Image Composition and Editing II

David Jacobs
CS 478 Winter 2012

Outline

- Alpha Matting
- Environment Matting
- Nearest Neighbor Search
- Image Collections
- Content-aware Retargeting and Resizing

Alpha Compositing

- Compositing Images
  - Include an “alpha” channel representing opacity
  - Compositing RGBA image A on top of B

A

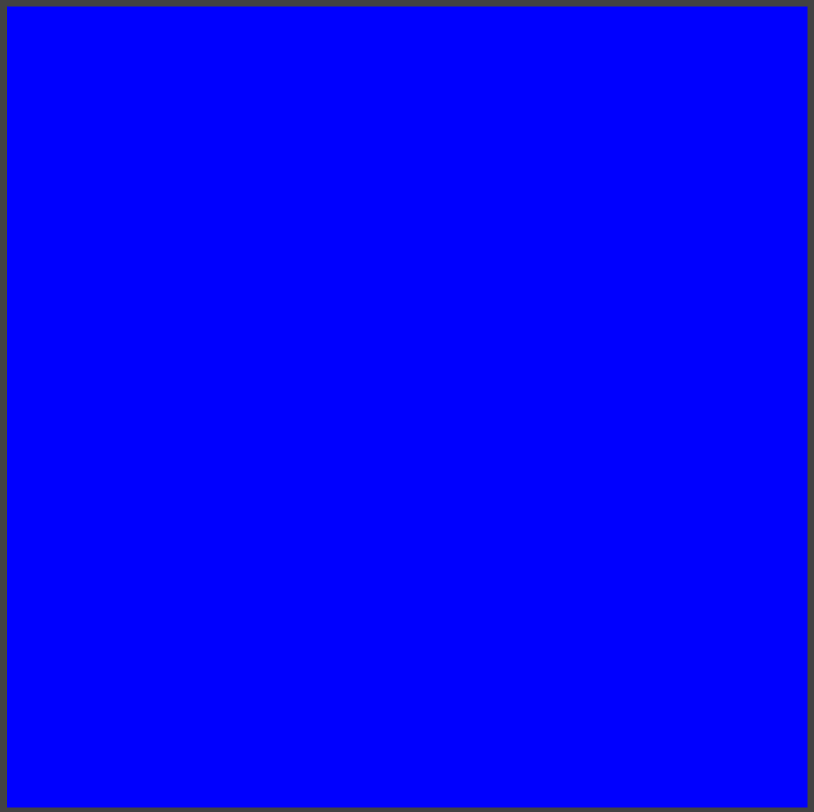

RGB

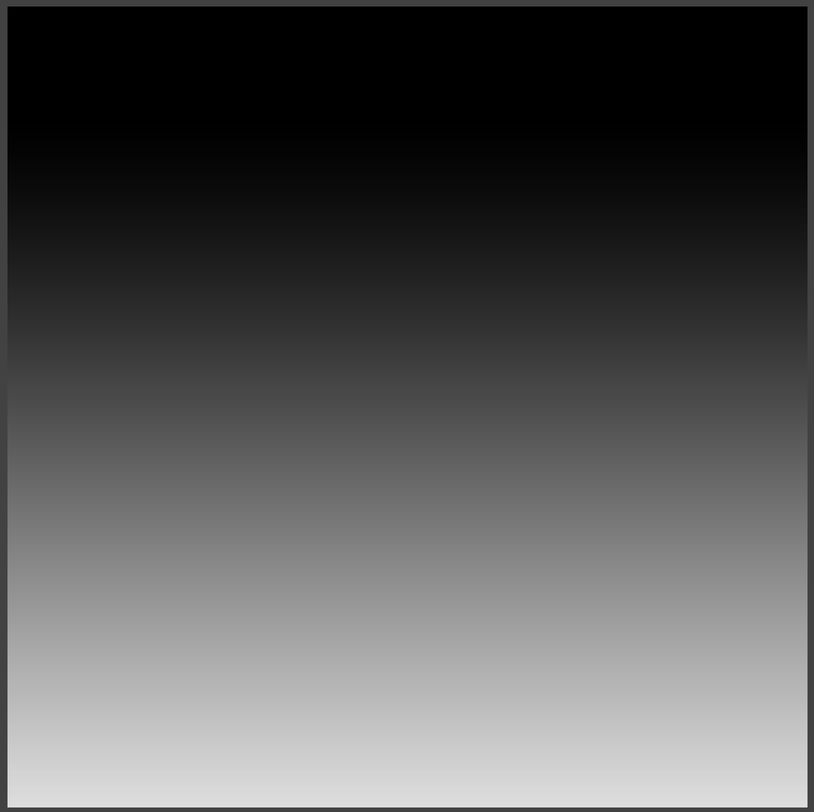

α

Alpha Compositing

- Compositing Images
  - Include an “alpha” channel representing opacity
  - Compositing RGBA image A on top of B

$$C = \alpha A + (1 - \alpha) B$$

A

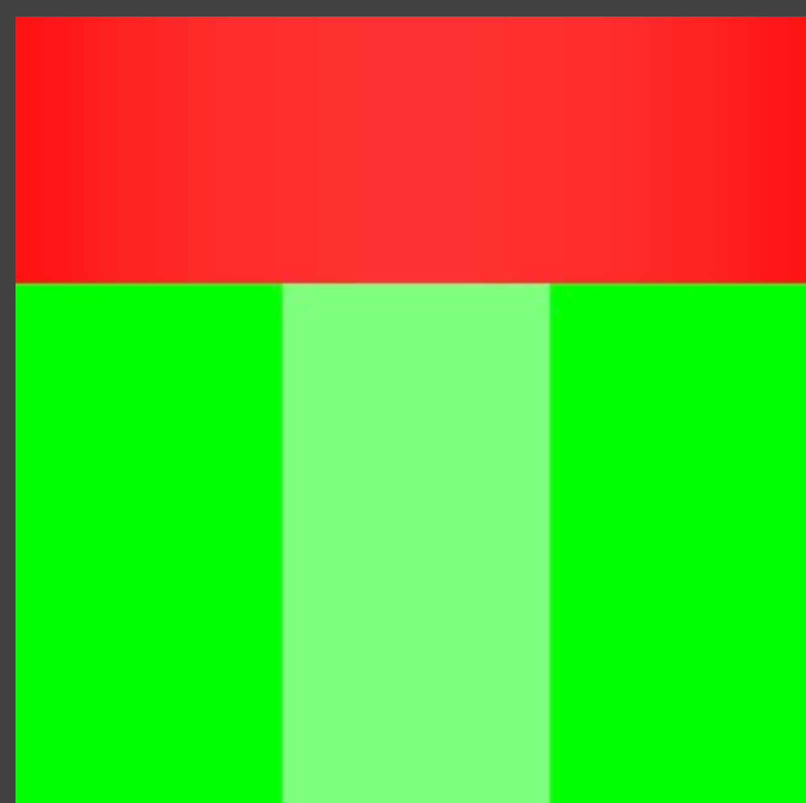

B

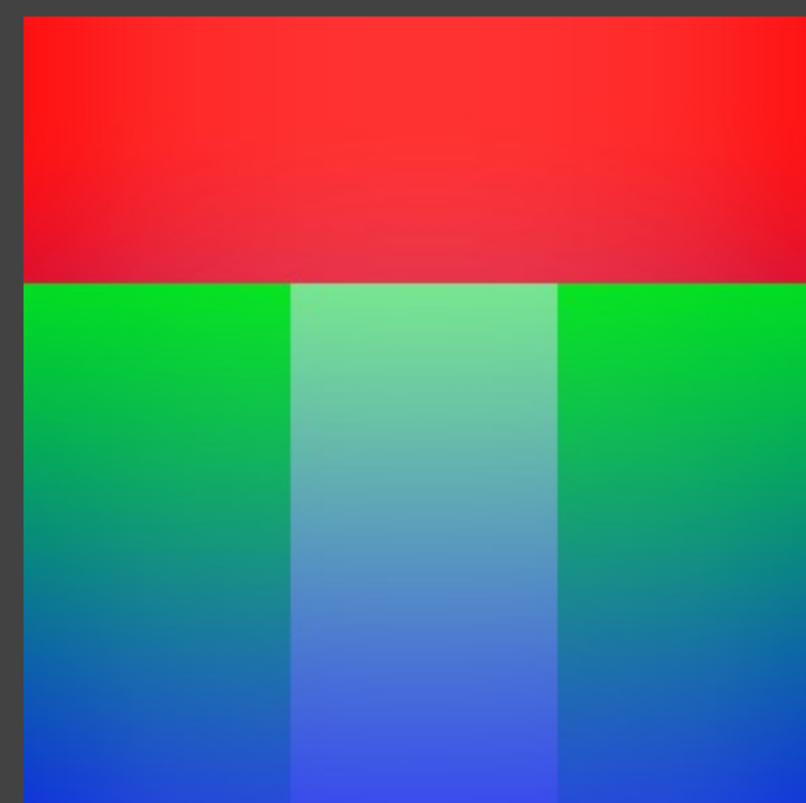

C

Alpha Matting

- Decomposing images into layers
  - From C, find A (including α) and B.
  - Much harder problem

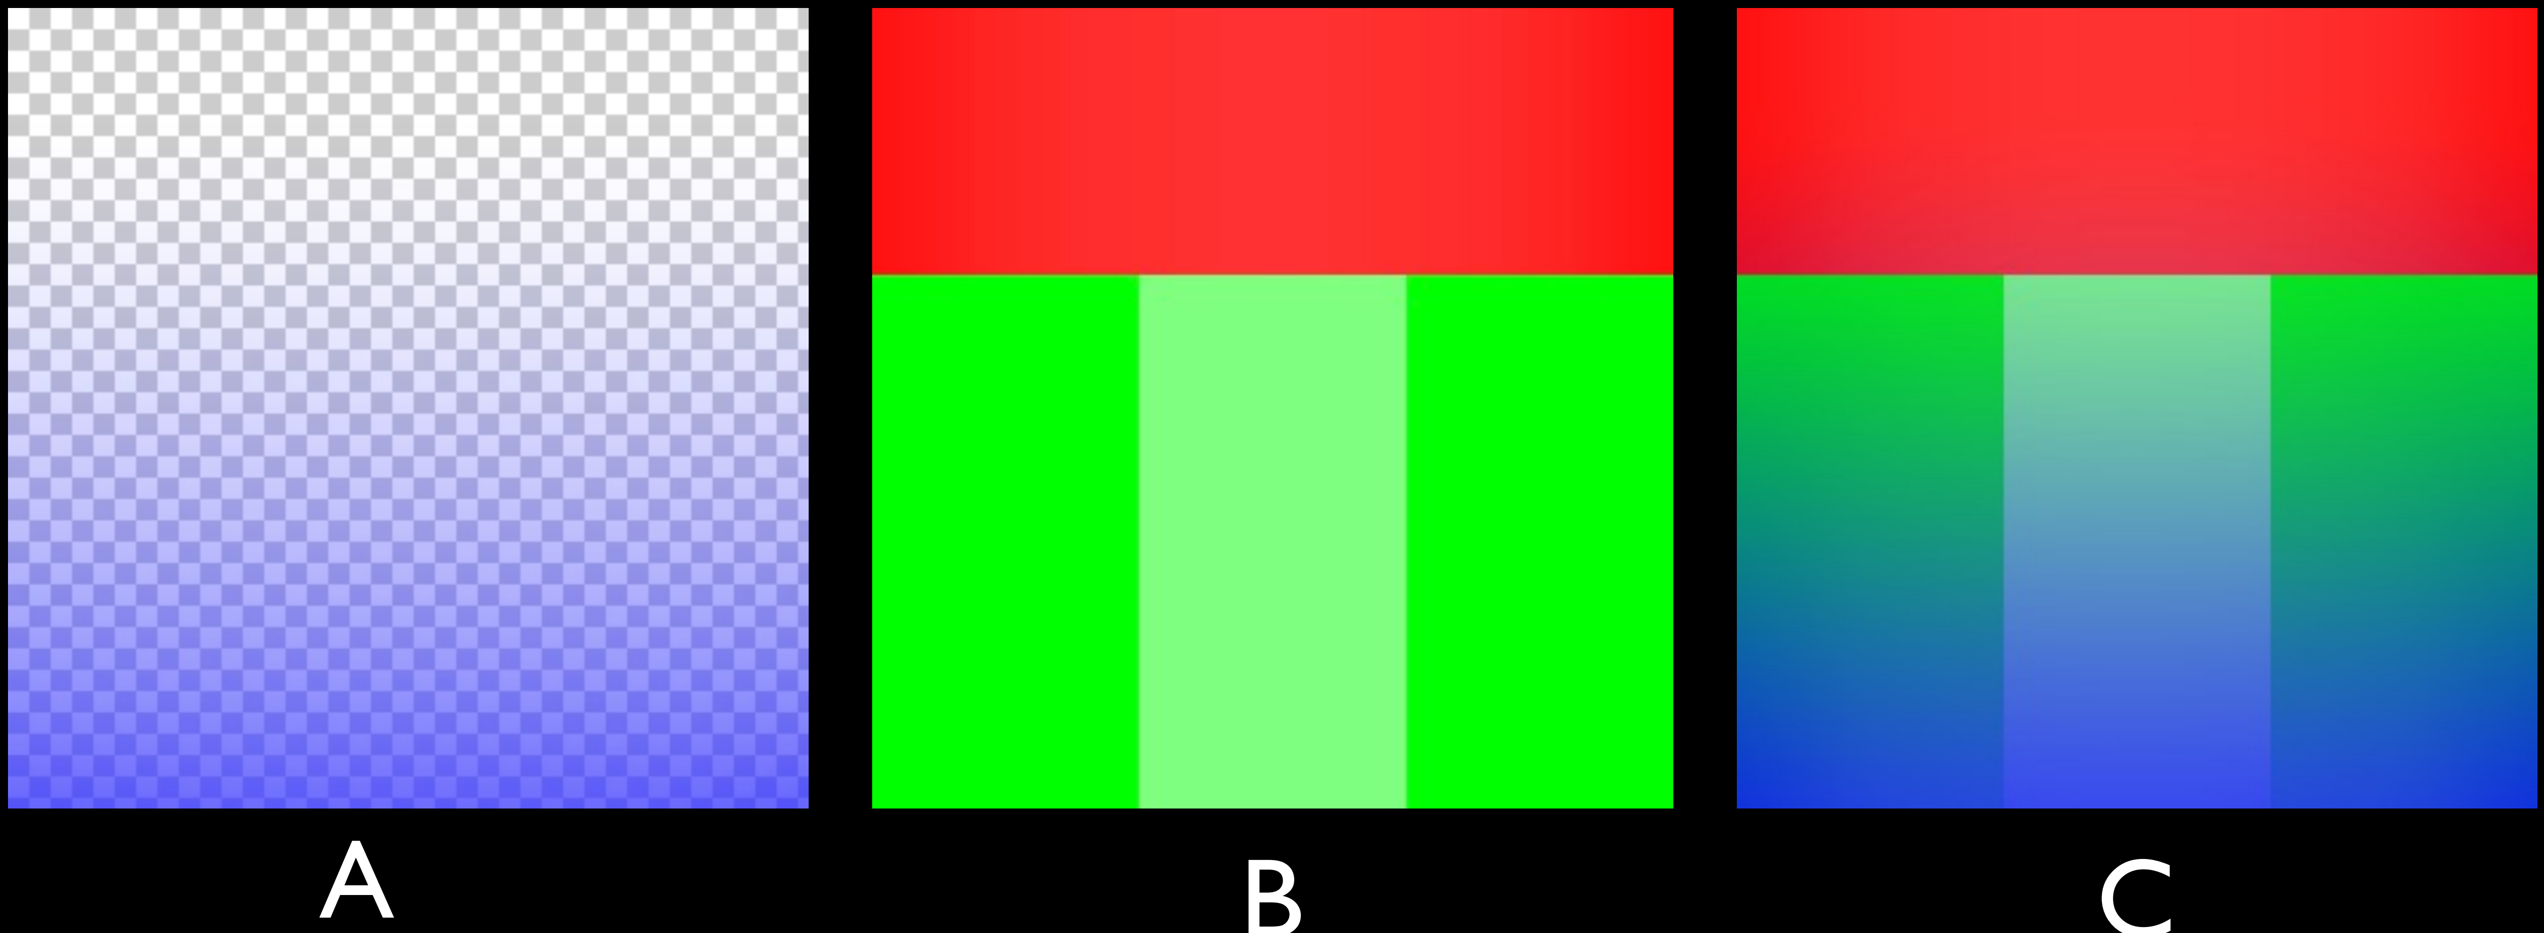

Soft Scissors

[Wang et al. 2007]

Soft Scissors : An Interactive Tool for Realtime High Quality Matting

Jue Wang
University of Washington

Maneesh Agrawala
University of California, Berkeley

Michael F. Cohen
Microsoft Research

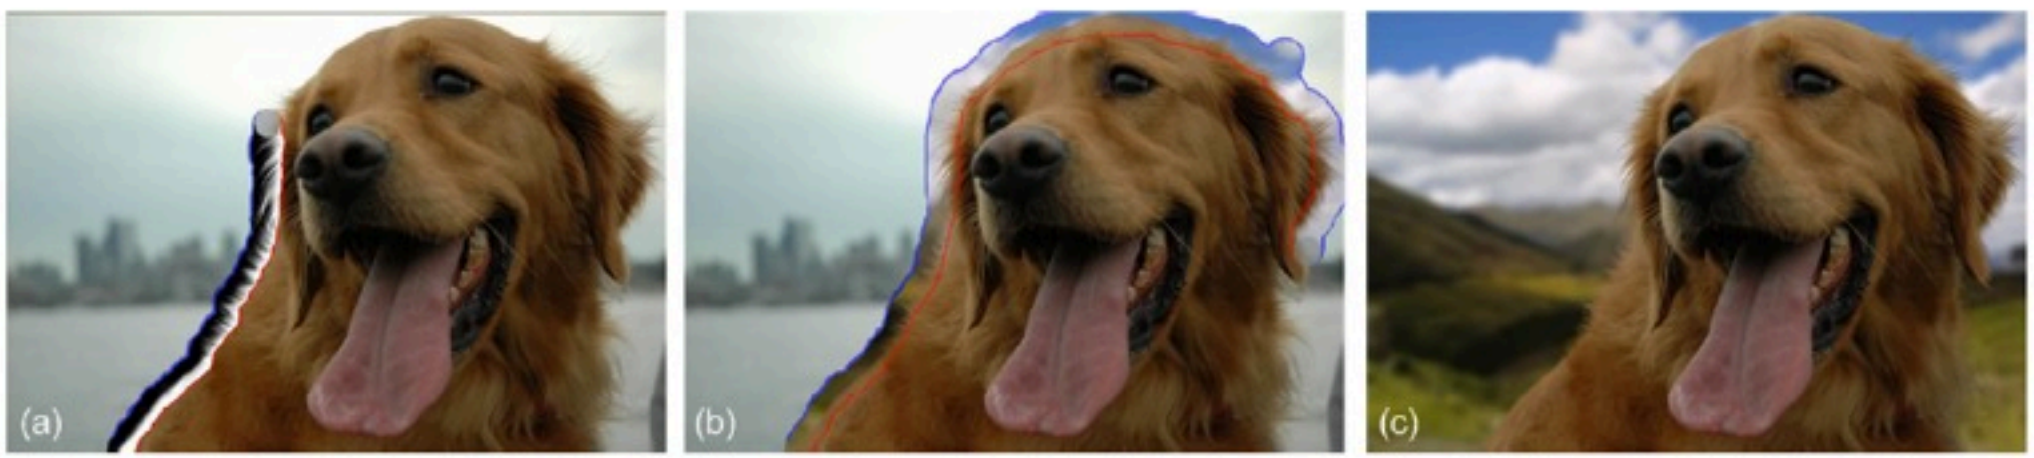

Figure 1: Our system computes a high quality matte (a) and a novel composite (b) in realtime as the user roughly paints the foreground boundary. Our system makes is easy to create new composites (c) very quickly.

