EBERHARD KARLS UNIVERSITÄT TÜBINGEN

FACULTY OF SCIENCE
Computer Graphics

Google Research

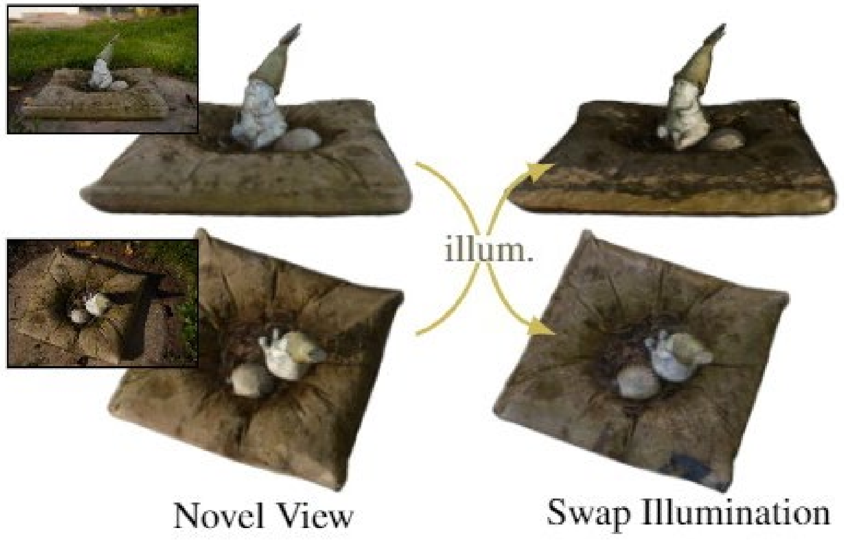

Neural-PIL: Neural Pre-Integrated Lighting for Reflectance Decomposition

Mark Boss[1], Varun Jampani[2], Raphael Braun[1], Jonathan T. Barron[2], Ce Liu[3*], Hendrik P. A. Lensch[1]

University of Tübingen [1], Google Research[2], Microsoft[3]

Conference on Neural Information Processing Systems (NeurIPS), 2021

*Work done during time at Google

Relightable 3D assets from image collections

- Learning shape, illumination and material properties (BRDF) from unconstrained image collection
- Enables relighting under any illumination

Samples from image collection
(Taken under unconstrained illumination)

