Real-Time Neural Appearance Models — Supplemental material

TIZIAN ZELTNER*, FABRICE ROUSSELLE*, ANDREA WEIDLICH*, PETRIK CLARBERG*, JAN NOVÁK*, BENEDIKT BITTERLI*, ALEX EVANS, TOMÁŠ DAVIDOVIČ, SIMON KALLWEIT, and AARON LEFOHN, NVIDIA, Global

Contents

1 TRAINING PROCEDURE DETAILS

The following list summarizes the parameters used for training the neural materials shown in the main paper. All our training takes place in PyTorch [Paszke et al. 2019].

Optimizer. We use the ADAM optimizer [Kingma and Ba 2014] ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e{-}7$, and zero weight decay) starting with a learning rate of 1e−3 that is progressively reduced based on a cosine decay until it reaches 1e−4. I.e. it is scaled by

$$\text{learning rate scale} = 0.1 + 0.5 * (1 - 0.1) * \left(1 + \cos\left(\frac{i\,\pi}{N}\right)\right), \tag{1}$$

where i and N are the current and maximum iteration count.

Mollification. Over the first $M = 20000$ iterations of training, the angle of our *mollification* cone is reduced to from 10 to 0 degrees with another cosine decay, i.e.

$$\text{mollification cone angle} = 0.5 * 10^\circ * \left(1 + \cos\left(\frac{i\,\pi}{M}\right)\right). \tag{2}$$

We use 256 sampled $\boldsymbol{\omega}_\text{o}$ directions within the cone in order to estimate the blurred version of the BRDF.

*Equal contribution. Order determined by a rock-paper-scissors tournament.

Authors' address: Tizian Zeltner; Fabrice Rousselle; Andrea Weidlich; Petrik Clarberg; Jan Novák; Benedikt Bitterli; Alex Evans; Tomáš Davidovič; Simon Kallweit; Aaron Lefohn, NVIDIA, Global.

2 REFERENCE MATERIALS SHADER NODE GRAPHS

The following MaterialX node graphs were used to design the reference materials and illustrate their complexity. The initial setup was created in Houdini whereas some nodes would be later replaced by custom nodes. Please note that since MaterialX is not aware of the concept of layer weights, all BRDFs are mixed with transparent nodes to account for this. In practice we do not evaluate such BRDFs.

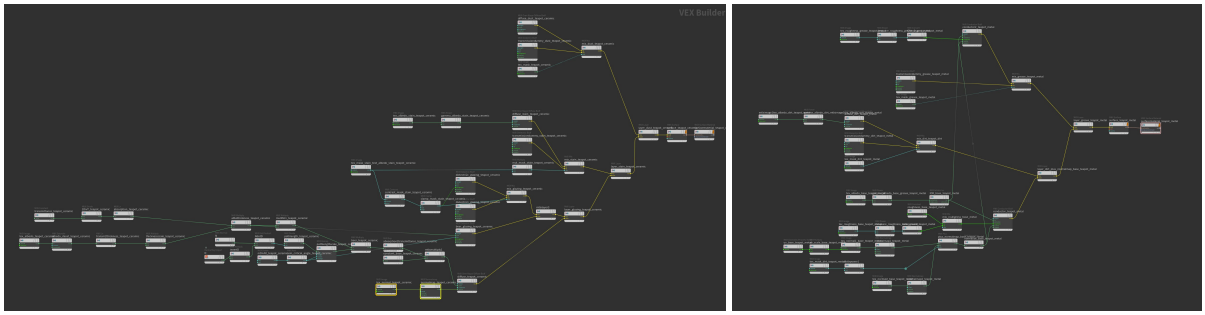

Fig. 1. MaterialX shader graph of the Teapot. Left: Ceramic glazing. Right: Metallic handle.

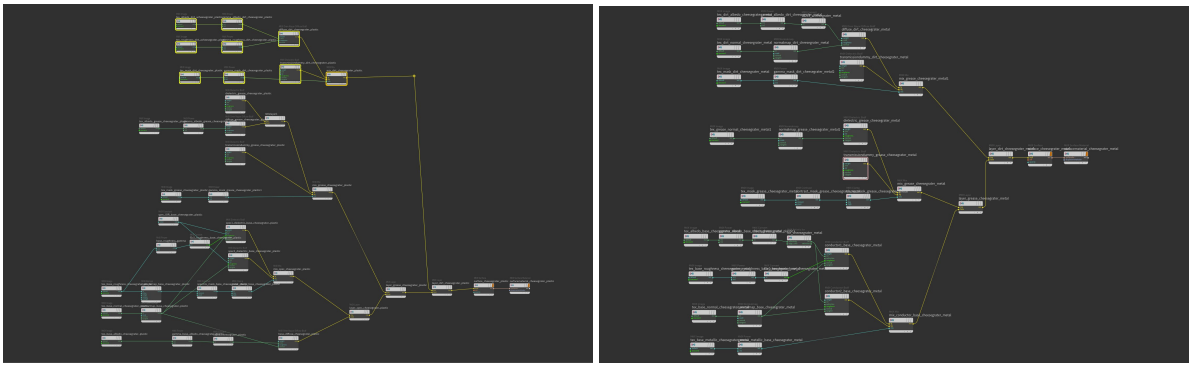

Fig. 2. MaterialX shader graph of the Cheese slicer. Left: Plastic handle. Right: Metallic blade.

Fig. 3. MaterialX shader graph of the Inkwell.

3 IMAGE ERROR METRICS

3.1 Ablation

Table 1. Image error metrics for the ablation study on the Inkwell metal body material (first row of Figure 7 in the main paper).

	Vanilla MLP	w/ encoder	w/ frame transform
Mean ꟻLIP	0.2589	0.1921	0.0703
Mean abs. error	0.0403	0.0284	0.0068
Mean sqr. error	0.0055	0.0029	0.0002
Mean rel. abs. error	0.2269	0.1443	0.0397
Mean rel. sqr. error	0.0676	0.0334	0.0043
SMAPE	0.2316	0.1700	0.0412

Table 2. Image error metrics for the ablation study on the Teapot ceramic body material (second row of Figure 7 in the main paper).

	Vanilla MLP	w/ encoder	w/ frame transform
Mean ꟻLIP	0.3343	0.2552	0.1362
Mean abs. error	0.1311	0.0979	0.0331
Mean sqr. error	0.1112	0.5560	0.0182
Mean rel. abs. error	0.2956	0.2767	0.1089
Mean rel. sqr. error	0.1369	0.1140	0.0176
SMAPE	0.3989	0.3999	0.1230

Table 3. Image error metrics for the ablation study on the Cheese slicer blade material (third row of Figure 7 in the main paper).

	Vanilla MLP	w/ encoder	w/ frame transform
Mean ꟻLIP	0.1381	0.1268	0.0774
Mean abs. error	0.0769	0.0522	0.0251
Mean sqr. error	0.0954	0.0402	0.0038
Mean rel. abs. error	0.1203	0.1027	0.0618
Mean rel. sqr. error	0.0444	0.0286	0.0107
SMAPE	0.1446	0.1183	0.0677

Table 4. Image error metrics for the ablation study on the Cheese slicer handle material (fourth row of Figure 7 in the main paper).

	Vanilla MLP	w/ encoder	w/ frame transform
Mean ꟻLIP	0.2245	0.2083	0.0423
Mean abs. error	0.0591	0.0822	0.0080
Mean sqr. error	0.0607	40.1738	0.0004
Mean rel. abs. error	0.2278	0.8520	0.0519
Mean rel. sqr. error	0.0704	1061.4310	0.0033
SMAPE	0.2931	0.2707	0.0532

3.2 Model quality

Table 5. Image error metrics for view 1 in the Inkwell scene (Figure 14 in the main paper).

