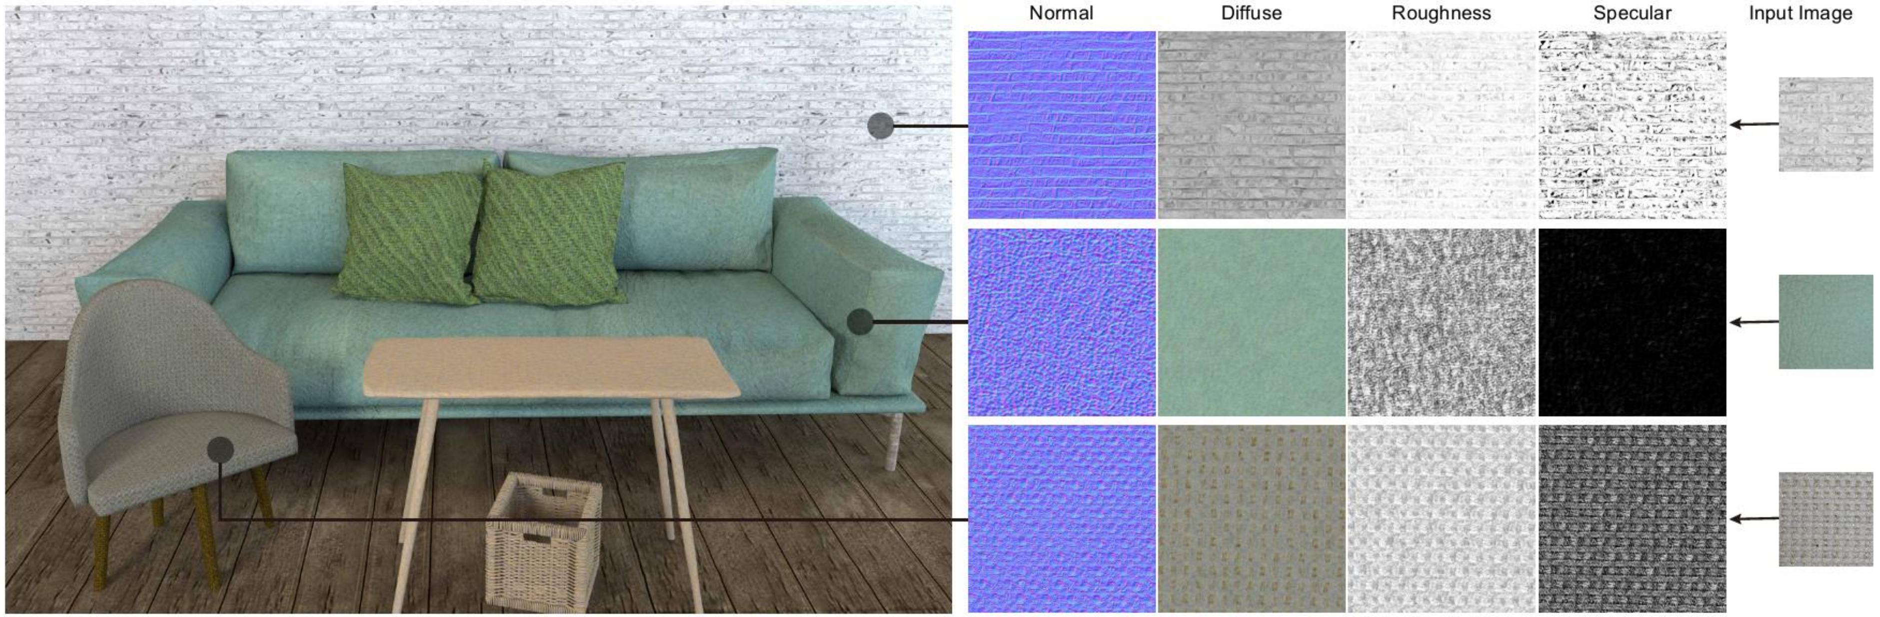

Joint SVBRDF Recovery and Synthesis From a Single Image using an Unsupervised Generative Adversarial Network

Yezi Zhao[1], Beibei Wang[2], Yanning Xu[1], Zheng Zeng[1], Lu Wang[1] and Nicolas Holzschuch[3]

[1] School of Software, Shandong University
[2] School of Computer Science and Engineering, Nanjing University of Science and Technology
[3] Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK

Motivation

normal

diffuse

roughness

specular

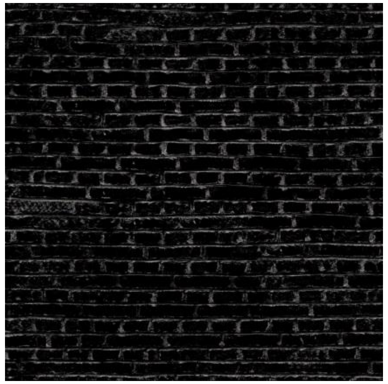

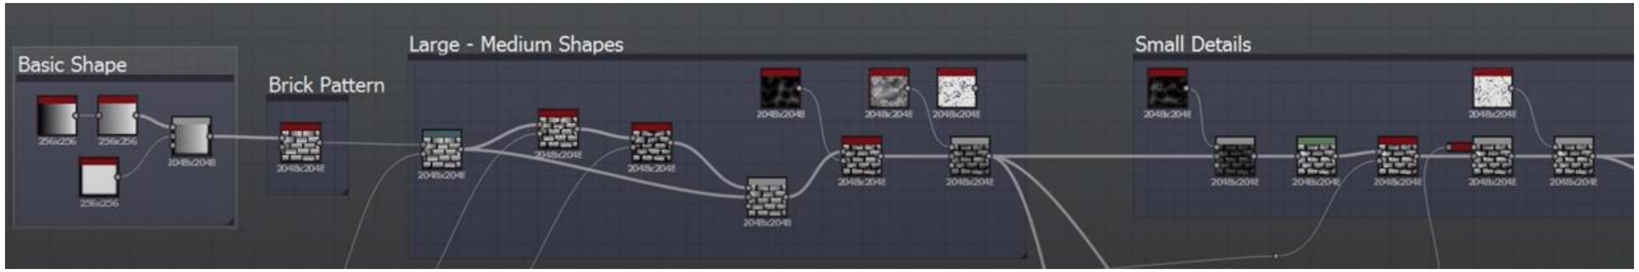

Substance by Adobe

Our Goal

A lightweight method for recovering real-world material

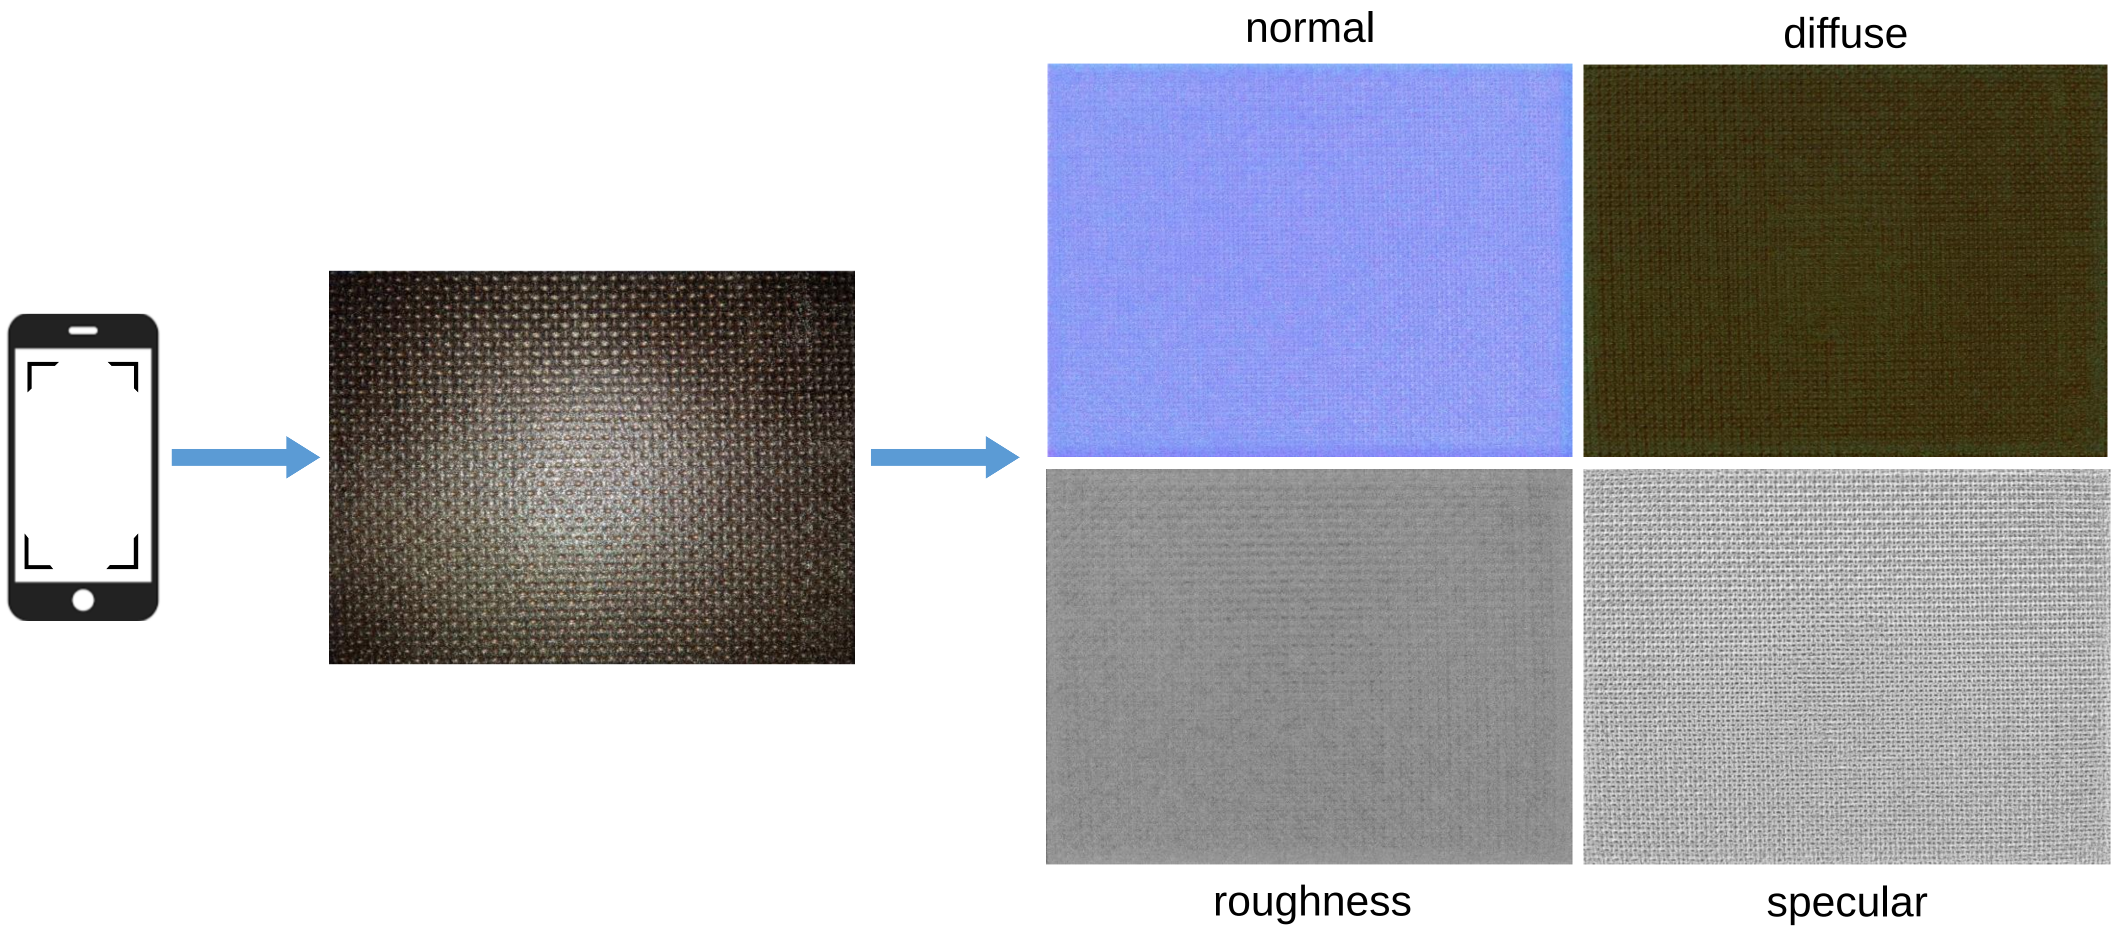

State of the art

Aittala et al., SIGGRAPH 2016
Reflectance Modeling by Neural Texture Synthesis

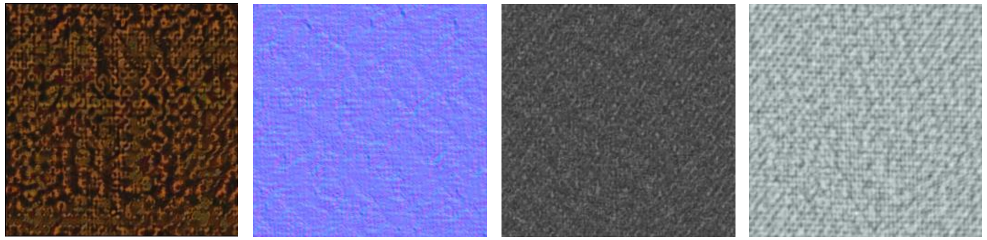

Li et al., SIGGRAPH 2017
Modeling surface appearance from a single photograph using self-augmented convolutional neural networks

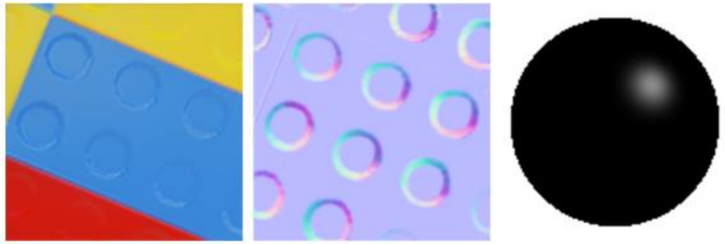

Deschaintre et al., SIGGRAPH 2018
Single-Image SVBRDF Capture with a Rendering-Aware Deep Network

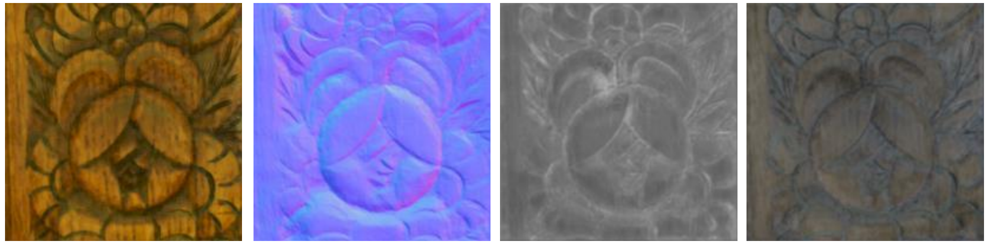

State of the art

Aittala et al., SIGGRAPH 2016
Reflectance Modeling by Neural Texture Synthesis

diffuse normal glossiness specular

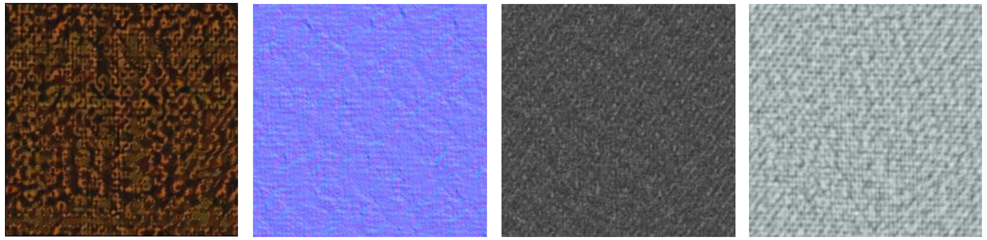

Li et al., SIGGRAPH 2017
Modeling surface appearance from a single photograph using self-augmented convolutional neural networks

diffuse normal specular

Deschaintre et al., SIGGRAPH 2018
Single-Image SVBRDF Capture with a Rendering-Aware Deep Network

diffuse normal roughness specular

State of the art

Aittala et al., SIGGRAPH 2016
Reflectance Modeling by Neural Texture Synthesis

diffuse normal glossiness specular

Li et al., SIGGRAPH 2017
Modeling surface appearance from a single photograph using self-augmented convolutional neural networks

diffuse normal specular

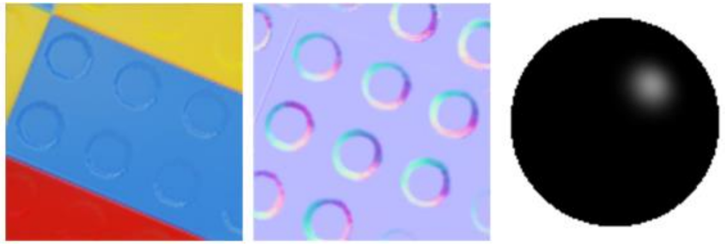

Deschaintre et al., SIGGRAPH 2018
Single-Image SVBRDF Capture with a Rendering-Aware Deep Network

diffuse normal roughness specular

State of the art

Aittala et al., SIGGRAPH 2016
Reflectance Modeling by Neural Texture Synthesis

Li et al., SIGGRAPH 2017
Modeling surface appearance from a single photograph using self-augmented convolutional neural networks

Deschaintre et al., SIGGRAPH 2018
Single-Image SVBRDF Capture with a Rendering-Aware Deep Network