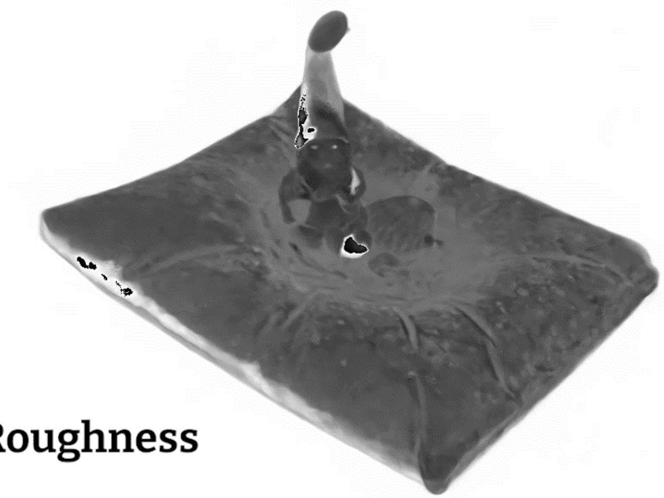

Neural Reflectance Volume

- Estimate density, BRDF and normals at a given 3D location

Sample Input Images

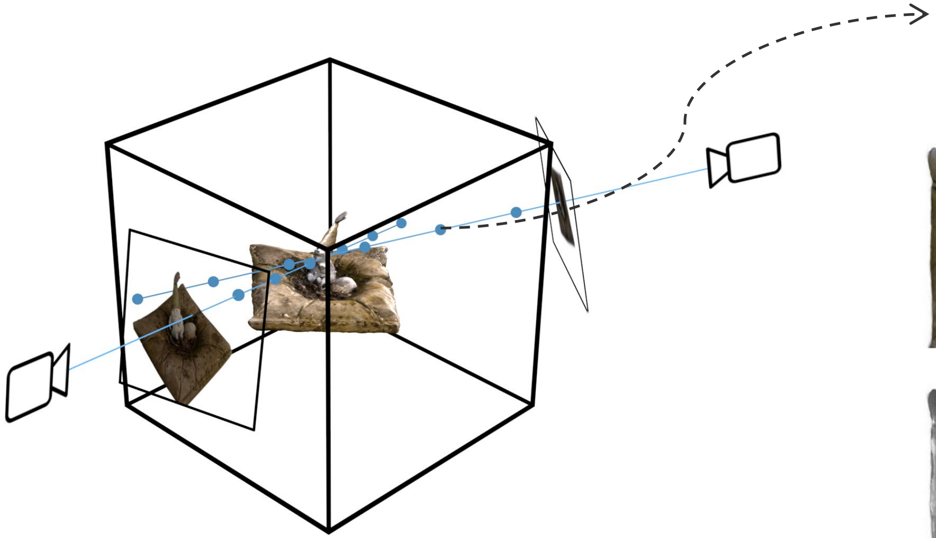

Density
Normals
Diffuse Albedo
Specular Albedo
Roughness

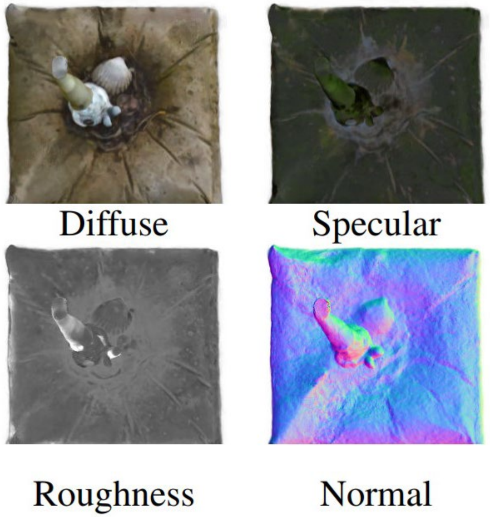

Two main challenges

1. Rendering with BRDF and illumination is expensive

$$L_o(\boldsymbol{x},\boldsymbol{\omega}_o) = \underbrace{\frac{\boldsymbol{b}_d}{\pi}\boxed{\int_\Omega L_i(\boldsymbol{x},\boldsymbol{\omega}_i)}(\boldsymbol{\omega}_i\cdot\boldsymbol{n})\,d\boldsymbol{\omega}_i}_{\text{diffuse}} + \underbrace{\boxed{\int_\Omega} f_s(\boldsymbol{x},\boldsymbol{\omega}_i,\boldsymbol{\omega}_o;\boldsymbol{b}_s,b_r)\boxed{L_i(\boldsymbol{x},\boldsymbol{\omega}_i)}(\boldsymbol{\omega}_i\cdot\boldsymbol{n})\,d\boldsymbol{\omega}_i}_{\text{specular}}$$

- Physically based rendering enables relighting under any illumination

2. Solution space is highly ambiguous

Neural Reflectance Volume – Previous Work

- NeRD[1] and PhySG[2] create neural volume jointly with illumination estimation
- Spherical Gaussian illumination model

