Copying and Editing Images

Bill Freeman,
Google Research and MIT
June 18, 2018

Research at Google

Copying and Editing Images

Photoscan
Ce Liu, Michael Rubinstein, Mike Krainin, Bill Freeman, 2016

Securing Visible Watermarks
Tali Dekel, Michael Rubinstein, Ce Liu, Bill Freeman, CVPR17

Smart, Sparse Contours to Represent and Edit Images
Tali Dekel, Chung Gan, Dilip Krishnan, Ce Liu, Bill Freeman, CVPR18

Research at Google

Google Cambridge

MIT

MIT Stata Center

Google offices, Cambridge

Google Cambridge Vision Team

VisCam Group members

team members:
Bill Freeman, Ce Liu, Miki Rubinstein, Dilip Krishnan, Inbar Mosseri, Forrester Cole, Aaron Sarna, Tali Dekel, Mike Krainin, Aaron Maschinot, Daniel Vlasic

We take summer interns!

Dereflection: from research paper to product

Summer, 2015

A Computational Approach for Obstruction-Free Photography. T. Xue, M. Rubinstein, C. Liu and W. Freeman. Siggraph 2015

https://youtu.be/xoyNiatRIh4?t=167

Google

Google Photos came to us with a critical problem: glare

January, 2016

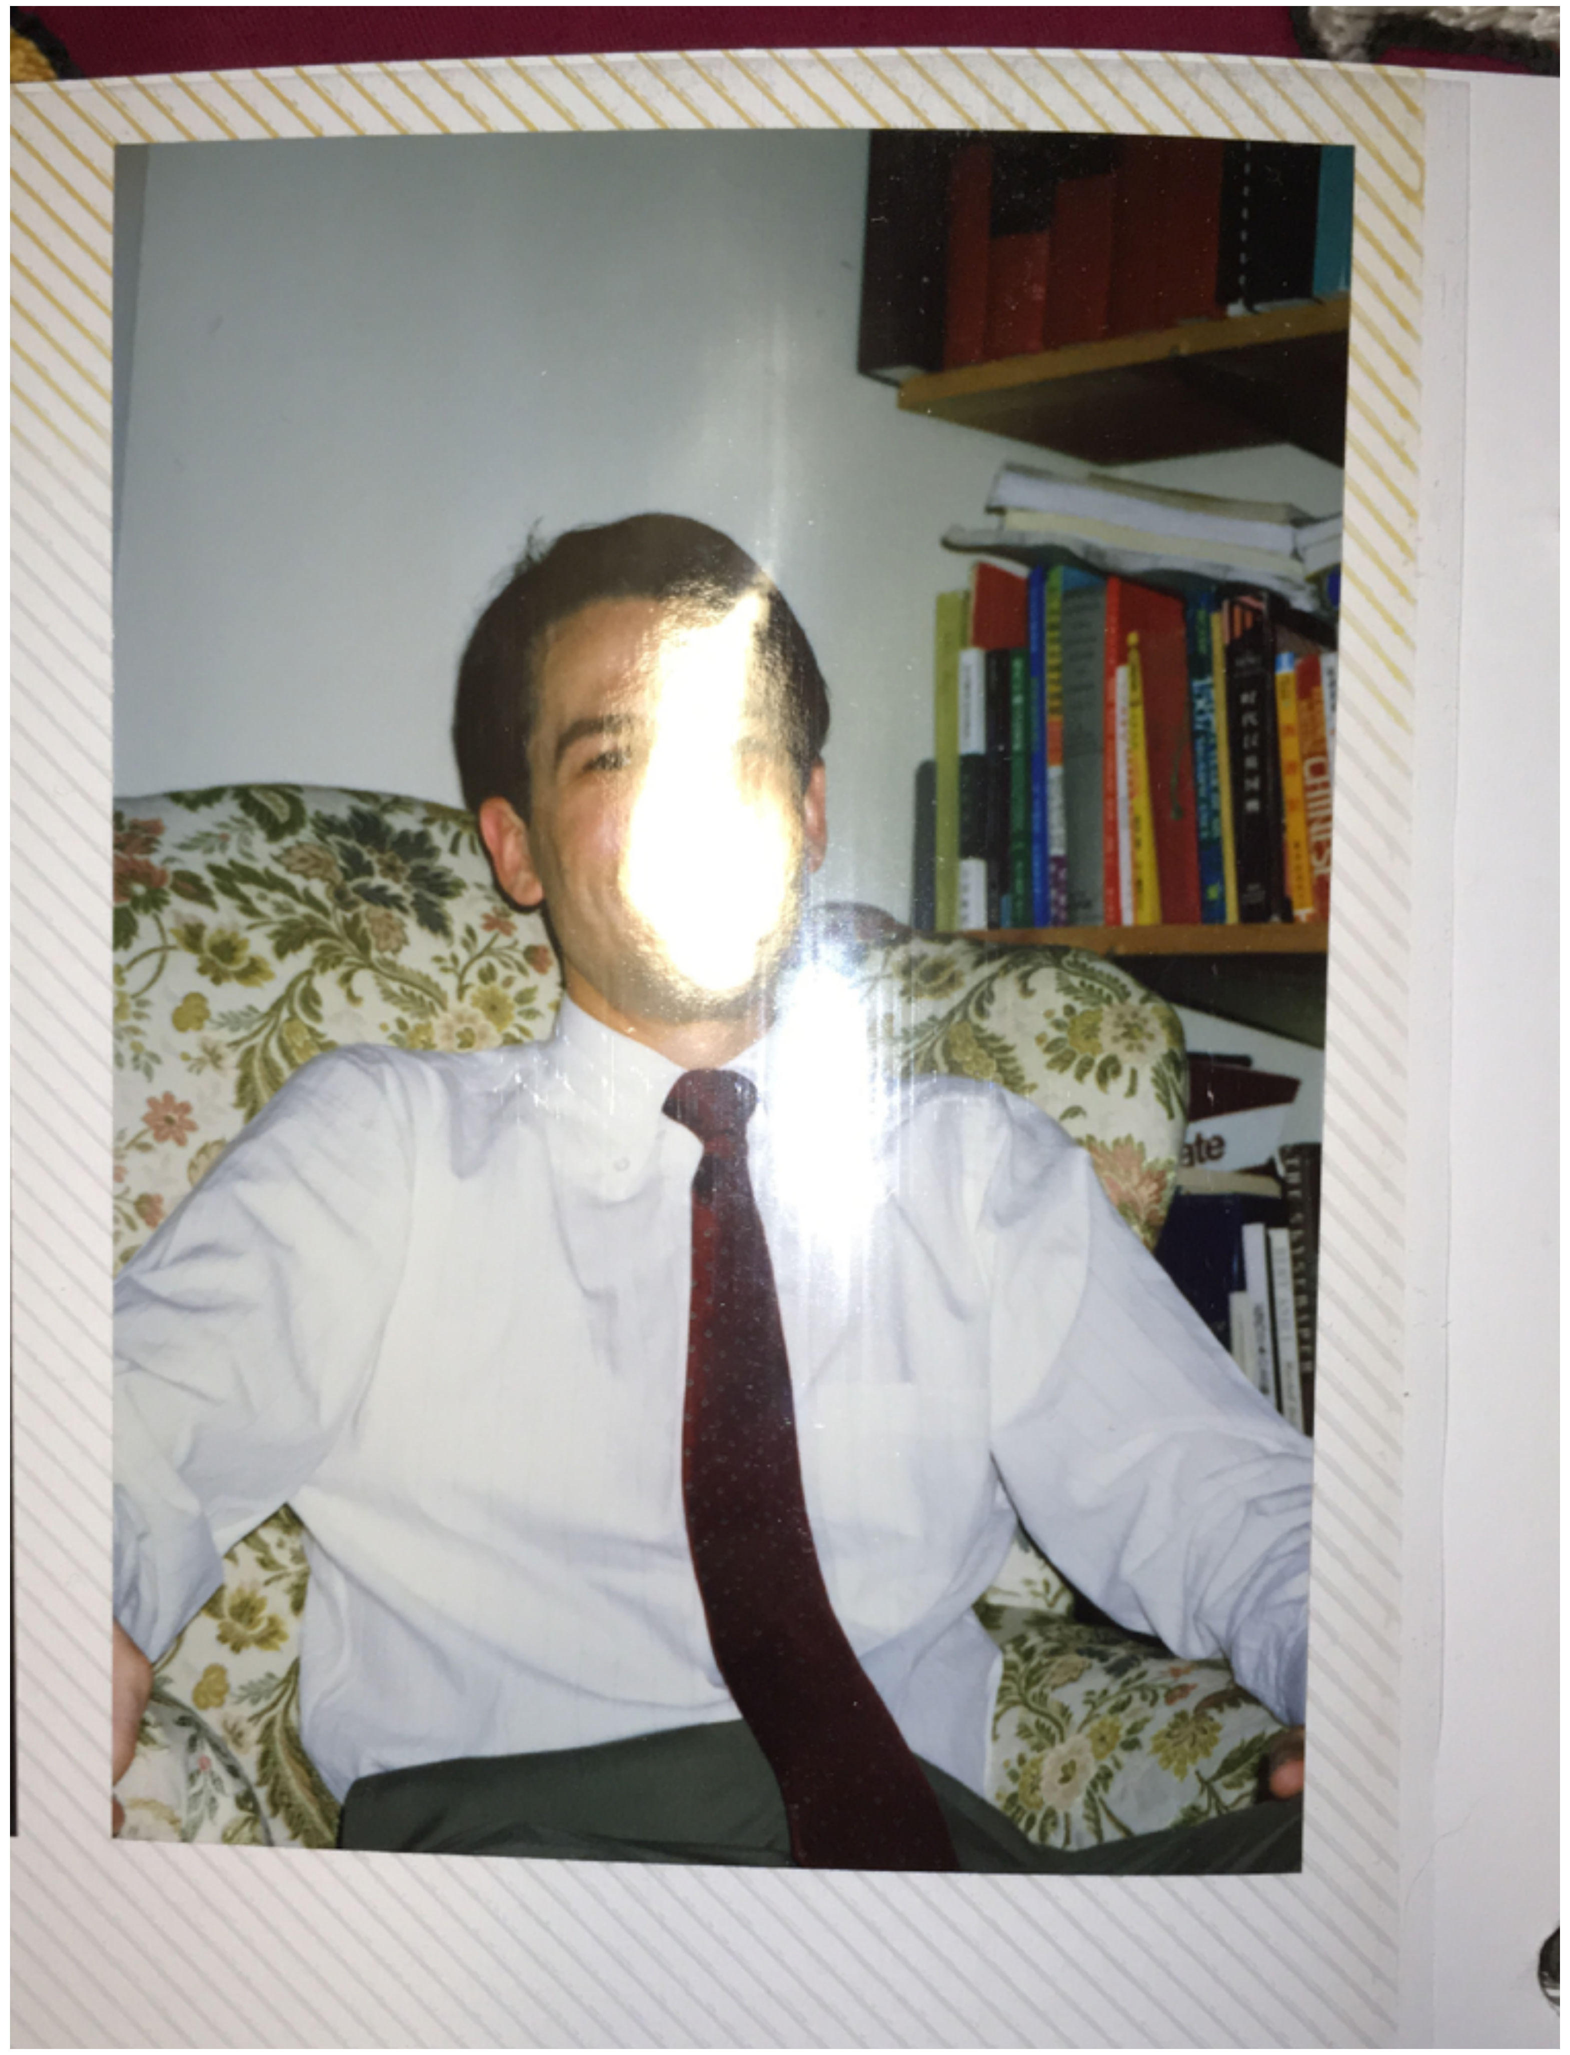

Photos in albums, in frames, or with glossy finishes suffer from glare.

We modified dereflection for on-phone glare removal

Input five frames

Output cropped glare-free image

Google

https://www.youtube.com/watch?v=MEyDt0DNjWU

How does the system work?

- Assume the glare-free version of a pixel is visible in one of the captures
- Register the five images
- Remove glare by treating it as outliers

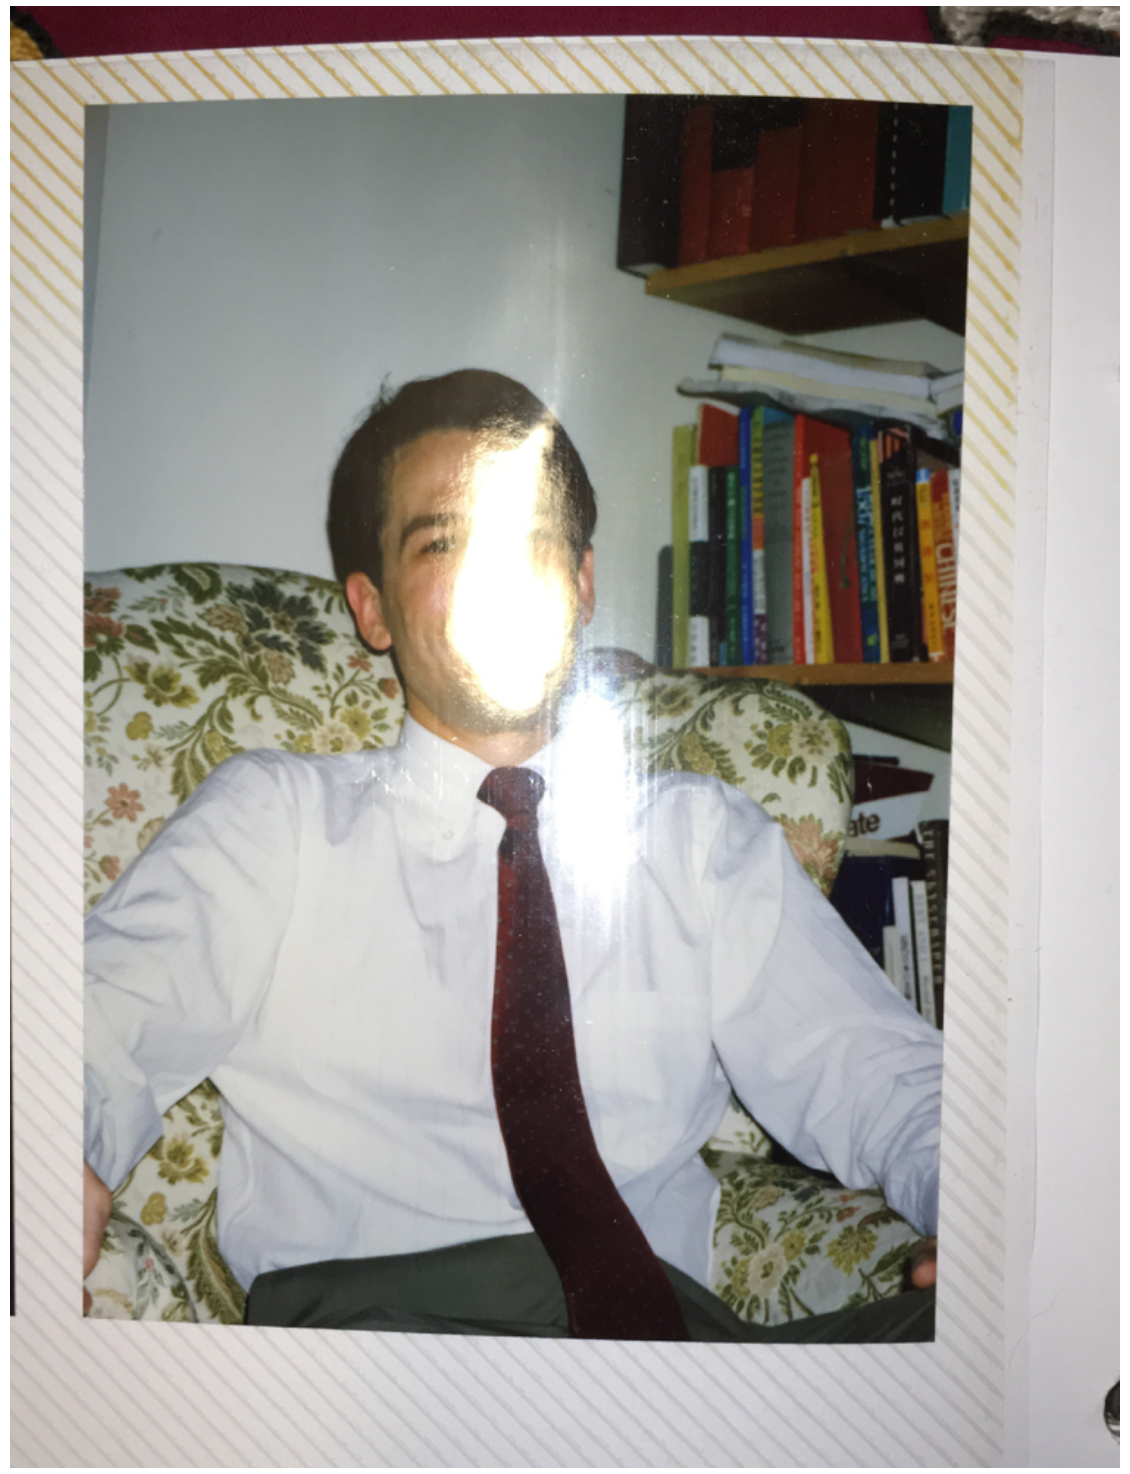

Reliable registration

Sparse feature points detection and matching

Refinement using optical flow

Flow color coding

Google

Accurate registration is key

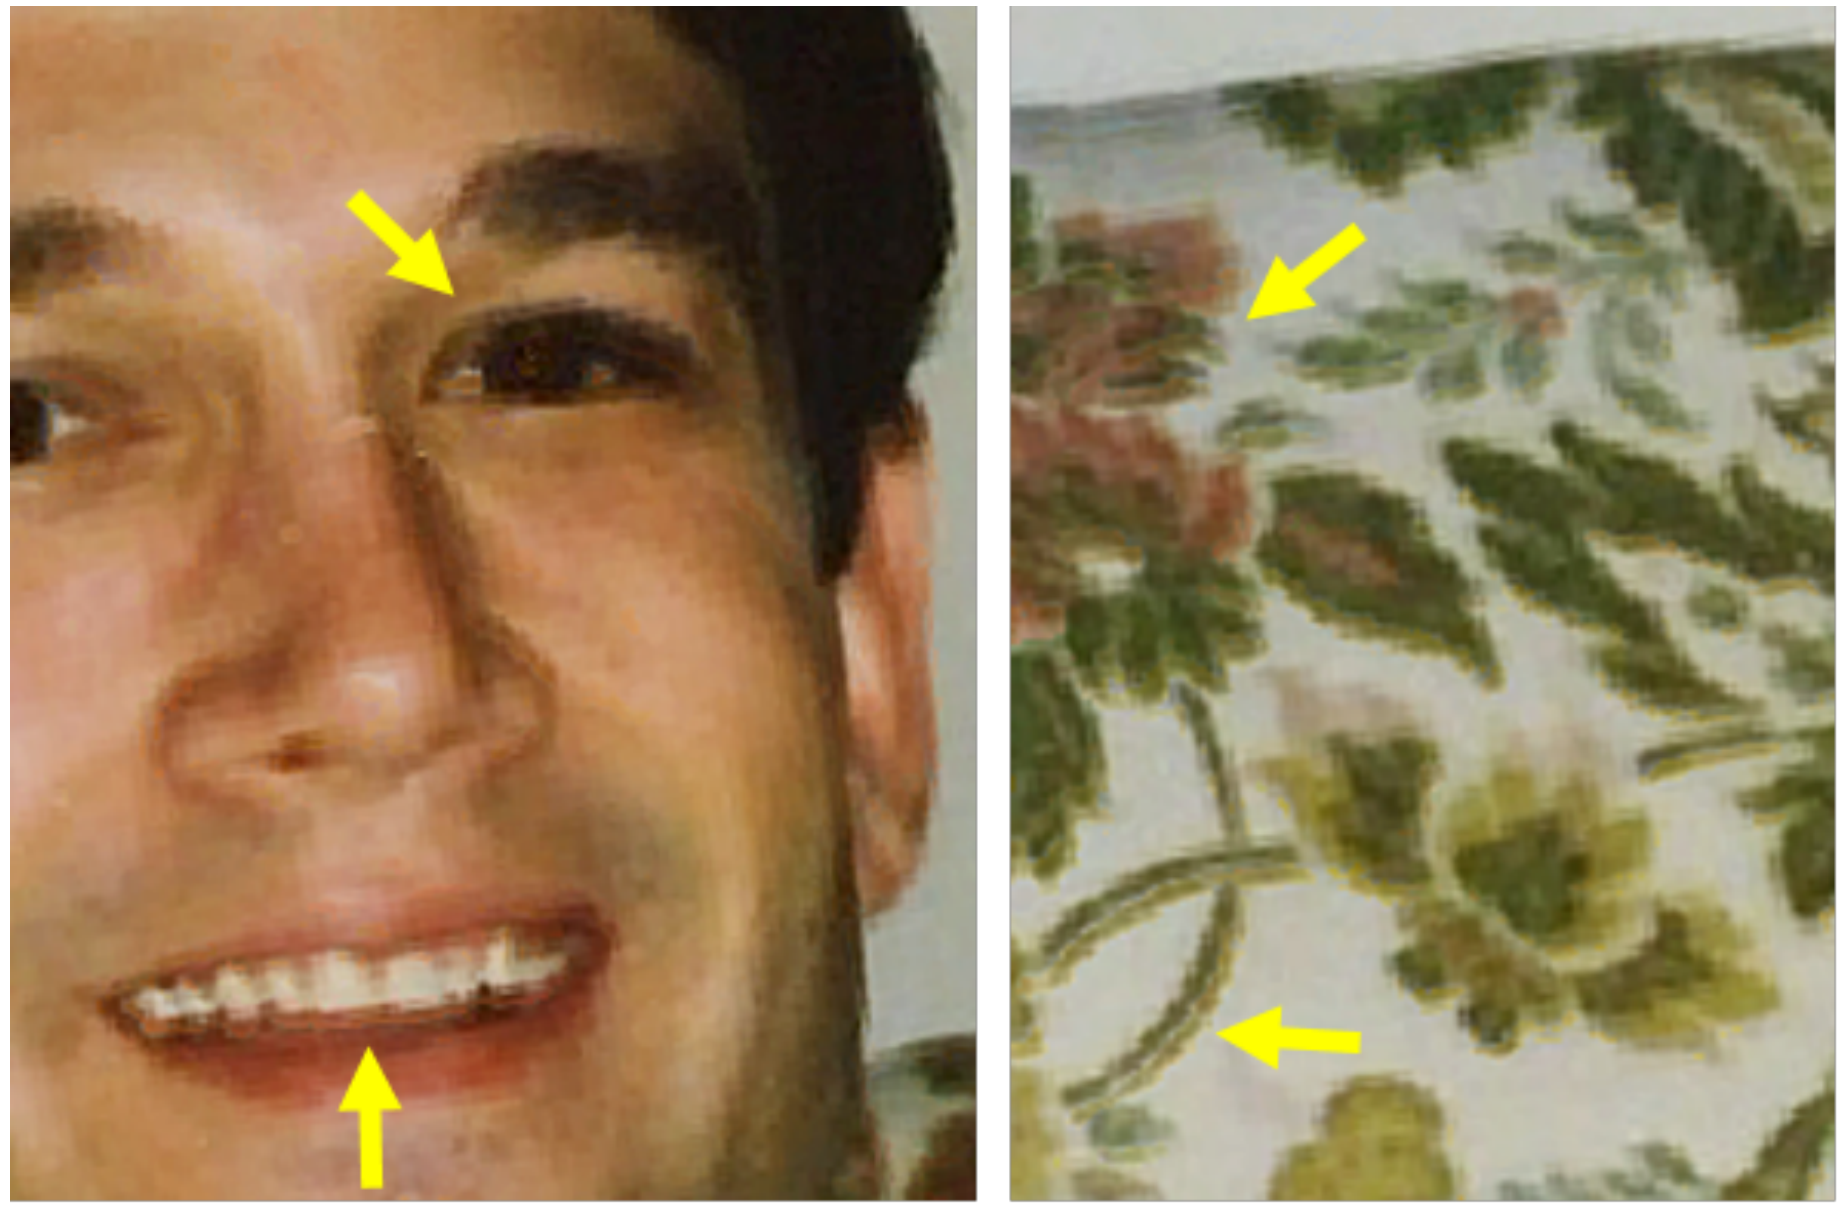

Homography-based registration only

With optical flow refinement

Some results

Google

VisCAM Contributors: Mike Krainin, Ce Liu, Miki Rubinstein, Bill Freeman

PhotoScan, released Nov. 2016, metrics and reviews

Glare removal was a prominent, enabling feature in the release.

