MaterialGAN: Reflectance Capture using a Generative SVBRDF Model

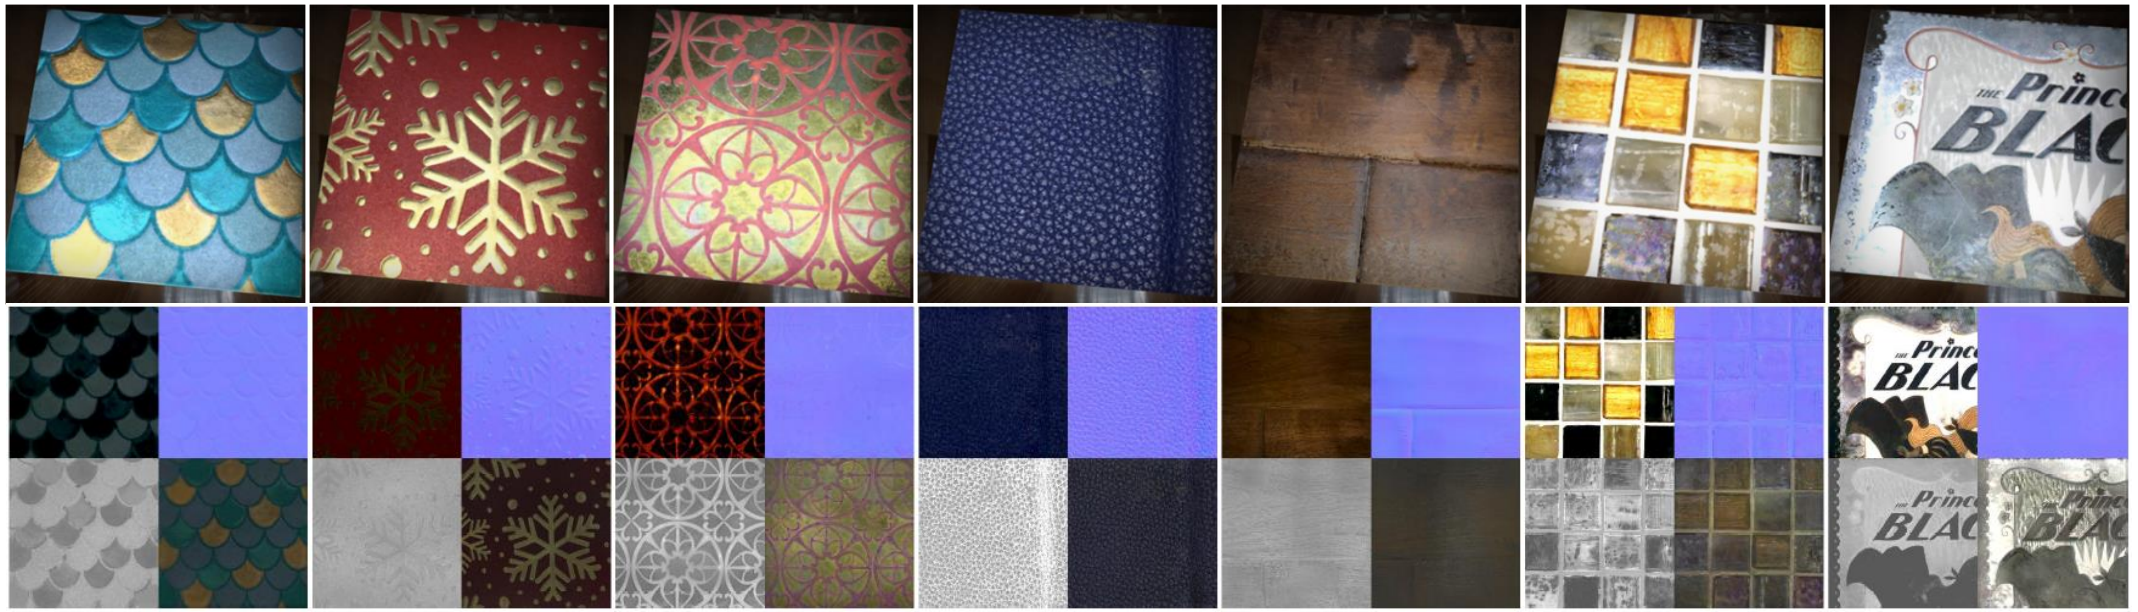

Yu Guo[1], Cameron Smith[2], Miloš Hašan[2],
Kalyan Sunkavalli[2] and Shuang Zhao[1]

