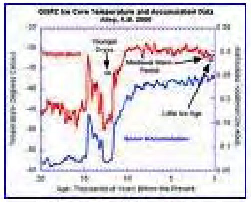

Richard B. **Alley**
450 x 361 - 29k - gif
www.ncdc.noaa.gov

Also, Chicken **Alley** is reported to **...**
450 x 337 - 82k
phidoux.typepad.com

Ego **Alley**
500 x 375 - 48k - jpg
dc.about.com

Scene Completion Result

The Algorithm

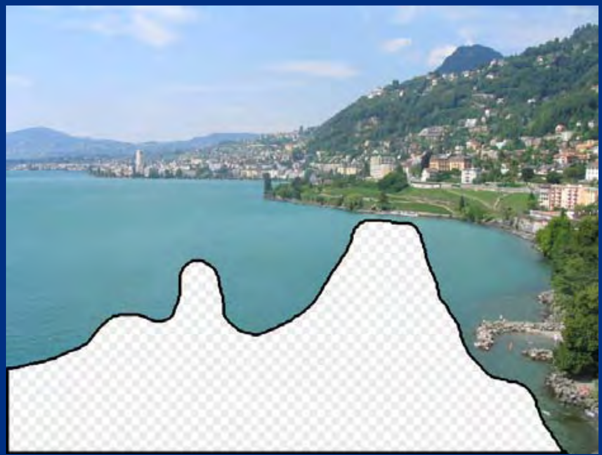

Input image

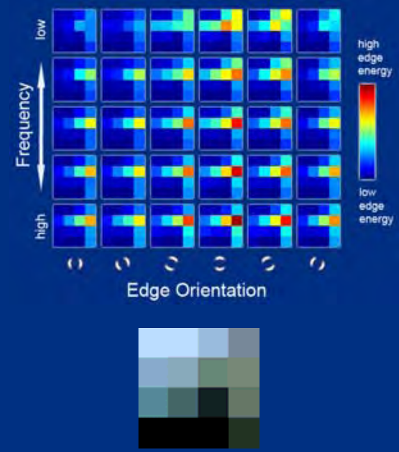

Scene Descriptor

...

Image Collection

...

200 matches

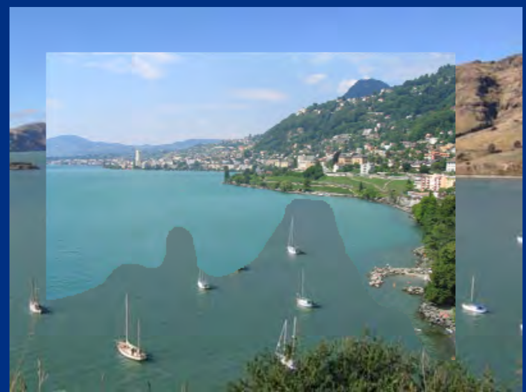

Context matching + blending

20 completions

Data

We downloaded **2.3 Million** unique images from Flickr groups and keyword searches.