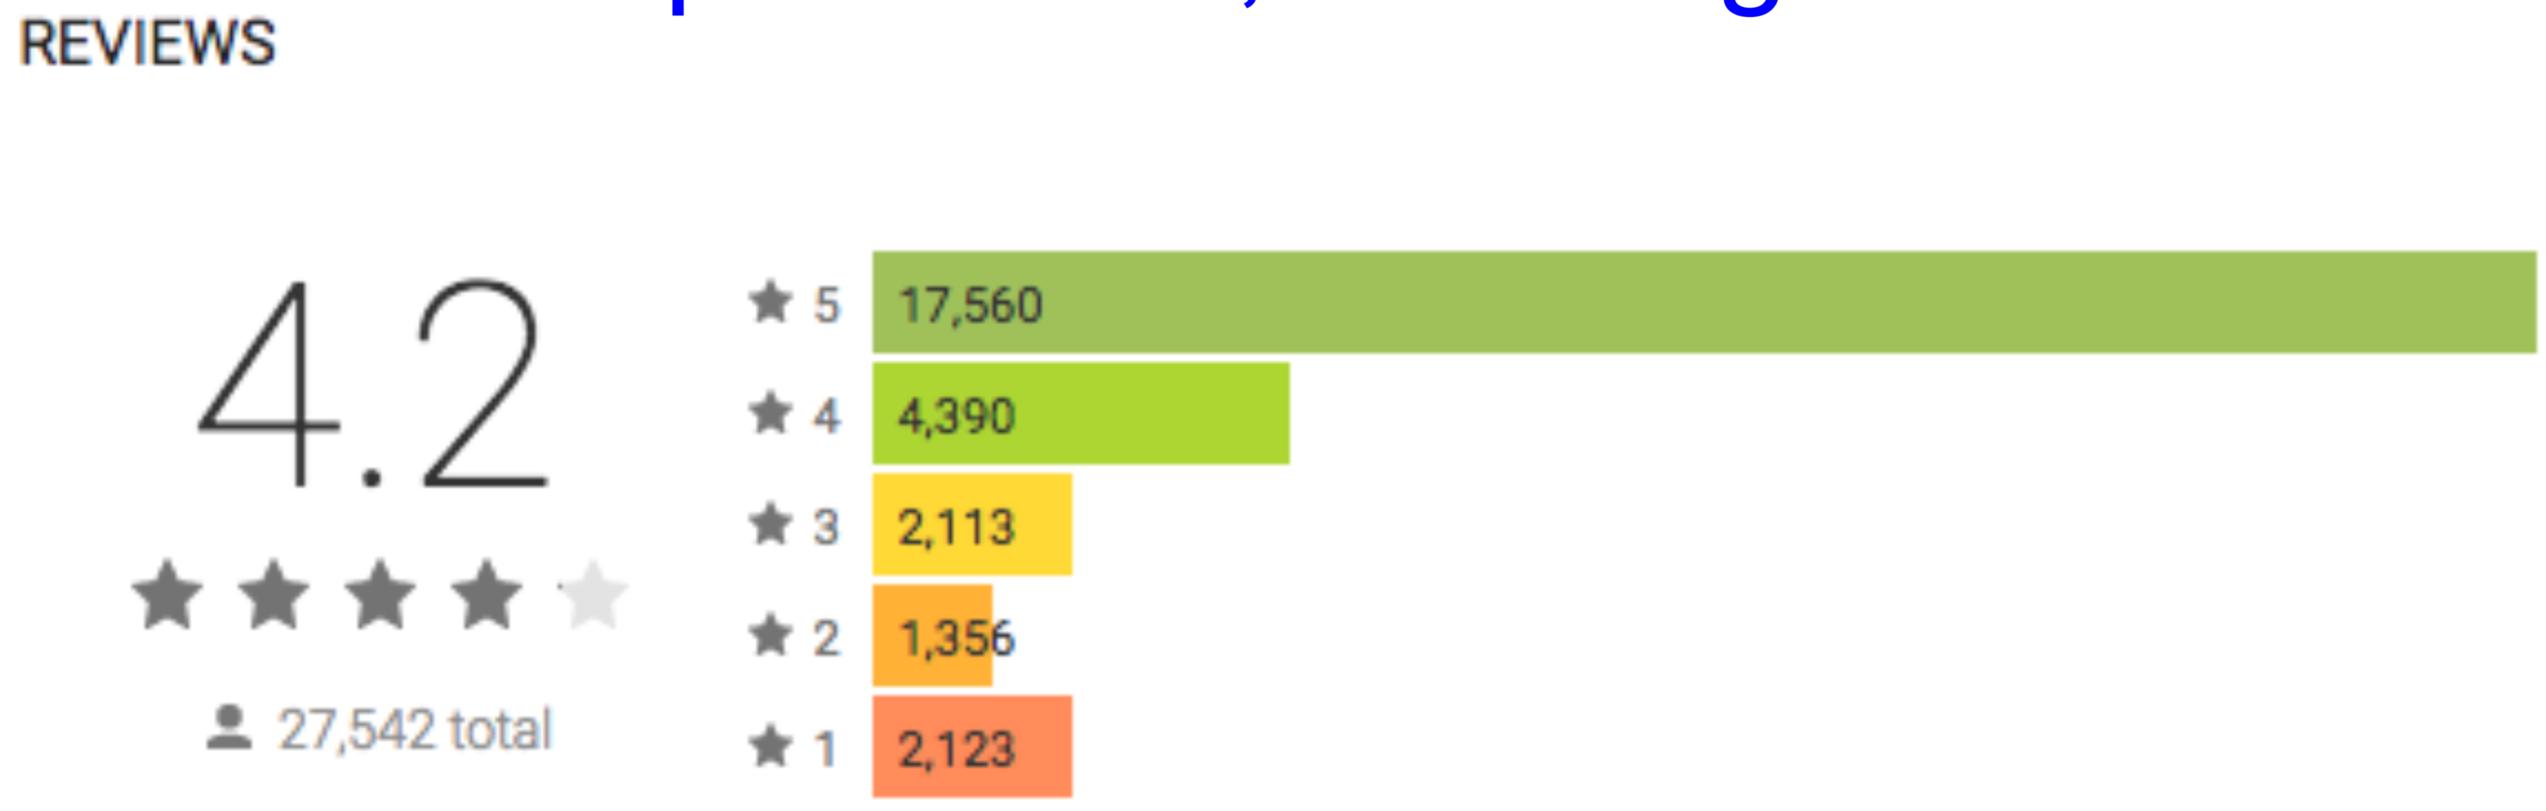

PhotoScan has created **tremendous impact**:

- As of 2/25/2017, the app has been downloaded **4.5 million** times. **32 million** photos scanned (more stats on this page).
- Received enthusiastic reviews, **5.0** on iOS (version 1.3) and **4.2** on Android.
- Much media coverage: CNN, NYTimes, Engadget, TheVerge
- Ranked among the top 10 apps of the year by Fast Co (#6) and Mashable (#9)

Press Coverage:

- *"Google is doing more interesting things with photography than Apple"* - AppAdvice
- "*By and large, PhotoScan is* ***simple and quick, with almost no learning curve****"* - *The New York Times*
- *"Goodbye to the glare and that awkward surface you used to shoot the photo on. Simply put,* ***this is a lifesaver****"* - *TechRadar*
- *"PhotoScan's* ***interface is admirably simple****"* - The Verge

Dereflection team (from Nat and Lo video)

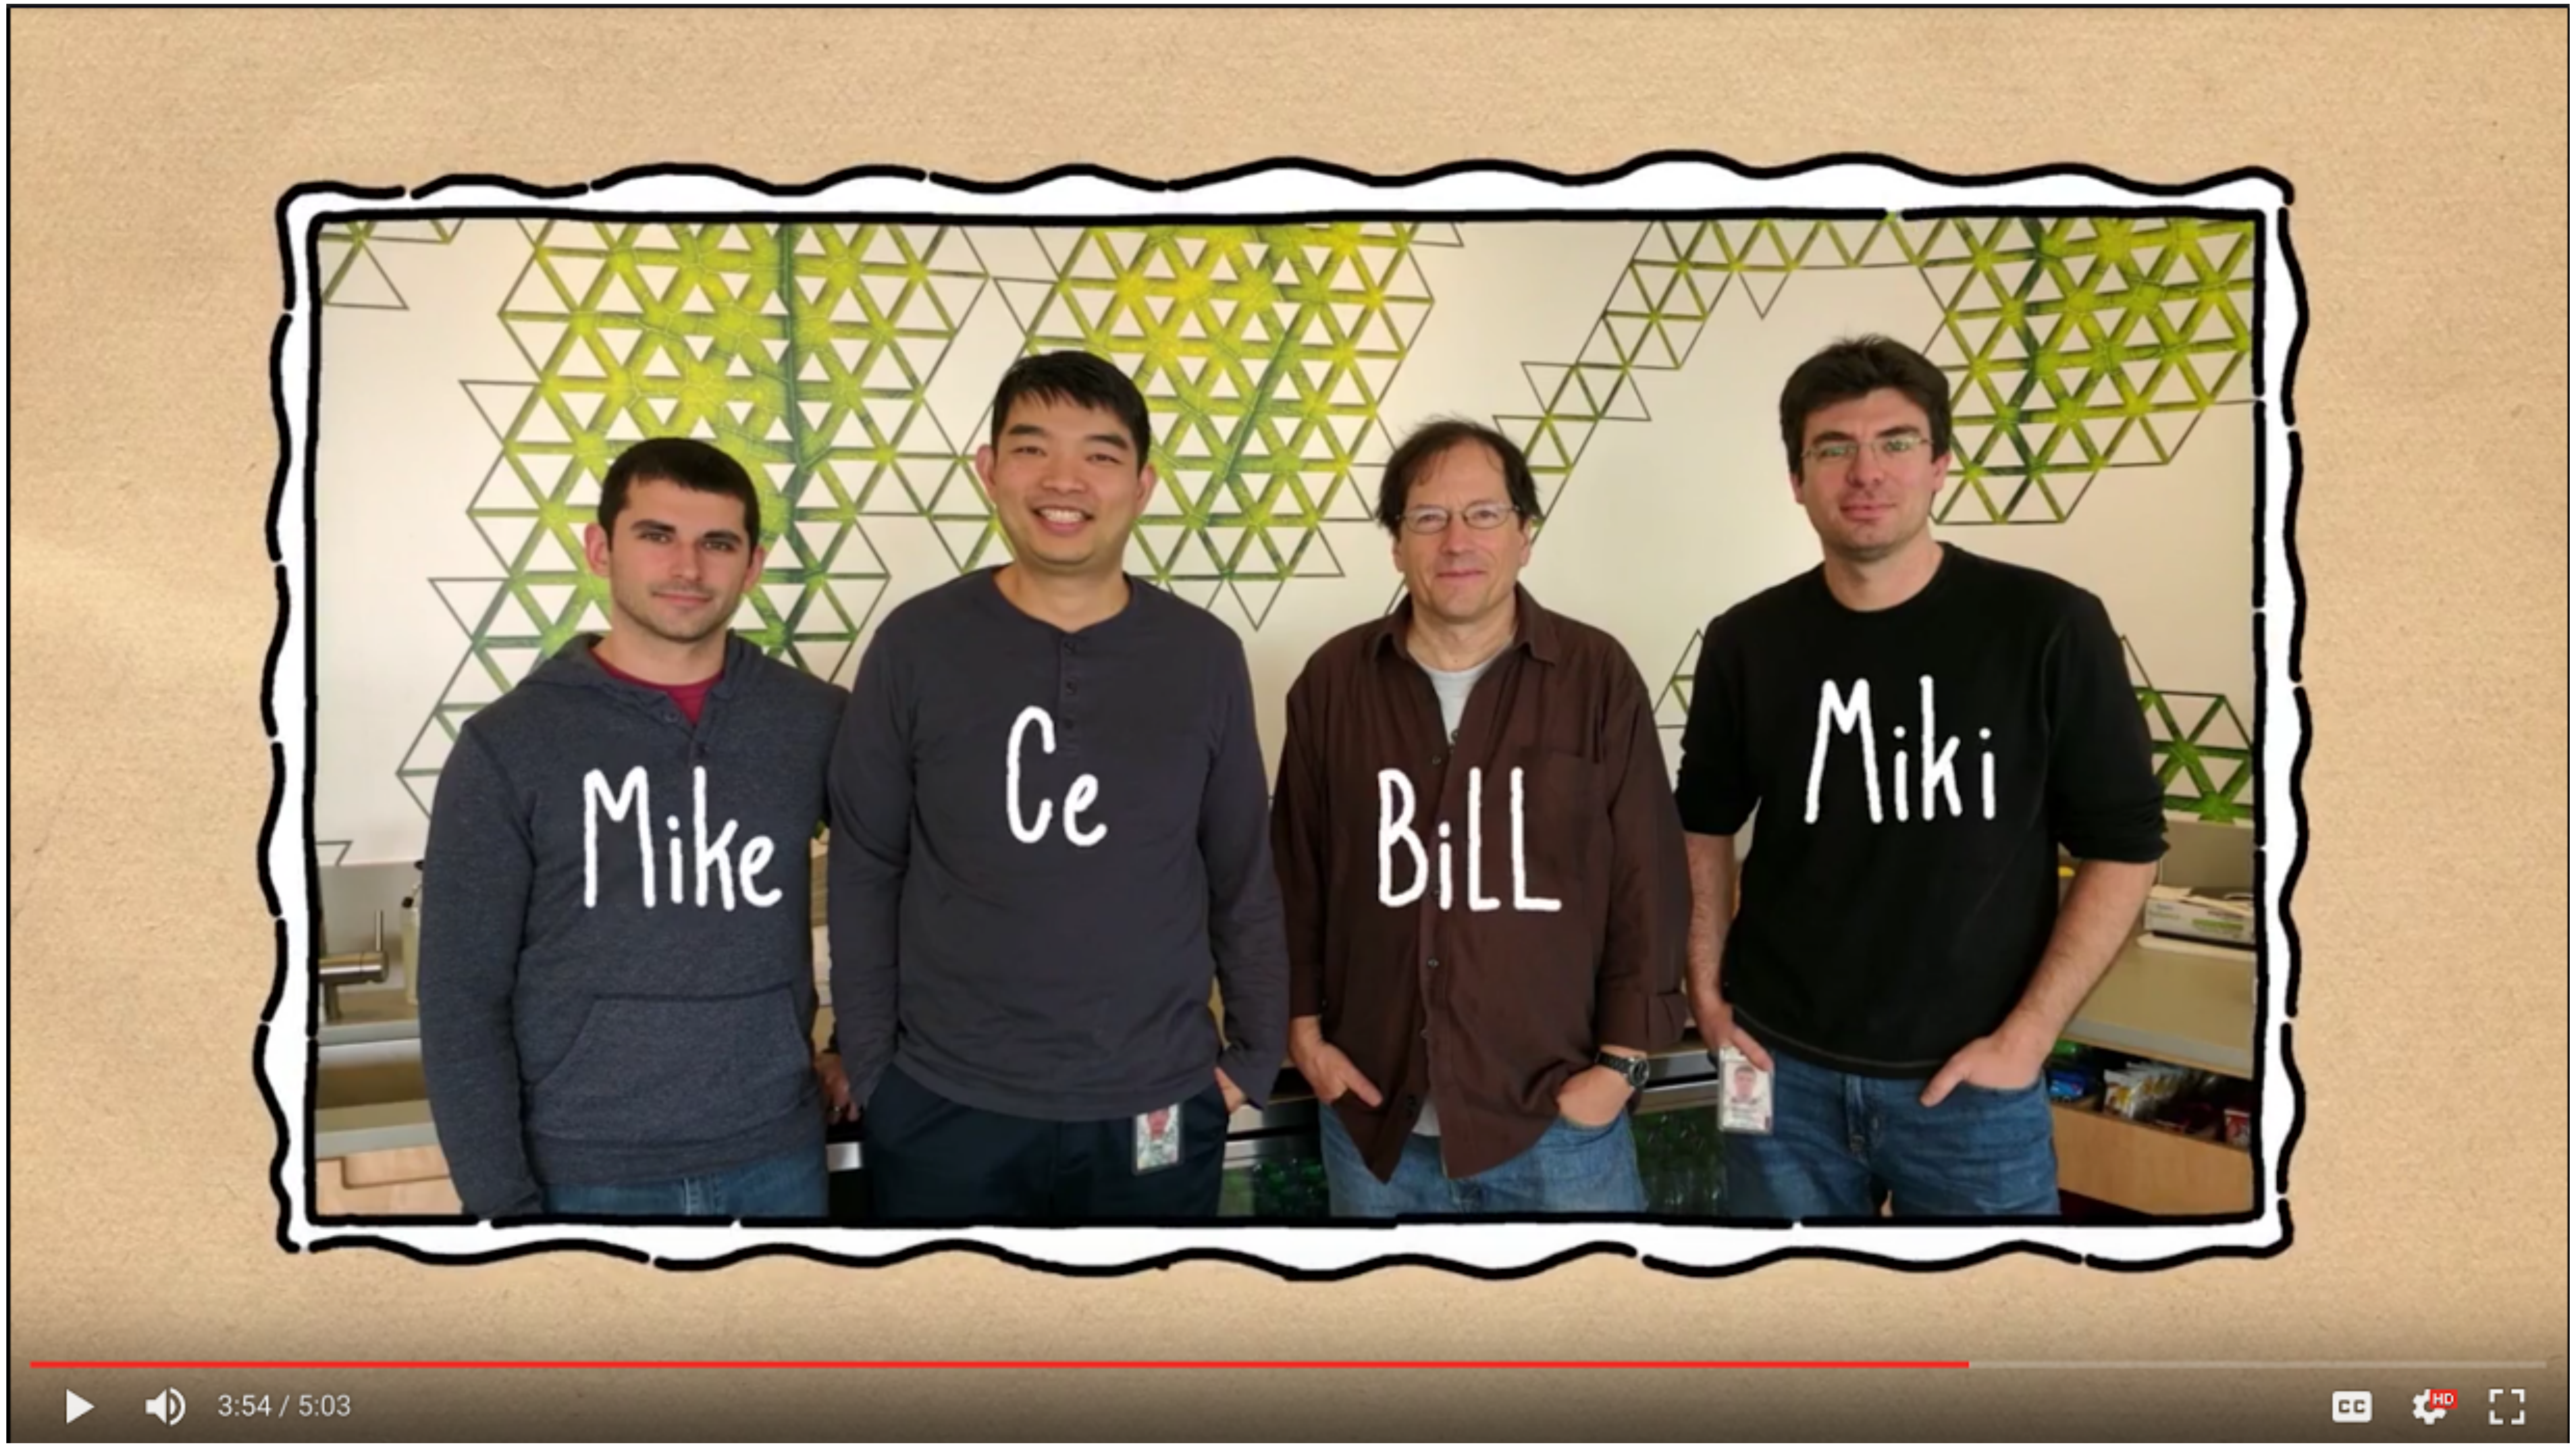

Google

Securing Visible Watermarks

Tali Dekel, Michael Rubinstein, Ce Liu, Bill Freeman, CVPR17

Tali Dekel

Miki Rubinstein

Ce Liu

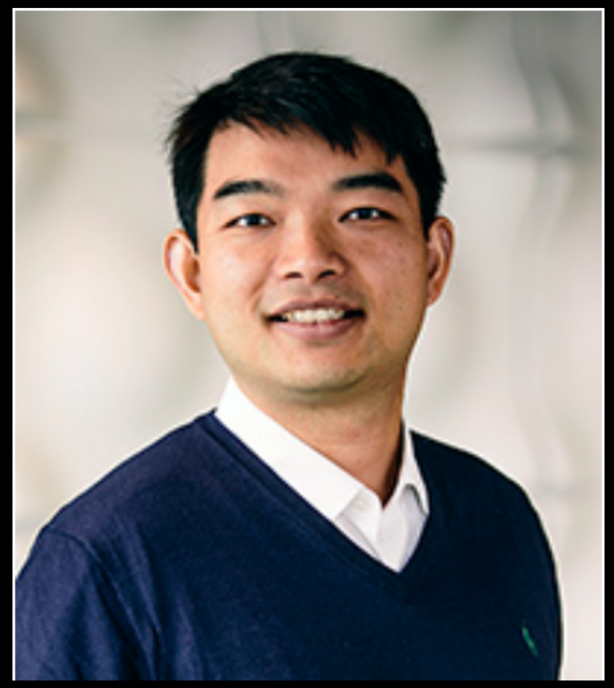

Visible watermarks are all over the Web

- Reveal the vulnerabilities of visible watermarks
- Propose solutions to improve their security

Watermarks are all over the Web

- Reveal the vulnerabilities of visible watermarks
- Propose solutions to improve their security

“…Google shows how to break watermarks…?”

Pretty similar to….

Reflections

Watermark

Tali Dekel

Miki Rubinstein

Ce Liu

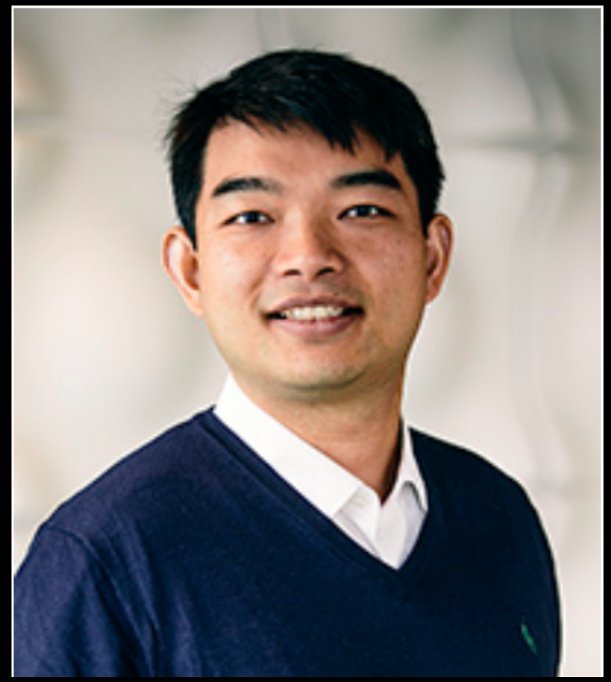

Removing a watermark in an image is hard!

Removing a watermark in an image is hard!