[1]University of California, Irvine
[2]Adobe Research

Material acquisition

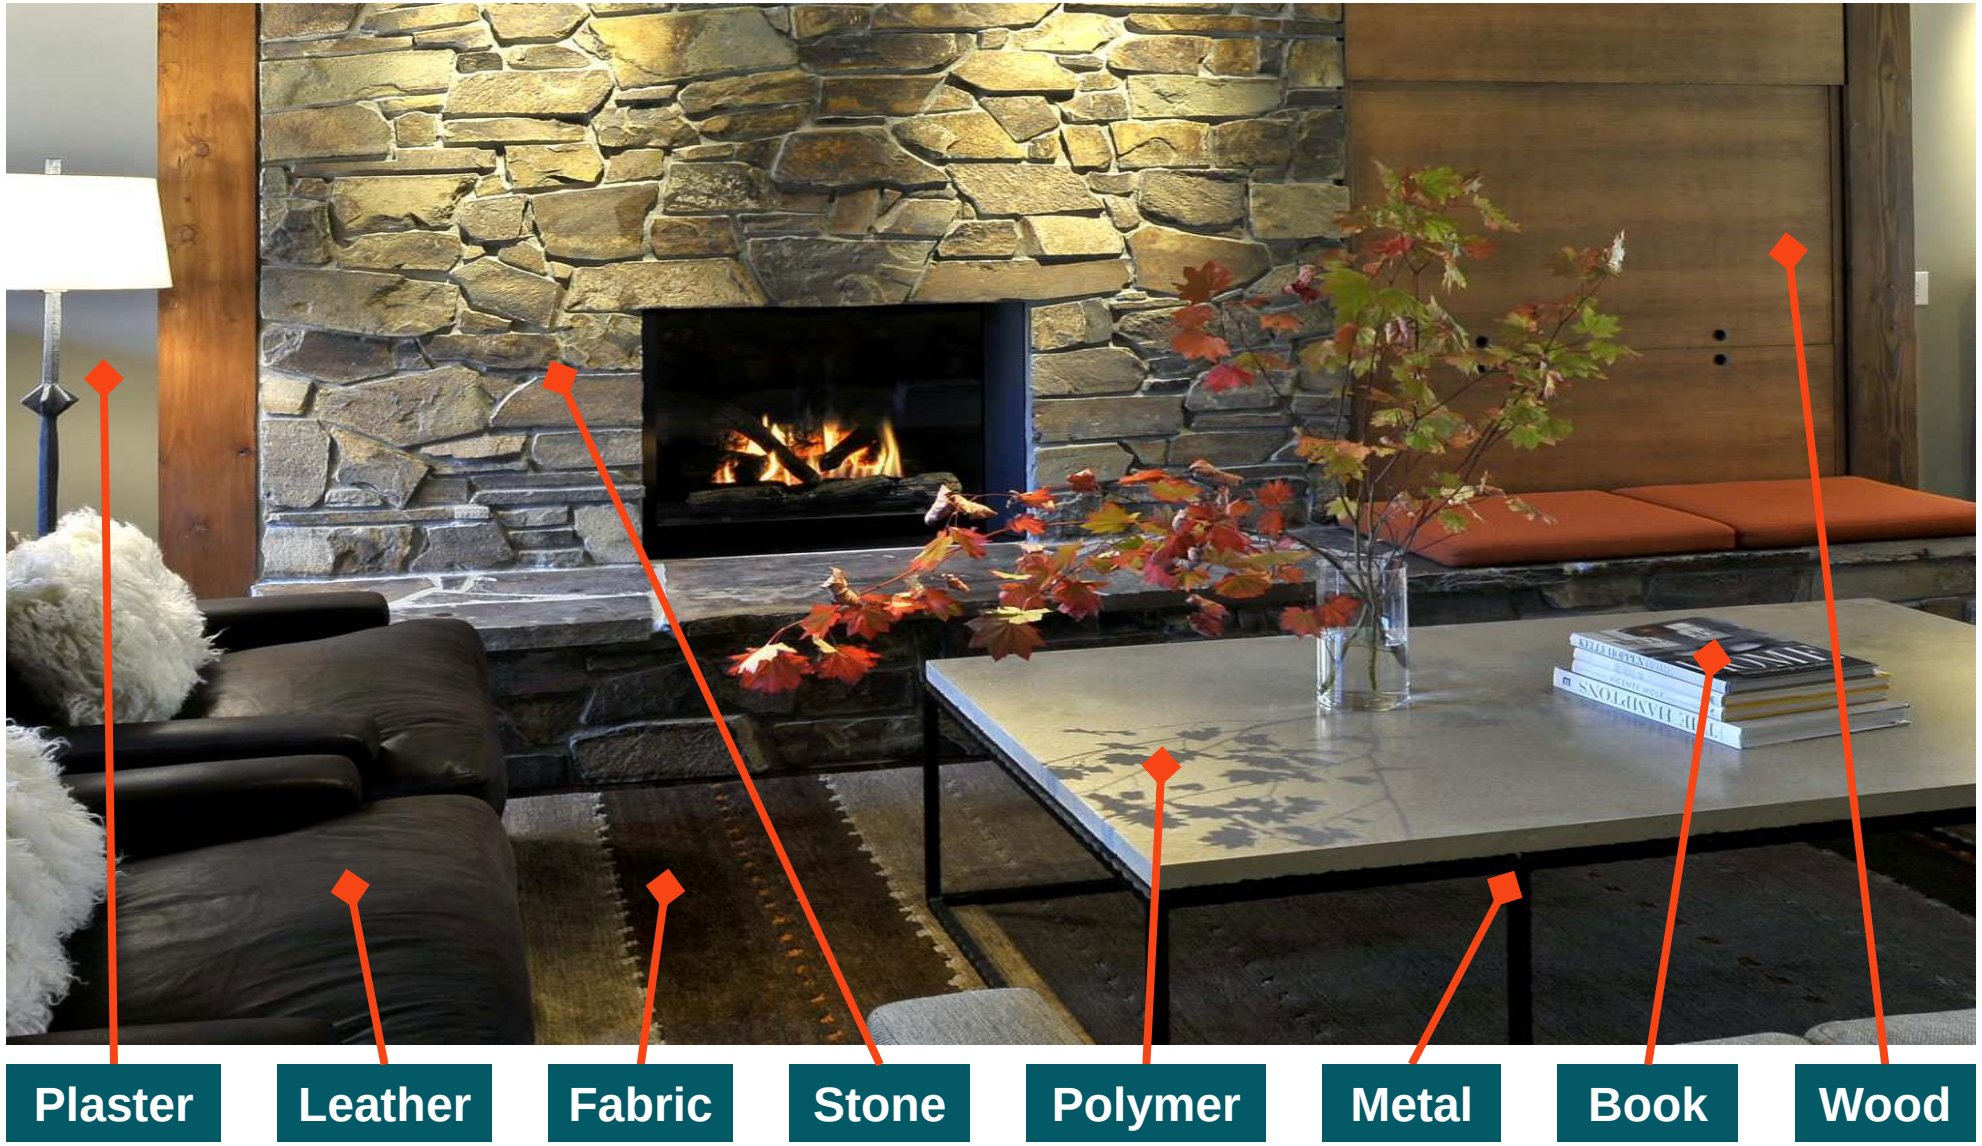

Related works

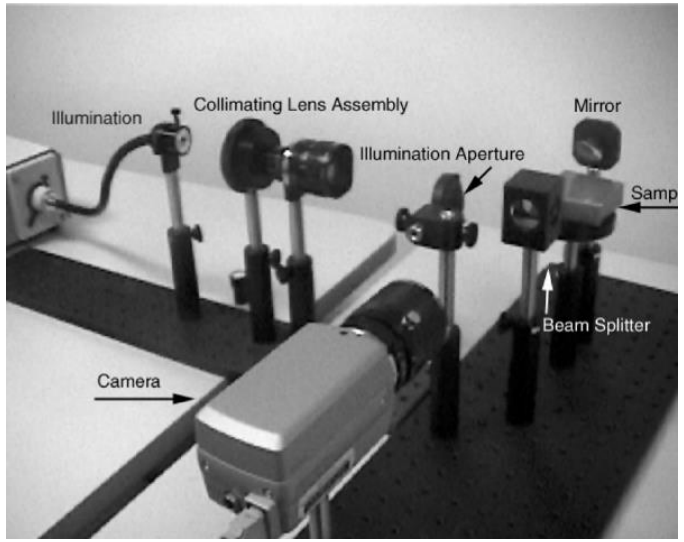

Dana and Wang 2004

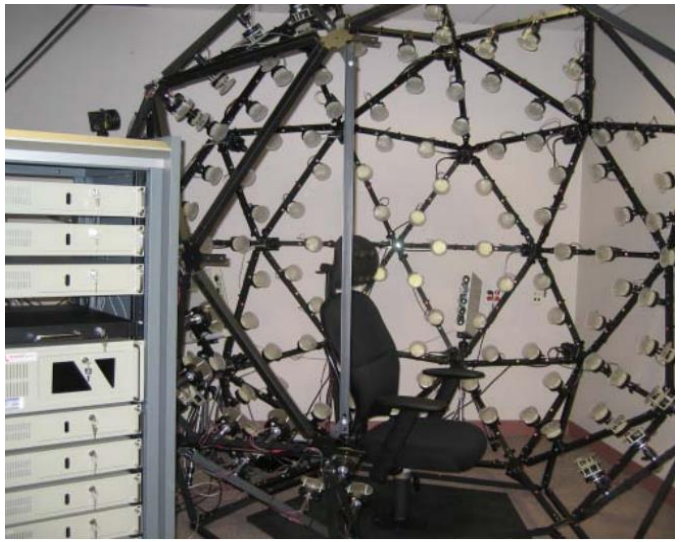

Weyrich et al. 2006

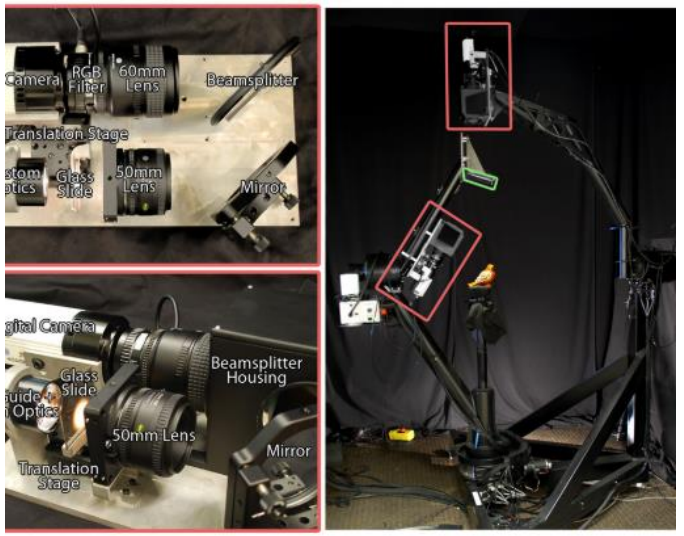

Spherical Gantry

High quality
Need dense captures with complex devices

Related works

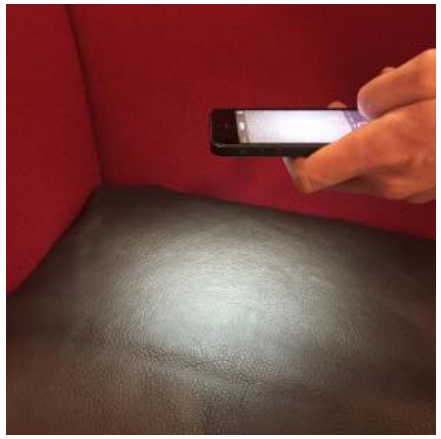

Aittala et al. 2015

Hui et al. 2017

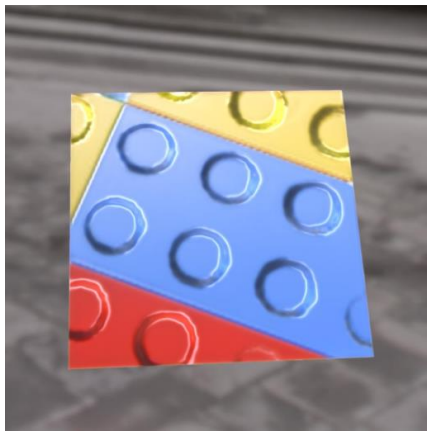

Li et al. 2017

Li et al. 2018

Deschaintre et al. 2018

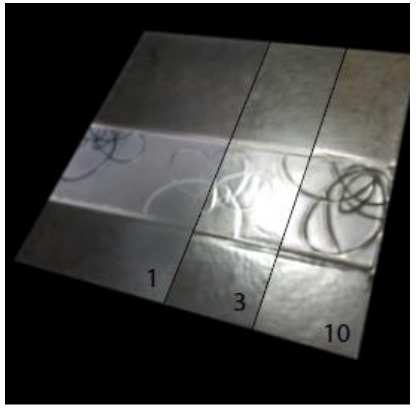

Deschaintre et al. 2019

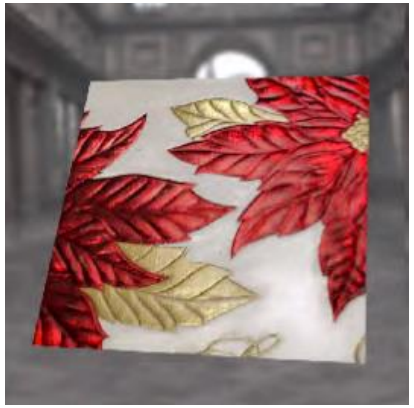

Gao et al. 2019

One or more images captured by handphone
Quality is not good as using specialized devices

SVBRDF

Inputs

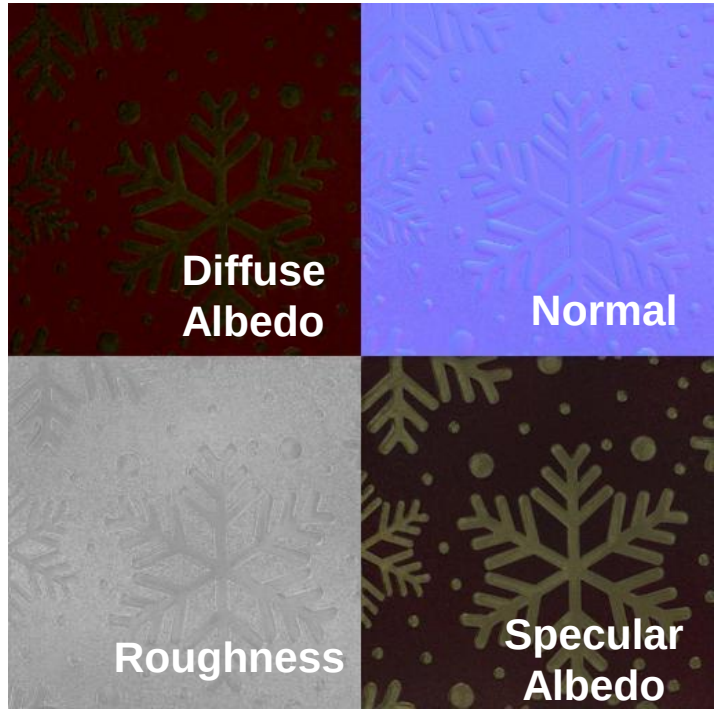

SVBRDF maps

Challenges

Small number of inputs,
lack of information

Inputs

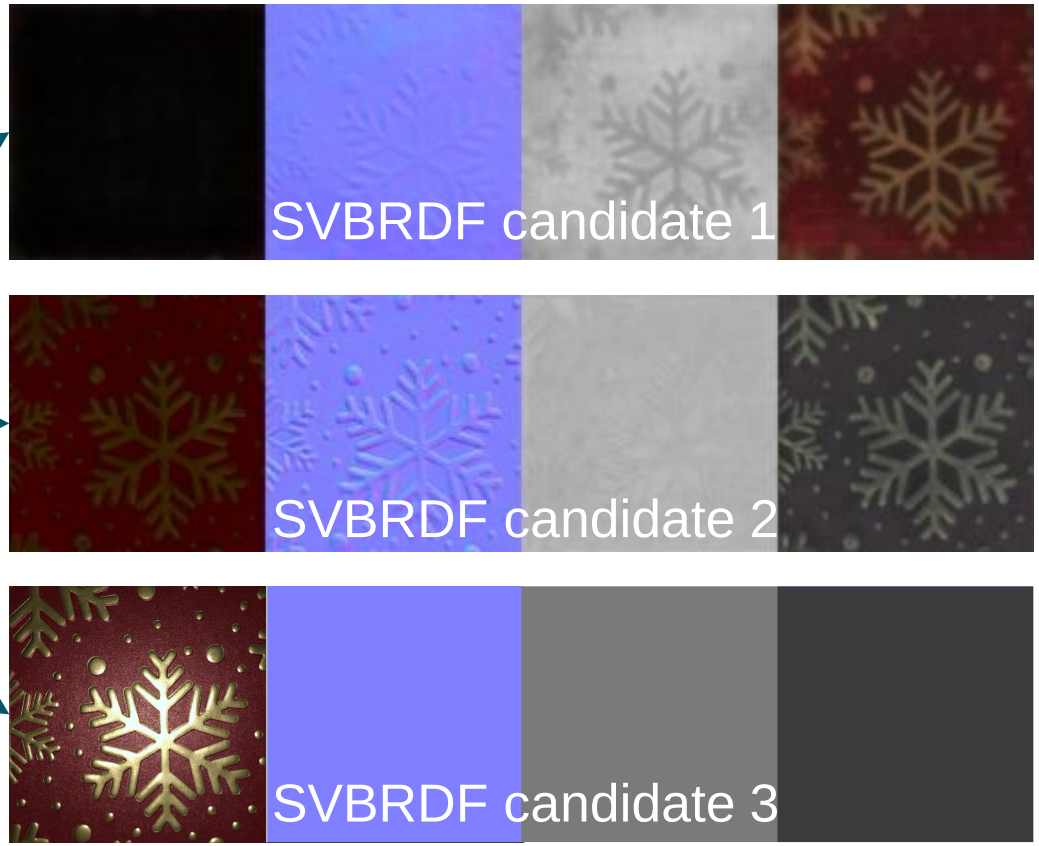

Under-constrained

Challenges

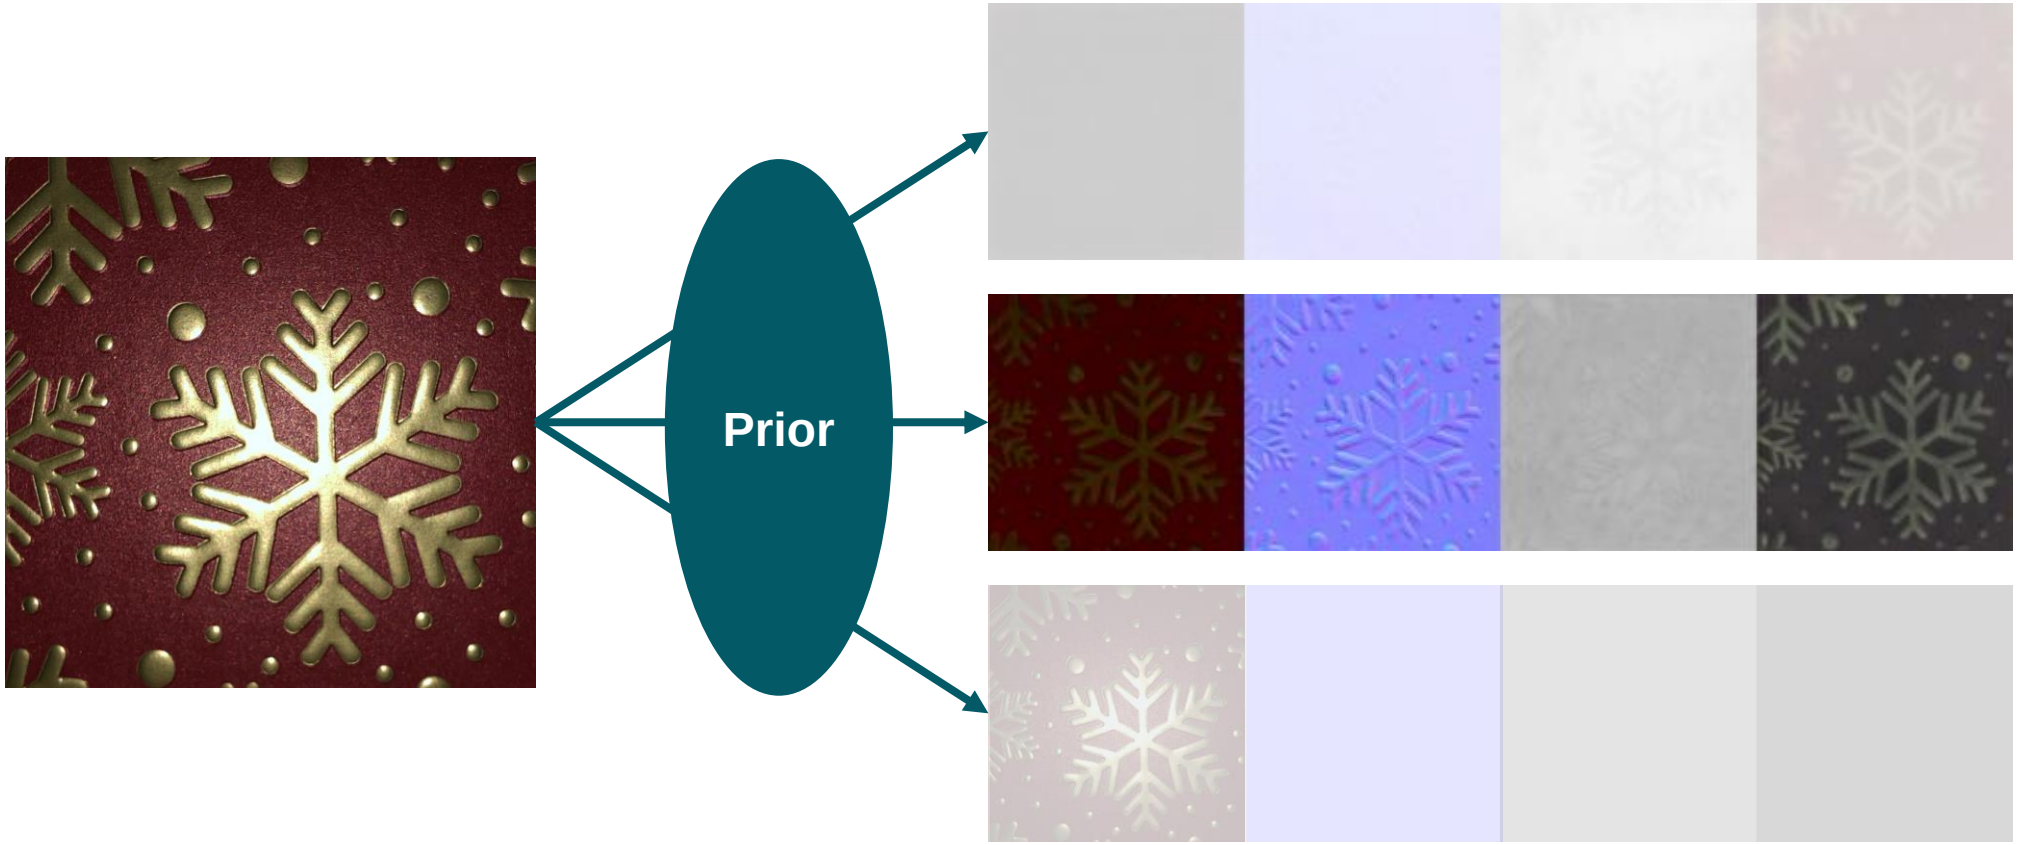

Prior :

- ❑ User interaction
- ❑ Linear low-dimensional BRDF models
- ❑ Stationary stochastic textures
- ❑ Learned with Neural Network
- ❑ …