	2×16	2×32	3×64
Mean ꟻLIP	0.0303	0.0245	0.0244
Mean abs. error	0.0033	0.0025	0.0025
Mean sqr. error	0.0001	0.0001	0.0021
Mean rel. abs. error	0.0200	0.0150	0.0140
Mean rel. sqr. error	0.0012	0.0019	0.0008
SMAPE	0.0206	0.0154	0.0145

Table 6. Image error metrics for view 2 in the Inkwell scene (Figure 14 in the main paper).

	2×16	2×32	3×64
Mean ꟻLIP	0.0617	0.0453	0.0430
Mean abs. error	0.0081	0.0054	0.0049
Mean sqr. error	0.0003	0.0002	0.0001
Mean rel. abs. error	0.0394	0.0274	0.0244
Mean rel. sqr. error	0.0030	0.0017	0.0014
SMAPE	0.0404	0.0282	0.0253

Table 7. Image error metrics for view 1 in the Stage scene (Figure 15 in the main paper).

	2×16	2×32	3×64
Mean ꟻLIP	0.0967	0.0430	0.0321
Mean abs. error	0.0402	0.0139	0.0116
Mean sqr. error	1.0302	0.0110	0.0057
Mean rel. abs. error	0.0909	0.0317	0.0247
Mean rel. sqr. error	0.0367	0.0044	0.0029
SMAPE	0.1338	0.0347	0.0259

Table 8. Image error metrics for view 2 in the Stage scene (Figure 15 in the main paper).

	2×16	2×32	3×64
Mean ꟻLIP	0.1658	0.0581	0.0454
Mean abs. error	0.1144	0.0165	0.0144
Mean sqr. error	8.1037	0.0436	0.0456
Mean rel. abs. error	0.1801	0.0466	0.0375
Mean rel. sqr. error	0.0651	0.0054	0.0037
SMAPE	0.2676	0.0514	0.0394

Table 9. Image error metrics for view 3 in the Stage scene (Figure 15 in the main paper).

	2×16	2×32	3×64
Mean ꟻLIP	0.2771	0.1363	0.0985
Mean abs. error	0.0971	0.0331	0.0242
Mean sqr. error	0.5584	0.0176	0.0152
Mean rel. abs. error	0.2929	0.1089	0.0805
Mean rel. sqr. error	0.1210	0.0175	0.0097
SMAPE	0.4325	0.1231	0.0855

Table 10. Image error metrics for view 4 in the Stage scene (Figure 15 in the main paper).

	2×16	2×32	3×64
Mean ꟻLIP	0.0895	0.0480	0.0392
Mean abs. error	0.0330	0.0200	0.0174
Mean sqr. error	0.0048	0.0019	0.0016
Mean rel. abs. error	0.0710	0.0401	0.0319
Mean rel. sqr. error	0.0147	0.0054	0.0032
SMAPE	0.0811	0.0429	0.0332

Table 11. Image error metrics for view 5 in the Stage scene (Figure 15 in the main paper).

	2×16	2×32	3×64
Mean ꟻLIP	0.0400	0.0302	0.0280
Mean abs. error	0.0114	0.0099	0.0095
Mean sqr. error	0.0011	0.0007	0.0006
Mean rel. abs. error	0.0349	0.0307	0.0296
Mean rel. sqr. error	0.0054	0.0032	0.0030
SMAPE	0.0383	0.0317	0.0302

4 MODEL EVALUATION PLOTS

4.1 MERL database

We additionally trained a neural material on the complete MERL database [Matusik et al. 2003] including 100 different BRDFs. As these are obtains from real-world measurements we do not have access to suitable surface parameters (albedo, roughness, etc.) that could be fed into our encoder. Instead, the encoder receives the material index (1–100) via one-hot encoding. Finally, all neural variants of the MERL materials can be evaluated by the combination of i) a unique eight-dimensional latent code and ii) the decoder MLP that is shared between all 100 BRDFs. We show (log-scale) polar plots for our evaluation and importance sampling network evaluations in Figure 4 and Figure 5.

Fig. 4. MERL database: evaluation network (3 layers w/ 64 neurons) plots.

Fig. 5. MERL database: importance sampling network plots.

4.2 High-fidelity neural material plots

The remainder of this document contains a set of similar (log-scale) polar plots for the neural materials discussed in the main paper. Each BRDF subplot corresponds to a fixed spatial location (obtained via regularly sampling UV space). These additionally visualize any unqique per-layer surface normal that deviates from the canonical ($[0, 0, 1]$) z-axis. PDF plots are scaled to compare against the learned BRDF up to a constant normalization factor.

Teapot ceramic body (3 layers w/ 64 neurons)

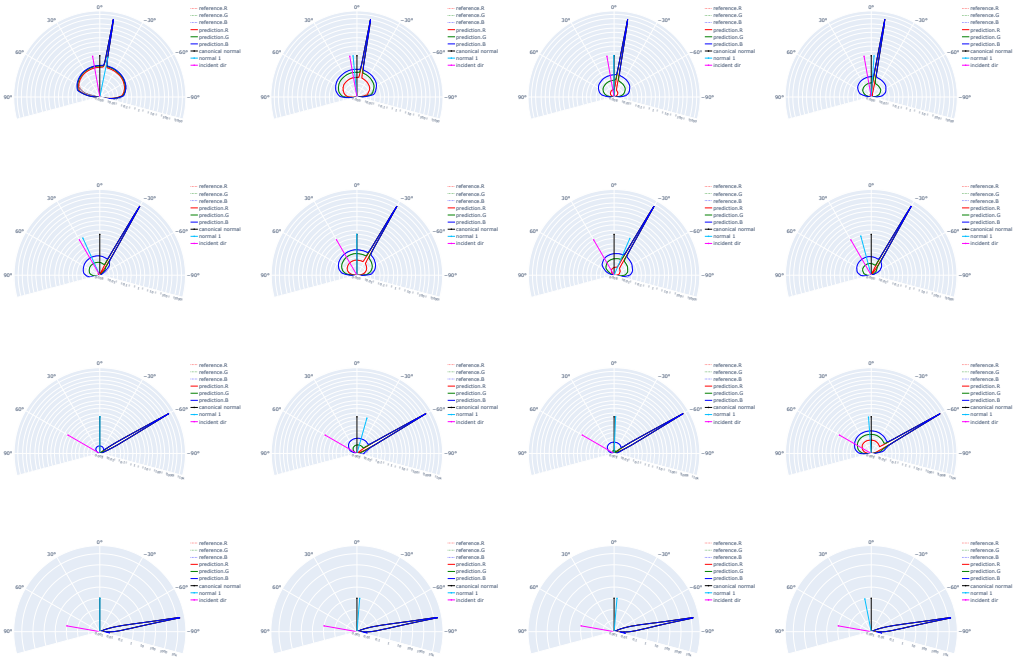

Fig. 6. Teapot ceramic body: evaluation network (3 layers w/ 64 neurons) plots.