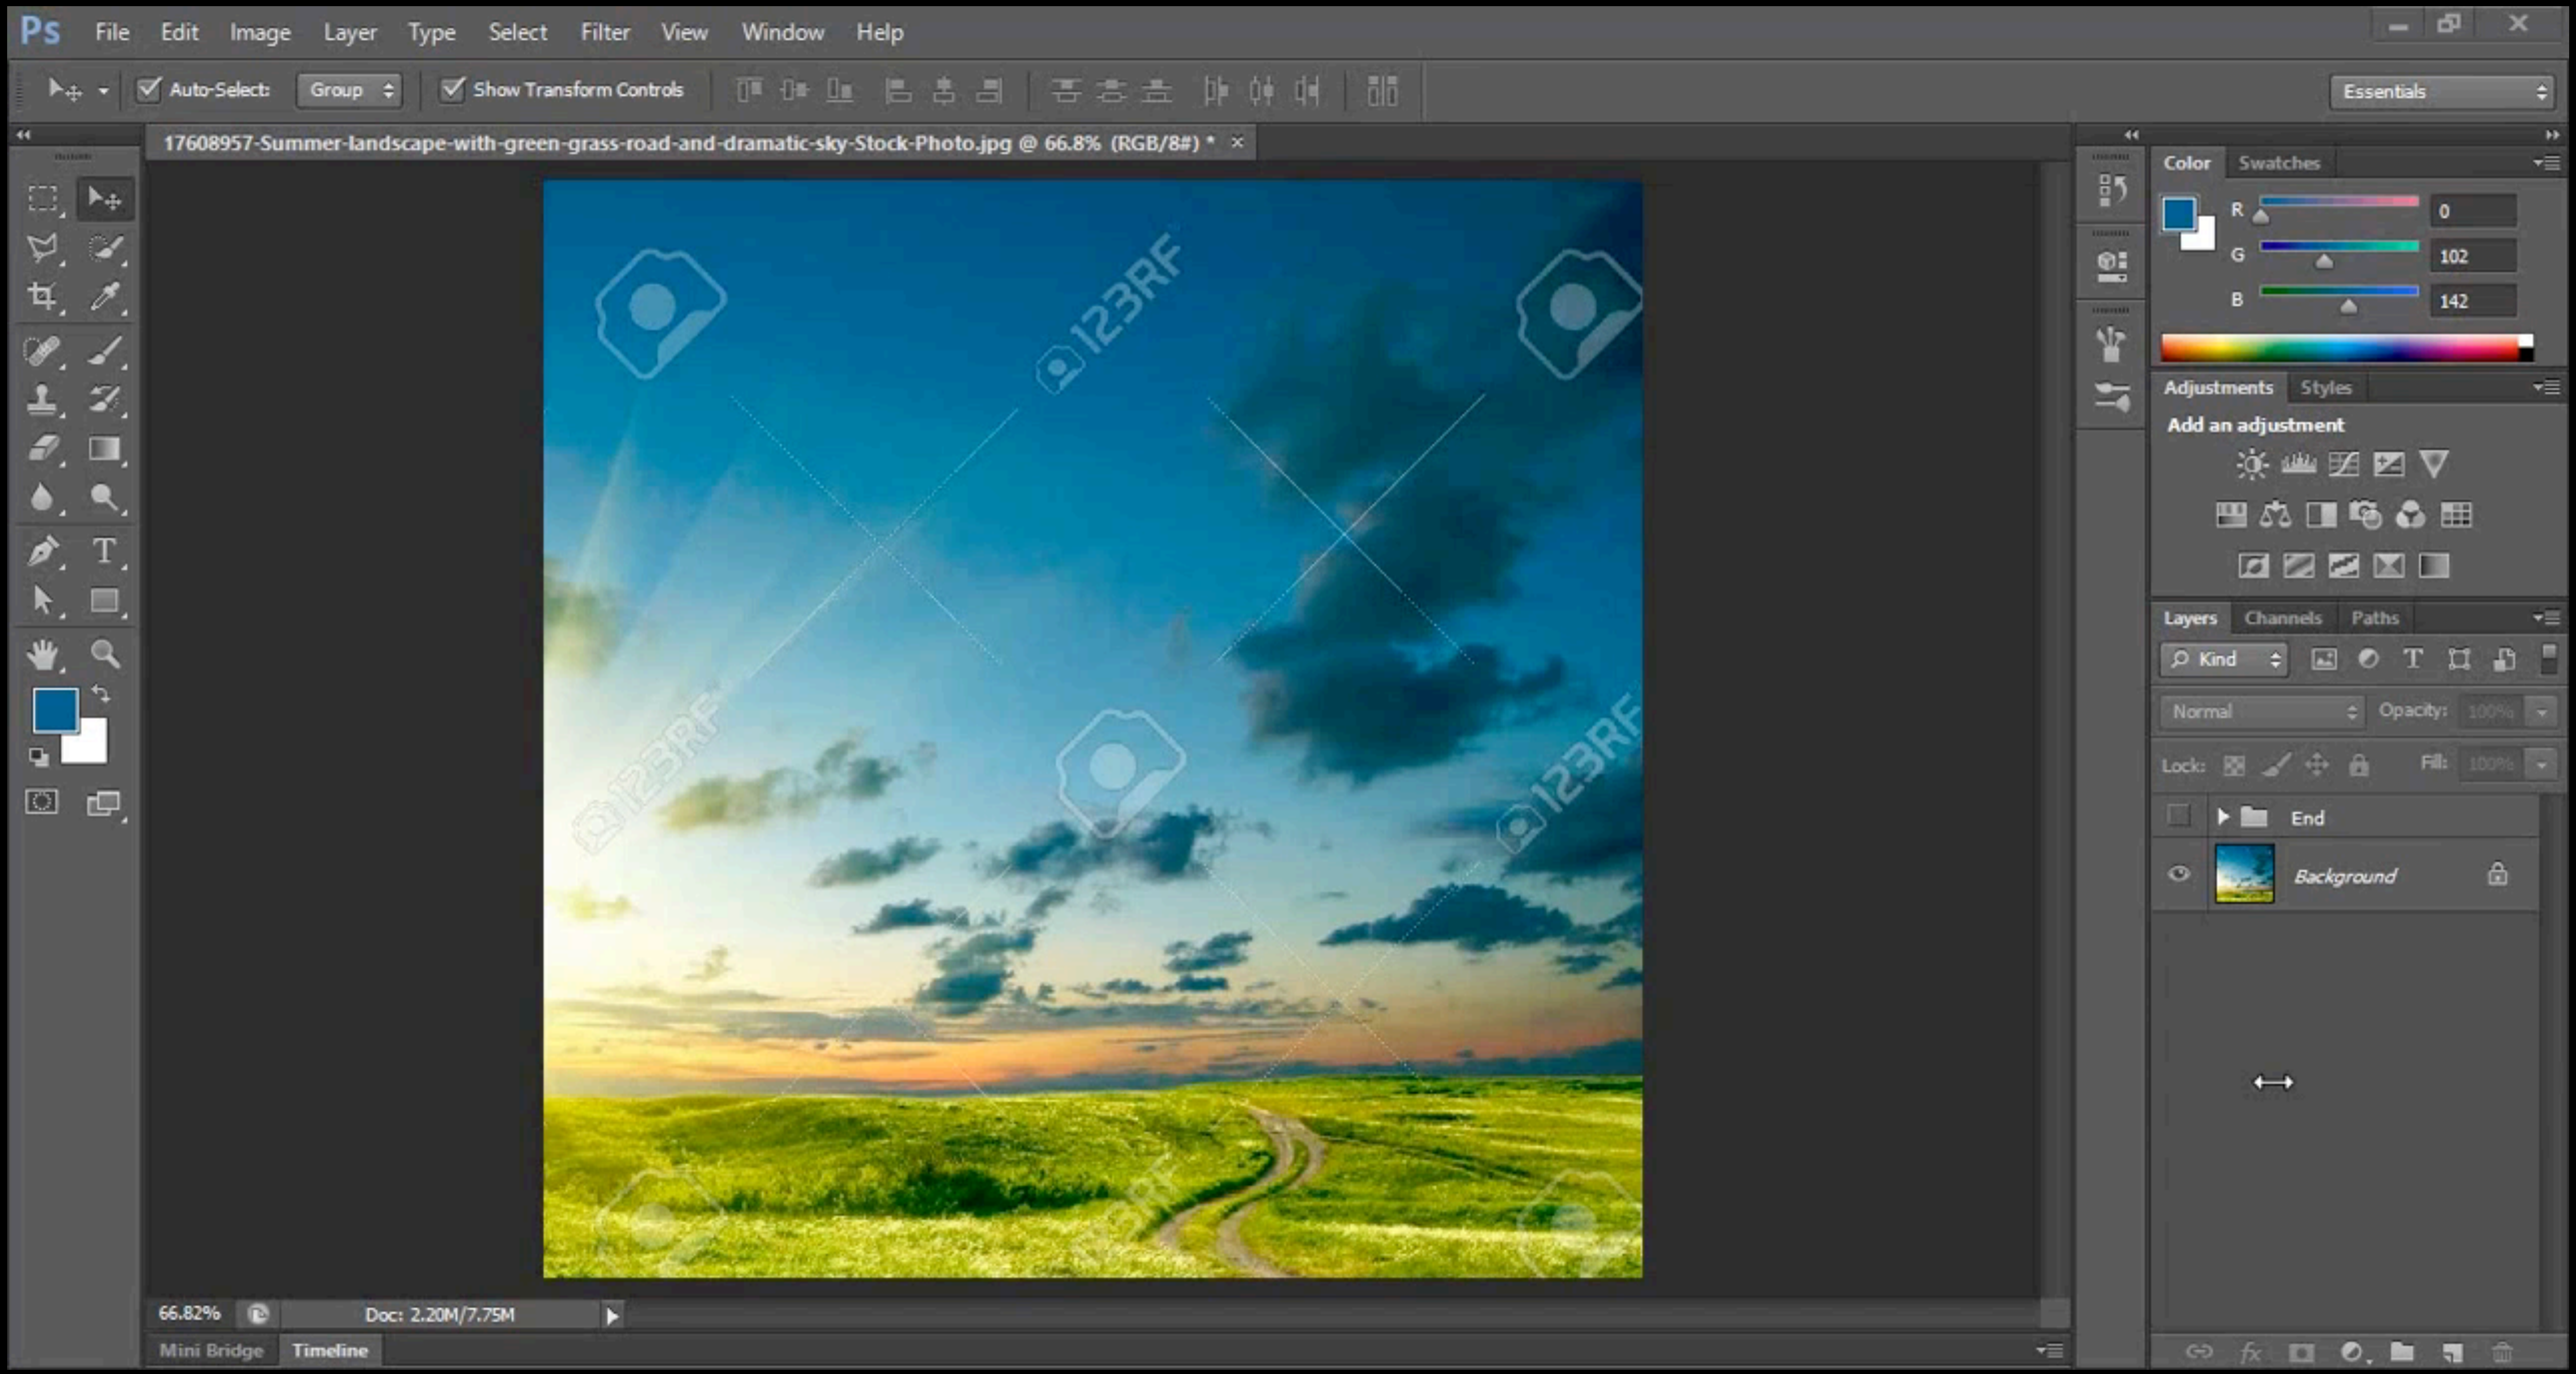

~9 minutes of editing (played at x10 speed)


Adobe Stock
#70306536

Watermarks are added **consistently** to **many** images

AdobeStock

Estimated (matted) watermark

Watermarked Image

AdobeStock

Estimated (matted) watermark

Reconstruction

AdobeStock

Estimated (matted) watermark

Watermarked Image

AdobeStock

Estimated (matted) watermark

Reconstruction

AdobeStock

Estimated (matted) watermark

Original Input Image

AdobeStock

Estimated (matted) watermark

Reconstruction

123RF

Estimated (matted) watermark

Reconstruction

123RF

Estimated (matted) watermark

Original Input Image

123RF

Estimated (matted) watermark

Reconstruction

123RF

Original Input Image

Estimated (matted) watermark

123RF

Estimated (matted) watermark

Reconstruction

Hatzalmania (הצלמניה)

Estimated (matted) watermark

Reconstruction

WARNING

NO DEEP NETWORKS

Formation Model

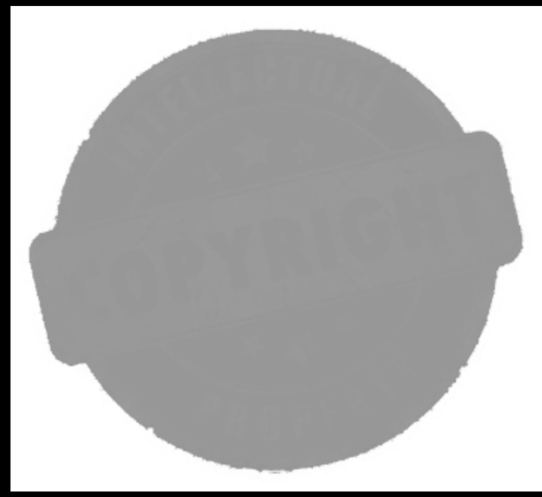

$$J(p) = W(p) \cdot \alpha(p) + (1 - \alpha(p)) \cdot I(p)$$

Under-determined — (W, α, I) are unknowns, single constraint

Compared to natural image matting:

- All pixels are either background or mixed
- Low opacity
- No user input

Formation Model

For a watermarked image collection:

<u>Still under-determined!</u>
3K equations
3K + 3 + 1 unknowns

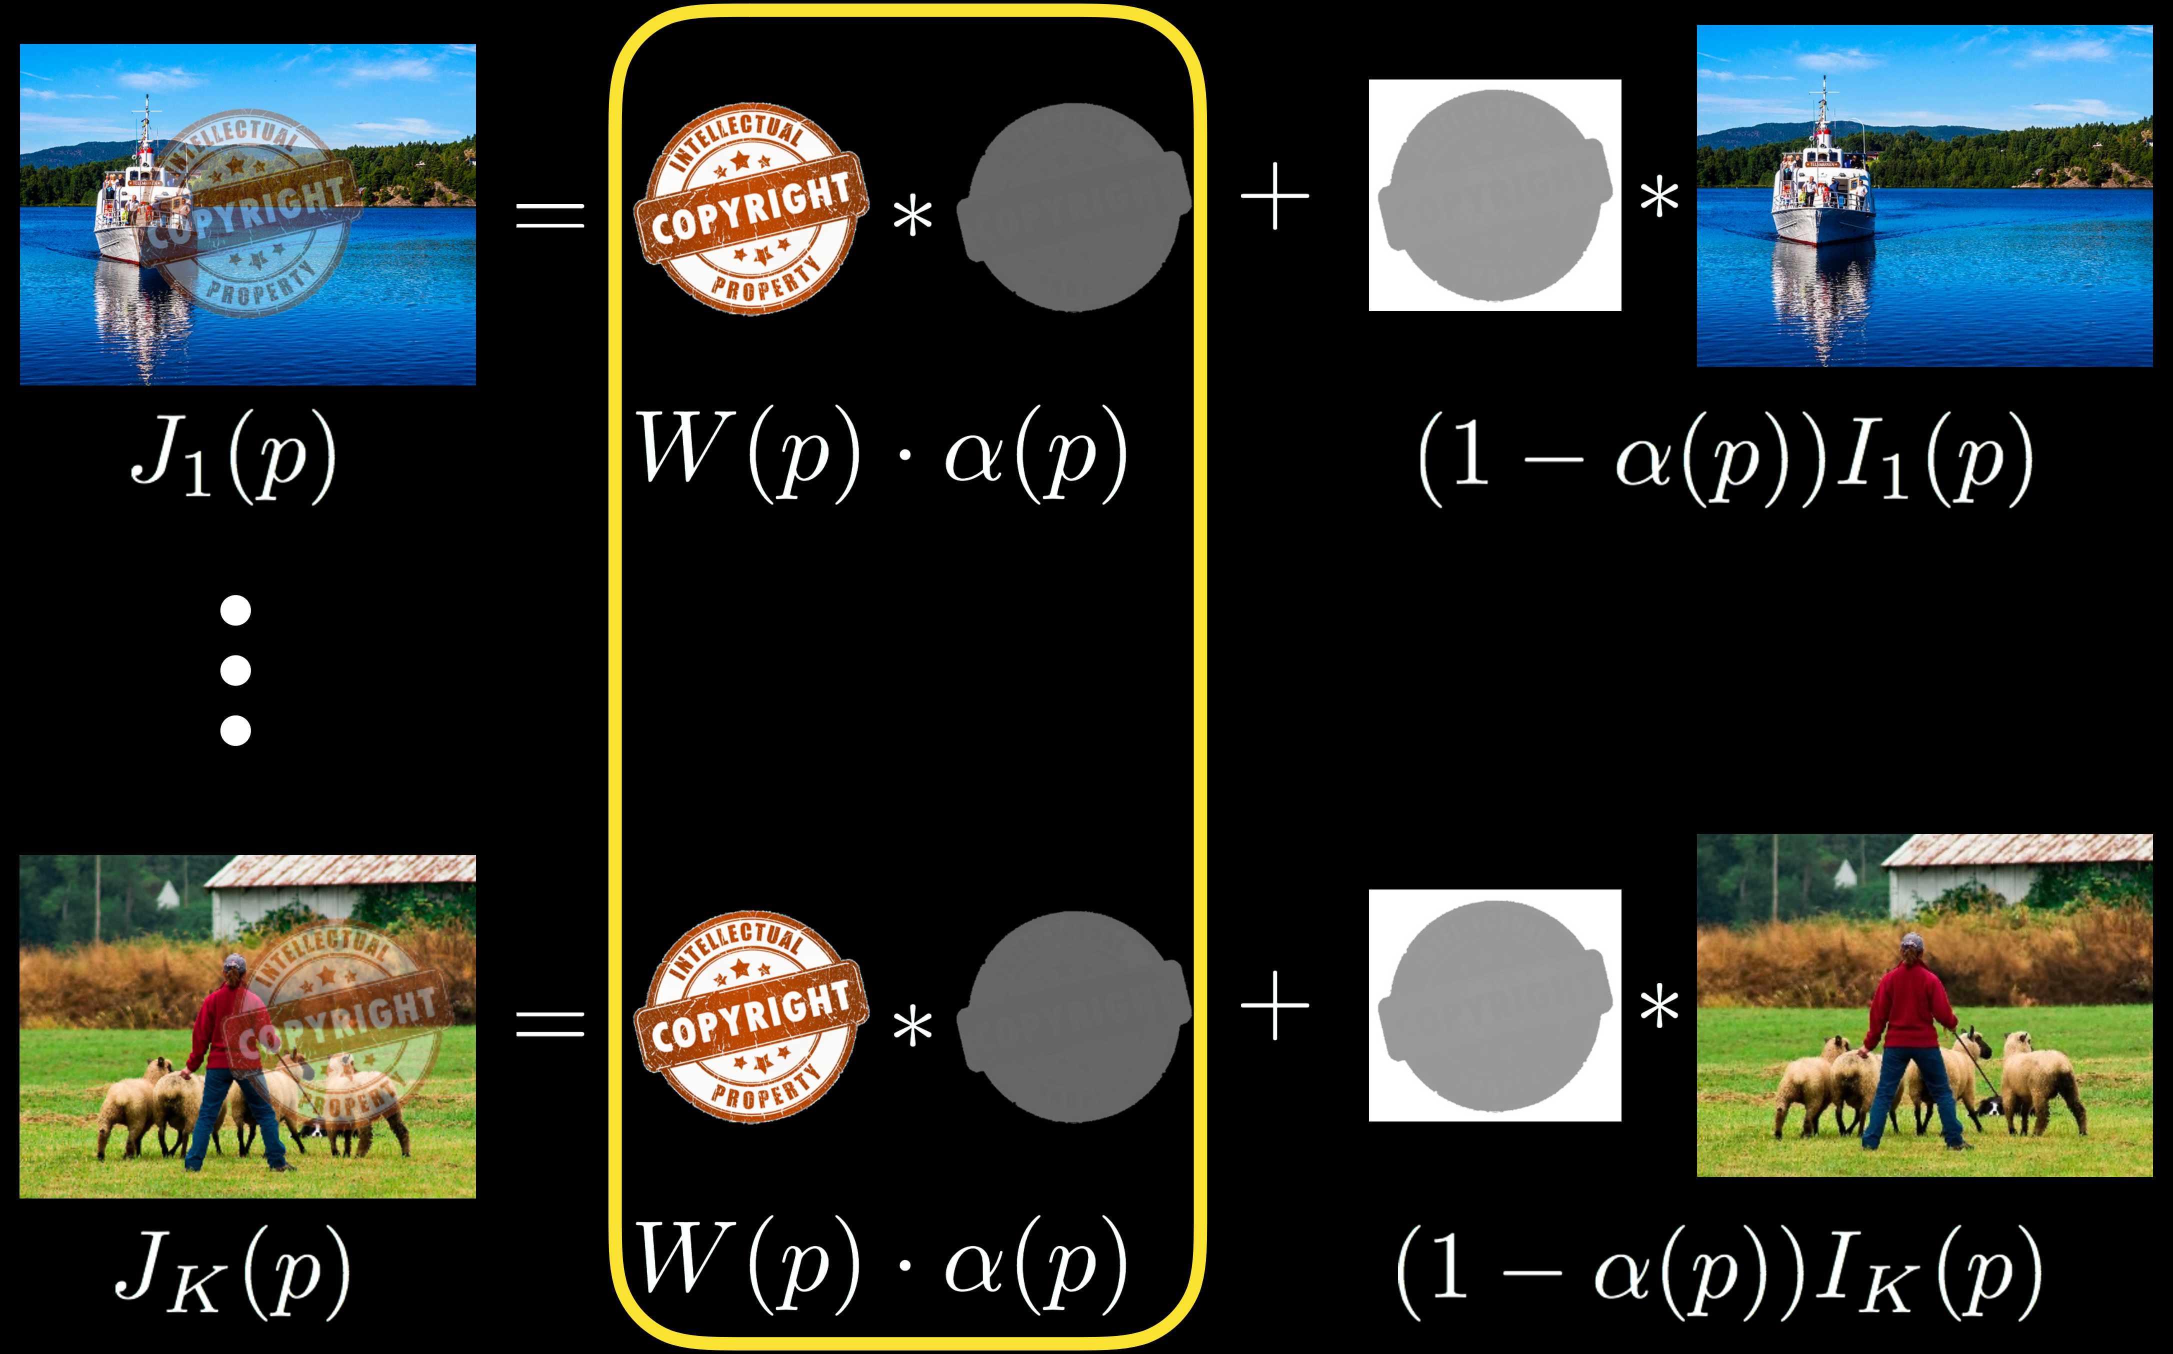

$$J_k = \alpha W + (1-\alpha)I_k, \quad k = 1, \dots, K$$

Initial Watermark Estimation

- Identify which image structures are repeating in the collection
- Assume for now consistent watermark (position, opacity, geometry)
- Estimate the median of gradients, apply Poisson Reconstruction

Image 1

Initial Estimation of Watermark

Direct Image Reconstruction

$$I(p) = \frac{J(p) - \alpha(p) W(p)}{1 - \alpha(p)}$$

Estimated, matted (!) watermark, αW

Input

Direct Subtraction

- Every little error in the W or alpha matte shows up as visual artifacts

Multi-Image Matting

- Formulate the inversion problem as a multi-image matting problem
- Solve jointly for $(W, \alpha, \{I_k\})$

$$+\beta E_f(\nabla(\alpha W)) + \lambda_w E_{\text{reg}}(\nabla W) + \lambda_\alpha E_{\text{reg}}(\nabla \alpha)$$

Fidelity Term

Standard Image piece wise smooth prior

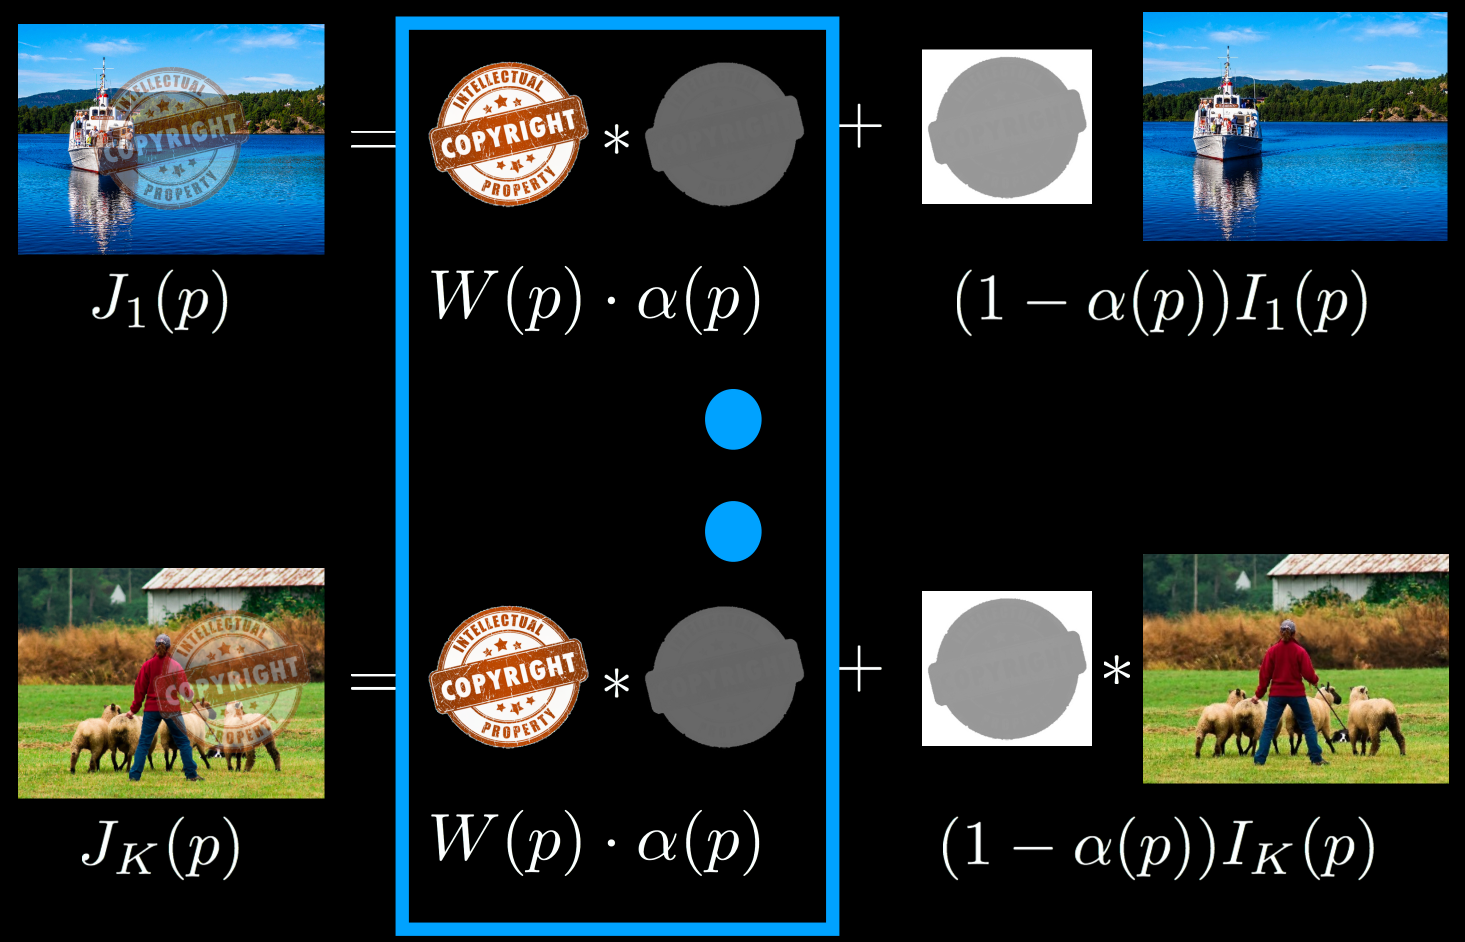

Same W, same alpha

Initial watermark grads

Optimization

- Non linear, many unknowns
- Alternating minimization — divide the problem into simple subproblems

Input: Watermarked images {J_k}, initial matted watermark
Output: W, alpha matte, watermark free images {I_k}