Challenges

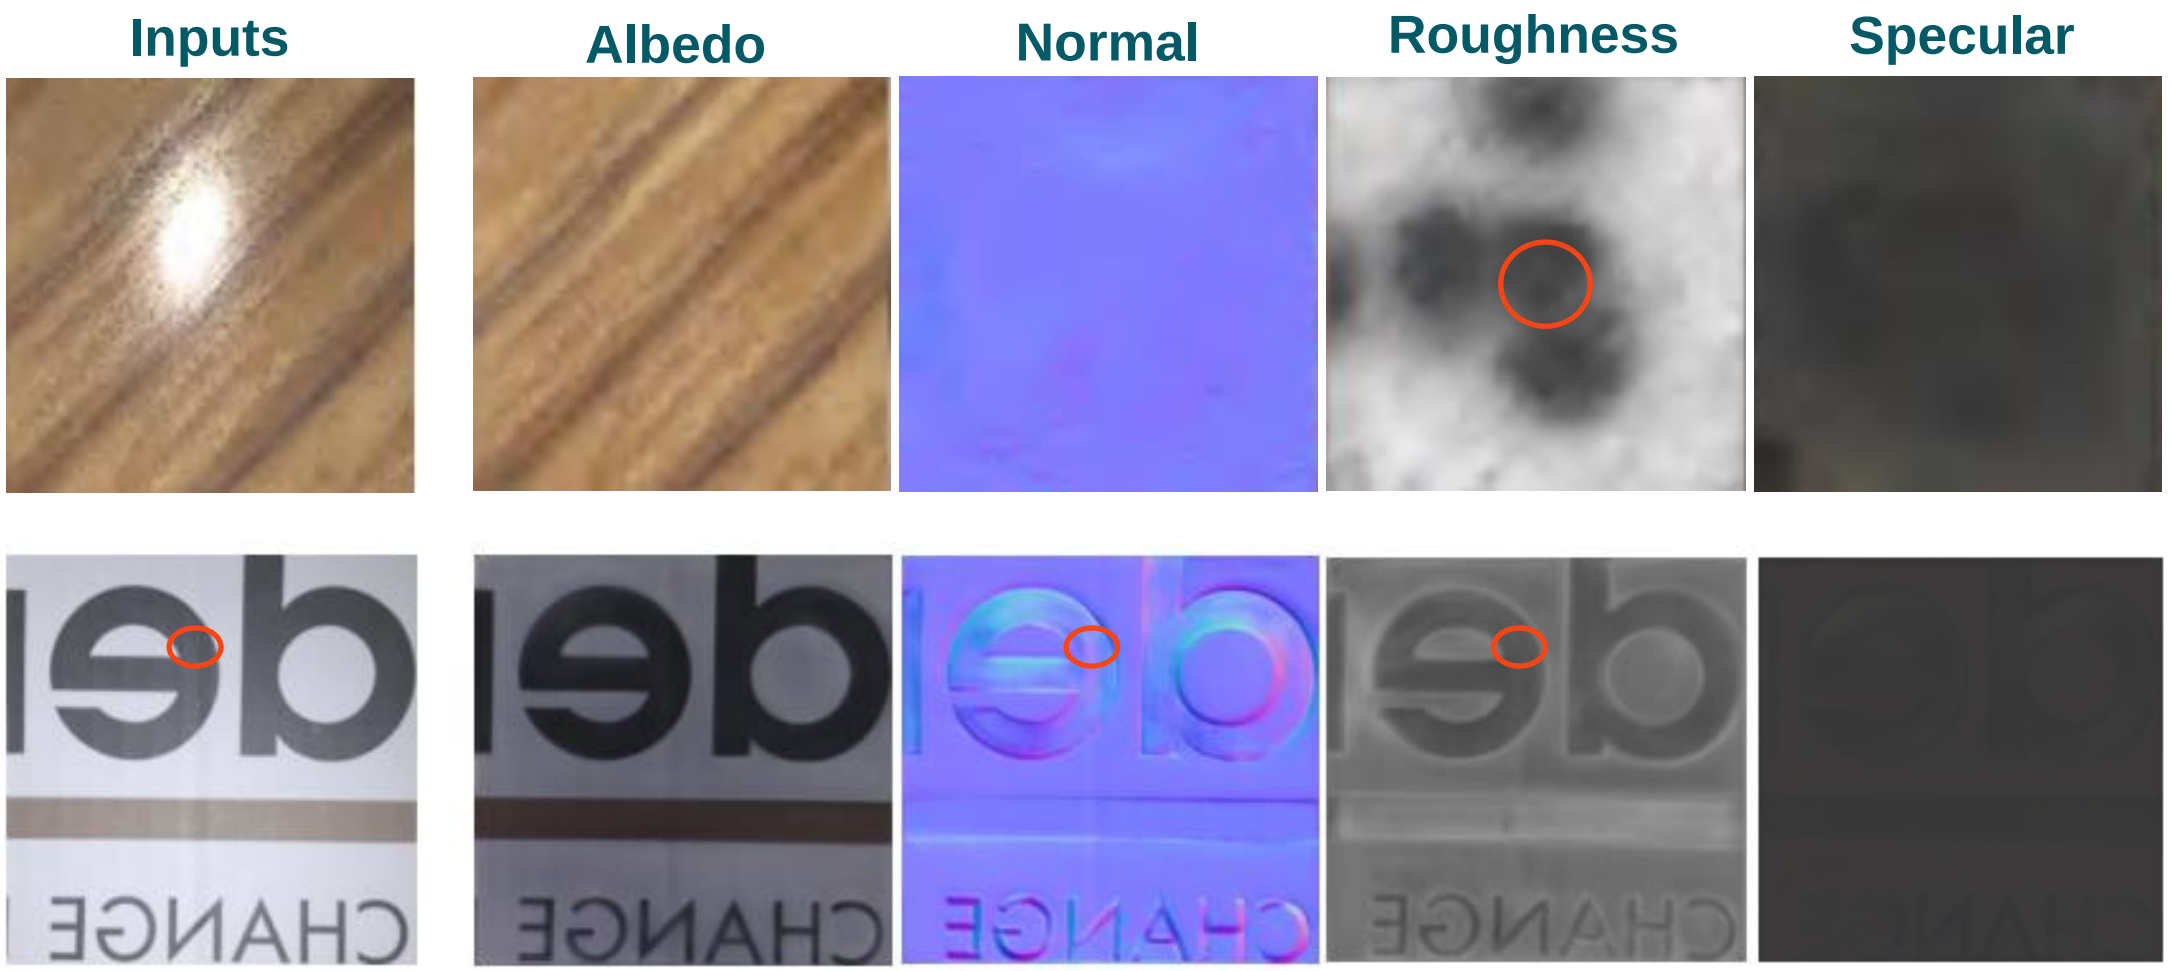

Unnatural SVBRDF maps

Our technique

Need a good prior to make SVBRDF maps look more natural

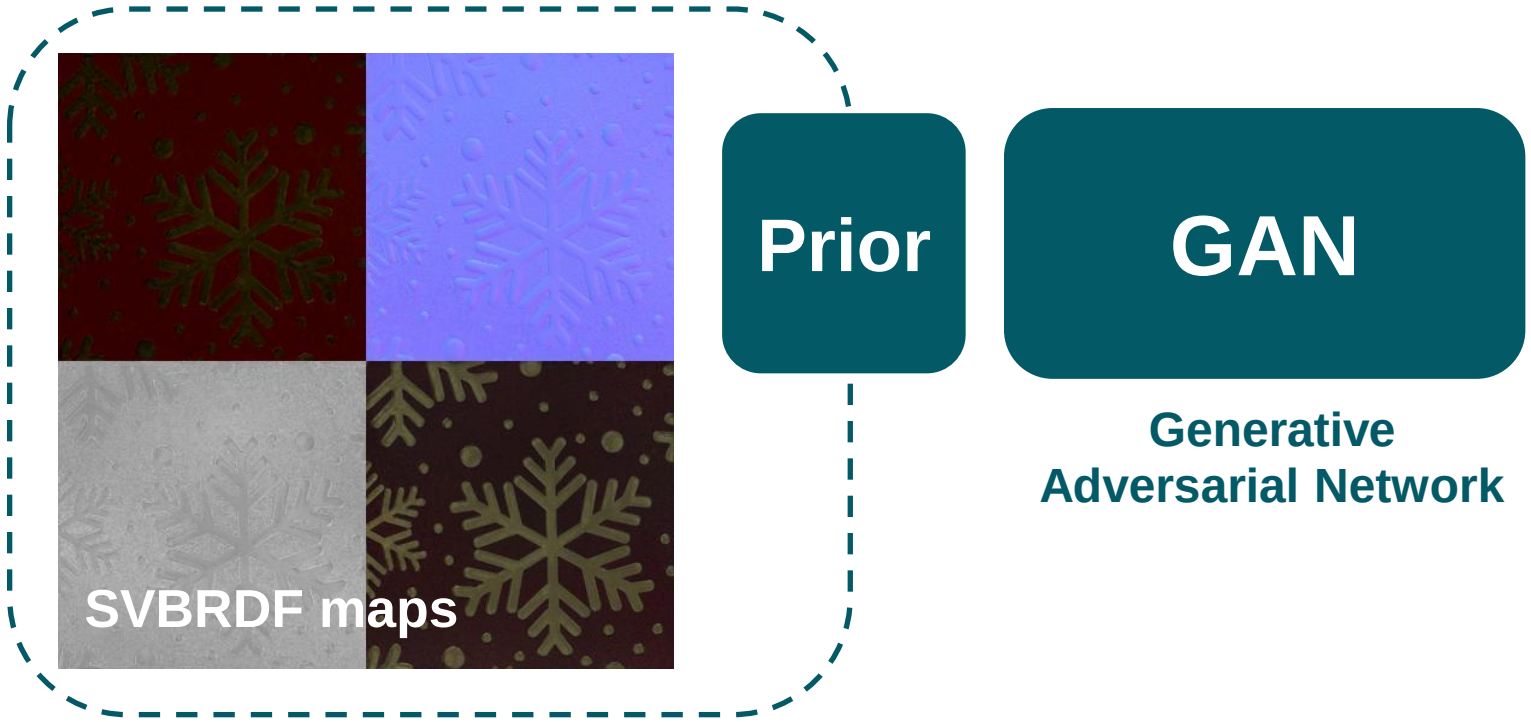

GAN review

Original

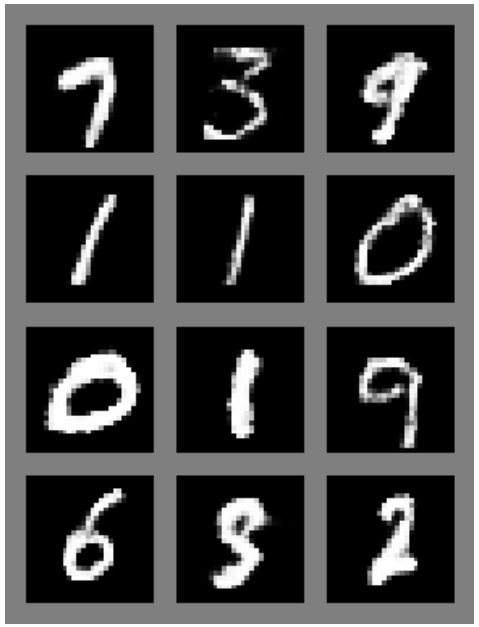

GAN
Goodfellow et al. 2014

Image

StyleGAN2
Karras et al. 2019

Video

MoCoGAN
Tulyakov et al. 2018

3D shape

MC-GAN
Li et al. 2019

Our method

Please check our main paper and supplemental materials for detailed analyses and comparisons