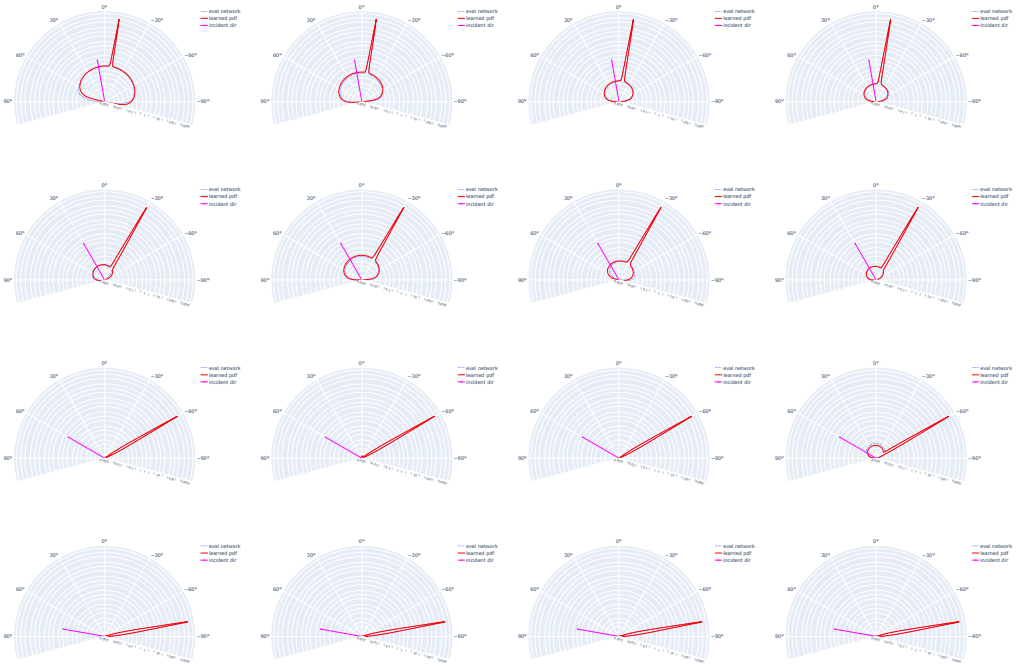

Fig. 7. Teapot ceramic body: importance sampling network plots.

Teapot ceramic body (2 layers w/ 32 neurons)

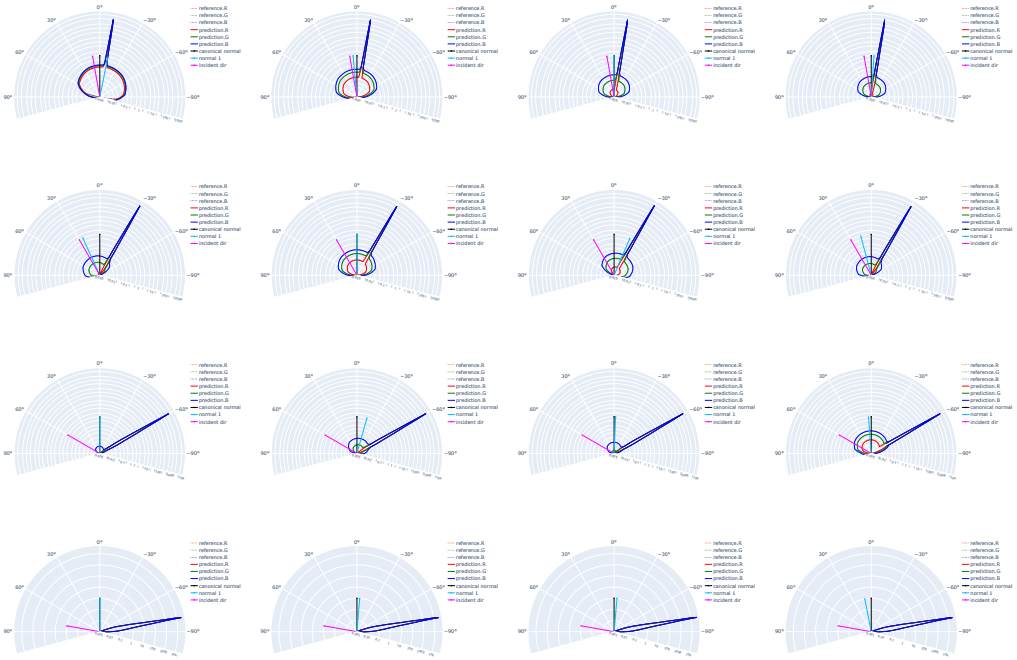

Fig. 8. Teapot ceramic body: evaluation network (2 layers w/ 32 neurons) plots.

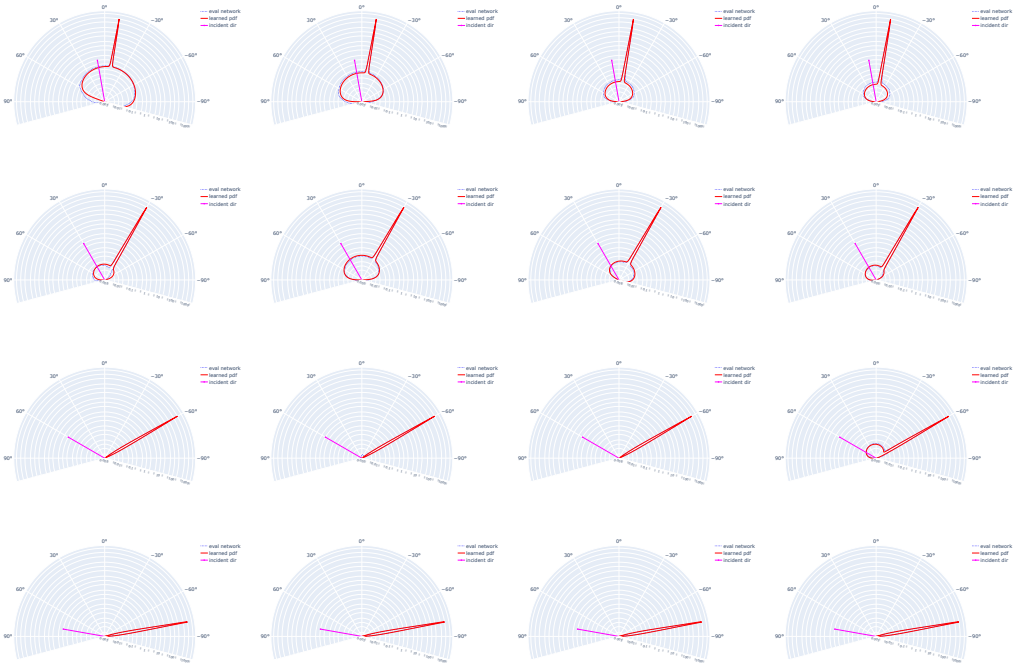

Fig. 9. Teapot ceramic body: importance sampling network plots.