Sample Input Images

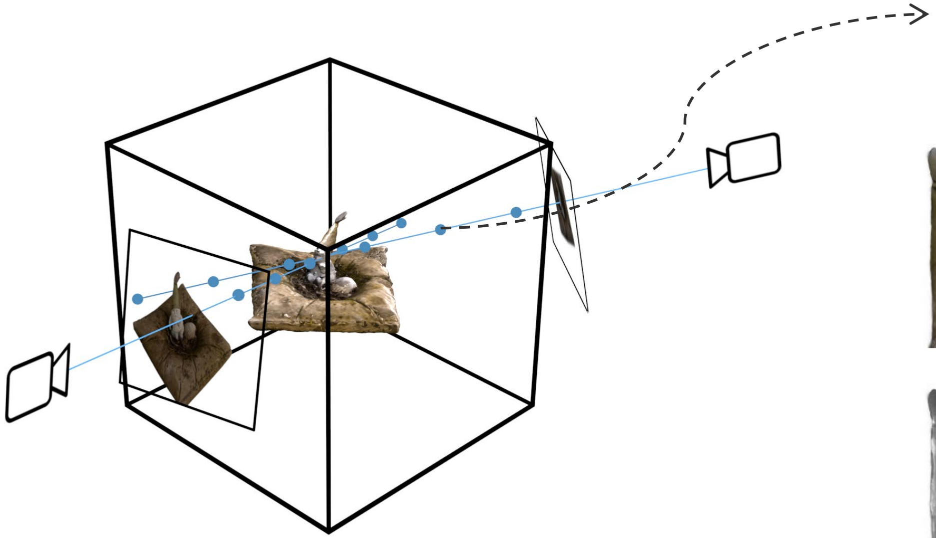

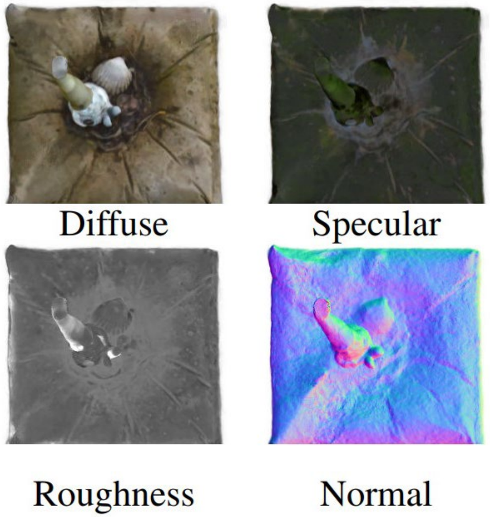

[1] Boss et al. "NeRD: Neural Reflectance Decomposition from Image Collections." In *ICCV 2021*
[2] Zhang et al. "PhySG: Inverse Rendering with Spherical Gaussians for Physics-based Material Editing and Relighting." In *CVPR* 2021

Issues with Spherical Gaussians

Fitted SGs

SGs Render

GT Environment Map

GT Render

Pre-Integrated Lighting [1]

- Pre-computing light integrals for faster rendering [4]

$$L_o(\boldsymbol{x},\boldsymbol{\omega}_o) = \underbrace{\frac{\boldsymbol{b}_d}{\pi}\boxed{\int_\Omega L_i(\boldsymbol{x},\boldsymbol{\omega}_i)}(\boldsymbol{\omega}_i\cdot\boldsymbol{n})\,d\boldsymbol{\omega}_i}_{\text{diffuse}} + \underbrace{\boxed{\int_\Omega} f_s(\boldsymbol{x},\boldsymbol{\omega}_i,\boldsymbol{\omega}_o;\boldsymbol{b}_s,b_r)\boxed{L_i(\boldsymbol{x},\boldsymbol{\omega}_i)}(\boldsymbol{\omega}_i\cdot\boldsymbol{n})\,d\boldsymbol{\omega}_i}_{\text{specular}}$$

$$\downarrow$$

$$L_o(\boldsymbol{x},\boldsymbol{\omega}_o) \approx \underbrace{(\boldsymbol{b}_d/\pi)\tilde{L}_i(\boldsymbol{n},1)}_{\text{diffuse}} + \underbrace{\boldsymbol{b}_s(F_0(\boldsymbol{\omega}_o,\boldsymbol{n})B_0(\boldsymbol{\omega}_o\cdot\boldsymbol{n},b_r) + B_1(\boldsymbol{\omega}_o\cdot\boldsymbol{n},b_r))\tilde{L}_i(\boldsymbol{\omega}_r,b_r)}_{\text{specular}}$$

$$\boxed{\tilde{L}_i(\boldsymbol{\omega}_r,b_r)} = \int_\Omega D(b_r,\boldsymbol{\omega}_i,\boldsymbol{\omega}_r)L_i(\boldsymbol{x},\boldsymbol{\omega}_i)d\boldsymbol{\omega}_i$$

[1] Karis et al. - Real Shading in Unreal Engine 4

Pre-Integrated Lighting [1]

- Pre-computing light integrals for faster rendering [4]

$$L_o(\boldsymbol{x}, \boldsymbol{\omega}_o) = \underbrace{\frac{\boldsymbol{b}_d}{\pi} \int_\Omega L_i(\boldsymbol{x}, \boldsymbol{\omega}_i)(\boldsymbol{\omega}_i \cdot \boldsymbol{n})\, d\boldsymbol{\omega}_i}_{\text{diffuse}} + \underbrace{\int_\Omega f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o; \boldsymbol{b}_s, b_r) L_i(\boldsymbol{x}, \boldsymbol{\omega}_i)(\boldsymbol{\omega}_i \cdot \boldsymbol{n})\, d\boldsymbol{\omega}_i}_{\text{specular}}$$

$$\downarrow$$

$$\boxed{L_o(\boldsymbol{x}, \boldsymbol{\omega}_o) \approx \underbrace{(\boldsymbol{b}_d/\pi)\tilde{L}_i(\boldsymbol{n}, 1)}_{\text{diffuse}} + \underbrace{\boldsymbol{b}_s(F_0(\boldsymbol{\omega}_o, \boldsymbol{n})B_0(\boldsymbol{\omega}_o \cdot \boldsymbol{n}, b_r) + B_1(\boldsymbol{\omega}_o \cdot \boldsymbol{n}, b_r))\tilde{L}_i(\boldsymbol{\omega}_r, b_r)}_{\text{specular}}}$$

$$\tilde{L}_i(\boldsymbol{\omega}_r, b_r) = \int_\Omega D(b_r, \boldsymbol{\omega}_i, \boldsymbol{\omega}_r) L_i(\boldsymbol{x}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i$$

