Supplementary Material - MIPNet: Neural Normal-to-Anisotropic-Roughness MIP mapping

ALBAN GAUTHIER, LTCI, Télécom Paris, Institut Polytechnique de Paris, France
ROBIN FAURY, Adobe Research, France
JÉRÉMY LEVALLOIS, Adobe Research, France
THÉO THONAT, Adobe Research, France
JEAN-MARC THIERY, Adobe Research, France
TAMY BOUBEKEUR, Adobe Research, France

CCS Concepts: • **Computing methodologies → Reflectance modeling**; **Texturing**.

Additional Key Words and Phrases: Material Appearance, Mipmaps, Machine Learning

ACM Reference Format:
Alban Gauthier, Robin Faury, Jérémy Levallois, Théo Thonat, Jean-Marc Thiery, and Tamy Boubekeur. 2022. Supplementary Material - MIPNet: Neural Normal-to-Anisotropic-Roughness MIP mapping. *ACM Trans. Graph.* 41, 6, Article 246 (December 2022), 6 pages. https://doi.org/10.1145/3550454.3555487

In this supplementary material, we present a comparison of different network architectures for our method, and an extended comparison of the generation of prefiltered materials between the baseline, our method and competitors.

1 SPECIFIC RESULTS

For materials which showcase examples on complex setup such as cross-shaped micro-structures, please refer to these examples:

- beaded sbai fabric
- synthetic flipflop topstitch square
- nickel diamond emboss pattern
- aluminium braided cable
- copper honeycomb pattern
- ferrofluid labyrinth
- pearl paint organic emboss

In addition, we point out to examples where out method greatly improves over the baseline:

- taurillon leather topstitch diamond
- pebble beach small
- urban rubble ground
- aluminium scrap
- fdm ceramic

2 ARCHITECTURE COMPARISONS

In Figure 1, we compare different architectures of our network, trained using the GGX and Beckmann distributions and Ashikhmin-Shirley model. First, we see similar behaviors for all distributions trained. We noticed that adding more layers in $\mathcal{H}_B$ has no influence on the resulting loss. In addition, there is a gap in loss values with $\mathcal{H}_B$ using hidden layers of size 512 compared to 1024. Finally, we did not notice any benefit to increase the size of hidden layers of $\mathcal{H}_A$ greater than 512.

3 ADDITIONNAL COMPARISONS

We show additional comparison with the GGX distribution in Figure 2, the Beckmann distribution in Figure 3, and the Ashikhmin-Shirley model Figure 4.

In Figure 5, we show additional qualitative comparison of renders with AutoLoD [Hasselgren et al. 2021] (middle right), and NeuMIP [Kuznetsov et al. 2021] (right). With respect to these two methods, we achieve reasonable material appearance preservation while being 3 to 4 orders of magnitude faster for pre-processing.

REFERENCES

Jon Hasselgren, Jacob Munkberg, Jaakko Lehtinen, Miika Aittala, and Samuli Laine. Appearance-driven automatic 3d model simplification. In *Eurographics Symposium on Rendering*, 2021.

Alexandr Kuznetsov, Krishna Mullia, Zexiang Xu, Miloš Hašan, and Ravi Ramamoorthi. Neumip: Multi-resolution neural materials. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 40(4), July 2021.

Authors' addresses: Alban Gauthier, LTCI, Télécom Paris, Institut Polytechnique de Paris, France, albangauthier25@gmail.com; Robin Faury, Adobe Research, France, faury@adobe.com; Jérémy Levallois, Adobe Research, France, levalloi@adobe.com; Théo Thonat, Adobe Research, France, thonat@adobe.com; Jean-Marc Thiery, Adobe Research, France, jthiery@adobe.com; Tamy Boubekeur, Adobe Research, France, boubek@adobe.com.