Teapot ceramic body (2 layers w/ 16 neurons)

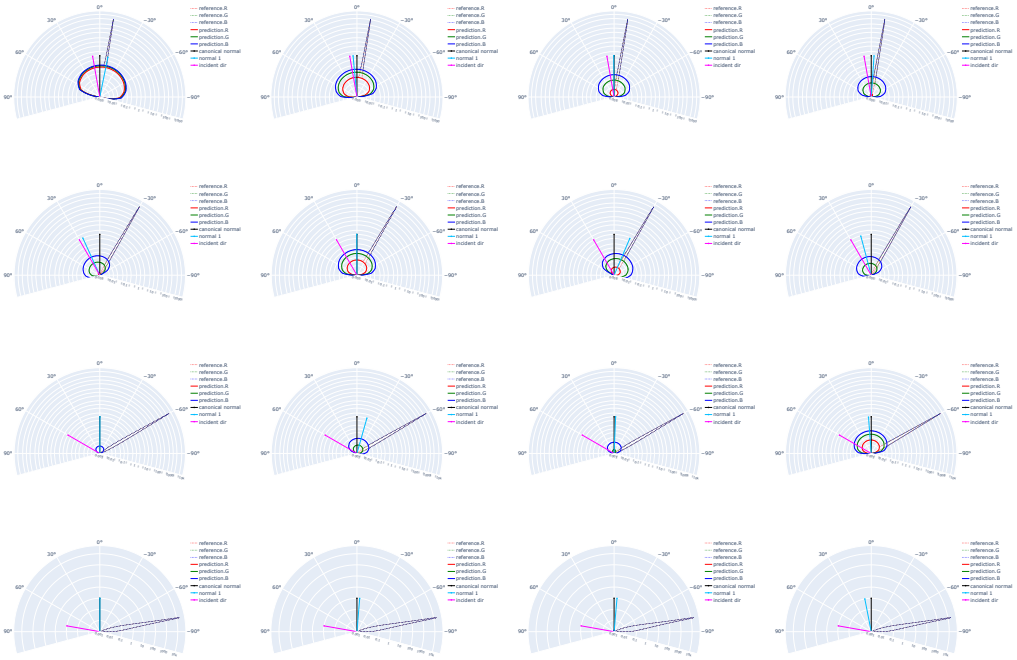

Fig. 10. Teapot ceramic body: evaluation network (2 layers w/ 16 neurons) plots.

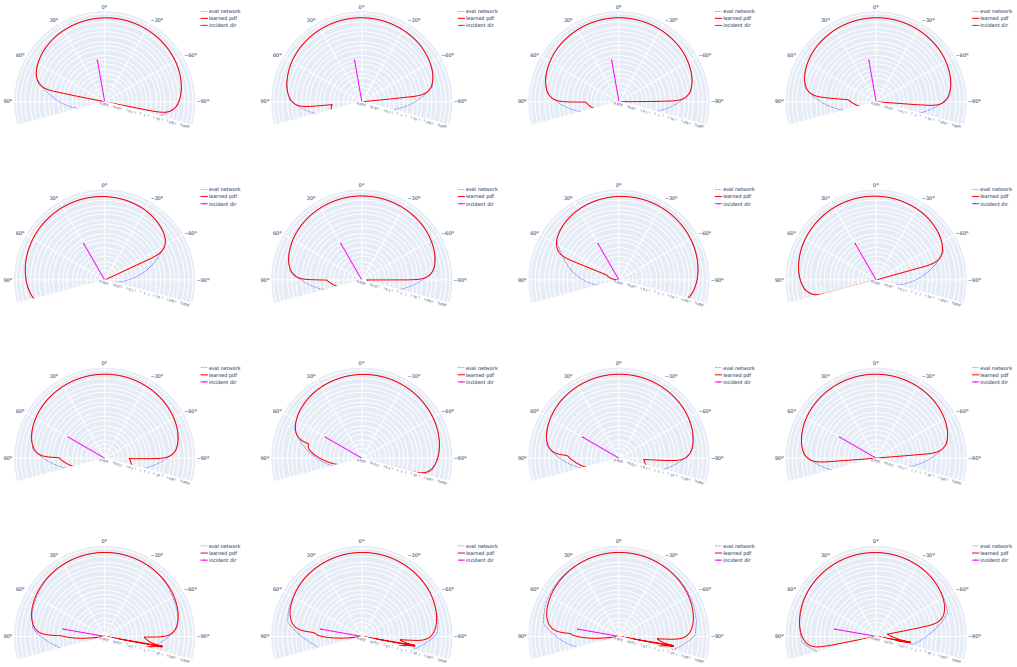

Fig. 11. Teapot ceramic body: importance sampling network plots.

Teapot metal handle (3 layers w/ 64 neurons)

Fig. 12. Teapot metal handle: evaluation network (3 layers w/ 64 neurons) plots.

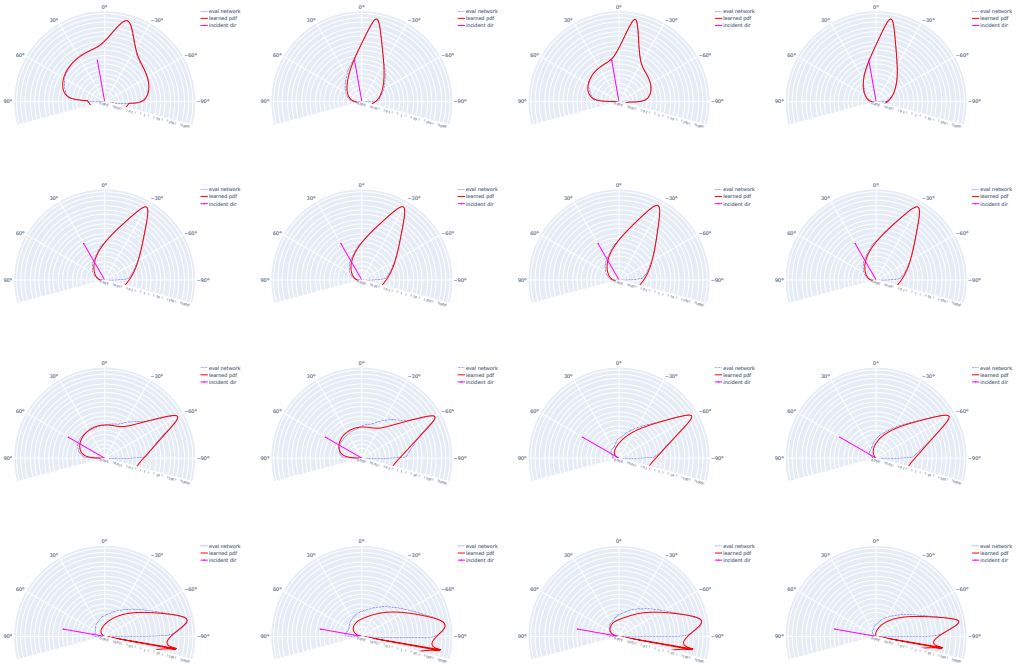

Fig. 13. Teapot metal handle: importance sampling network plots.

Teapot metal handle (2 layers w/ 32 neurons)

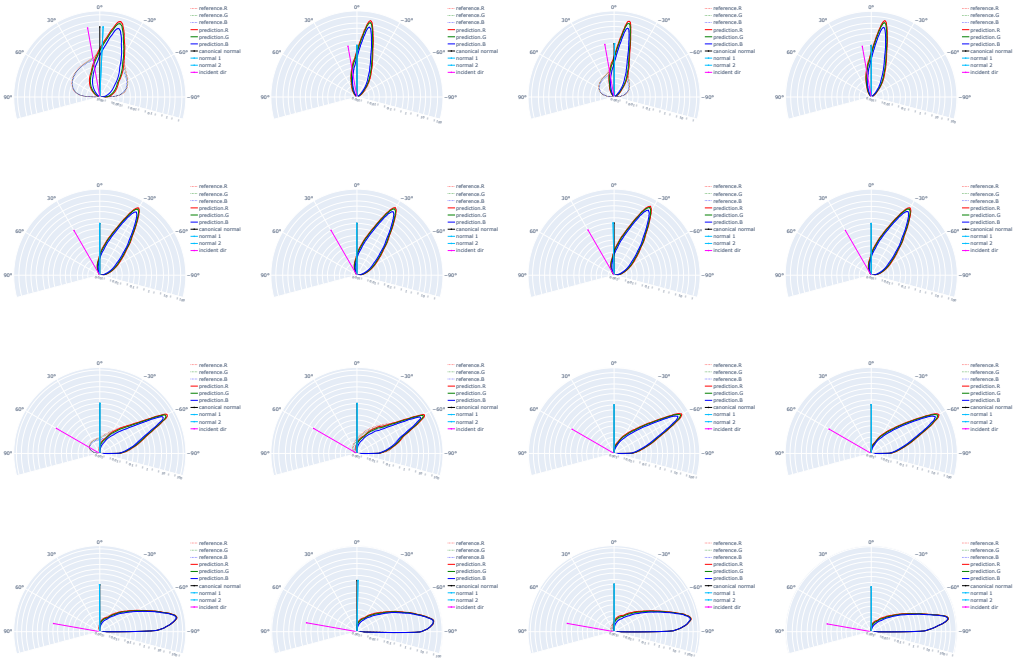

Fig. 14. Teapot metal handle: evaluation network (2 layers w/ 32 neurons) plots.