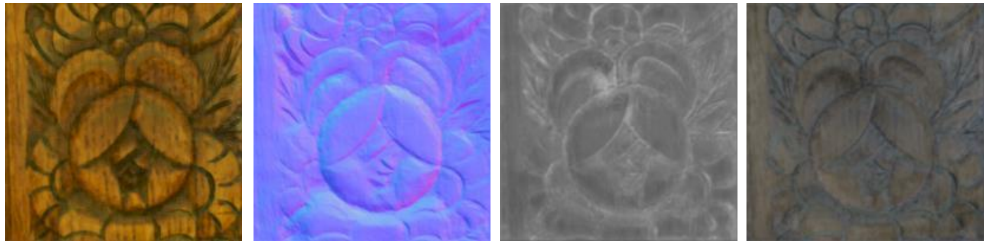

State of the art

Aittala et al., SIGGRAPH 2016
Reflectance Modeling by Neural Texture Synthesis

Low-resolution ☹

Li et al., SIGGRAPH 2017
Modeling surface appearance from a single photograph using self-augmented convolutional neural networks

Deschaintre et al., SIGGRAPH 2018
Single-Image SVBRDF Capture with a Rendering-Aware Deep Network

State of the art

Gao et al., SIGGRAPH 2019
Deep Inverse Rendering for High-resolution SVBRDF Estimation from an Arbitrary Number of Images

input

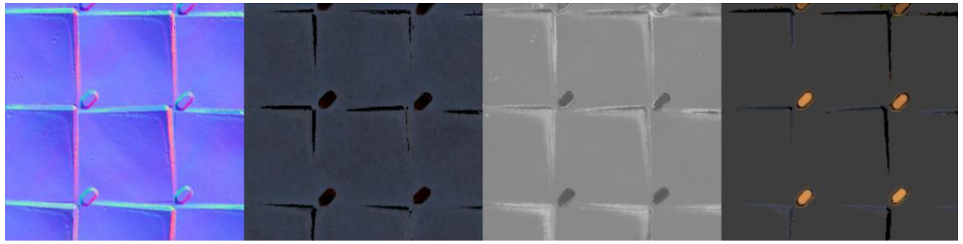

normal diffuse roughness specular

State of the art

Gao et al., SIGGRAPH 2019
Deep Inverse Rendering for High-resolution SVBRDF Estimation from an Arbitrary Number of Images

input

Rely on a plausible starting point of optimizing ☹

normal diffuse roughness specular

State of the art

Zhou et al., SIGGRAPH 2018
Non-stationary texture synthesis by adversarial expansion

Synthesize

[Zhou et al. 2018]

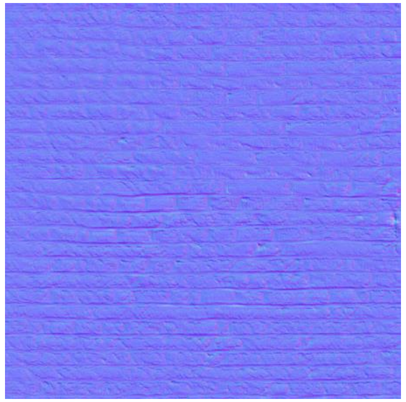

State of the art

Zhou et al., SIGGRAPH 2018
Non-stationary texture synthesis by adversarial expansion

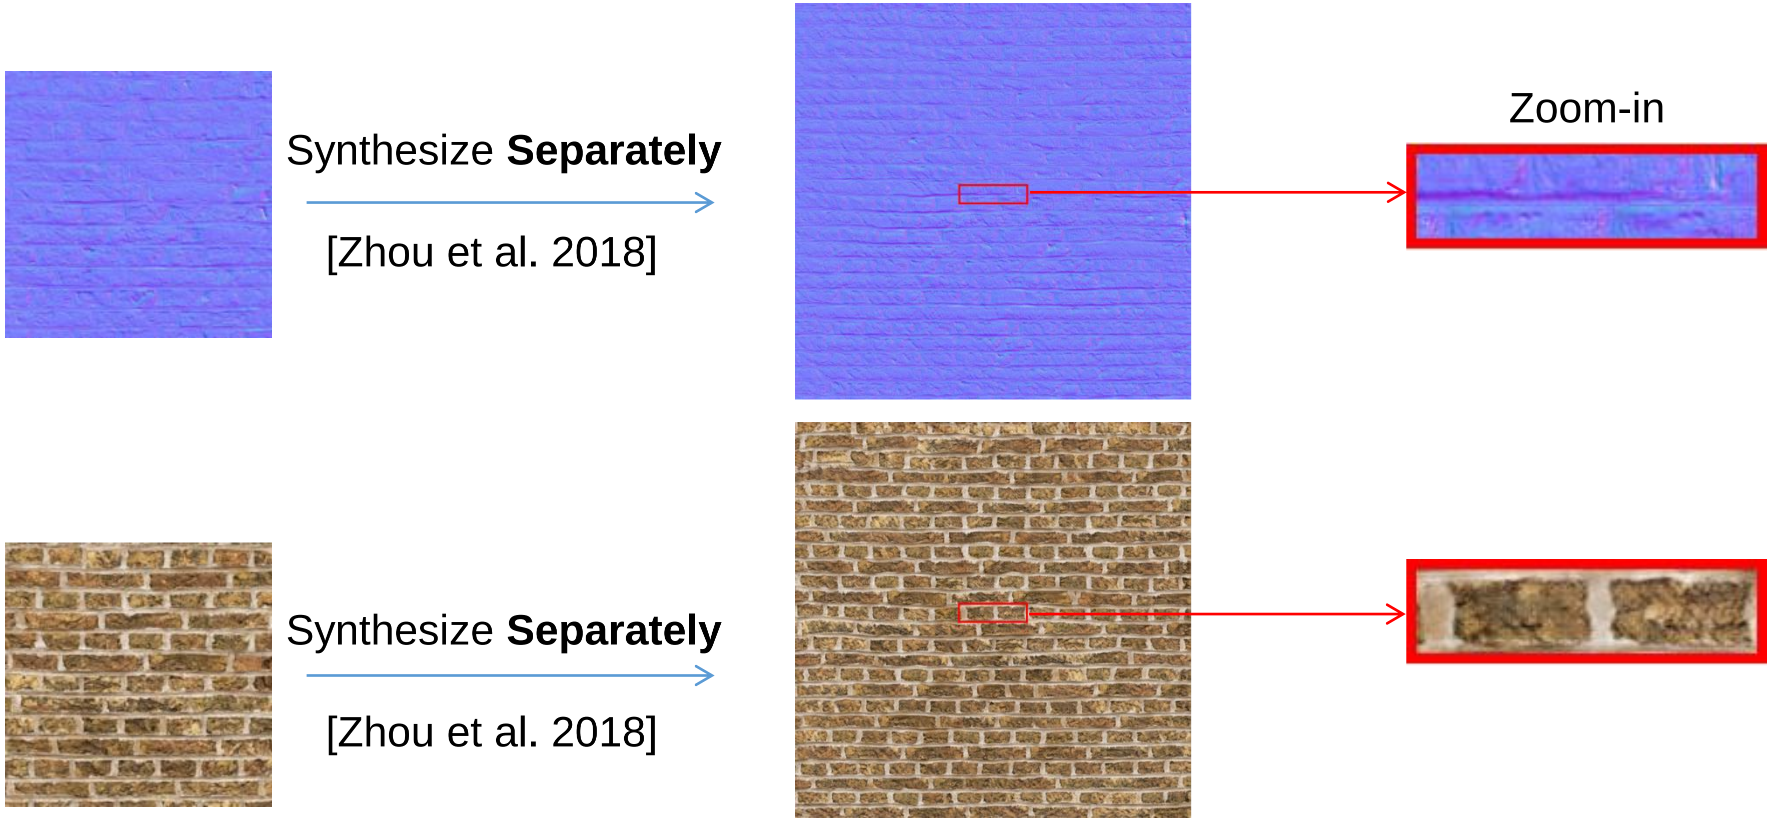

State of the art

Zhou et al., SIGGRAPH 2018
Non-stationary texture synthesis by adversarial expansion

Synthesisze **Seperatly**

[Zhou et al. 2018]

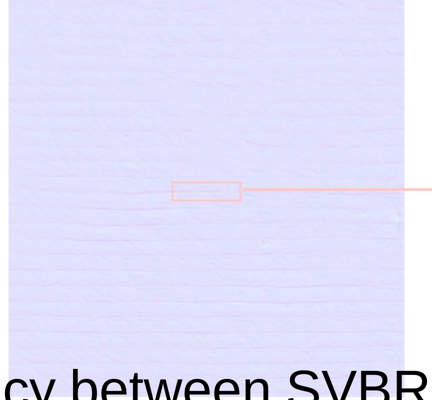

Zoom-in

Inconsistency between SVBRDF maps ☹

Synthesisze **Seperatly**

[Zhou et al. 2018]