Soft Scissors

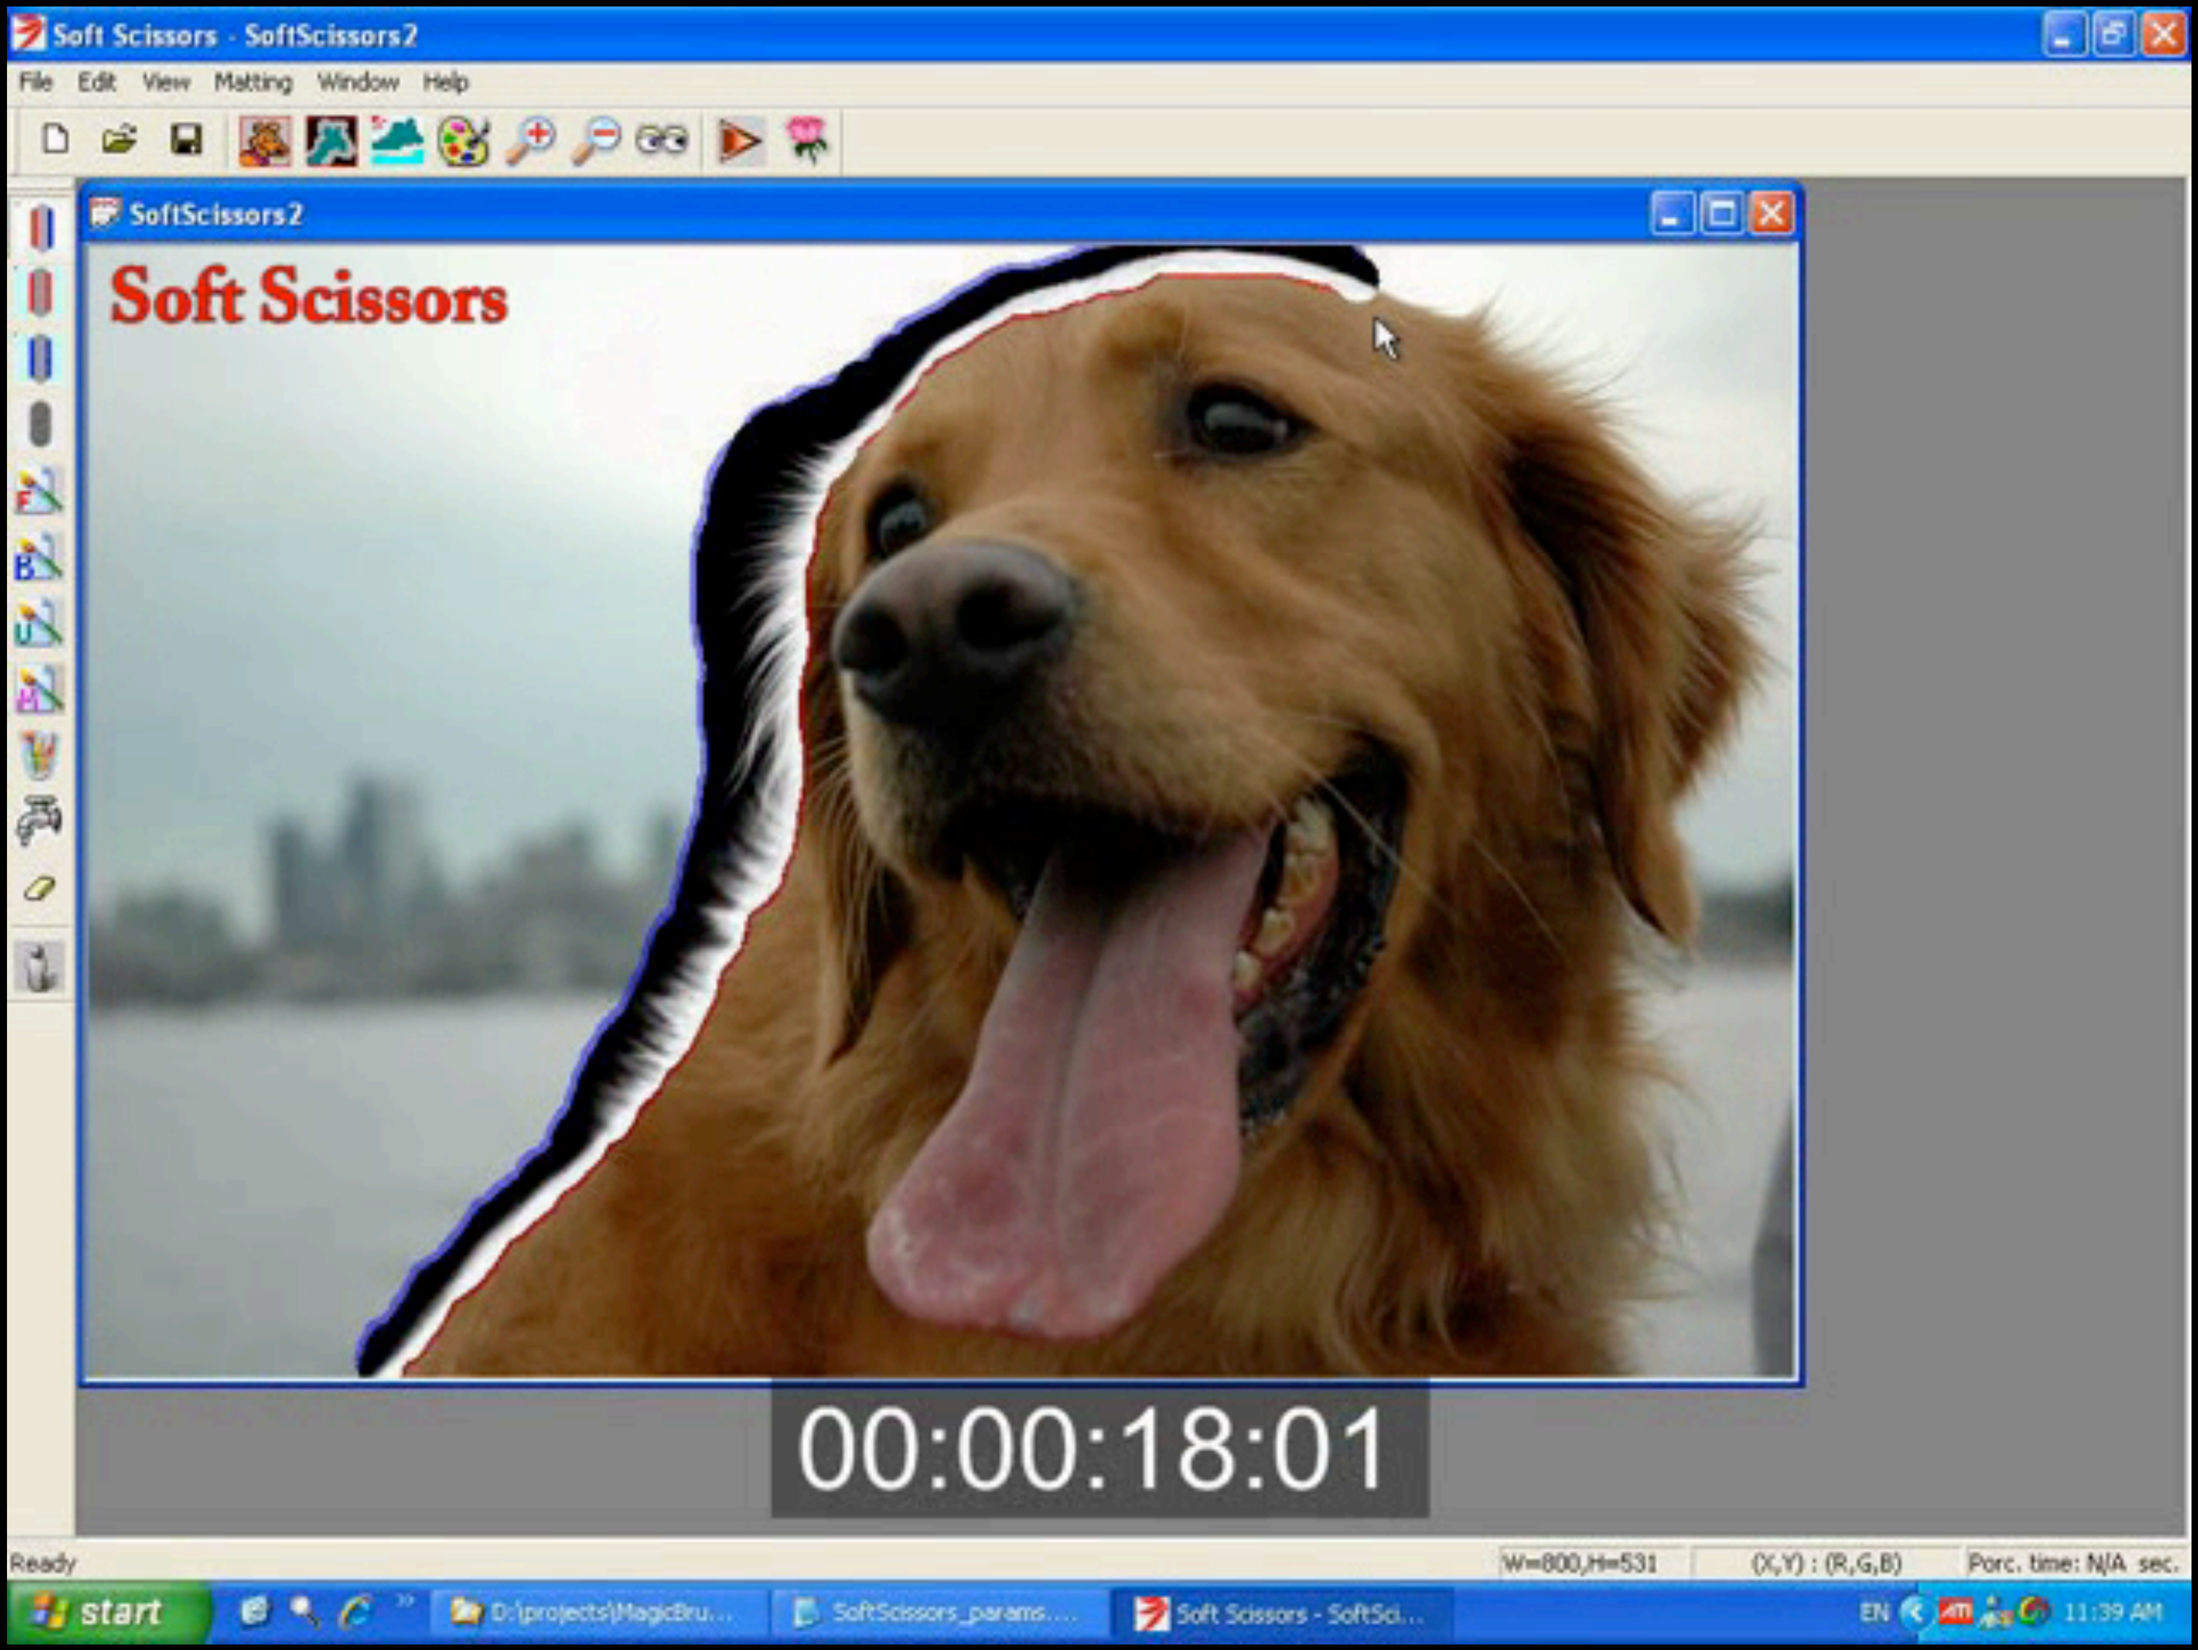

Soft Scissors

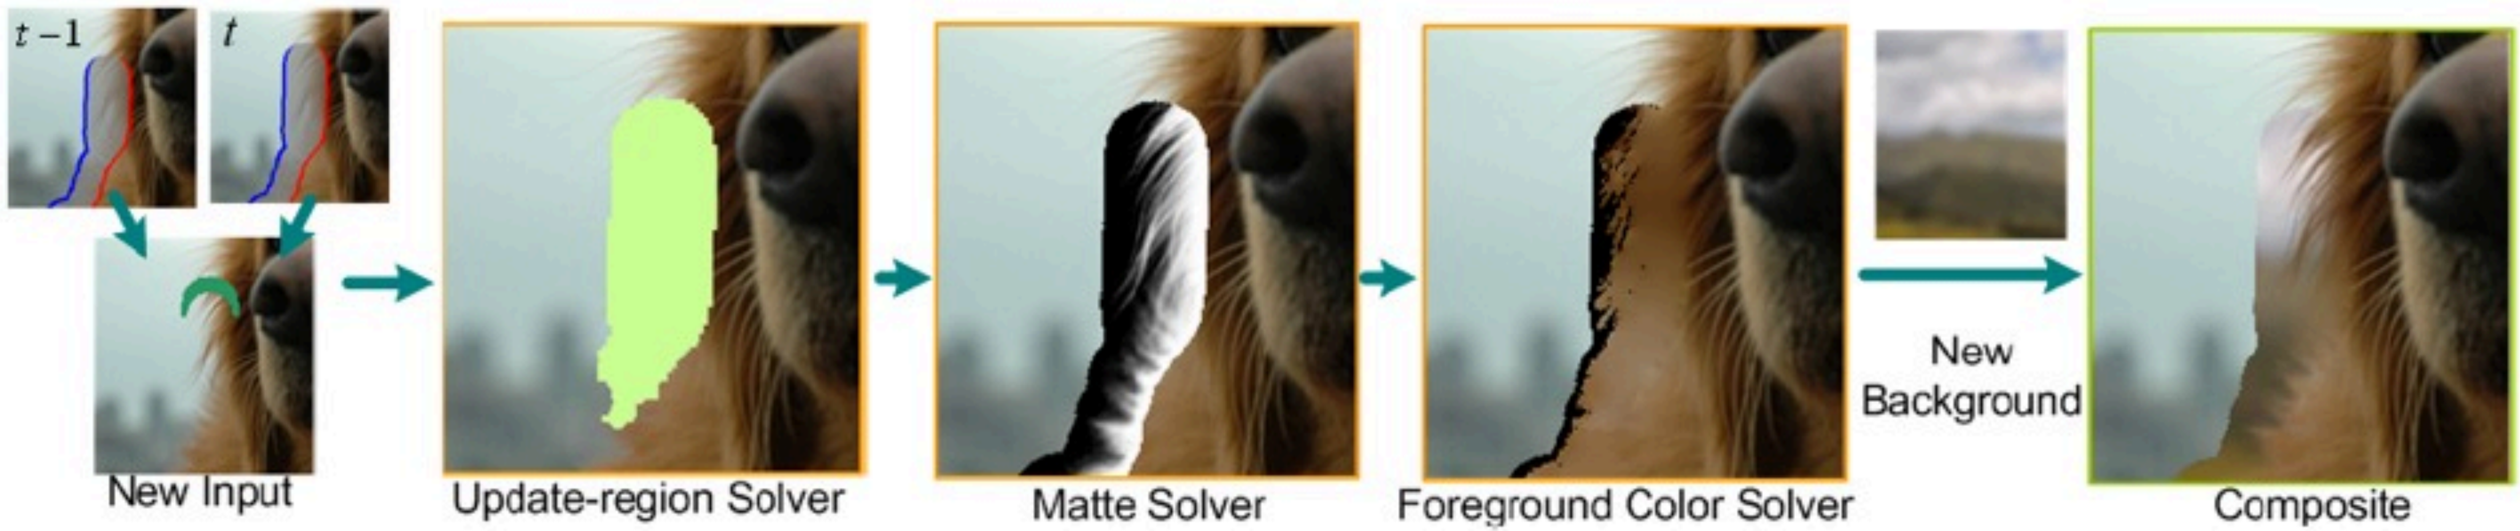

Figure 2: A flowchart of our system.

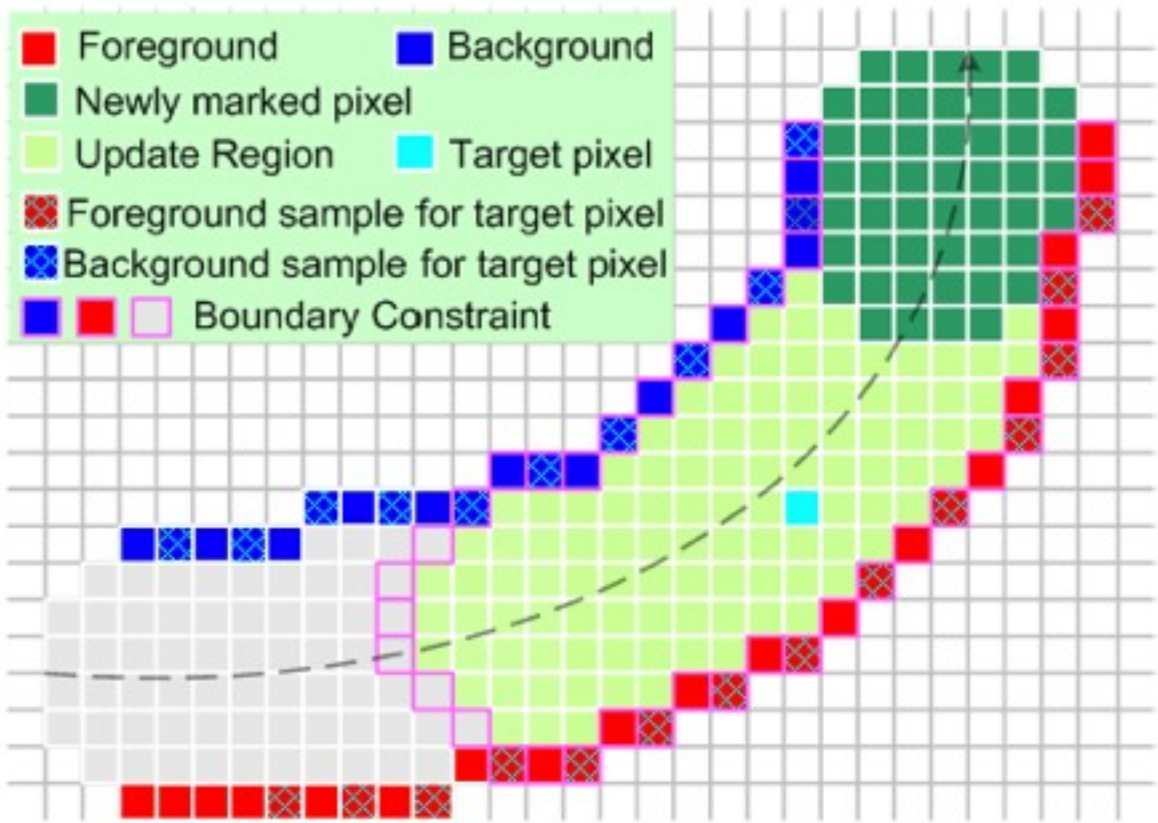

Figure 3: Our system quickly solves the matte under the leading edge of the soft scissors, constrained by boundary pixels.

Soft Scissors

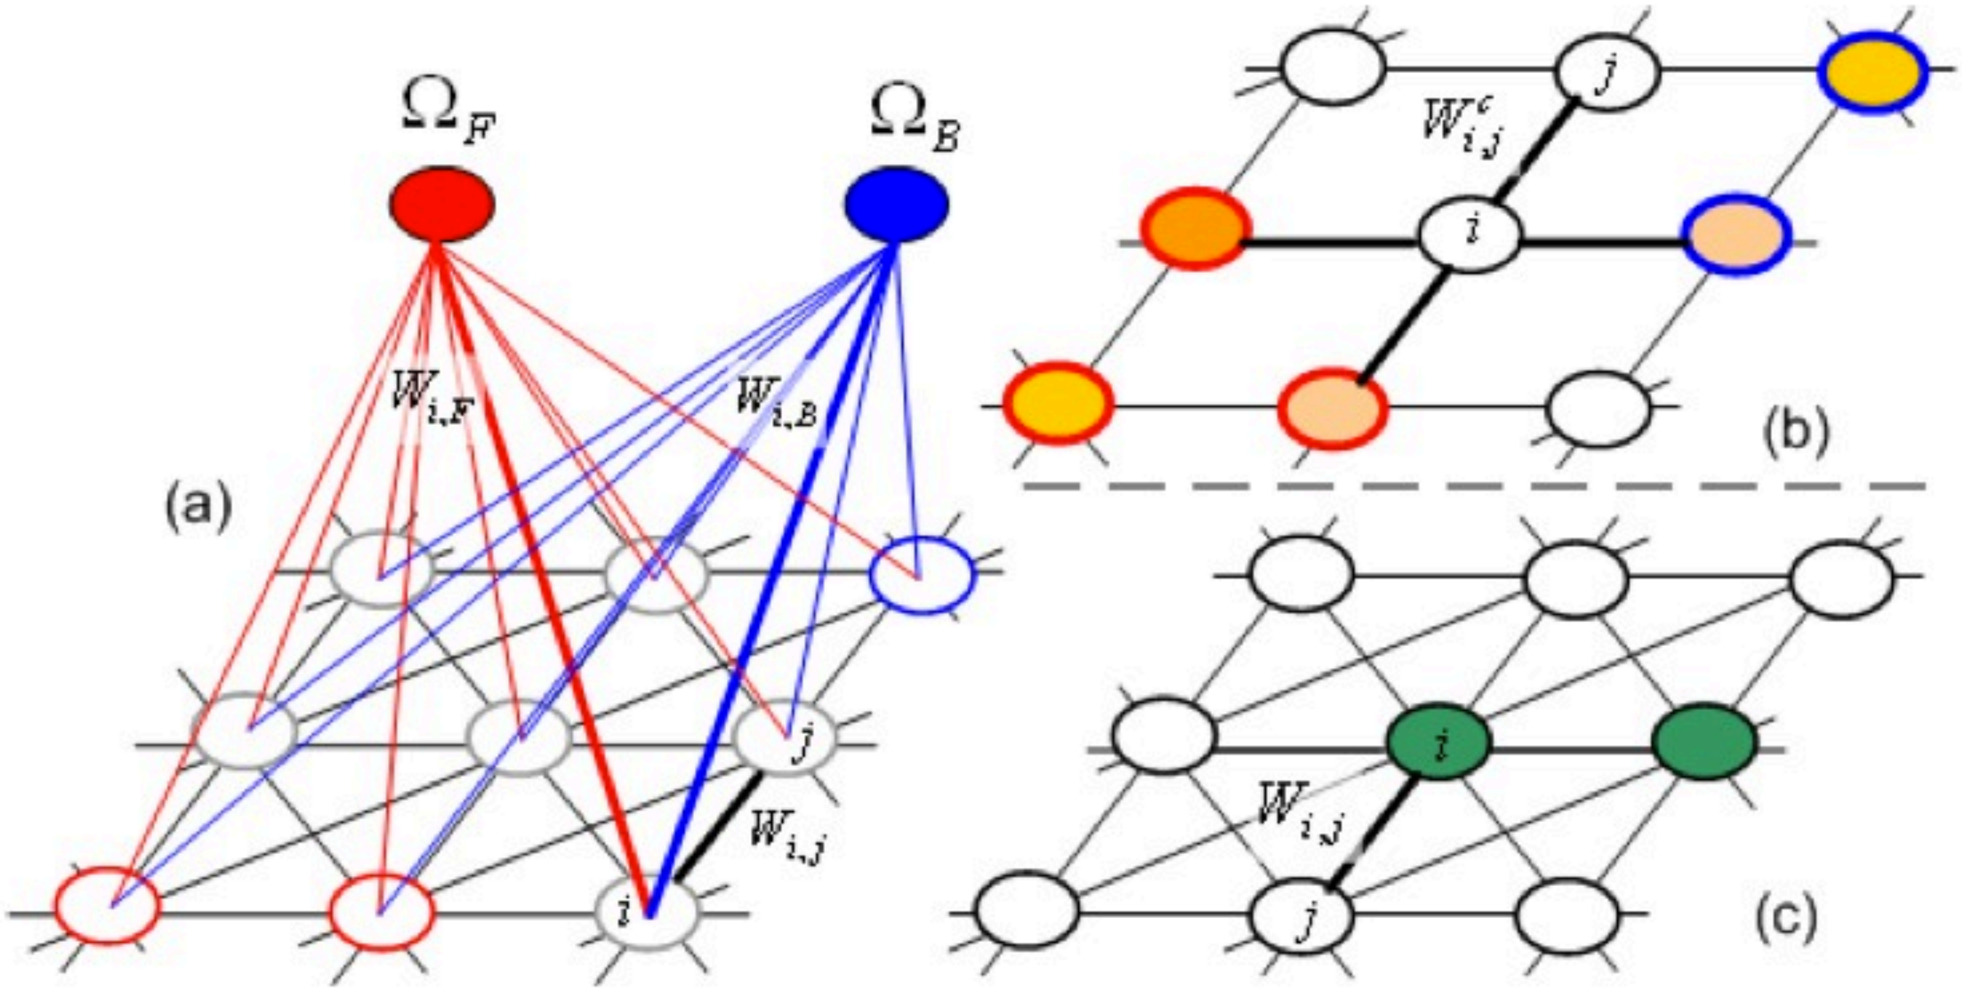

Figure 4: The matte (a), foreground colors (b) and the update region (c) are solved as soft graph-labeling problems.