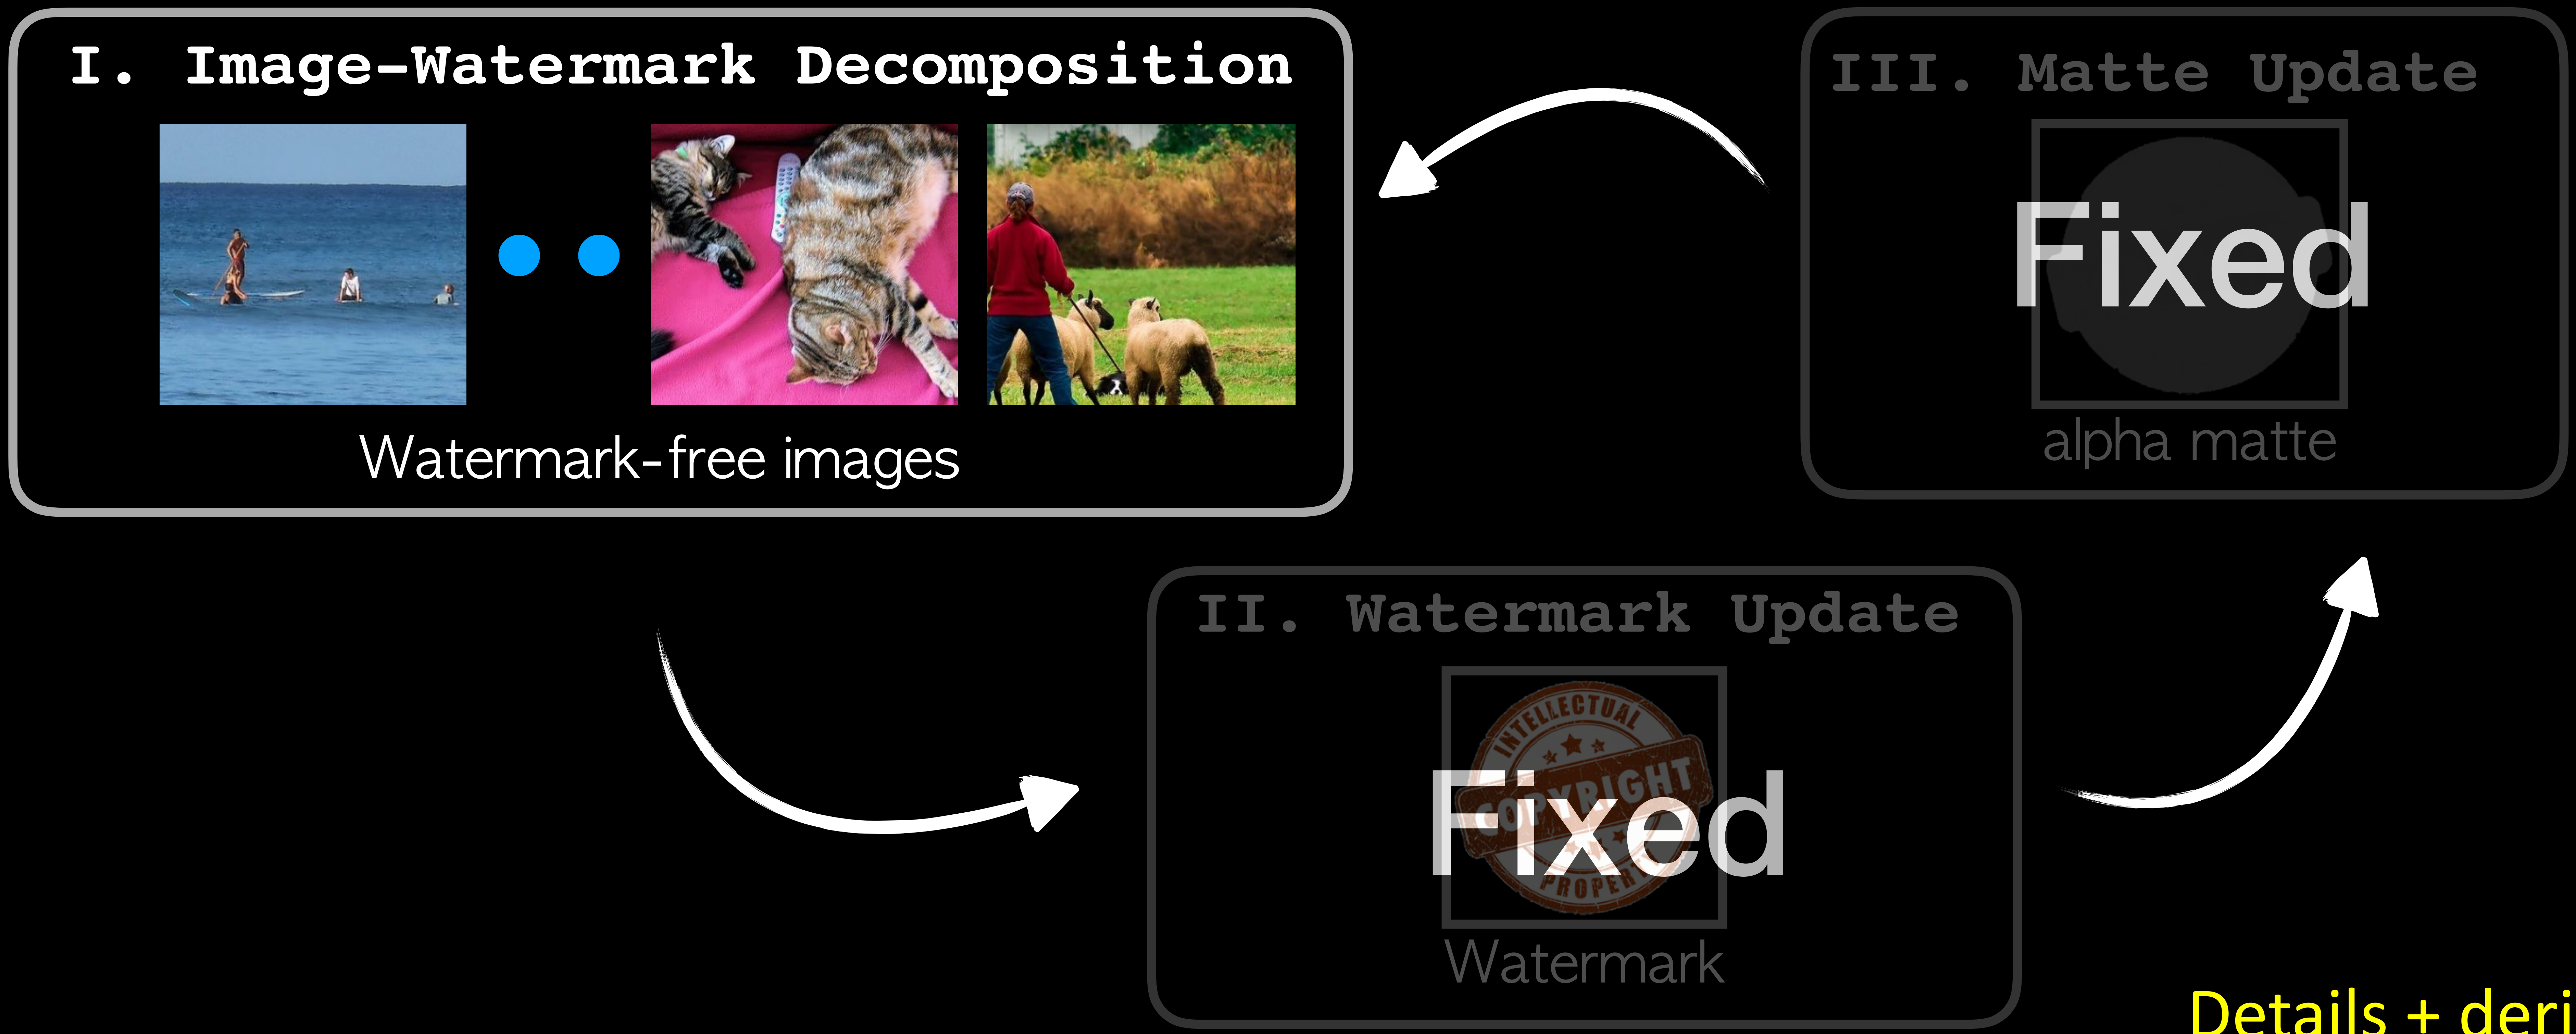

Details + derivation in the paper!

Optimization

- Non linear, many unknowns
- Alternating minimization — divide the problem into simple subproblems

Input: `Watermarked images {J_k}, initial matted watermark`
Output: `W, alpha matte, watermark free images {I_k}`

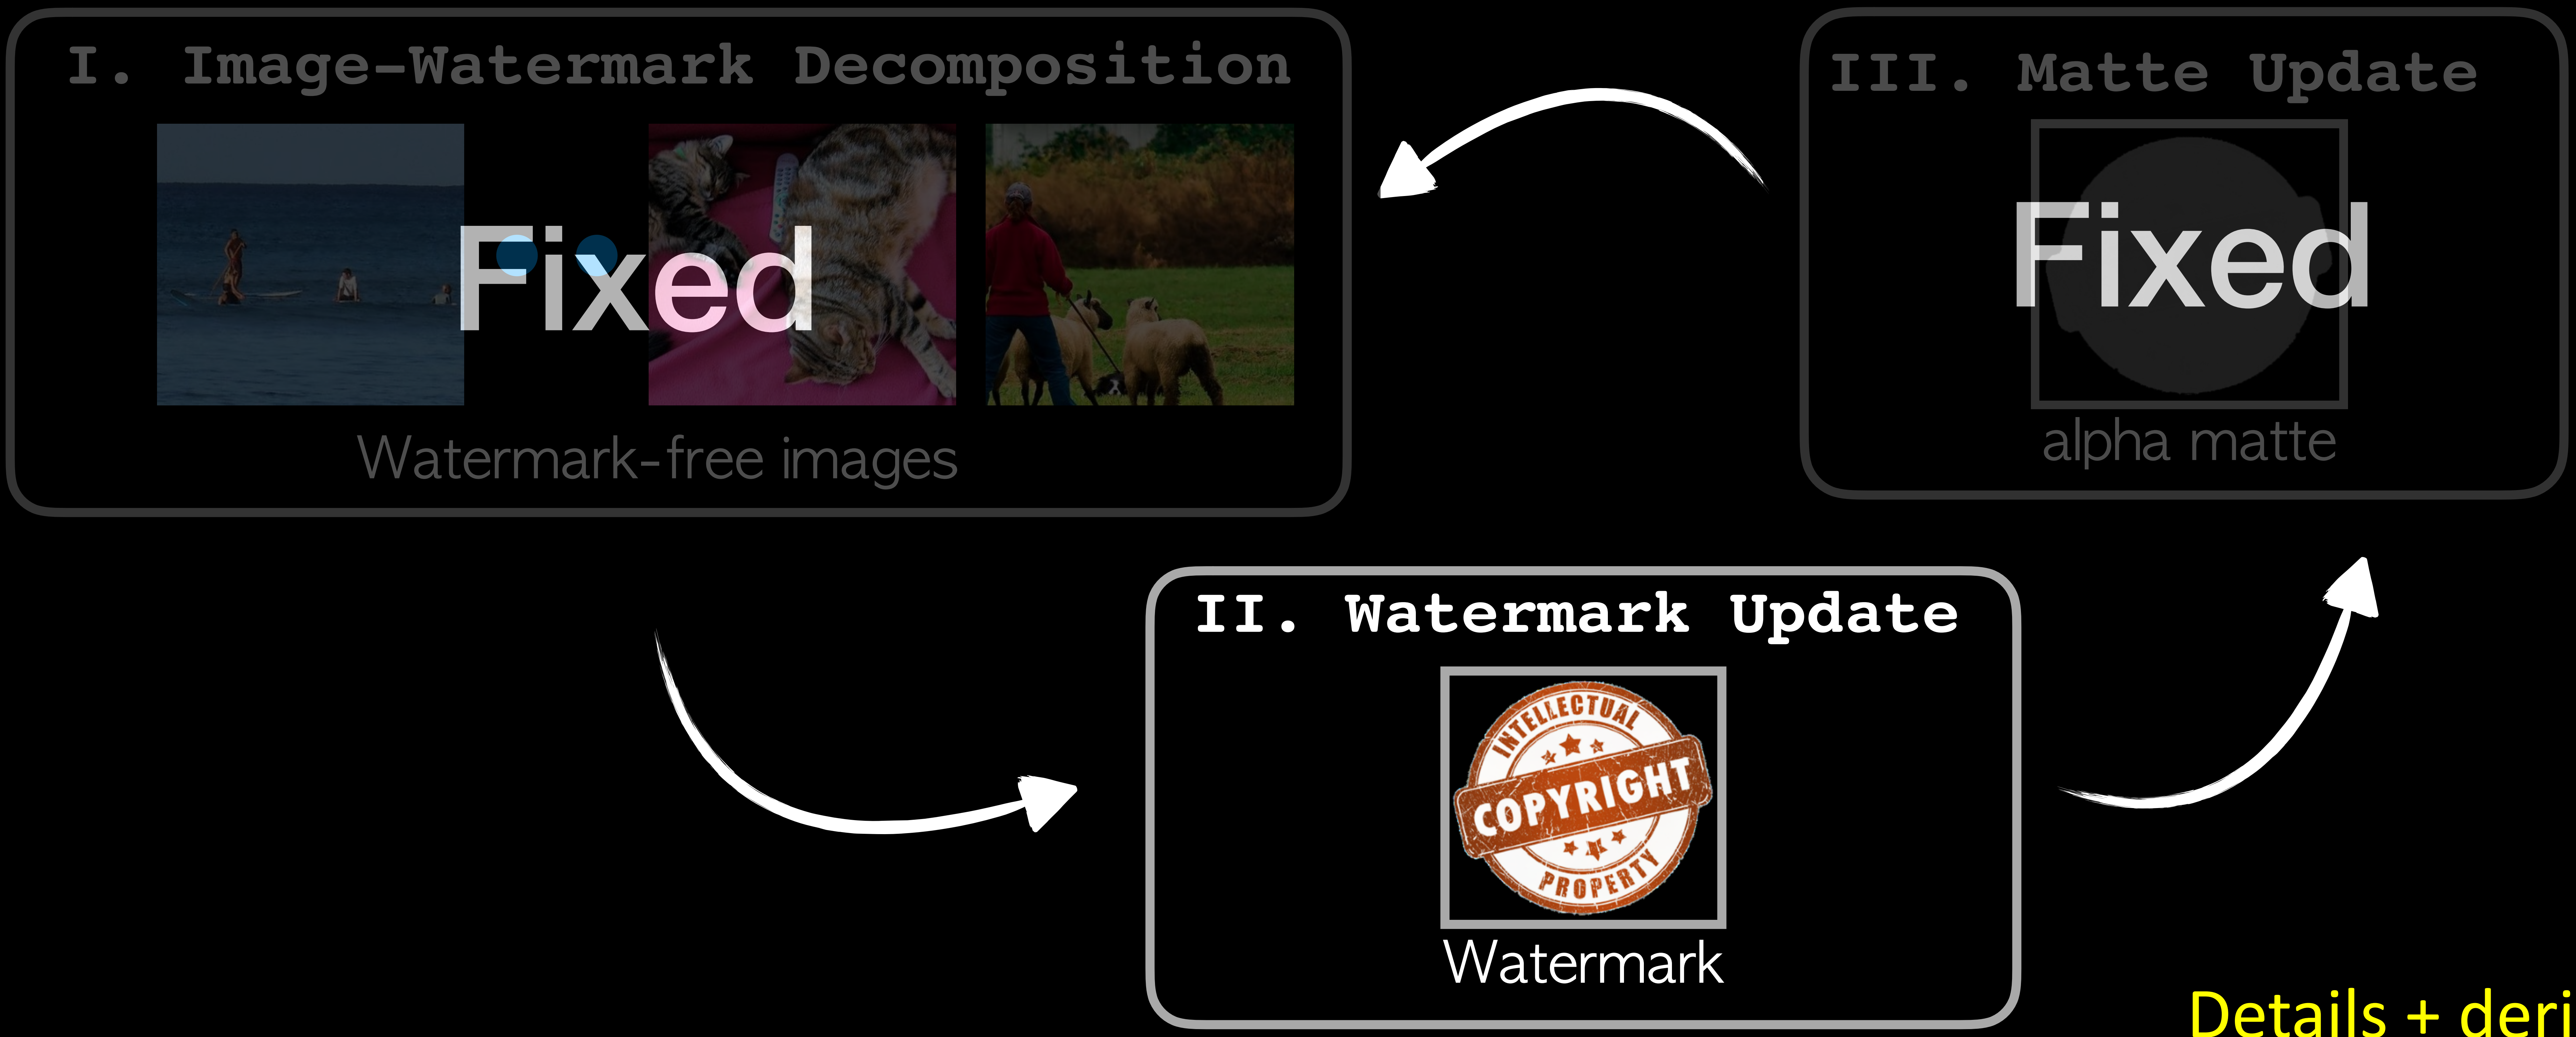

Details + derivation in the paper!

Optimization

- Non linear, many unknowns
- Alternating minimization — divide the problem into simple subproblems

Input: Watermarked images {J_k}, initial matted watermark
Output: W, alpha matte, watermark free images {I_k}

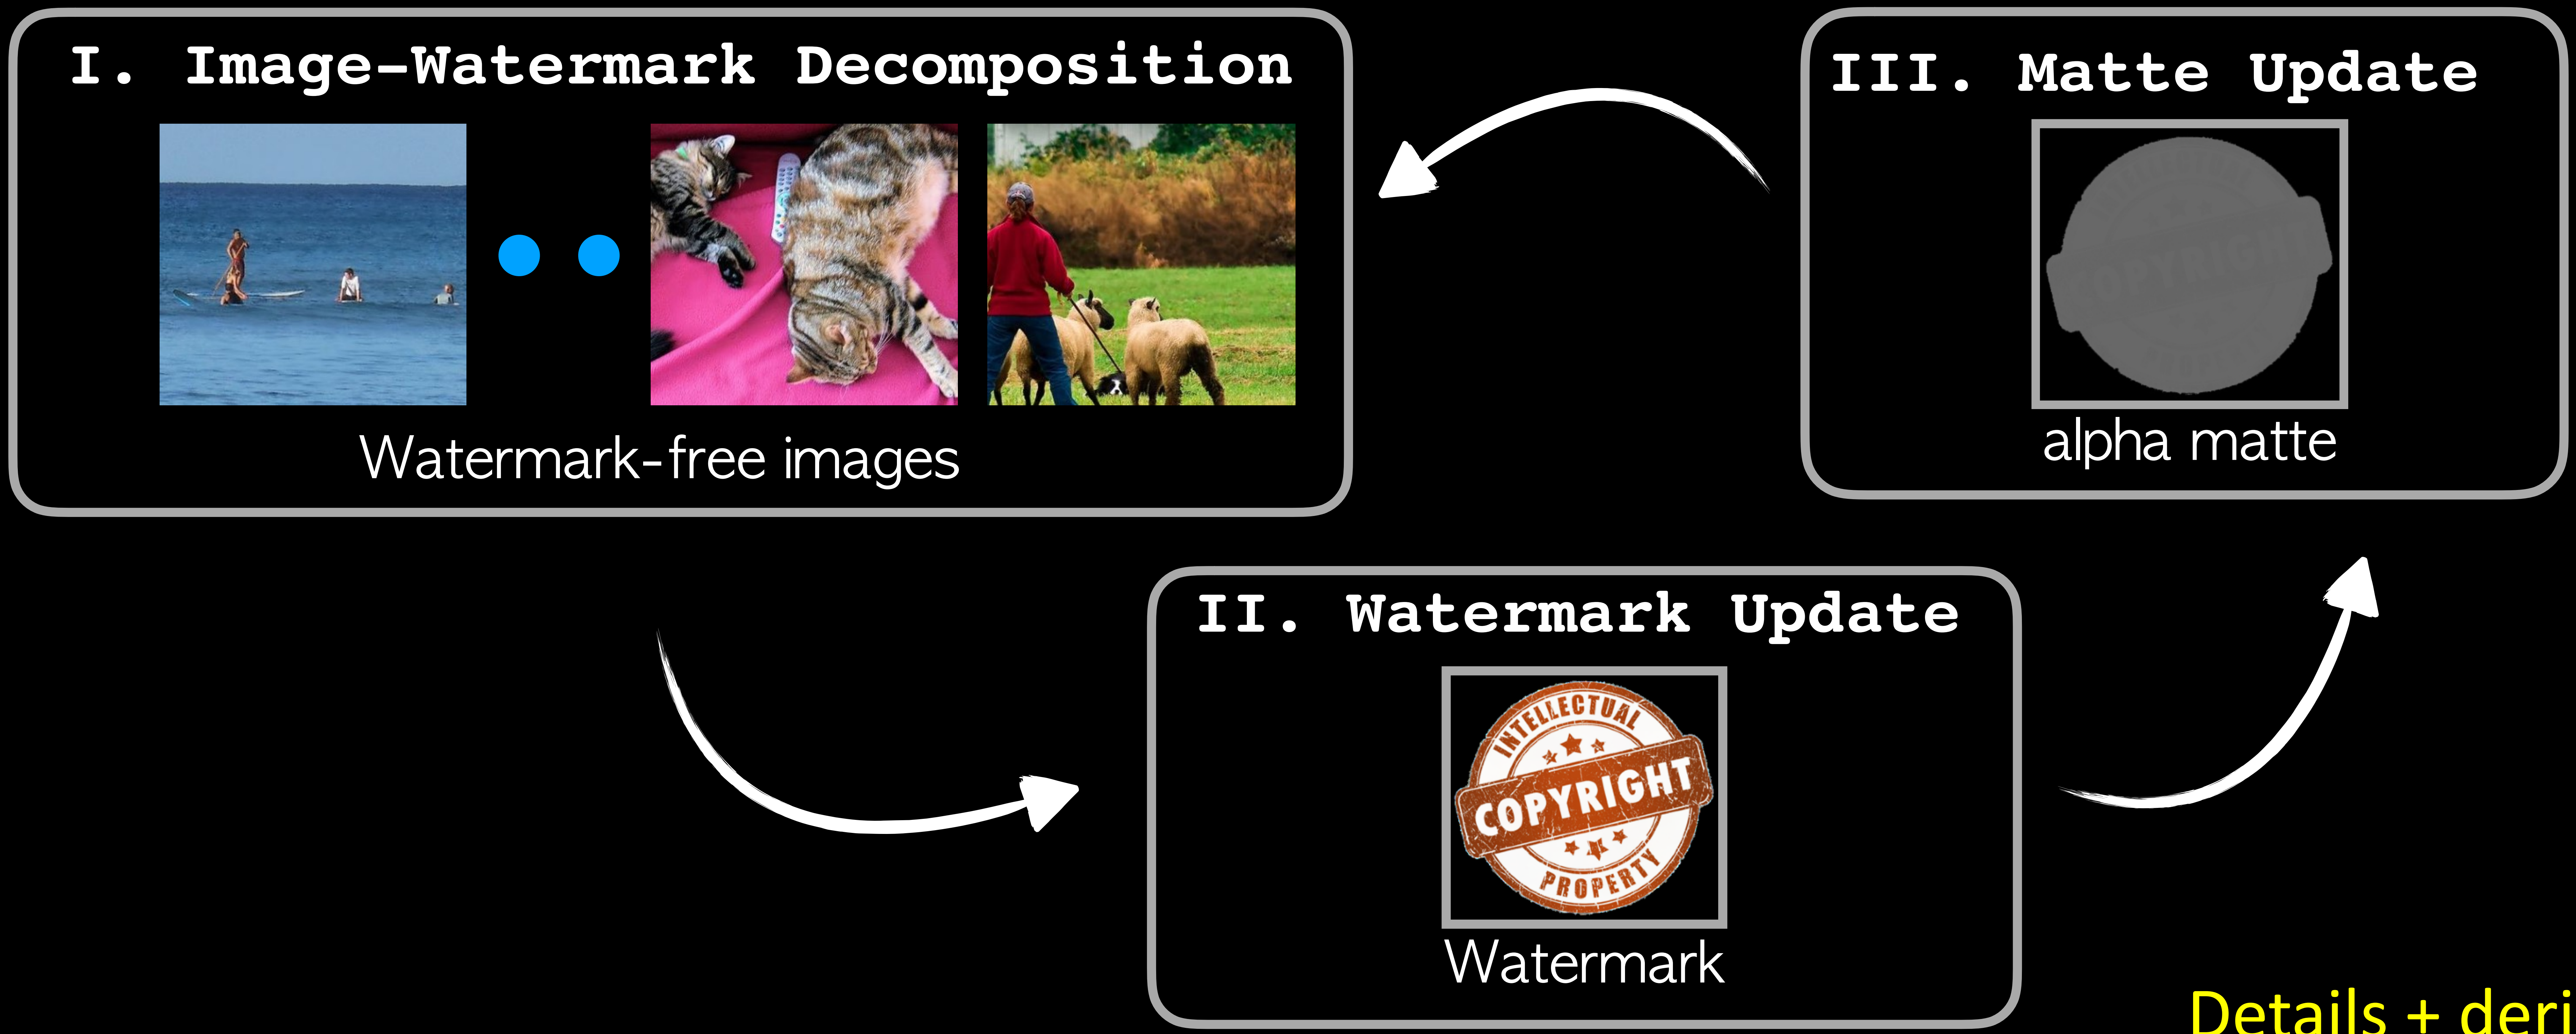

Details + derivation in the paper!

Optimization

Input

Iteration 1

Iteration 2

Iteration 3

Iteration 4

CanStock

Estimated watermark from collection (automatically):

Watermarked image

Watermark removed (automatic)

Adobe Stock

Estimated watermark from collection (automatically):

Watermarked image

Watermark removed (automatic)

Not always perfect… but generally very good

Securing Visible Watermarks

Attack relies on consistency —> break consistency

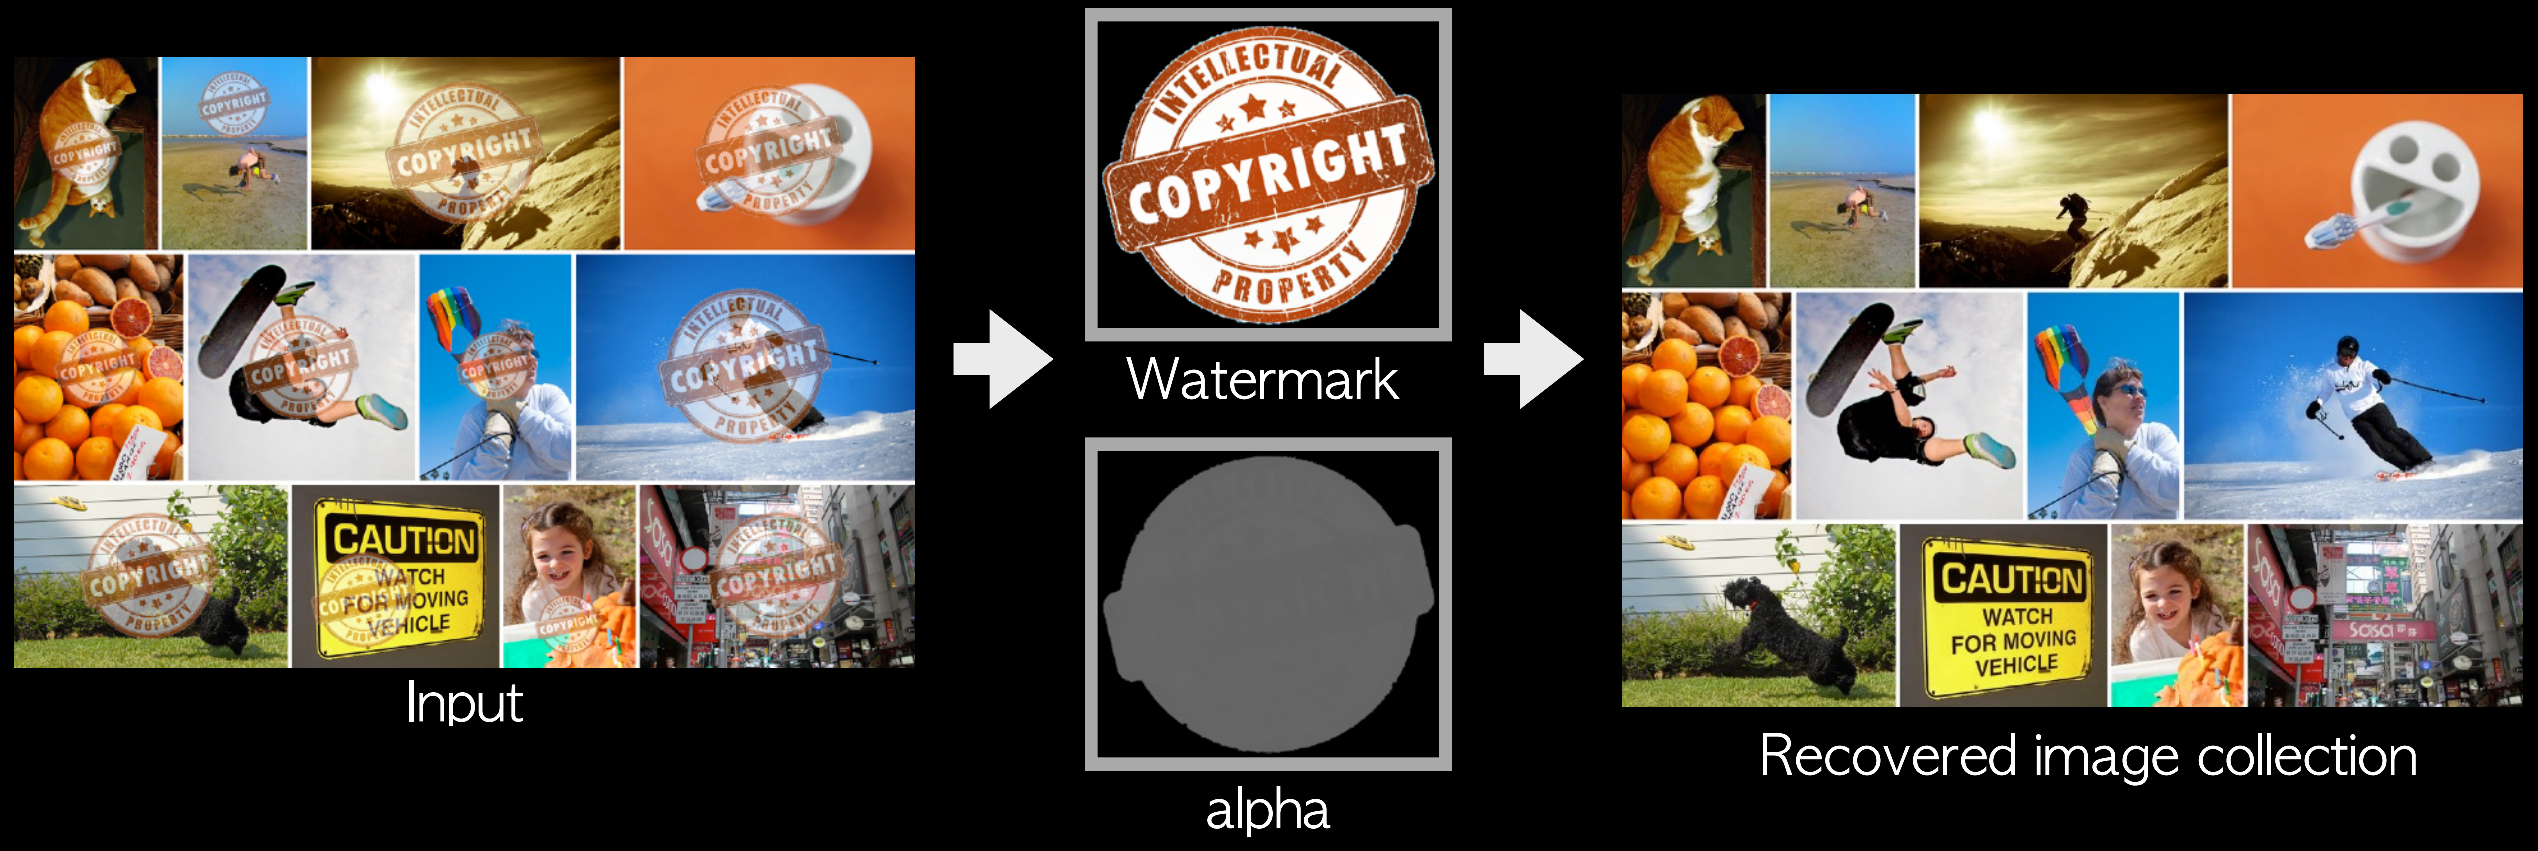

- Rules: need to use substantially the same watermark design for each image.
- Introduce per-image variation:
  - Random location
  - Random opacity
  - Random geometric deformation