Method overview

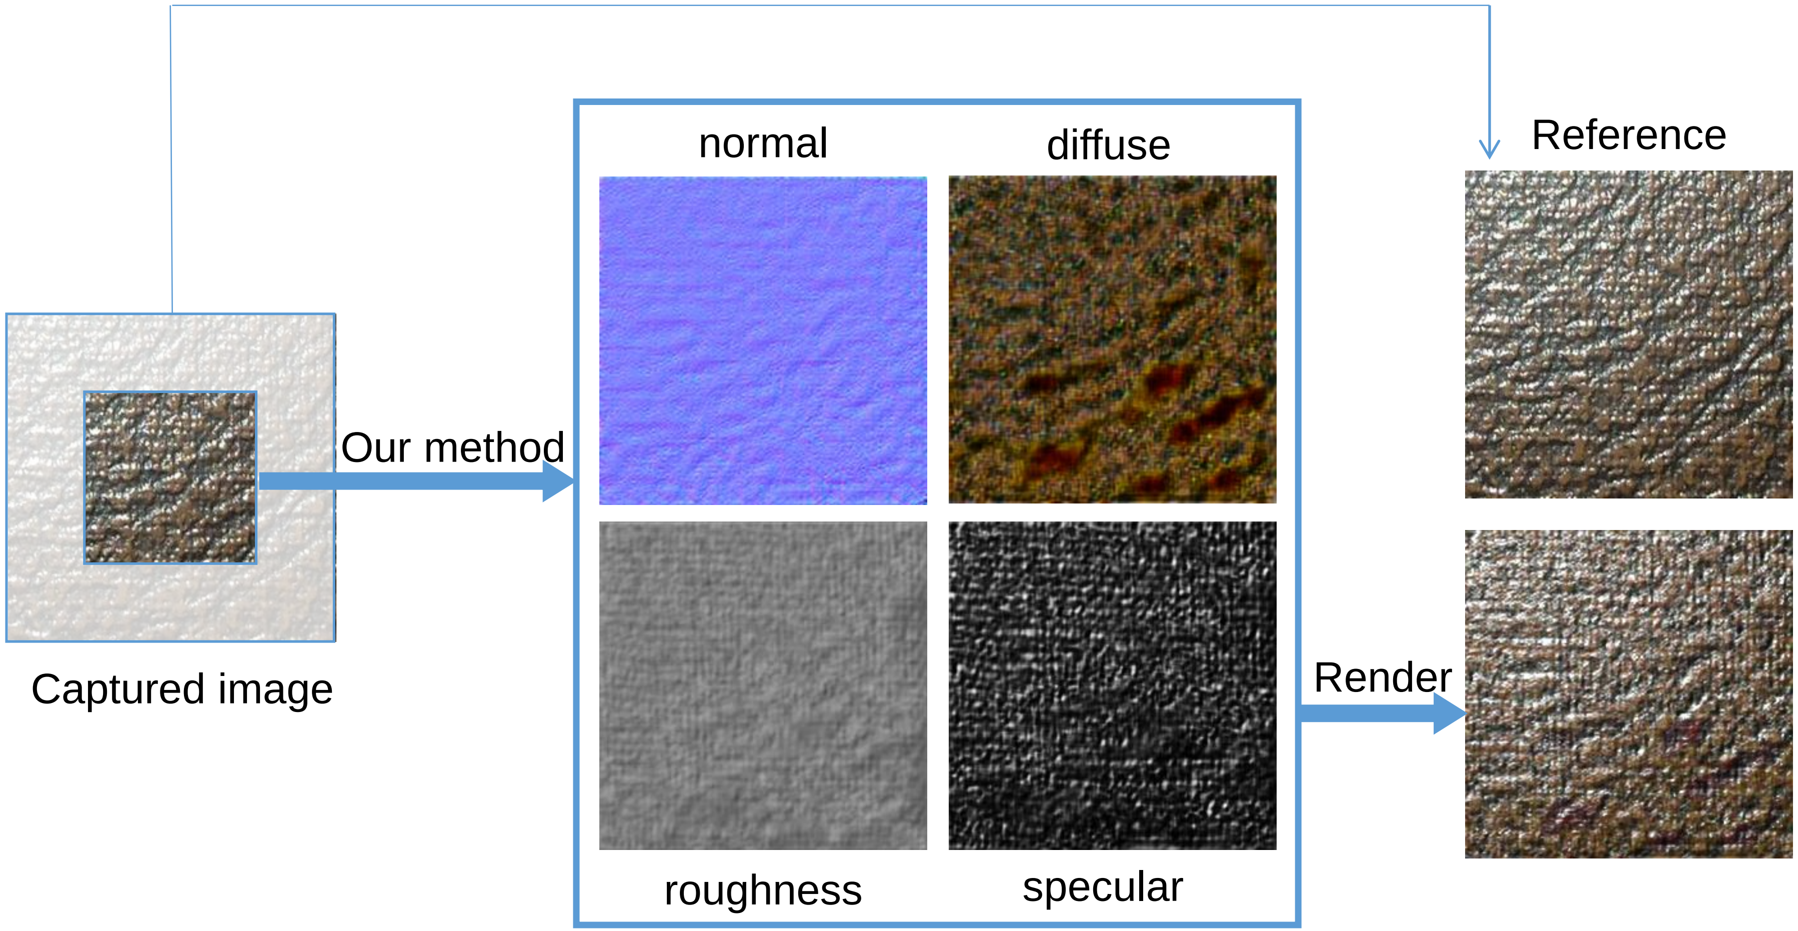

Imaging setup

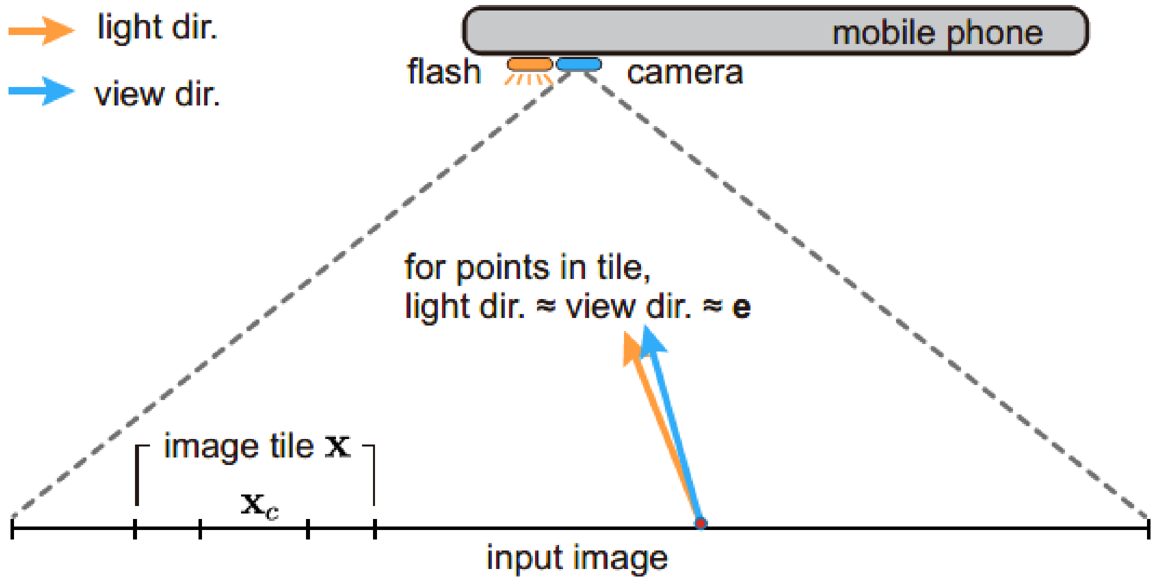

Imaging setup

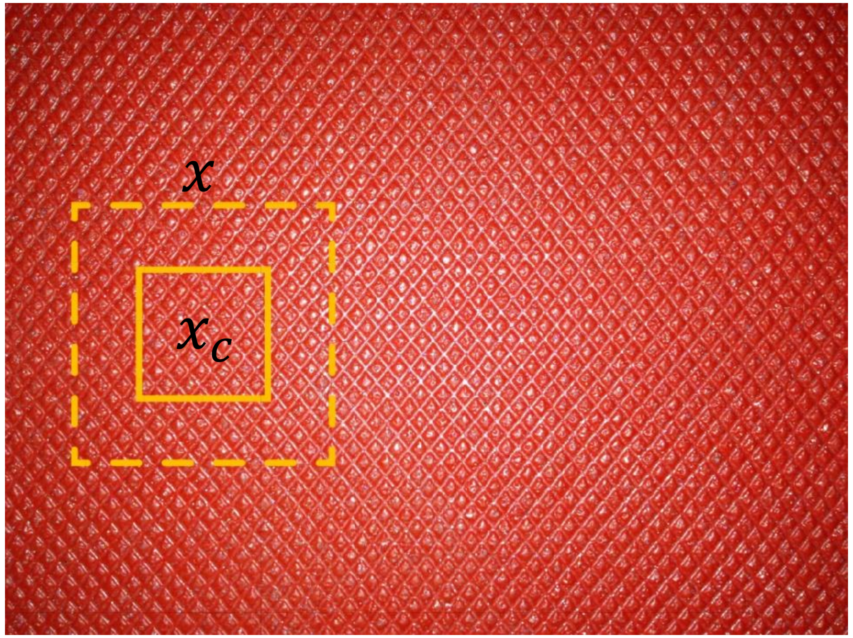

Captured image

Generator

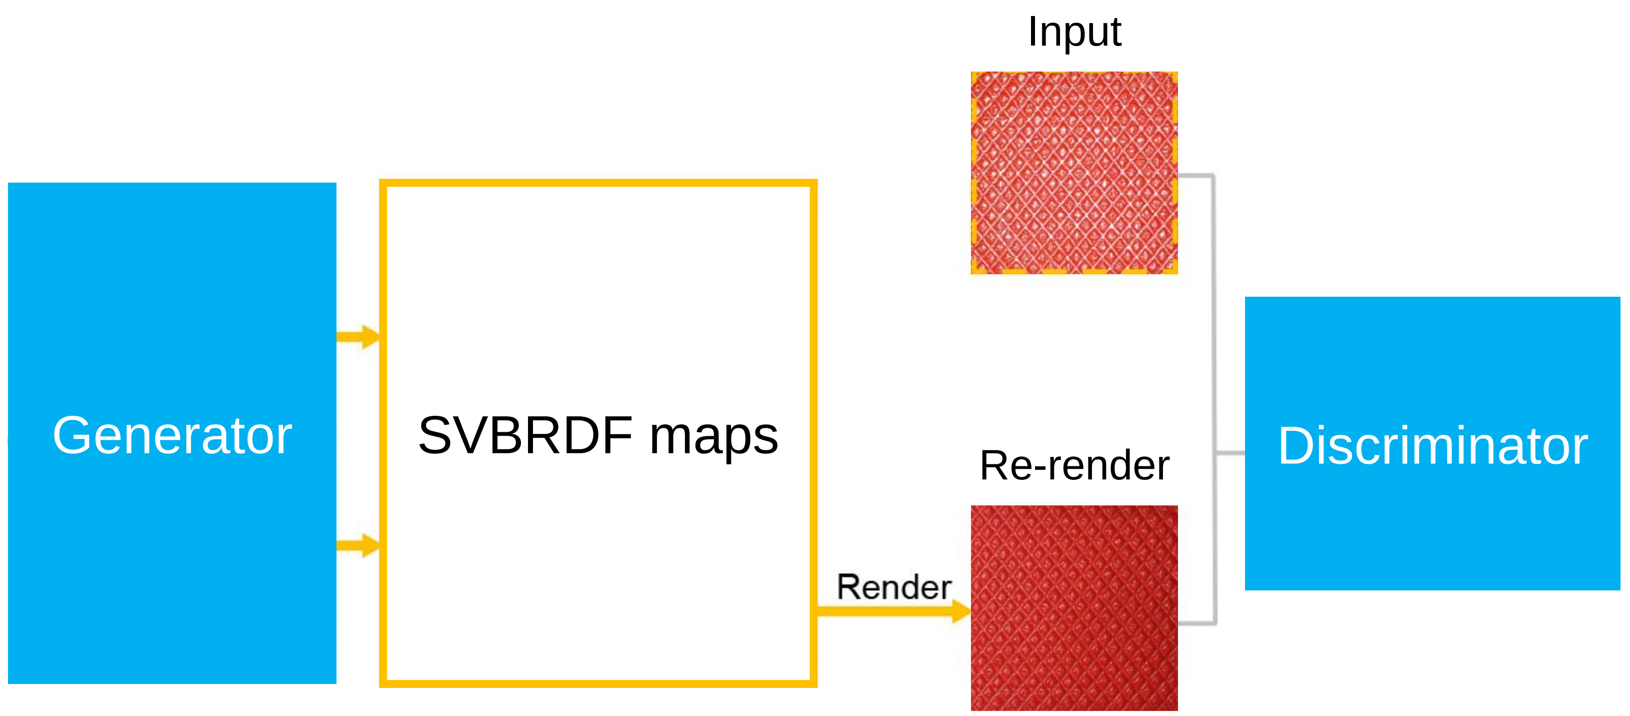

Generative Adversarial Network (GAN)

Generator

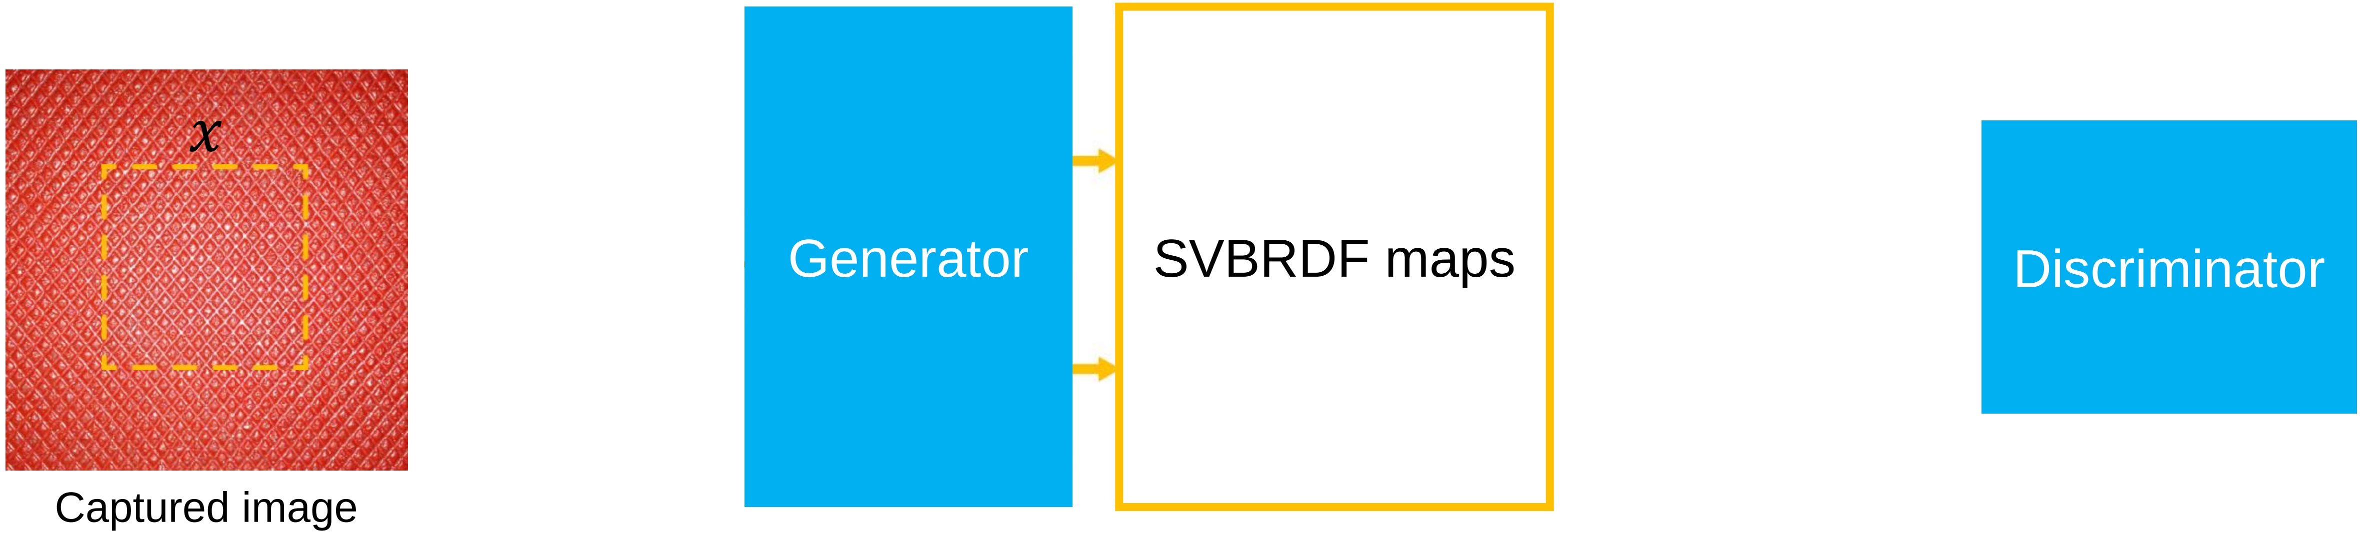

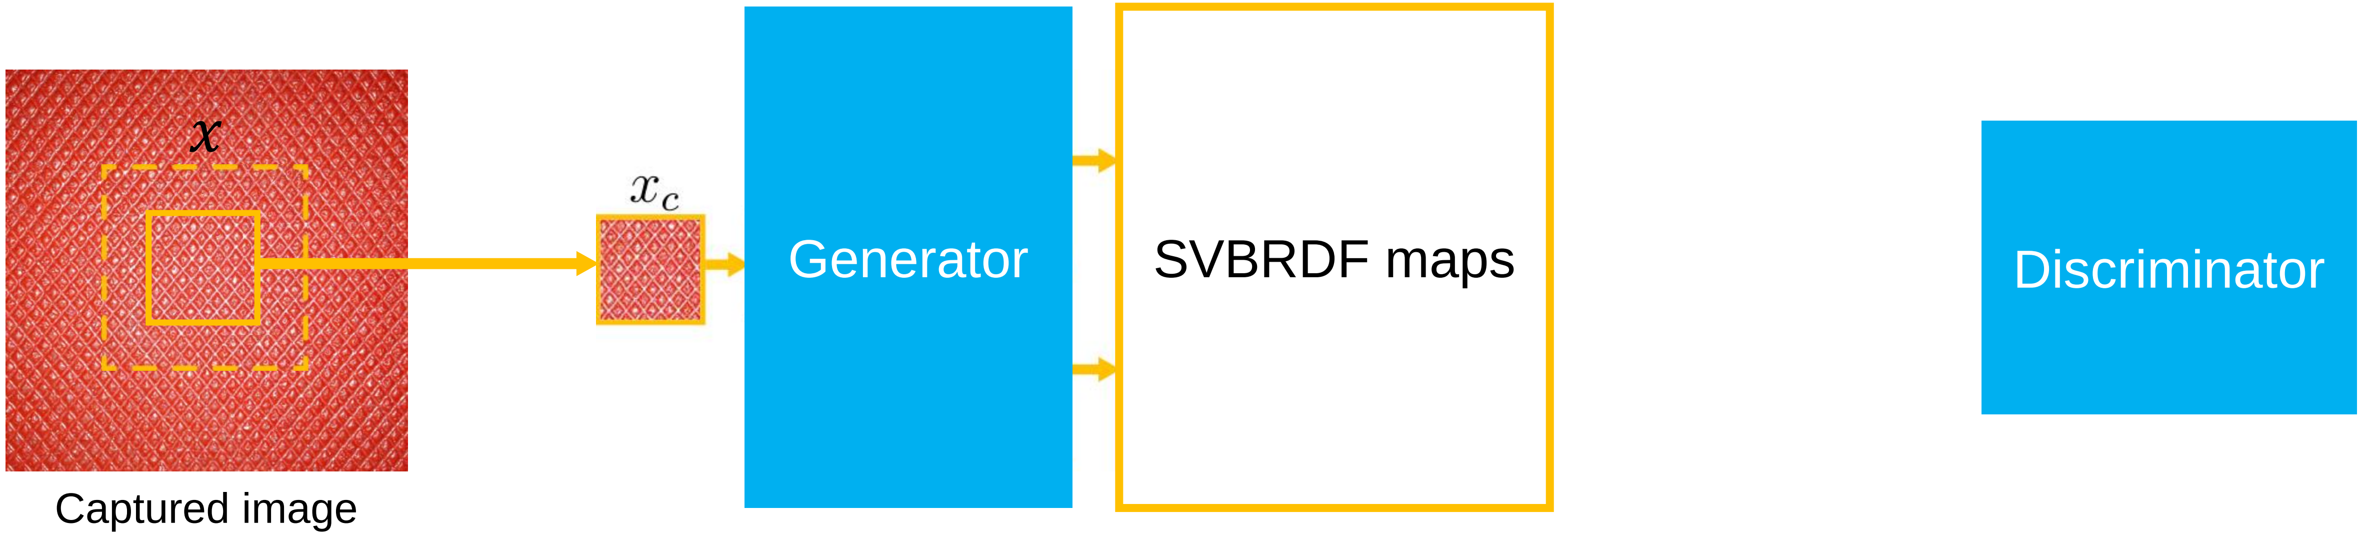
x
x_c
Generator
SVBRDF maps
Discriminator
Captured image

