

6 Ablation Studies

6.1 Latent Dimension

We experimented with the dimension size of our latent representations Z on the test dataset for the reconstruction with all available samples and observed that 40-dimensional embeddings give the minimum average mean squared error (See the first row of Table 4). Therefore, we chose our latent space size to be 40. However, we also observed that these can change with the number of sparse samples. For consistency, we kept it same for all our reconstruction results.

6.2 Log Relative Mapping (LRM)

We apply Log Relative Mapping (LRM) [38] to the BRDF data before feeding it the hypernetwork, as described in Section 3.1. We illustrate that LRM improves the reconstruction quality for both PCA and hypernetwork results in Table 7 and Figure 11. Please note that IPCA is the mapping-applied version of PCA, and "Ours (No LRM)" refers to our results without LRM.

6.3 Cosine Weighting

Similar to [35], we also weigh the BRDFs with a cosine term based on the assumption of uniform incoming radiance and observed that it offers higher accuracy:

Table 3: Average metric results over the renderings of the test set (20 MERL materials).

	No Cosine	Cosine
PSNR \uparrow	33.077	33.166
Delta E \downarrow	2.233	2.117
SSIM \uparrow	0.974	0.979
MAE \downarrow	3.908	3.492
RMSE \downarrow	7.370	6.507
RAE \downarrow	0.098	0.089

6.4 Principal Component Analysis (PCA)

In Section 4.2, we briefly discussed the ablation study we ran to choose the number of principal components for the best-performing IPCA results. Here, we elaborate on the effect of the number of principal components on the performance of IPCA.

Number of Principal Components: We analyzed the optimal number of principal components by initially choosing the same numbers as we did for the ablation study of the latent dimension. Looking at Table 4, we observe that N_{PC} gives the minimum error for sparse cases, where the number of samples N is 40. We further validated our choice $N_{PC} = 8$ by running IPCA with the number of principal components ranging from 1 to 16 with an incremental change of 1. Figure 8 illustrates that $N_{PC} = 8$ still gives the minimum mean squared error (MSE) averaged over the test set.

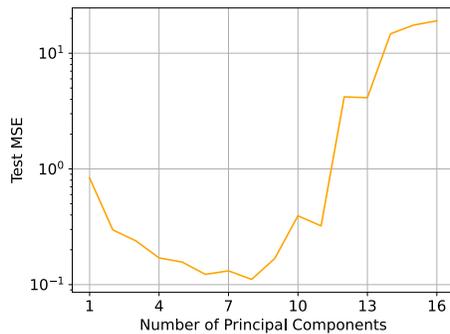


Fig. 8: IPCA optimization for the number of principal components.

Table 4: Average Mean squared errors for varying latent space dimensions (first row) and number of principal components (second row)

Methods/Dimension	8	20	30	40	50	60	70	80
Ours	0.073	0.046	0.062	0.045	0.051	0.053	0.049	0.046
IPCA	0.131	0.278	12.288	0.507	0.878	0.883	0.965	1.124

6.5 GGX Fittings

We also fitted an analytic model, GGX [52], to the sparse number of measured BRDF samples using a non-linear optimisation (L-BFGS-B method) with Log l_1 loss (same loss used in NBRDF). The average PSNR values across our test dataset for varying sample sizes can be found in Table 5. For the results with all samples, we used the dj_brdf mitsuba plug-in [11]. Rendering results are also included in this supplementary ("ggx.pdf").

Table 5: GGX - Average metric results over the renderings of our test dataset. We highlight **best** results.

N	PSNR \uparrow	Delta E \downarrow	SSIM \uparrow	MAE \downarrow	RMSE \downarrow	RAE \downarrow
40	29.198	2.821	0.924	6.724	12.778	0.141
160	28.893	2.831	0.923	6.906	13.004	0.147
400	29.193	2.771	0.924	6.722	12.798	0.141
2000	29.040	2.794	0.923	6.795	12.919	0.143
4000	29.035	2.792	0.923	6.786	12.914	0.143
All samples	30.334	2.241	0.967	4.784	8.787	0.123

7 Additional Results

7.1 Sparse Reconstruction

In addition to Figure 5, we provide MAE, RMSE, and RAE metric plots (Figure 9) along with the hypernetwork’s average metric results across varying sample sizes (Table 1).

