Single Image Portrait Relighting

Tiancheng Sun[1], Jonathan T. Barron[2], Yun-Ta Tsai[2], Zexiang Xu[1],
Xueming Yu[3], Graham Fyffe[3], Christoph Rhemann[3], Jay Busch[3],
Paul Debevec[3], Ravi Ramamoorthi[1]

[1]University of California, San Diego, [2]Google Research, [3]Google

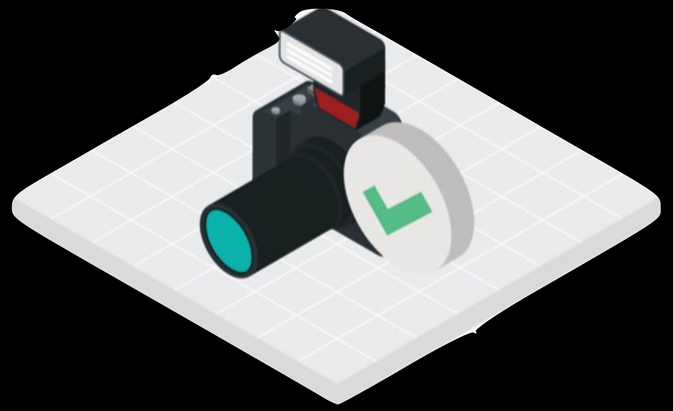

Photography & Recording Allowed

Overview

Overview

light from the back

Overview

light from the back

shadow on the face

Overview

light from the back

shadow on the face

want to add
dramatic lighting

Overview

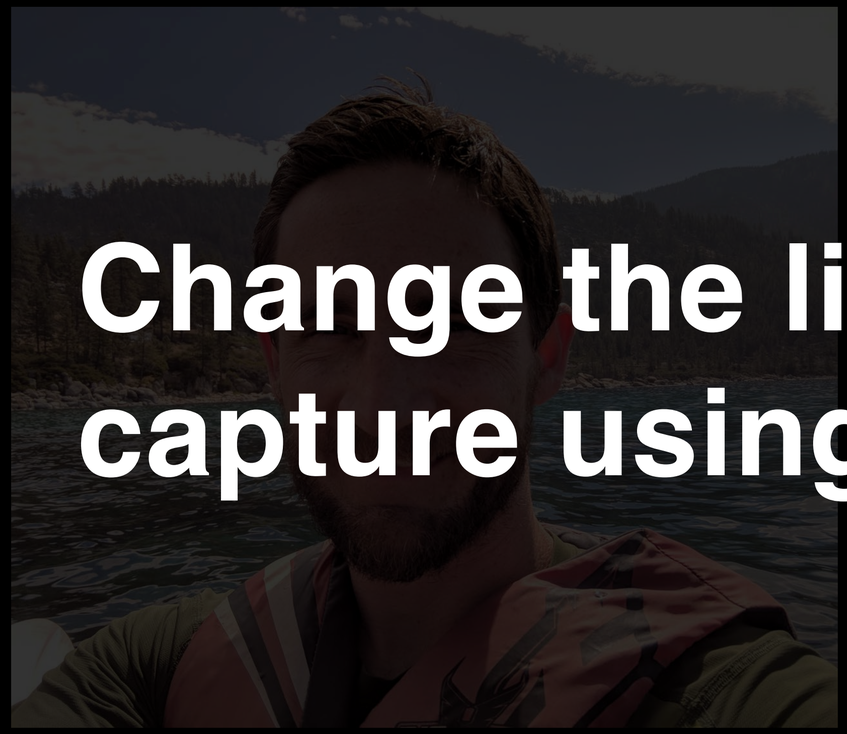

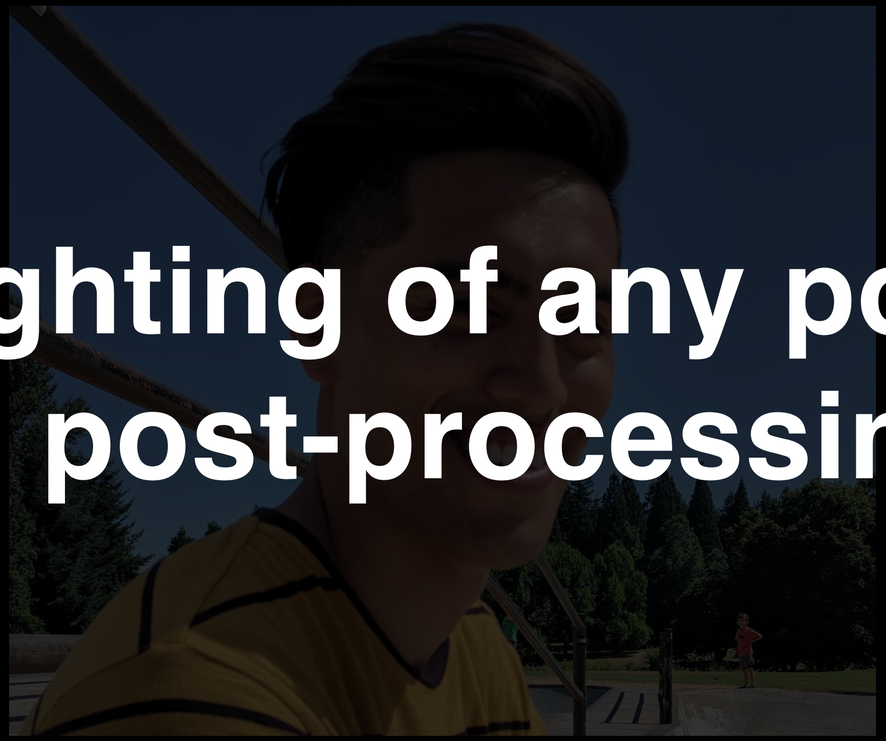

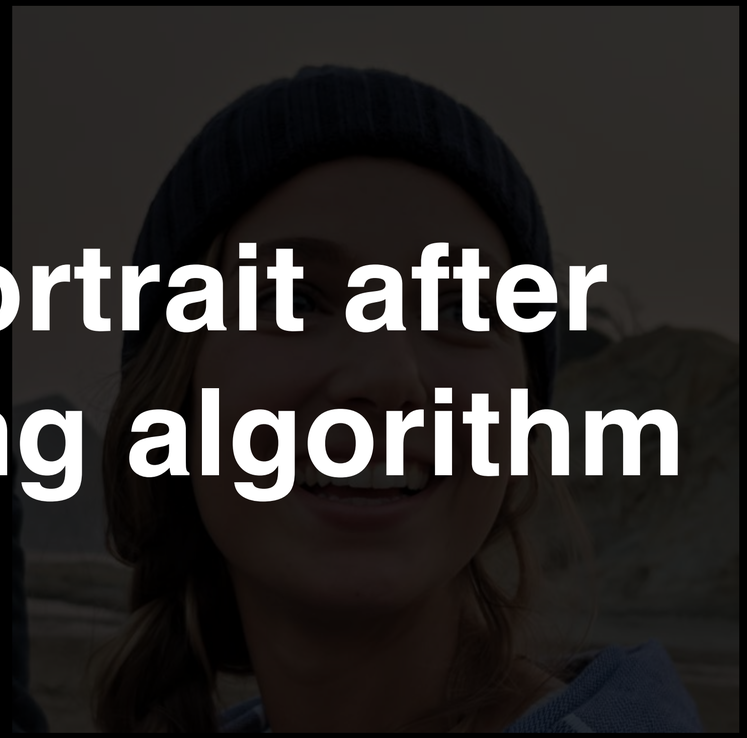

Change the lighting of any portrait after capture using post-processing algorithm

light from the back

shadow on the face

want to add dramatic lighting

Overview

Overview

Overview

Portrait Relighting System

Overview

Portrait Relighting System

another lighting

Overview

Portrait Relighting System

another lighting

Overview

Portrait Relighting System

Overview

Portrait Relighting System

Overview

Portrait Relighting System

I want to rotate the lighting a little bit.

Overview

Portrait Relighting System

I want to rotate the lighting a little bit.

Previous work

Previous work

- Light Stage

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Previous work

- Light Stage

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Previous work

- Light Stage

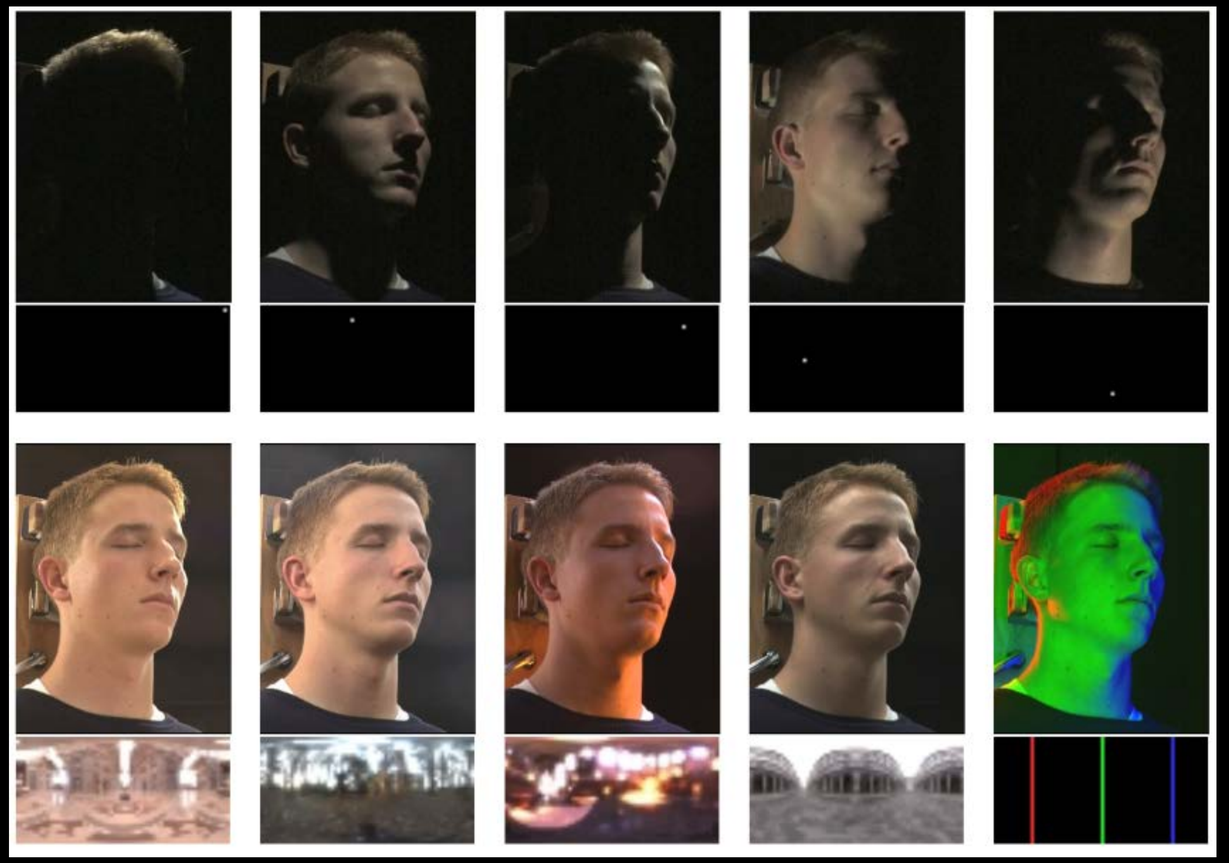

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Previous work

- Light Stage: **capture ~100 images and do image-based relighting**

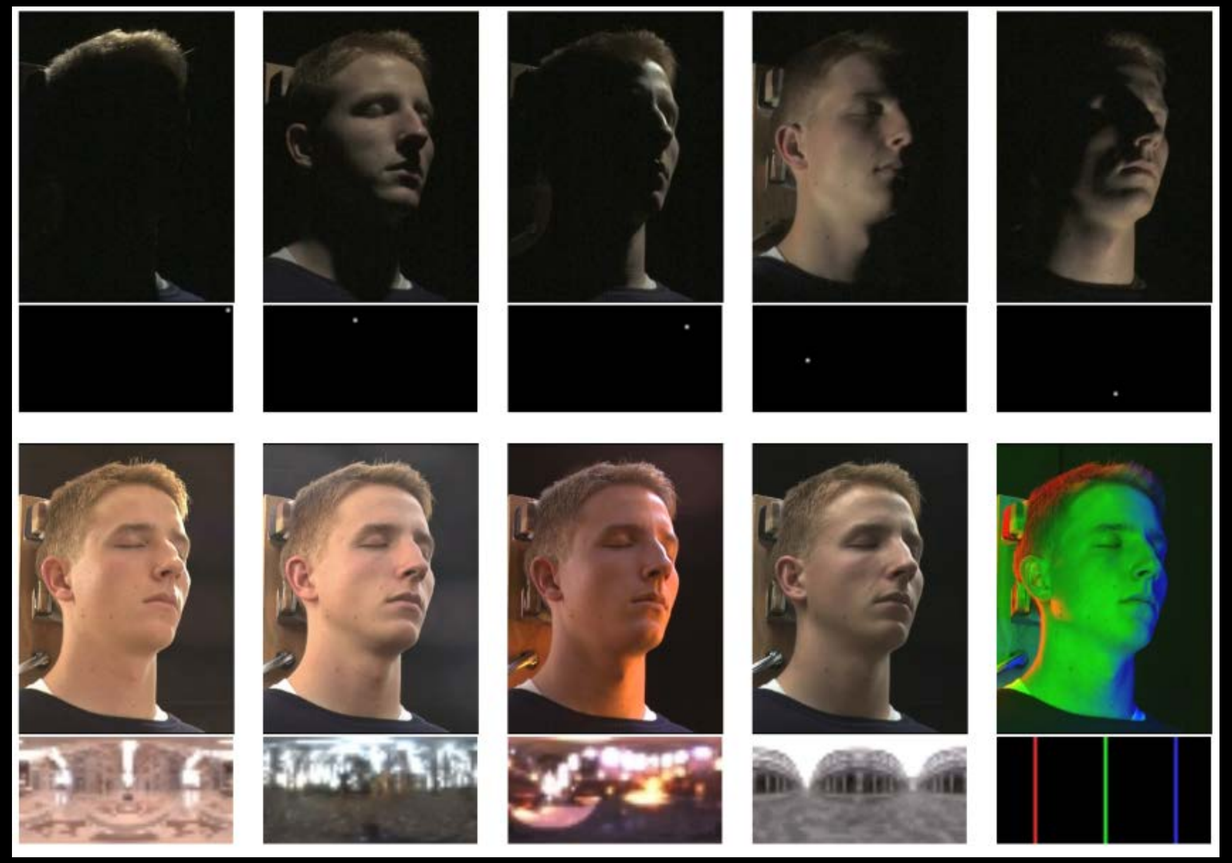

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Previous work

Previous work

- Deep image-based relighting

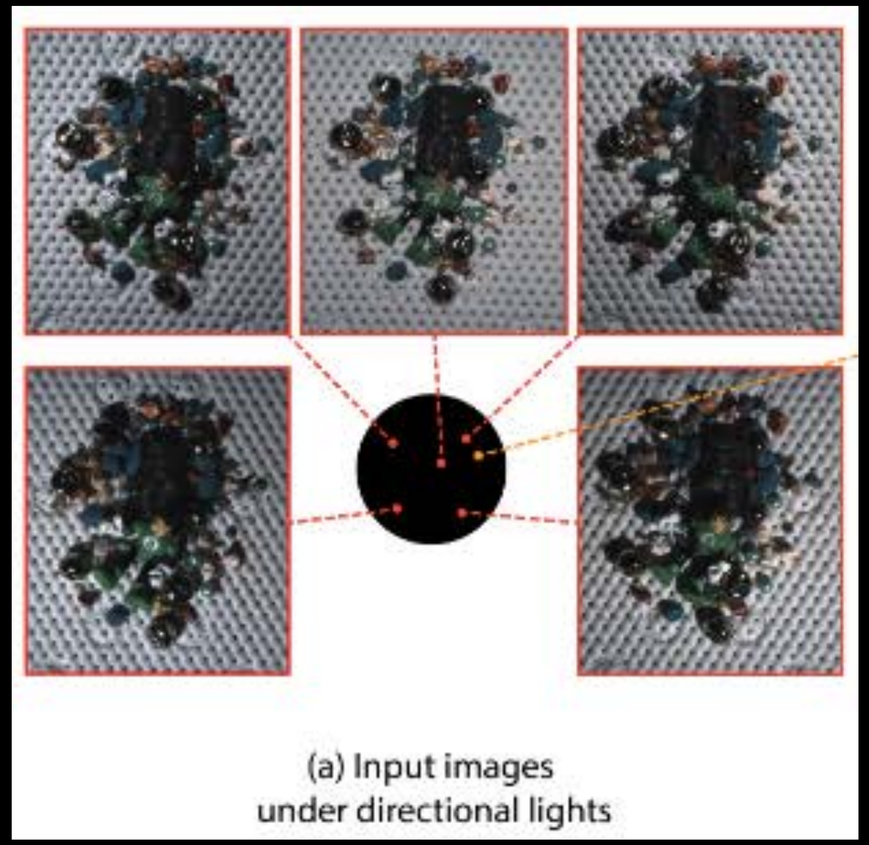

Xu, Zexiang, et al. "Deep image-based relighting from optimal sparse samples." SIGGRAPH 2018

Previous work

- Deep image-based relighting

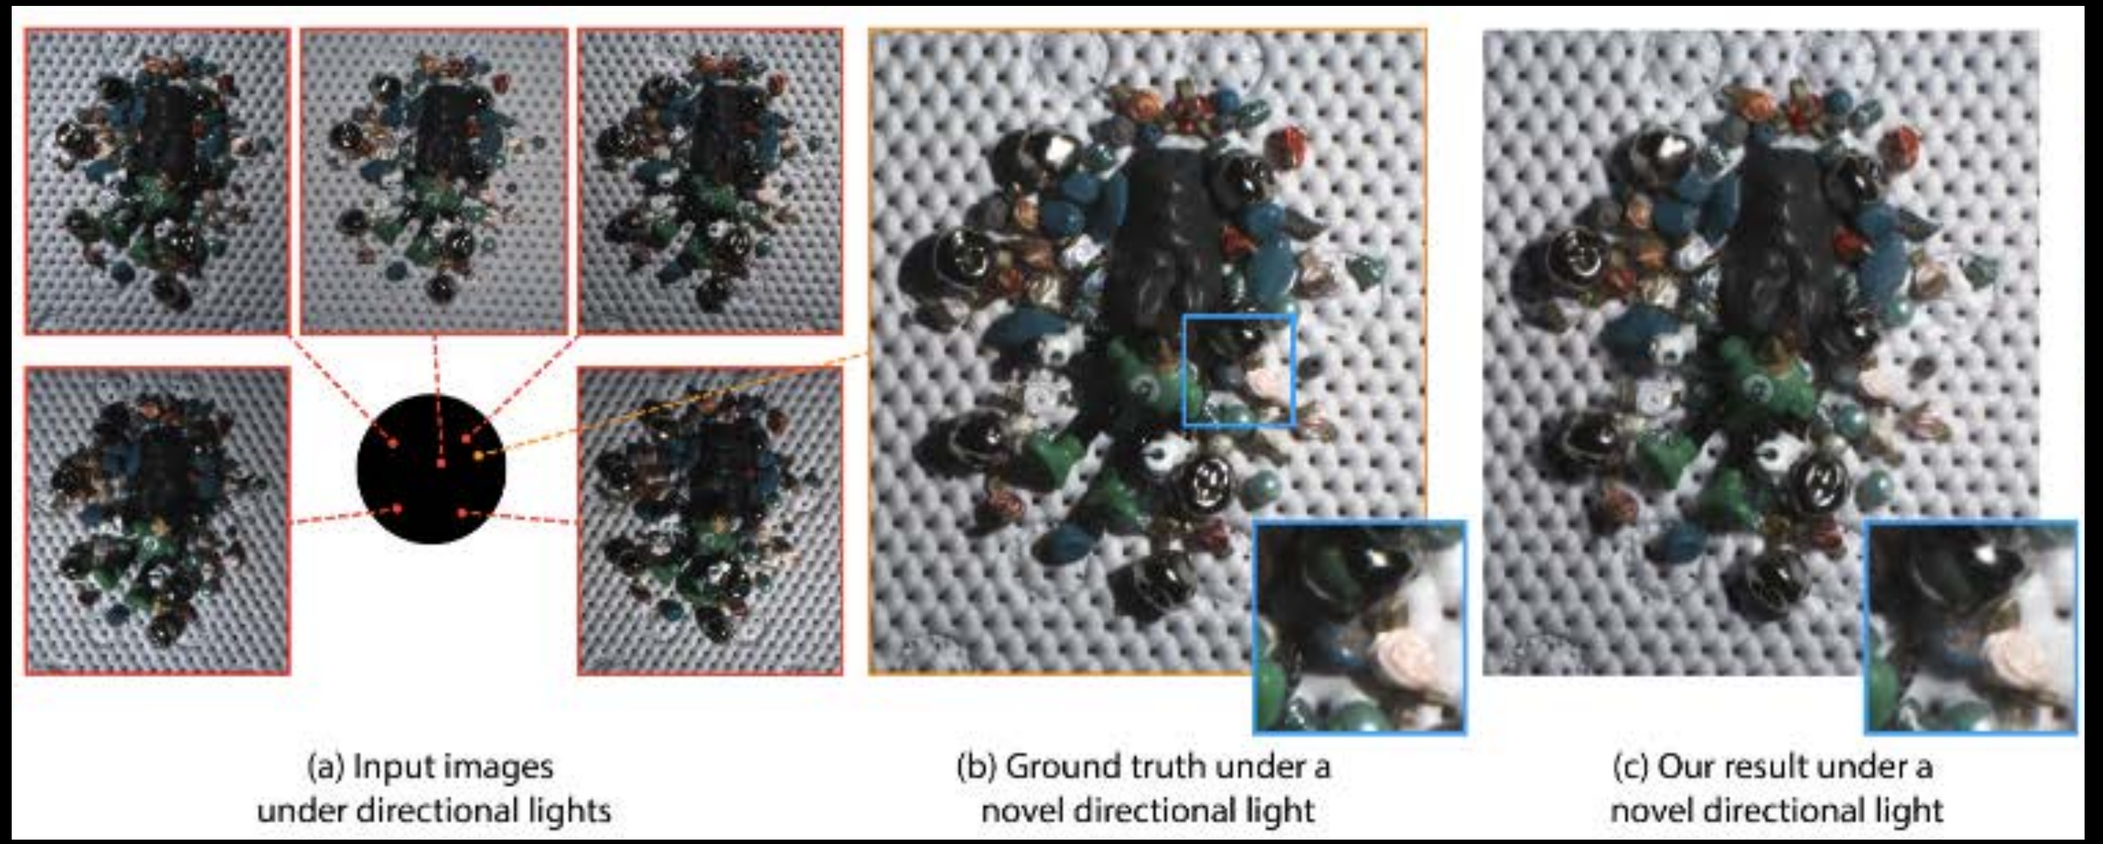

Xu, Zexiang, et al. "Deep image-based relighting from optimal sparse samples." SIGGRAPH 2018

Previous work

- Deep image-based relighting

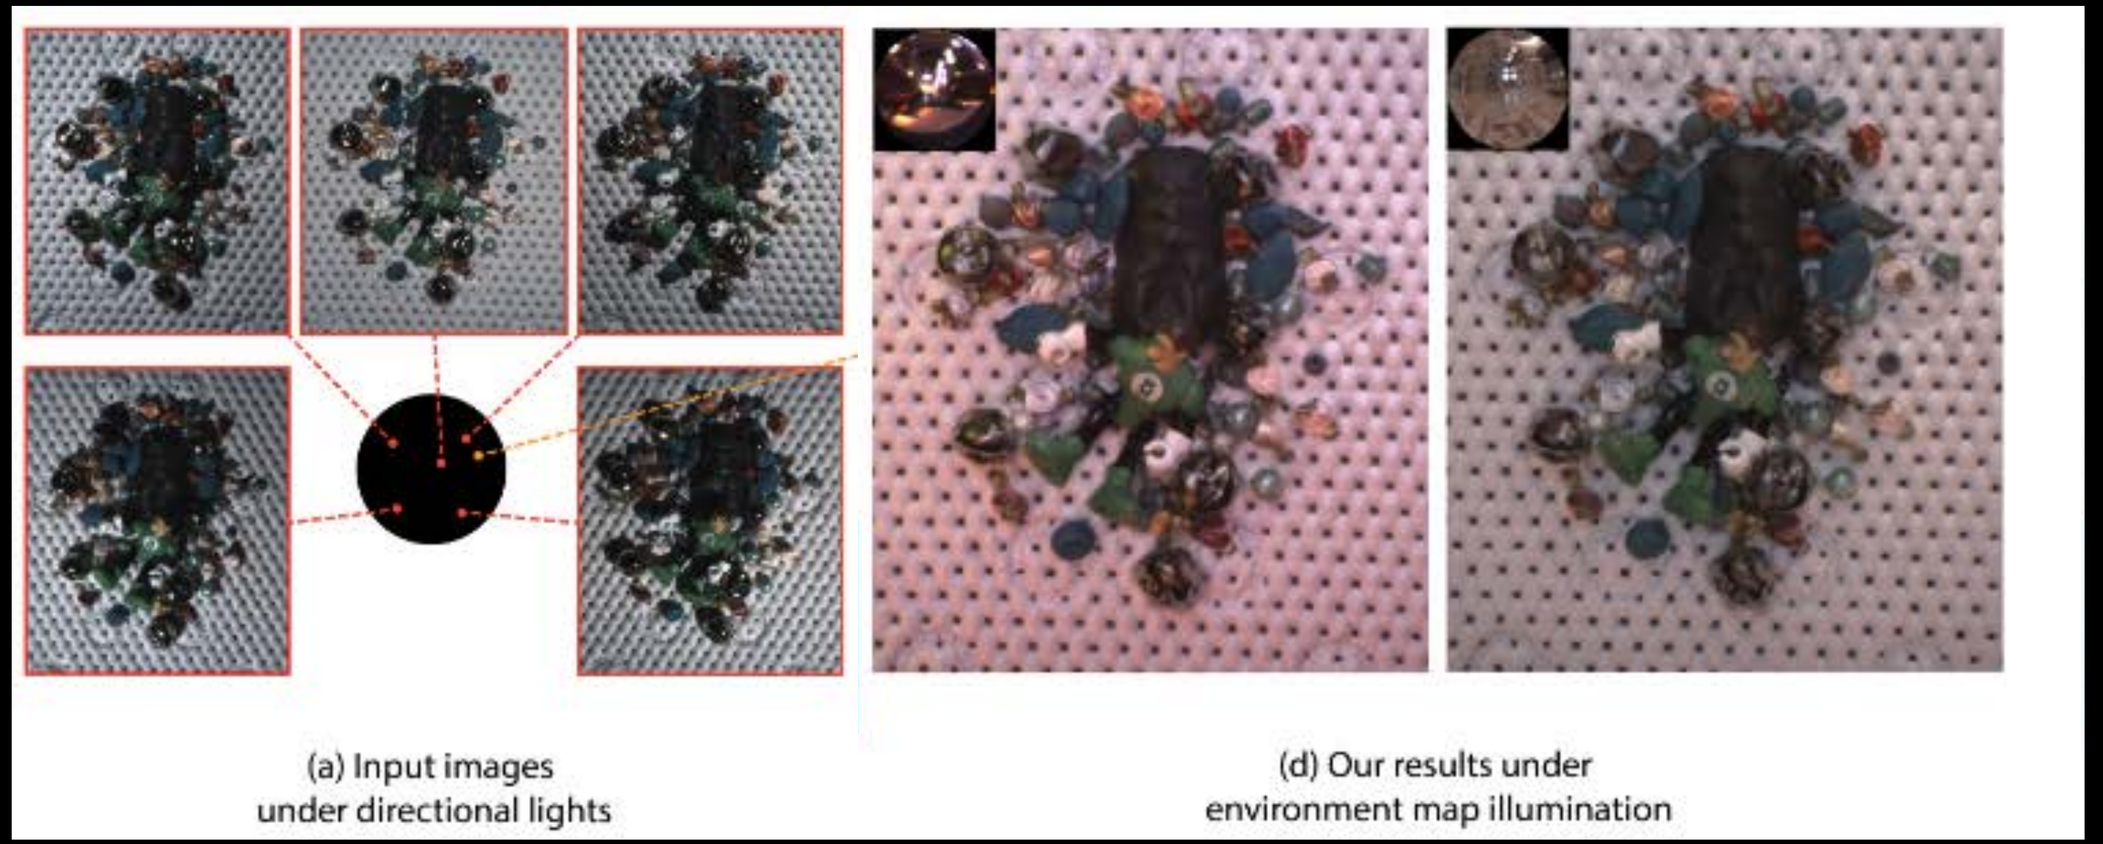

Xu, Zexiang, et al. "Deep image-based relighting from optimal sparse samples." SIGGRAPH 2018