0730-0301/2022/12-ART246 $15.00
https://doi.org/10.1145/3550454.3555487

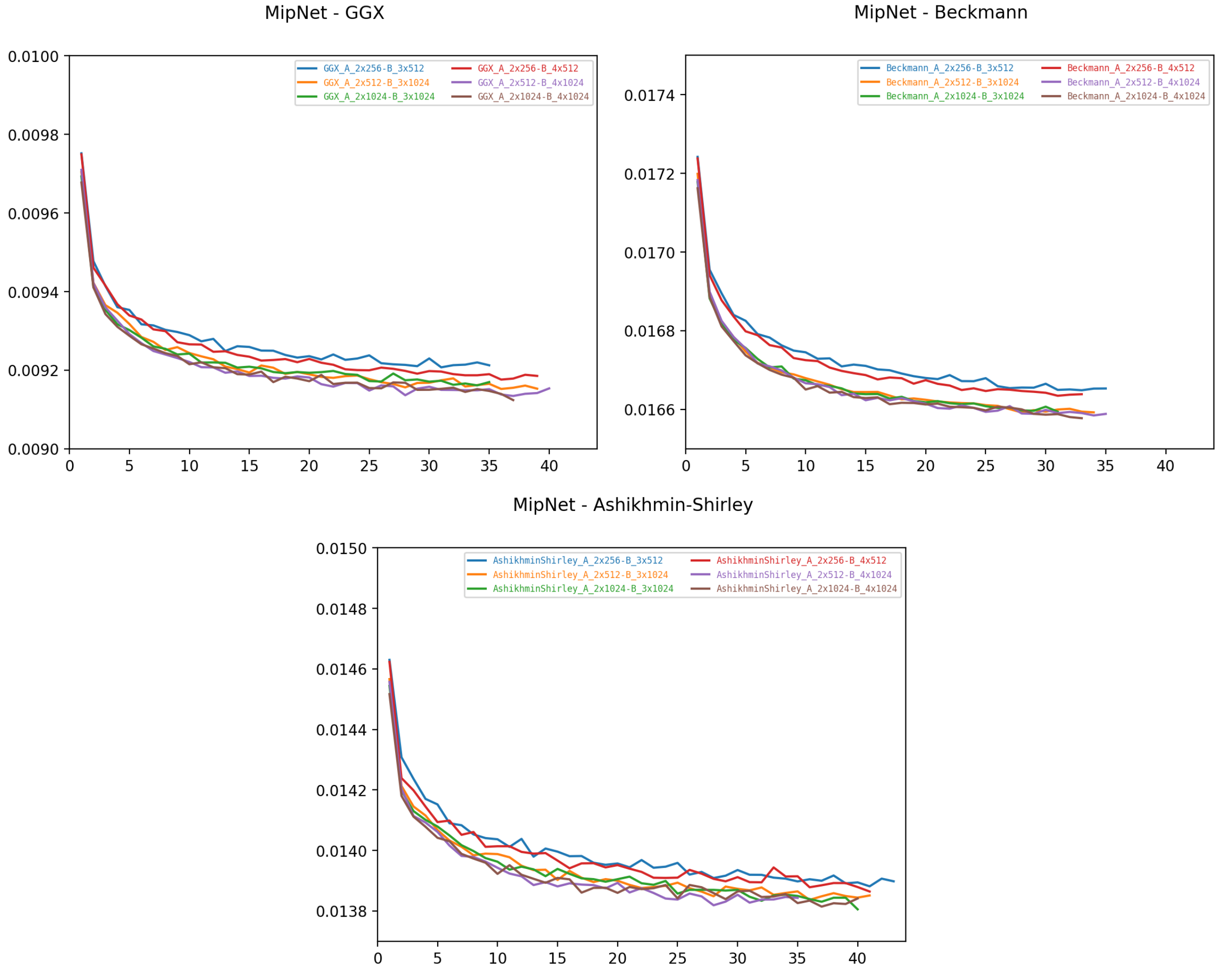

Fig. 1. Mean L1 loss (Y-axis) of our network computed with different architectures depending to the number of epochs (X-axis).
Top left: GGX distribution, top right: Beckmann distribution, bottom: Ashikhmin-Shirley distribution.
For all graphs:
XXX_A_2x256_B_3x512: $\mathcal{H}_A$ has 2 hidden layers of size 256, $\mathcal{H}_B$ has 3 hidden layers of size 512.
XXX_A_2x256_B_4x512: $\mathcal{H}_A$ has 2 hidden layers of size 256, $\mathcal{H}_B$ has 4 hidden layers of size 512.
XXX_A_2x512_B_3x1024: $\mathcal{H}_A$ has 2 hidden layers of size 512, $\mathcal{H}_B$ has 3 hidden layers of size 1024.
XXX_A_2x512_B_4x1024: $\mathcal{H}_A$ has 2 hidden layers of size 512, $\mathcal{H}_B$ has 4 hidden layers of size 1024.
XXX_A_2x1024_B_3x1024: $\mathcal{H}_A$ has 2 hidden layers of size 1024, $\mathcal{H}_B$ has 3 hidden layers of size 1024.
XXX_A_2x1024_B_4x1024: $\mathcal{H}_A$ has 2 hidden layers of size 1024, $\mathcal{H}_B$ has 4 hidden layers of size 1024.

