Learning to Reconstruct Shape and Spatially-Varying Reflectance from a Single Image

ZHENGQIN LI, University of California, San Diego
ZEXIANG XU, University of California, San Diego
RAVI RAMAMOORTHI, University of California, San Diego
KALYAN SUNKAVALLI, Adobe Research
MANMOHAN CHANDRAKER, University of California, San Diego

ACM Reference Format:
Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. 2018. Learning to Reconstruct Shape and Spatially-Varying Reflectance from a Single Image. *ACM Trans. Graph.* 37, 6, Article 269 (November 2018), 6 pages. https://doi.org/10.1145/3197517.3201313

1 VIDEO

We render video clips to better demonstrate the SVBRDF estimation results. For synthetic data, we have a side-by-side comparison of videos rendered with ground-truth BRDF parameters and estimated BRDF parameters. The appearances in the two videos are observed to be very similar. We can see the change of specular highlights when rotating the environment maps, which shows that spatially varying roughness has been successfully captured by the network. We do not have ground truth BRDF parameters for real data. But the change of shading and specular highlights when rotating the environment maps looks realistic. In order to observe more high frequency specular highlights, we also render a video with a moving point light source for real data. Note that the above results are rendered without depth maps. To show that our method can successfully recover the geometry of objects, we also render two video clips of novel view synthesis using the estimated depths, with and without the flash light. We use median filter to make the original depth estimations smoother. However, when rendered without flash light, some artifacts can still be observed due to incorrect depth estimation. Given that the depth estimation is less accurate than normal estimation, one potential future improvement will be to use normal prediction to refine depth estimation [Nehab et al. 2005].

2 CASCADE NETWORK

In Figures 1 and 2, we show the effect of using cascade structure for SVBRDF estimation. For synthetic data, we add an error visualization so that the differences between estimated BRDF parameters and ground truth can be clearly observed. We observe that the cascade structure significantly reduces the error. Interestingly, the cascade network implicitly learns a piece-wise constant prior for spatially varying roughness prediction. Such a prior is likely to be important for this challenging problem, since with such limited inputs, the network may not have enough information to determine roughness values for large portions of the surface. For real data, even without ground truth, the improvements across different levels of cascades are still very evident.

3 BRDF MODEL

We use the microfacet BRDF model in [Karis and Games 2013]. Following the notation in the main paper, let A, N, R be the diffuse albedo, normal and roughness respectively. Let L and V be light and view direction and $H = \frac{V+H}{2}$ be their half vector. Our BRDF model is defined as

$$f(A, N, R, L, V) = \frac{A}{\pi} + \frac{D(H,R)F(V,H)G(L,V,H,R)}{4(N \cdot L)(N \cdot V)} \qquad (1)$$

$D(H,R)$, $F(V,H)$ and $G(L,V,H,R)$ are the distribution, Fresnel and geometric term respectively, which are defined as

$$\begin{aligned}
D(H,R) &= \frac{\alpha^2}{\pi\left[(N\cdot H)^s(\alpha^2-1)+1\right]^2} \\
\alpha &= R^2 \\
F(V,H) &= (1-F_0)2^{-[5.55473(V\cdot H)+6.8316](V\cdot H)} \\
G(L,V,R) &= G_1(V,N)G_1(L,N) \\
G_1(V,N) &= \frac{N\cdot V}{(N\cdot V)(1-k)+k} \\
G_1(L,N) &= \frac{N\cdot L}{(N\cdot L)(1-k)+k} \\
k &= \frac{(R+1)^2}{8}
\end{aligned}$$

We set $F_0 = 0.05$ as suggested in [Karis and Games 2013].

REFERENCES

Brian Karis and Epic Games. 2013. Real shading in unreal engine 4. *Proc. Physically Based Shading Theory Practice* (2013).

Diego Nehab, Szymon Rusinkiewicz, James Davis, and Ravi Ramamoorthi. 2005. Efficiently combining positions and normals for precise 3D geometry. In *ACM transactions on graphics (TOG)*, Vol. 24. ACM, 536–543.