Previous work

- Deep image-based relighting

capture 5 images and do relighting via neural network

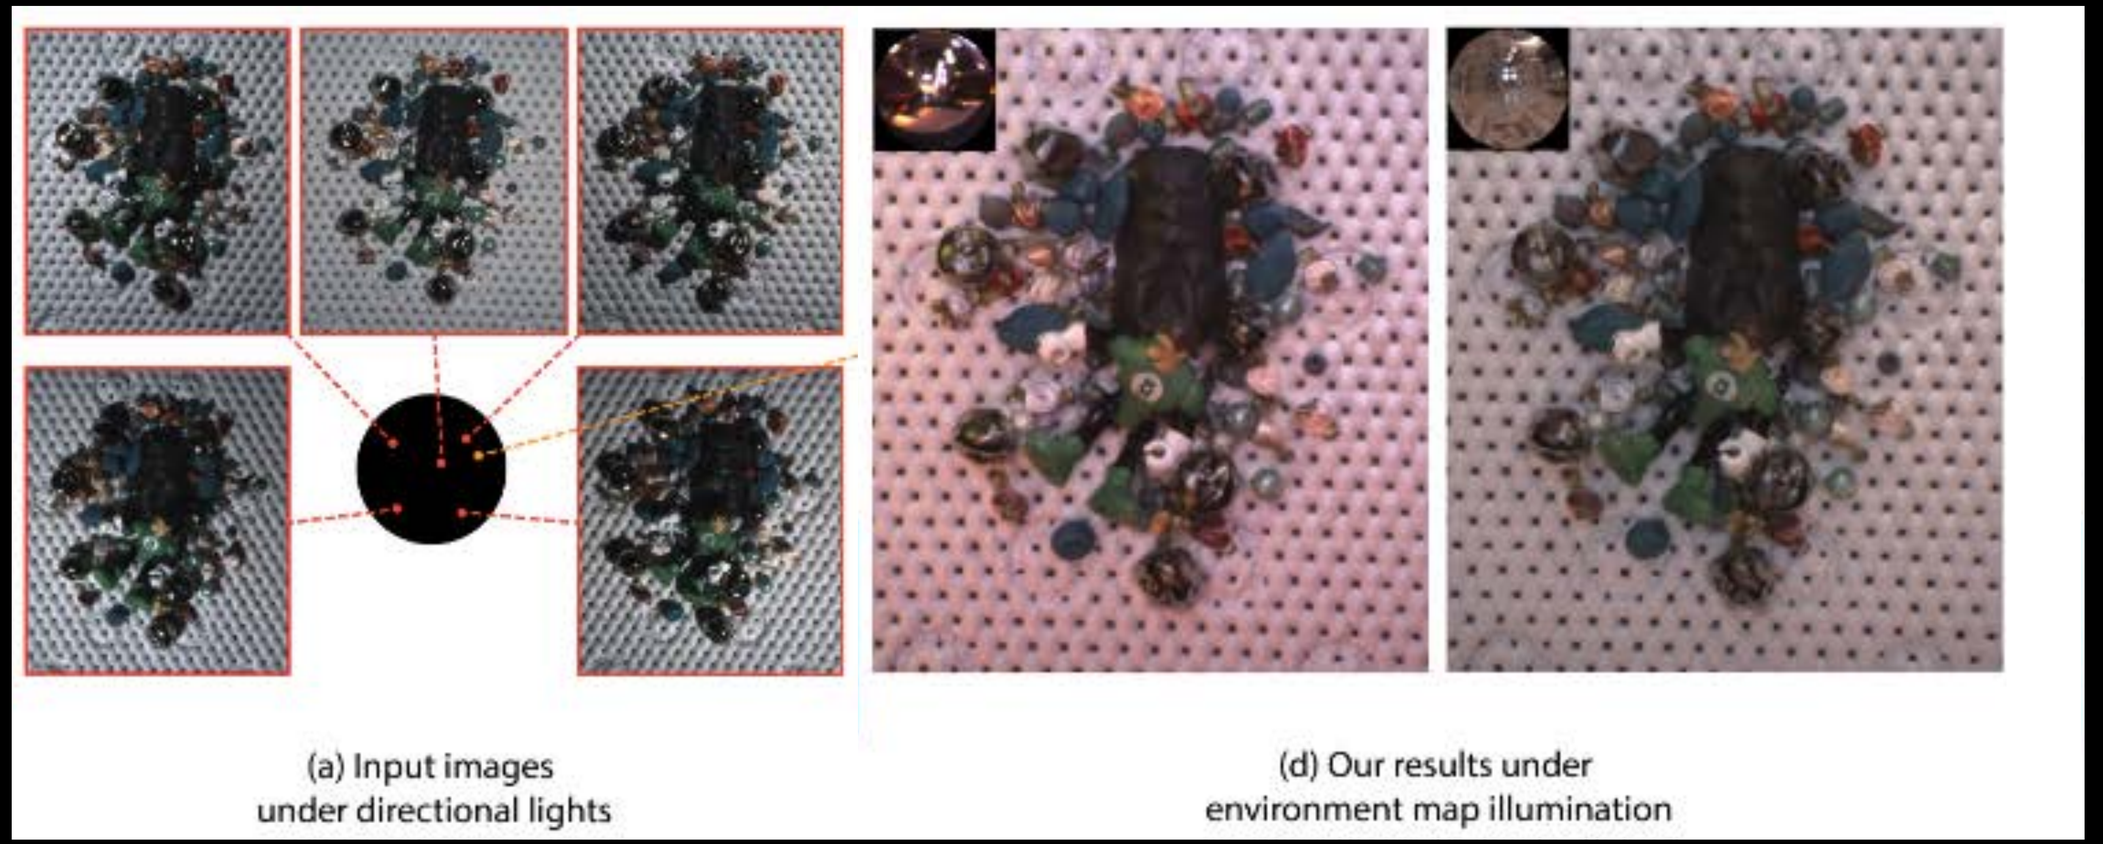

Xu, Zexiang, et al. "Deep image-based relighting from optimal sparse samples." SIGGRAPH 2018

Previous work

Previous work

- Portrait lighting transfer

Shu, Zhixin, et al. "Portrait lighting transfer using a mass transport approach." SIGGRAPH 2018

Previous work

- Portrait lighting transfer

Shu, Zhixin, et al. "Portrait lighting transfer using a mass transport approach." SIGGRAPH 2018

Previous work

- Portrait lighting transfer

  transfer lighting from one portrait to another

Shu, Zhixin, et al. "Portrait lighting transfer using a mass transport approach." SIGGRAPH 2018

Previous work

Previous work

- SfSNet

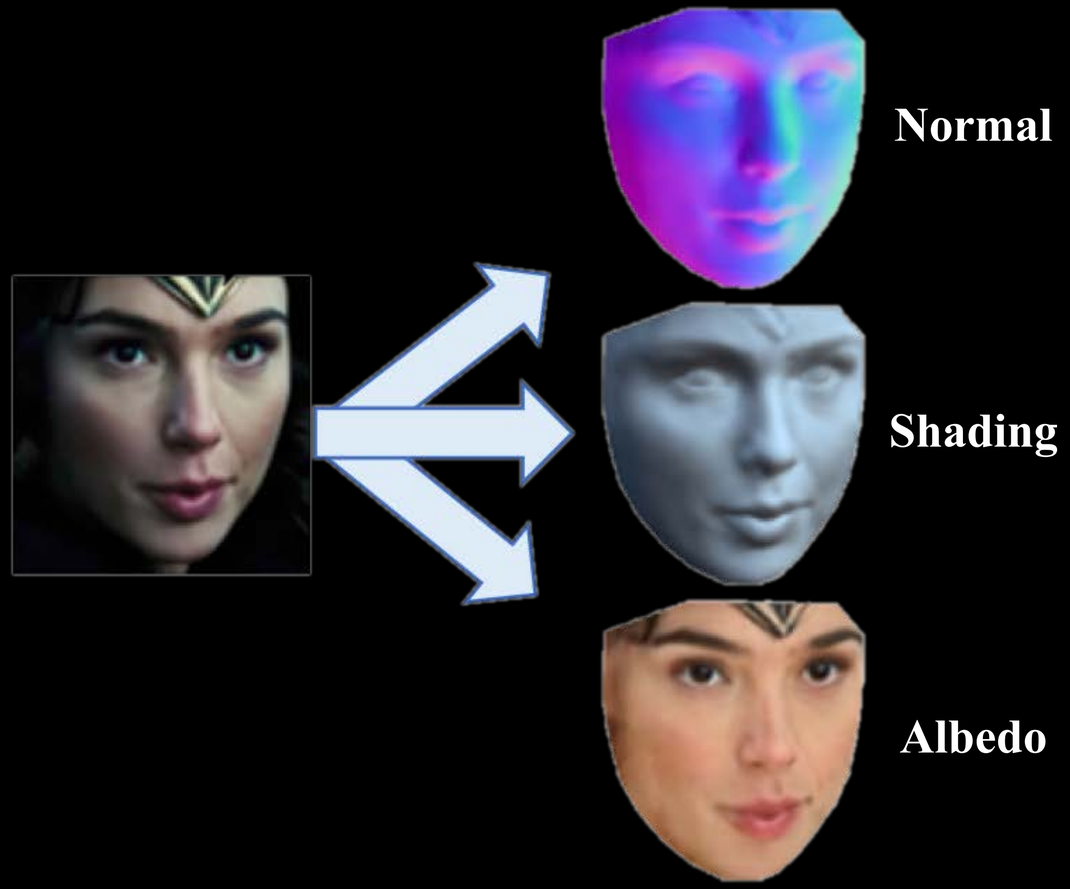

Sengupta, Soumyadip, et al. "SfSNet: Learning Shape, Reflectance and Illuminance of Faces in the Wild'." CVPR 2018.

Previous work

- SfSNet: **deep intrinsic decomposition mostly trained on synthetic faces**

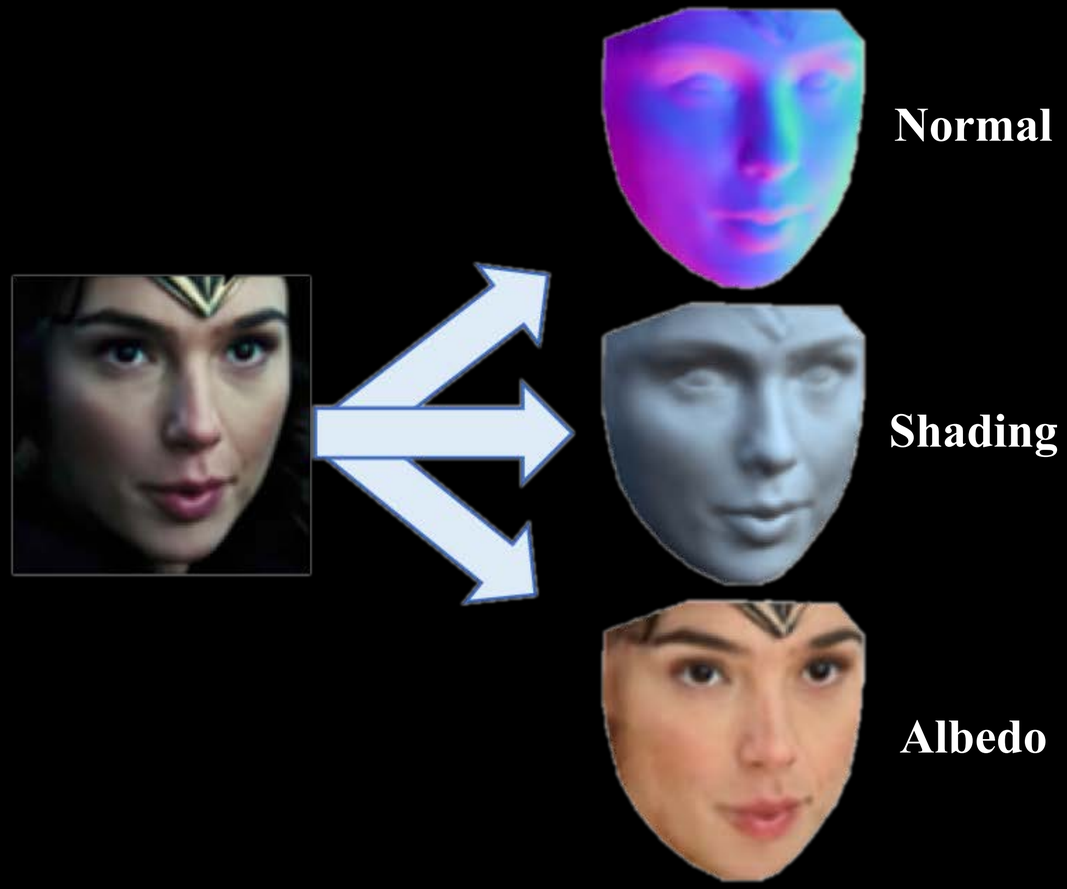

Sengupta, Soumyadip, et al. "SfSNet: Learning Shape, Reflectance and Illuminance of Faces in the Wild'." CVPR 2018.

Previous work

- SfSNet: **deep intrinsic decomposition mostly trained on synthetic faces**

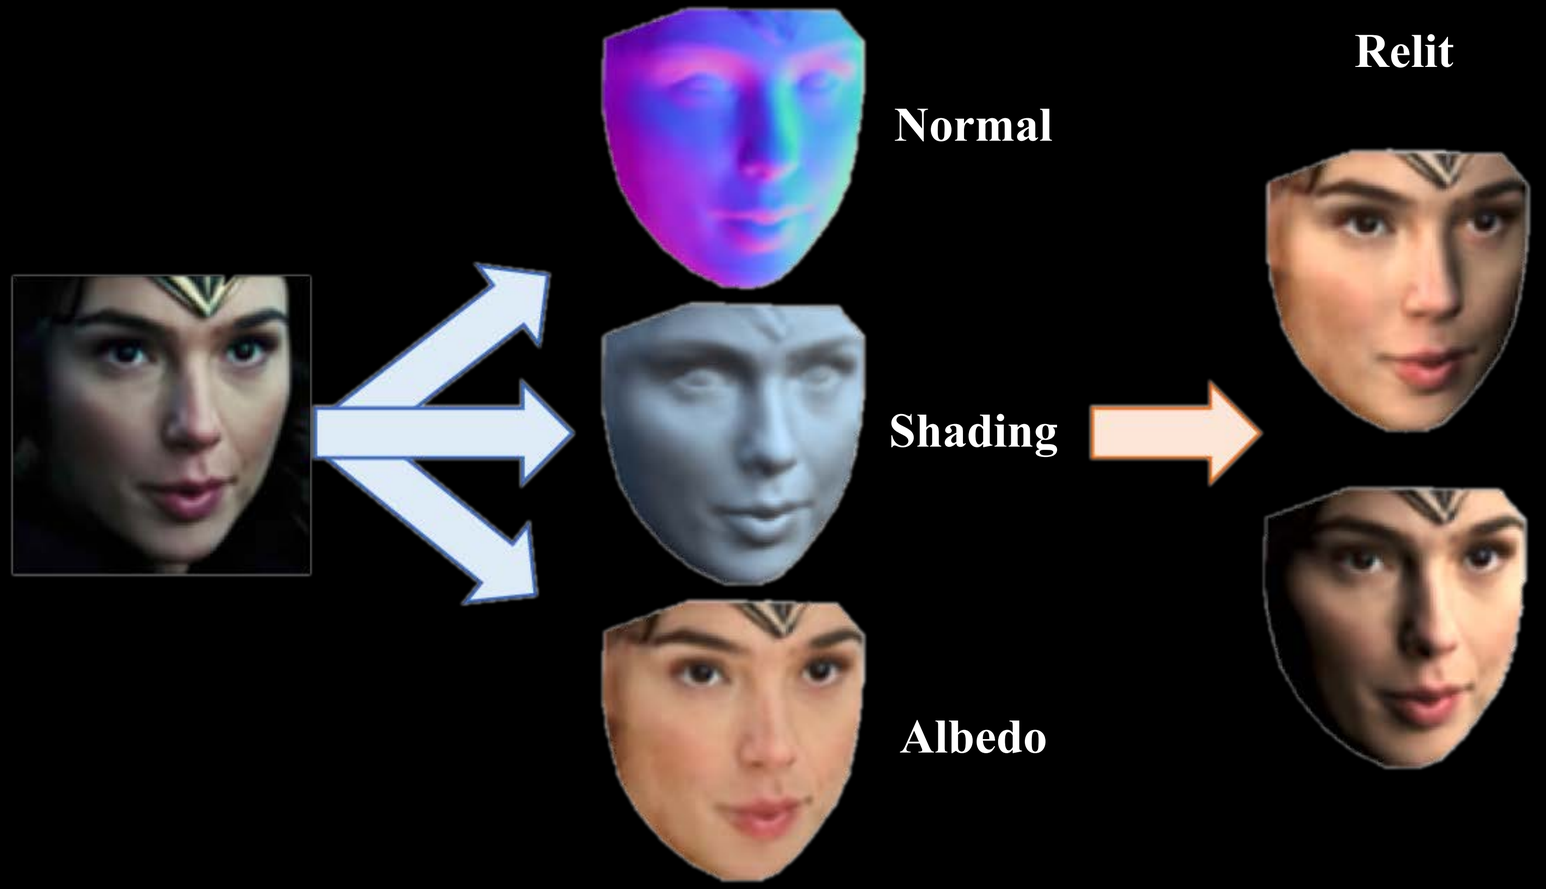

Sengupta, Soumyadip, et al. "SfSNet: Learning Shape, Reflectance and Illuminance of Faces in the Wild'." CVPR 2018.

Overview

Overview

- Goal: **practical relighting on single portrait image**

Overview

- Goal: **practical relighting on single portrait image**
- Practical in detail:
  - Robust to the pose and camera view
  - Work well on natural lightings
  - Adapt to high-resolution images
  - Run at interactive rate

Overview

- Goal: **practical relighting on single portrait image**
- Practical in detail:
  - Robust to the pose and camera view
  - Work well on natural lightings
  - Adapt to high-resolution images
  - Run at interactive rate
- Solution: Deep Neural Network + Real Face Data.

Method

Method

portrait under
lighting A

Method

portrait under lighting A

Portrait Relighting System (Neural Network)

Method

portrait under lighting A

Portrait Relighting System (Neural Network)

lighting B

Method

portrait under lighting A

Portrait Relighting System (Neural Network)

lighting B

portrait under lighting B

Method

portrait under lighting A

Portrait Relighting System (Neural Network)

lighting A

lighting B

portrait under lighting B

Method

portrait under lighting A

Portrait Relighting System (Neural Network)

lighting A

lighting B

portrait under lighting B

Method

portrait under lighting A

Portrait Relighting System (Neural Network)

lighting A

lighting B

portrait under lighting B

Method

portrait under lighting A

Portrait Relighting System (Neural Network)

lighting A

lighting B

portrait under lighting B

Method

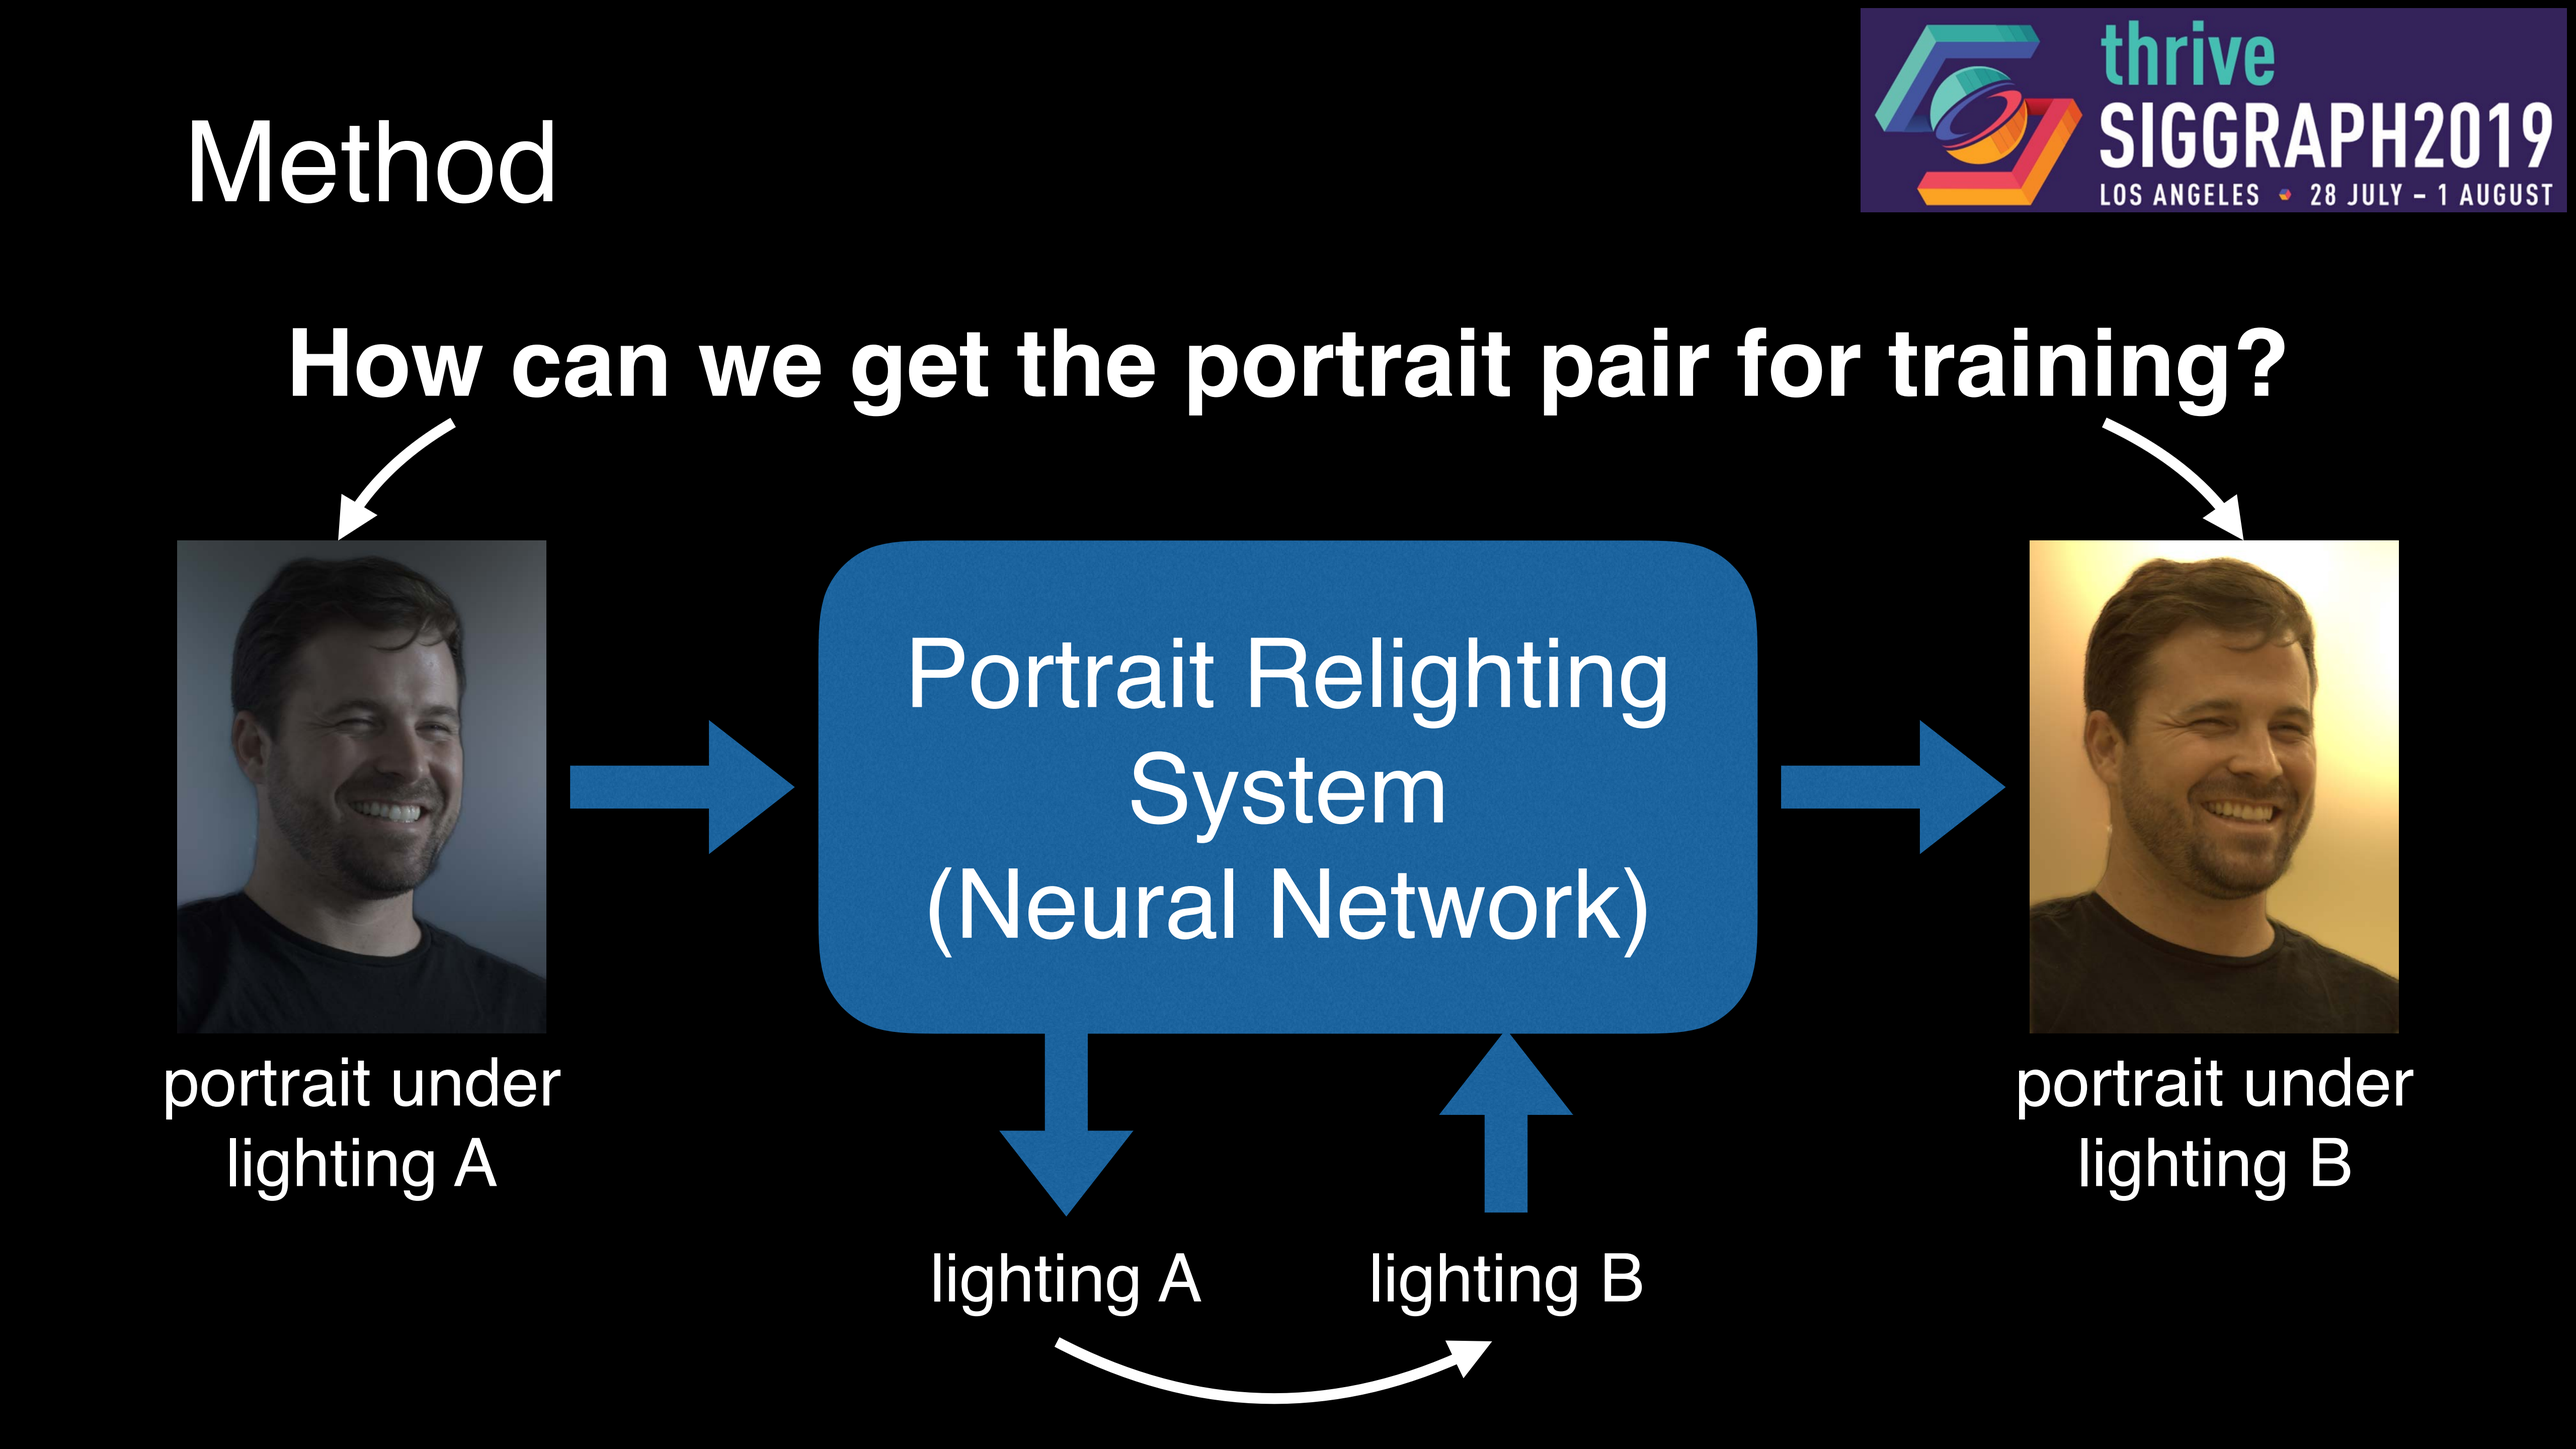
How can we get the portrait pair for training?
Portrait Relighting System (Neural Network)
portrait under lighting A
portrait under lighting B
lighting A
lighting B

