Neural BRDF Representation and Importance Sampling

Alejandro Sztrajman Gilles Rainer Tobias Ritschel Tim Weyrich

University College London, UK

EGSR 2022 Prague, 6 July 2022

BRDF representations

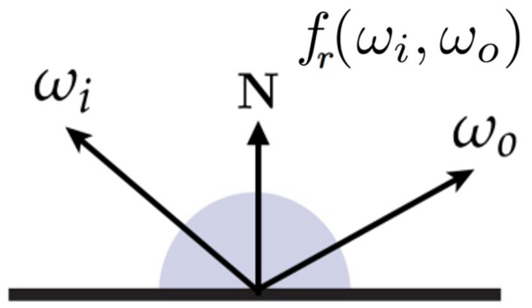

Despite its limitations, the BRDF is the most widely used reflectance parameterisation.

Traditionally, two prevailing modelling paradigms:

- Analytic BRDFs
  - closed-form reflectance functions
- Data-driven BRDFs
  - measured
  - discrete reflectance values

MERL Dataset [Matusik et al. 2003]

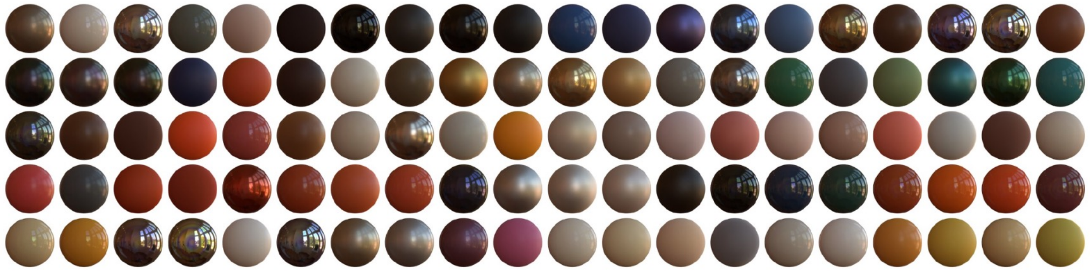

Why data-driven representations?

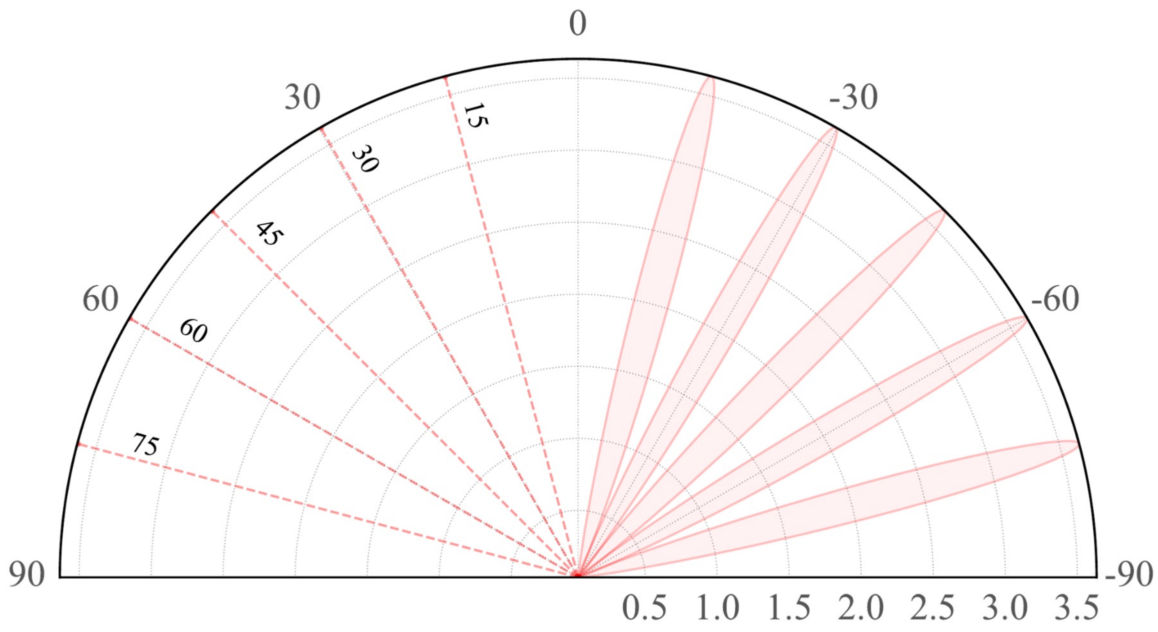

Polar plot of a Phong BRDF for multiple fixed incident azimuth angles (15, 30, 45, 60, 75).

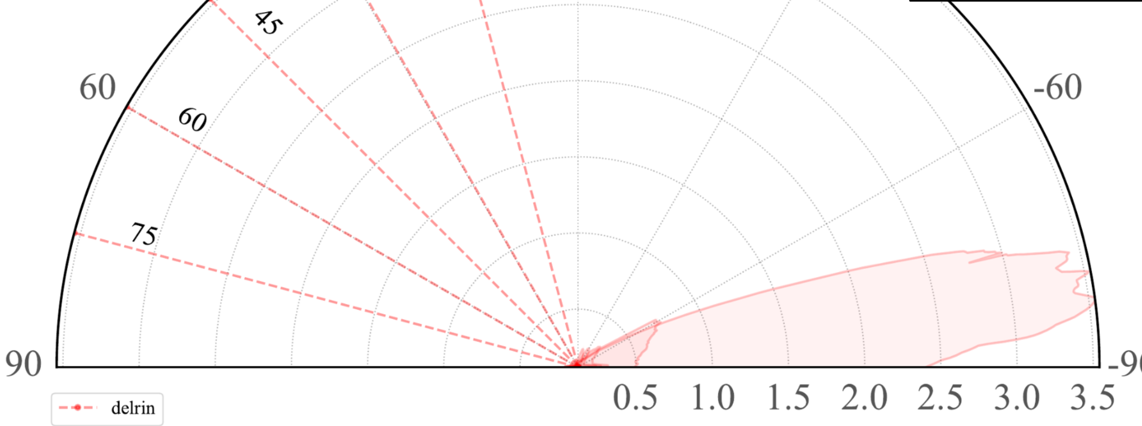

Delrin

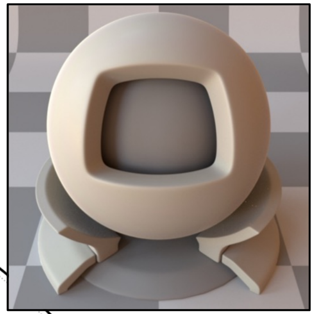

Polar plot of a real-world measured BRDF from the MERL dataset, for multiple fixed incident azimuth angles (15, 30, 45, 60, 75).

Representing measured reflectance data

GT (Tabular)

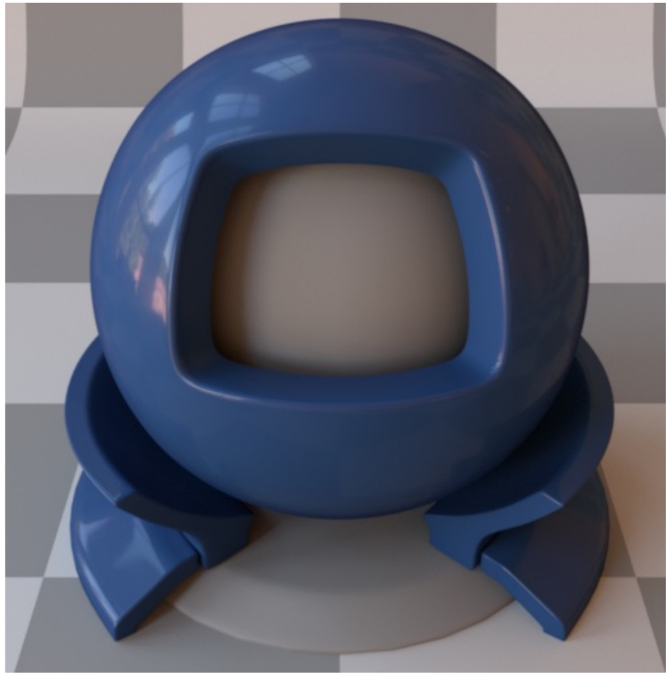

- Accurate
- Large storage (~34 MB)
- Requires interpolation

Analytic Model (here: GGX)

