SIGGRAPH2017
AT THE ♥ of COMPUTER GRAPHICS & INTERACTIVE TECHNIQUES

Convergence Analysis for Anisotropic Monte Carlo Sampling Spectra

Gurprit Singh **Wojciech Jarosz**

DARTMOUTH
VISUAL COMPUTING LAB

Thank you for being here. Today, I will be presenting our work on convergence analysis for Monte Carlo integration that is done in collaboration with Wojciech Jarosz in the visual computing lab at Dartmouth.

Monte Carlo Integration

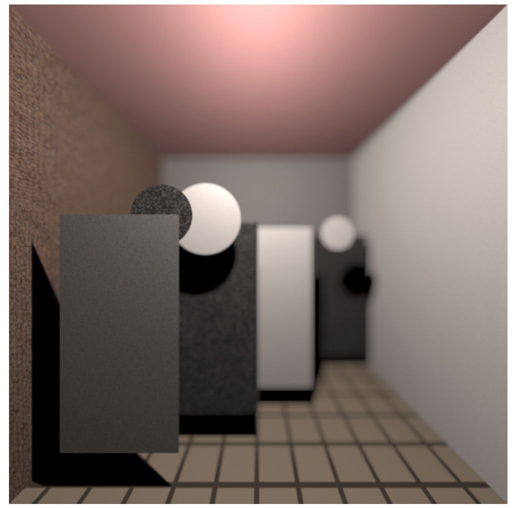

2

Monte Carlo integration is a numerical method employed to solve multi-dimensional integrals that we normally encounter in light transport problems during rendering. We are showing one such rendering example here, which is lit with a point light source, as a result…

Monte Carlo Integration

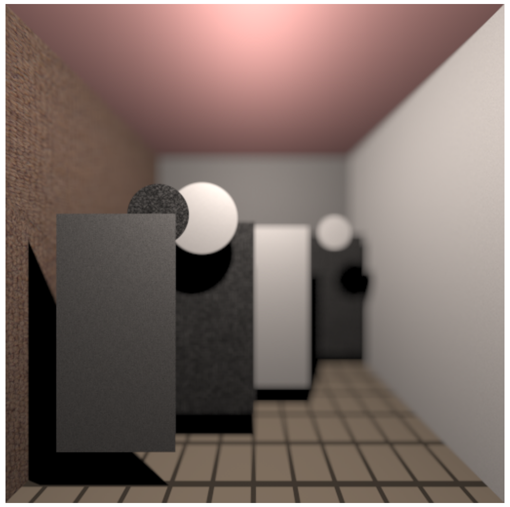

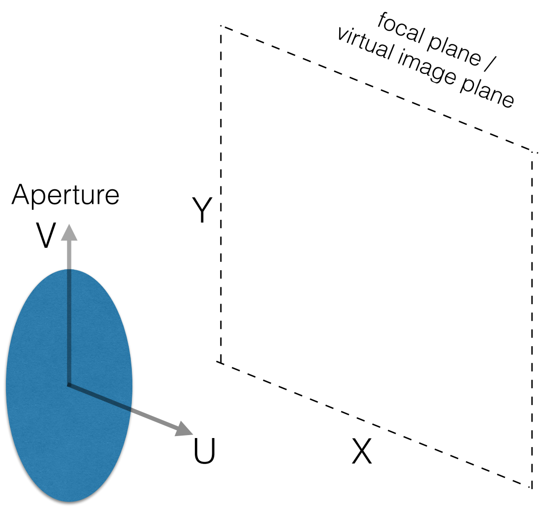

Monte Carlo Integration

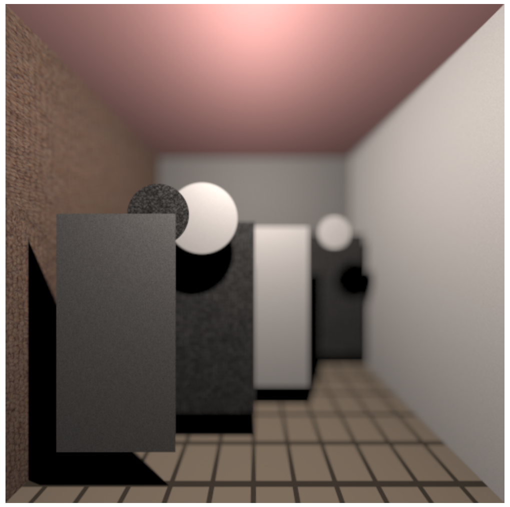

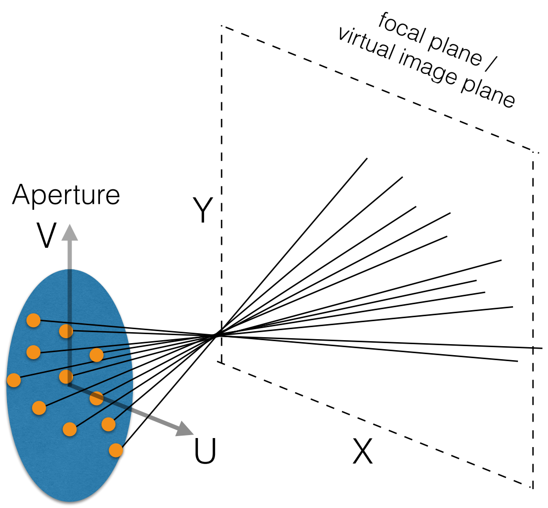

Monte Carlo Integration

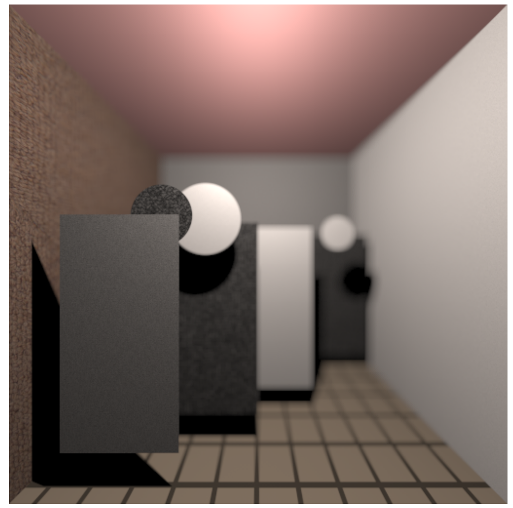

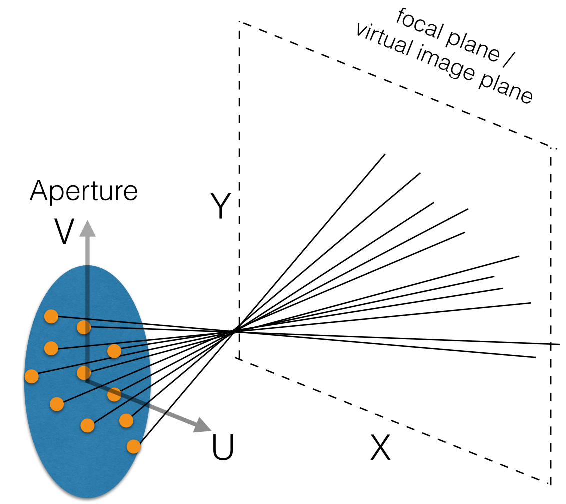

$$\int_x \int_y \int_u \int_v f(x, y, u, v)\; dv\; du\; dy\; dx$$

Monte Carlo Integration

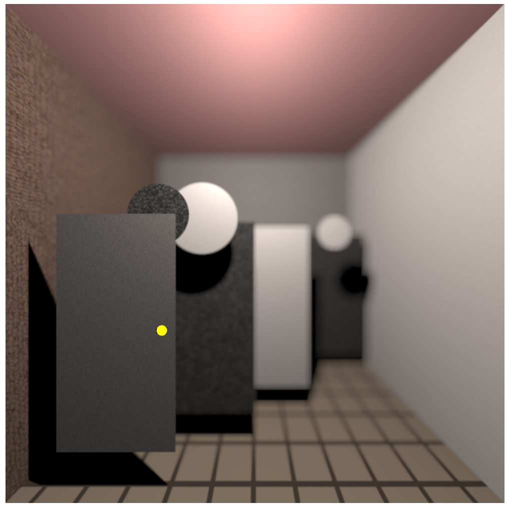

… an in-focus pixel has an underlying…

Monte Carlo Integration

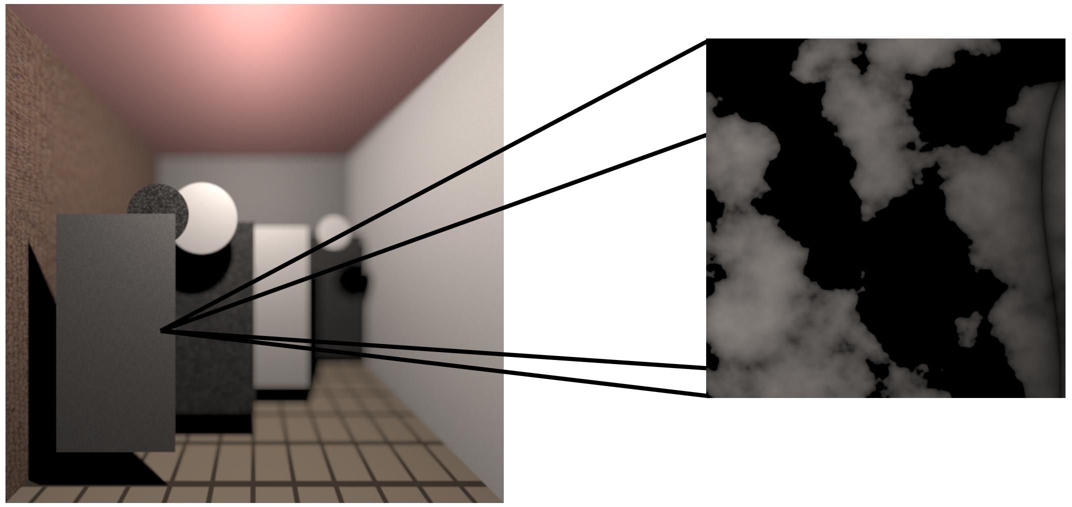

…2D function f(x).

Monte Carlo Integration

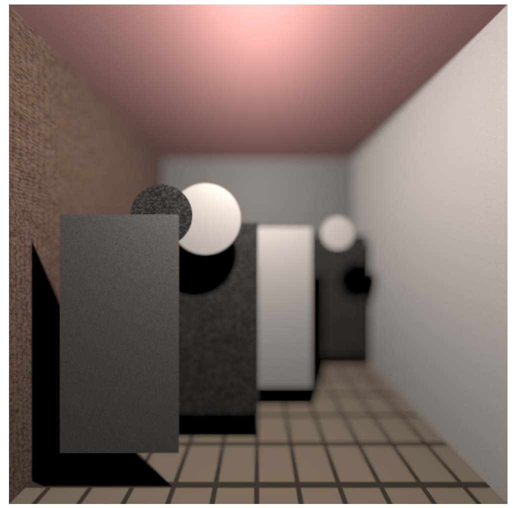

$f(\vec{x})$

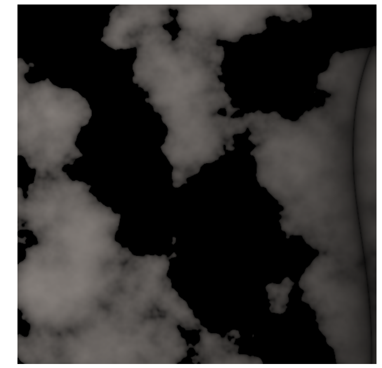

We are interested in computing [CLICK] the integral of this function. In practice, it is not always possible to compute this integral analytically. This is where we sought to Monte Carlo integration which involves point sampling…

Monte Carlo Integration

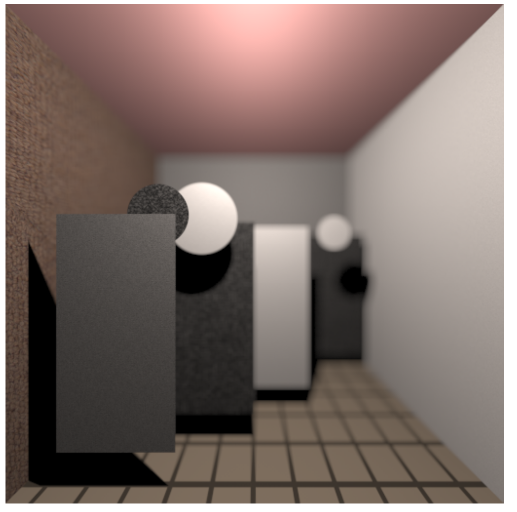

$f(\vec{x})$

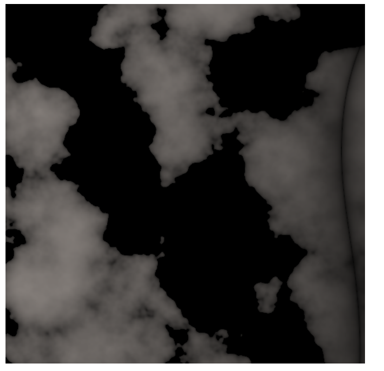

$$I = \int_0^1 f(\vec{x}) d\vec{x}$$

Monte Carlo Integration

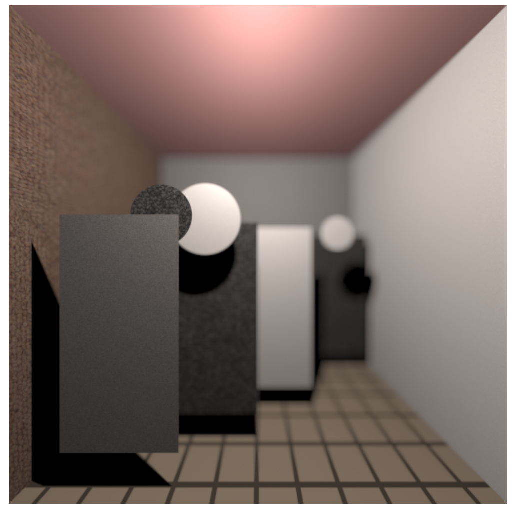

$f(\vec{x})$

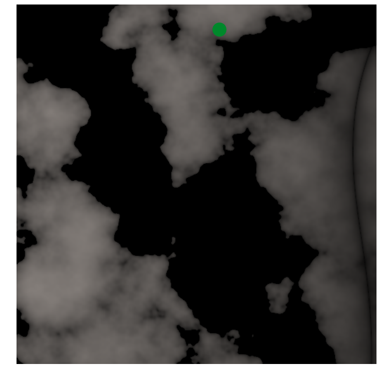

$$I = \int_0^1 f(\vec{x}) d\vec{x}$$

…the underlying function, as a result we can represent this integral as a Monte Carlo estimator…

Monte Carlo Integration

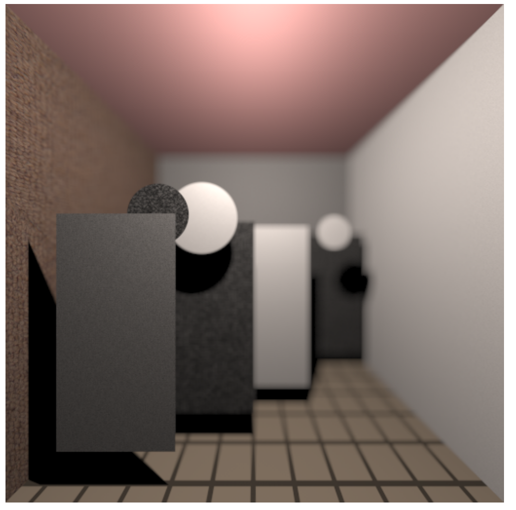

$f(\vec{x})$

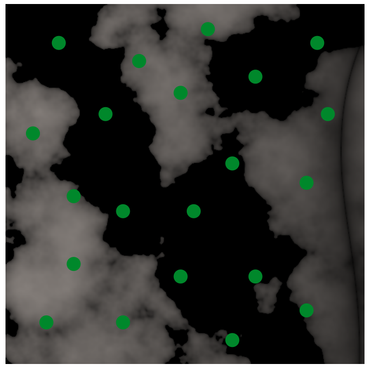

$$I = \int_0^1 f(\vec{x}) d\vec{x}$$

Monte Carlo Integration

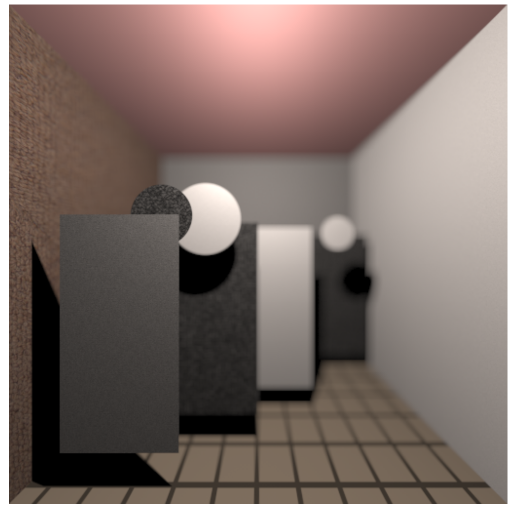

$f(\vec{x})$

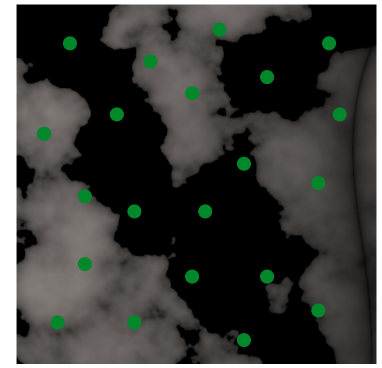

$$I_N = \frac{1}{N} \sum_{k=1}^{N} \frac{f(\vec{x}_k)}{p(\vec{x}_k)}$$

… where p [CLICK] is the probability density function used to distribute these point samples. This process is highly noise prone and if we look at an out of focus region…

Monte Carlo Integration

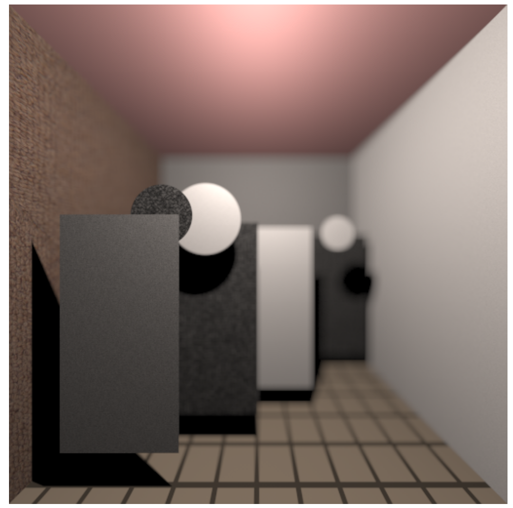

$f(\vec{x})$

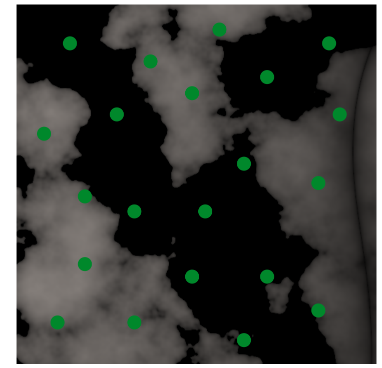

$$I_N = \frac{1}{N}\sum_{k=1}^{N}\frac{f(\vec{x}_k)}{\boxed{p(\vec{x}_k)}}$$

Variance

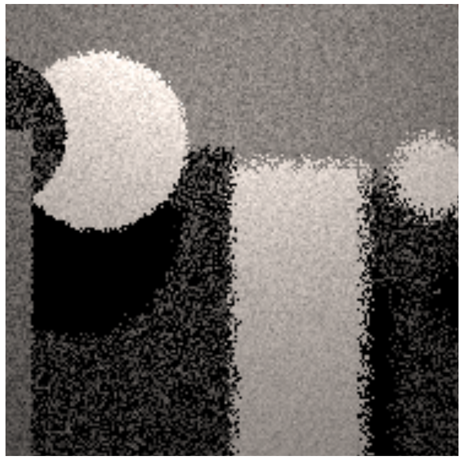

…which has an underlying 4D integral, it appears to be very noise. One way to reduce this noise is to increase…

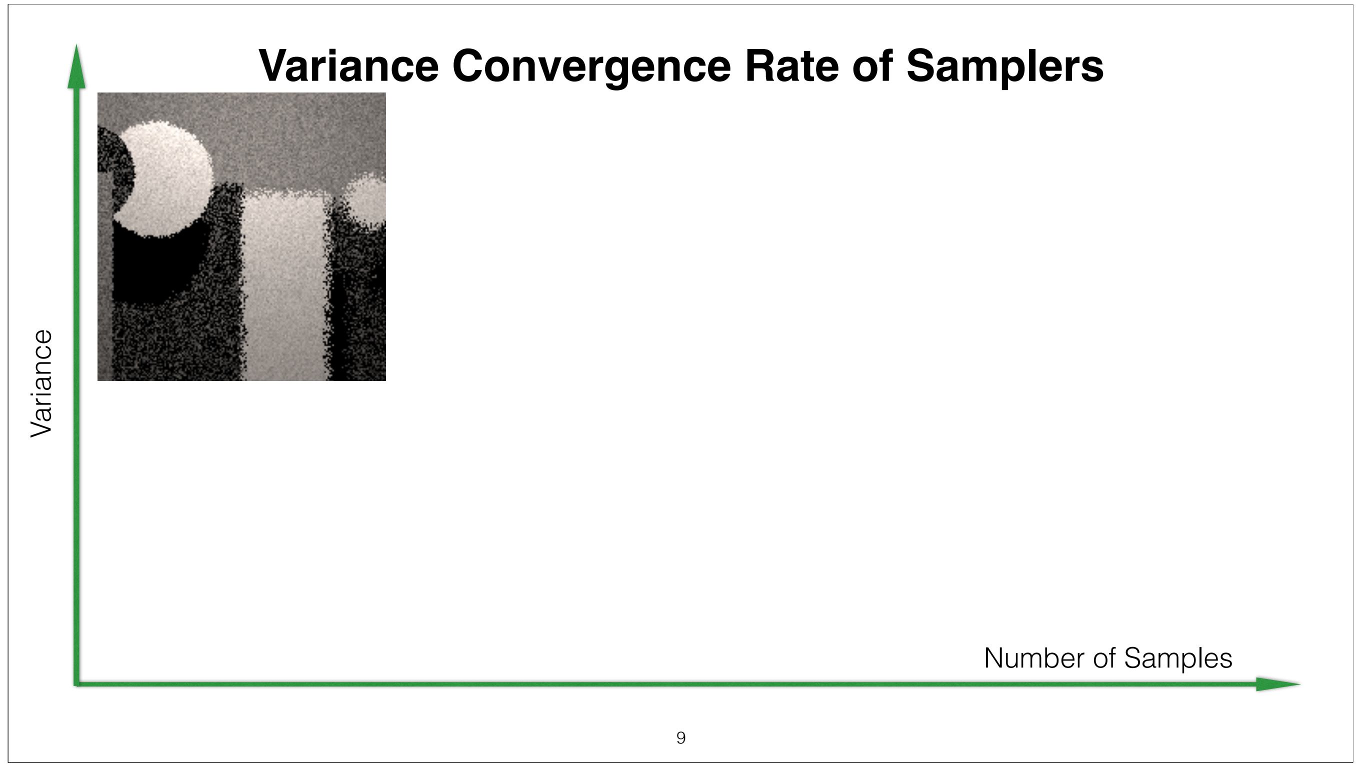

…increase the sample count till the image becomes noise-free (converge). The rate at which this image converges depend on the underlying sampling pattern used. For example, with 4D jittered samples, we would obtain a 4D convergence rate of O(N^-1.25) whereas with Poisson disk…