Method: Data

Method: Data

Light Stage

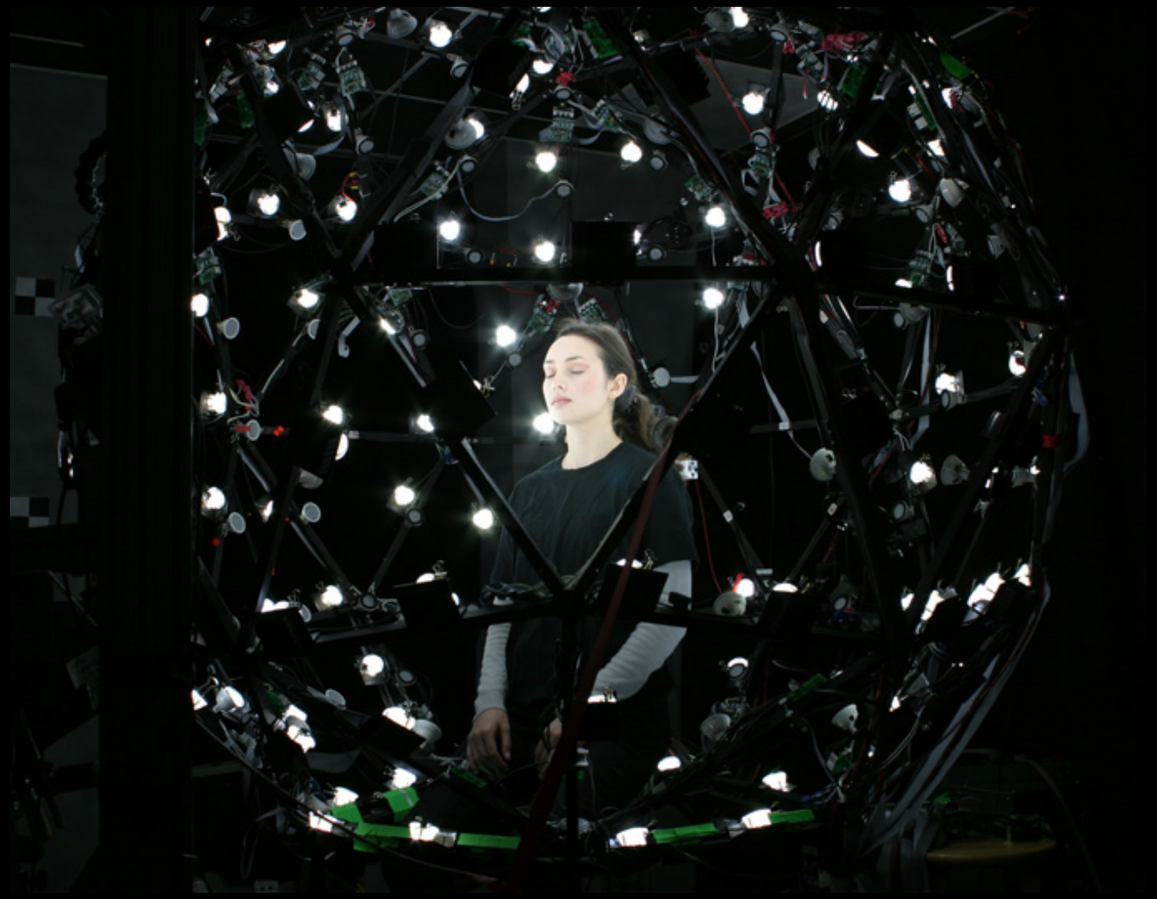

Method: Data

Light Stage

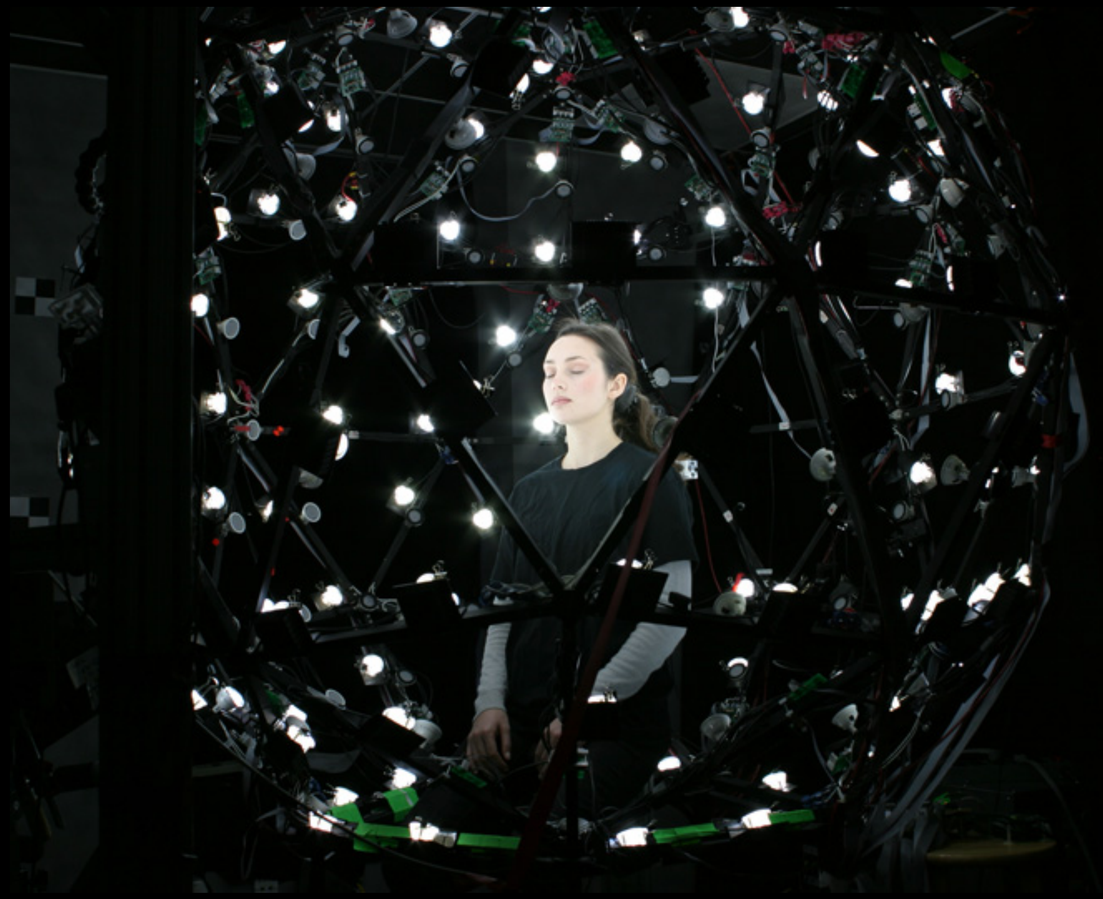

One-Light-At-a-Time scans (OLAT)

Method: Data

Method: Data

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

captured OLAT

captured OLAT

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

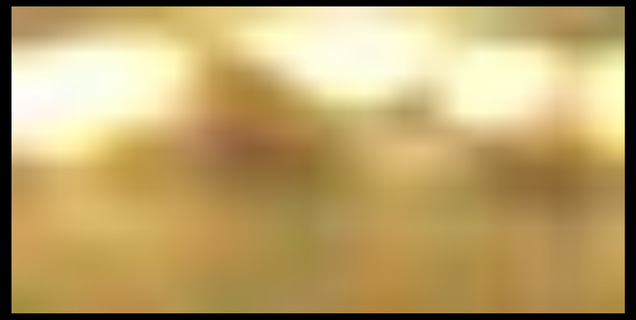

lighting

captured OLAT

captured OLAT

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

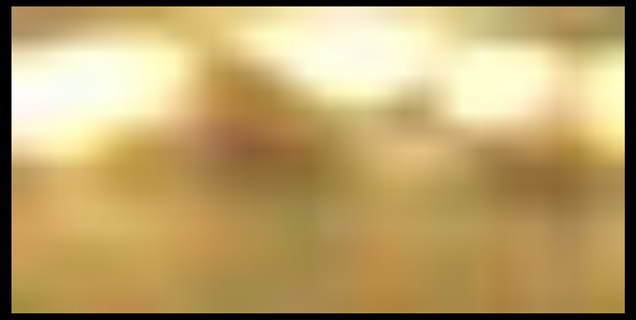

lighting

latitude-longitude representation

captured OLAT

captured OLAT

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

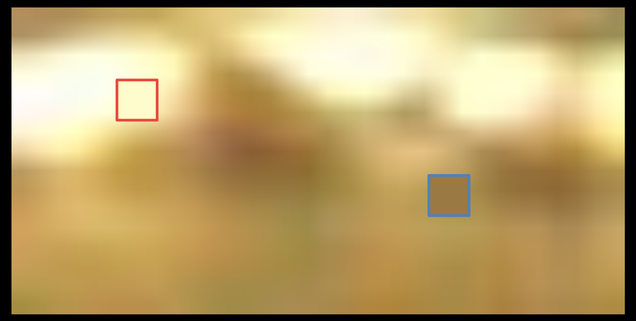

lighting

latitude-longitude representation

captured OLAT

captured OLAT

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

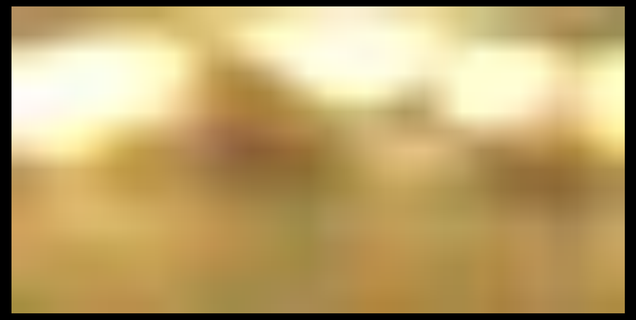

lighting

latitude-longitude representation

captured OLAT

captured OLAT

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

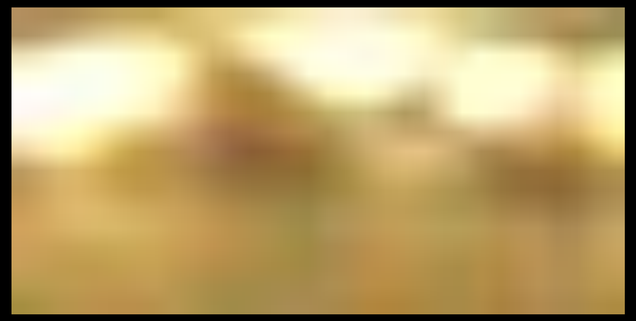

lighting

latitude-longitude representation

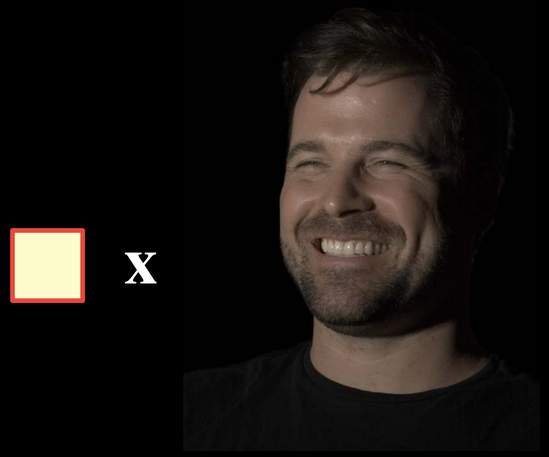

captured OLAT

captured OLAT

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

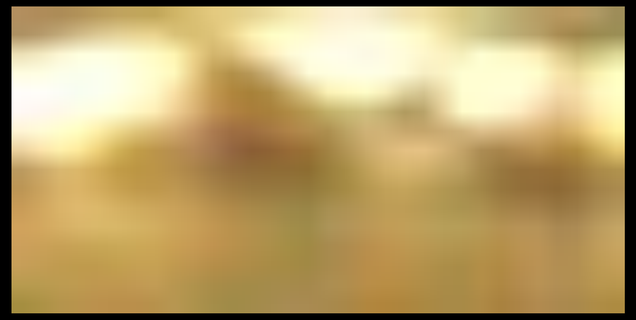

lighting

latitude-longitude representation

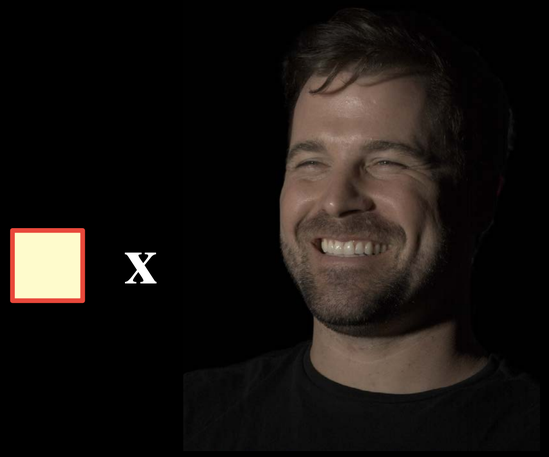

captured OLAT

captured OLAT

relit image (background removed)

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.

Method: Data

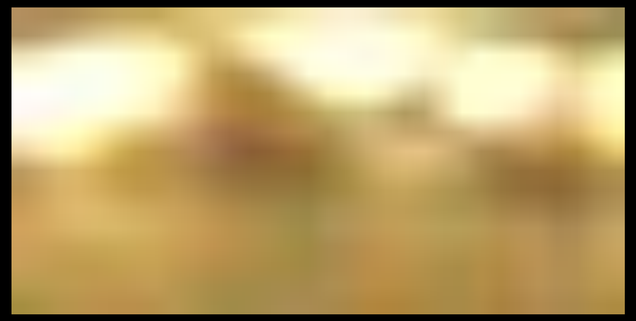

lighting

latitude-longitude representation

captured OLAT

captured OLAT

relit image (background removed)

Debevec, Paul, et al. "Acquiring the reflectance field of a human face." SIGGRAPH 2000.
Wadhwa, Neal, et al. "Synthetic depth-of-field with a single-camera mobile phone." SIGGRAPH 2018

Method: Data

Method: Data

- OLAT images
  - 22 people (18 training, 4 validation), each 3~5 facial expressions

Method: Data

- OLAT images
  - 22 people (18 training, 4 validation), each 3~5 facial expressions
  - Each OLAT captured with 7 cameras in 6 seconds.

Method: Data

- OLAT images
  - 22 people (18 training, 4 validation), each 3~5 facial expressions
  - Each OLAT captured with 7 cameras in 6 seconds.
- HDR lighting environments

Method: Data

- OLAT images
  - 22 people (18 training, 4 validation), each 3~5 facial expressions
  - Each OLAT captured with 7 cameras in 6 seconds.
- HDR lighting environments
  - ~2000 indoor HDR lighting from Laval Dataset

Method: Data

- OLAT images
  - 22 people (18 training, 4 validation), each 3~5 facial expressions
  - Each OLAT captured with 7 cameras in 6 seconds.
- HDR lighting environments
  - ~2000 indoor HDR lighting from Laval Dataset
  - ~1000 outdoor HDR lighting from the web

Method: Data

- OLAT images
  - 22 people (18 training, 4 validation), each 3~5 facial expressions
  - Each OLAT captured with 7 cameras in 6 seconds.
- HDR lighting environments
  - ~2000 indoor HDR lighting from Laval Dataset
  - ~1000 outdoor HDR lighting from the web