[1] Karis et al. - Real Shading in Unreal Engine 4

Pre-Integrated Lighting [1]

[1] Karis et al. - Real Shading in Unreal Engine 4

Pre-Integrated Lighting [1]

[1] Karis et al. - Real Shading in Unreal Engine 4

Pre-Integrated Lighting [1]

$$\tilde{L}_i(\omega_r, b_r) = \int_\Omega D(b_r, \omega_i, \omega_r) L_i(x, \omega_i) d\omega_i$$

[1] Karis et al. - Real Shading in Unreal Engine 4

Pre-Integrated Lighting [1]

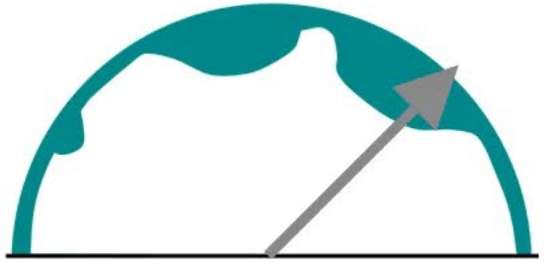

Pre-integrated Environment

Rendered Sphere

$$\tilde{L}_i(\omega_r, \boxed{b_r}) = \int_\Omega D(b_r, \omega_i, \omega_r) L_i(x, \omega_i) d\omega_i$$

[1] Karis et al. - Real Shading in Unreal Engine 4

NEURAL PIL

Challenge 1: Rendering with BRDF and illumination is expensive

Neural-PIL

- Light pre-integration is still expensive.
- We need to do the pre-integration on the fly as we also estimate lighting.
- We propose Neural-PIL that converts light pre-integration into a simple network query.
- Architecture based on Pi-GAN[1].

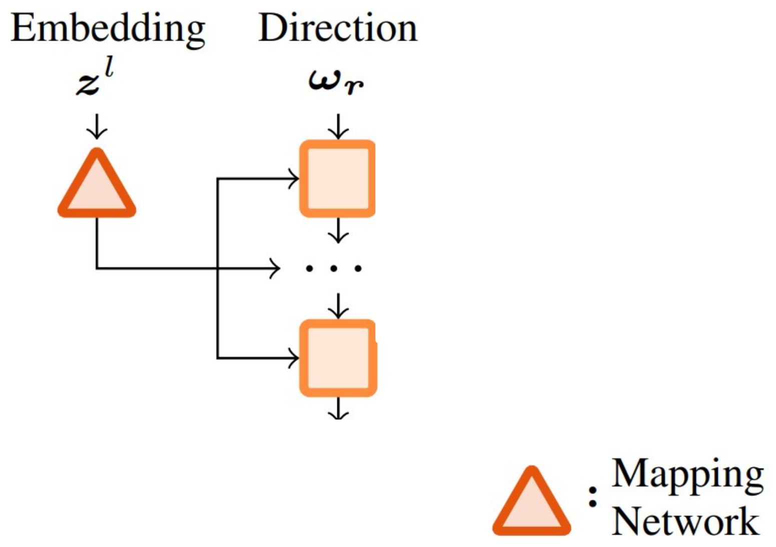

[1] Chan et al. – "pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis" In *CVPR* 2021

Neural-PIL

- Light pre-integration is still expensive.
- We need to do the pre-integration on the fly as we also estimate lighting.
- We propose Neural-PIL that converts light pre-integration into a simple network query.
- Architecture based on Pi-GAN[1].

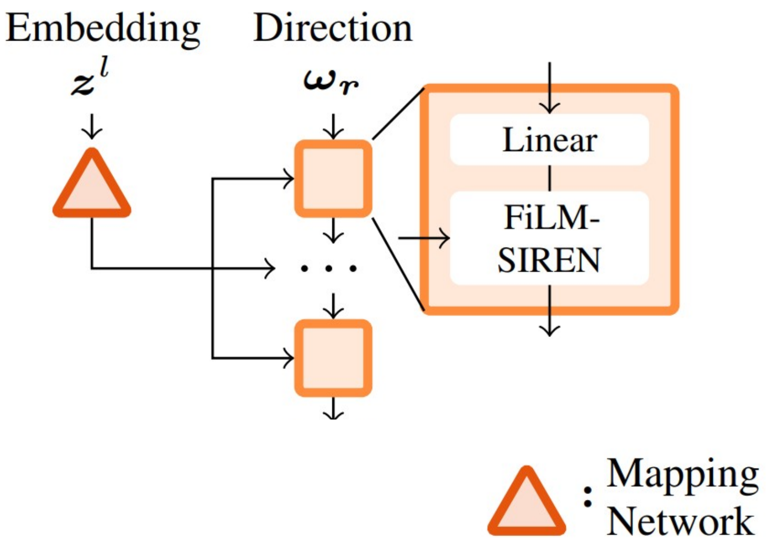

[1] Chan et al. – “pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis” In *CVPR* 2021

Neural-PIL

- Light pre-integration is still expensive.
- We need to do the pre-integration on the fly as we also estimate lighting.
- We propose Neural-PIL that converts light pre-integration into a simple network query.
- Architecture based on Pi-GAN[1].