Generator

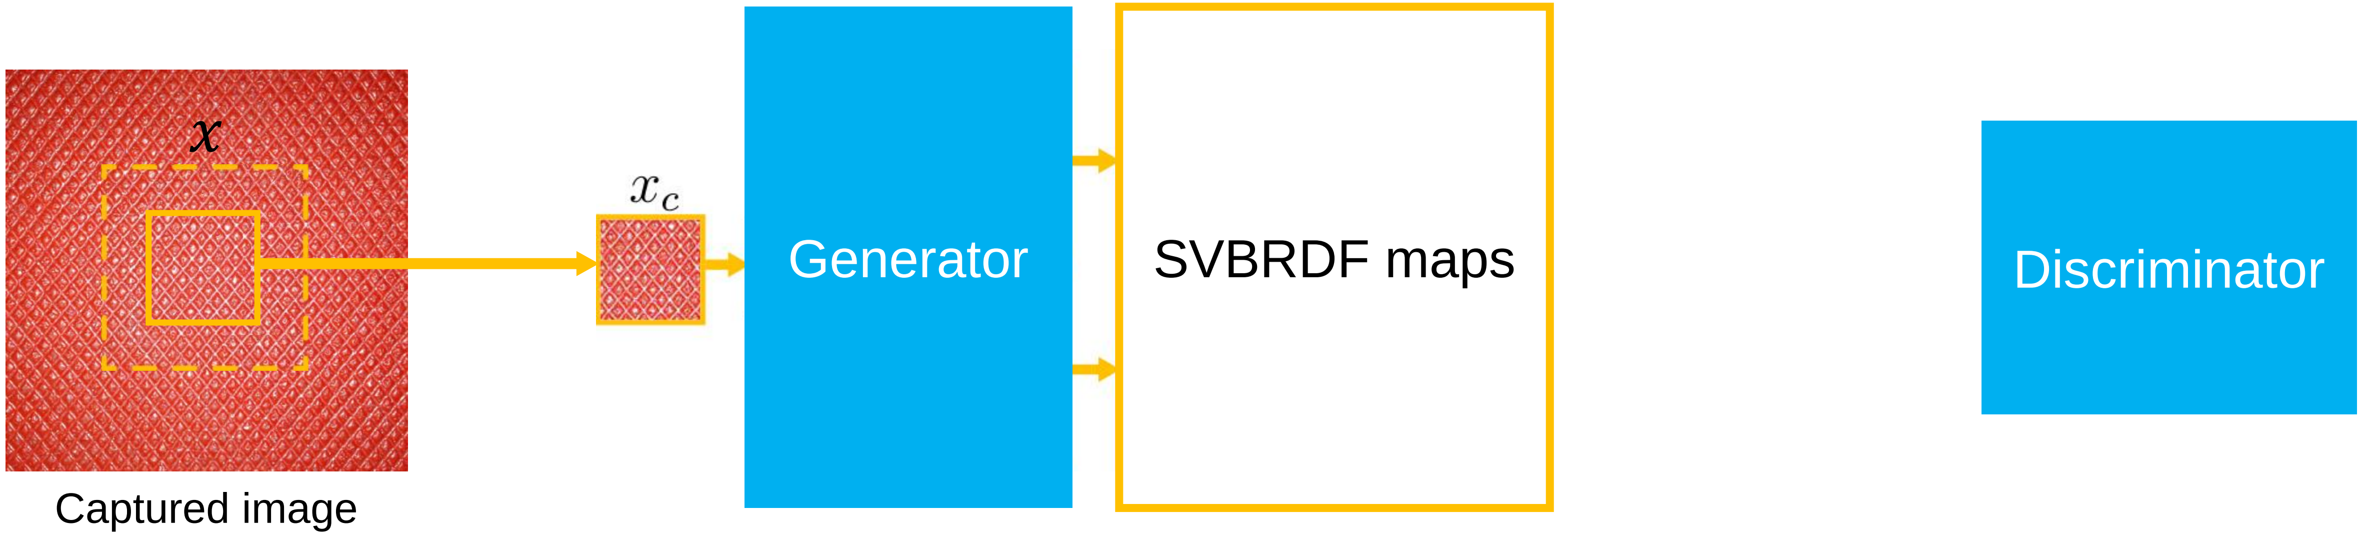

Generator

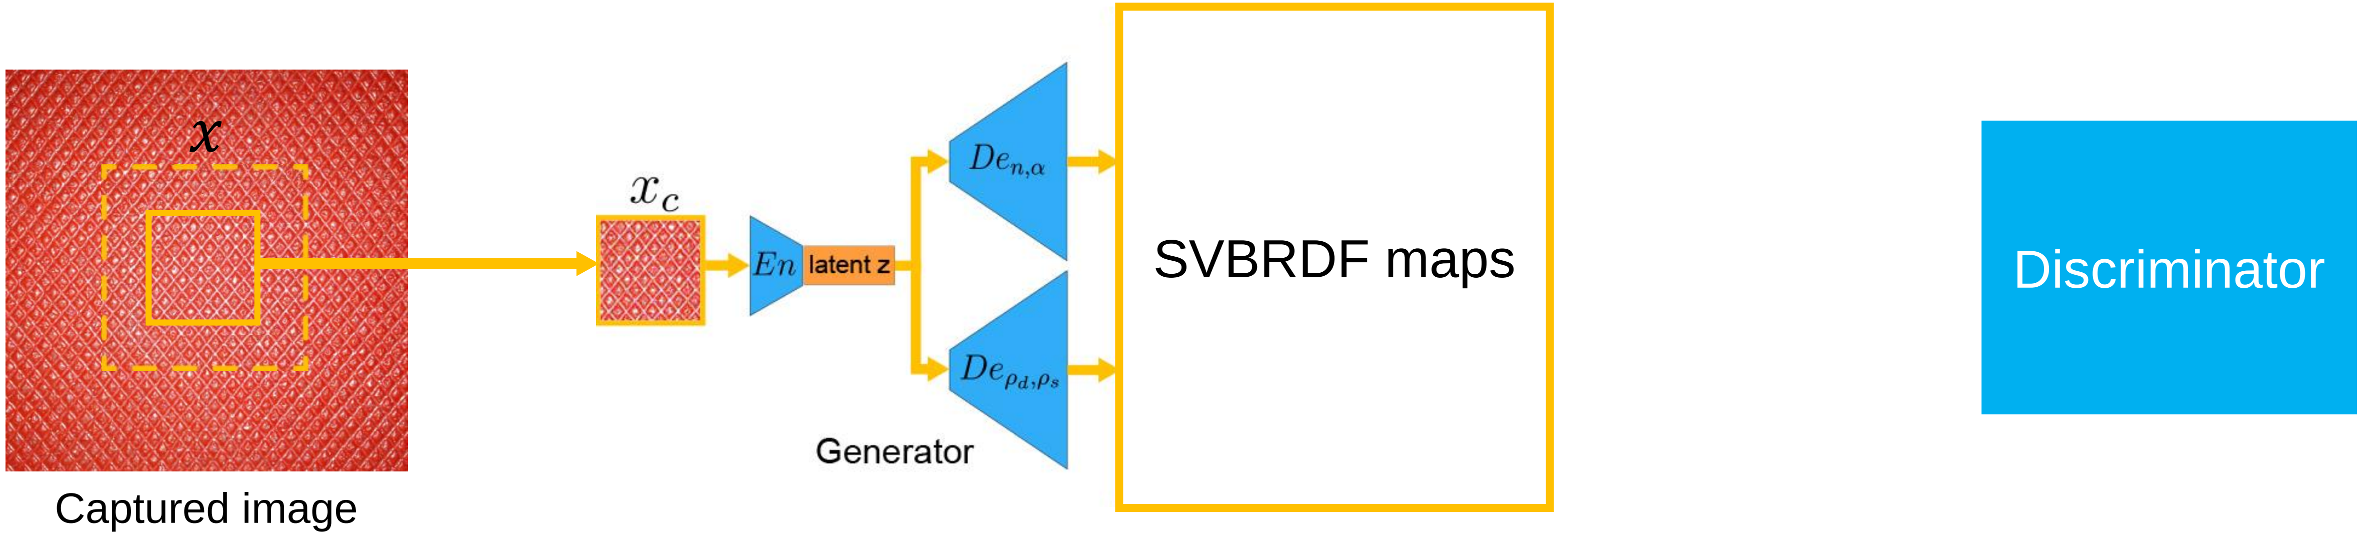

x
x_c
En
latent z
$De_{n,\alpha}$
De_{ρ_d,ρ_s}
Generator
SVBRDF maps
Discriminator
Captured image

Generator

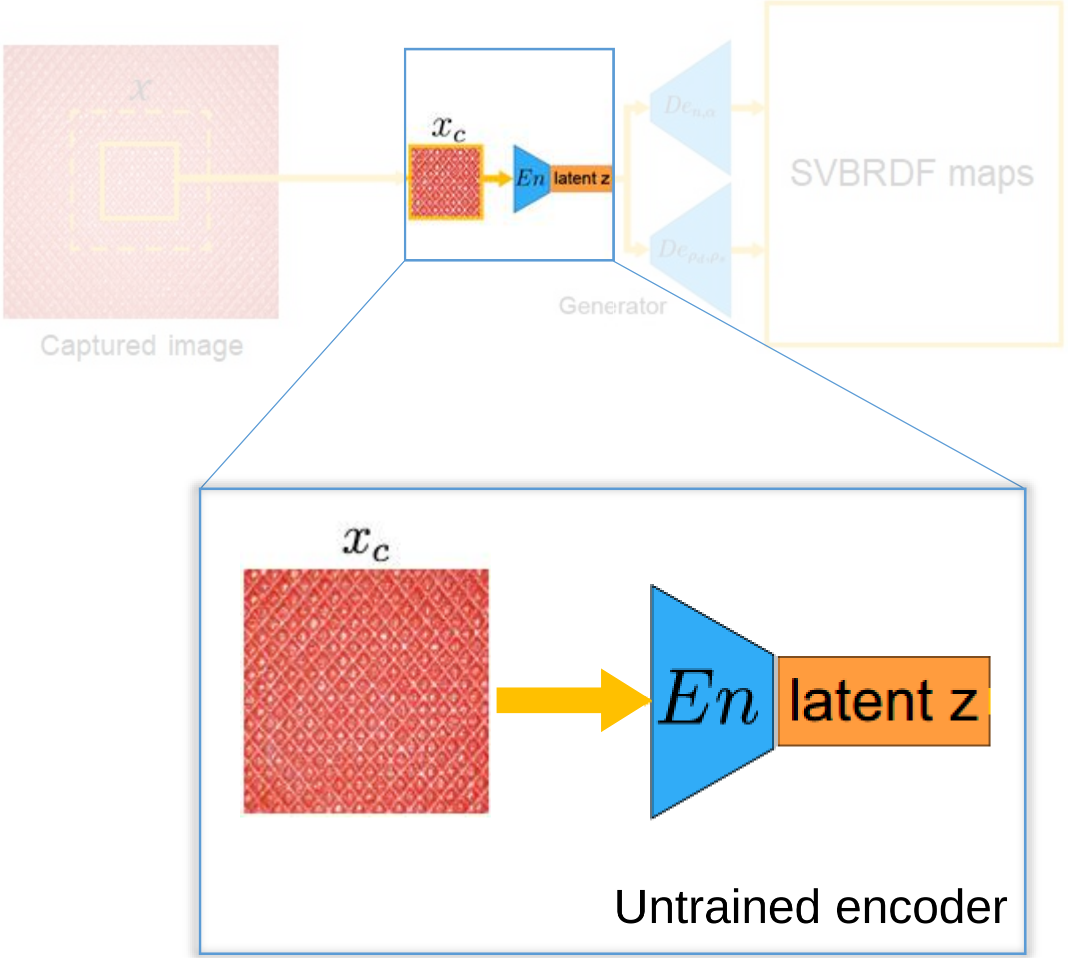

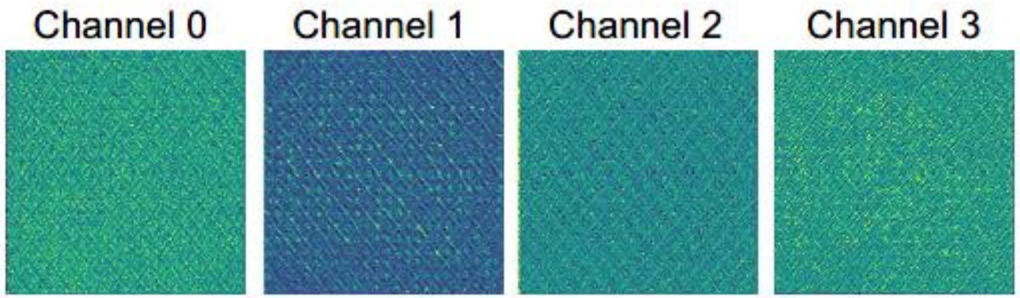

Visualization of the first 4 layers latent vector from encoder

Generator

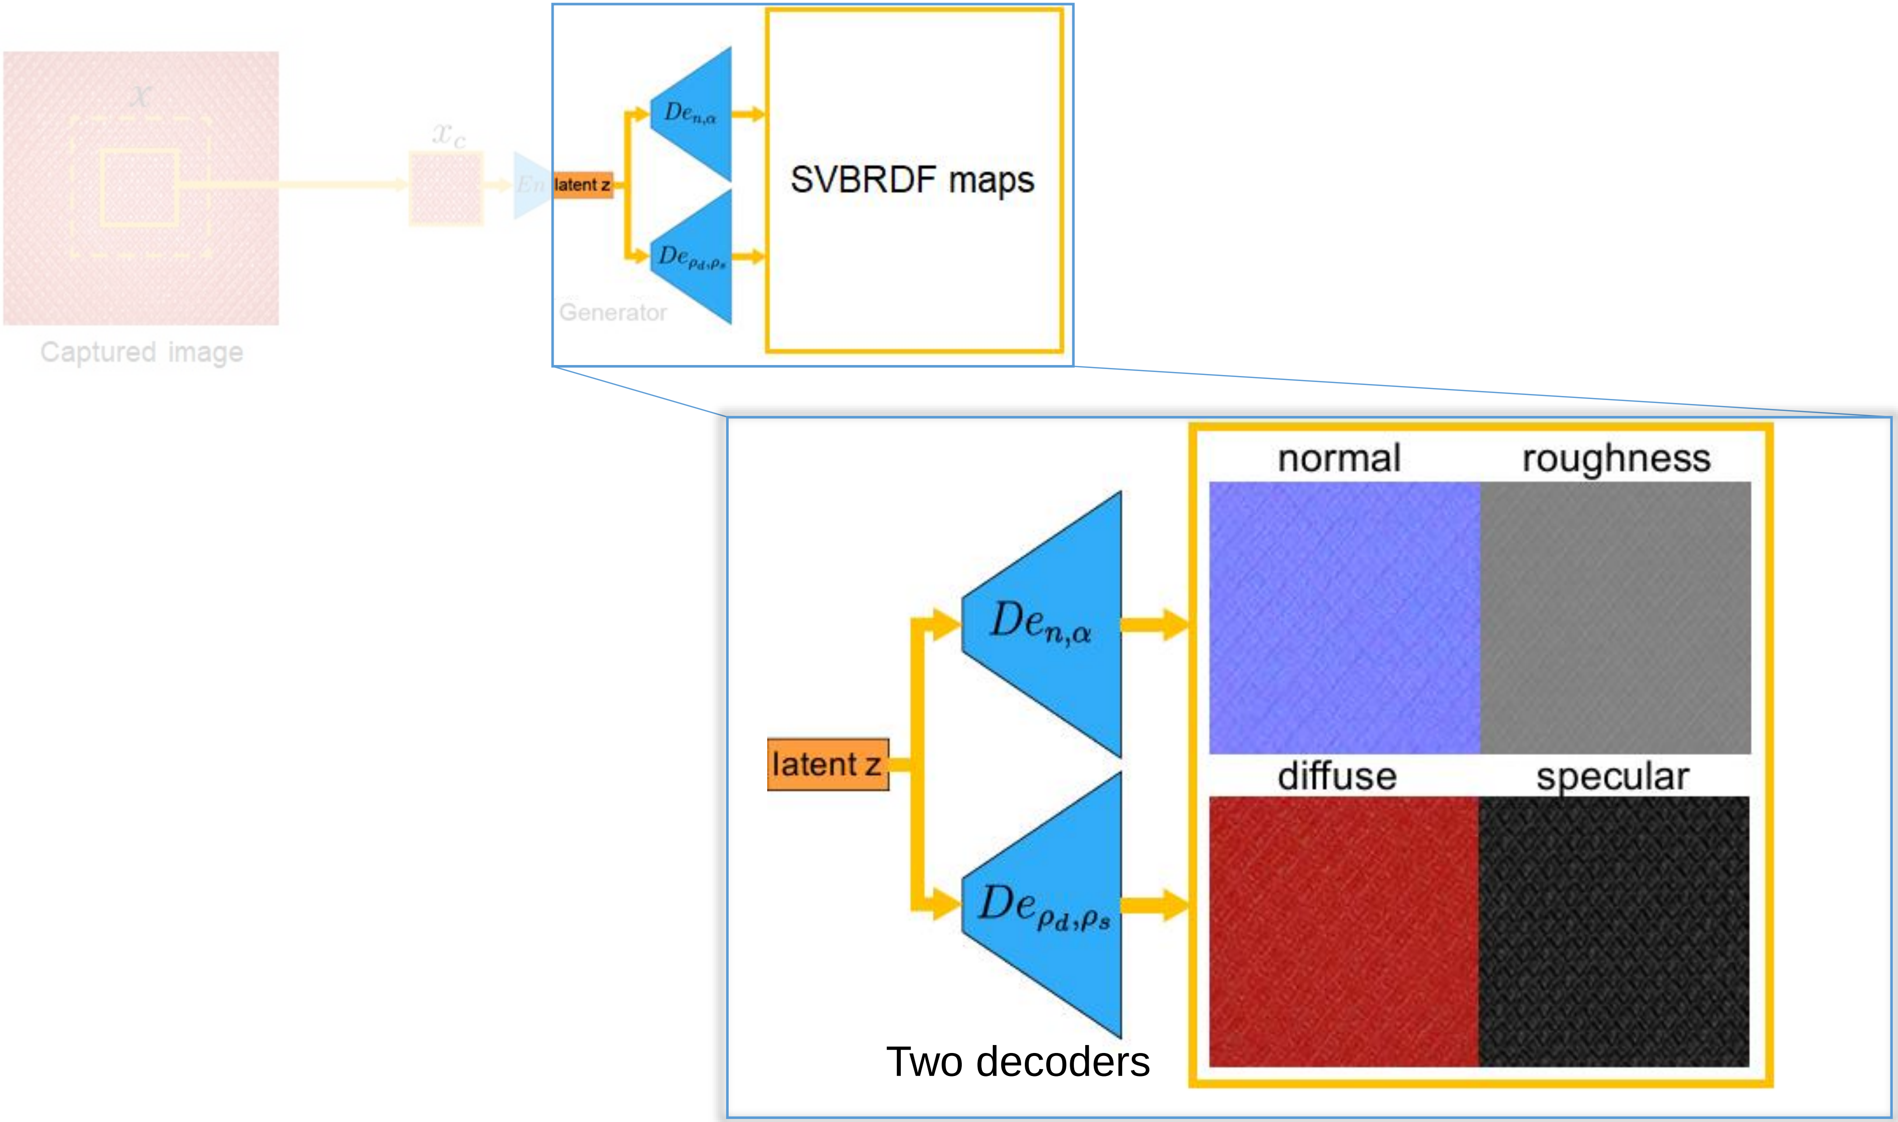

Captured image
latent z
SVBRDF maps
Generator
latent z
normal
roughness
diffuse
specular
Two decoders