7.2 Full Reconstruction

We first evaluated the full reconstruction capacity of our method by using all available samples during the inference time. Figure 11 shows our results without/with LRM (Sec. 6.2) along with the results for Hyper autoencoder (Hyper-AE) [48], PCA, and IPCA. For all methods, we do the same train/test split and compare the methods both qualitatively (Figure 11) and quantitatively (Table 7) over the test set. Please note that Hyper-AE is not included in the results for reconstruction with sparse samples since it only works with a fixed grid of input samples, unable to reconstruct materials with varying sizes.

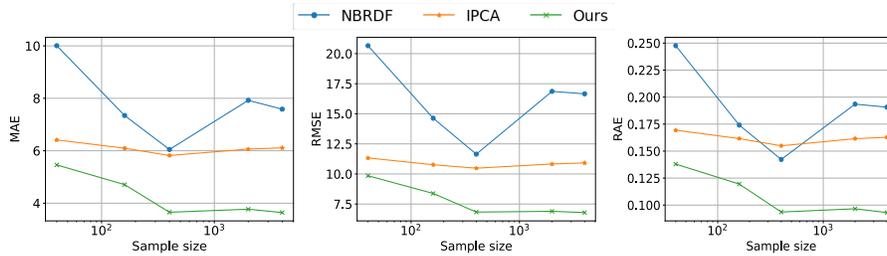
Figure 11 shows that our method attains superior results in terms of the reconstruction quality of measured BRDFs. The reconstructed materials with PCA or Hyper-AE cannot capture the material properties, such as the diffuse components, effectively. They can also cause some artifacts (chrome, green-metallic-paint). The linear structure of PCA cannot handle the high dynamic range of the data. Although IPCA [38] improves the results substantially, they still have difficulty capturing the right colors. The NBRDF [48] reconstruction results can be found in their supplementary. Since our hypernetwork model offers a generalizable approach to NBRDF, it is expected that an overfitting NBRDF model performs better in reconstruction with full samples.

7.3 Compression

We report the rendering results of 100 MERL materials for compression (Table 2) as a pdf file; see "Compression.pdf".

Table 6: Hypernetwork - Average metric results across varying sample sizes (N) over the test set (Sparse and full reconstruction of unseen materials).

Metrics/ N	8	40	4000	40 000	640 000	1 458 000
PSNR \uparrow	23.090	29.822	33.170	33.019	33.166	33.128
Delta E \downarrow	5.948	3.086	2.138	2.118	2.117	2.181
SSIM \uparrow	0.927	0.969	0.977	0.979	0.979	0.978
MAE \downarrow	13.167	5.306	3.642	3.568	3.492	3.574
RMSE \downarrow	22.293	9.601	6.791	6.618	6.507	6.740
RAE \downarrow	0.339	0.134	0.093	0.091	0.089	0.091

**Fig. 9:** Average MAE, RMSE, and RAE results across different sample sizes.**Table 7:** Quantitative comparison results for full reconstruction over the renderings of the test set.

	Hyper-AE	PCA	IPCA	Ours (No LRM)	Ours
PSNR \uparrow	21.696	16.407	28.471	20.696	33.128
Delta E \downarrow	7.761	13.233	4.023	8.970	2.181
SSIM \uparrow	0.873	0.818	0.975	0.896	0.978
MAE \downarrow	13.058	30.949	5.940	16.927	3.574
RMSE \downarrow	23.051	47.018	10.276	28.314	6.740
RAE \downarrow	0.329	0.780	0.158	0.421	0.091