- Total: 226,800 portrait and lighting pairs for training

Method: Training

Method: Training

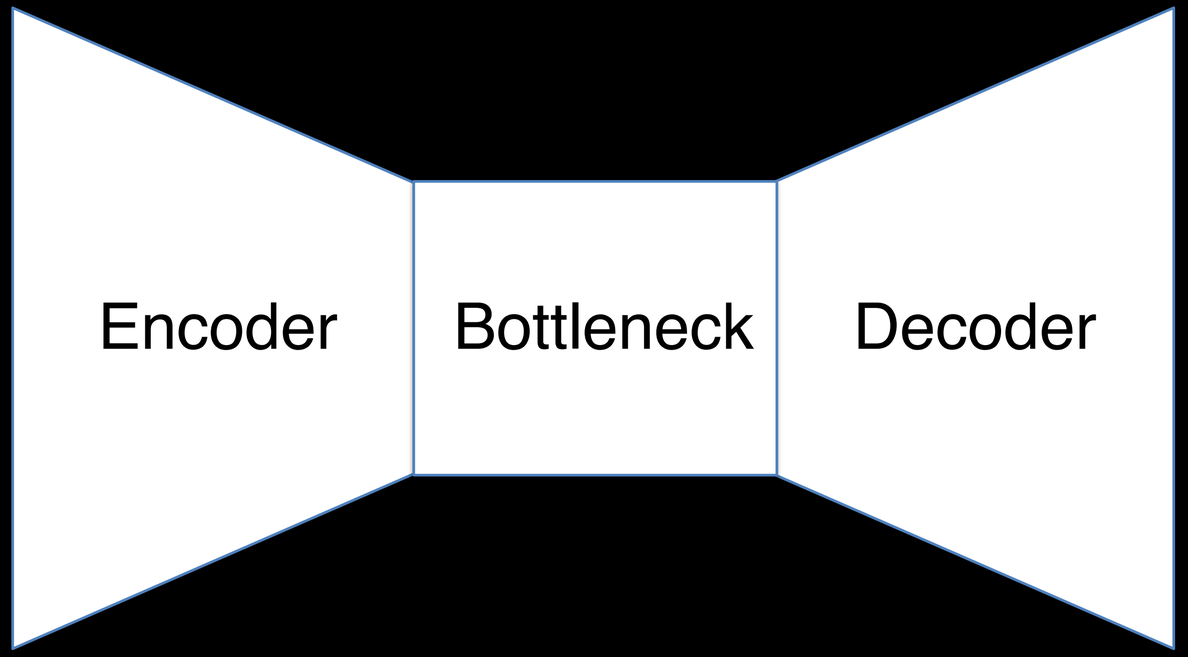

Method: Training

- Task 1: Complete relighting

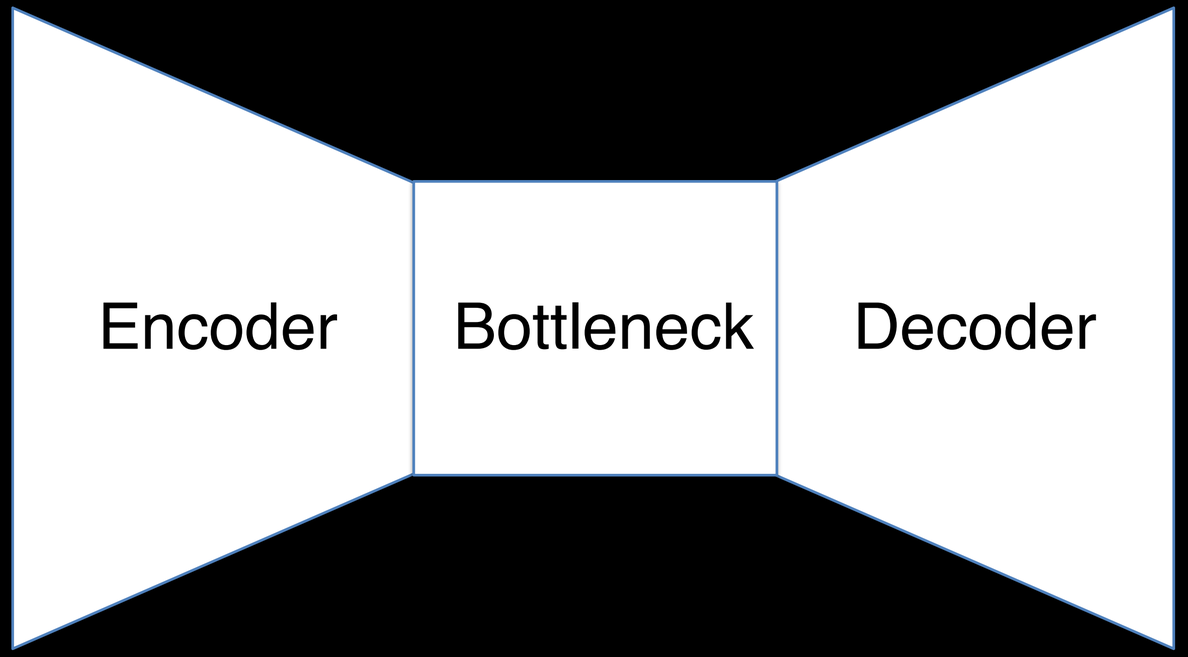

Method: Training

- Task 1: Complete relighting

source image

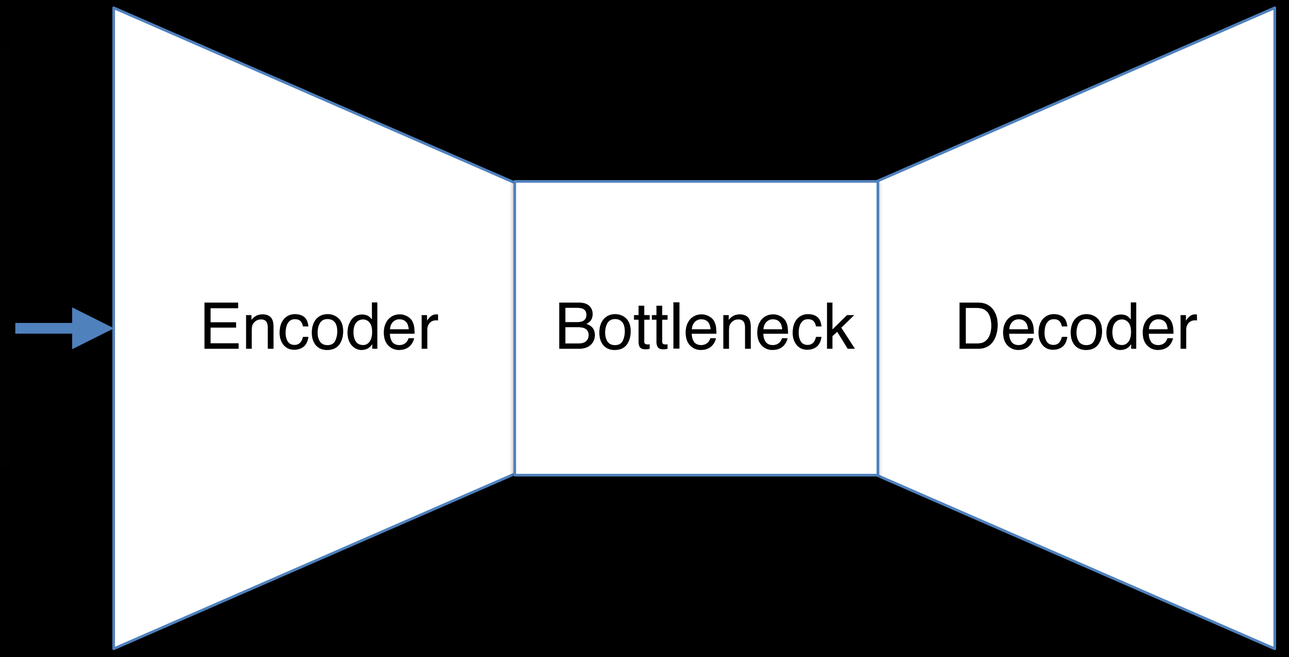

Method: Training

- Task 1: Complete relighting

source image

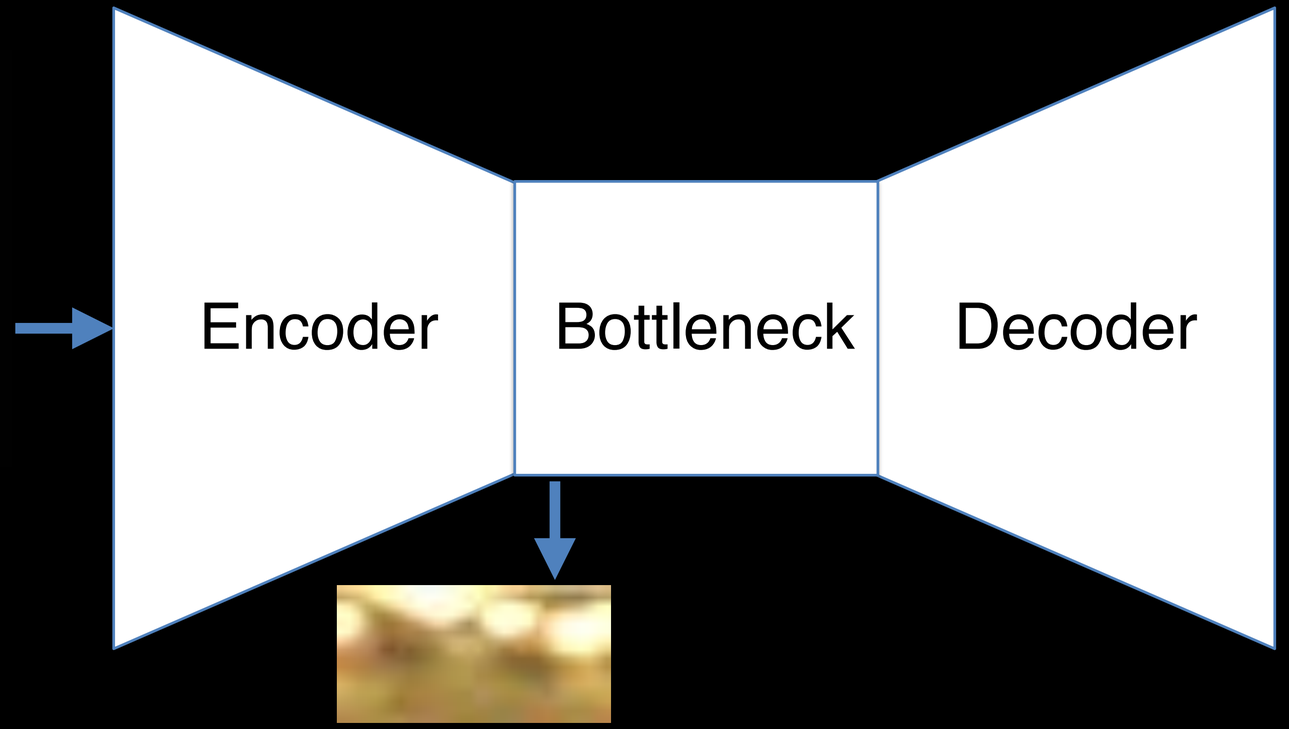

source light

Method: Training

- Task 1: Complete relighting

source image

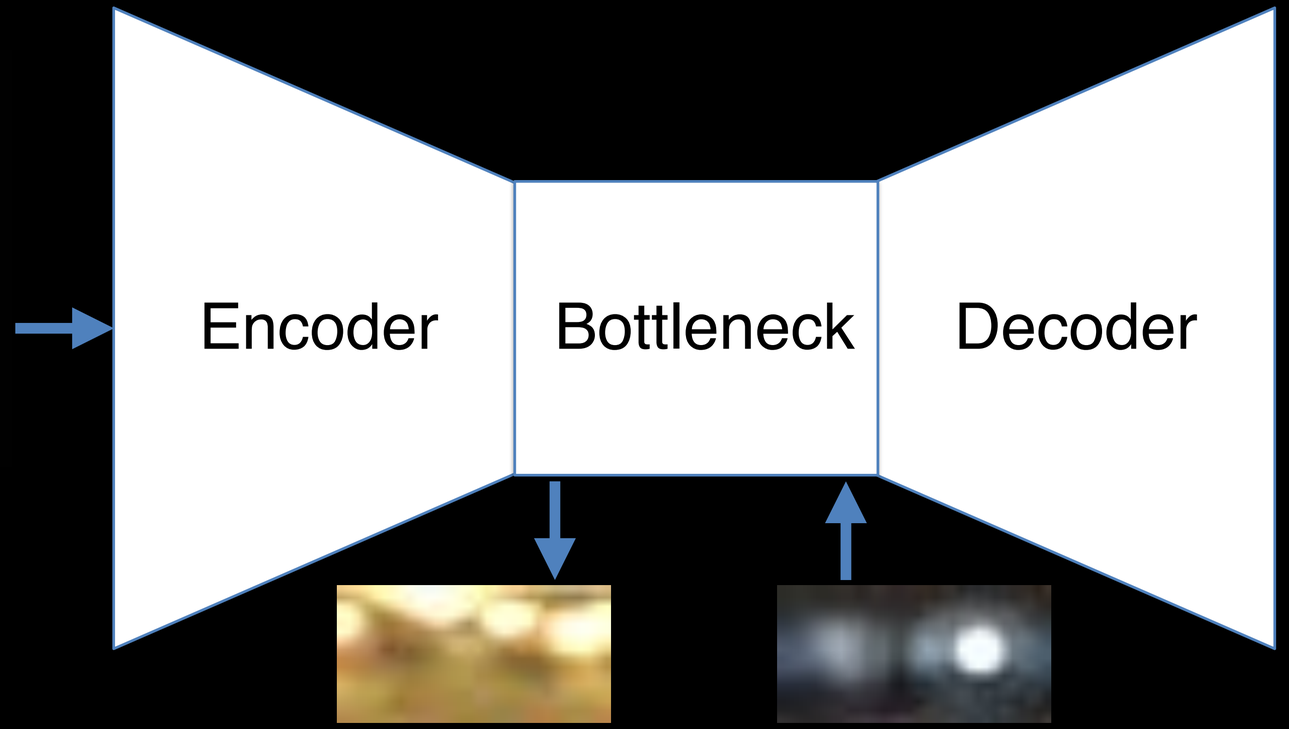

source light

target light

Method: Training

- Task 1: Complete relighting

source image

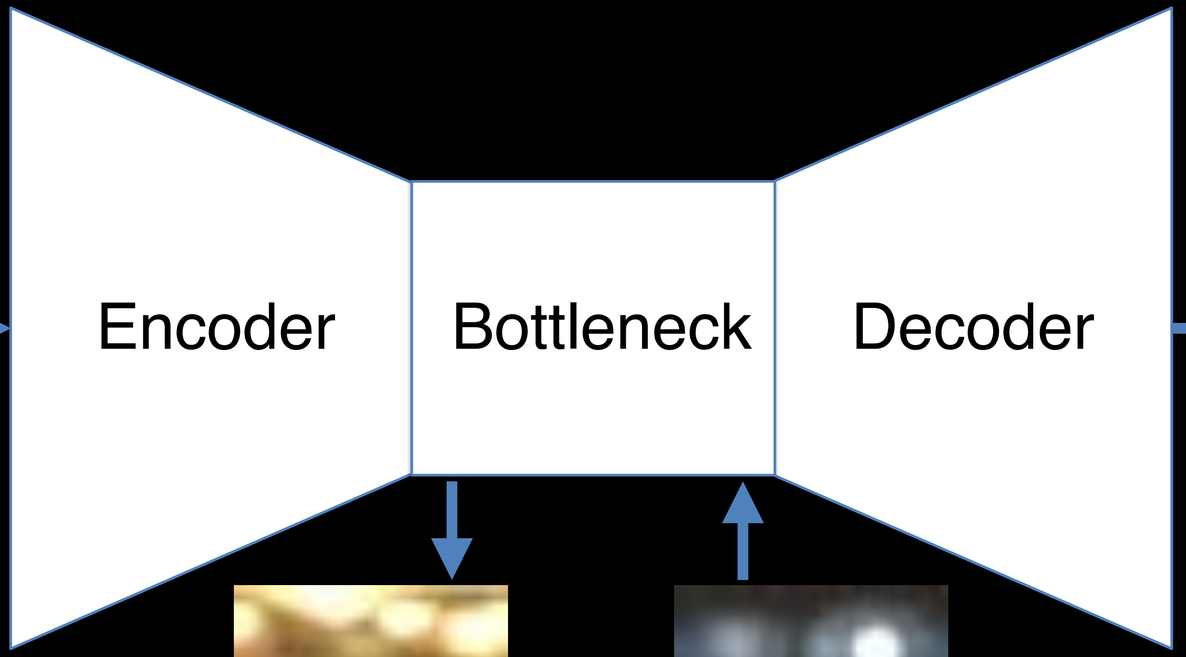

source light

target light

target image

Method: Training

- Task 1: Complete relighting

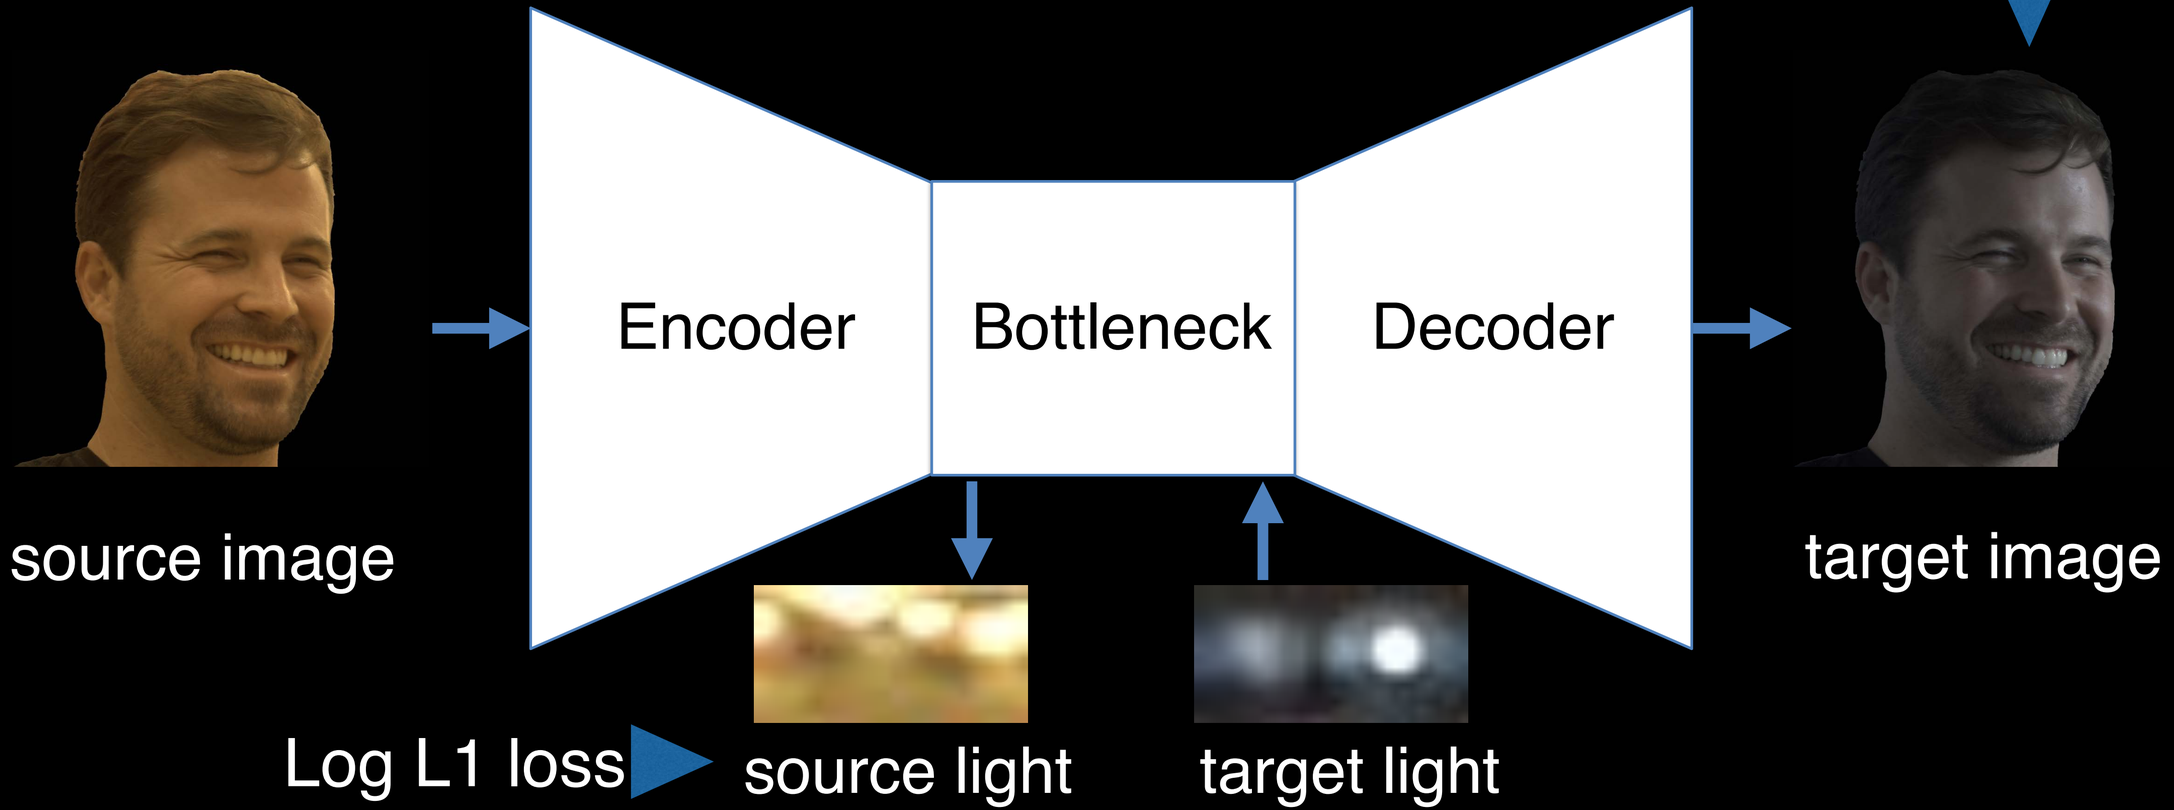

Method: Training

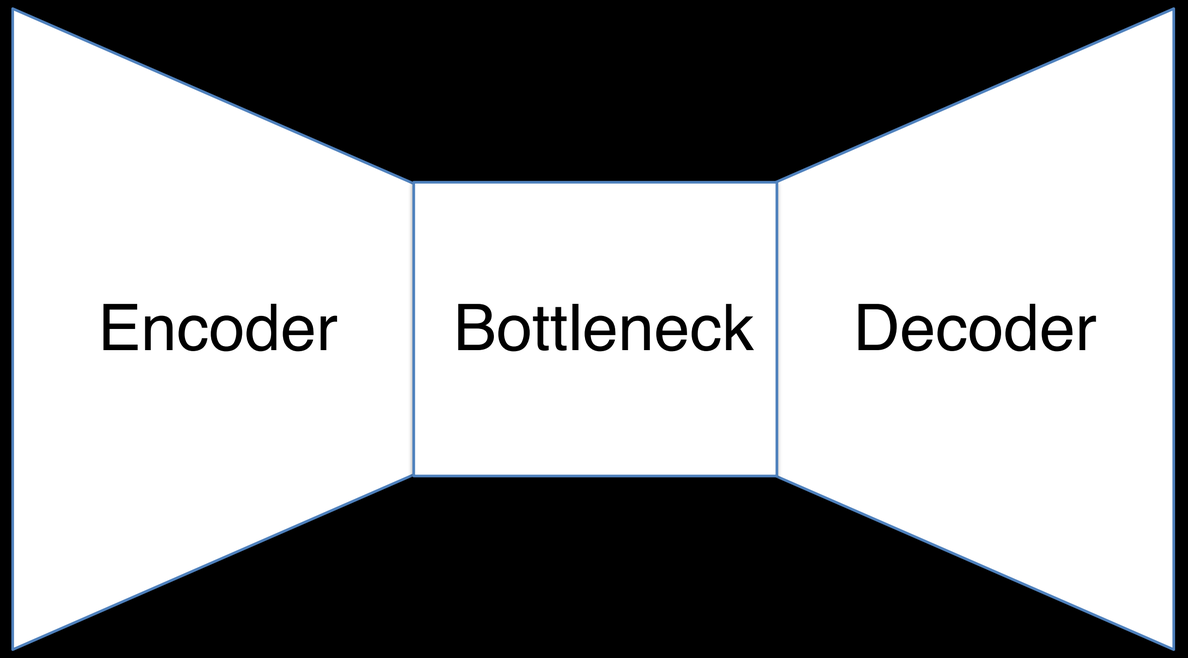

Method: Training

- Task 2: Illumination retargeting

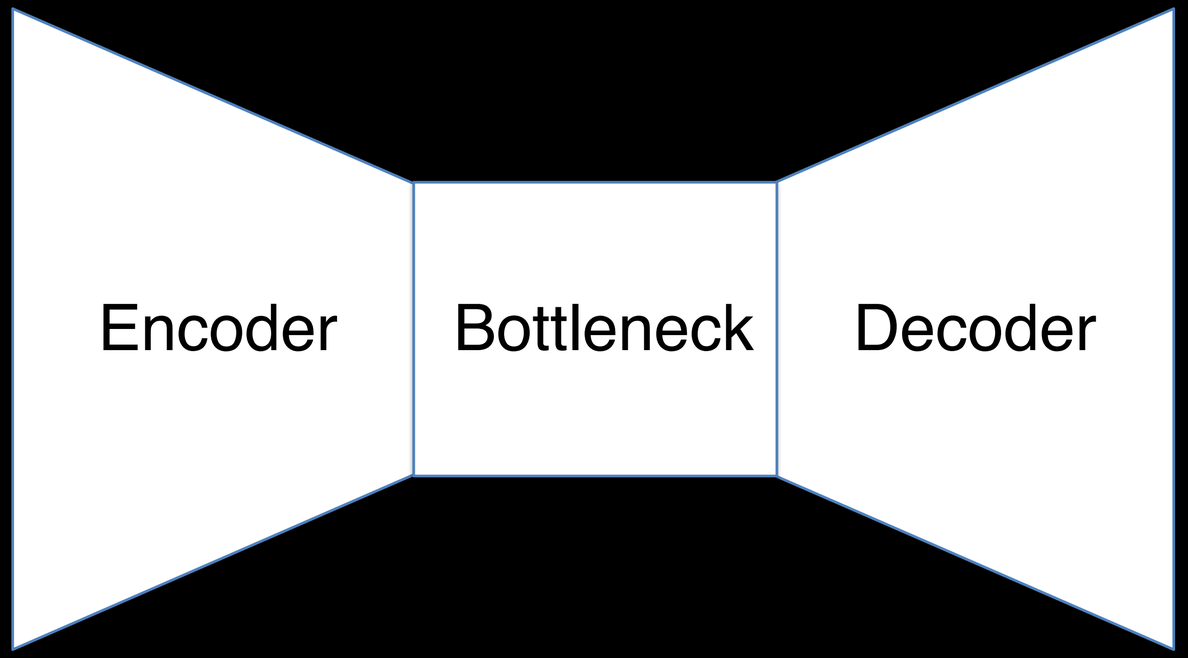

Method: Training

- Task 2: Illumination retargeting

source image

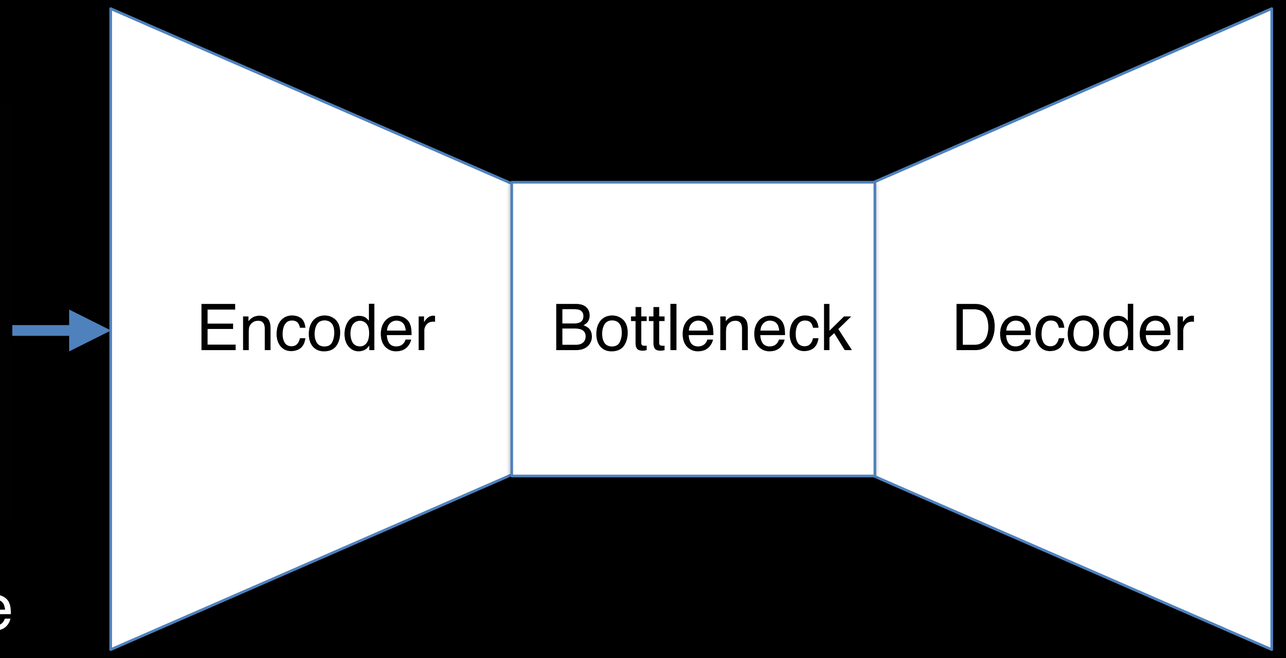

Method: Training