Environment Matting

- Alpha matting fails for complex light transport
- Let's explicitly model refraction and reflection

Figure 4 From left to right: an alpha matte composite, an environment matte composite, and a photograph of an object in front of a background image. The top row shows a ribbed glass candle holder; the bottom row shows a rough-surfaced glass bookend.

Environment Matting

[Zongker et al. 1999]

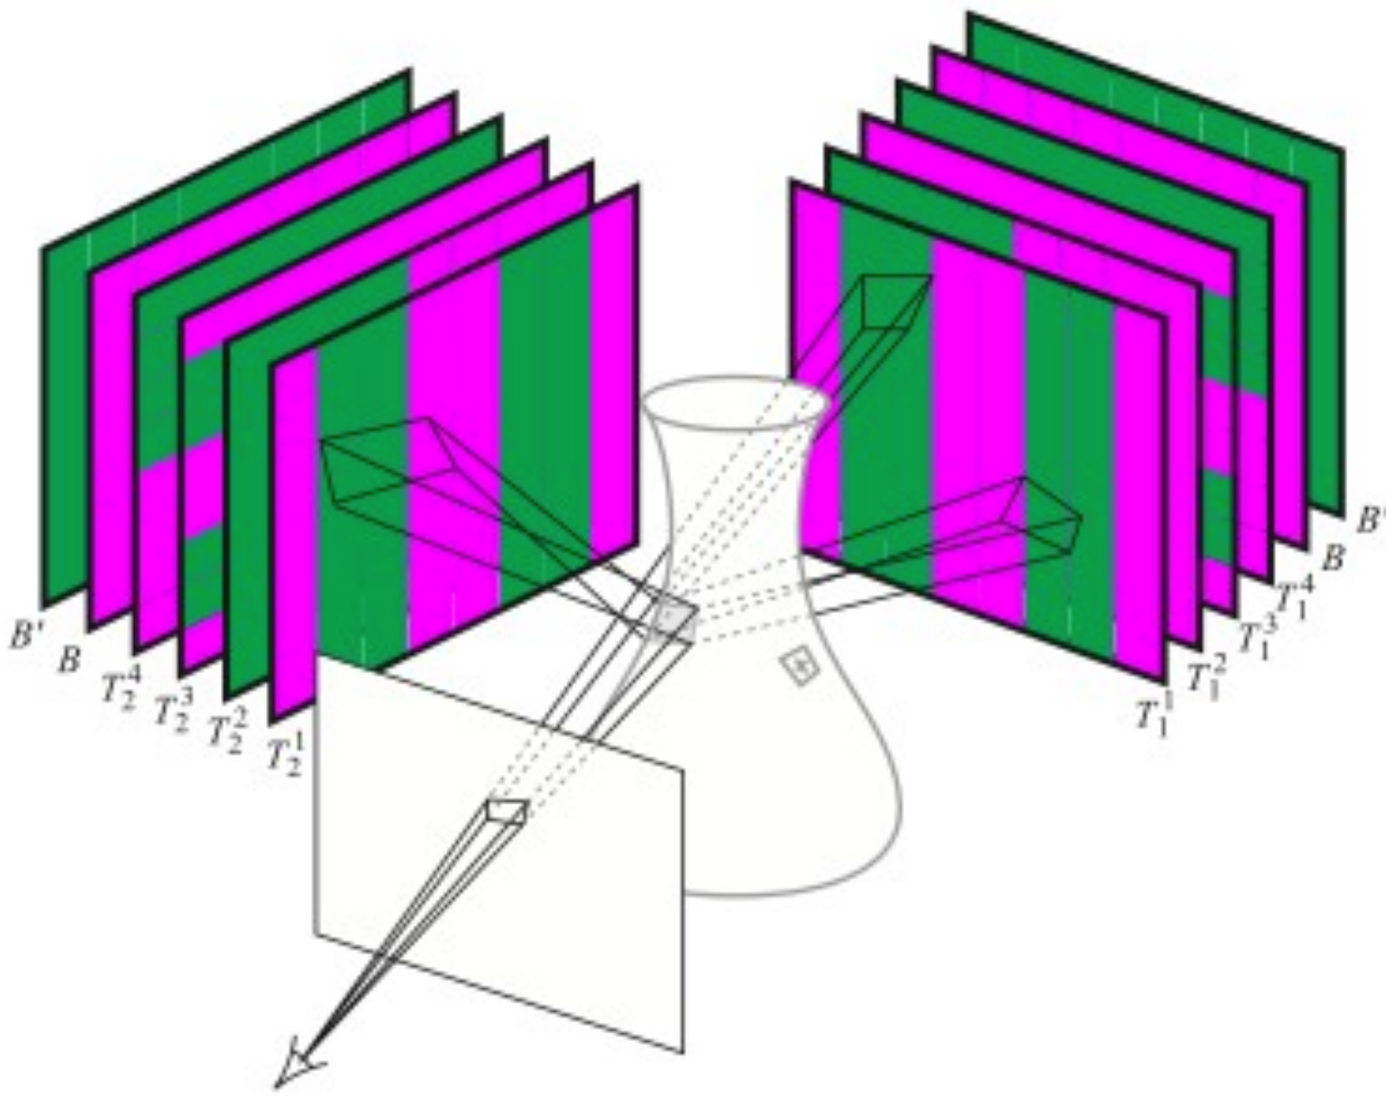

Figure 2 The environment matting process uses structured textures to capture how light is reflected and refracted from a backdrop (right shaft), as well as from various sidedrops (left shaft). The process also captures light coming from the backdrop that is seen through uncovered portions of a pixel (center shaft).

Attenuation-Refraction Mattes

[Yeung et al. 2011]

- Humans are quite bad at evaluating light transport
- So lets cheat on the refraction

Fig. 1. Attenuation-Refraction Matte (ARM). From left to right: input image (courtesy of Francesco Dazzi), casual markup, ARM, and the composite completed with simulated Fresnel effect and caustic shadow.

Attenuation-Refraction Mattes

$$C_M(\mathbf{x}) = \alpha(\mathbf{x})S(\mathbf{x}) + (1 - \alpha(\mathbf{x}))\beta(\mathbf{x})B(G(\mathbf{x})).$$

Specular Alpha

Specular Color

Transmission Color

Background Color

Refraction Warp

Attenuation-Refraction Mattes

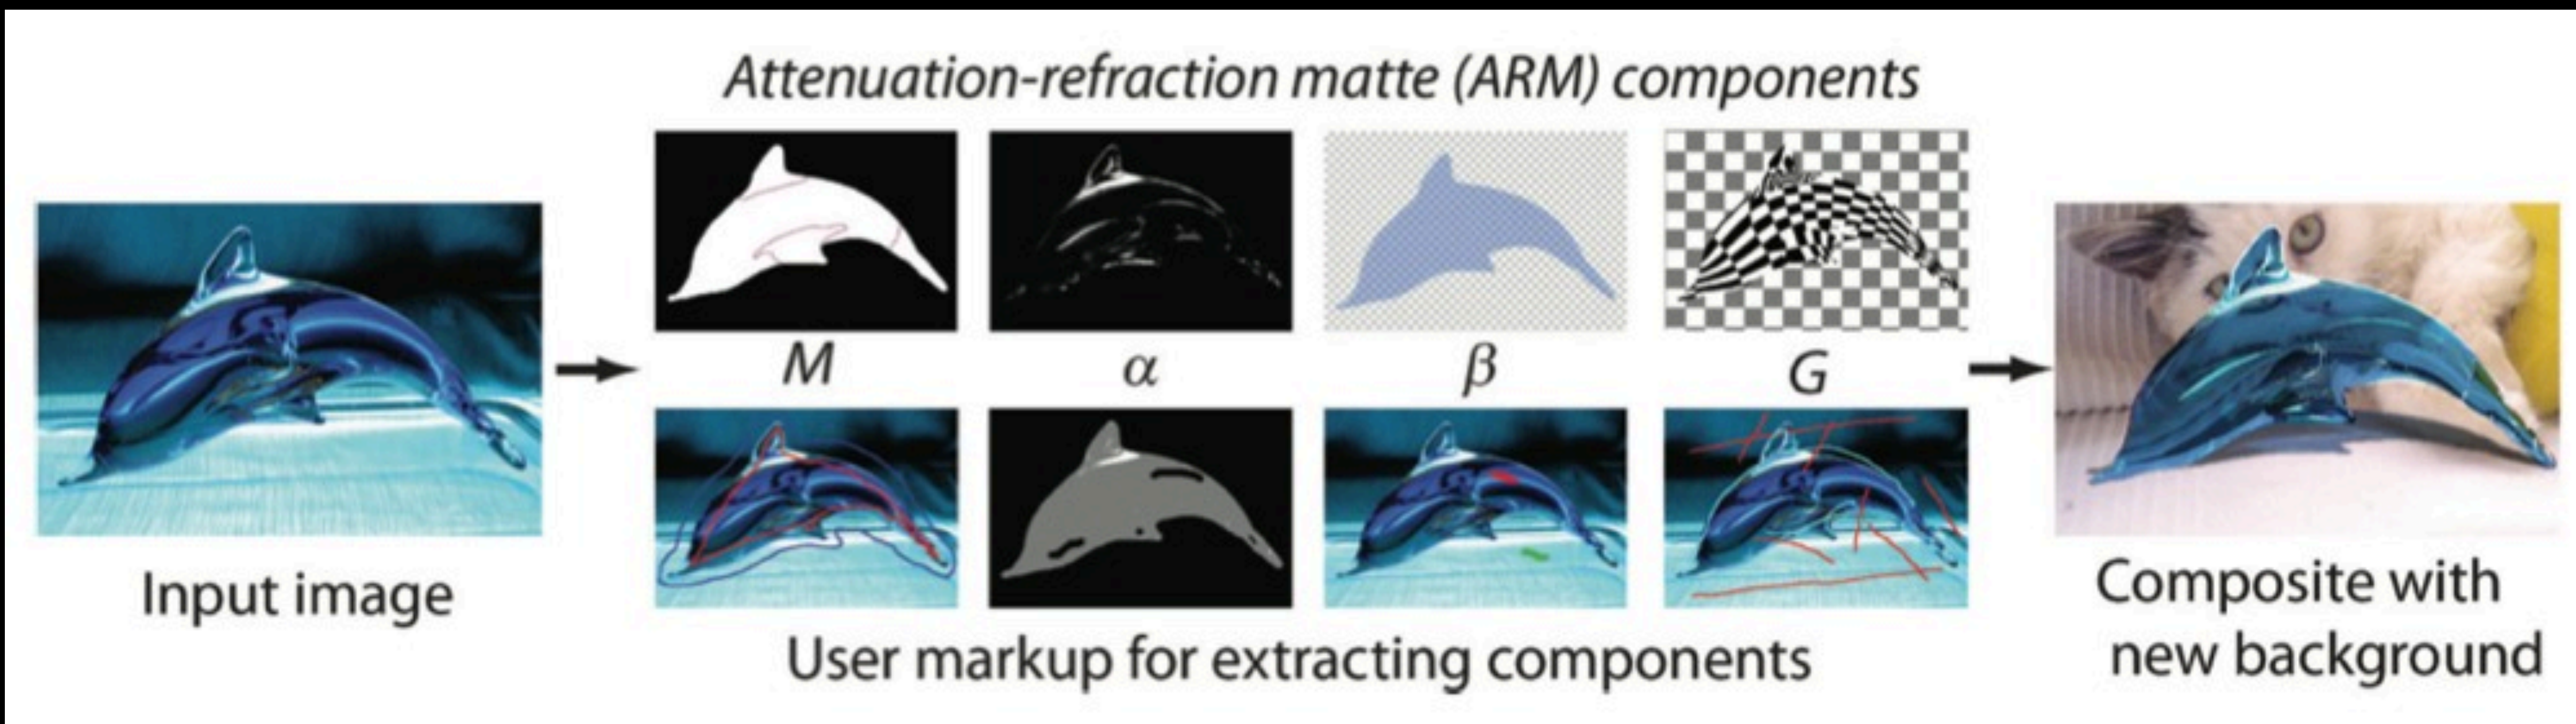

Fig. 3. Overview of Attenuation-Refraction Matte (ARM) extraction.

Attenuation-Refraction Mattes

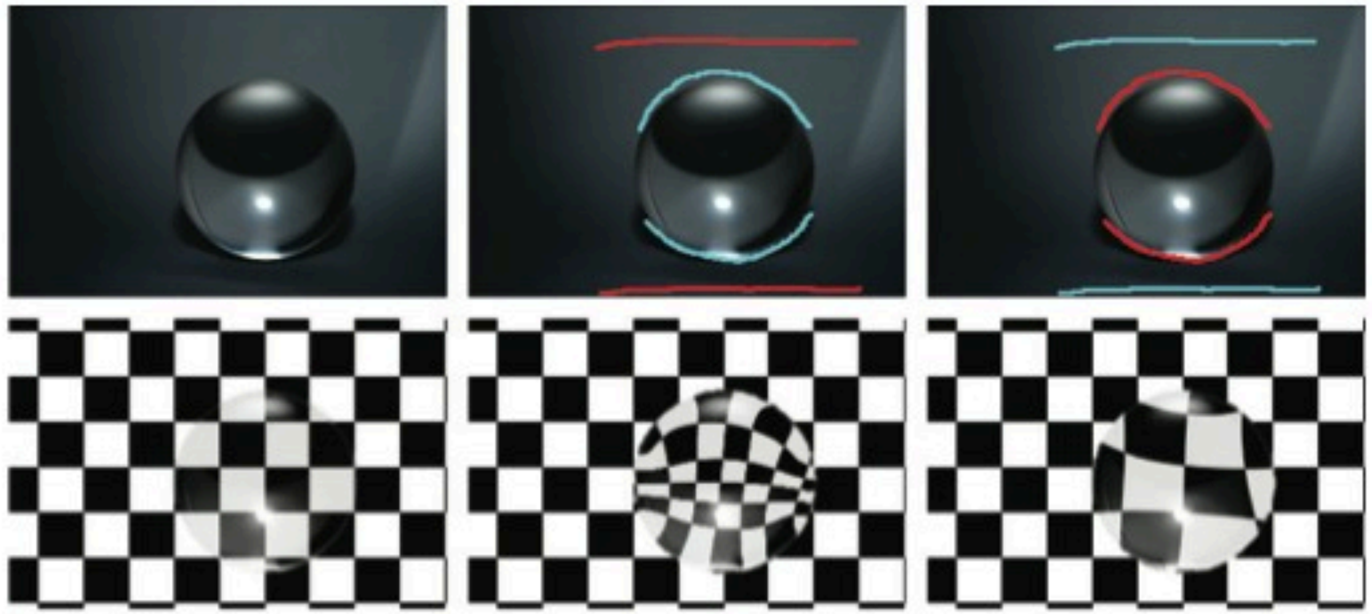

(a) (b) (c)

Fig. 4. Object cue. Three basic markups of refractive light-transport: (a) no markup, (b) markup for simulating light convergence, (c) markup for simulating light divergence. c_{ref} is in red and c_{target} is in cyan. c_{target} is where c_{ref} is perceived to distort.

(a) (b)

Fig. 5. Background cue. (a) Markup drawn using the background as a cue. As with object cue, c_{target} (cyan) is where c_{ref} (red) is perceived to distort. (b) Example of the resulting deformation.

Attenuation-Refraction Mattes

Can you spot the real glass?