Discriminator

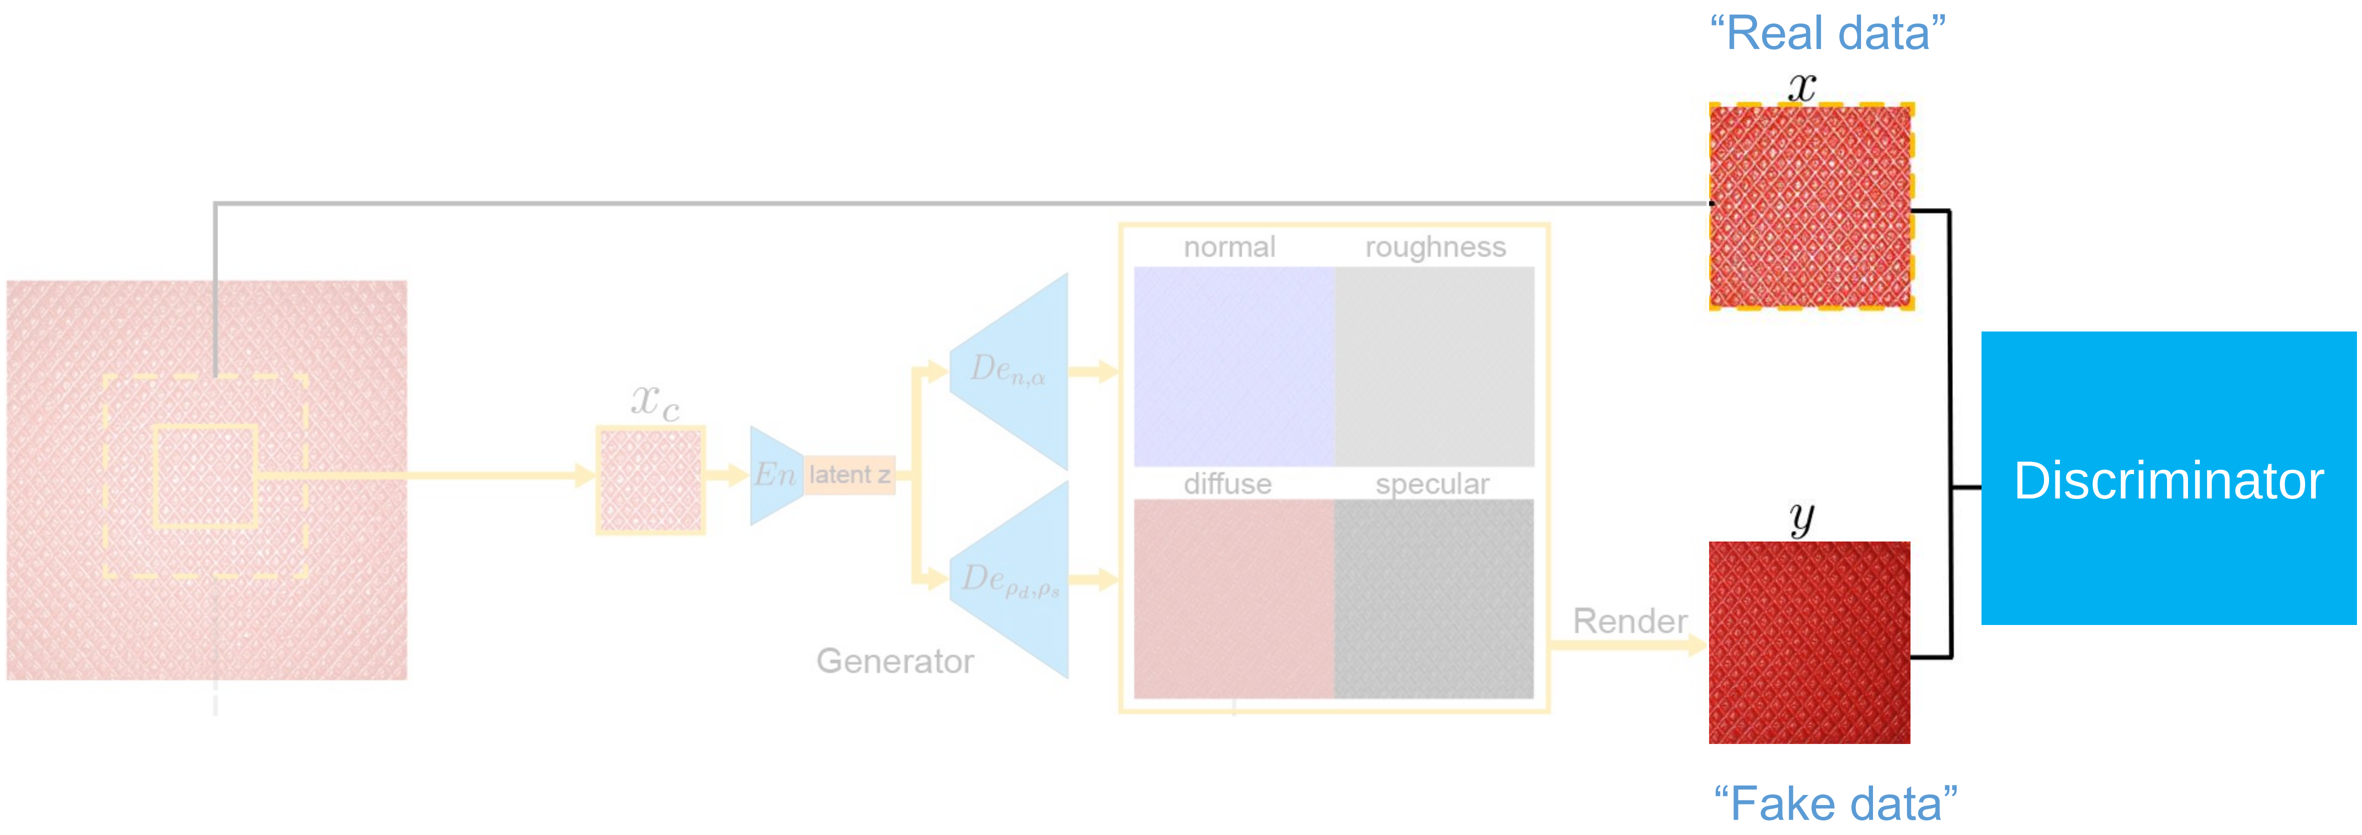

Loss function

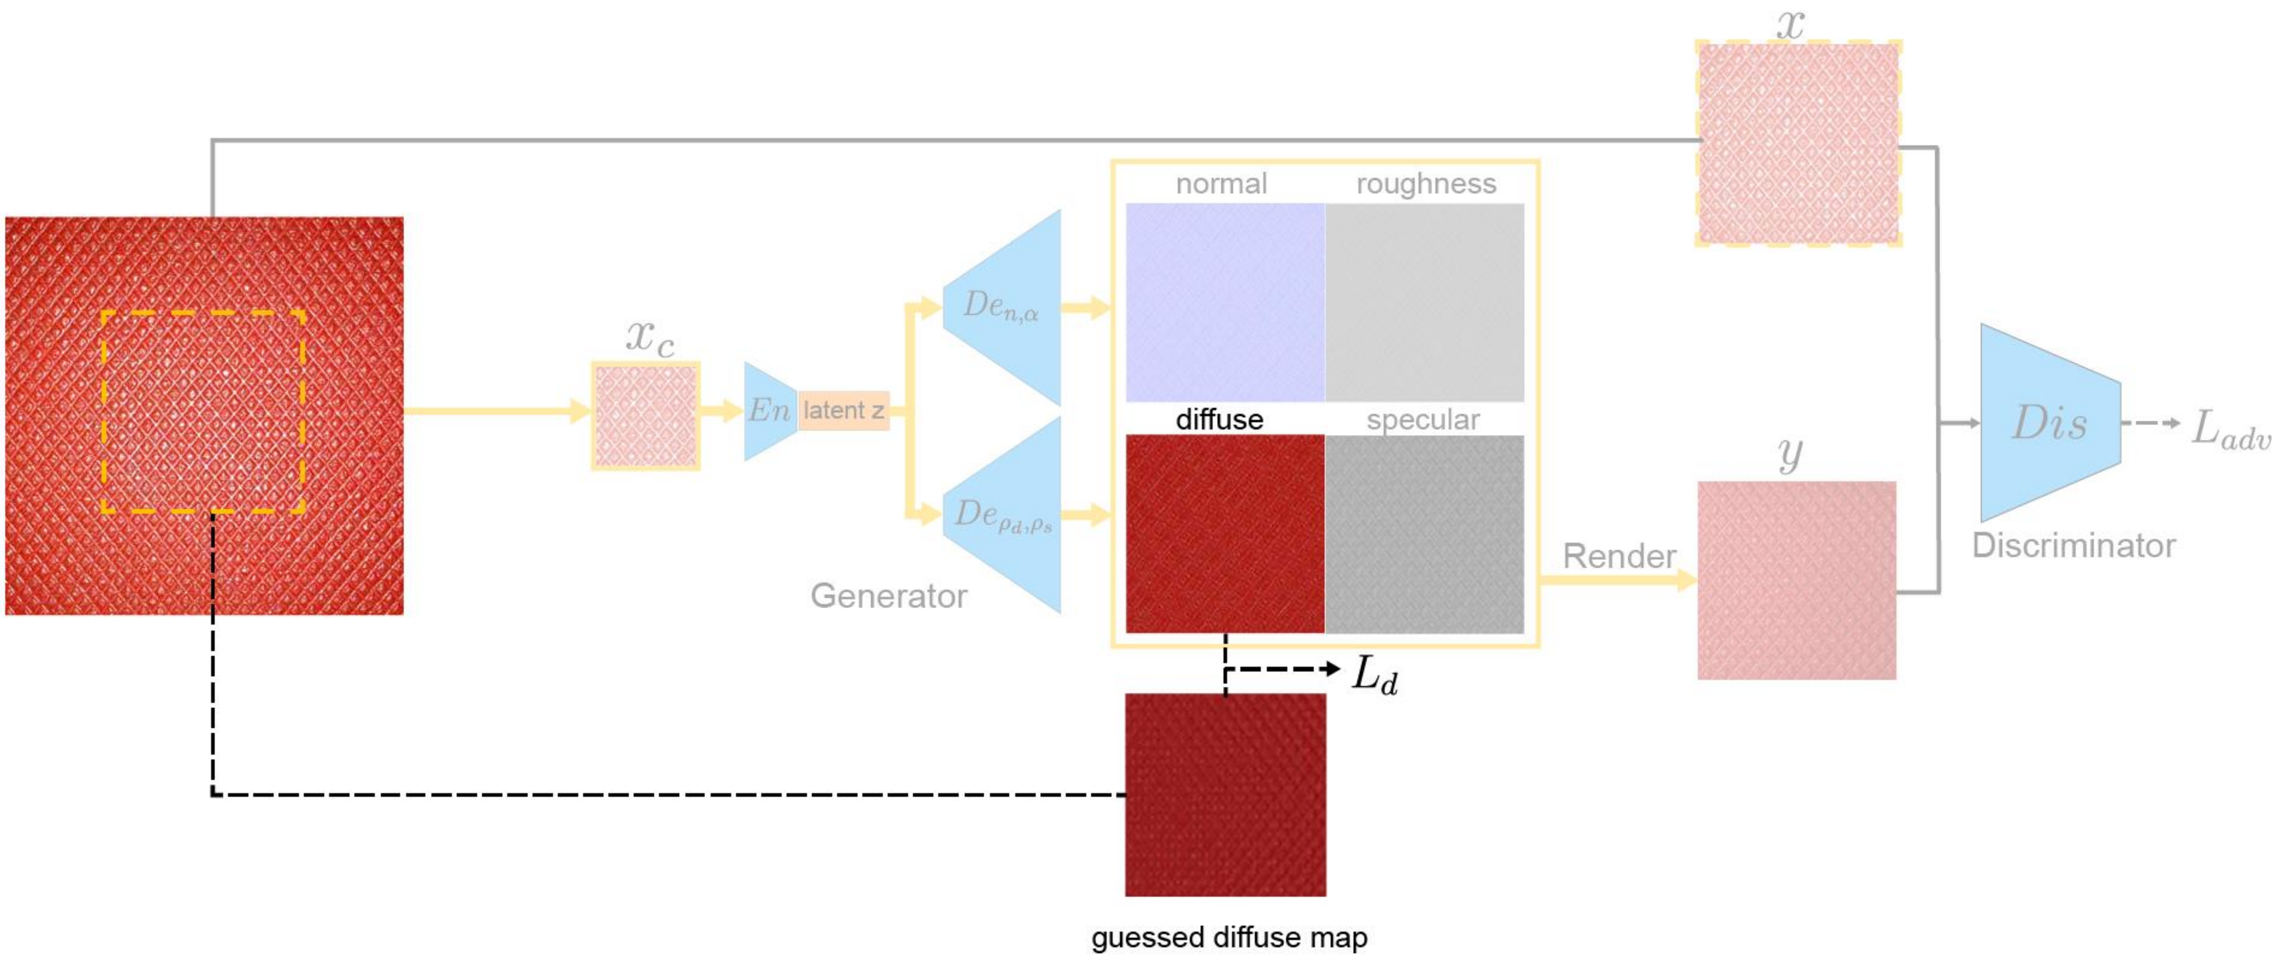

Guessed diffuse map

Input image

Guessed diffuse map

Computed as in [AAL16]

[AAL16] Aittala et al.
Reflectance modeling by neural texture synthesis.

Loss function

$$\mathcal{L}_{final} = \lambda \mathcal{L}_{GAN}(G,D) + \mathcal{L}_d(G),$$

Adversarial loss

$$\mathcal{L}_{GAN}(G,D) = \mathbb{E}[\log Dis(\mathbf{x})] + \mathbb{E}[\log(1 - Dis(\mathbf{y})],$$

Tile x Re-render y

Discriminator( , )

L1 loss

$$\mathcal{L}_d(G) = \mathbb{E}\left[||\tilde{\rho_{\mathbf{d}}} - \rho_{\mathbf{d}}||_1\right].$$

Guessed diffuse Generated diffuse

|  —  $|_1$

Loss function

$$\mathcal{L}_{final} = \lambda \mathcal{L}_{GAN}(G,D) + \mathcal{L}_{render}(G)$$

Adversarial loss

$$\mathcal{L}_{GAN}(G,D) = \mathbb{E}[\log Dis(\mathbf{x})] + \mathbb{E}[\log(1 - Dis(\mathbf{y})],$$

Tile x Re-render y

Discriminator(,)

L1 loss

$$\mathcal{L}_{render}(G) = \mathbb{E}\left[||\mathbf{x} - \mathbf{y}||_1\right]$$

Tile x Re-render y

|  —  |$_1$

☹

Loss function

$$\mathcal{L}_{final} = \lambda \mathcal{L}_{GAN}(G,D) + \mathcal{L}_d(G),$$

Input

Diffuse

Highlights

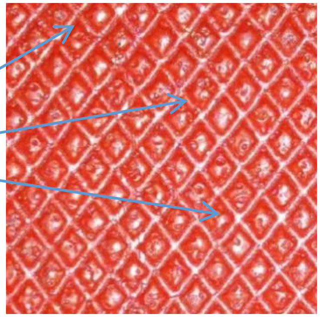

$$\mathcal{L}_{final} = \lambda \mathcal{L}_{GAN}(G,D) + \mathcal{L}_{render}(G)$$