Scene Matching

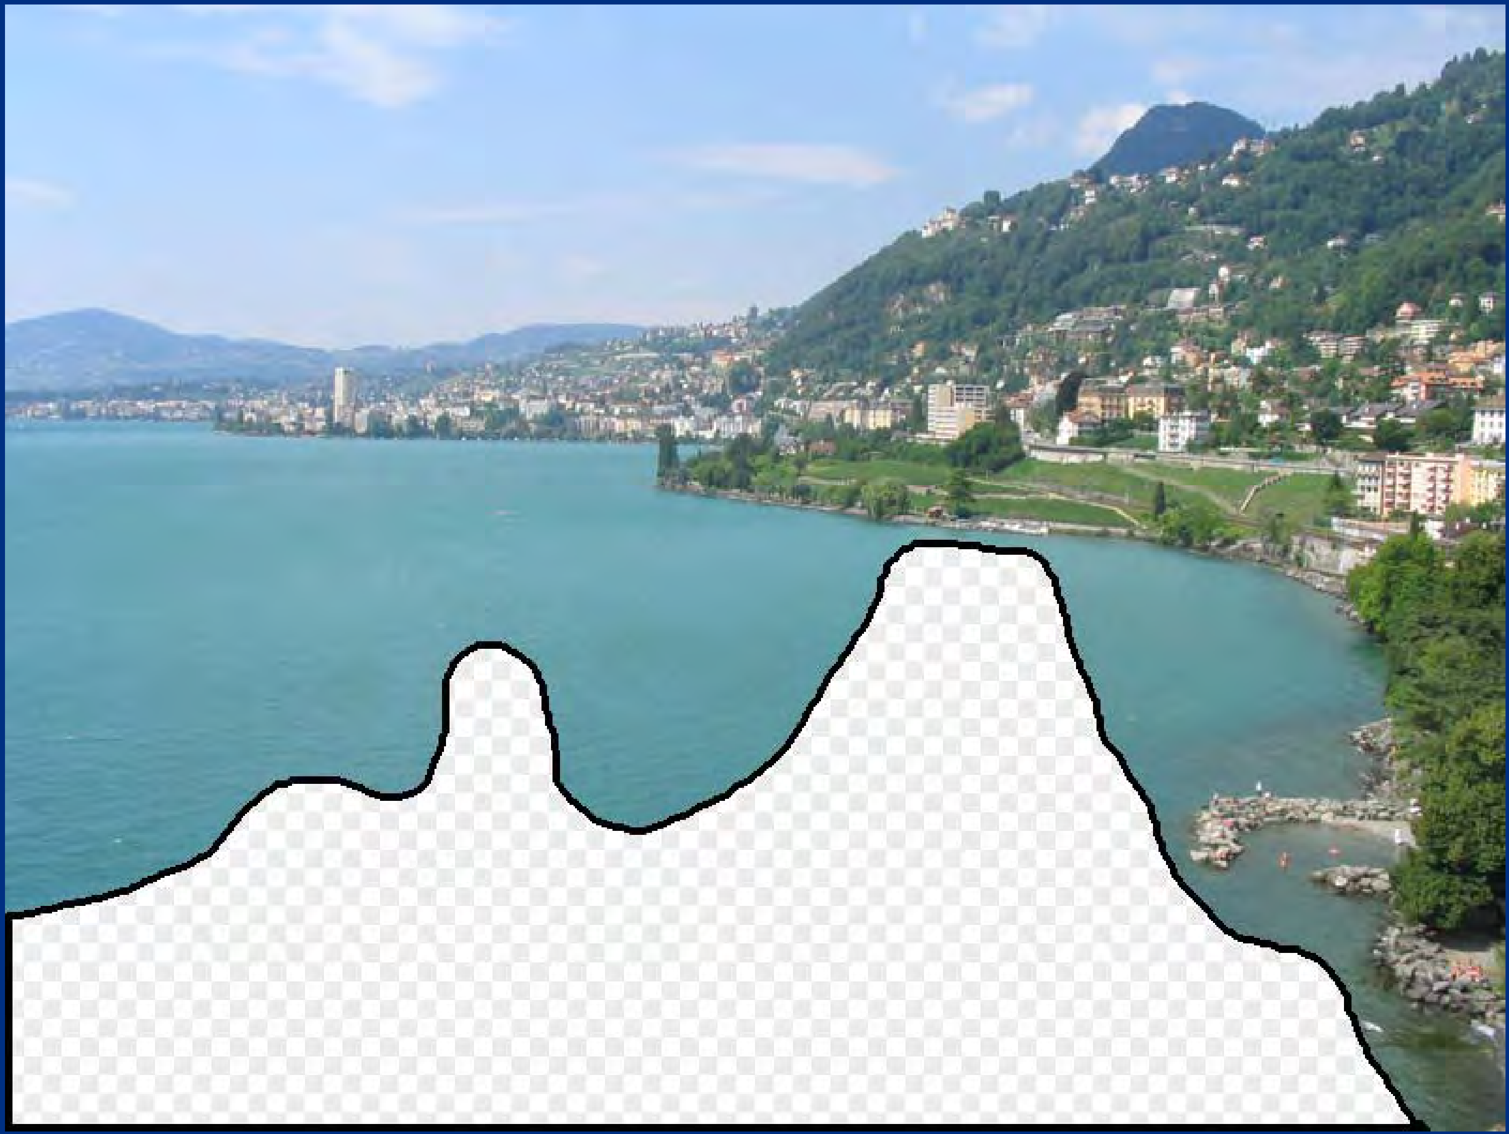

Scene Descriptor

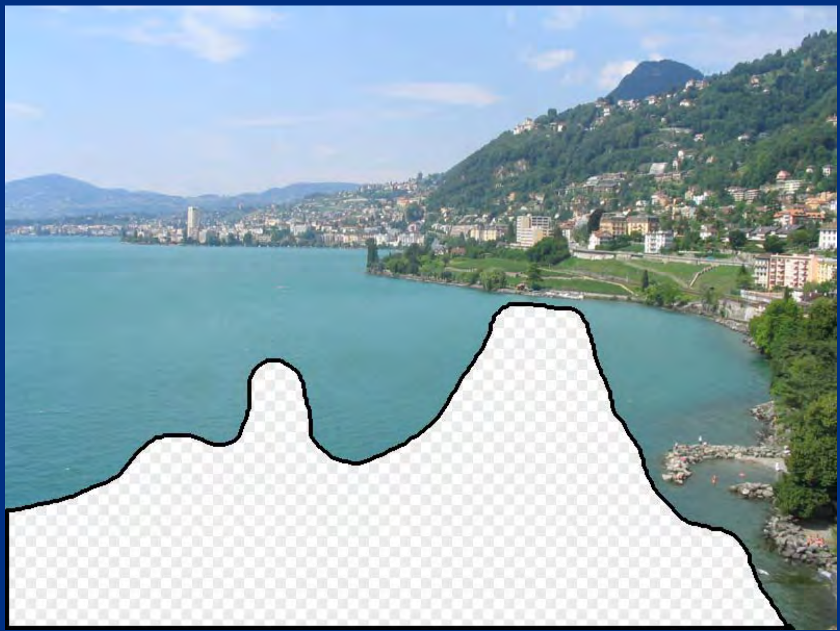

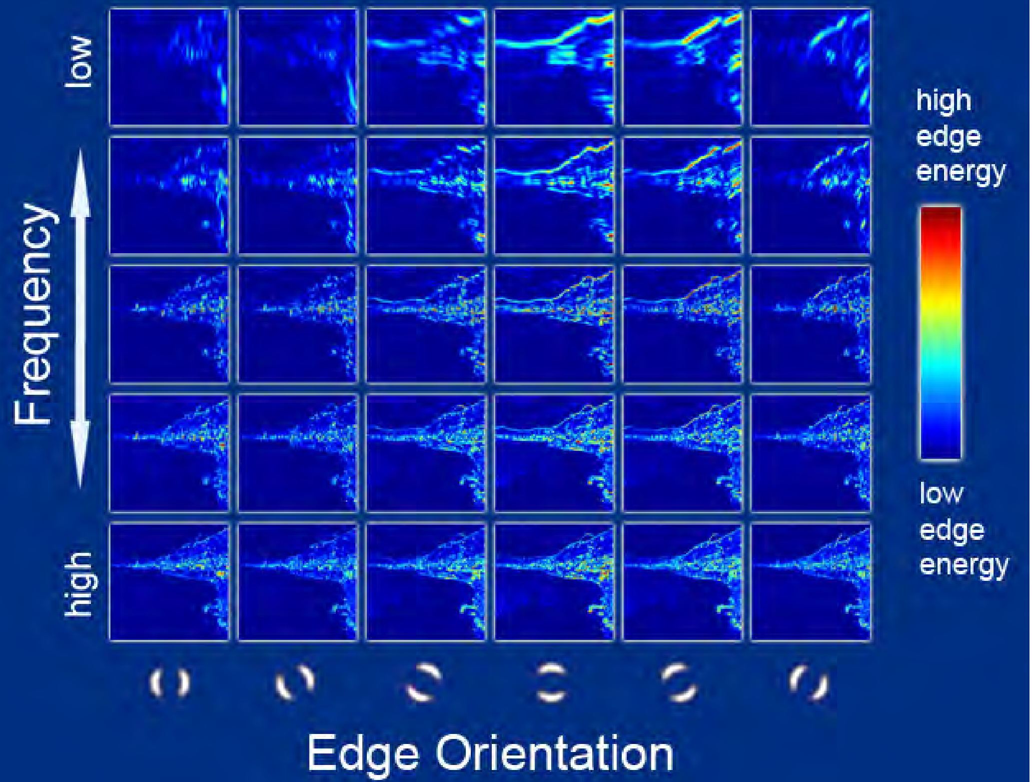

Scene Descriptor

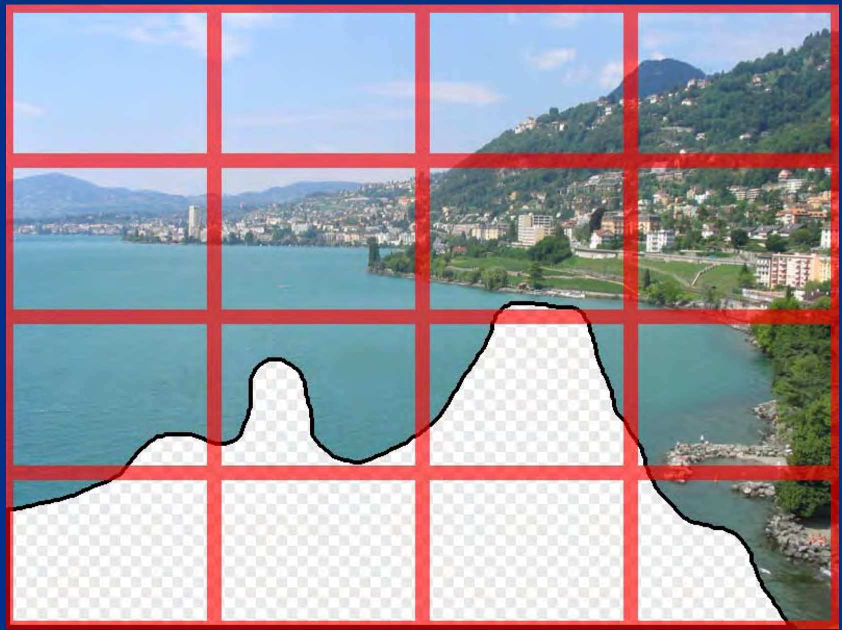

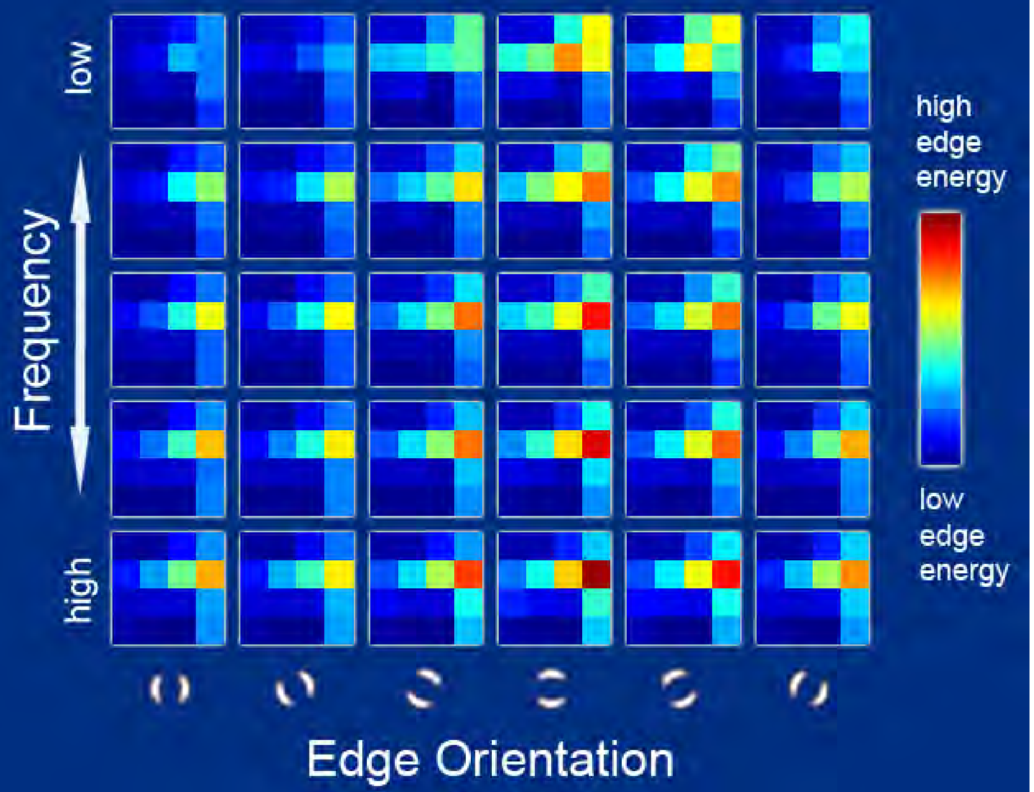

Gist scene descriptor
(Oliva and Torralba 2001)

Scene Descriptor

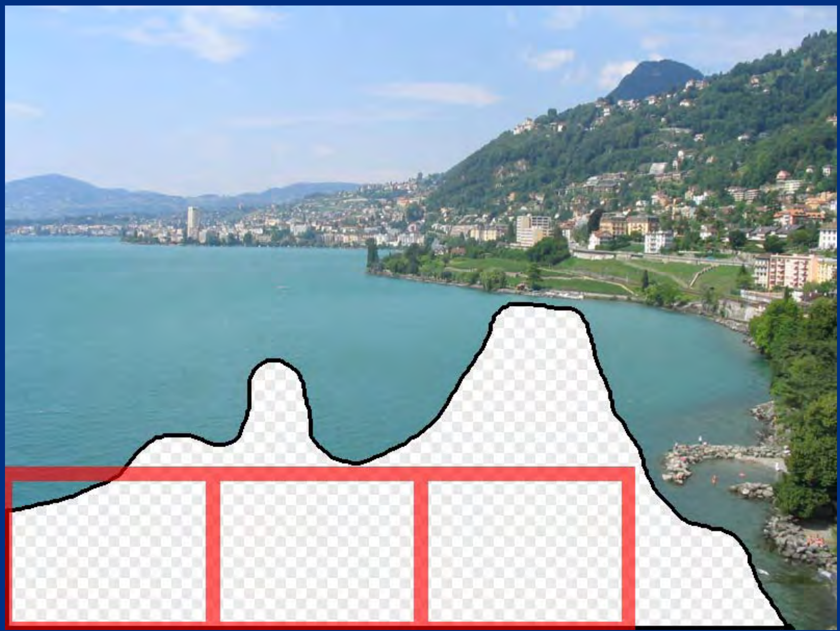

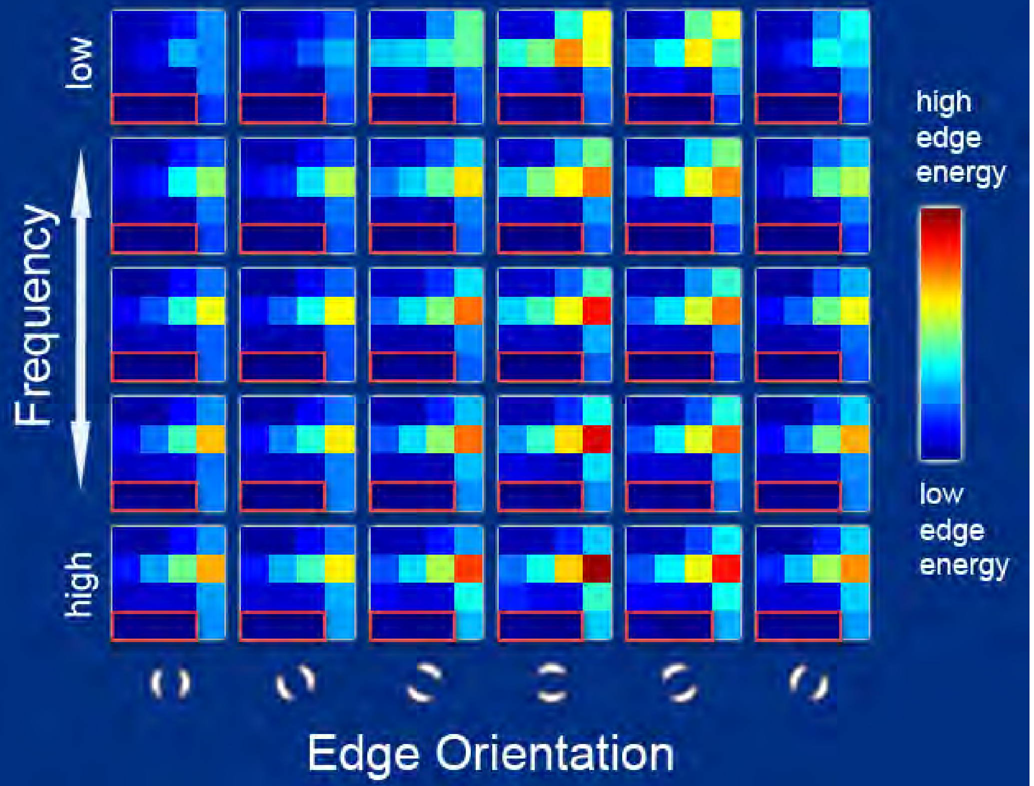

Gist scene descriptor
(Oliva and Torralba 2001)

Scene Descriptor

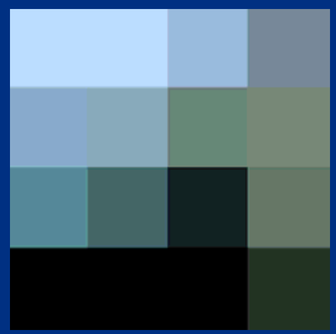

+

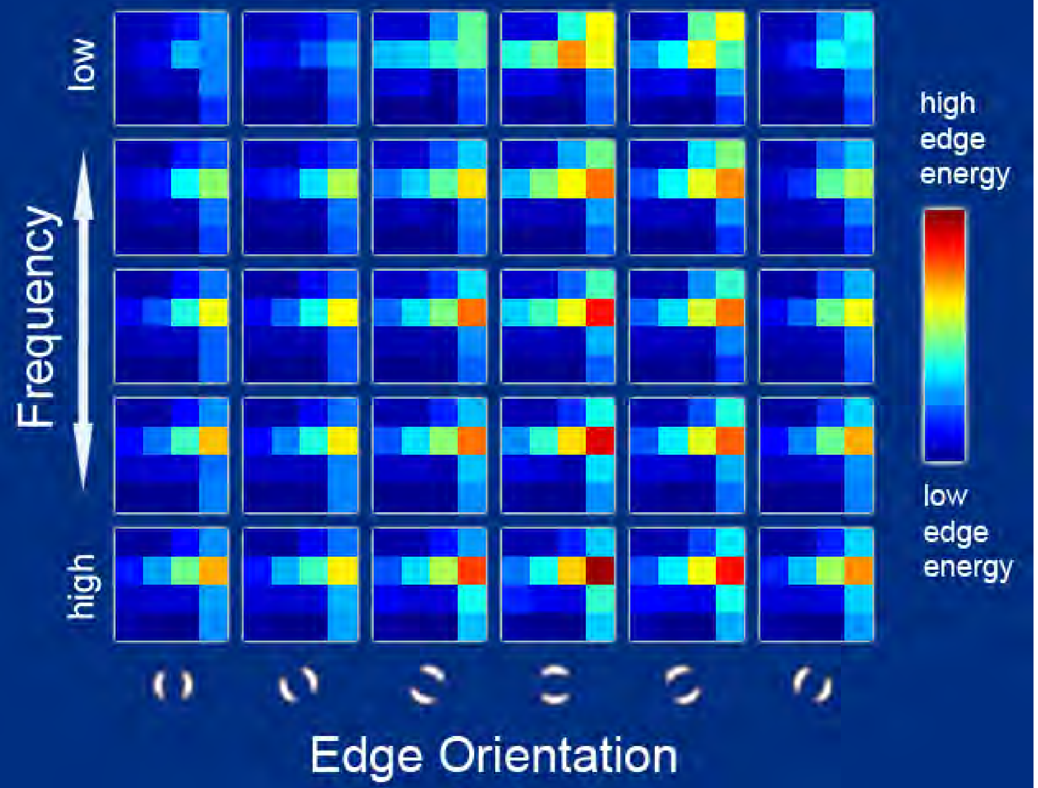

Gist scene descriptor
(Oliva and Torralba 2001)