Variance Convergence Rate of Samplers

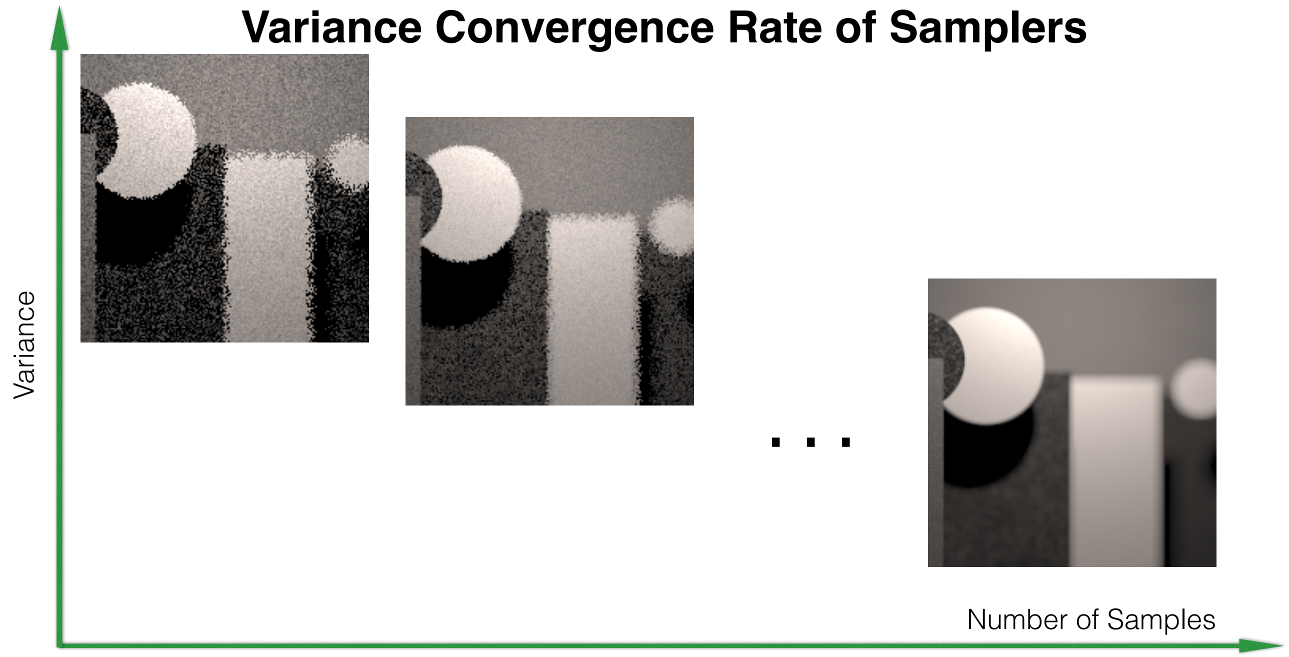

Variance Convergence Rate of Samplers

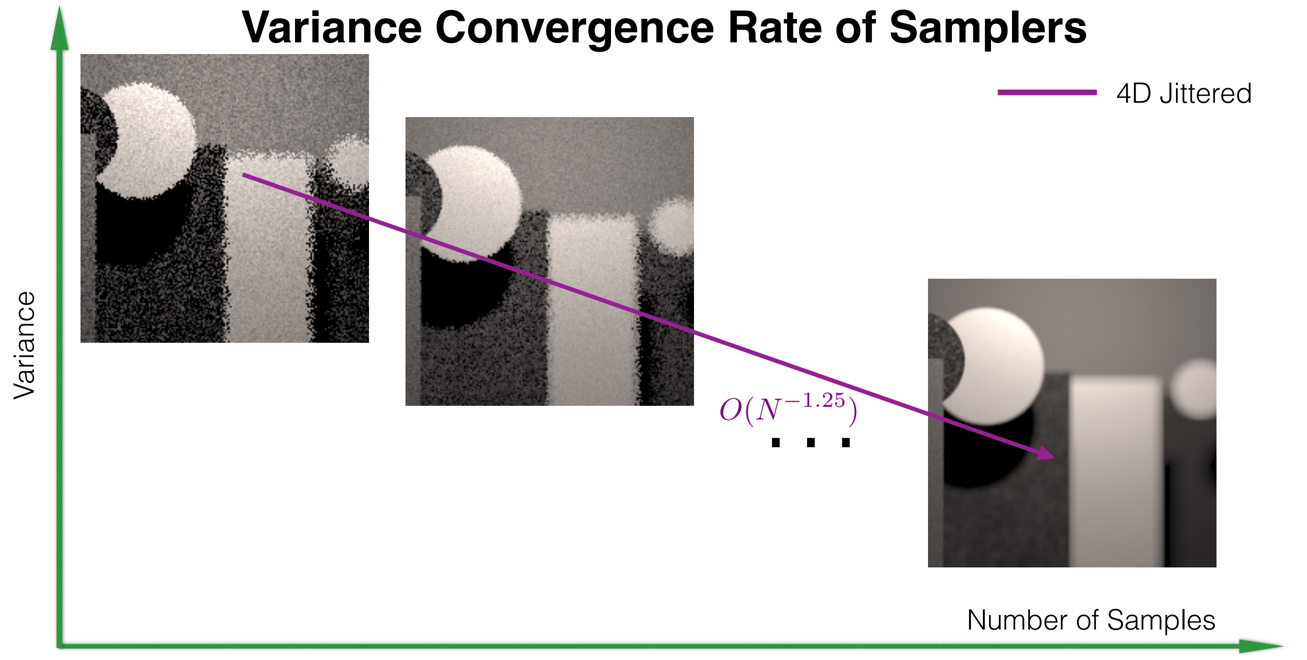

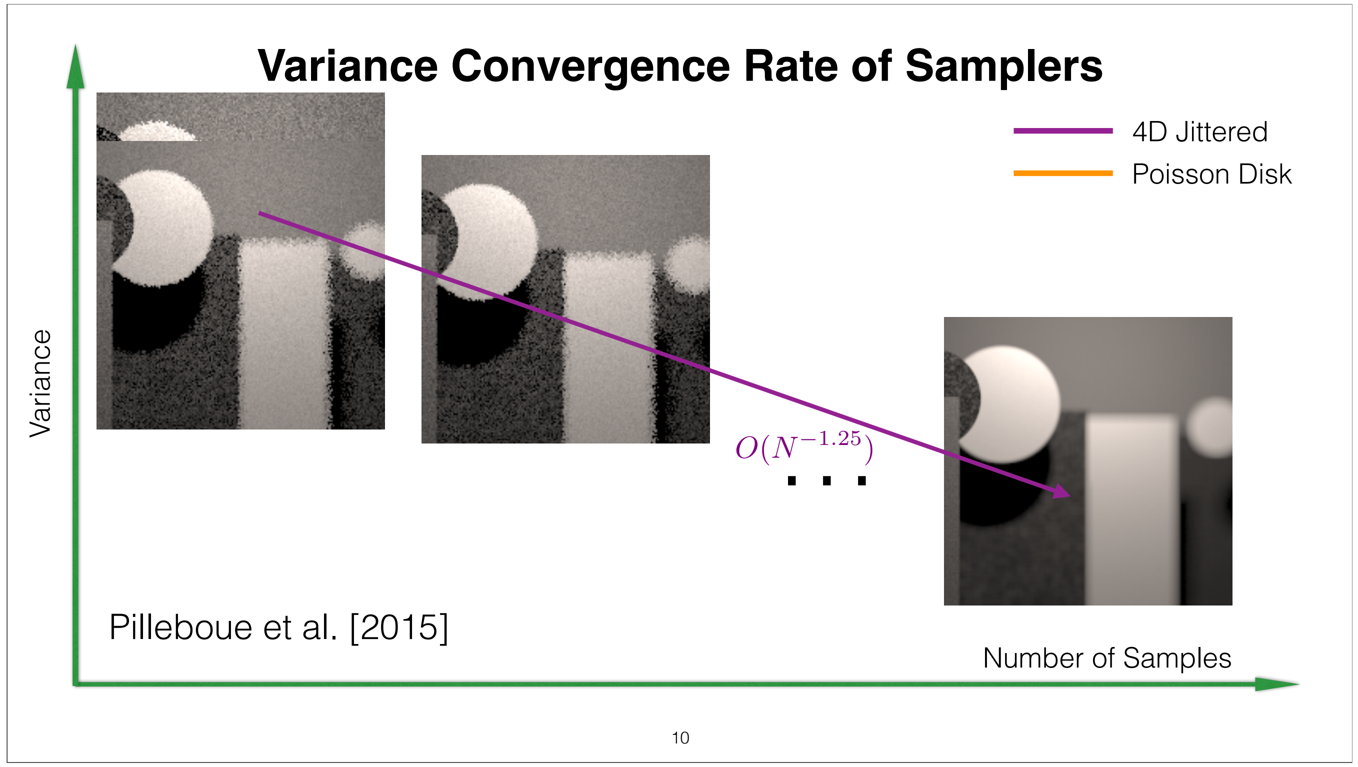

… we can obtain less noisy images at small sample count but [CLICK] as we increase the samples the convergence rate obtained is O(N^-1). These convergence rates can be empirically computed using the sample variance but recently Pilleboue and Colleagues in [2015] proposed a variance formulation in the Fourier domain, which allows to theoretically derive these convergence rates even for blue noise samples. Their work was developed following the work by [CLICK] Fredo Durand and Subr & Kautz.

We will briefly revisit this work first.

Variance Convergence Rate of Samplers

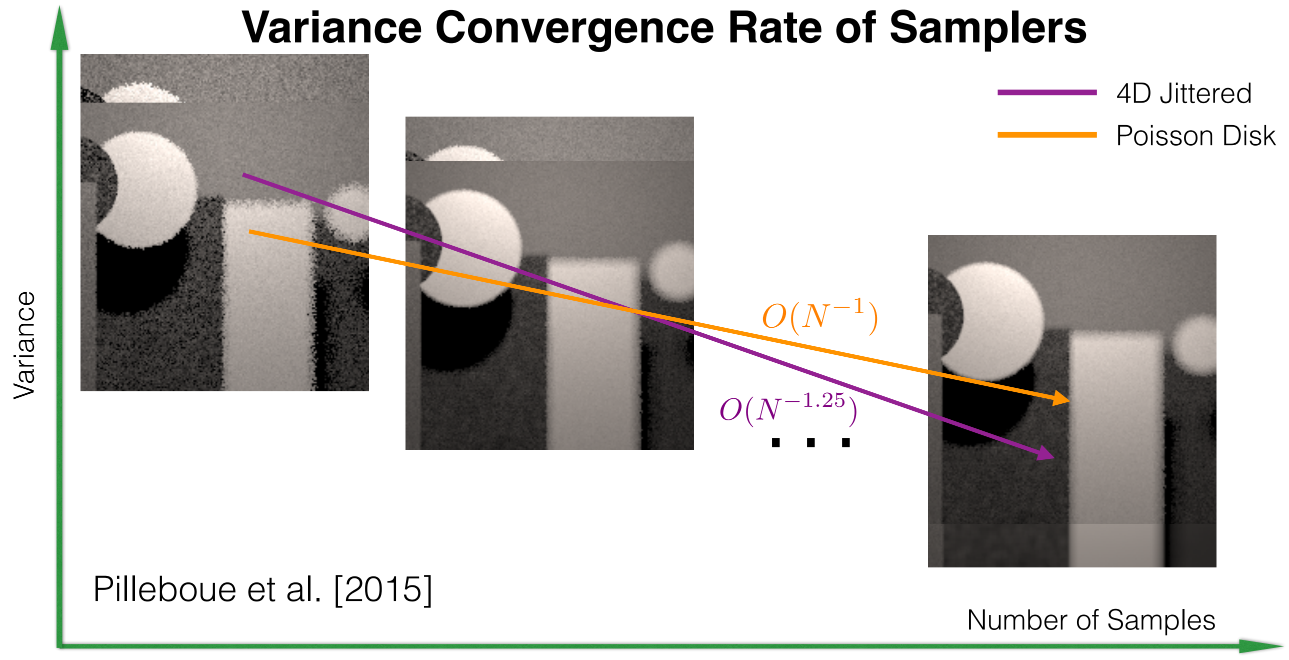

Variance Convergence Rate of Samplers

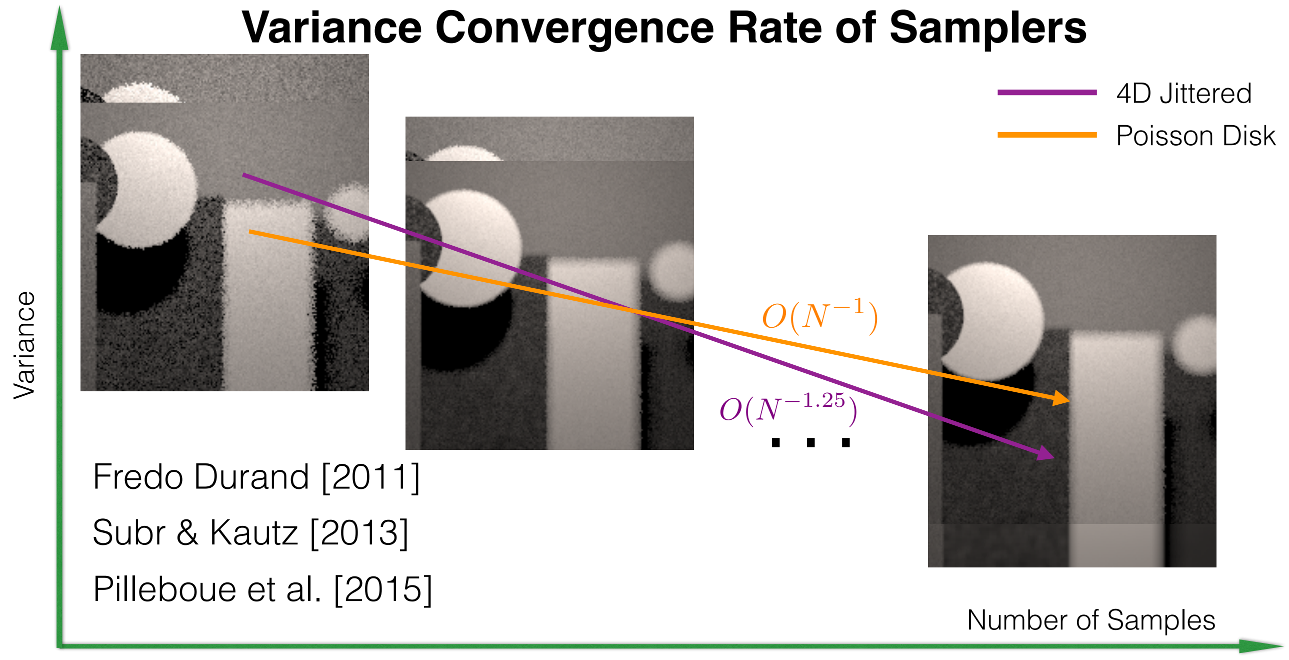

Monte Carlo Estimator

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

Fredo Durand [2011]

We start with a Monte Carlo estimator with a unit pdf, which can be written [CLICK] in continuous form using the [CLICK] dirac-delta functions. This integral can be rewritten in the [CLICK] following form where S(x) represents the sum of diracs [CLICK] as a sampling function. In this sampling function, x_k [CLICK] represents the sample locations shown as these green dots. To represent variance in the Fourier domain, we need the expected power spectrum of these samples. This is computed by first…

Monte Carlo Estimator

$$I_N = \frac{1}{N}\sum_{k=1}^{N} f(\vec{x}_k) = \int_0^1 \frac{1}{N}\sum_{k=1}^{N}\delta(\vec{x}-\vec{x}_k)\, f(\vec{x})\, d\vec{x}$$

Fredo Durand [2011]

Monte Carlo Estimator

$$I_N = \frac{1}{N}\sum_{k=1}^{N} f(\vec{x}_k) = \int_0^1 \boxed{\frac{1}{N}\sum_{k=1}^{N}\delta(\vec{x}-\vec{x}_k)}\, f(\vec{x})\, d\vec{x}$$

Fredo Durand [2011]

Monte Carlo Estimator

$$I_N = \frac{1}{N}\sum_{k=1}^{N} f(\vec{x}_k) = \int_0^1 \boxed{\frac{1}{N}\sum_{k=1}^{N}\delta(\vec{x}-\vec{x}_k)} f(\vec{x})\,d\vec{x} = \int_0^1 S_N(\vec{x})f(\vec{x})\,d\vec{x}$$

Fredo Durand [2011]

Monte Carlo Estimator

$$I_N = \frac{1}{N}\sum_{k=1}^{N} f(\vec{x}_k) = \int_0^1 \boxed{\frac{1}{N}\sum_{k=1}^{N}\delta(\vec{x}-\vec{x}_k)} f(\vec{x})\, d\vec{x} = \int_0^1 \boxed{S_N(\vec{x})} f(\vec{x})\, d\vec{x}$$

Fredo Durand [2011]

Monte Carlo Estimator

$$I_N = \frac{1}{N}\sum_{k=1}^{N} f(\vec{x}_k) = \int_0^1 \frac{1}{N}\sum_{k=1}^{N}\delta(\vec{x}-\vec{x}_k)\, f(\vec{x})\, d\vec{x} = \int_0^1 S_N(\vec{x})\, f(\vec{x})\, d\vec{x}$$

$$S_N(\vec{x}) = \frac{1}{N}\sum_{k=1}^{N}\delta(\vec{x}-\vec{x}_k)$$

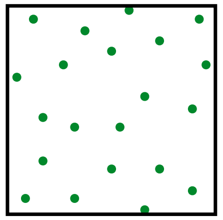

Fredo Durand [2011]

Monte Carlo Estimator

$$I_N = \frac{1}{N}\sum_{k=1}^{N} f(\vec{x}_k) = \int_0^1 \frac{1}{N}\sum_{k=1}^{N} \delta(\vec{x} - \vec{x}_k)\, f(\vec{x})\, d\vec{x} = \int_0^1 S_N(\vec{x})\, f(\vec{x})\, d\vec{x}$$

$$S_N(\vec{x}) = \frac{1}{N}\sum_{k=1}^{N} \delta(\vec{x} - \boxed{\vec{x}_k})$$

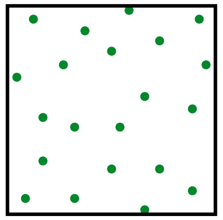

Fredo Durand [2011]

Samples Power Spectrum

$$I_N = \int_0^1 S_N(\vec{x}) f(\vec{x})\, d\vec{x}$$

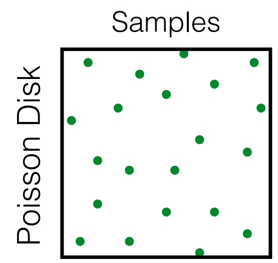

$$S_N(\vec{x}) = \frac{1}{N} \sum_{k=1}^{N} \delta(\vec{x} - \vec{x}_k)$$

… computing the Power spectrum of the samples, followed by…

Samples Power Spectrum

$$I_N = \int_0^1 S_N(\vec{x}) f(\vec{x})\, d\vec{x}$$

Samples

Poisson Disk

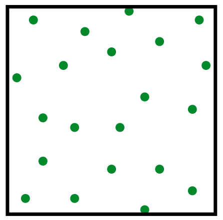

Spectrum

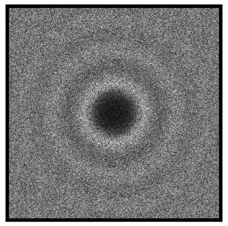

$$S_N(\vec{x}) = \frac{1}{N}\sum_{k=1}^{N} \delta(\vec{x} - \vec{x}_k)$$

$$\mathcal{P}_{S_N}(\nu) = \left| \frac{1}{N}\sum_{k=1}^{N} e^{-i2\pi\nu\cdot\vec{x}_k} \right|^2$$

Expected Sampling Power Spectra

$$I_N = \int_0^1 S_N(\vec{x}) f(\vec{x})\, d\vec{x}$$

Samples

Poisson Disk

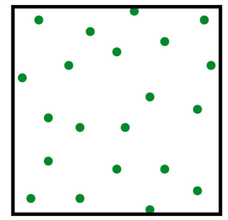

Spectrum

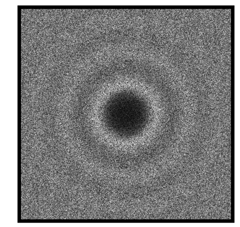

$$S_N(\vec{x}) = \frac{1}{N} \sum_{k=1}^{N} \delta(\vec{x} - \vec{x}_k)$$

…generating multiple realizations of these samples and their corresponding power spectra. We then take the expectation…

Expected Sampling Power Spectra

$$I_N = \int_0^1 S_N(\vec{x}) f(\vec{x})\, d\vec{x}$$

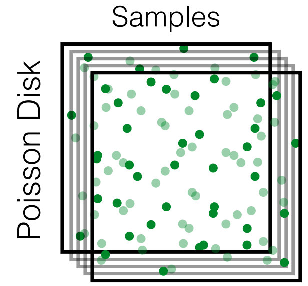

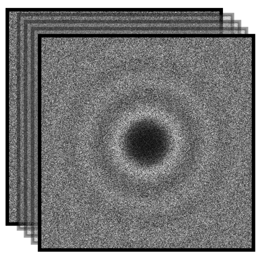

$$S_N(\vec{x}) = \frac{1}{N}\sum_{k=1}^{N} \delta(\vec{x} - \vec{x}_k)$$

$$\mathcal{P}_{S_N}(\nu) = \left| \frac{1}{N}\sum_{k=1}^{N} e^{-i2\pi\nu\cdot\vec{x}_k} \right|^2$$

Expected Sampling Power Spectra

$$I_N = \int_0^1 S_N(\vec{x}) f(\vec{x})\, d\vec{x}$$

Samples

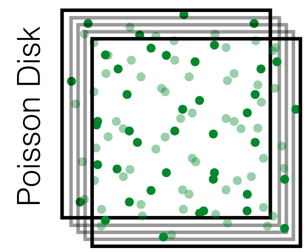

$$S_N(\vec{x}) = \frac{1}{N}\sum_{k=1}^{N} \delta(\vec{x} - \vec{x}_k)$$

Spectrum

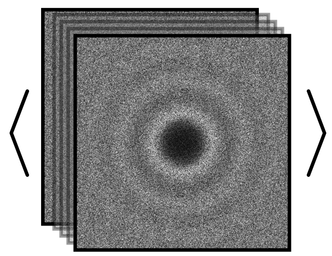

$$\left\langle \mathcal{P}_{S_N}(\nu) \right\rangle = \left\langle \left| \frac{1}{N} \sum_{k=1}^{N} e^{-i2\pi\nu\cdot\vec{x}_k} \right|^2 \right\rangle$$

