Neural Fields for Data Representation and Generation

Dr. Alejandro Sztrajman

Brown University

Oct 17, 2022

UCL

Content

1. Brief introduction of my background.

2. Summary of research projects.

3. Neural BRDF Representation and Importance Sampling.

4. HyperTime: Neural Fields for Interpretable Time Series Generation

Background

BSc in Physics — *University of Buenos Aires (UBA), Argentina*

Visiting Student — *Columbia University, New York, USA*
Worked on physics-based animation of fluids under the supervision of Profs. *Eitan Grinspun* and *Christopher Batty*.

PhD in Computer Science — *University College London (UCL), UK*
Machine learning applications in appearance modelling. Combining methods from neural networks and physics-based rendering on tasks such as material and light representation, appearance transfer and light estimation.
Supervisors: Profs. *Tim Weyrich* and *Tobias Ritschel*.

Marie-Curie Fellowship - *European Commission*
Funding provided by the Marie-Curie Actions Programme, through the DISTRO innovative training network.

Visiting Student — *Charles University, Prague, Czech Republic*
Worked on material appearance remapping with Profs. *Jaroslav Krivanek* and *Alexander Wilkie*.

Research Intern — *Adobe Substance 3D, Clermont-Ferrand, France*
Worked on material appearance transfer between renderers with Dr. *Cyrille Damez*.

Research Intern — *Microsoft, Reading, UK*
Worked on HDR light representation and estimation with Drs. *Eric Sommerlade* and *Alexandros Neophytou*.

Research Projects

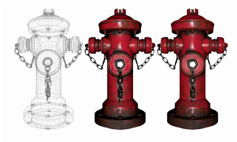

Remapping of Material Appearance

Neural Blue Noise Generation

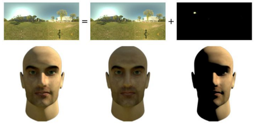

HDR Lighting Representation and Estimation

Neural Fields for Material Appearance Representation and Generation

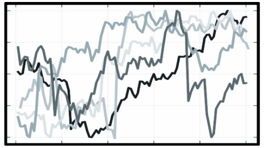

Neural Fields for Time-Series Interpretable Representation and Generation

Research Projects

In collaboration with: Charles University   Adobe Substance 3D

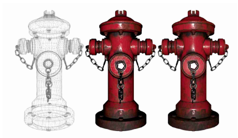

Remapping of Material Appearance

Neural Blue Noise Generation

HDR Lighting Representation and Estimation

Neural Fields for Material Appearance

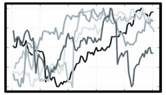

Neural Fields for Time-Series

Motivation:

- Material creation and editing is hard and requires expert knowledge of the parameters used by each specific software.
- Pipelines for content creation often rely on multiple software, with different shader implementations, and materials cannot be interchanged between them.

Figure: Robot created in Adobe Substance 3D (left) and then rendered in Unity 5 (right). [unity3D forum]

Research Projects

In collaboration with: Charles University   Adobe Substance 3D

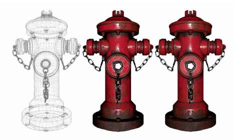

Remapping of Material Appearance

Neural Blue Noise Generation

HDR Lighting Representation and Estimation

Neural Fields for Material Appearance

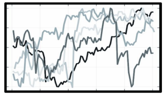

Neural Fields for Time-Series

Remapping of uniform material appearance

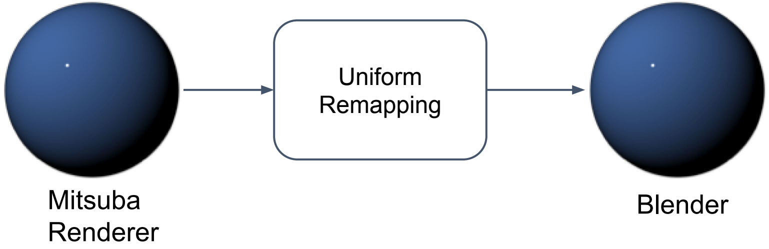

- TRF nonlinear optimisation with image-based loss.
- No need to access the shading code. Only requirement is a rendering of a sphere.

Spatially-Varying Material Appearance (SVBRDF)

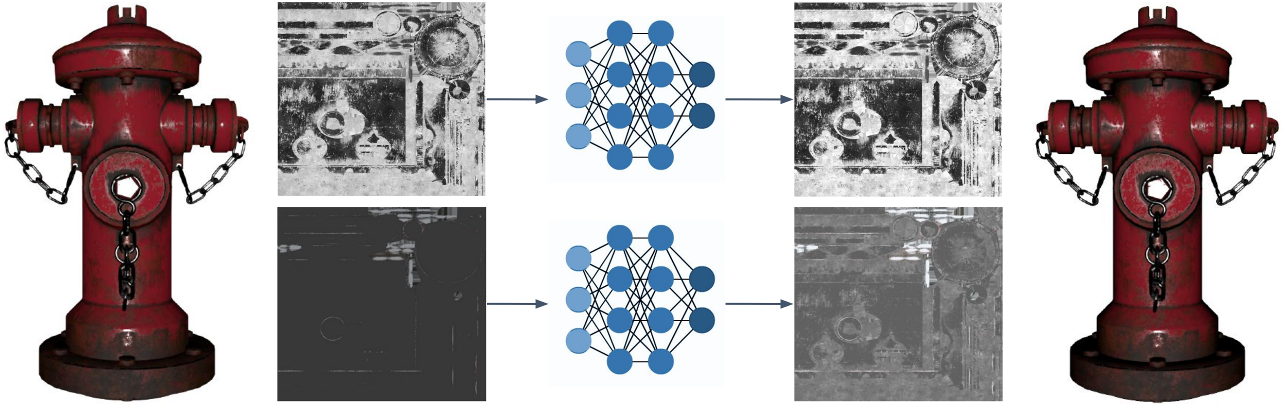

UCL

Research Projects

In collaboration with: Charles University Adobe Substance 3D

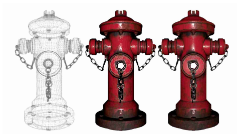

Remapping of Material Appearance

Neural Blue Noise Generation

HDR Lighting Representation and Estimation

Neural Fields for Material Appearance

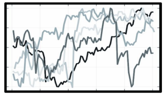

Neural Fields for Time-Series

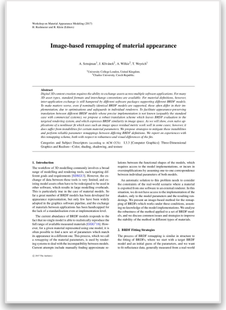

Image-based remapping of material appearance

Image-based remapping of material appearance
A. Sztrajman, J. Krivanek, A. Wilkie, T. Weyrich
Eurographics Workshop on Material Appearance Modeling (2017)

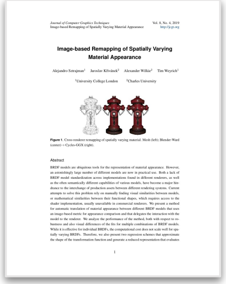

Image-based Remapping of Spatially Varying Material Appearance

Image-based remapping of spatially-varying material appearance
A. Sztrajman, J. Krivanek, A. Wilkie, T. Weyrich
Journal of Computer Graphics Techniques (2019)

Research Projects

In collaboration with:

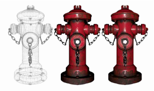

Remapping of
Material Appearance

Neural Blue Noise Generation

HDR Lighting Representation
and Estimation

Neural Fields for
Material Appearance

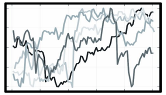

Neural Fields for Time-Series

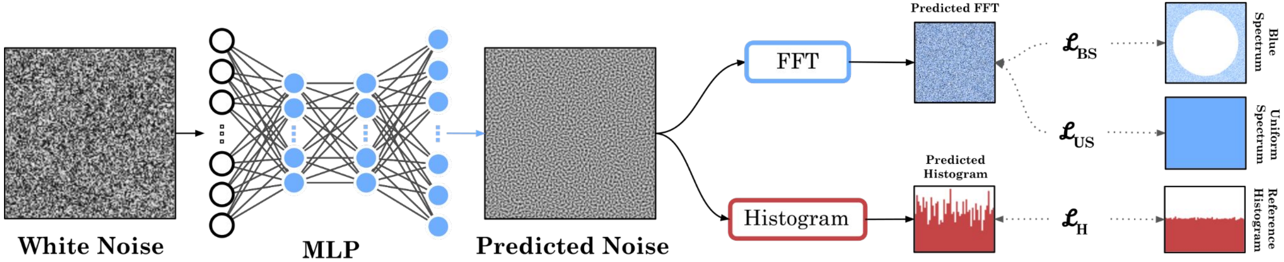

GT | Blue Noise | White Noise

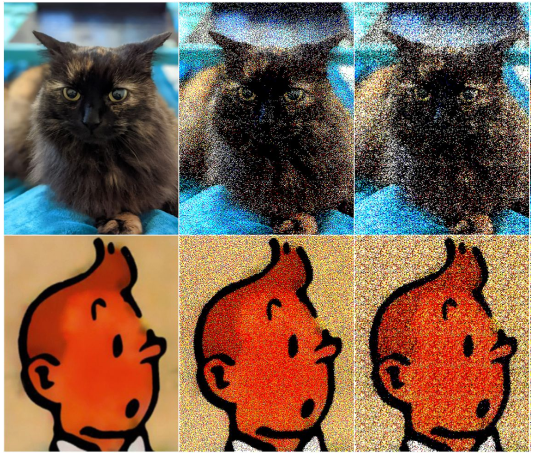

Figure: Image dithering comparison (1 bit).