Authors' addresses: Zhengqin Li, University of California, San Diego, zhl378@eng.ucsd.edu; Zexiang Xu, University of California, San Diego, zexiangxu@cs.ucsd.edu; Ravi Ramamoorthi, University of California, San Diego, ravir@cs.ucsd.edu; Kalyan Sunkavalli, Adobe Research, sunkaval@adobe.com; Manmohan Chandraker, University of California, San Diego, mkchandraker@eng.ucsd.edu.

0730-0301/2018/11-ART269 $15.00
https://doi.org/10.1145/3197517.3201313

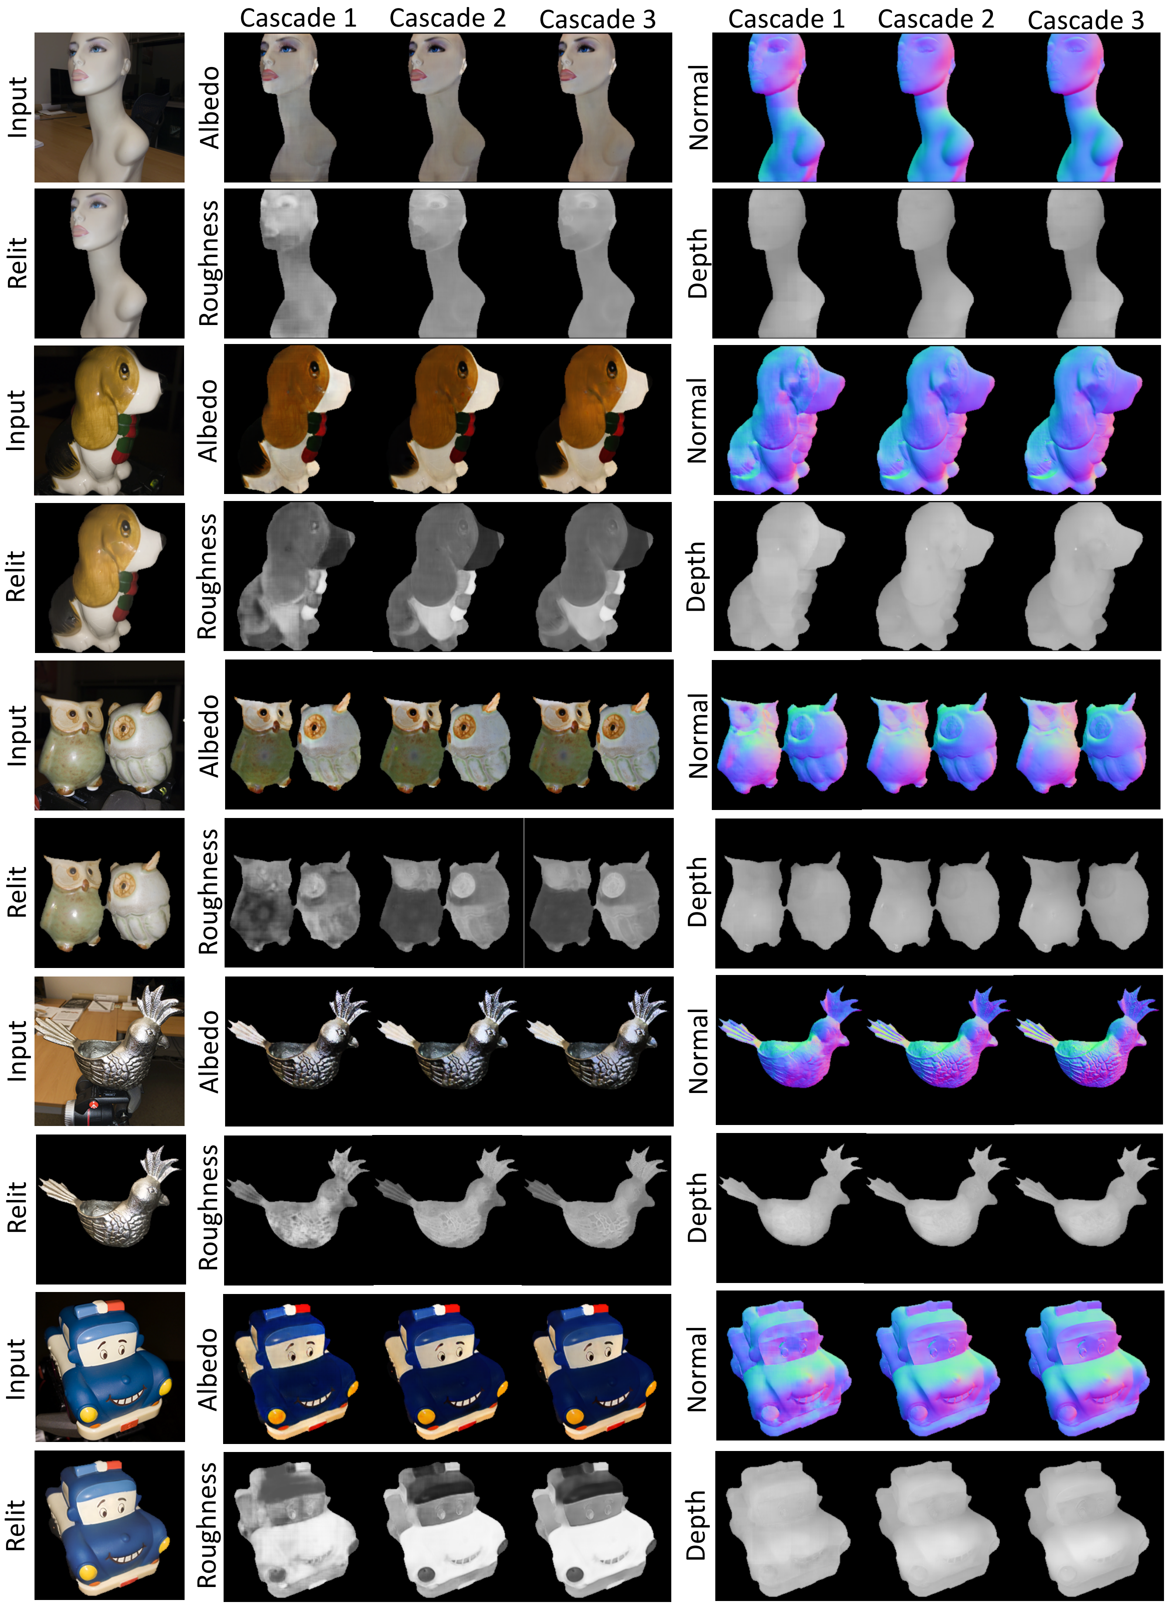
Cascade 1
Cascade 2
Cascade 3
Cascade 1
Cascade 2
Cascade 3
Input
Relit
Albedo
Roughness
Normal
Depth