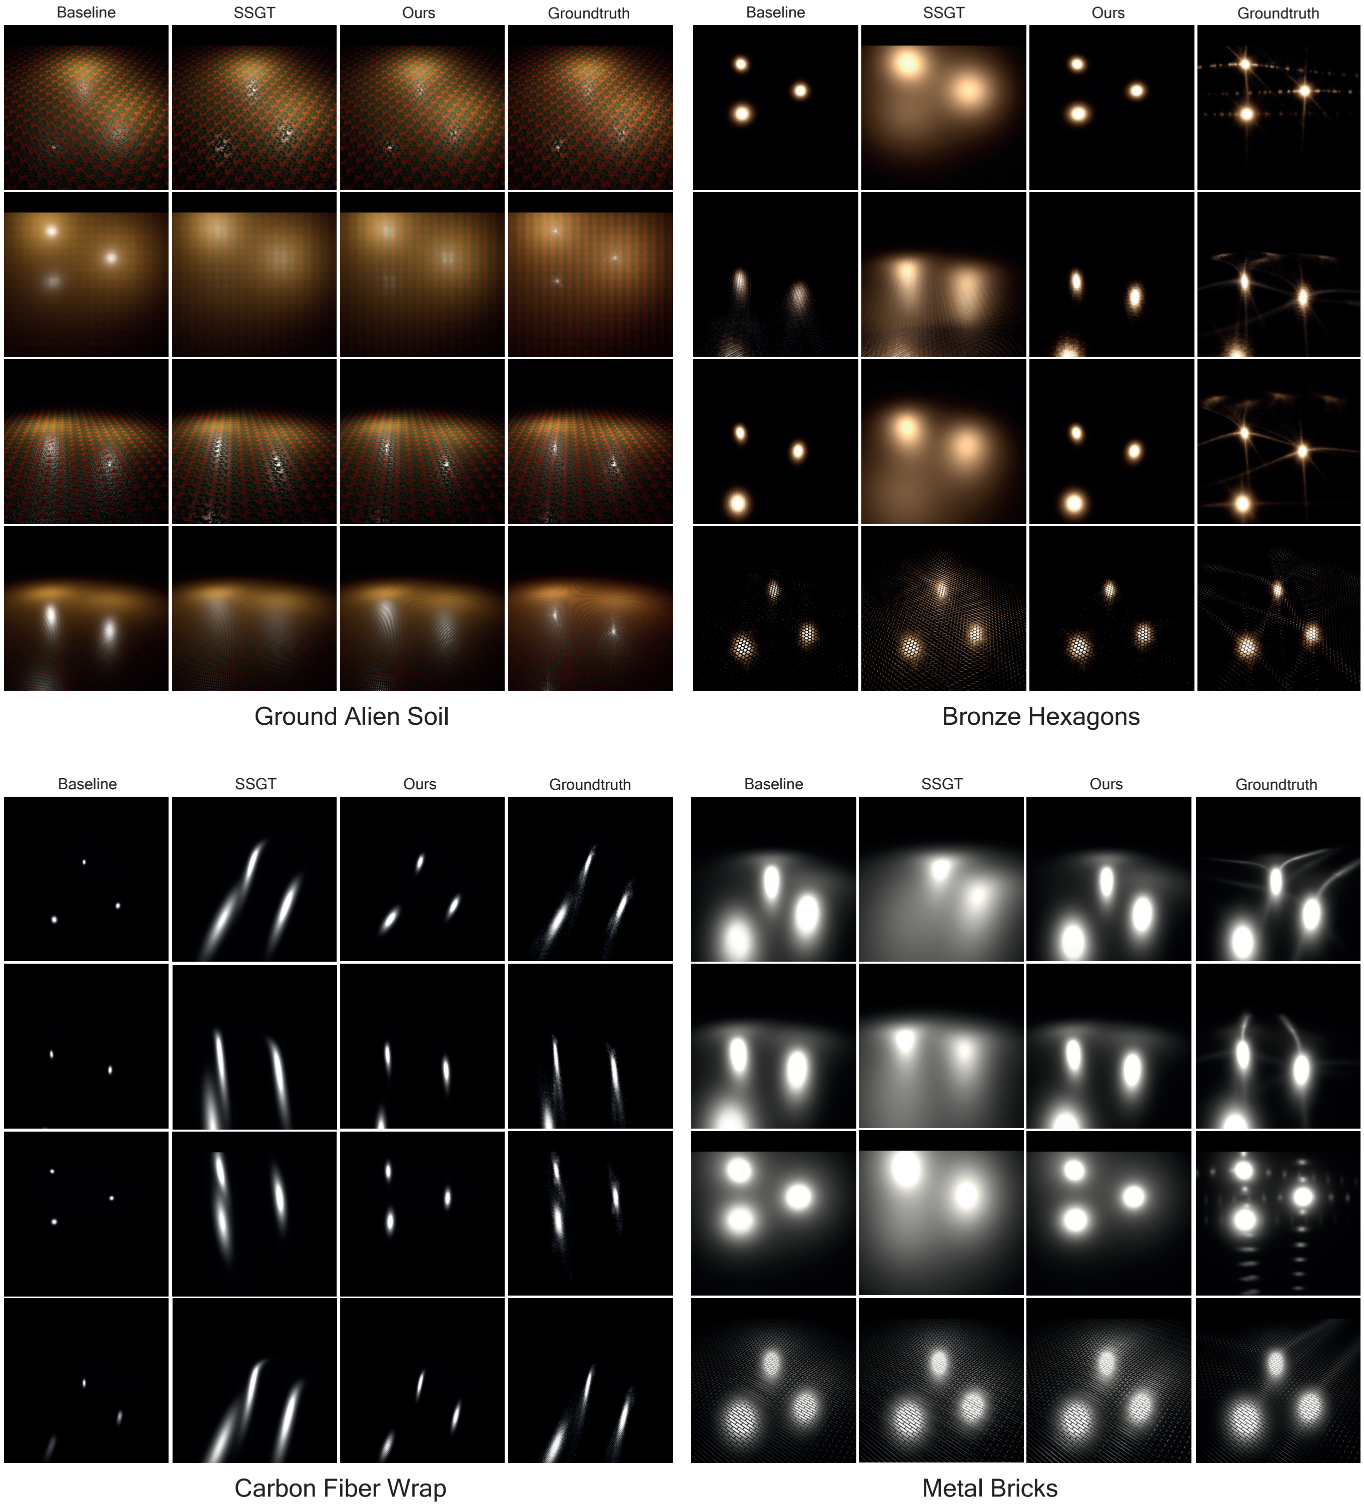

Fig. 2. Additional renderings using the GGX model.