7.4 Scene Renderings

We also rendered multiple scenes [5] using our reconstructed materials, including sparse reconstruction, compression and interpolation. The details about the objects covered with our reconstructed materials are as follows (Figure 10): **teapot**: steel. **dragons**: interpolation of delrin and green-metallic-paint, white marble (dragon paint), silver-metallic-paint3 (ground), dark-red-paint (cloth). **cars**: chrome steel, gold metallic paint3, specular-red-phenolic (car paint), exterior car parts (aluminium), inner wheel (alumina-oxide). **kitchen**: cupboards (natural-209), utensils (chrome), handles, pot, microwave and cooker (two-layer-silver), extractor hood (aluminium), pot and kettle handles (black-obsidian), kettle paint (dark-red-paint), tea towel and cushions (yellow-paint), radio and lamp (polyethylene). **living room**: sofa, coffee table, side tables, wall book shelf (pure-rubber), cushions (green-metallic-paint), twigs (natural-209), legs (dark-specular-fabric). **sofas**: violet rubber and green latex (sofa cover).

Renderings can be found in png format under the "scene-renderings" folder.

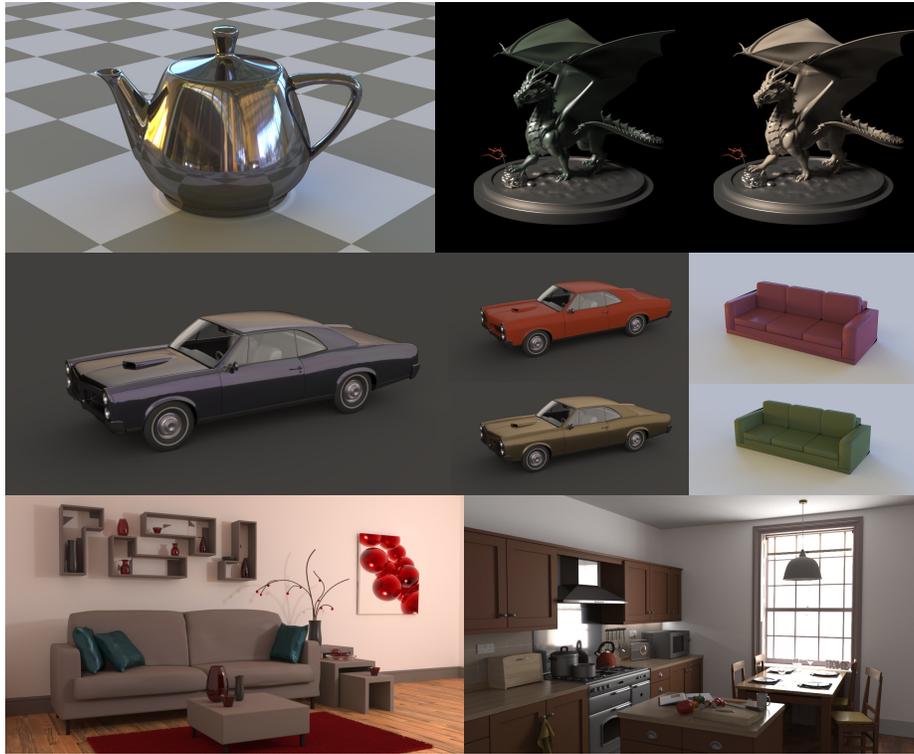


Fig. 10: Scene renderings with our reconstructed materials.