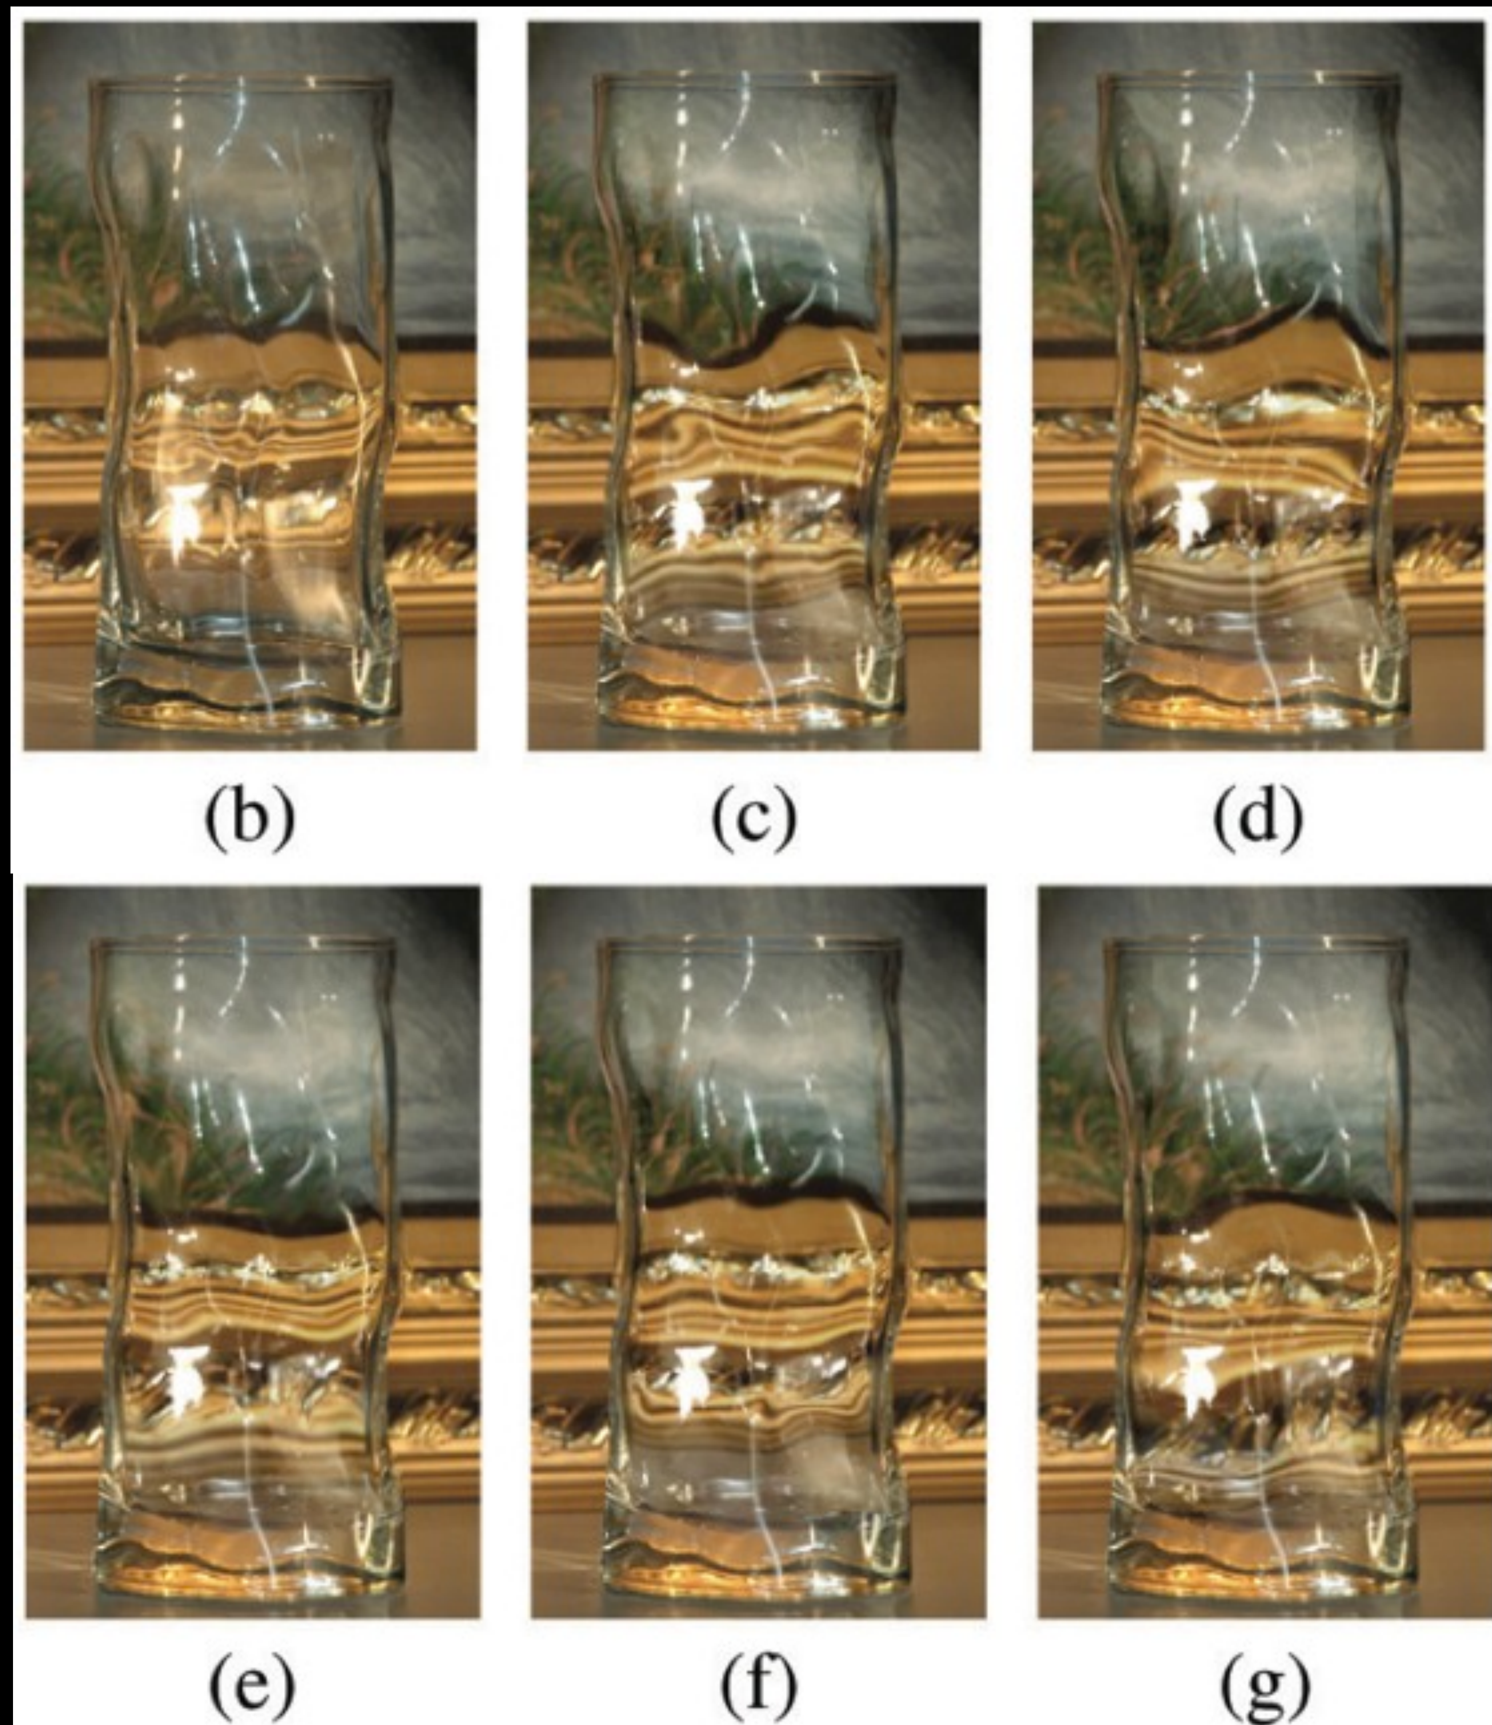

Nearest Neighbor Search

- Given an image patch, find a similar one elsewhere
- Useful for many tasks:
  - Texture synthesis
  - Hole filling
  - Retargeting
  - Denoising
  - Optical flow

PatchMatch

[Barnes et al. 2009]

PatchMatch: A Randomized Correspondence Algorithm for Structural Image Editing

Connelly Barnes[1] Eli Shechtman[2,3] Adam Finkelstein[1] Dan B Goldman[2]

[1]Princeton University [2]Adobe Systems [3]University of Washington

(a) original (b) hole+constraints (c) hole filled (d) constraints (e) constrained retarget (f) reshuffle

Figure 1: *Structural image editing. Left to right: (a) the original image; (b) a hole is marked (magenta) and we use line constraints (red/green/blue) to improve the continuity of the roofline; (c) the hole is filled in; (d) user-supplied line constraints for retargeting; (e) retargeting using constraints eliminates two columns automatically; and (f) user translates the roof upward using reshuffling.*

PatchMatch

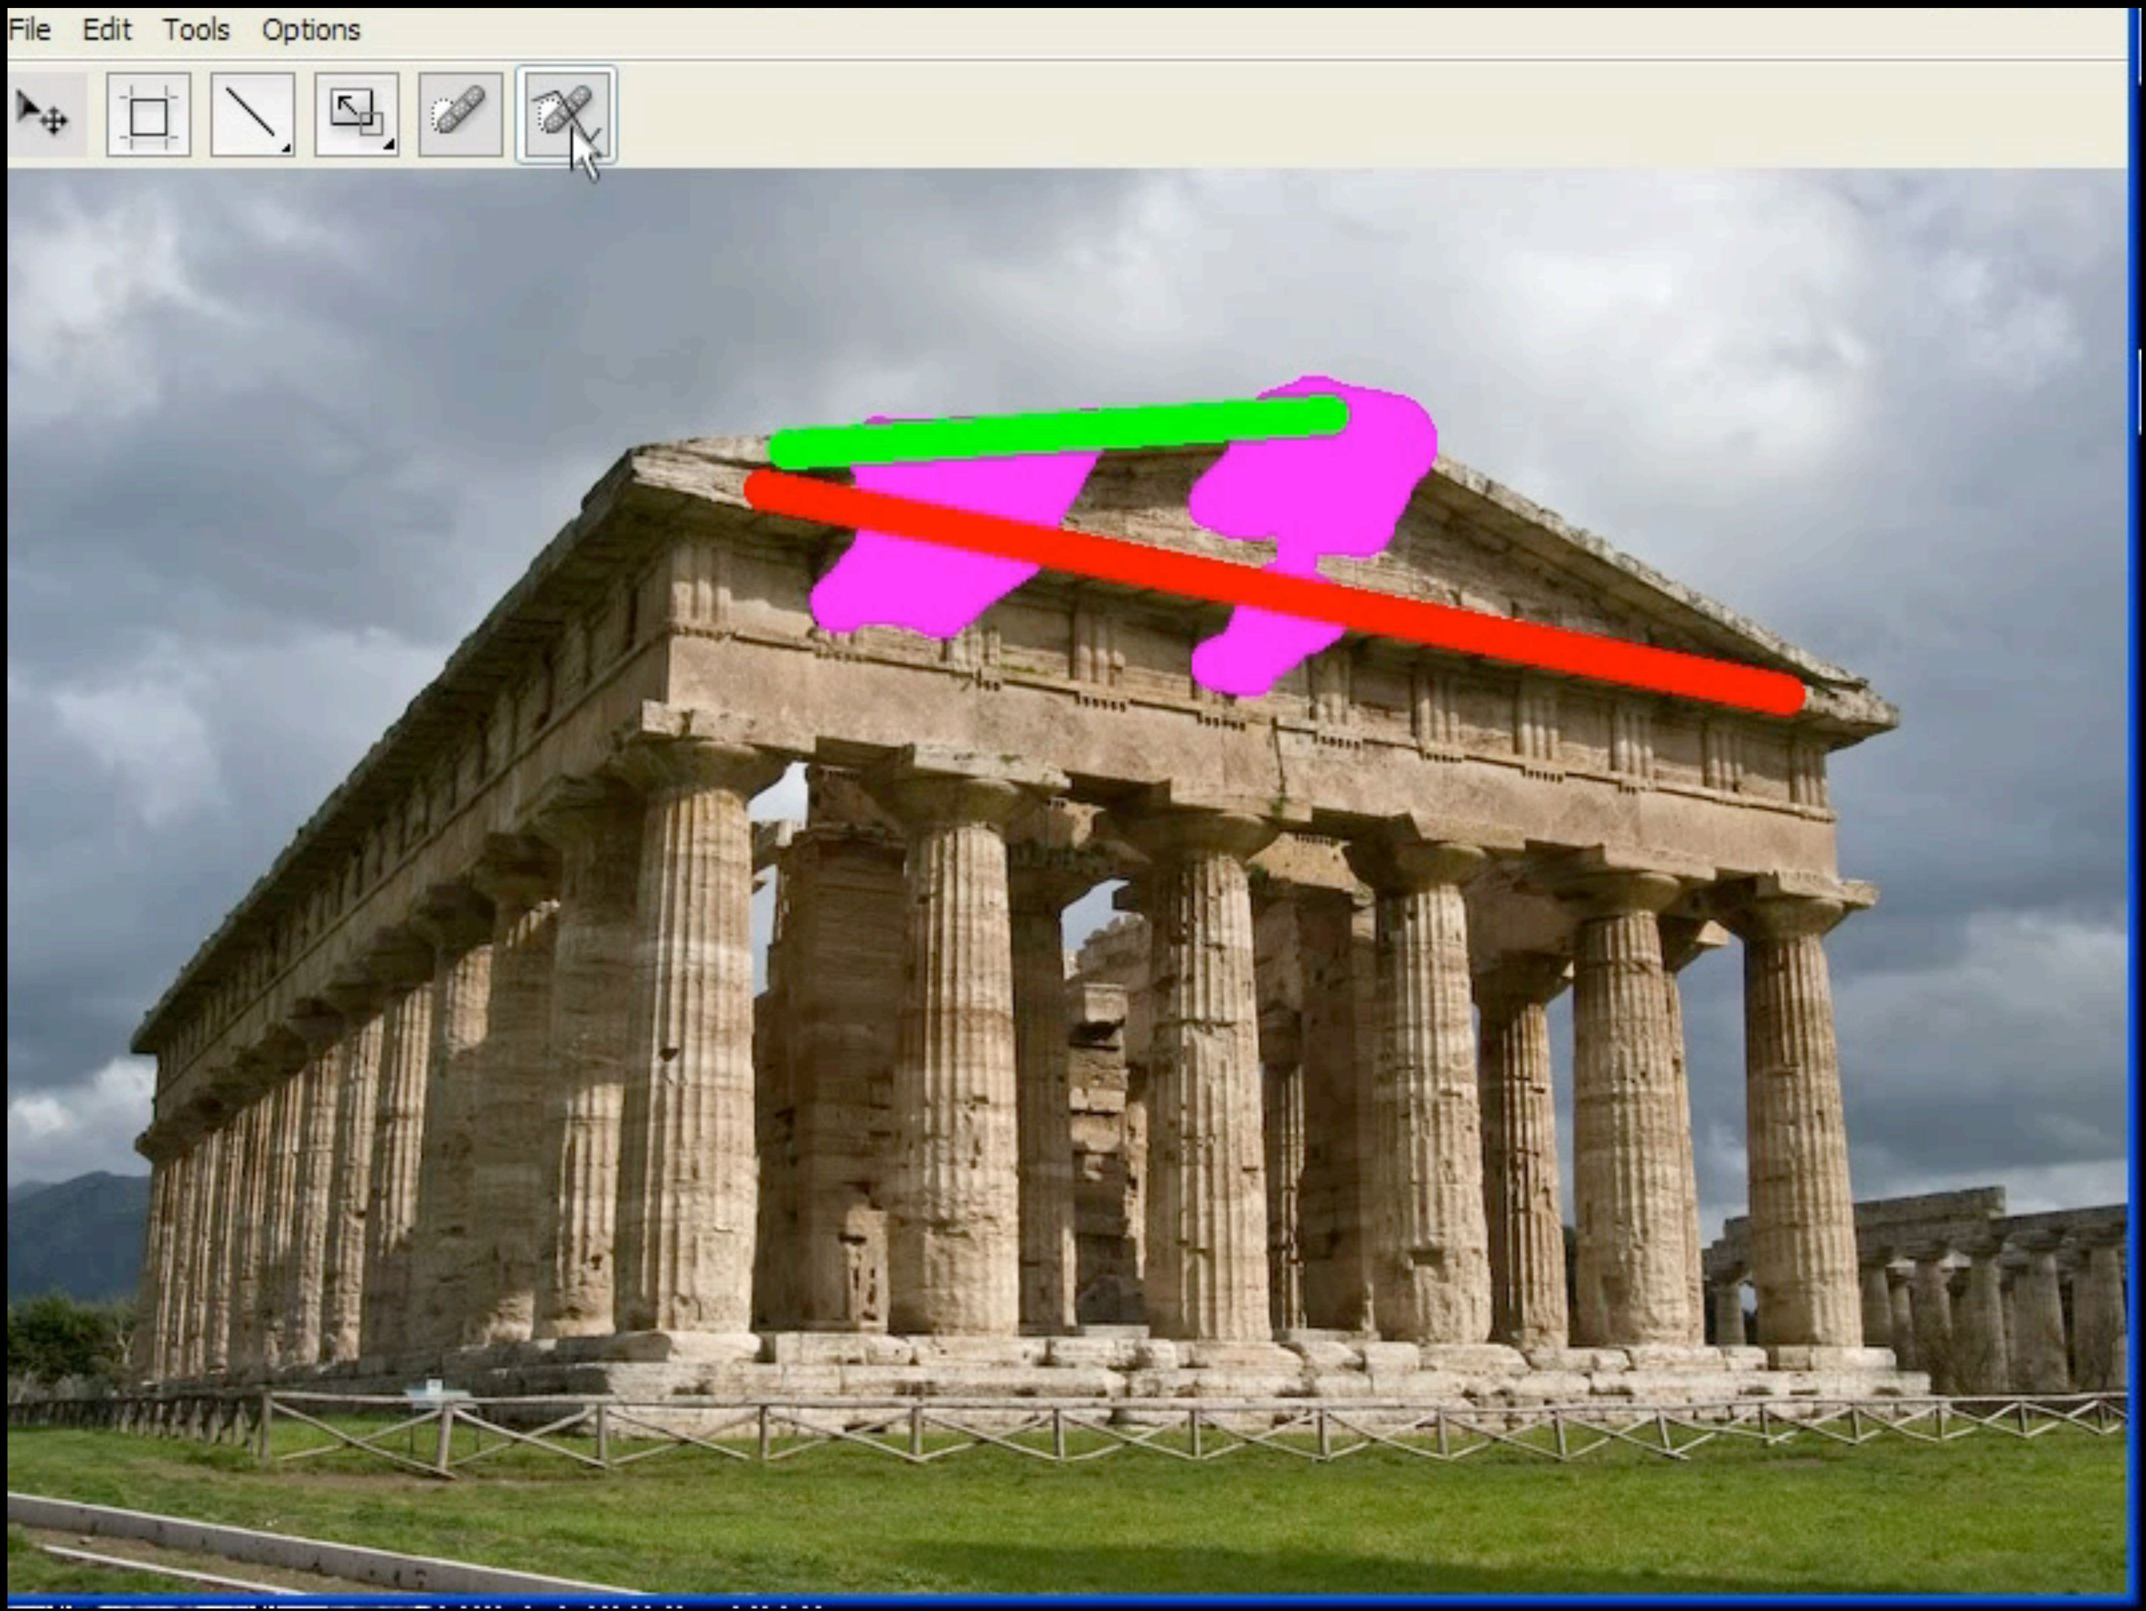

PatchMatch

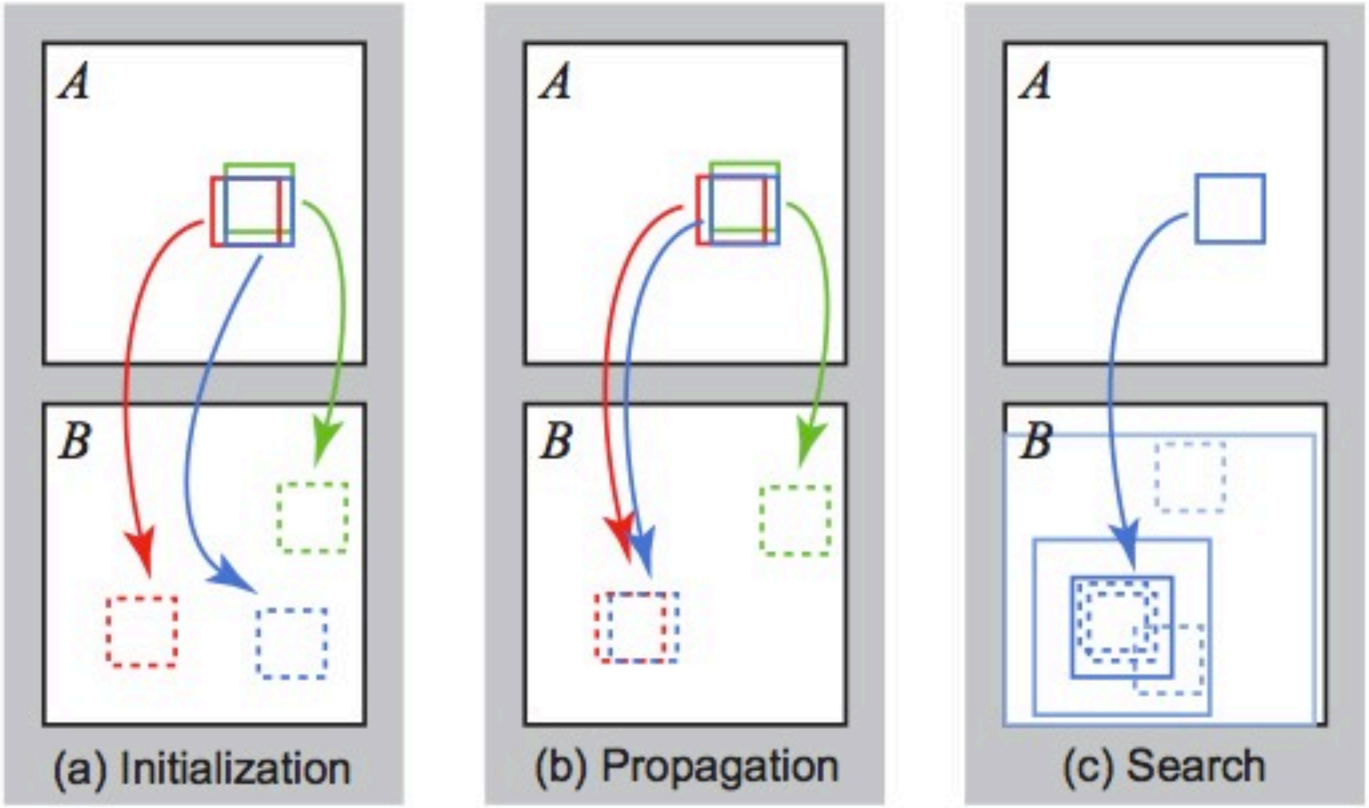

Figure 2: *Phases of the randomized nearest neighbor algorithm: (a) patches initially have random assignments; (b) the blue patch checks above/green and left/red neighbors to see if they will improve the blue mapping, propagating good matches; (c) the patch searches randomly for improvements in concentric neighborhoods.*