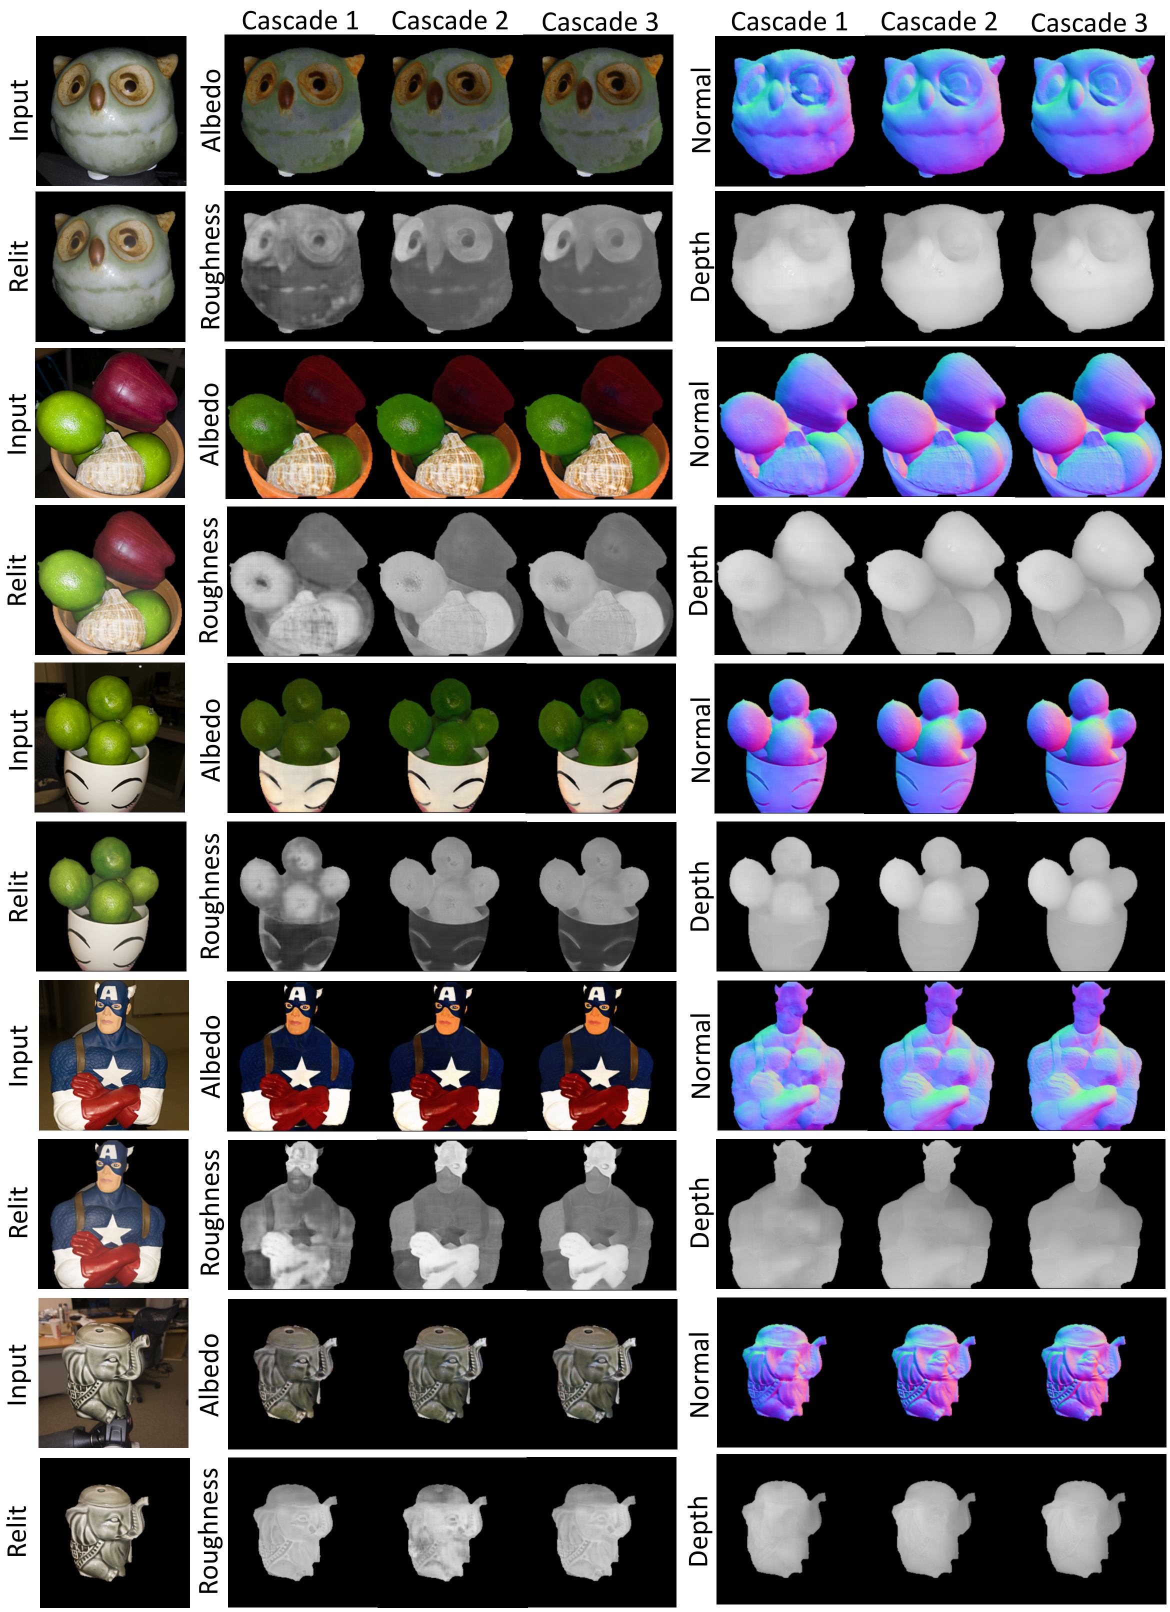
Cascade 1
Cascade 2
Cascade 3
Cascade 1
Cascade 2
Cascade 3
Input
Albedo
Normal
Relit
Roughness
Depth

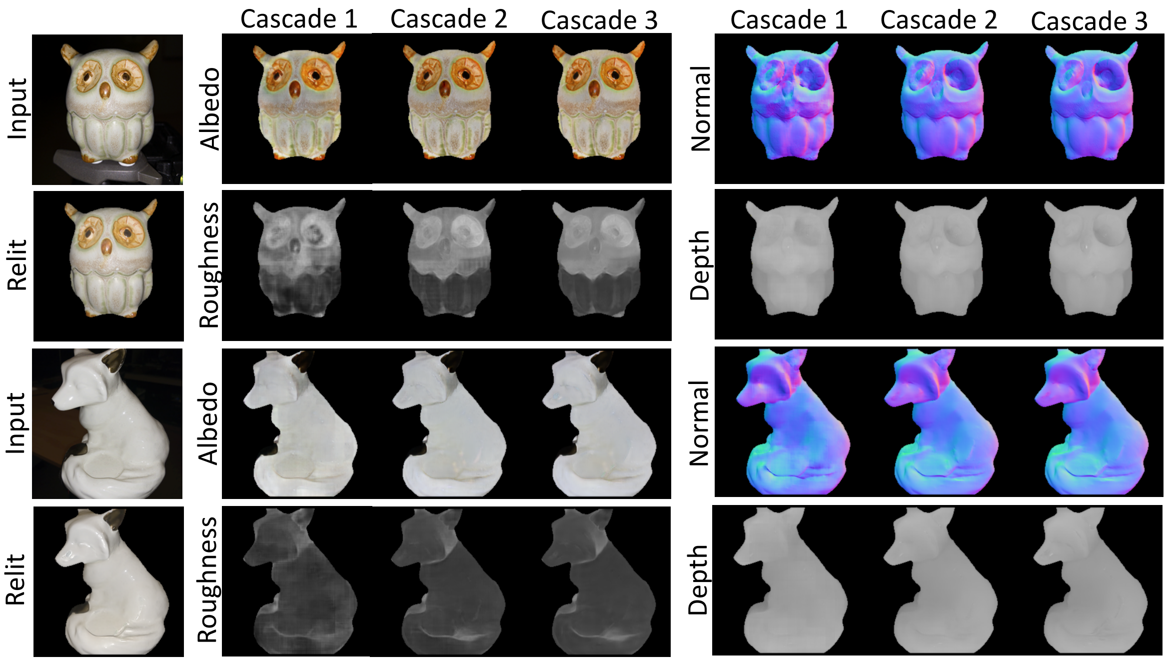

Fig. 1. Effect of cascade structure of real data. While ground truth is not available for these examples, qualitative improvements due to the cascade structure may still be clearly observed.