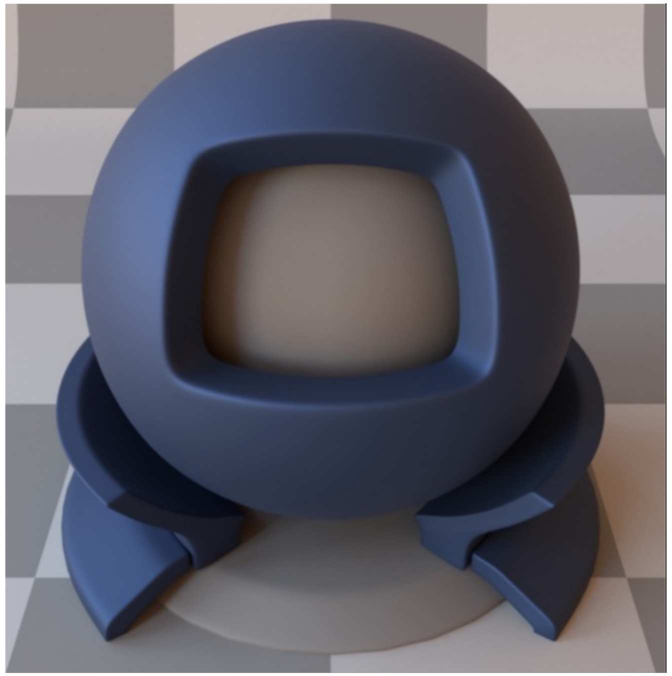

- Requires costly and unstable optimisation
- Often inaccurate
- Very low storage (0.03 KB)
- Fast built-in interpolation

Representation objectives

- Expressive enough *for measured data*
- Compactness (low storage)
- Practical for rendering
  - fast evaluation
  - no angular interpolation artefacts
  - no spatial interpolation artefacts (SVBRDFs)
  - suitable for importance sampling

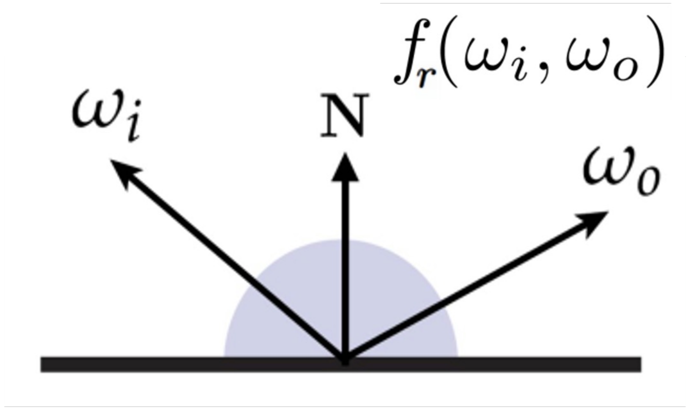

Past representations for measured data

- Analytic model fits
  - [Marschner et al. 1999], [Ngan et al. 2005], [Bagher et al. 2012], [Löw et al. 2012], ...
- Data volume compression
  - PCA [Matusik et al. 2003]
  - matrix factorisation [Lawrence et al. 2004], [Ngan et al. 2006], [Nielsen et al. 2015]
- Non-parametric
  - [Bagher et al. 2016], [Dupuy & Jakob 2018]
- Neural
  - [Maximov et al. 2019], [Hu et al. 2020], [Rainer et al. 2019/2020]

Our neural BRDF representation

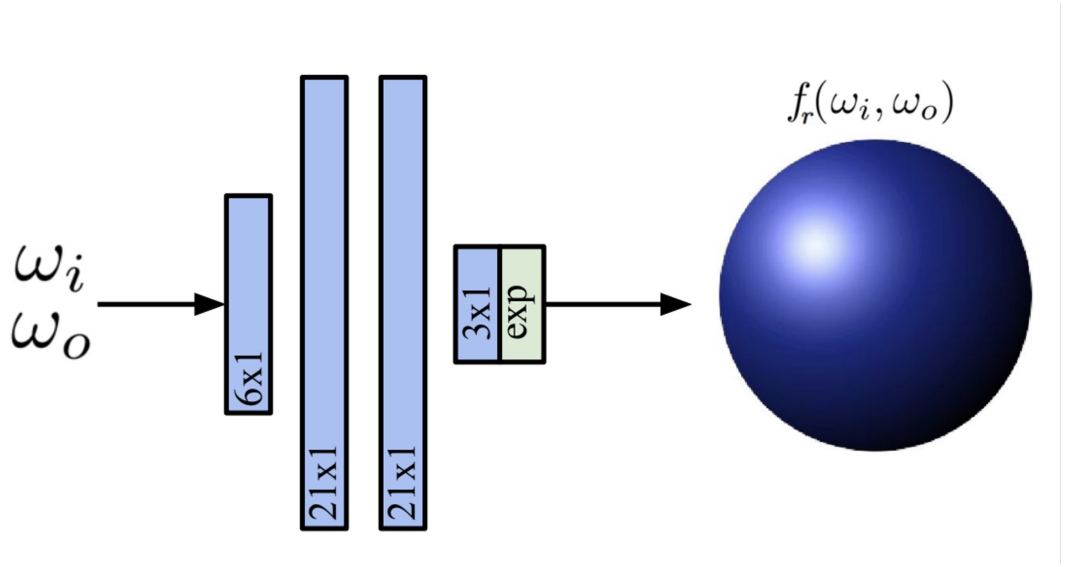

- Directly encodes a BRDF
  - maps hemispherical directions to reflectance
  - Rusinkiewicz parameterisation
  - exponential activation
- Training:
  - image-based loss (cosine-weighted)
  - sampled uniformly in Rusinkiewicz space (denser near highlights)
  - convergence in 10 secs to 3 mins

Neural BRDF

GT (Tabular)

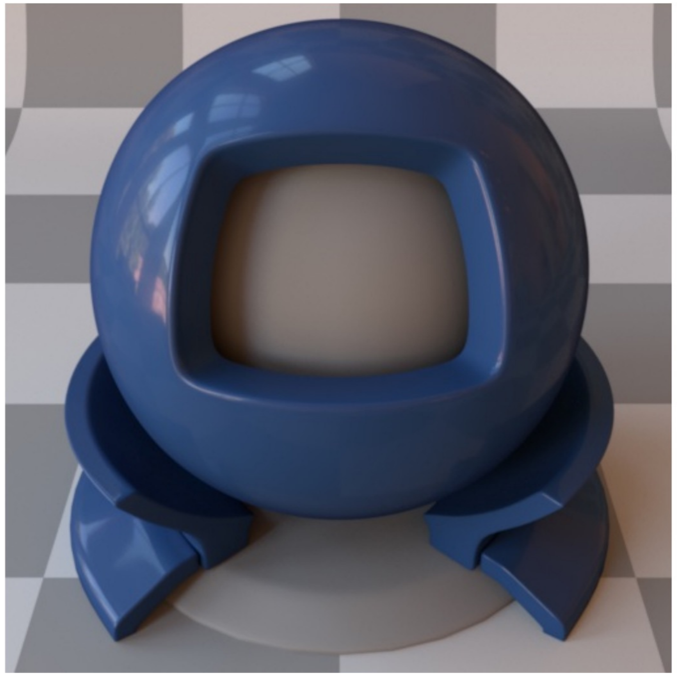

Analytic Model (GGX)

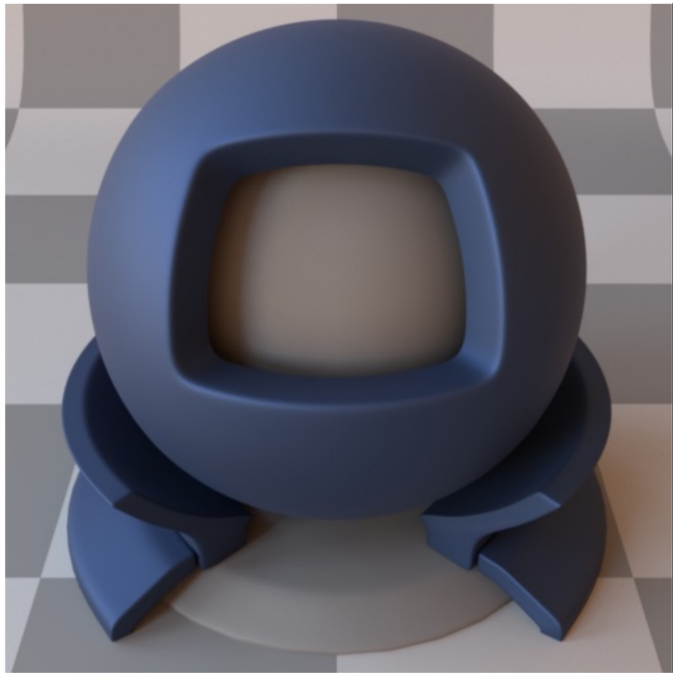

NBRDF (Ours)

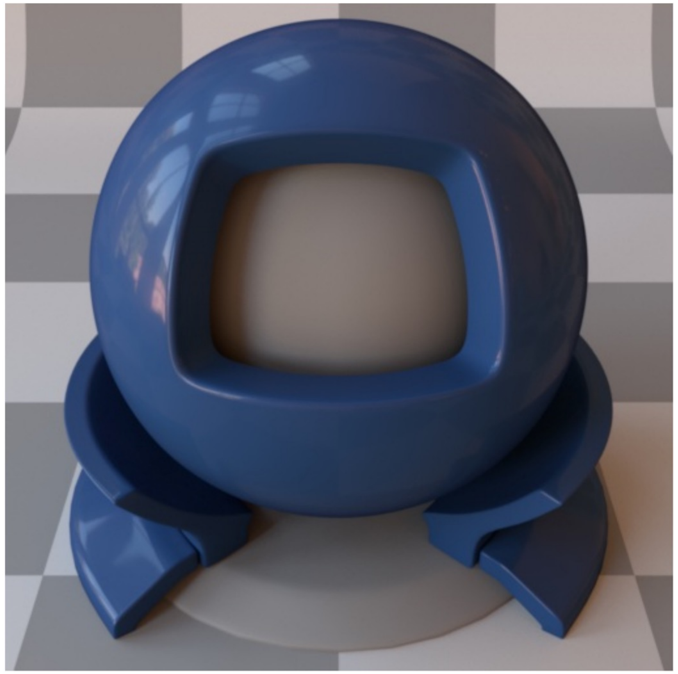

GT (Tabular)
- Accurate
- Large storage (34 MB)
- Requires interpolation

Analytic Model (GGX)
- Requires costly and unstable optimisation
- Often inaccurate
- Very low storage (0.03 KB)
- Fast built-in interpolation

NBRDF (Ours)
- Costly but stable training
- Accurate
- Very low storage (2.7 KB)
- Fast built-in interpolation

MERL materials in different BRDF representations

Inset values: SSIM

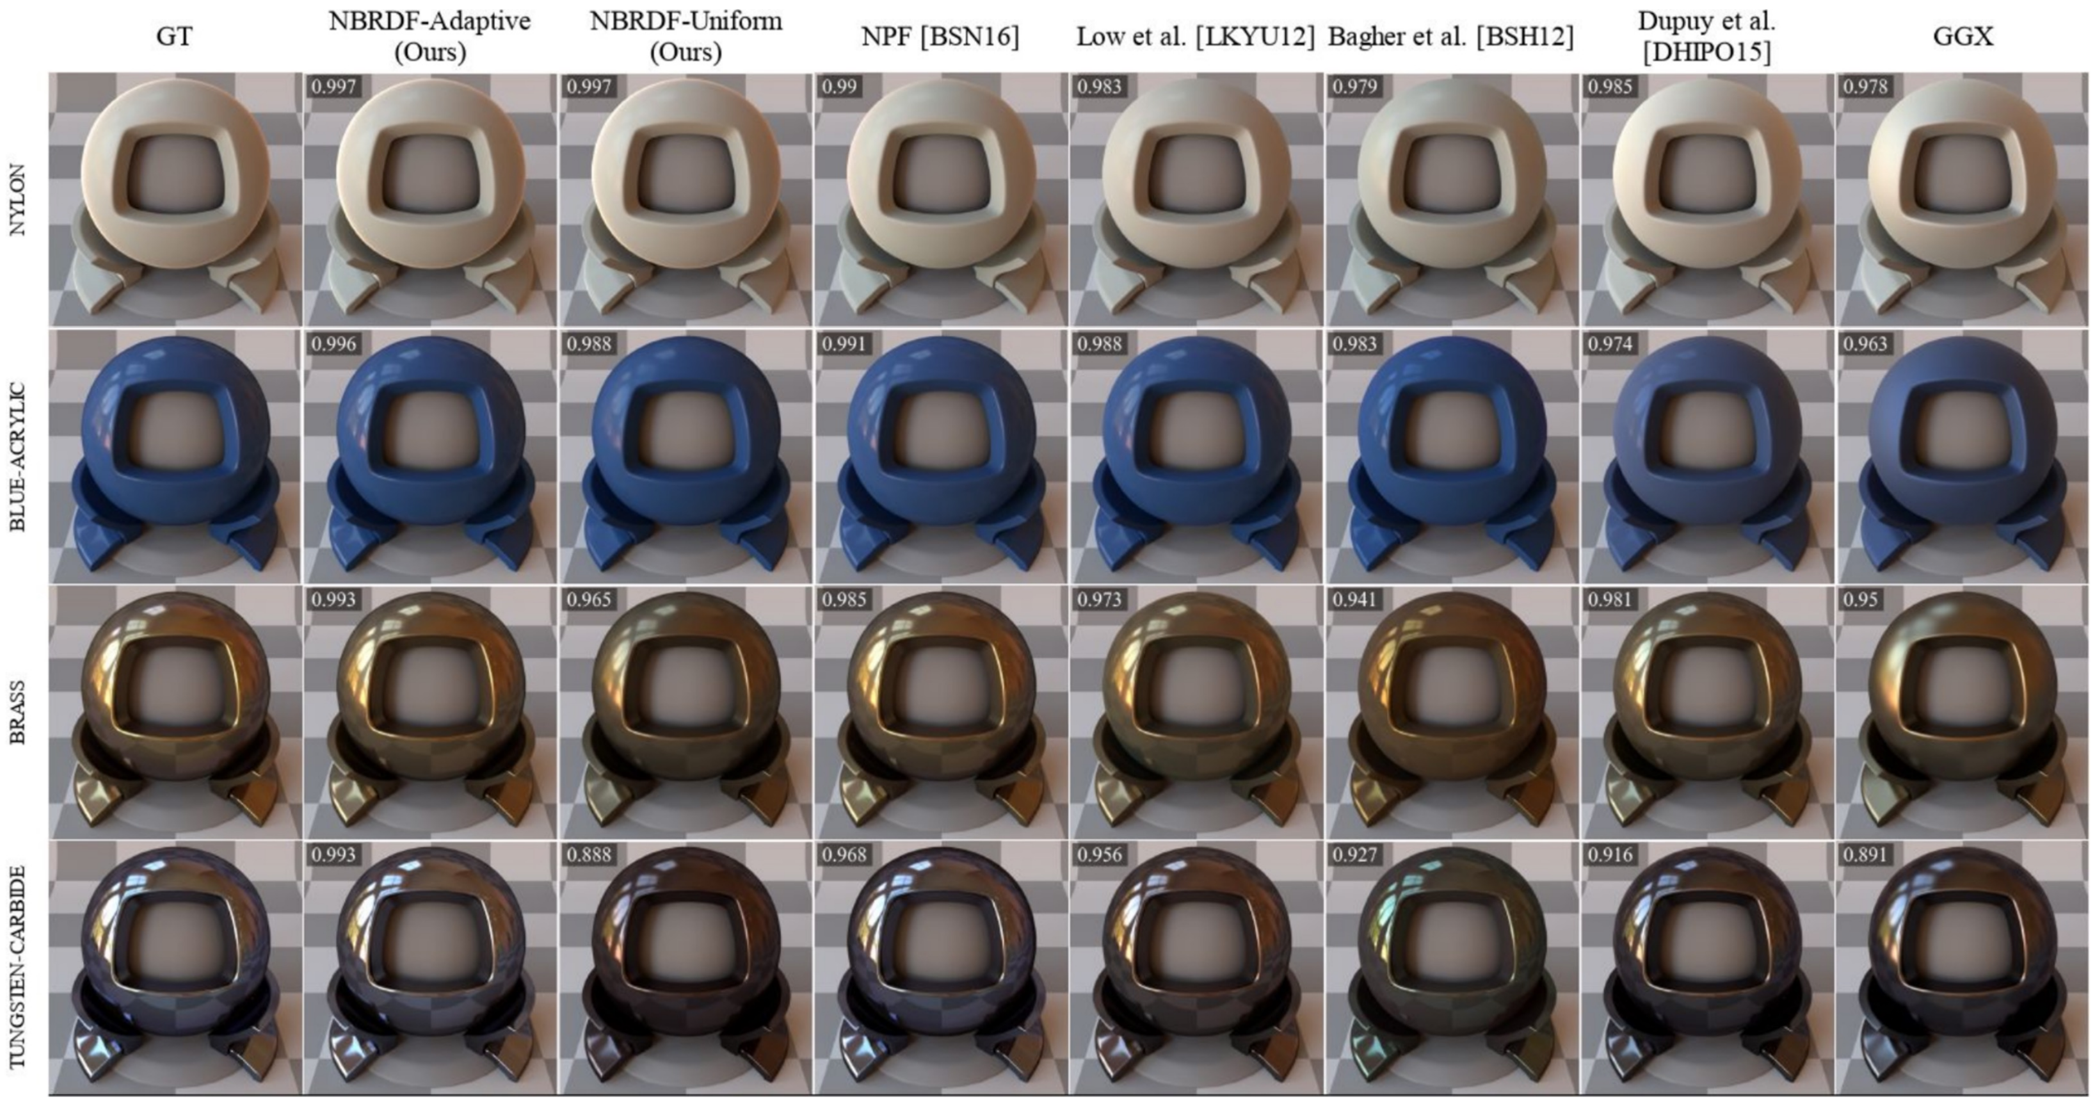

Average SSIM over all MERL materials:

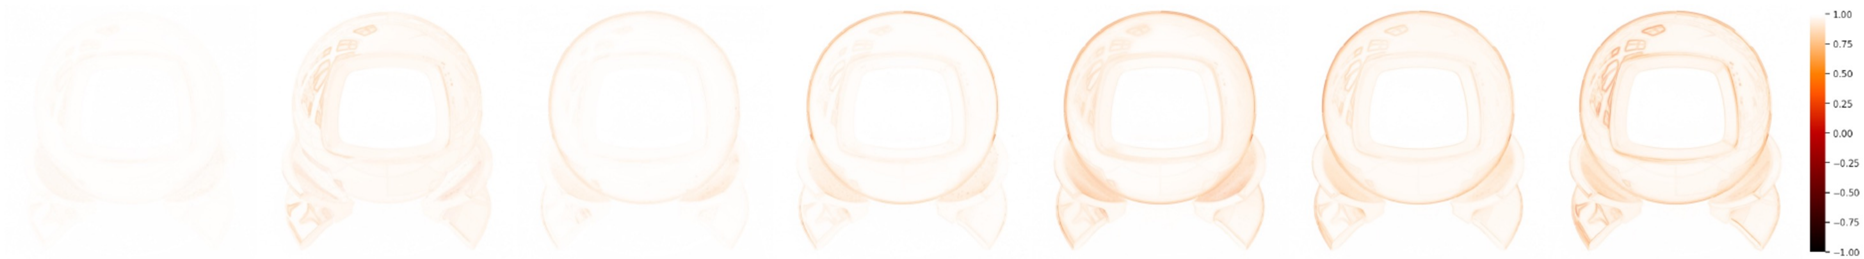

Reconstruction error

Average image-based losses of BRDF representations for all MERL materials:

	MAE	RMSE	SSIM
NBRDF Adaptive Sampling	$\mathbf{0.0028 \pm 0.0034}$	$\mathbf{0.0033 \pm 0.0038}$	$\mathbf{0.995 \pm 0.008}$
NBRDF Uniform Sampling	0.0072 ± 0.0129	0.0078 ± 0.0134	0.984 ± 0.029
NPF [BSN16]	0.0056 ± 0.0046	0.0062 ± 0.0047	0.990 ± 0.008
Low *et al.* [LKYU12] (ABC)	0.0080 ± 0.0070	0.0088 ± 0.0075	0.986 ± 0.012
Bagher *et al.* [BSH12] (SGD)	0.0157 ± 0.0137	0.0169 ± 0.0145	0.974 ± 0.027
Dupuy *et al.* [DHI+15]	0.0174 ± 0.0143	0.0190 ± 0.0151	0.976 ± 0.021
GGX	0.0189 ± 0.0118	0.0206 ± 0.0126	0.969 ± 0.024

Compression and speed

Reconstruction Error vs Representation Size

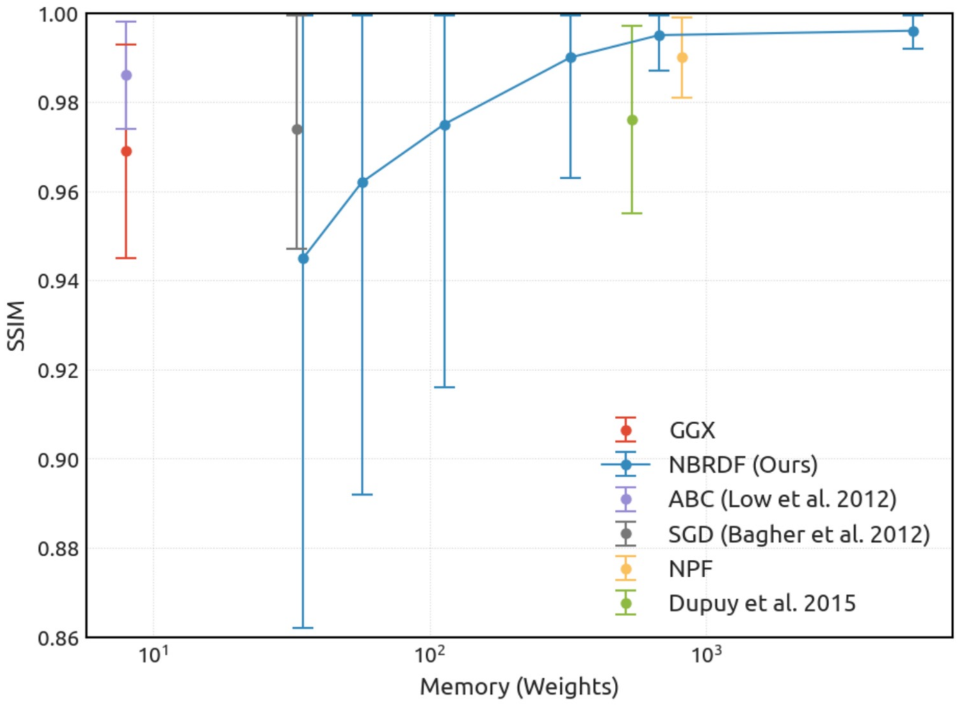

Average SSIM error vs Memory footprint (log scale) for multiple BRDF representations. NBRDFs (in blue) shown for multiple network sizes. *(675 is second from the right)*

High Compression and Fast Evaluation

	Rays/sec ($\times 10^6$)	Memory (KB)
Bagher et al. [BSH12]	10.64	0.13
RGL [DJ18]	10.66	48.0
NBRDF + PhongIS (Ours)	12.50	2.70
Cook-Torrance	13.59	0.03
Dupuy et al. [DHI+15]	14.05	2.16
Low et al. [LKYU12]	15.13	0.03
GGX	16.82	0.03
NPF [BSN16]	–	3.20

Rays traced per second in Mitsuba renderer, and memory footprint, for different material representations.

Anisotropic materials

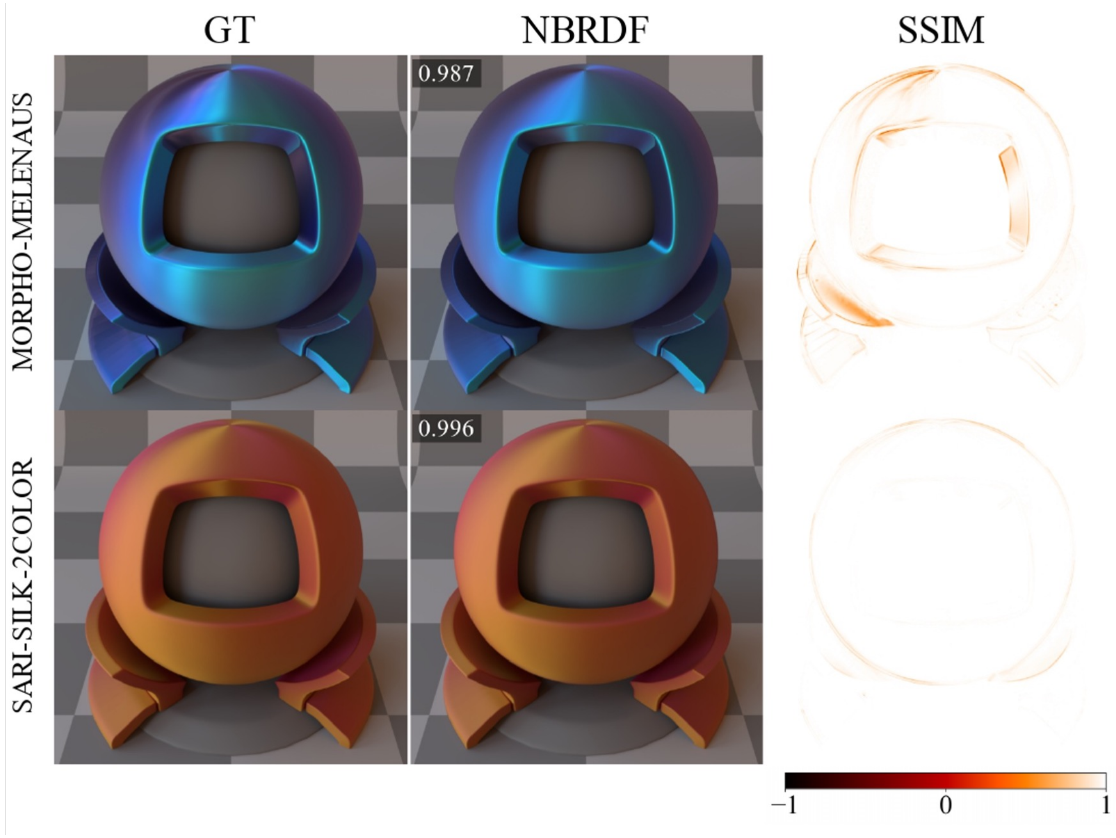

- Neural BRDF reconstruction of materials from the EPFL/RGL dataset [Dupuy and Jakob 2018]
  - additional DOF requires 5× sample count for training
  - slight increase in visual differences (average SSIM of 0.981 ± 0.016)

Representation objectives

- ☑ Expressive enough for measured data
- (☑) Compactness (low storage)
- ☐ Practical for rendering
  - ☑ fast evaluation
  - ☑ no angular interpolation artefacts
  - ☐ no spatial interpolation artefacts (SVBRDFs)
  - ☐ suitable for importance sampling

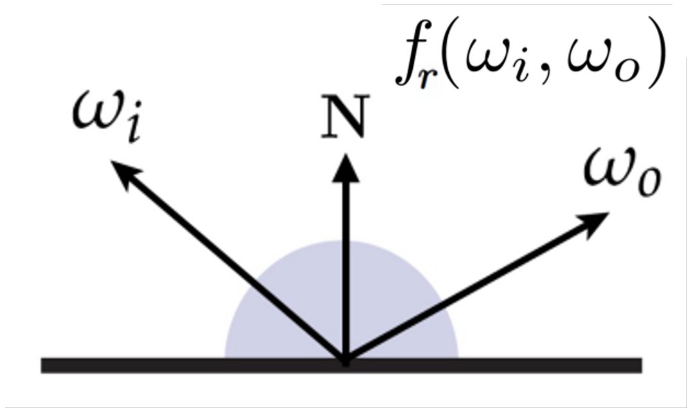

Hyper-network: NBRDF autoencoder

- Input and output are Neural BRDF network weights
- Latent representation are 32-value vectors
  - a more compact NBRDF parameterisation
  - ideally, suited for NBRDF interpolation
- Training
  - with NBRDFs of the MERL database
  - image-based loss in NBRDF output domain
    - evaluates GT and predicts output NBRDF's output
    - implemented as differentiable rendering loss

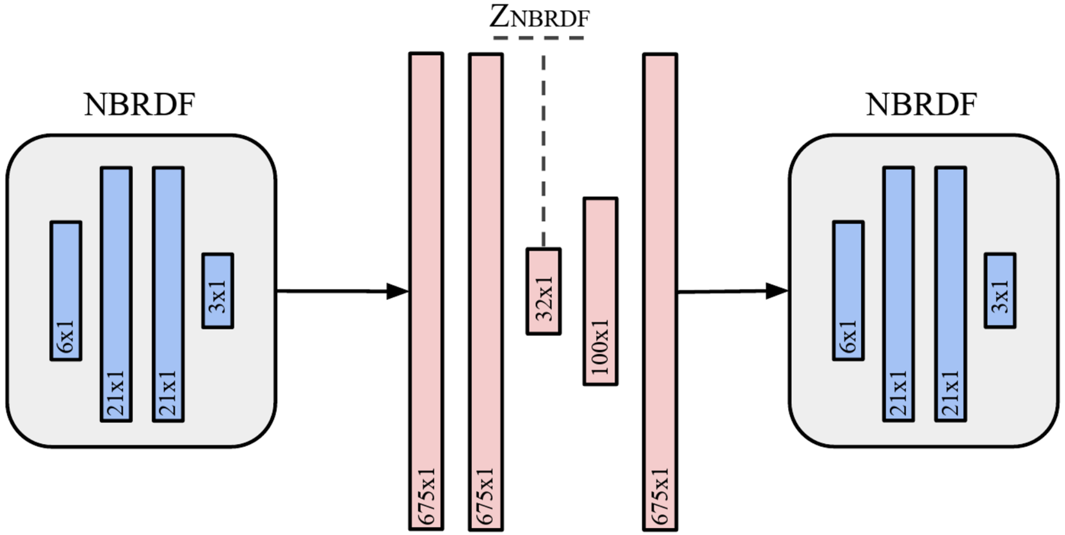

NBRDF embedding in latent space

- Evaluation by t-SNE clustering
  - of MERL materials encoded by hyper-network
  - test-set materials outlined in red
- Materials cluster according to common reflectance properties
  - suggests favourable outcome

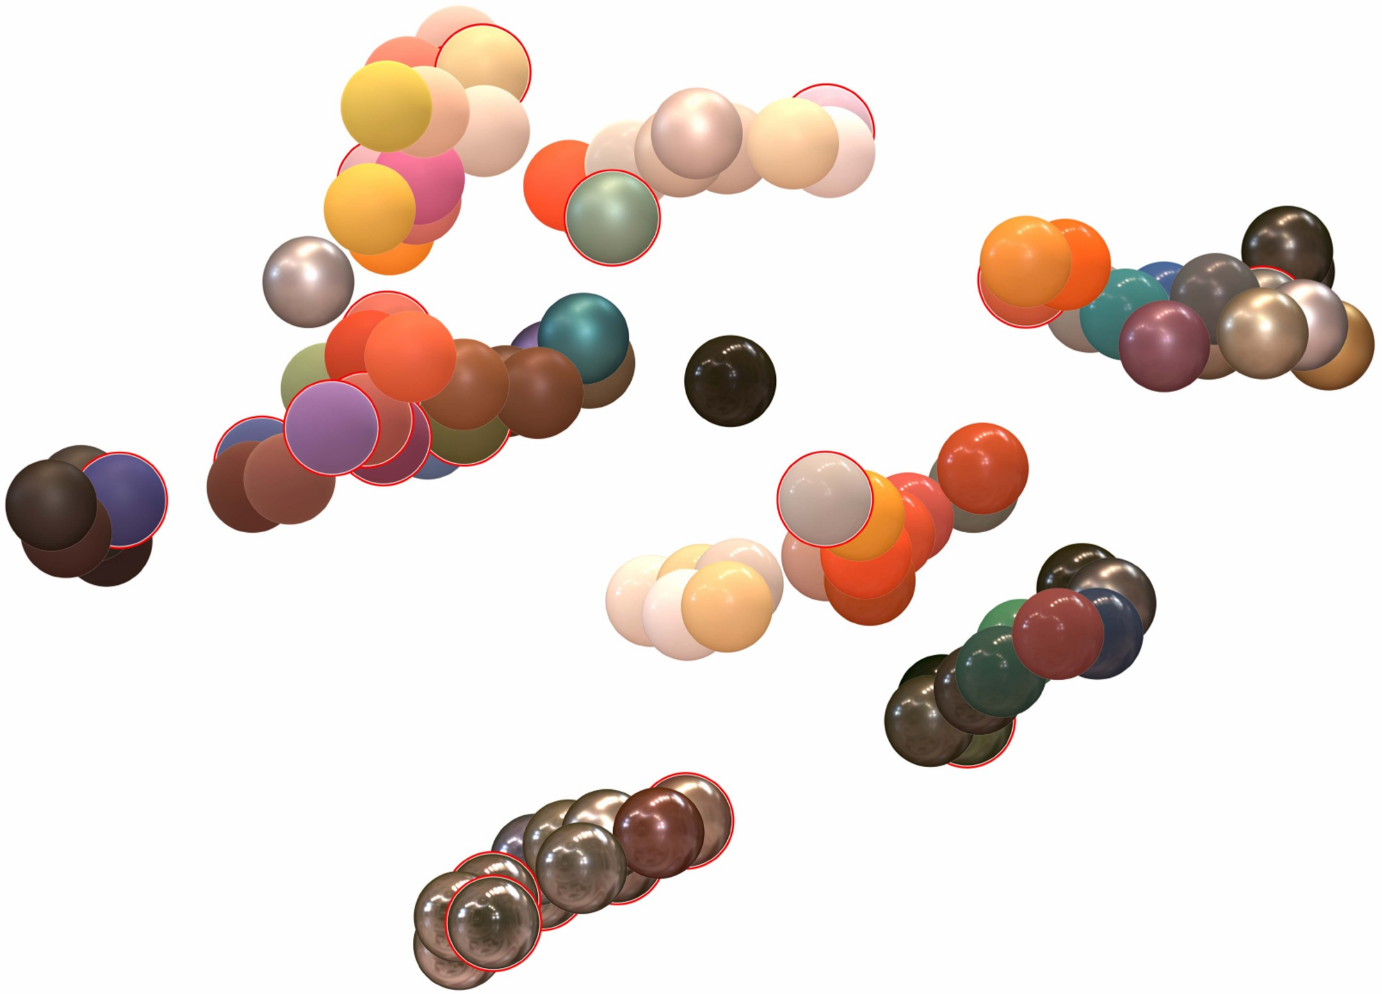

NBRDF interpolation

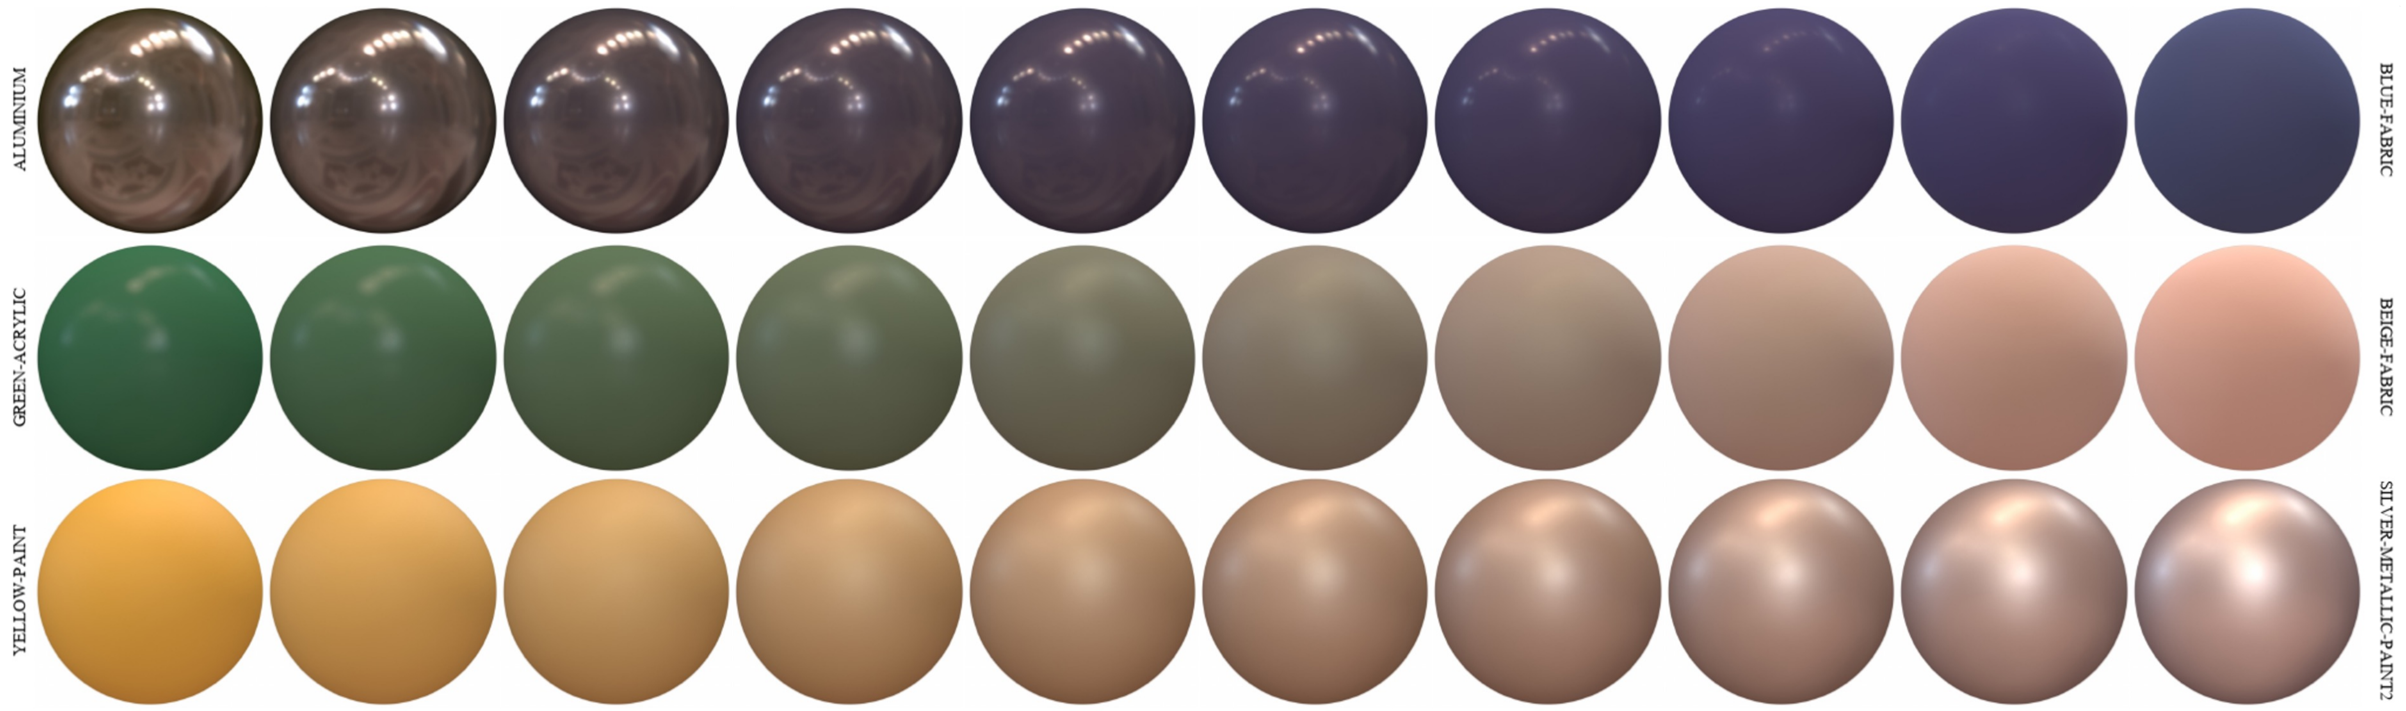

Plausible interpolation between NBRDF embeddings

- enables creation of new materials
- desirable property for extension to Neural SVBRDFs

Representation objectives

- ☑ Expressive enough *for measured data*
- (☑) Compactness (low storage)
- ☐ Practical for rendering
  - ☑ fast evaluation
  - ☑ no angular interpolation artefacts
  - ☑ no spatial interpolation artefacts (SVBRDFs)
  - ☐ suitable for importance sampling

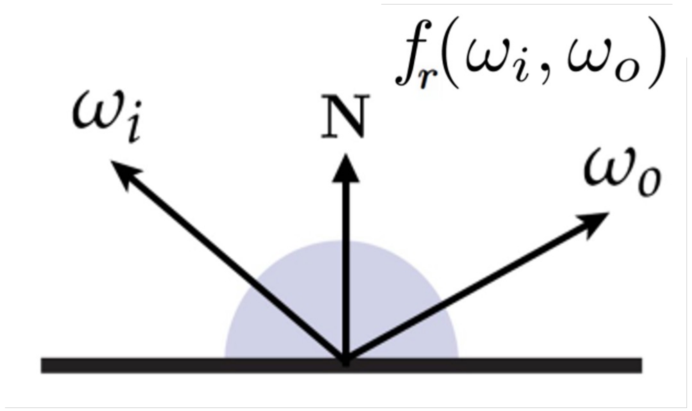

Importance sampling

- Indispensable for efficient path tracing
- Requires sampling from a PDF
  - via uniform sampling of its CDF^{-1}...
  - ... which would not be readily available for measured data and/or an NBRDF :-(
- Key insight [Lawrence et al. 2004]
  - importance sampling converges even if the PDF differs from the BRDF
  - provides room to *pick a PDF whose* CDF^{-1} *is known* :-)

Importance sampling

- General approach
  - choose any parametric BRDF model with known CDF^{-1}
  - fit that model to the NBRDF
  - choose its CDF^{-1} for importance sampling
- How to do so efficiently?
- Neural implementation
  - network to predict analytic parameters from (embedded) NBRDF
  - only predicting parameters relevant for IS
  - we tested Phong and GGX
  - Phong performed best; CDF^{-1} defined by two parameters

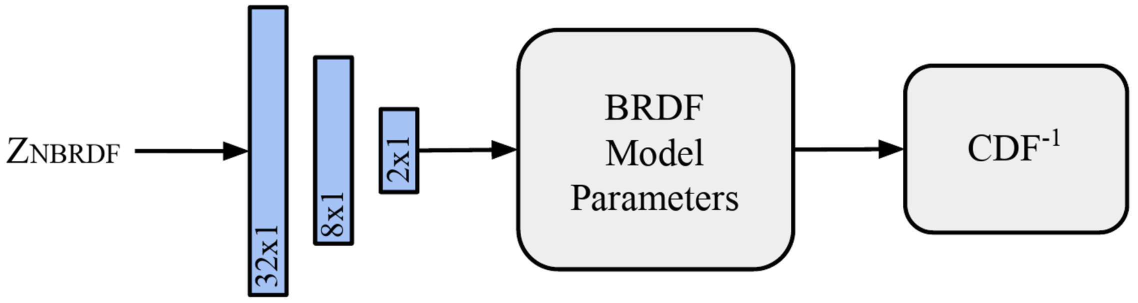

Importance sampling

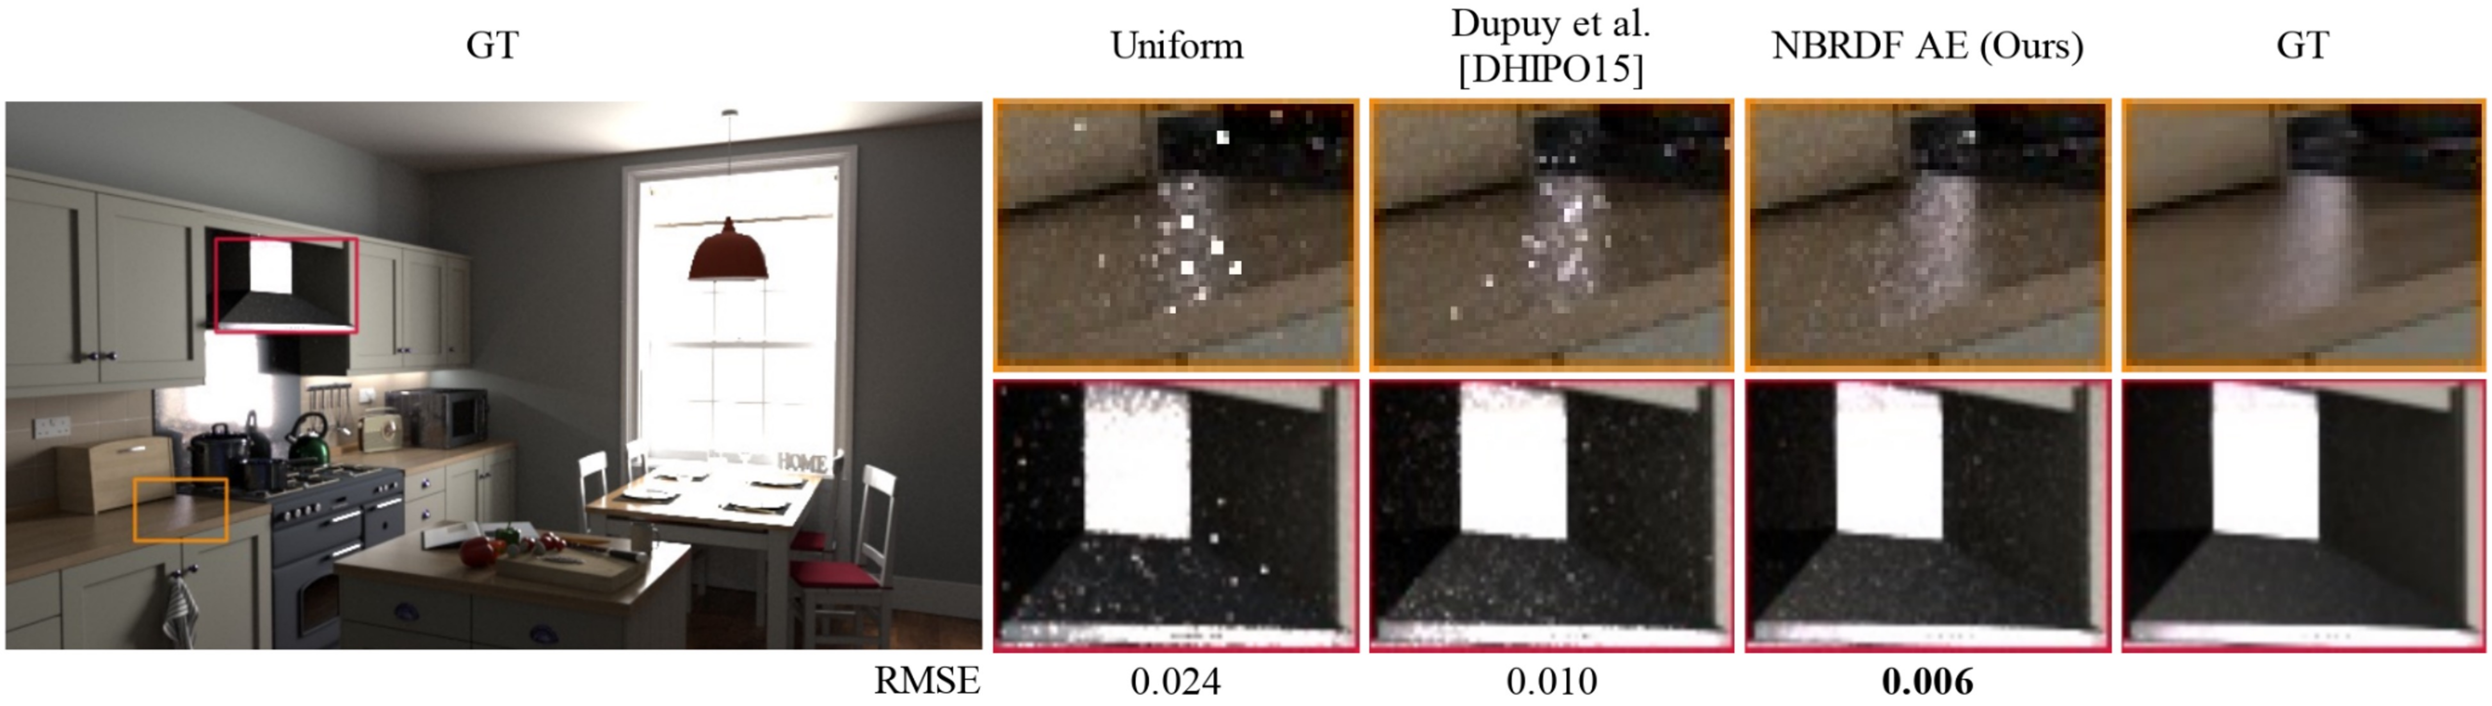

Importance sampling of a kitchen scene using 64 SPP. Most materials in the scene have been replaced by MERL materials within our test set.

Importance sampling

Average RMSE errors (log scale)
vs SPP/render time.