PatchMatch

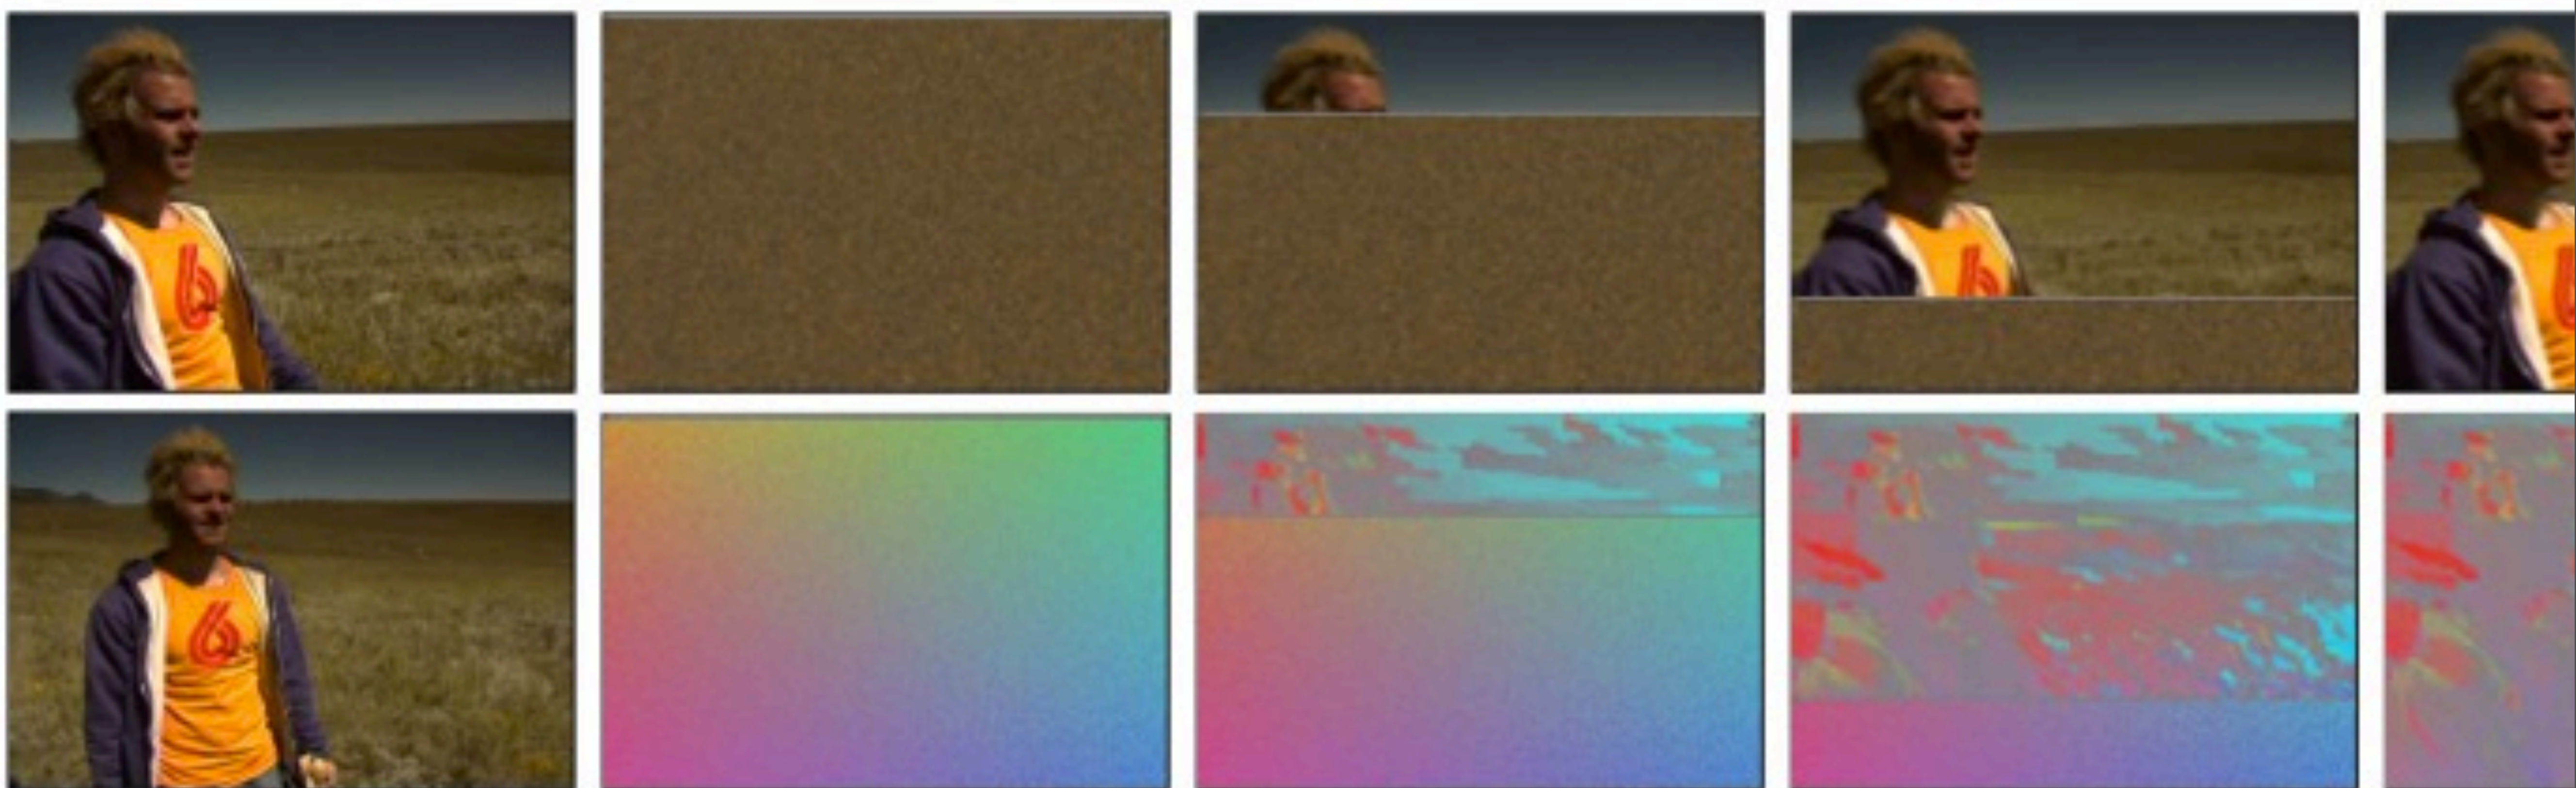

(a) originals (b) random (c) $\frac{1}{4}$ iteration (d) $\frac{3}{4}$ iteration (e)

Figure 3: *Illustration of convergence. (a) The top image is reconstructed using o reconstruction by the patch "voting" described in Section 4, below: a random initial and angle visualized as hue. (c)* 1/4 *of the way through the first iteration, high-qual the current scan line (denoted with the horizontal bar). (d)* 3/4 *of the way through th iterations. (g) After 5 iterations, almost all patches have stopped changing. The tiny later iterations.*