Subtle Geometric Perturbation

Generalized Watermarking Model

$$\boxed{J_k = c_k \alpha(\omega_k) W(\omega_k) + (1 - c_k \alpha(\omega_k)) I_k}$$

subtle random spatial perturbation

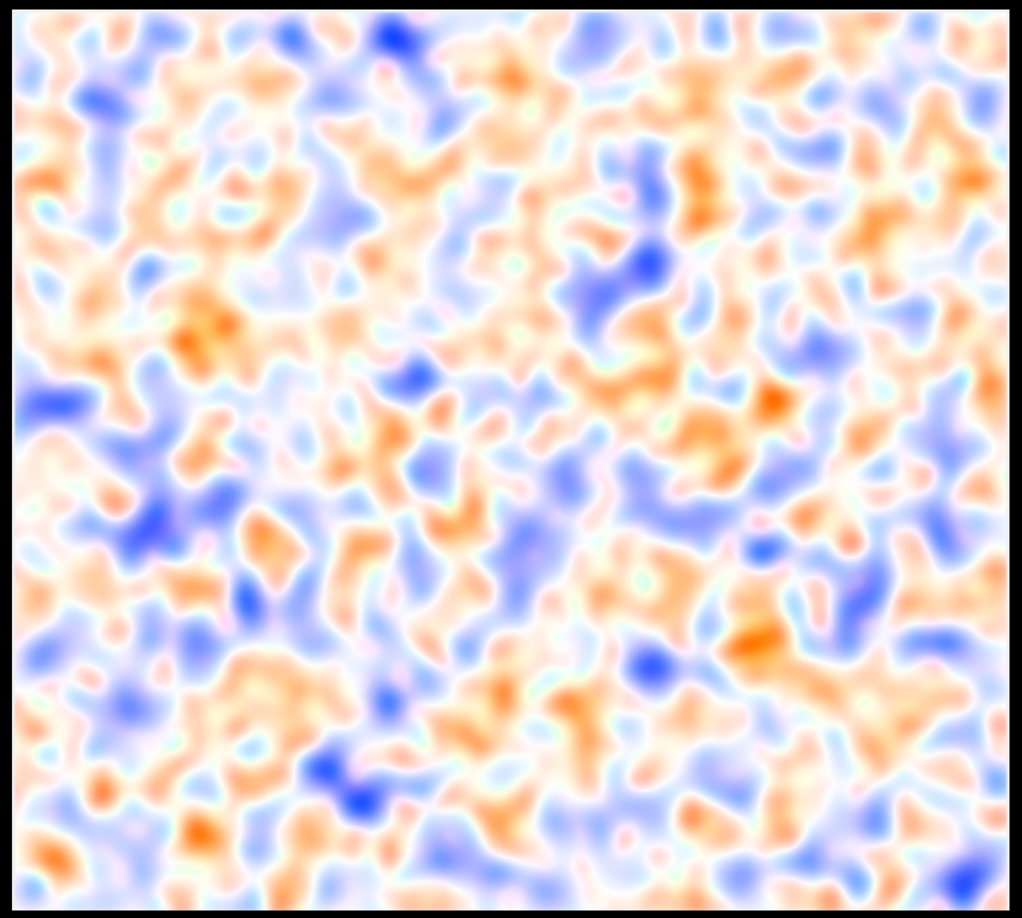

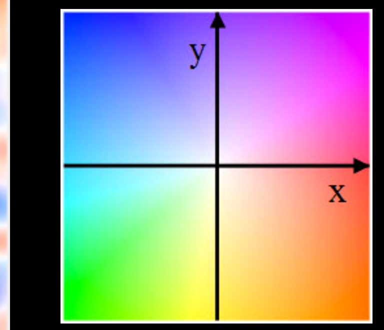

Generated warp field
max deformation of 1px

Original Watermark

Deformed Watermark

Random Spatial Perturbation

Consistent Watermark

Subtle spatial perturbation


Reconstruction


Reconstruction w/o flow estimation

Random Spatial Perturbation

Reconstruction w/o flow estimation

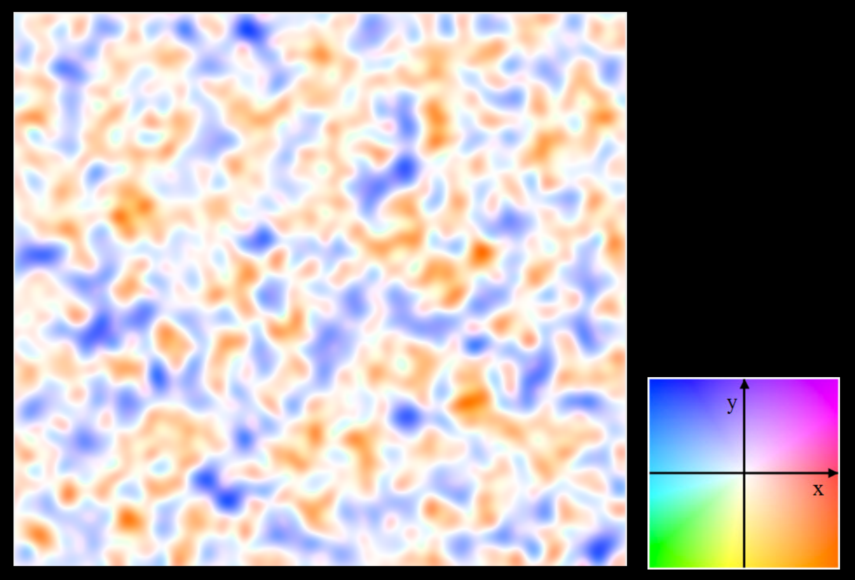

Generated warp field

Subtle Geometric Perturbation

Generalized Watermarking Model

$$\boxed{J_k = c_k \alpha_k(\omega_k) W(\omega_k) + (1 - c_k \alpha_k(\omega_k)) I_k}$$

subtle random spatial perturbation

I. Image-Watermark Decomposition

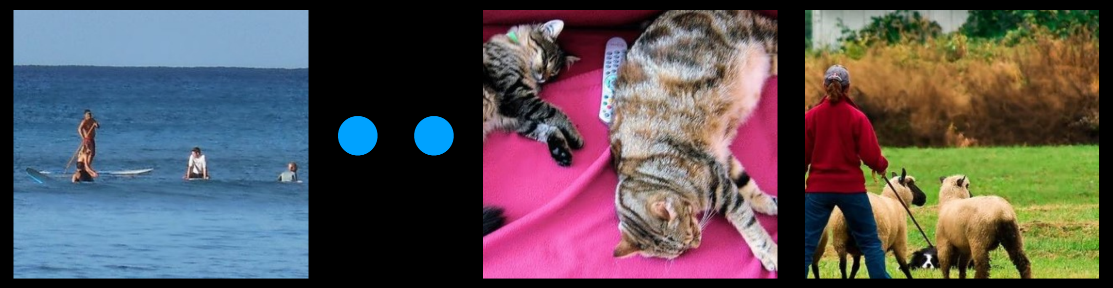

II. Watermark Update

III. Matte Update (global)

IV. Opacity Estimation (per Image)

V. Flow Estimation (per Image)

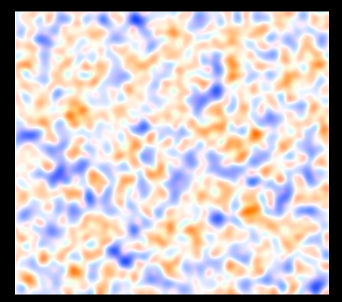

Random Spatial Perturbation

Reconstruction w/o flow estimation

Random Spatial Perturbation

Reconstruction w/ flow estimation

Random Spatial Perturbation

Reconstruction w/o flow estimation

Random Spatial Perturbation

Reconstruction w/ flow estimation

Deployed! (Shutterstock, >150M images)

https://image.shutterstock.com/z/stock-photo--years-old-boy-walking-on-beach-wooden-pier-ready-to-swim-in-a-sea-summer-outdoor-activities-with-382879477.jpg

- **Watermarks as used today are breakable because of their consistency across image collection**
- **Study how robust the attack is per image inconsistencies**
- **Takeaway message:**

Watermarks should be designed to also be hard to remove from
Image collections

Tali Dekel

Miki Rubinstein

Ce Liu

Smart, Sparse Contours to Represent and Edit Images

Tali Dekel, Chung Gan, Dilip Krishnan, Ce Liu, Bill Freeman, CVPR18

Tali Dekel

Dilip Krishnan

Ce Liu

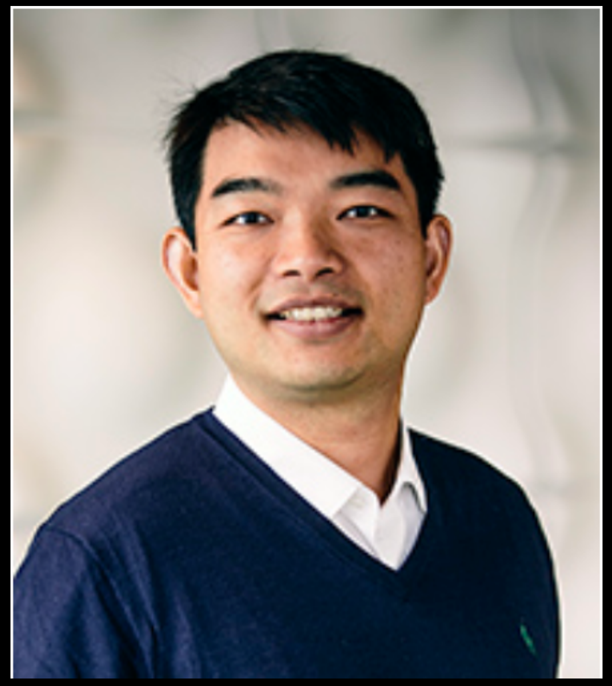

Chung Gan

Research at Google

Image Editing in the Contour Domain

James H. Elder
Department of Psychology

Rick M. Goldberg
Department of Computer Science

Centre for Vision Research
Human Performance Laboratory, CRESTech
York University, Toronto, Canada M3J 1P3

IEEE PAMI March 2001

Motivation


Original image

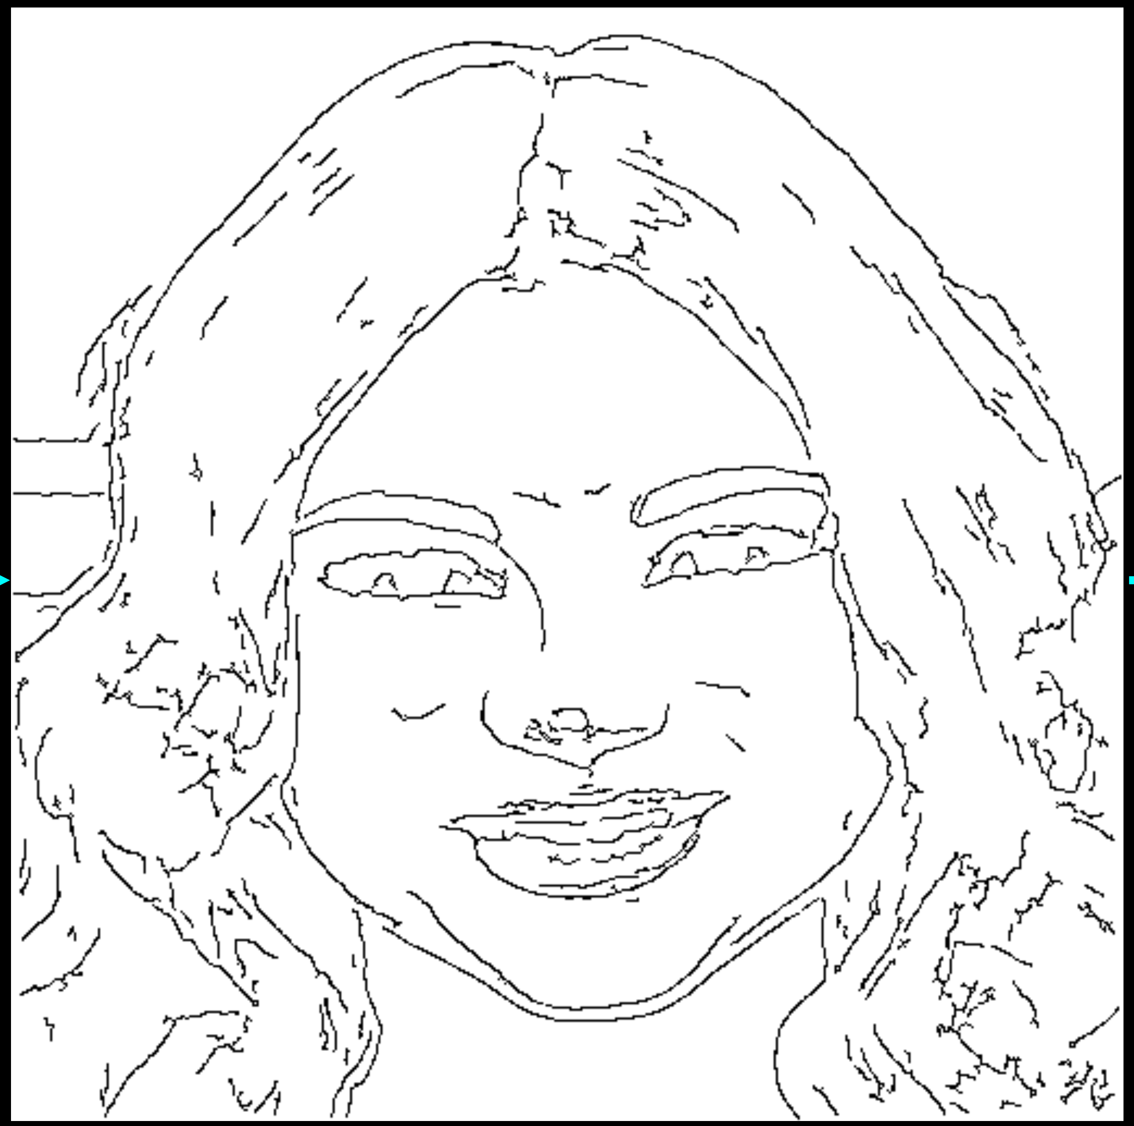
Extracted contours

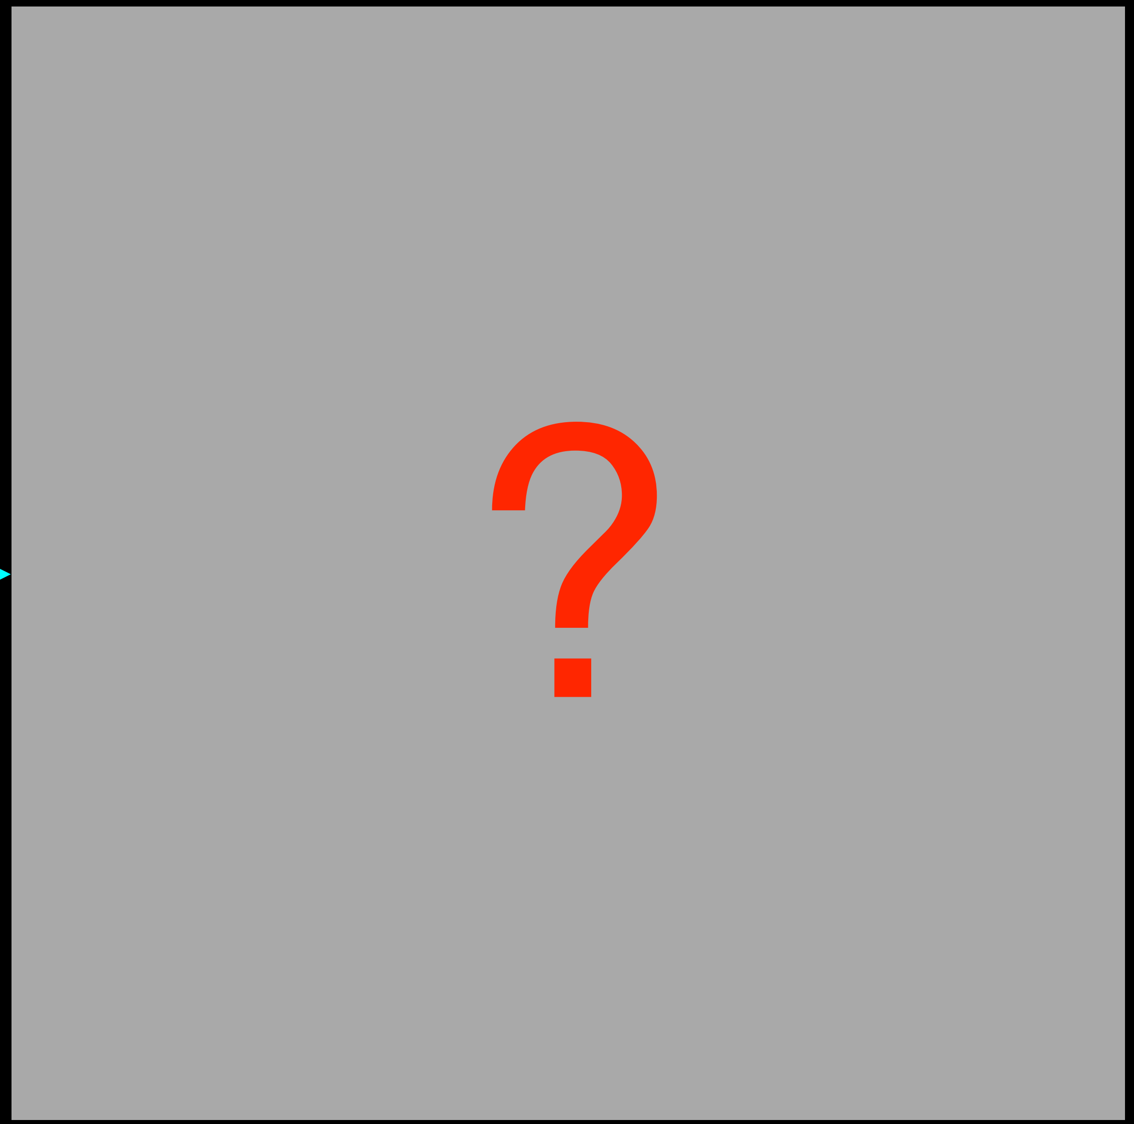

Reconstruction

Motivation

Original image | Extracted contours | Contours + gradient | Our reconstruction

Compare with PDE approach

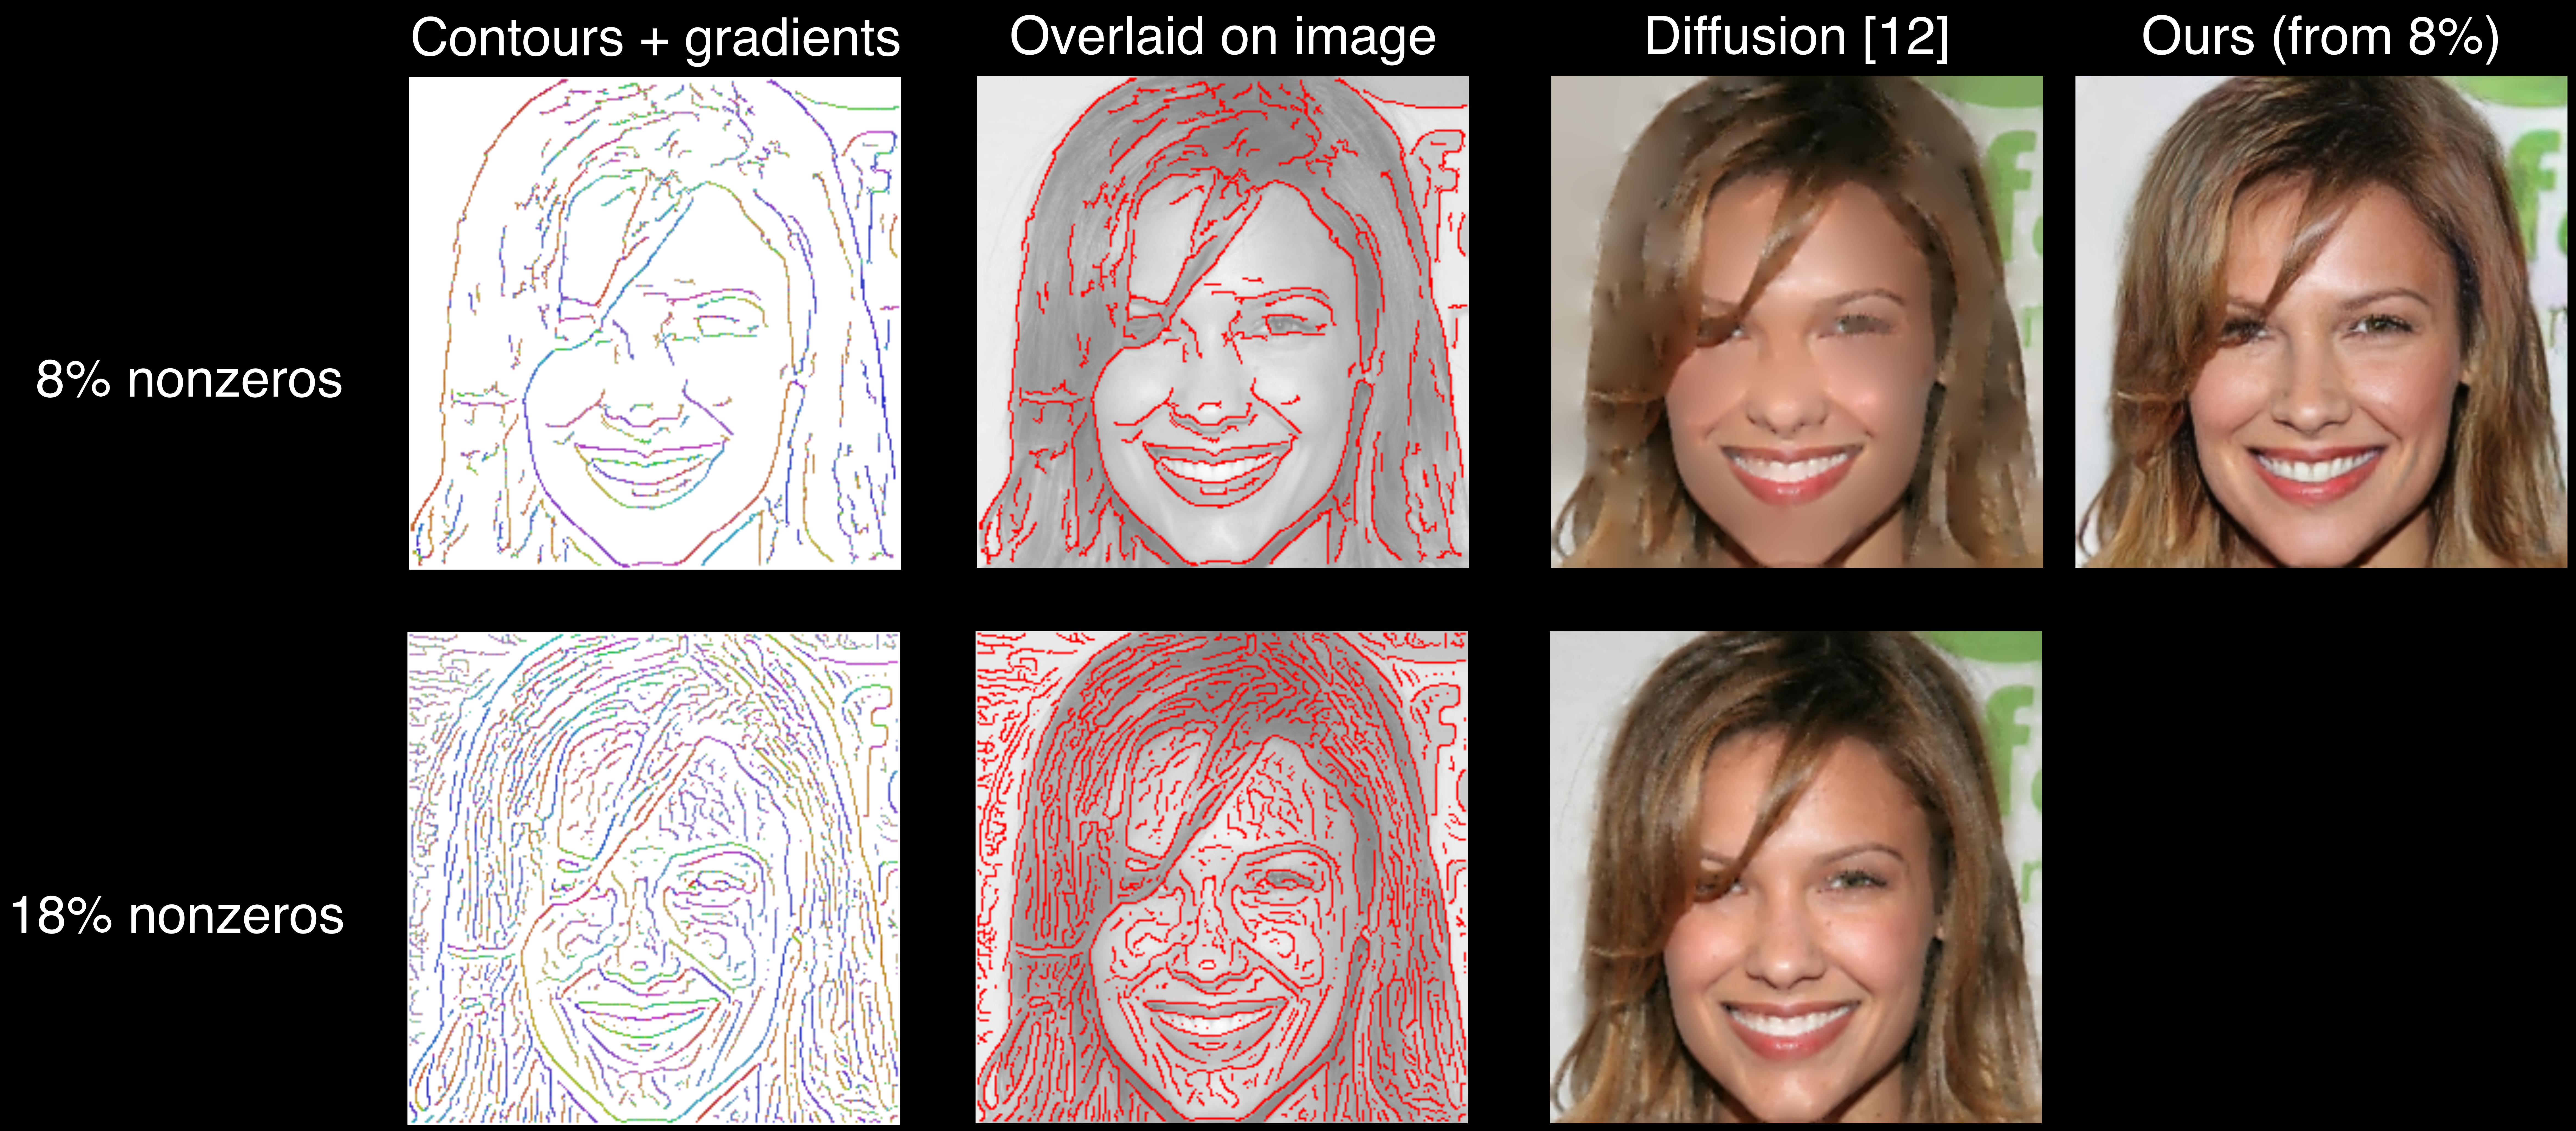

Close-up comparison

Original image

18% nonzeros

Diffusion [12]