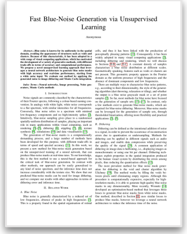

Fast Blue Noise Generation via Unsupervised Learning
*D. Giunchi, *A. Sztrajman, A. Steed
International Joint Conference on Neural Networks (IJCNN), 2022.

Research Projects

In collaboration with: Microsoft

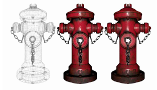

Remapping of Material Appearance

Neural Blue Noise Generation

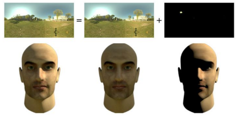

HDR Lighting Representation and Estimation

Neural Fields for Material Appearance

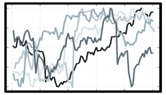

Neural Fields for Time-Series

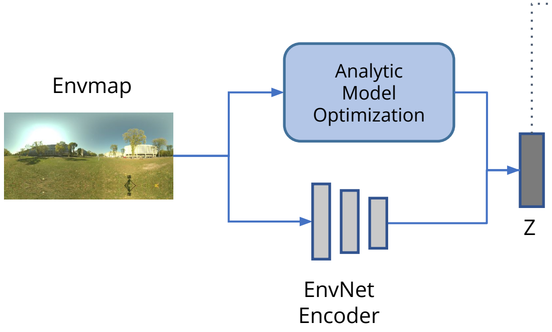

Figure: Hybrid encoding of HDR light.

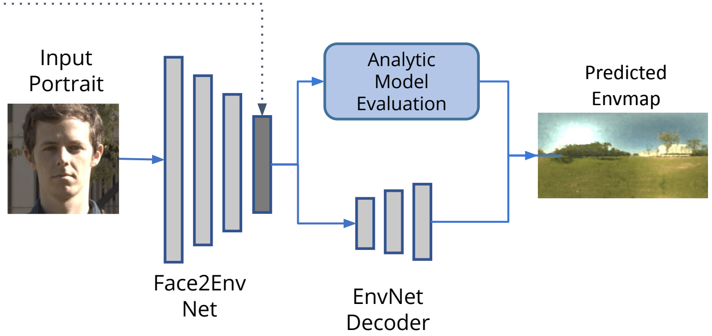

Figure: HDR light prediction from portrait image.

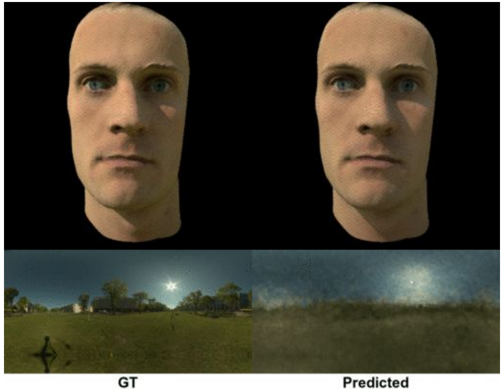

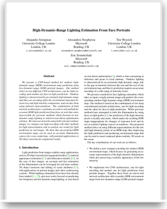

High-Dynamic-Range Lighting Estimation from Face Portraits
A. Sztrajman, A. Neophytou, T. Weyrich, E. Sommerlade
International Conference on 3D Vision (3DV), 2020 (Oral Presentation).

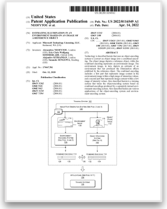

Estimating Illumination in an Environment Based on an Image of a Reference Object
A. Neophytou, E. Sommerlade, A. Sztrajman, S. Sengupta
US Patent 2022/0116549 A1.

UCL

Research Projects

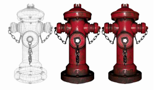

Remapping of
Material Appearance

Neural Blue Noise
Generation

HDR Lighting Representation
and Estimation

Neural Fields for Material Appearance

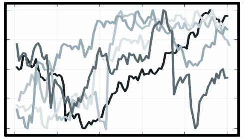

Neural Fields for Time-Series

Neural BRDF Representation and Importance Sampling

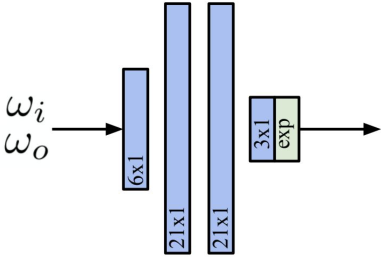

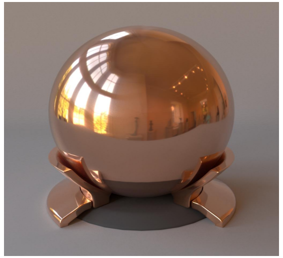

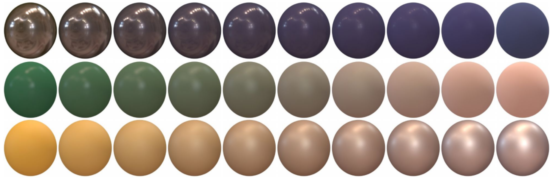

Figure: new realistic materials generated by interpolation.

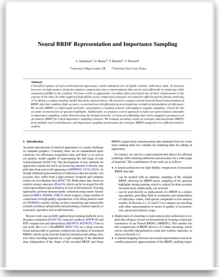

Neural BRDF Representation and Importance Sampling
A. Sztrajman, G. Rainer, T. Ritschel, T. Weyrich
Computer Graphics Forum (CGF), 2021 (EGSR 2022).

Research Projects

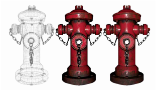

Remapping of Material Appearance

Neural Blue Noise Generation

HDR Lighting Representation and Estimation

Neural Fields for Material Appearance

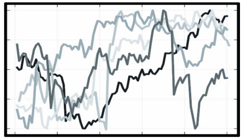

Neural Fields for Time-Series

HyperTime: Implicit Neural Representations for Interpretable Time-Series Generation

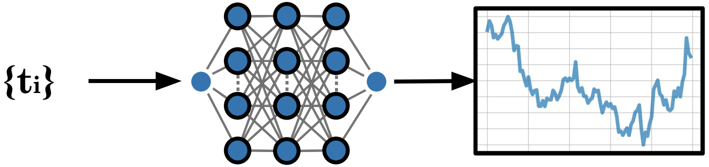

Under review as a conference paper at ICLR 2023

HYPERTIME: IMPLICIT NEURAL REPRESENTATIONS FOR TIME SERIES GENERATION

Anonymous authors

Paper under double-blind review

ABSTRACT

1 INTRODUCTION

HyperTime: Implicit Neural Representations for Interpretable Time-Series Generation
E. Fons, A. Sztrajman, Y. El-Laham, A. Iosifidis, S. Vyetrenko
In Review (2022).