PatchMatch

- Hole Filling

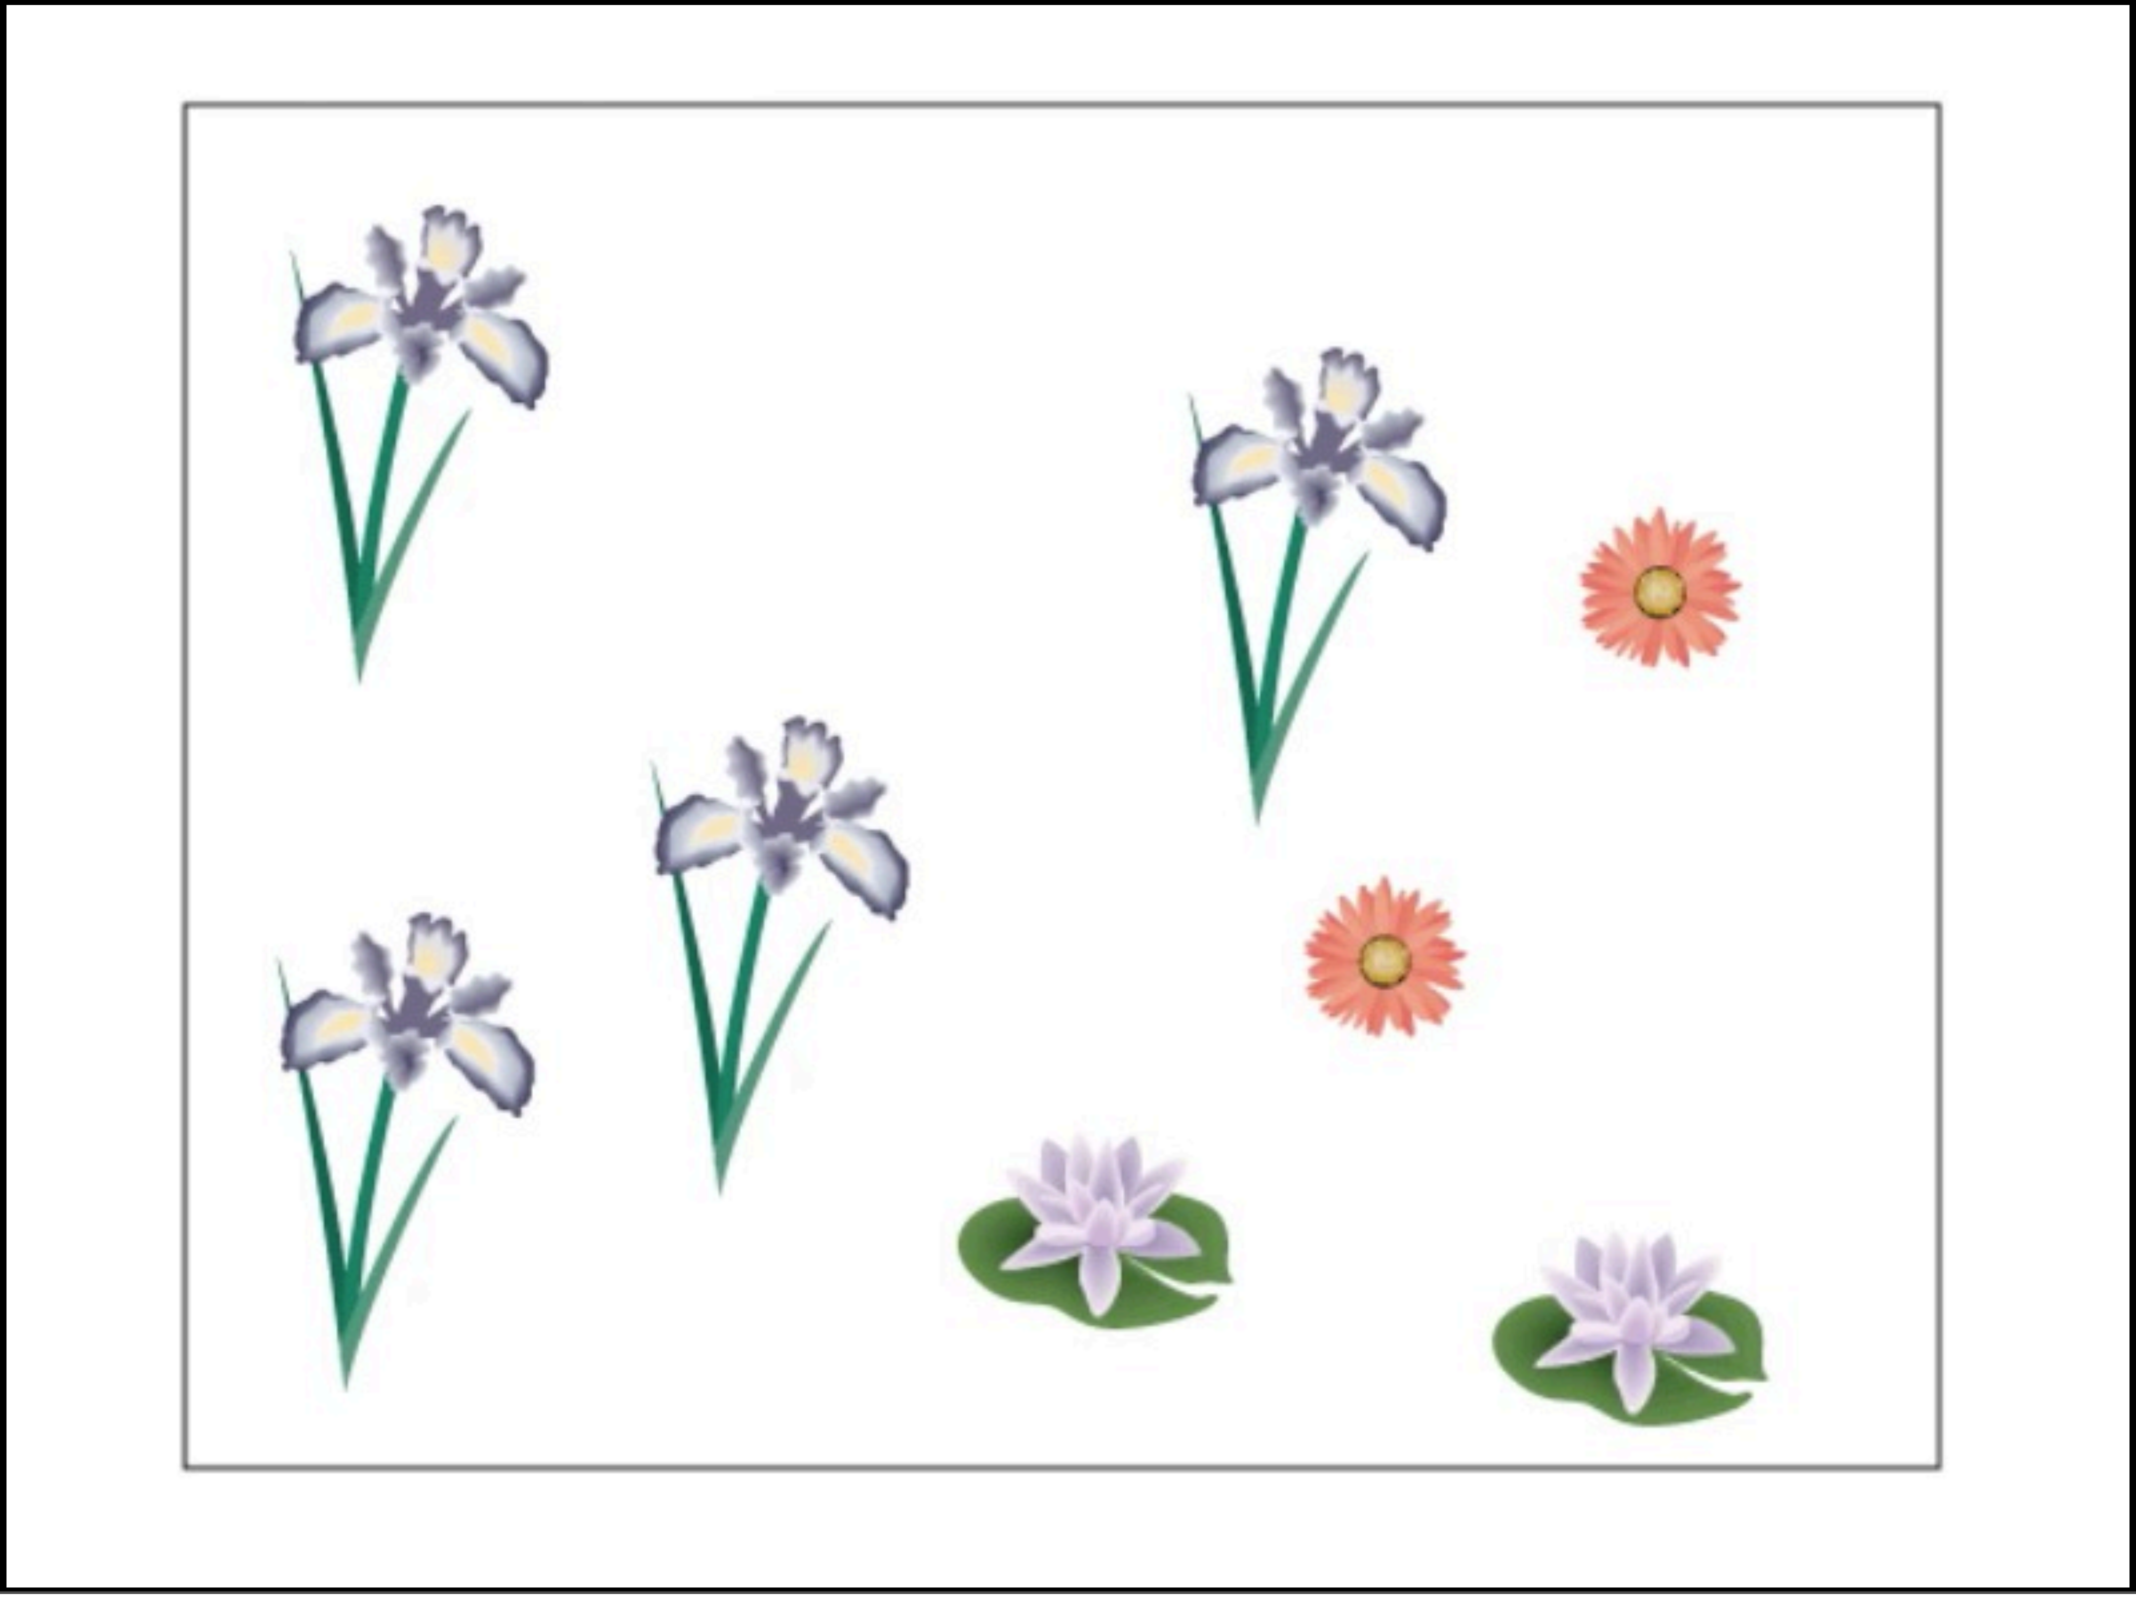

PatchMatch

- Hole Filling

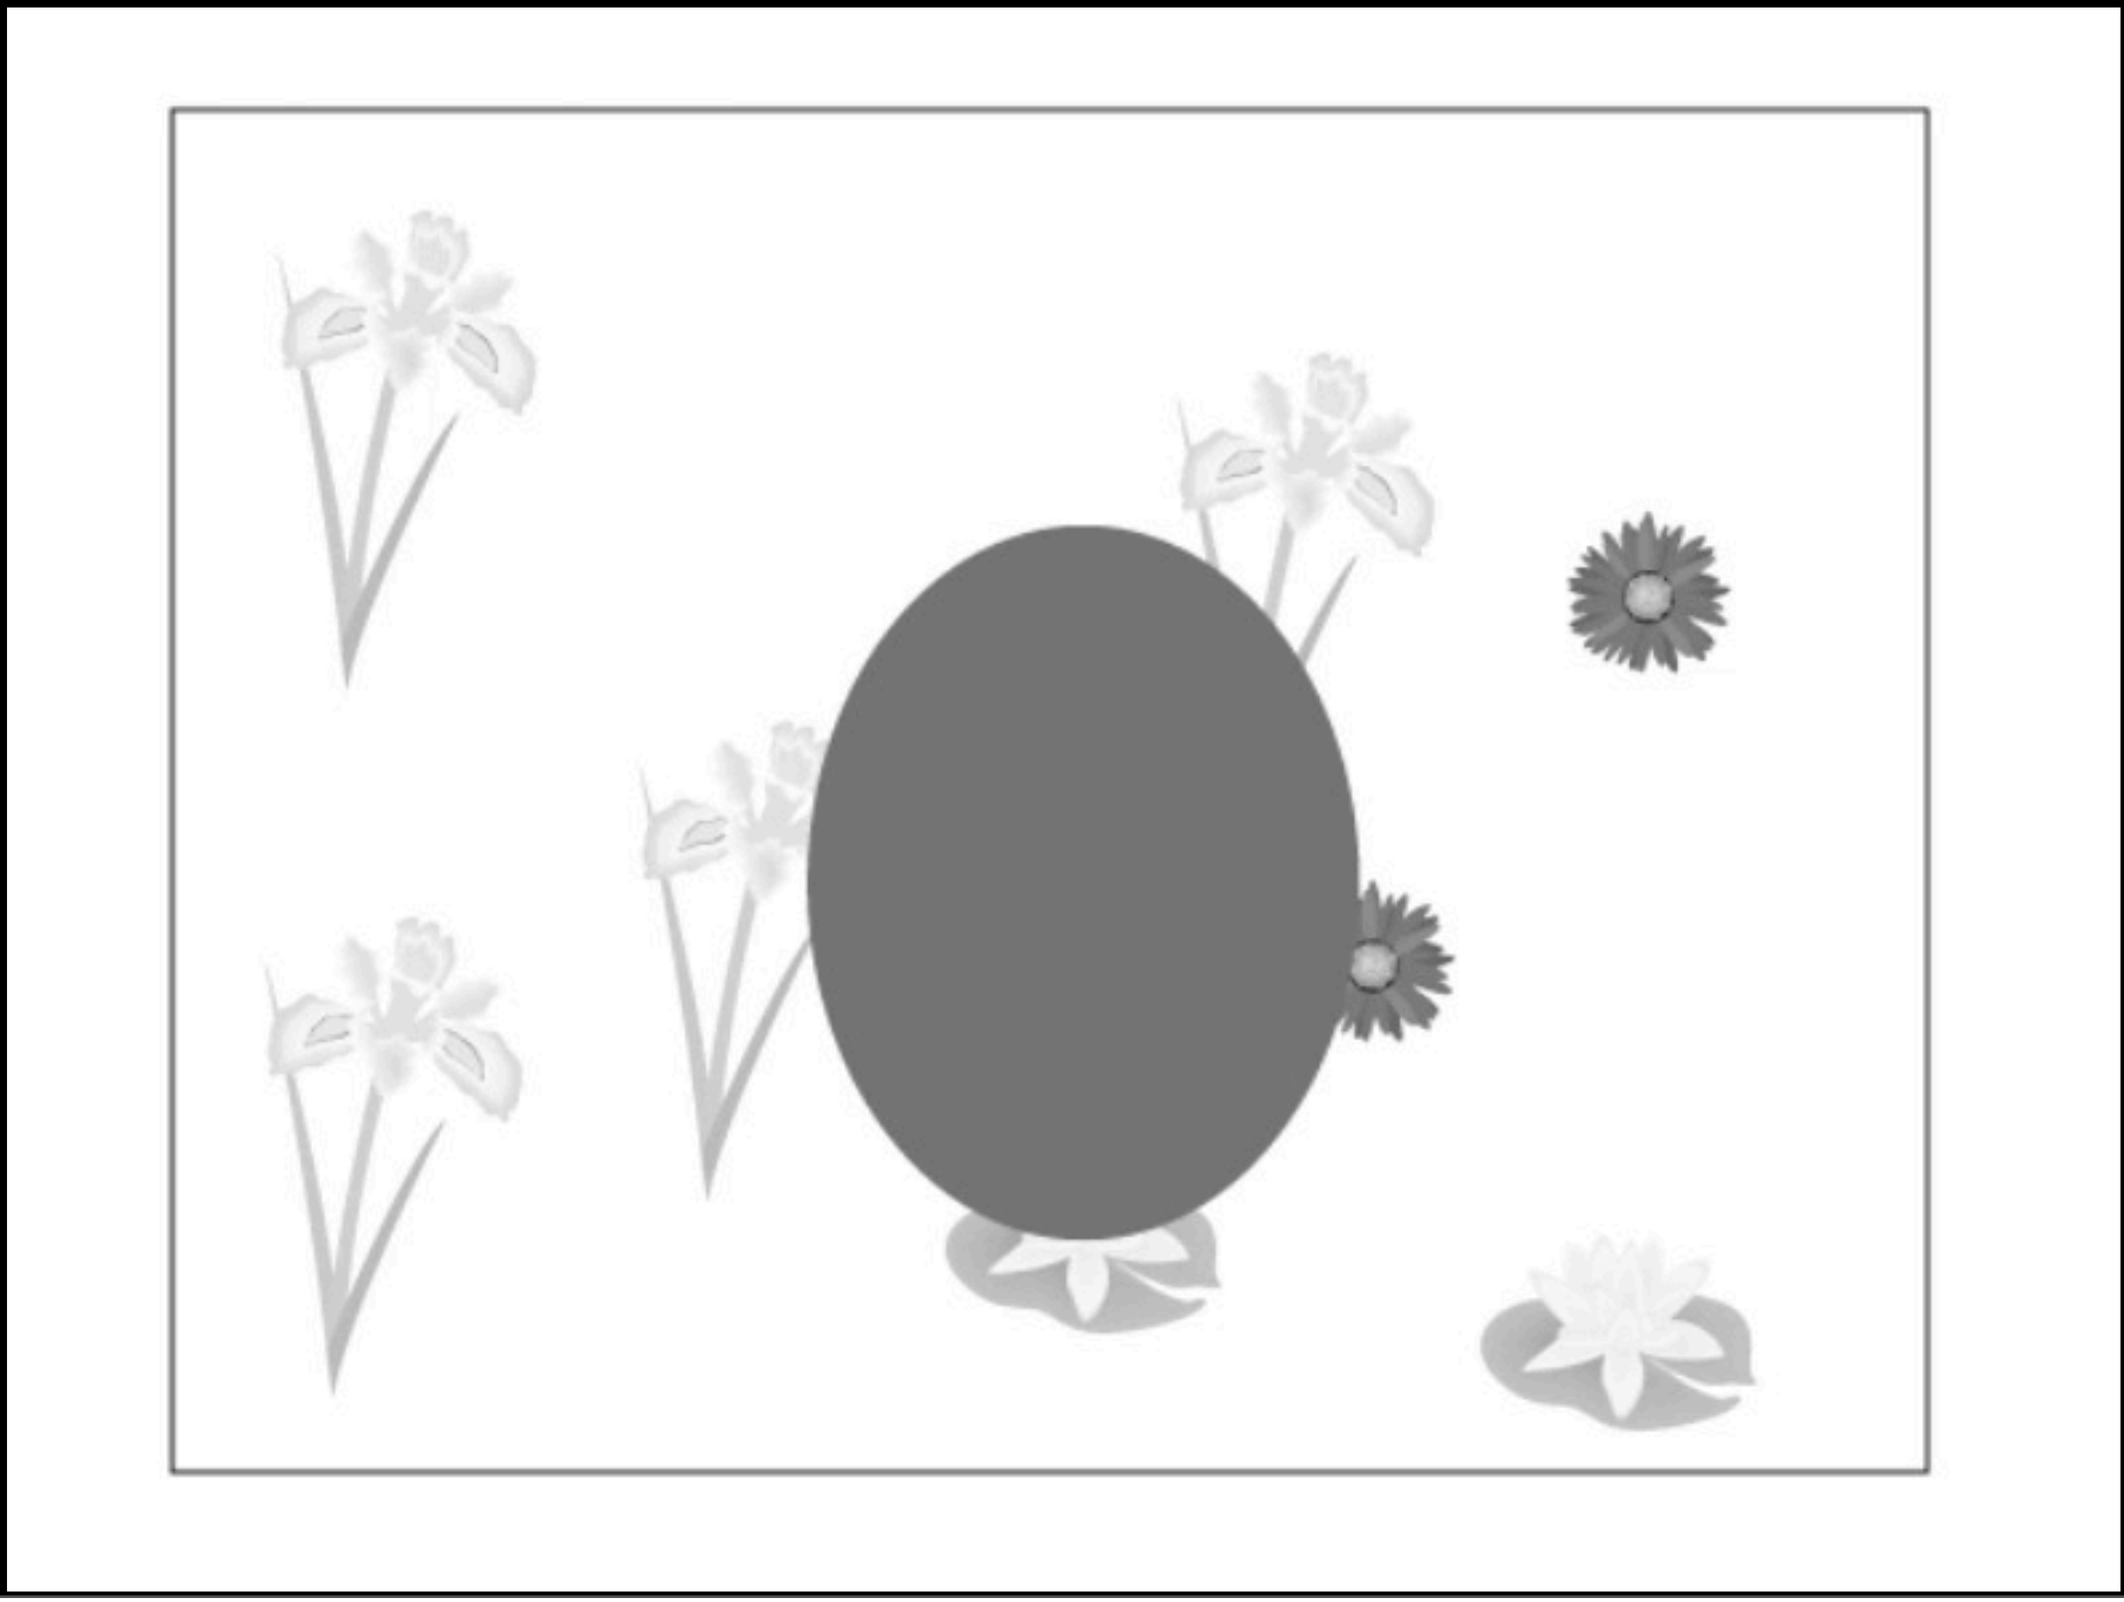

PatchMatch

- Hole Filling

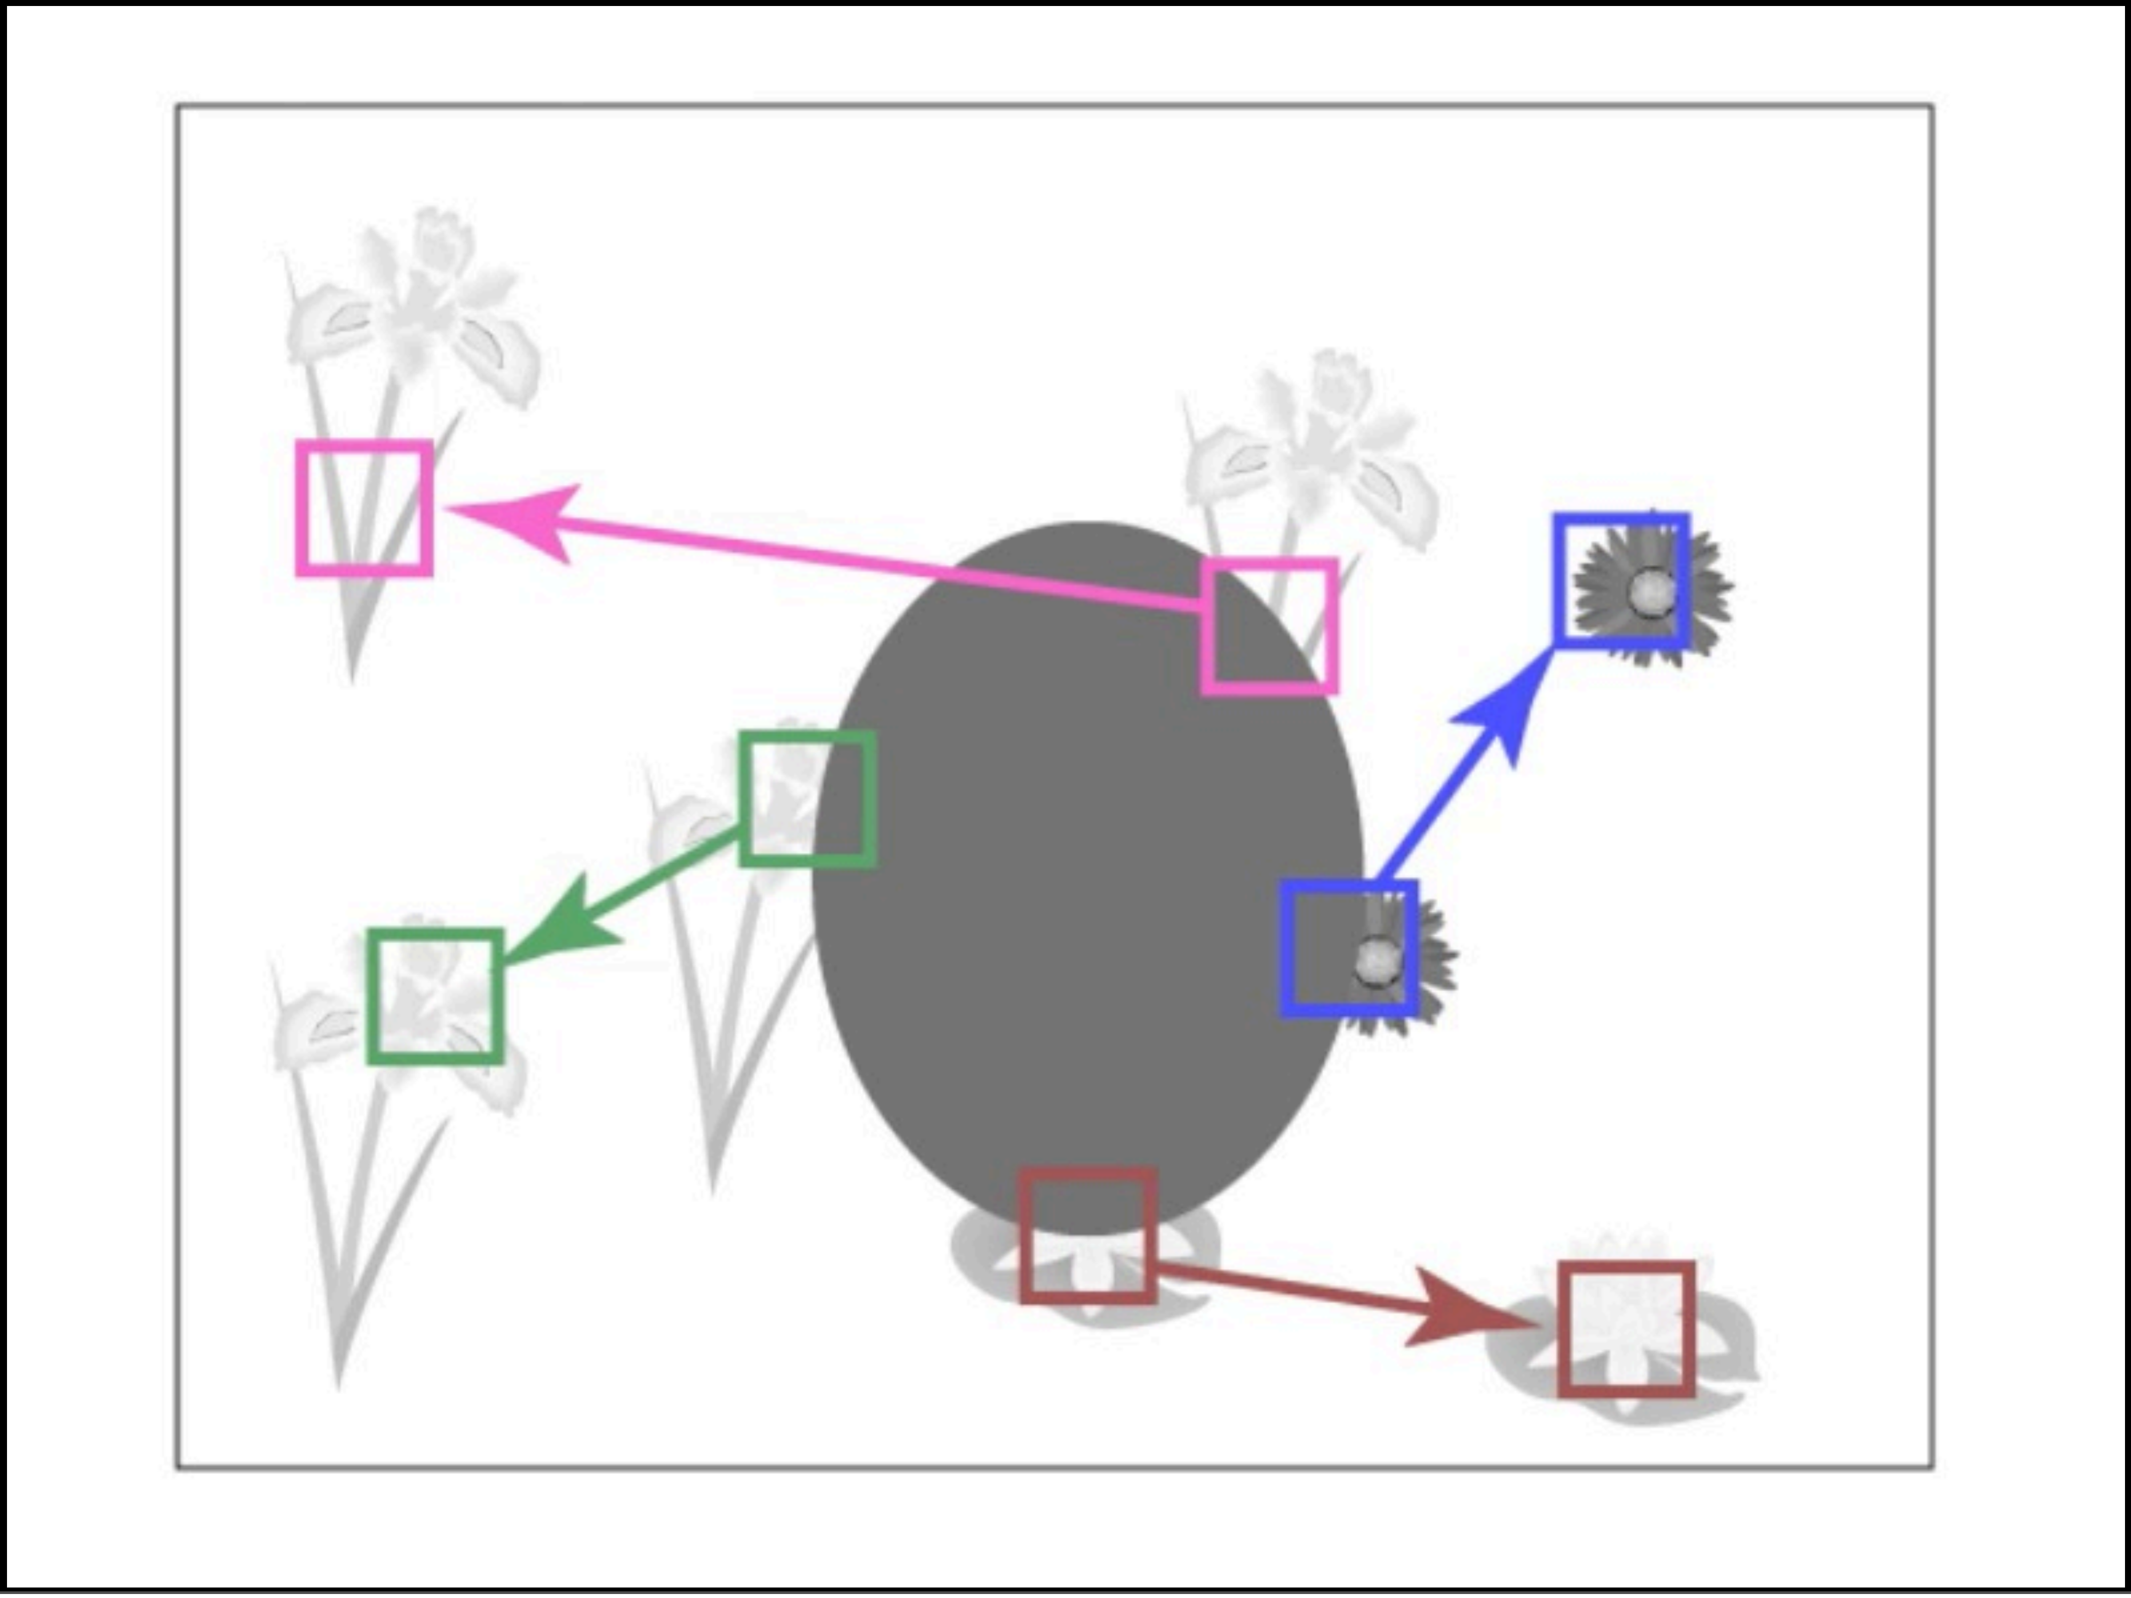

PatchMatch

- To Photoshop!

PatchMatch

- To Photoshop!

Image Collections

- What happens when there just isn't enough data?
- Instead, let's steal data from other photos
- Likelihood of a good match increases with dataset size

Scene Completion

[Hays and Efros 2007]

Scene Completion Using Millions of Photographs

James Hays Alexei A. Efros
Carnegie Mellon University

Original Image Input Scene Matches Output

Figure 1: Given an input image with a missing region, we use matching scenes from a large collection of photographs to complete the image.

Scene Completion

- Not as straightforward as it might seem at first.

  "Existing image completion methods might produce sterile images but they don't risk putting an elephant in someone's back yard or a submarine in a parking lot."

- Must limit possible sources by semantic similarity to target

Scene Completion

- GIST [Oliva and Torralba 2006] encodes scene semantics
- Histograms of oriented edge filter responses in coarse spatial bins at multiple scales
- Only works for semantic matching with HUGE datasets