…of these power spectra to obtain…

Expected Sampling Power Spectra

$$I_N = \int_0^1 S_N(\vec{x}) f(\vec{x})\, d\vec{x}$$

Samples

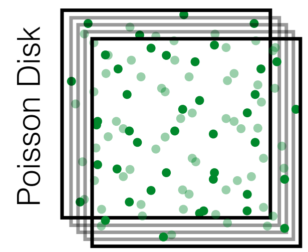

$$S_N(\vec{x}) = \frac{1}{N} \sum_{k=1}^{N} \delta(\vec{x} - \vec{x}_k)$$

Expected Spectrum

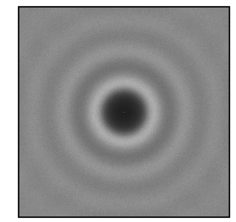

$$\langle \mathcal{P}_{S_N}(\nu) \rangle = \left\langle \left| \frac{1}{N} \sum_{k=1}^{N} e^{-i2\pi\nu\cdot\vec{x}_k} \right|^2 \right\rangle$$

… the expected spectrum. Given the sampling expected power spectrum,…

Variance of Monte Carlo Estimator

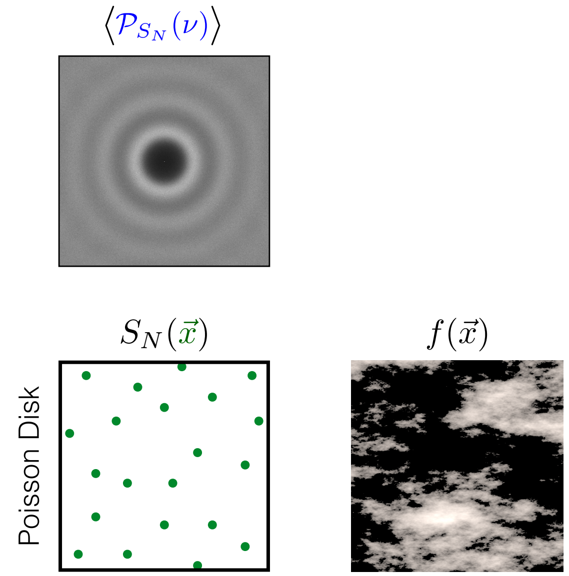

Fredo Durand [2011]

Subr & Kautz [2013]

16

…and the integrand f(x) that we are interested to evaluate, if we can compute the power spectrum [CLICK] of f(x), then we can obtain the variance by simply [CLICK] taking the integral over the product of these spectra. Pilleboue and colleagues [CLICK] extended this formulation to different samplers using homogenization. They rewrote this formulation…

Variance of Monte Carlo Estimator

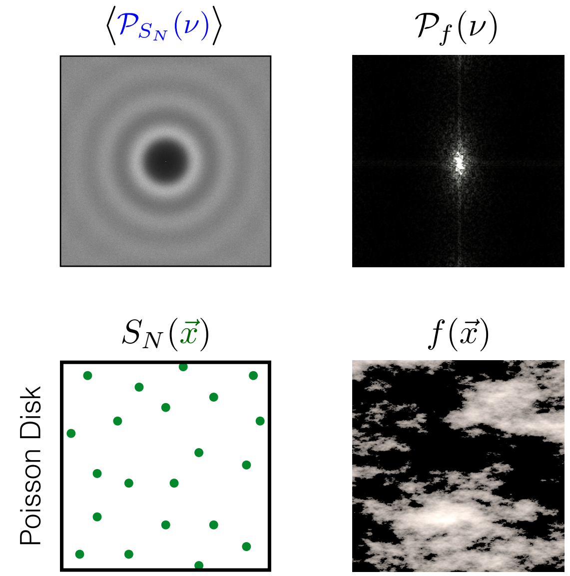

Fredo Durand [2011]

Subr & Kautz [2013]

Variance of Monte Carlo Estimator

$\left\langle \mathcal{P}_{S_N}(\nu) \right\rangle$ $\quad$ $\mathcal{P}_f(\nu)$

$\mathrm{Var}(I_N) = \int_\Omega$ 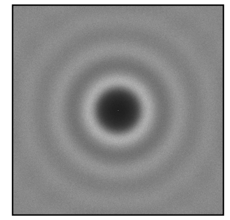 $\times$ 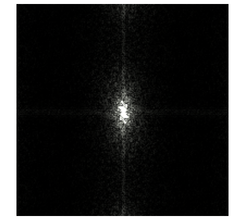 $d\nu$

Fredo Durand [2011]

Subr & Kautz [2013]

$S_N(\vec{x})$ $\quad$ $f(\vec{x})$

Poisson Disk 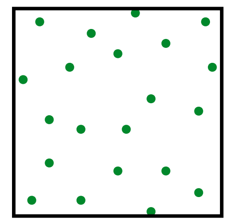 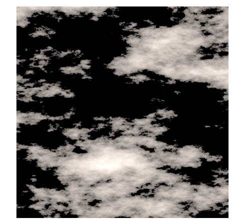

Variance of Monte Carlo Estimator

$\langle \mathcal{P}_{S_N}(\nu) \rangle$

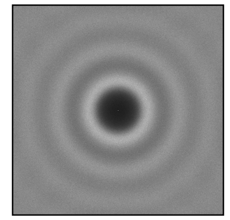

$\mathcal{P}_f(\nu)$

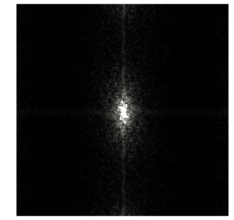

$$\mathrm{Var}(I_N) = \int_\Omega \langle \mathcal{P}_{S_N}(\nu) \rangle \times \mathcal{P}_f(\nu) \, d\nu$$

$S_N(\vec{x})$

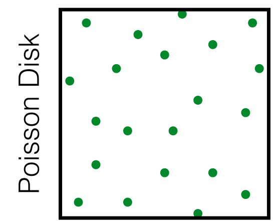

$f(\vec{x})$

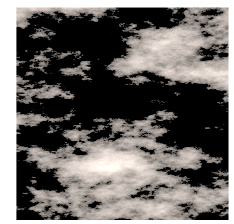

Fredo Durand [2011]

Subr & Kautz [2013]

Pilleboue et al. [2015]

Variance of Monte Carlo Estimator in Polar Coordinates

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1} \int_{\mathcal{S}^{d-1}} \langle \tilde{\mathcal{P}}_{S_N}(\rho\mathbf{n}) \rangle \times \mathcal{P}_f(\rho\mathbf{n}) \, d\mathbf{n} \, d\rho$$

$\langle \tilde{\mathcal{P}}_{S_N}(\rho\mathbf{n}) \rangle$

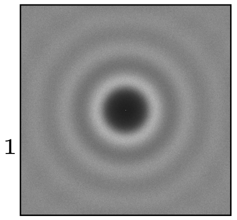

$\times$

$\mathcal{P}_f(\rho\mathbf{n})$

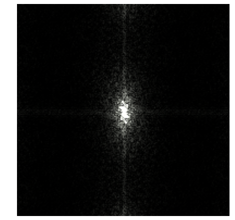

Pilleboue et al. [2015]

…in polar coordinates, as a double integral where the inner integral [CLICK] is over all the directions and the outer integral is over all the [CLICK] radial frequencies. They simplified this formulation further for isotropic sampling spectra [CLICK] which has the same energy in all directions for a given radial frequency. As a result, the sampling spectrum does not depend on directions, and we can safely take it out of the inner integral…

Variance of Monte Carlo Estimator in Polar Coordinates

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1} \boxed{\int_{\mathcal{S}^{d-1}} \langle \tilde{\mathcal{P}}_{S_N}(\rho\mathbf{n})\rangle \times \mathcal{P}_f(\rho\mathbf{n}) \, d\mathbf{n}}\, d\rho$$

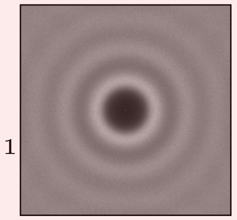

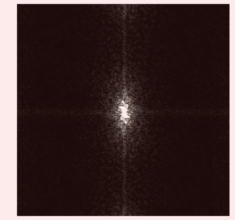

Pilleboue et al. [2015]

Variance of Monte Carlo Estimator in Polar Coordinates

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1} \int_{\mathcal{S}^{d-1}}$$

$\left\langle \tilde{\mathcal{P}}_{S_N}(\rho\mathbf{n}) \right\rangle$

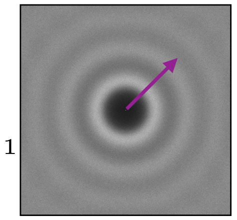

$\times$

$\mathcal{P}_f(\rho\mathbf{n})$

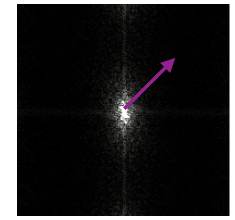

$$d\mathbf{n}\, d\rho$$

Pilleboue et al. [2015]

Variance of Monte Carlo Estimator in Polar Coordinates

$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1} \int_{\mathcal{S}^{d-1}}$

$\left\langle \tilde{\mathcal{P}}_{S_N}(\rho\mathbf{n}) \right\rangle$

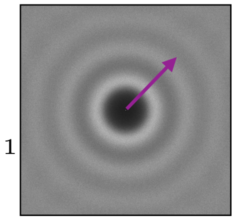

$\times$

$\mathcal{P}_f(\rho\mathbf{n})$

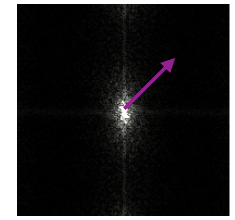

$d\mathbf{n}\, d\rho$

Pilleboue et al. [2015]

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$\left\langle \tilde{\mathcal{P}}_{S_N}(\rho\mathbf{n}) \right\rangle$ 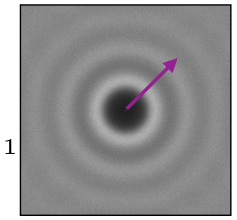 $\mathcal{P}_f(\rho\mathbf{n})$ 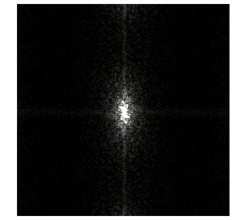

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1} \int_{\mathcal{S}^{d-1}} \left\langle \tilde{\mathcal{P}}_{S_N}(\rho\mathbf{n}) \right\rangle \times \mathcal{P}_f(\rho\mathbf{n}) \, d\mathbf{n} \, d\rho$$

Pilleboue et al. [2015]

…in polar coordinates, where the inner integral [CLICK] is over all the directions and the outer integral is over all the radial frequencies. They simplified this formulation further for [CLICK] isotropic sampling spectra [CLICK] which has the same energy in all directions for a given radial frequency,. As a result, the sampling spectrum does not depend on directions, and we can safely take it out of the inner integral…

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$\tilde{\mathcal{P}}_{\mathcal{S}_N}(\rho)$ $\quad$ $\mathcal{P}_f(\rho\mathbf{n})$

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1}$$

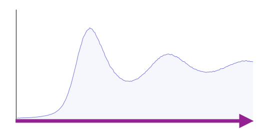

$$\int_{\mathcal{S}^{d-1}}$$

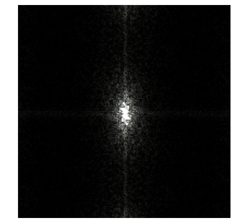

$$d\mathbf{n}\, d\rho$$

Pilleboue et al. [2015]

…for any one direction. Now, The integral [CLICK] over the integrand spectrum can be rewritten…

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1}$$

$\tilde{\mathcal{P}}_{\mathcal{S}_N}(\rho)$

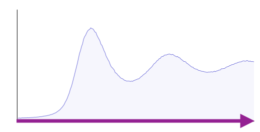

$$\int_{\mathcal{S}^{d-1}}$$

$\mathcal{P}_f(\rho\mathbf{n})$

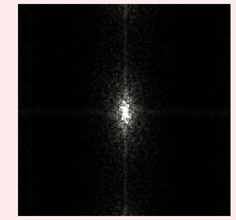

$$d\mathbf{n}\, d\rho$$

Pilleboue et al. [2015]

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$\tilde{\mathcal{P}}_{S_N}(\rho)$ $\qquad$ $\mathcal{P}_f(\rho\mathbf{n})$

$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1}$ 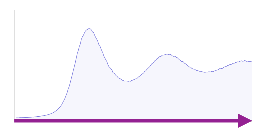 $\times$ 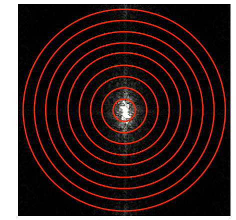 $d\rho$

Pilleboue et al. [2015]

… as a radial average over all the directions, which results in a…

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1}\, \tilde{\mathcal{P}}_{S_N}(\rho) \times \mathcal{P}_f(\rho)\, d\rho$$

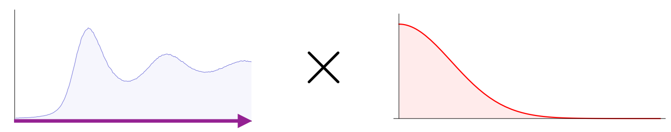

Pilleboue et al. [2015]

…1D radial function. This allowed Pilleboue and colleagues to represent variance as 1D integral just over the radial frequencies. They used this formulation to derive convergence rates…

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$\tilde{\mathcal{P}}_{S_N}(\rho)$ $\quad\quad$ $\mathcal{P}_f(\rho)$

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1}$$

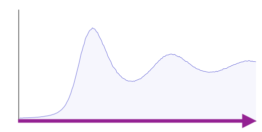

$$\times$$

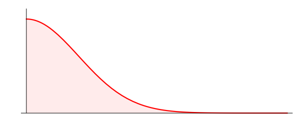

$$d\rho$$

Samplers	Worst Case	Best Case
Random		
Poisson Disk		
CCVT		

Pilleboue et al. [2015]

…of different samplers [CLICK], including random, poisson disk and CCVT. This was done under the assumption that the sampling spectra are [CLICK] isotropic. In practice, however, the samplers that we use are highly anisotropic in nature. For example,..

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$\tilde{\mathcal{P}}_{S_N}(\rho)$ $\mathcal{P}_f(\rho)$

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1}$$

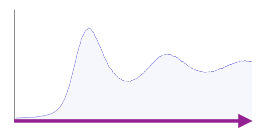

$\times$

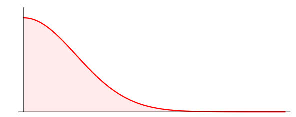

$d\rho$

Samplers	Worst Case	Best Case
Random	$\mathcal{O}(N^{-1})$	$\mathcal{O}(N^{-1})$
Poisson Disk	$\mathcal{O}(N^{-1})$	$\mathcal{O}(N^{-1})$
CCVT	$\mathcal{O}(N^{-1.5})$	$\mathcal{O}(N^{-3})$

Pilleboue et al. [2015]

Variance of Monte Carlo Estimator for Isotropic Sampling Spectra

$\tilde{\mathcal{P}}_{S_N}(\rho)$ $\mathcal{P}_f(\rho)$

$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1}$ 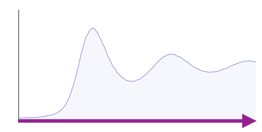 $\times$ 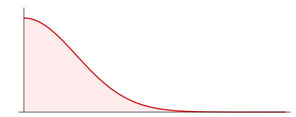 $d\rho$

Samplers	Worst Case	Best Case
Random	$\mathcal{O}(N^{-1})$	$\mathcal{O}(N^{-1})$
Poisson Disk	$\mathcal{O}(N^{-1})$	$\mathcal{O}(N^{-1})$
CCVT	$\mathcal{O}(N^{-1.5})$	$\mathcal{O}(N^{-3})$

Isotropic Spectrum
Poisson Disk

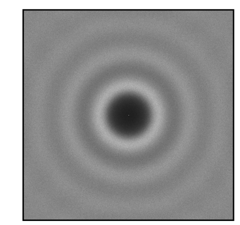

Pilleboue et al. [2015]

Latin Hypercube Sampler (N-rooks)

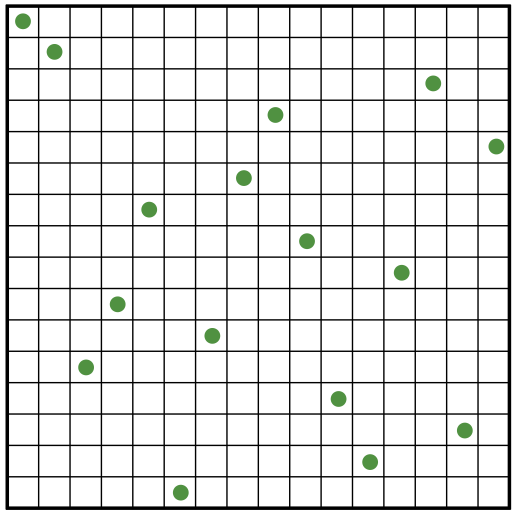

…a Latin hypercube sampler, that uses well 1D stratified samples, which are then randomly permuted to form [CLICK] 2D samples, has an…

Latin Hypercube Sampler (N-rooks)

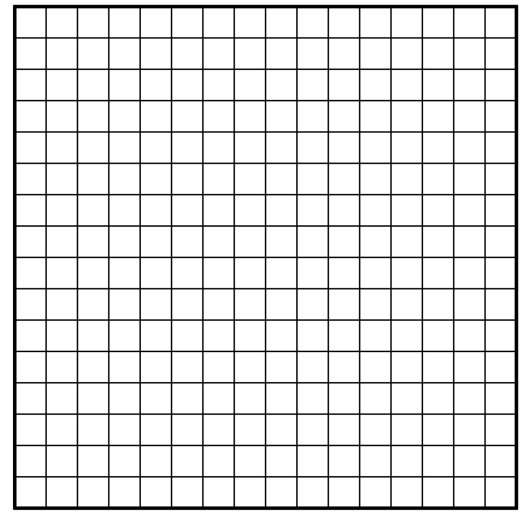

…a Latin hypercube sampler, that uses well 1D stratified samples, which are then randomly permuted to form [CLICK] 2D samples, has an…

Latin Hypercube Sampler (N-rooks)

Initialize

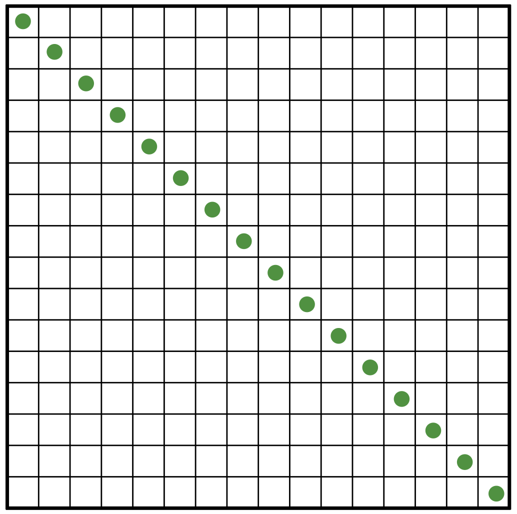

Latin Hypercube Sampler (N-rooks)

Shuffle rows

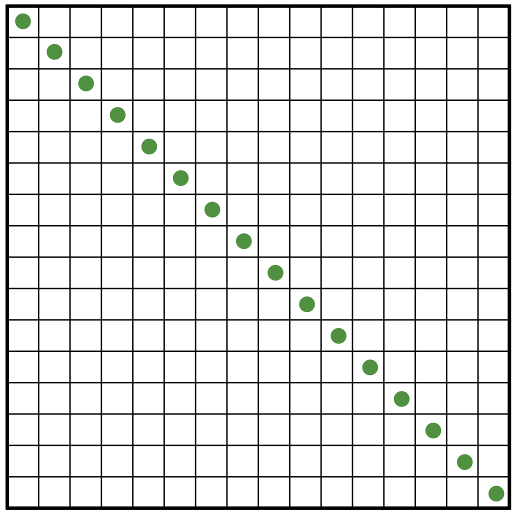

25

…a Latin hypercube sampler, that uses well 1D stratified samples, which are then randomly permuted to form [CLICK] 2D samples, has an…

Latin Hypercube Sampler (N-rooks)

Shuffle rows

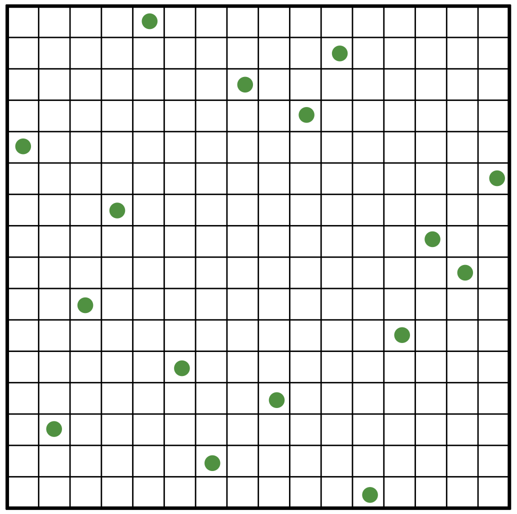

Latin Hypercube Sampler (N-rooks)

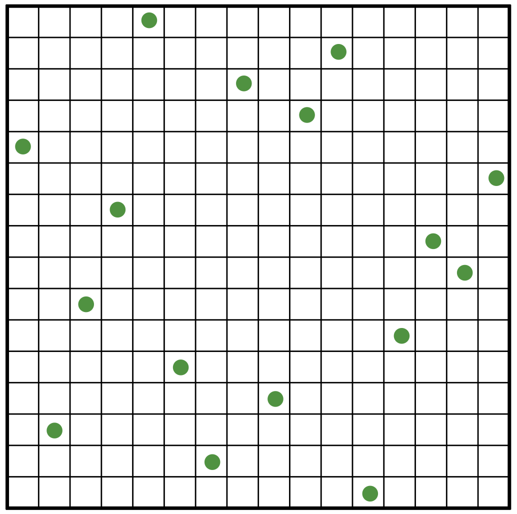

26

…a Latin hypercube sampler, that uses well 1D stratified samples, which are then randomly permuted to form [CLICK] 2D samples, has an…

Latin Hypercube Sampler (N-rooks)