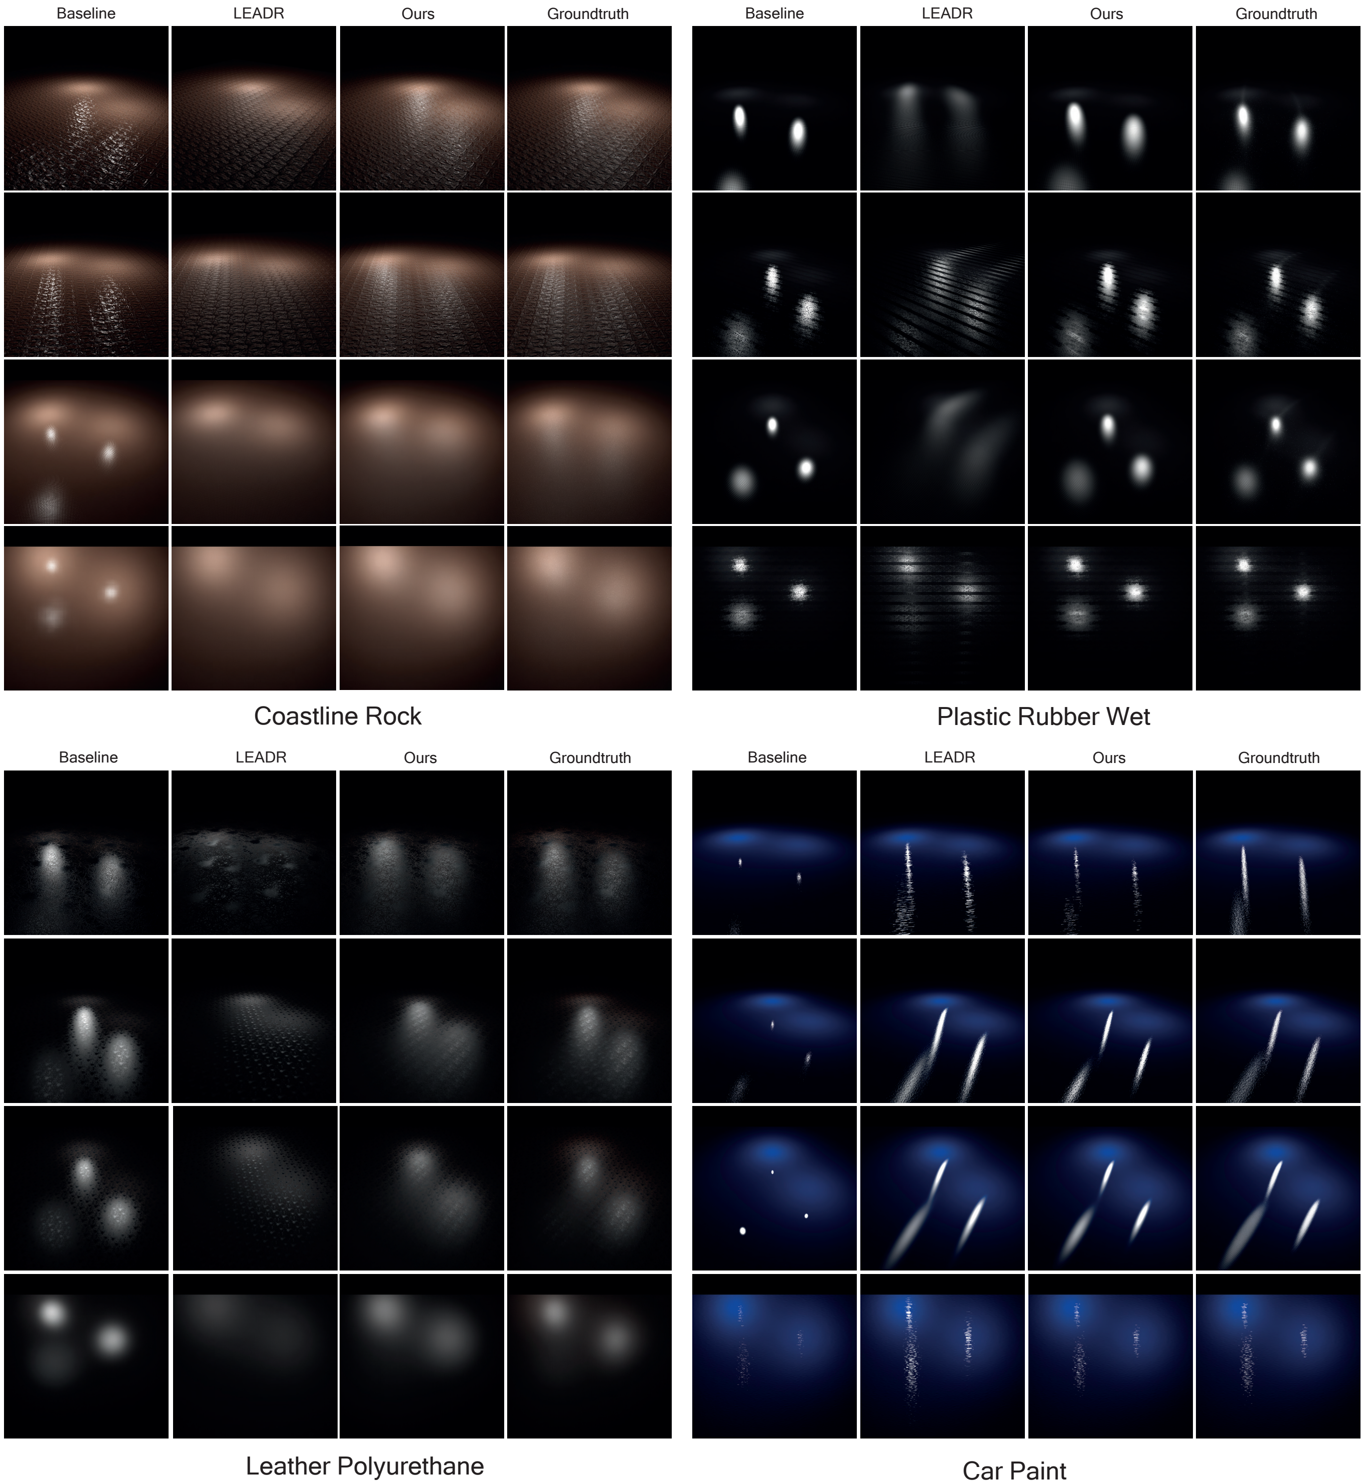

Fig. 3. Additional renderings using the Beckmann model.