- Task 2: Illumination retargeting

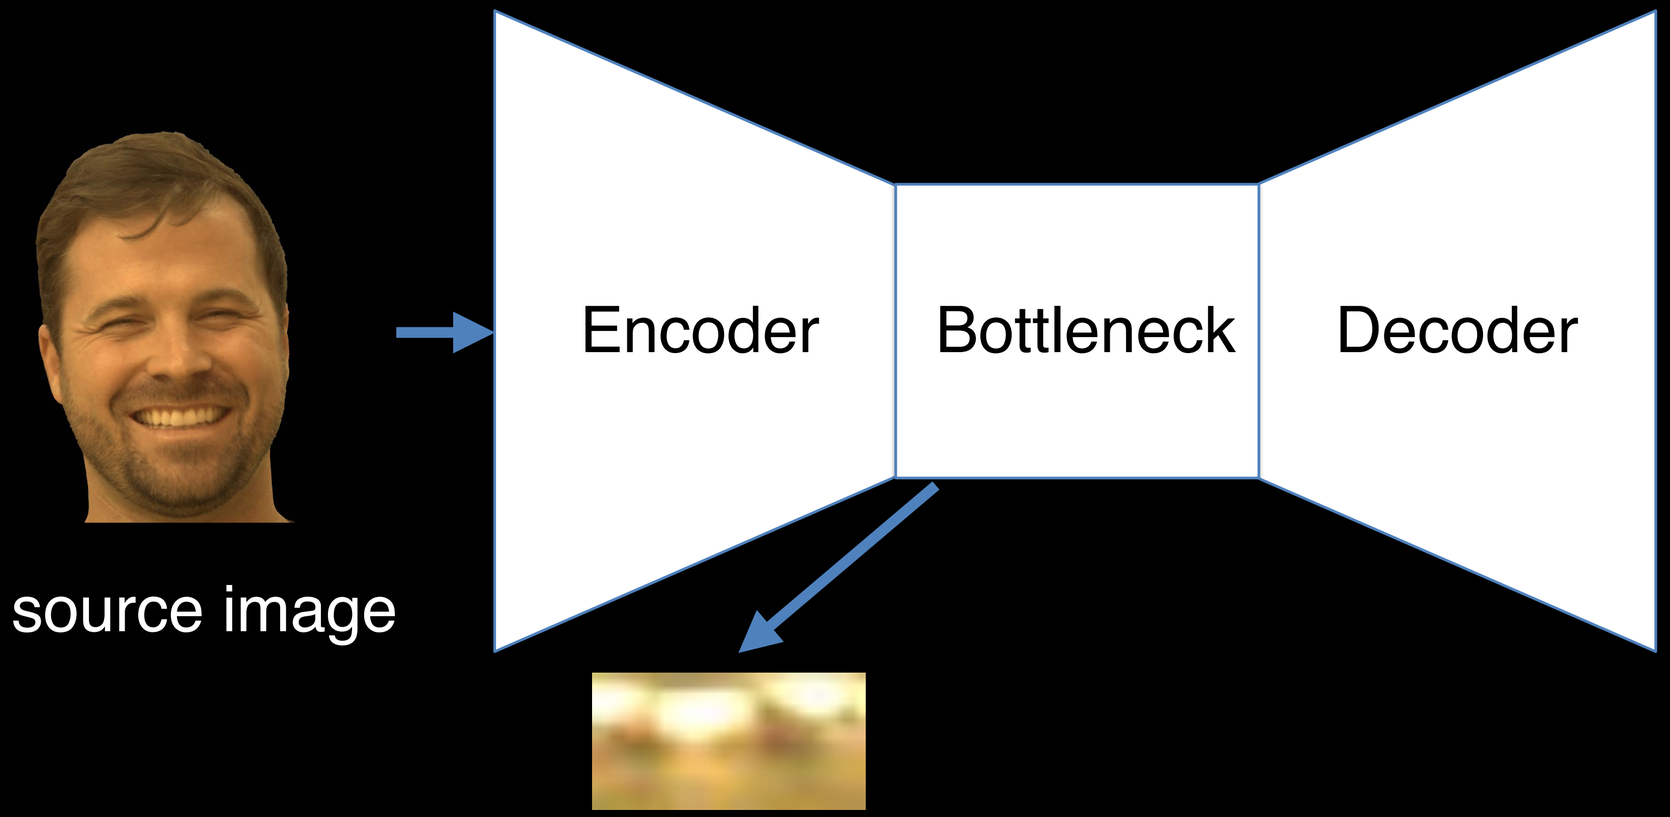

Method: Training

- Task 2: Illumination retargeting

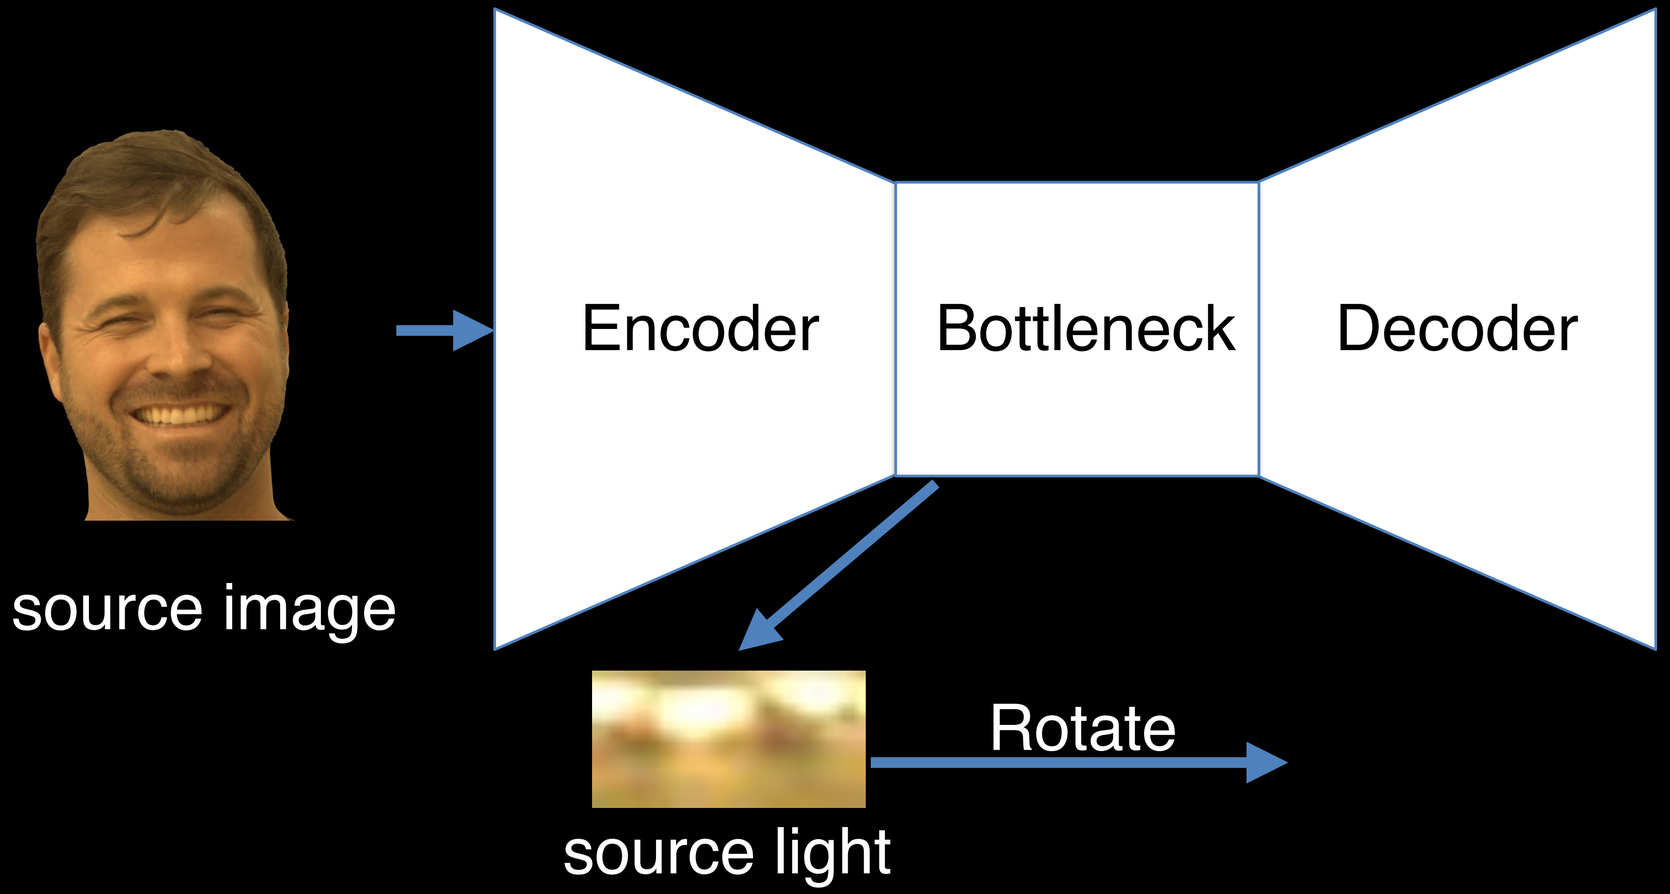

Method: Training

- Task 2: Illumination retargeting

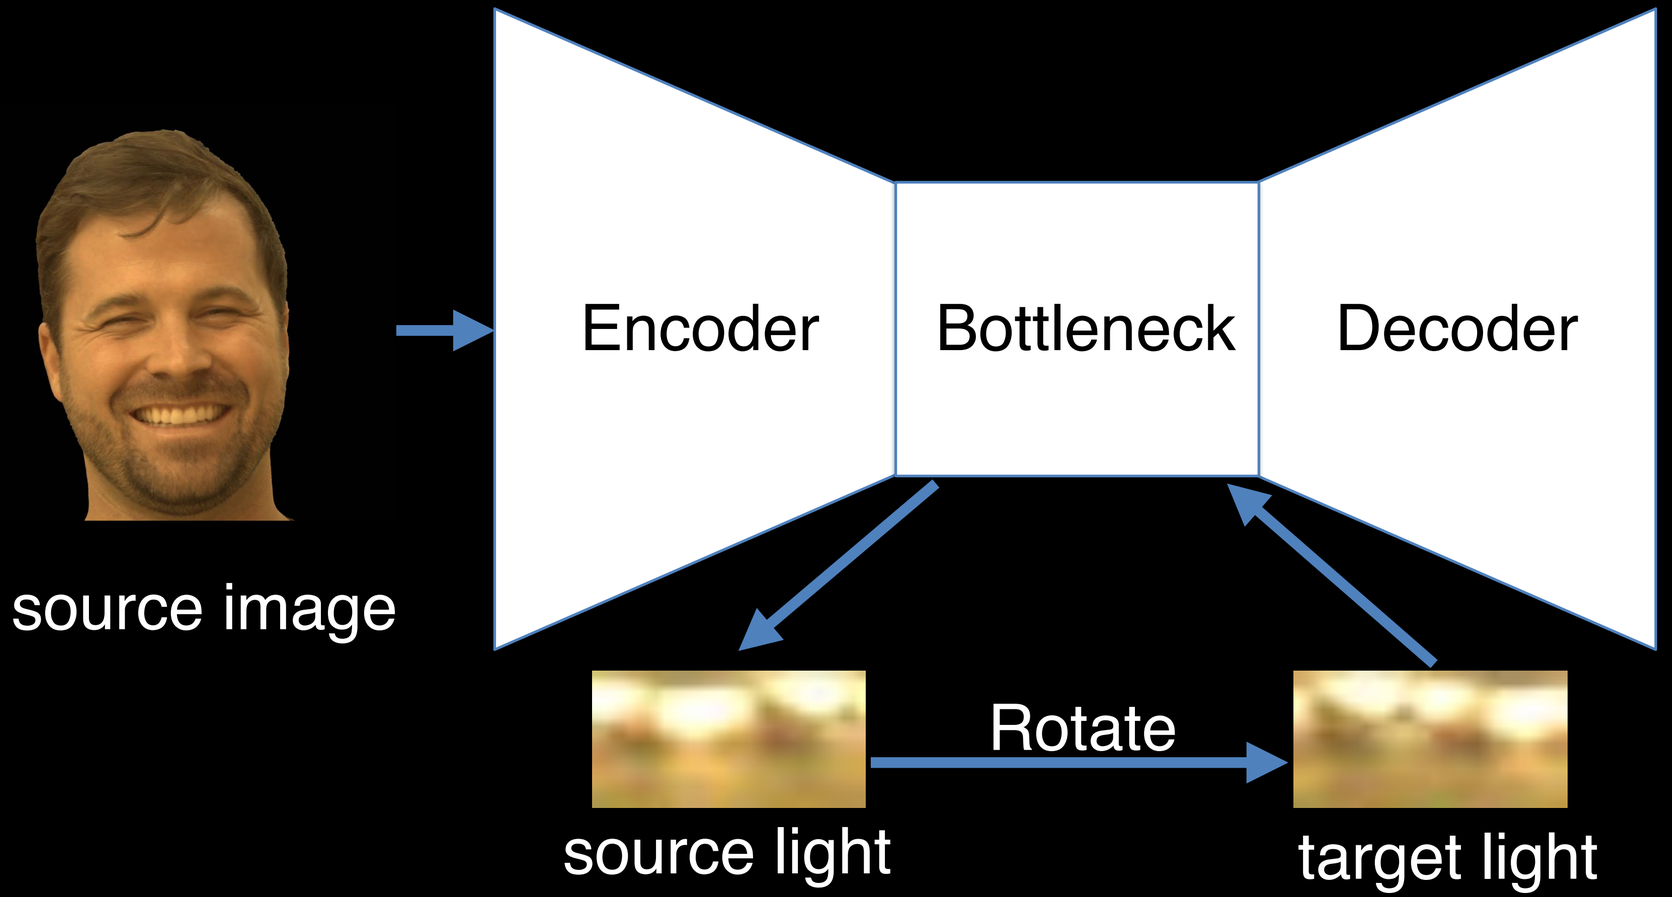

Method: Training

- Task 2: Illumination retargeting

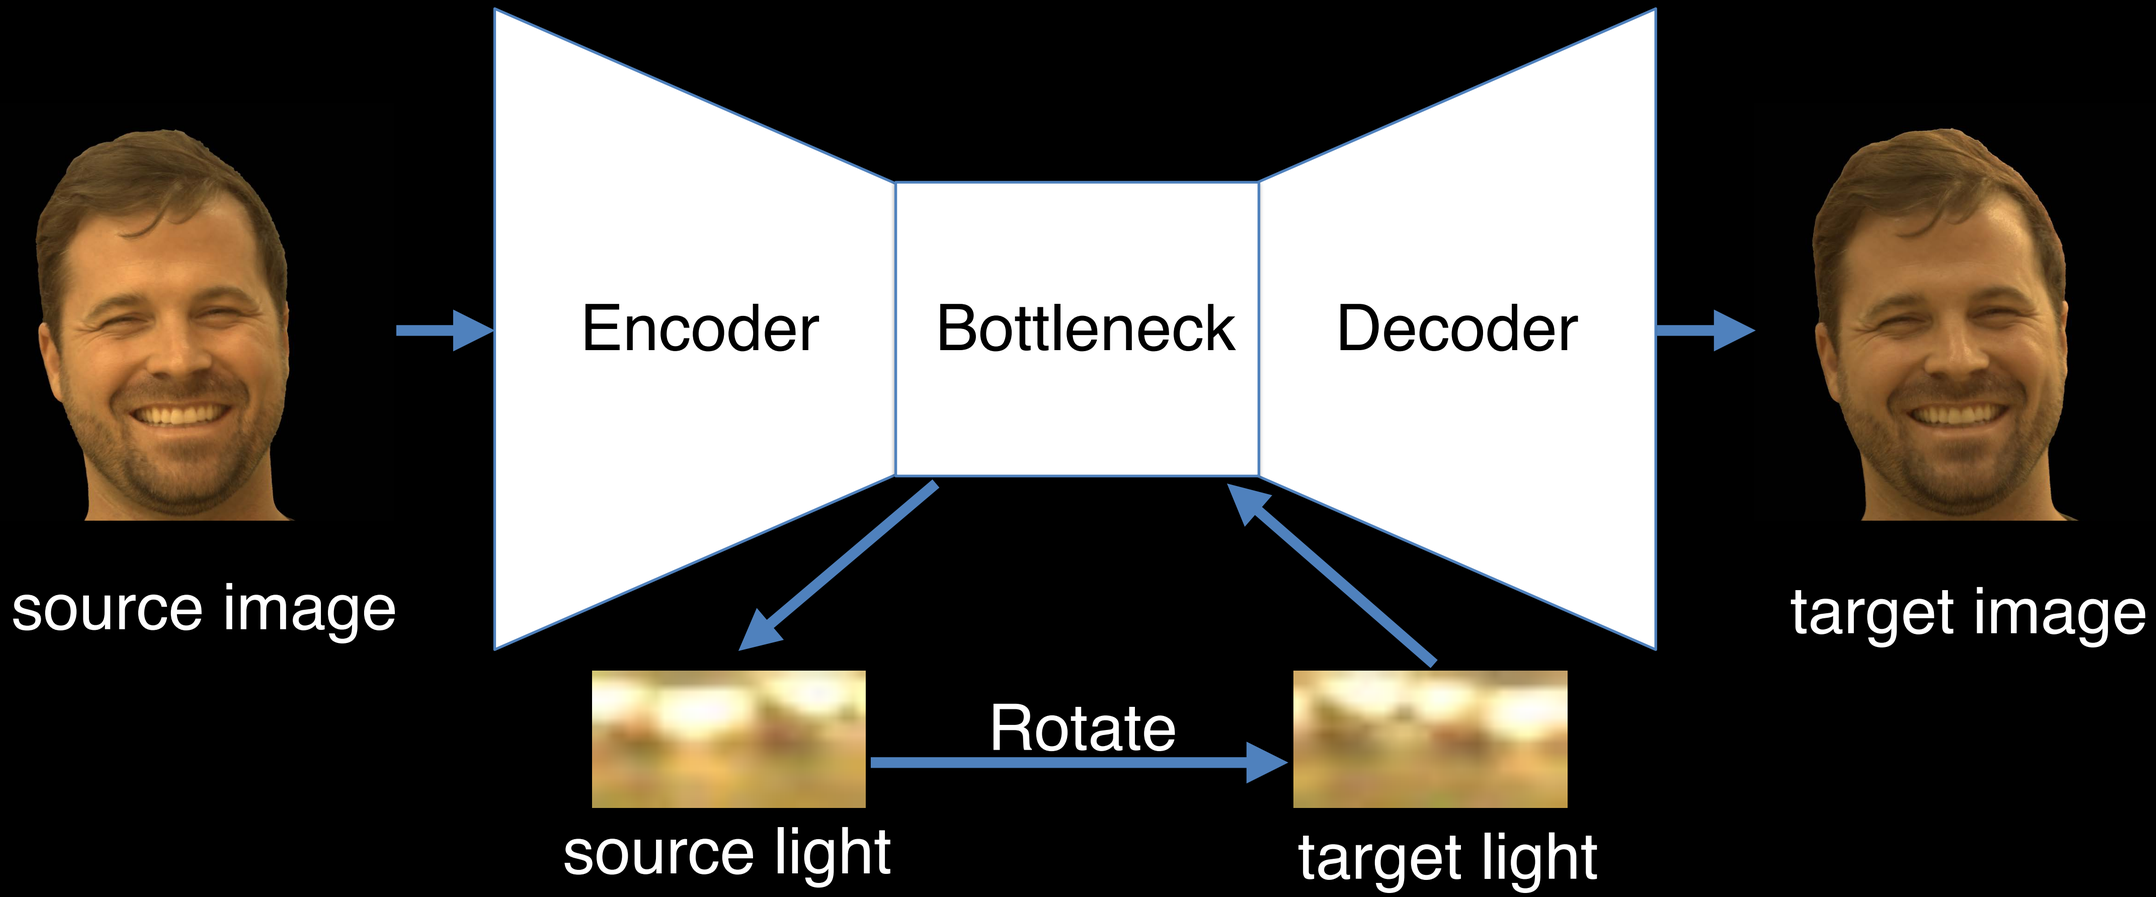

Method: Training

- Task 2: Illumination retargeting

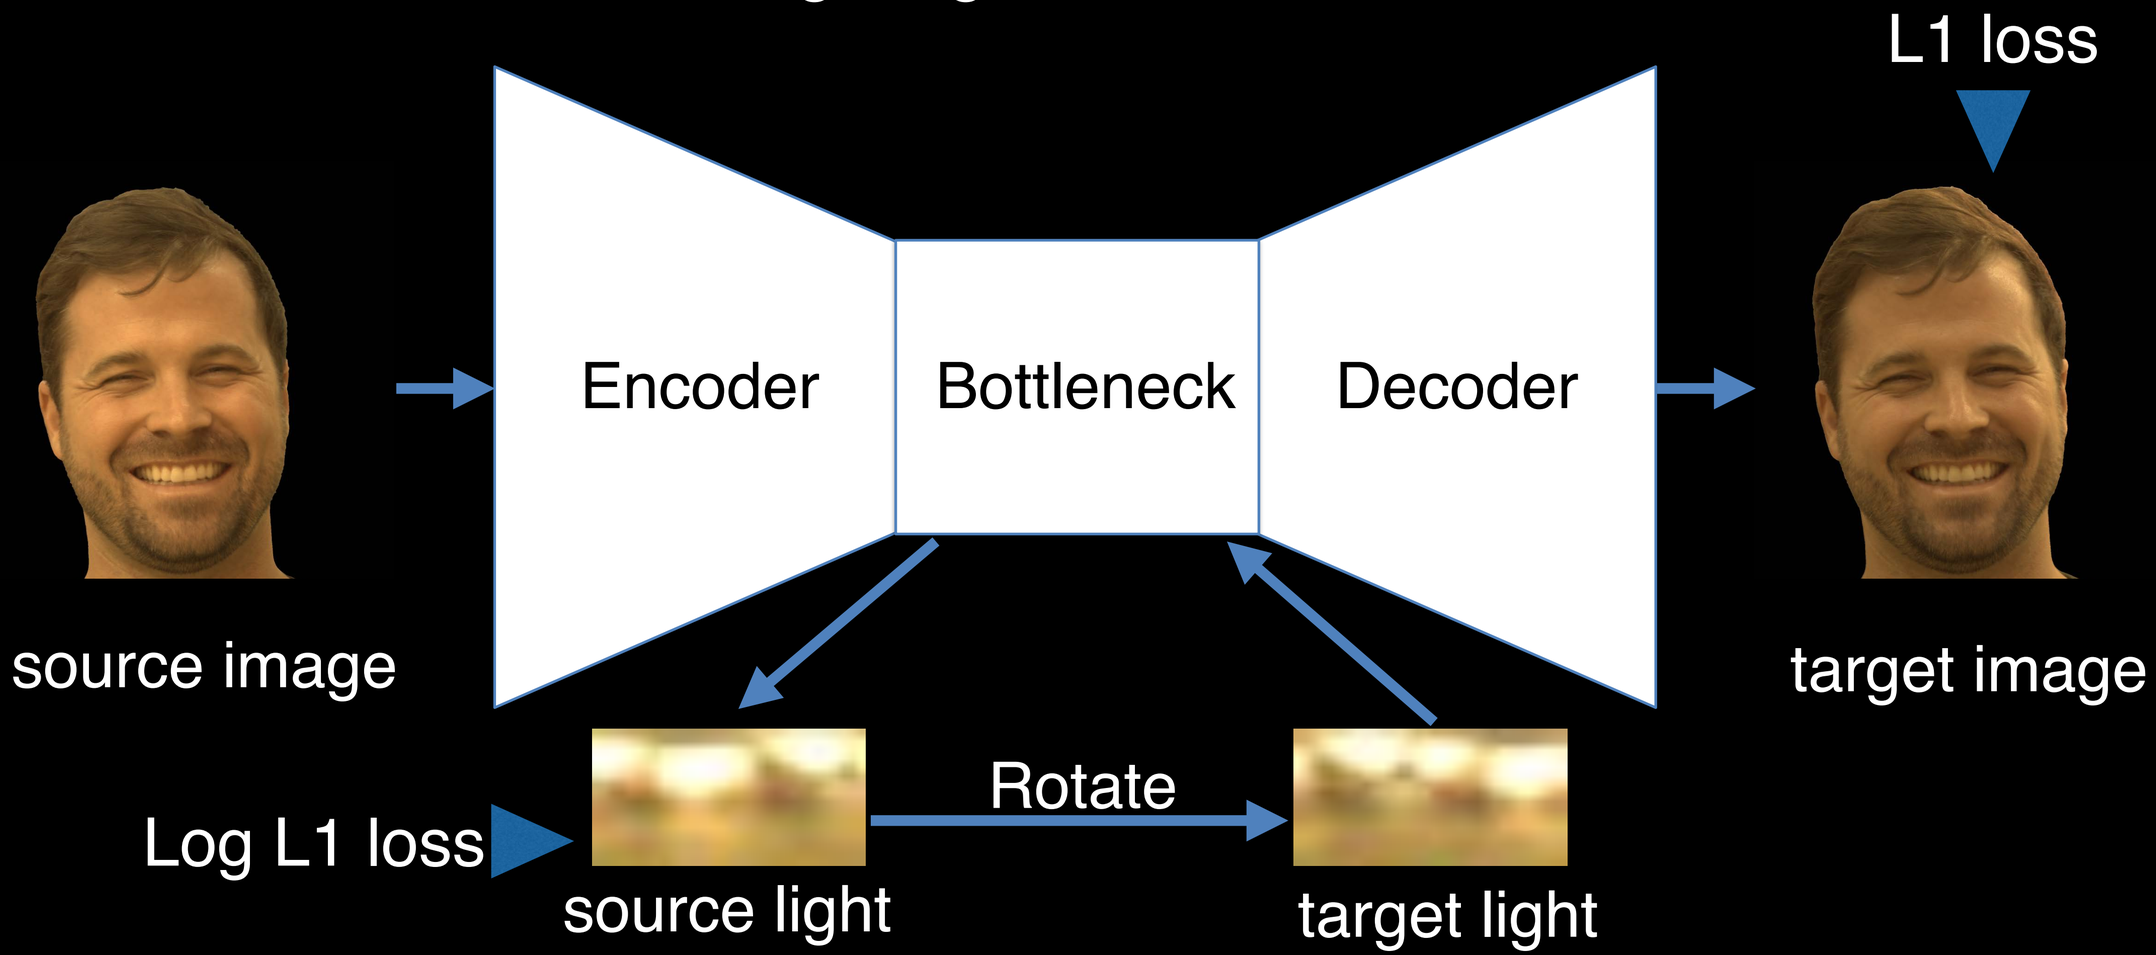

Method: Training

Method: Training

- Network structure
  - U-Net

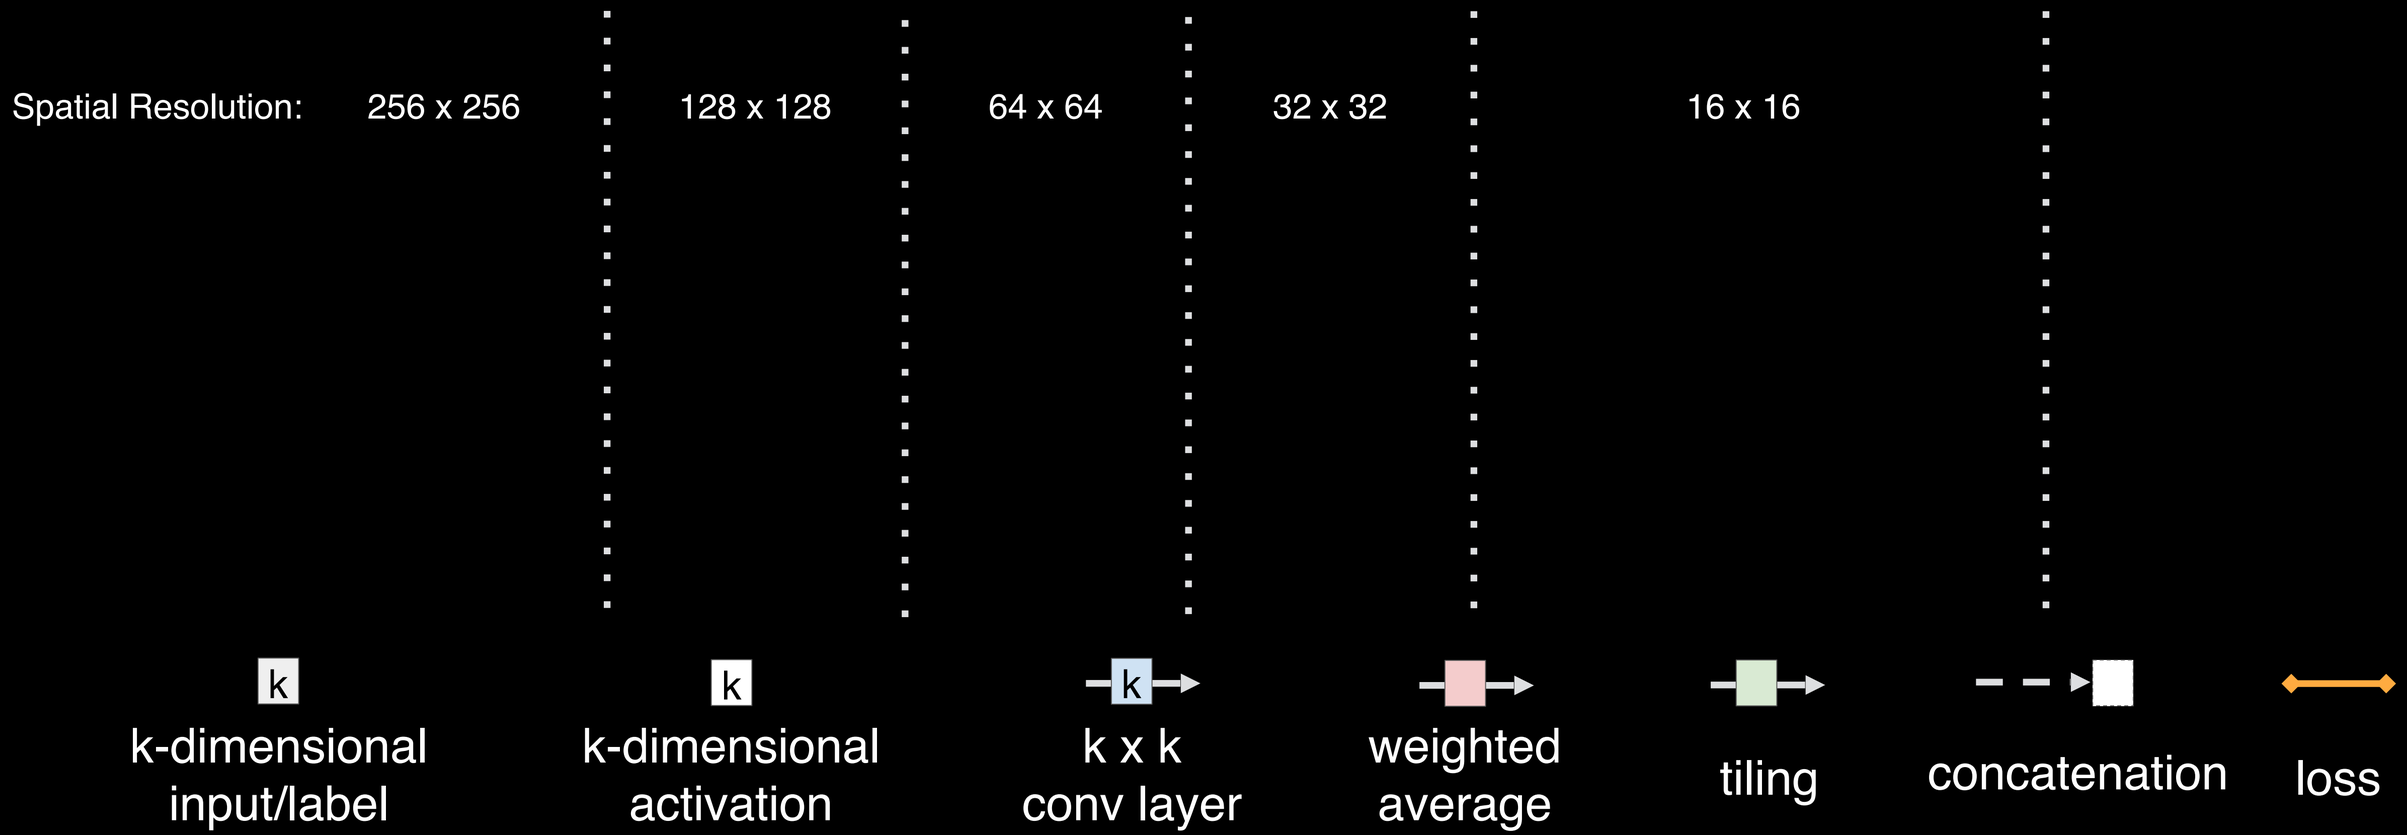

Method: Training

- Network structure
  - U-Net

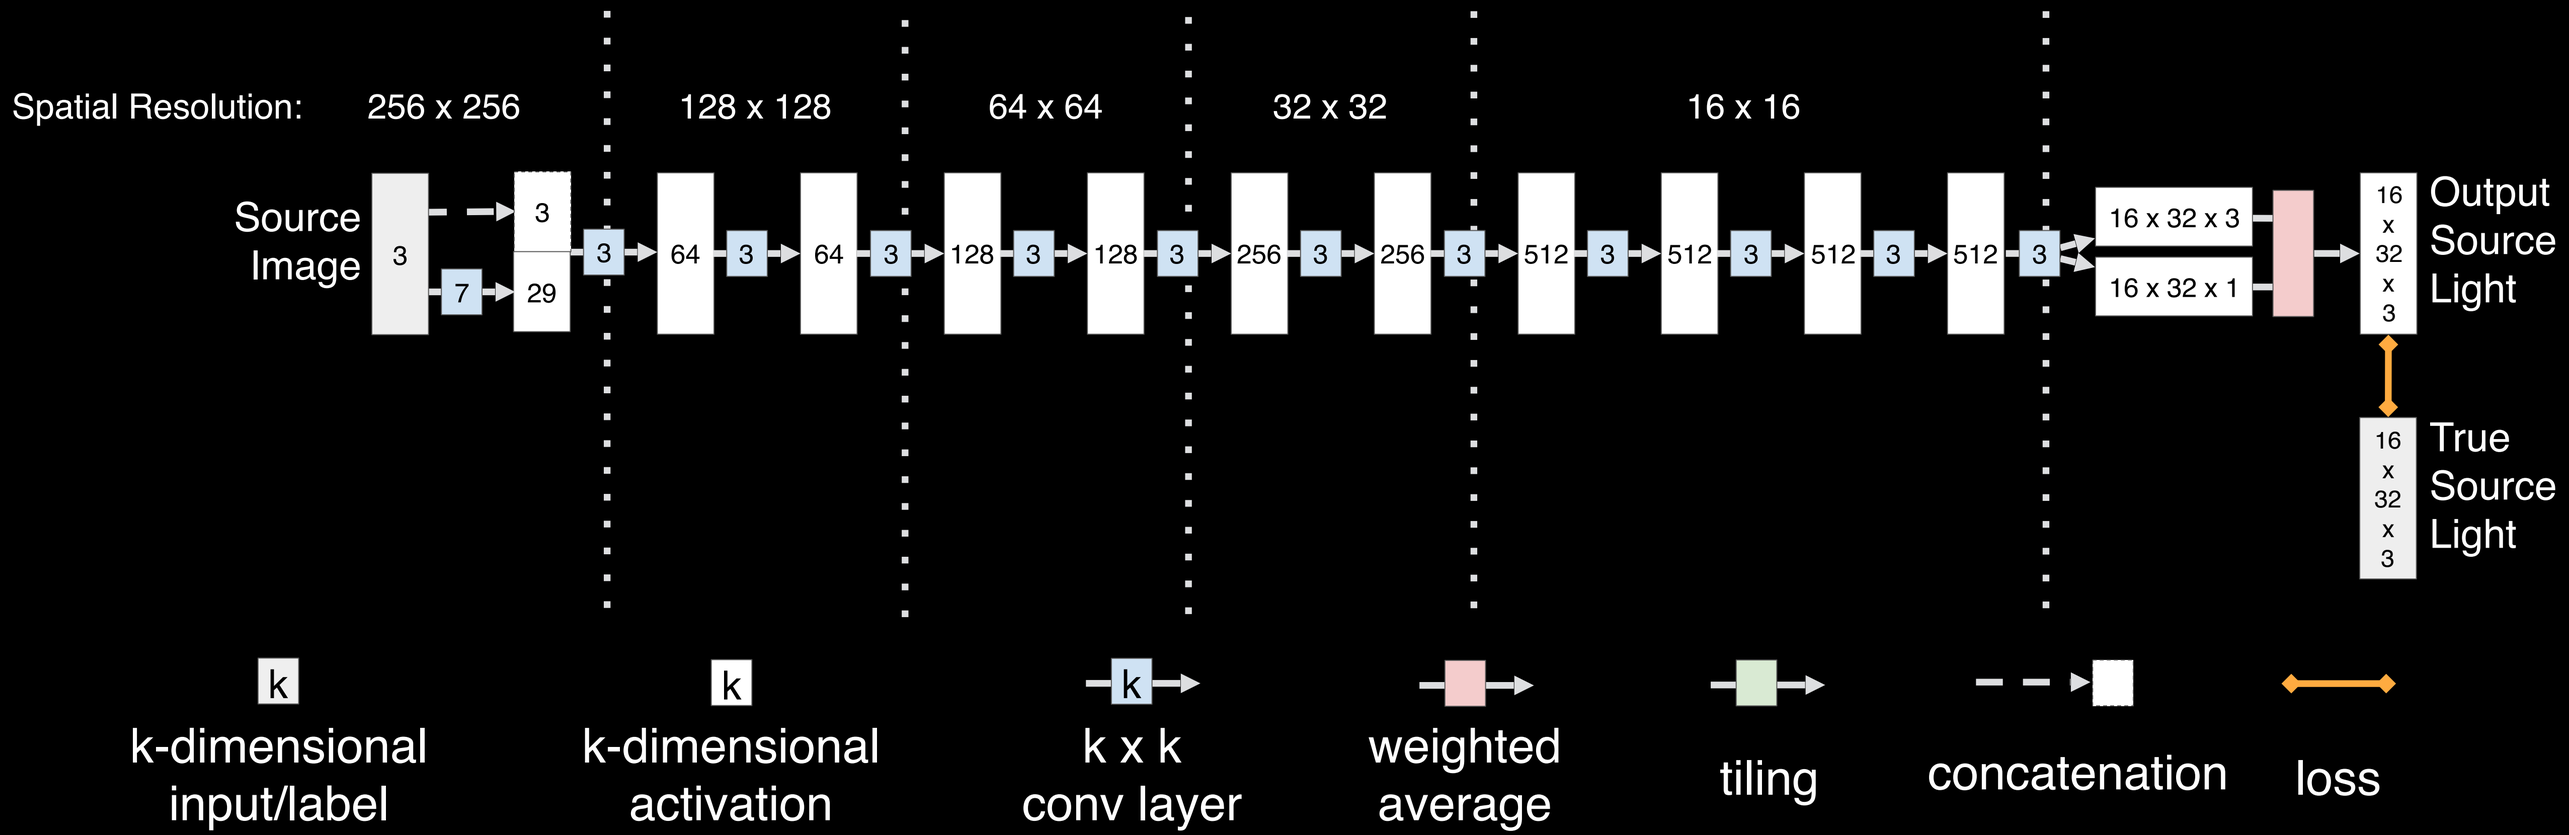

Method: Training

- Network structure
  - U-Net
  - Predict and feed in light at bottleneck

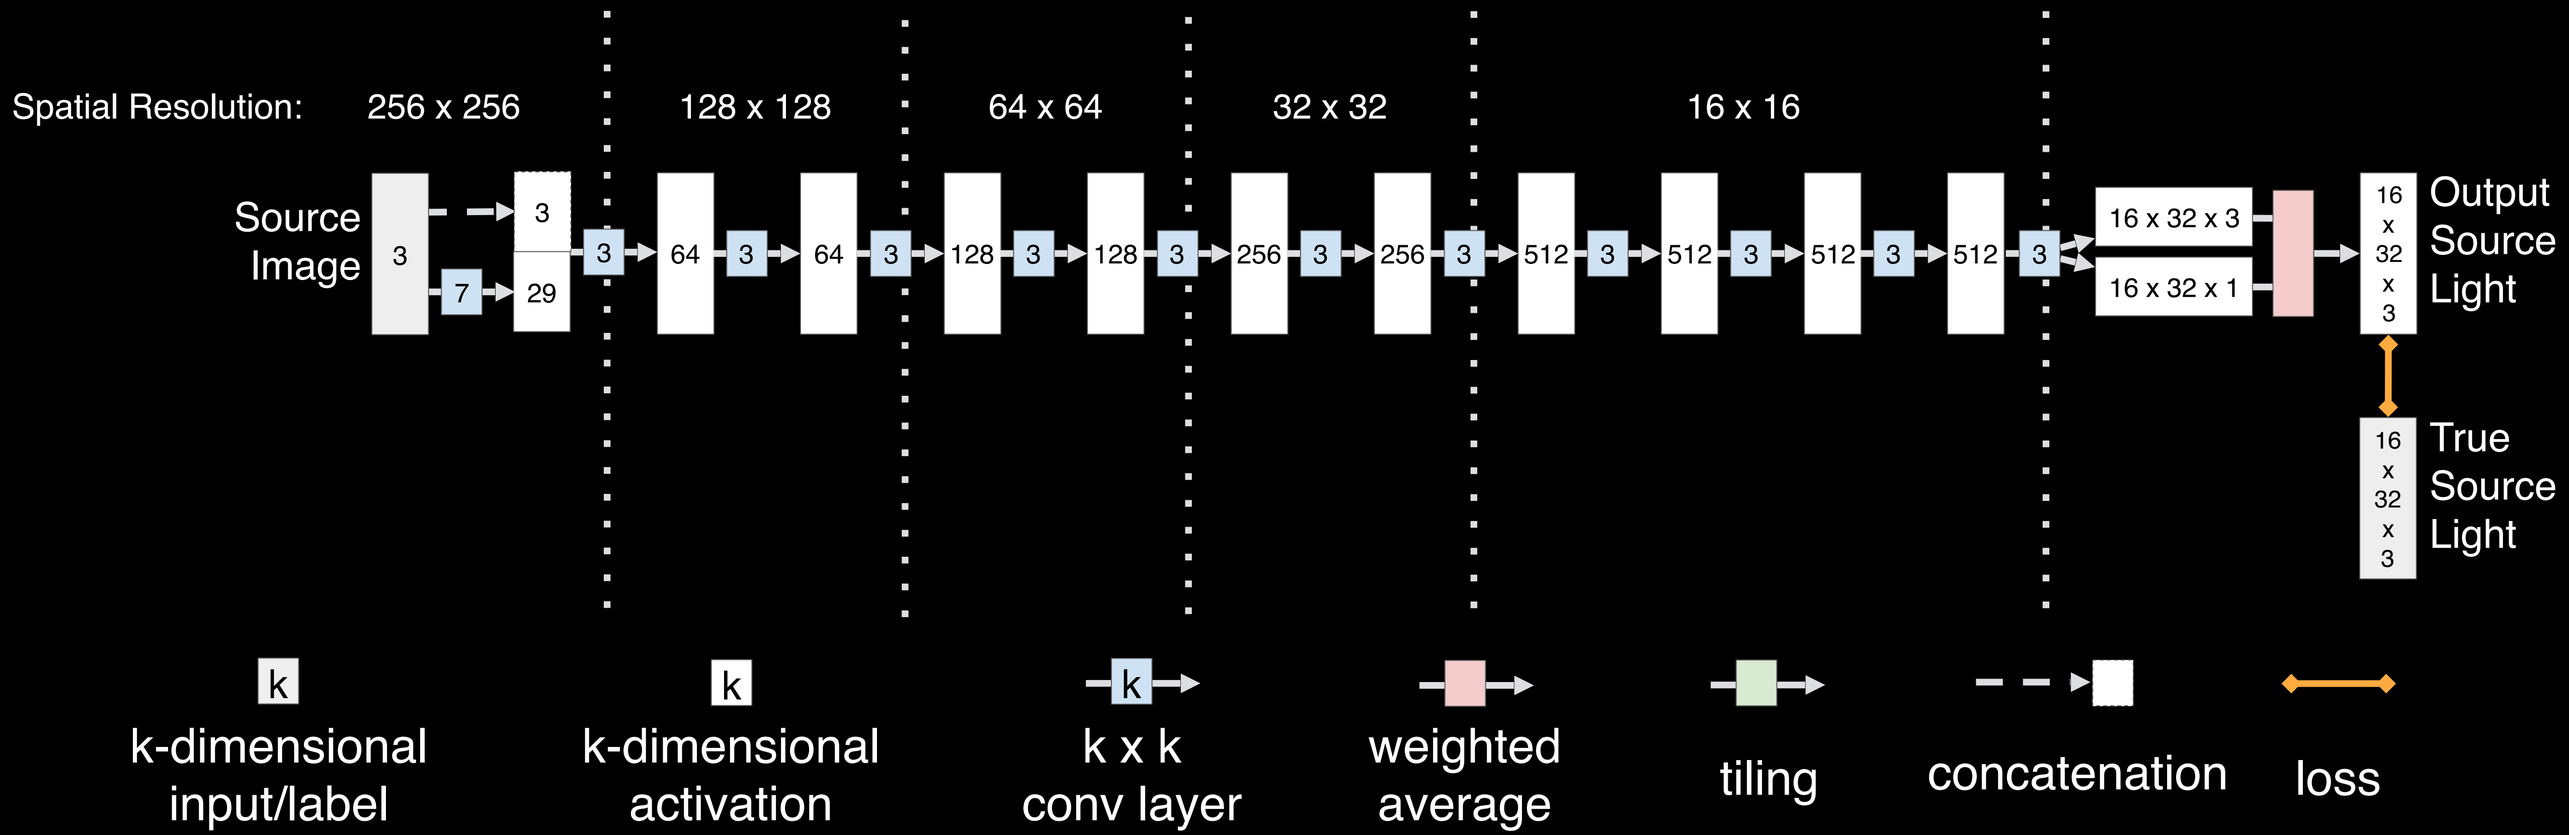

Method: Training

- Network structure
  - U-Net
  - Predict and feed in light at bottleneck

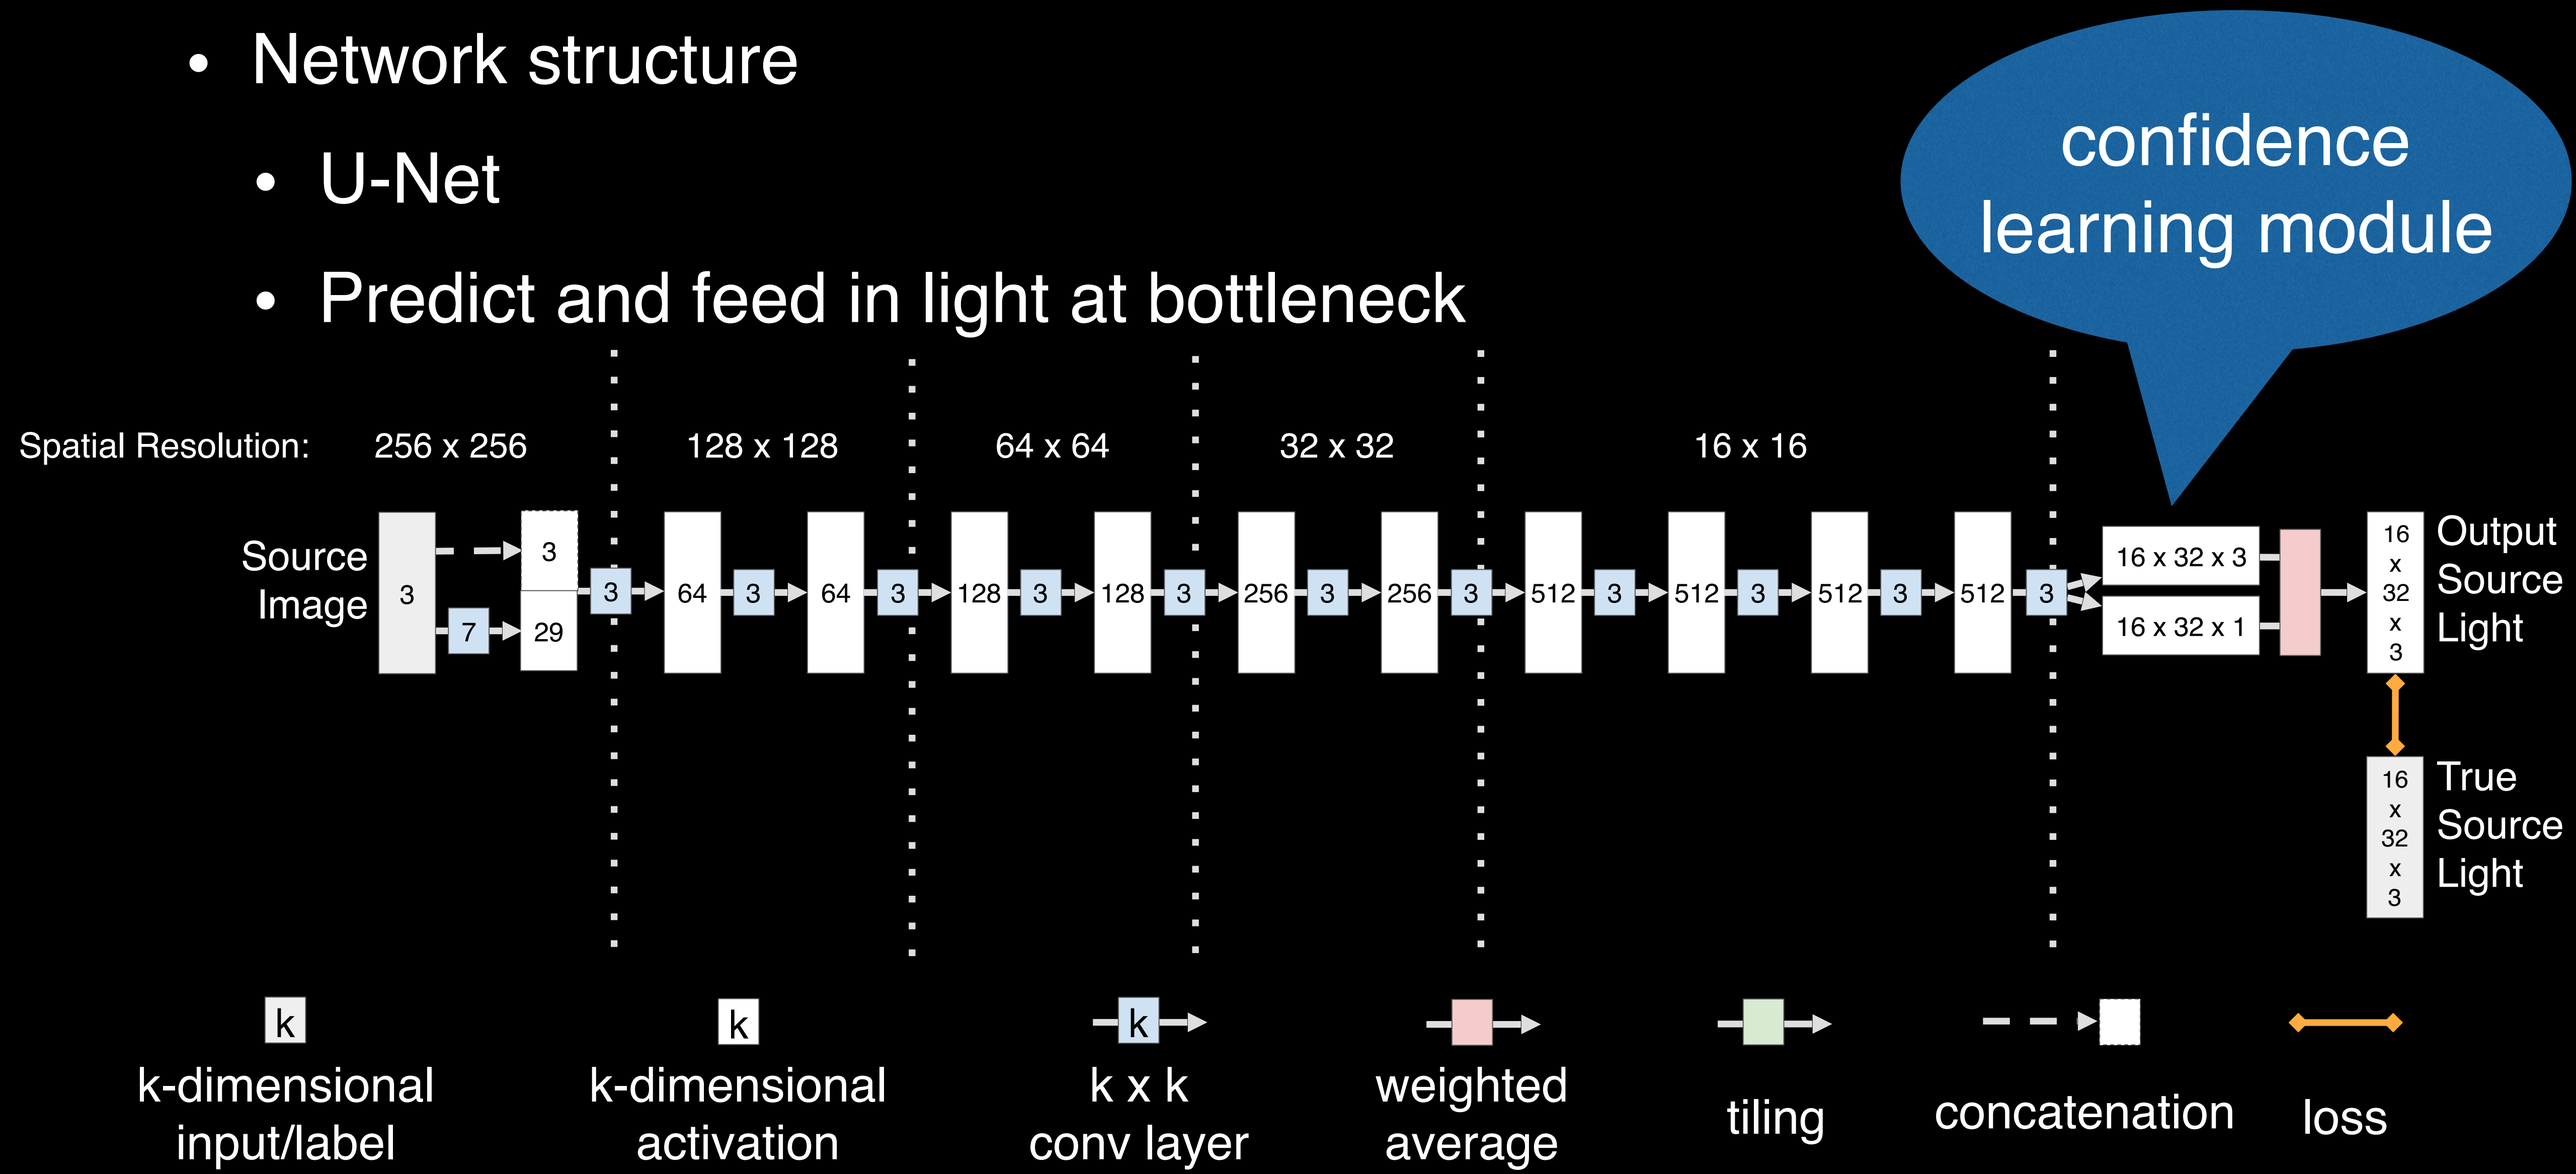

Method: Training

- Network structure
  - U-Net
  - Predict and feed in light at bottleneck

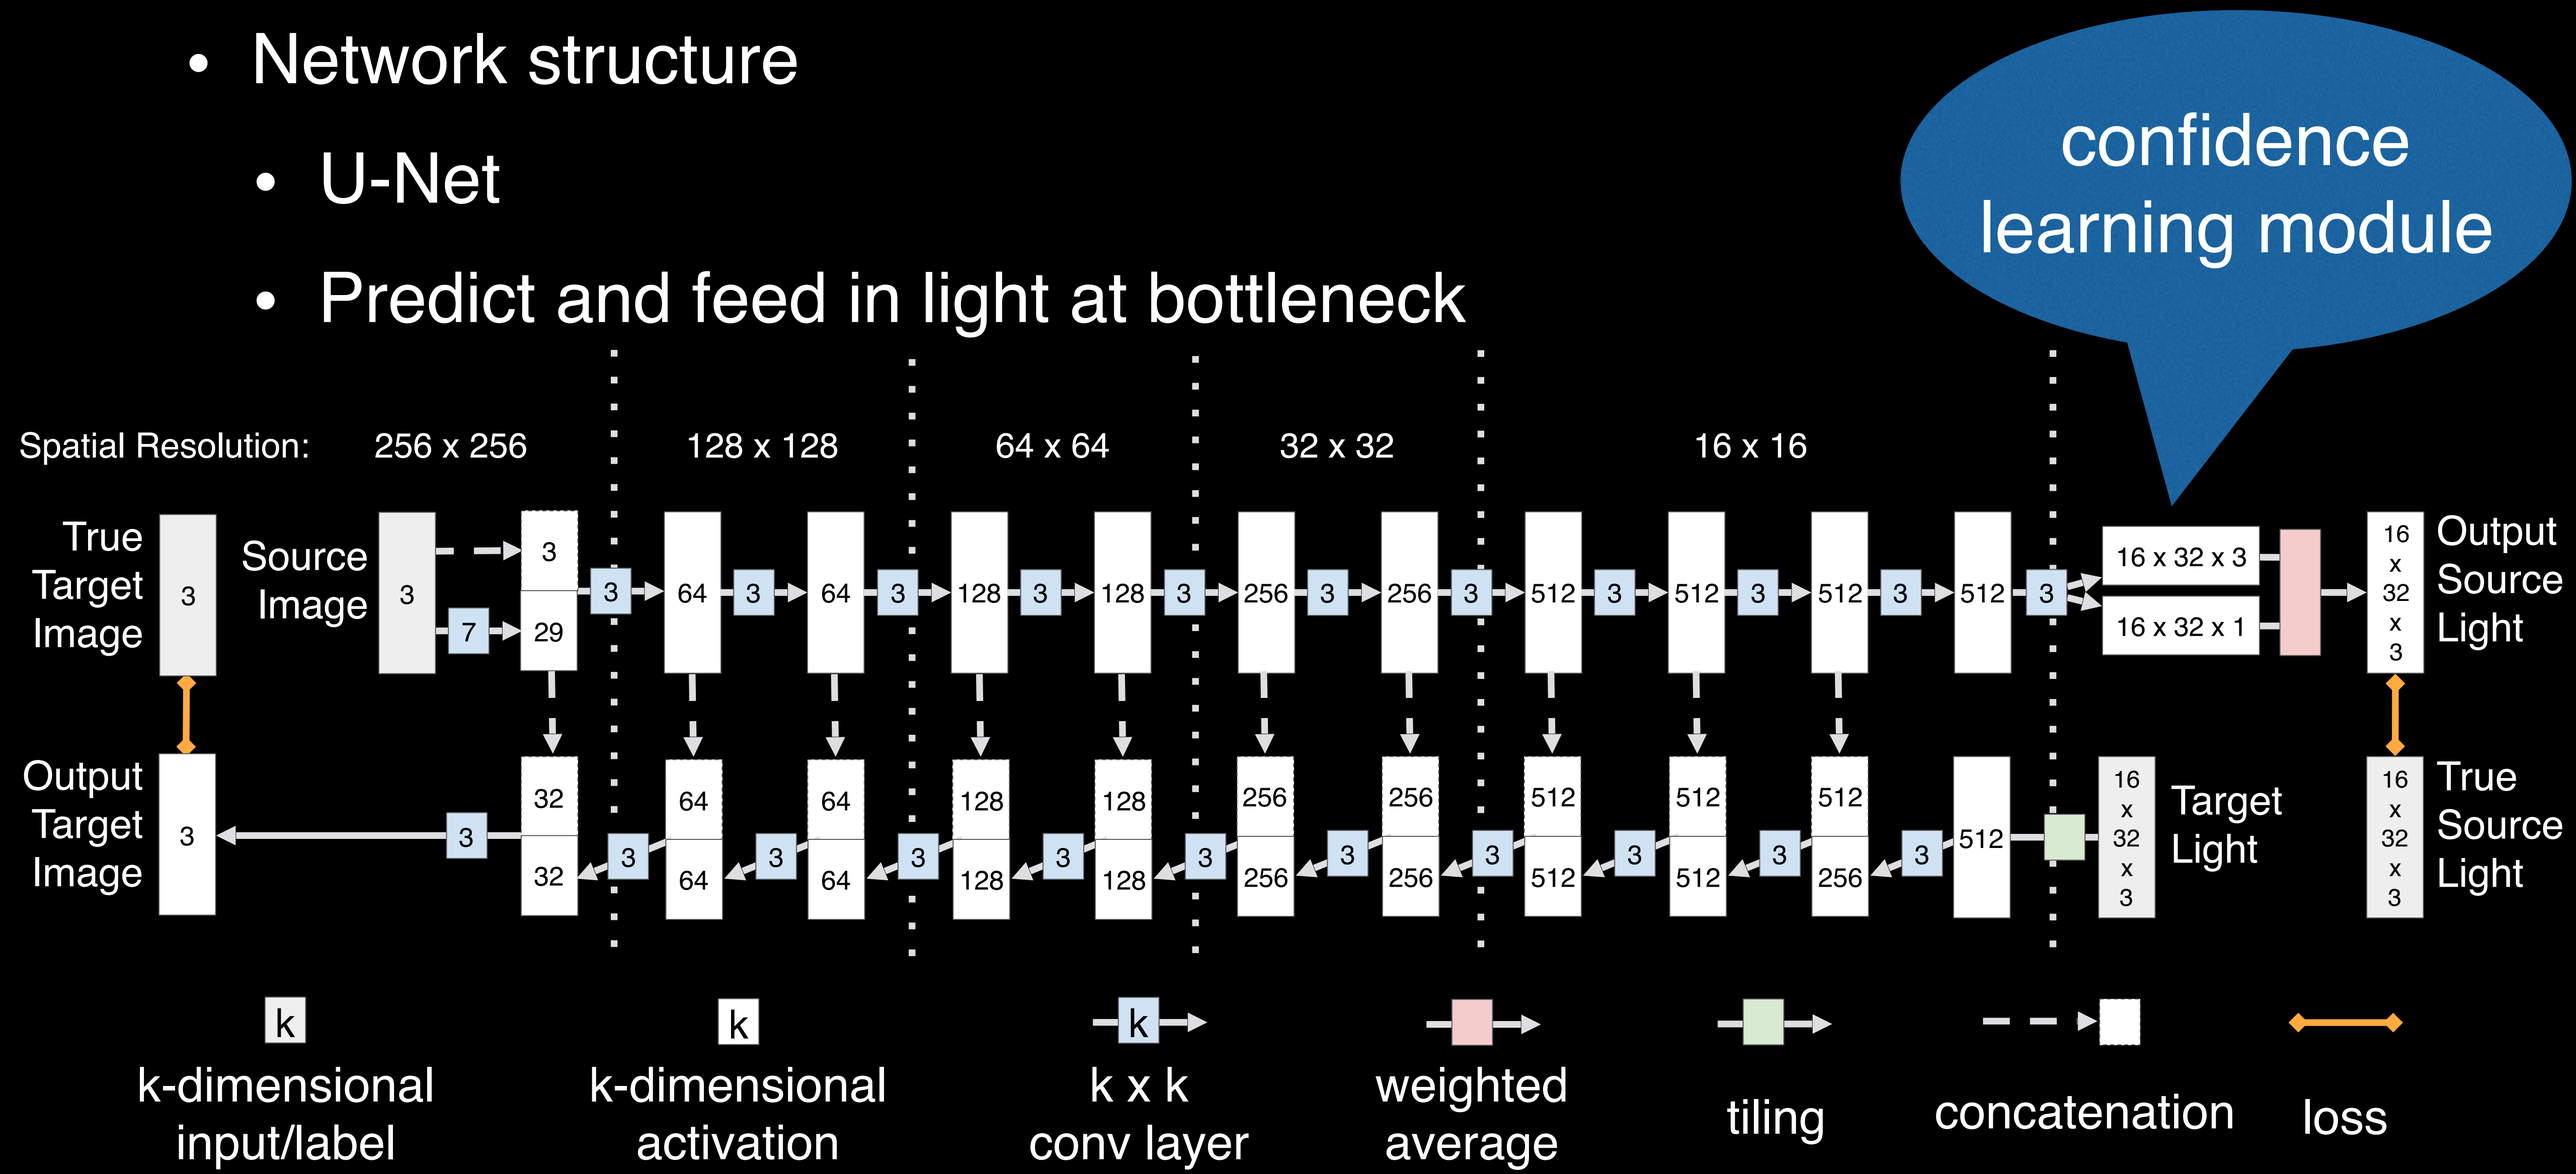

Method: Training

Method: Training

- Confidence learning