Our reconstruction

8% nonzeros

Compare with Pix2Pix

Original image

Binary contours

Pix2pix [18]

Our reconstruction

Contour + gradient

System pipeline

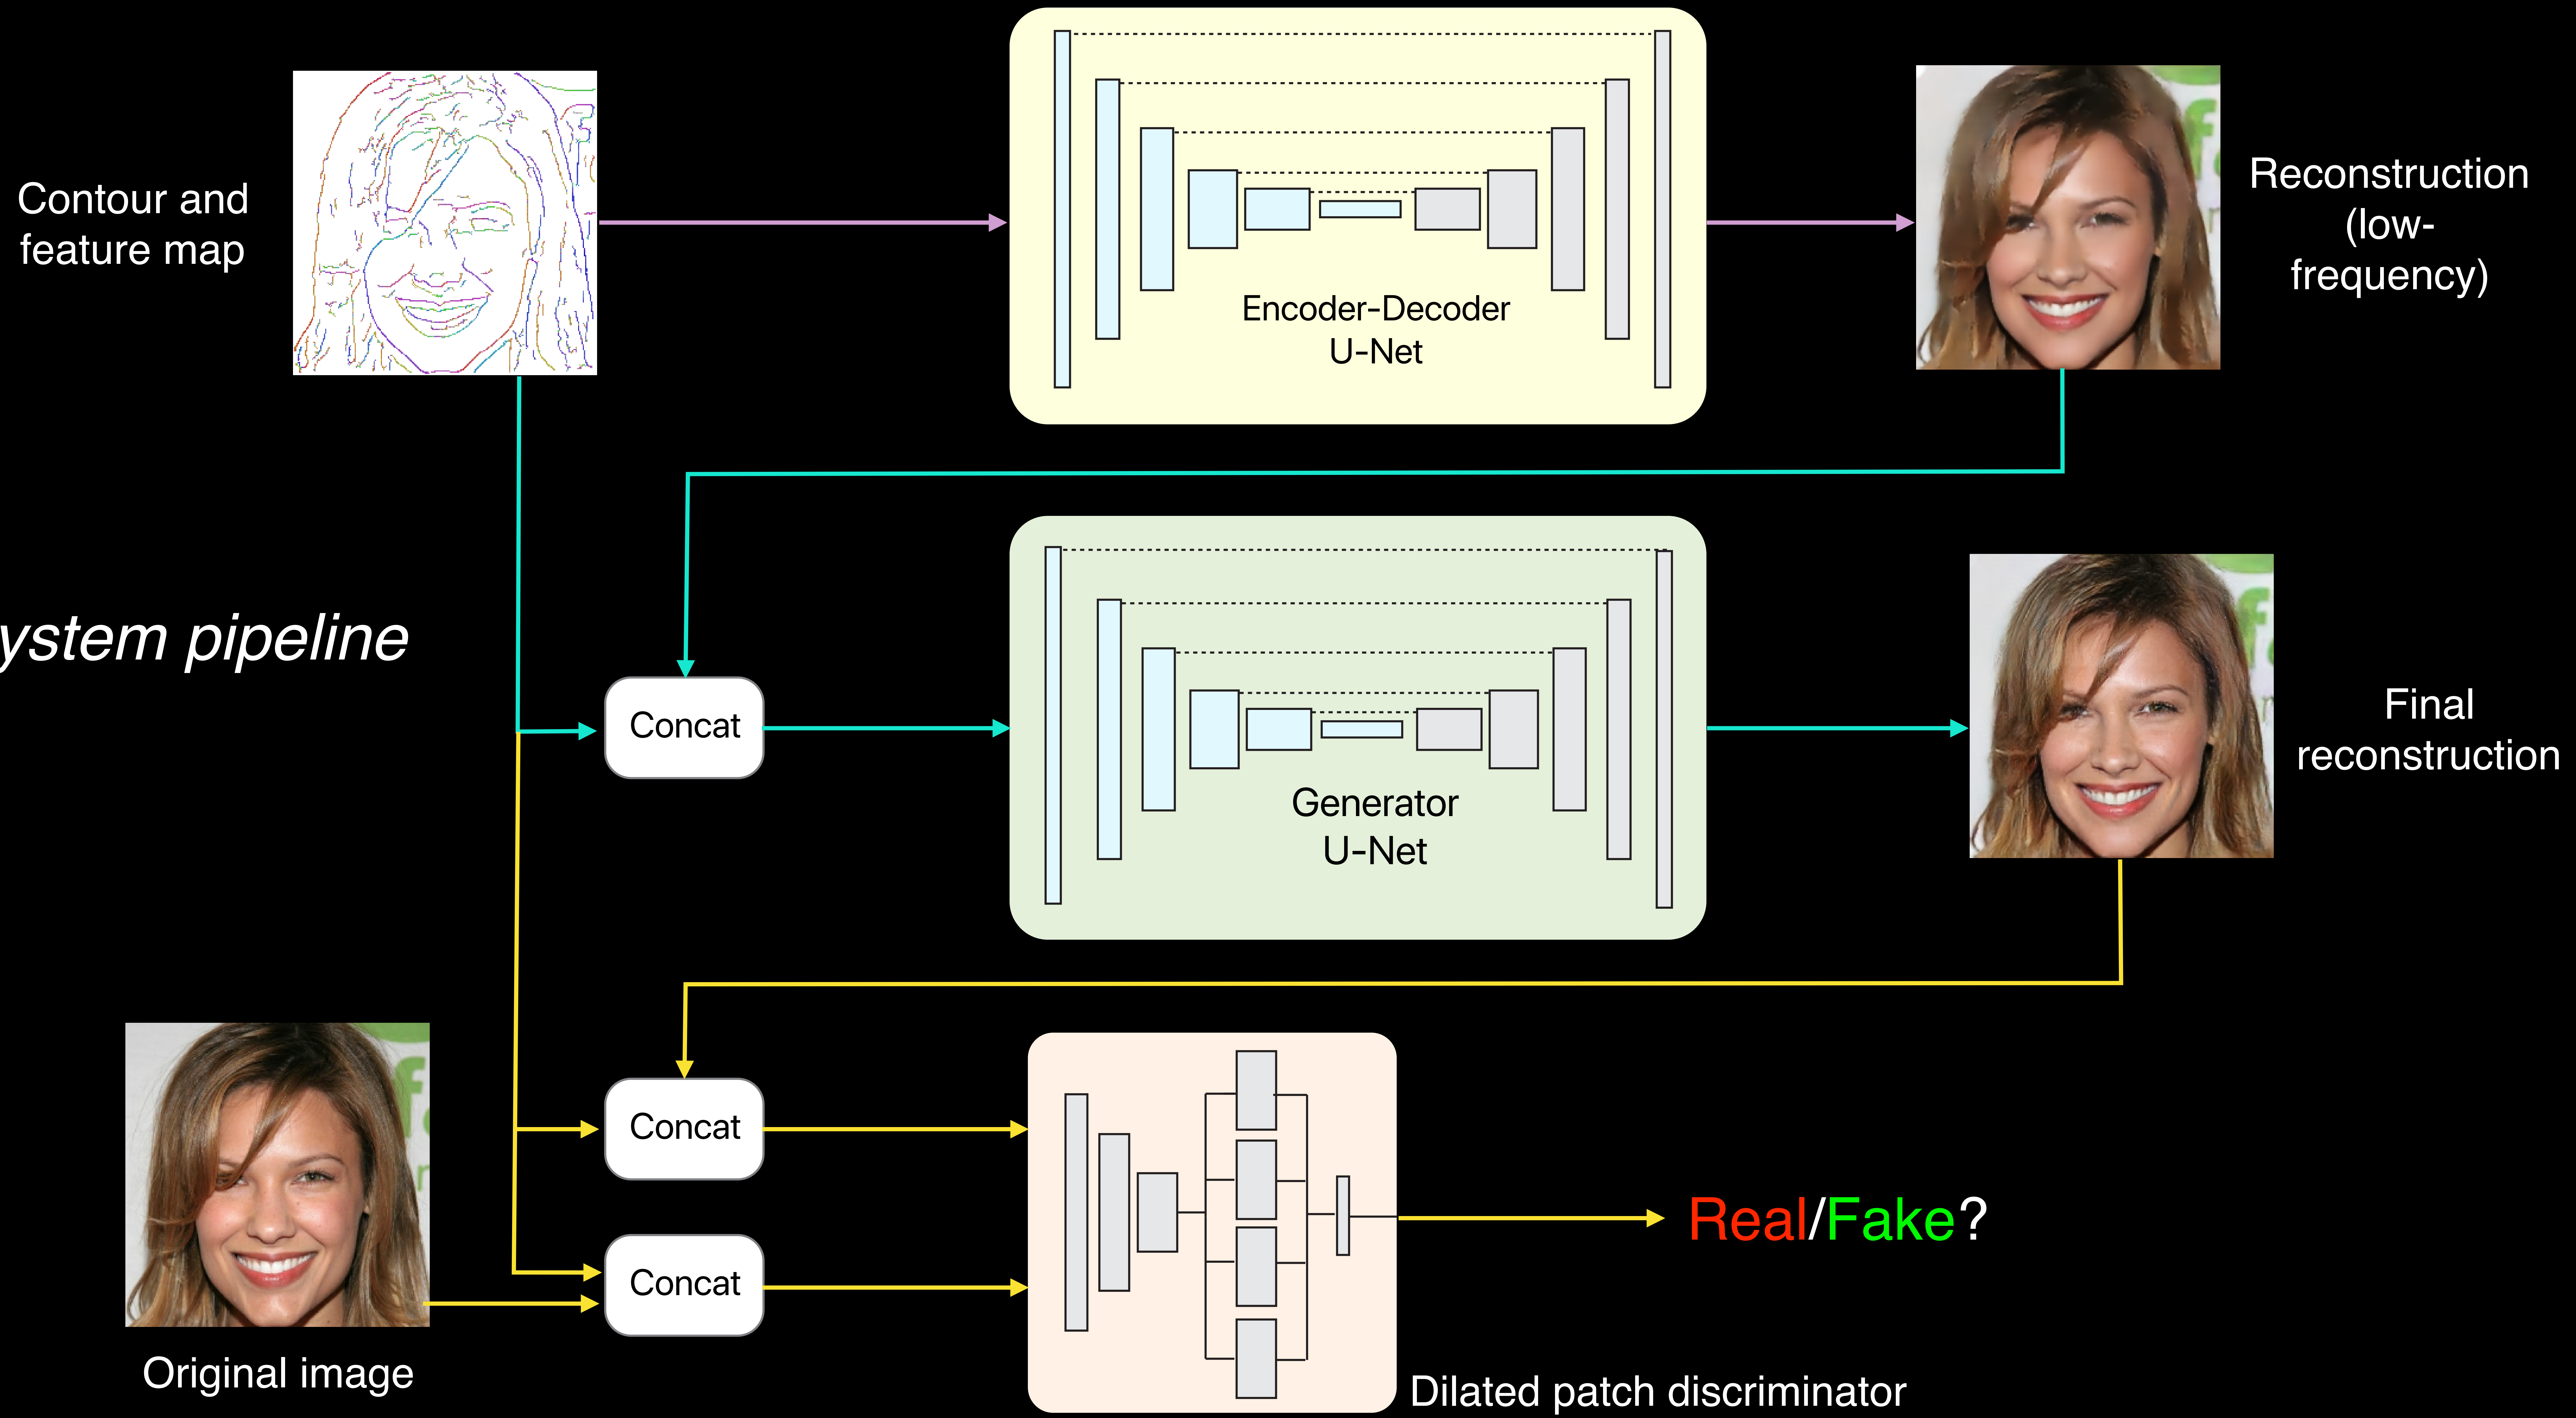

Reconstruction as a Function of Sparsity Level

Original image

1% 2% 3% 5% 10%

Input contours (overlaid on image)