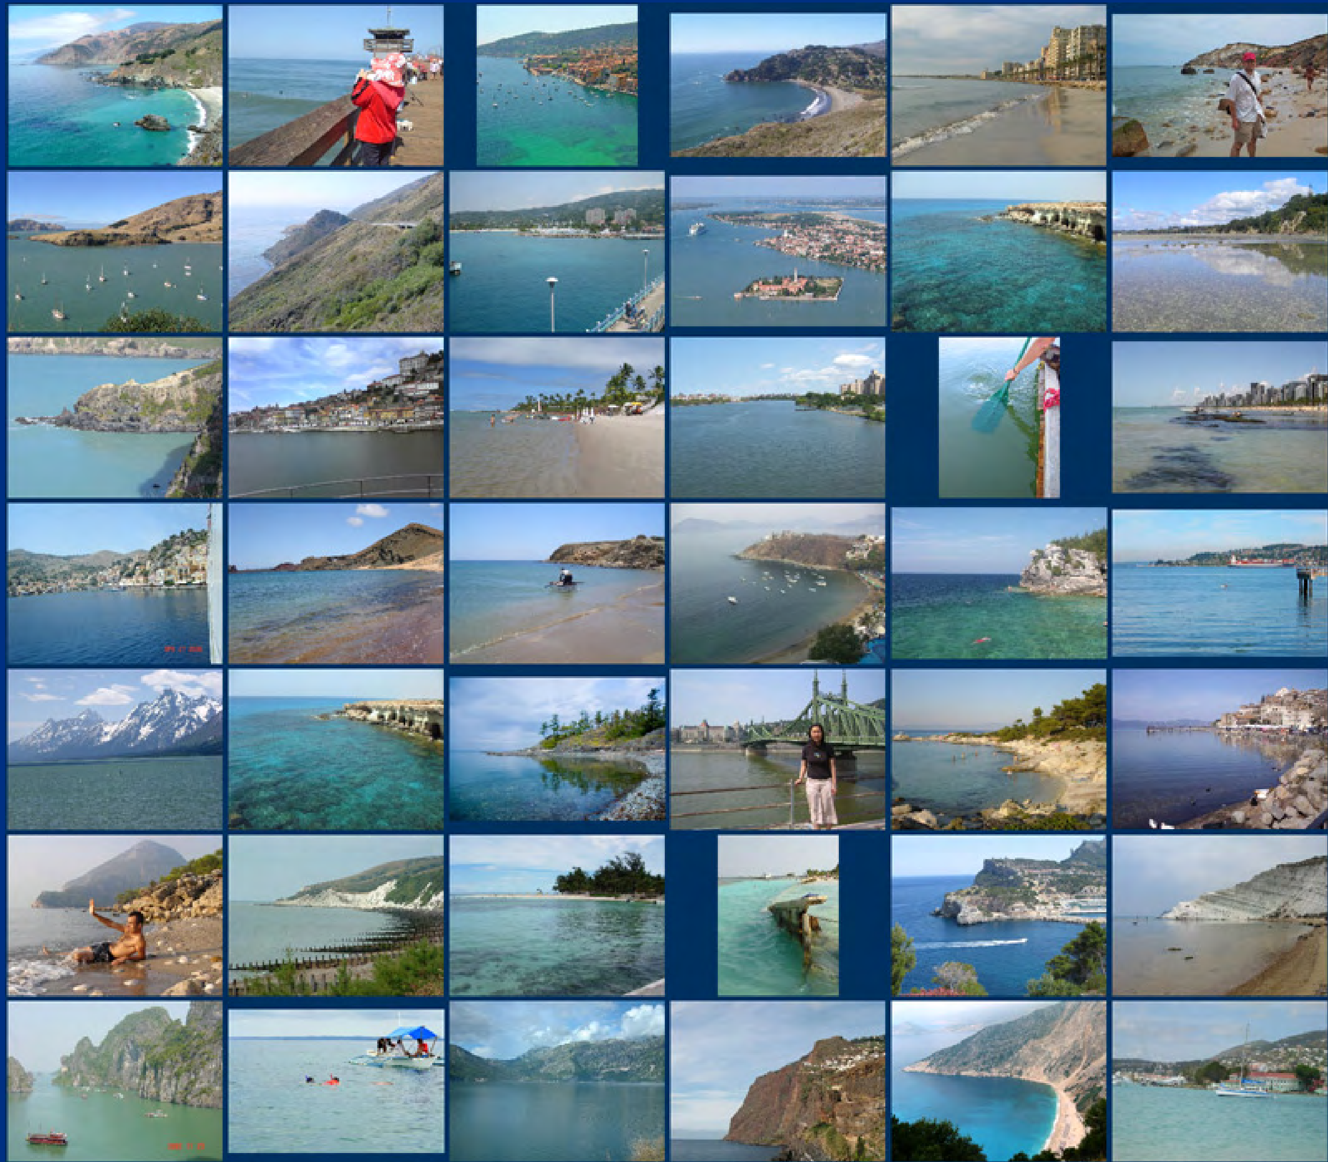

… 200 total

Context Matching

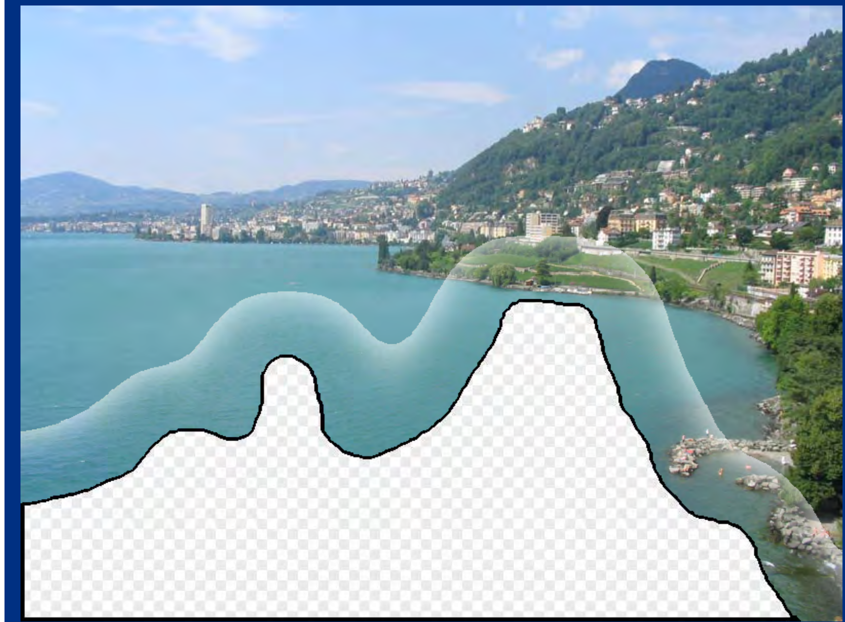

Result Ranking

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

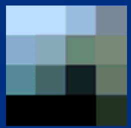
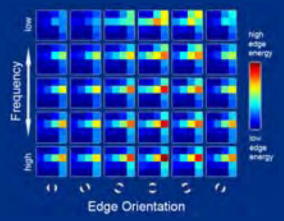

The scene matching distance

The context matching distance
(color + texture)