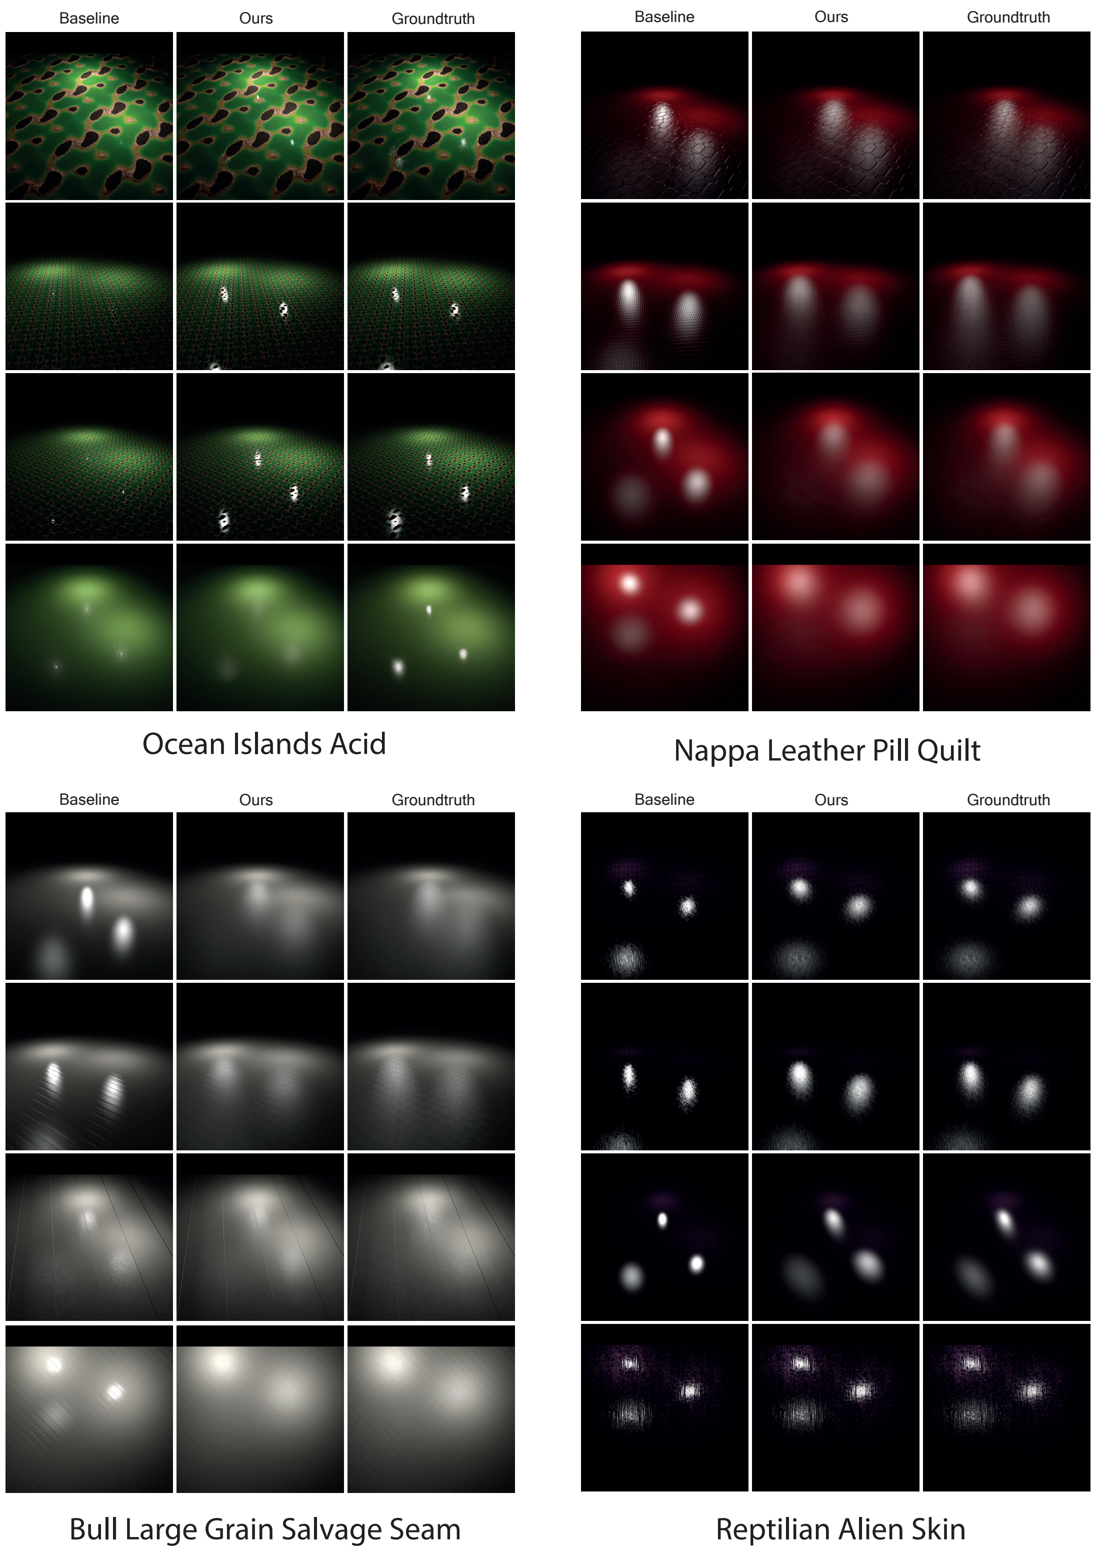

Fig. 4. Additional renderings using the Ashikhmin-Shirley model.