MaterialGAN

MaterialGAN = StyleGAN2 + SVBRDF

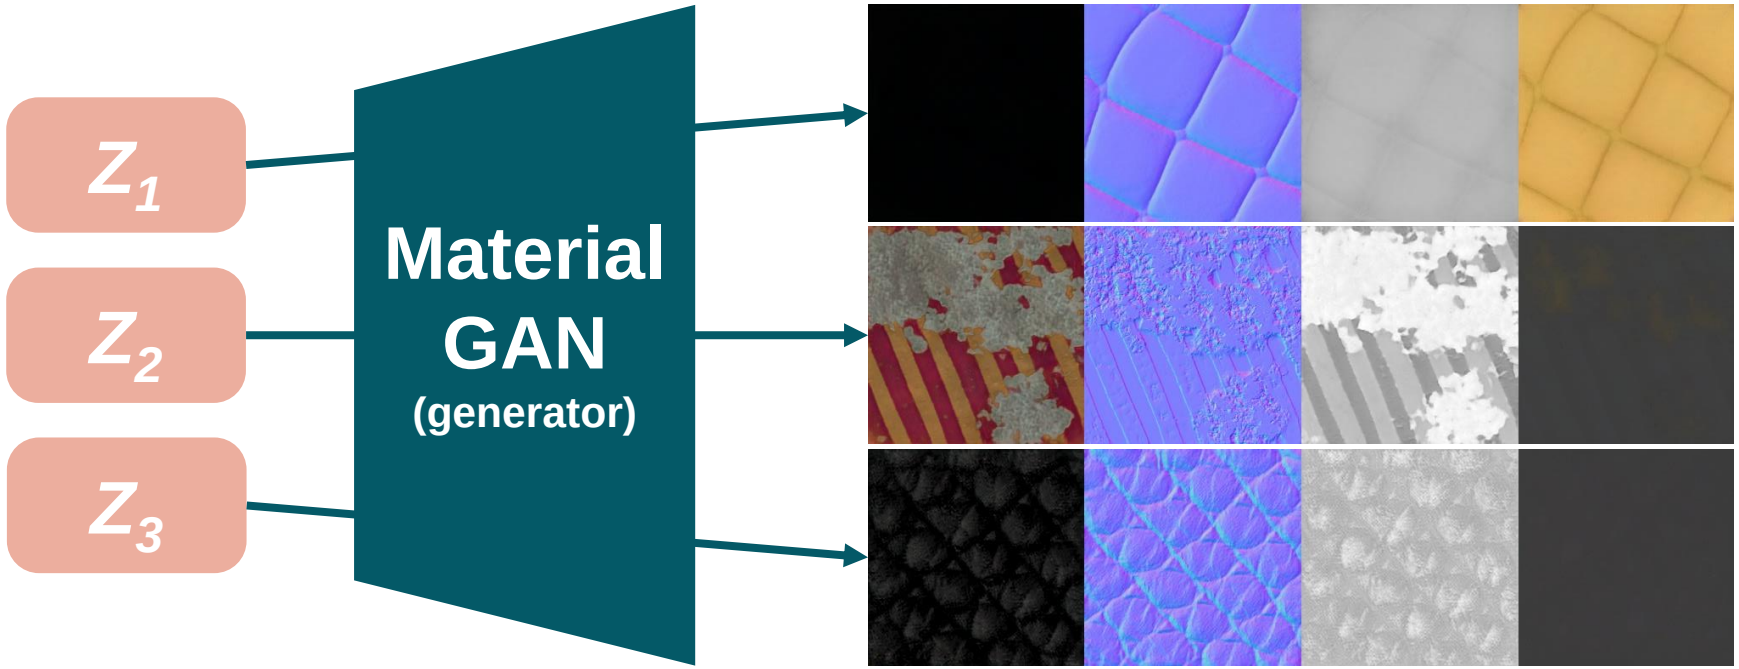

Z: latent vector of MaterialGAN

MaterialGAN

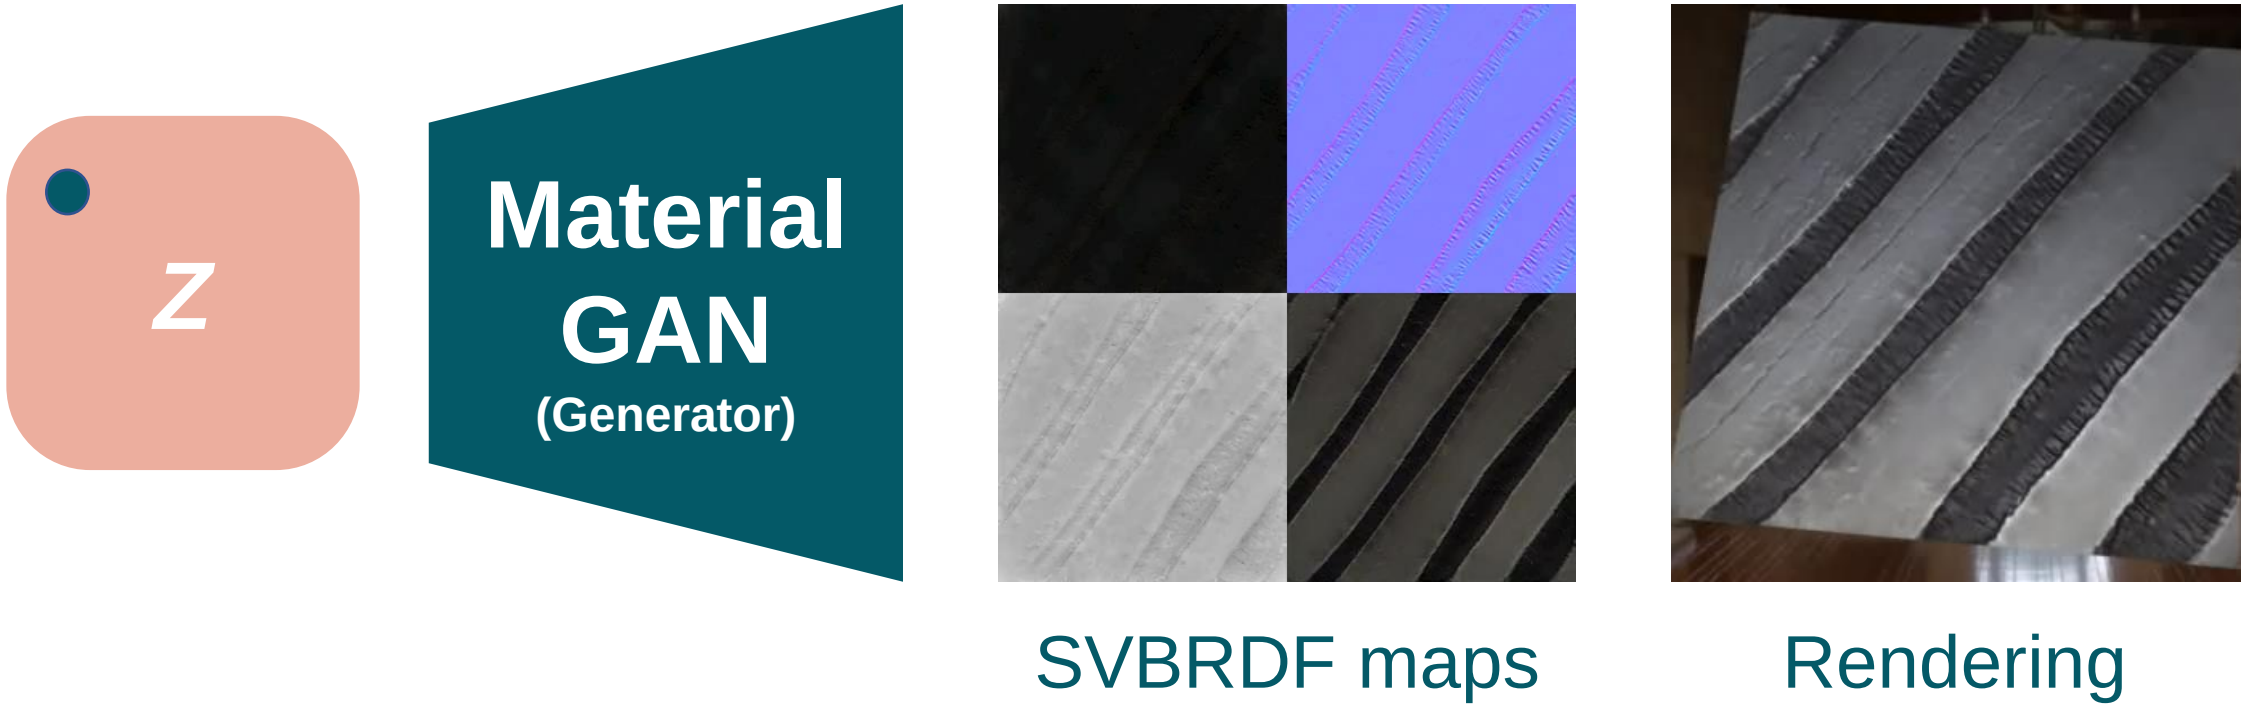

Training details of MaterialGAN

- Treat SVBRDF maps as 9-channel “image”
- 100,000 training data (including augmentation) from Deschaintre et al. 2018
- Resolution 256x256
- Tensorflow
- 8× Nvidia Tesla V100, 5 days