The graph cut cost

Top 20 Results


OCT 30 2003

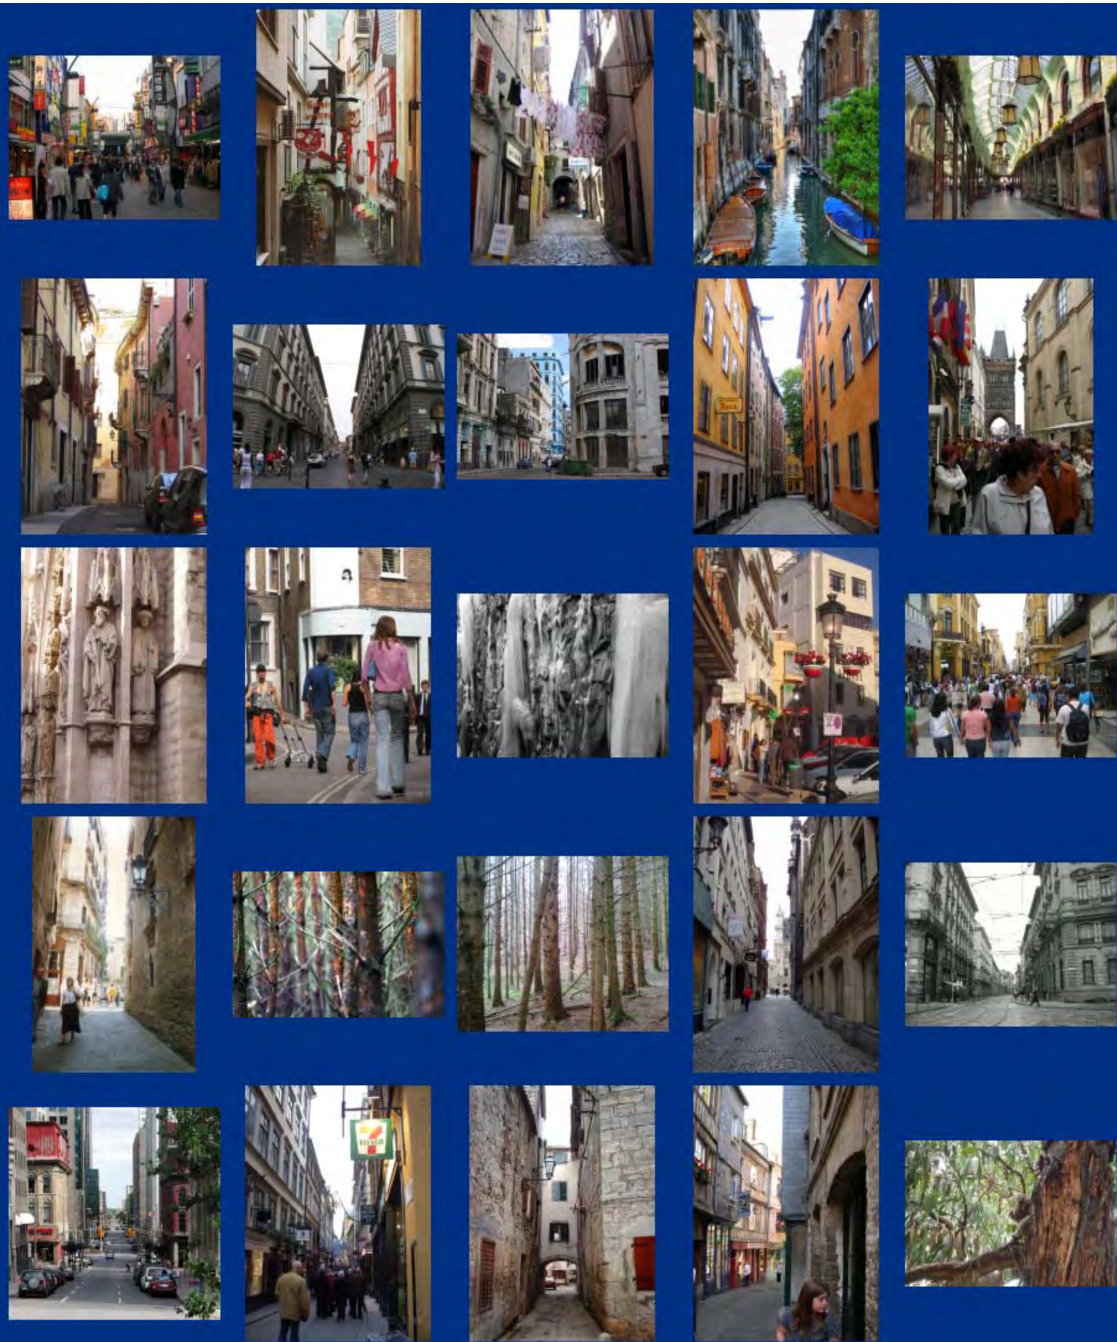

… 200 scene matches

… 200 scene matches


HOTEL
Grand Place
Jupiler

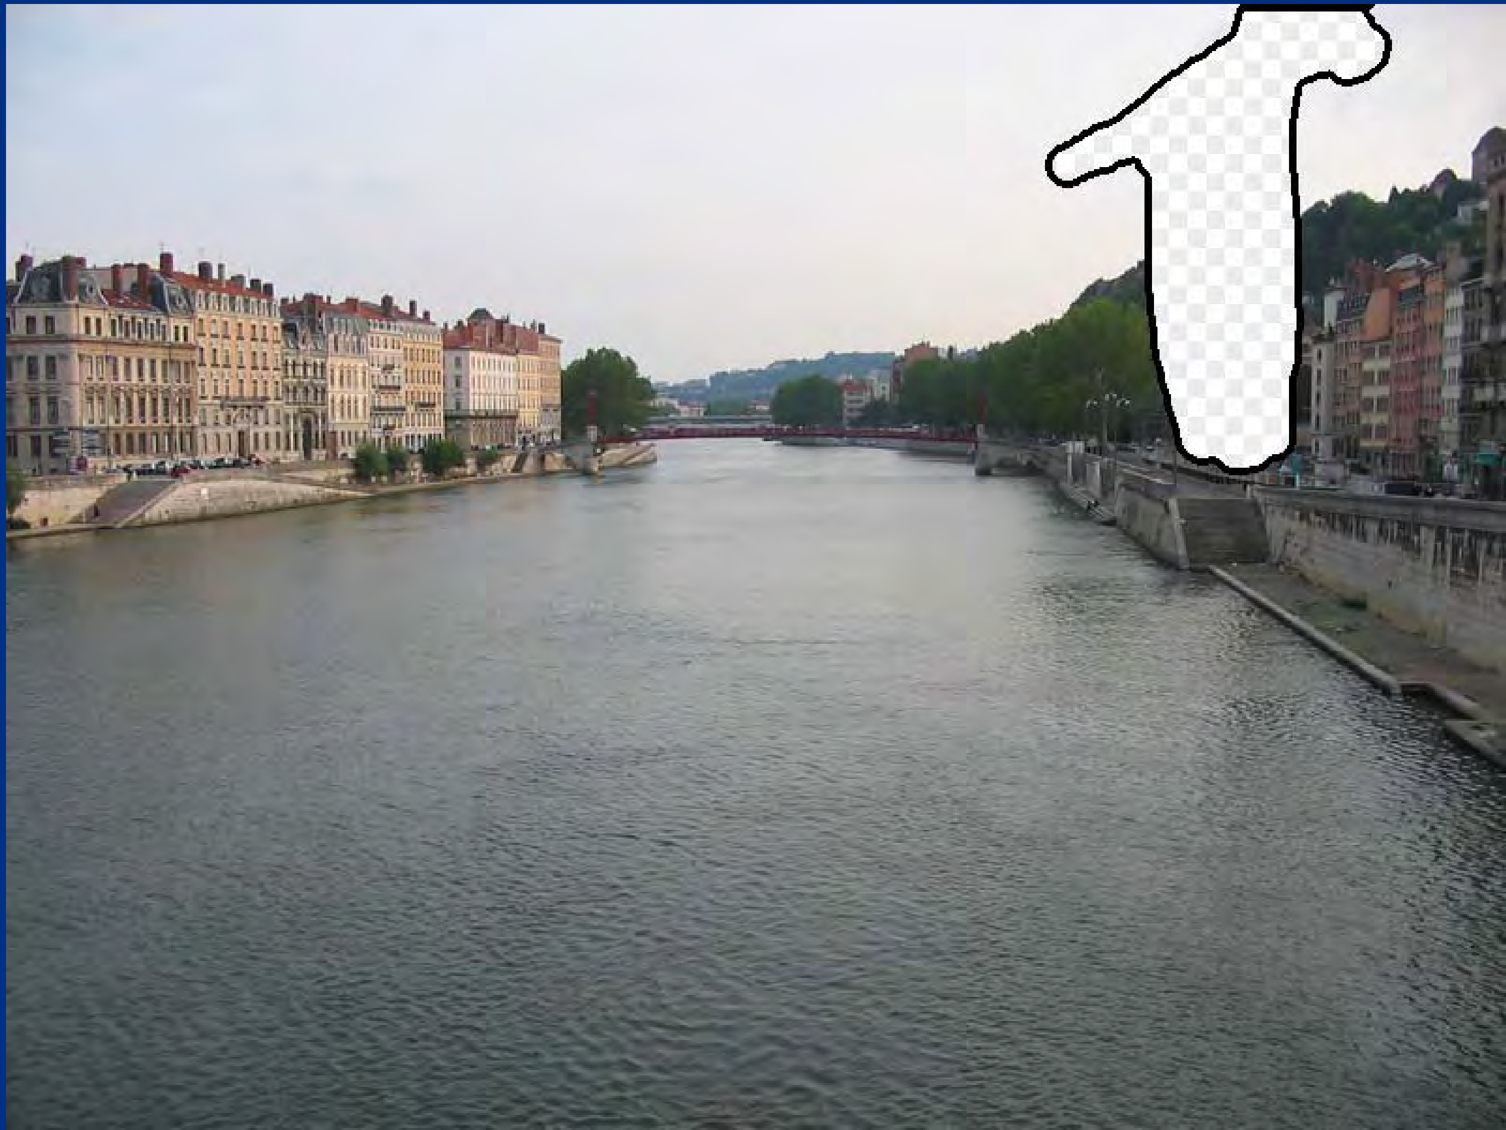

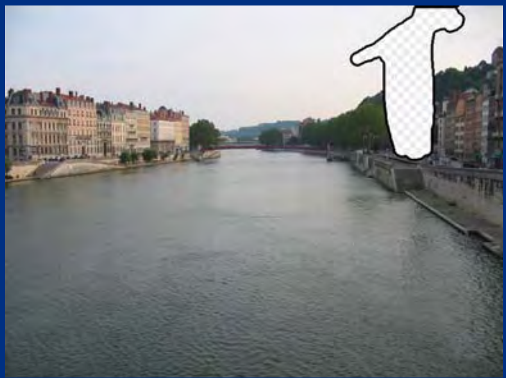

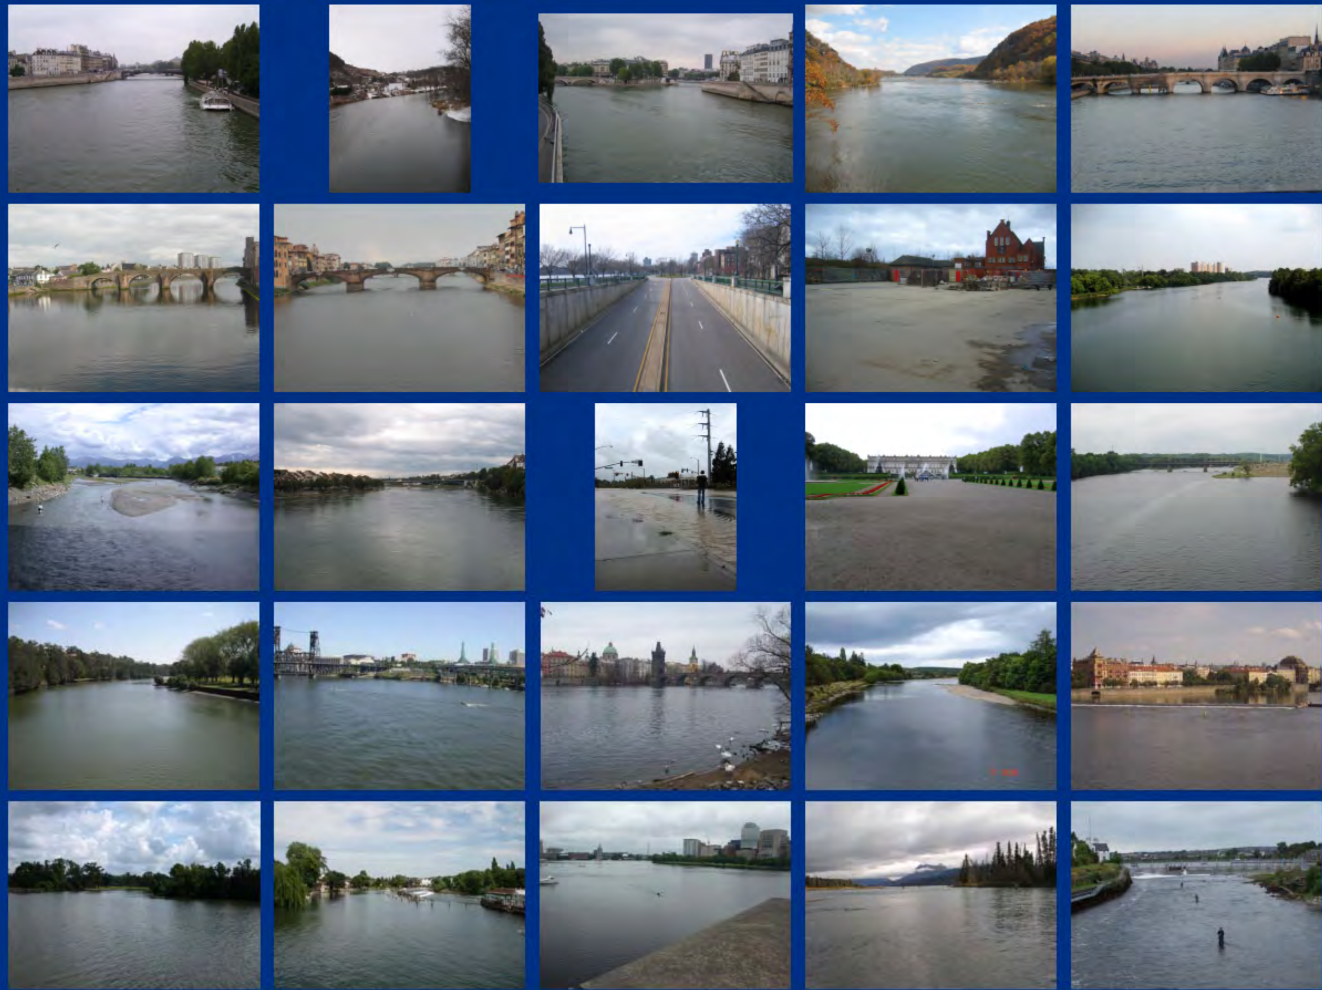

… 200 scene matches

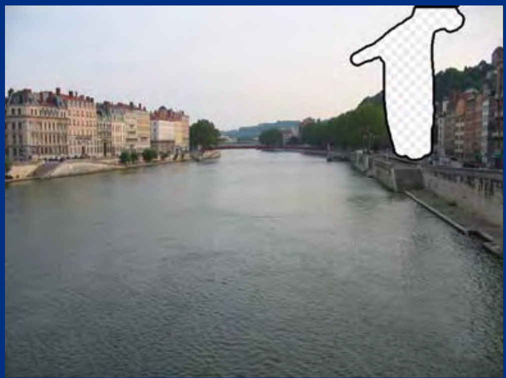

… 200 scene matches

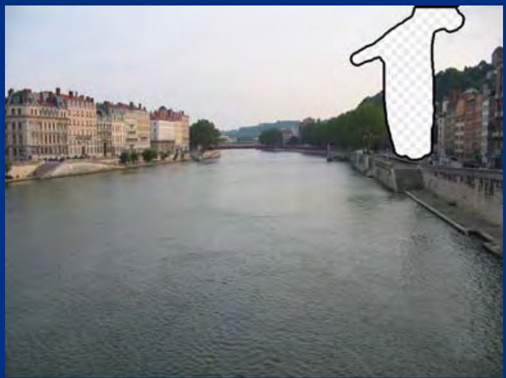