Shuffle columns

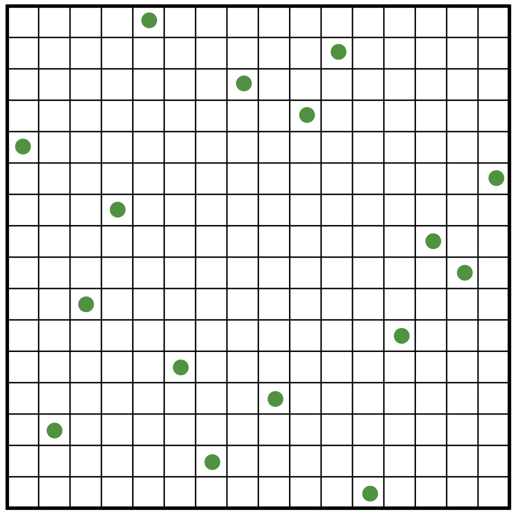

Latin Hypercube Sampler (N-rooks)

Shuffle columns

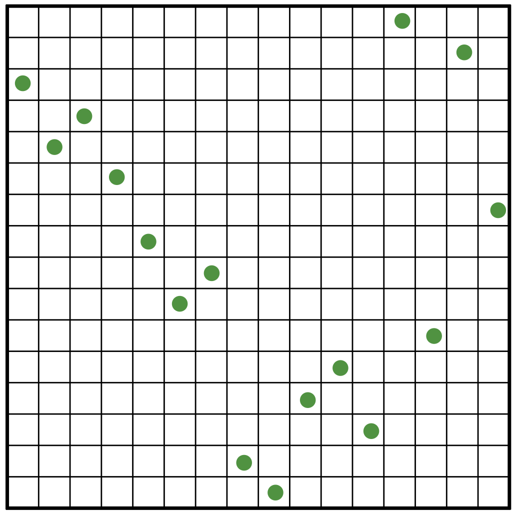

Latin Hypercube Sampler (N-rooks)

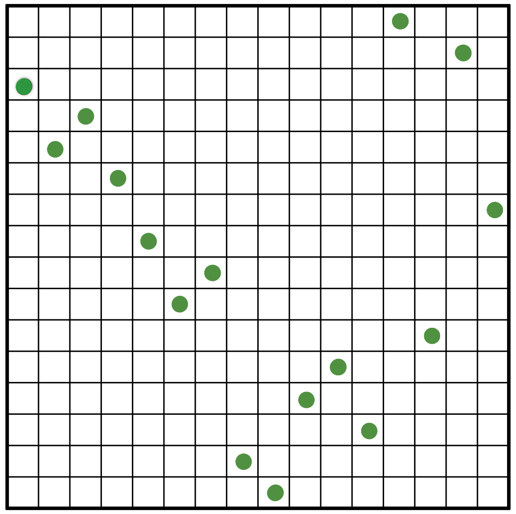

Latin Hypercube Sampler (N-rooks)

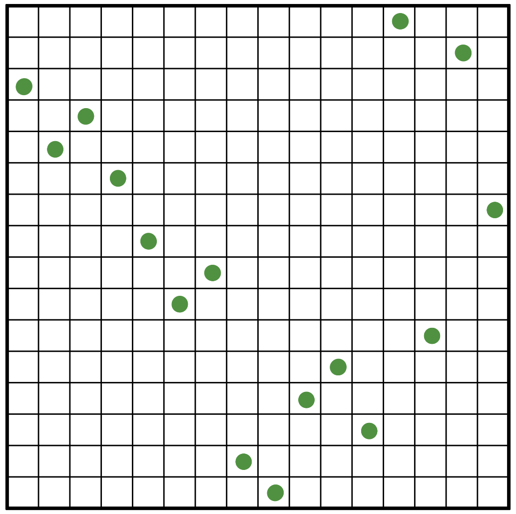

Anisotropic Sampling Power Spectra

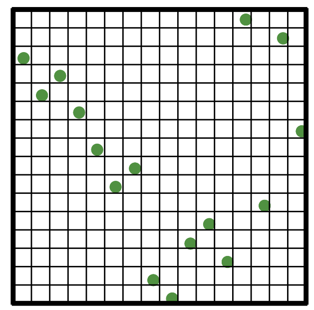

N-rooks /
Latin Hypercube

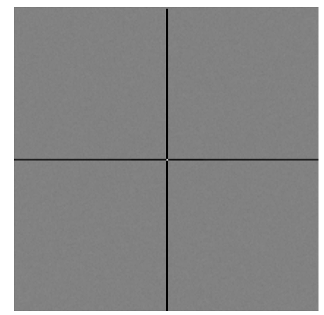

N-rooks
Spectrum

… anisotropic power spectrum with hairline structures visible as a dark cross in the middle. These hairline anisotropies are there due to the denser stratification along the X…

Anisotropic Sampling Power Spectra

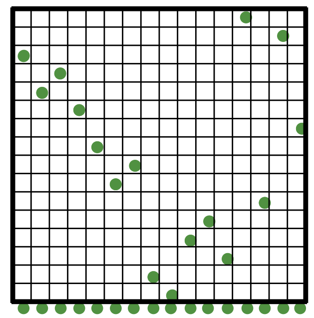

N-rooks /
Latin Hypercube

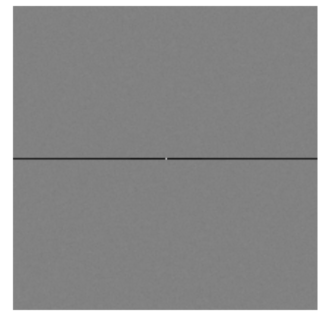

Spectrum

… anisotropic power spectrum with a dark cross in the middle. These hairline anisotropies are there due to the denser stratification along the X…

Anisotropic Sampling Power Spectra

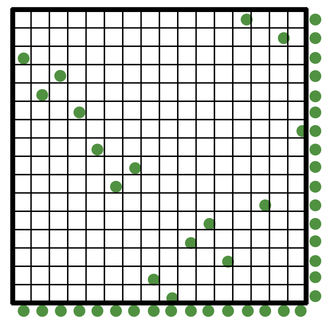

N-rooks /
Latin Hypercube

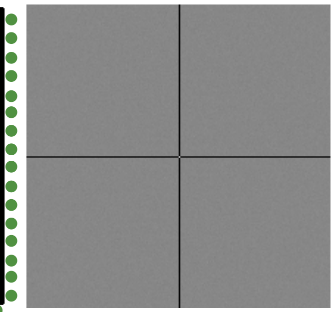

N-rooks
Spectrum

…and the Y-axis. It is also possible to directly obtain good 2D stratified samples…

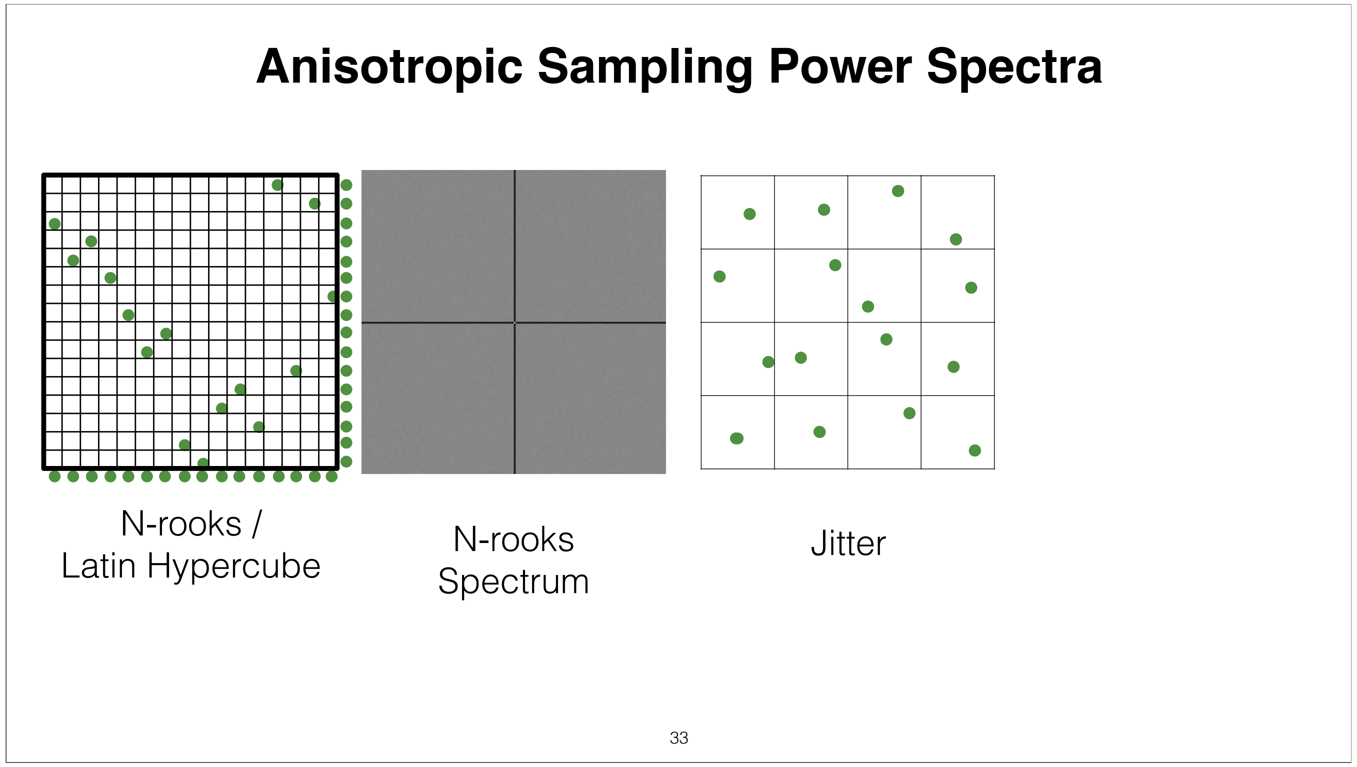

…which has a power spectrum [CLICK] with a dark region around the center. Chiu and colleagues, optimized these samples…

Anisotropic Sampling Power Spectra

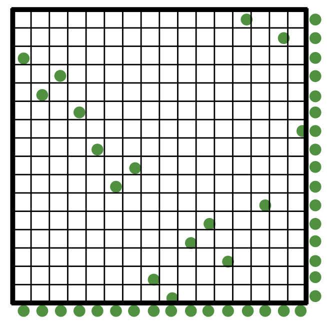

N-rooks /
Latin Hypercube

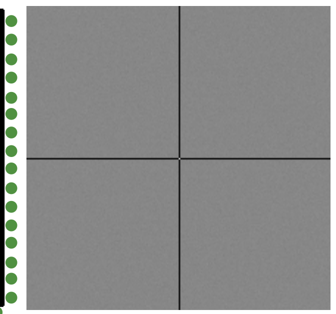

N-rooks
Spectrum

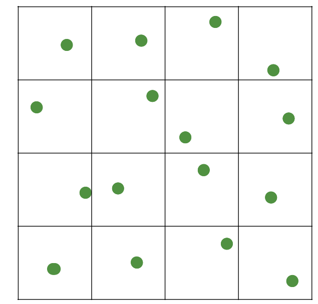

Jitter

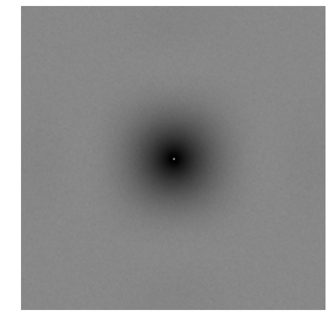

Jitter
Spectrum

Anisotropic Sampling Power Spectra

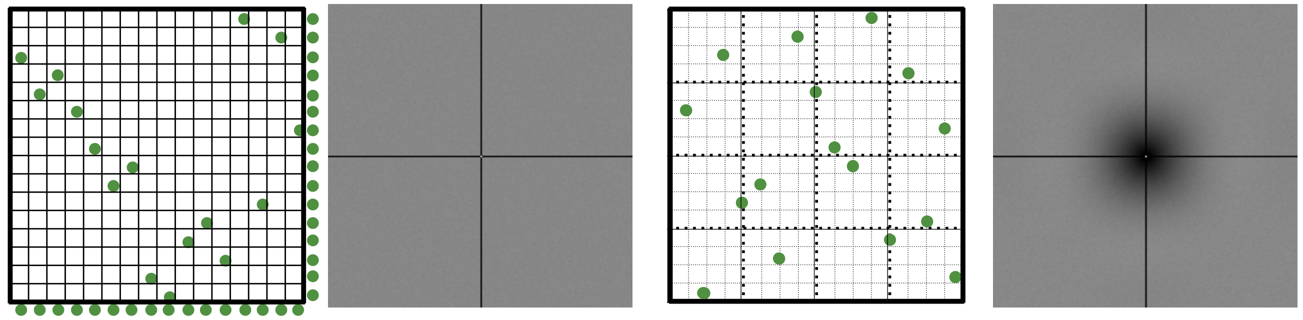

N-rooks / Latin Hypercube

N-rooks Spectrum

Multi-Jitter

Multi-Jitter Spectrum

Chiu et al. [1993]

…to obtain denser stratification…

Anisotropic Sampling Power Spectra

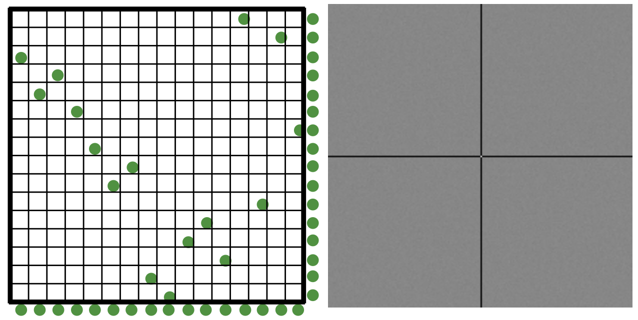

N-rooks /
Latin Hypercube

N-rooks
Spectrum

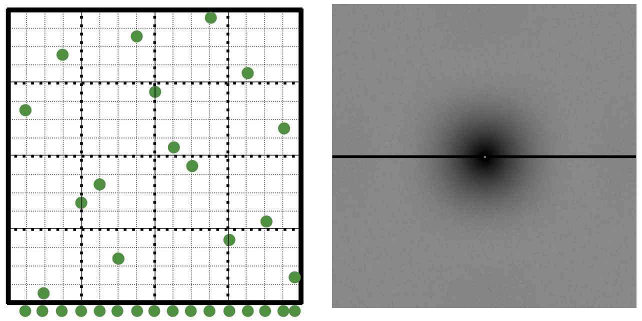

Multi-jitter

Multi-Jitter
Spectrum

Chiu et al. [1993]

…along the horizontal…

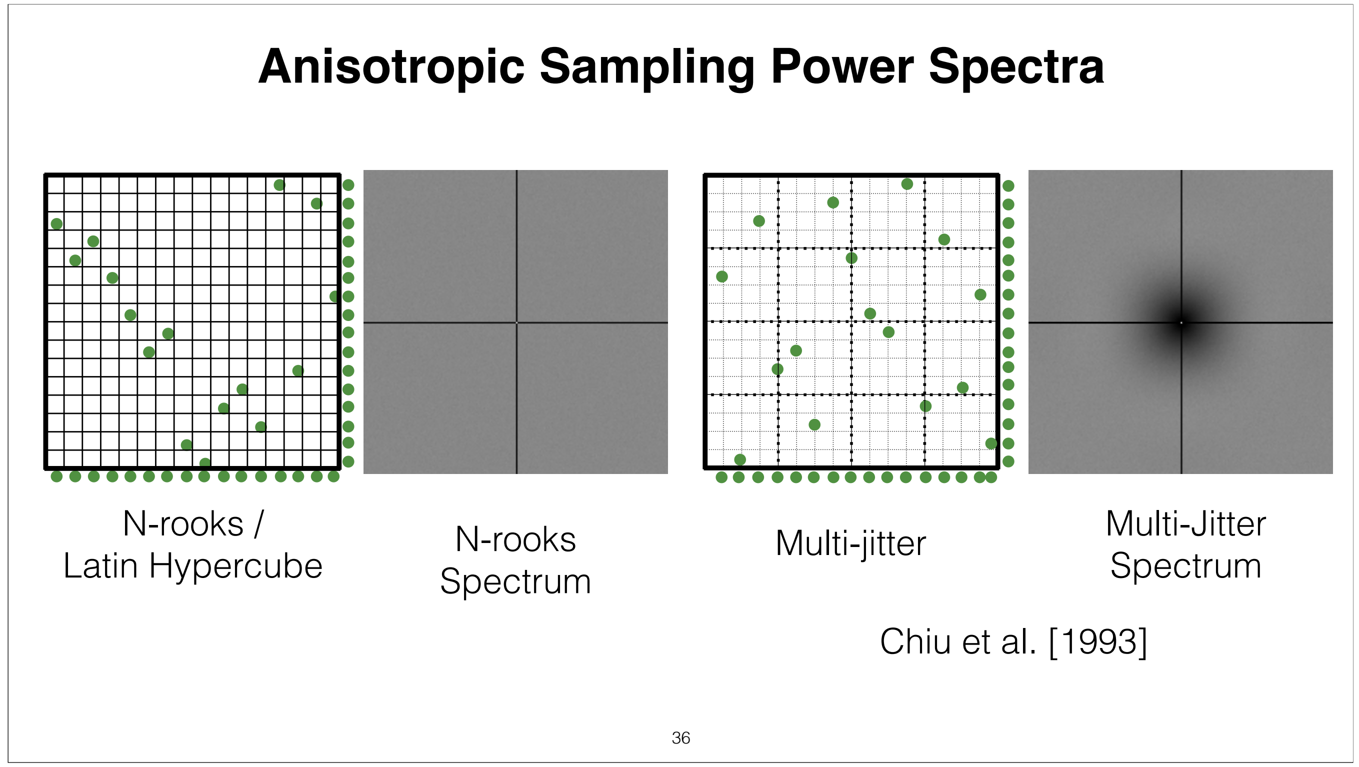

…and vertical axis, on top of 2D stratification, which results in multi-jittered samples with a hairline anisotropy along the canonical axes that is visible as a cross in the middle of it’s spectrum. The same ideas extend to…

Sampling in Higher Dimensions

…to higher dimensions. For example, in 4D…

4D Sampling

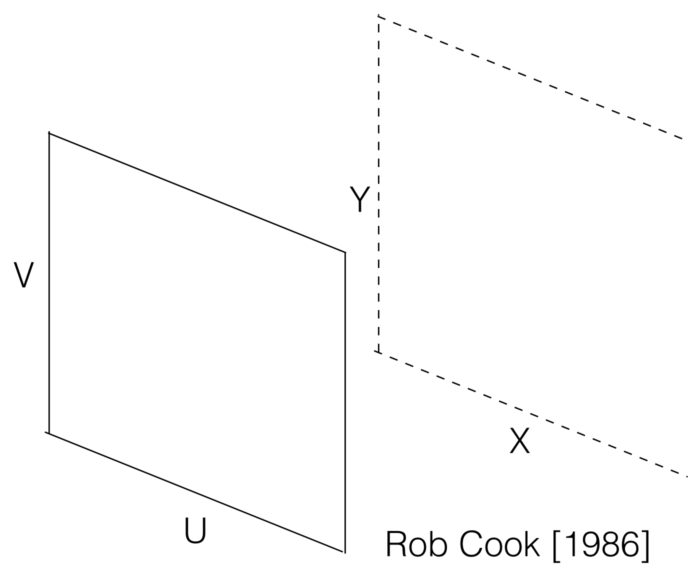

Rob Cook [1986]

…instead of directly sampling the full 4D space, Rob Cook in [1986] proposed to sample [CLICK] the lower 2D subspaces first, UV and XY here, and then randomly permute these 2D samples to form [CLICK] 4D tuples, which can then be used to evaluate an underlying 4D integrand. In practice…

4D Sampling

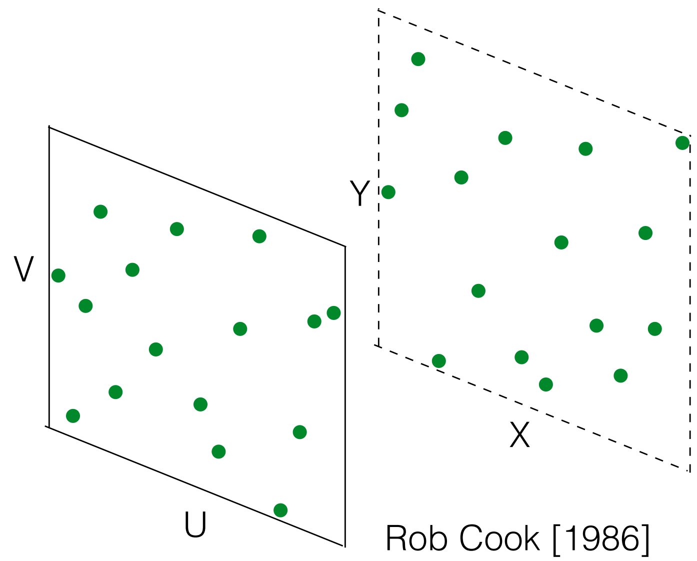

2D	2D
(x_1, y_1)	(u_1, v_1)
(x_2, y_2)	(u_2, v_2)
(x_3, y_3)	(u_3, v_3)
(x_4, y_4)	(u_4, v_4)
$\vdots$	$\vdots$

Rob Cook [1986]

4D Sampling

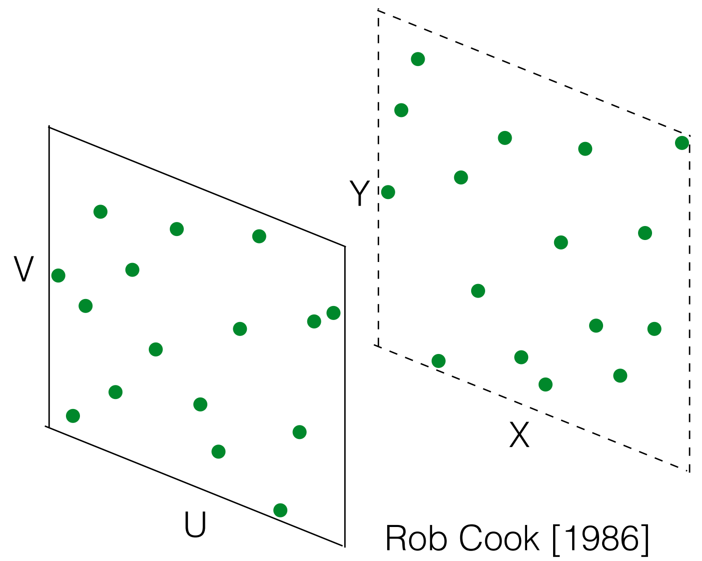

Rob Cook [1986]

2D	2D
(x_1, y_1)	(u_1, v_1)
(x_2, y_2)	(u_2, v_2)
(x_3, y_3)	(u_3, v_3)
(x_4, y_4)	(u_4, v_4)
$\vdots$	$\vdots$

4D

$$(x_1, y_1, u_3, v_3)$$
$$(x_2, y_2, u_1, v_1)$$
$$(x_3, y_3, u_4, v_4)$$
$$(x_4, y_4, u_2, v_2)$$
$$\vdots$$

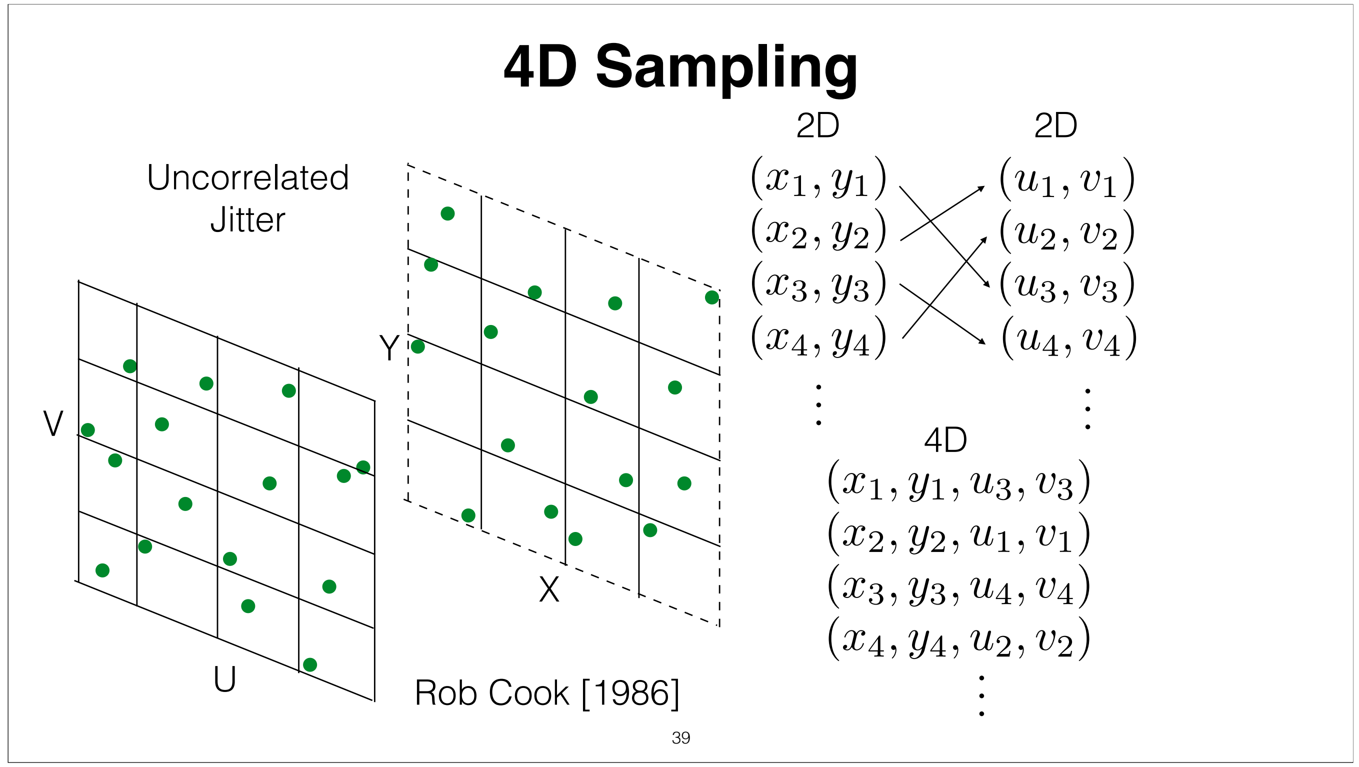

…rendering systems tend to use jittered samples on these 2D subspaces. It is, however, beneficial to use…

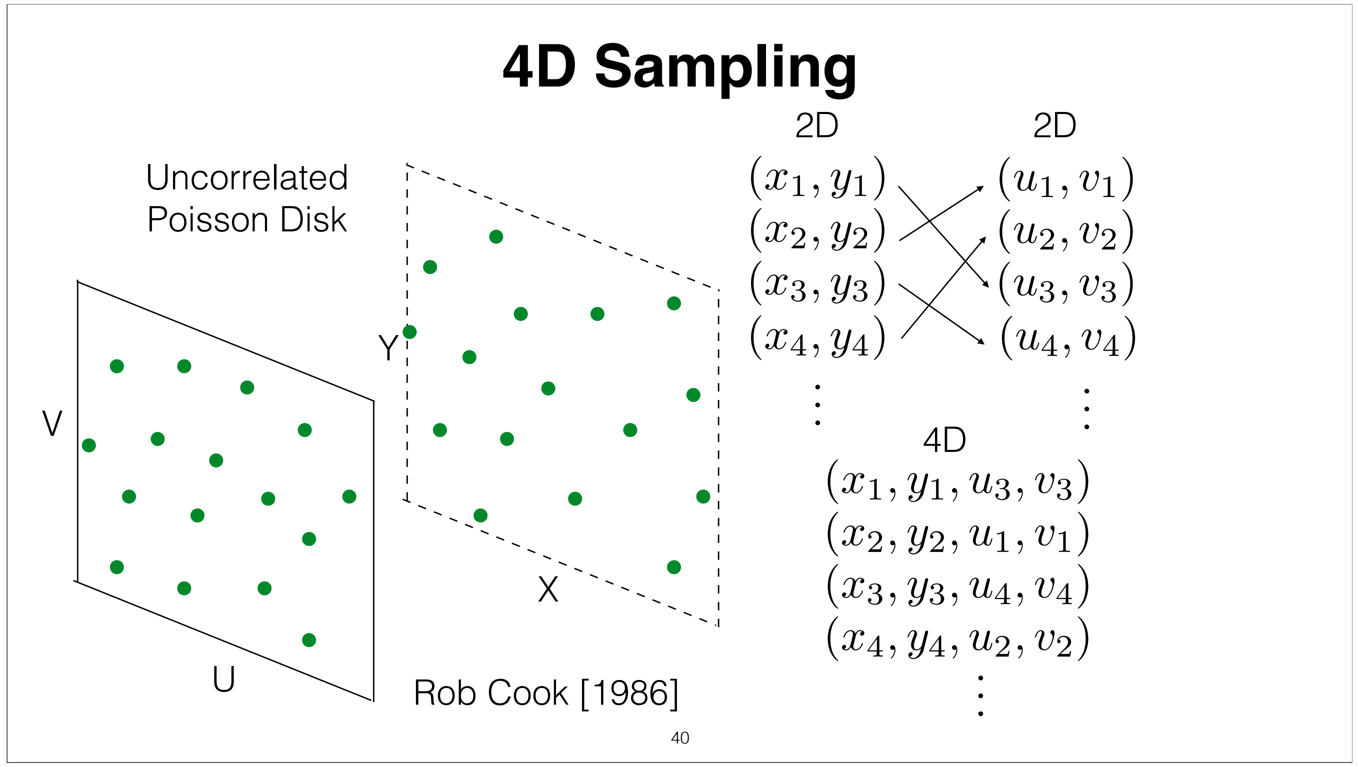

…rendering systems tend to use jittered samples on these 2D subspaces. It is, however, beneficial to use…

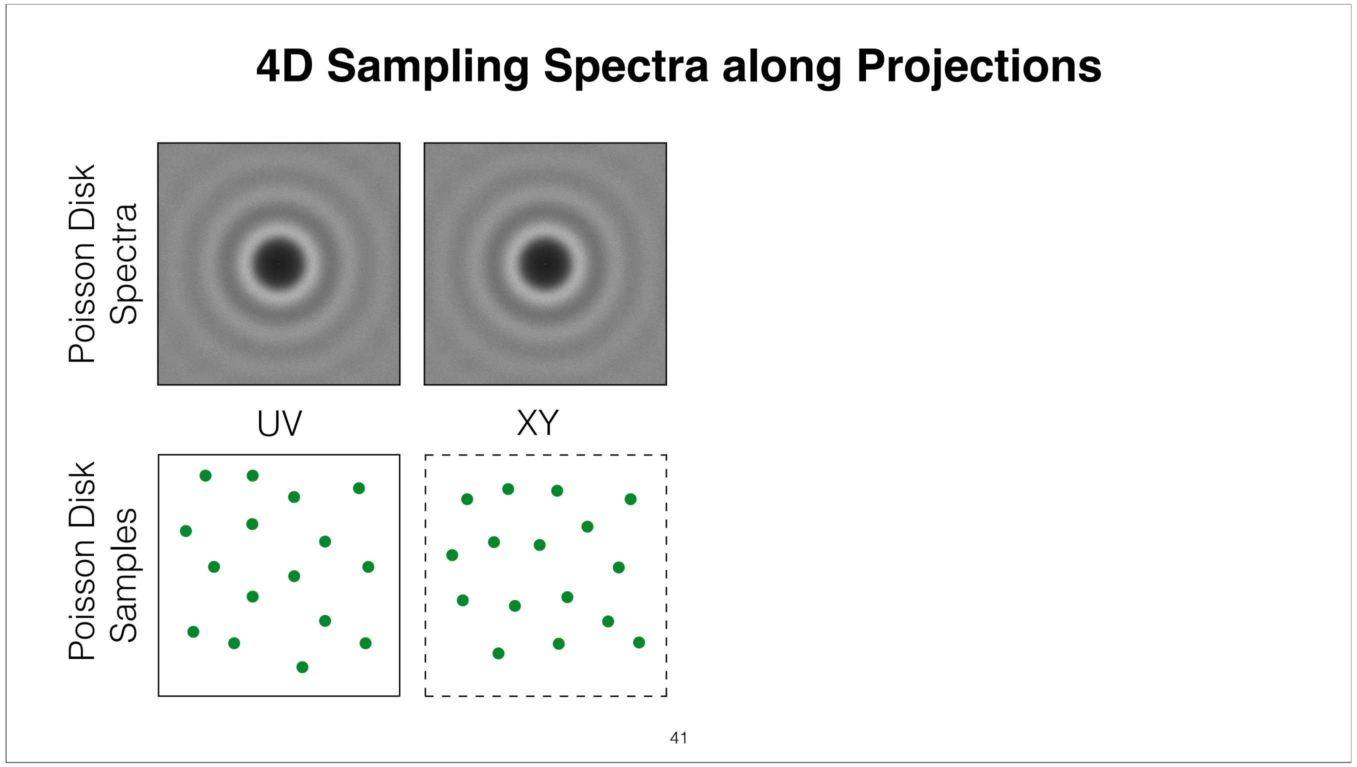

…power spectrum in these 2D projections, it has hairline anisotropy along the [CLICK] canonical axes on top of the [CLICK] big dark region that corresponds to 2D jittered sampling. However, If we look at the XU projection…

4D Sampling Spectra along Projections

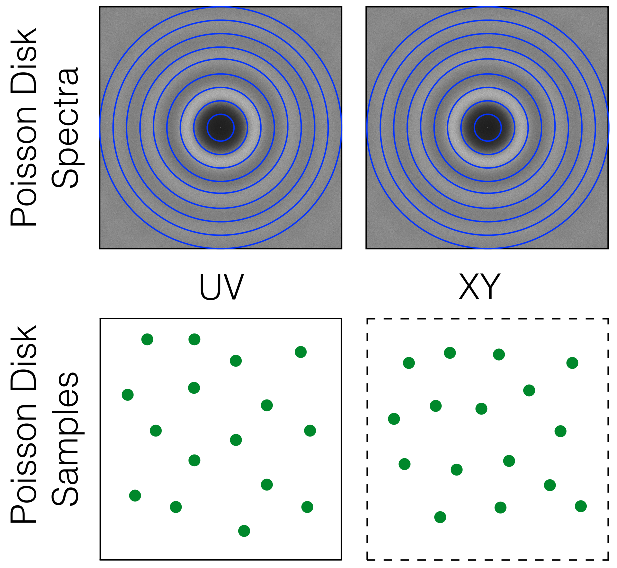

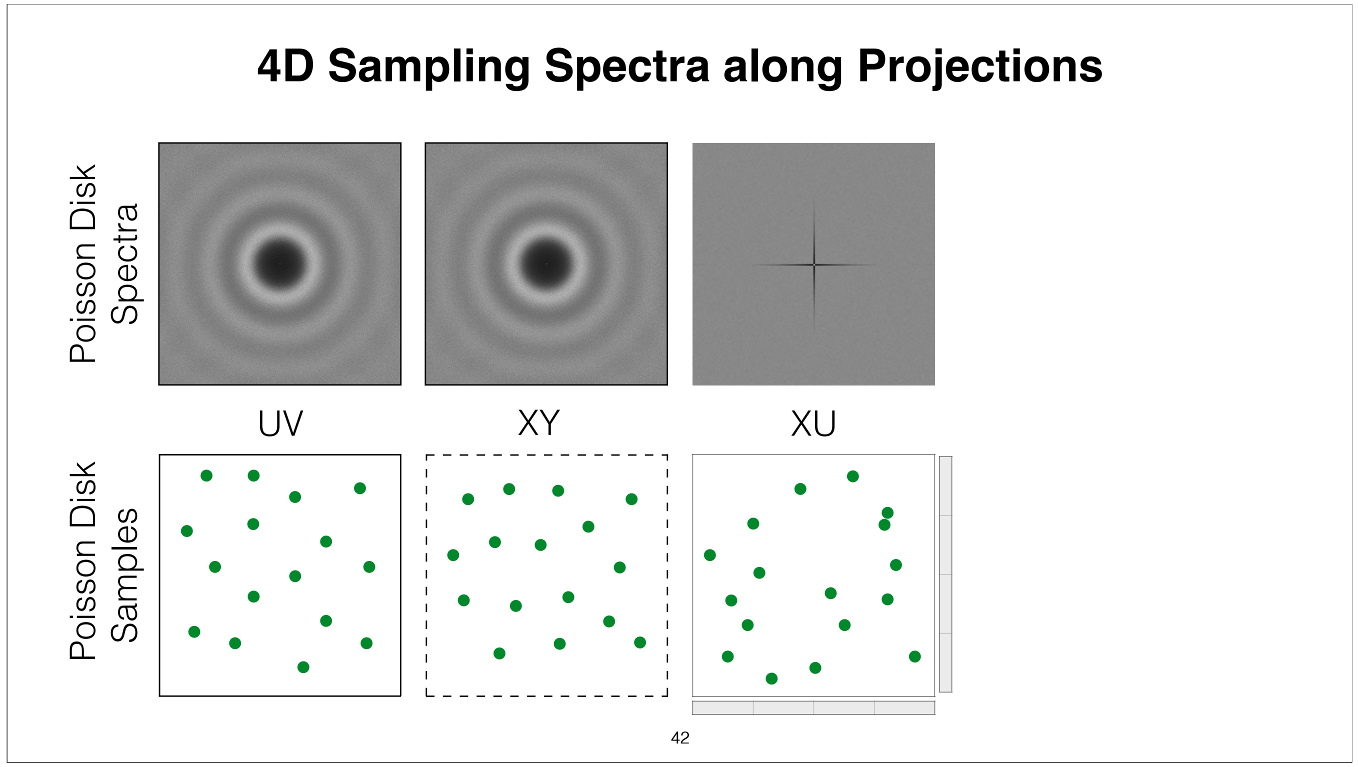

…power spectrum in these 2D projections, it has hairline anisotropy along the [CLICK] canonical axes on top of the [CLICK] big dark region that corresponds to 2D jittered sampling. However, If we look at the XU projection…

4D Sampling Spectra along Projections

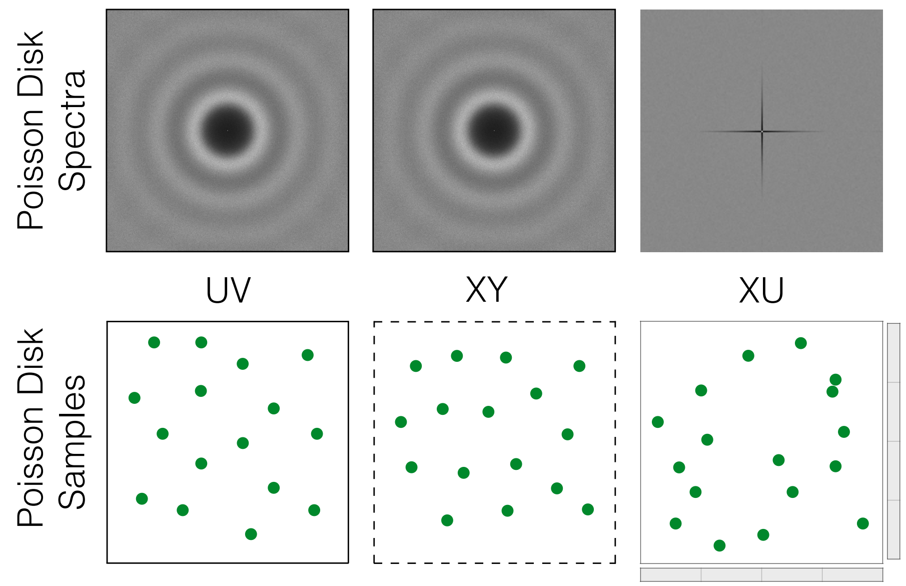

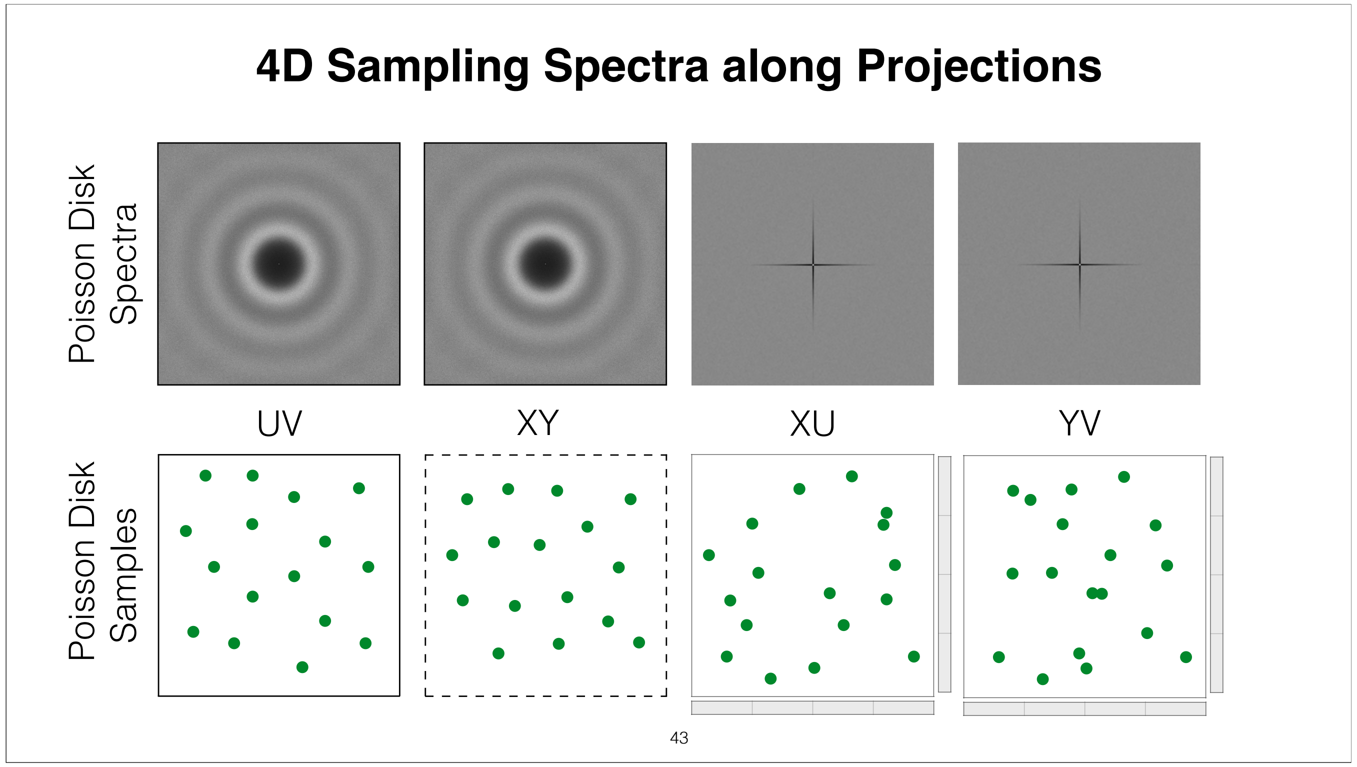

…power spectrum in these 2D projections, it has hairline anisotropy along the [CLICK] canonical axes on top of the [CLICK] big dark region that corresponds to 2D jittered sampling. However, If we look at the XU projection…

How can we perform Convergence Analysis for Anisotropic Sampling Spectra ?