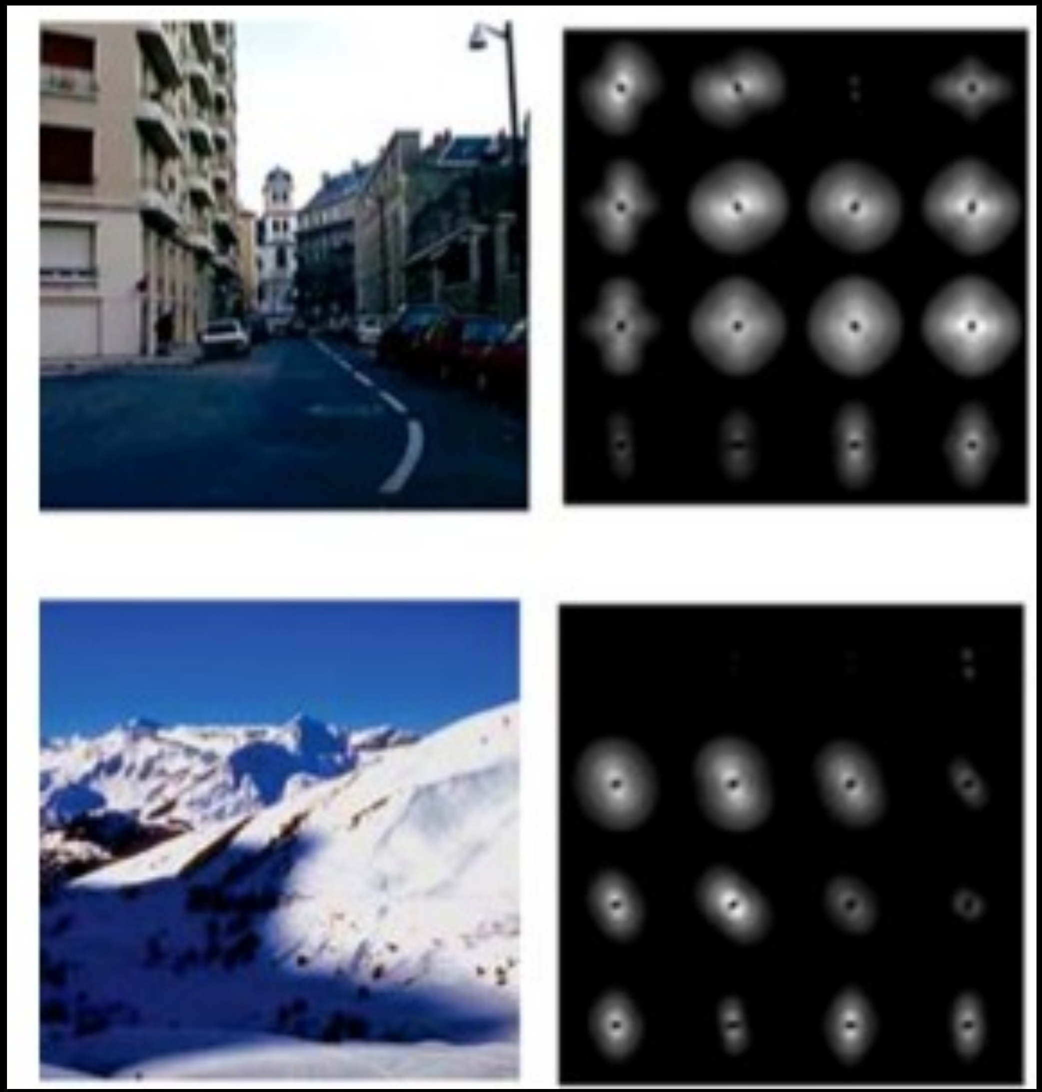

Scene Completion

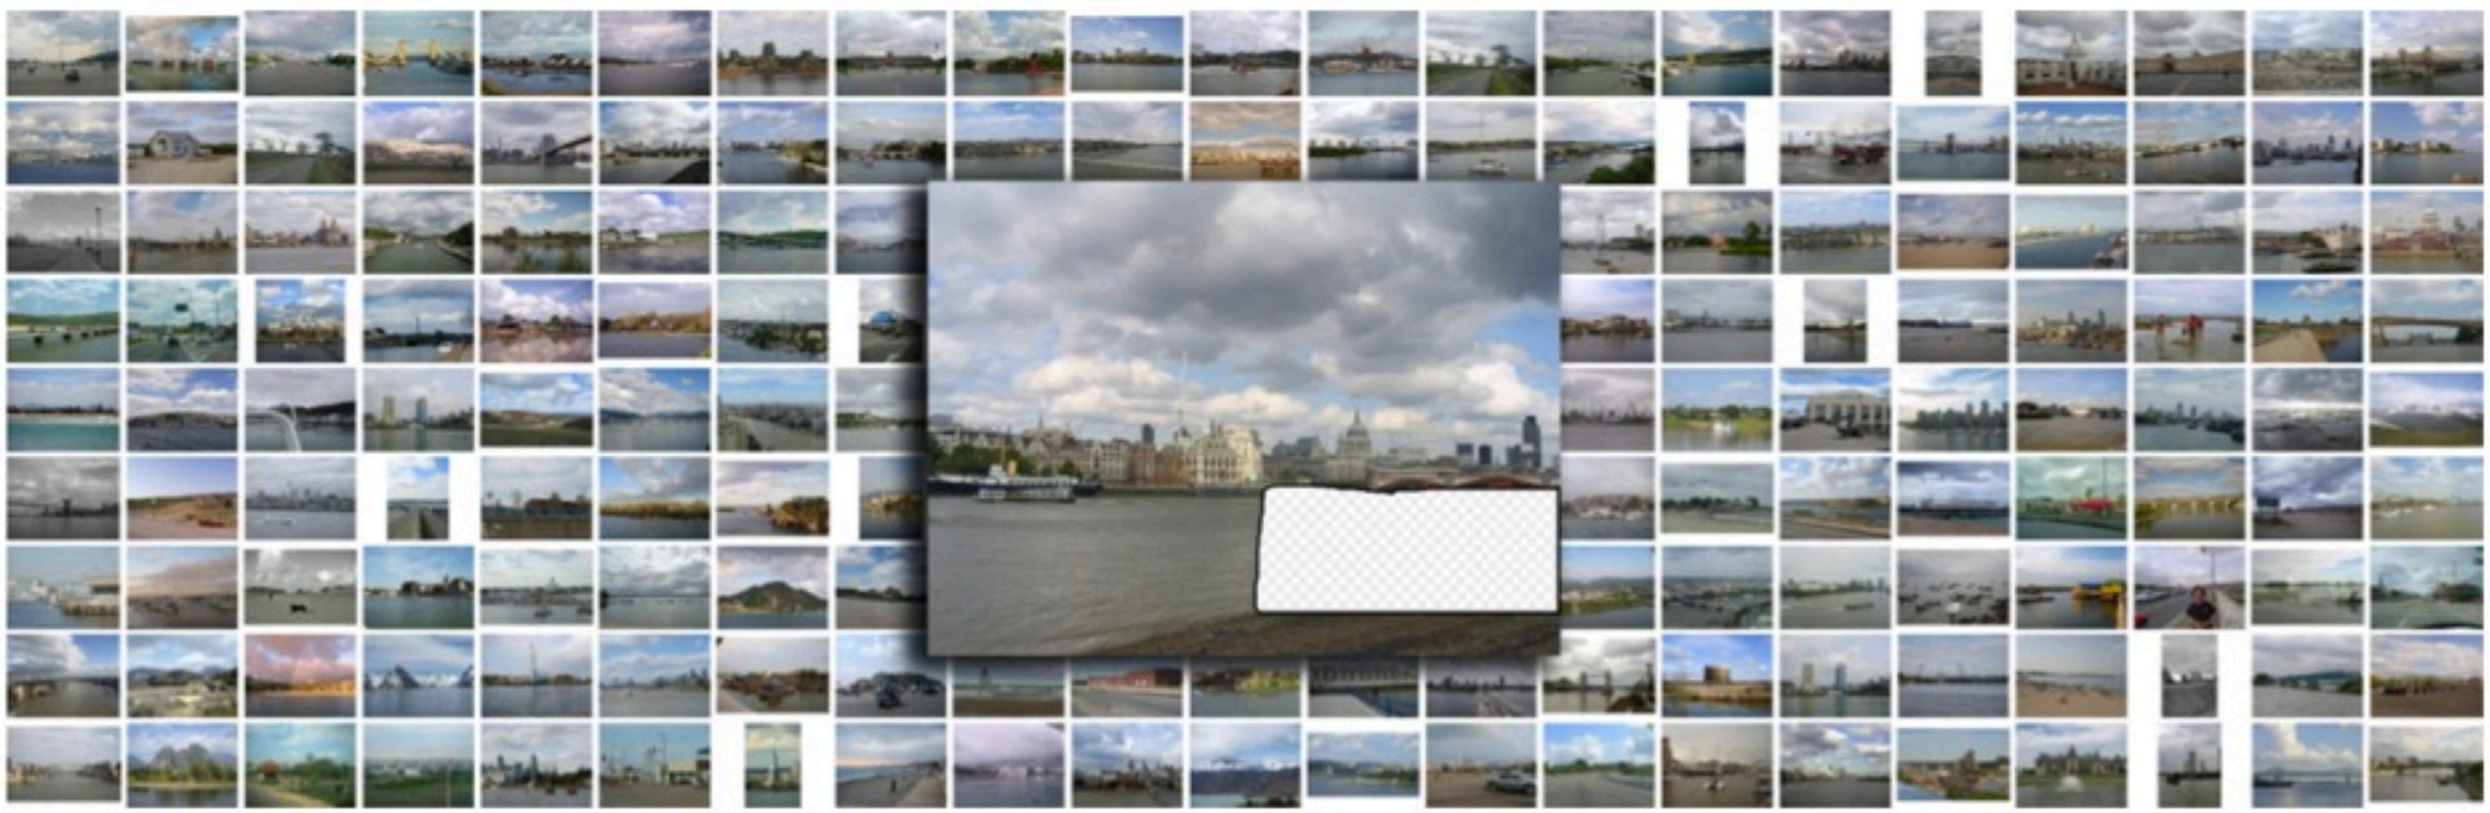

Figure 3: The 164 closest scenes to the incomplete image in the center. Most of the scenes are semantically and structurally similar; many are even from the same city (London).

- Show top N choices to user
- Composite using Graphcut and Poisson blending

Scene Completion

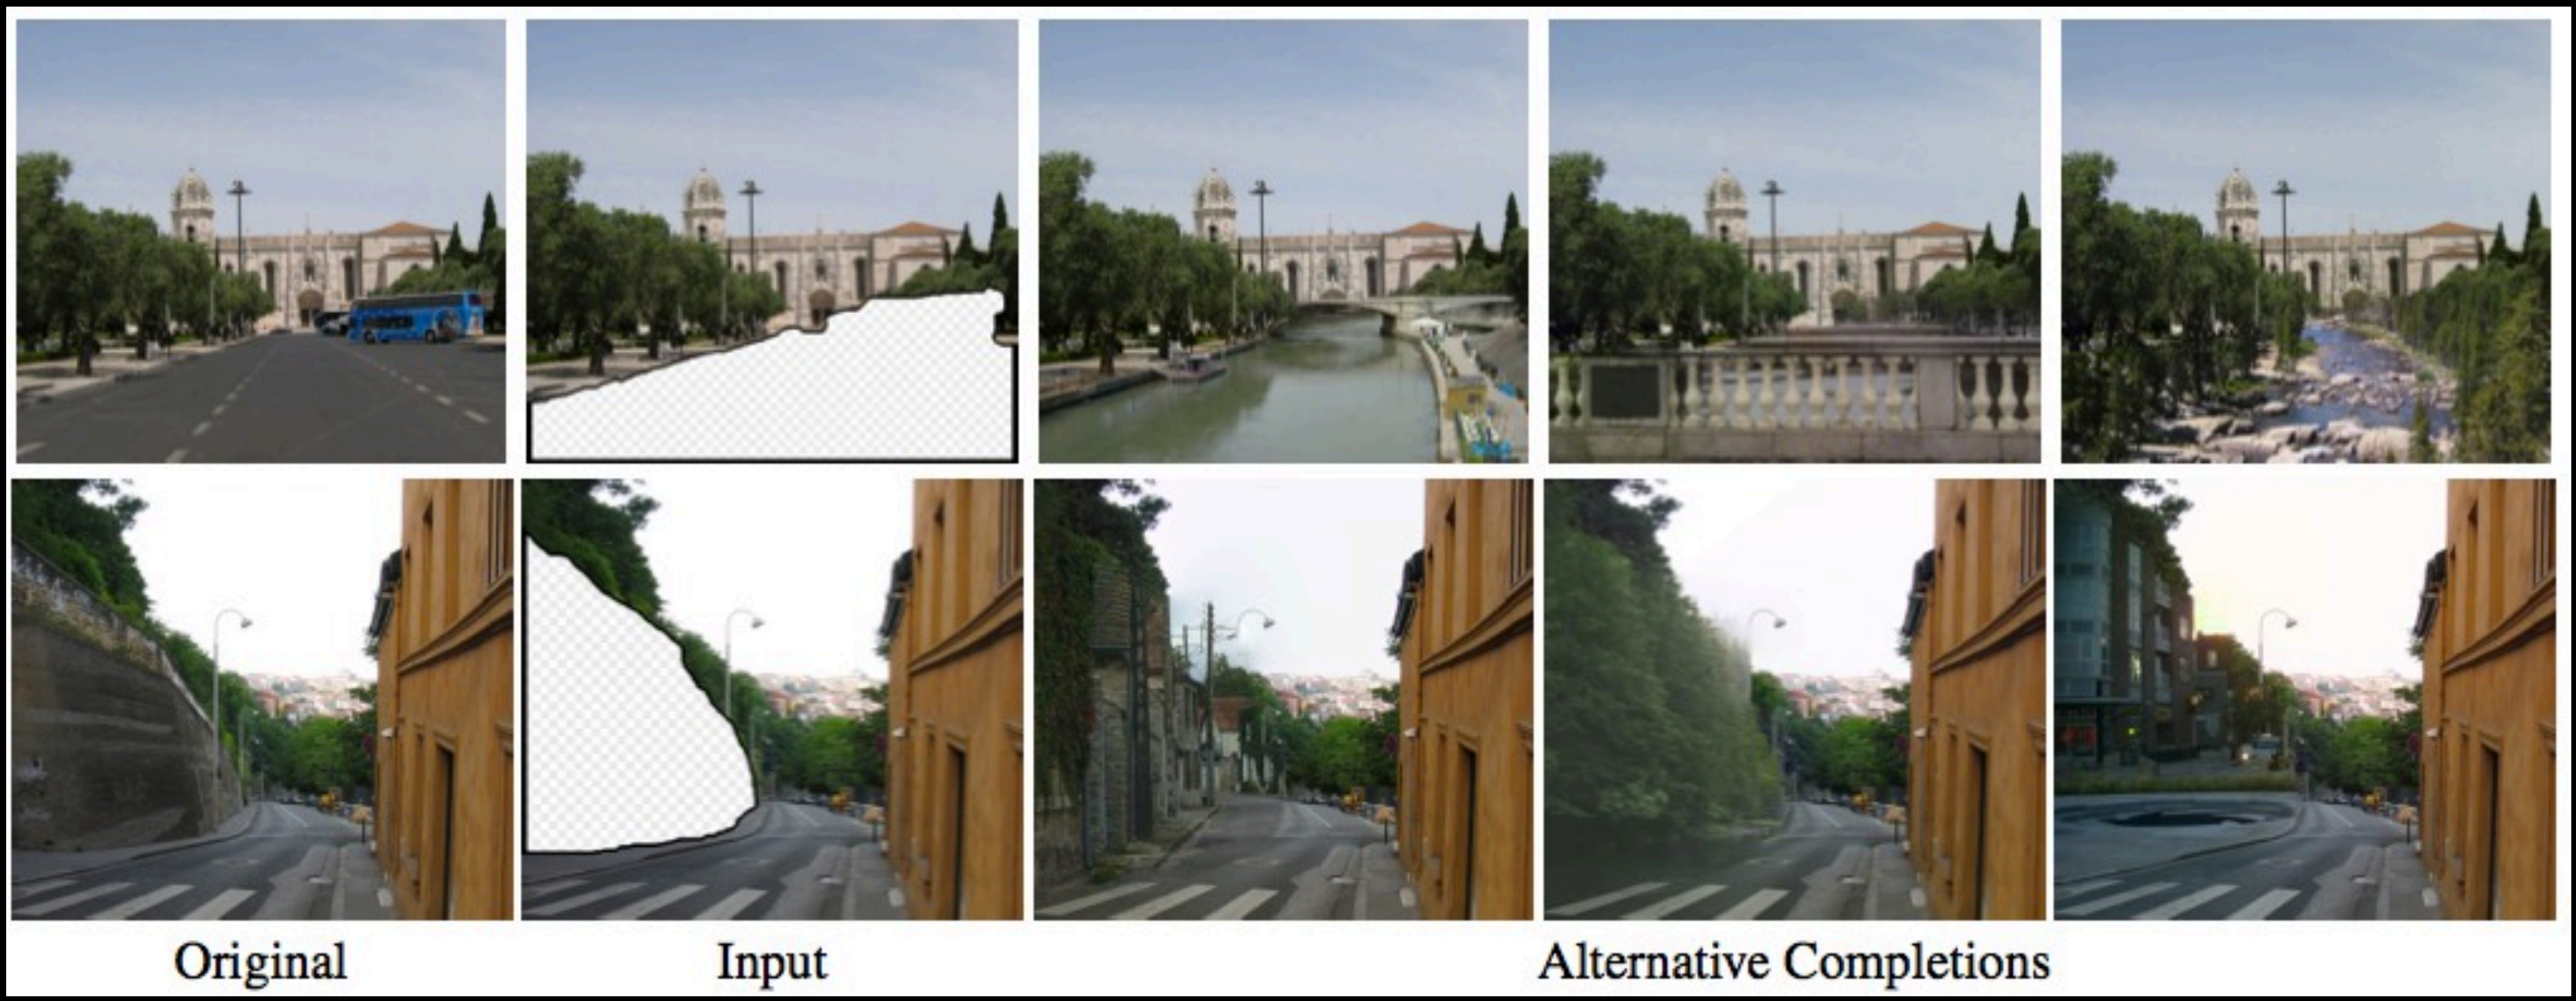

Aside: Phototourism

[Snavely et al. 2006]

- Navigating large image collections is hard
- Make it easy by leveraging spatial structure present in large datasets of landmarks

Aside: Phototourism

Photo Tourism

Exploring photo collections in 3D

Noah Snavely   Steven M. Seitz   Richard Szeliski

University of Washington   *Microsoft Research*

SIGGRAPH 2006

Retargeting

- Display devices have different aspect ratios

Retargeting

- Media (photos, video, webpages, etc.) have different aspect ratios.
- How do you show a 1.3:1 image on a 1:1 display?
  - Aspect Fill (crop)
  - Aspect Fit (letter box)
  - Non-uniform scale

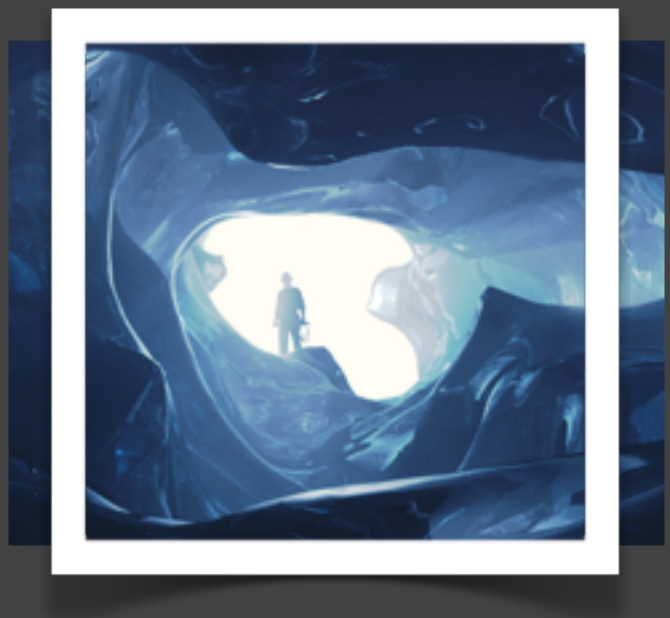

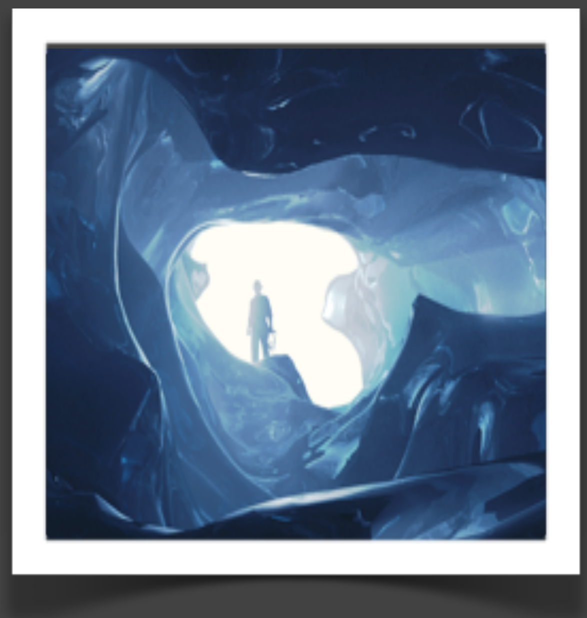

Seam Carving

[Avidan and Shamir 2007]

Seam Carving for Content-Aware Image Resizing

Shai Avidan
Mitsubishi Electric Research Labs

Ariel Shamir
The Interdisciplinary Center & MERL

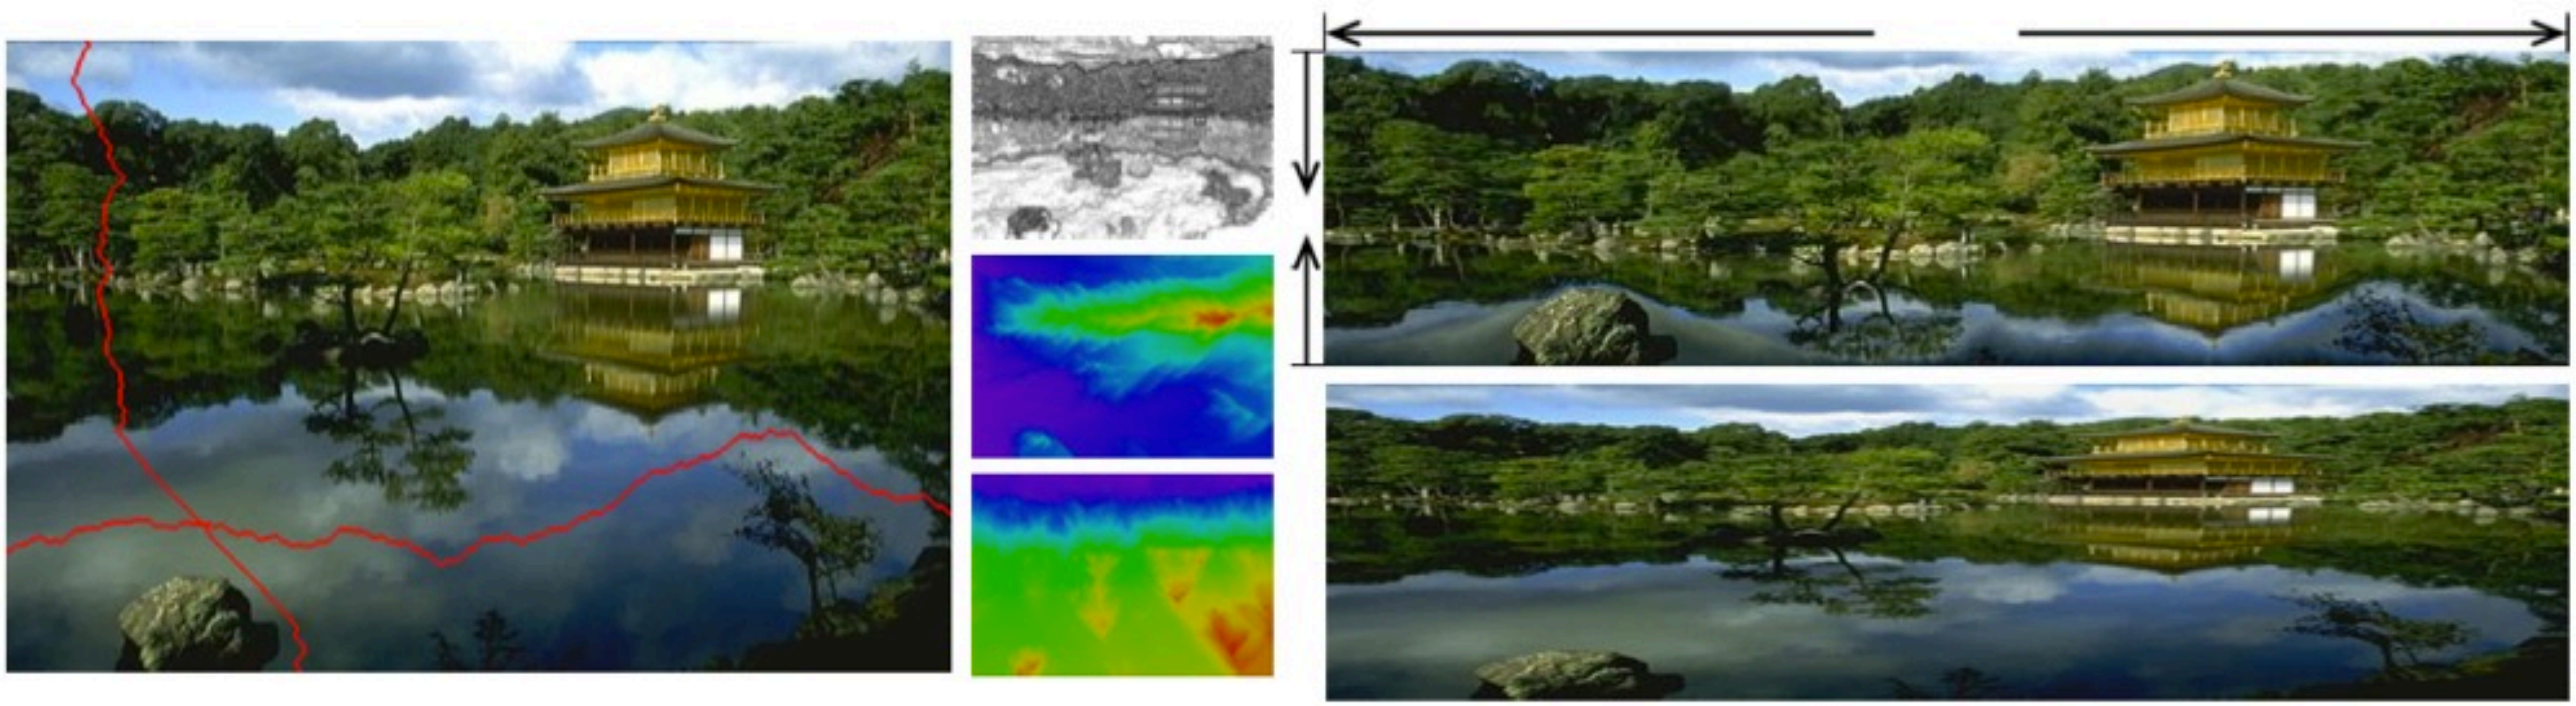

Figure 1: A seam is a connected path of low energy pixels in an image. On the left is the original image with one horizontal and one vertical seam. In the middle the energy function used in this example is shown (the magnitude of the gradient), along with the vertical and horizontal path maps used to calculate the seams. By automatically carving out seams to reduce image size, and inserting seams to extend it, we achieve *content-aware resizing*. The example on the top right shows our result of extending in one dimension and reducing in the other, compared to standard scaling on the bottom right.