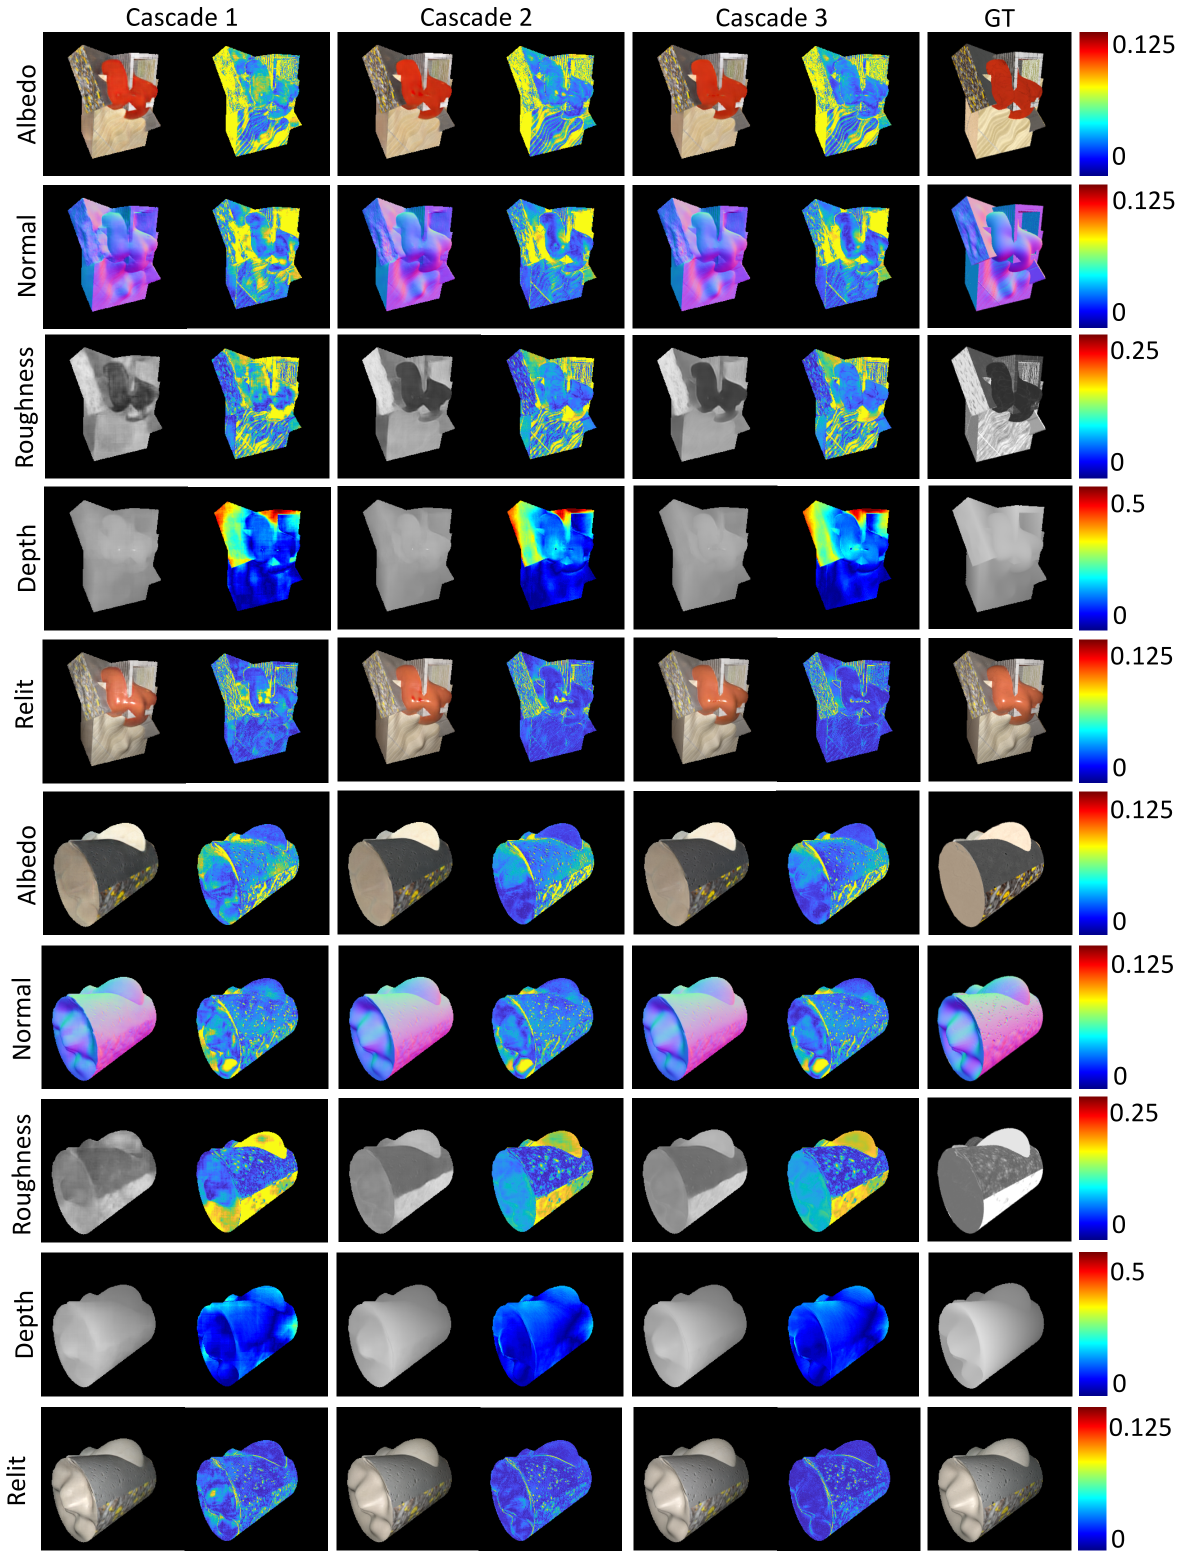
Cascade 1
Cascade 2
Cascade 3
GT
Albedo
Normal
Roughness
Depth
Relit
Albedo
Normal
Roughness
Depth
Relit
0.125
0
0.125
0
0.25
0
0.5
0
0.125
0
0.125
0
0.125
0
0.25
0
0.5
0
0.125
0