Input

Diffuse

Highlights

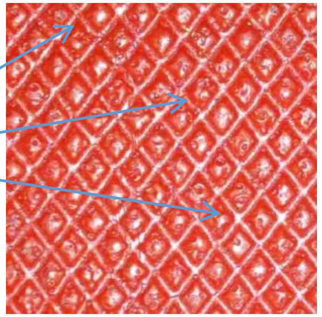

Highlights

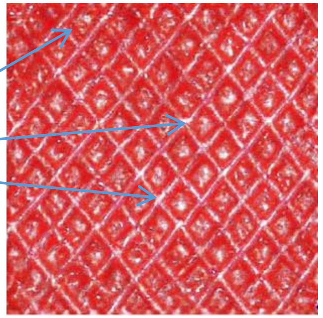

☹

Results

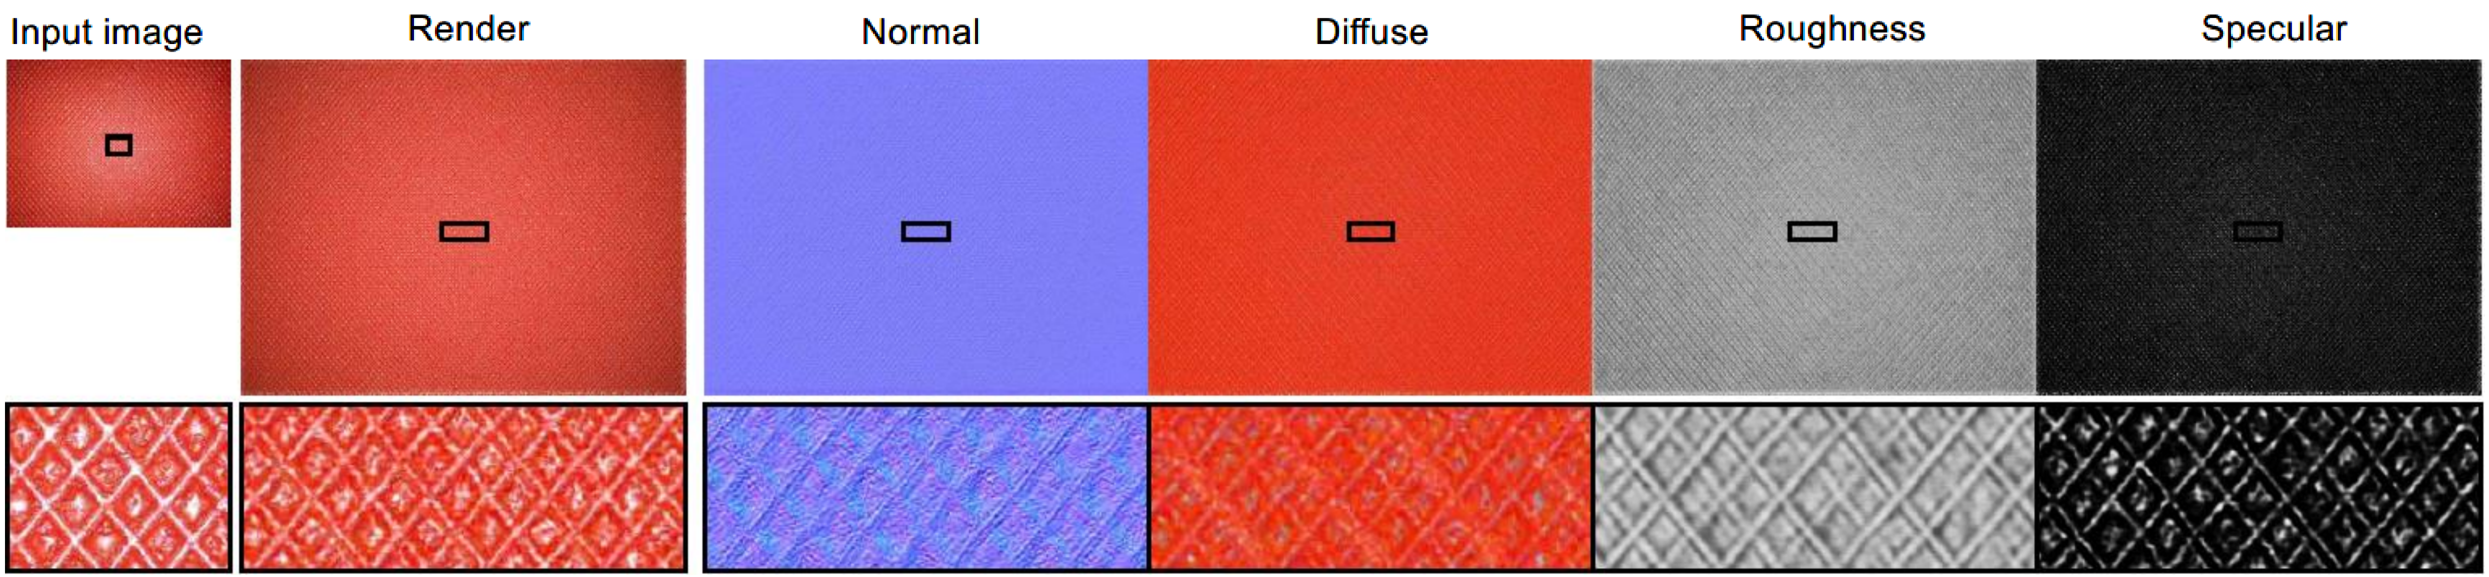

Input image: 1632×1224

SVBRDF & Render: 3264×2448

Results

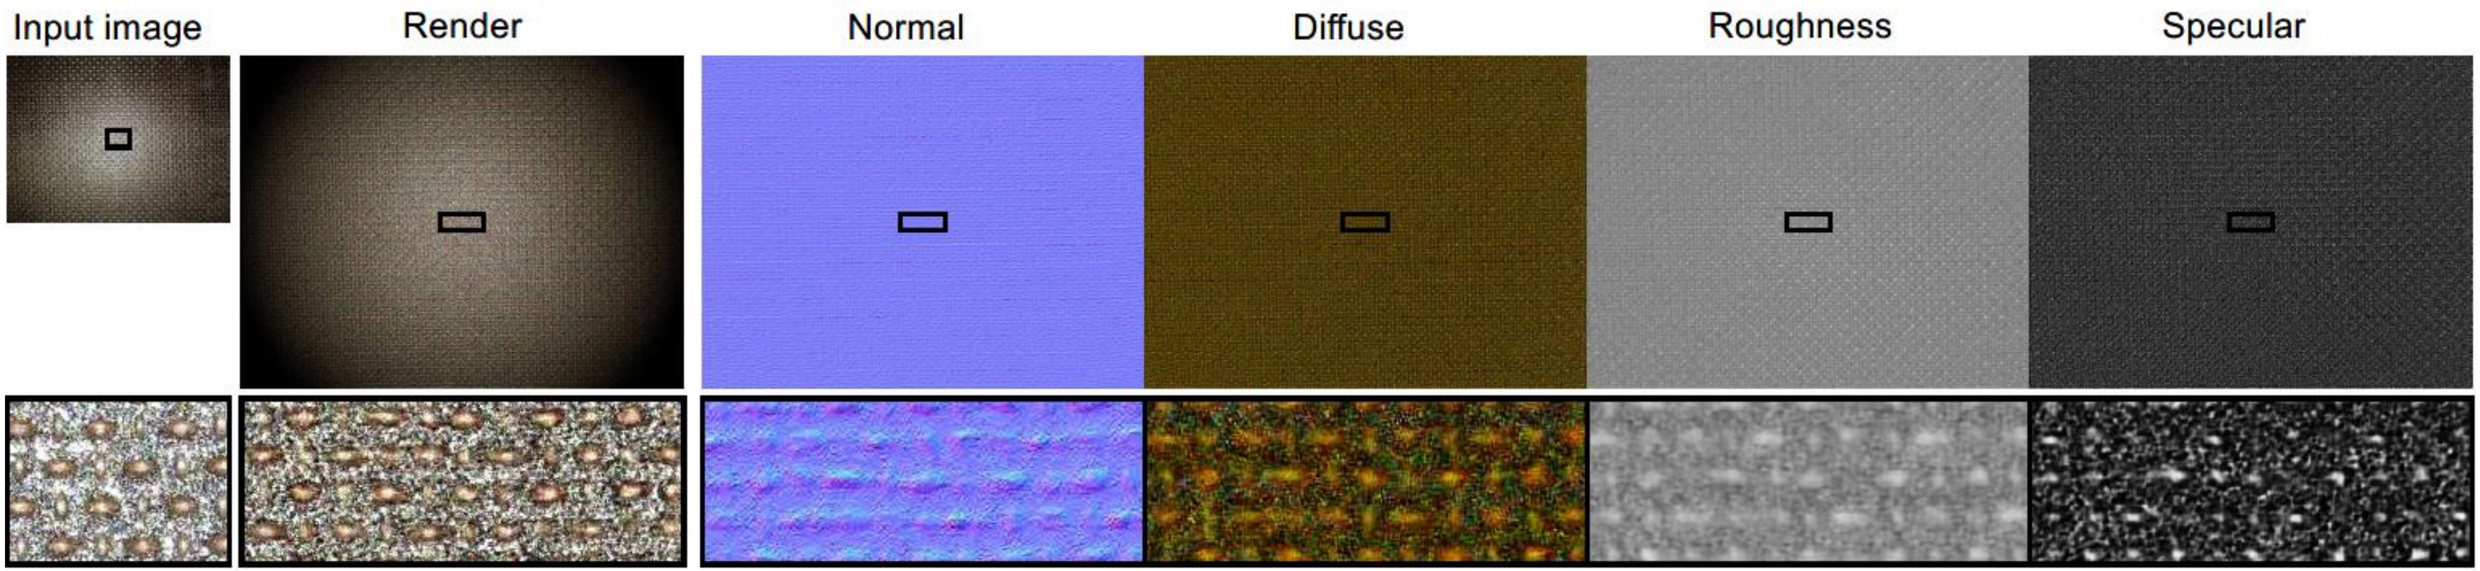

Input image: 1632 × 1224

SVBRDF & Render: 3264 × 2448

Results

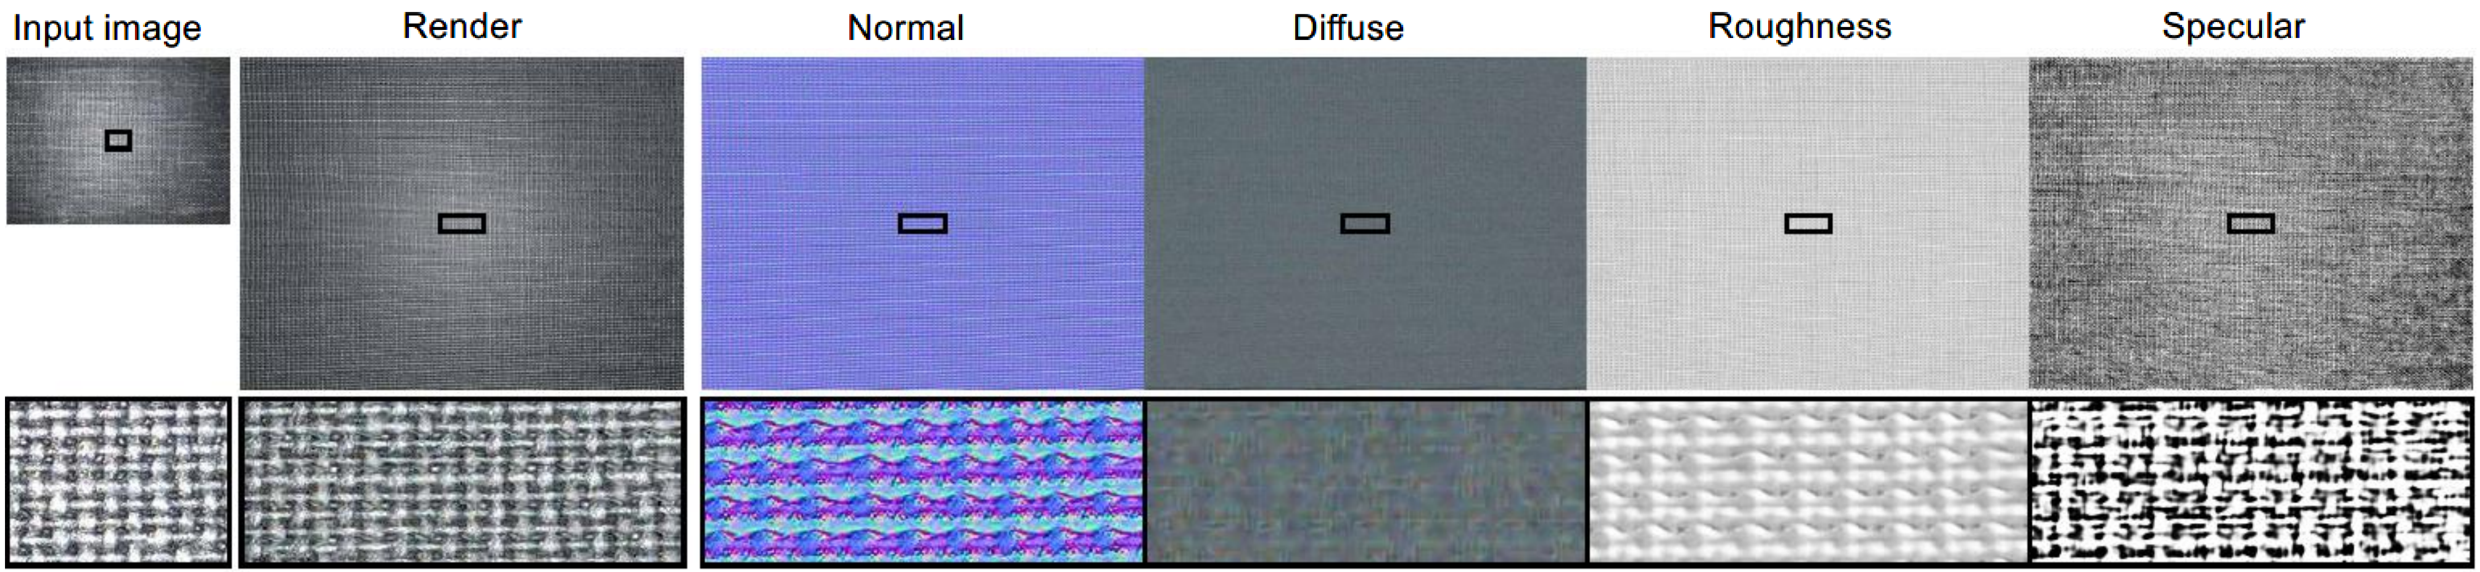

Input image: 1632×1224

SVBRDF & Render: 3264×2448

Results

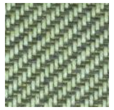

×1

×2

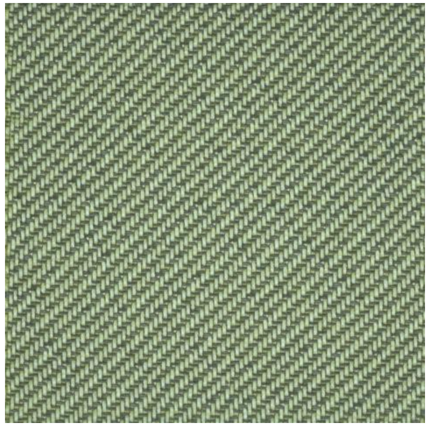

×4

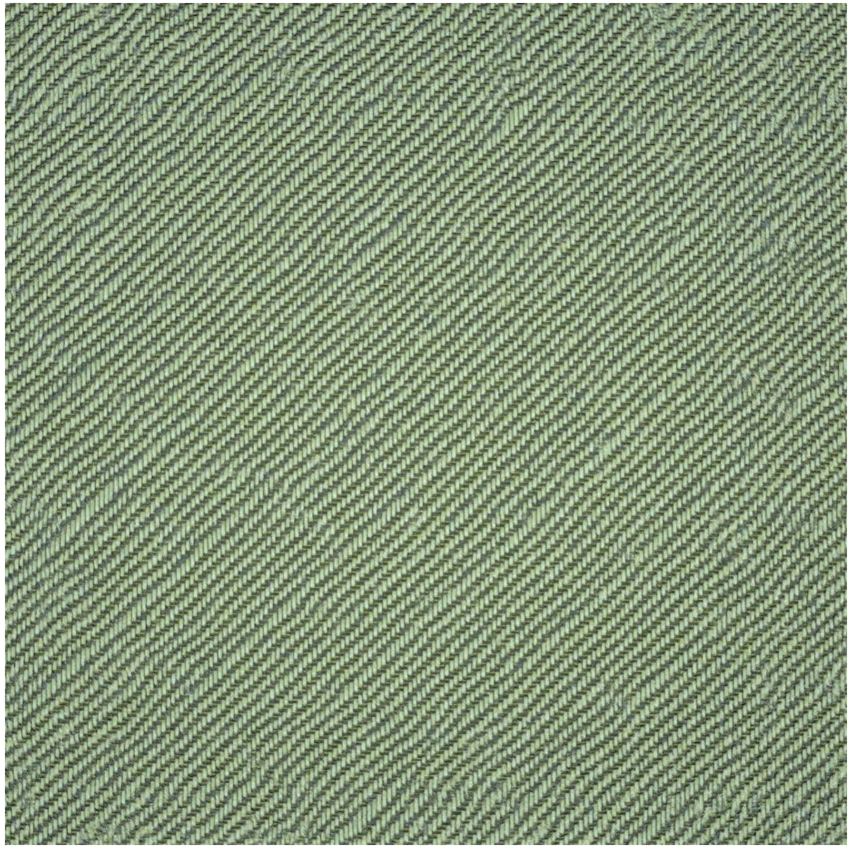

×8

Results

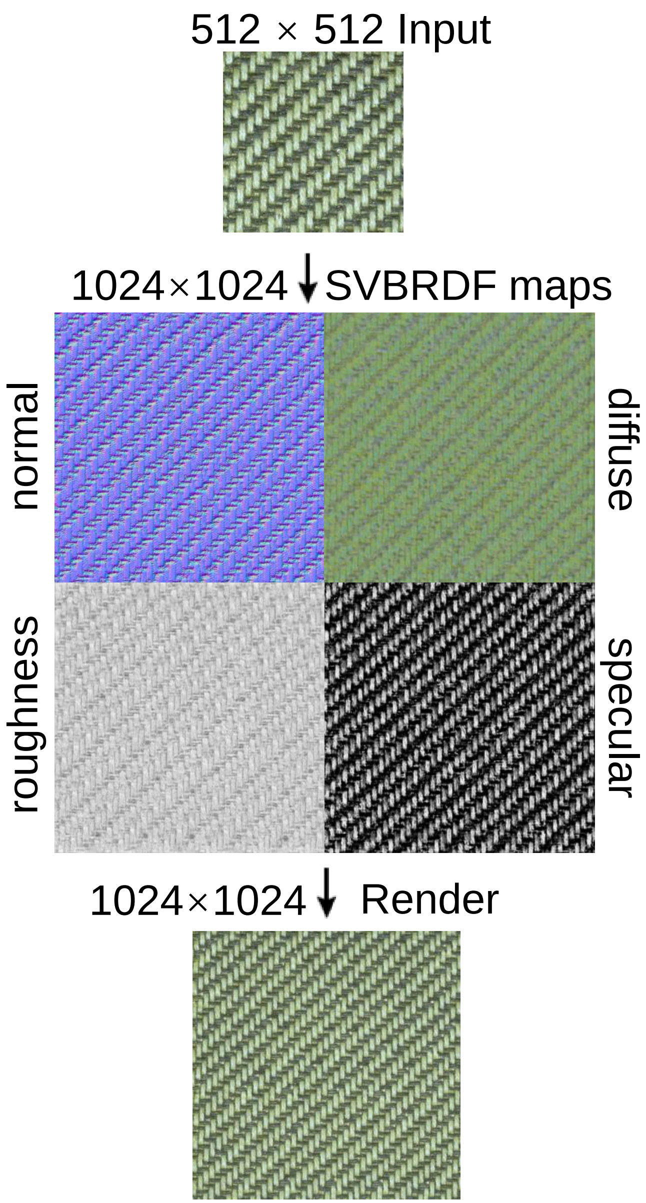
512 × 512 Input
1024×1024 SVBRDF maps
normal
diffuse
roughness
specular
1024×1024 Render

Results

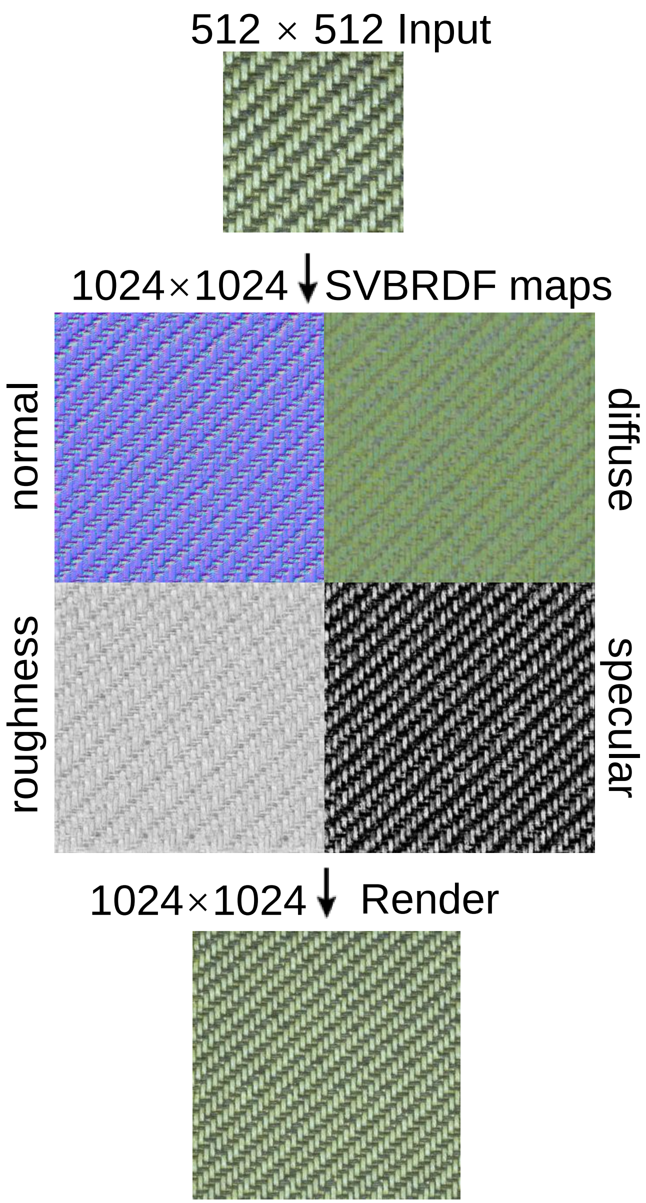
512 × 512 Input
1024×1024 SVBRDF maps
normal
diffuse
roughness
specular
1024×1024 Render