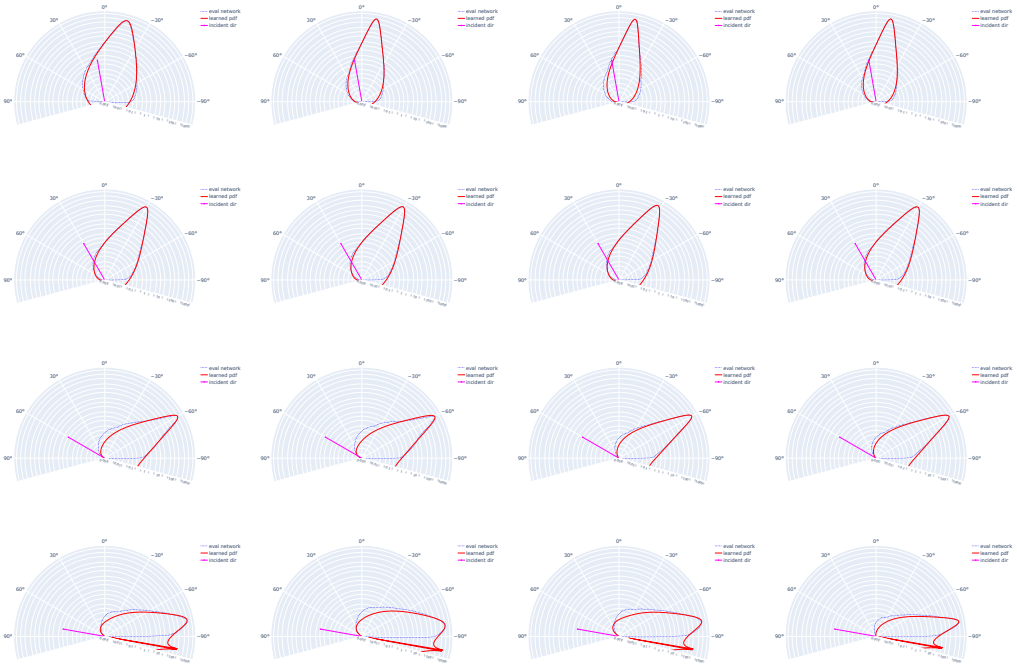

Fig. 15. Teapot metal handle: importance sampling network plots.

Teapot metal handle (2 layers w/ 16 neurons)

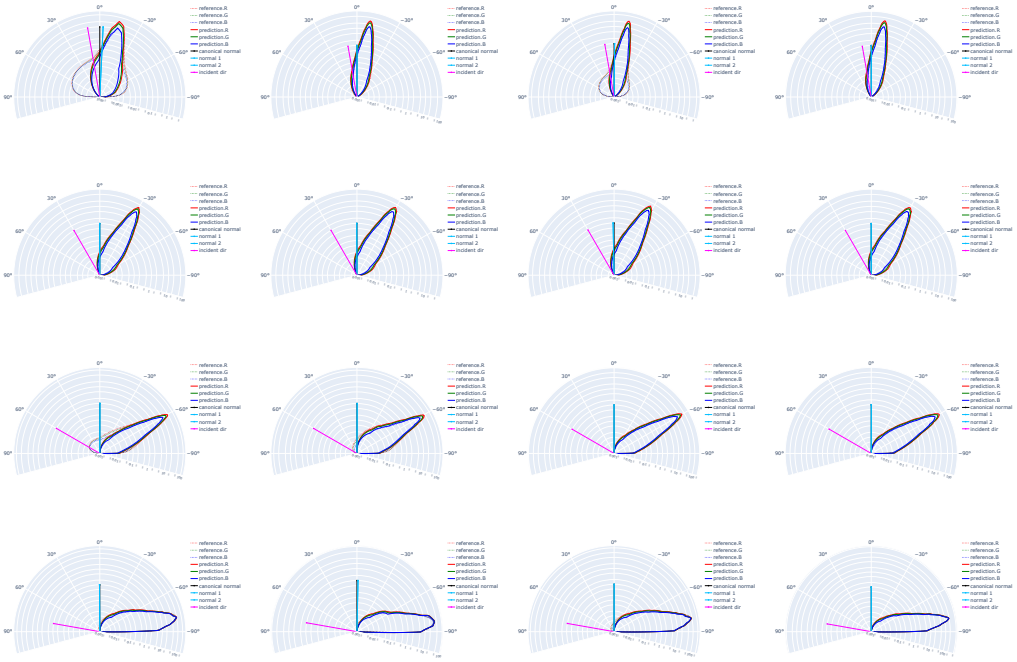

Fig. 16. Teapot metal handle: evaluation network (2 layers w/ 16 neurons) plots.

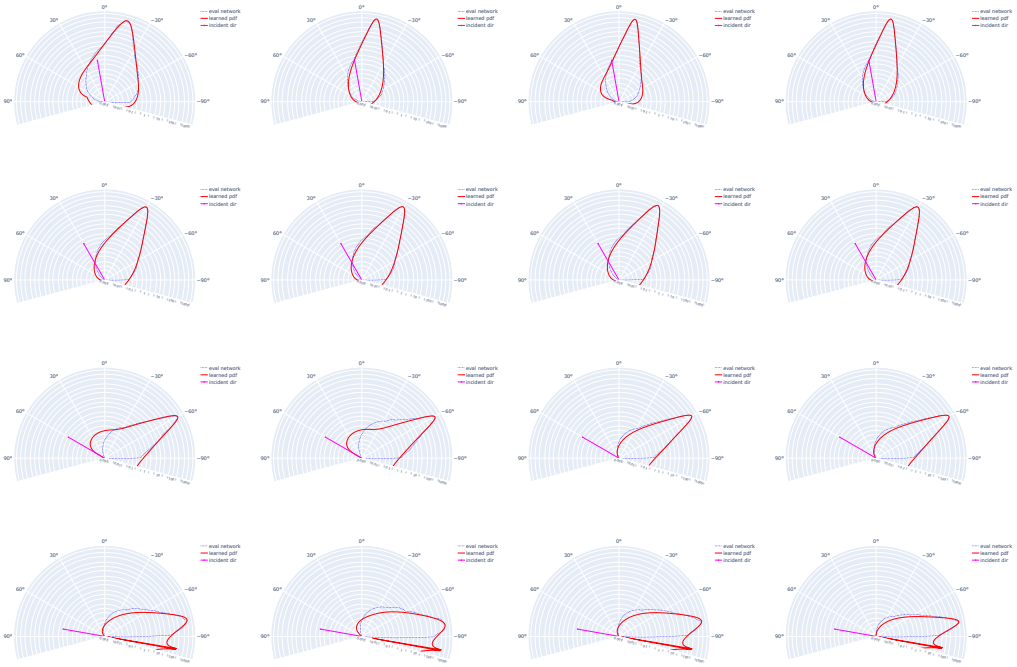

Fig. 17. Teapot metal handle: importance sampling network plots.