…to higher dimensions. For example, in 4D…

Variance Formulation for Anisotropic Sampling Spectra

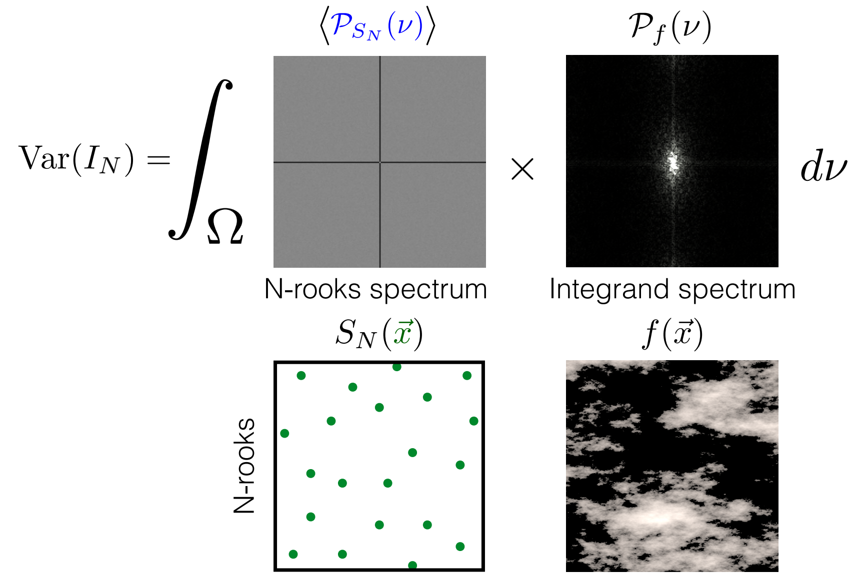

…for N-rooks sampler as a product of sampling and integrand spectra. As before, we rewrite this formulation in polar coordinates…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \int_0^\infty \rho^{d-1} \int_{\mathcal{S}^{d-1}} \left\langle \mathcal{P}_{S_N}(\rho\mathbf{n}) \right\rangle \times \mathcal{P}_f(\rho\mathbf{n}) \, d\mathbf{n}\, d\rho$$

$\left\langle \mathcal{P}_{S_N}(\rho\mathbf{n}) \right\rangle$

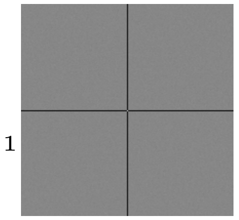

$\mathcal{P}_f(\rho\mathbf{n})$

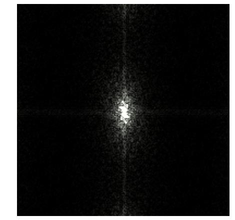

…as a double integral. By switching the order of the integration…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \int_{\mathcal{S}^{d-1}} \int_0^{\infty} \rho^{d-1}$$

$\left\langle \mathcal{P}_{S_N}(\rho\mathbf{n}) \right\rangle$

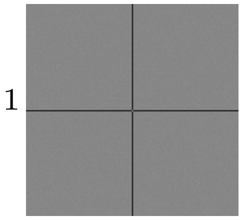

$\times$

$\mathcal{P}_f(\rho\mathbf{n})$

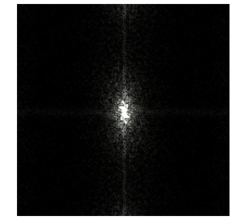

$d\rho \, d\mathbf{n}$

…the inner integral now represents the integration over the radial frequencies…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \int_{\mathcal{S}^{d-1}} \int_0^{\infty} \rho^{d-1}$$

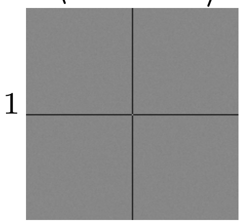

$$\times$$

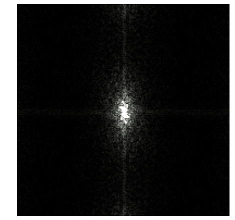

$$d\rho \, d\mathbf{n}$$

… for a given k-th direction. Since the outer integral…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \int_{\mathcal{S}^{d-1}} \int_0^{\infty} \rho^{d-1}$$

$\left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle$

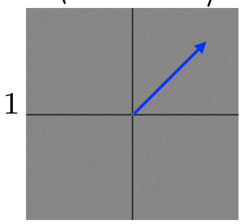

$\times$

$\mathcal{P}_f(\rho_k \mathbf{n}_k)$

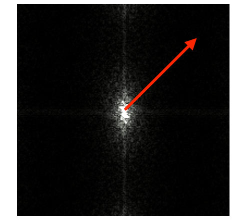

$d\rho \, d\mathbf{n}$

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \int_{\mathcal{S}^{d-1}} \int_0^\infty \rho^{d-1} \left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle \times \mathcal{P}_f(\rho_k \mathbf{n}_k) \, d\rho \, d\mathbf{n}$$

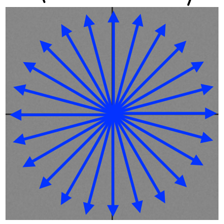

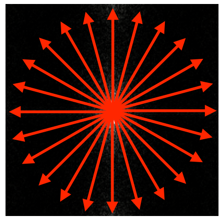

… is over all the directions, using Reimann summation, we can rewrite the outer integral…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \lim_{m\to\infty} \sum_{k=1}^{m} \int_0^\infty \rho^{d-1} \left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle \times \mathcal{P}_f(\rho_k \mathbf{n}_k) \, d\rho \; \Delta\mathbf{n}_k$$

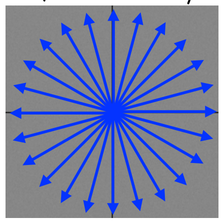

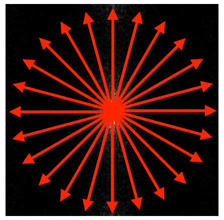

…as a summation over an infinite directional cones. After slight rearrangement, we obtain the variance formulation…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \lim_{m\to\infty} \sum_{k=1}^{m} \int_0^{\infty} \rho^{d-1} \quad \left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle \quad \times \quad \mathcal{P}_f(\rho_k \mathbf{n}_k) \quad d\rho \; \Delta\mathbf{n}_k$$

…as a summation over an infinite directional cones. After slight rearrangement, we obtain the variance formulation…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \lim_{m\to\infty} \sum_{k=1}^{m} \int_0^{\infty} \rho^{d-1} \left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle \mathcal{P}_f(\rho_k \mathbf{n}_k)\, d\rho \, \Delta\mathbf{n}_k$$

… for anisotropic samplers. One thing to note about this formulation is that, the inner integral…

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \lim_{m\to\infty} \sum_{k=1}^{m} \boxed{\int_0^\infty \rho^{d-1} \left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle \mathcal{P}_f(\rho_k \mathbf{n}_k)\, d\rho} \, \Delta\mathbf{n}_k$$

$\left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle$

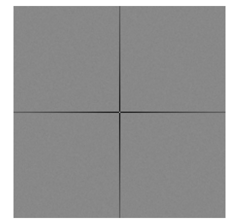

$\mathcal{P}_f(\rho_k \mathbf{n}_k)$

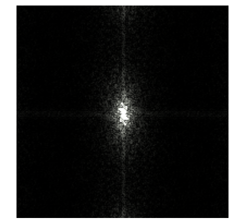

…considers the radial behavior of both the sampling and the integrand spectra for a given k-th direction, which when [CLICK] added up for all the directional cones, gives the variance. This shows that, unlike previous work, our formulation not only [CLICK] handles anisotropic sampling spectra but also intimately couples the anisotropic structures present in the integrand spectrum with that of the sampling spectrum. We will see shortly how this impacts the convergence rate but first, lets analyze the anisotropic structures of N-rooks spectrum more closely.

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \lim_{m\to\infty}\sum_{k=1}^{m}\boxed{\int_0^\infty \rho^{d-1}\left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k})\right\rangle \mathcal{P}_f(\rho_k\mathbf{n}_k)\,d\rho}\,\Delta\mathbf{n}_k$$

$\left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k})\right\rangle$

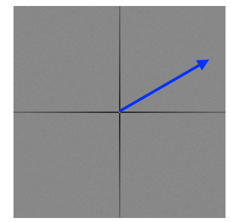

$\mathcal{P}_f(\rho_k\mathbf{n}_k)$

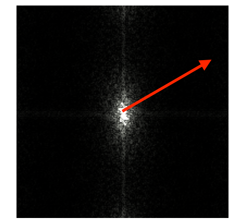

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \lim_{m\to\infty} \sum_{k=1}^{m} \boxed{\int_0^{\infty} \rho^{d-1} \left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle \mathcal{P}_f(\rho_k \mathbf{n}_k)\, d\rho} \,\Delta \mathbf{n}_k$$

$\left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle$

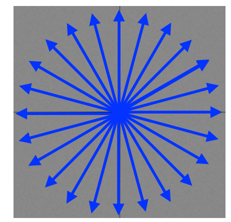

$\mathcal{P}_f(\rho_k \mathbf{n}_k)$

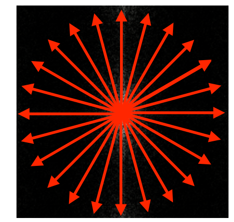

Variance Formulation for Anisotropic Sampling Spectra

$$\mathrm{Var}(I_N) = \lim_{m\to\infty} \sum_{k=1}^{m} \boxed{\int_0^\infty \rho^{d-1} \left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle \mathcal{P}_f(\rho_k \mathbf{n}_k)\, d\rho} \, \Delta \mathbf{n}_k$$

$\left\langle \mathcal{P}_{S_N}(\rho_k \mathbf{n_k}) \right\rangle$

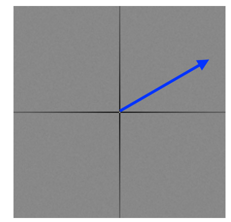

$\mathcal{P}_f(\rho_k \mathbf{n}_k)$

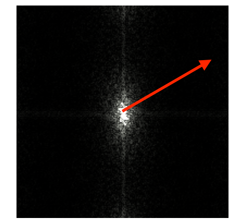

Convergence Analysis for Anisotropic Sampling Spectra

Power Spectrum

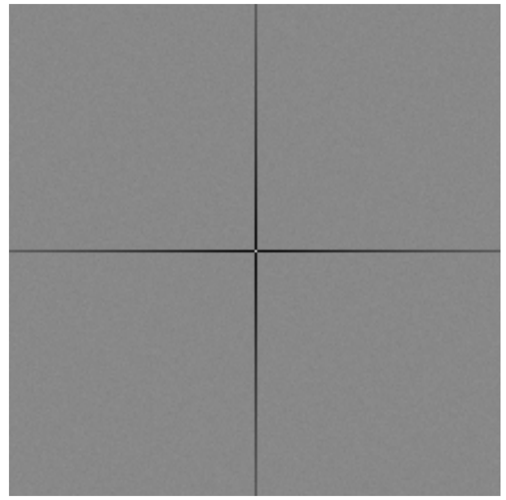

Radial Power Spectrum

N-rooks power spectrum has [CLICK] jittered radial profile along the canonical axes, and [CLICK] a constant radial profile along all other directions. For the convergence rate, we only need…

Convergence Analysis for Anisotropic Sampling Spectra

Power Spectrum

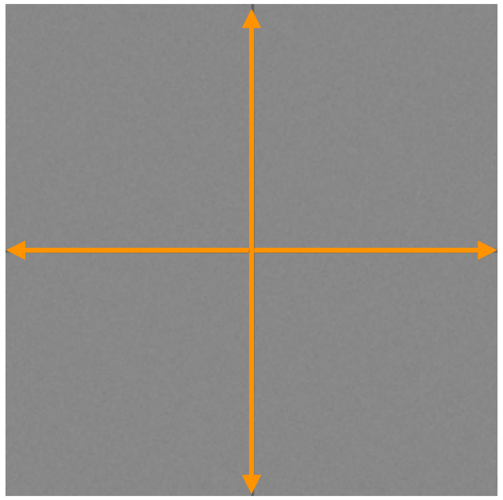

Radial Power Spectrum

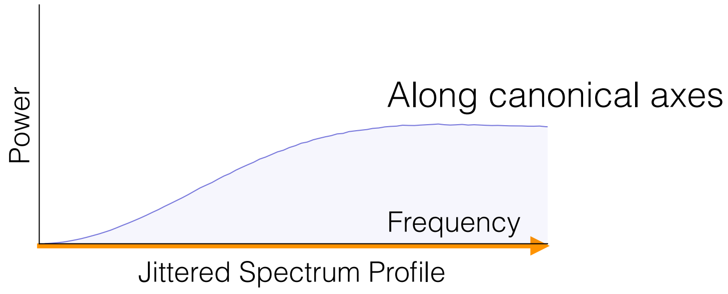

Convergence Analysis for Anisotropic Sampling Spectra

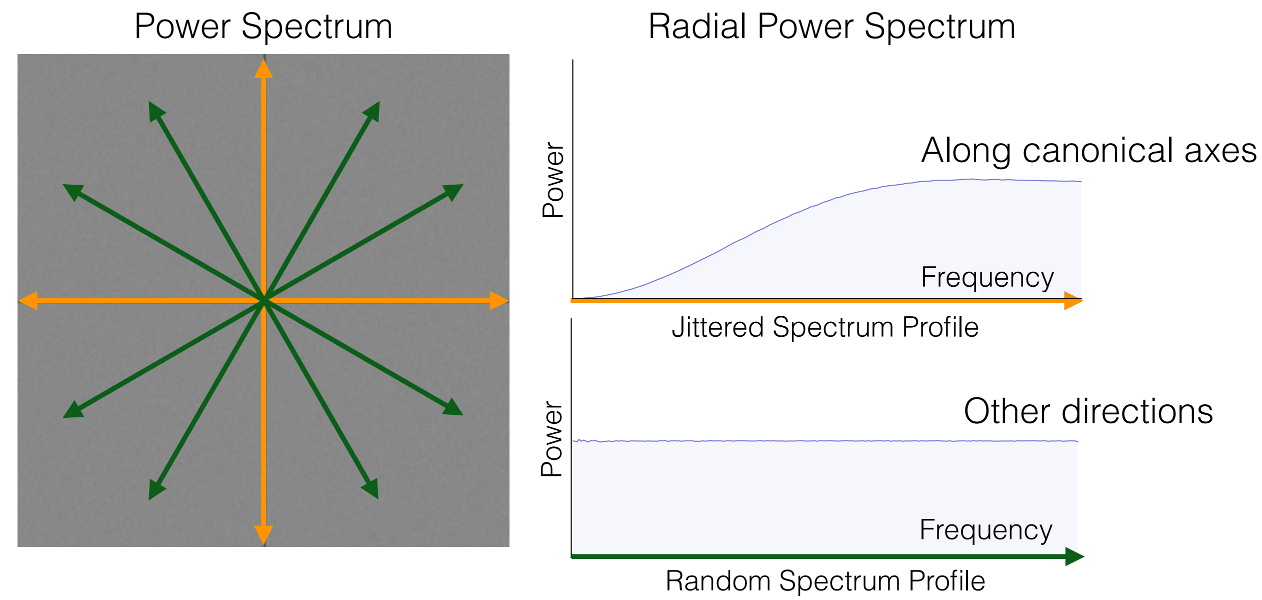

Convergence Analysis for Anisotropic Sampling Spectra

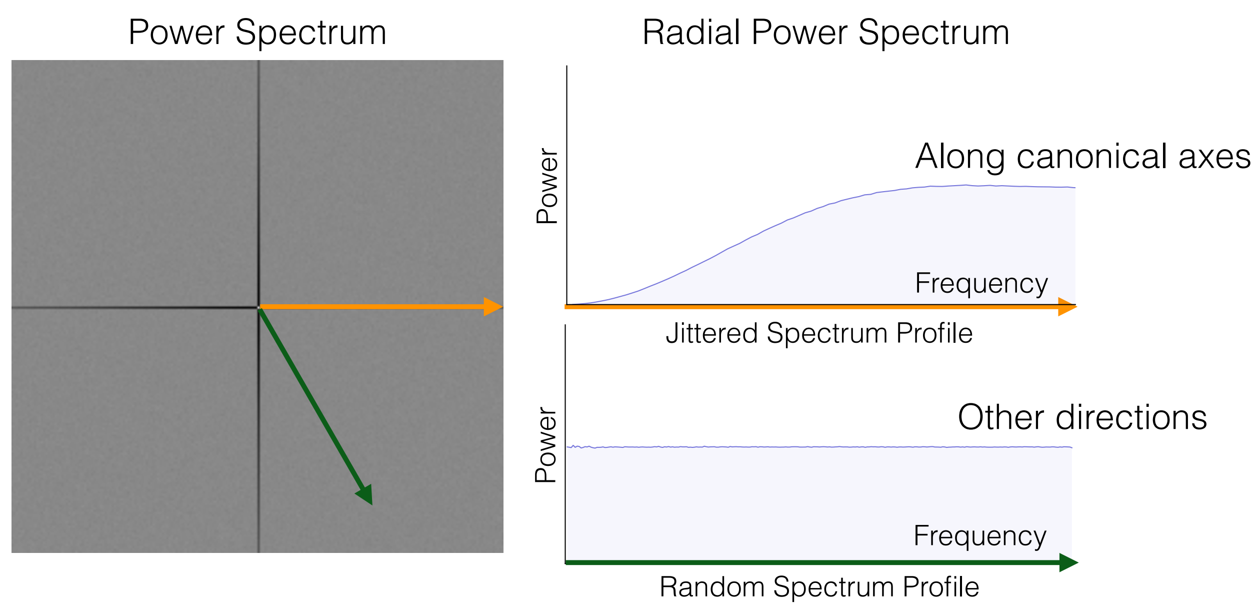

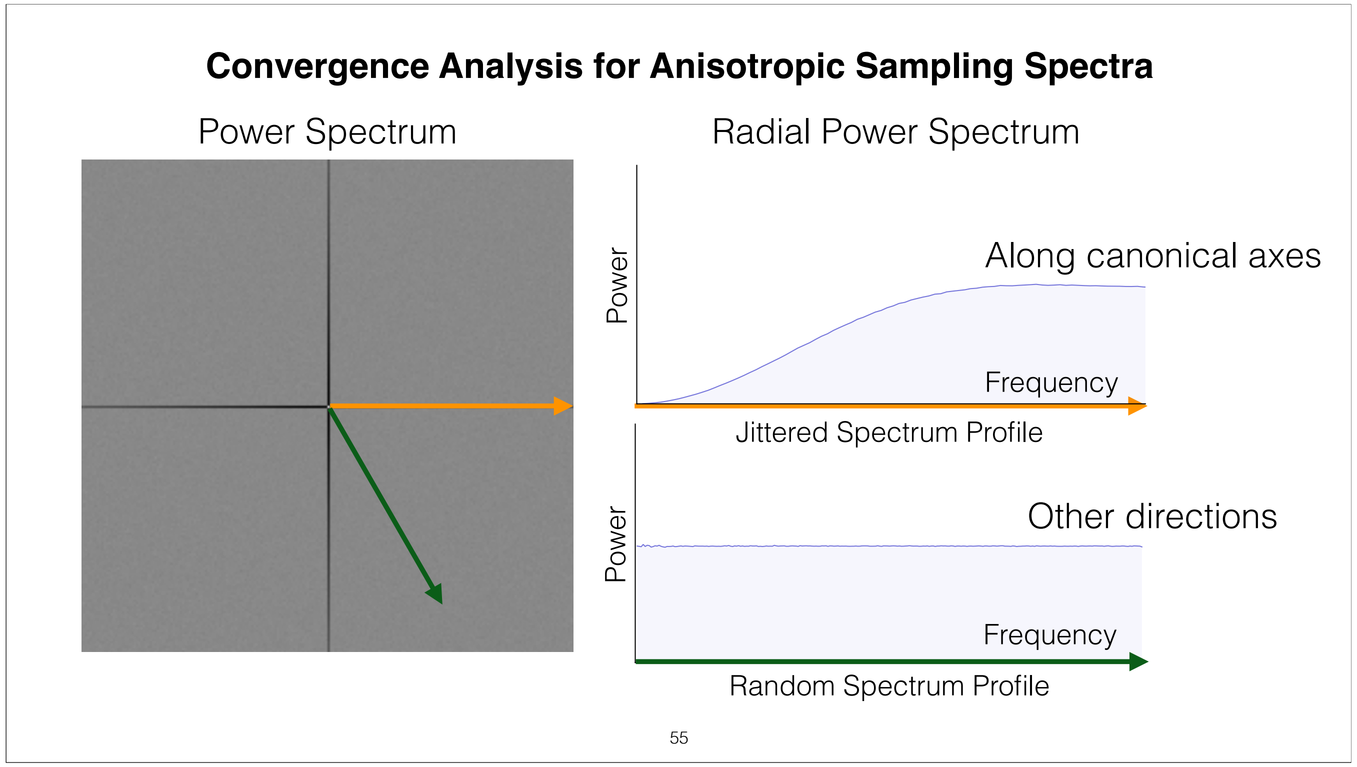

… one of the canonical direction (shown in purple) and one of the direction from the rest of the spectrum (shown in green) since the behavior is the same in all other directions. Now, depending on the integrands…

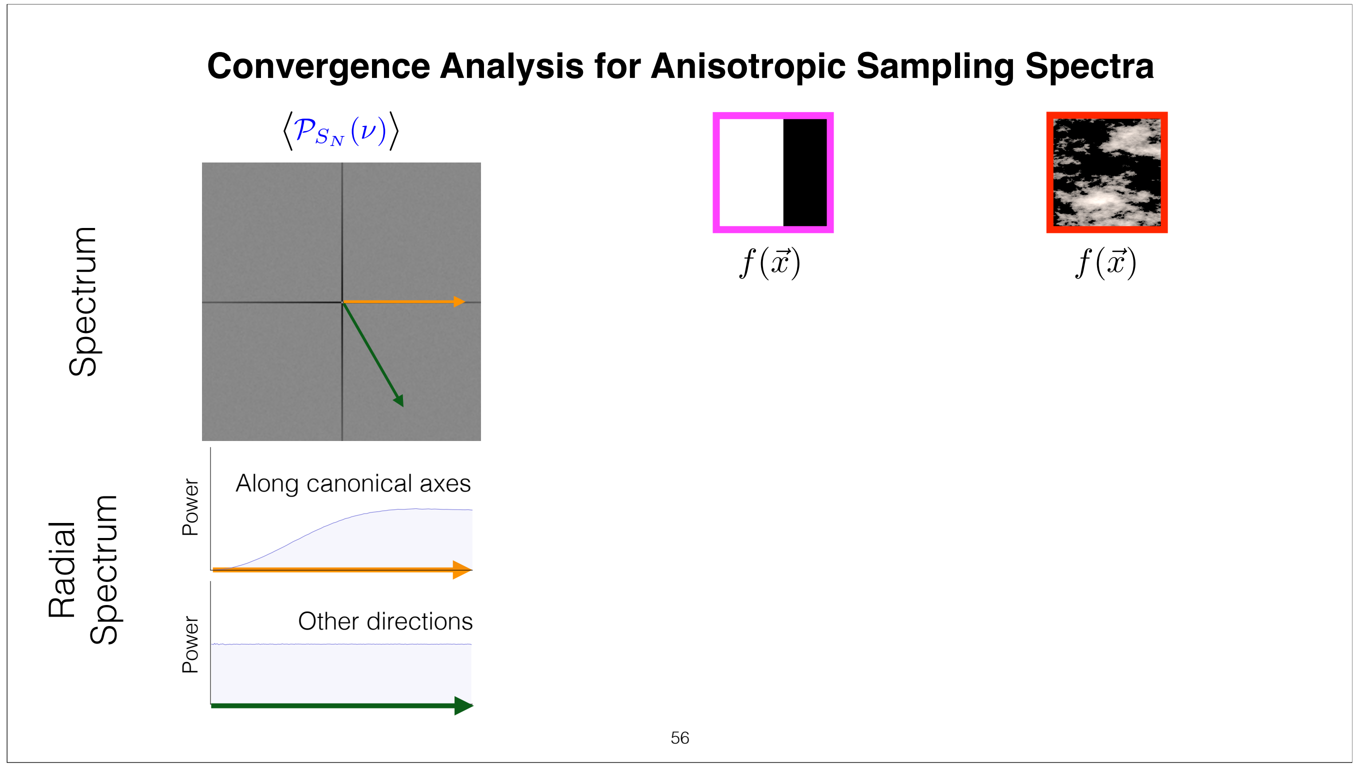

…we can get different convergence rates from the same sampler. For example, the step function in magenta box has [CLICK] a power spectrum with all its energy along the horizontal axis. As a result, only the horizontal axis [CLICK] with jittered profile would overlap with the integrand spectrum and will result in a convergence rate of [CLICK] O(N^-2). Since the other directions doesn't overlap with this integrand spectrum, they won't [CLICK] impact the convergence behavior.

However, if we have an integrand with a [CLICK] power spectrum having energy along all the directions, we may see [CLICK] two different convergence behavior in its variance plot as we go toward higher sample count. However, asymptotically only the worse of the two would dominate, and we will see a convergence rate of O(N^-1).