… 200 scene matches

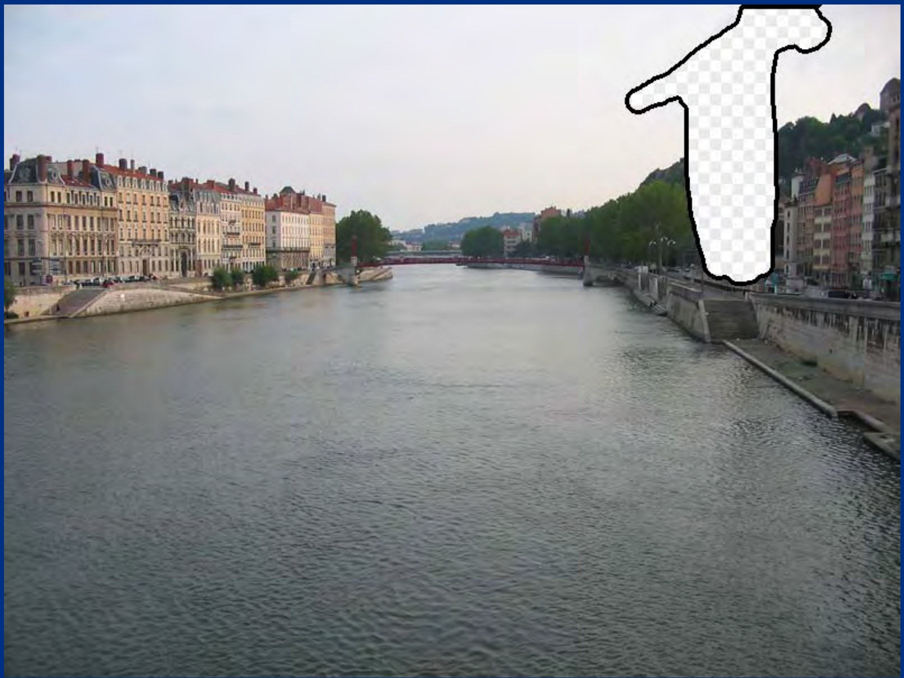


七五三詣り
入口

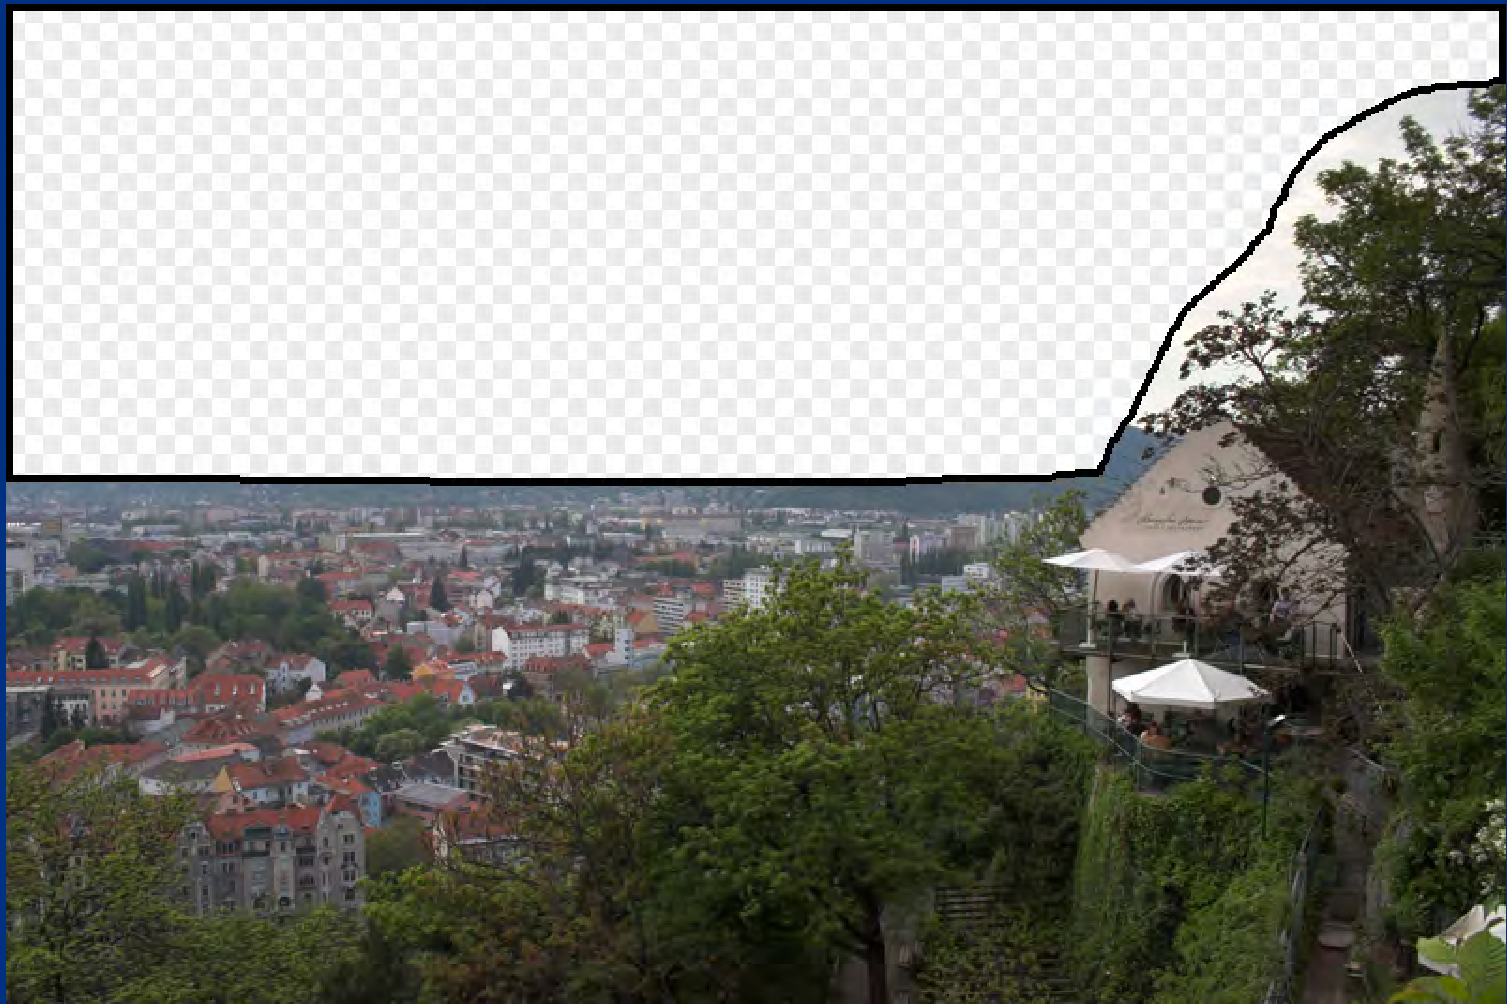

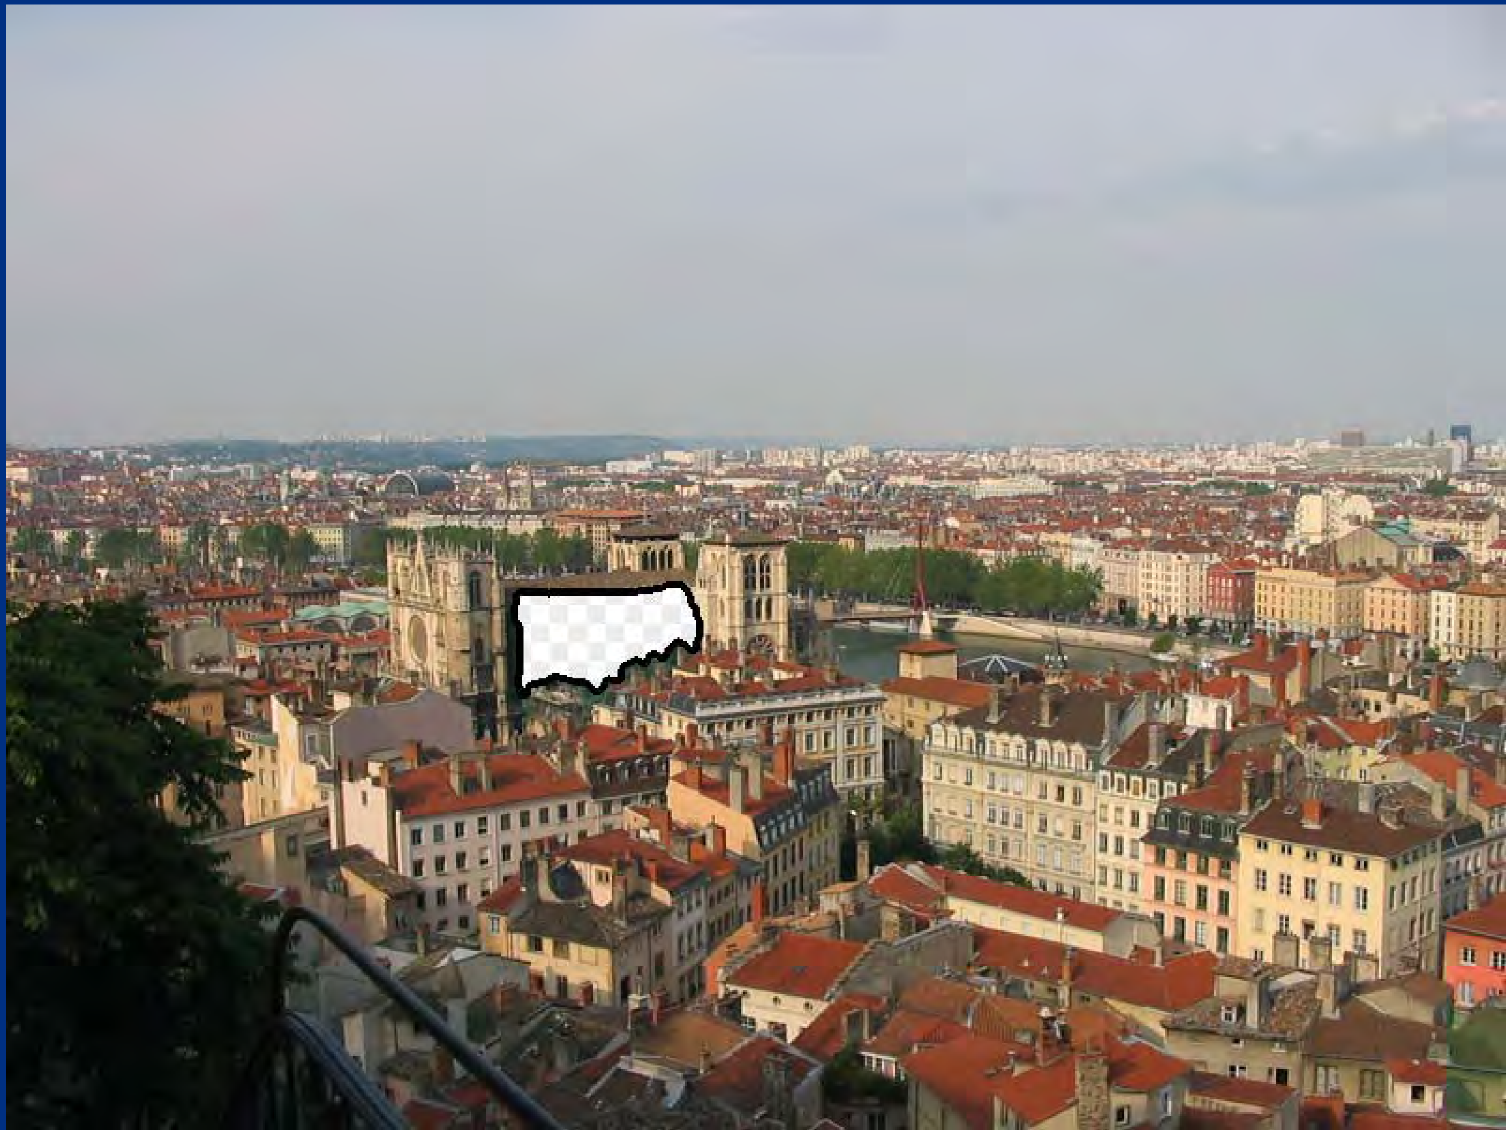

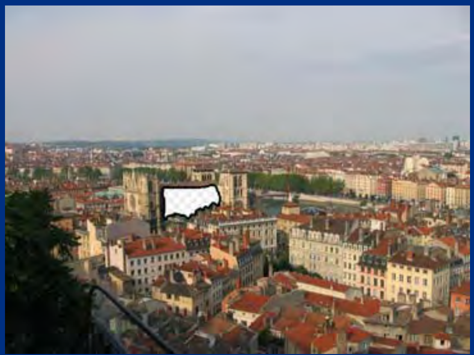

… 200 scene matches

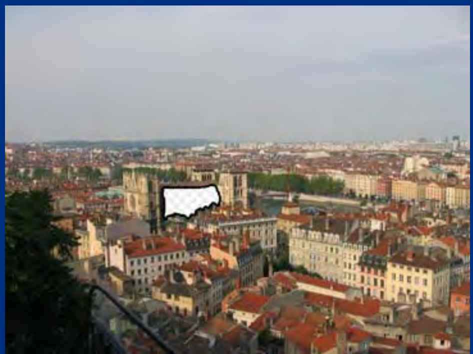

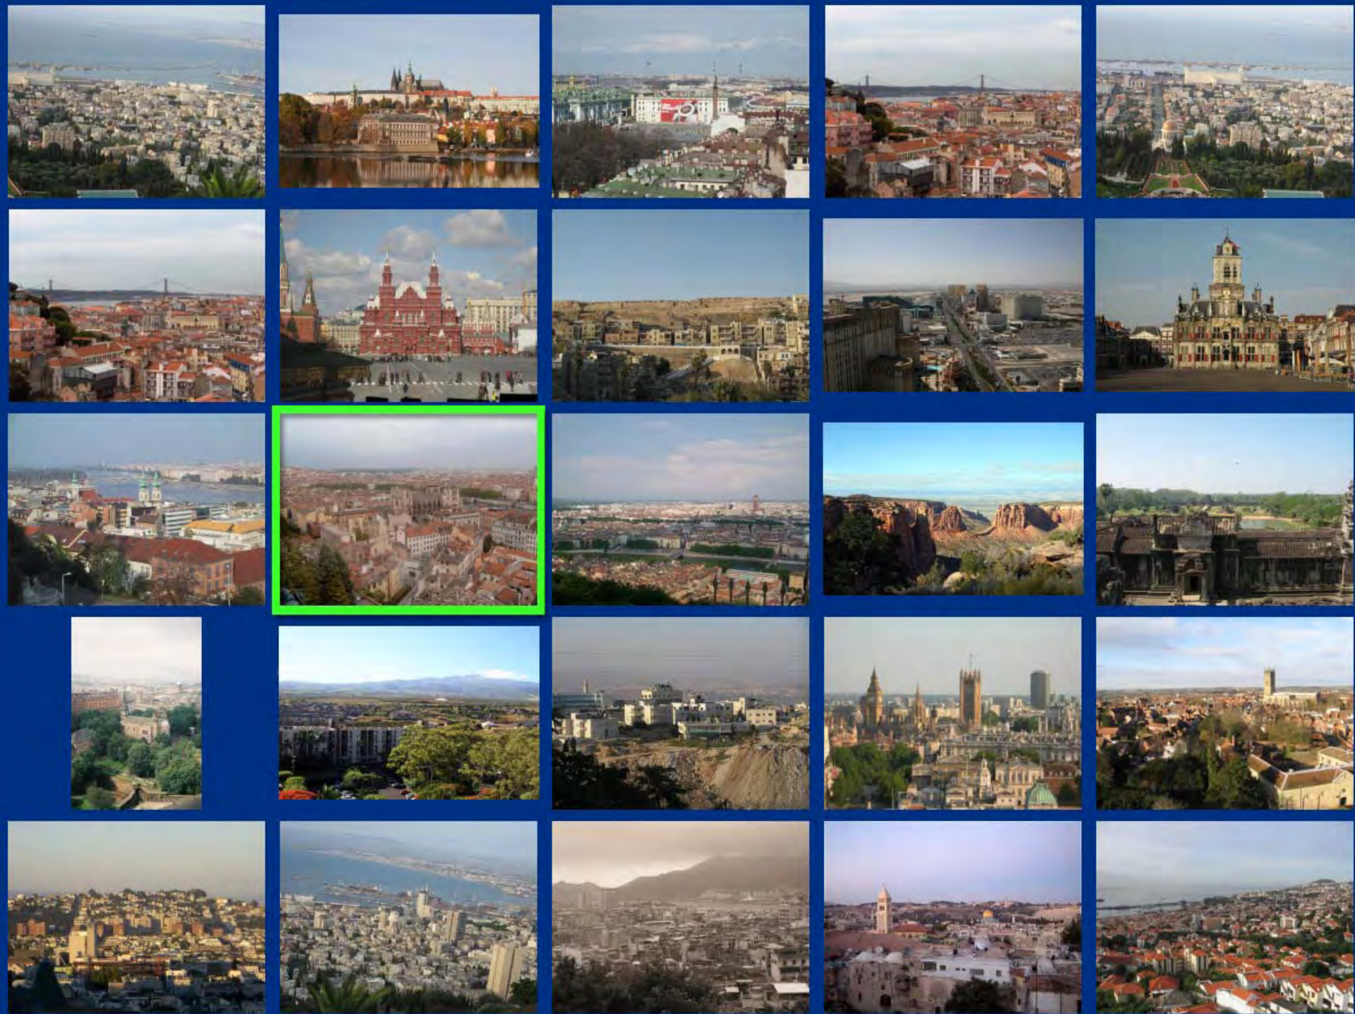

… 200 scene matches

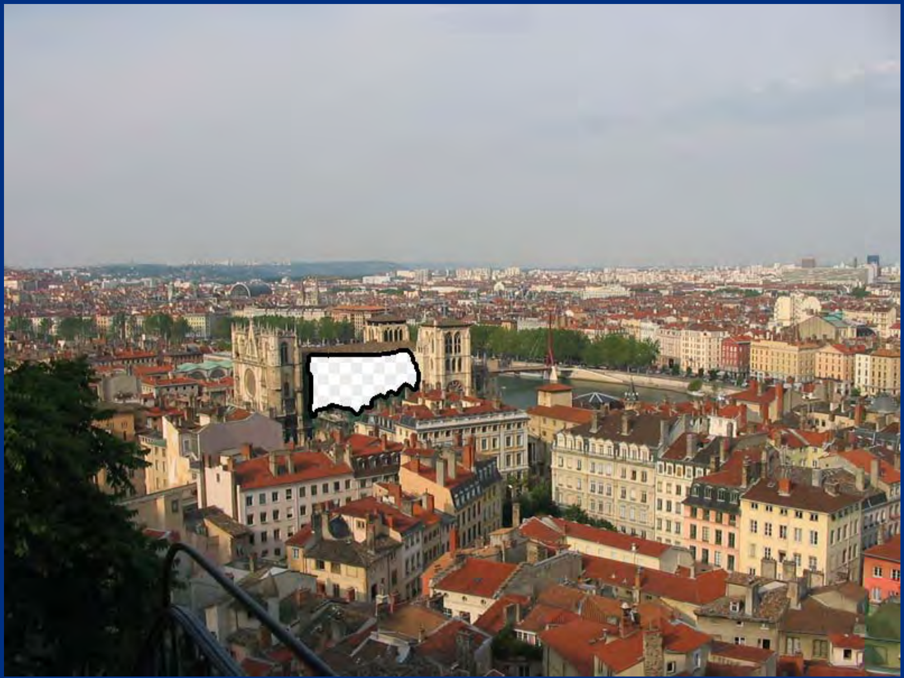

Failures

Failures