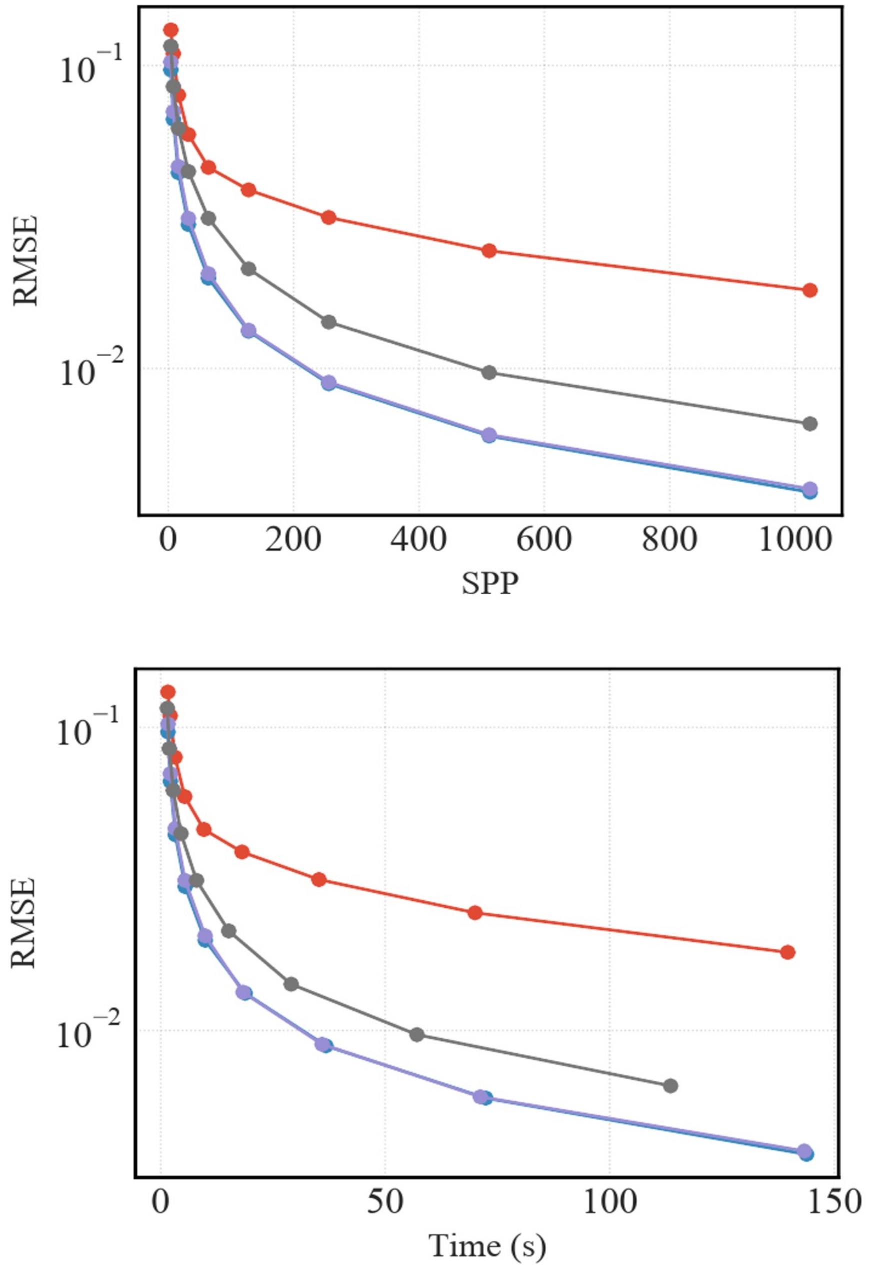

Representation objectives

- ☑ Expressive enough *for measured data*
- ☑ Compactness (low storage)
- ☑ Practical for rendering
  - ☑ fast evaluation
  - ☑ no angular interpolation artefacts
  - ☑ no spatial interpolation artefacts (SVBRDFs)
  - ☑ suitable for importance sampling

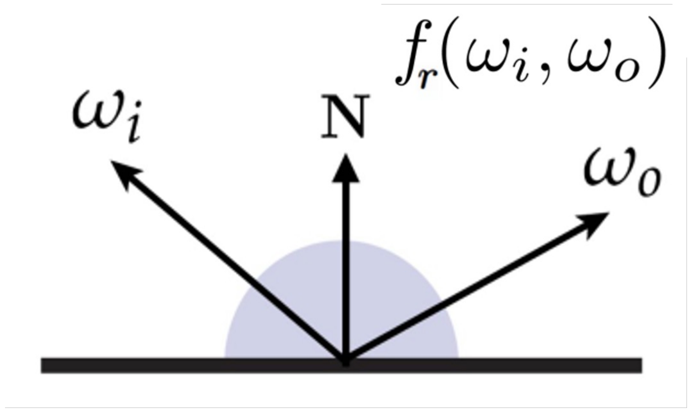

Summary

- Neural representation for measured BRDF data (NBRDF)
  - isotropic + anisotropic
  - higher fidelity than other representations
  - storage- and compute-efficient
- Hyper-network autoencoder with a differentiable rendering loss
  - creates compact embedding of NBRDFs with good interpolation properties
- Learnt mapping between embedded NBDRFs and an invertible analytic BRDF/CDF, enabling importance sampling
- Improves viability of measured BRDFs for practical applications

Subsequent / concurrent work

- “A compact representation of measured BRDFs using neural processes” [Zheng et al. 2022; concurrent]
  - autoencoder representation for BRDFs
  - (lower) 7-dimensional representation, but much larger decoder

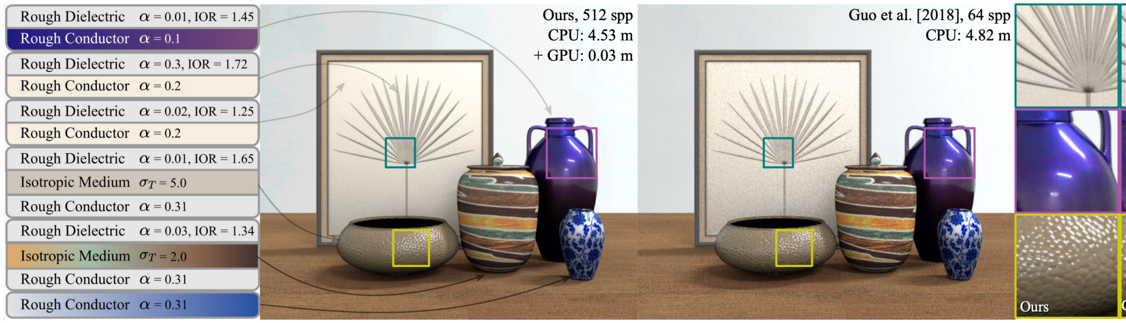

- “Neural layered BRDFs” [Fan et al. 2022]
  - also directly trains a latent space of BRDFs that share one decoder

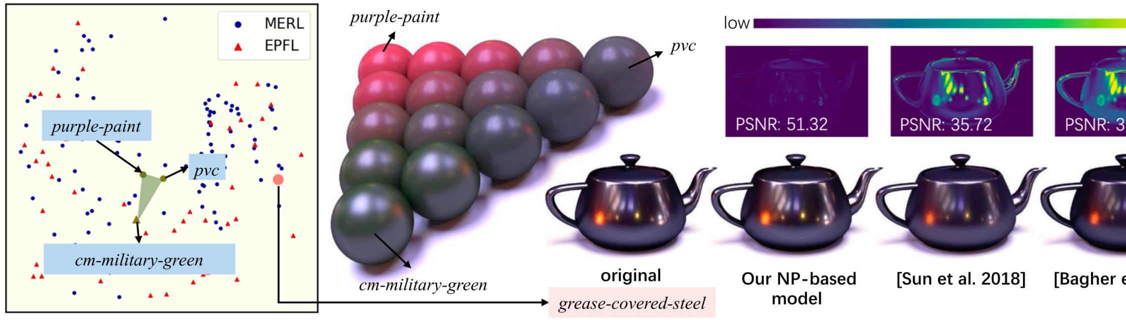

Supplemental material

See https://reality.cs.ucl.ac.uk/projects/reflectance-remapping/sztrajman2021neural.html for...

- reconstruction results for both MERL and EPFL/RGL databases
- our NBRDF training implementation (Keras)
- a Mitsuba plugin to render NBRDFs
- a dataset of pretrained NBRDFs for materials from the MERL, EPFL/RGL and Nielsen et al. databases
- an interactive WebGL demo

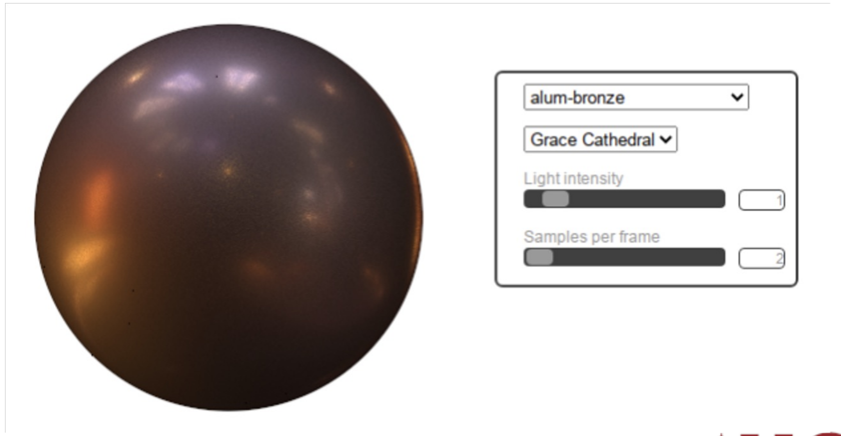

Neural BRDF Representation and Importance Sampling

Alejandro Sztrajman Gilles Rainer Tobias Ritschel Tim Weyrich

University College London, UK

EGSR 2022 Prague, 6 July 2022

UCL