Reconstruction

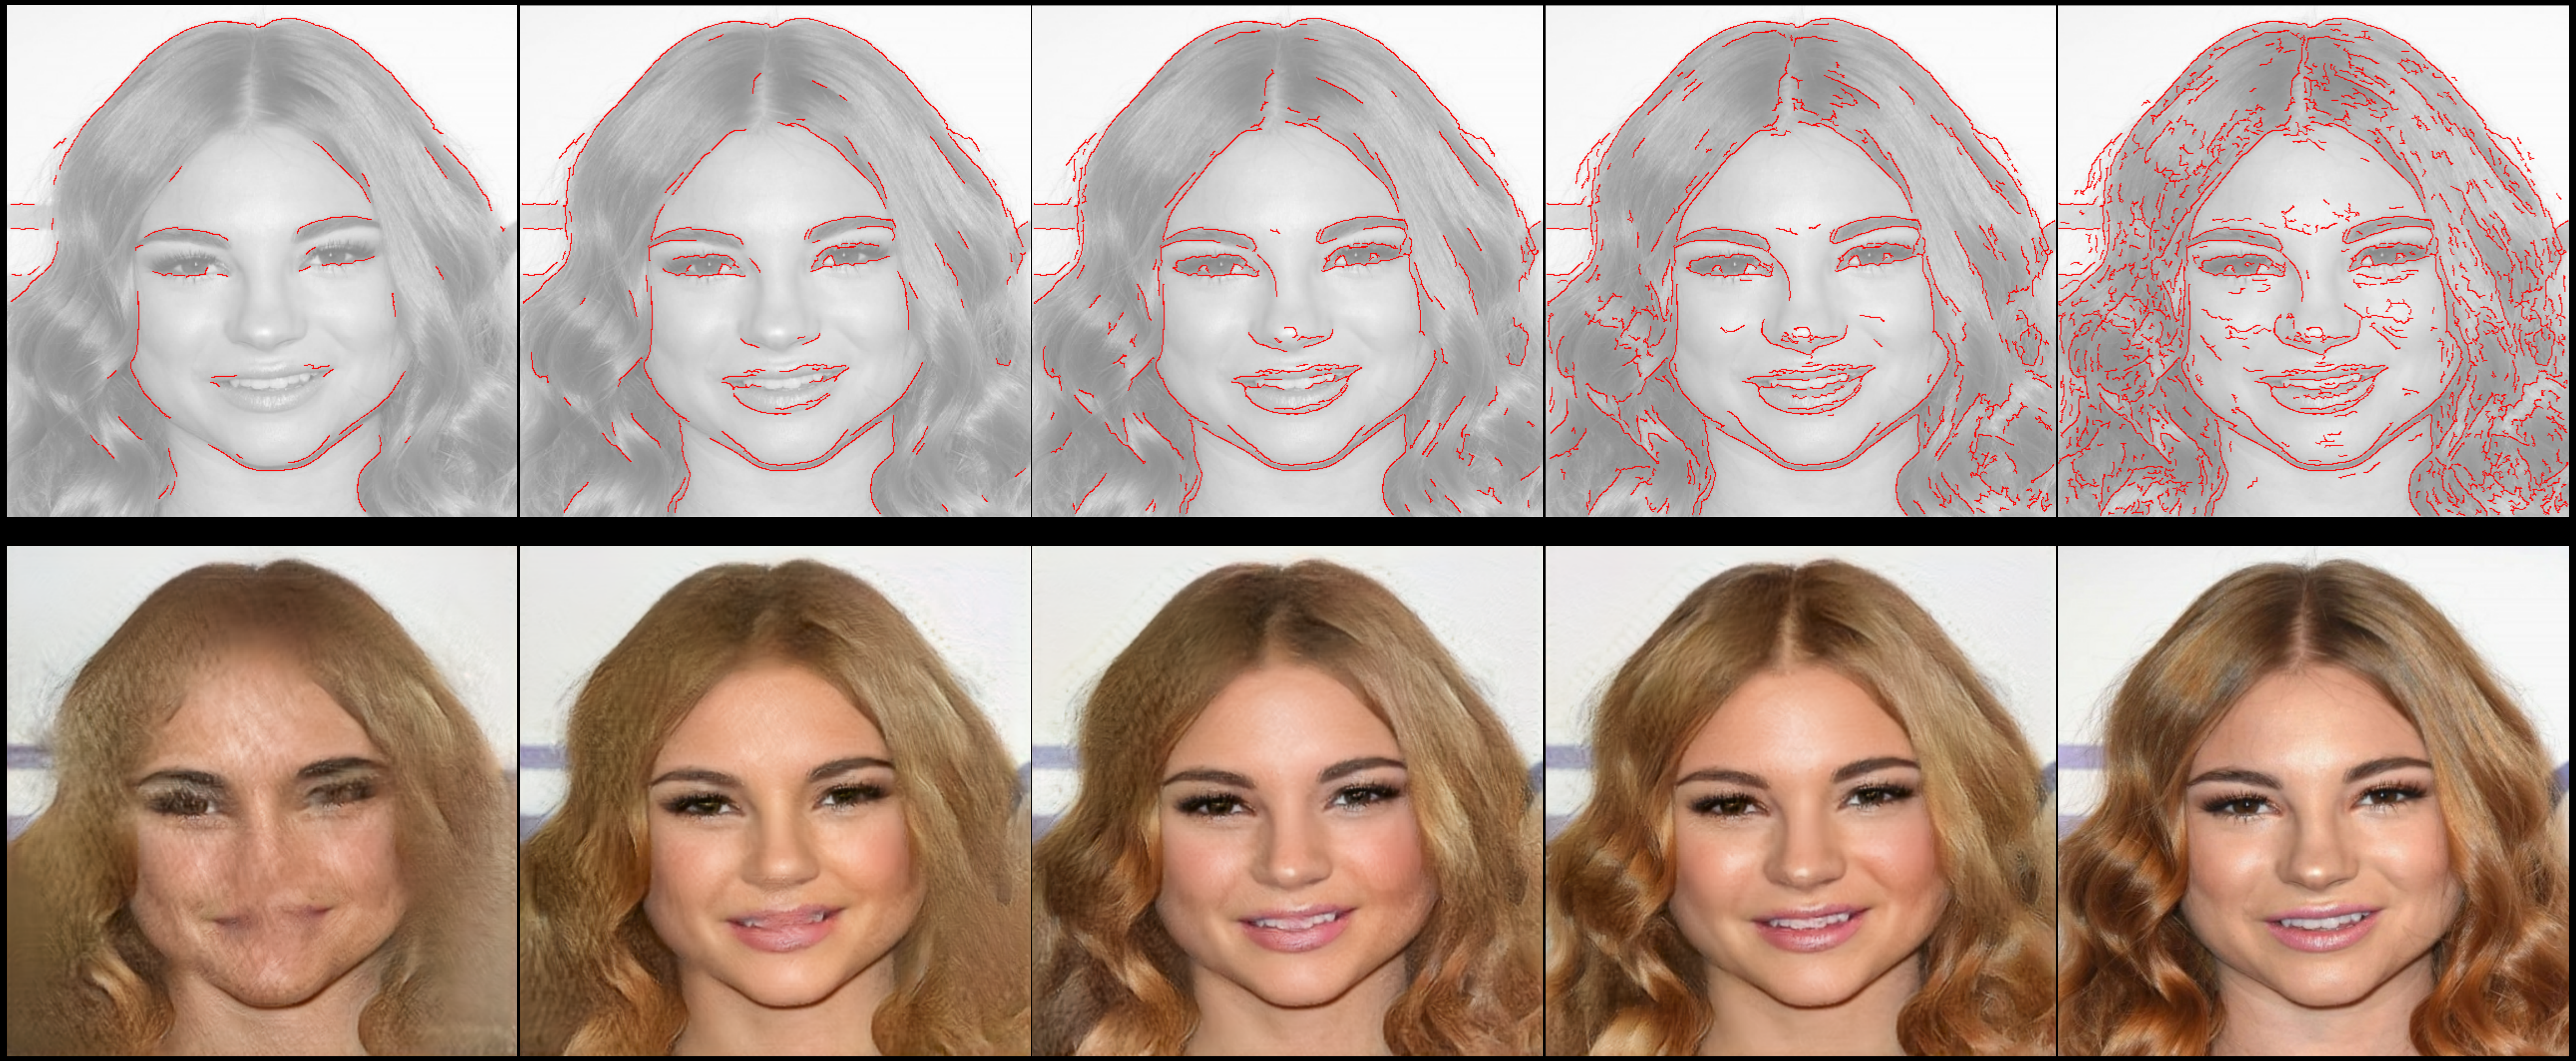

Animation of the Reconstructions

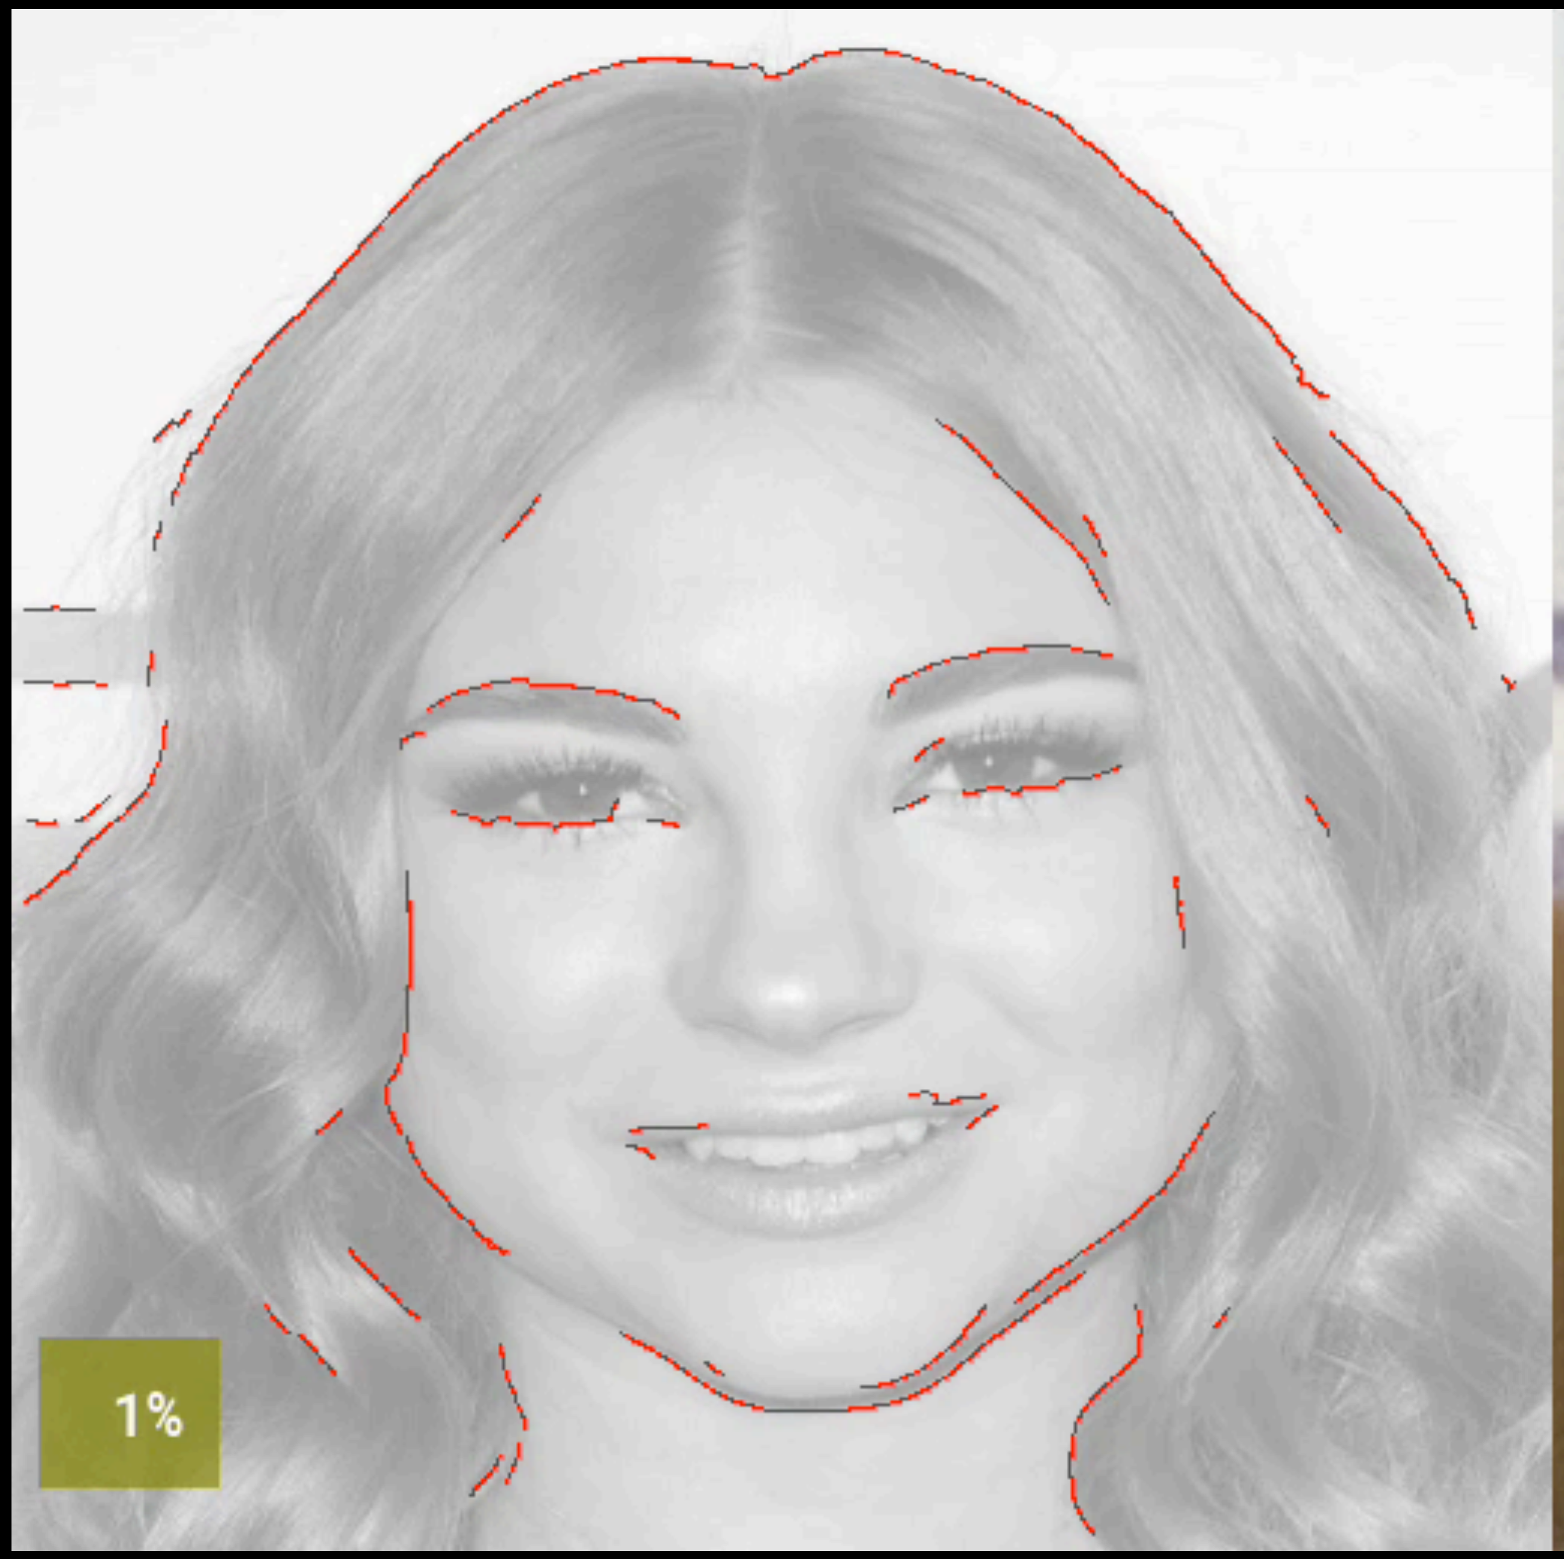

Input contours (overlaid on image)