Failures

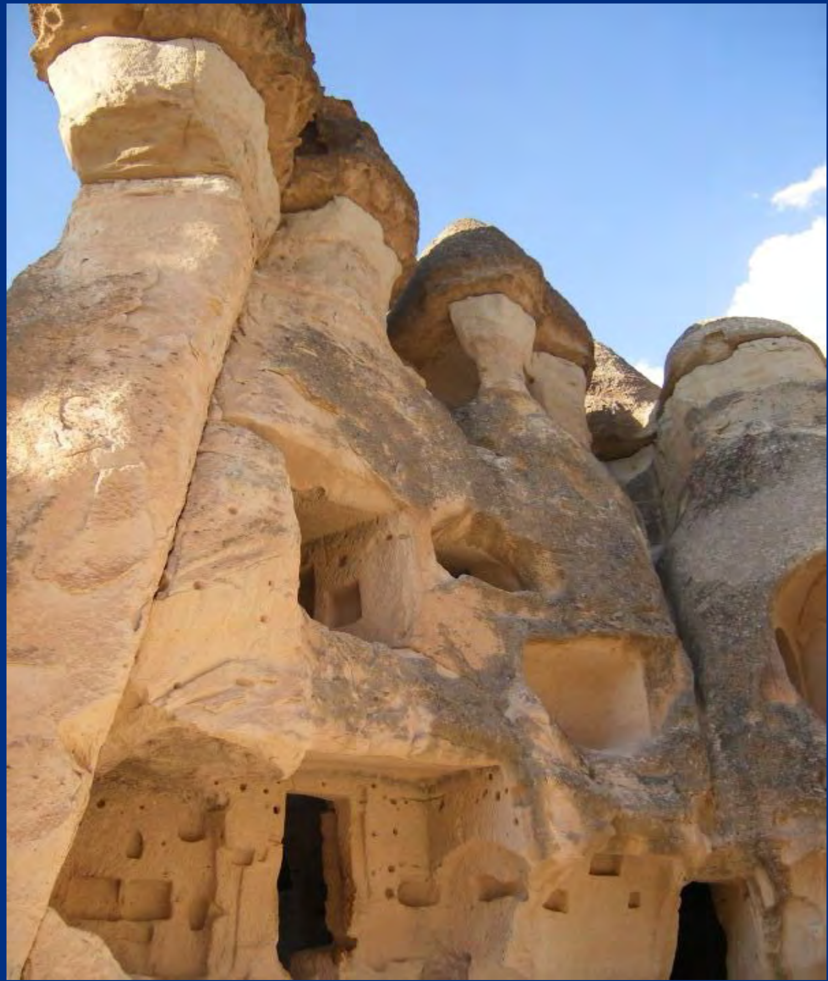

Failures

Failures

Failures

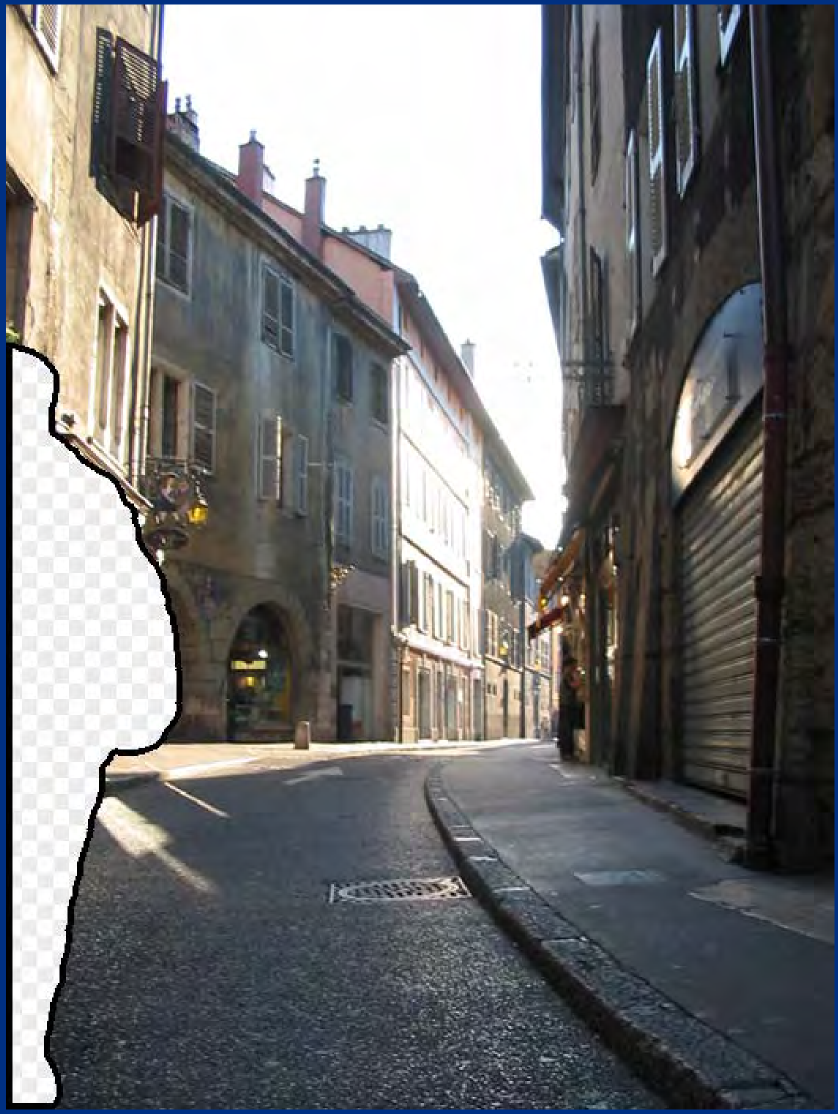

Failures

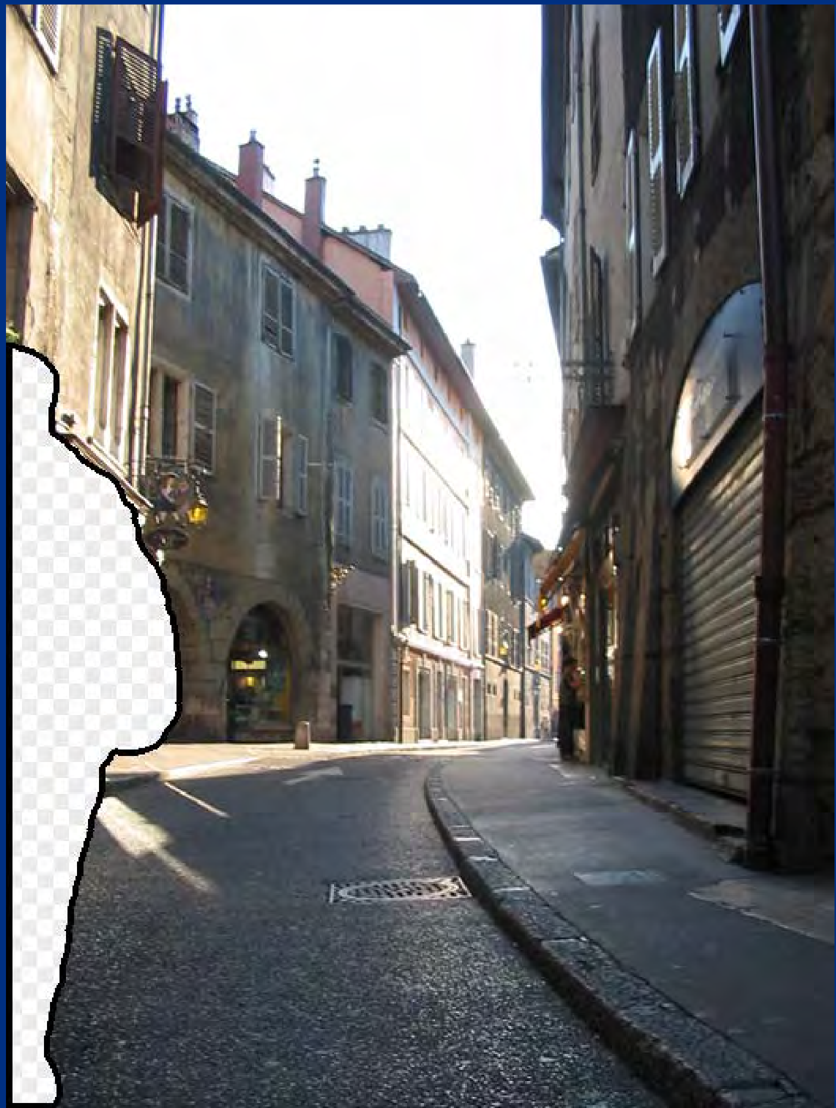

Failures

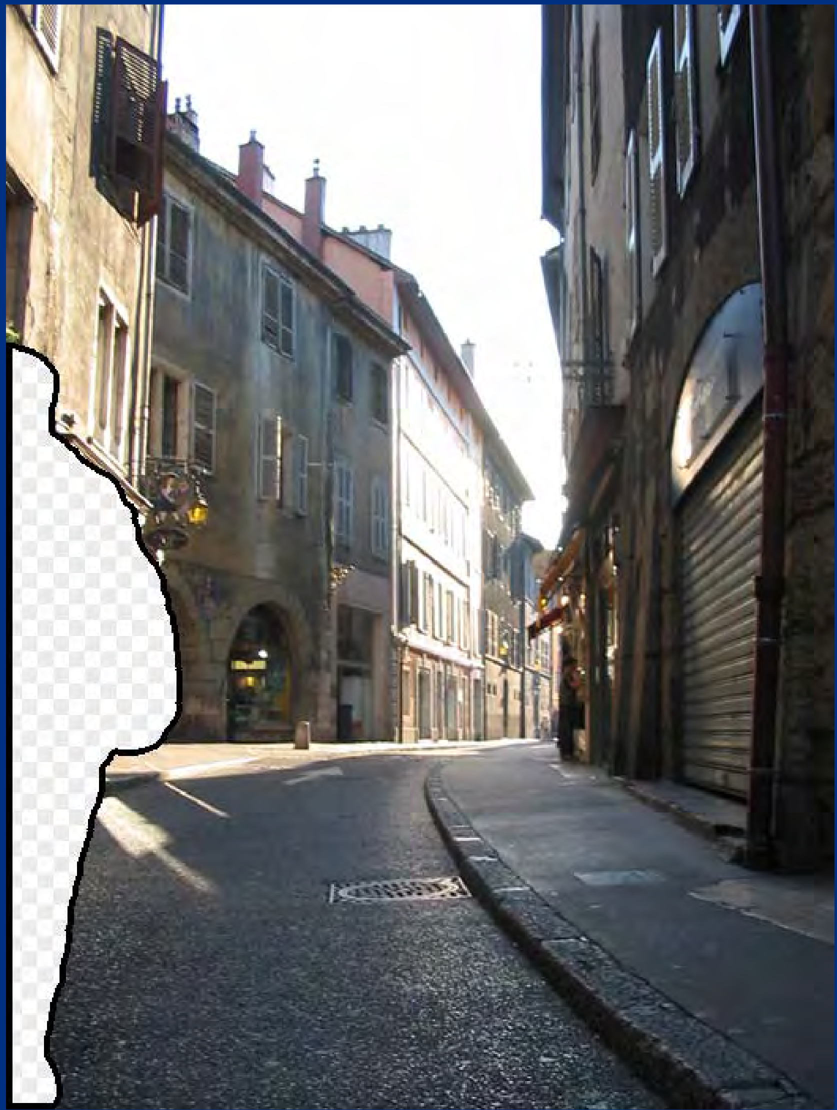

Failures

Failures

Failures

Failures

Evaluation

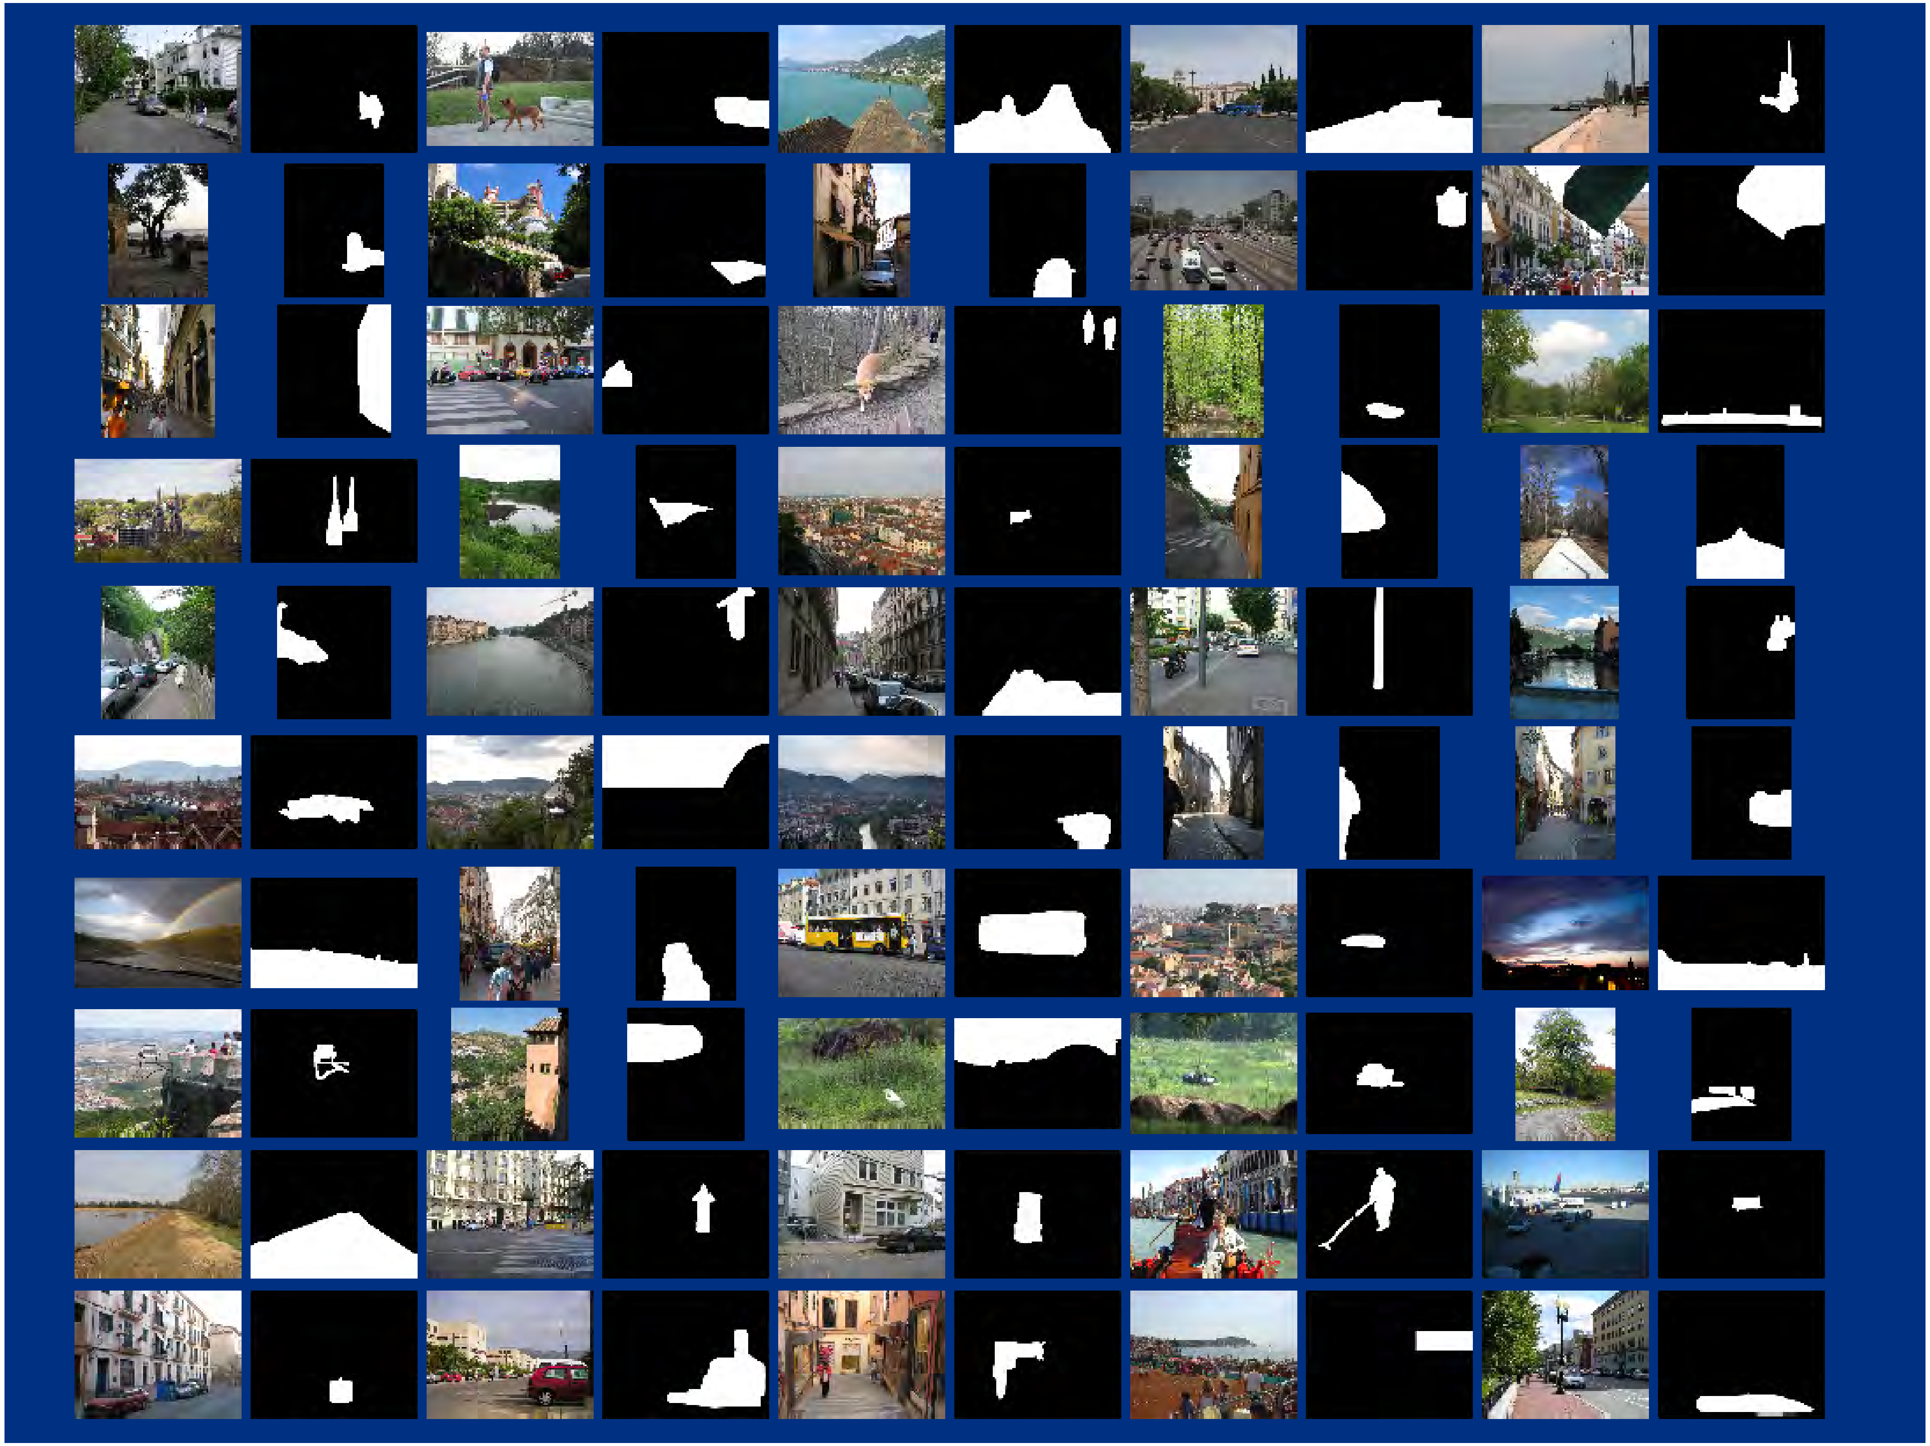

Original Images

Criminisi et al.

Single result

Scene Completion

Each result