Results

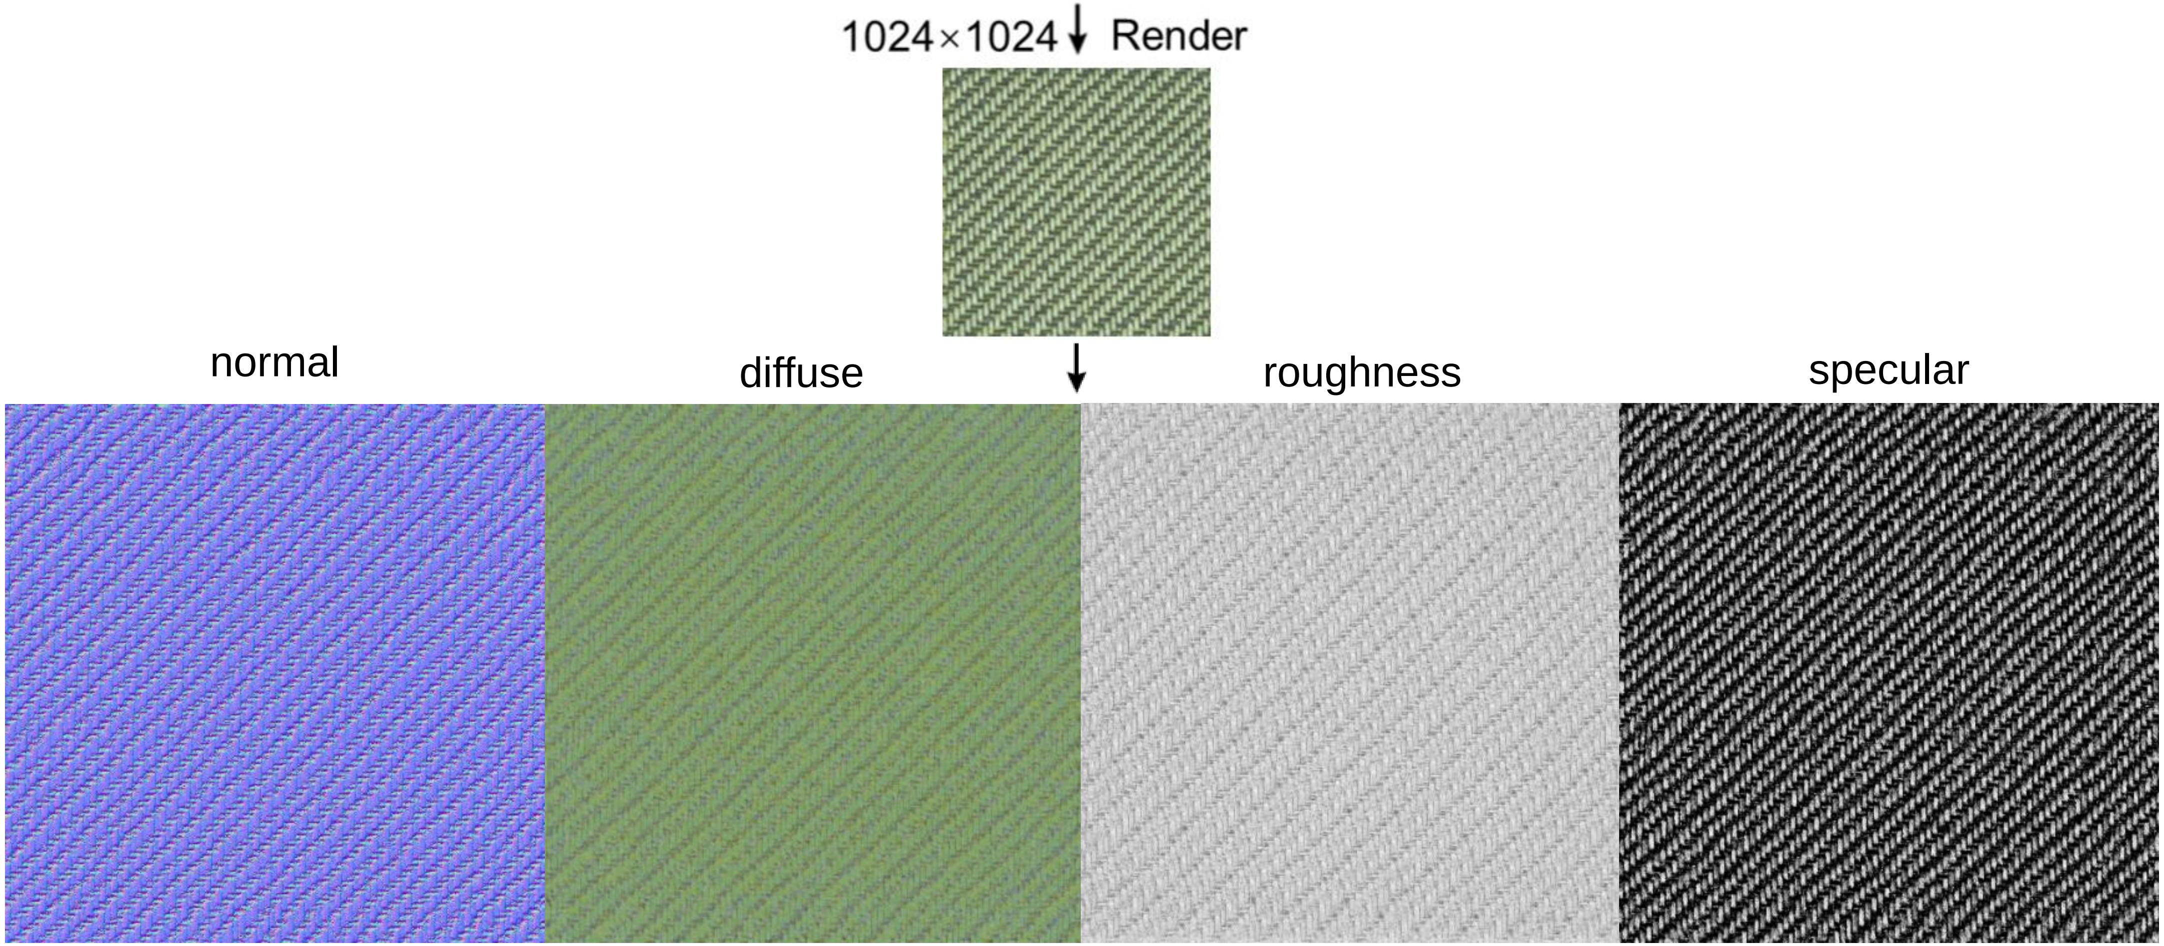

2048×2048 SVBRDF maps

Results

1024×1024 Render

normal diffuse roughness specular

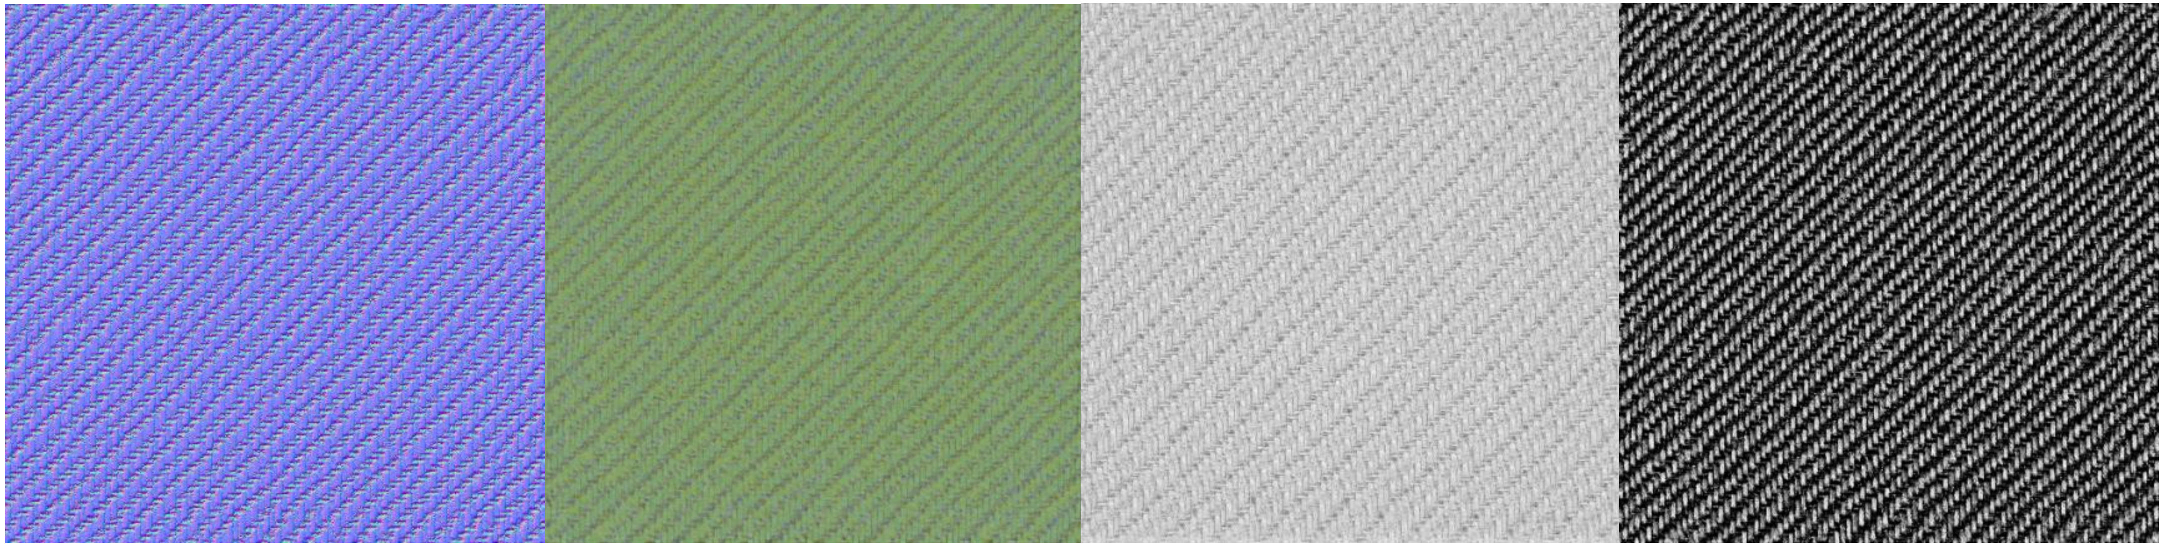

2048×2048 SVBRDF maps

2048×2048 Render (omit)

Results

2048×2048 Render (omit)