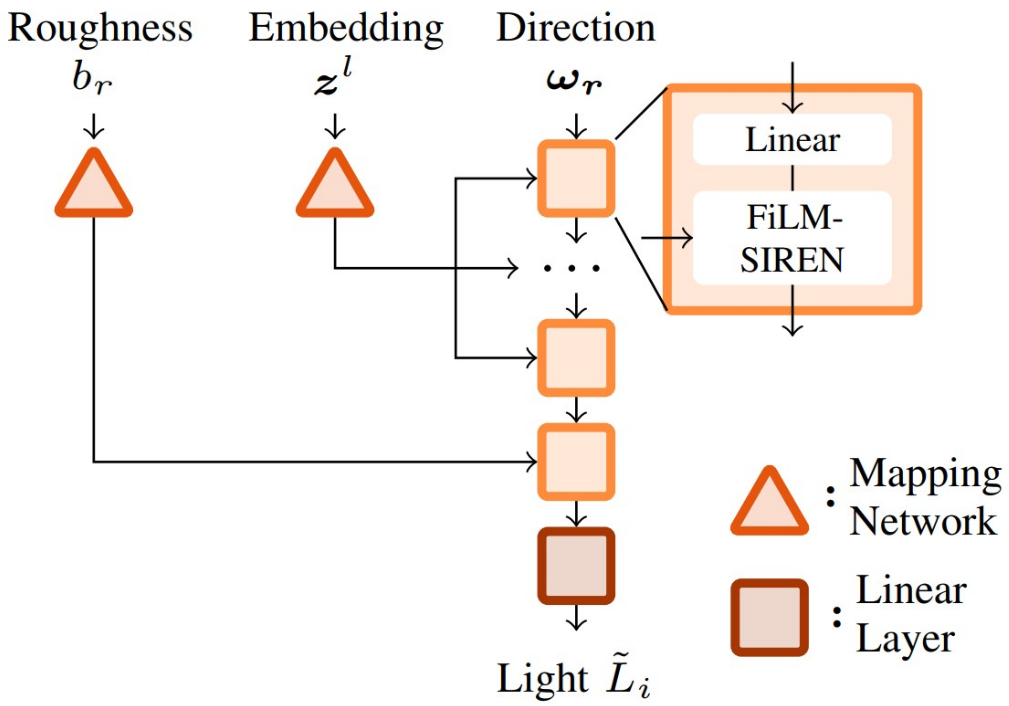

[1] Chan et al. – "pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis" In *CVPR* 2021

Neural-PIL

- Light pre-integration is still expensive.
- We need to do the pre-integration on the fly as we also estimate lighting.
- We propose Neural-PIL that converts light pre-integration into a simple network query.
- Architecture based on Pi-GAN[1].