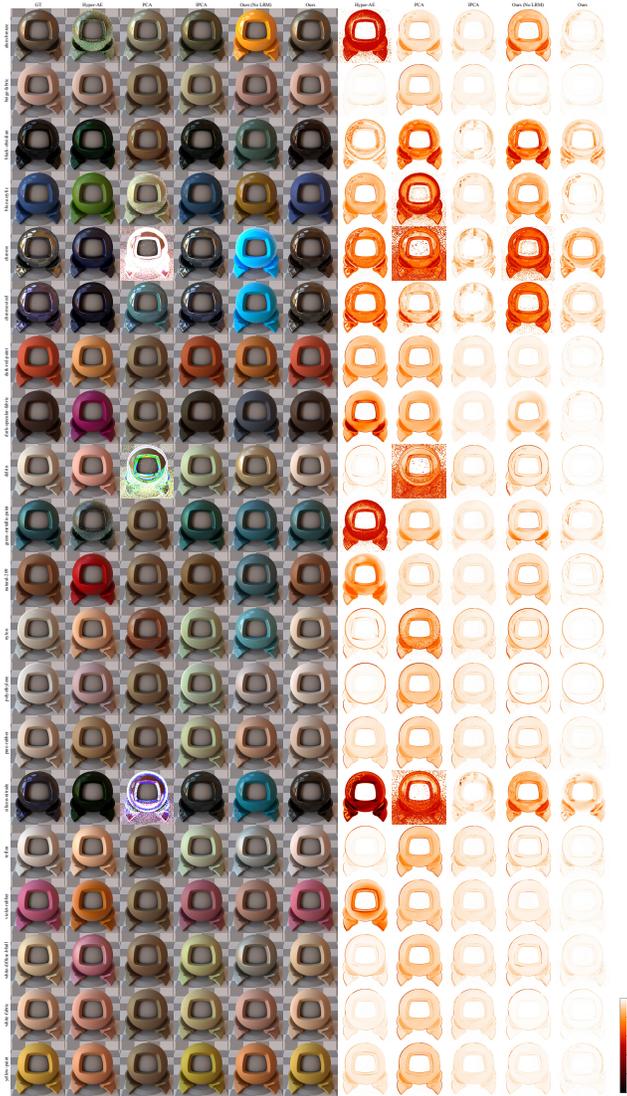


Fig. 11: Qualitative comparison for full reconstruction capacity on the test dataset.

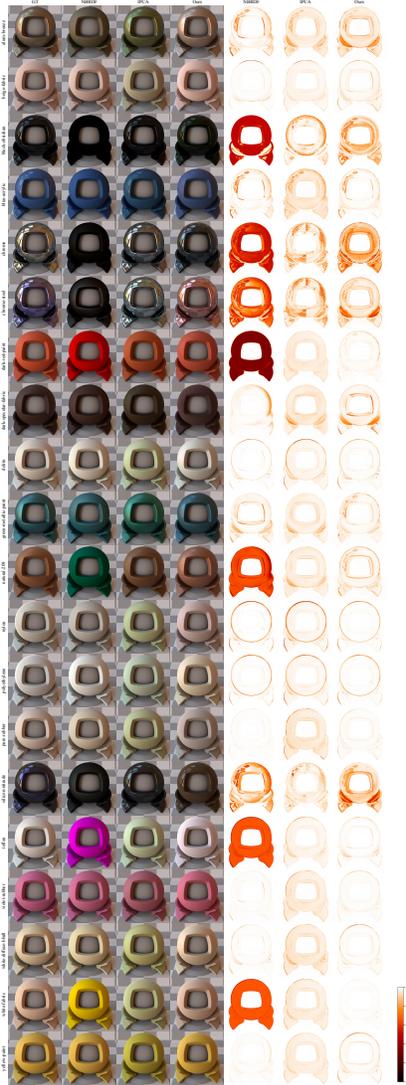


Fig. 12: Sparse reconstruction results, $N = 40$.



Fig. 13: Sparse reconstruction results, $N = 160$.