Method: Training

- Confidence learning

Method: Training

- Confidence learning

Several
conv layers

Method: Training

- Confidence learning

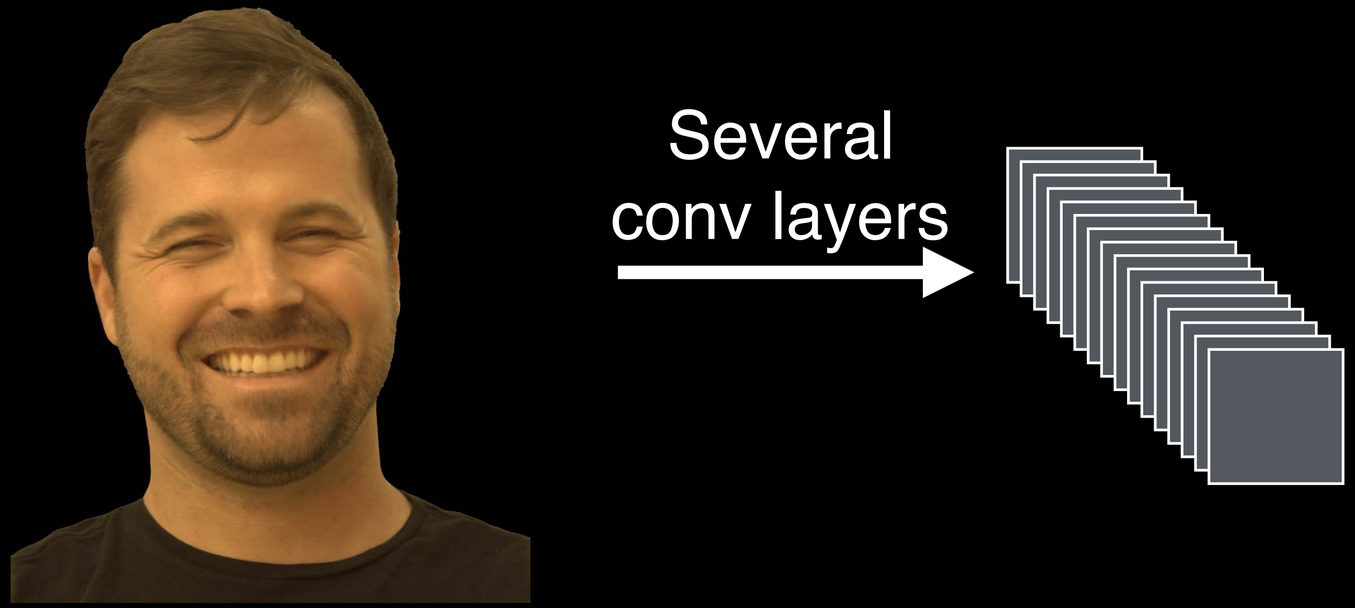

Method: Training

- Confidence learning

Several
conv layers

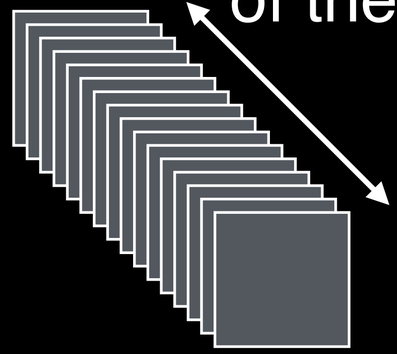

resolution
of the light

Method: Training

- Confidence learning

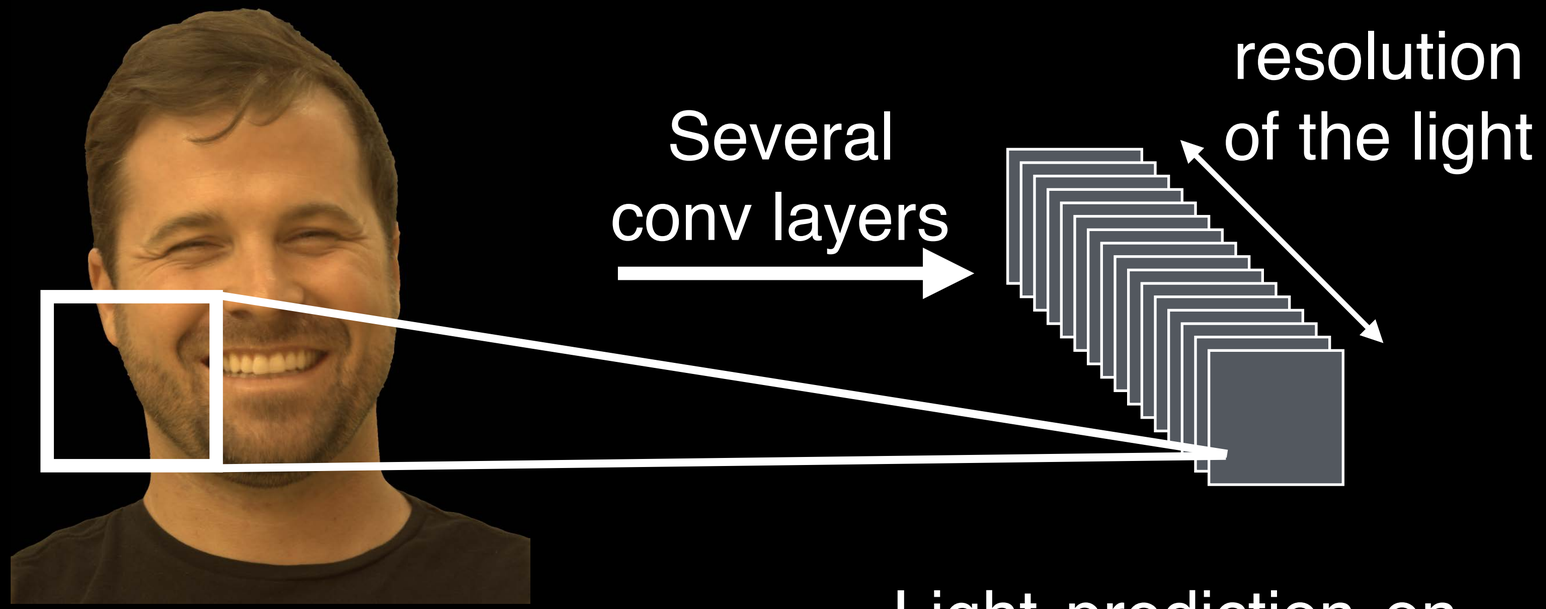

Method: Training

- Confidence learning
  - Predict the confidence of light prediction

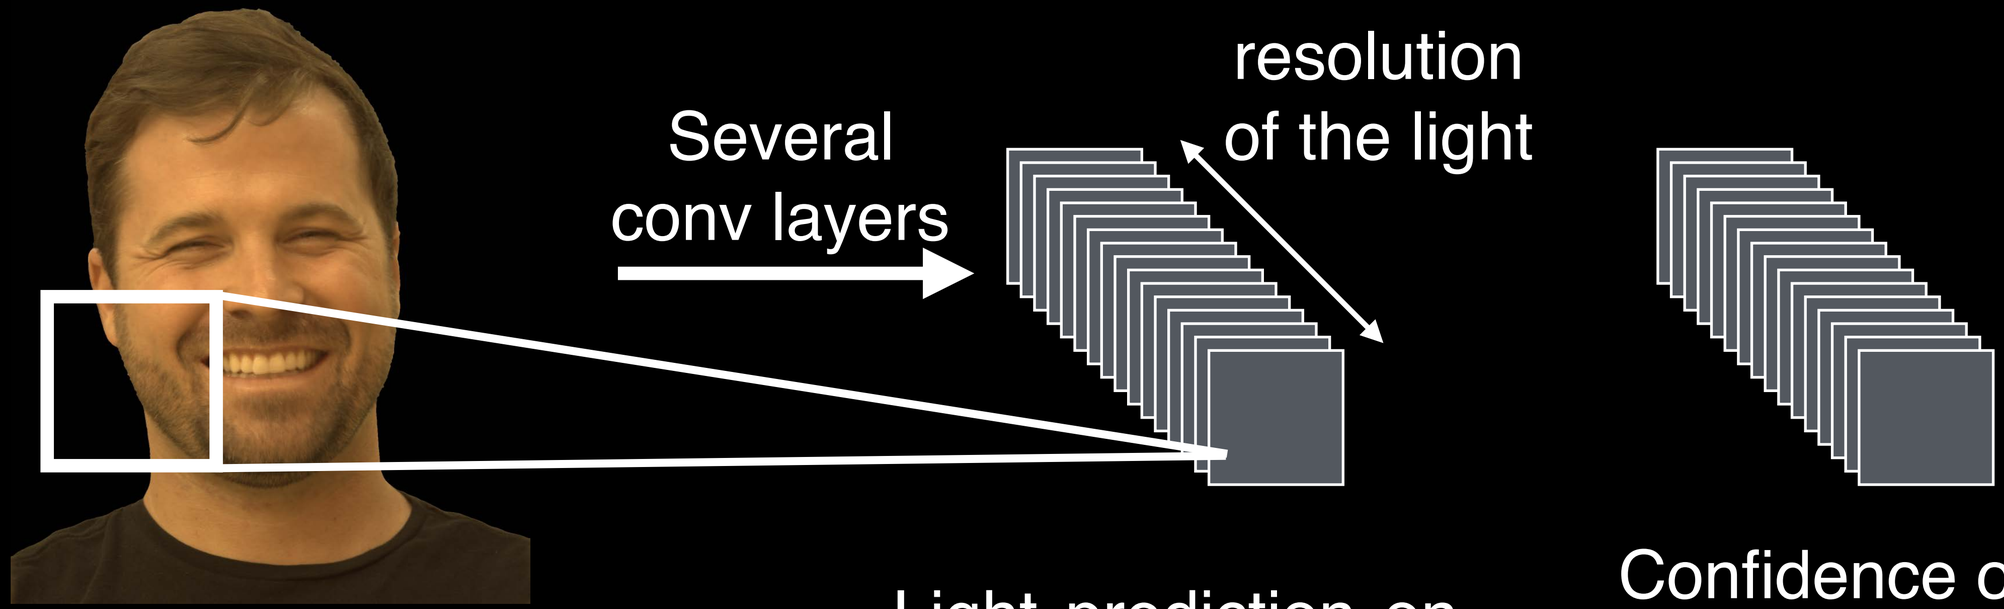

Method: Training

- Confidence learning
  - Predict the confidence of light prediction

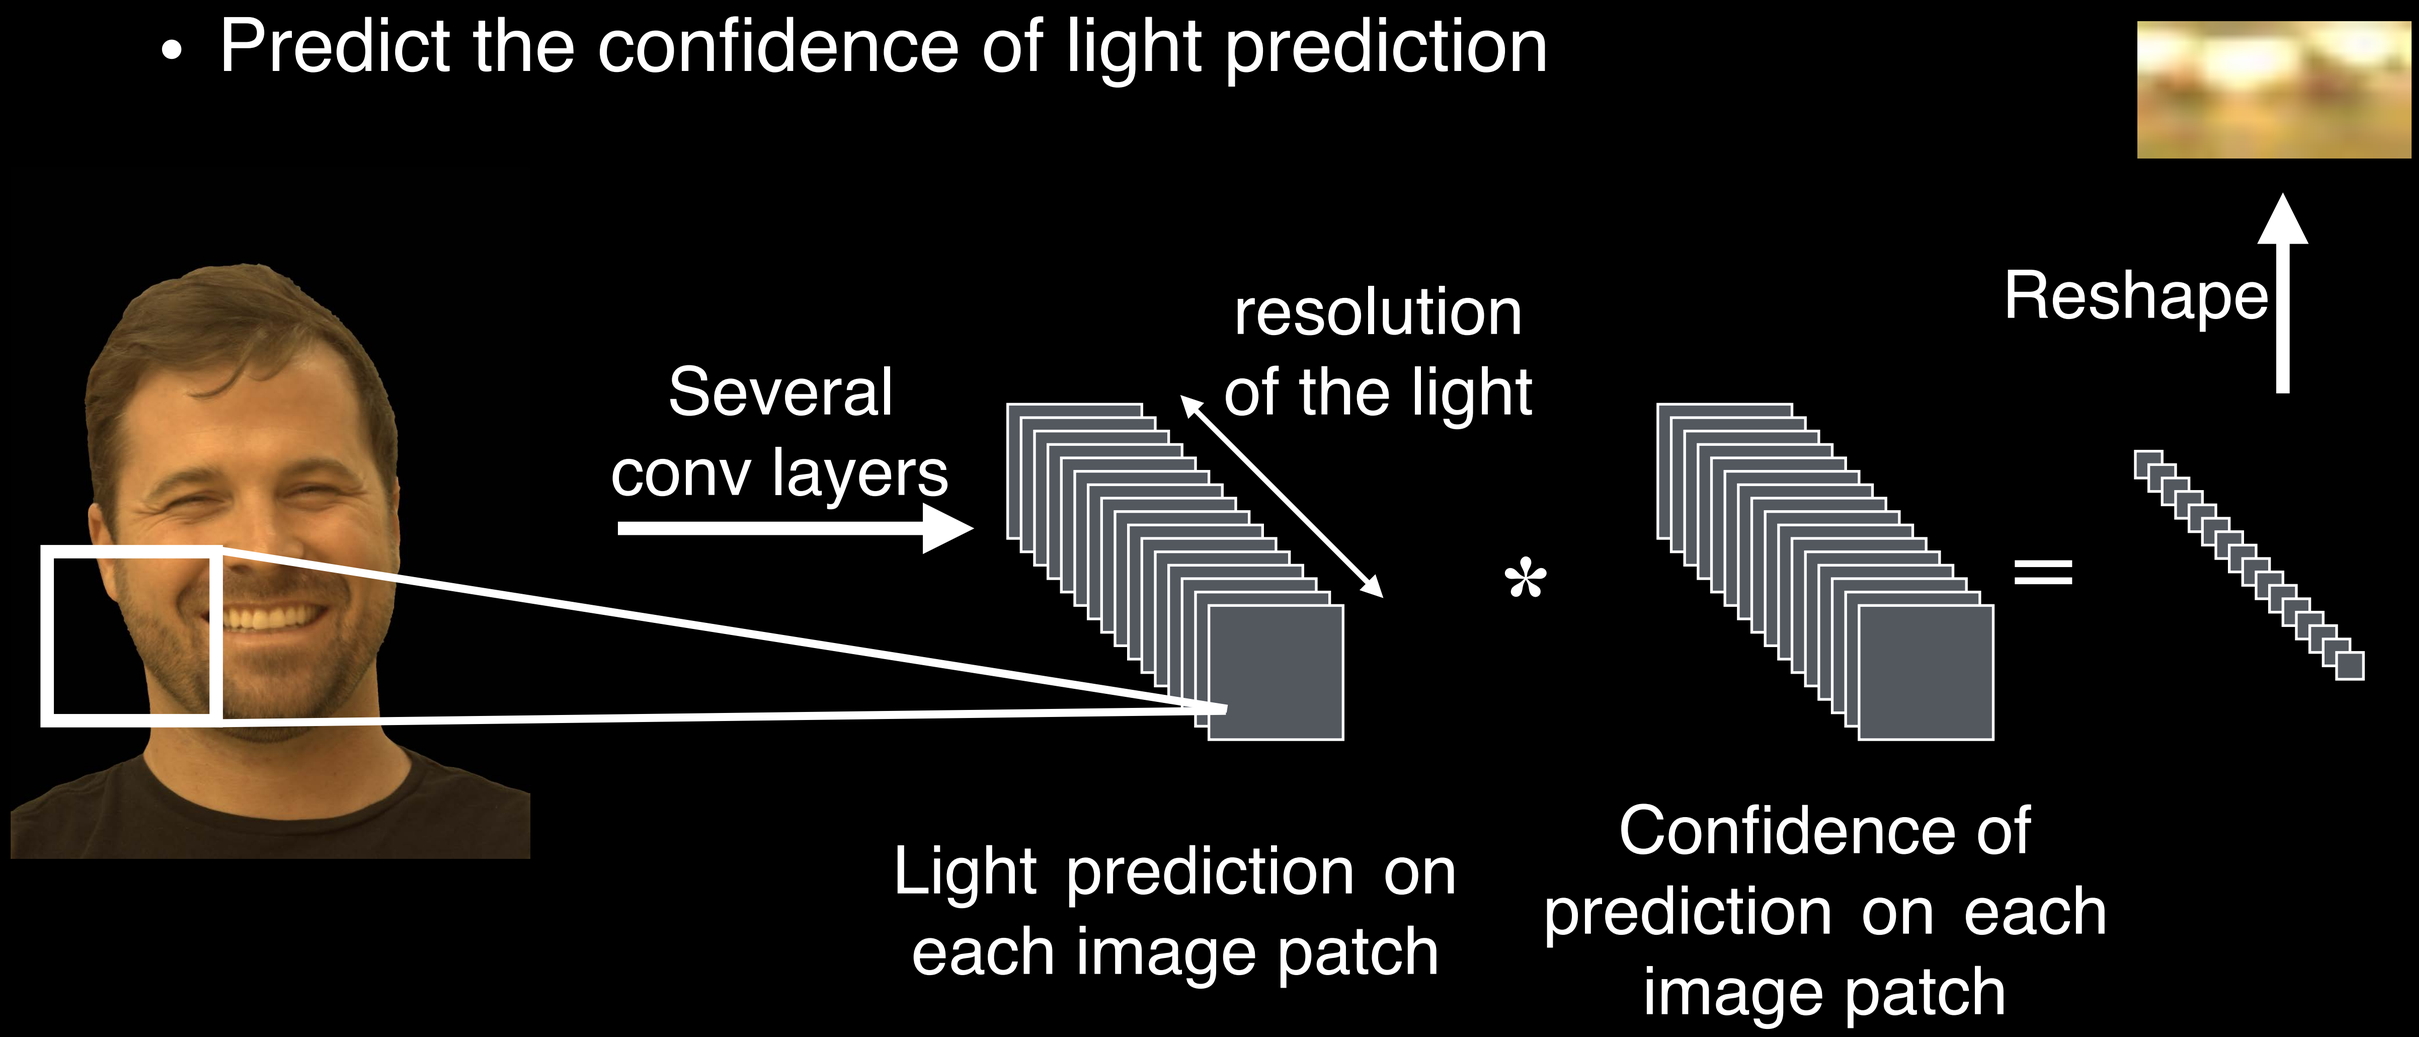

Method: Training

- Confidence learning
  - Predict the confidence of light prediction
  - Allow network to say “I don’t know”

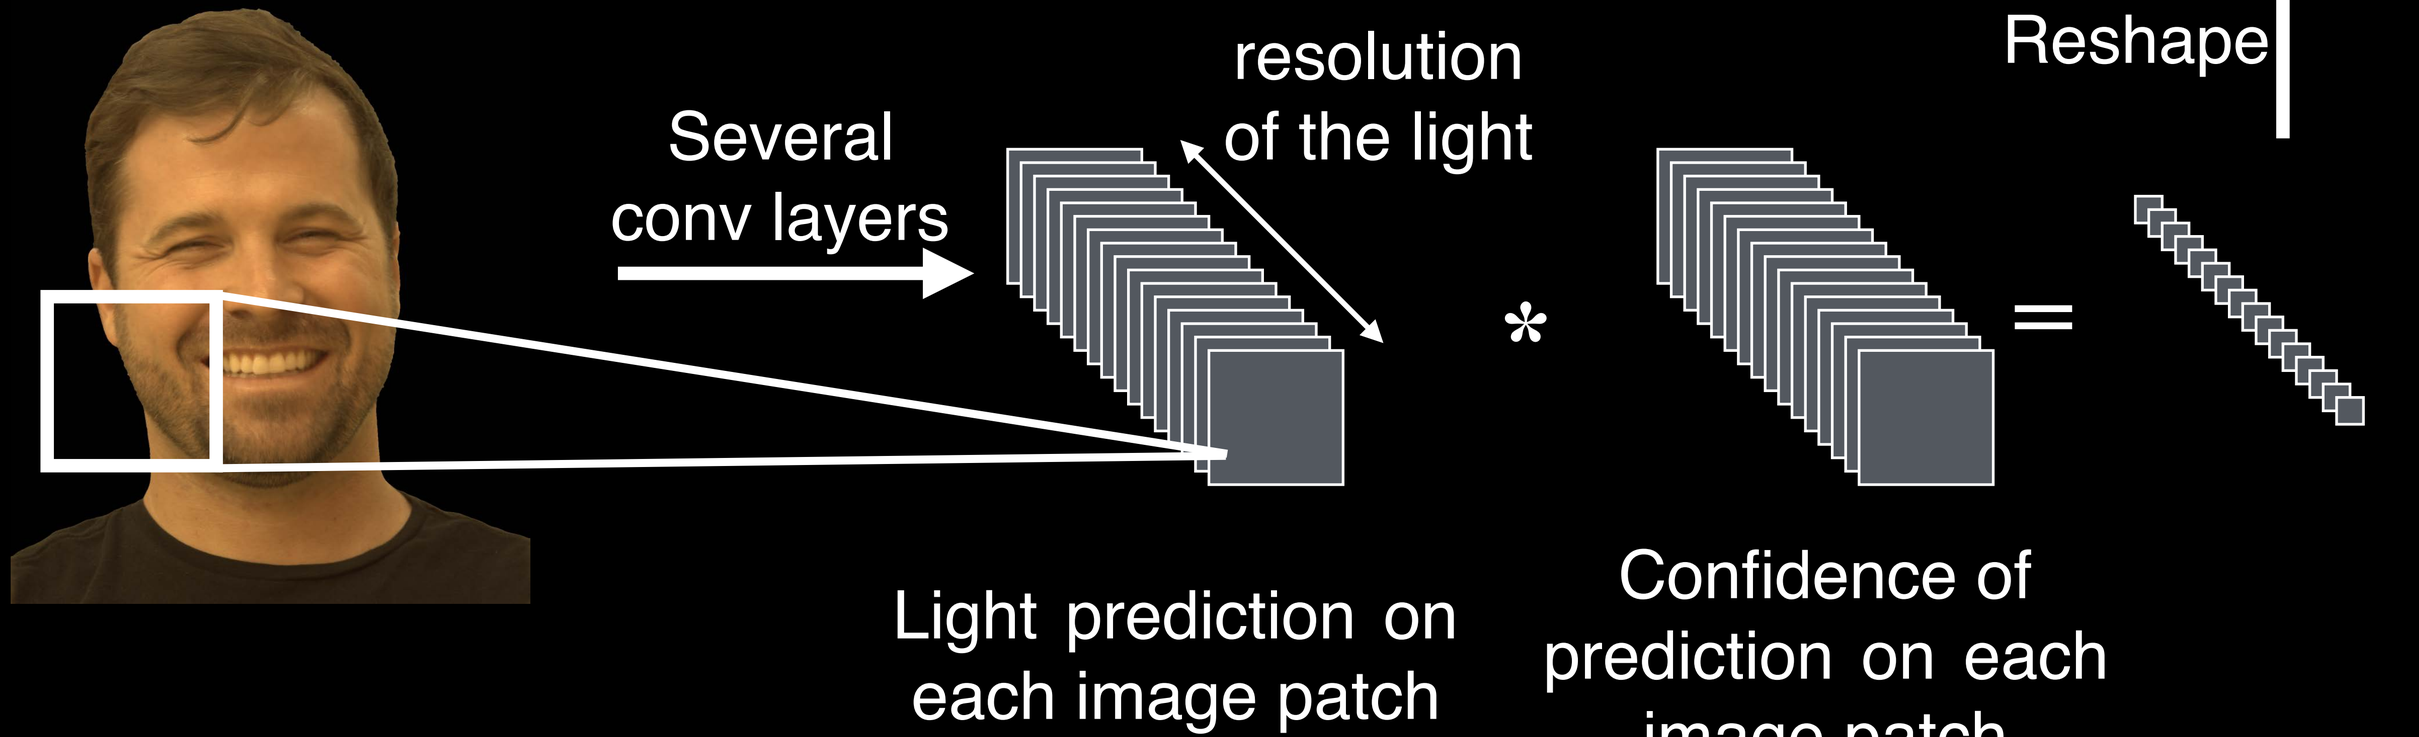

Results: Validation

Results: Validation

- Relit images for complete relighting

Results: Validation

- Relit images for complete relighting

source image

Results: Validation

- Relit images for complete relighting

source image

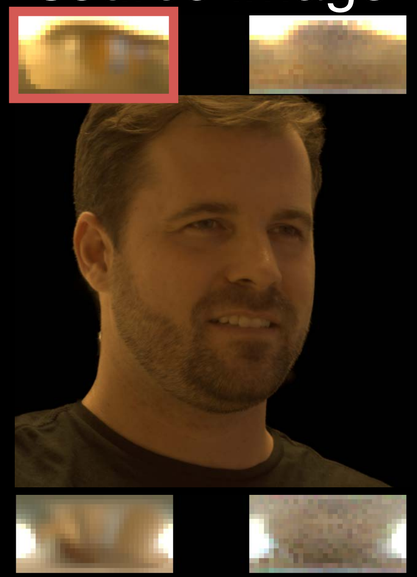

Results: Validation