HyperTime: Implicit Neural Representation for Time Series

Abstract

1 Introduction

Generating Interpretable Time-Series by Meta-Learning with Implicit Neural Representations
E. Fons, A. Sztrajman, Y. El-Laham, A. Iosifidis, S. Vyetrenko
Patent Pending

Neural BRDF Representation and Importance Sampling

Neural BRDF: Introduction

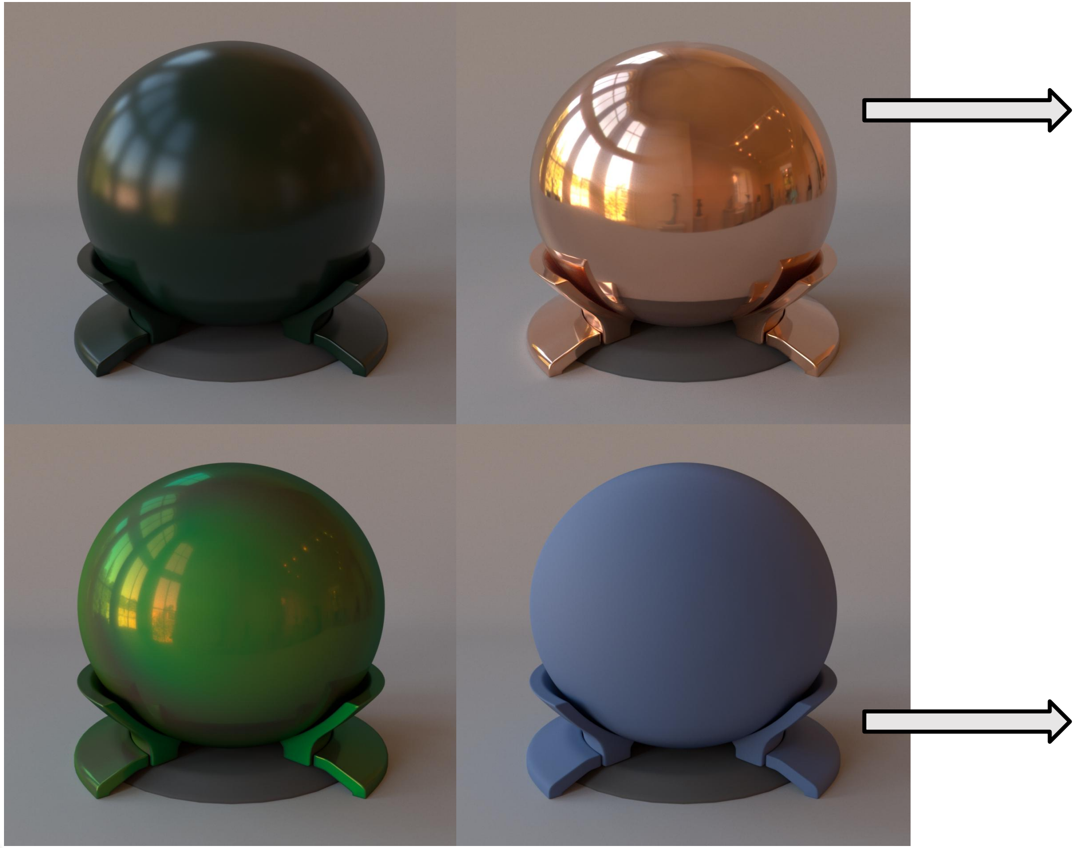

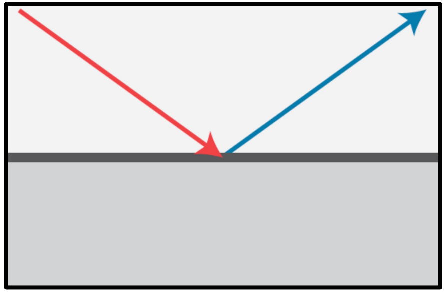

Figure: Scheme of light scattering in a smooth specular surface.

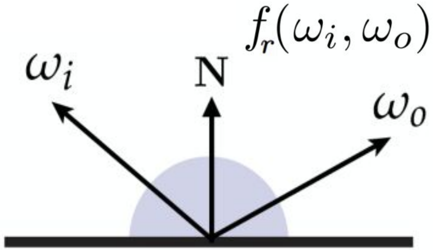

Figure: Diagram of incident and outgoing light directions on the surface. The BRDF is defined as a function of these two directions.

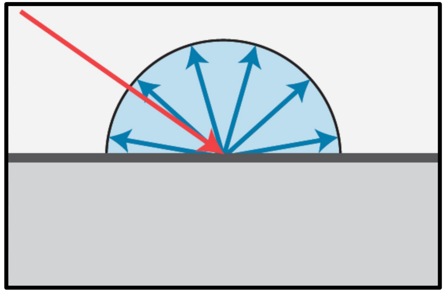

Figure: Scheme of light scattering in a diffuse material.

Neural BRDF: Introduction

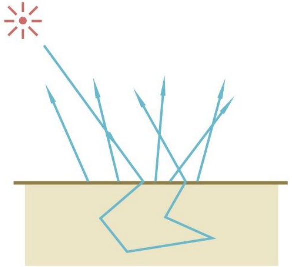

Subsurface Scattering (8D)

Milk rendered with (left) and without (right) subsurface scattering. [H. W. Jensen et al. 2001]

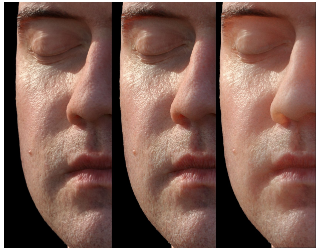

[M. Seymour 2017, PIXAR Deep Dive on SSS]

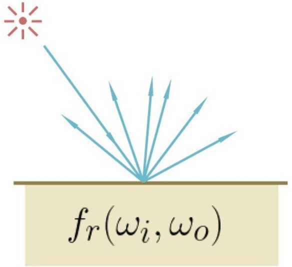

BRDF

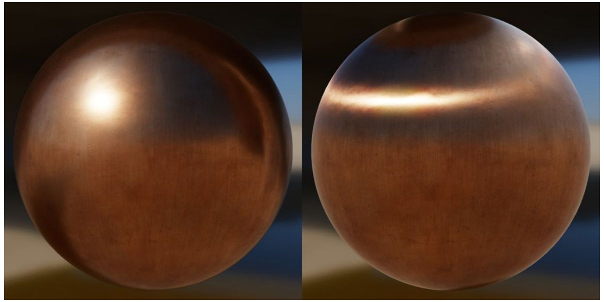

Isotropic (left) and Anisotropic (right) materials.

Anisotropic brushed metal. [blenderartists.org]

Neural BRDF: Introduction

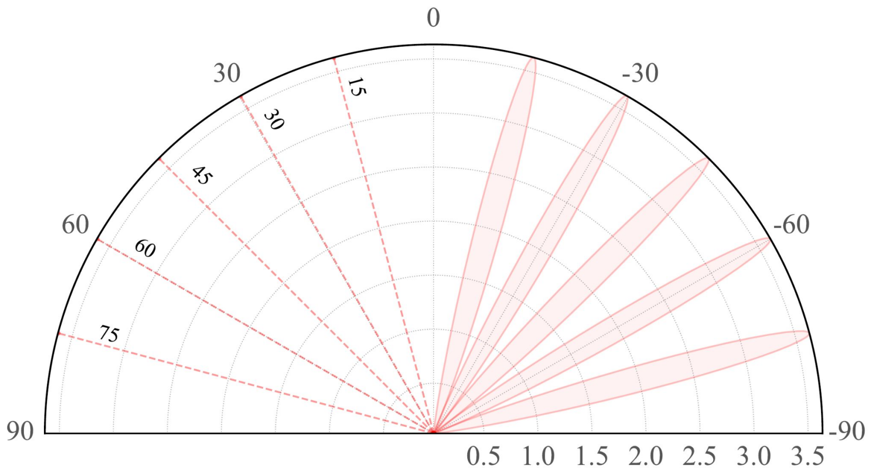

Figure: Polar plot of a Phong BRDF for multiple fixed incident azimuth angles (15, 30, 45, 60, 75).

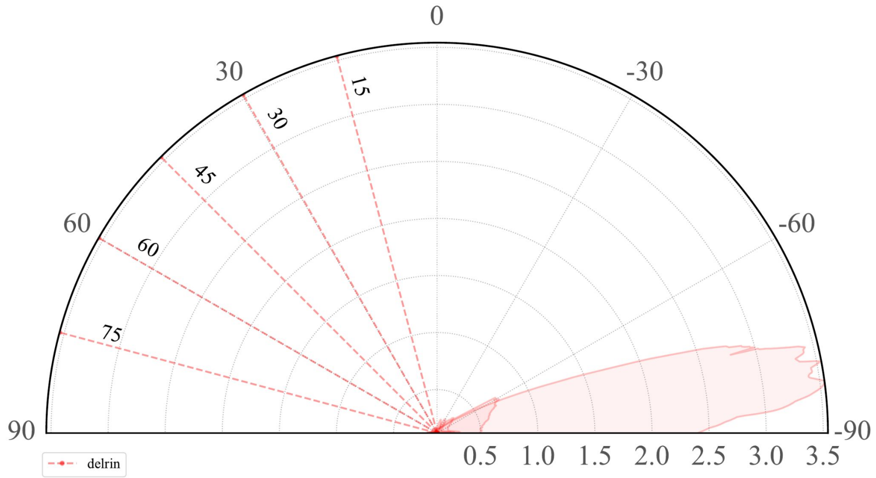

Figure: Polar plot of a real-world measured BRDF from the MERL dataset, for multiple fixed incident azimuth angles (15, 30, 45, 60, 75).

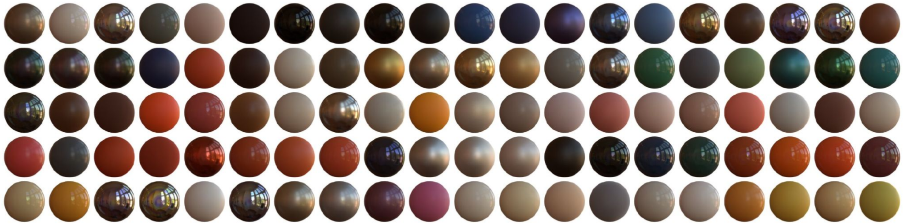

Figure: The MERL database contains 100 real-world measured isotropic materials.

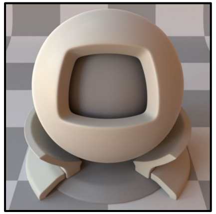

Delrin

Neural BRDF: Introduction

GT (Tabular)

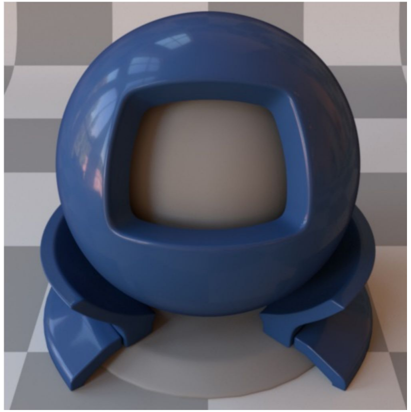

- Accurate
- Large storage (~34 MB)
- Manual interpolation

Analytic Model (GGX)

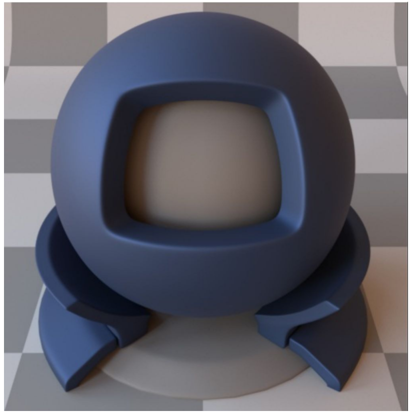

- Very inaccurate
- Very low storage (0.03 KB)
- Fast built-in interpolation
- Costly and unstable optimisation required

Neural BRDF: Representation

Neural BRDF Representation

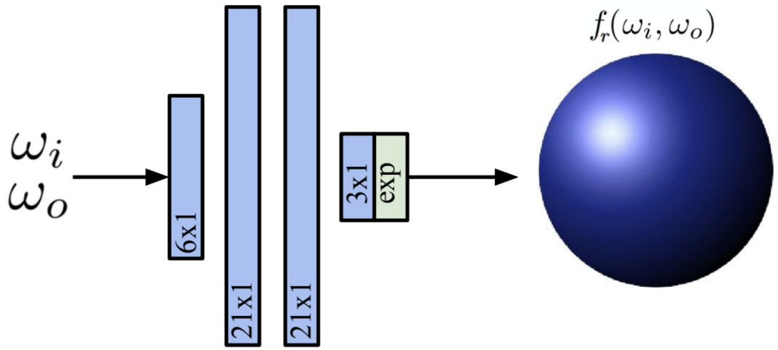

Figure: coordinate-based neural network (Neural BRDF). After training, the network encodes the BRDF function $f_r(\omega_i, \omega_o)$.

- Exponential activation.
- Rusinkiewicz parameterization of input.
- Rendering-based loss function.
- BRDF-aware sampling of light directions during training.

Neural BRDF

GT (Tabular)

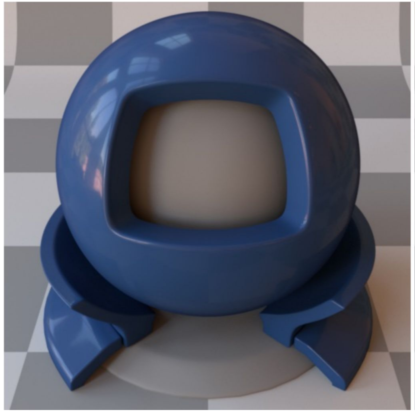

- Accurate
- Large storage (34 MB)
- Manual interpolation

Analytic Model (GGX)

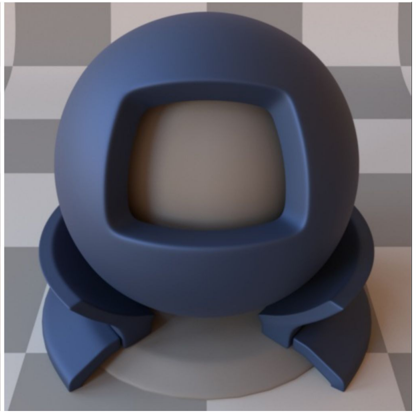

- Very inaccurate
- Very low storage (0.03 KB)
- Fast built-in interpolation
- Costly and unstable optimisation required

Neural-BRDF (Our)

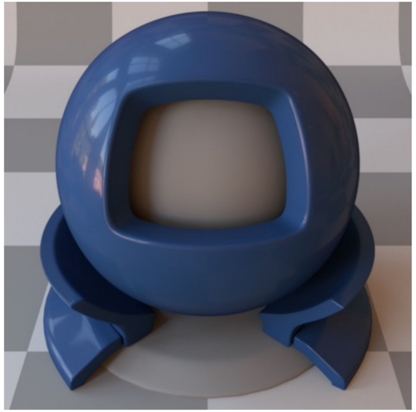

- Accurate
- Very low storage (2.7 KB)
- Fast built-in interpolation
- Costly but stable training

Neural BRDF: Reconstruction Accuracy

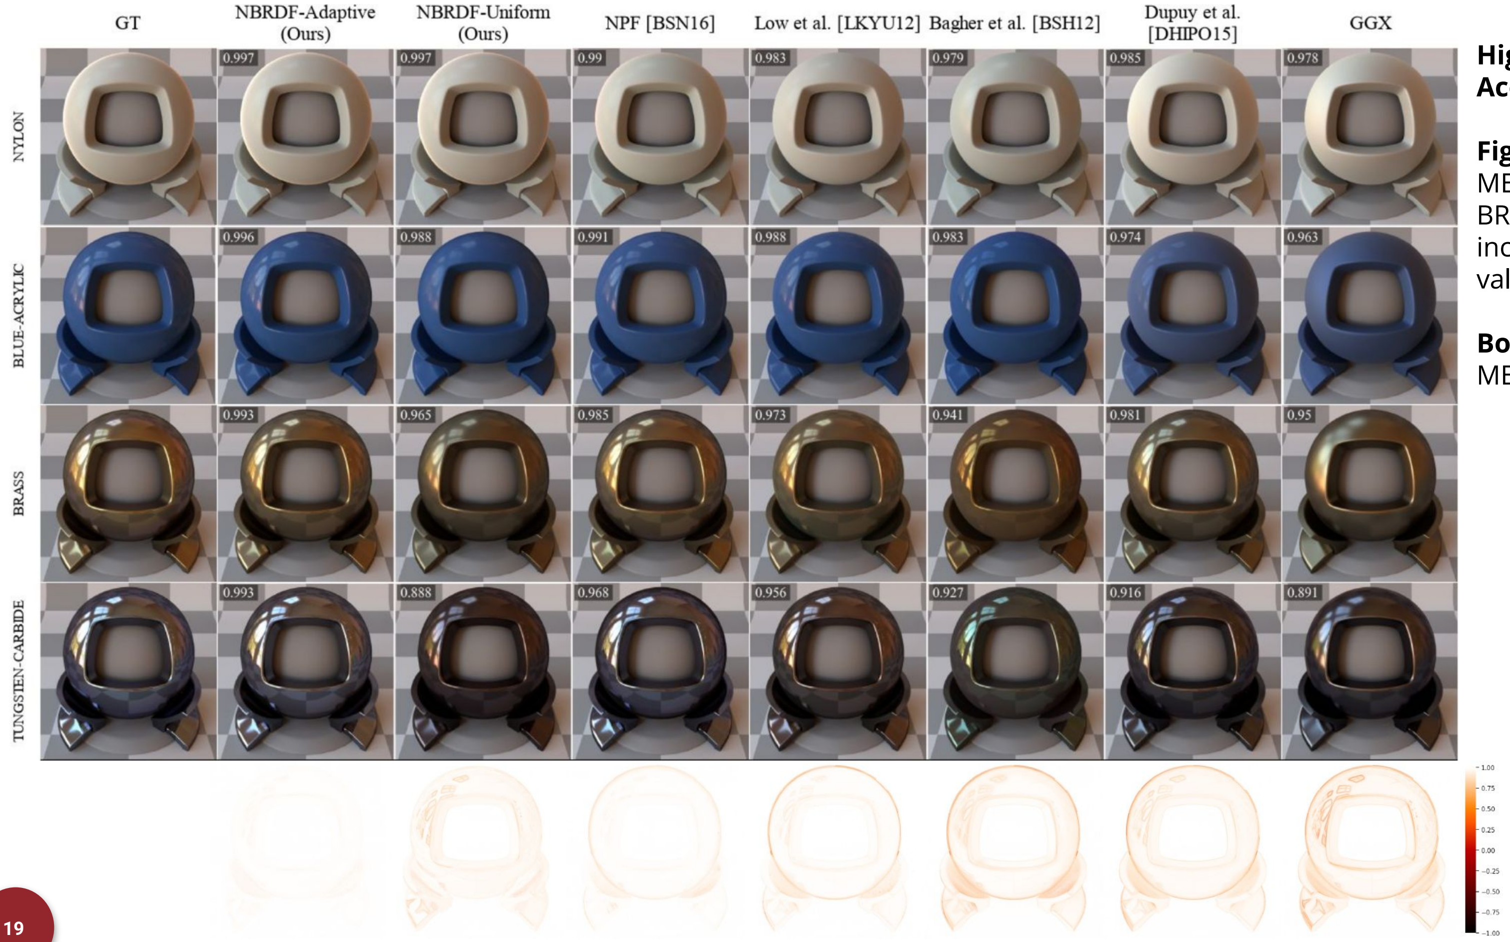

High Reconstruction Accuracy

Figure: Reconstruction of MERL materials using different BRDF representations, including the average SSIM value for each image.

Bottom: Average SSIM over all MERL materials.

Neural BRDF: Reconstruction Error

	MAE	RMSE	SSIM
NBRDF Adaptive Sampling	**0.0028 ± 0.0034**	**0.0033 ± 0.0038**	**0.995 ± 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

Table: Average image-based losses of BRDF representation methods over all MERL materials.

Neural BRDF: Compression and Speed

Reconstruction Error vs Representation Size

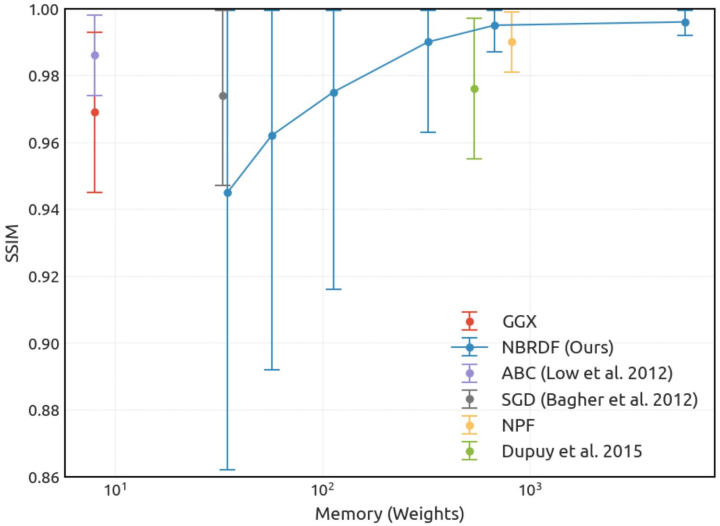

Figure: Average SSIM error vs Memory footprint (log scale) for multiple BRDF representations, with standard deviations. For NBRDFs (in blue), the reconstruction accuracy can be adjusted by modifying the network size.

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

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

Neural BRDF: Anisotropy

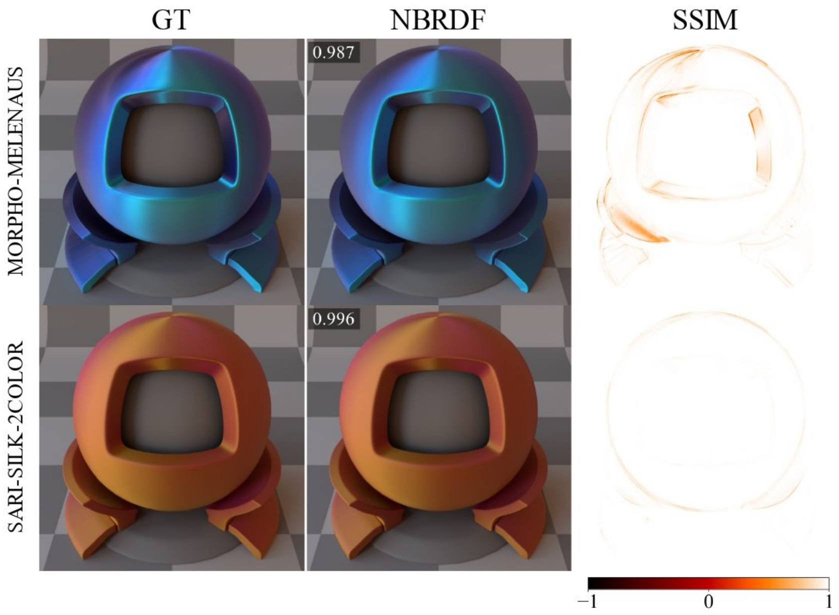

Support for Anisotropic Materials

Figure: Neural BRDF reconstruction of anisotropic materials from the RGL [Dupuy and Jakob 2018].

Neural BRDF: Hyper-Network

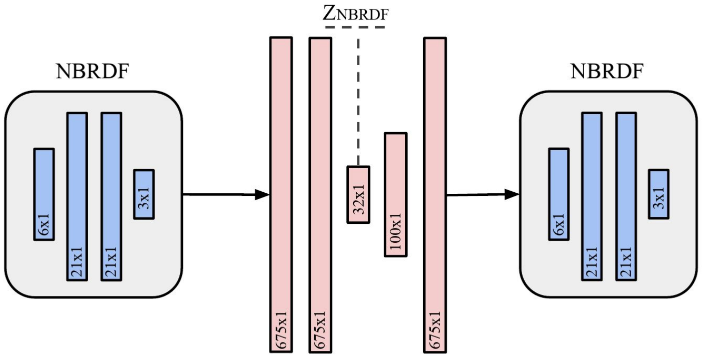

Figure: Neural BRDF autoencoder hyper-network.
Input and output are Neural BRDF network weights.

NBRDF Autoencoder (HyperNetwork)

- Training is done with NBRDF networks of the MERL database.
- Training loss: Instead of comparing NBRDF parameters, we implement a ***differentiable rendering loss*** that evaluates the GT and predicted Neural BRDFs, to generate renderings of a scene. The loss is then computed with an image-based error metric.
- Materials are encoded as 32-values vectors.

Neural BRDF: Generation of new materials.

Figure: t-SNE clustering of encodings for MERL materials, produced by the Neural BRDF hyper-network. Materials with similar reflectance properties cluster together. Test-set materials are indicated in red.

The generation of a unified encoding for the space of materials opens up multiple possible applications.

We show results for two applications:

1) Generation of new realistic materials through interpolation of the embeddings generated by the Neural BRDF hyper-network.

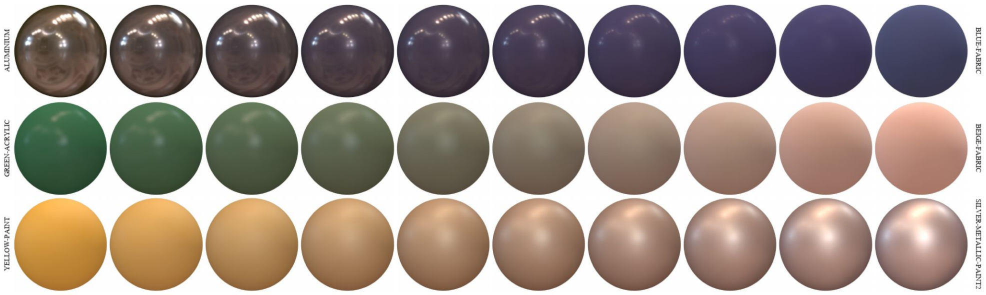

Figure: new materials generated by interpolation of encodings of MERL materials.

Neural BRDF: Importance Sampling

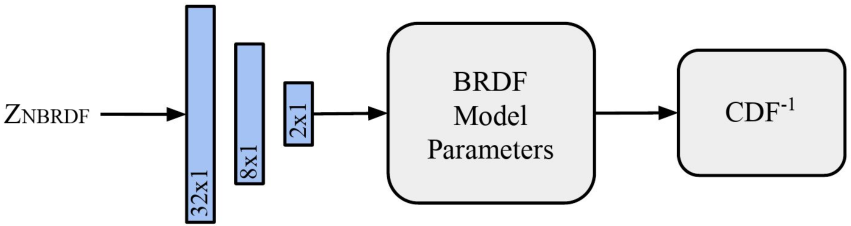

Figure: Scheme for computation of inverse CDF from an NBRDF.

2) BRDF Importance Sampling.

We train a small neural network to predict analytic BRDF model parameters, using the NBRDF embeddings as input.

This is essentially a neural-based BRDF fitting, but we only predict a limited number of model parameters, required for importance sampling. The target analytic model can be arbitrary, as long as its CDF is invertible.

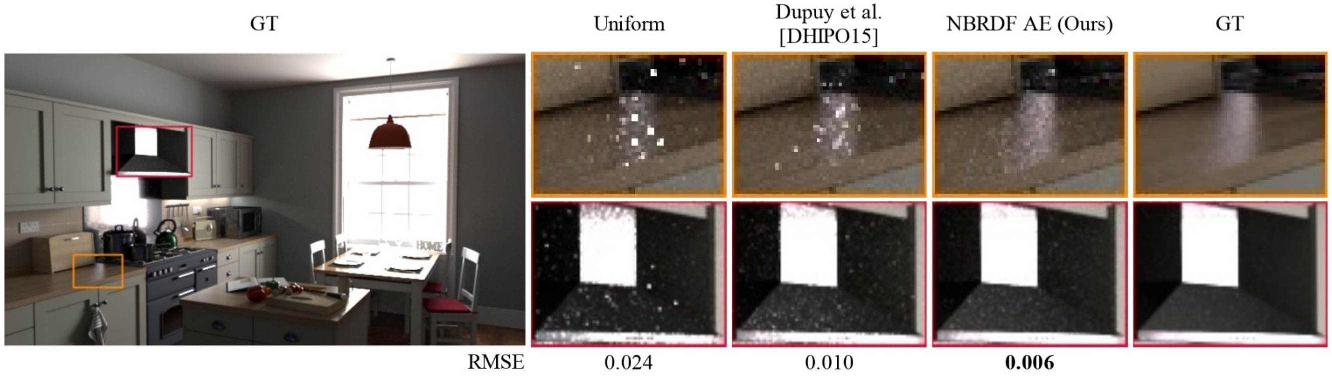

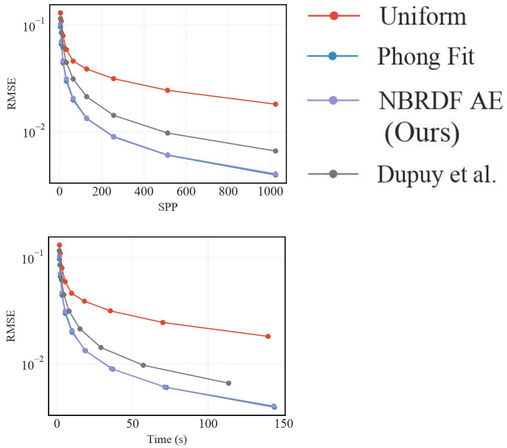

Figure: Importance sampling of kitchen scene using 64 SPP. Most materials in the scene have been replaced by test-set MERL materials.
Right: Average RMSE errors (log scale) vs SPP/render time.

Neural BRDF: Summary

- New neural-based representation for high-fidelity compression of measured BRDF data, supporting isotropic and anisotropic materials.

- Comparison of Neural BRDF with other BRDF representations, in terms of reconstruction accuracy, evaluation speed and memory footprint.

- Implementation of a hyper-network autoencoder architecture with a differentiable rendering loss to explore the space of real-world materials by learning latent representations of the Neural BRDF networks.

- Further compression of the BRDF data to 32-values encodings, which can be smoothly interpolated to create new realistic materials.

- Learned a mapping between our neural representation and an invertible analytic BRDF model, enabling the importance sampling of Neural BRDFs for efficient rendering.

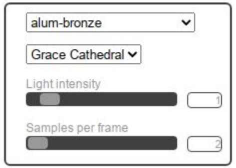

HyperTime: Neural Fields for Interpretable Time Series Generation

HyperTime: Motivation

Time-Series have many applications

- Climate, Biology, Medicine, Finance, etc.
- Many works using neural networks on time-series data (classification, forecasting, etc.).
- In particular, synthetic generation of time-series is used to augment training datasets and improve performance on downstream tasks.

Neural Fields are a great match for time-series data:

- Good for periodic signals.
- Good for representing a wide spectrum of frequencies.
- Grid-free representation: good for missing data and irregularly sampled datasets.

HyperTime: Univariate and Multivariate Time-Series

Univariate Time-Series Network

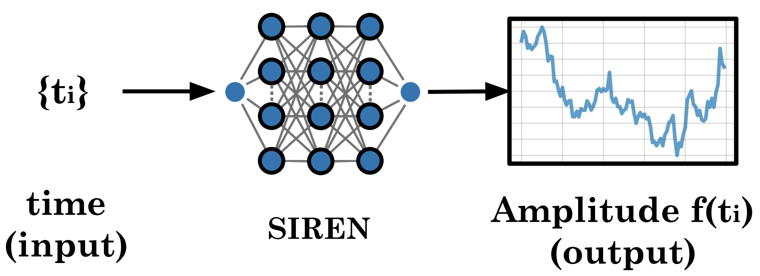

Figure: Diagram of network for univariate time-series representation. The network is composed of fully-connected layers with sine activations.

Multivariate Time-Series Network

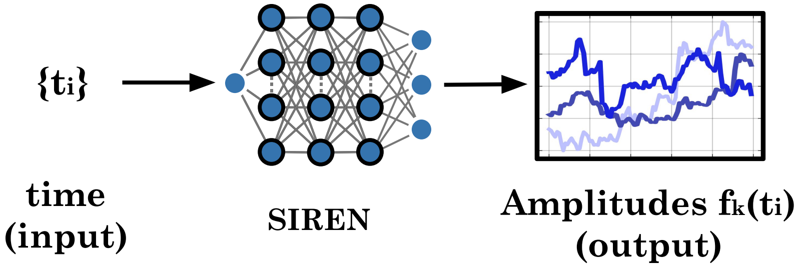

Figure: Diagram of network for multivariate time-series representation. The number of output neurons matches the number of channels of the time-series (3).

HyperTime: Interpretable Decomposition (iSIREN)

However, we want to introduce interpretability into the model.

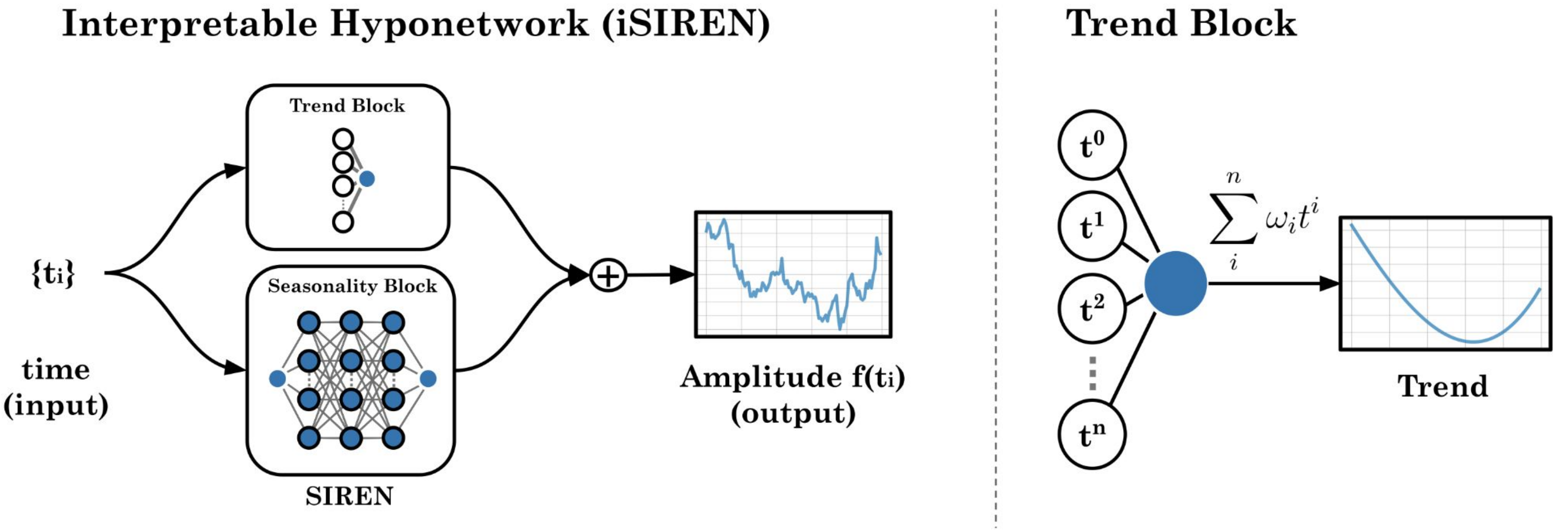

UCL

HyperTime: Reconstruction Results

Table 1: Comparison using MSE on time space and MAE in frequency space (FFT) of implicit networks using different activation functions and of iSIREN on univariate and multivariate datasets.

Dataset	iSIREN (Ours)		SIREN		P.E.		ReLU		Tanh	
	FFT	MSE	FFT	MSE	FFT	MSE	FFT	MSE	FFT	MSE
Univariate										
Crop	1.4e-3	5.6e-6	1.4e-3	1.6e-6	**6.8e-4**	**7.3e-7**	5.4e-1	2.1e-2	8.6e-1	6.0e-2
Energy	**4.1e-3**	**5.3e-6**	1.8e-2	1.2e-5	1.3e-1	7.7e-4	1.5e+0	4.9e-2	1.9e+0	8.3e-2
FordA	**1.7e-2**	**4.9e-6**	1.9e-2	6.2e-6	3.1e-1	2.1e-3	2.5e+0	1.3e-1	2.8e+0	1.4e-1
NonInv	**3.6e-2**	**1.2e-5**	4.0e-2	1.3e-5	1.1e-1	1.3e-4	1.0e+0	2.2e-2	1.3e+0	4.6e-2
Phalanges	**1.4e-3**	**2.1e-6**	1.8e-6	7.6e-3	1.2e-5	2.4e-1	3.8e-3	7.5e-1	8.4e-2	3.4e-1
Stock	**2.5e-3**	5.1e-6	4.4e-3	**1.4e-6**	4.3e-2	1.2e-4	6.2e-1	1.2e-2	8.9e-1	3.8e-2
Multivariate										
Cricket	**3.9e-1**	**4.1e-4**	4.5e-1	4.2e-4	1.7e+0	3.7e-3	3.5e+0	1.7e-2	3.9e+0	3.1e-2
MotorImagery	**5.1e+0**	**2.1e-3**	7.2e+0	6.2e-3	1.1e+1	2.4e-2	1.0e+1	2.6e-2	1.1e+1	3.0e-2
PhonemeSpectra	**2.9e-2**	**2.1e-6**	4.2e-1	2.7e-4	1.8e+0	5.9e-3	3.0e+0	1.5e-2	3.4e+0	2.0e-2

HyperTime: Trend-Seasonality Decomposition

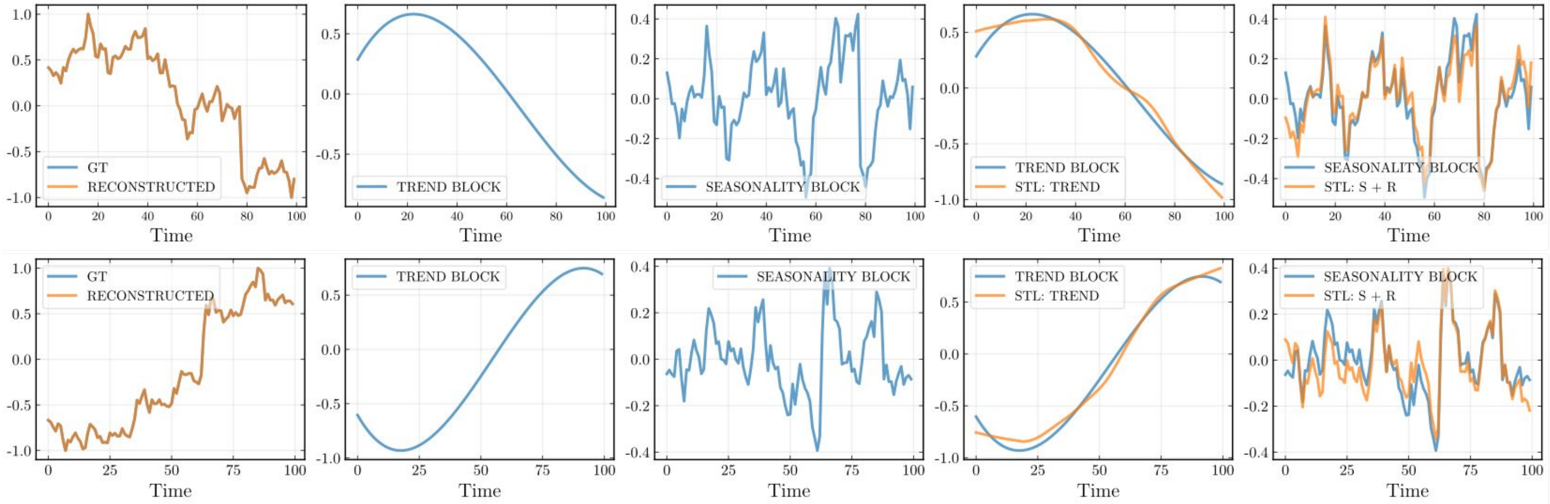

(a) Reconstruction (b) Trend (c) Seasonality (d) STL: trend (e) STL: seas

Figure Interpretable decomposition of iSIREN on two time series from the `Stock` dataset. (a) Ground truth and iSIREN reconstruction. (b) Trend Block output. (c) Seasonality Block output. Columns (d) and (e) compare the output of iSIREN blocks with classic STL decomposition.

HyperTime: Interpretable Generation of Time-Series

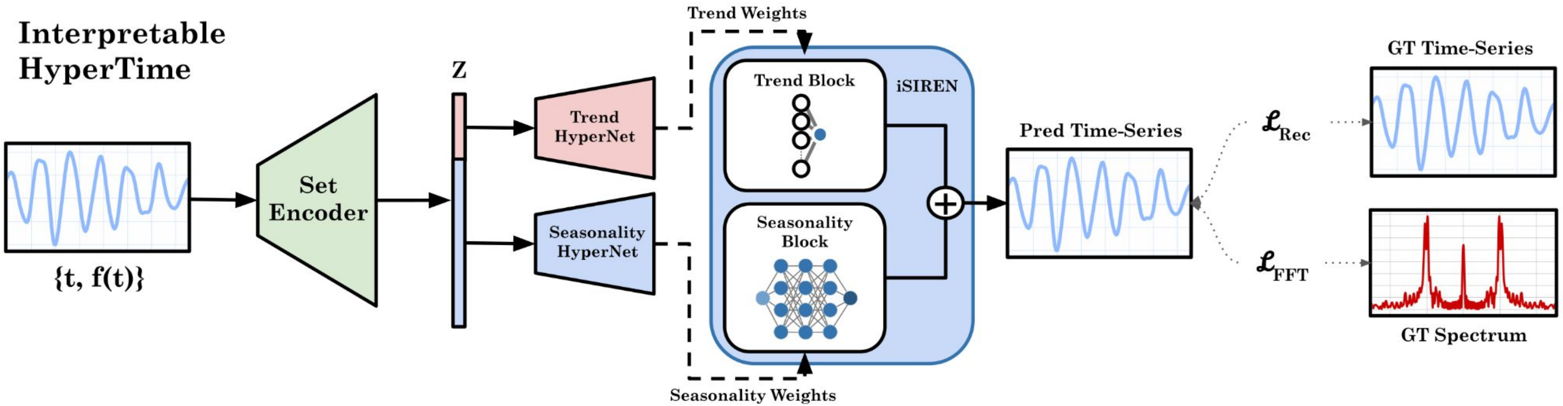

$$\mathcal{L} = \underbrace{\frac{1}{N}\sum_{i=1}^{N}\left\|f(t_i) - \hat{f}(t_i)\right\|^2}_{\mathcal{L}_{\text{rec}}} + \lambda_1 \underbrace{\frac{1}{N}\sum_{i=1}^{N}\left\|\text{FFT}[f(t)]_i - \text{FFT}[\hat{f}(t)]_i\right\|}_{\mathcal{L}_{\text{FFT}}}$$