Cheese slicer blade(3 layers w/ 64 neurons)

Fig. 18. Cheese slicer blade: evaluation network (3 layers w/ 64 neurons) plots.

Fig. 19. Cheese slicer blade: importance sampling network plots.

Cheese slicer blade(2 layers w/ 32 neurons)

Fig. 20. Cheese slicer blade: evaluation network (2 layers w/ 32 neurons) plots.

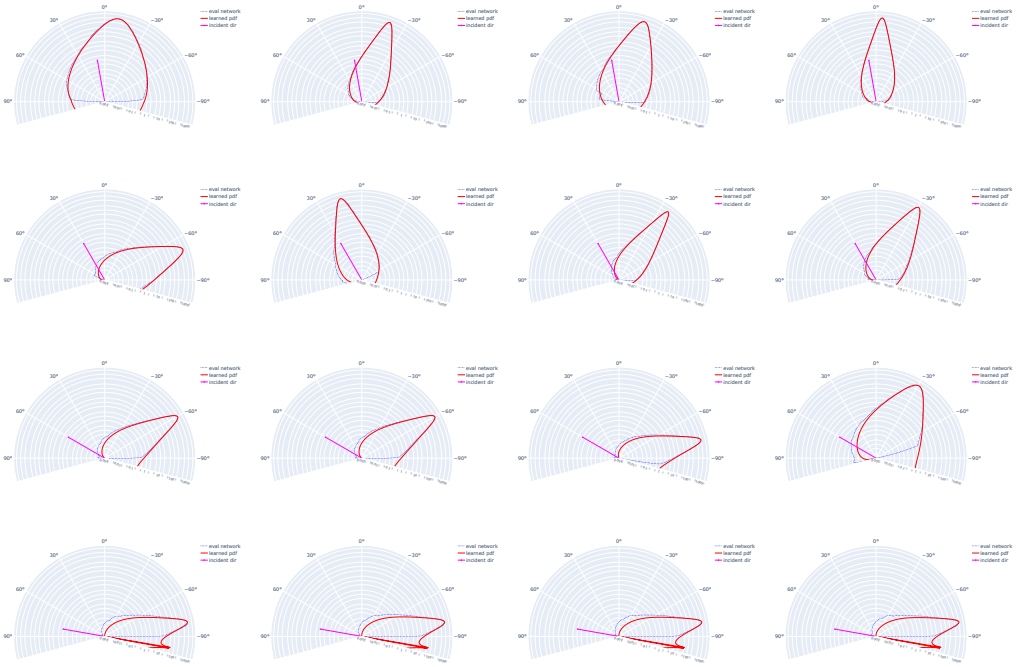

Fig. 21. Cheese slicer blade: importance sampling network plots.

Cheese slicer blade(2 layers w/ 16 neurons)

Fig. 22. Cheese slicer blade: evaluation network (2 layers w/ 16 neurons) plots.

Fig. 23. Cheese slicer blade: importance sampling network plots.

Cheese slicer handle(3 layers w/ 64 neurons)

Fig. 24. Cheese slicer handle: evaluation network (3 layers w/ 64 neurons) plots.

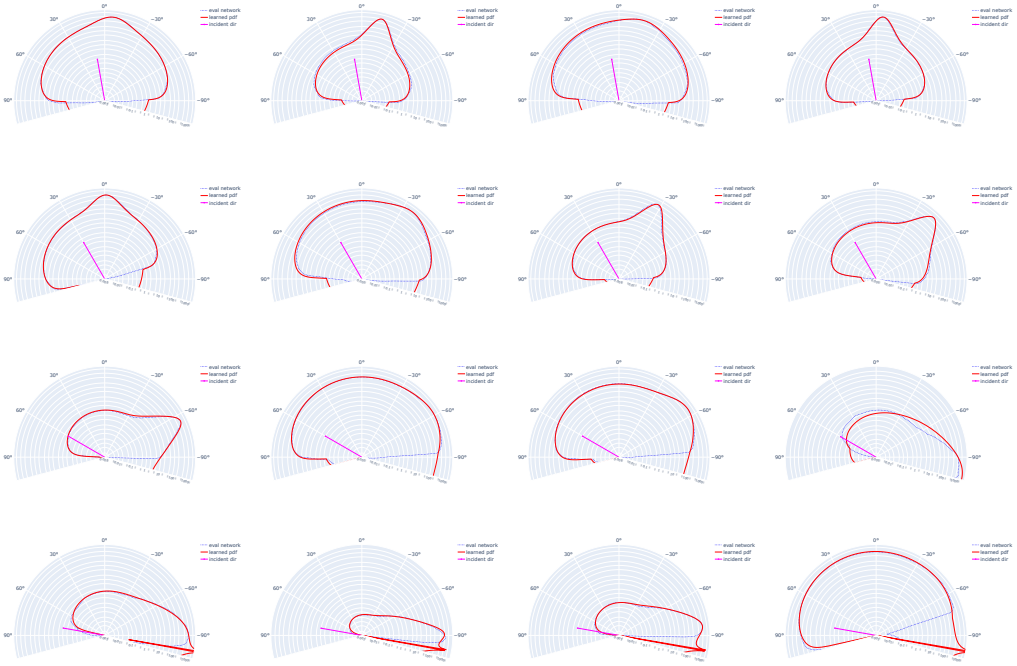

Fig. 25. Cheese slicer handle: importance sampling network plots.

Cheese slicer handle(2 layers w/ 32 neurons)

Fig. 26. Cheese slicer handle: evaluation network (2 layers w/ 32 neurons) plots.

Fig. 27. Cheese slicer handle: importance sampling network plots.

Cheese slicer handle(2 layers w/ 16 neurons)

Fig. 28. Cheese slicer handle: evaluation network (2 layers w/ 16 neurons) plots.

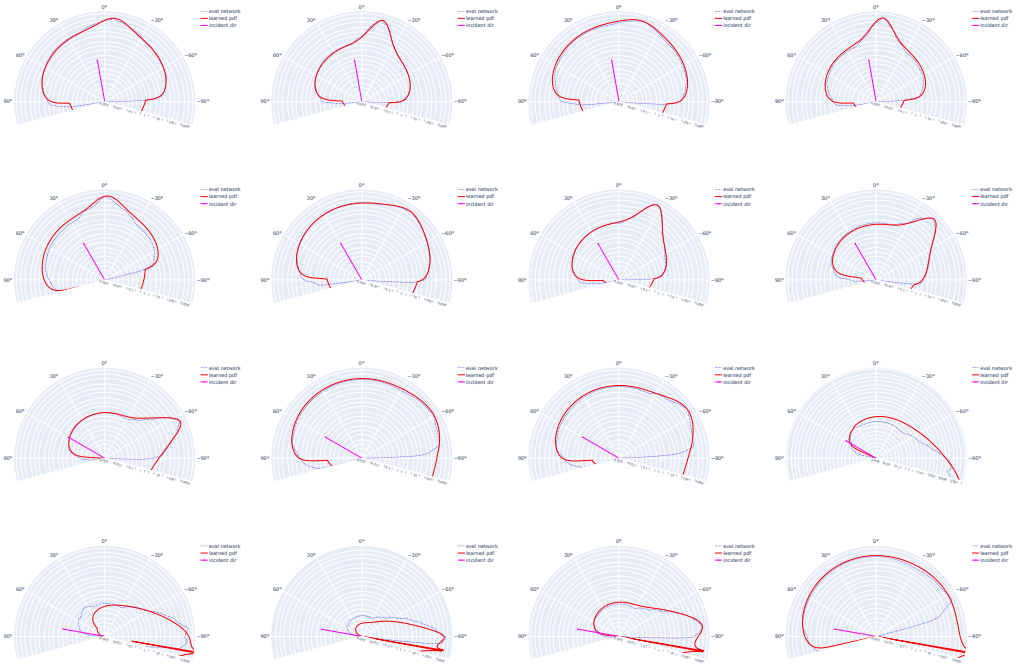

Fig. 29. Cheese slicer handle: importance sampling network plots.