- Relit images for complete relighting

source image

Results: Validation

- Relit images for complete relighting

source image target image

Results: Validation

- Relit images for complete relighting

Results: Validation

- Relit images for complete relighting

Results: Validation

Complete Relighting

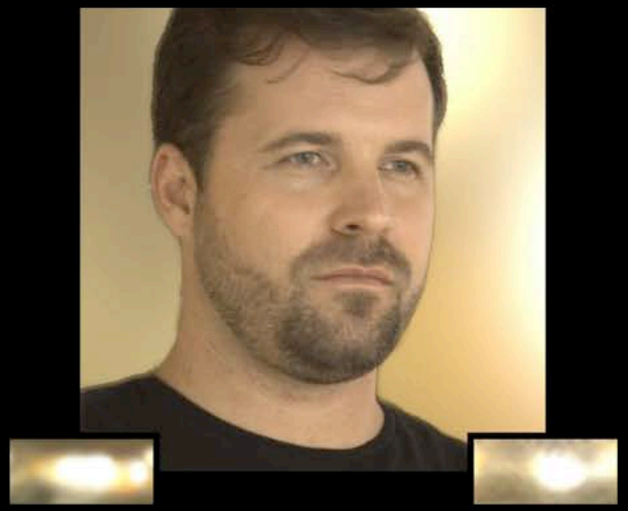

Input
(Lower-left: GT source light)
(Lower-right: predicted source light)

Results: Validation

Complete Relighting

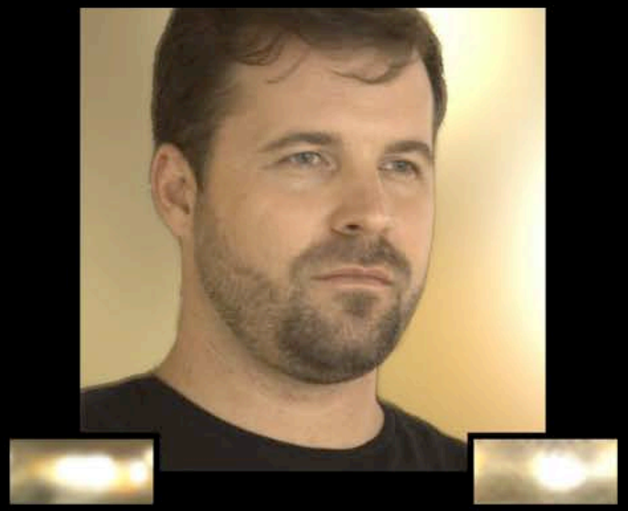

Input
(Lower-left: GT source light)
(Lower-right: predicted source light)

Results: Validation

Results: Validation

- Relight images with predicted light as target light

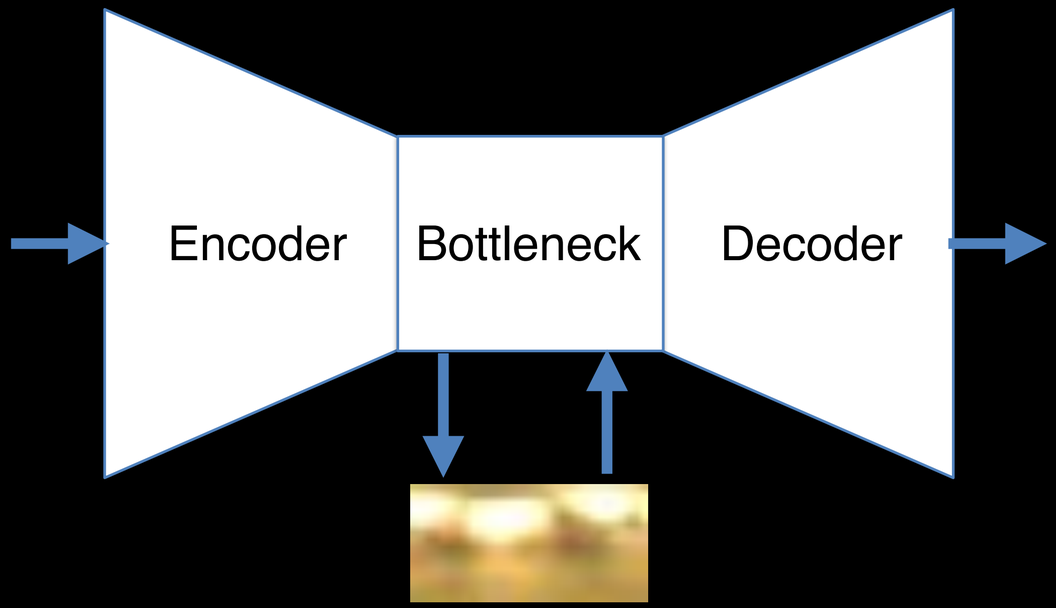

Results: Validation

- Relight images with predicted light as target light

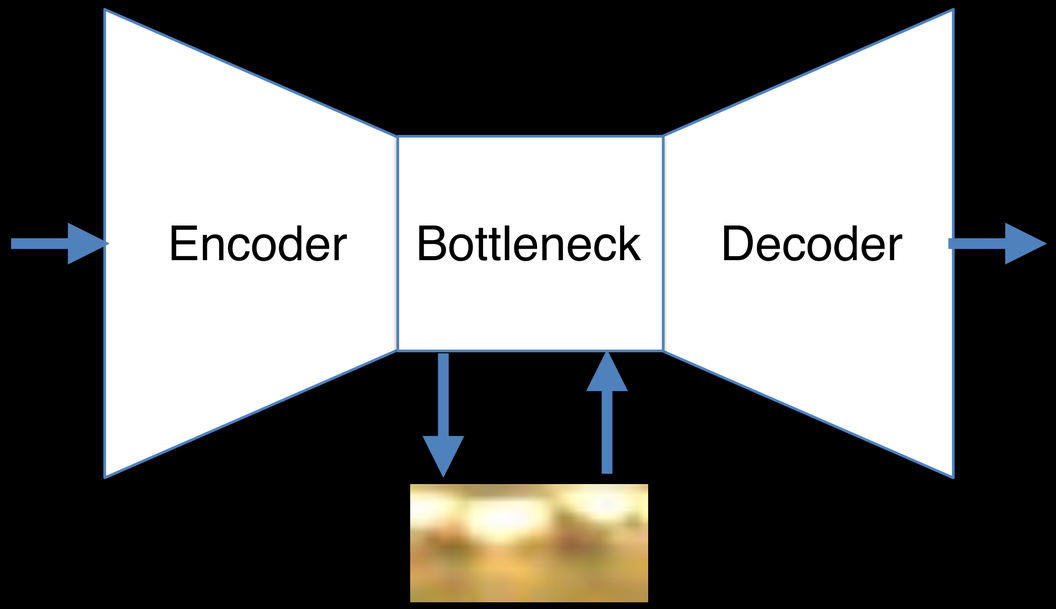

source image

Results: Validation

- Relight images with predicted light as target light

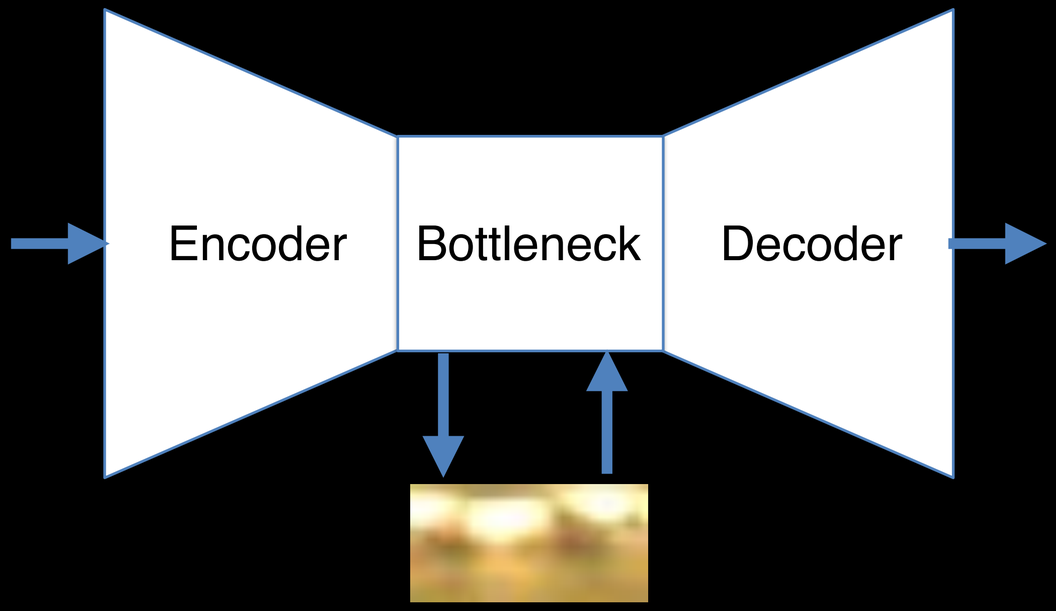

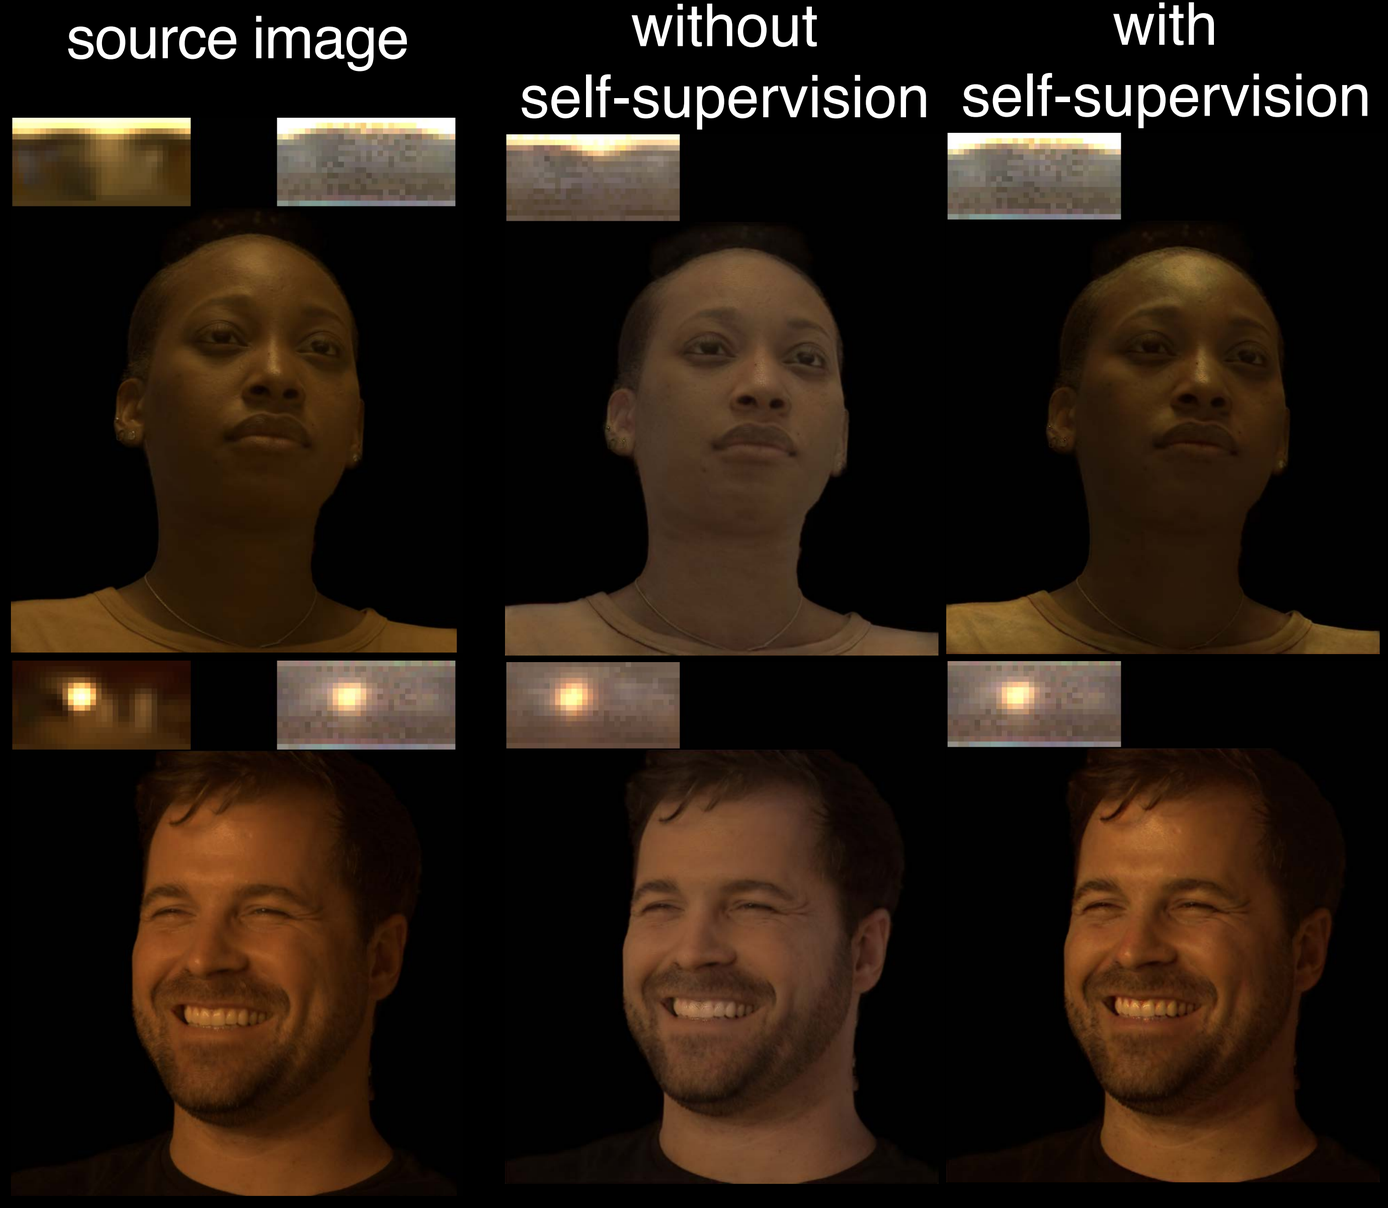

Results: Validation

Illumination Retargeting

Input

Ground Truth

(Relit by Ground Truth Source Light)