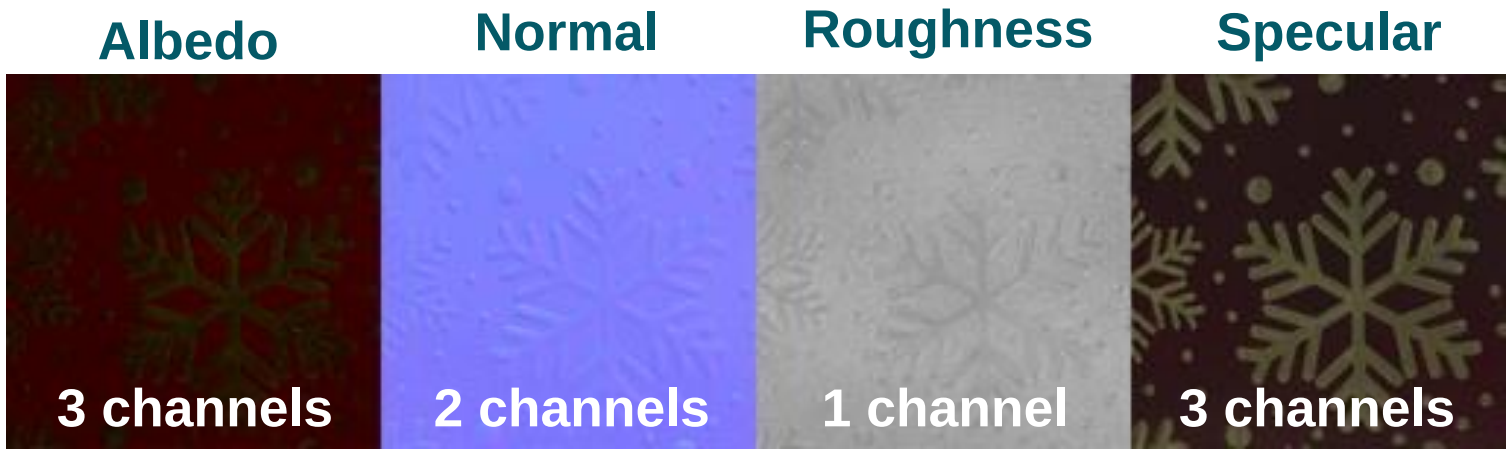

Optimization

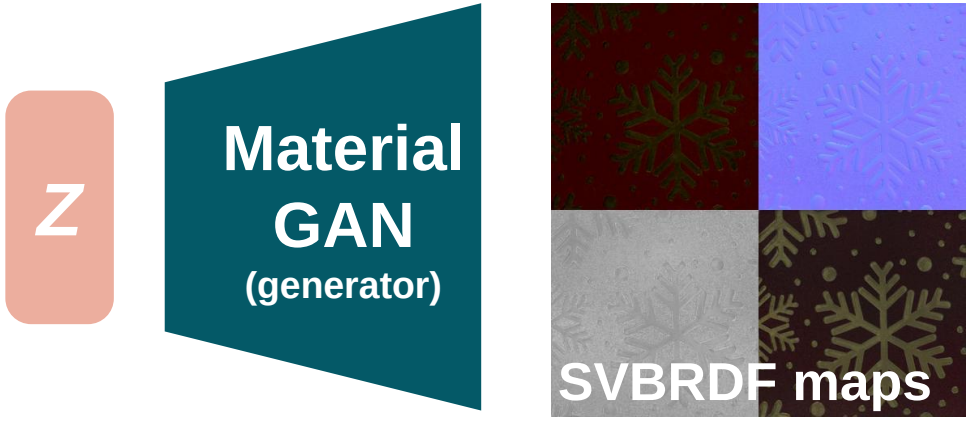

Optimization

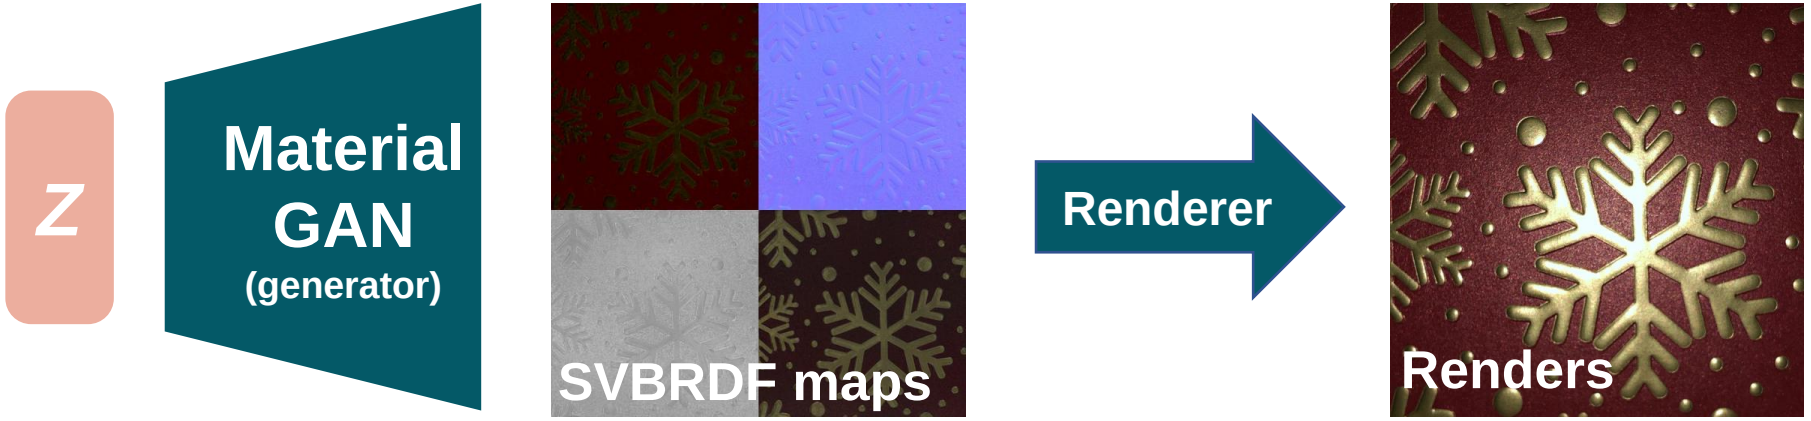

Optimization

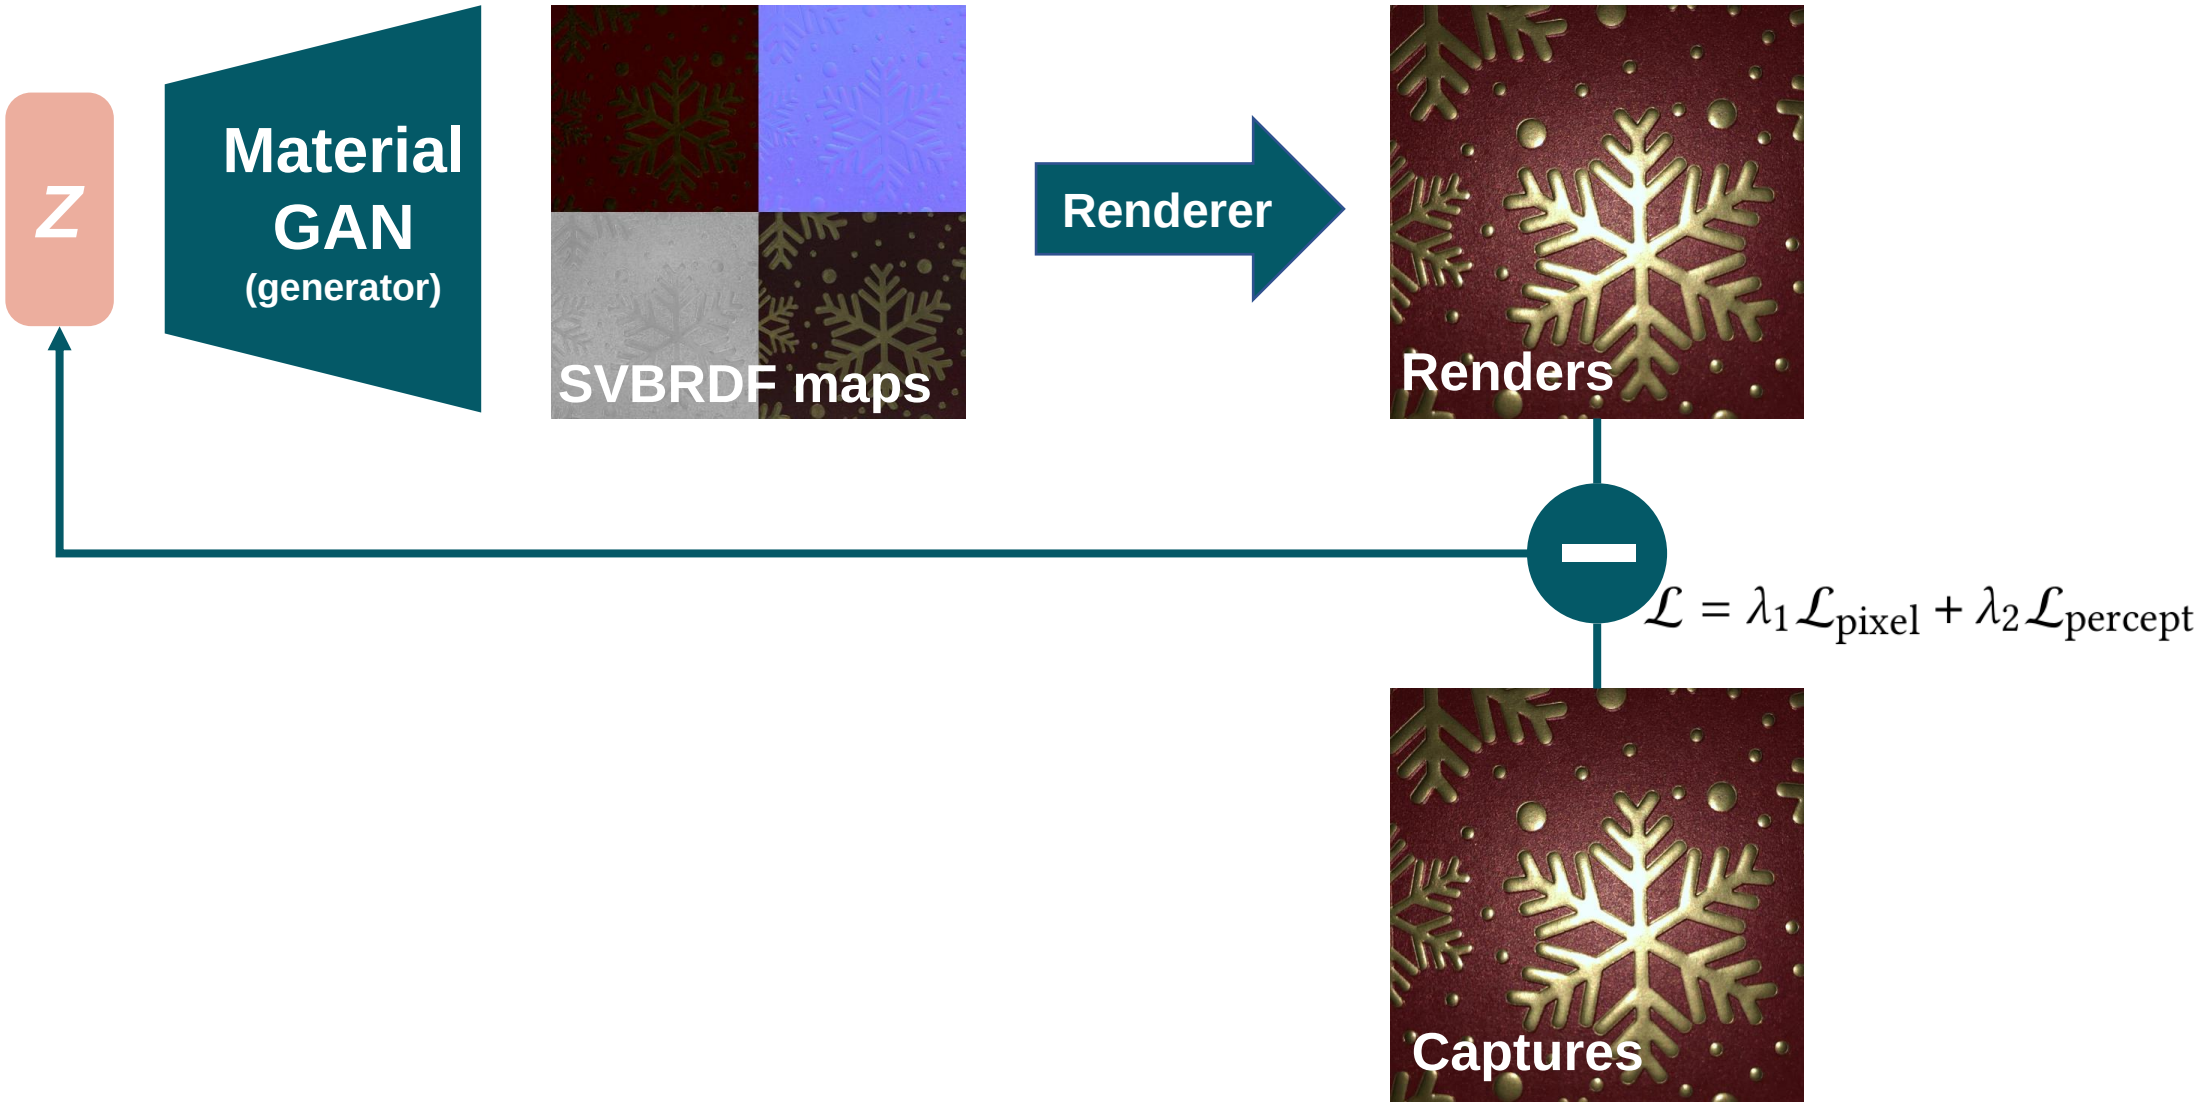

Optimization

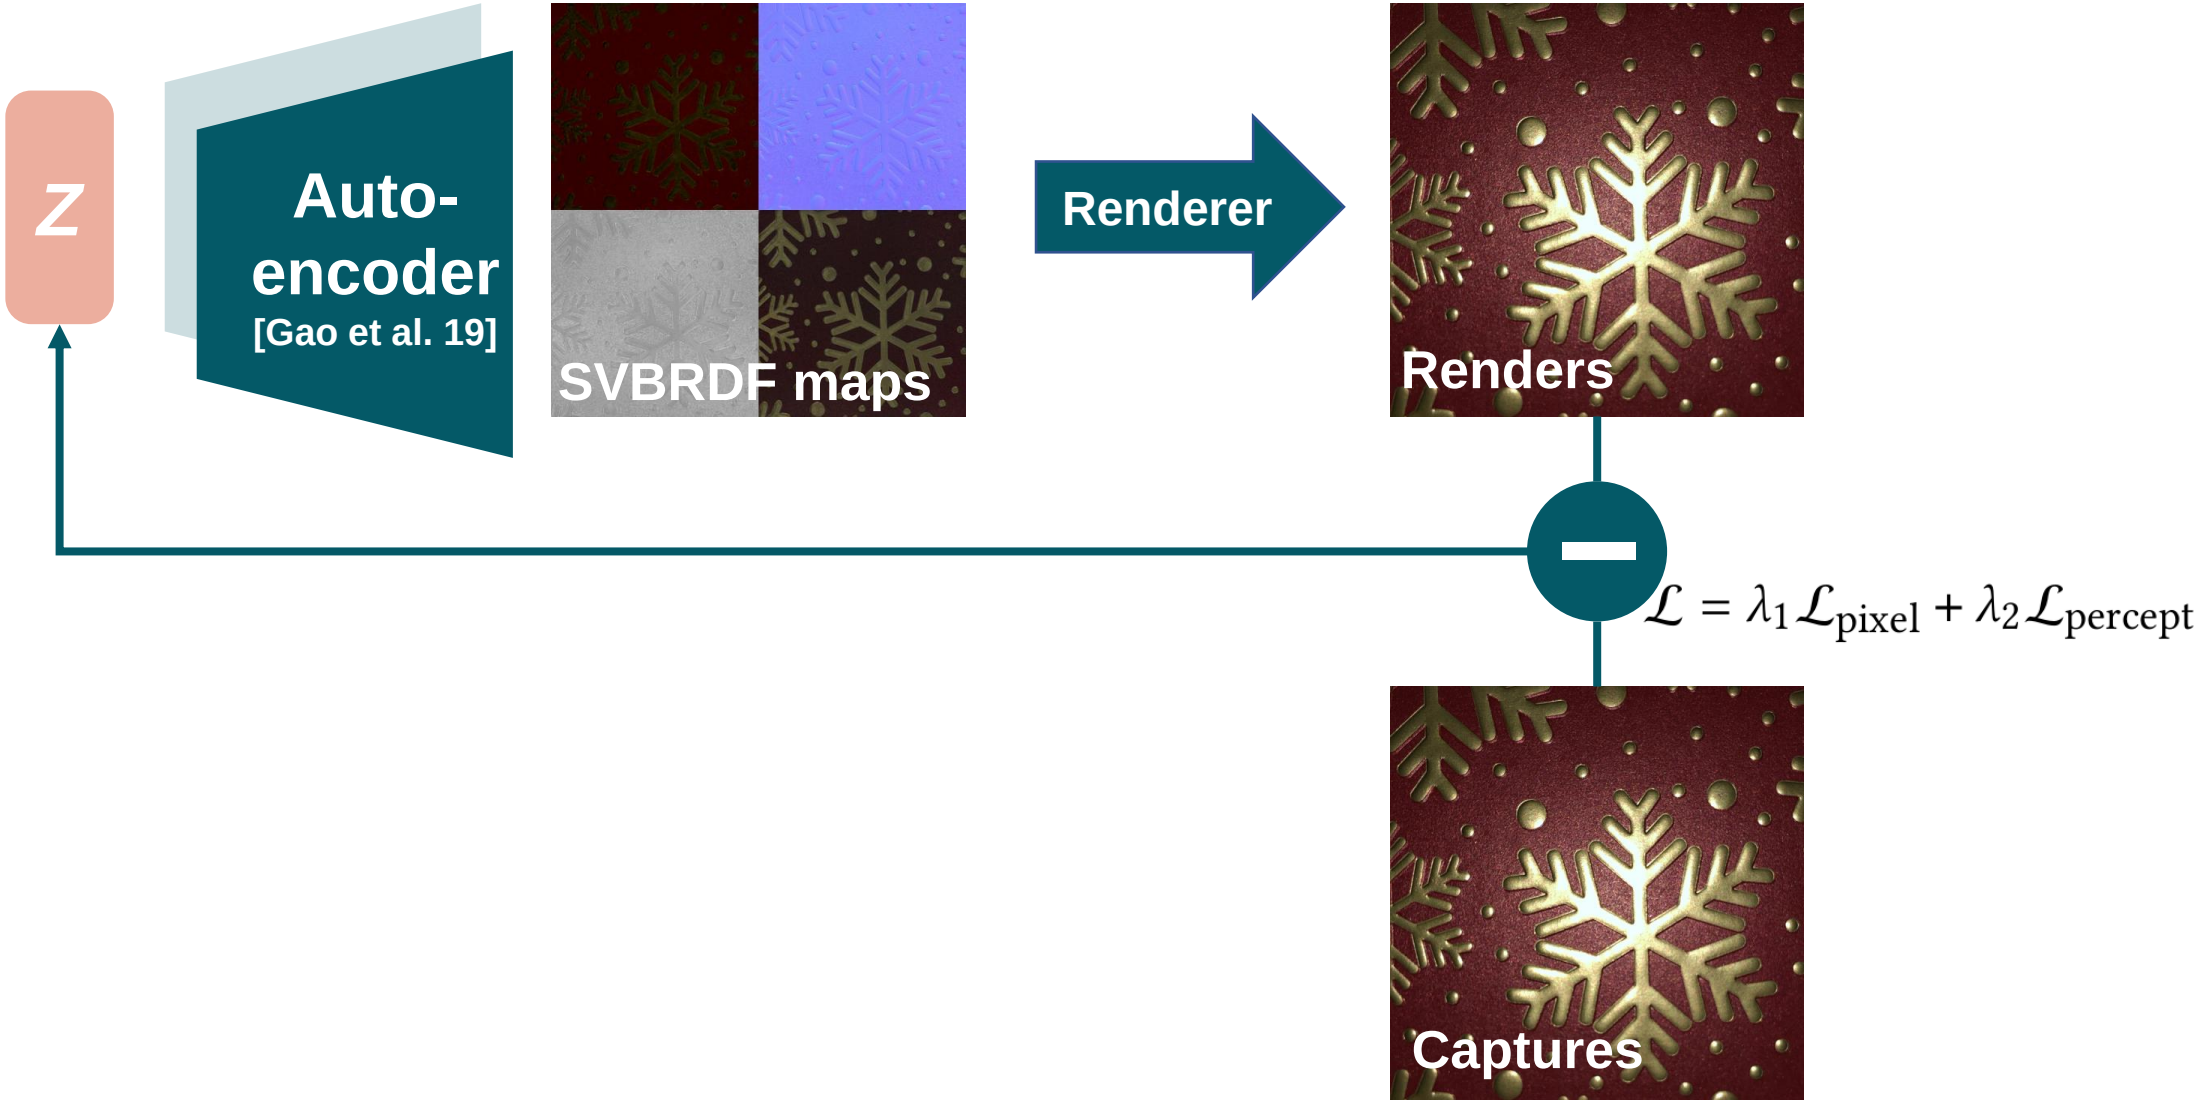

MaterialGAN latent space

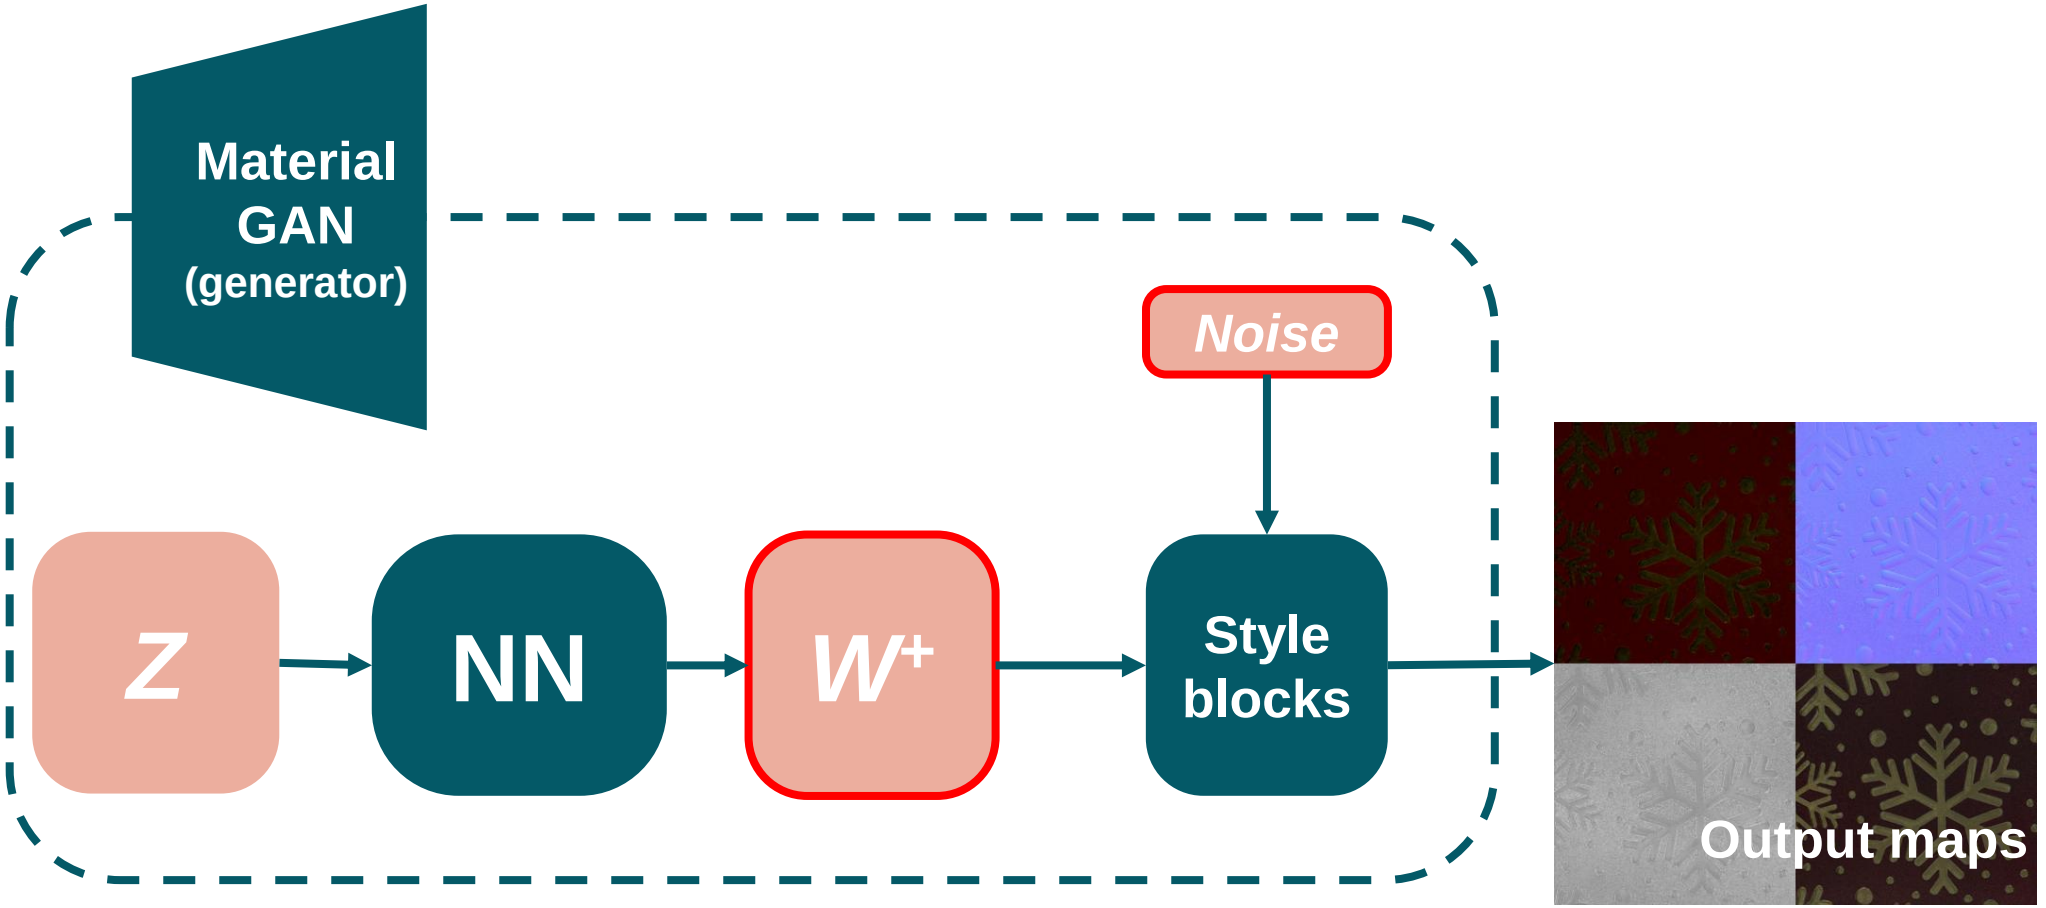

Final optimization pipeline

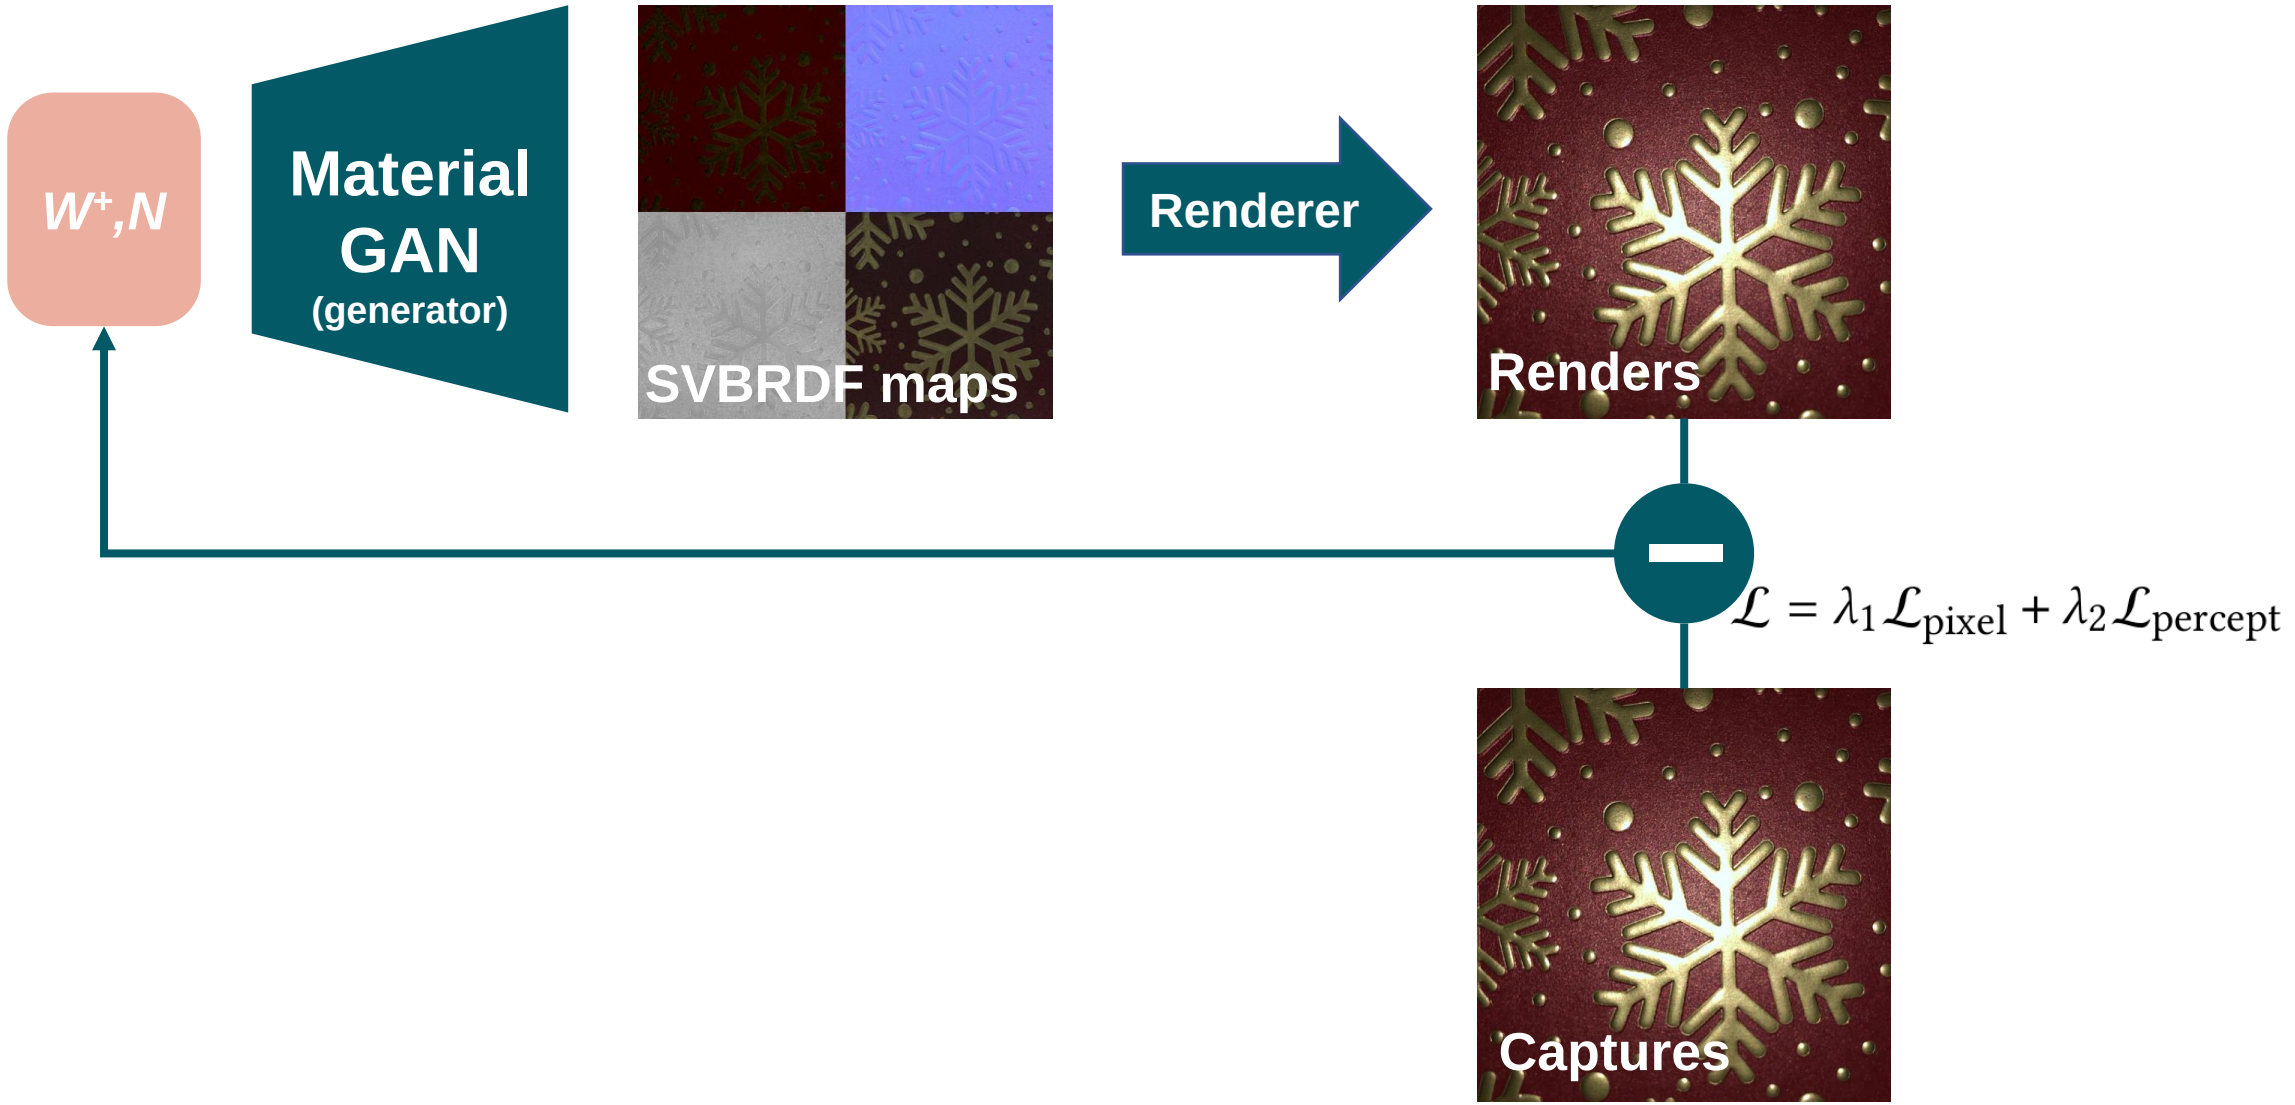

Results

Please check our main paper and supplemental materials for detailed analyses and comparisons