Inkwell metal body (3 layers w/ 64 neurons)

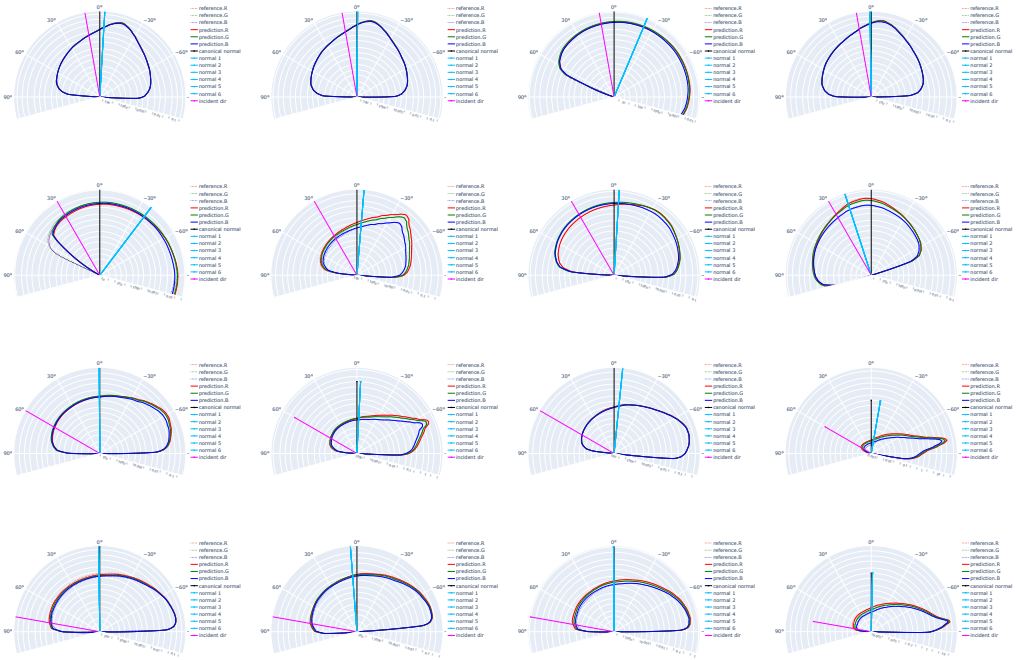

Fig. 30. Inkwell metal body: evaluation network (3 layers w/ 64 neurons) plots.

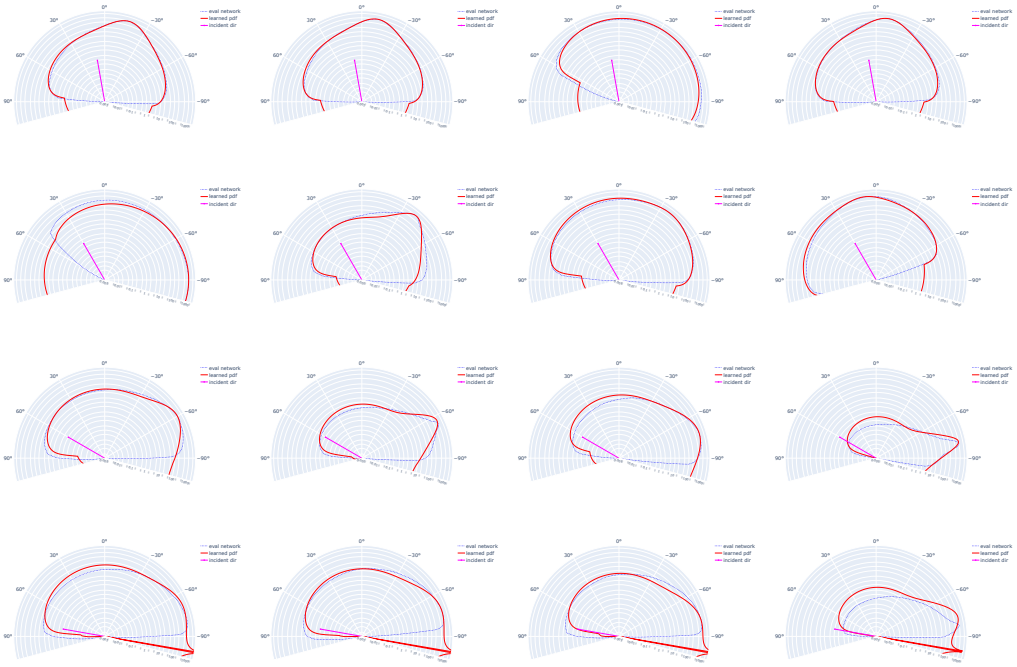

Fig. 31. Inkwell metal body: importance sampling network plots.

Inkwell metal body (2 layers w/ 32 neurons)

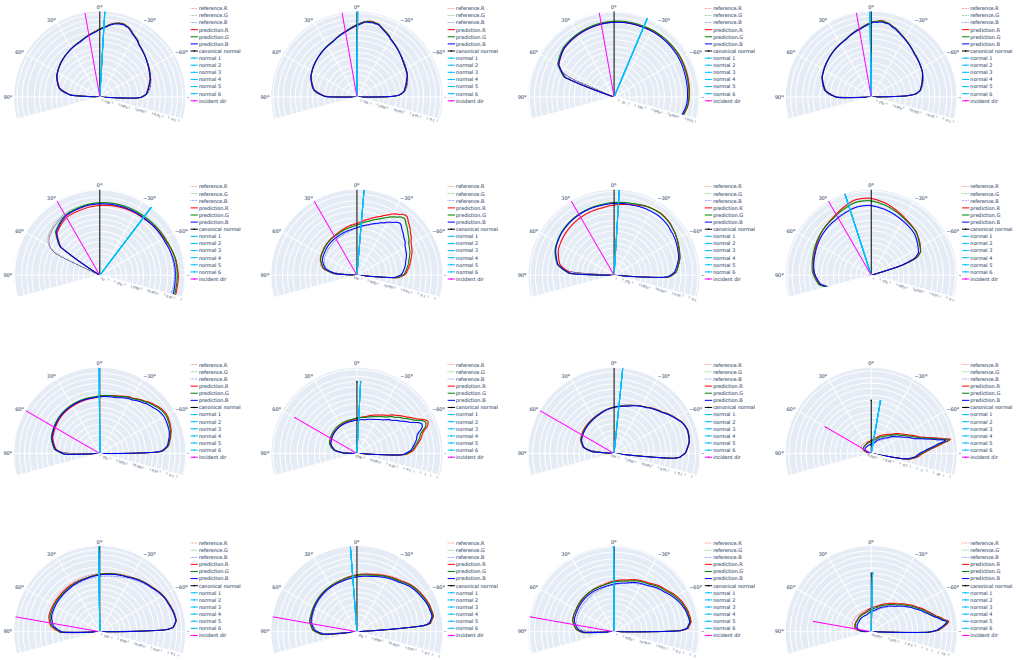

Fig. 32. Inkwell metal body: evaluation network (2 layers w/ 32 neurons) plots.