To understand this mathematically, lets look at the variance formulation, which is the product…

Convergence Analysis for Anisotropic Sampling Spectra

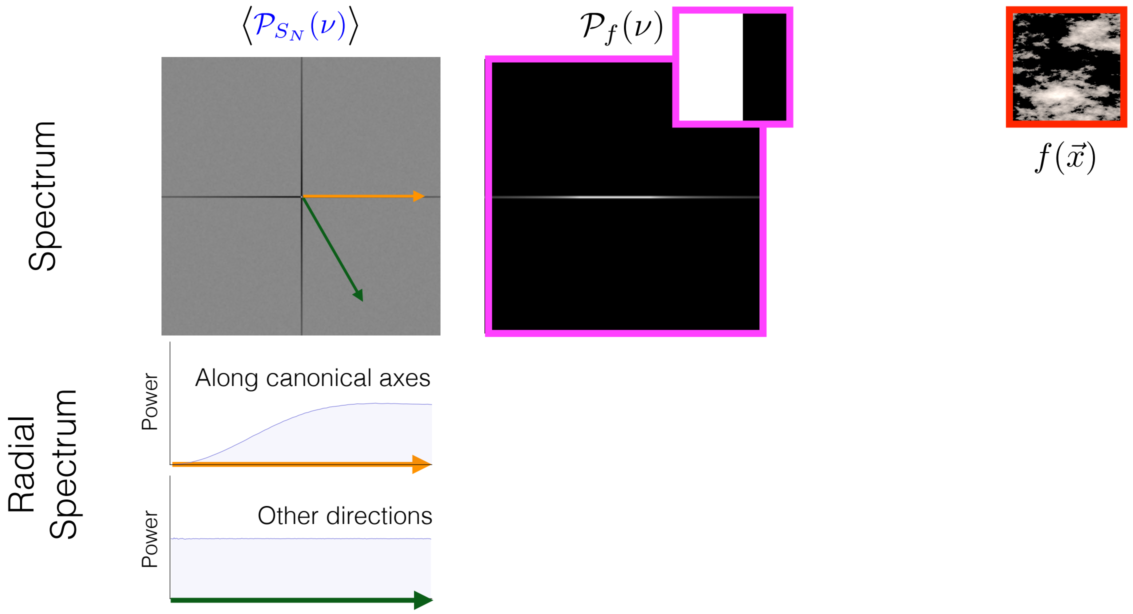

Convergence Analysis for Anisotropic Sampling Spectra

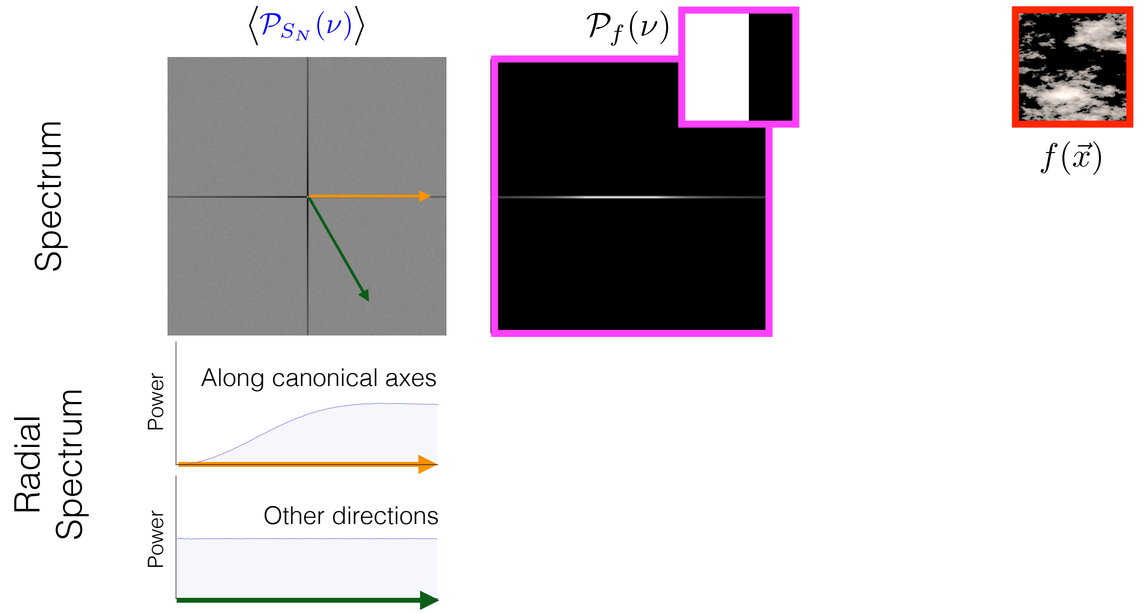

Convergence Analysis for Anisotropic Sampling Spectra

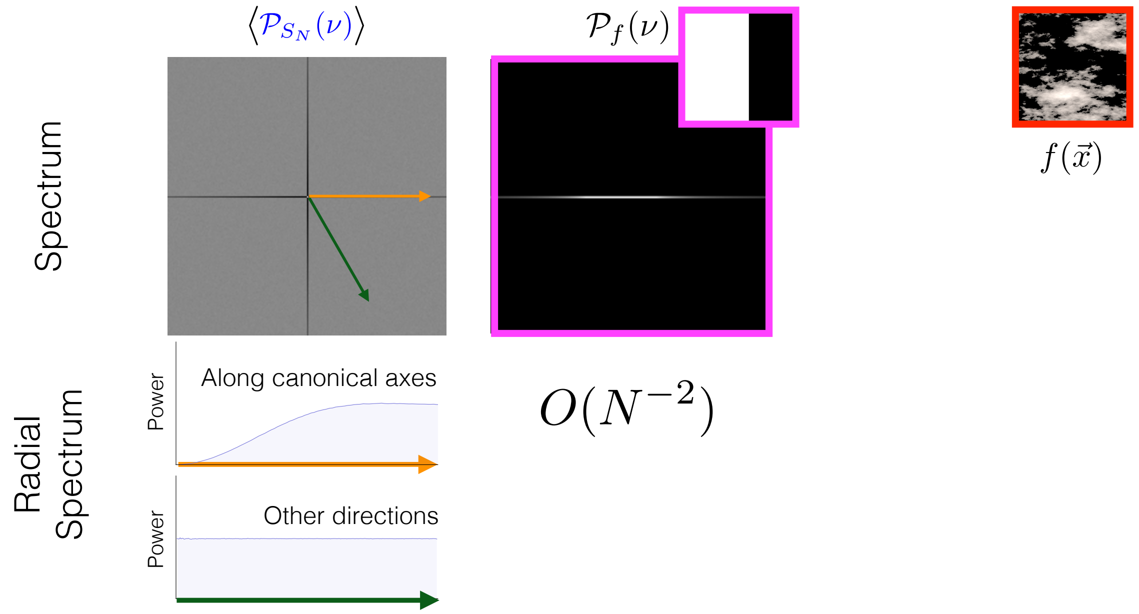

Convergence Analysis for Anisotropic Sampling Spectra

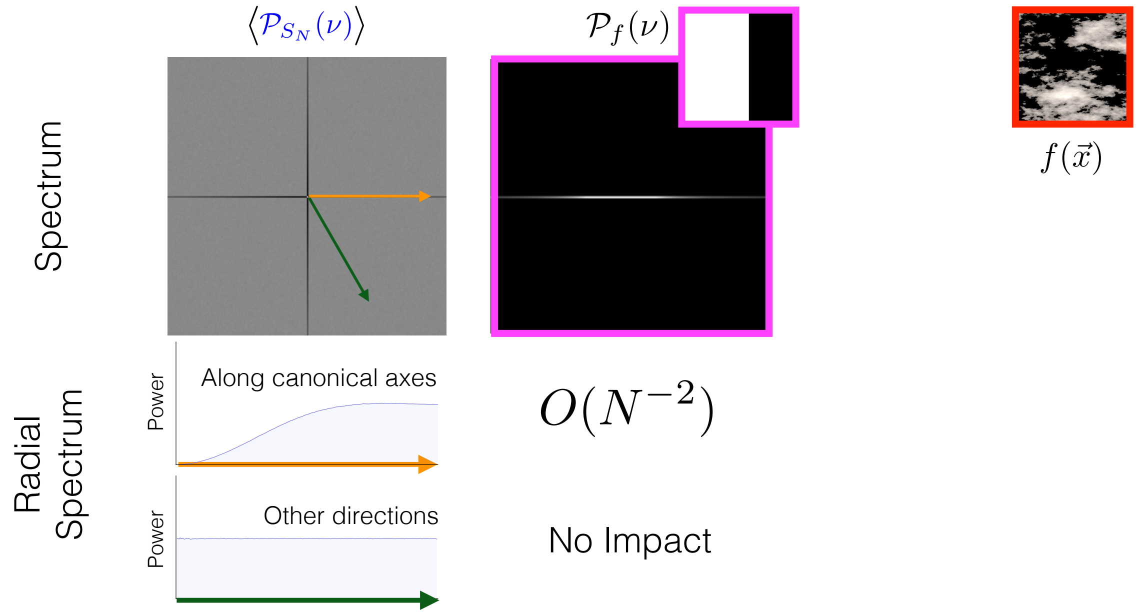

$O(N^{-2})$

No Impact

Convergence Analysis for Anisotropic Sampling Spectra

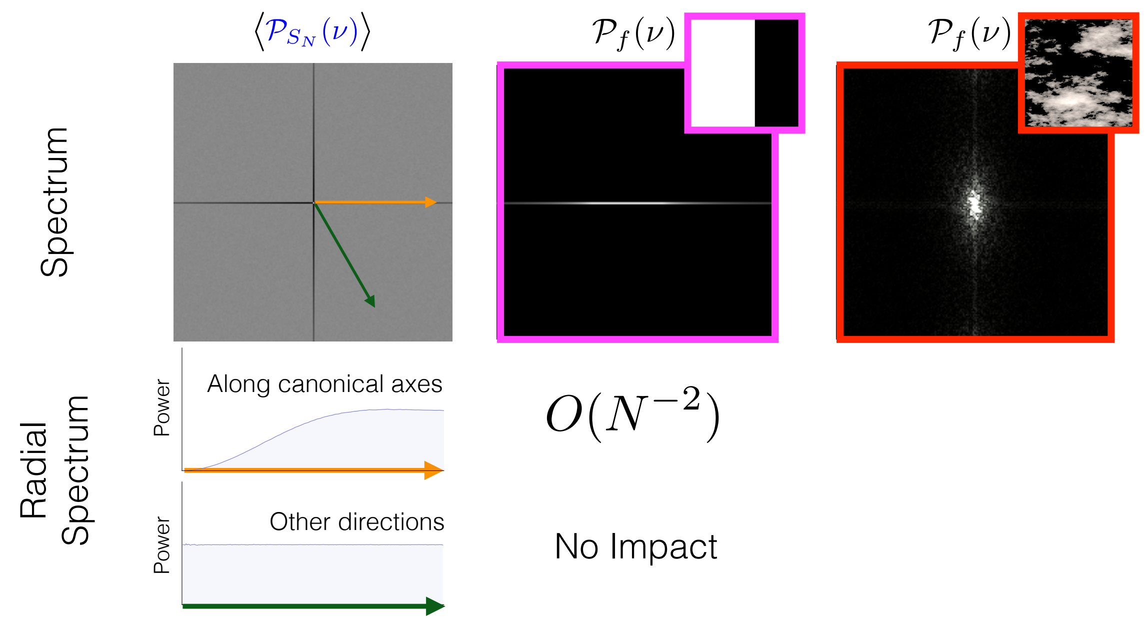

Convergence Analysis for Anisotropic Sampling Spectra

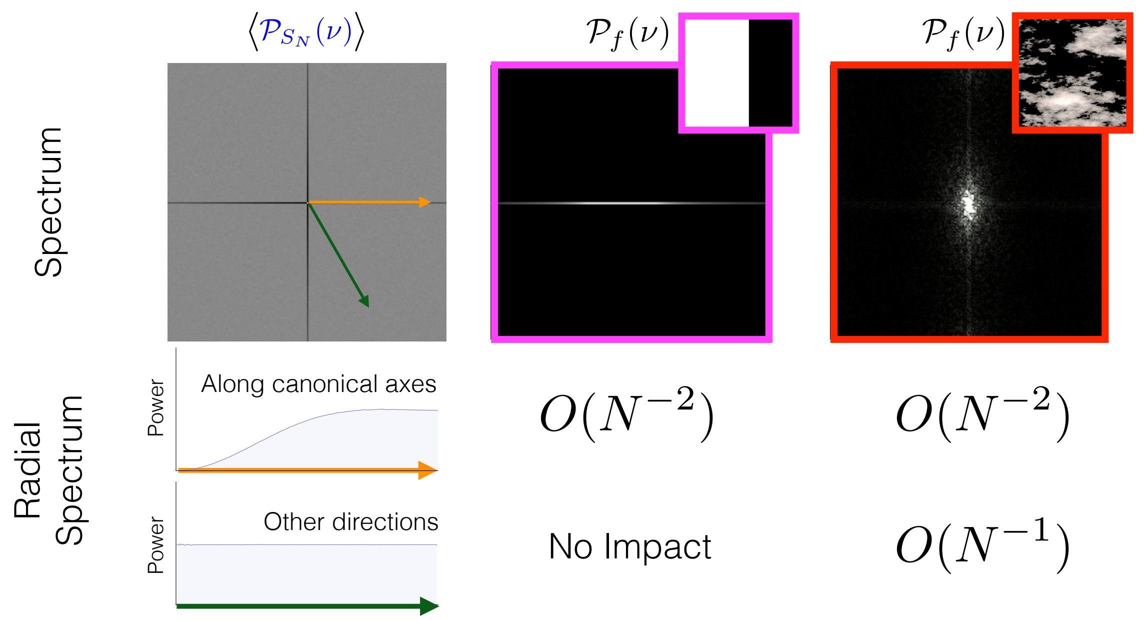

Convergence Analysis for Anisotropic Sampling Spectra

Spectrum

$\langle \mathcal{P}_{S_N}(\nu) \rangle$

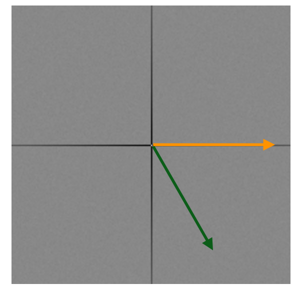

$\mathcal{P}_f(\nu)$

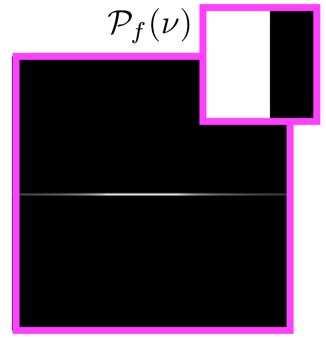

$\mathcal{P}_f(\nu)$

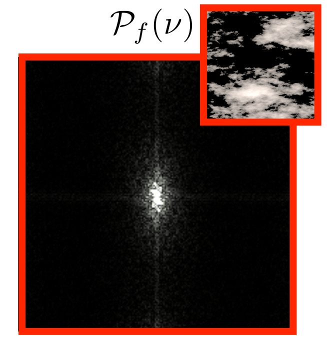

Radial Spectrum

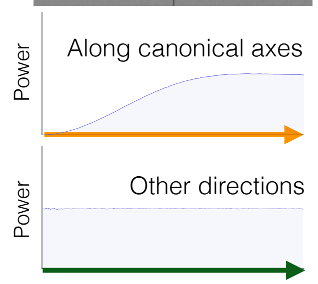

$O(N^{-2})$

No Impact

$$\left.\begin{matrix} O(N^{-2}) \\ \\ O(N^{-1}) \end{matrix}\right\} O(N^{-1})$$

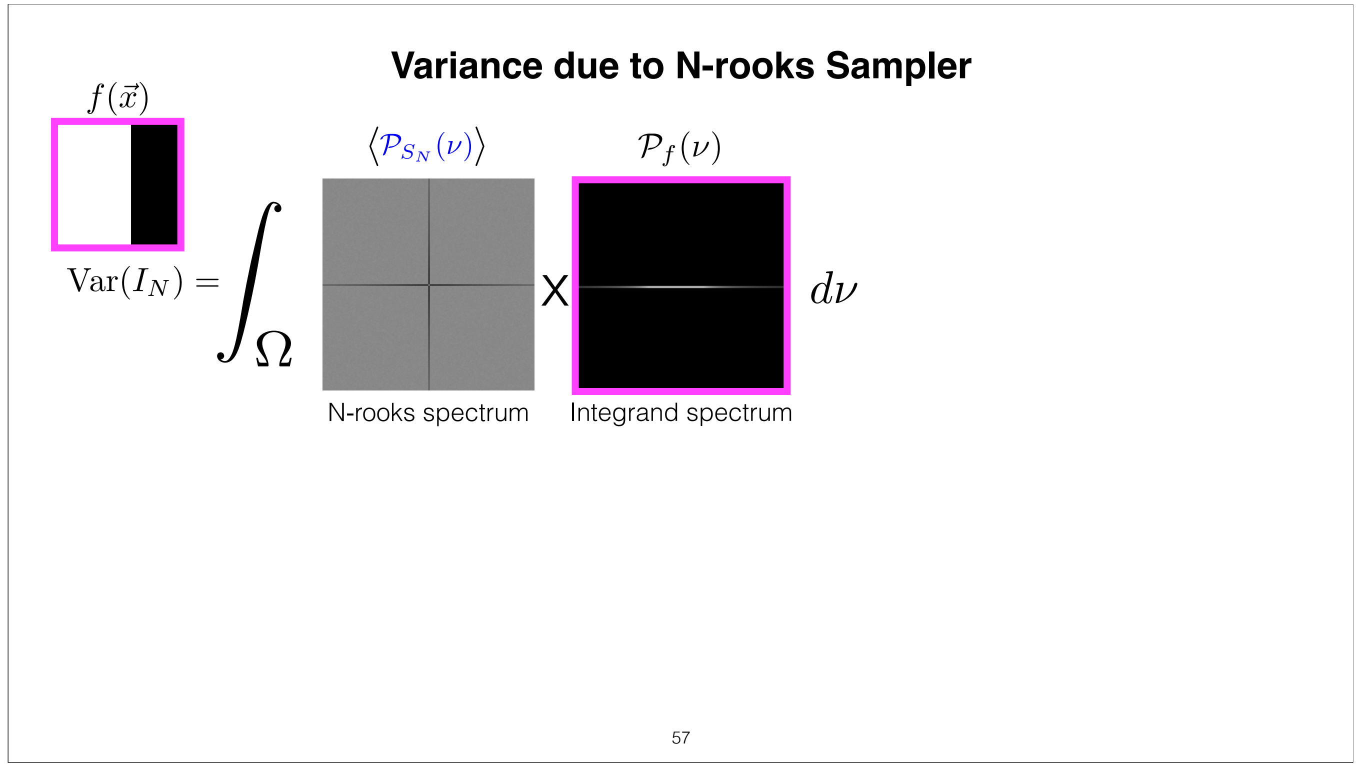

… of N-rooks sampling spectrum and the integrand spectrum. Due to the dark hairline anisotropy [CLICK] present in the sampling spectrum, their [CLICK] product goes down very quickly, resulting in huge variance reduction and good asymptotic convergence. However, for the second pixel…

Variance due to N-rooks Sampler

$f(\vec{x})$

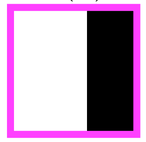

$$\mathrm{Var}(I_N) = \int_\Omega$$

$\langle \mathcal{P}_{S_N}(\nu) \rangle$

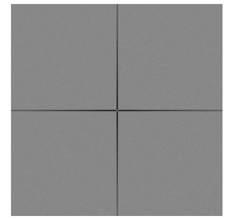

N-rooks spectrum

X

$\mathcal{P}_f(\nu)$

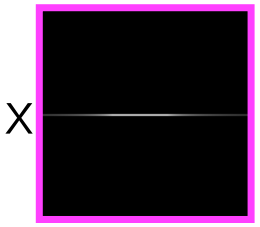

Integrand spectrum

$d\nu$

Variance due to N-rooks Sampler

$f(\vec{x})$

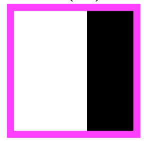

$\mathrm{Var}(I_N) = \int_\Omega$

$\langle \mathcal{P}_{S_N}(\nu) \rangle$

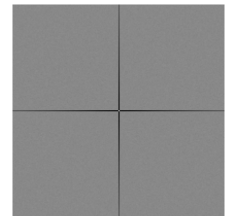

N-rooks spectrum

X

$\mathcal{P}_f(\nu)$

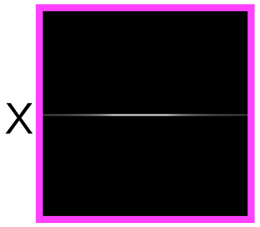

Integrand spectrum

$d\nu = \int_\Omega$

$d\nu$

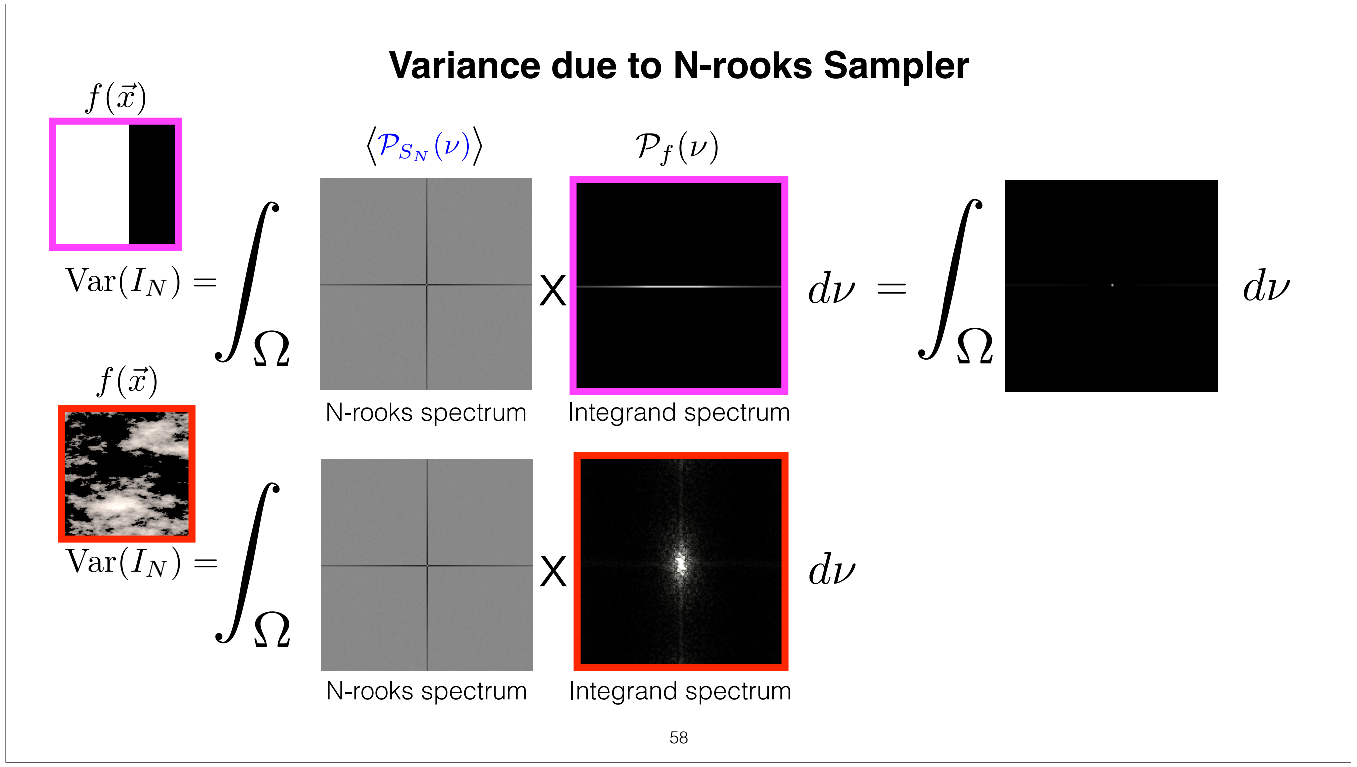

… since the integrand spectrum has energy spread over all the directions, the [CLICK] hairline anisotropy of the sampling spectrum [CLICK] does not significantly reduce the product, resulting in higher variance. We further verified this..