Rendering	SGs	Neural-PIL
1 Million Samples	0.21s	**0.00186s**

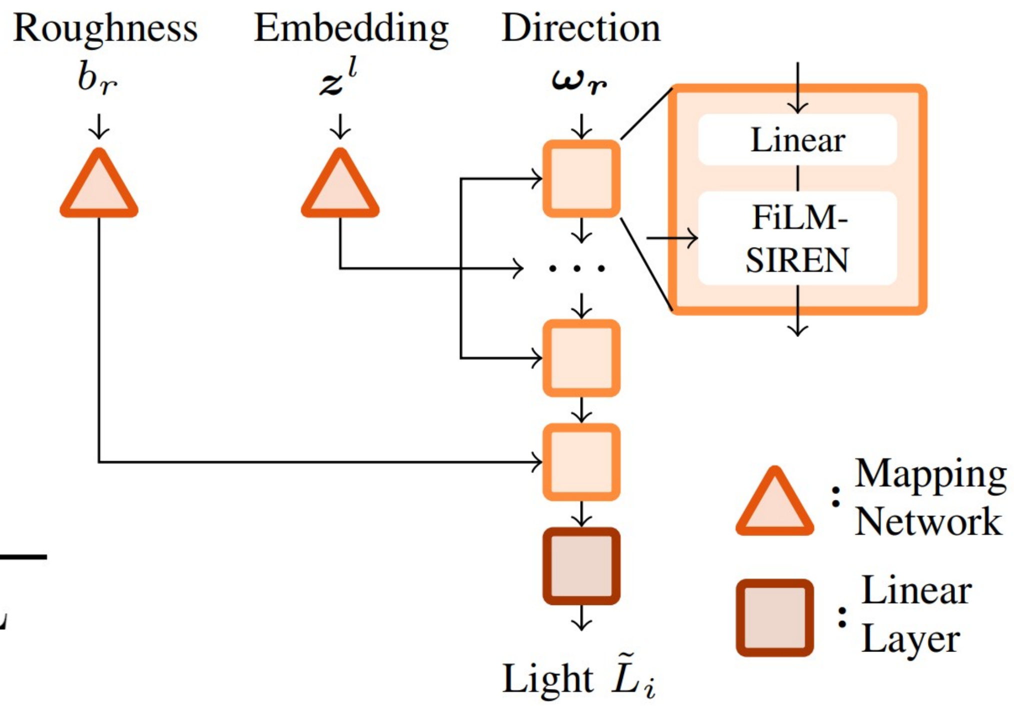

[1] Chan et al. – “pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis” In *CVPR* 2021