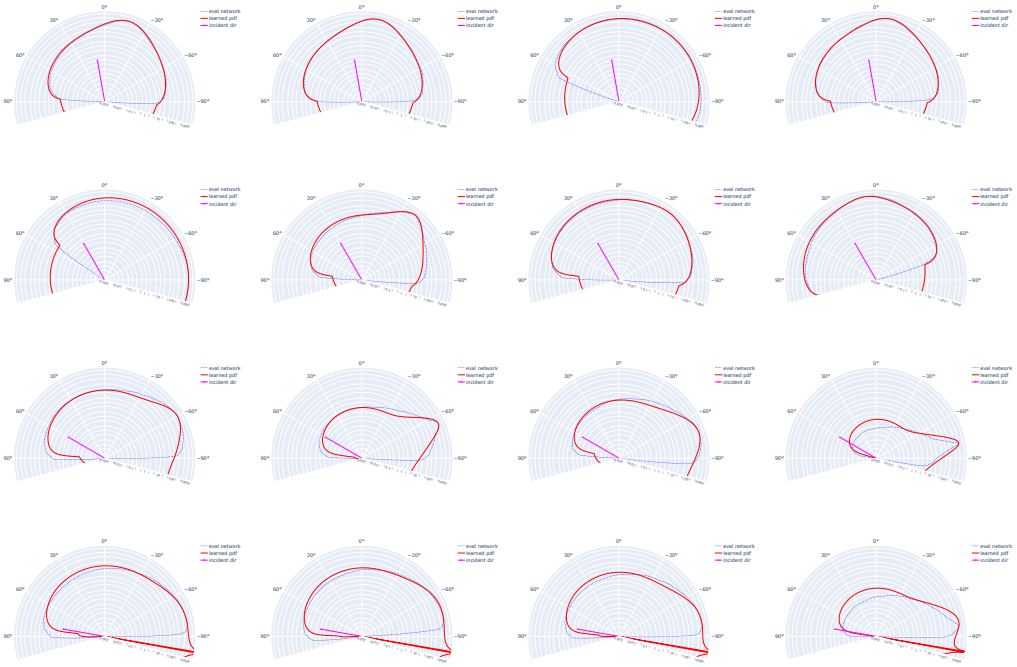

Fig. 33. Inkwell metal body: importance sampling network plots.

Inkwell metal body (2 layers w/ 16 neurons)

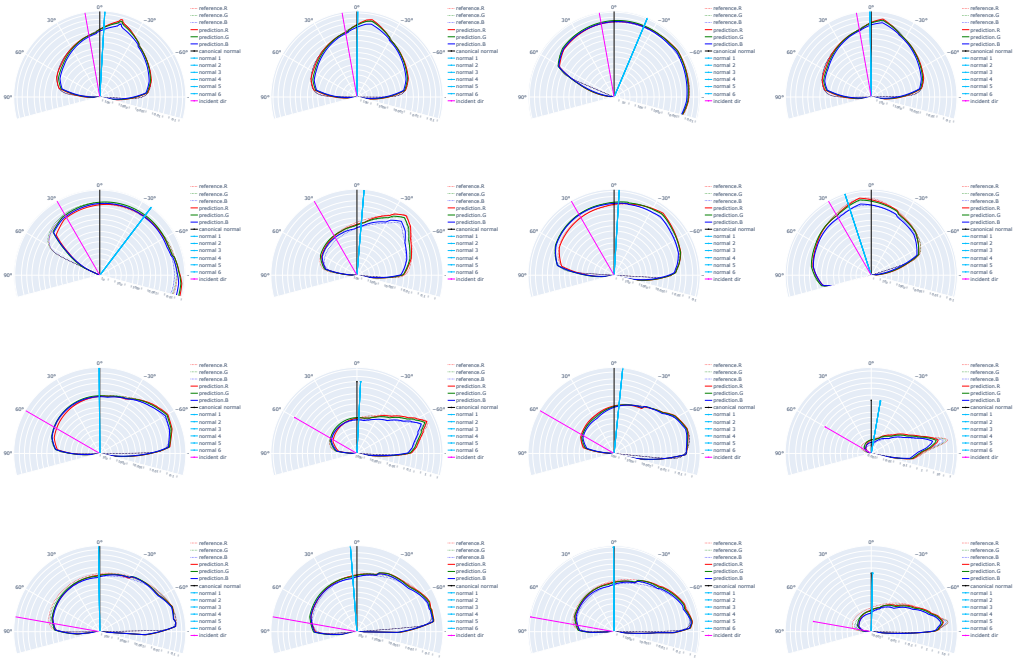

Fig. 34. Inkwell metal body: evaluation network (2 layers w/ 16 neurons) plots.

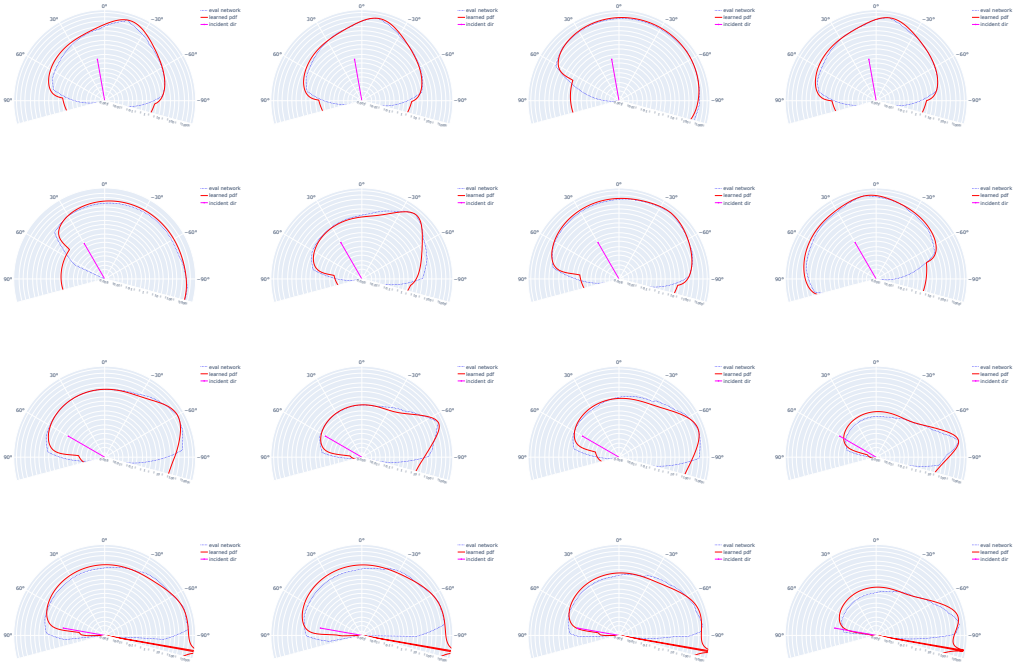

Fig. 35. Inkwell metal body: importance sampling network plots.

REFERENCES

Diederik Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations* (12 2014).

Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2003. A Data-driven Reflectance Model. *ACM Transactions on Graphics* 22, 3 (2003), 759–769.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 8024–8035. http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf