

Path to Neural Networks II



Image courtesy Vogel et al. [2018], Gharbi et al. [2019]

Today's Menu

Sample-based denoising

CNN-based approach to generate blue-noise samples

Normalizing Flows

Path guiding using Normalizing Flows


Recap




input Monte Carlo (8 samples/pixel)




after RPF (8 samples/pixel)



(a) Input MC (8 spp)

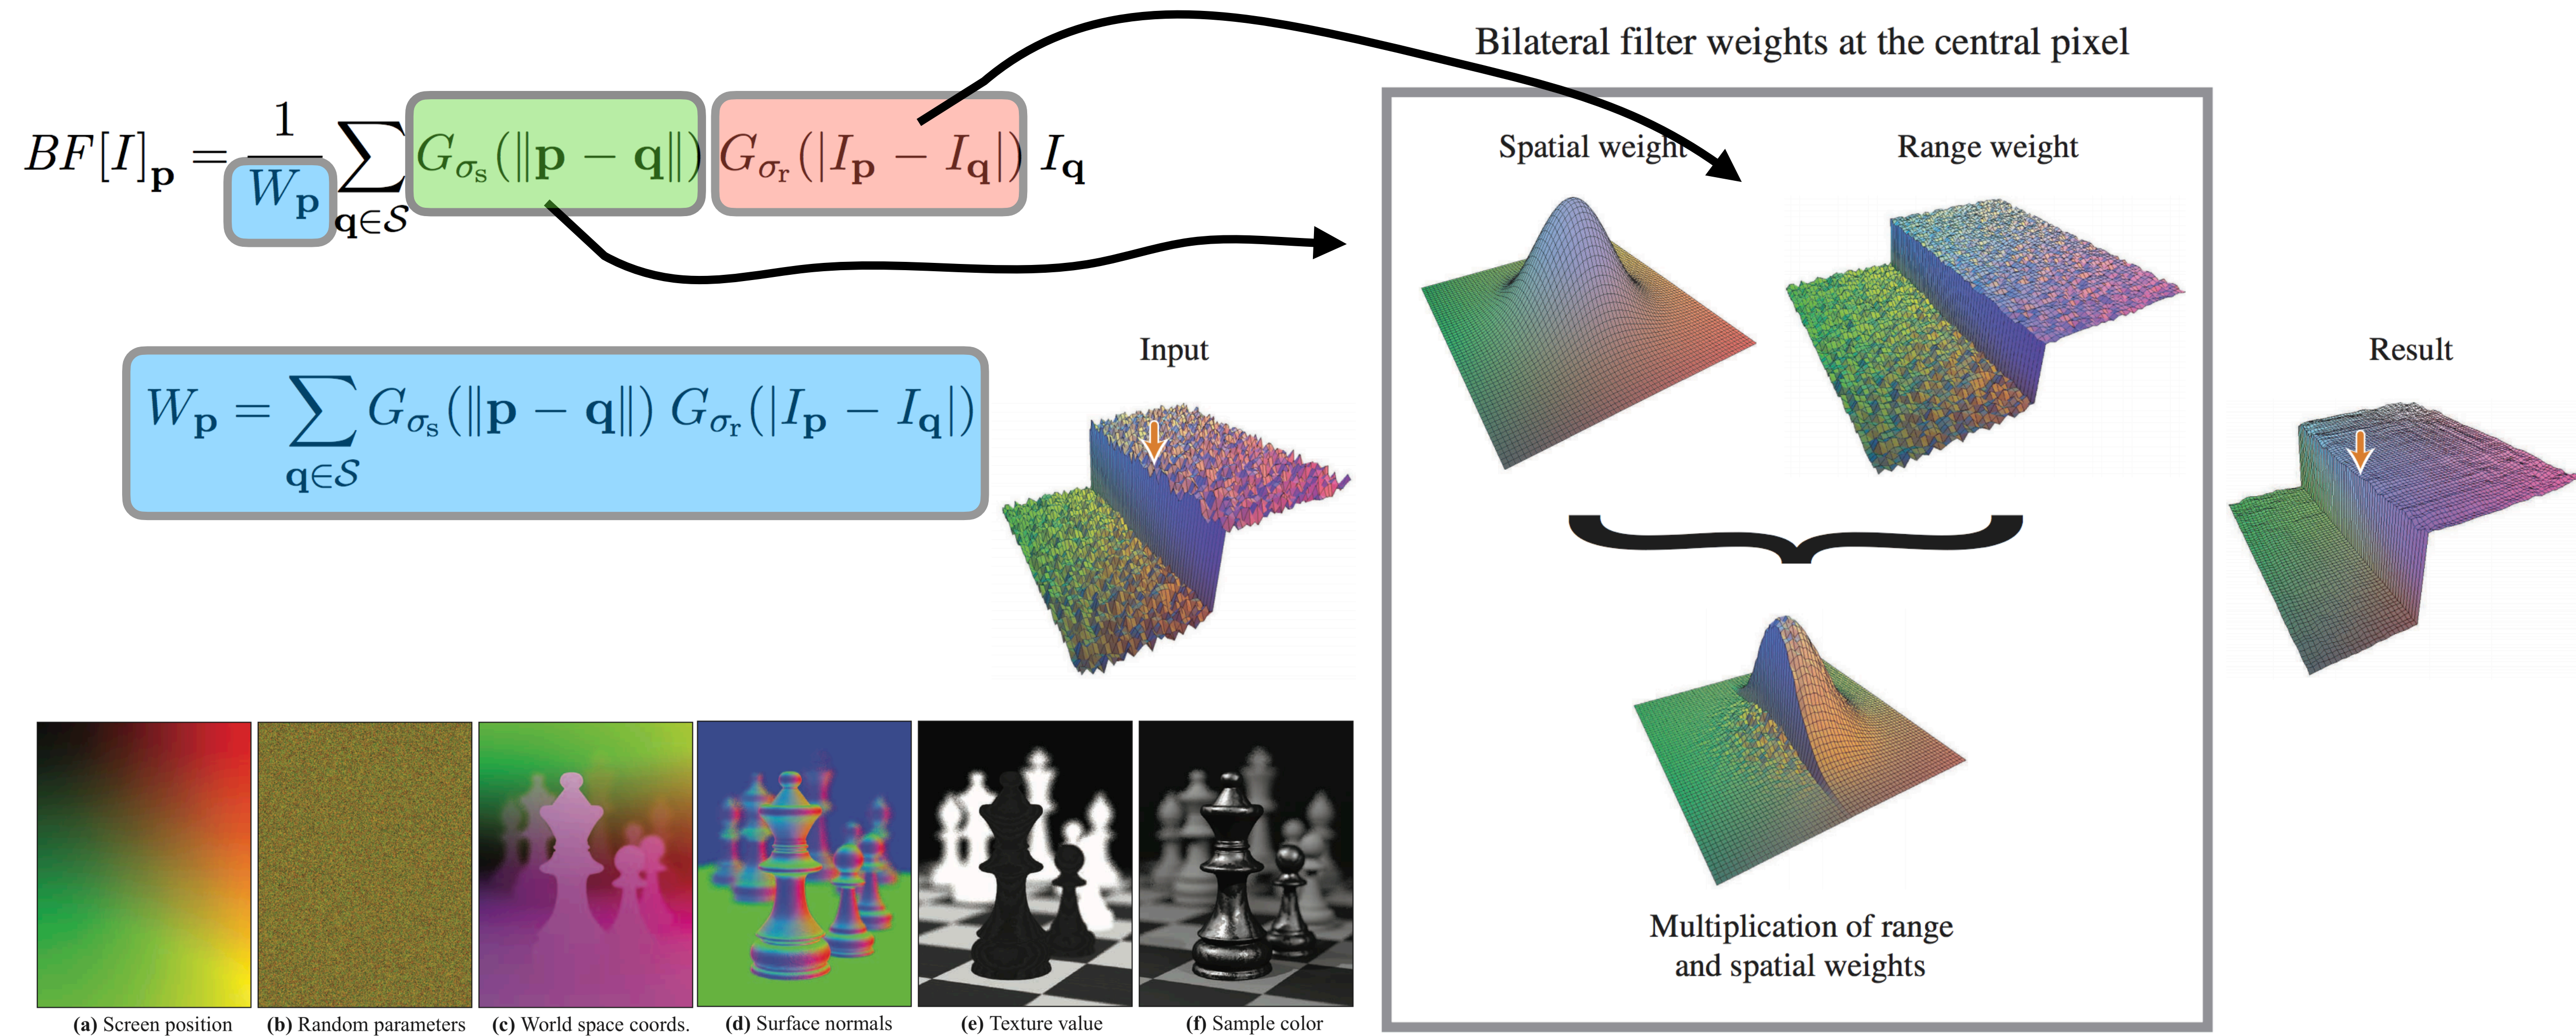


(b) Dependency on (u, v)



(c) Our approach (RPF)

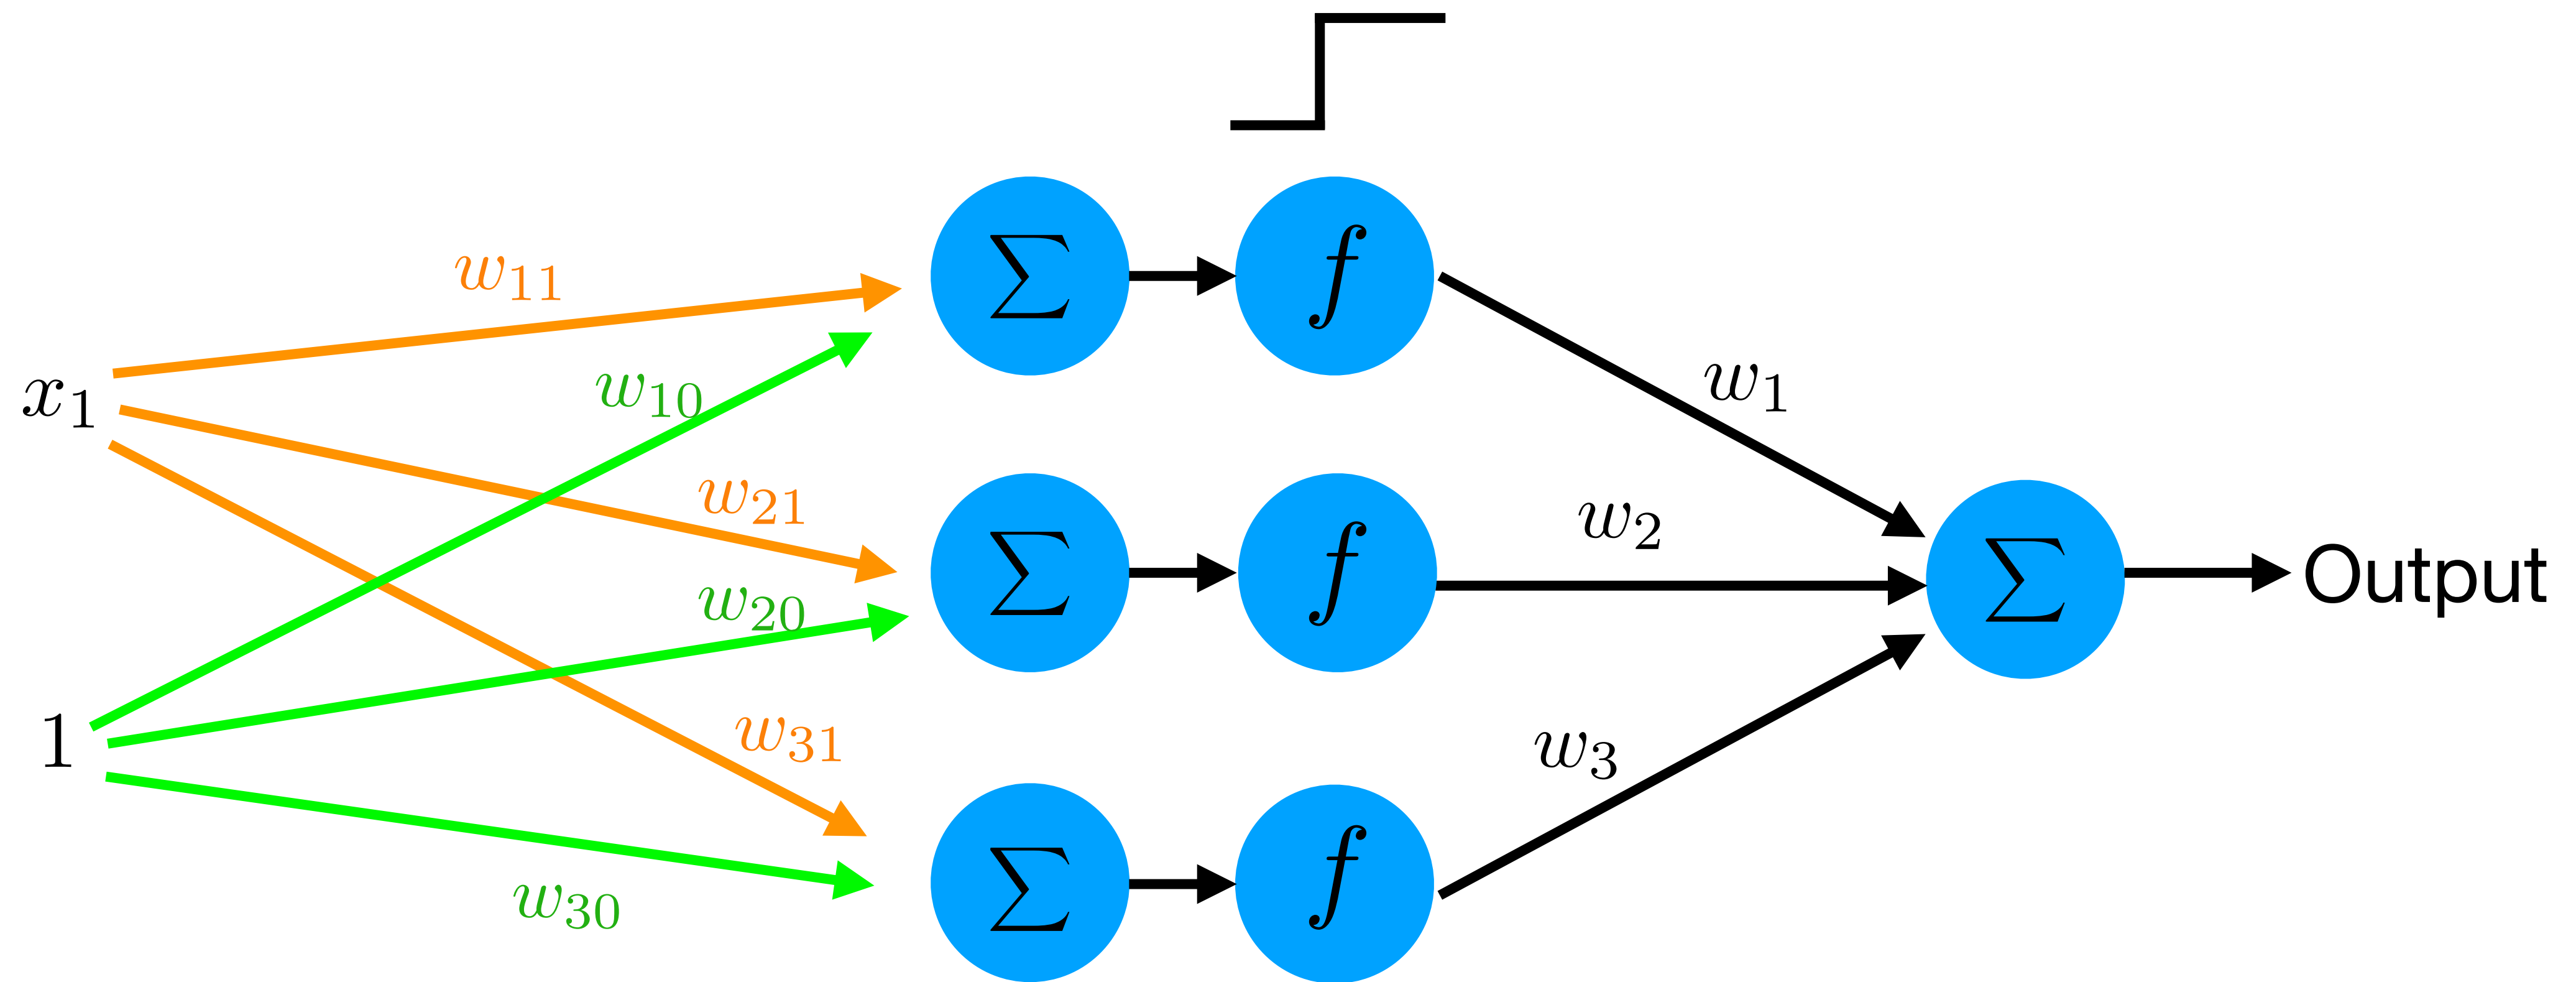
Bilateral Filtering



Bilateral Filtering of Features

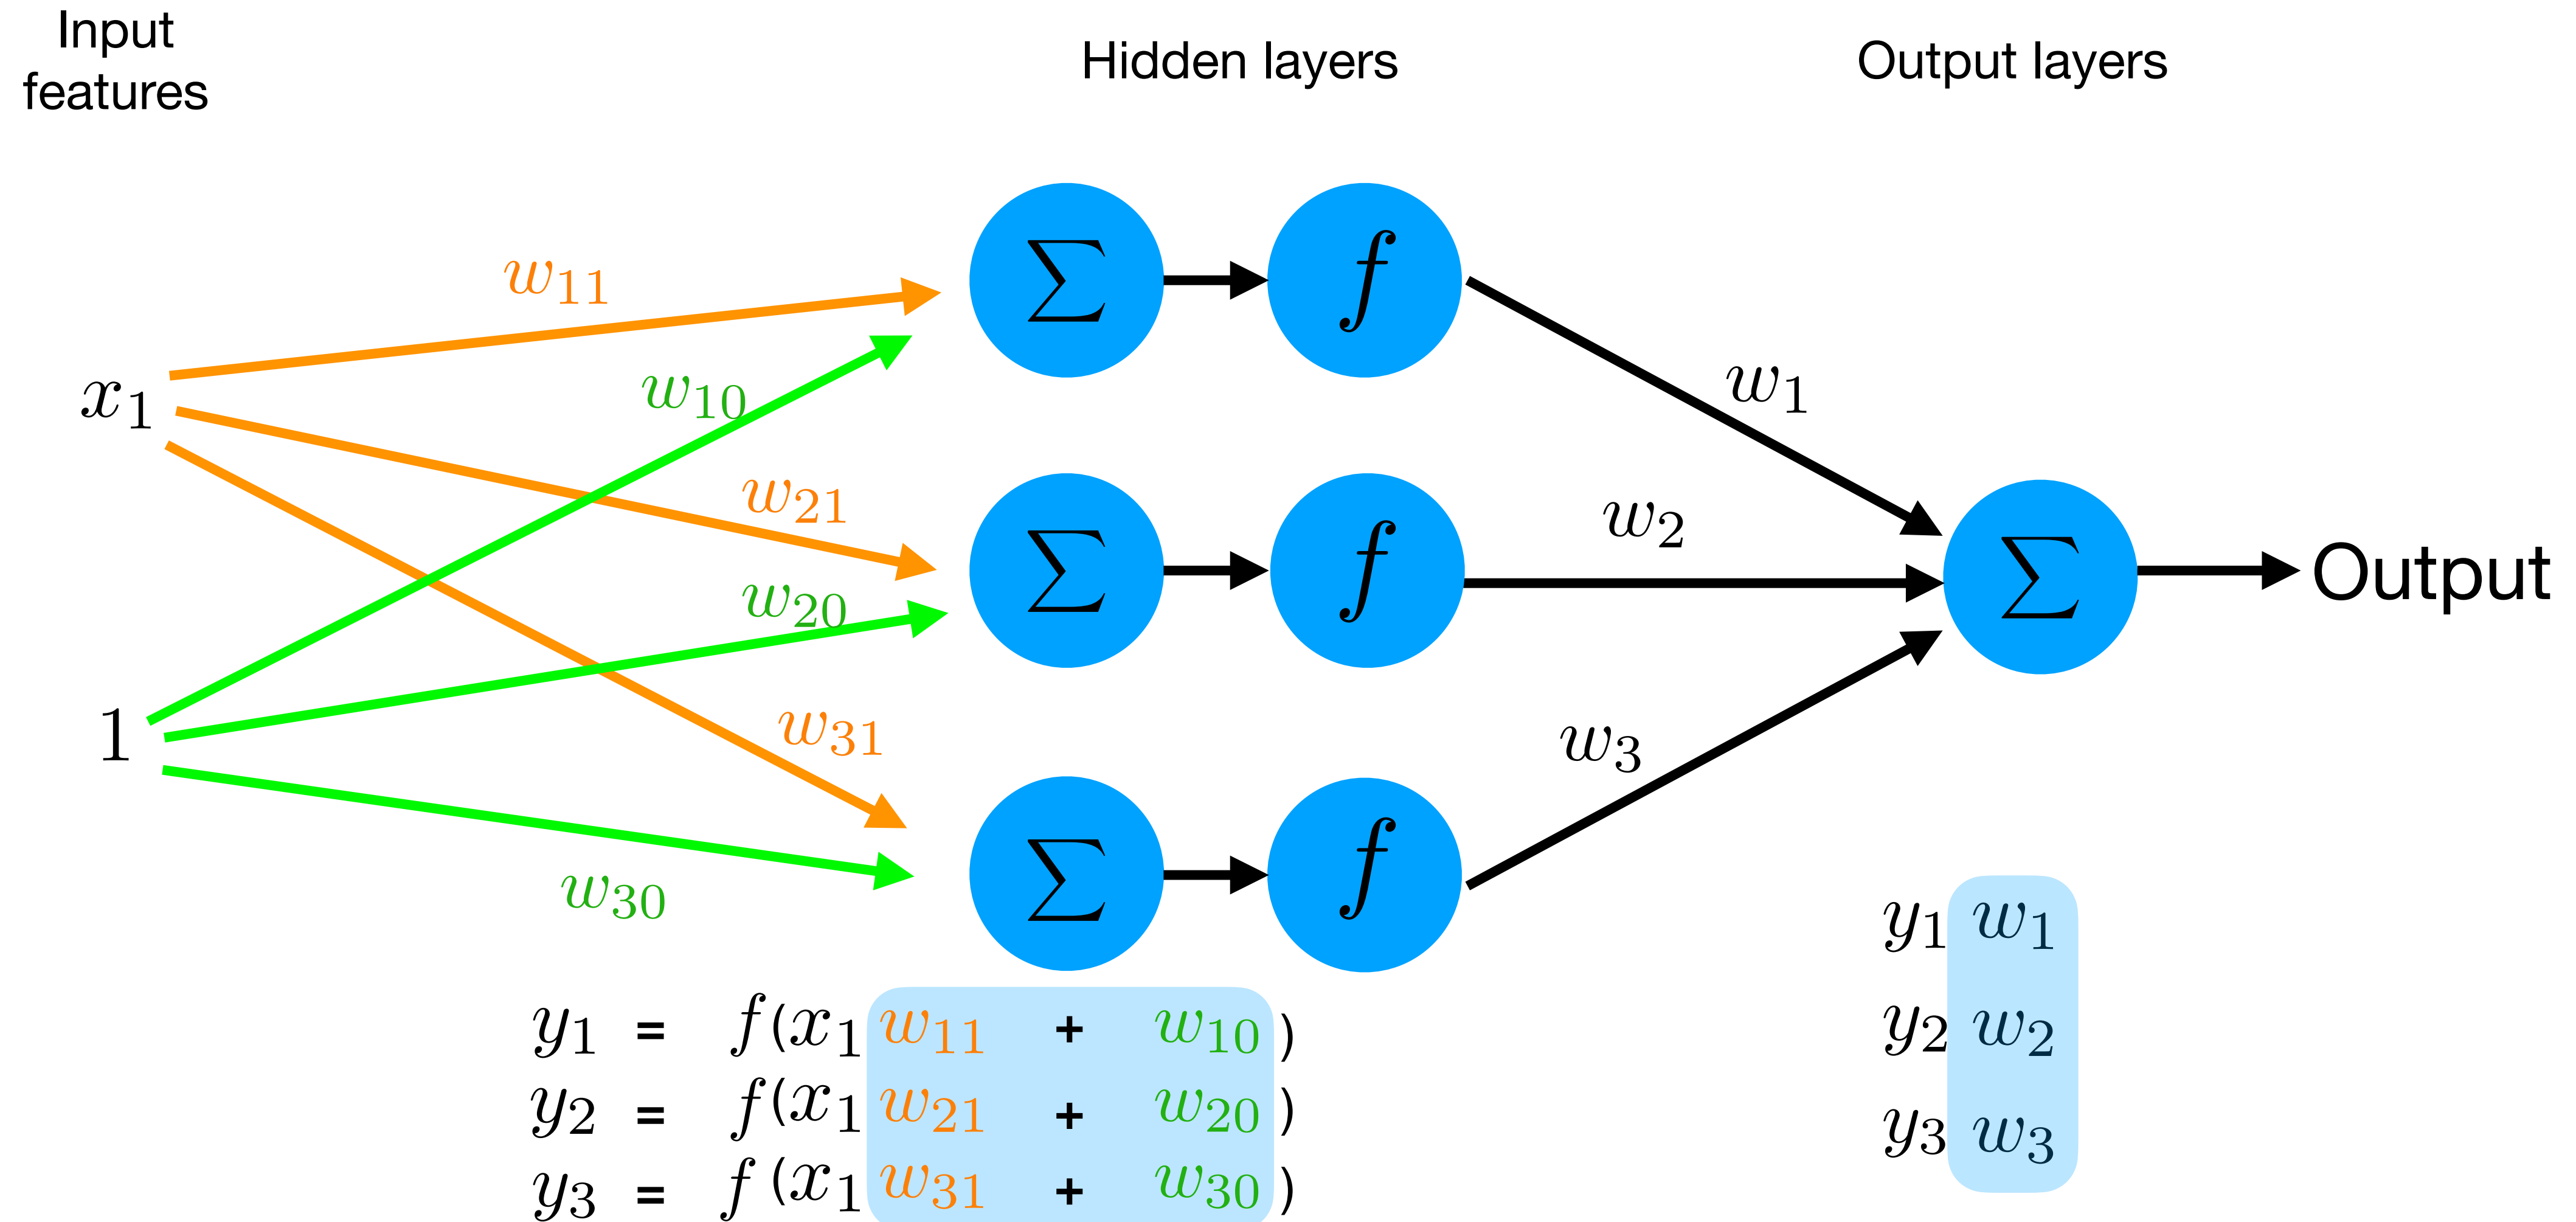
$$w_{ij} = \exp\left[-\frac{1}{2\sigma_{\mathbf{p}}^2} \sum_{1 \leq k \leq 2} (\bar{\mathbf{p}}_{i,k} - \bar{\mathbf{p}}_{j,k})^2\right] \times$$
$$\exp\left[-\frac{1}{2\sigma_{\mathbf{c}}^2} \sum_{1 \leq k \leq 3} \alpha_k (\bar{\mathbf{c}}_{i,k} - \bar{\mathbf{c}}_{j,k})^2\right] \times$$
$$\exp\left[-\frac{1}{2\sigma_{\mathbf{f}}^2} \sum_{1 \leq k \leq m} \beta_k (\bar{\mathbf{f}}_{i,k} - \bar{\mathbf{f}}_{j,k})^2\right],$$

Multi-layer Perceptron



$$\begin{aligned}
 y_1 &= f(x_1 w_{11} + w_{10}) \\
 y_2 &= f(x_1 w_{21} + w_{20}) \\
 y_3 &= f(x_1 w_{31} + w_{30})
 \end{aligned}$$

Multi-layer Perceptron



Filter weights


For cross Bilateral filters:

$$d_{i,j} = \exp \left[- \frac{\|\bar{\mathbf{p}}_i - \bar{\mathbf{p}}_j\|^2}{2\alpha_i^2} \right] \times \exp \left[- \frac{D(\bar{\mathbf{c}}_i, \bar{\mathbf{c}}_j)}{2\beta_i^2} \right] \times \prod_{k=1}^K \exp \left[- \frac{D_k(\bar{\mathbf{f}}_{i,k}, \bar{\mathbf{f}}_{j,k})}{2\gamma_{k,i}^2} \right],$$

Pixel screen coordinates

Mean sample color value

Scene features



(a) Screen position



(b) Random parameters

(c) World space coords.

(d) Surface normals

(e) Texture value

(f) Sample color



Our result with a cross-bilateral filter (4 spp)

Overview on Convolutional Neural Networks (CNNs)

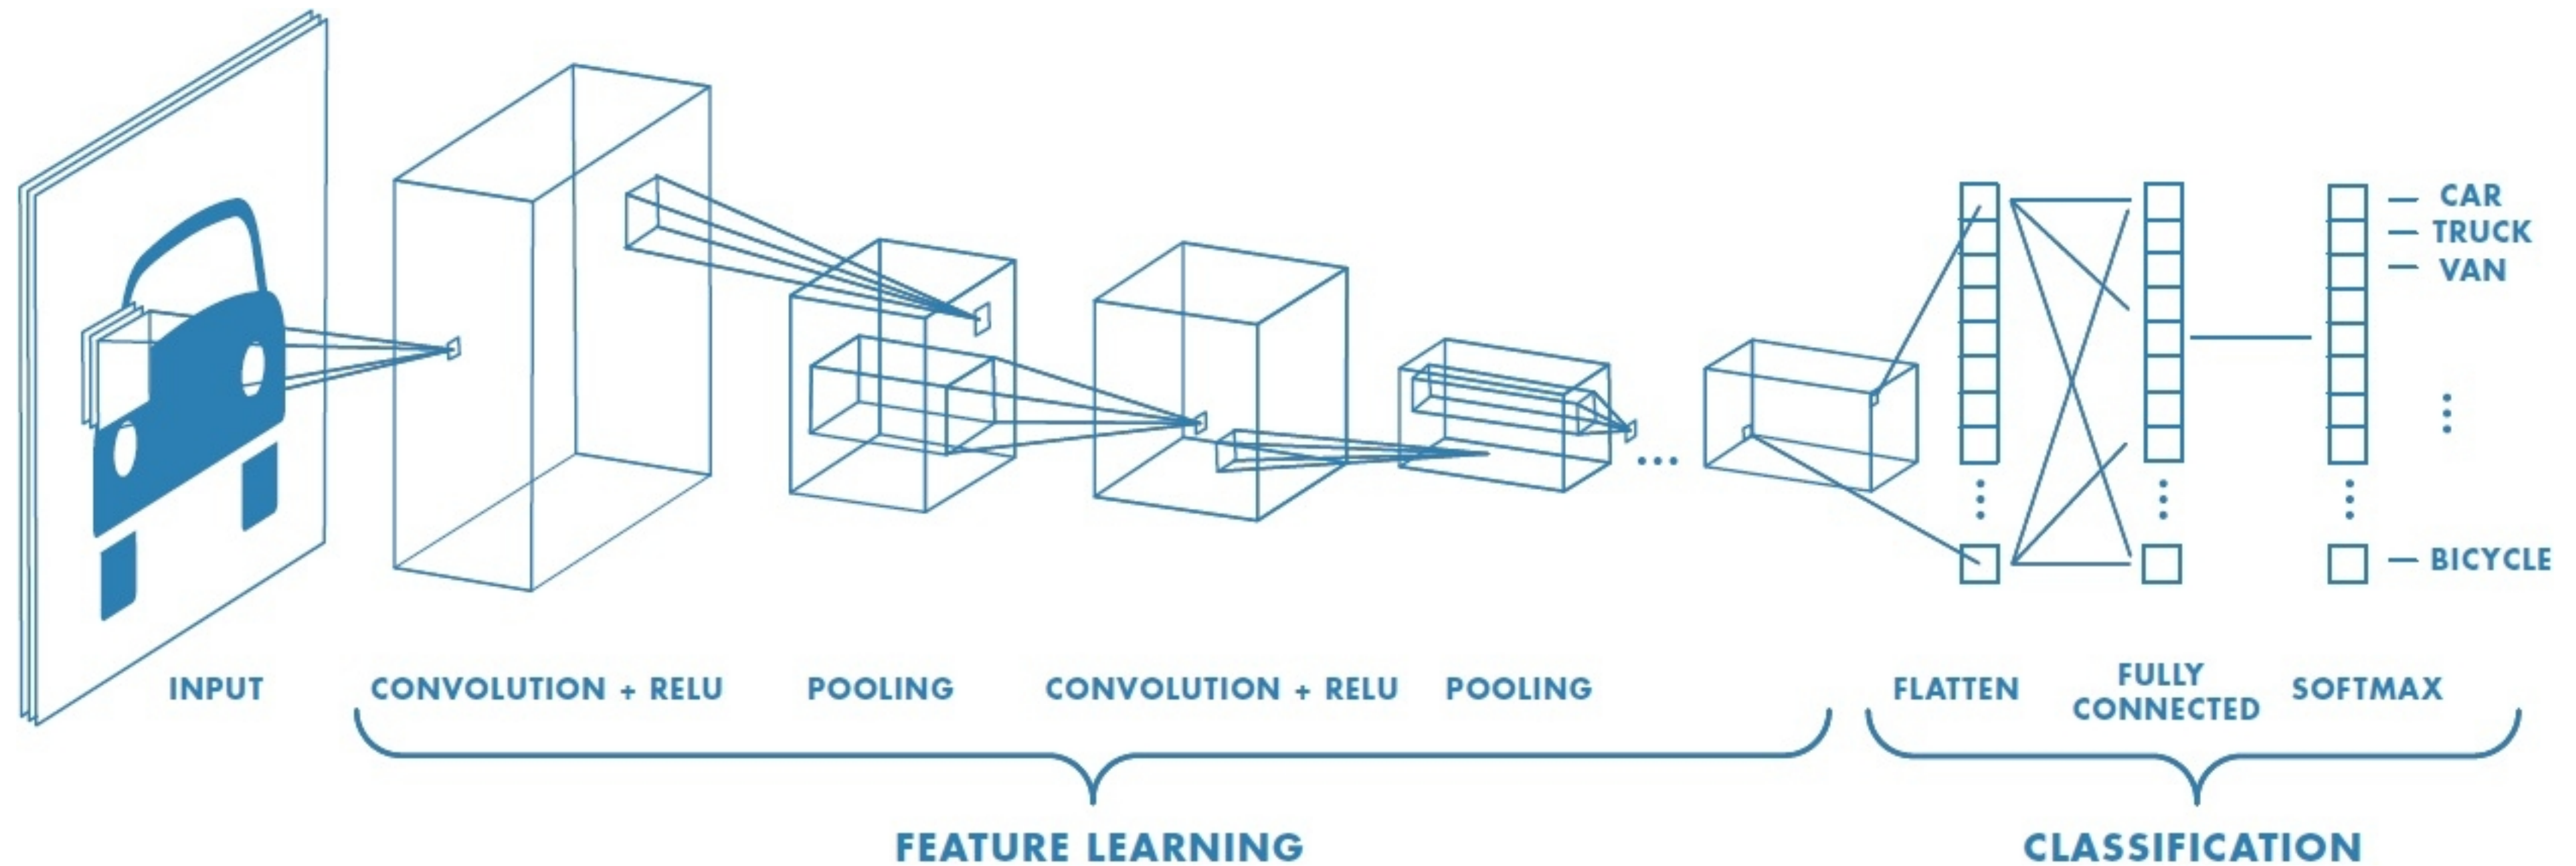



Image Courtesy: Mathworks (online tutorial)

Multi-layer Perceptron vs. CNNs

Multi-layer perceptron




All nodes are fully connected in all layers

In theory, should be able to achieve good quality results in small number of layers.

Number of weights to be learnt are very high

CNNs




Weights are shared across layers

Requires significant number of layers to capture all the features (e.g. Deep CNNs)

Relatively small number of weights required

Kernel-Predicting Networks for Denoising Monte-Carlo Renderings



Recurrent AutoEncoder for Interactive Reconstruction






Fig. 2. Architecture of our recurrent autoencoder. The input is 7 scalar values per pixel (noisy RGB, normal vector, depth, roughness). Each encoder stage has a convolution and 2×2 max pooling. A decoder stage applies a 2×2 nearest neighbor upsampling, concatenates the per-pixel feature maps from a skip connection (the spatial resolutions agree), and applies two sets of convolution and pooling. All convolutions have a 3×3 -pixel spatial support. On the right we visualize the internal structure of the recurrent RCNN connections. I is the new input and h refers to the hidden, recurrent state that persists between animation frames.

Recurrent Neural Networks vs. Simple Feed-Forward NN

[Source link](#)



Recurrent Neural Network



Feed-Forward Neural Network

Loss Functions

Spatial Loss to emphasize more the dark regions

$$L_s = \frac{1}{N} \sum_i^N |P_i - T_i|$$

Temporal loss

$$L_t = \frac{1}{N} \sum_i^N \left(\left| \frac{\partial P_i}{\partial t} - \frac{\partial T_i}{\partial t} \right| \right)$$

High frequency error norm loss for stable edges

$$L_g = \frac{1}{N} \sum_i^N |\nabla P_i - \nabla T_i|$$

Final Loss is a weighted averaged of above losses

$$L = w_s L_s + w_g L_g + w_t L_t$$

**Pixel-space
Kernel Predicting
Denoising**

**#Learnable
Parameters?**

How to compute "learnable" parameters?

**Sample-based
MC Denoising**

How to compute "learnable" parameters?




Image Source: Google

How to compute "learnable" parameters?




Image Source: Google

Feed-Forward Neural Network





Image Source: towards-data-science

Feed-Forward Neural Network



$$(3 \times 5) + (5 \times 2) + (5 + 2) = 17 \text{ parameters}$$

weights biases

Image Source: towards-data-science

Feed-Forward Neural Network




Image Source: towards-data-science

**Pixel-space
Kernel Predicting
Denoising**

**#Learnable
Parameters?**


**Sample-based
MC Denoising**

Sample-based Denoising Network


Michael Gharbi, Tzu-Mao Li, Miika Aittala, Jakko Lehtinen, Fredo Durand

SIGGRAPH 2019


Multimodal distribution of sample features



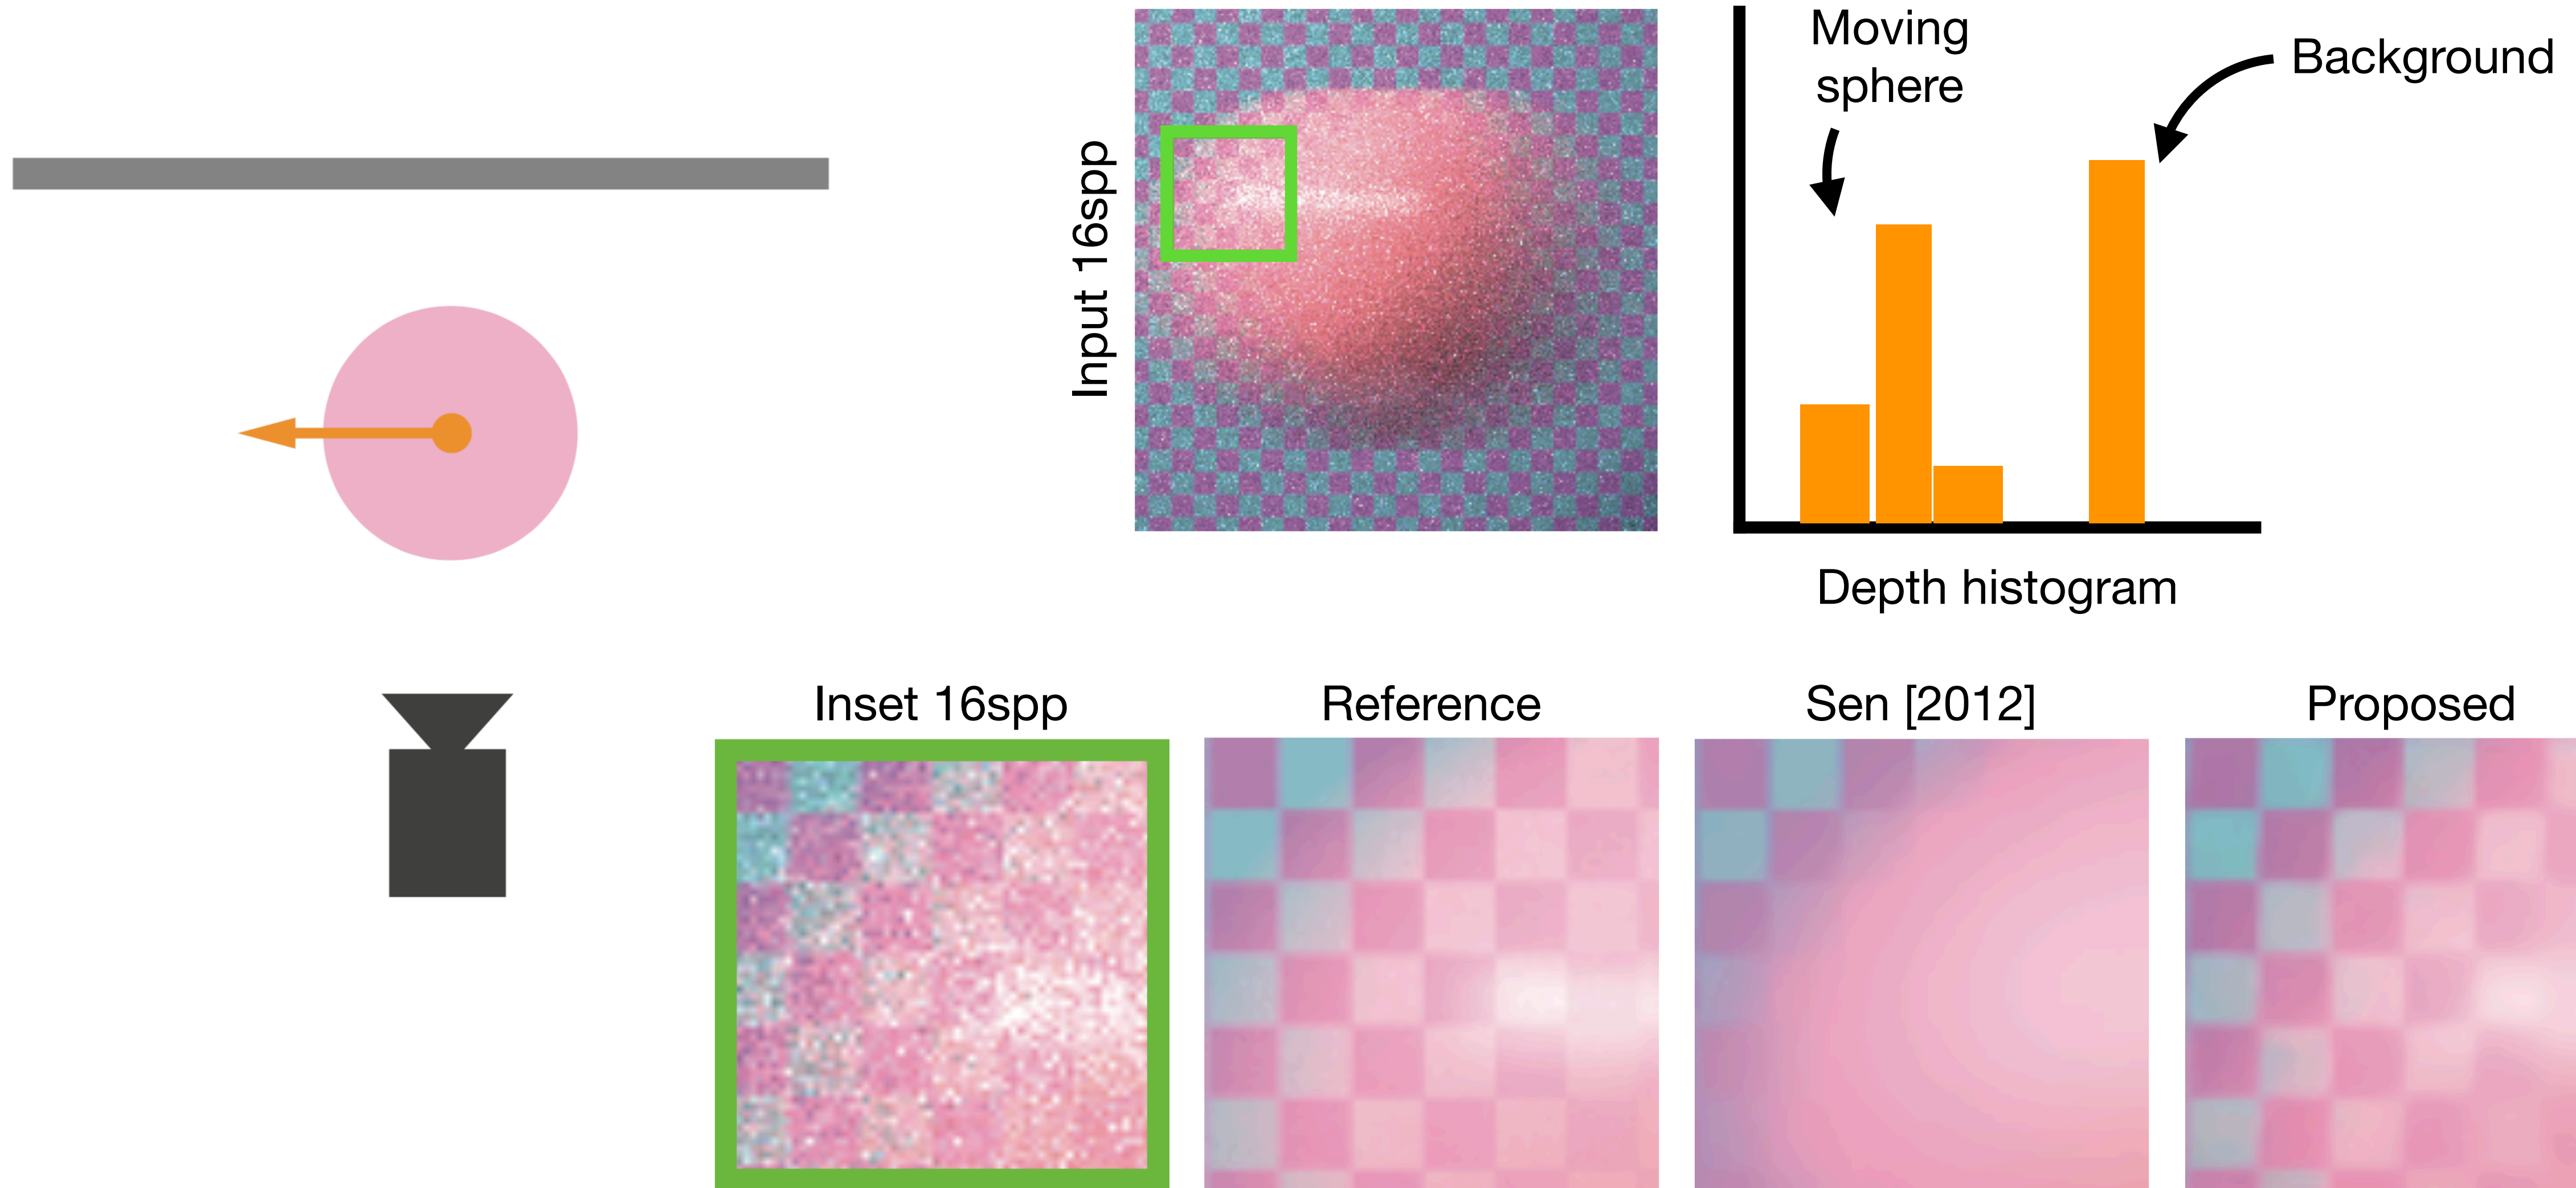
Input 16spp




Multimodal distribution of sample features




Multimodal distribution of sample features




Reconstruction: Kernel Gather



Reconstruction: Kernel Gather




Reconstruction: Kernel Gather



Reconstruction: Kernel Gather

Kernel gather




2D example

How should nearby samples influence me?


Reconstruction: Kernel Splatting

Kernel gather



2D example

Kernel Splatting




2D example

How should nearby samples influence me?


Reconstruction: Kernel Splatting

Kernel gather



2D example

Kernel Splatting




1D example

How should nearby samples influence me?


Reconstruction: Kernel Splatting

Kernel gather



2D example

Kernel Splatting




1D example

How should nearby samples influence me?


Reconstruction: Kernel Splatting

Kernel gather



2D example

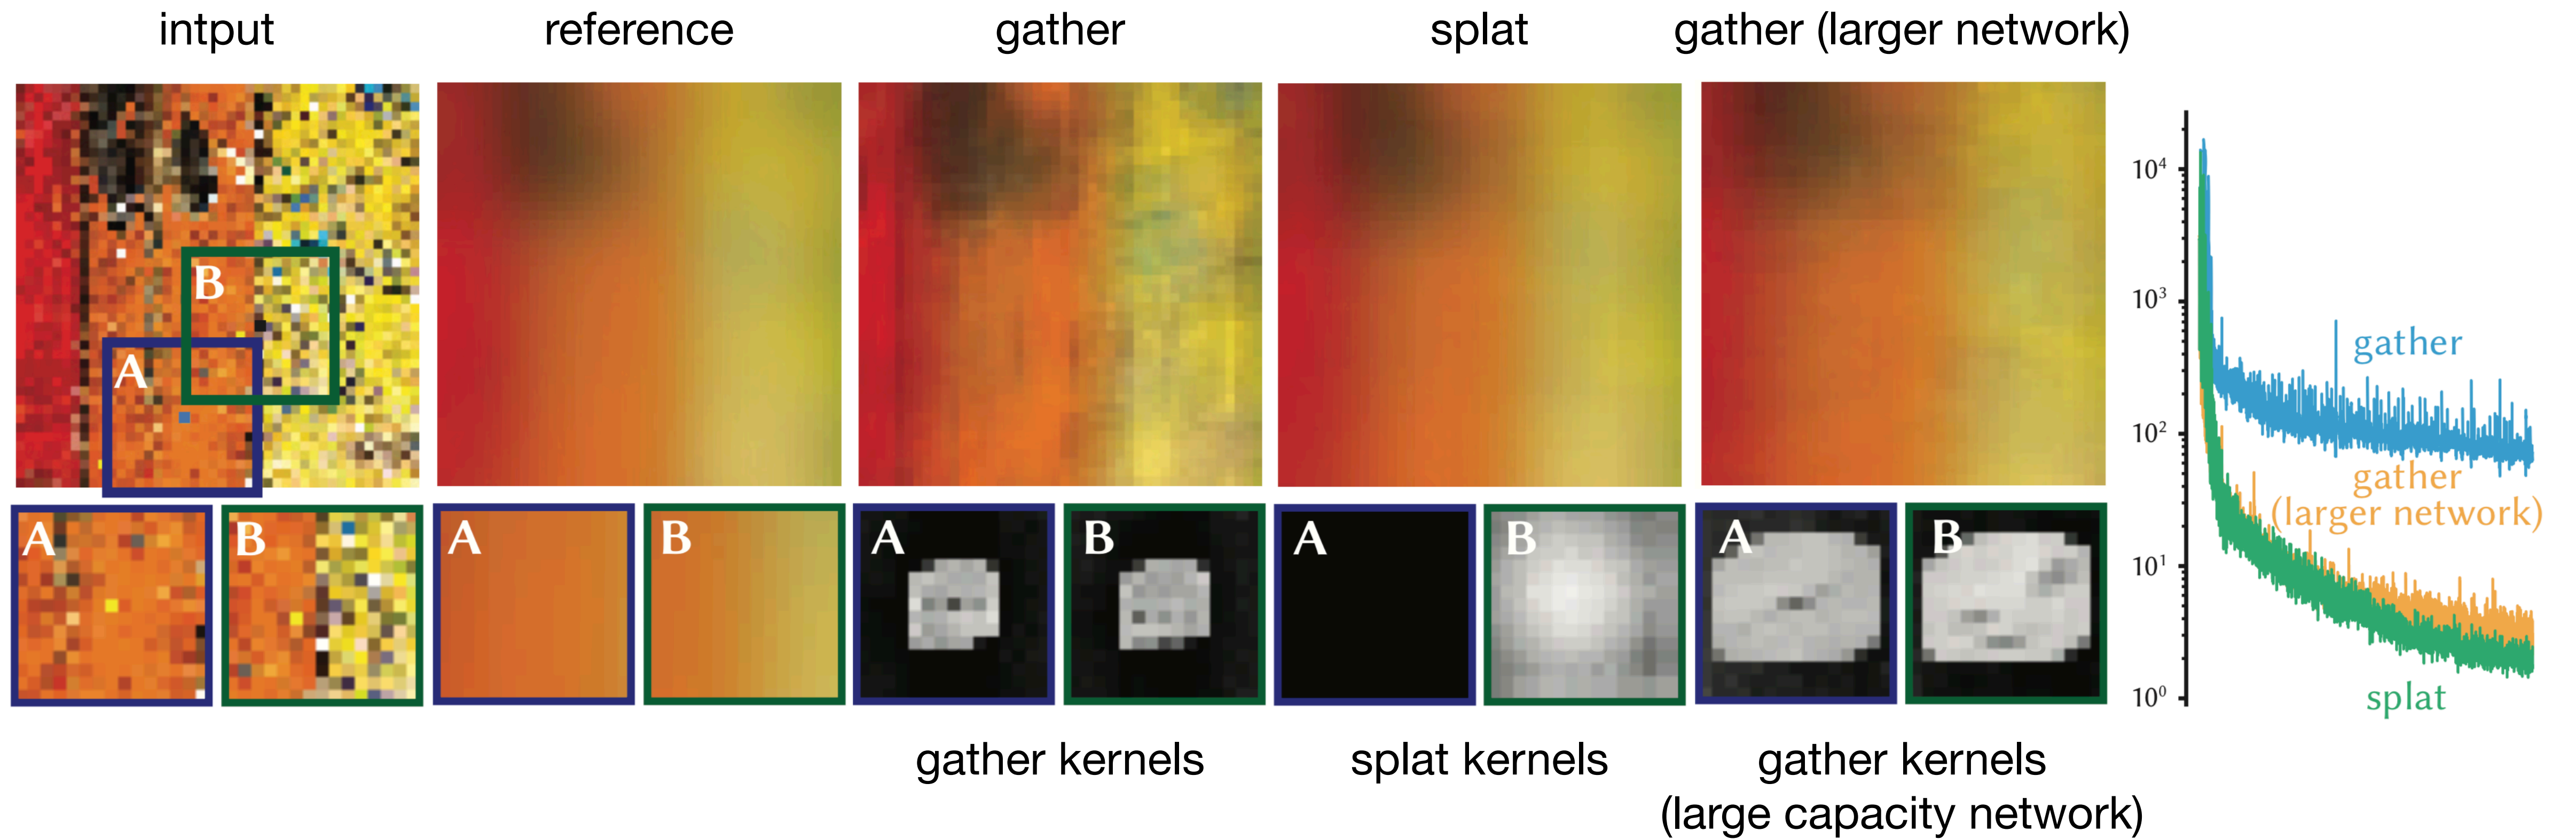
Kernel Splatting



How should nearby samples influence me?

*How do I contribute to nearby pixels,
given all the samples around me?*

Network: Kernel Gather vs Splatting



Permutation Invariance

Permutation Invariance

A model that produces the same output regardless of the order of elements in the input vector

Permutation Invariance: Example



*



=

Permutation Invariance: Example



Permutation Invariance: Example

Not Permutation Invariance



*



=



*



=

Permutation Invariance: Example

Not Permutation Invariance



*



=



*




=



Permutation Invariance: Architectures

- A standard feedforward neural net such as multilayer perceptron (MLP) is insensitive to order of elements in input vector - so it is inherently permutation insensitive
- However, both a Convnet and RNNs for instance make full use of input ordering - they are permutation sensitive.

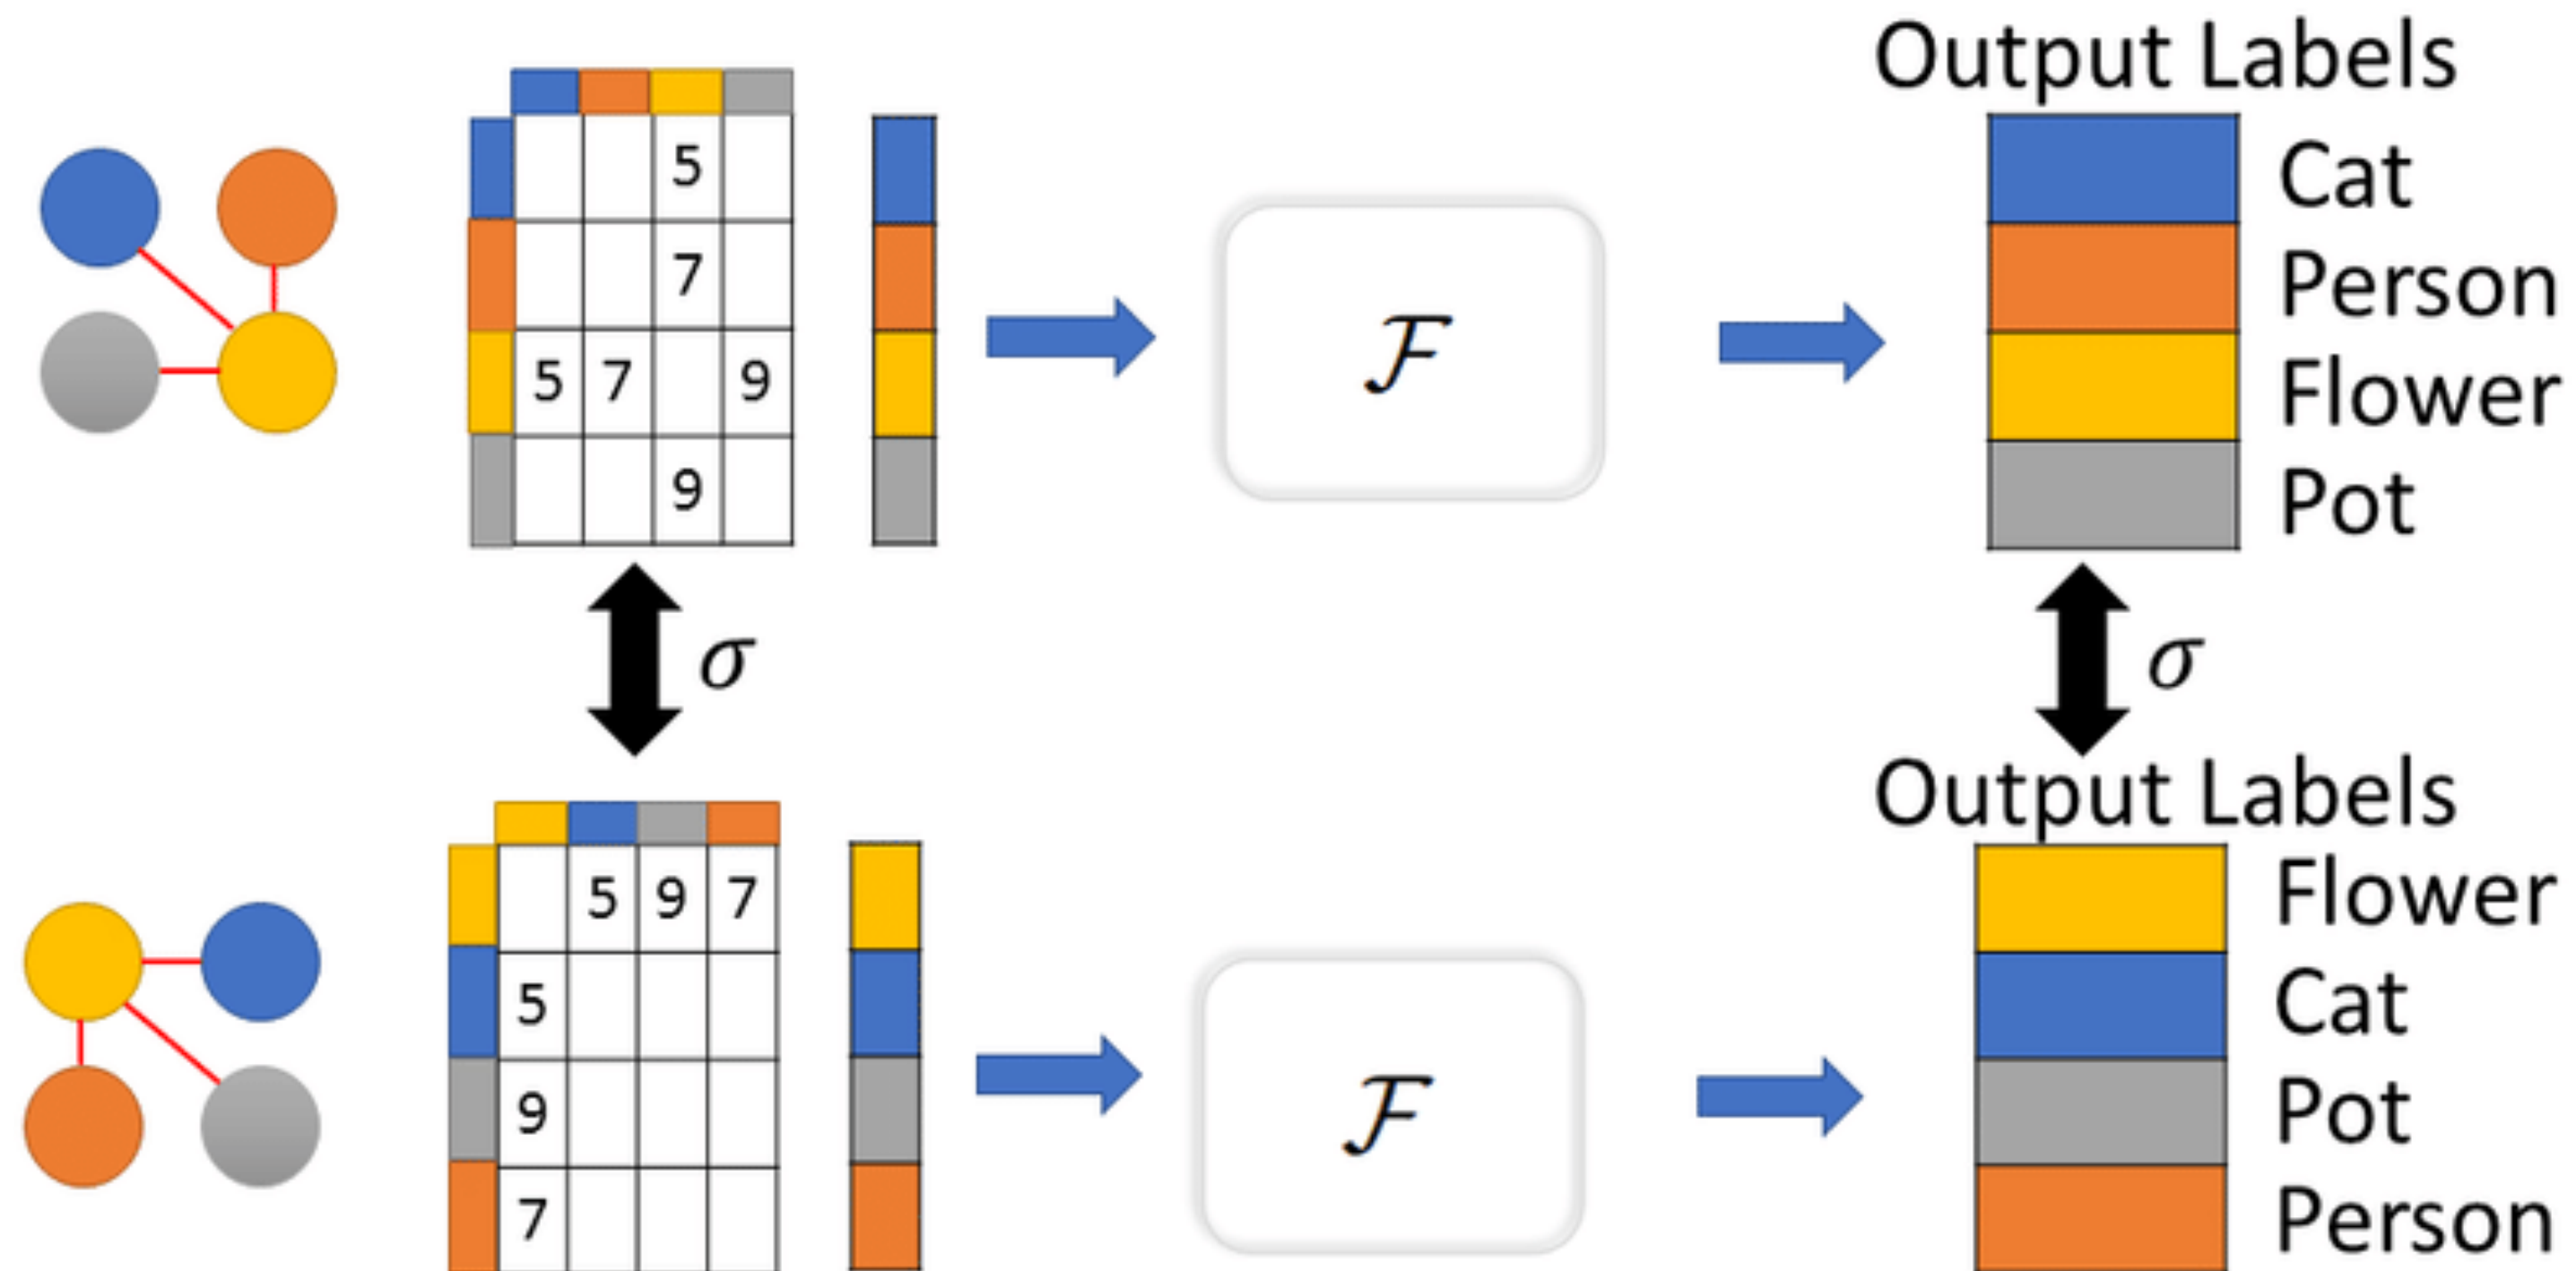
Permutation Invariance: Example



Permuting pixels makes it difficult for humans to understand the images.

However, permutation invariant networks like MLP can detect digits irrespective of the order of pixels

Permutation Invariance: Example

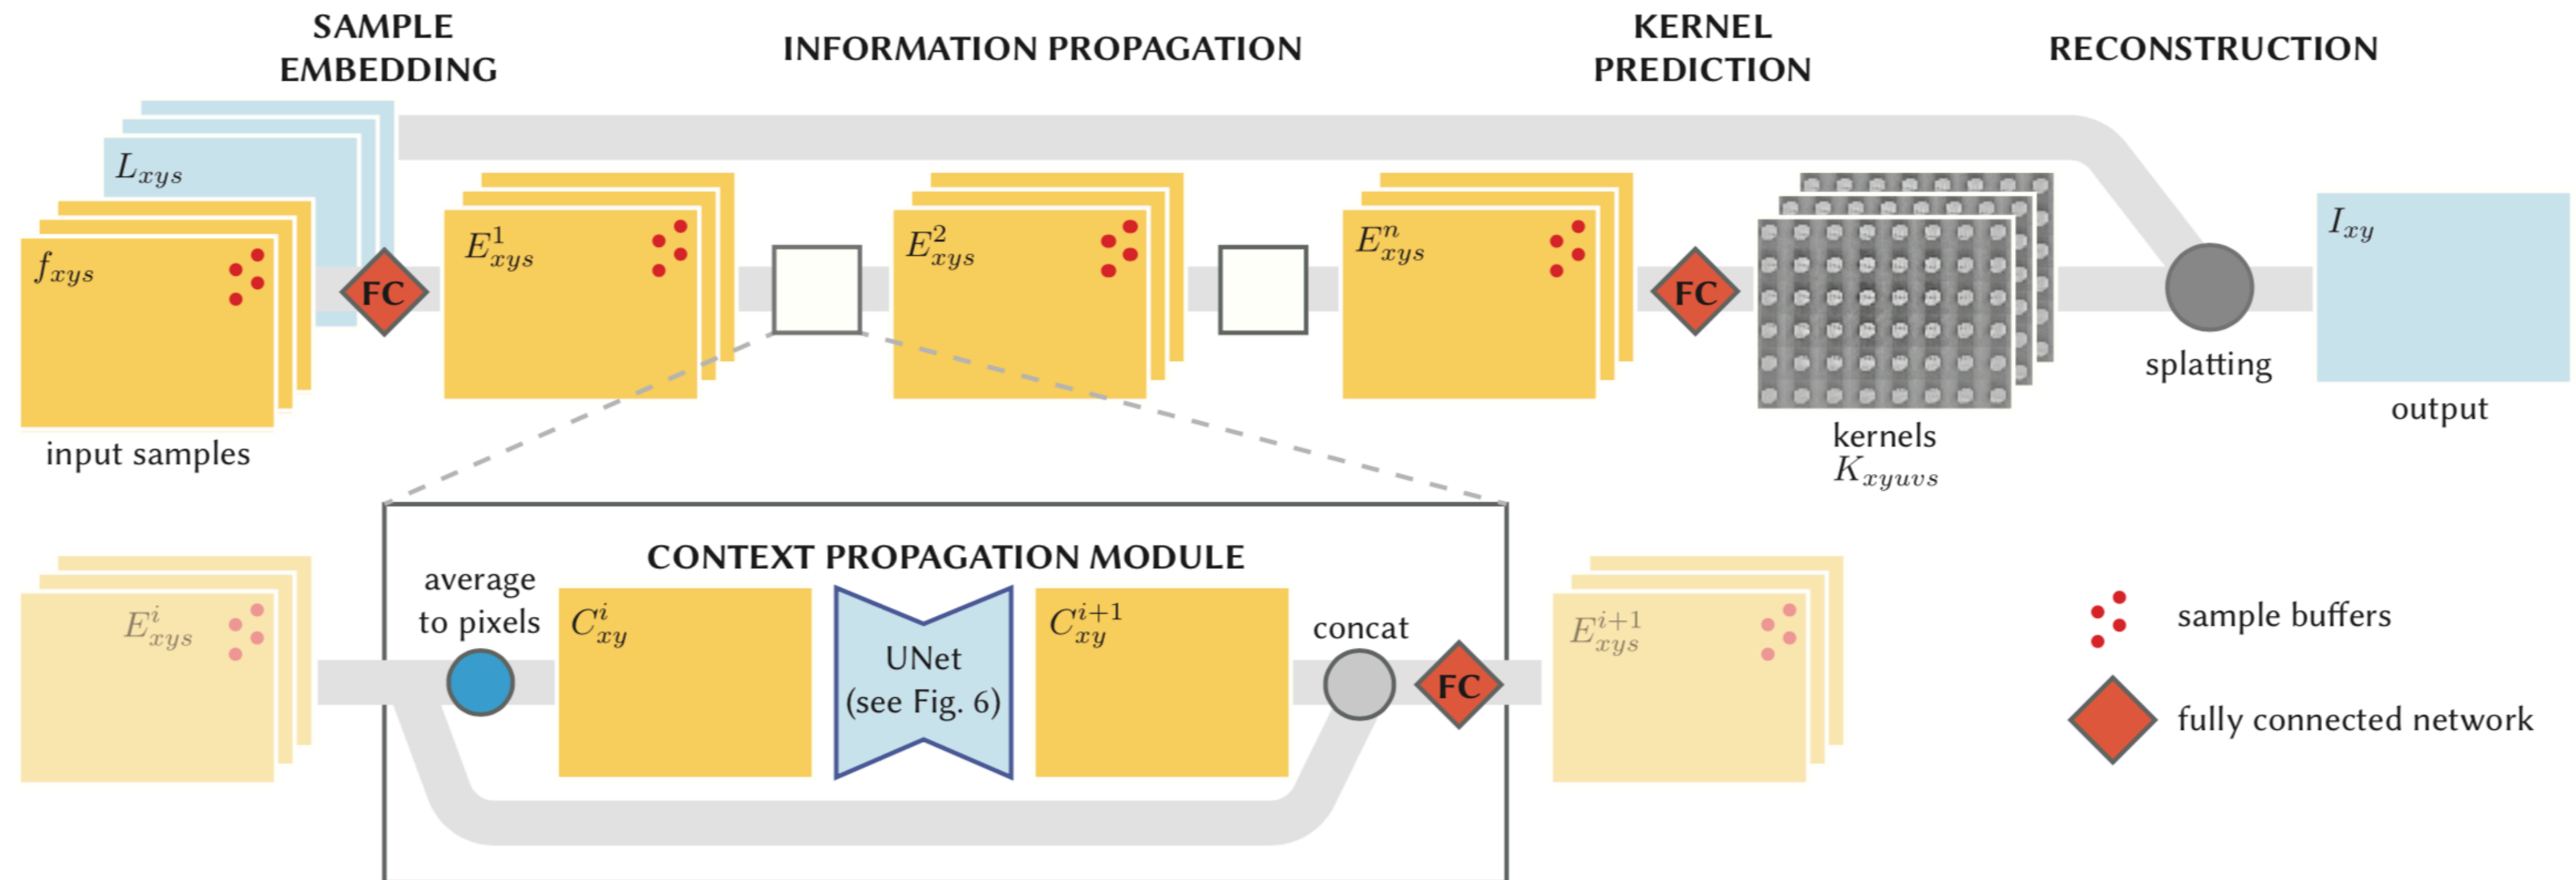


A graph labeling function F is graph permutation invariant (GPI) if permuting the names of nodes maintains the output. Herzig et al.[2018]

Permutation Invariance

- In MLPs, since each component is connected to each other, the order does not matter
- In structured convolutions, the order matters and therefore, it is not permutation invariant.

Proposed Network Architecture




Dataset and Training Procedure

Procedurally generated dataset: 300,000 renderings with 128x128 resolution

Also generated input buffer (4, 32 spp), but this time also maintained auxiliary features


Reference was generated for 4096 samples

Splat vs Gather





Input


per sample gather




per pixel gather



per sample splat




per pixel splat




Reference

Results


input 4spp







[Bako 2017]



ours



ref. 8192spp



Network Architecture Comparisons


reference 8192spp

input

finetuned [Bako2017]

ours


reference 8192spp



(Deep) Convolutional Neural Networks

Based on Convolutional Neural Networks


Unstructured data




N number of point samples

Based on Convolutional Neural Networks

Unstructured data




Convolution \otimes






N number of point samples

Based on Convolutional Neural Networks

Unstructured data




Convolution \otimes



N number of point samples

Based on Convolutional Neural Networks

Unstructured data






Convolution \otimes




\otimes

\otimes



Loss function




Back-propagate


N number of point samples

Based on Convolutional Neural Networks

Unstructured data




Convolution \otimes



N number of point samples





Based on Convolutional Neural Networks

Unstructured data



■■■ Keep training the network!


Convolution \otimes



N number of point samples


Based on Convolutional Neural Networks

Unstructured data




■■■ Keep training the network!

Convolution \otimes



Which Loss function can we use?





N number of point samples

Spectral Loss Function

Spectral Loss at i -th training iteration


$$L_{\text{spectral}} = ||\underline{\langle \mathcal{P}_i(\nu) \rangle} - \langle \mathcal{P}(\nu) \rangle||^2$$




Radially averaged power

Training Process


Points




Power Spectrum



56x Slowdown

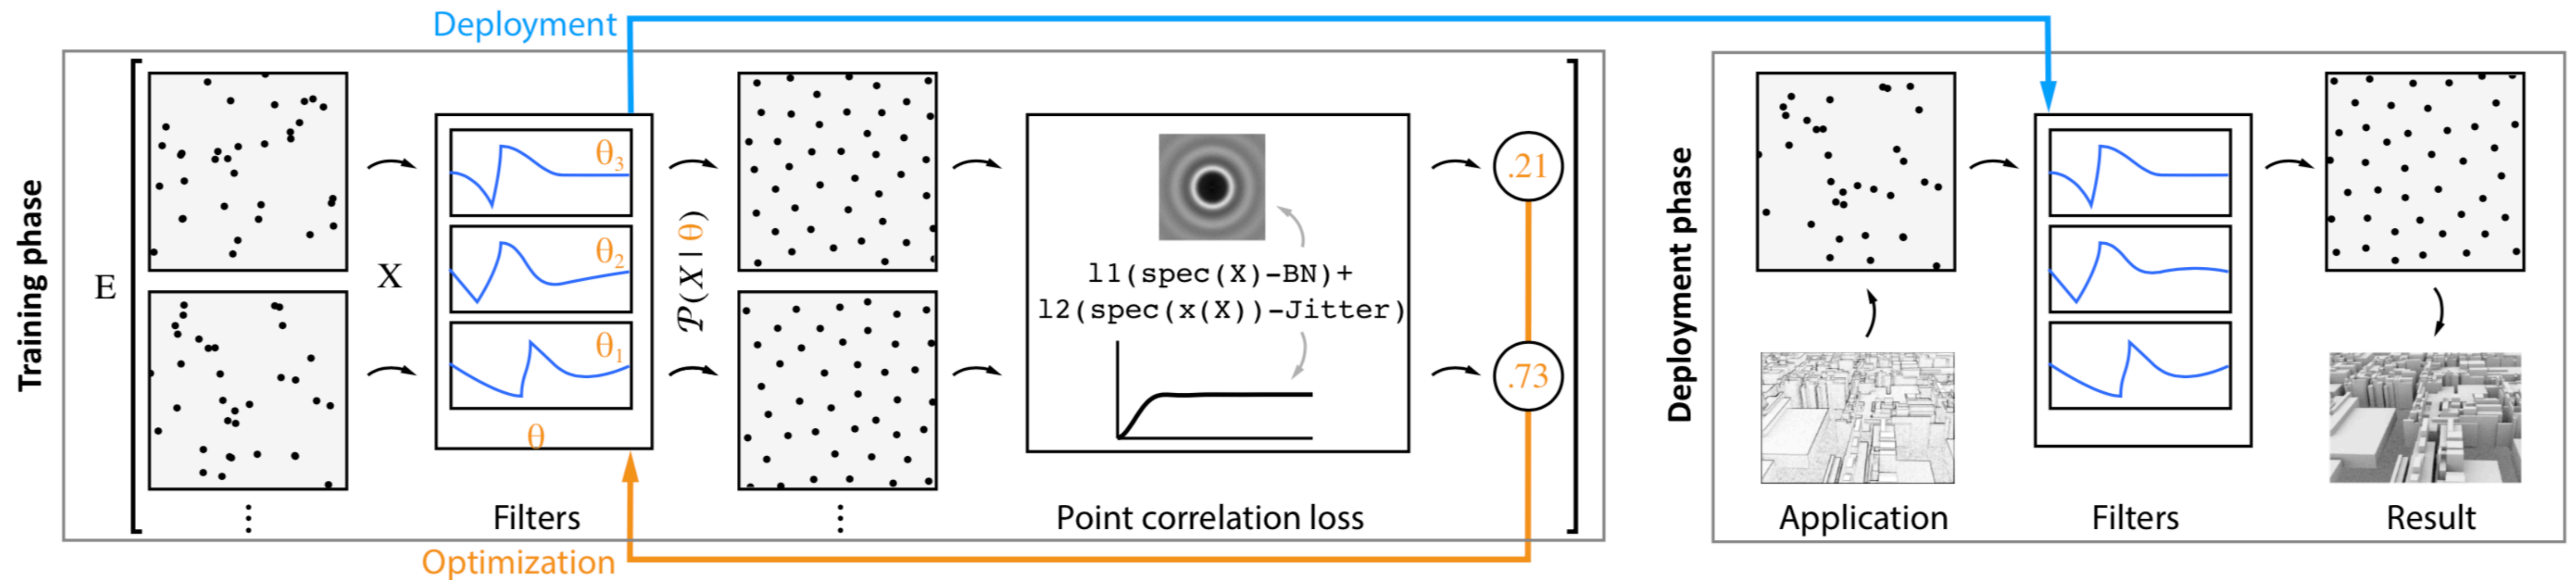


Kernels for BNOT
(de Goes et al.[2012])

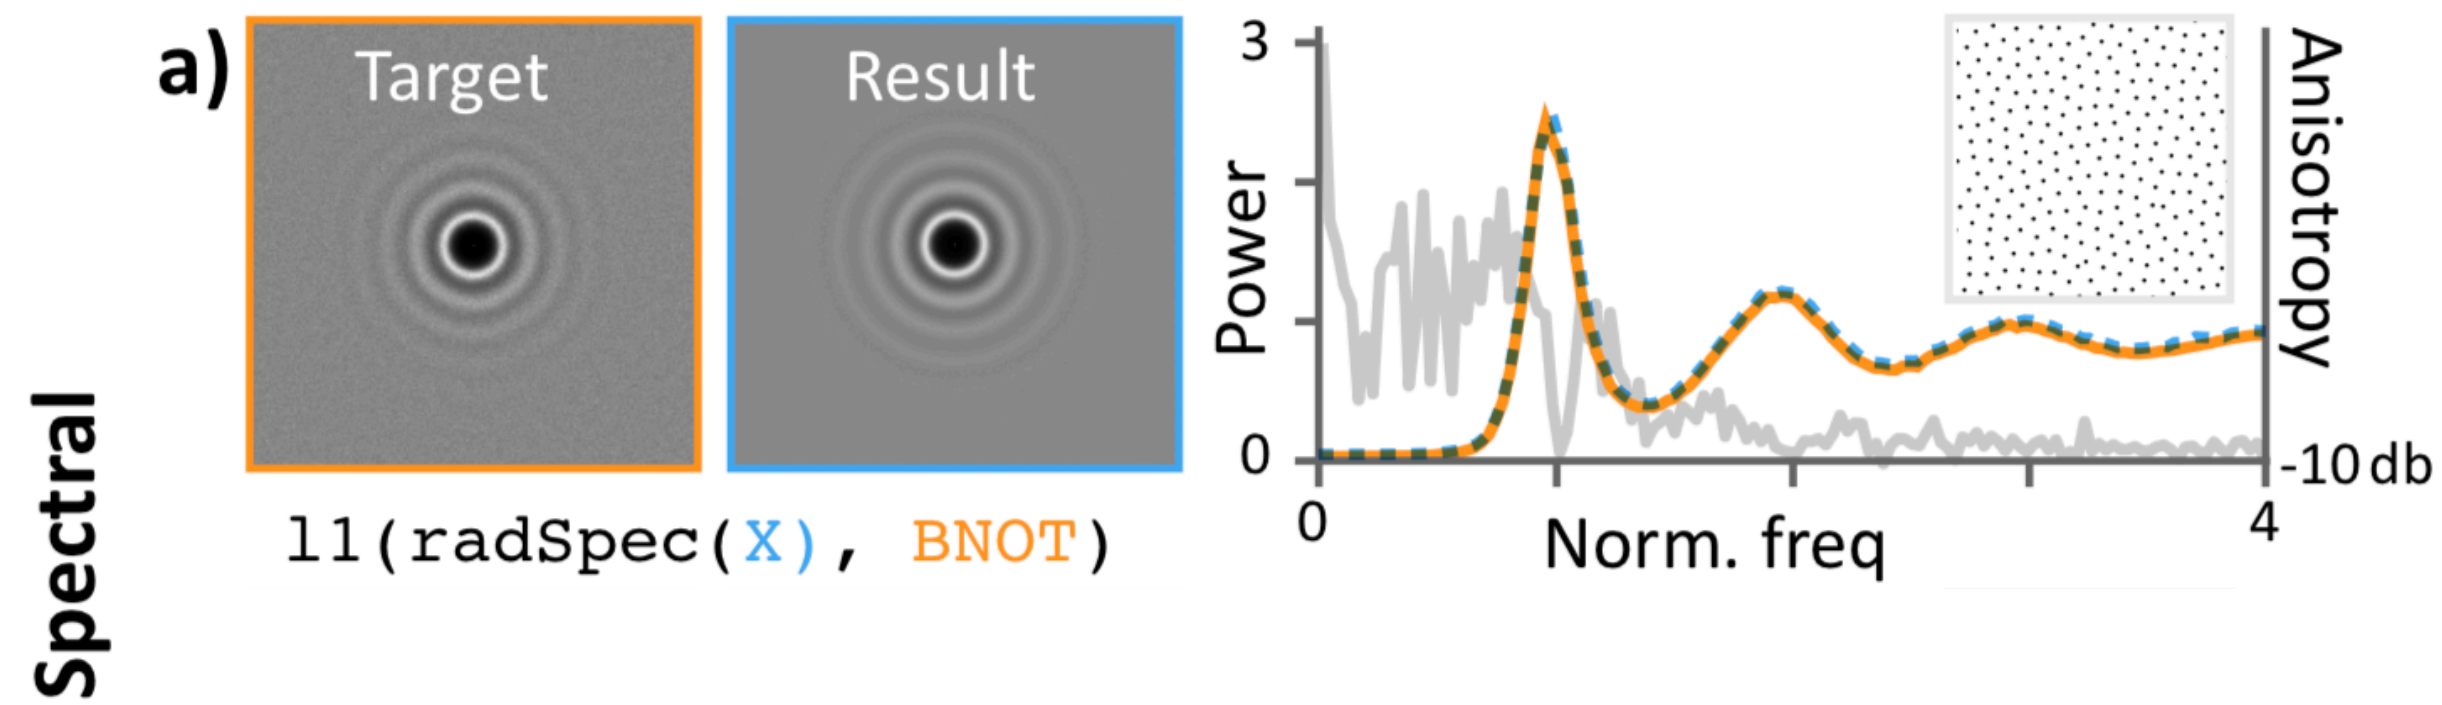


Kernels for Step
(de Heck et al.[2013])

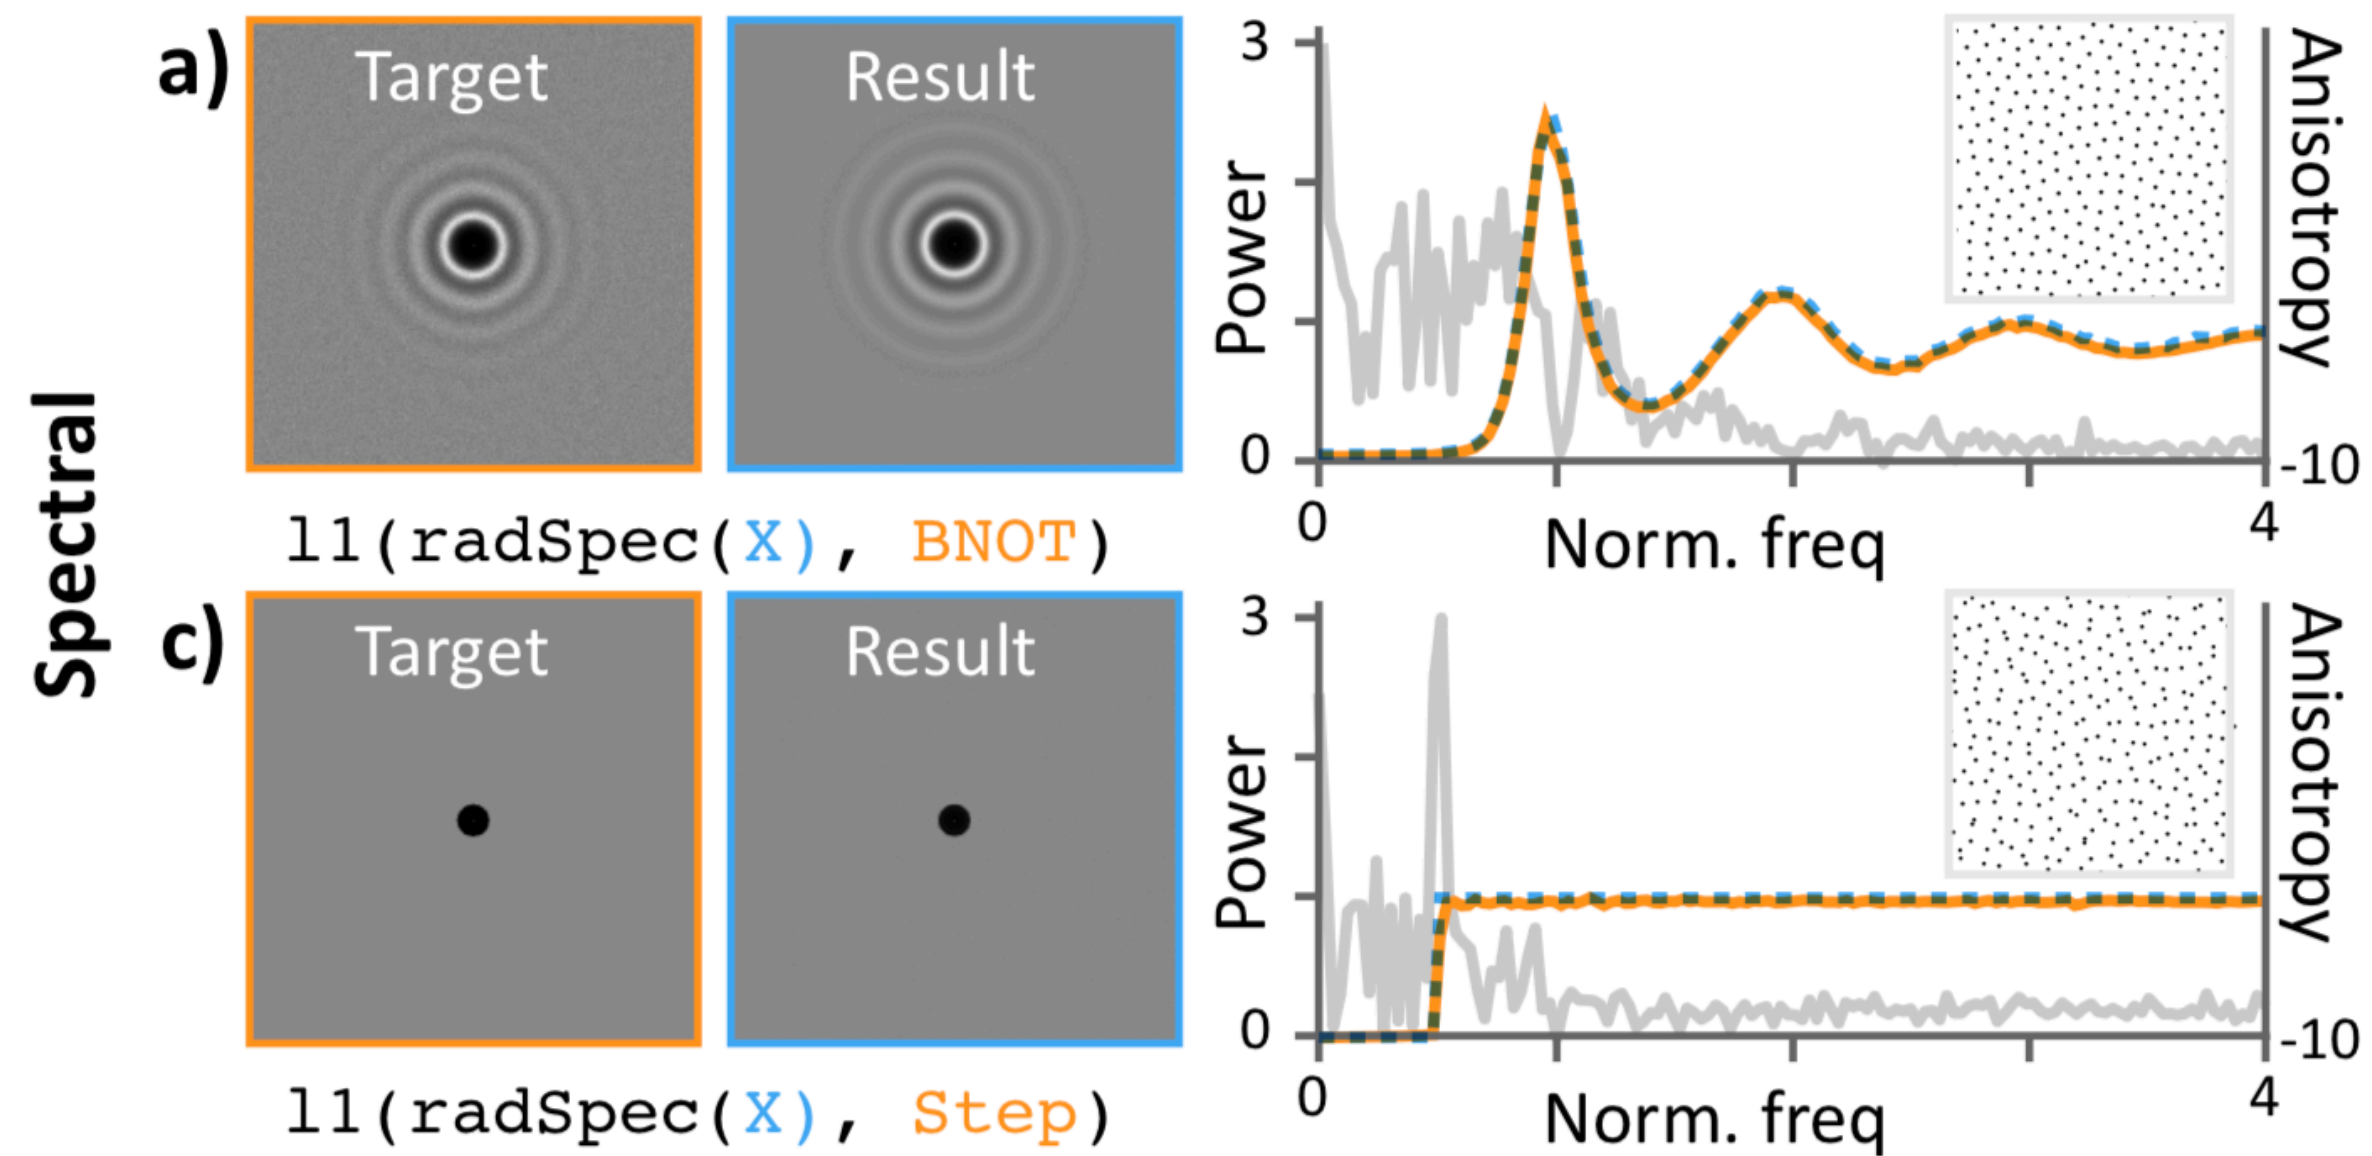
Architecture: Full pipeline



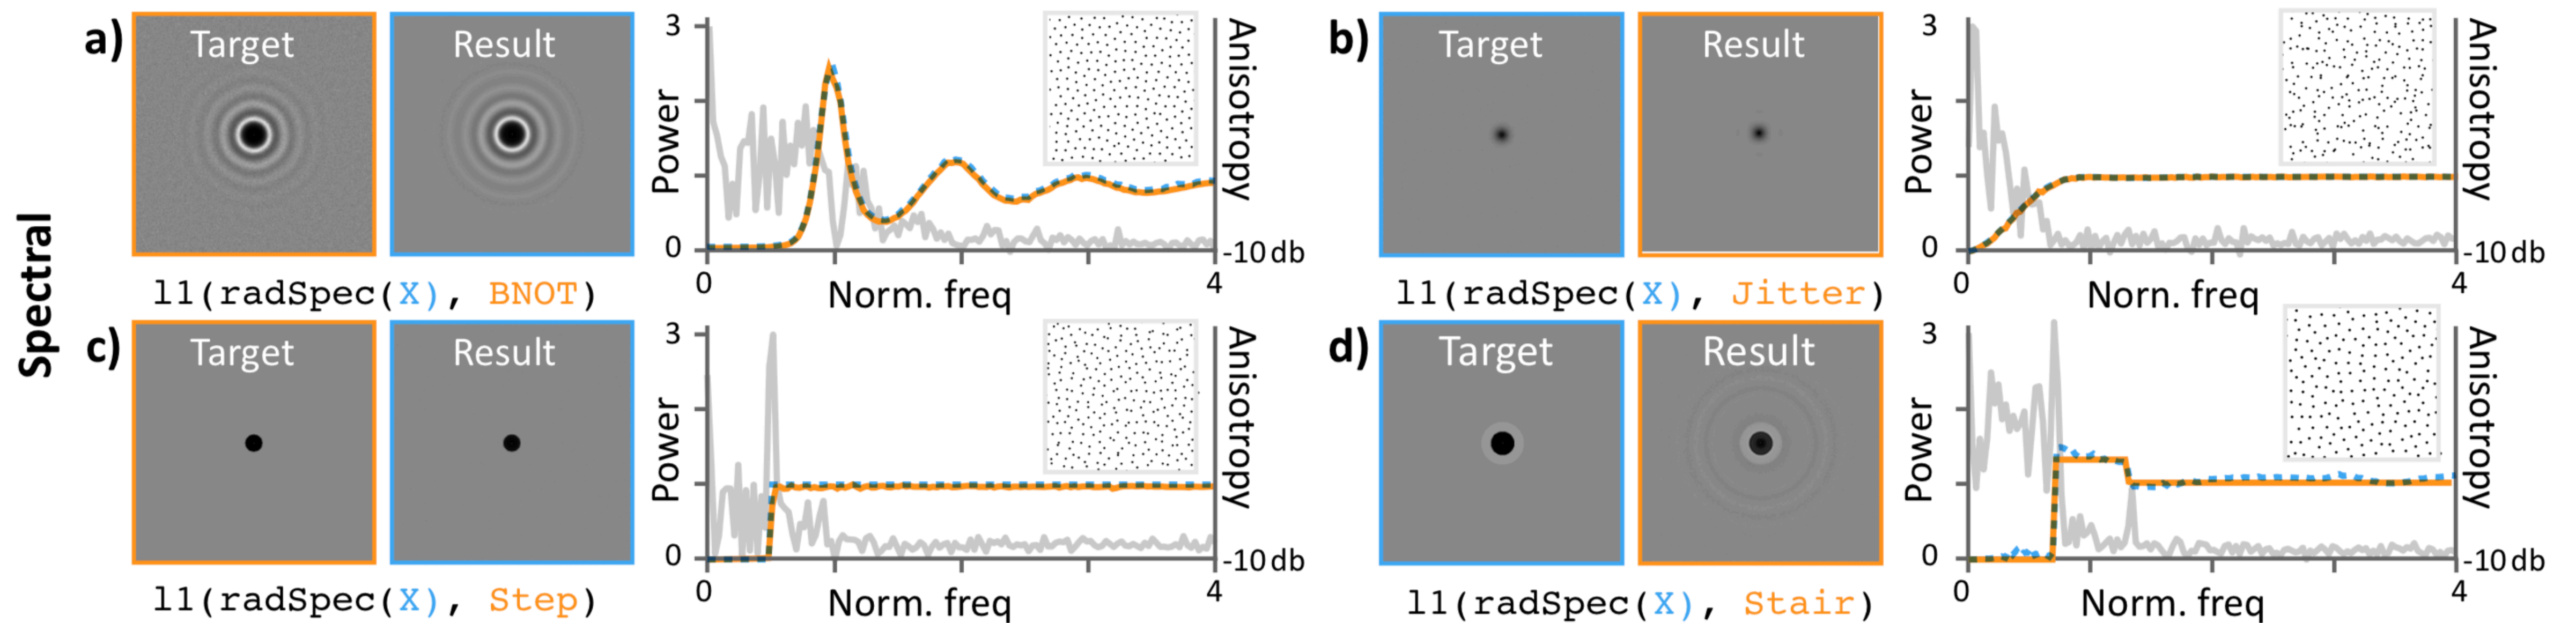
Results: Spectral Target Spectra




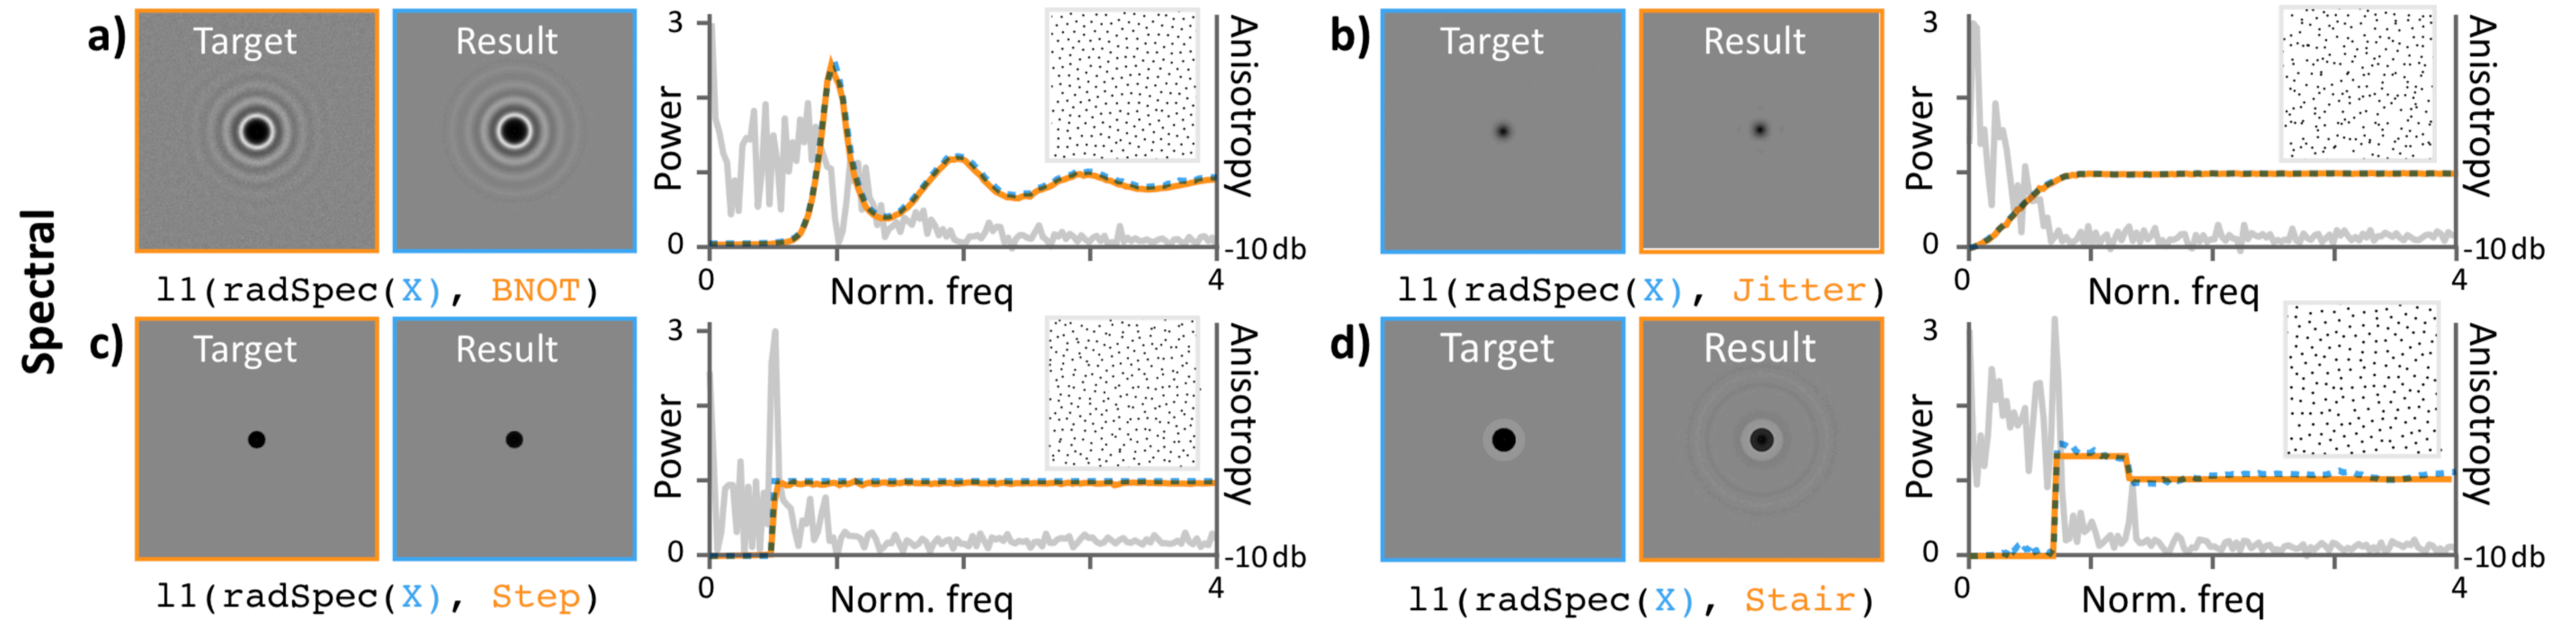
Results: Spectral Target Spectra



Results: Spectral Target Spectra



Results: Spectral Target Spectra




Spatial Loss Function



PCF Loss at i -th training iteration

Spatial Domain

Blue Noise




Samples



Loss Functions

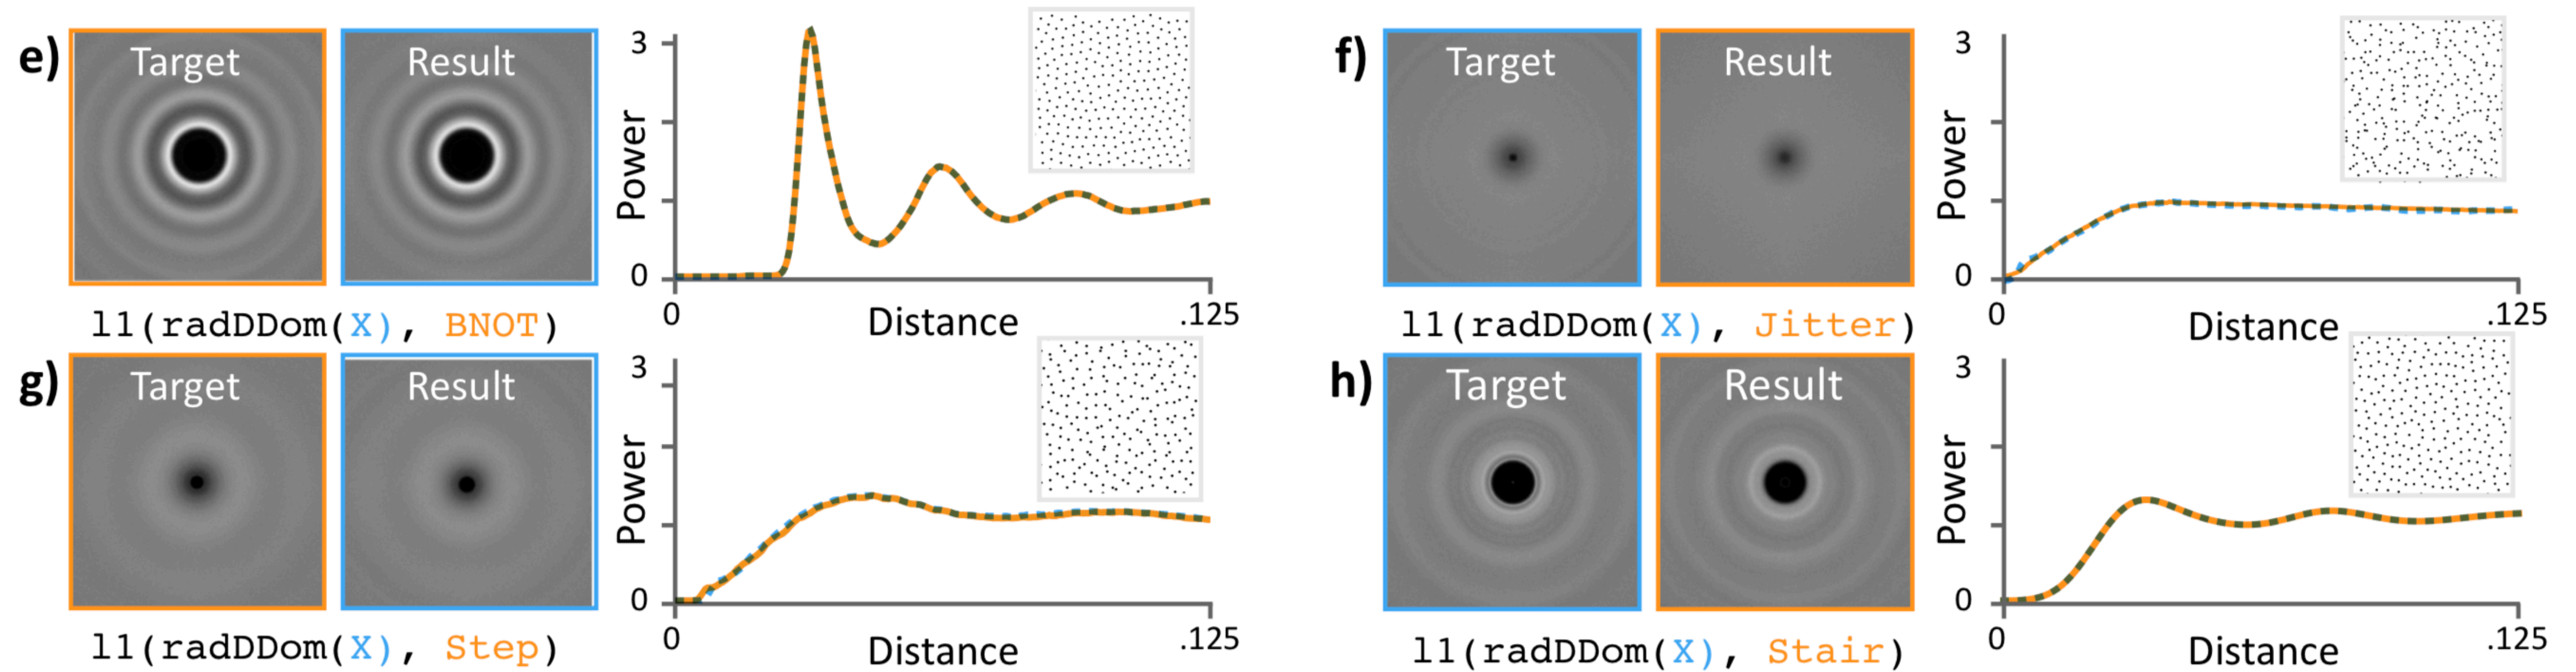
PCF Loss at i -th training iteration

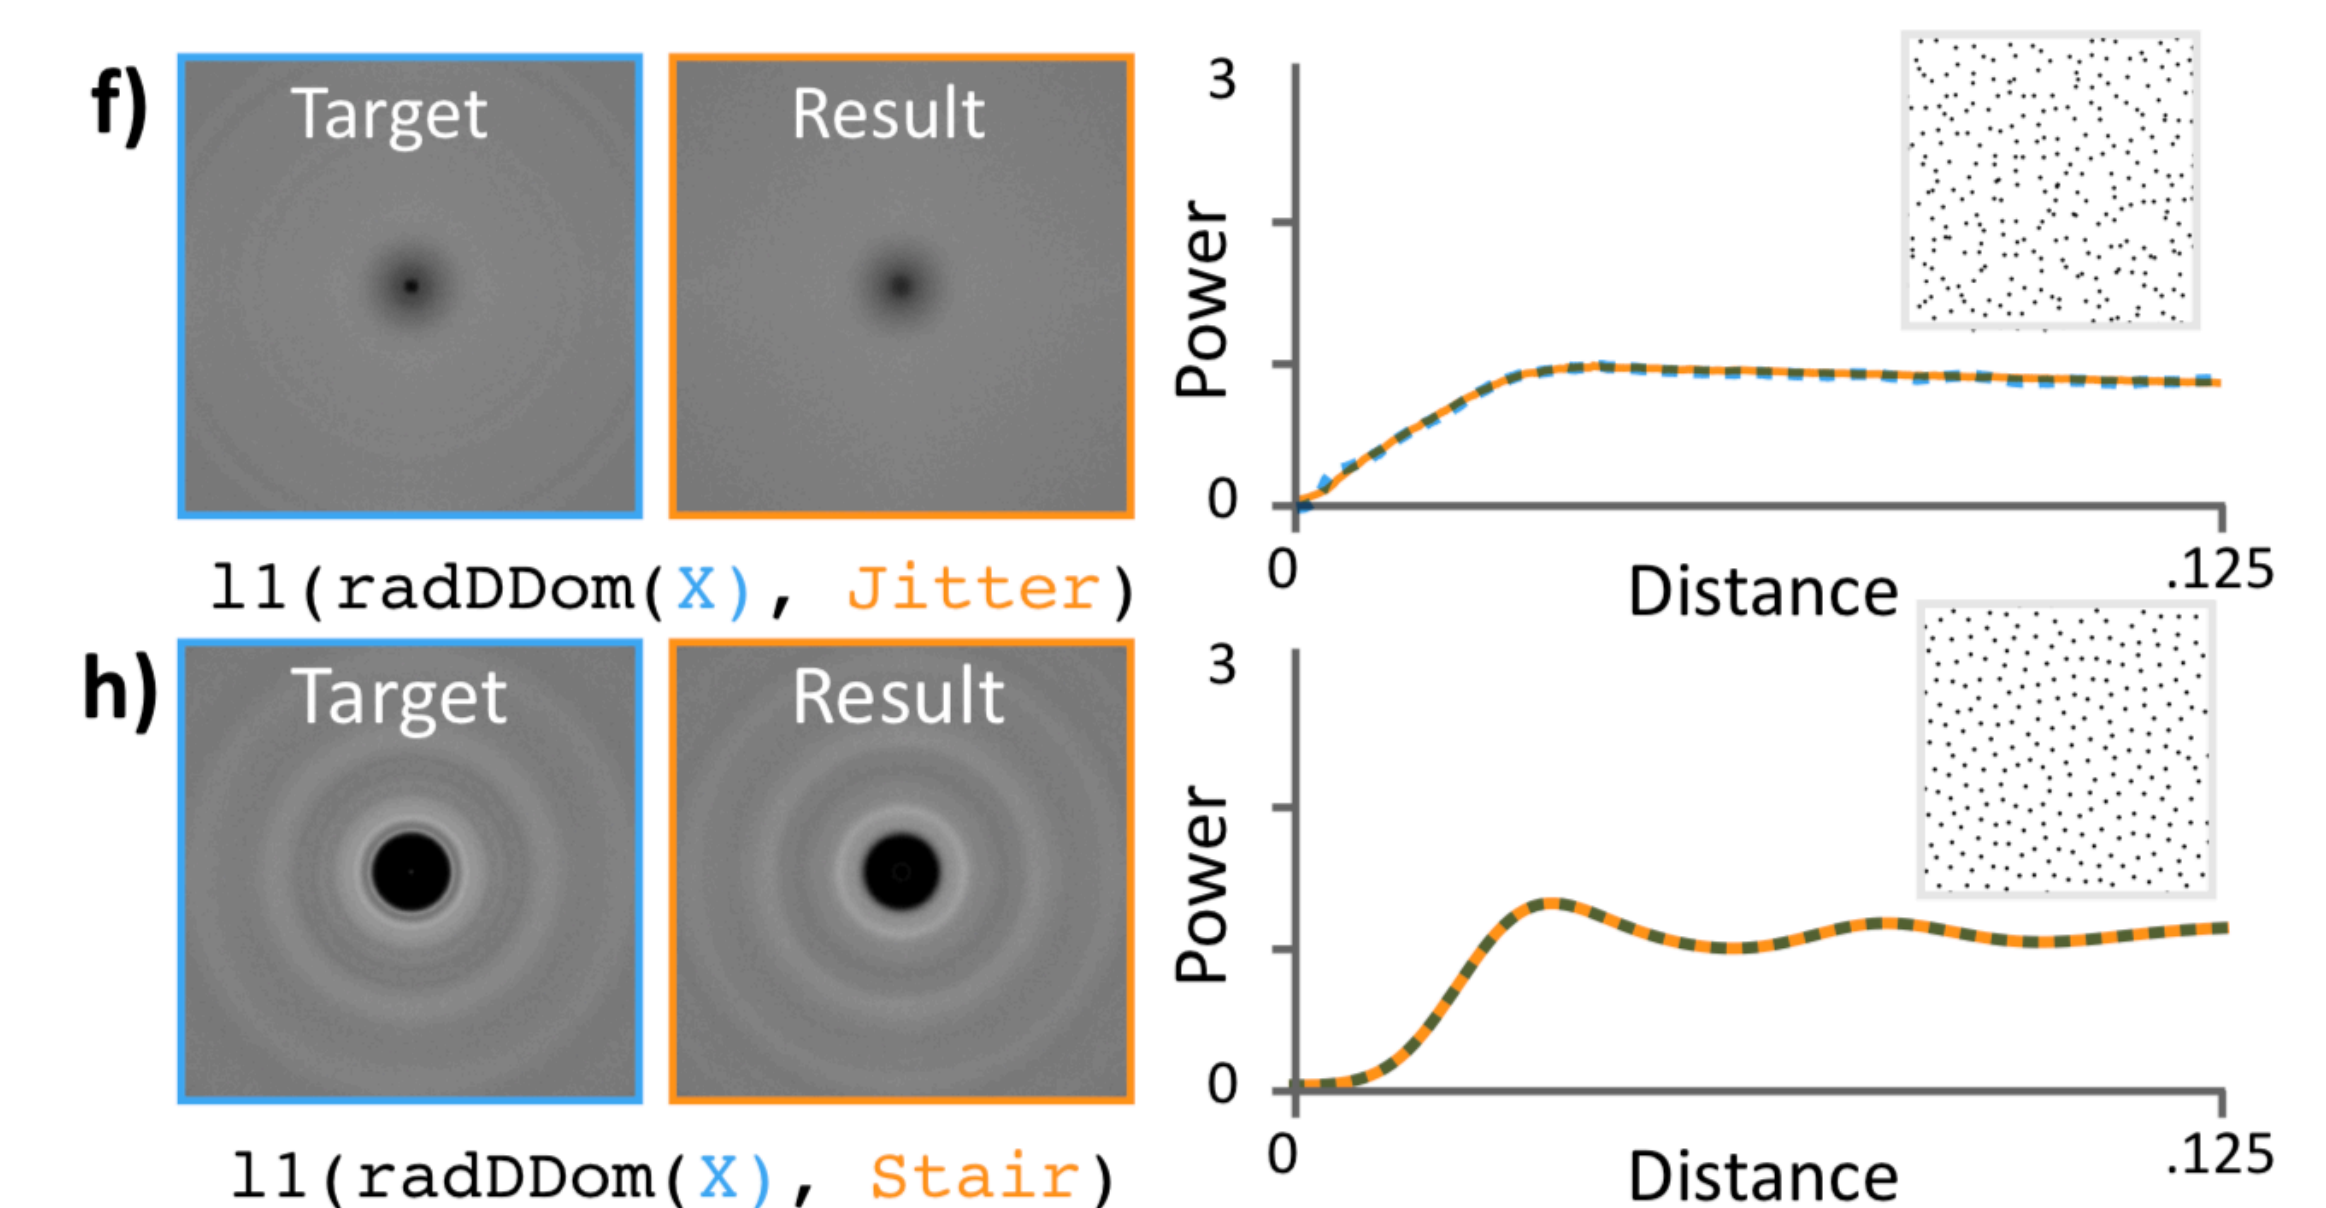
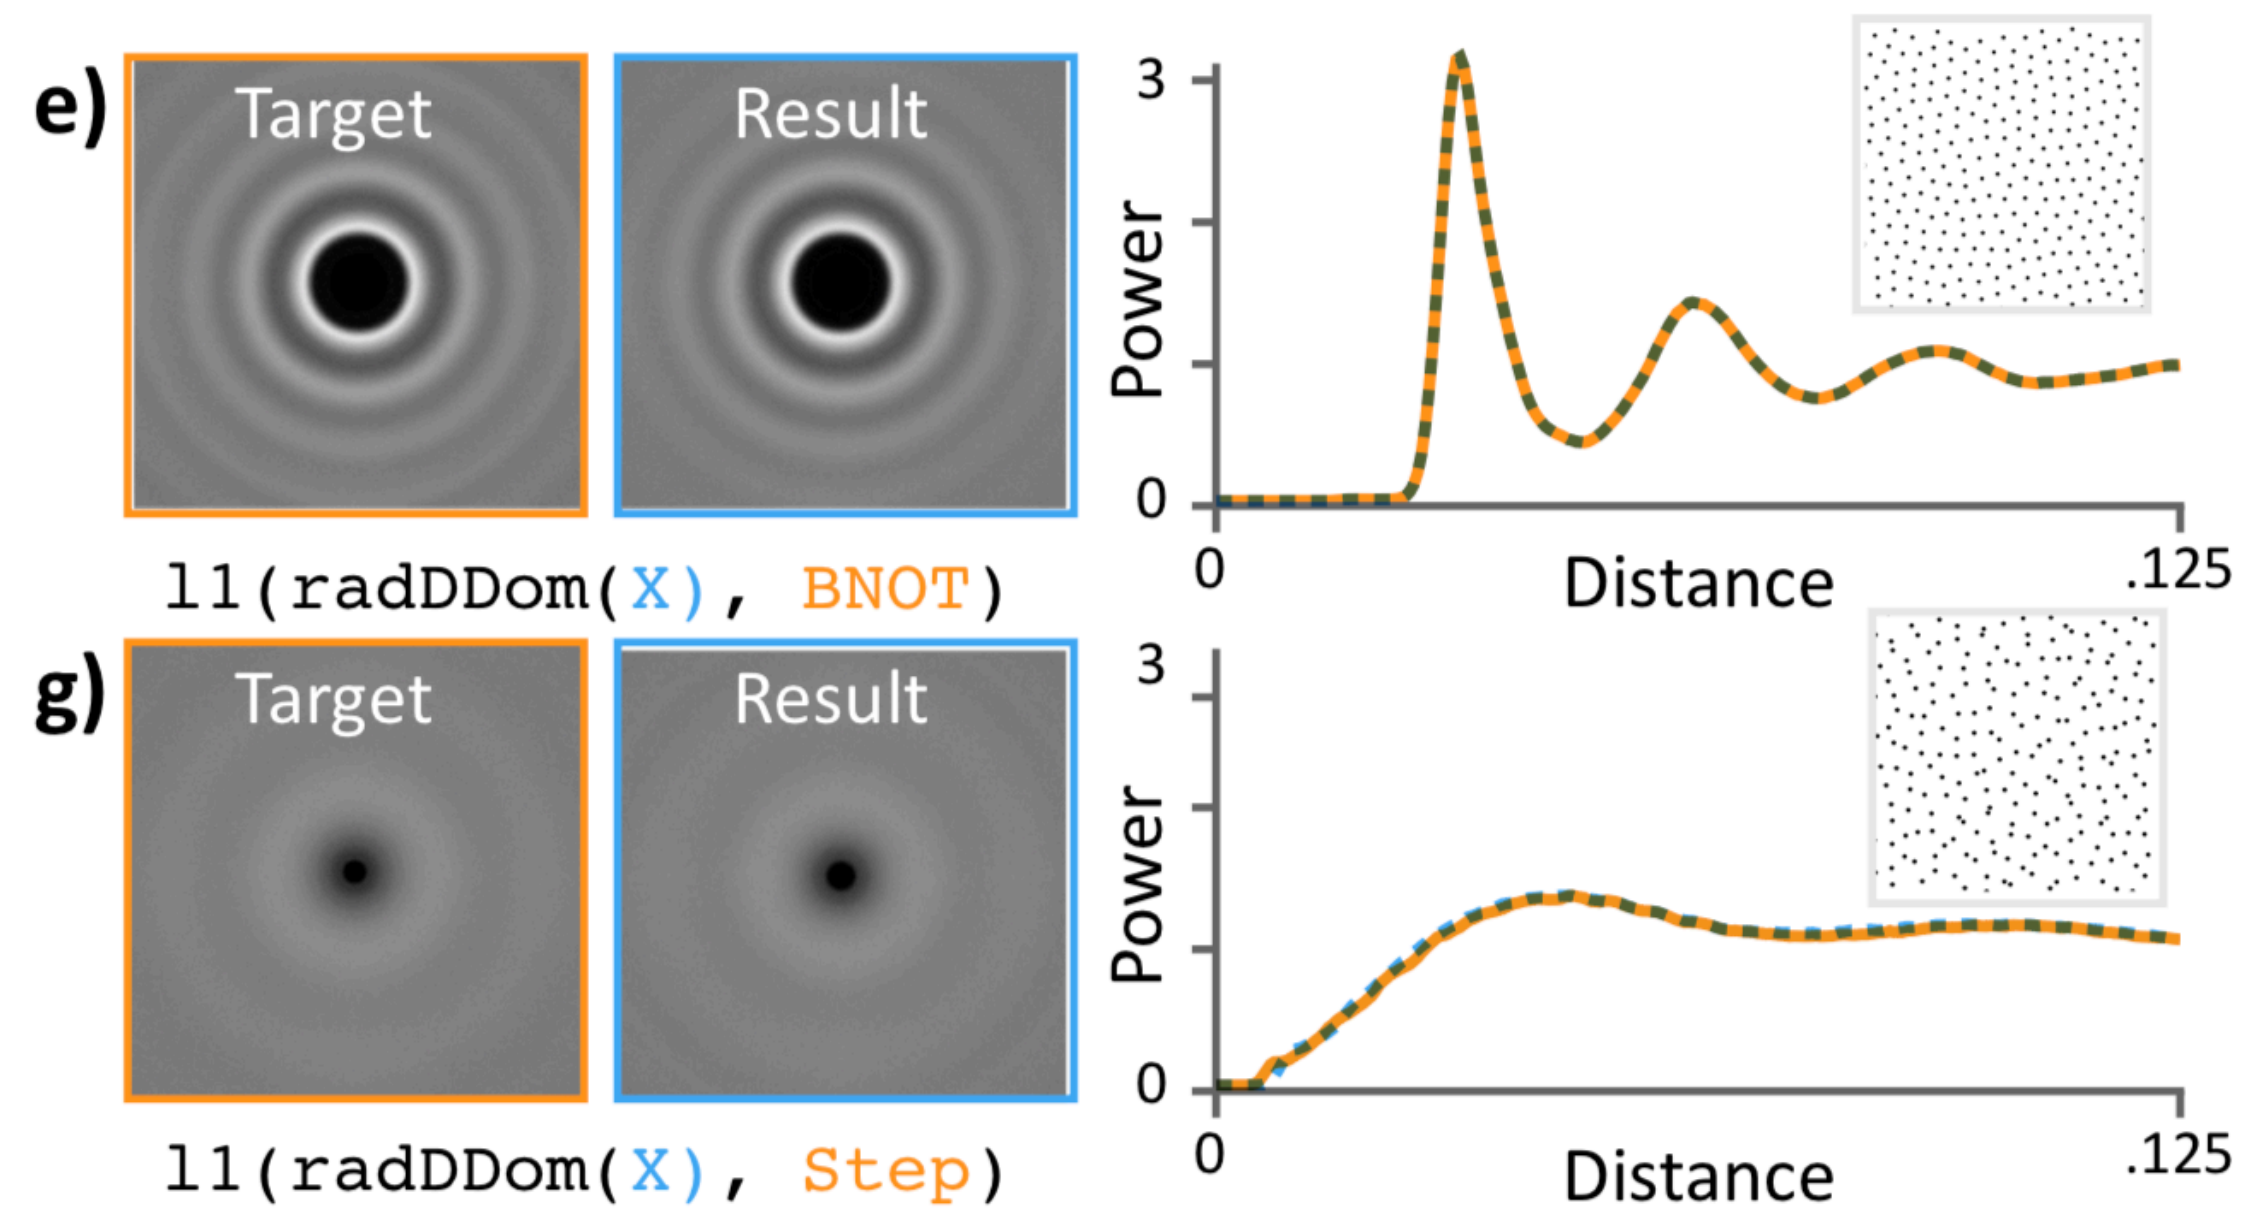
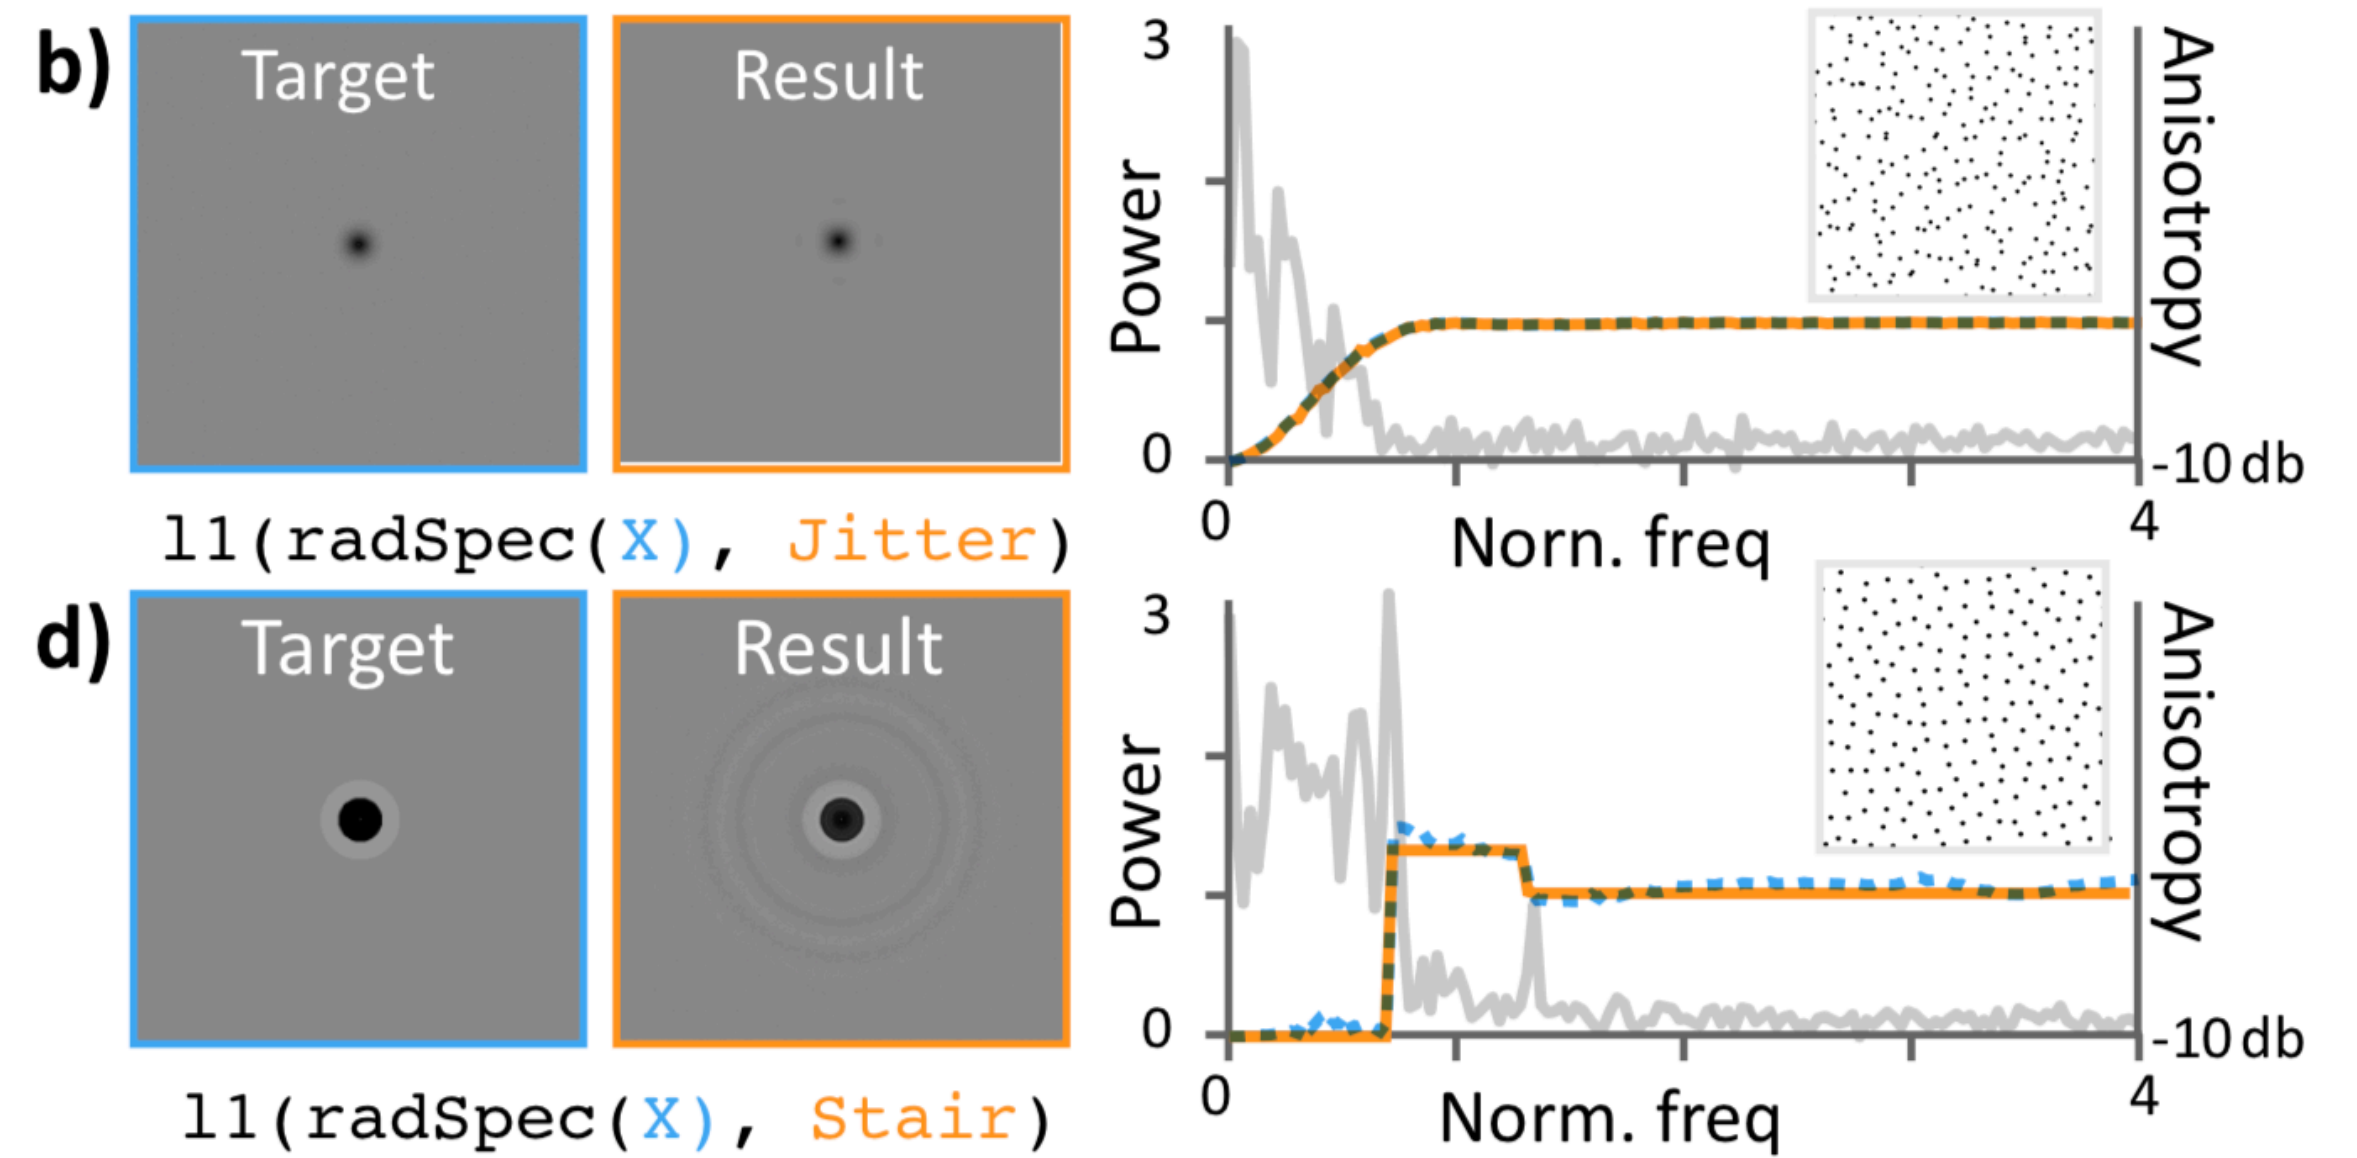
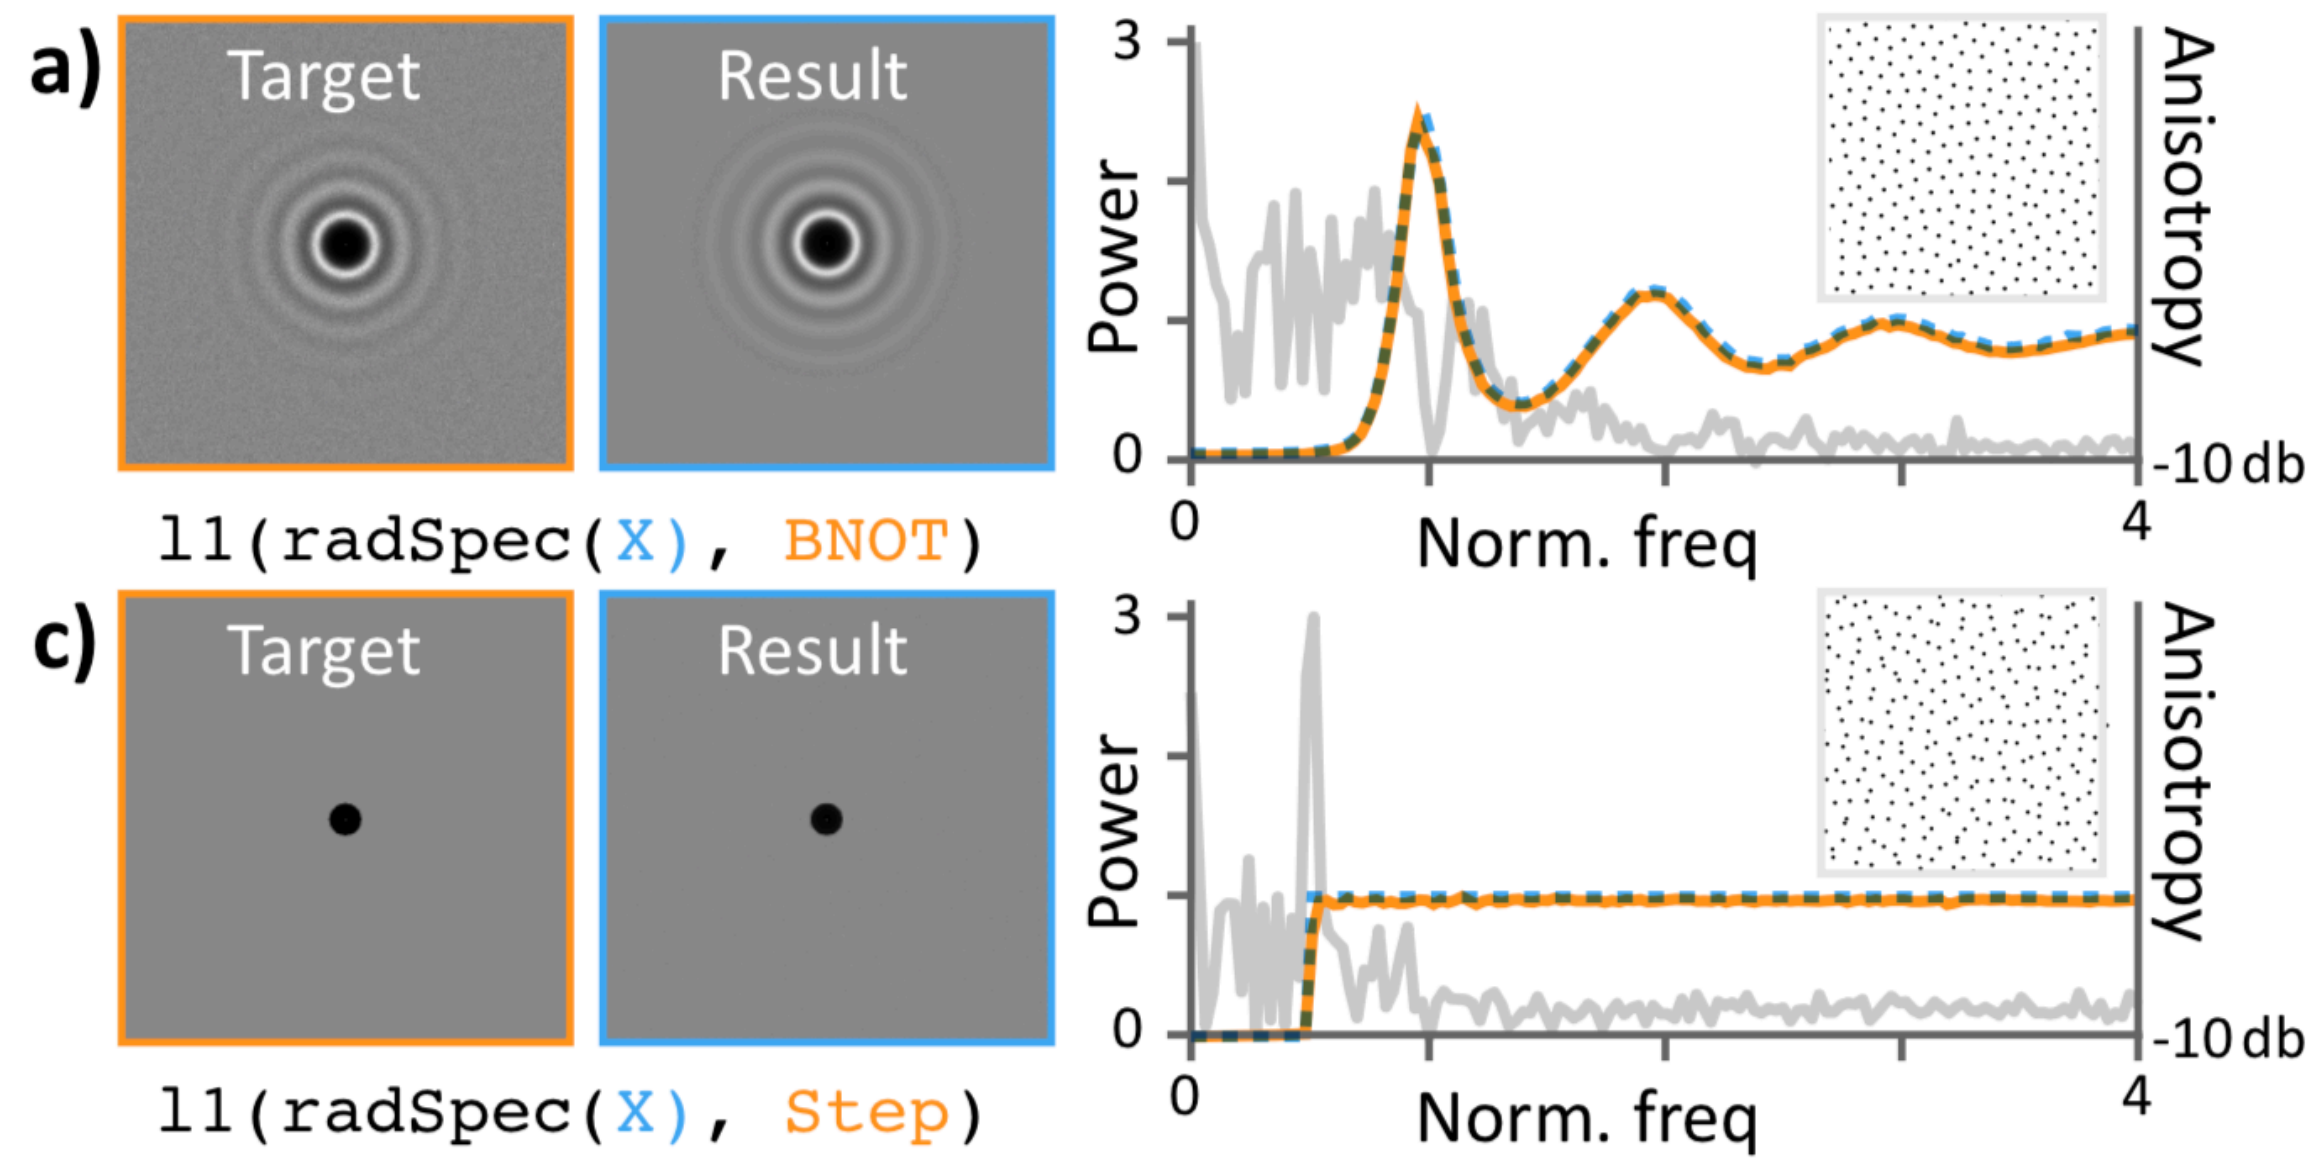
$$L_{\text{PCF}} = ||\underbrace{\langle r_i(\text{dist}) \rangle}_{\text{green bar}} - \langle r(\text{dist}) \rangle||^2$$




Spatial Target PCFs

Differential






3D Point Samples (Different Projection Targets)




3D Point Samples (Different Projection Targets)



XY
(BNOT)


Target Spectra





YZ
(Jitter)





XZ
(Step)




3D Point Samples (Different Projection Targets)



Point set: Projections are Preserved




Novel Sampling Patterns



Leimkuhler et al. [SIGGRAPH Asia 2019]


Novel Sampling Patterns



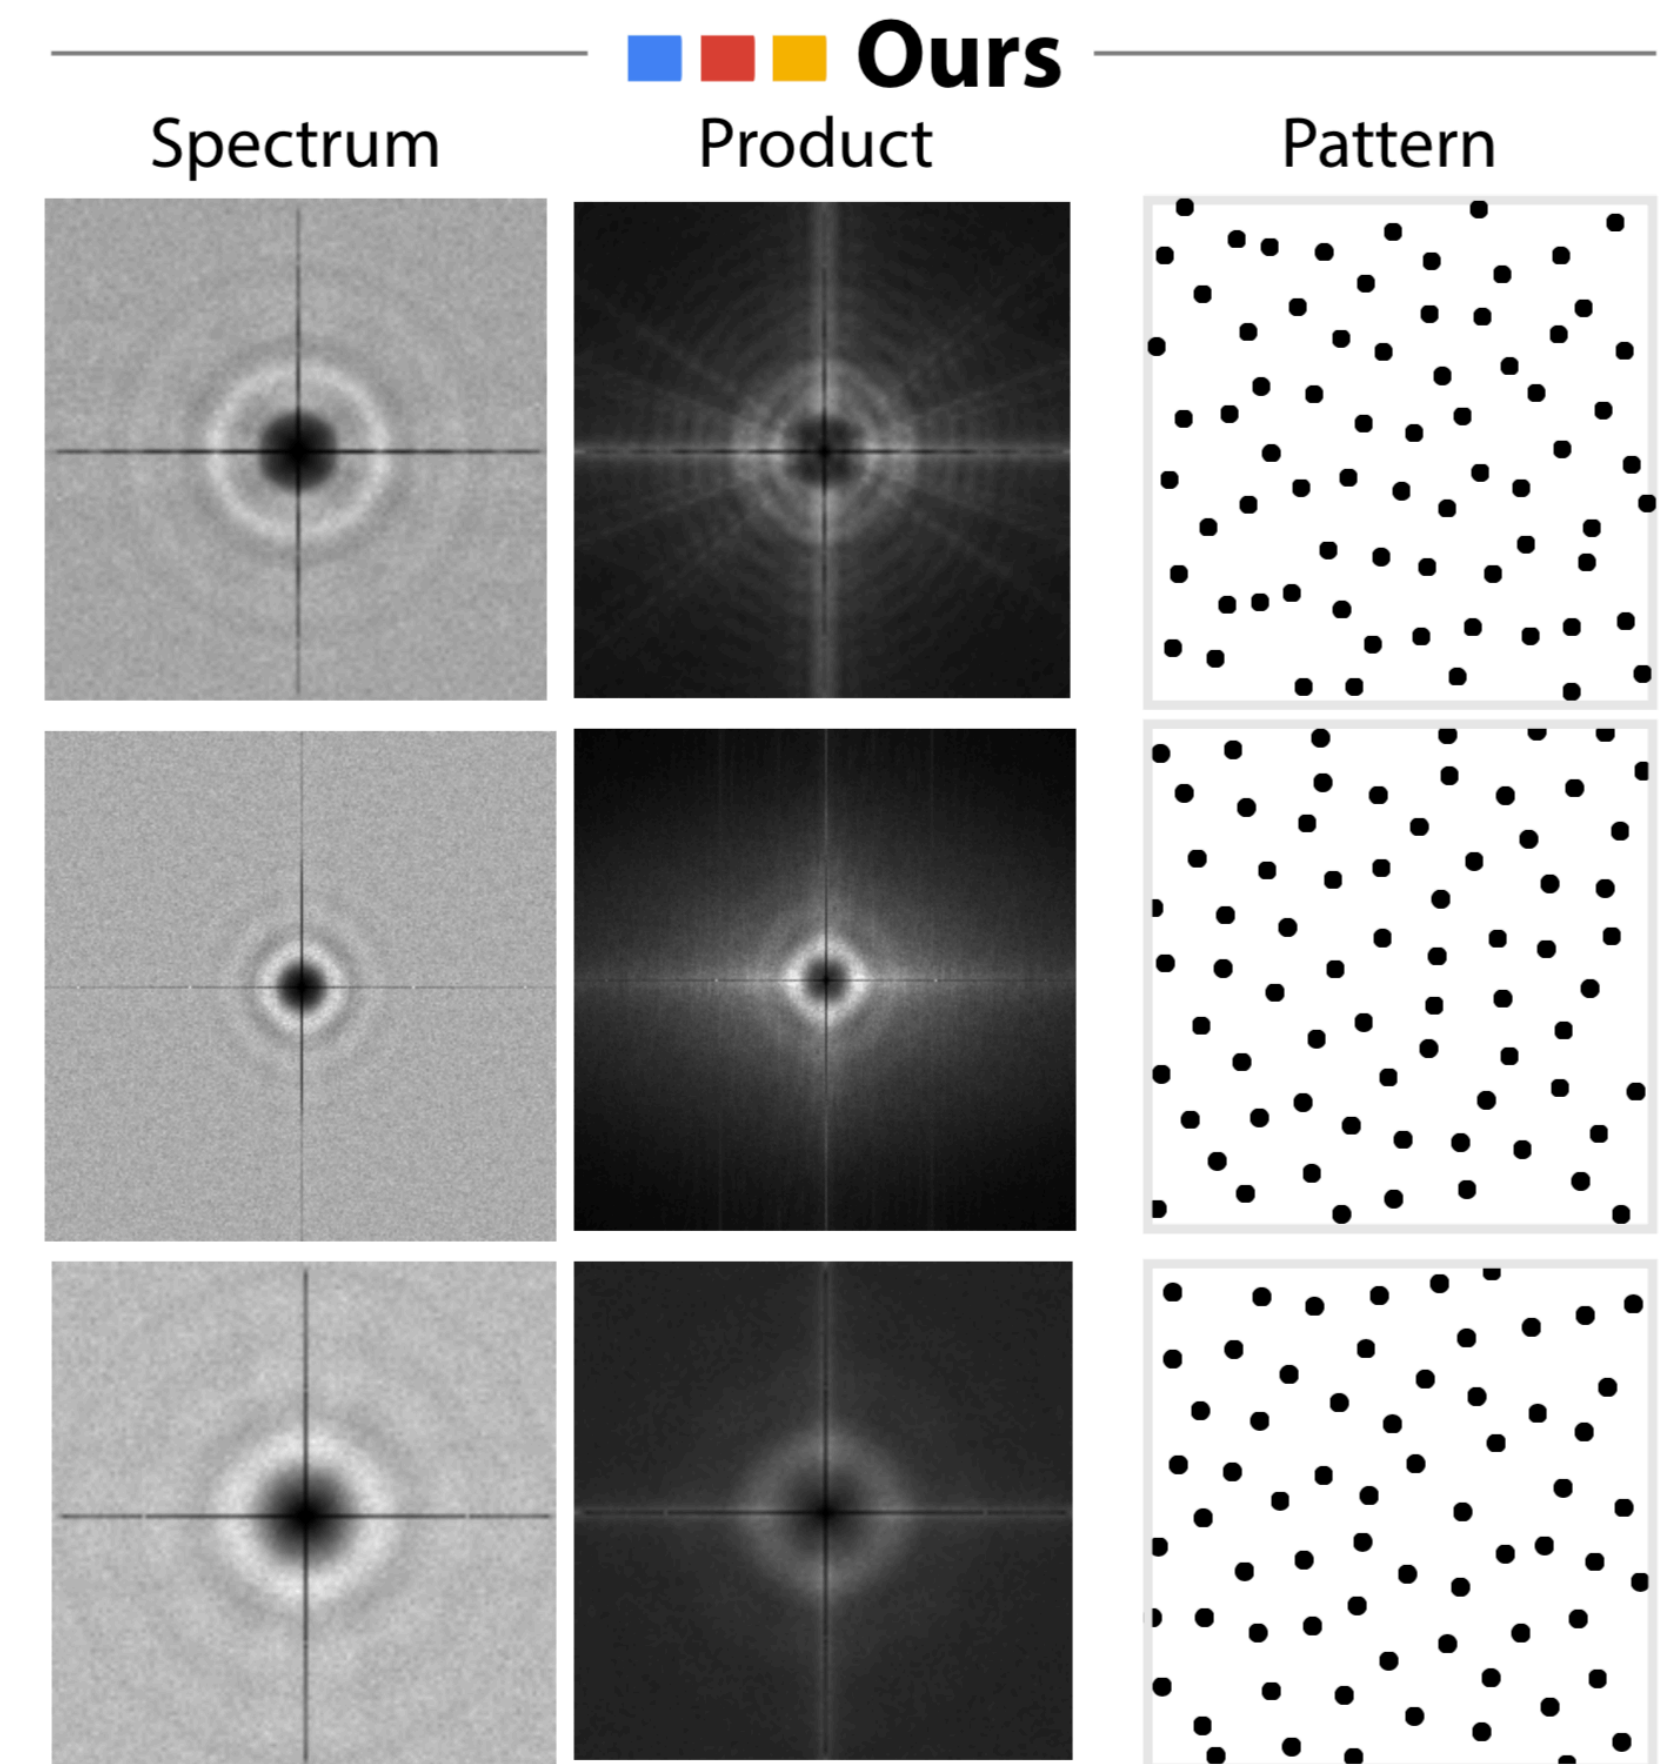
$$\text{Var}(I_N) = \sum_{\Omega} \text{Sampling Power spectrum} \times \text{Integrand Power spectrum}$$

Leimkuhler et al. [SIGGRAPH Asia 2019]

Novel Sampling Patterns




$$\text{Var}(I_N) = \sum_{\Omega} \text{[Spectrum]} \times \text{[Spectrum]}$$



Leimkuhler et al. [SIGGRAPH Asia 2019]

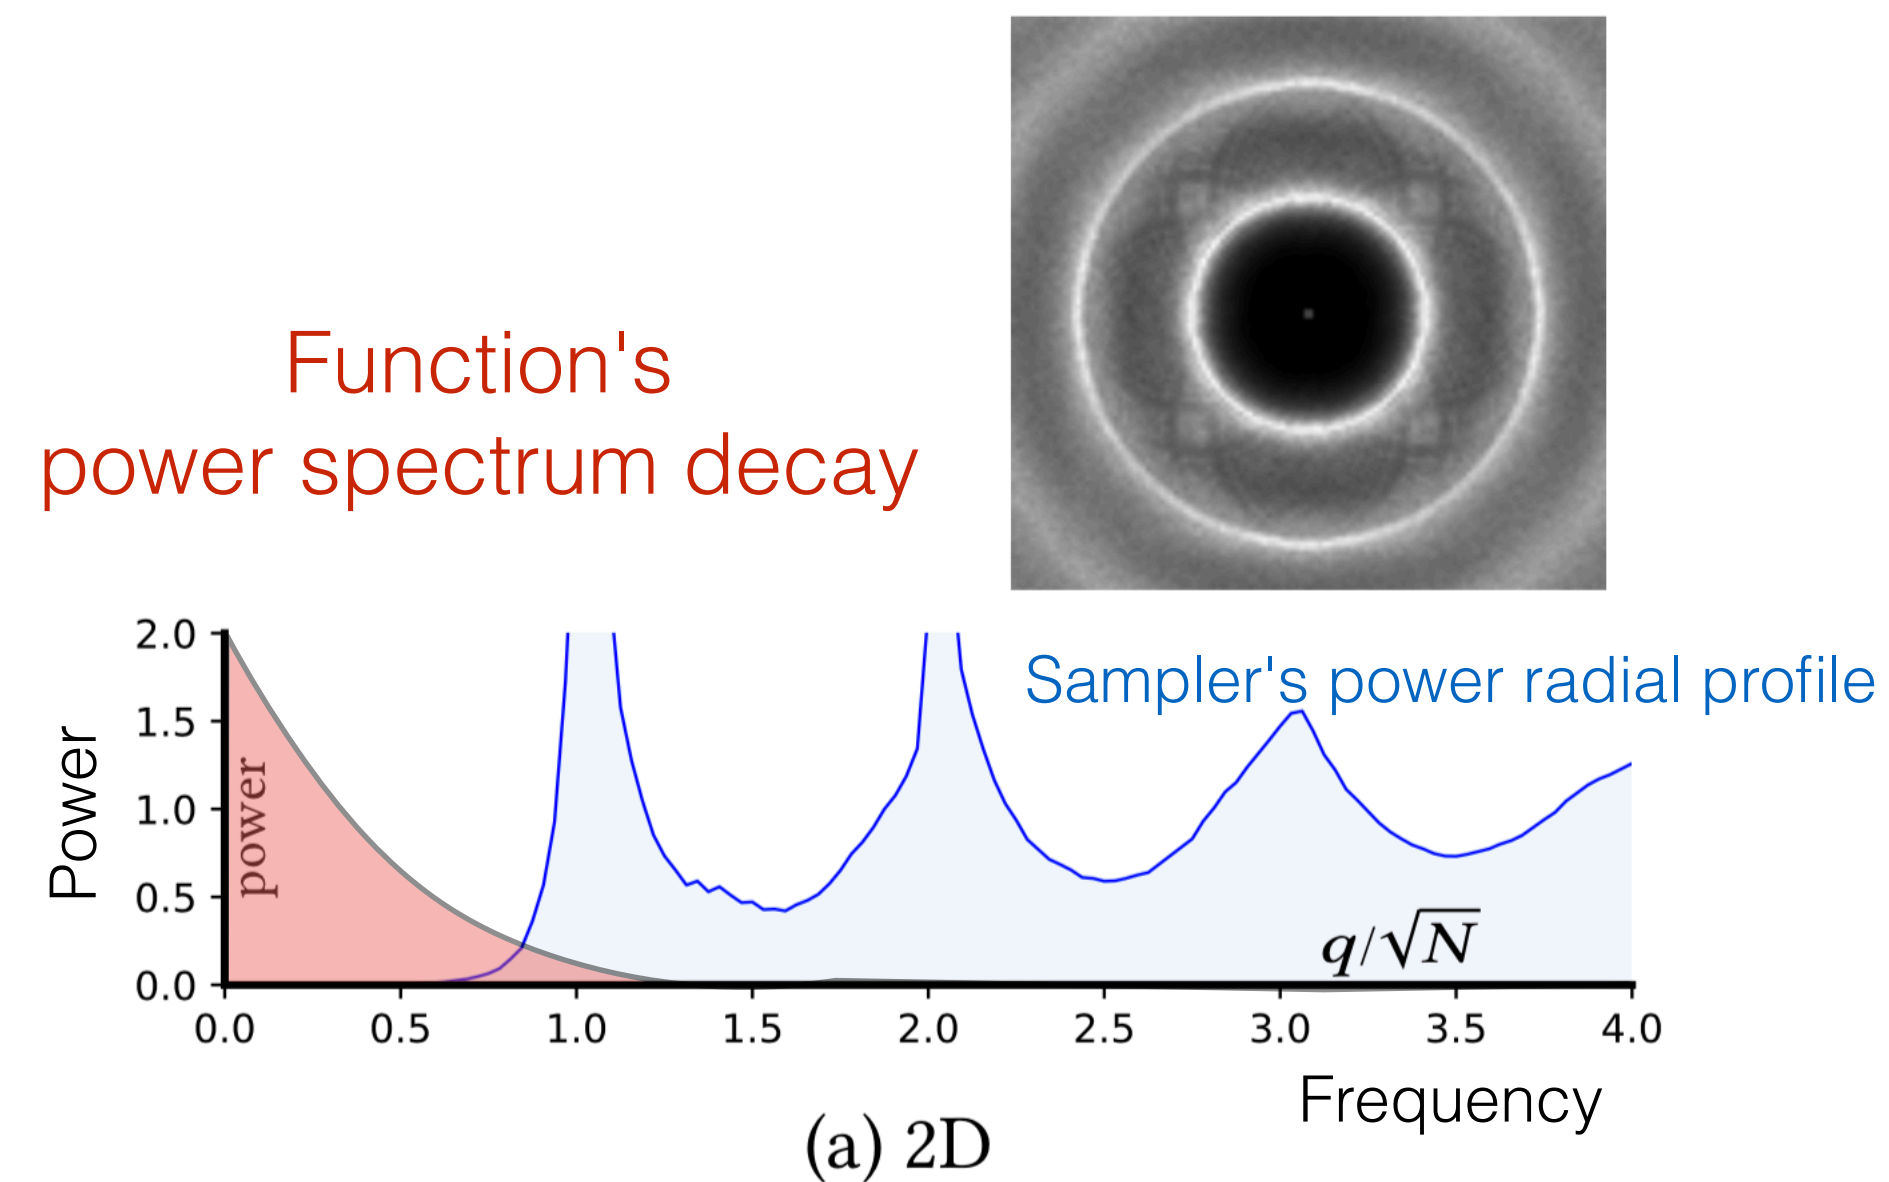
Blue Noise Dithering



Leimkuhler et al. [SIGGRAPH Asia 2019]

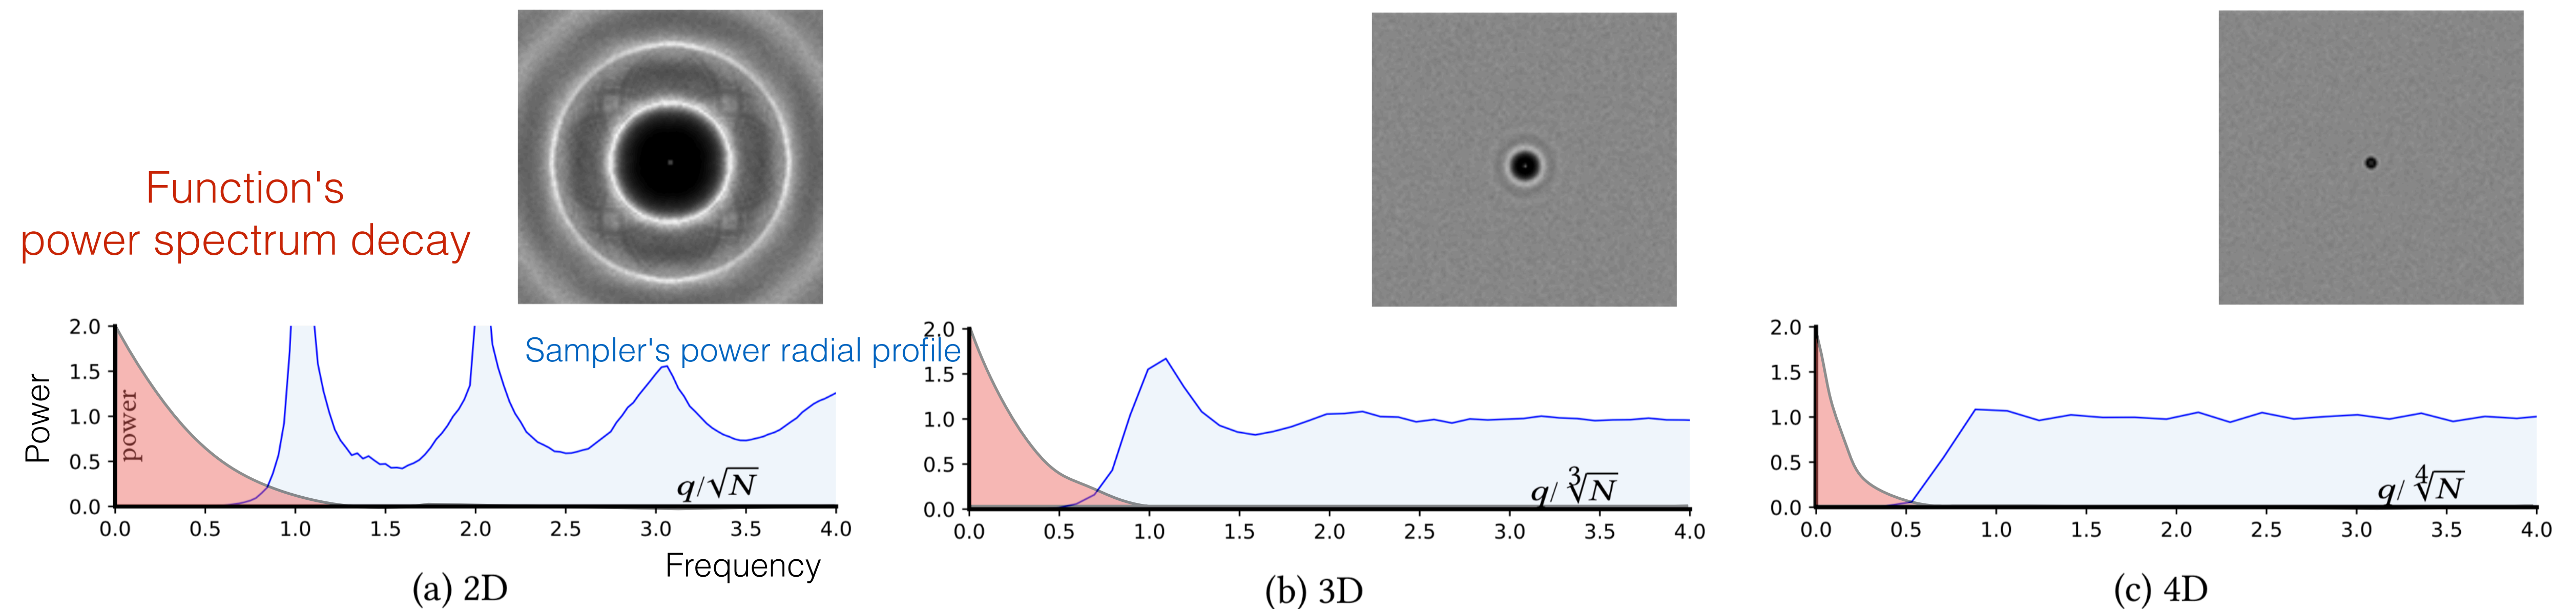
Novel Sampling Patterns

using radially averaged loss





Novel Sampling Patterns

using radially averaged loss




using radially averaged loss

Blue Noise Dithering




Normalizing Flows

Importance Sampling



$$I_N = \frac{1}{N} \sum_{k=1}^N \frac{f(x)}{p(x)}$$

Importance Sampling



$$I_N = \frac{1}{N} \sum_{k=1}^N \frac{f(x)}{p(x)}$$

$p(x) = ???$

Normalizing Flows



Normalizing Flows

Technique used in Machine learning to build complex probability distributions by transforming simple ones


Used in the context of generative modeling

Generative modeling: learning without any target (unsupervised)

Complex Probability distributions from simple ones



Complex Probability distributions from simple ones



Normalizing Flows:


Basic mathematical framework

$z \sim p_\theta(z)$ Given a continuous variable with a distribution


$x = f_\theta(z) = f_k \circ \dots \circ f_2 \circ f_1(z)$ New distribution obtained

each f_i is invertible (bijective)

Distributions



Distributions



Distributions

$z \sim p_\theta(z)$ Given a continuous variable with a distribution

$$x = f_\theta(z) = f_k \circ \dots \circ f_2 \circ f_1(z)$$

each f_i is invertible (bijective)

$$p(x) \neq p(f^{-1}(x))$$

Change of Variables

$f : Z \rightarrow X$, f is invertible

$p(z)$ defined over $z \in Z$

Change of variable formula says that:

$$p(x) = p(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}(x)}{\partial x} \right) \right|$$

Change of Variables

$f : Z \rightarrow X$, f is invertible

$p(z)$ defined over $z \in Z$

$$p(x) = p(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}(x)}{\partial x} \right) \right|$$

$$p(x) = p(z) \left| \det \left(\frac{\partial z}{\partial x} \right) \right|$$


Jacobian Matrix

$$\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$$

$$\mathbf{J} = \begin{bmatrix} \frac{\partial \mathbf{f}}{\partial x_1} & \dots & \frac{\partial \mathbf{f}}{\partial x_n} \end{bmatrix} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \dots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \dots & \frac{\partial f_m}{\partial x_n} \end{bmatrix}$$

Jacobian Matrix

$$\mathbf{f} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$$



Jacobian determinant gives the ratio of the area of the approximating parallelogram to that of the original square.

Jacobian Matrix

$f : Z \rightarrow X$, f is invertible

$p(z)$ defined over $z \in Z$



$$p(x) = p(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}(x)}{\partial x} \right) \right|$$

$$p(x) = p(z) \left| \det \left(\frac{\partial z}{\partial x} \right) \right|$$

Invertible and Differentiable Mapping

$f : Z \rightarrow X$, f is invertible
 $p(z)$ defined over $z \in Z$

$$\det \left(\frac{\partial x}{\partial z} \right) = 4$$



$$p(x) = p(z) = \left| \frac{1}{\det \left(\frac{\partial x}{\partial z} \right)} \right|$$

Maximize Log-likelihood

$f : Z \rightarrow X$, f is invertible
 $p(z)$ defined over $z \in Z$

$$\log p(x) = \log p(z) + \log \left| \det \left(\frac{\partial f^{-1}(x)}{\partial x} \right) \right|$$

Maximize Log-likelihood

$f : Z \rightarrow X$, f is invertible
 $p(z)$ defined over $z \in Z$

$$\log p(x) = \log p(z) + \log \left| \det \left(\frac{\partial f^{-1}(x)}{\partial x} \right) \right|$$

$$\log p(x) = \log p(z) + \sum_{i=1}^K \log \left| \det \left(\frac{\partial f^{-1}(x)}{\partial x} \right) \right|$$

Jacobian: Lower Triangular Matrix

$f : Z \rightarrow X$, f is invertible
 $p(z)$ defined over $z \in Z$

$$\mathbf{J} = \begin{bmatrix} \ell_{1,1} & & & & 0 \\ \ell_{2,1} & \ell_{2,2} & & & \\ \ell_{3,1} & \ell_{3,2} & \ddots & & \\ \vdots & \vdots & \ddots & \ddots & \\ \ell_{n,1} & \ell_{n,2} & \dots & \ell_{n,n-1} & \ell_{n,n} \end{bmatrix}$$

How to ensure lower-triangular Jacobian matrix?

$$z \in \mathbb{R}^D$$

$$z_{1:d}$$

$$z_{d+1:D}$$

NICE: NON-LINEAR INDEPENDENT COMPONENTS ESTIMATION

Laurent Dinh David Krueger Yoshua Bengio*
Département d'informatique et de recherche opérationnelle
Université de Montréal
Montréal, QC H3C 3J7

ABSTRACT

We propose a deep learning framework for modeling complex high-dimensional densities called Non-linear Independent Component Estimation (NICE). It is based on the idea that a good representation is one in which the data has a distribution that is easy to model. For this purpose, a non-linear deterministic transformation of the data is learned that maps it to a latent space so as to make the transformed data conform to a factorized distribution, i.e., resulting in independent latent variables. We parametrize this transformation so that computing the determinant of the Jacobian and inverse Jacobian is trivial, yet we maintain the ability to learn complex non-linear transformations, via a composition of simple building blocks, each based on a deep neural network. The training criterion is simply the exact log-likelihood, which is tractable. Unbiased ancestral sampling is also easy. We show that this approach yields good generative models on four image datasets and can be used for inpainting.

1 INTRODUCTION

One of the central questions in unsupervised learning is how to capture complex data distributions that have unknown structure. Deep learning approaches (Bengio, 2009) rely on the learning representation of the data that would capture its most important factors of variation. This raises the question: *what is a good representation?* Like in recent work (Kingma and Welling, 2014; Real et al., 2014; Ozair and Bengio, 2014), we take the view that a good representation is one in which the distribution of the data is easy to model. In this paper, we consider the special case where the learner to find a transformation $h = f(x)$ of the data into a new space such that the distribution factorizes, i.e., the components h_d are independent:

$$p_H(h) = \prod_d p_{H_d}(h_d).$$

The proposed training criterion is directly derived from the log-likelihood. Moreover, we consider a change of variables $h = f(x)$, which assumes that f is invertible and its inverse is the same as the dimension of x , in order to fit a distribution p_H . The change of variables is given by:

$$p_X(x) = p_H(f(x)) \left| \det \frac{\partial f(x)}{\partial x} \right|.$$


where $\frac{\partial f(x)}{\partial x}$ is the Jacobian matrix of function f at x . In this paper, we consider a sample from $p_X(x)$ easily as follows:

$$h \sim p_H(h)$$
$$x = f^{-1}(h)$$

of such a transformation, its inverse, while a

How to ensure lower-triangular Jacobian matrix?

Coupling layer



$$z \in \mathbb{R}^D$$

NICE: NON-LINEAR INDEPENDENT COMPONENTS ESTIMATION

Laurent Dinh, David Krueger, Yoshua Bengio*
Département d'informatique et de recherche opérationnelle
Université de Montréal
Montréal, QC H3C 3J7

ABSTRACT

We propose a deep learning framework for modeling complex high-dimensional densities called Non-linear Independent Component Estimation (NICE). It is based on the idea that a good representation is one in which the data has a distribution that is easy to model. For this purpose, a non-linear deterministic transformation of the data is learned that maps it to a latent space so as to make the transformed data conform to a factorized distribution, i.e., resulting in independent latent variables. We parametrize this transformation so that computing the determinant of the Jacobian and inverse Jacobian is trivial, yet we maintain the ability to learn complex non-linear transformations, via a composition of simple building blocks, each based on a deep neural network. The training criterion is simply the exact log-likelihood, which is tractable. Unbiased ancestral sampling is also easy. We show that this approach yields good generative models on four image datasets and can be used for inpainting.

1 INTRODUCTION

One of the central questions in unsupervised learning is how to capture complex data distributions that have unknown structure. Deep learning approaches (Bengio, 2009) rely on the learning representation of the data that would capture its most important factors of variation. This raises the question: *what is a good representation?* Like in recent work (Kingma and Welling, 2014; Real et al., 2014; Ozair and Bengio, 2014), we take the view that a *good representation* is one in which the distribution of the data is *easy to model*. In this paper, we consider the special case where we ask the learner to find a transformation $h = f(x)$ of the data into a new space such that the distribution factorizes, i.e., the components h_d are independent:

$$p_H(h) = \prod_d p_{H_d}(h_d).$$

The proposed training criterion is directly derived from the log-likelihood. Moreover, we consider a change of variables $h = f(x)$, which assumes that f is invertible and its inverse is the same as the dimension of x , in order to fit a distribution p_H . The change of variables is:

$$p_X(x) = p_H(f(x)) \left| \det \frac{\partial f(x)}{\partial x} \right|.$$

where $\frac{\partial f(x)}{\partial x}$ is the Jacobian matrix of function f at x . In this paper, we consider a sample from $p_X(x)$ easily as follows:

$$h \sim p_H(h)$$


us: $x = f^{-1}(h)$

Neural Importance Sampling

Thomas Müller
Brian McWilliams
Fabrice Rousselle
Markus Gross
Jan Novák


Affiliation:  **NVIDIA.**


Work done while at:  **ETH** zürich  **Disney** Research  **cgl** computer graphics laboratory






Render time: sometimes >100 cpu-hours






512 paths per pixel



Path tracing

Neural path guiding


Path tracing: BSDF sampling




Path tracing: direct-illumination sampling

Multiple Importance Sampling

[Veach and Guibas 1995]




Where is path guiding useful?




Where is path guiding useful?


Goal: Sample proportional to incident radiance.



Where is path guiding useful?




Learning incident radiance in a Cornell box




Neural networks as function approximators


Reference




SD-tree [Müller et al. 2017]



Neural Network




GMM [Vorba et al. 2014]






Directional distribution


Neural network




Reference




SD-tree




Gaussian mixture



Neural path guiding overview




Path tracer



Neural network

Neural path guiding overview



Path tracer

Feedback loop


Sample



Optimize



How?




Neural network

How?



How to draw samples?



Path tracer


Sample



Feedback loop



Optimize



Neural network

Goal: warp random numbers to good distribution with NN

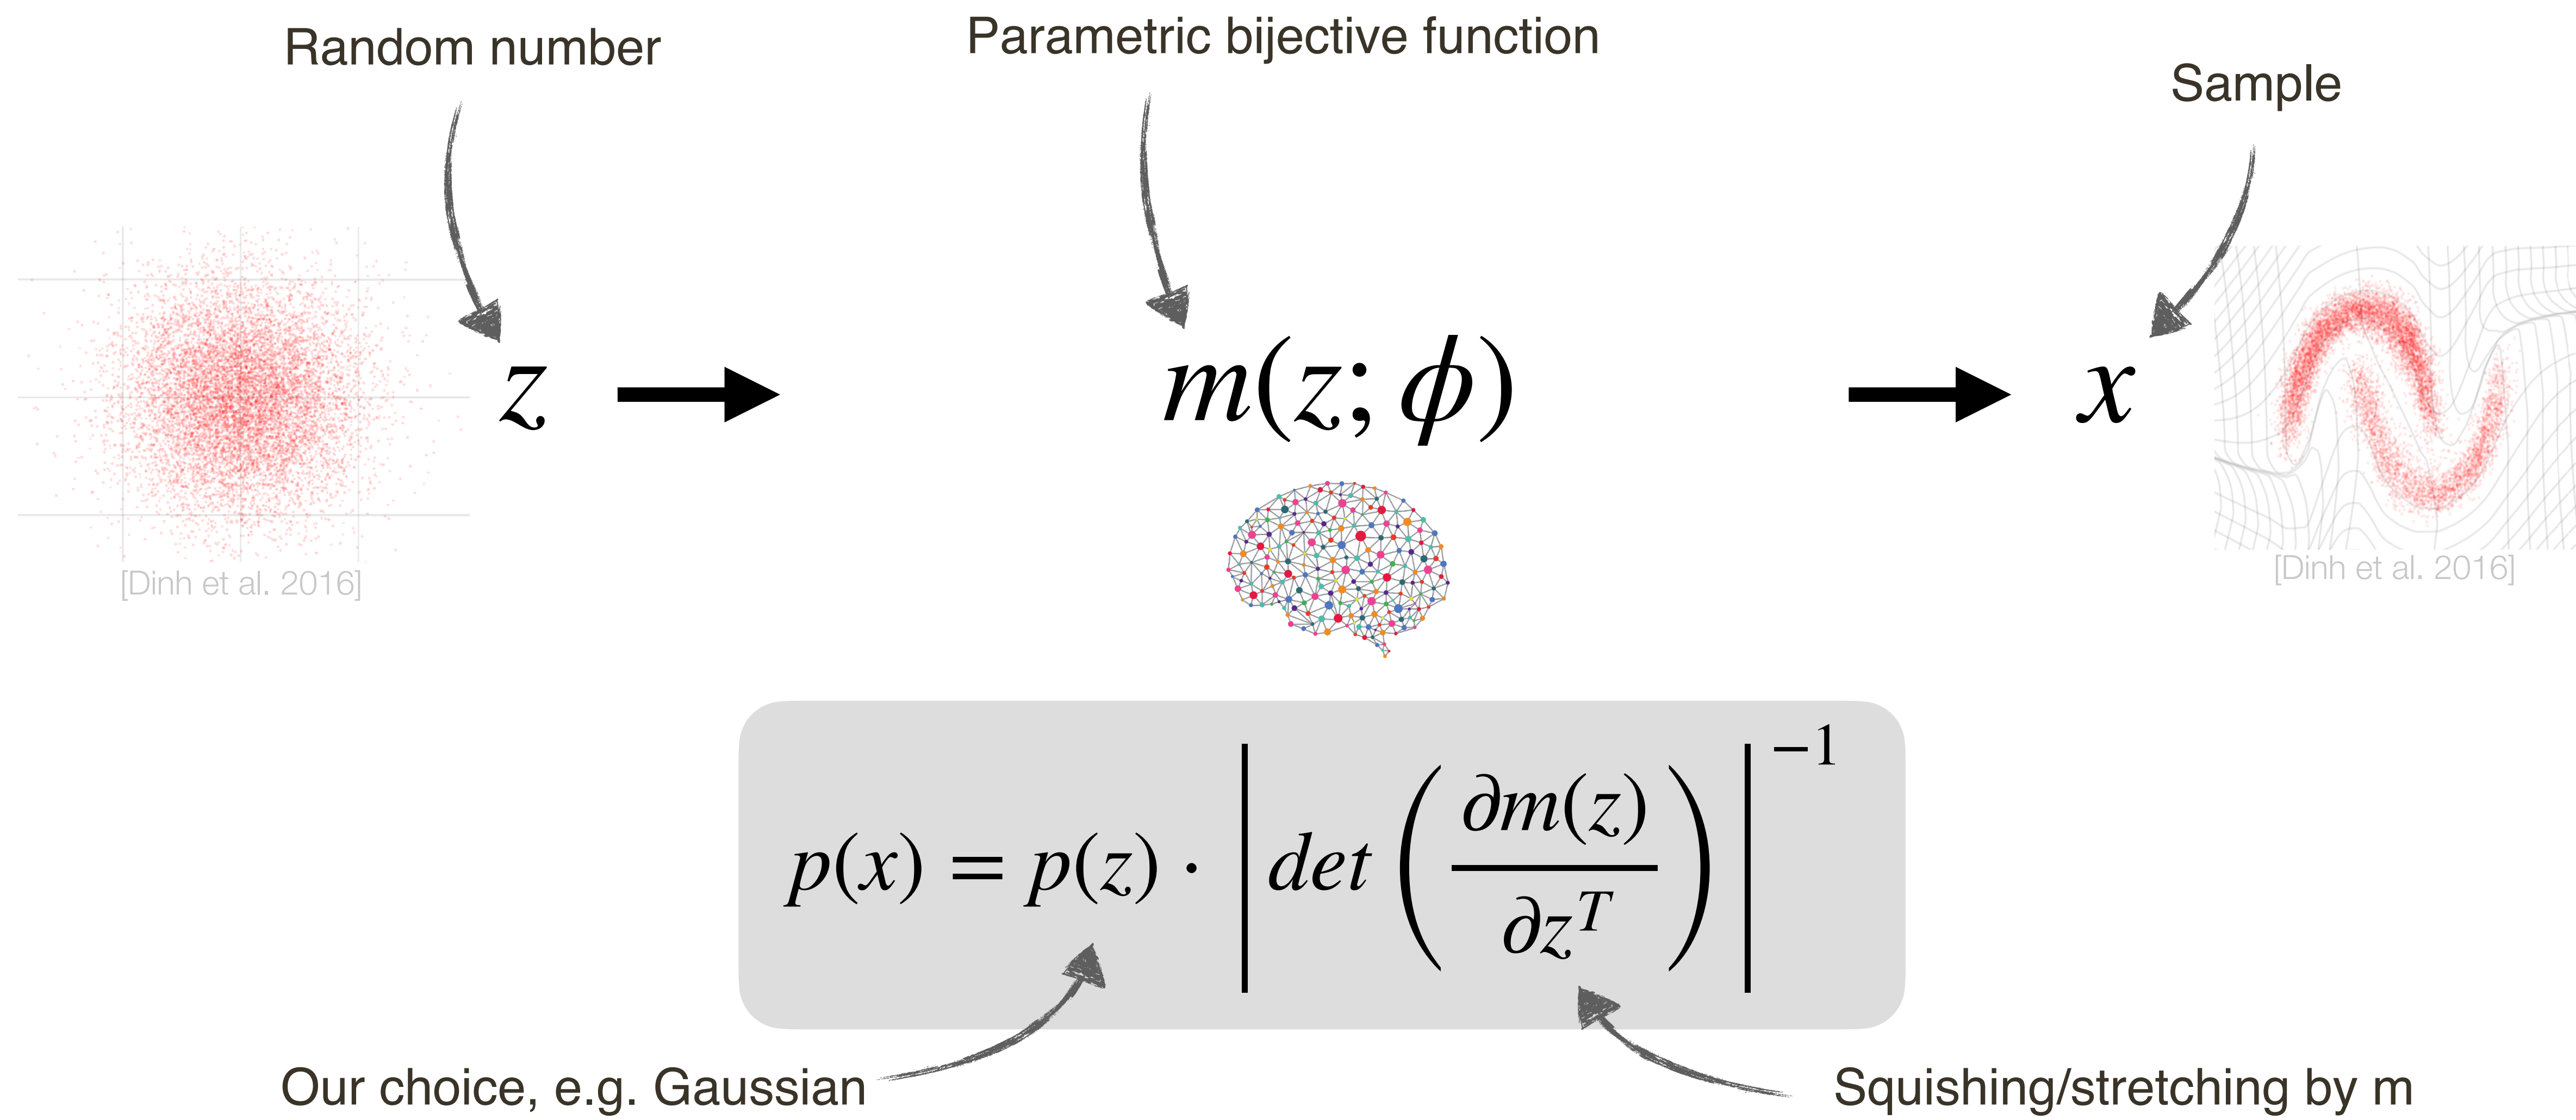


Monte Carlo estimator


$$F \approx \frac{1}{N} \sum_{i=1}^N \frac{f(X_i)}{p(X_i)}$$

Need p in closed form!
Addressed by "normalizing flows"

Parameterizing a bijection allows using the change-of-variable formula



A chain of simple bijections can model complicated functions




$$p(x) = p(z) \cdot \prod_{i=1}^L \left| \det \left(\frac{\partial m_i(z)}{\partial z^T} \right) \right|^{-1}$$

Our choice, e.g. Gaussian

Squishing/stretching by m

How to optimize?




Path tracer

Sample

Feedback loop

Optimize




Neural network

Training with data from the correct distribution is simple




Training data

Desired distribution



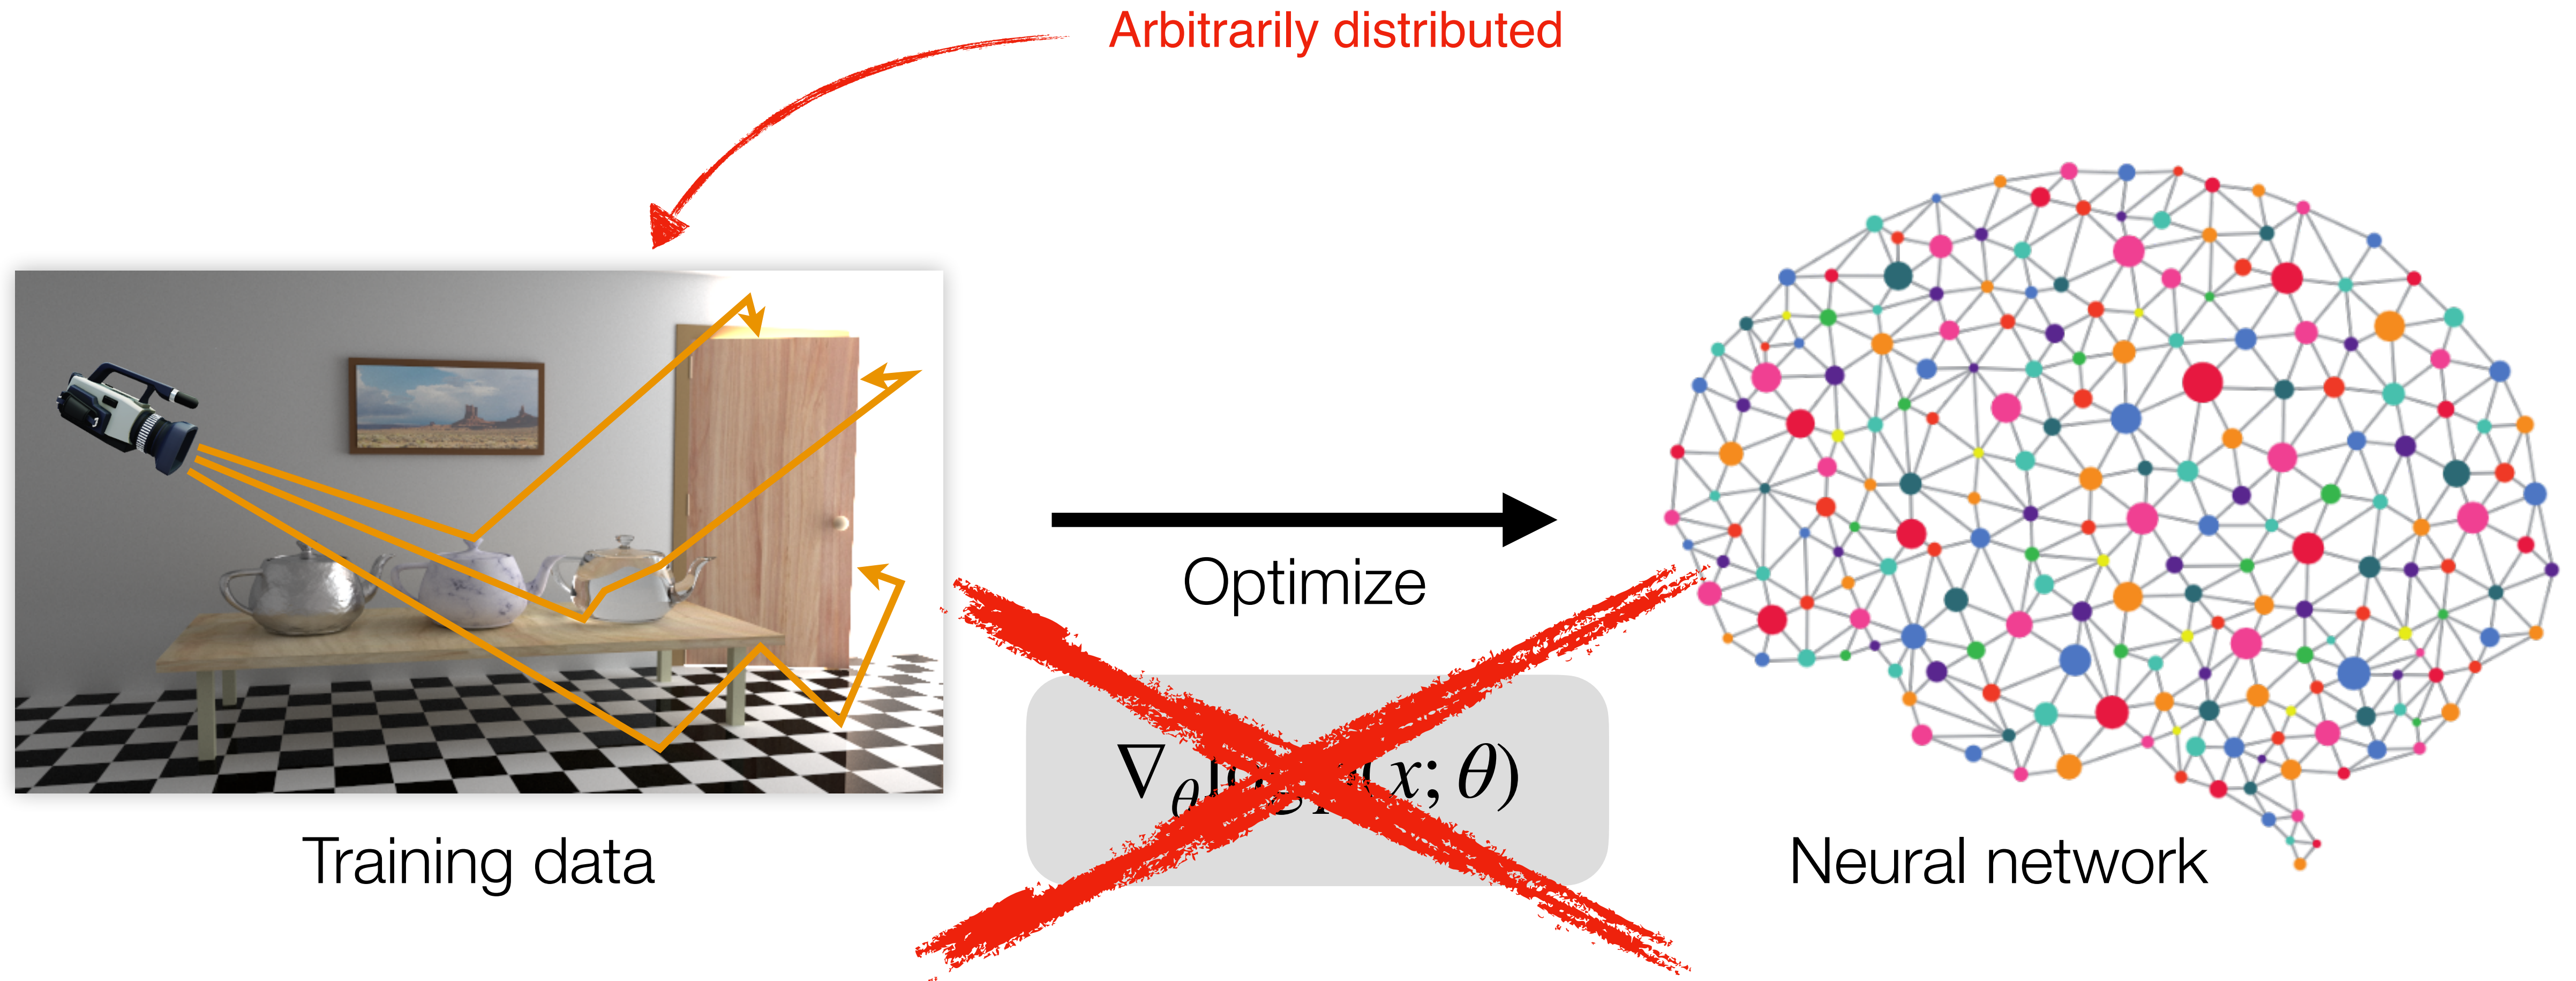
Optimize

$$\nabla_{\theta} \log p(x; \theta)$$



Neural network

Training from Monte Carlo samples requires careful weighting




Training with data from the correct distribution is simple




Training data

Desired distribution



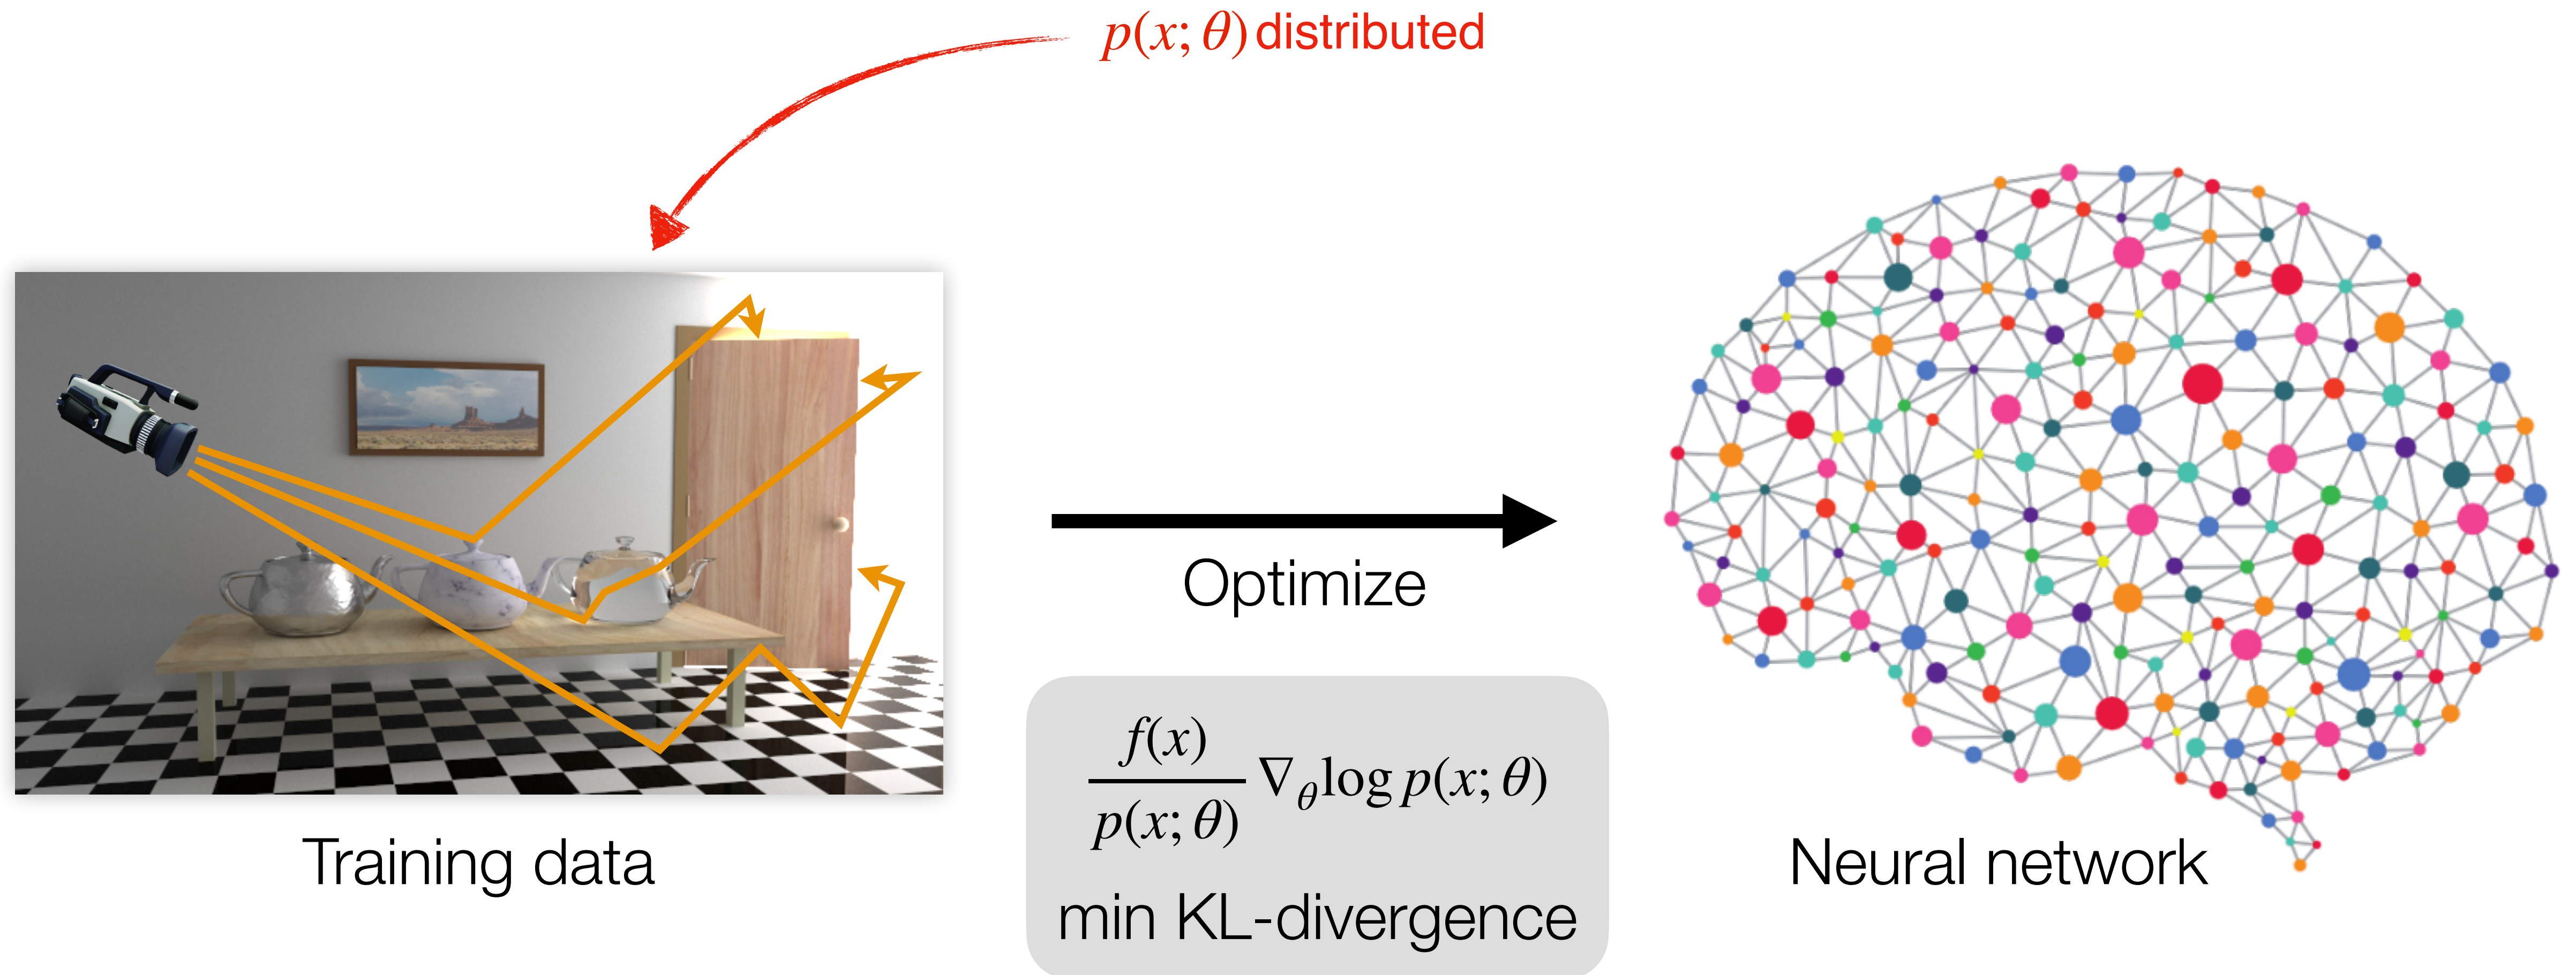
Optimize

$\nabla_{\theta} \log p(x; \theta)$
min KL-divergence

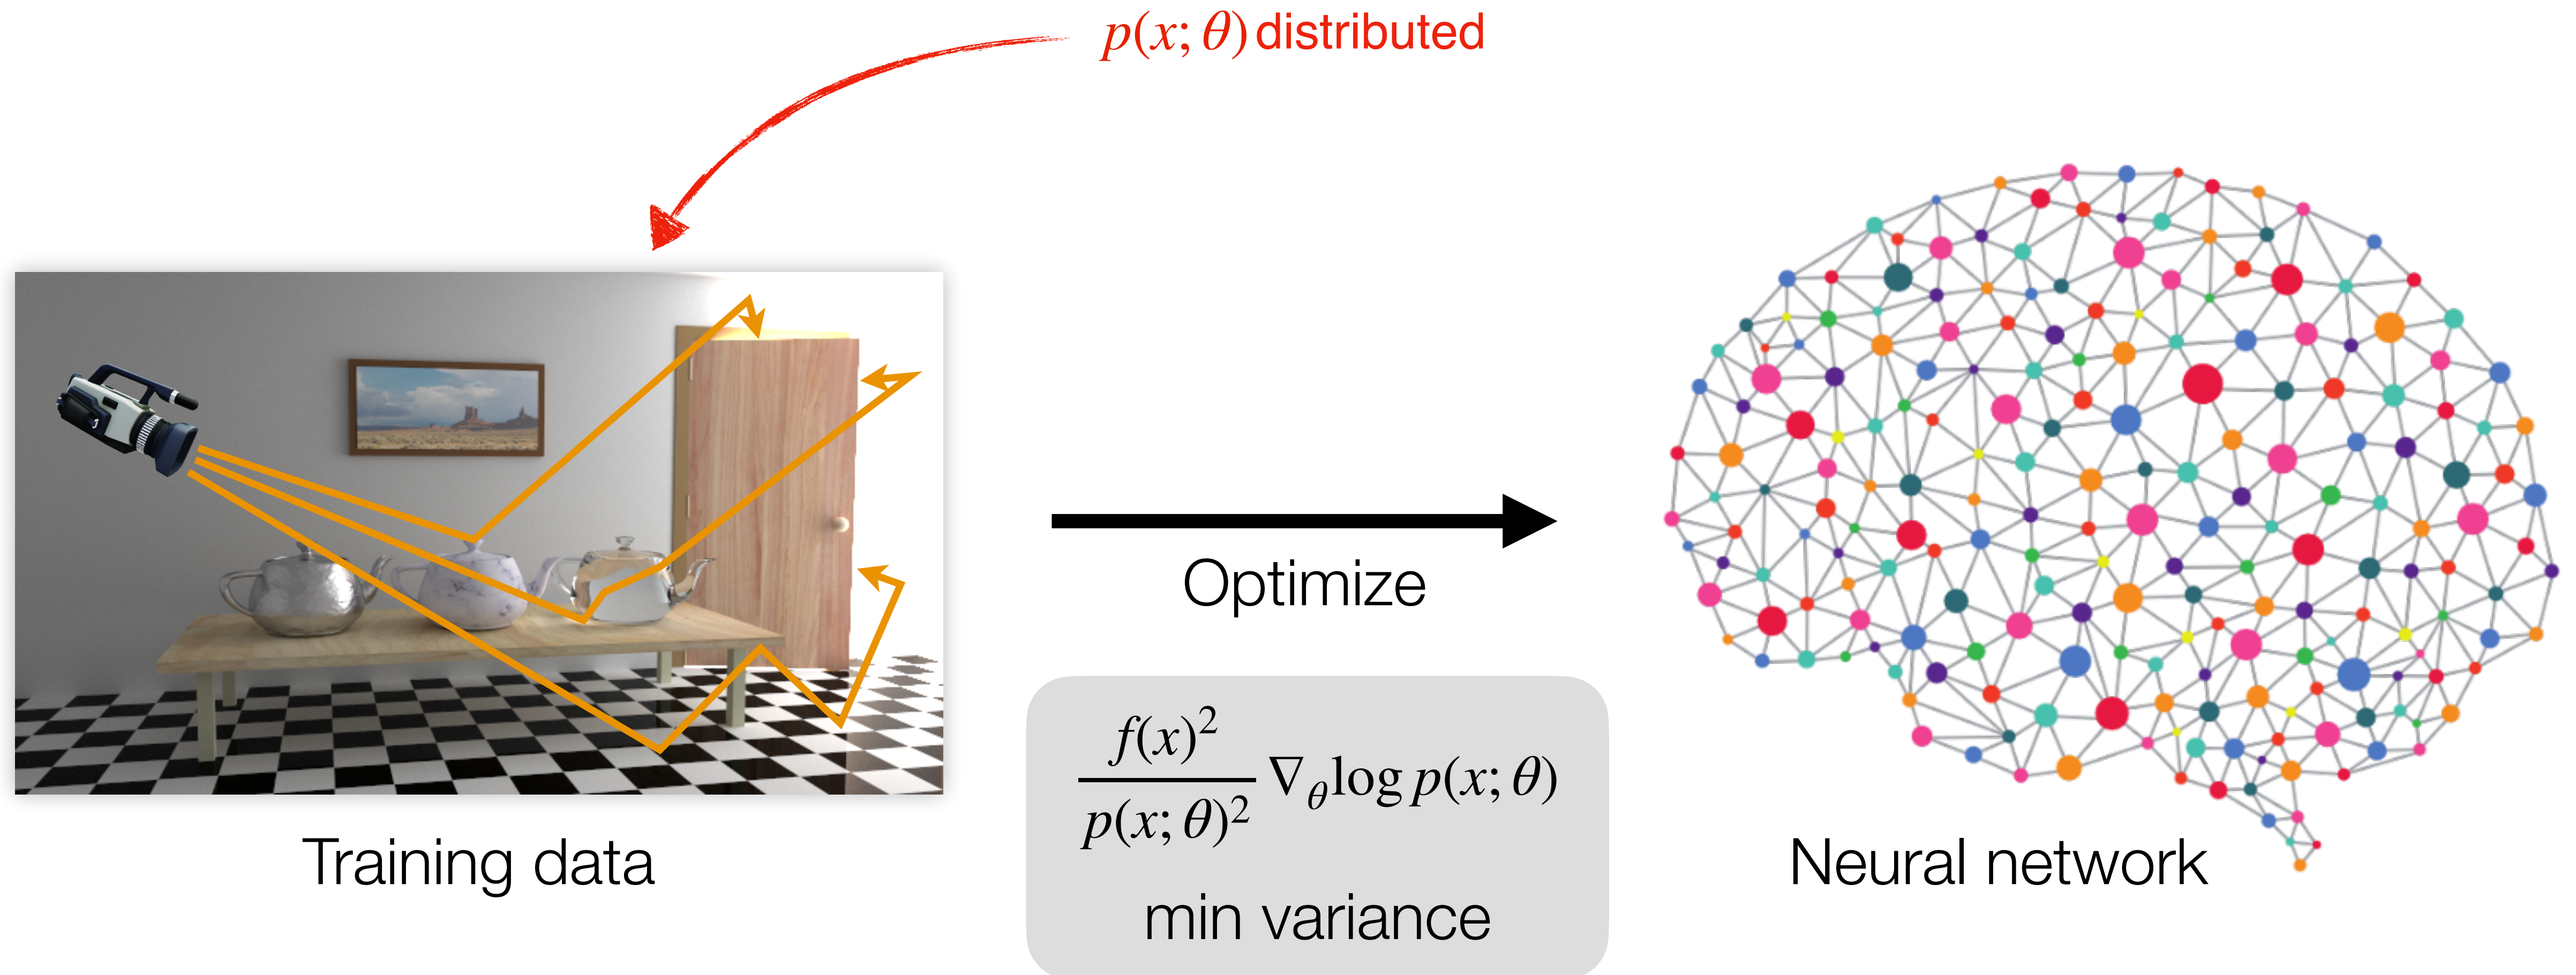


Neural network

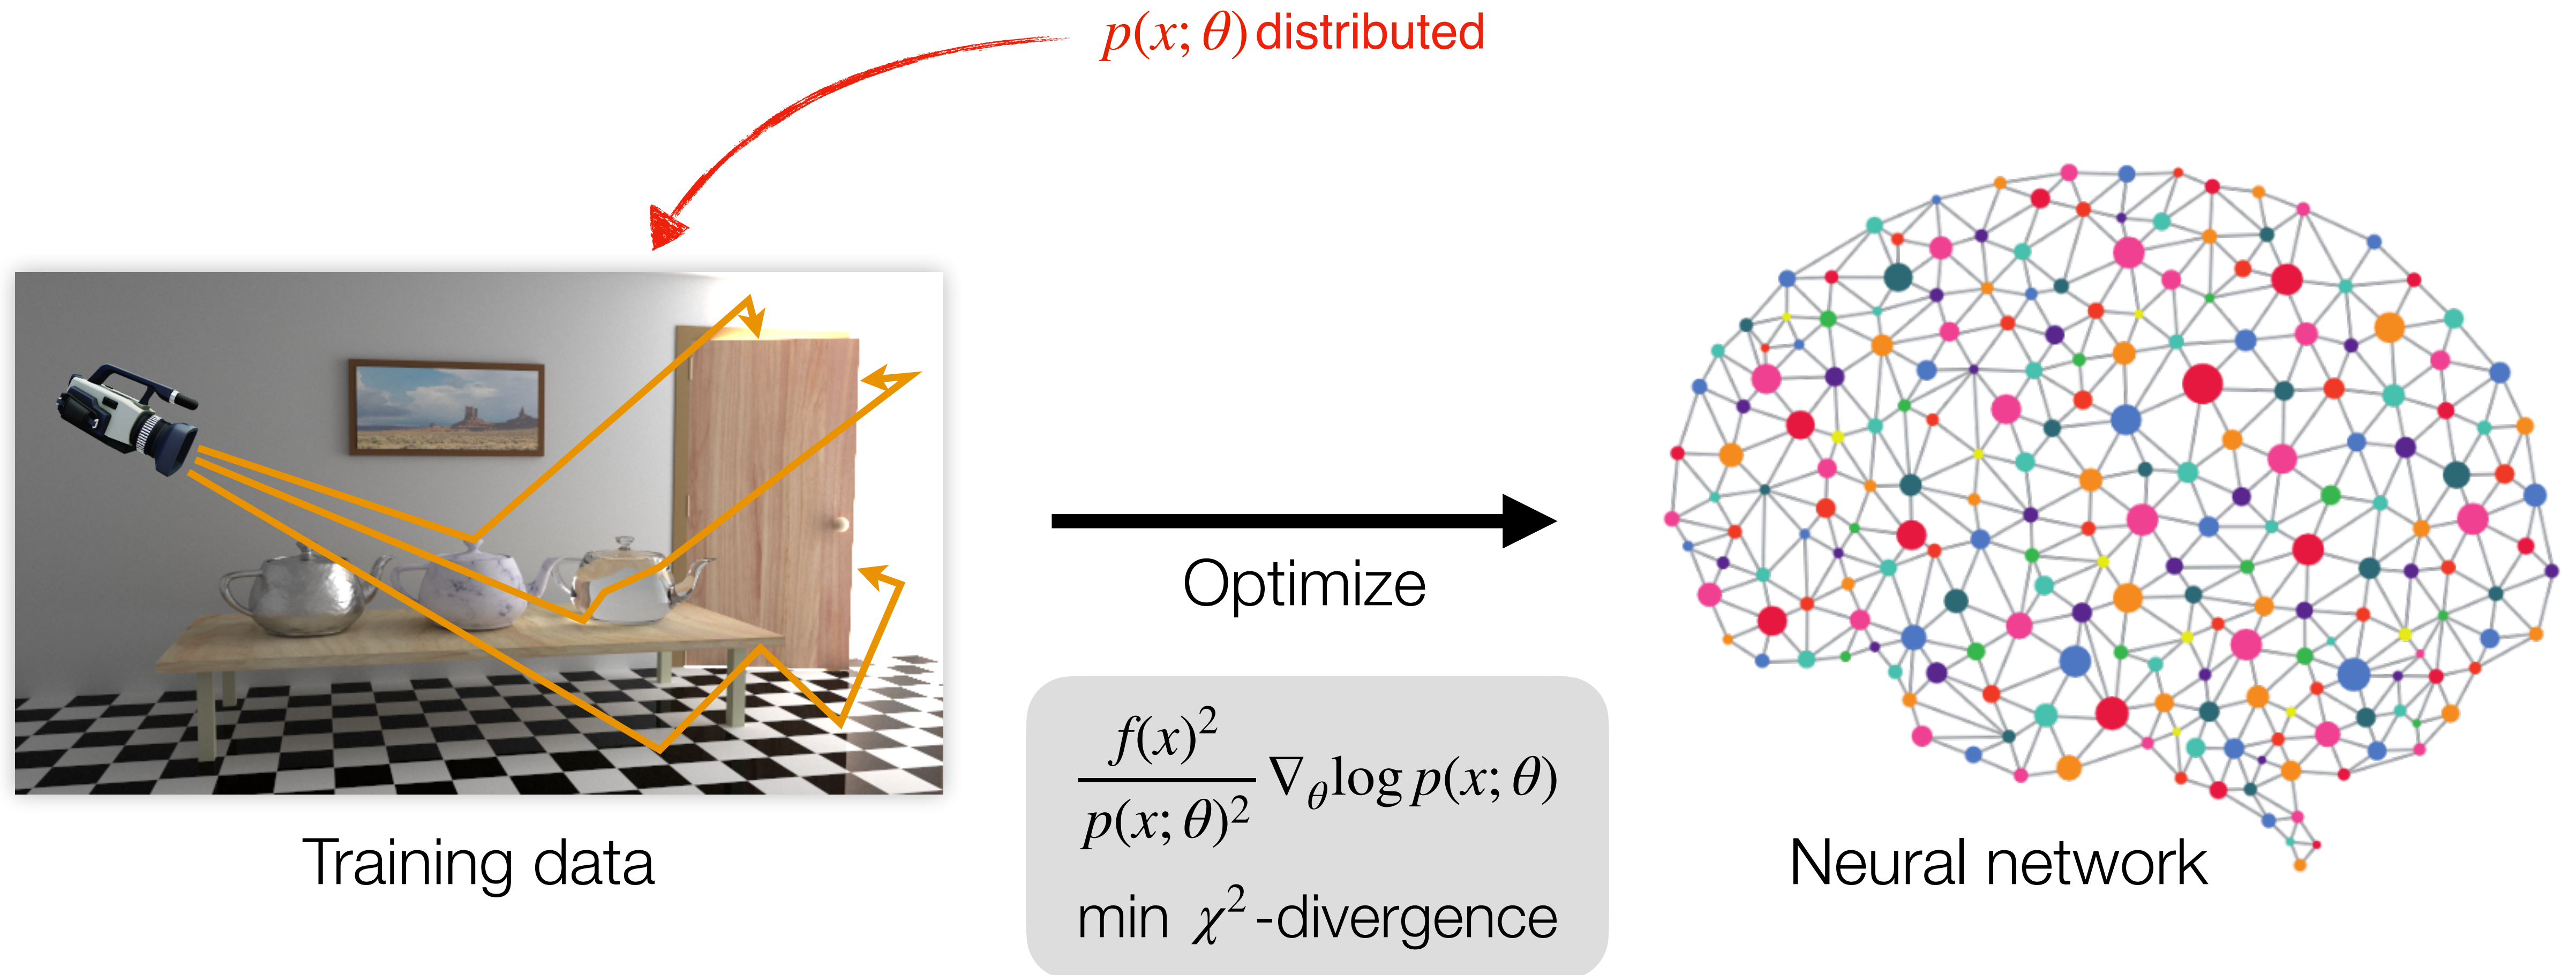
Training from Monte Carlo samples requires careful weighting



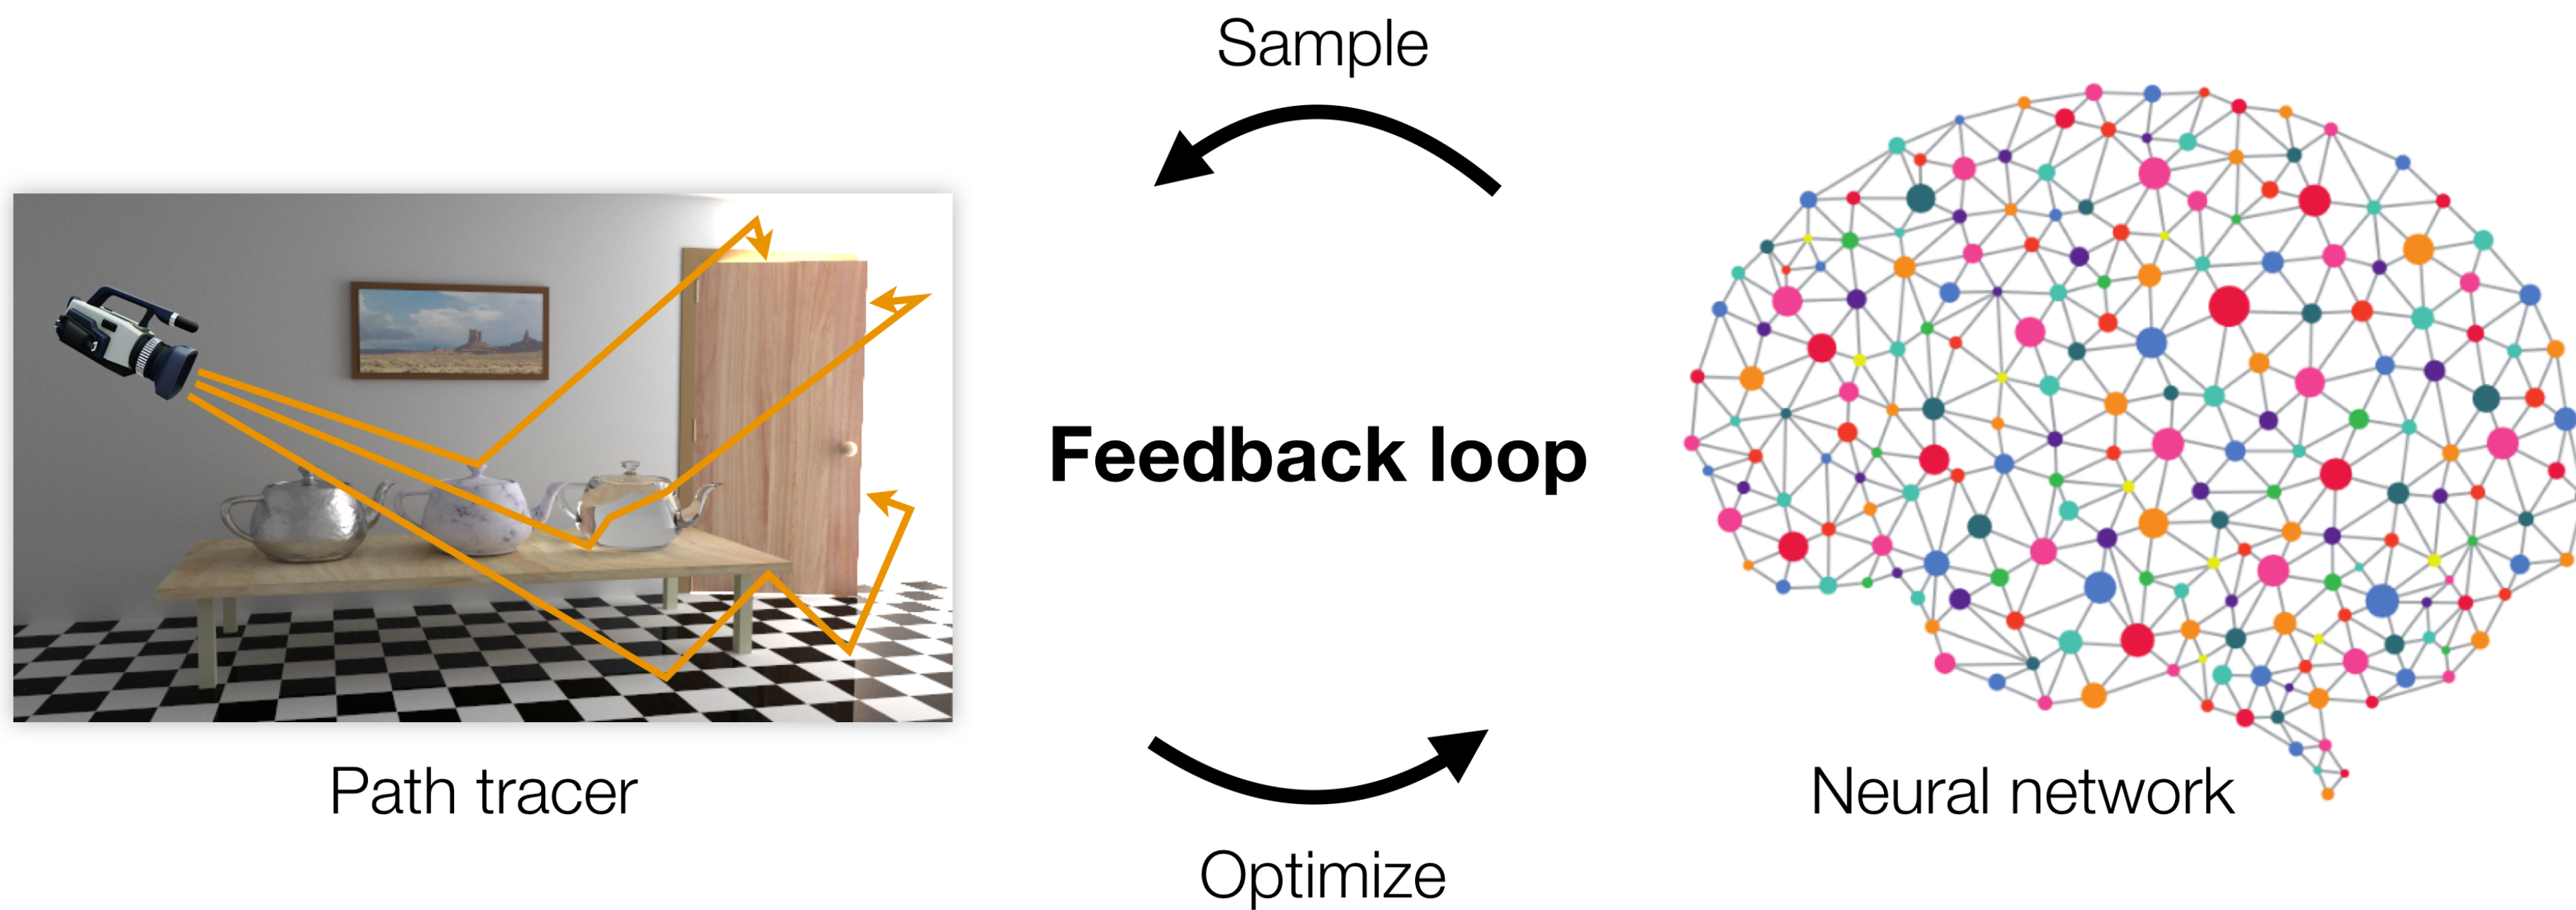
Training from Monte Carlo samples requires careful weighting



Training from Monte Carlo samples requires careful weighting



Putting it together...





1 path per pixel

Path tracing

Neural path guiding

2 paths per pixel

Path tracing


Neural path guiding

4 paths per pixel

Path tracing

Neural path guiding

8 paths per pixel




Path tracing

Neural path guiding

16 paths per pixel



Path tracing



Neural path guiding

32 paths per pixel

Path tracing

Neural path guiding

64 paths per pixel

Path tracing


Neural path guiding

128 paths per pixel

Path tracing

Neural path guiding

256 paths per pixel




Path tracing



Neural path guiding

512 paths per pixel




Path tracing

Neural path guiding

Product guiding


Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




Product path guiding

$$L_r(\mathbf{x}, \omega_o) = \int L_i(\mathbf{x}, \omega_i) f(\mathbf{x}, \omega_i, \omega_o) \cos \theta \, d\omega_i$$




MIS optimization


MIS-aware optimization



MIS-aware optimization



MIS-aware optimization



Results


Equal time

Path tracing


Müller et al. [2017]

Neural path guiding


Equal spp




Path tracing




Müller et al. [2017]



Neural path guiding



Path tracing



Müller et al. [2017]



Neural path guiding

Conclusion

- Neural networks can drive unbiased MC integration
- Complicated integrands (e.g. product path guiding)
- Computational cost of neural path guiding is high, but quality is state of the art

References

A frequency analysis of light transport, Durand et al. SIGGRAPH 2005

Frequency Analysis and Sheared Reconstruction for Rendering Motion Blur, Egan et al. SIGGRAPH 2009

Temporal Light Field Reconstruction for Rendering Distribution Effects, Lehtinen et al. SIGGRAPH 2011

On Filtering the Noise from the Random Parameters in Monte Carlo Rendering, Sen and Darabi 2012

A Machine Learning Approach for Filtering Monte Carlo Noise, Kalantari et al. SIGGRAPH 2015

Interactive Reconstruction of Monte Carlo Image Sequences using a Recurrent Denoising Autoencoder, Chaitanya et al. SIGGRAPH 2017

Kernel-Predicting Convolutional Networks for Denoising Monte Carlo Renderings, Bako et al. SIGGRAPH 2017

Sample-based Monte Carlo Denoising using a Kernel-Splatting Network, Gharbi et al. SIGGRAPH 2019

NICE: Non-linear Independent Components Estimation

Normalizing Flows: An Introduction and Review of Current Methods

Neural Importance Sampling SIGGRAPH 2019

Acknowledgements

I would like to thank Thomas Muller and colleagues to make their slides available online