Neural-PIL vs. Spherical Gaussians

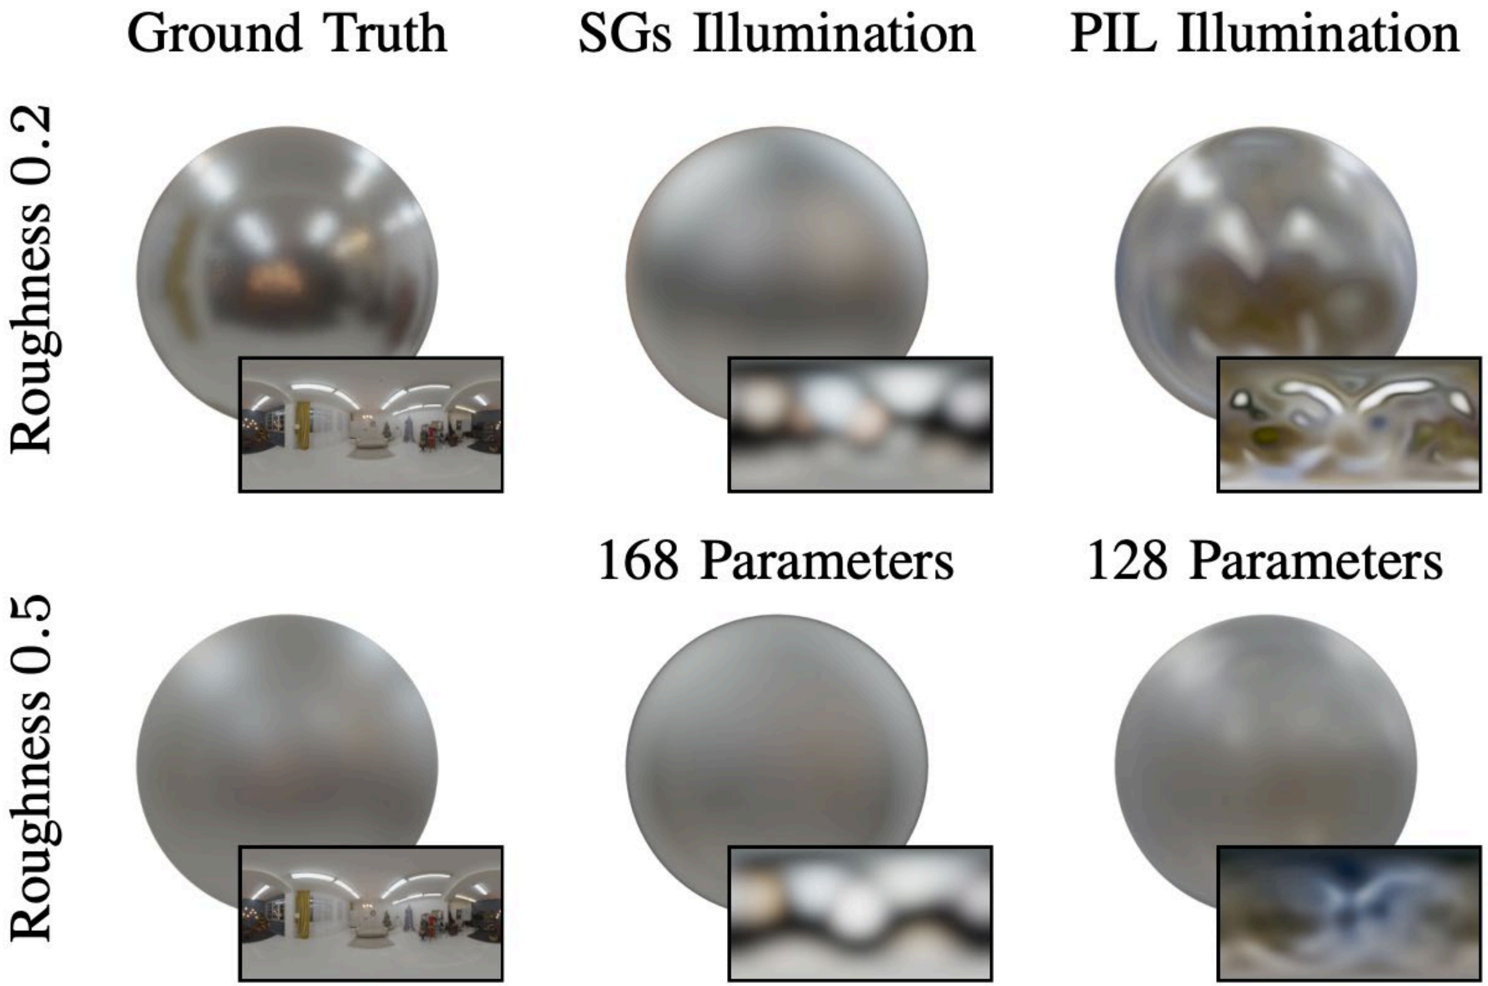

DEEP PRIORS WITH SMOOTH MANIFOLD AUTO-ENCODERS

Challenge 2: Solution space is highly ambiguous

Smooth Manifold Auto-Encoders (SMAE)

- Learn a smooth low-dimensional manifold to represent BRDF and lighting.
- Auto-encoder learning with interpolated latent space.