Implementation details

- Single point source light, collocated with the camera
- Support single input to 25 inputs (GPU with 16GB memories)
- 39 synthetic testing data from Deschaintre et al. 2018 and Adobe Stocks dataset
- 39 cellphone captures for testing
- Pytorch
- Titan RTX, 2000 iterations takes about 2 minutes

Optimization results

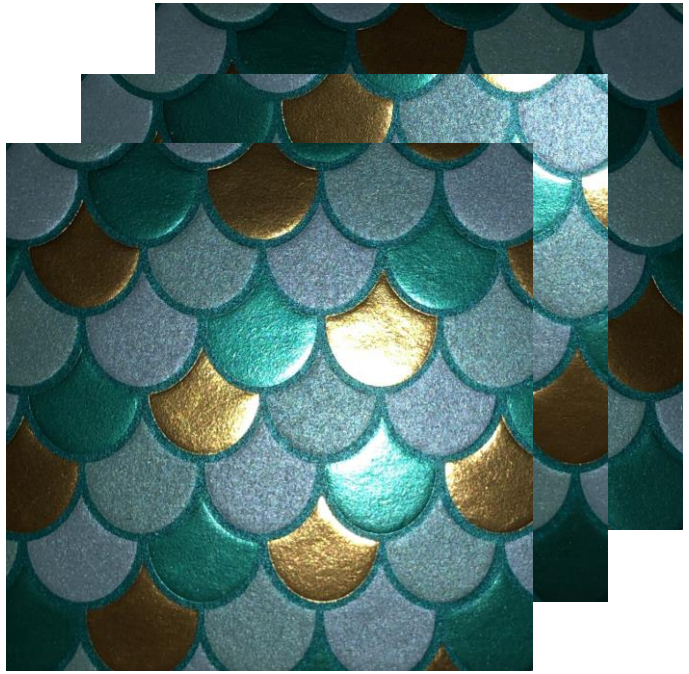

Inputs

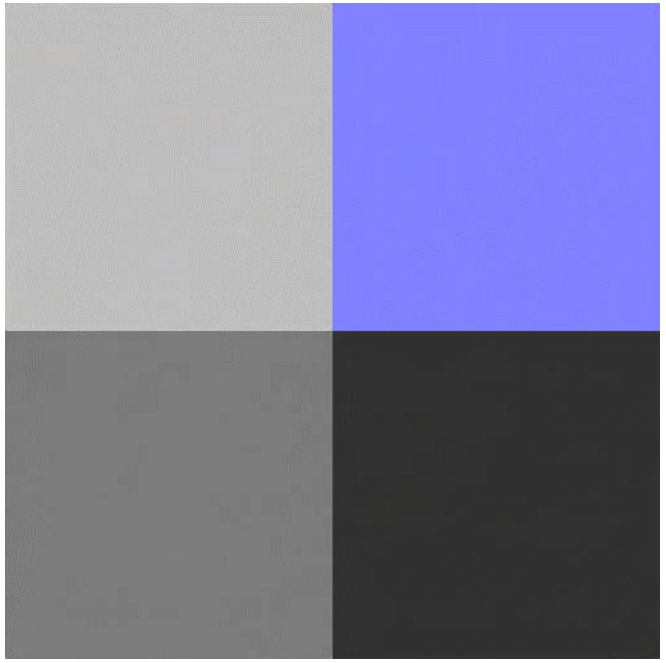

Optimized maps

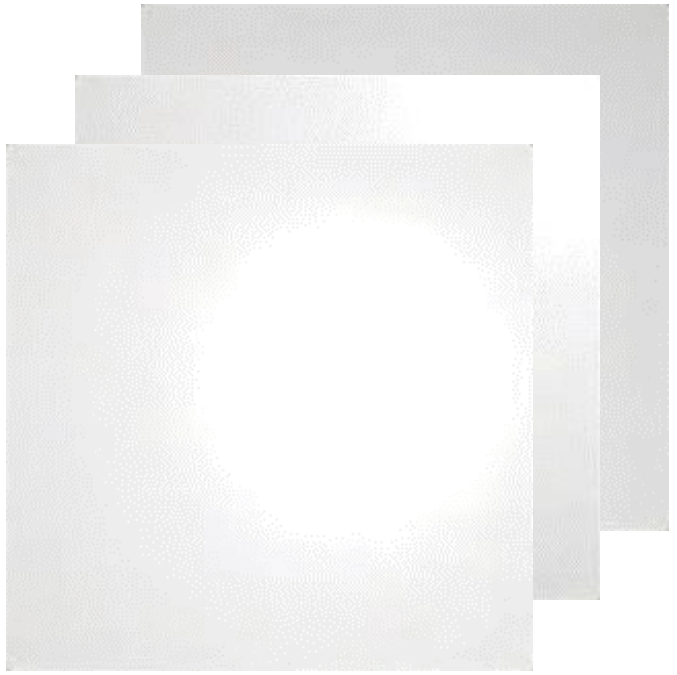

Renderings

Card

Reference (Inputs)