Reconstruction

Original image

Able to Synthesize Texture Where Contours Are Absent

Original image

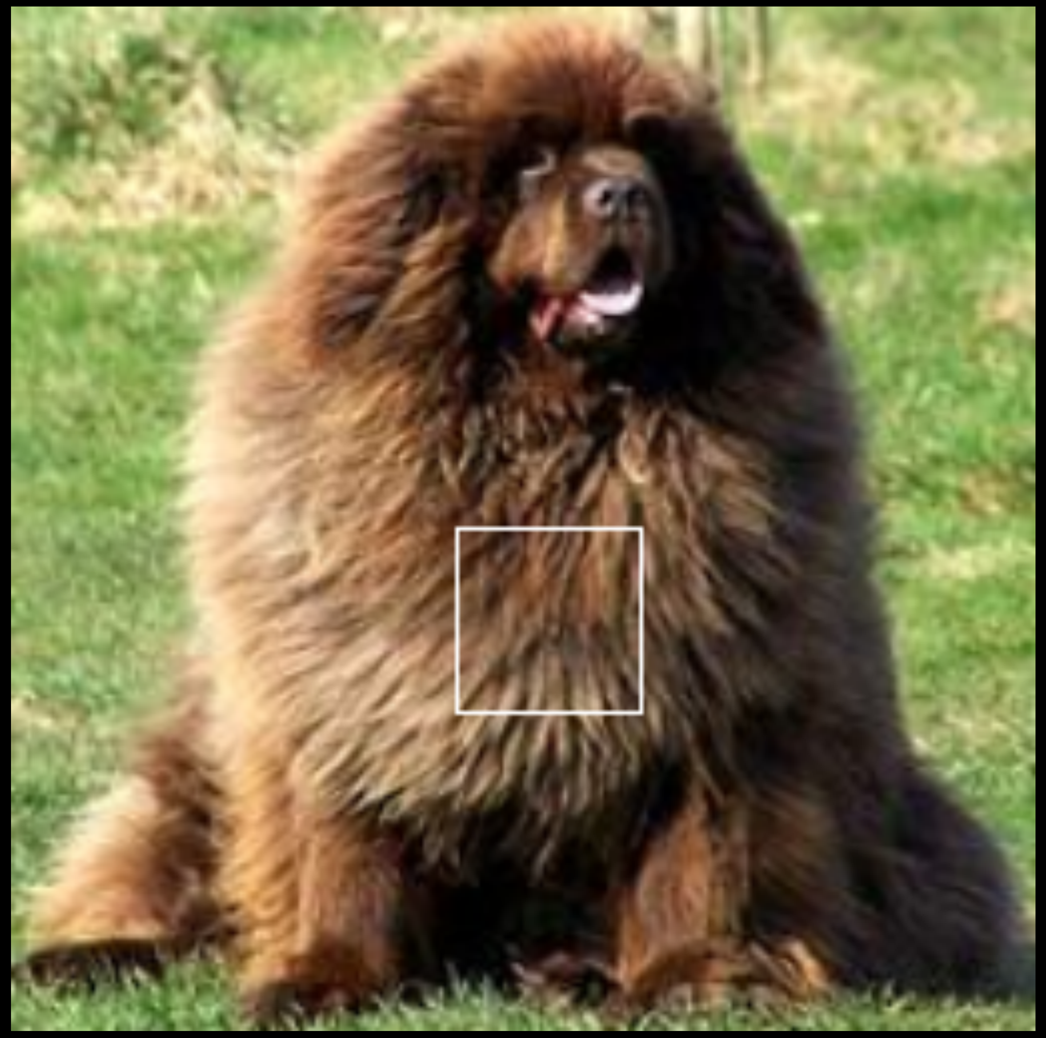

Contours overlaid

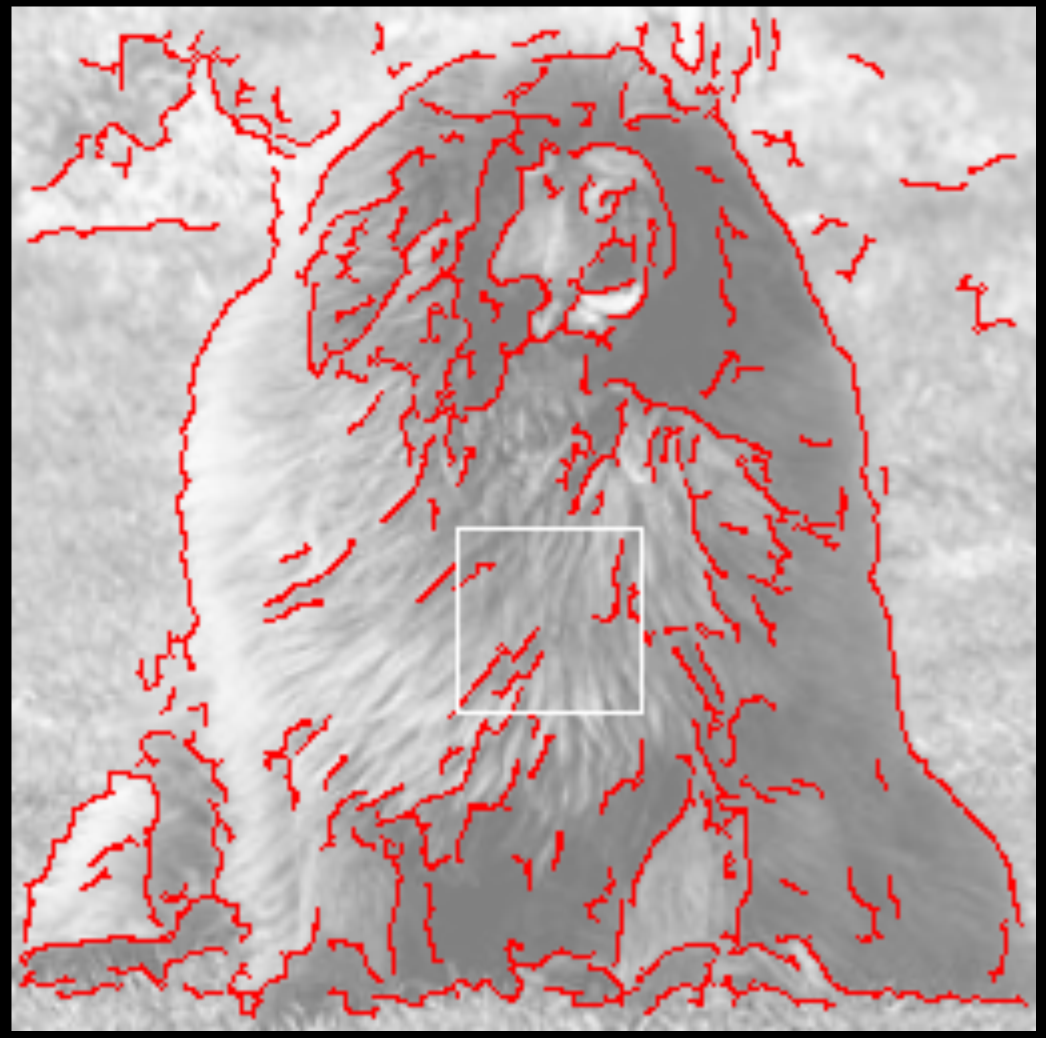

Our reconstruction

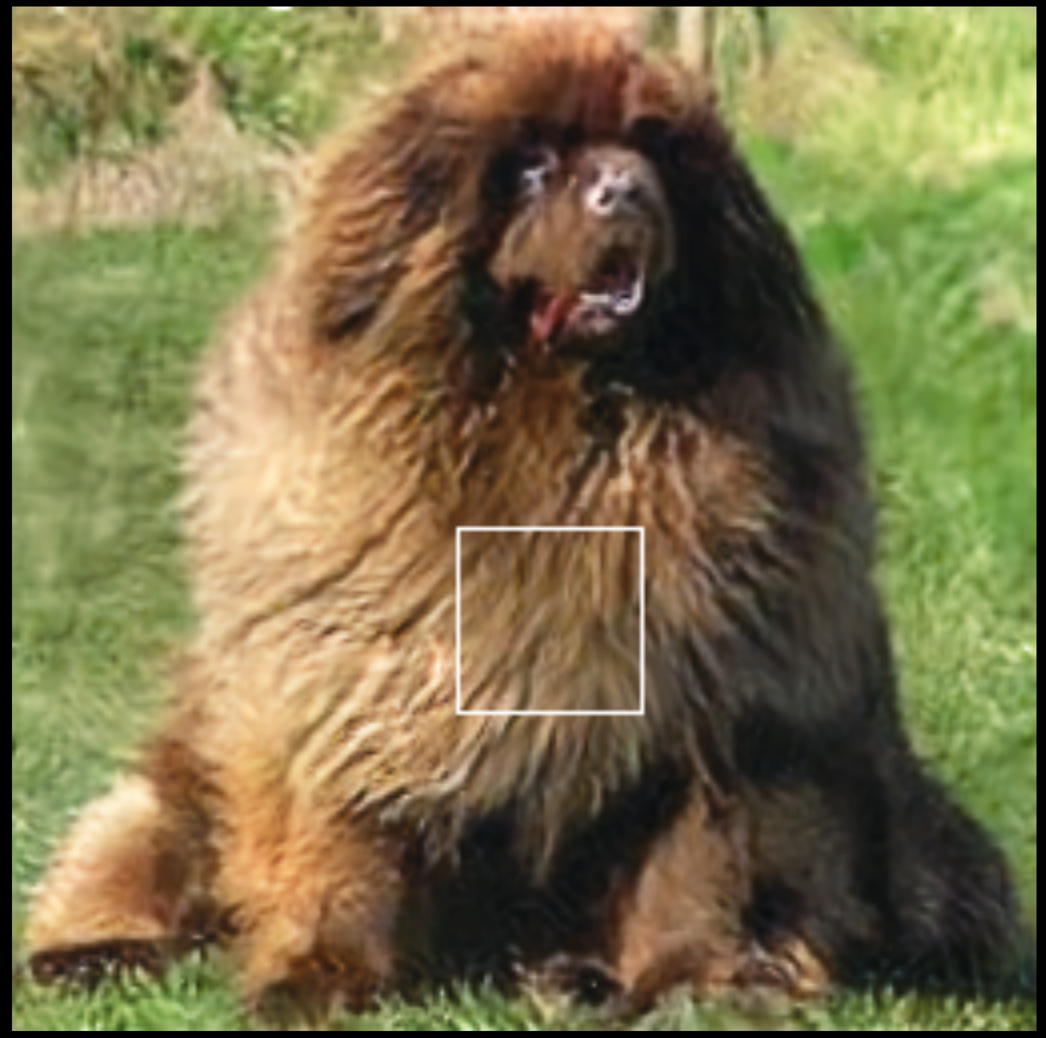

Close-up

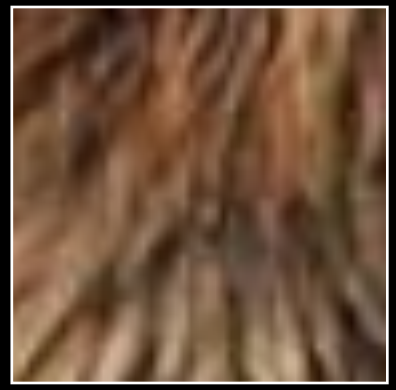

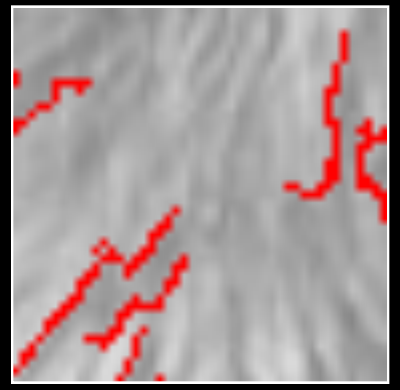

Able to Synthesize Texture Where Contours Are Absent

Original image

Contours overlaid

Our reconstruction

Close-up

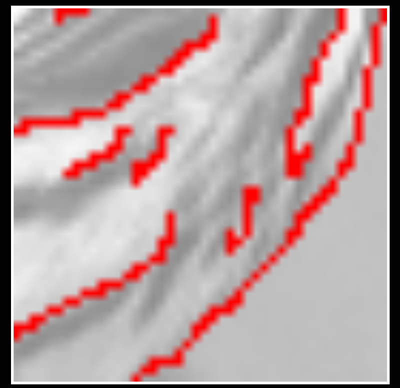

Semantically Aware Editing in the Contour Domain

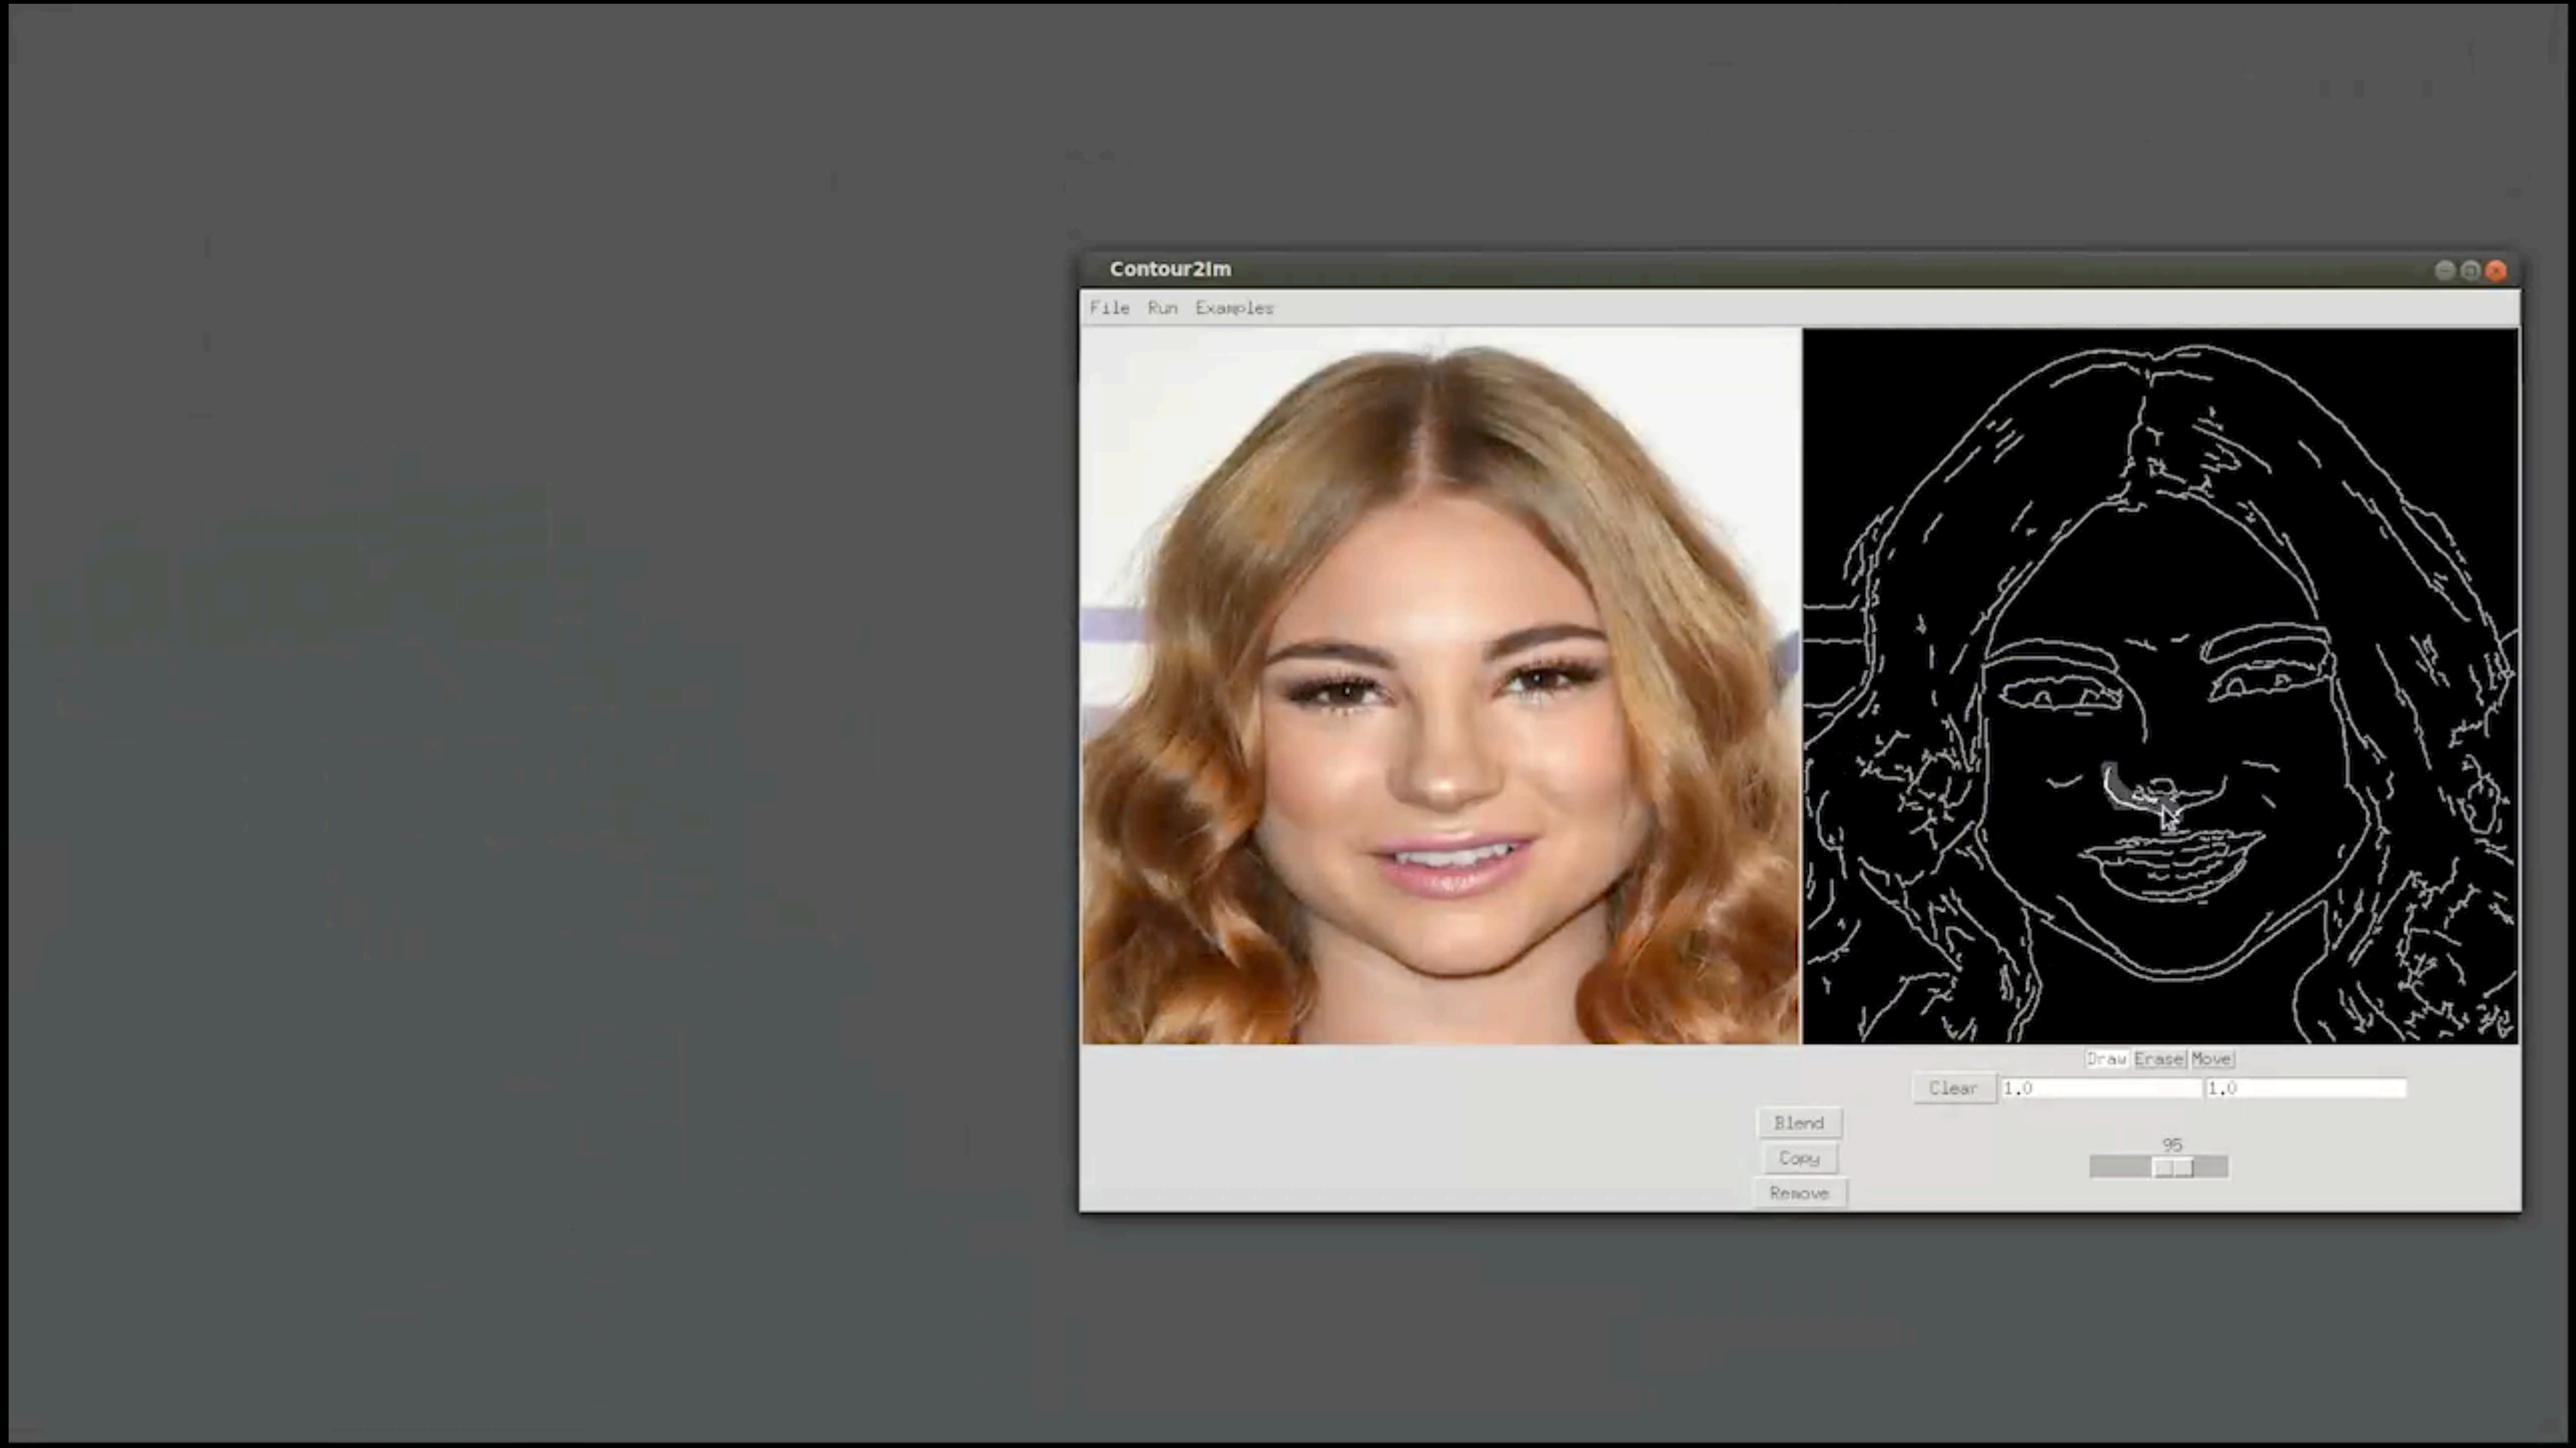

Semantically Aware Editing in the Contour Domain

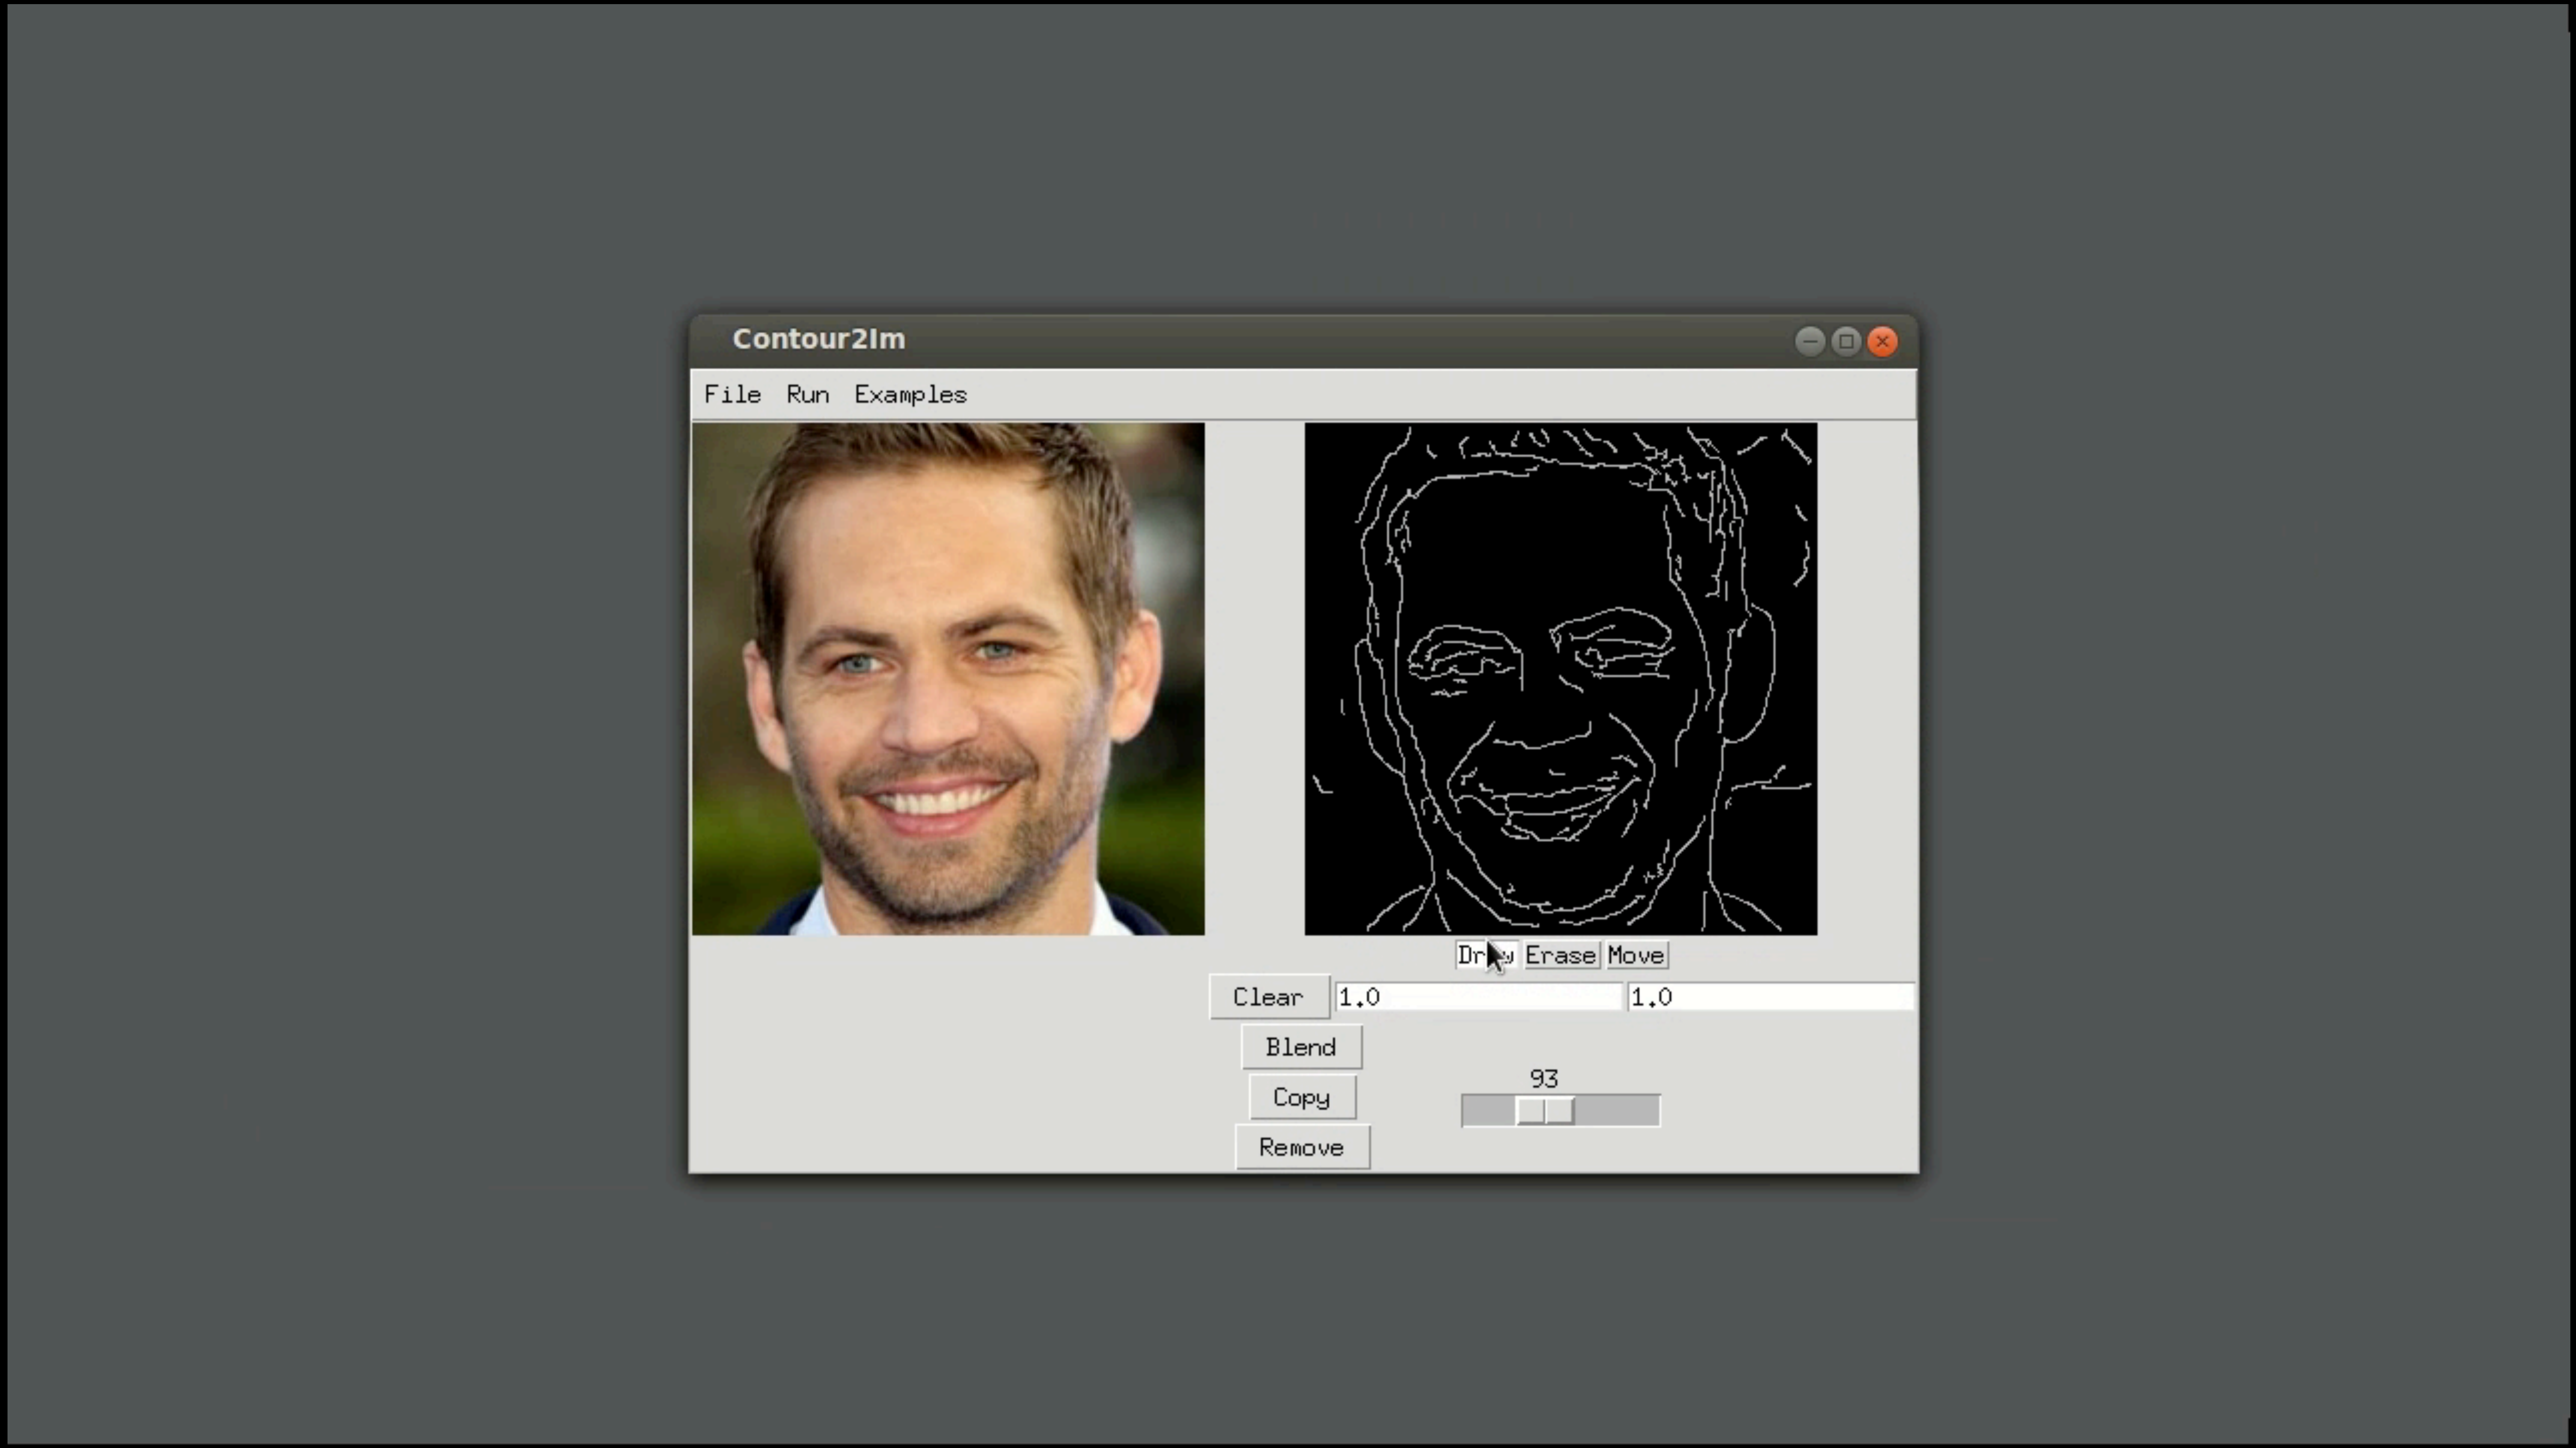

Limitations

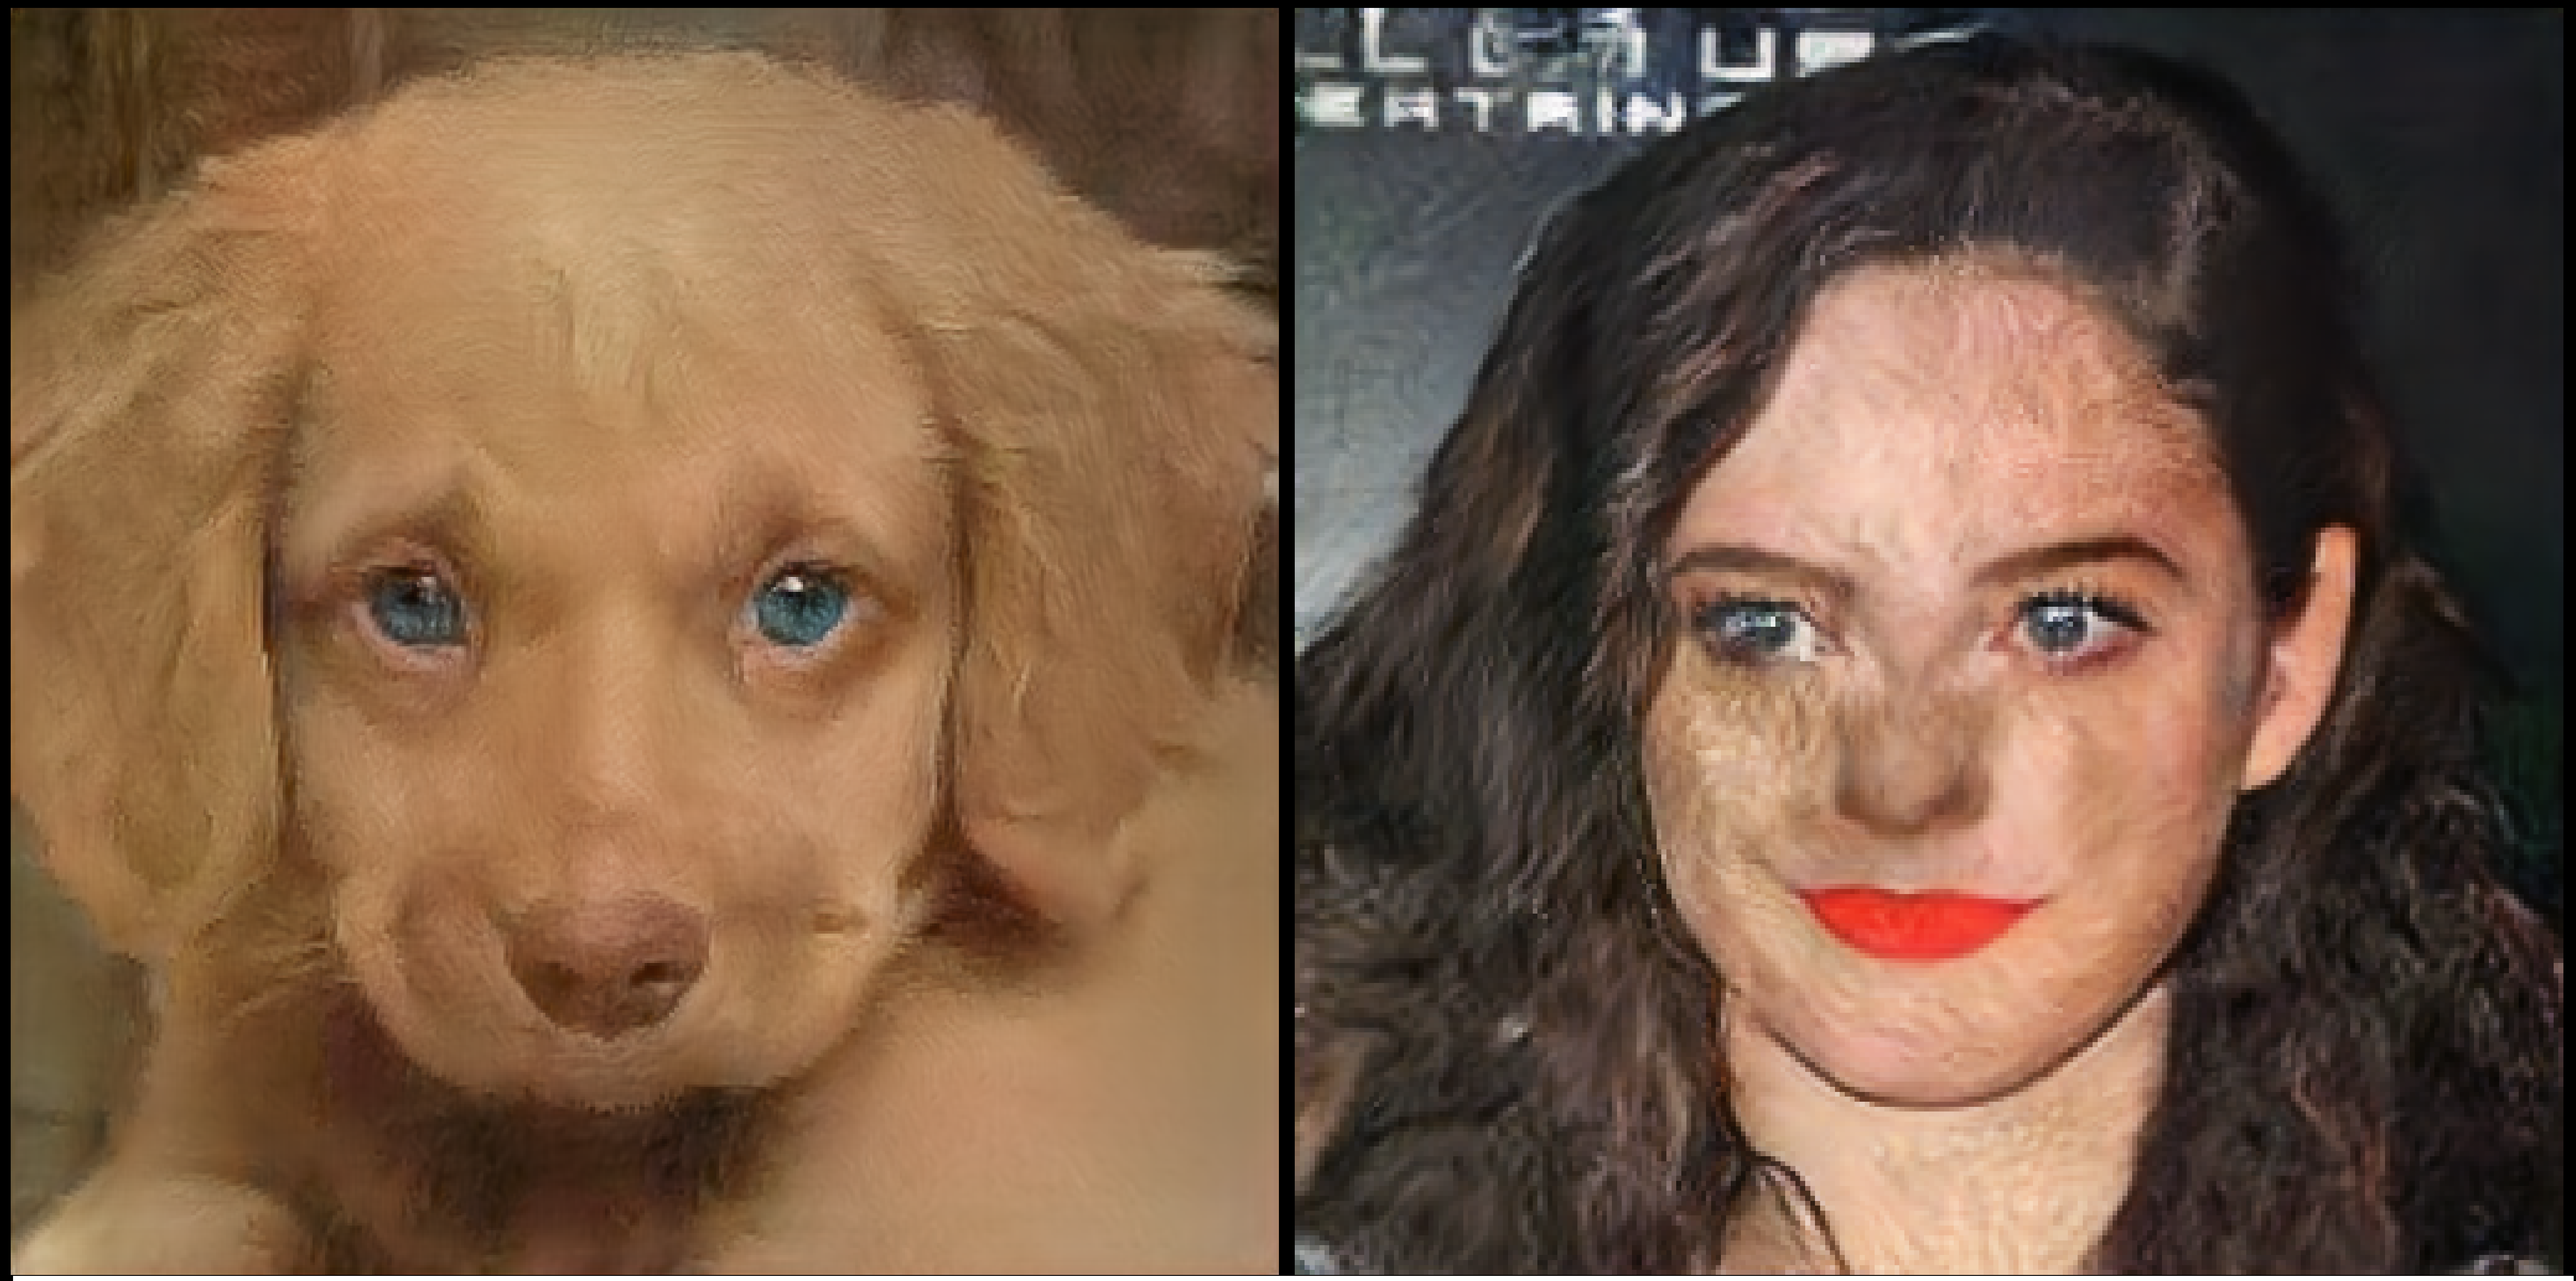

(a) Recon. using face model (b) Recon. using dog model

CVPR 2018

Poster 85, Tuesday 4:30-6:30 Halls C-E

Smart, Sparse Contours to Represent and Edit Images

Tali Dekel, Chung Gan, Dilip Krishnan, Ce Liu, Bill Freeman, CVPR18

Tali Dekel

Dilip Krishnan

Ce Liu

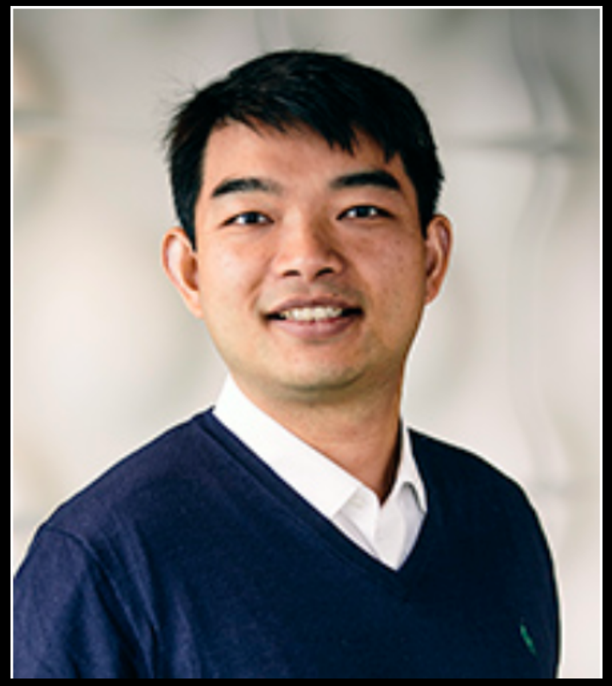

Chung Gan

Research at Google

Copying and Editing Images

Photoscan

Ce Liu, Michael Rubinstein, Mike Krainin, Bill Freeman, 2016

Securing Visible Watermarks

Tali Dekel, Michael Rubinstein, Ce Liu, Bill Freeman, CVPR17

Smart, Sparse Contours to Represent and Edit Images

Tali Dekel, Chung Gan, Dilip Krishnan, Ce Liu, Bill Freeman, CVPR18

Research at Google