HyperTime: Trend-Seasonality Split Generation

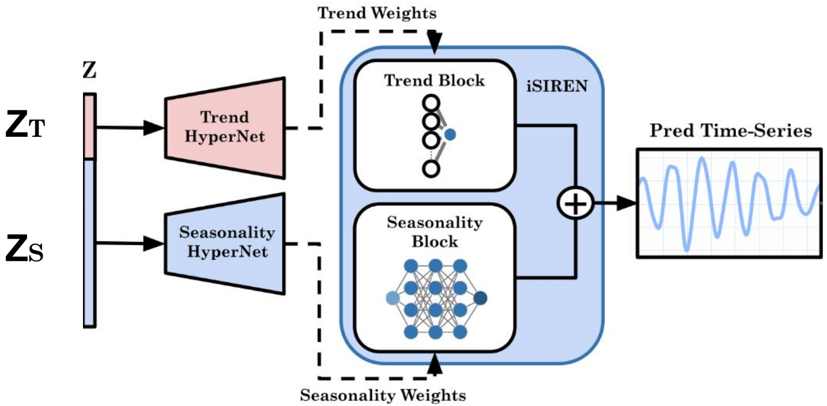

HyperTime: Trend-Seasonality Split Generation

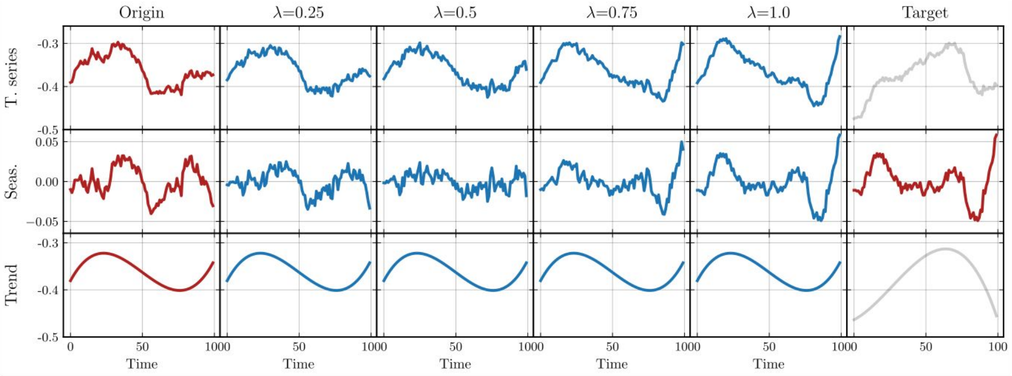

(a) Interpolation in seasonality, leaving trend fixed.

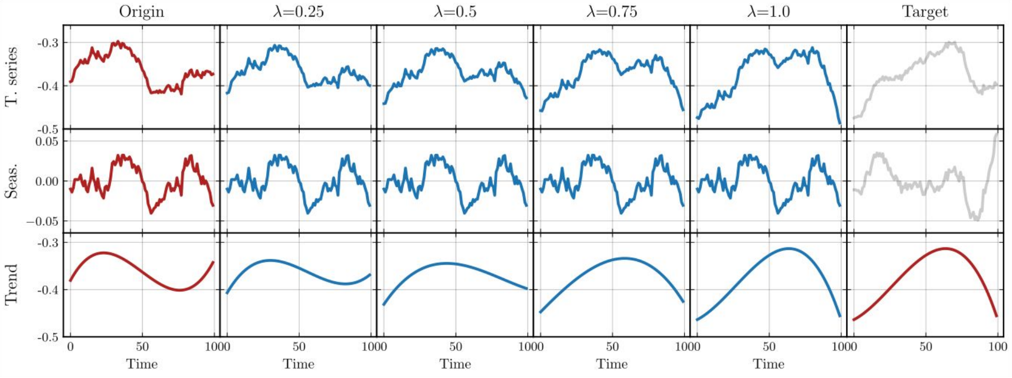

(b) Interpolation in trend, leaving seasonality fixed.

Figure (a) Interpolation of seasonality, with fixed trend. (b) Interpolation in trend, with fixed seasonality. In red: original TS (1st column) and target seasonality/trend (last column).

HyperTime: Qualitative Evaluation

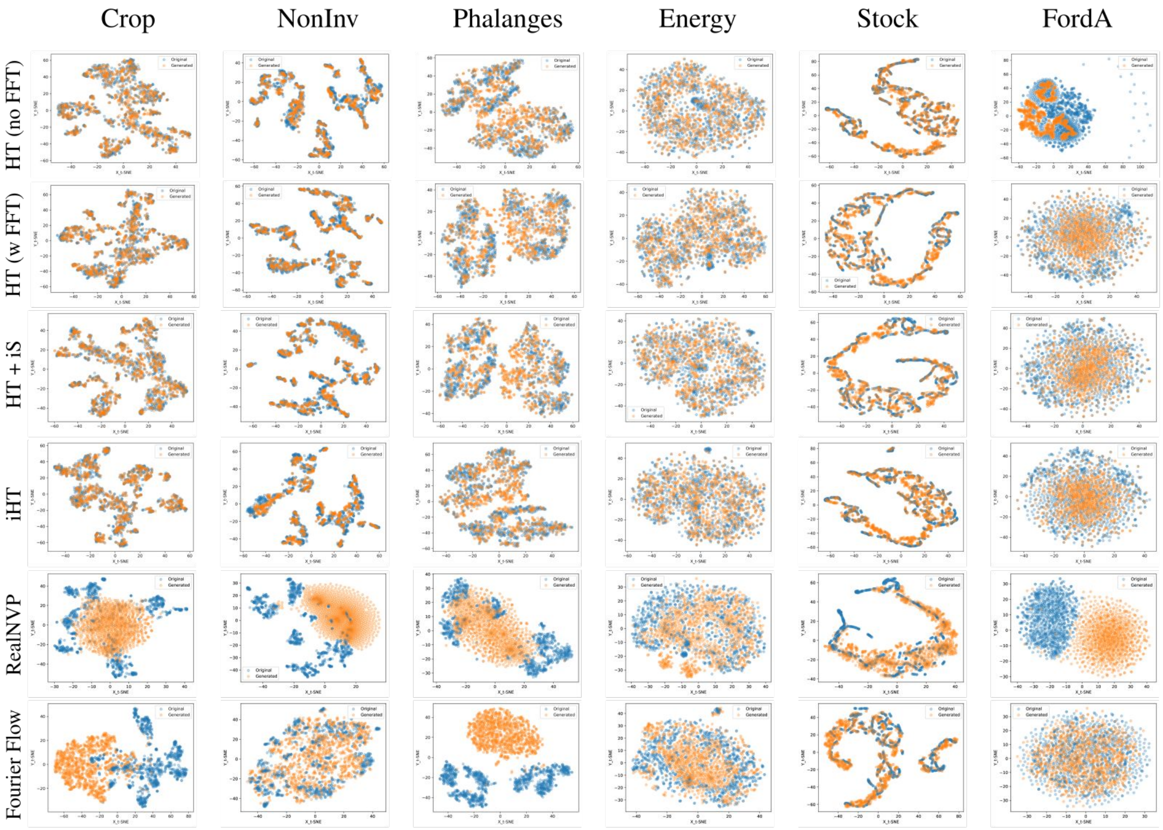

Figure t-SNE visualization of real (blue) and synthetic (orange) data for all univariate datasets (in columns), using different time series generation methods (in rows).

HyperTime: Quantitative Evaluation

	Crop	NonInv	Phalan.	Energy	Stock	FordA
RealNVP						
MAE	0.170	0.038	0.073	0.036	0.019	0.115
F1 Score	0.981	0.986	0.976	0.964	0.977	0.999
TimeGAN						
MAE	0.048	–	0.108	0.056	0.173	–
F1 Score	0.831	–	0.960	0.479	0.938	–
Fourier Flows						
MAE	0.040	0.018	0.056	**0.030**	**0.010**	0.024
F1 Score	0.991	0.990	0.992	0.936	0.990	0.998
HyperTime						
HT (no FFT)						
MAE	0.040	0.005	**0.023**	0.058	0.012	0.17
F1 Score	**0.999**	0.996	0.996	**0.998**	**0.995**	0.084
HT (w/ FFT)						
MAE	0.040	0.005	**0.023**	0.057	0.013	**0.007**
F1 Score	**0.999**	**0.997**	**0.999**	0.997	0.994	**0.998**
HT (iSiren)						
MAE	**0.039**	**0.004**	0.024	0.057	0.013	0.008
F1 Score	**0.999**	**0.997**	**0.999**	0.997	**0.995**	0.997
iHT						
MAE	**0.039**	**0.004**	0.024	0.056	0.011	0.009
F1 Score	**0.999**	**0.997**	0.997	0.997	**0.995**	0.996

Quantitative Metrics:

- MAE (Predictive Score)
- F1-Score (Quality)

Table 2: Performance scores for data generation using baselines (TimeGAN, Fourier Flows, RealNVP) and multiple hypernet models: **HT (no FFT):** SIREN hyponetwork, trained without spectral loss. **HT (w/FFT):** SIREN hyponetwork. **HT (iSIREN):** iSIREN hyponetwork. **iHT:** interpretable HT.

UCL

HyperTime: Summary

- **Introduced an interpretable NF architecture for univariate and multivariate time-series representation.**

- **Compared iSIREN with other models in terms of reconstruction performance.**

- **Proposed HyperTime, a hypernetwork architecture that allows learning a prior from an entire dataset of time series.**

- **Introduced a spectral loss to guide HyperTime training.**

- **Introduced a modification of HyperTime to introduce interpretability into the latent representation, enabling the potential injection of expert knowledge into the generation process.**

Thank you for listening!

Dr. Alejandro Sztrajman

Oct 17, 2022