Reference (Novel view)

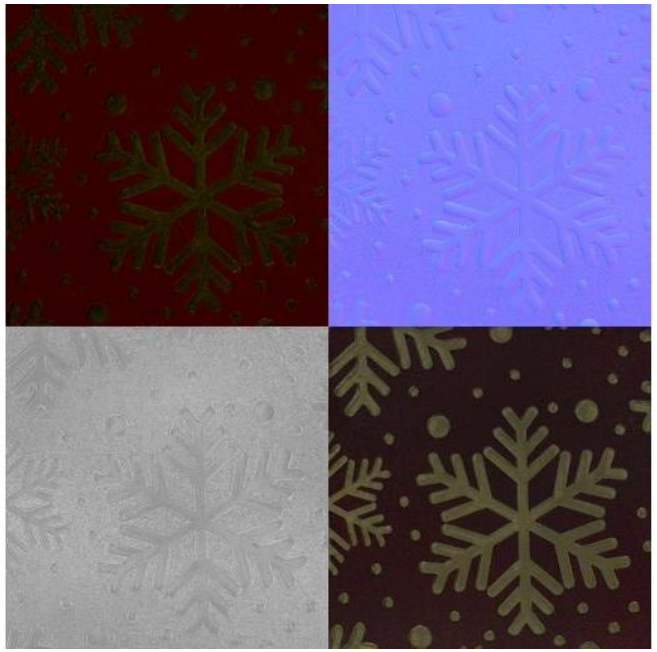

Estimated Maps

Ours (Optimized)

Ours (Novel view)

Input
(7 used, 4 shown)

Ours

SVBRDF
maps

Wood

Reference (Inputs)

Reference (Novel view)

Estimated Maps

Ours (Optimized)

Ours (Novel view)

Ours

Input
(7 used, 4 shown)

SVBRDF
maps

Leather

Reference (Inputs)

Reference (Novel view)

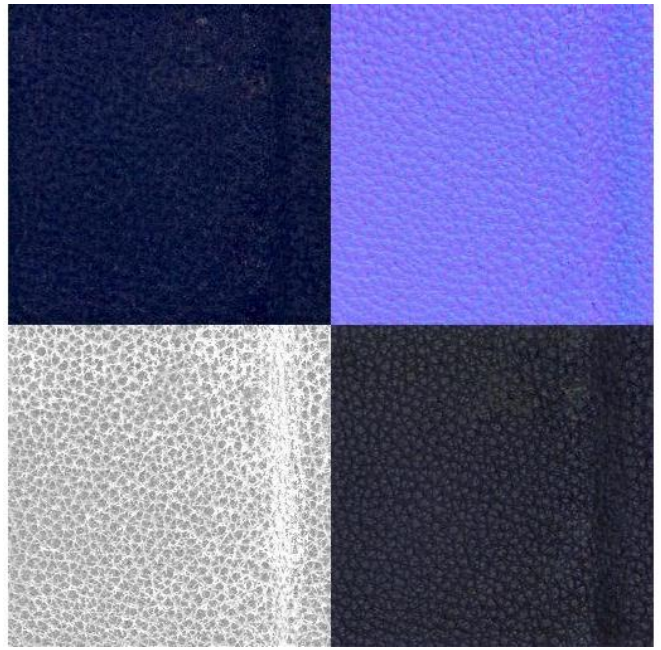

Estimated Maps

Ours (Optimized)

Ours (Novel view)

Ours

Input
(7 used, 4 shown)

SVBRDF
maps

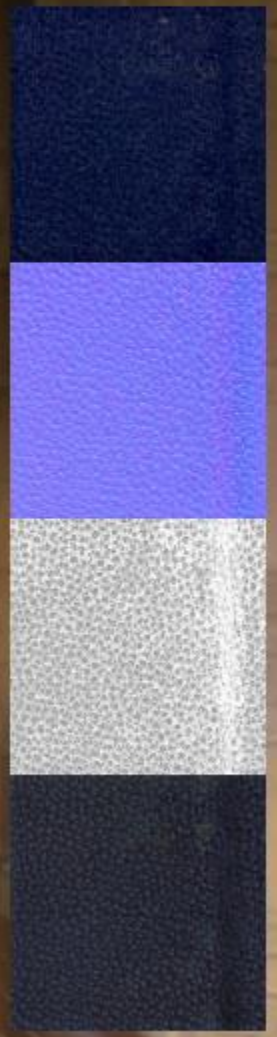

Tiles

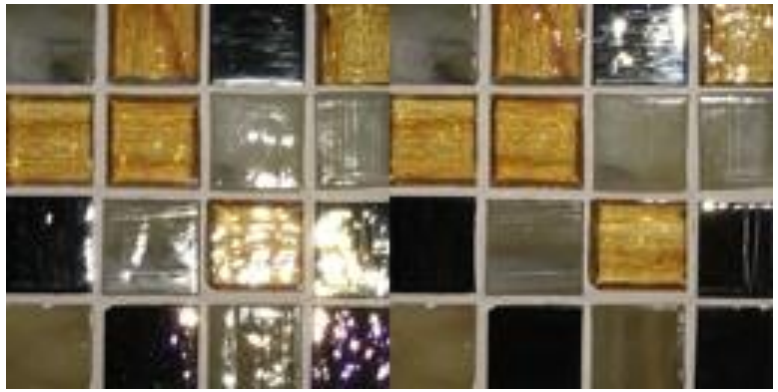

Reference (Inputs)

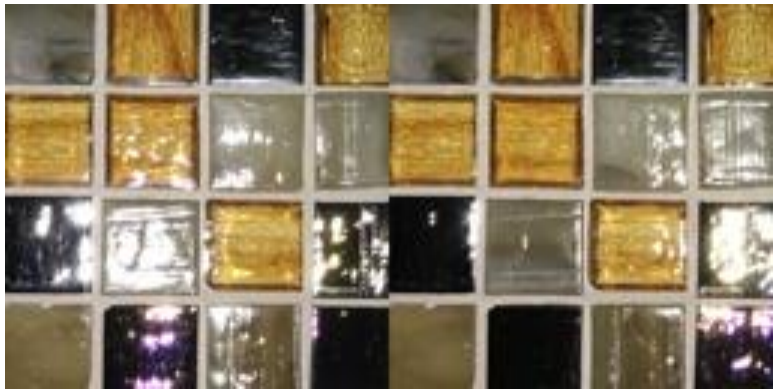

Reference (Novel view)

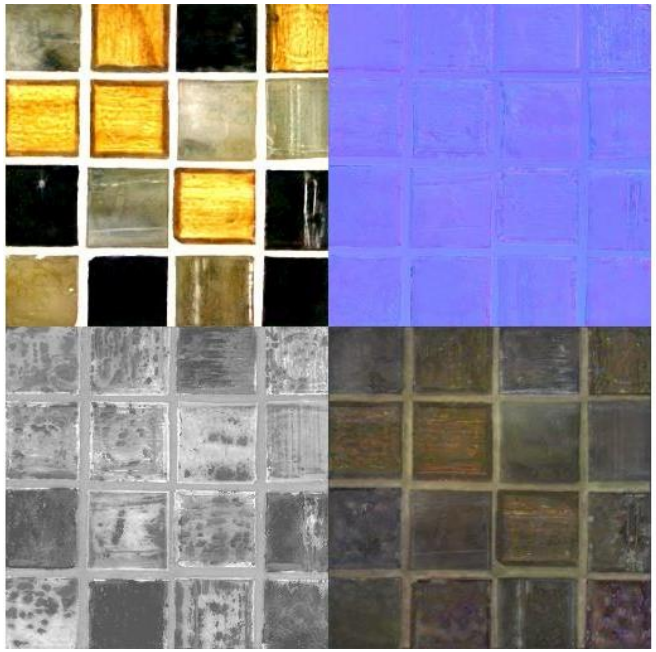

Estimated Maps

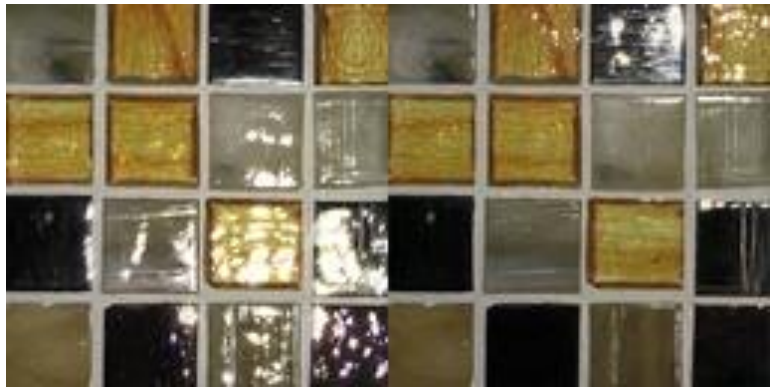

Ours (Optimized)

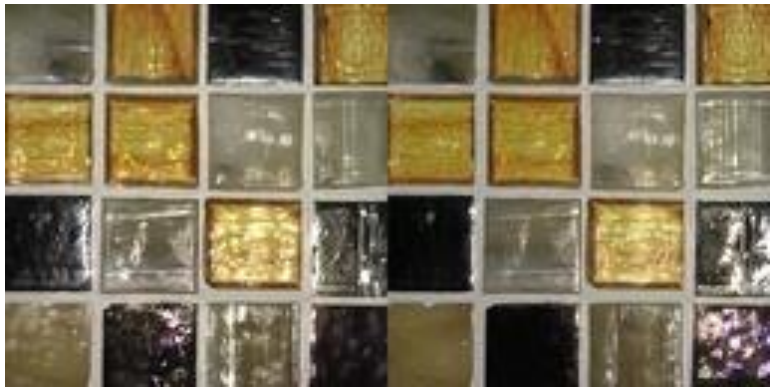

Ours (Novel view)

Input
(7 used, 4 shown)