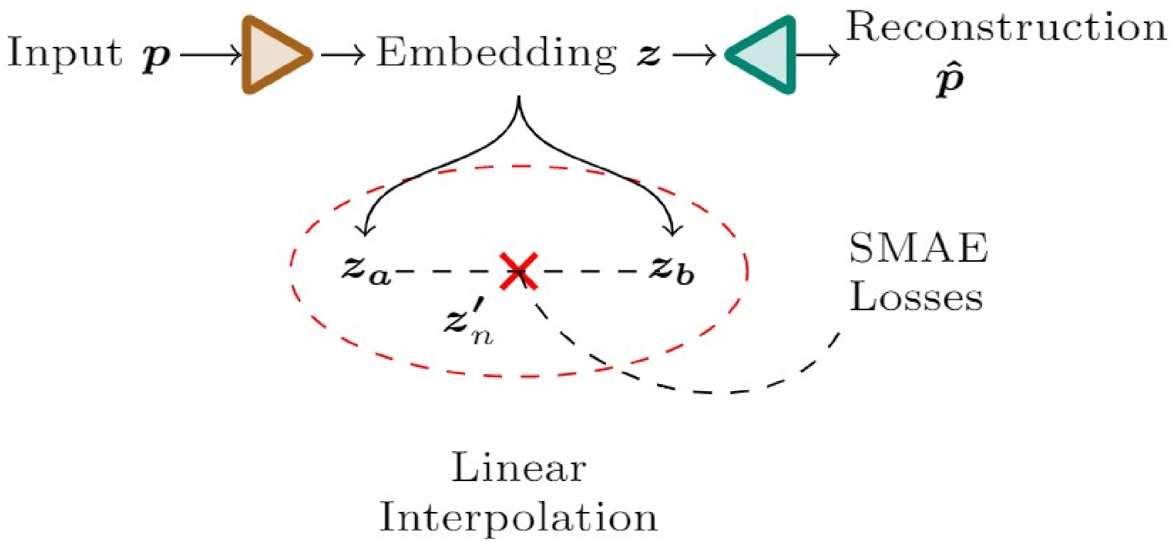

BRDF Manifold with SMAE

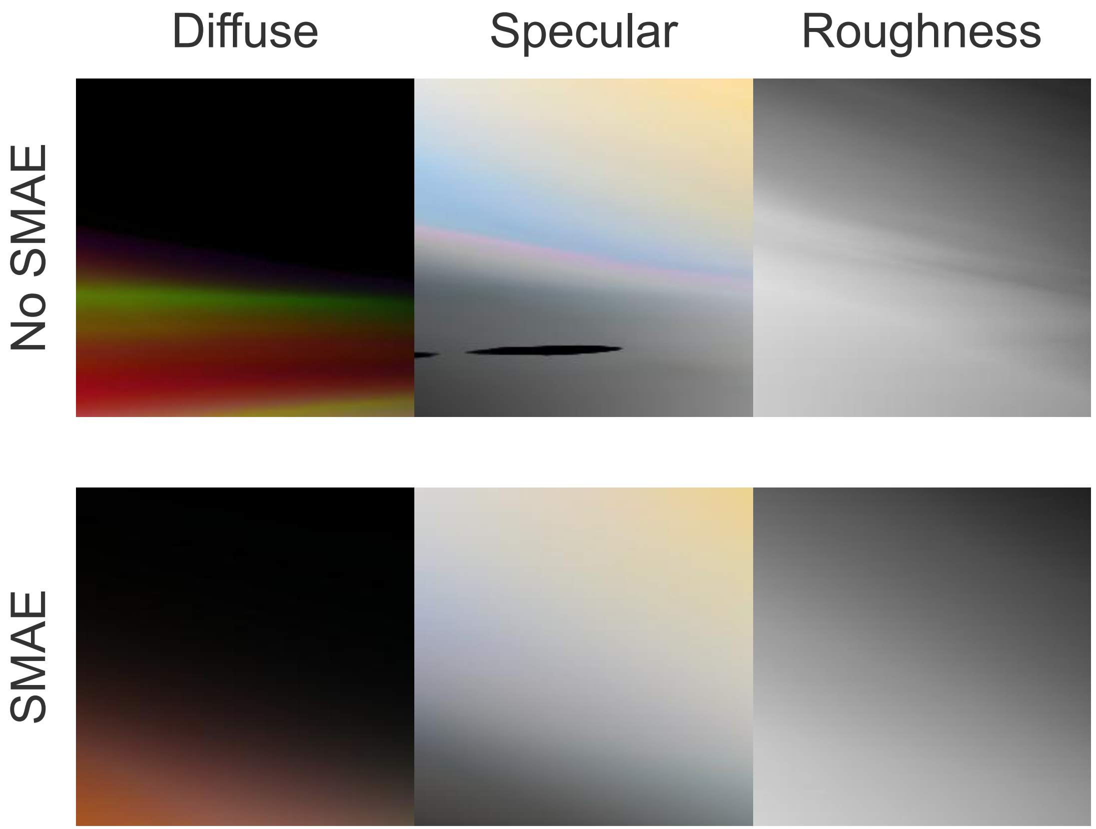

BRDF-SMAE Latent Space Interpolation (Corners)

RESULTS

Novel View and Relighting

Compared to NeRF[1] and NeRD[2]

Ground Truth

NeRF

[1] Mildenhall et al. “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis.” In *ECCV 2020*
[2] Boss et al. “NeRD: Neural Reflectance Decomposition from Image Collections.” In *ICCV 2021*

BRDF and Light Estimation

Parameter [PSNR↑]	Li *et al.*[1]	Li *et al.*[1] + NeRF[2]	NeRD[3]	Ours
Diffuse	1.06	1.15	18.24	**20.22**
Specular	—	—	**25.70**	16.84
Roughness	17.18	17.28	15.00	**24.82**

[1] Li et al. "Learning to Reconstruct Shape and Spatially-Varying Reflectance from a Single Image." In *Siggraph Asia* 2018
[2] Mildenhall et al. "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis." In *ECCV 2020*
[3] Boss et al. "NeRD: Neural Reflectance Decomposition from Image Collections." In *ICCV 2021*

View Synthesis and Relighting

View synthesis (Single illumination dataset)

	Synthetic		Real-World	
Method	PSNR↑	SSIM↑	PSNR↑	SSIM↑
NeRF[1]	**34.24**	**0.97**	23.34	0.85
NeRD[2]	30.07	0.95	23.86	0.88
Ours	30.08	0.95	**23.95**	**0.90**

Both view synthesis and relighting

	Synthetic		Real-World	
Method	PSNR↑	SSIM↑	PSNR↑	SSIM↑
NeRF[1]	21.05	0.89	20.11	0.87
NeRD[2]	27.96	0.95	25.81	**0.95**
Ours	**29.24**	**0.96**	**26.23**	**0.95**

[1] Mildenhall et al. “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis.” In *ECCV 2020*
[2] Boss et al. “NeRD: Neural Reflectance Decomposition from Image Collections.” In *ICCV 2021*

Neural-PIL: Summary

- Estimate shape, BRDF and illumination from images taken under different lightings.

- Addresses two main challenges
  - Expensive rendering → Neural-PIL for fast and effective light integration.
  - Ambiguous solution space → Smooth manifold auto-encoders.

- Experiments on both synthetic and real-world objects.
- State-of-the-art decomposition, view synthesis and relighting results.

Neural-PIL: Neural Pre-Integrated Lighting for Reflectance Decomposition

Thanks for watching

Visit the project page at:
https://markboss.me/publication/2021-neural-pil/