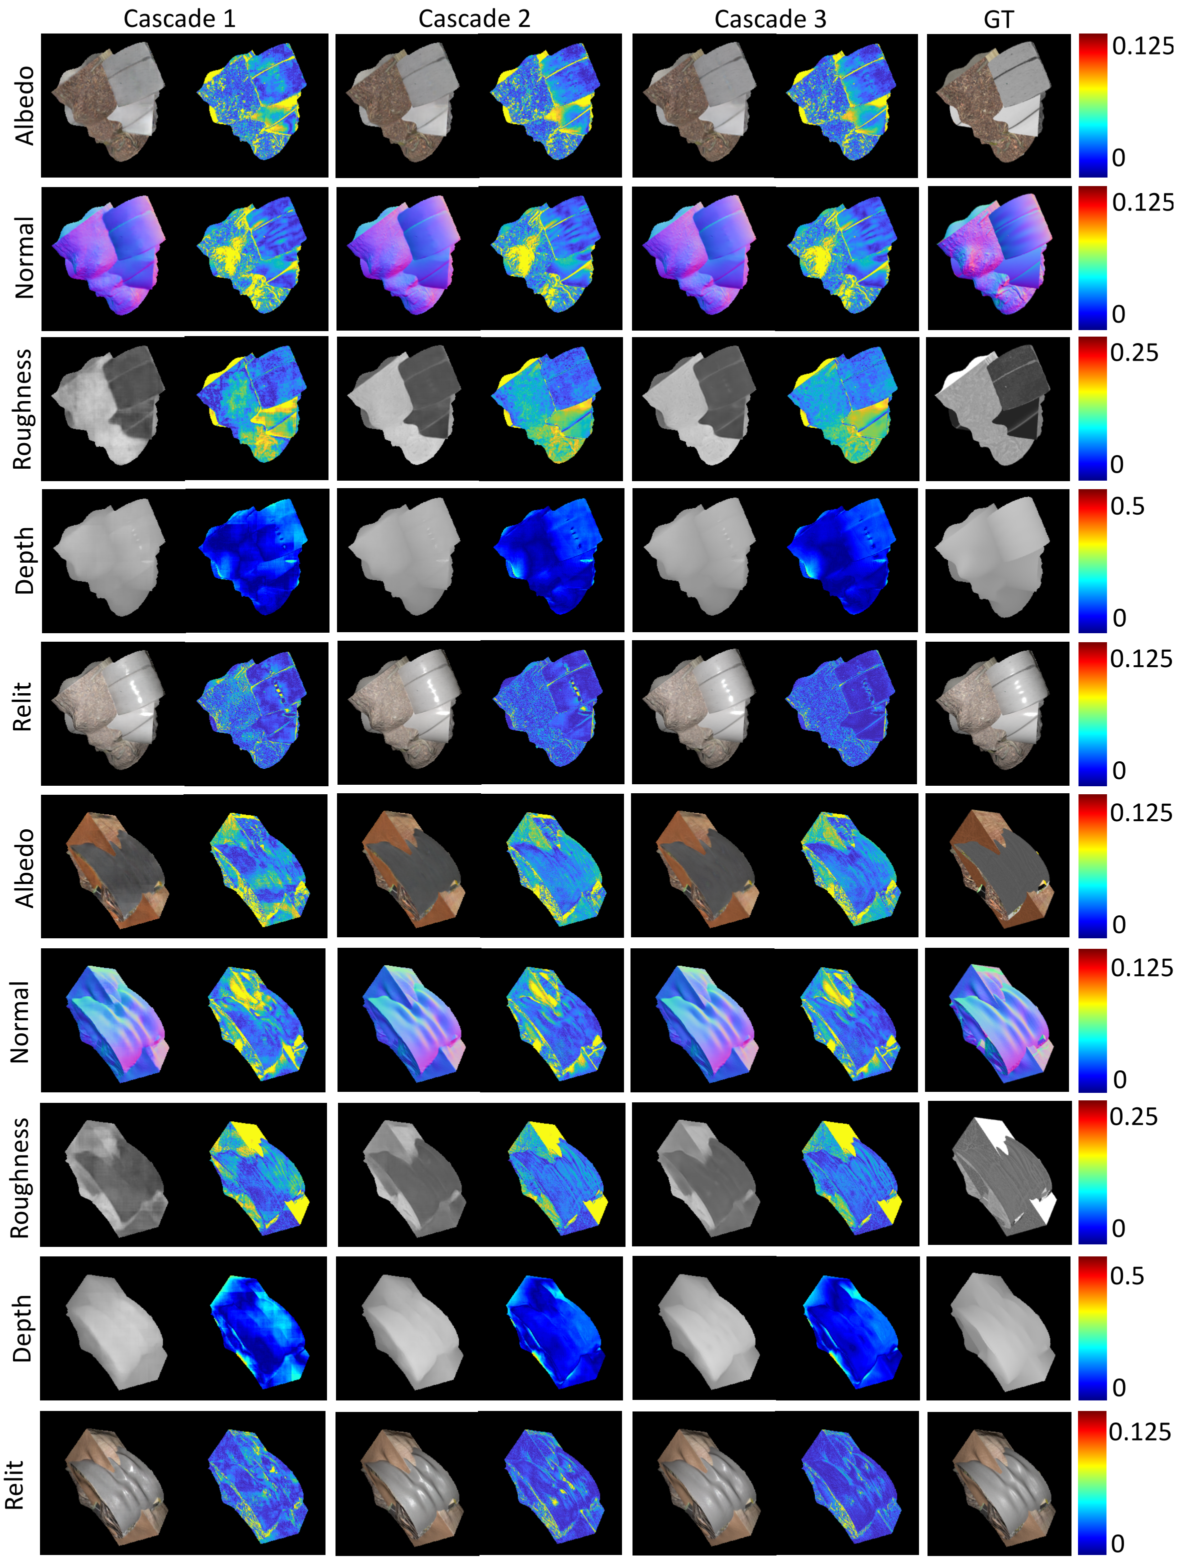

Fig. 2. Effect of cascade structure on synthetic data. We add an error visualization so that the improvements due to different cascade levels can be clearly observed. We show absolute errors for all three BRDF parameters, while the depth map is normalized such that the ground truth value is in range of 1 unit.