selected from 20

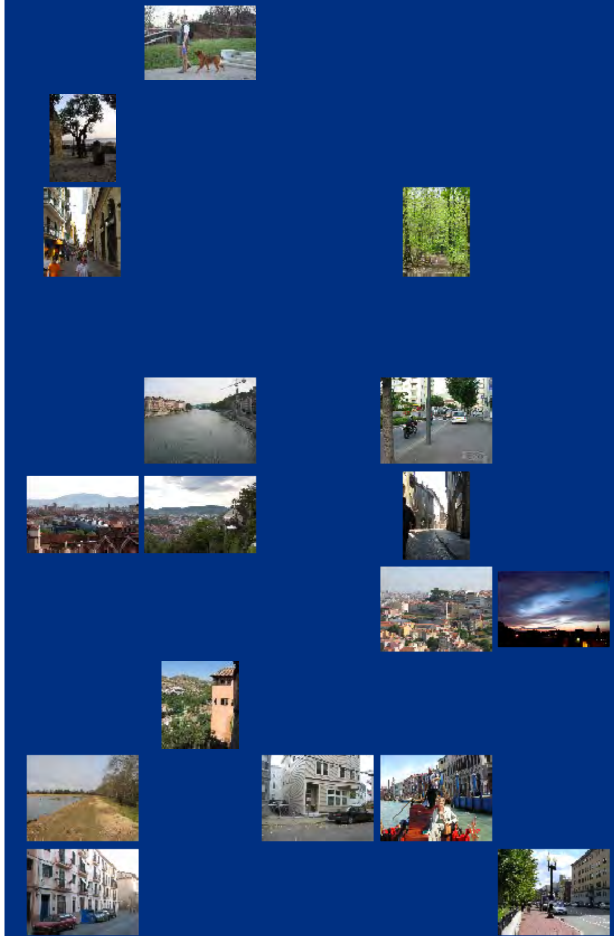

Original Images

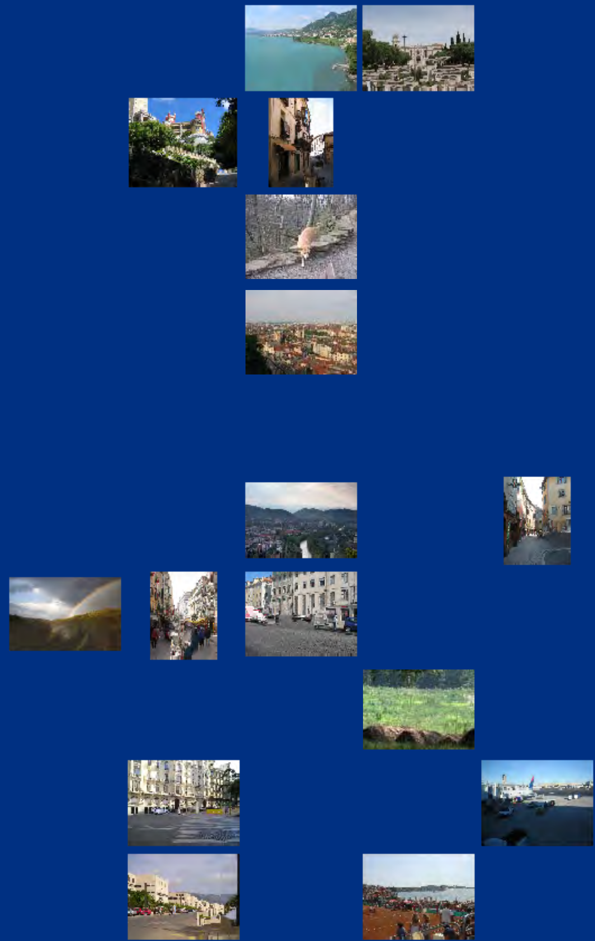

Criminisi et al.

Single result

Scene Completion

Each result
selected from 20

Real Image. This image has not been manipulated

or

Fake Image. This image has been manipulated

User Study Results - 20 Participants

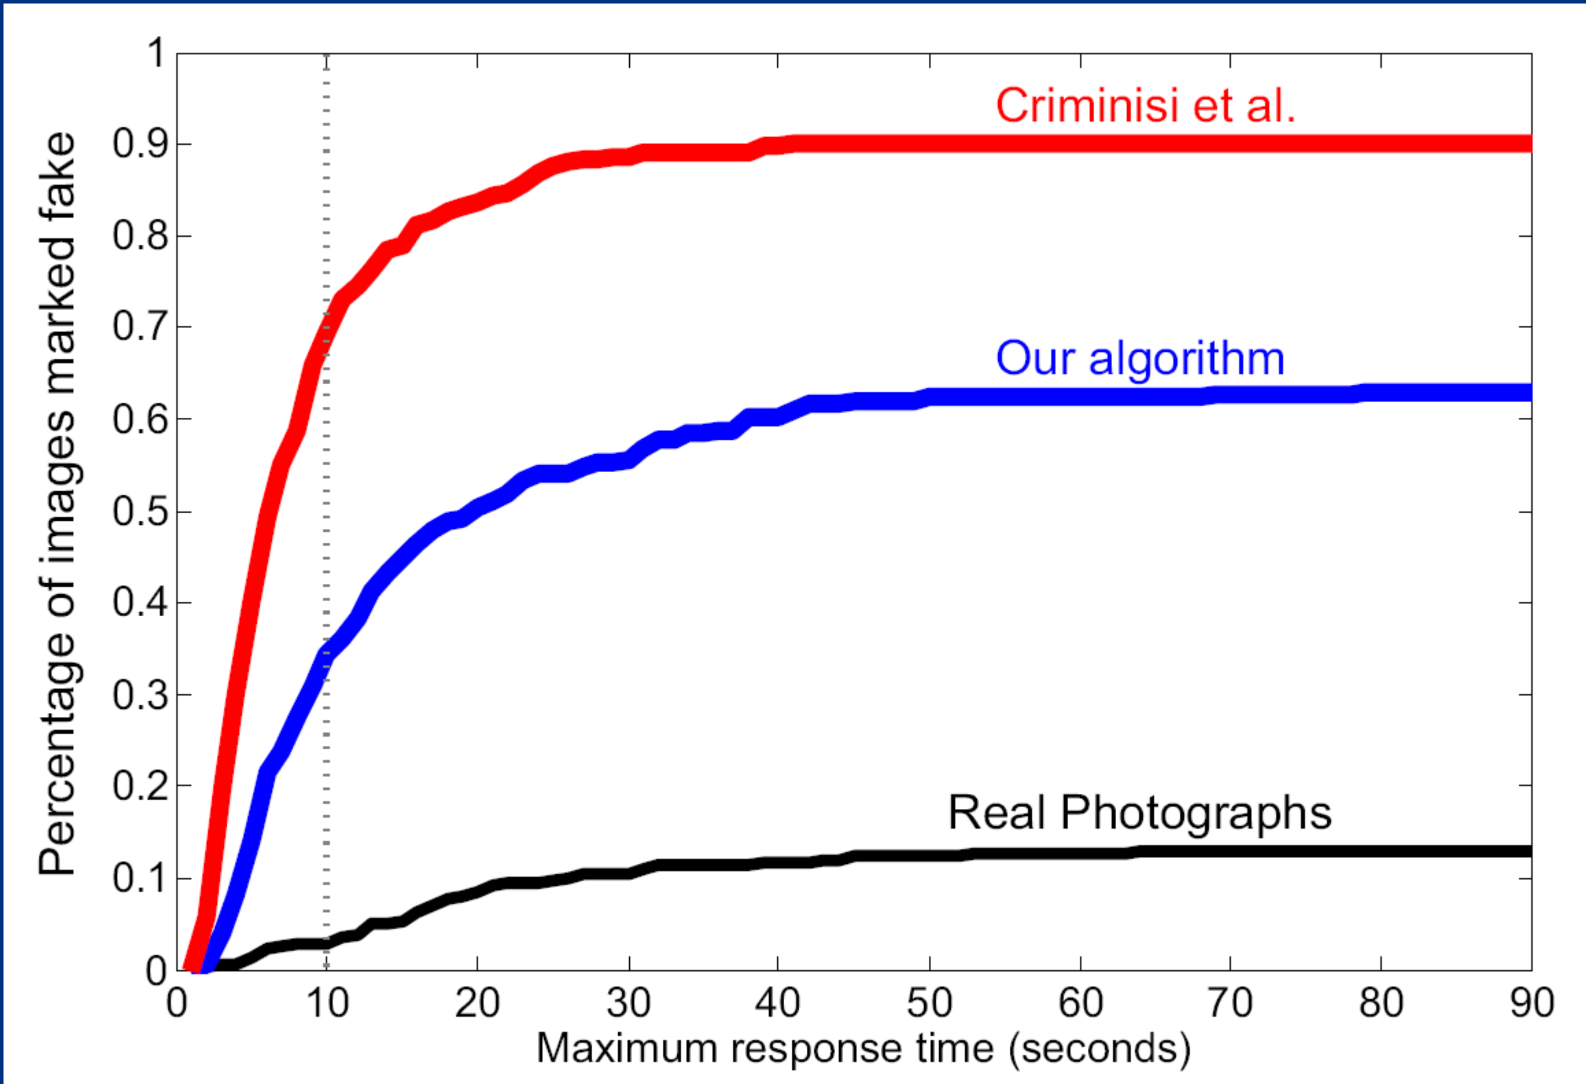

Why does it work?