Results: Validation

Illumination Retargeting

Input

Ground Truth

(Relit by Ground Truth Source Light)

Results: Validation

Results: Validation

- Comparison with portrait lighting transfer

Results: Validation

- Comparison with portrait lighting transfer

reference image

Results: Validation

- Comparison with portrait lighting transfer

reference image

source image

Results: Validation

- Comparison with portrait lighting transfer

reference image

Extract light
from reference

source image

Results: Validation

- Comparison with portrait lighting transfer

reference image

Extract light from reference

Apply to source image

source image

Results: Validation

- Comparison with portrait lighting transfer

reference image | groundtruth | **ours**

Extract light from reference

Apply to source image

source image

Results: Validation

- Comparison with portrait lighting transfer

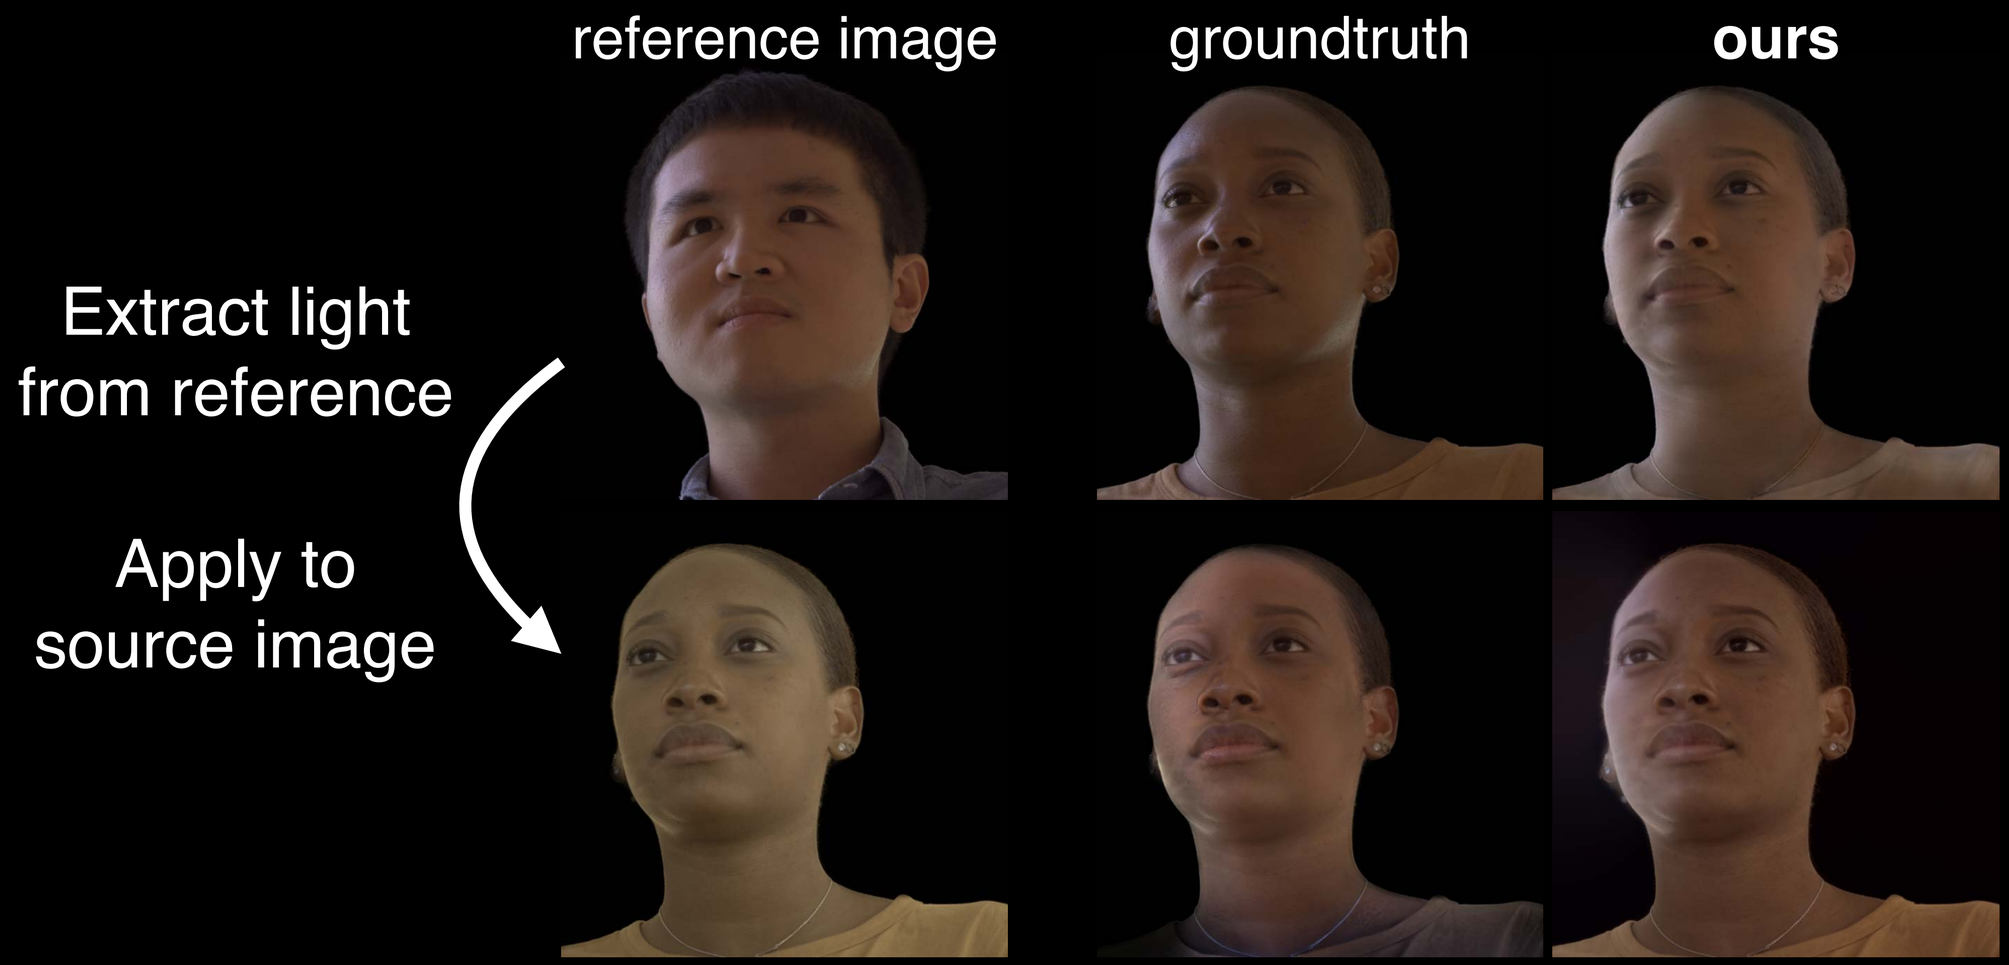

Results: Validation

Results: Validation

- Evaluation on lighting prediction

Results: Validation

- Evaluation on lighting prediction

source image

ground truth

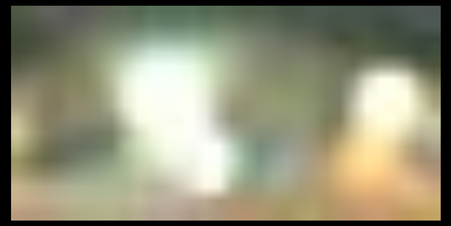

Results: Validation

- Evaluation on lighting prediction

source image

ground truth

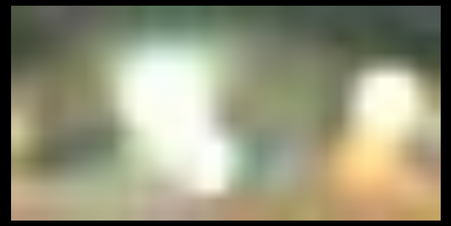

ours

ours w/o
confidence learning

Results: Validation

- Evaluation on lighting prediction

source image

ground truth

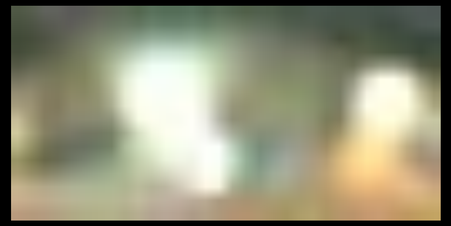

ours

[Barron & Malik 2015]

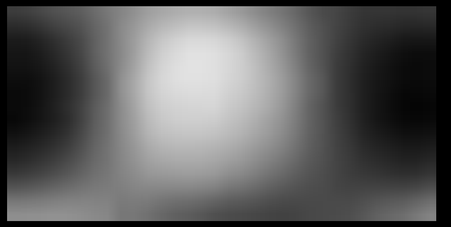

ours w/o confidence learning

[Sengupta et al. 2018]

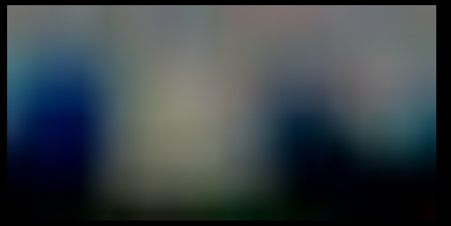

Results: Images in the wild

Results: Images in the wild

Input Image

Relit Image

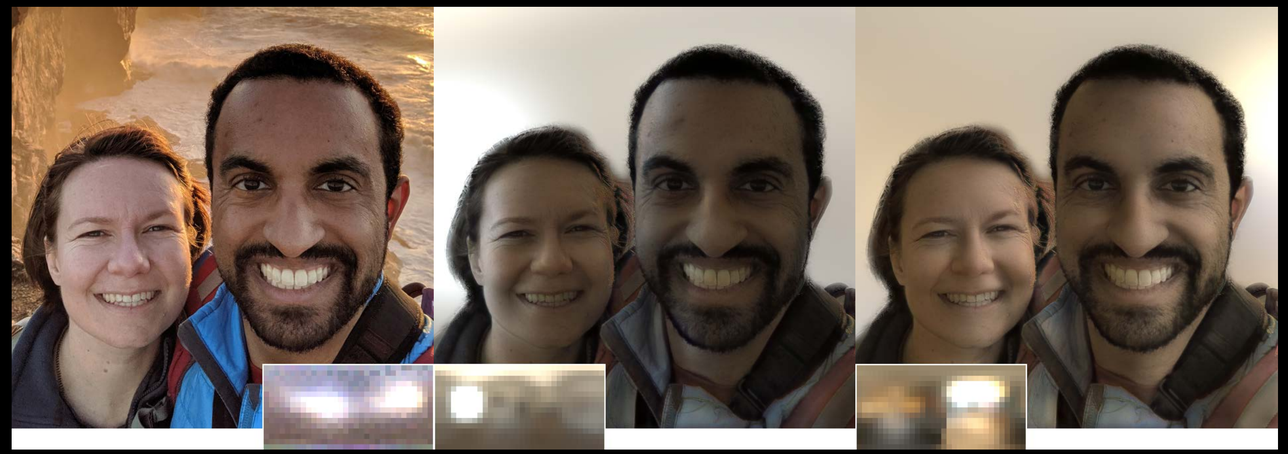

Results: Images in the wild

Input Image

Relit Image

Results: Images in the wild

Input Image Relit Image Input Image Relit Image

Results: Images in the wild

Input Image Relit Image Input Image Relit Image

Results: Images in the wild

Complete Relighting

Input Image

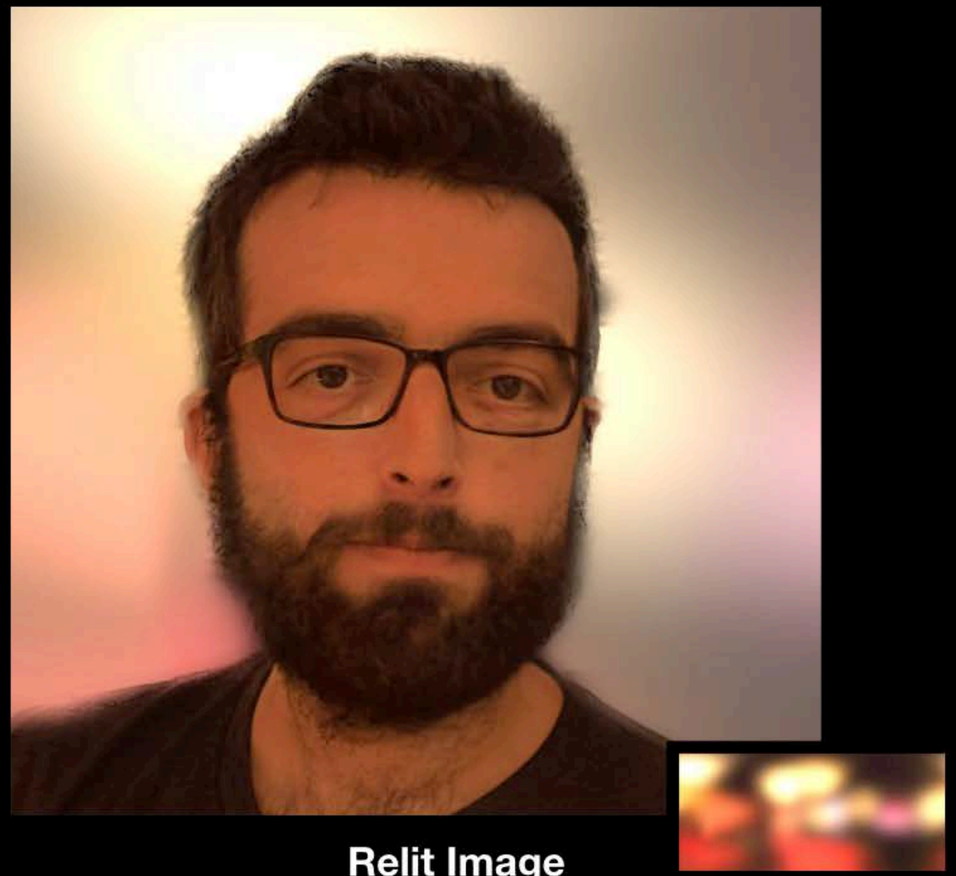

Relit Image
(Lower-right: input target light)

Results: Images in the wild

Complete Relighting

Input Image

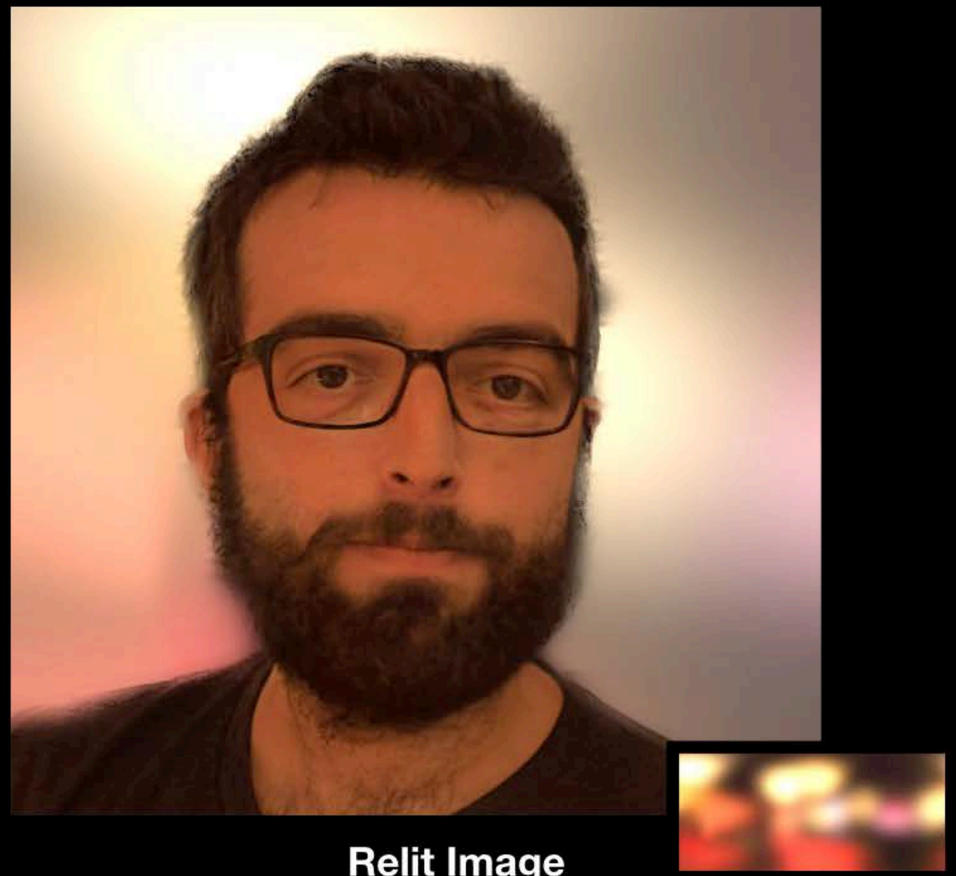

Relit Image
(Lower-right: input target light)

Results: Images in the wild

Illumination Retargeting

Input Image

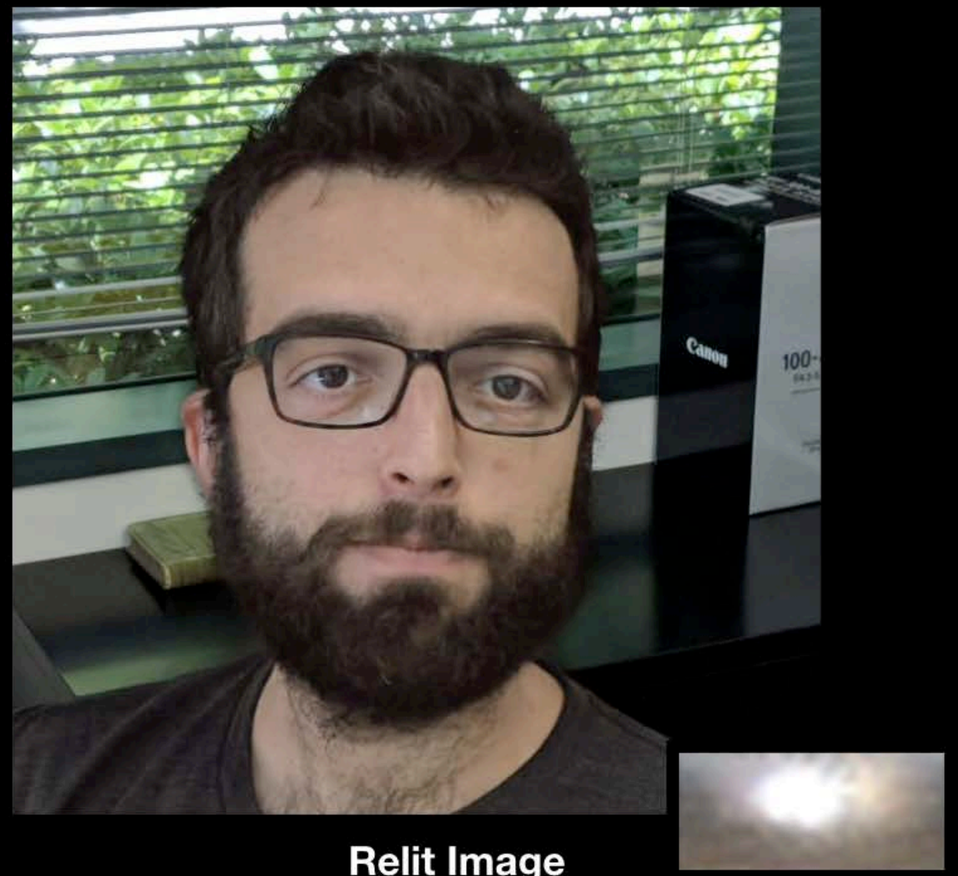

Relit Image
(Lower-right: predicted source light)

Results: Images in the wild

Illumination Retargeting

Input Image

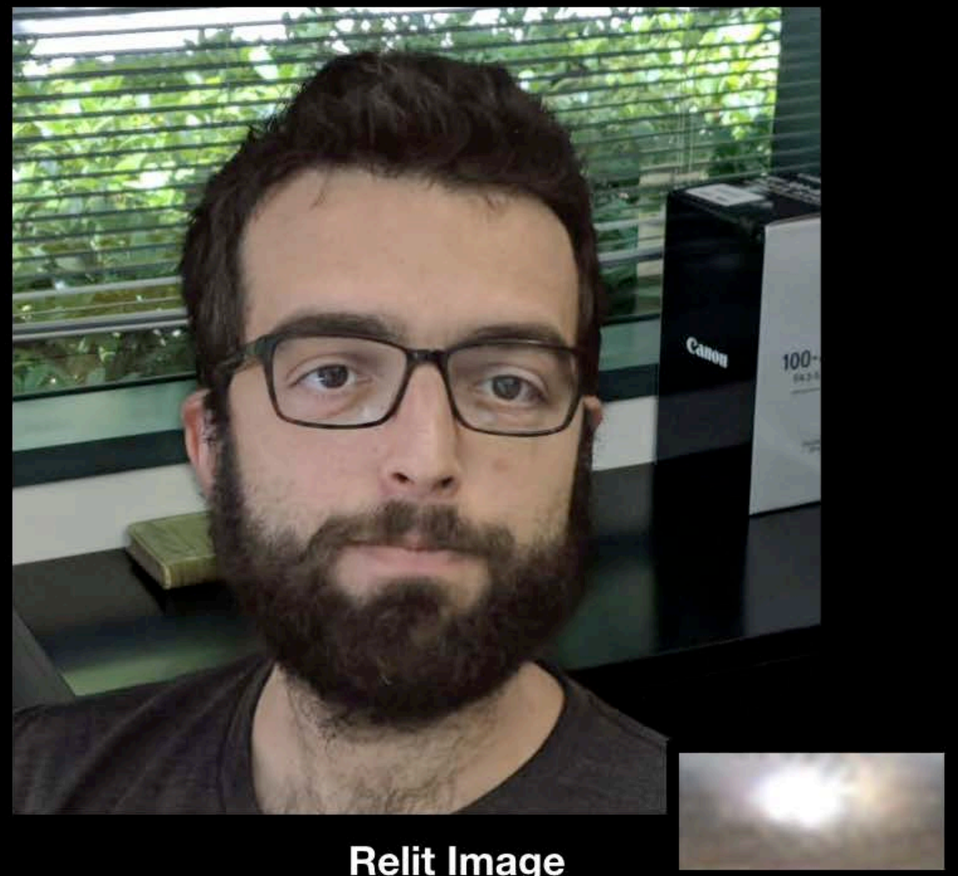

Relit Image
(Lower-right: predicted source light)

Limitations

Limitations

- Complex shadows
- Specular highlights
- Overexposed pixels

Input Image

Relit Image

Limitations

- Complex shadows
- Specular highlights
- Overexposed pixels
- Over-smoothing

Input Image

Relit Image

Limitations

- Complex shadows
- Specular highlights
- Overexposed pixels
- Over-smoothing
- Unseen high-saturation color

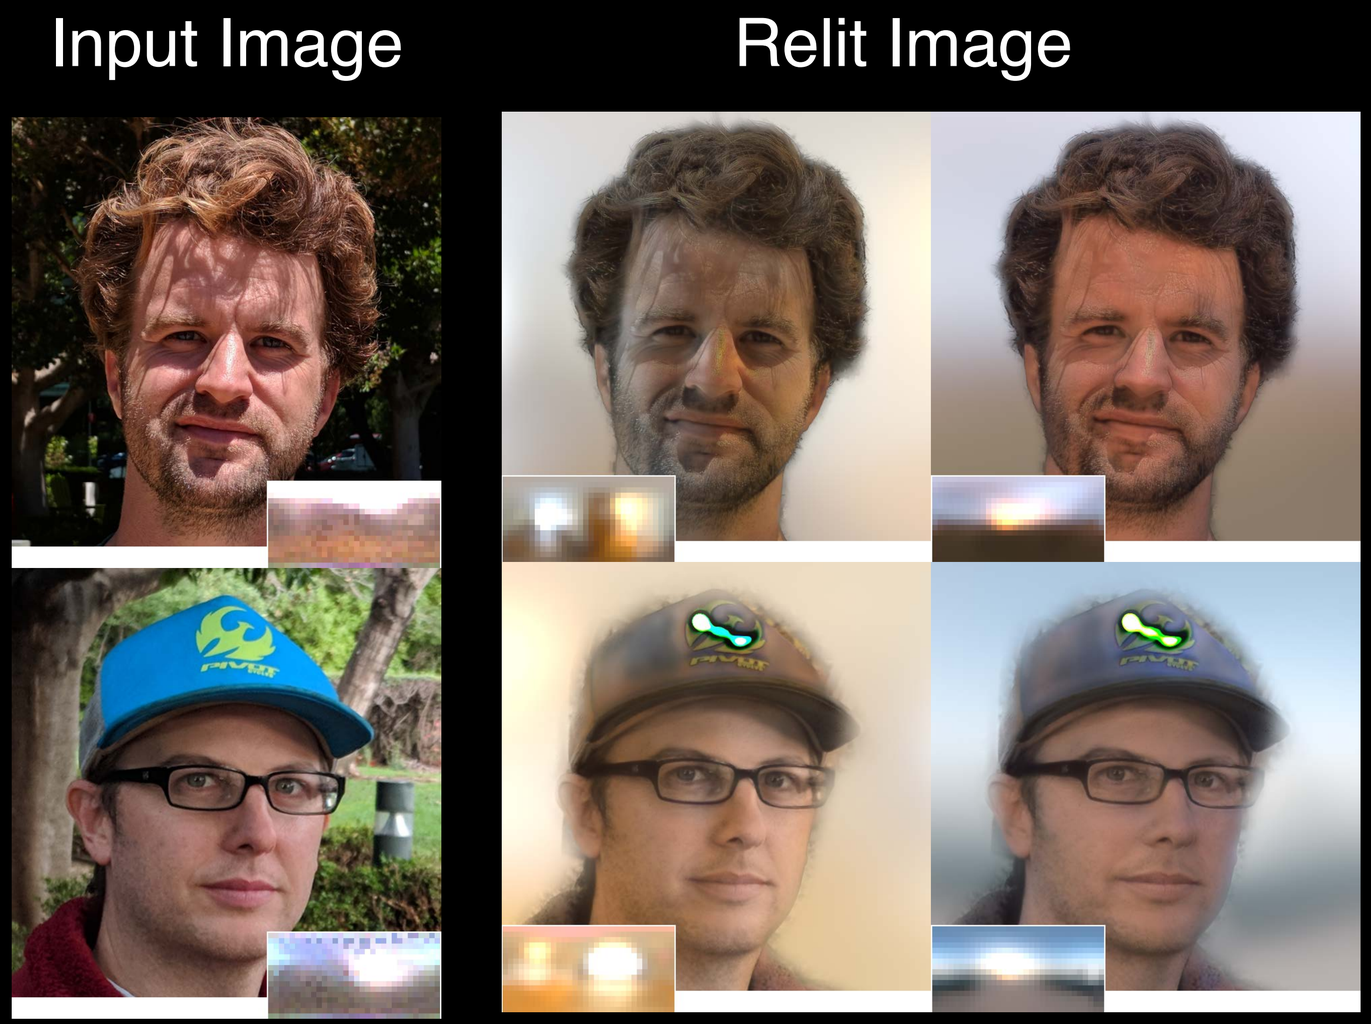

Conclusion

Conclusion

- Learn the relighting function on portraits using Light Stage data

Conclusion

- Learn the relighting function on portraits using Light Stage data
- Take home message:

Conclusion

- Learn the relighting function on portraits using Light Stage data
- Take home message:
  - For human faces, use real data.

Conclusion

- Learn the relighting function on portraits using Light Stage data
- Take home message:
  - For human faces, use real data.
  - End-to-end training vs assuming models.

Acknowledgement

- This work was funded in part by a Jacobs Fellowship, the Ronald L. Graham Chair, and the UC San Diego Center for Visual Computing.
- Thanks to Zhixin Shu and Yichang Shih for the help on baseline algorithms
- Thanks to Jean-François Lalonde for providing the indoor lighting dataset.
- Thanks to Peter Denny for coordinating dataset capture.
- **Thanks to all the anonymous volunteers in the dataset from Google, UCSD and UCLA**

Single Image
Portrait Relighting

Also presenting in poster session today at 12:15-1:15.

Single Image Portrait Relighting

input image

generated portraits under new illuminations

Single Image Portrait Relighting

input image

generated portraits under new illuminations