Seam Carving

Energy Function (gradient magnitude)

Seam Carving

- To shrink an image, remove pixels with low energy

(b) Crop (c) Column (d) Seam (e) Pixel (f) Optimal

Seam Carving

- Find minimum energy seams by dynamic programming
- Suppose we want to reduce image width by one pixel
  - Define an image C such that C(x,y) is the minimum cost of a seam passing through (x,y)
  - $C(x,y) = e(x,y) + \min(C(x-1,y-1), C(x,y-1), C(x+1,y-1))$
  - Best seam found by backtracking from y_{max}

Seam Carving

- To enlarge an image, insert pixels where they are unlikely to be noticed (low energy seams)

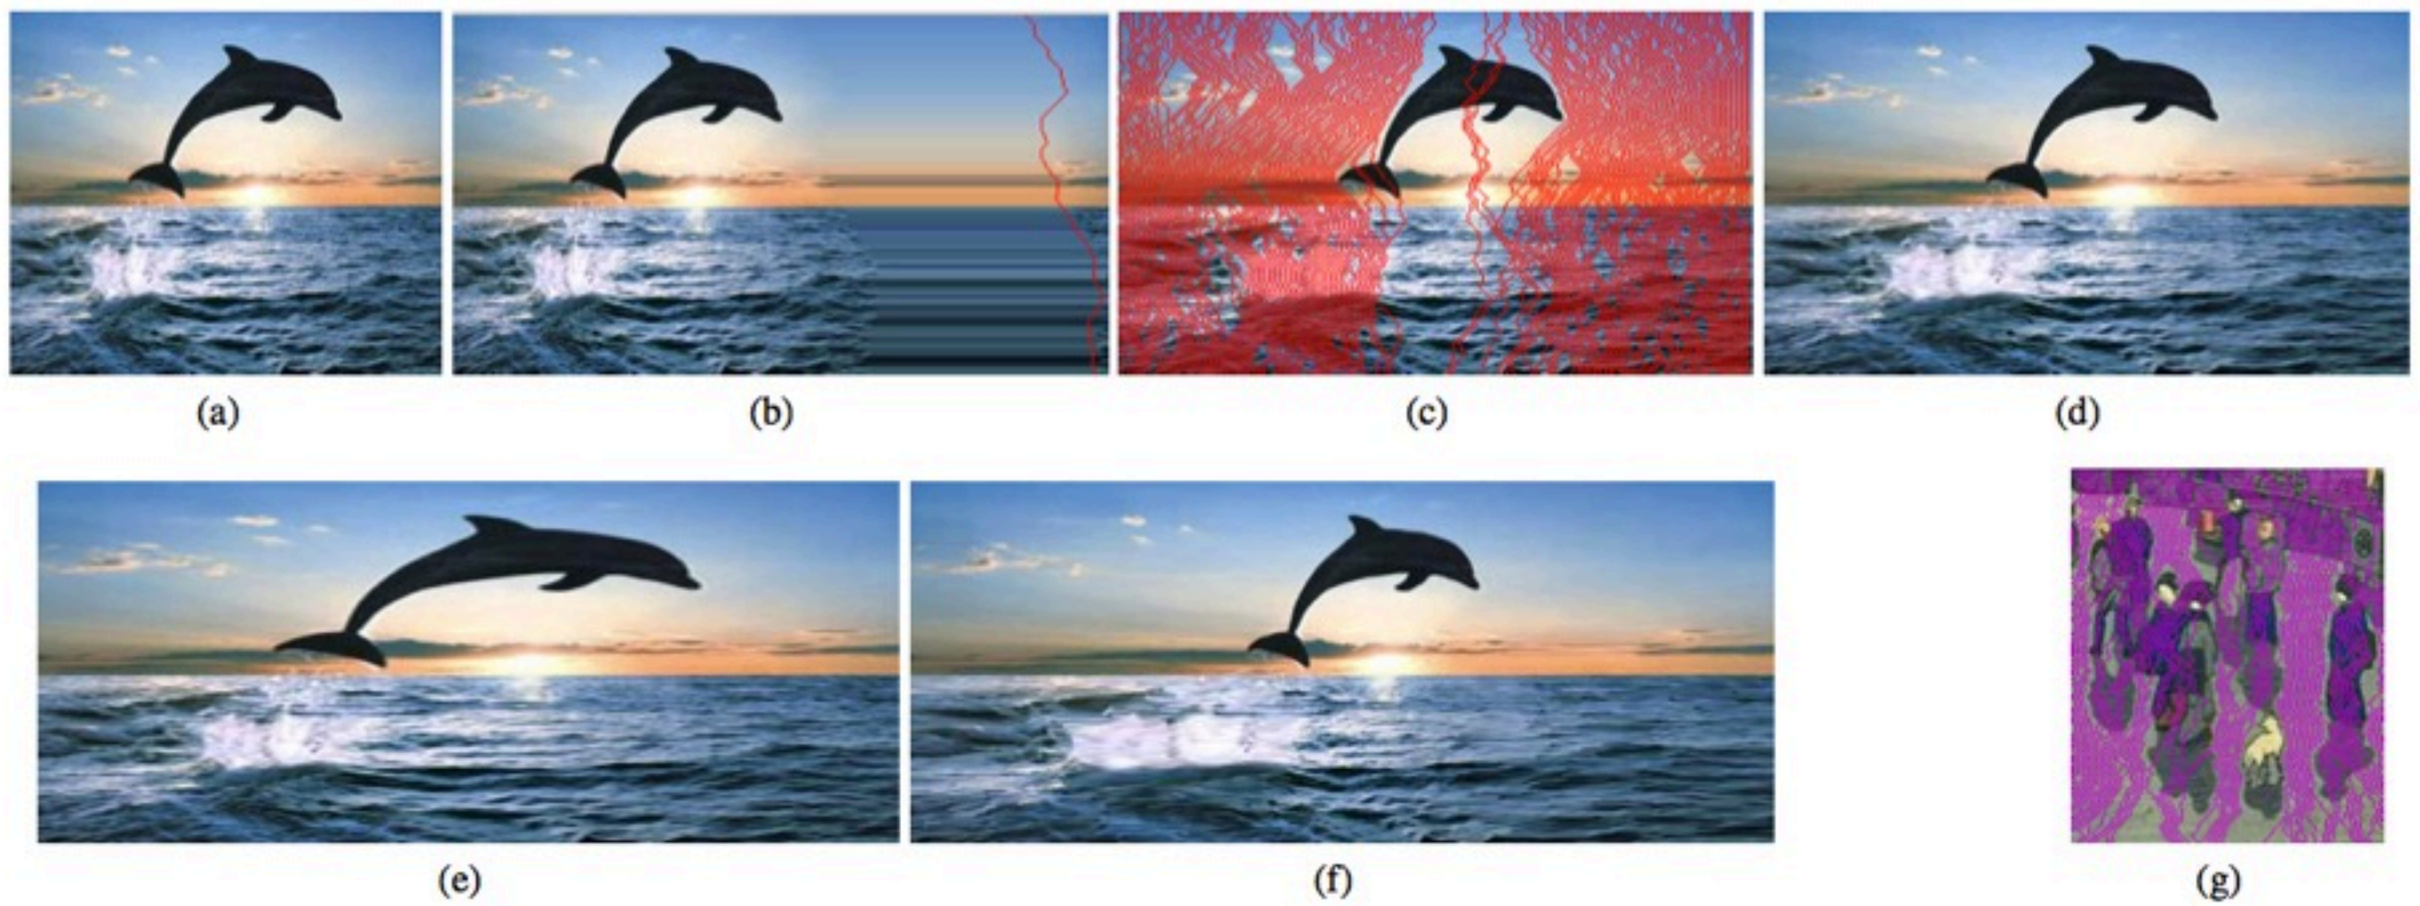

Figure 8: Seam insertion: finding and inserting the optimum seam on an enlarged image will most likely insert the same seam again and again as in (b). Inserting the seams in order of removal (c) achieves the desired 50% enlargement (d). Using two steps of seam insertions of 50% in (f) achieves better results than scaling (e). In (g), a close view of the seams inserted to expand figure 6 is shown.

Seam Carving

- Can also use seam carving for hole-filling*

*The entire image changes, so it's not the typical meaning of hole-filling

Seam Carving

- Fails when good seams don't exist

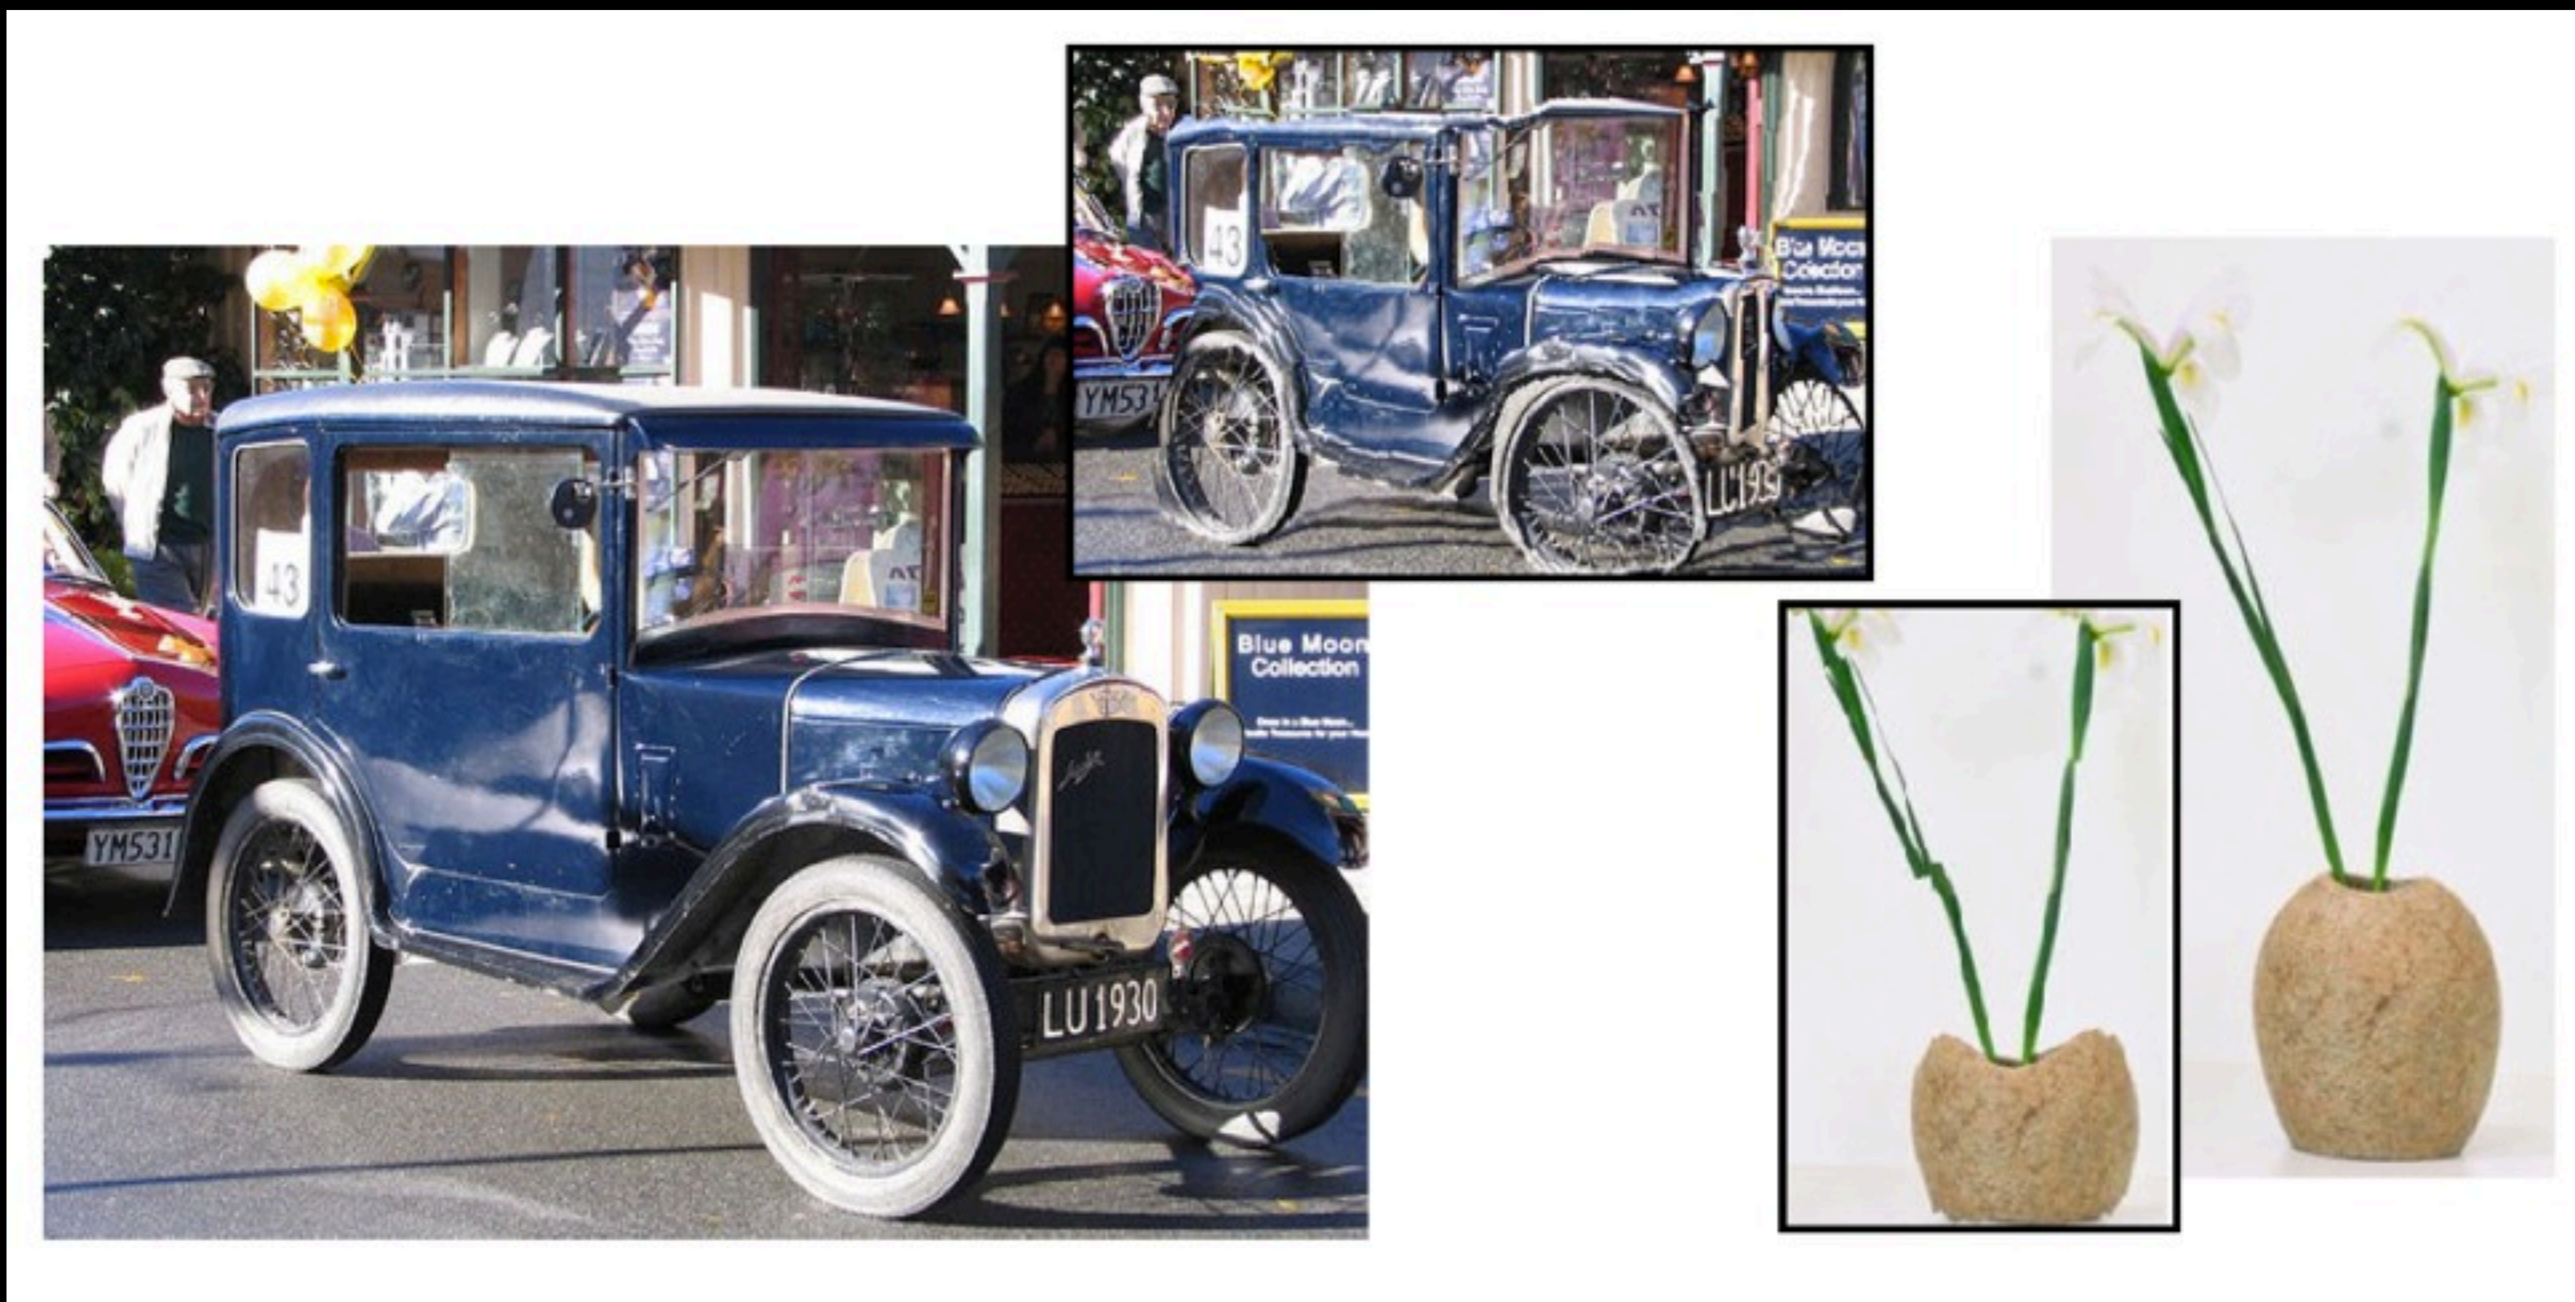

Now we watch paper videos until we run out of time

- Matting
  - Attenuation-Refraction Matting [Yeung et al. 2011]
- Nearest Neighbor Search
  - Non-Rigid Dense Correspondence [HaCohen et al. 2011]
- Image Collections
  - Finding paths through the worlds photos [Snavely et al. 2008]
- Retargeting and resizing
  - Video seam carving [Rubinstein et al. 2008]
  - Multi-operator retargeting [Rubinstein et al. 2009]
  - Upscaling using Local Self-Examples [Freeman and Fattal 2011]