normal diffuse roughness specular

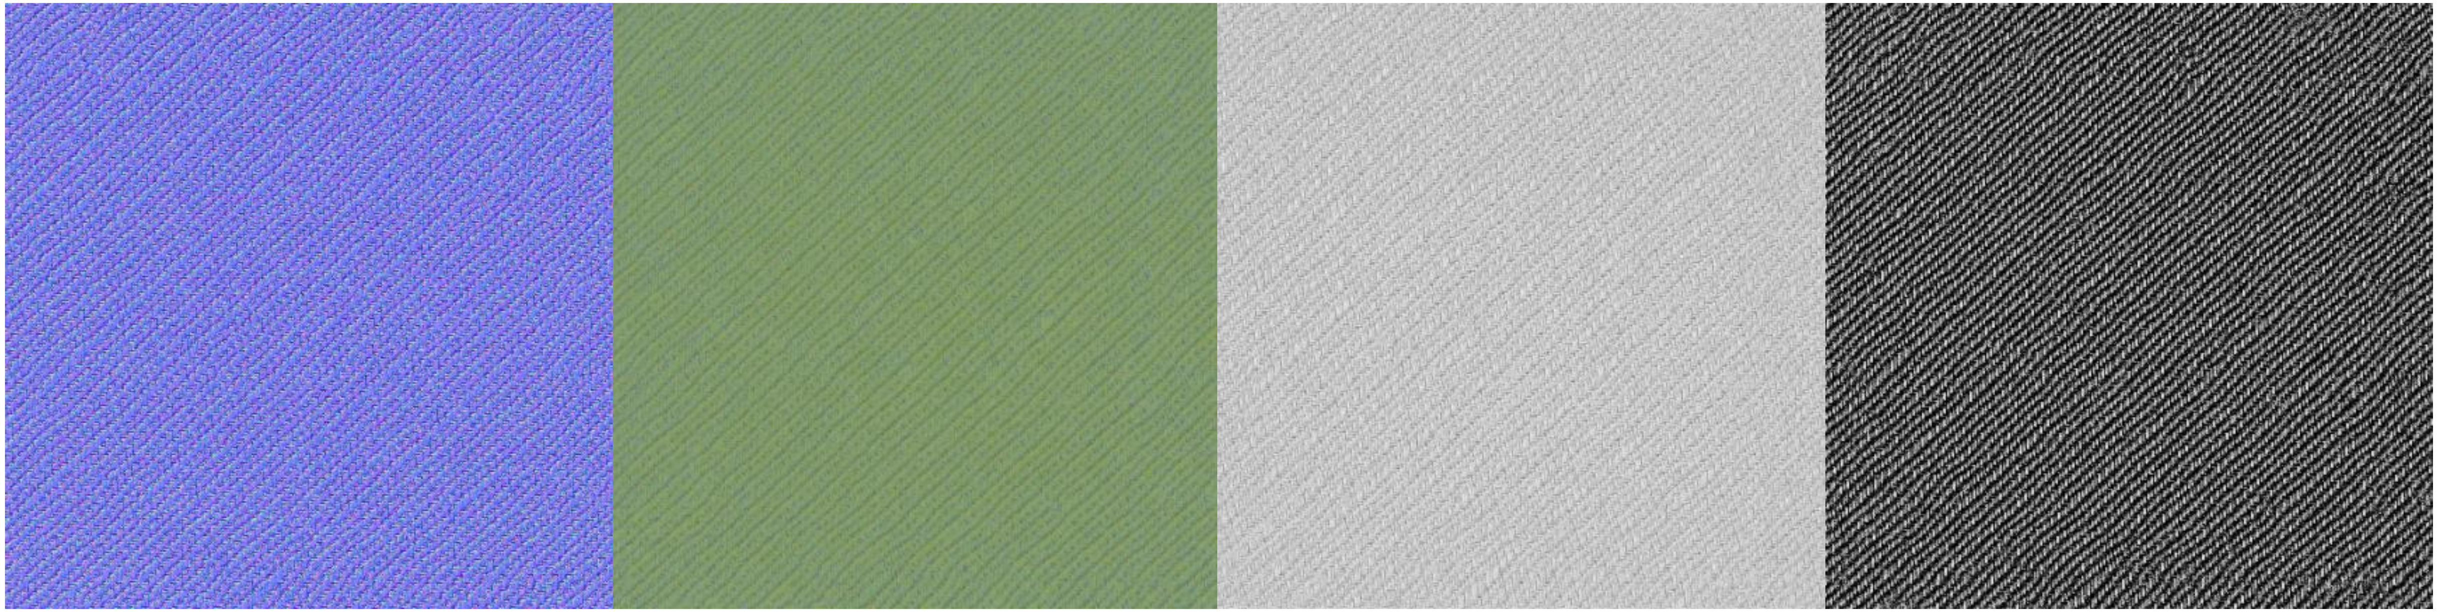

4096×4096 SVBRDF maps

Results

Captured images

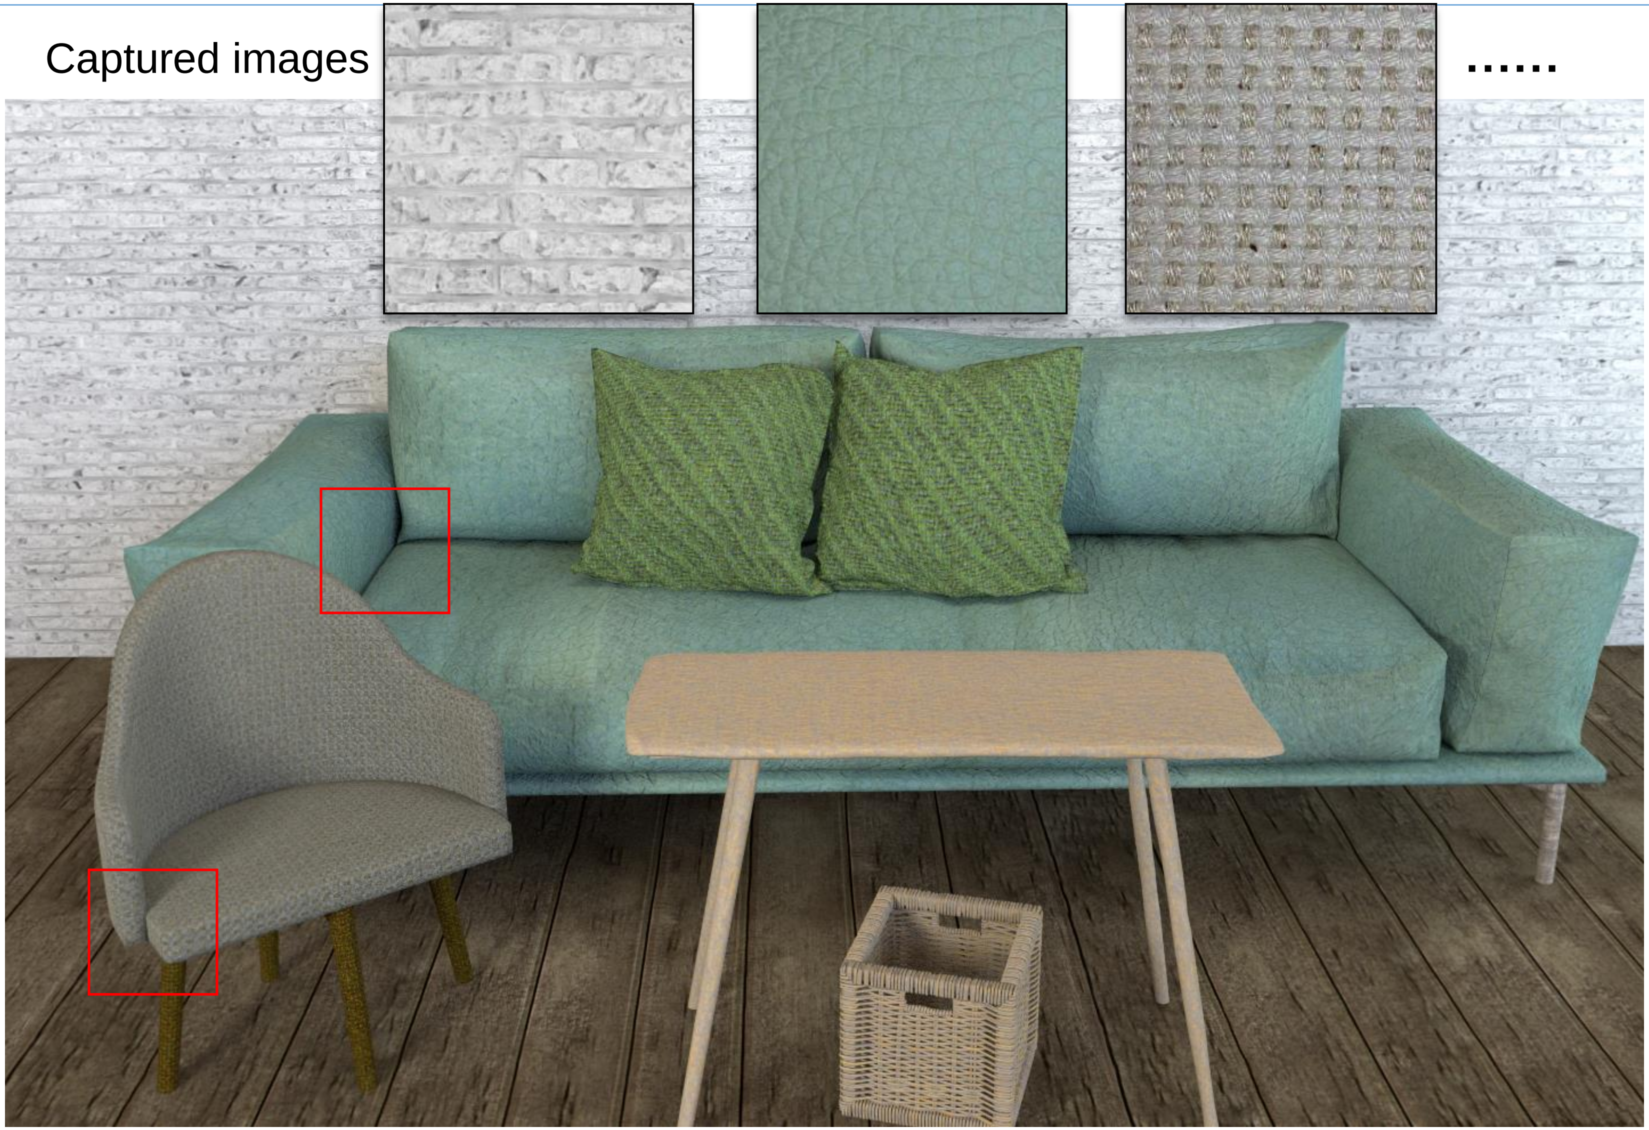

......

Rendered with Arnold

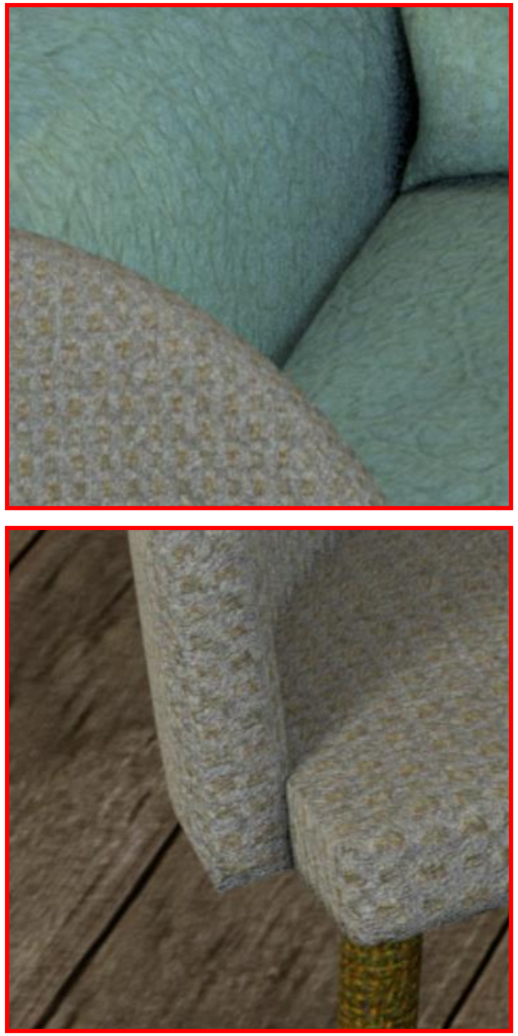

Network Analysis – loss function

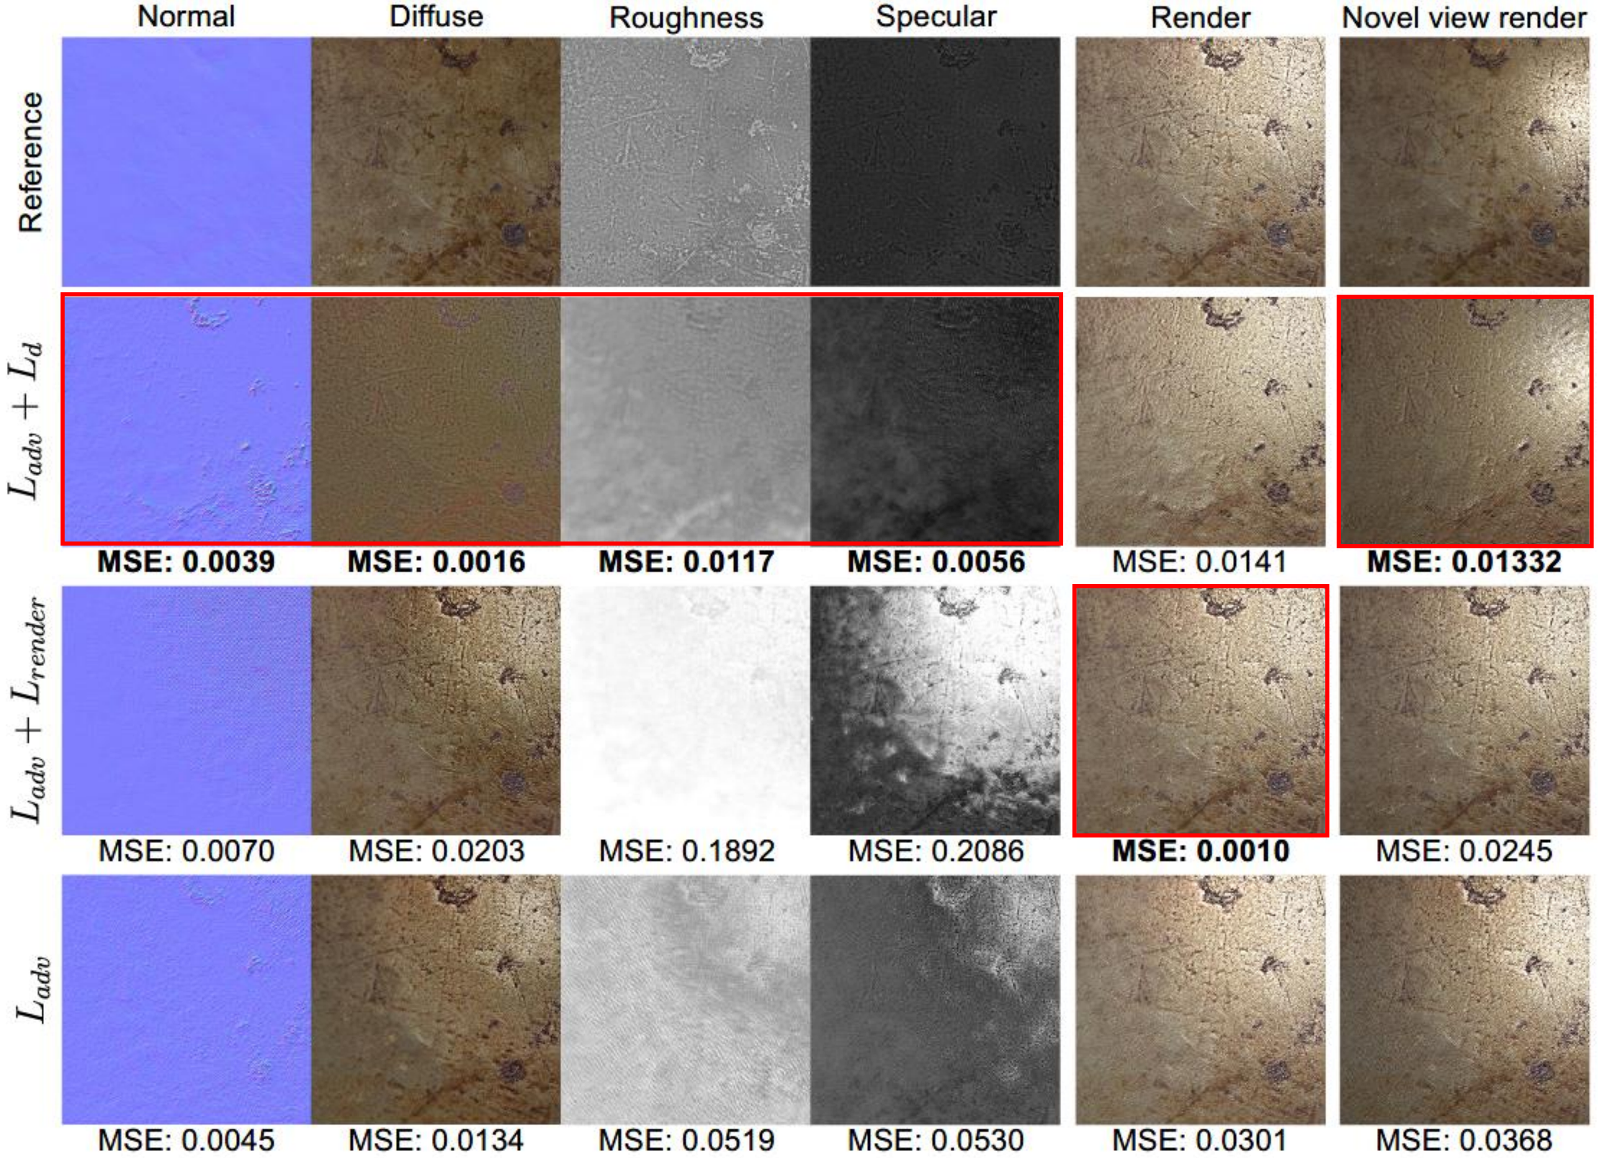

Network Analysis – generator

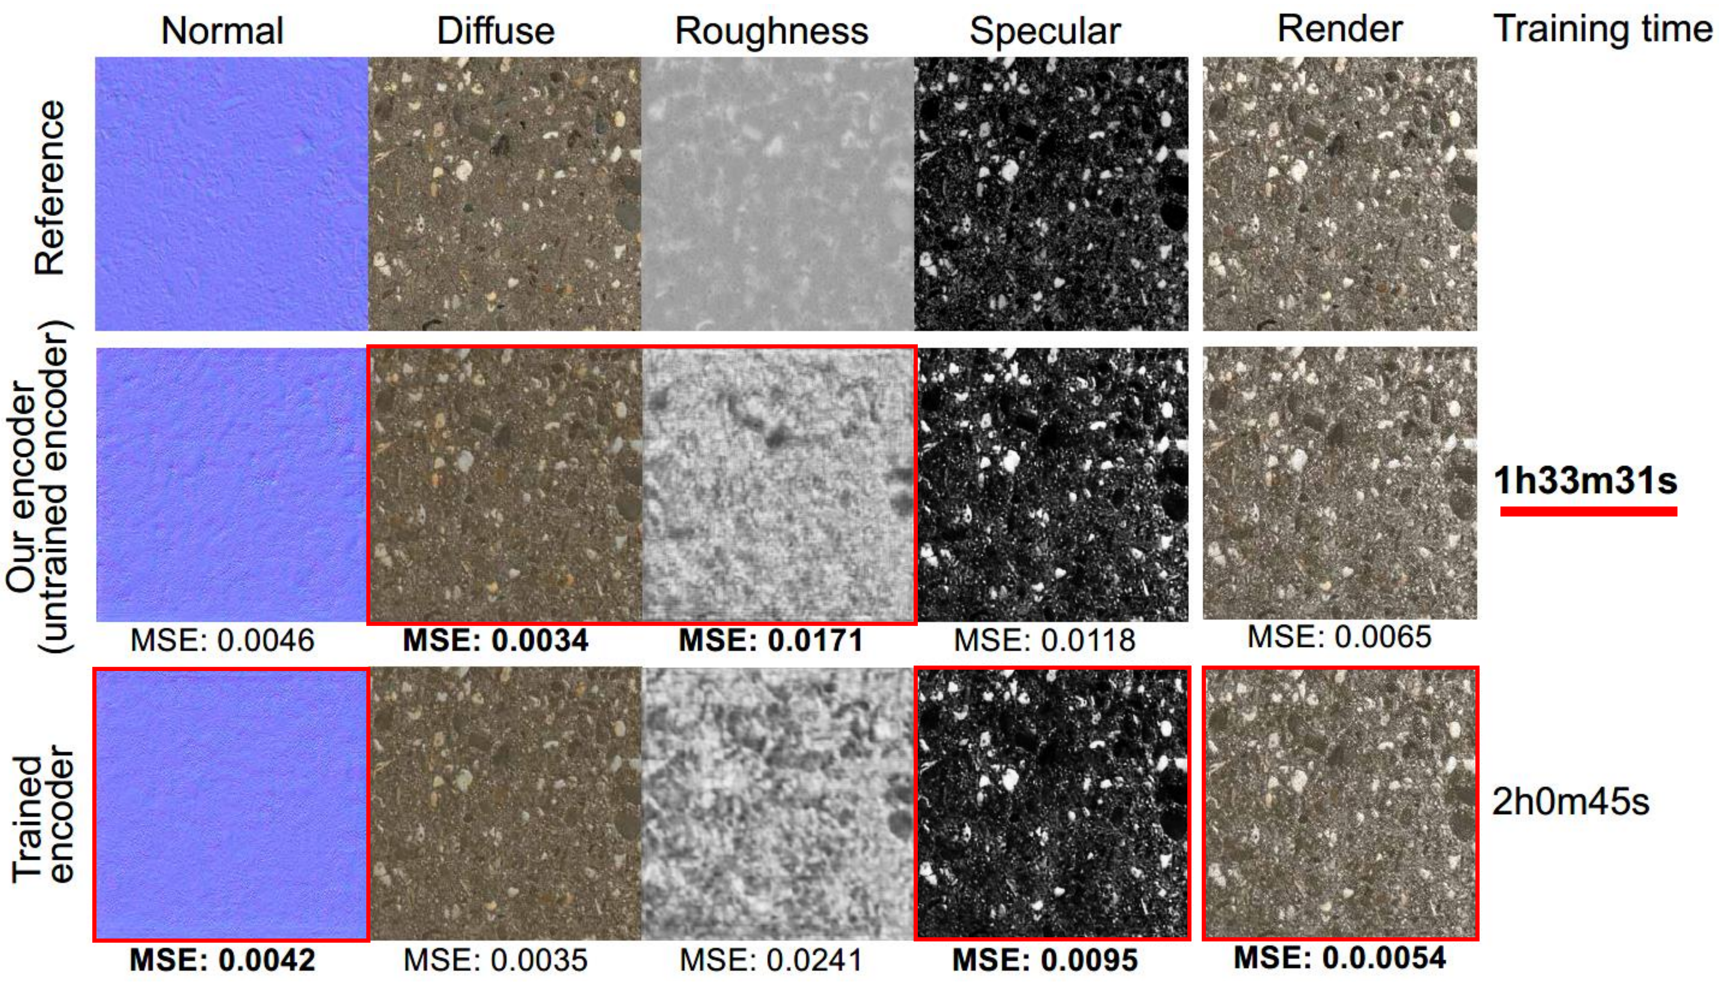

Limitation

- Our method failed to synthesis textures when the input image has a global structure.

Limitation

- Our method failed to synthesis textures when the input image has a global structure.

- Each input for our method requires individual training, which costs about 3 hours.

Conclusion

- An unsupervised GAN for joint SVBRDF recovery and synthesis without a large training dataset.
- A two-stream generator to enhance specular component.
- A novel joint loss function for high-quality novel view renderings.

Future work

- Introduce existing knowledge about the material.

Thank you for your attention

The code is available: https://github.com/mengshu1996/SVBRDF-GAN

Yezi Zhao, Beibei Wang, Yanning Xu, Zheng Zeng, Lu Wang and Nicolas Holzschuch