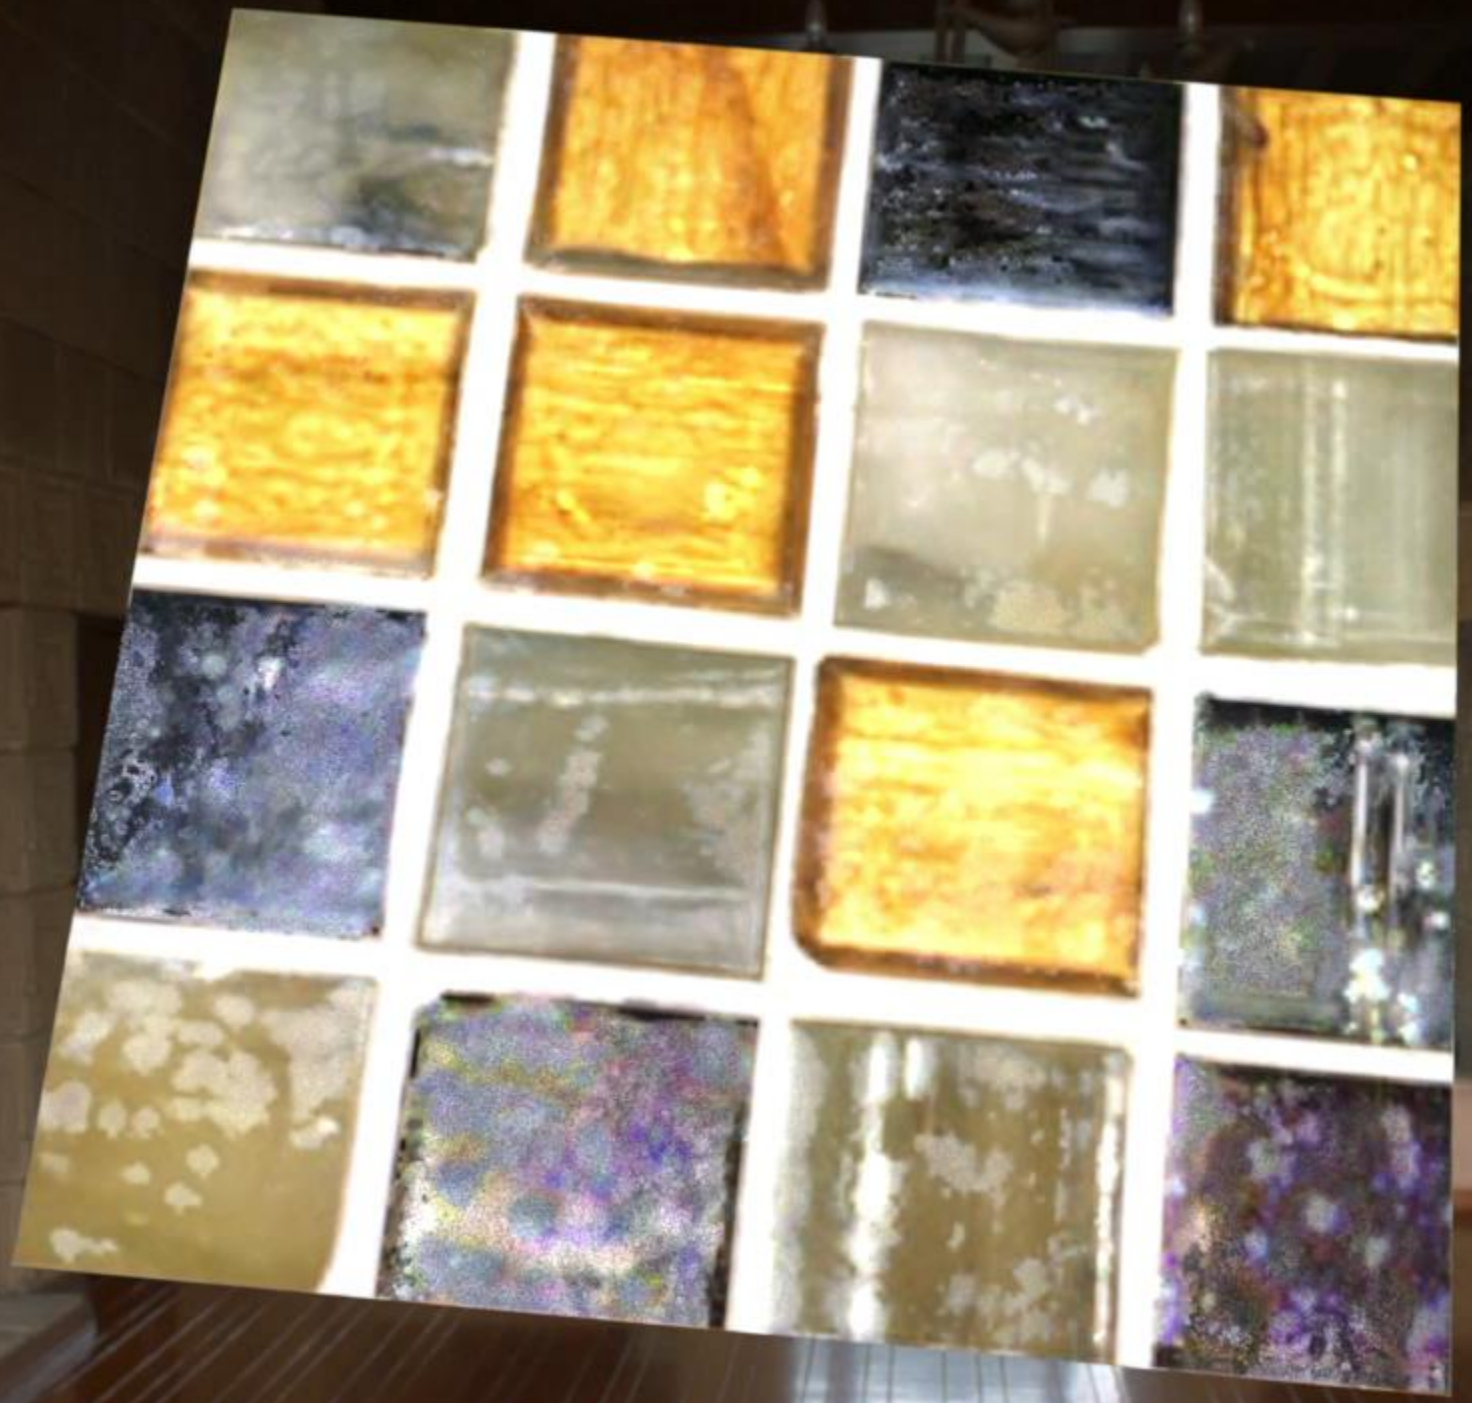

Ours

SVBRDF
maps

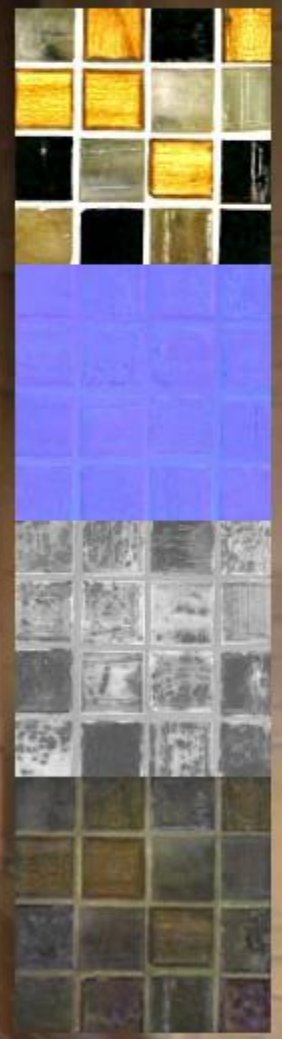

Book

Reference (Inputs)

Reference (Novel view)

Estimated Maps

Ours (Optimized)

Ours (Novel view)

Ours

Input
(7 used, 4 shown)

SVBRDF
maps

Comparison with [Gao et al. 2019]

Please check our main paper and supplemental materials for detailed analyses and comparisons

Comparison – Initialization

Ours is less sensitive to initialization

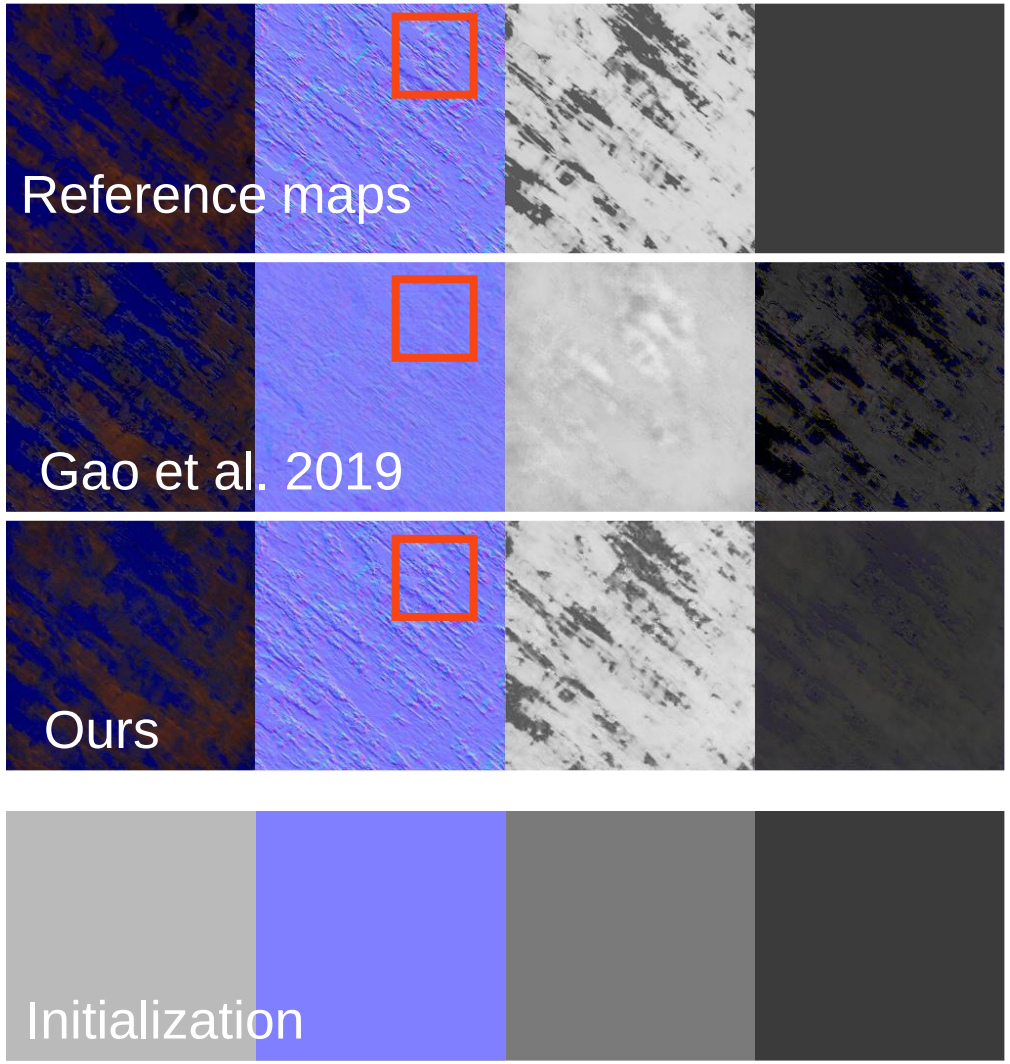

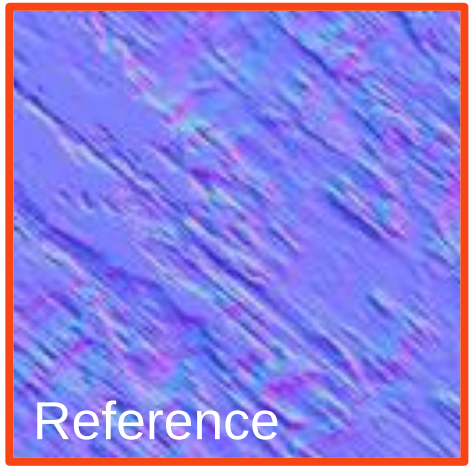

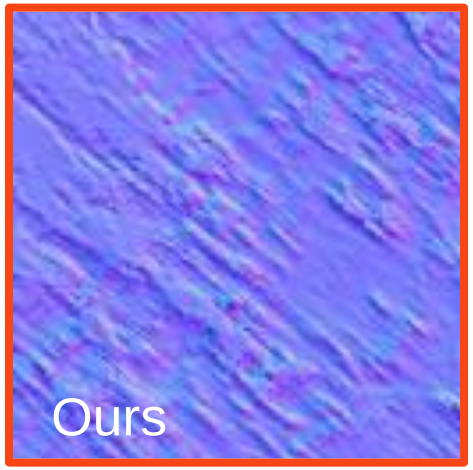

Comparison – Post-refinement

Ours is not require post-refinement to get sharp maps

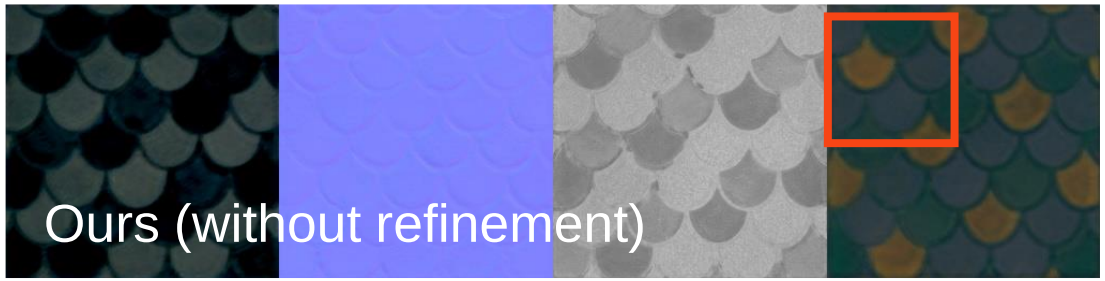

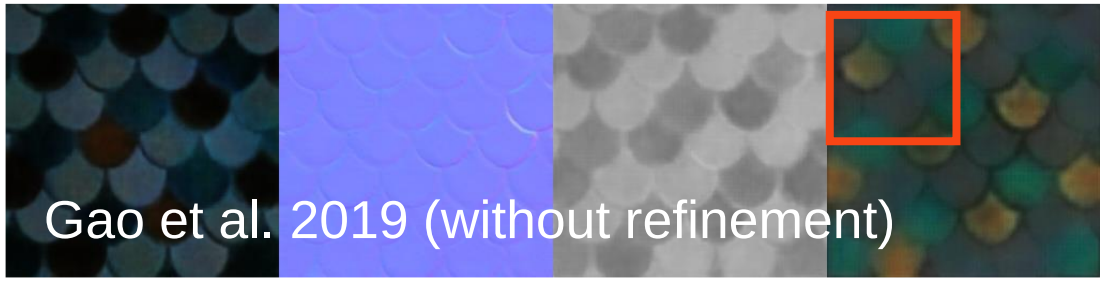

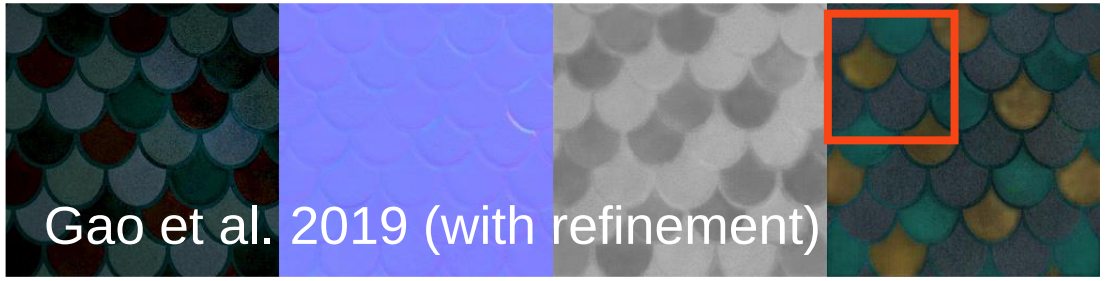

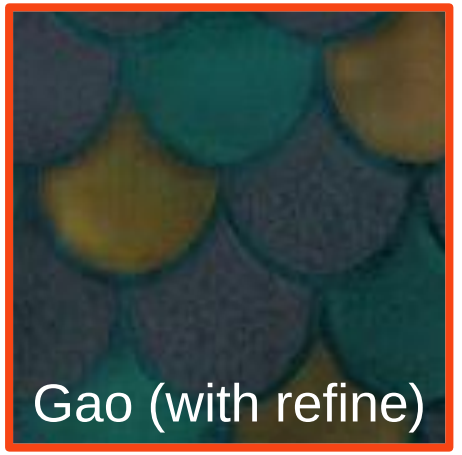

Latent Space Interpolation

Ours

Interpolation between

SVBRDF maps

Limitations and future work

- Our model relies on simple BRDF model
- Flat surface only
- Relies on known illumination
- Dataset is not general enough

Conclusion

- SVBRDF acquisition from a small number of input captures with smartphone

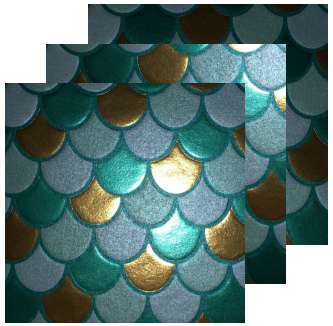

- An optimization framework with a powerful material prior (MaterialGAN)

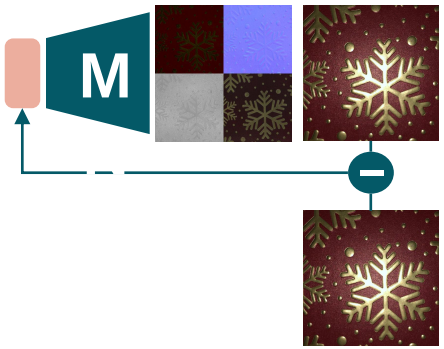

- High quality SVBRDF reconstruction without any good initialization

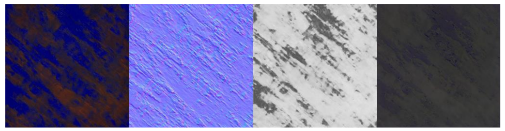

Thank you!

- Anonymous reviewers
- TJ Rhodes from Adobe Research for help with material capture hardware setup
- NSF IIS-1813553

Project page

Github repo