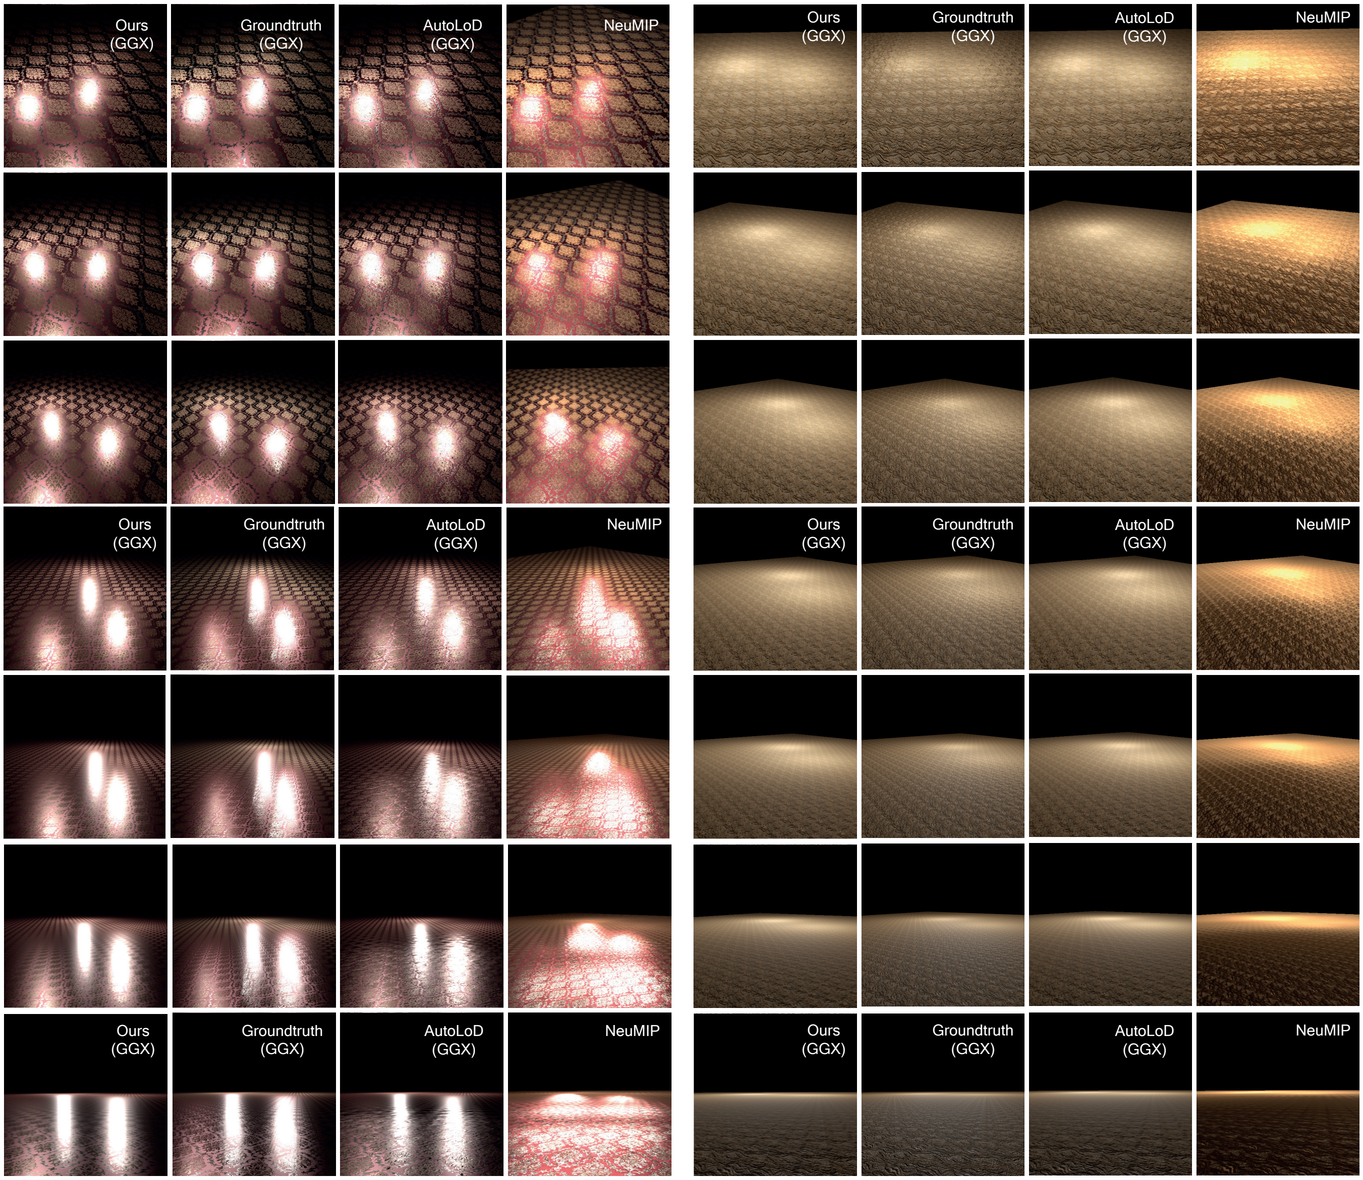

Fig. 5. Extended qualitative comparison of renders between our method (left), the groundtruth (middle left), AutoLoD [Hasselgren et al. 2021] (middle right), and NeuMIP [Kuznetsov et al. 2021] (right), with different viewpoints.