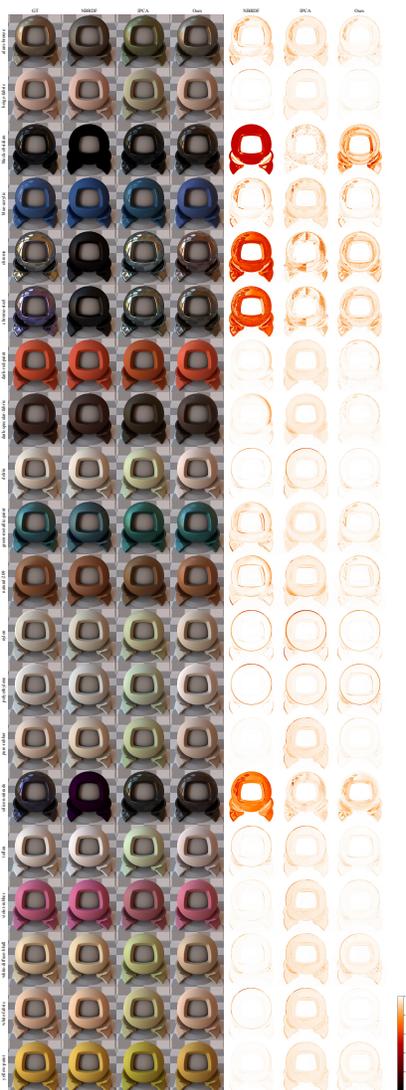


Fig. 14: Sparse reconstruction results, $N = 400$.

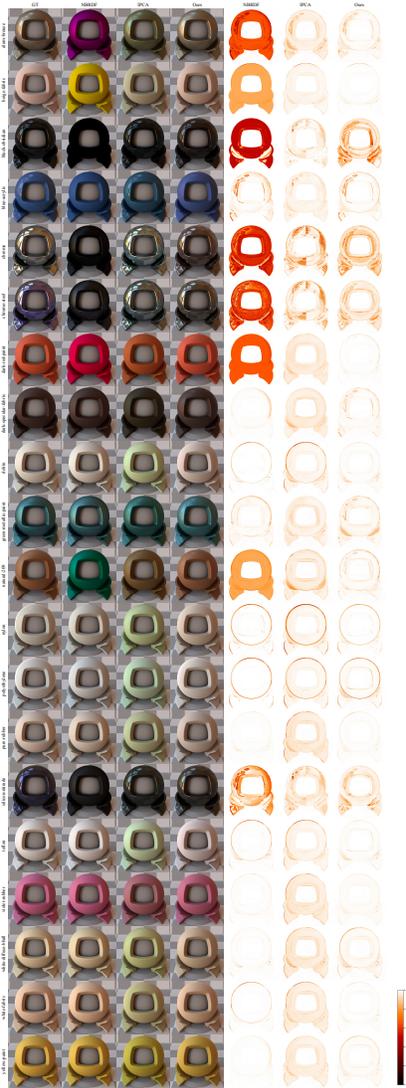


Fig. 15: Sparse reconstruction results, $N = 2000$.

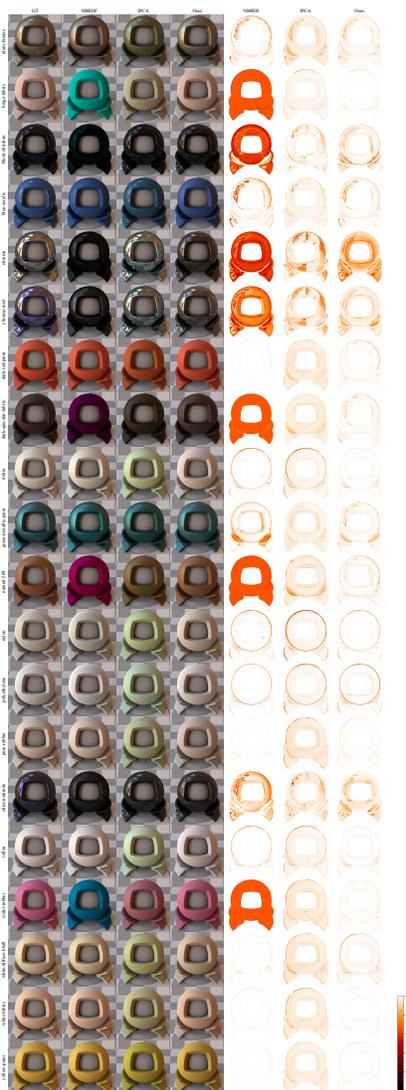


Fig. 16: Sparse reconstruction results, $N = 4000$.