10 nearest neighbors from a
collection of 20,000 images

10 nearest neighbors from a
collection of 2 million images

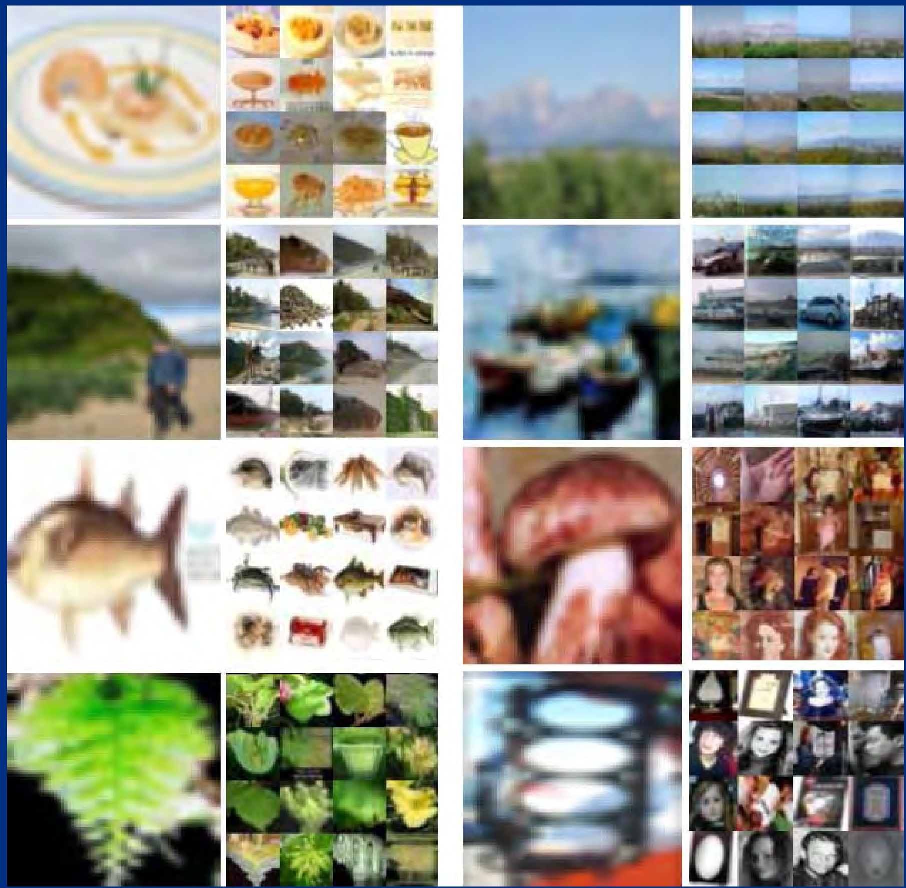

Database of 70 Million 32x32 images

Torralba, Fergus, and Freeman. Tiny Images. MIT-CSAIL-TR-2007-024. 2007.

The Small Picture

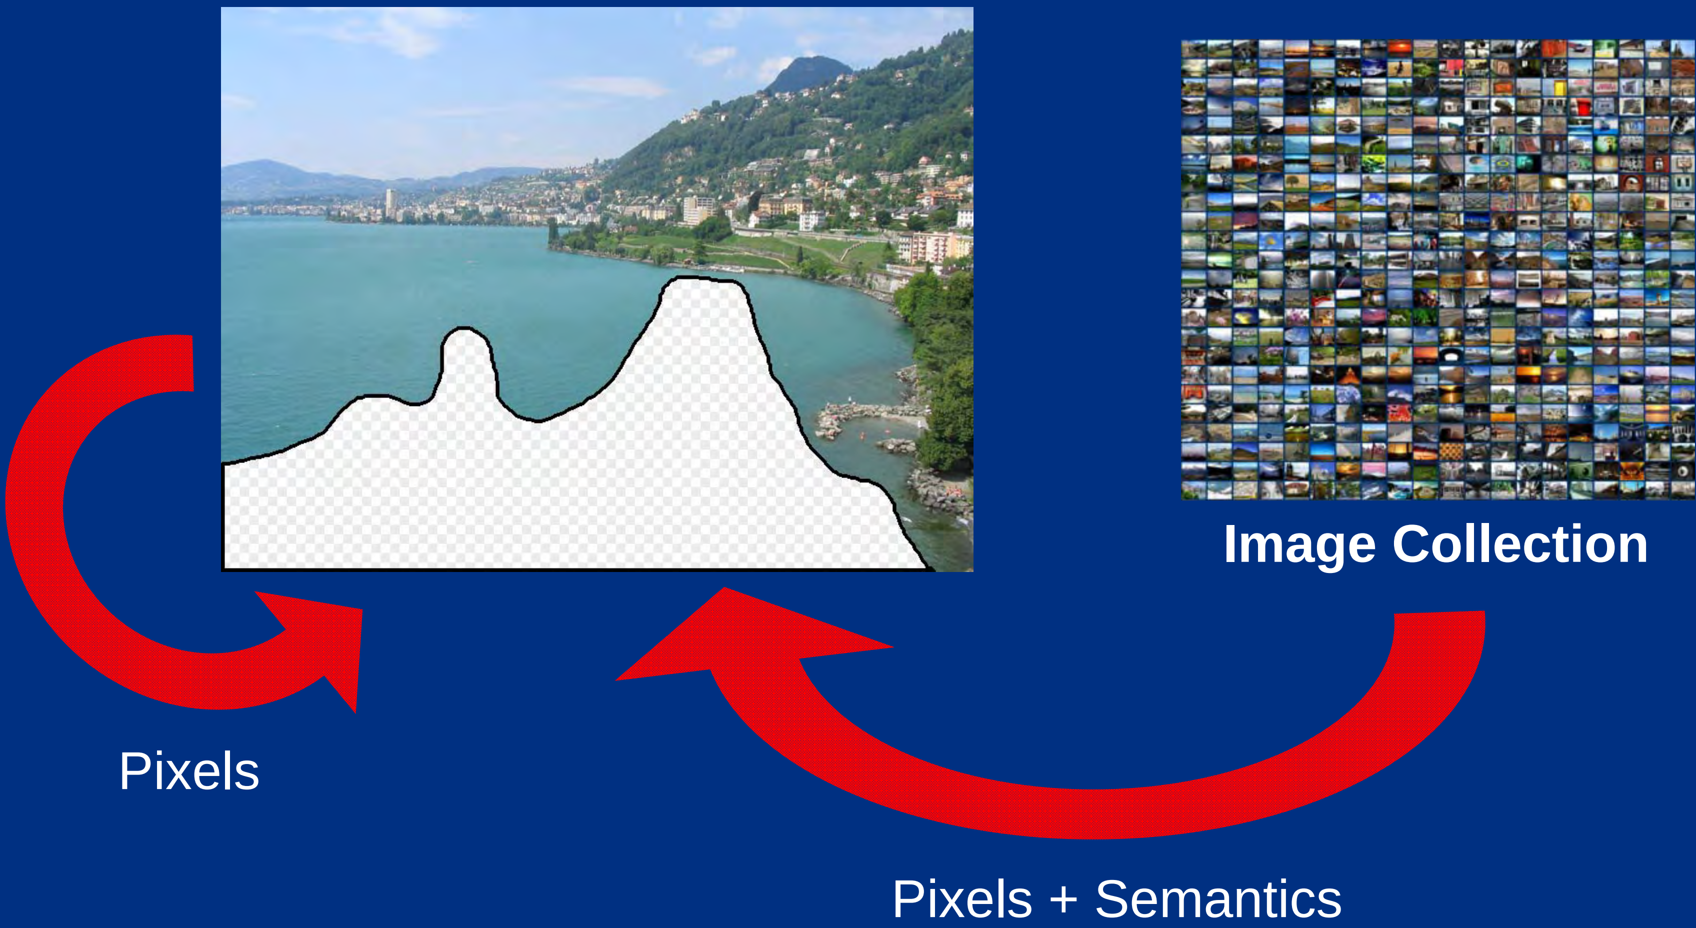

What Next? Small Picture

- Add outline of what we just presented
- We presented a very different hole-filling technique
- Sometimes works better than old stuff, but not always.
- Value in reusing original material. Just need semantics.

- Hybrid solution.

Hybrid Solution?

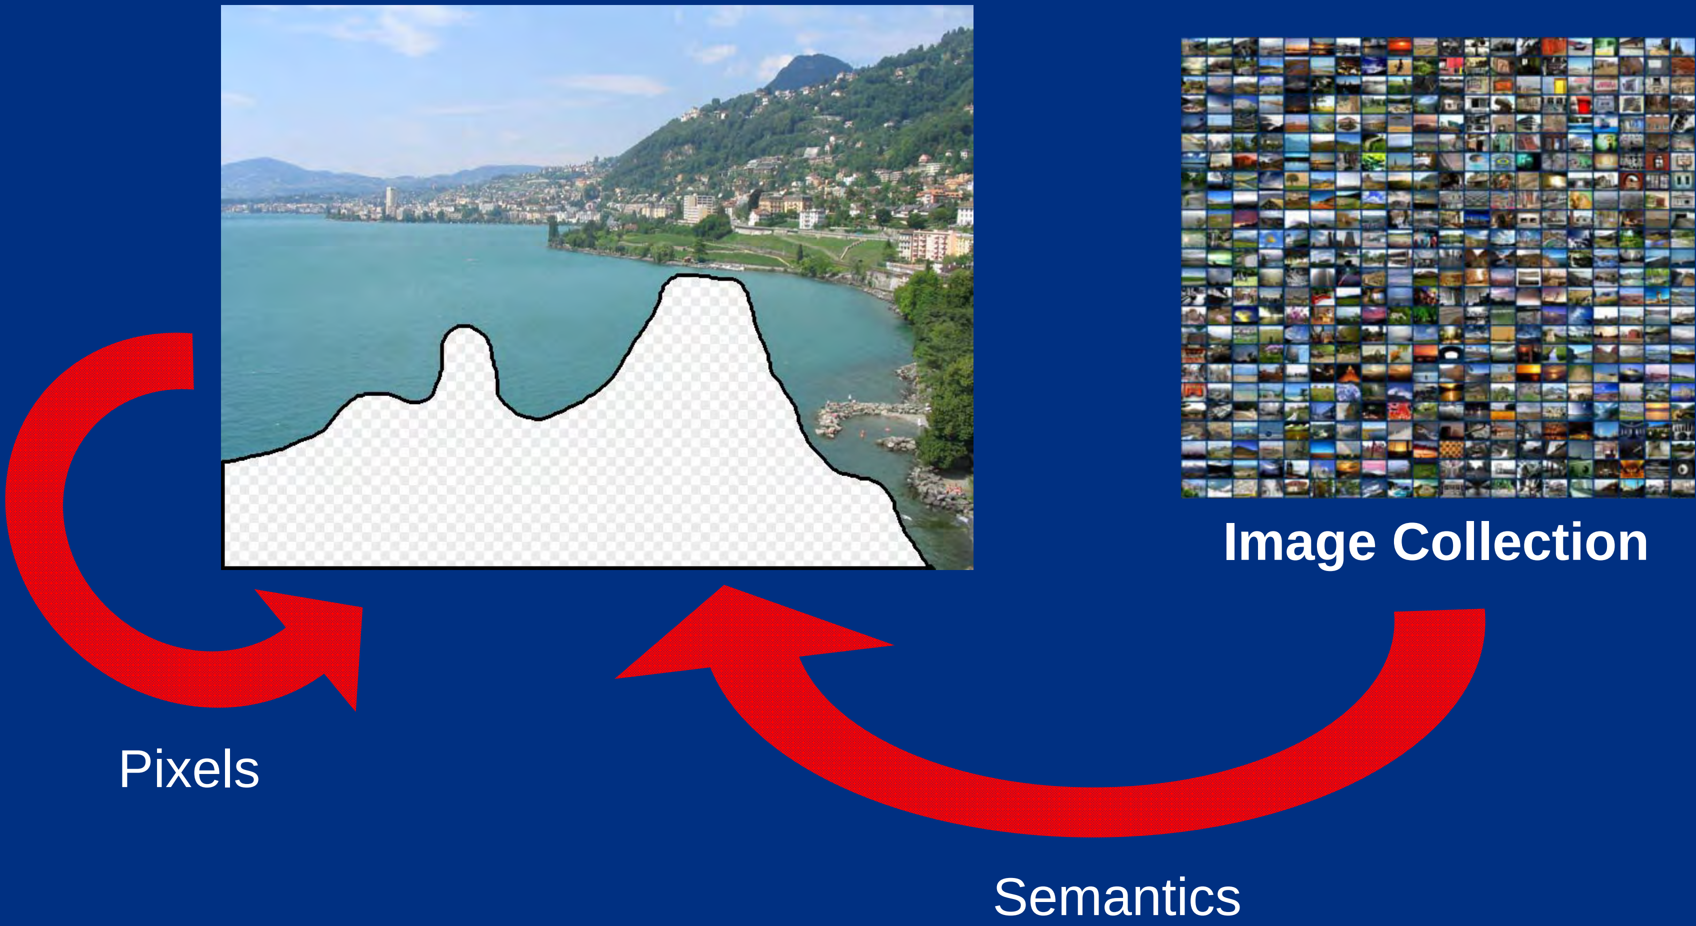

The Big Picture

Sky, Water, Hills, Beach,
Sunny, mid-day

Brute-force Image Understanding