Variance due to N-rooks Sampler

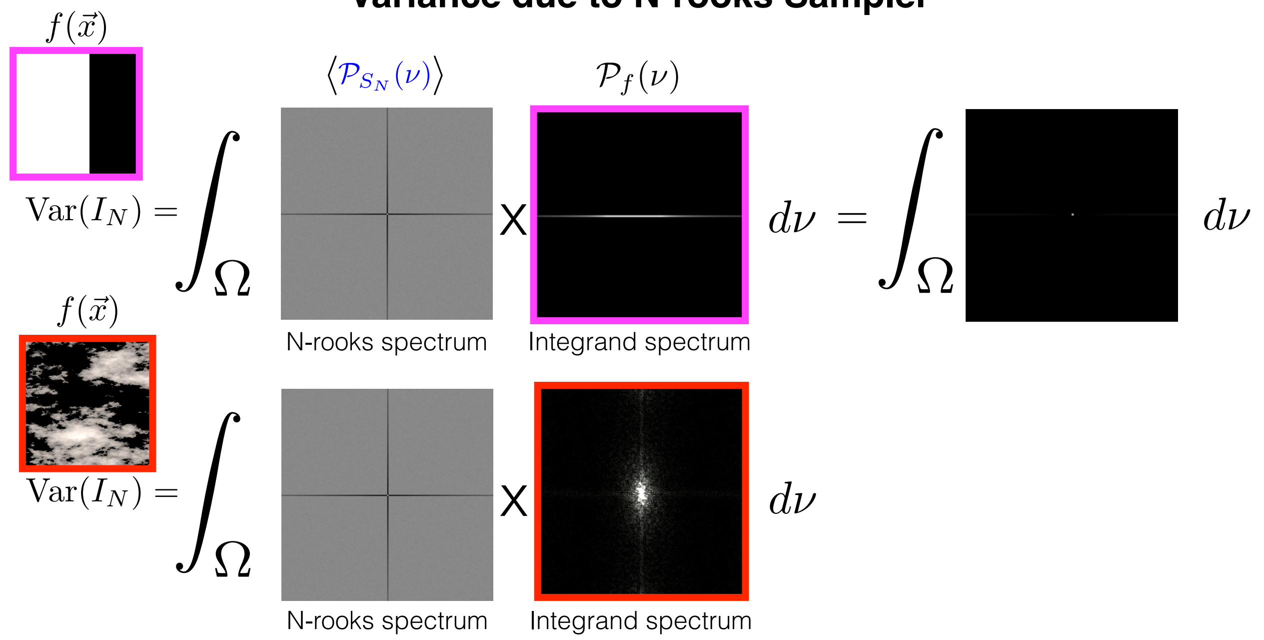

Variance due to N-rooks Sampler

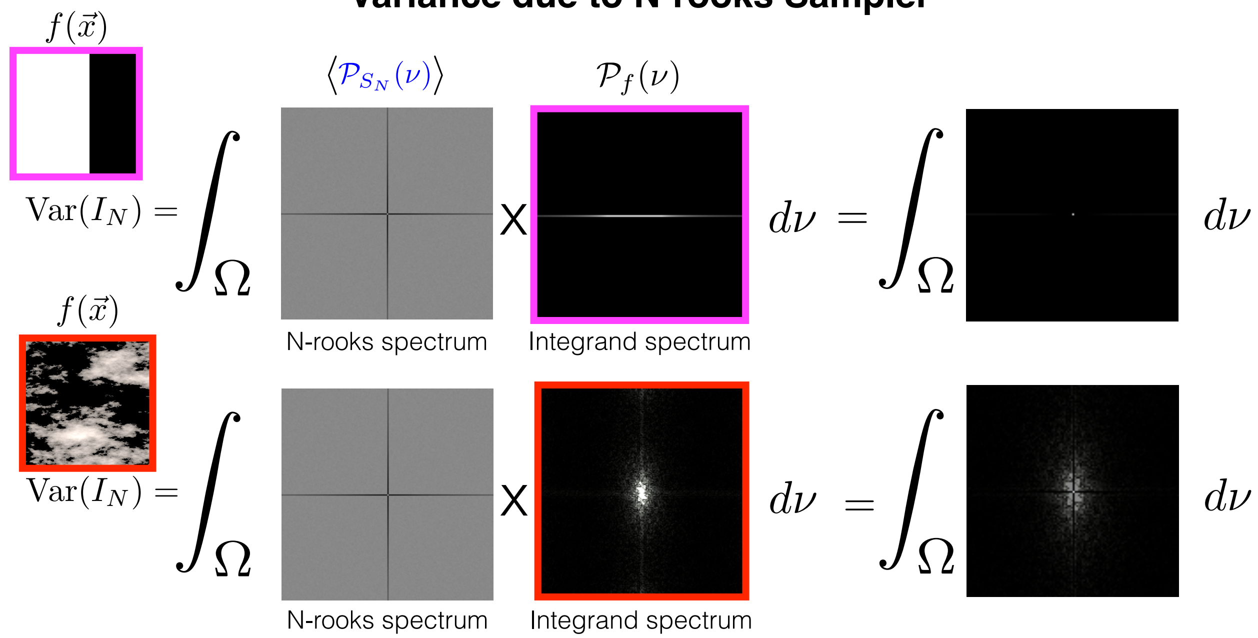

Variance Convergence of Latin Hypercube (N-rooks)

$\langle \mathcal{P}_{S_N}(\nu) \rangle$ $\mathcal{P}_f(\nu)$

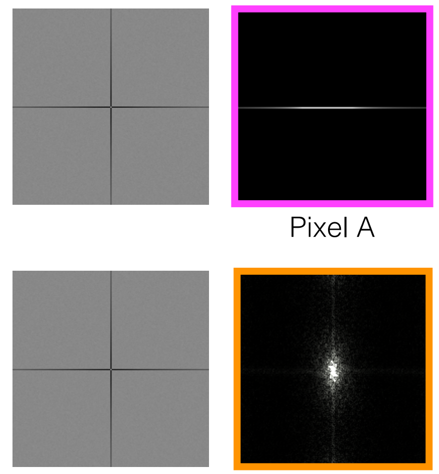

Pixel A

Pixel B

…experimentally, where we plot variance with increasing sample count. This shows that if we can align the anisotropic structures of the sampling spectrum Ps with that of the integrand spectrum Pf, we can gain huge variance reductions, as shown with the magenta curve. But in most scenarios…

Variance Convergence of Latin Hypercube (N-rooks)

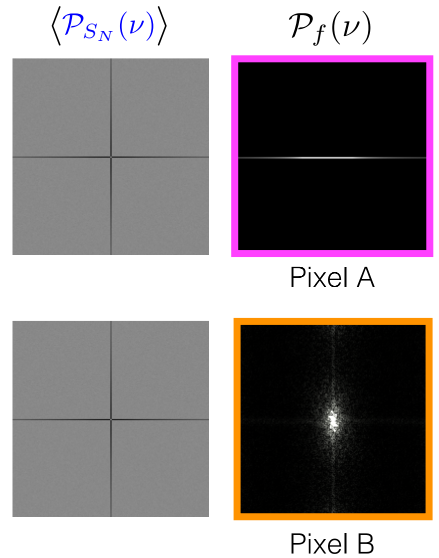

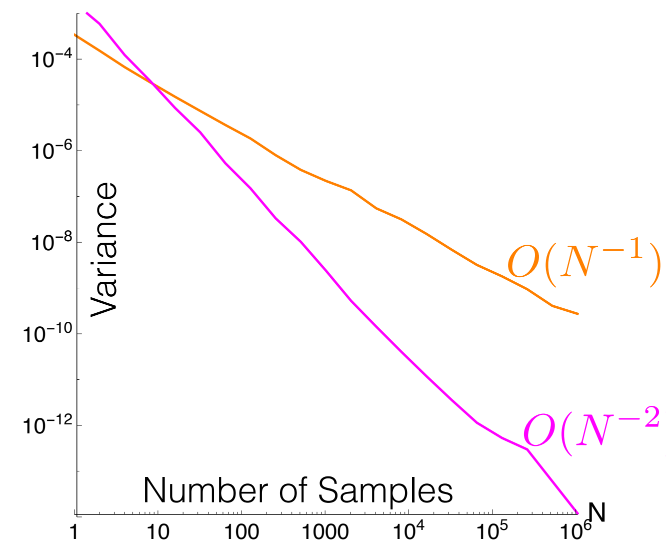

Non-Axis Aligned Integrand Spectra

$\mathcal{P}_f(\nu)$

Integrand Spectrum

…the underlying integrand spectrum has arbitrary orientation. If we choose to sample this function…

Non-Axis Aligned Integrand Spectra

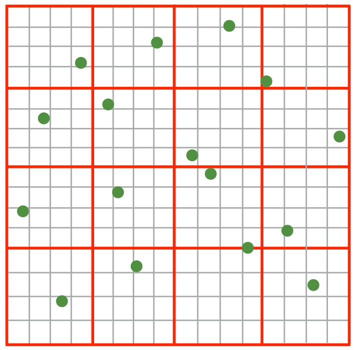

Multi-jittered Samples

$\mathcal{P}_f(\nu)$

Integrand Spectrum

…with multi-jittered samples which has [CLICK] the following power spectrum, we won't be able to benefit from these hairline anisotropic structures since they are axis-aligned. To solve this issue, we propose to shear…

Non-Axis Aligned Integrand Spectra

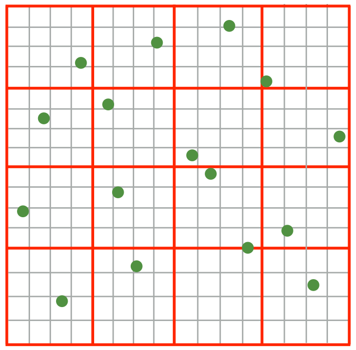

Multi-jittered Samples

$\langle \mathcal{P}_{S_N}(\nu) \rangle$

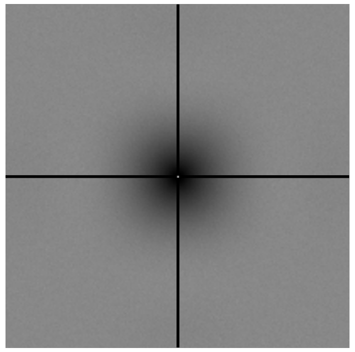

Sampling Spectrum

$\mathcal{P}_f(\nu)$

Integrand Spectrum

Shearing Multi-Jittered Samples

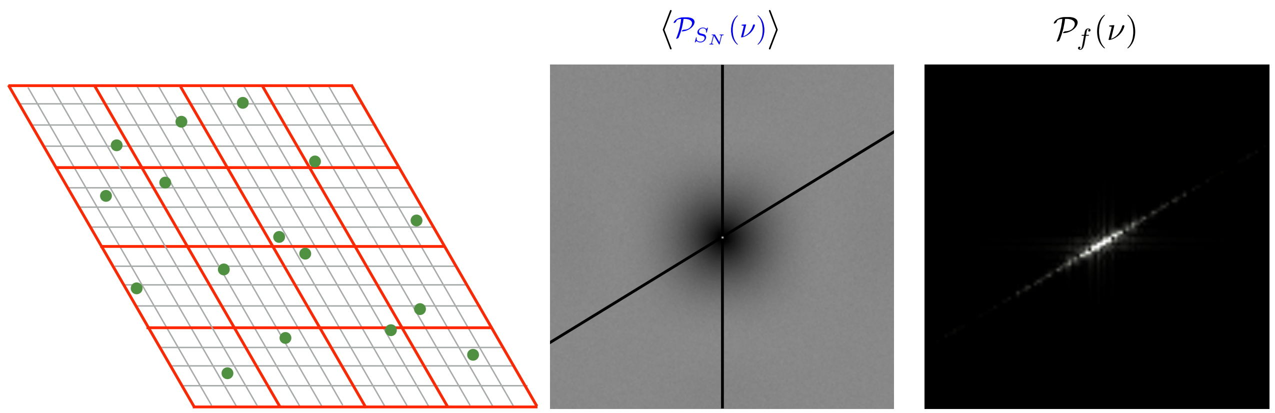

Sheared Samples Sheared Spectrum Integrand Spectrum

…the samples in such a way that we can align the sampling spectrum with that of the integrand spectrum. But, the key question here is…

How can we determine the sample shearing parameters ?

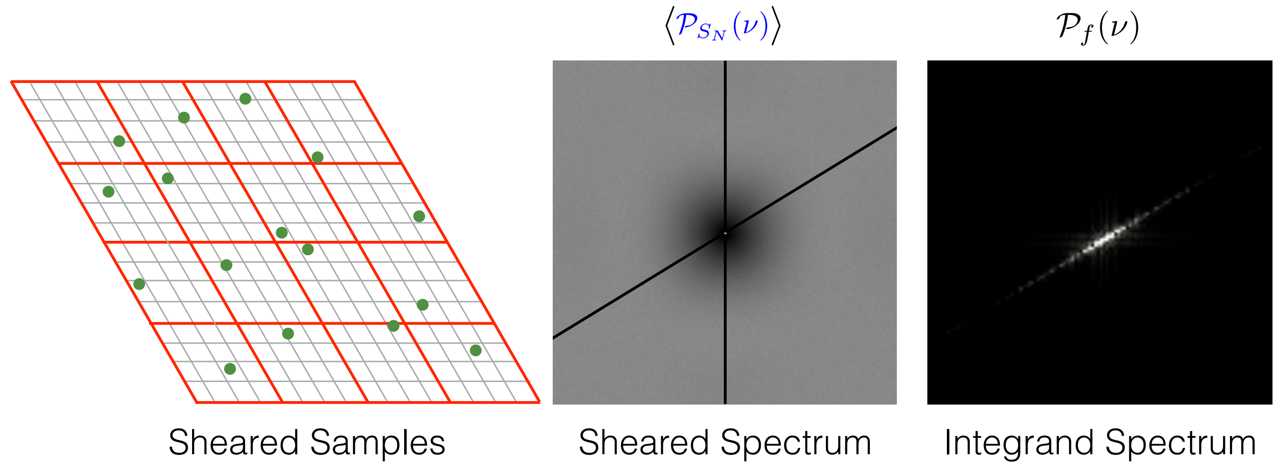

…how can we determine these shear parameters ? The answer to this question requires some preprocessing to know the frequency content of the integrand. Therefore, we propose the following steps.

Our Algorithm

64

In the first step, we
[CLICK] leverage the light transport frequency analysis, developed over more than a decade, to create an oracle that can give us the shear parameters of the integrand spectrum.
[CLICK] We then use these shear parameters given by the oracle to shear the samples.
[CLICK] After that, we use these sheared samples to perform Monte Carlo Integration.
Lets go over this algorithm starting from the….

Our Algorithm

1) Develop an oracle using the Frequency Analysis of Light Transport

Our Algorithm

1) Develop an oracle using the Frequency Analysis of Light Transport

2) Use this oracle to shear the samples

Our Algorithm

1) Develop an oracle using the Frequency Analysis of Light Transport

2) Use this oracle to shear the samples

3) Perform Monte Carlo integration using the sheared samples

Frequency Analysis of Light Transport

…frequency analysis of light transport. In 2005,…

Related Work

- Frequency Analysis of Light Transport Durand et al. [2005]

Durand and colleagues proposed a Fourier domain framework to study the light transport. Later on, [CLICK] this analysis was leveraged for depth of field, motion blur, soft shadows, ambient occlusion and many other effects. [CLICK] All this previous work has been extensively…

Related Work

- Frequency Analysis of Light Transport Durand et al. [2005]
- Depth of Field Soler et al. [2009]
- Motion Blur Egan et al. [2009]
- Ambient Occlusion Egan et al. [2011] and more…

Related Work

- Frequency Analysis of Light Transport Durand et al. [2005]
- Depth of Field Soler et al. [2009]
- Motion Blur Egan et al. [2009]
- Ambient Occlusion Egan et al. [2011] and more…

Durand and colleagues proposed a Fourier domain framework to study the light transport. Later on, [CLICK] this analysis was leveraged for depth of field, motion blur, soft shadows, ambient occlusion and many other effects. [CLICK] All this previous work has been extensively…

Related Work

- Frequency Analysis of Light Transport Durand et al. [2005]
- Depth of Field Soler et al. [2009]
- Motion Blur Egan et al. [2009]
- Ambient Occlusion Egan et al. [2011] and more…

Reconstruction

…used for reconstruction purposes. In this work [CLICK], we leverage the light transport frequency analysis [CLICK] for Integration purposes. We demonstrate our approach for a depth of field setup but our algorithm directly applies to other distribution effects like motion blur. In our setup…

Related Work

- Frequency Analysis of Light Transport Durand et al. [2005]
- Depth of Field Soler et al. [2009]
- Motion Blur Egan et al. [2009]
- Ambient Occlusion Egan et al. [2011] and more…

Reconstruction

Our Work

Related Work

- Frequency Analysis of Light Transport Durand et al. [2005]
- Depth of Field Soler et al. [2009]
- Motion Blur Egan et al. [2009]
- Ambient Occlusion Egan et al. [2011] and more…

Reconstruction

Our Work

Integration

Depth of Field Analysis

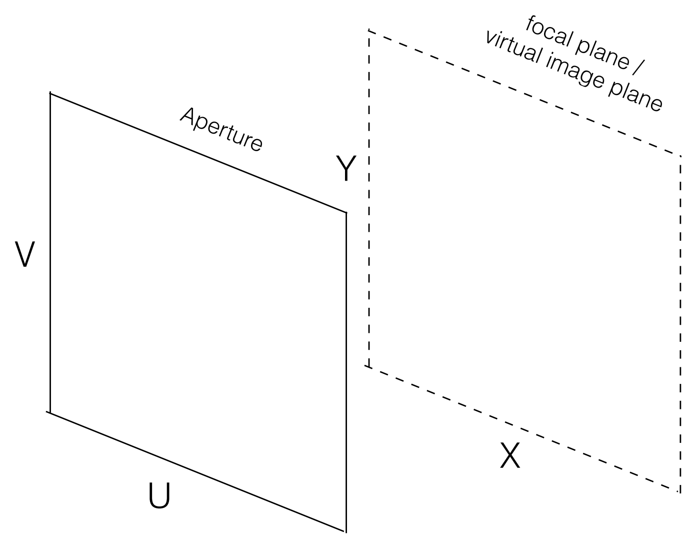

…we have a virtual image plane XY and a square aperture UV to simplify the analysis. We render a cornell box scene…

Depth of Field Analysis

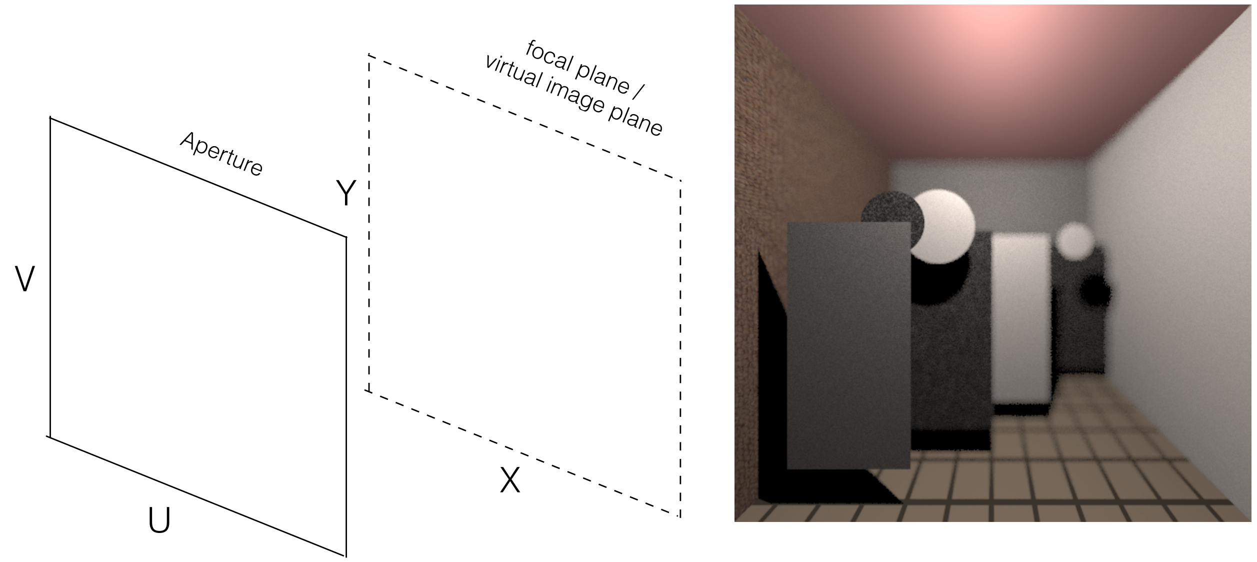

…with a defocus blur. Objects in this cornell box are placed at an increasing depth from your view point. Lets look at [CLICK] one pixel of this image [CLICK] which has [CLICK] the following underlying texture, and see how the light field is changing. To simplify the setup, we consider a 1D aperture…

Depth of Field Analysis

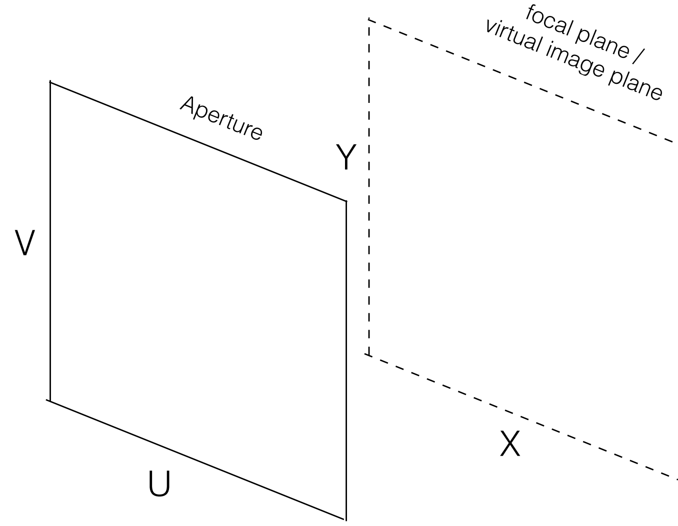

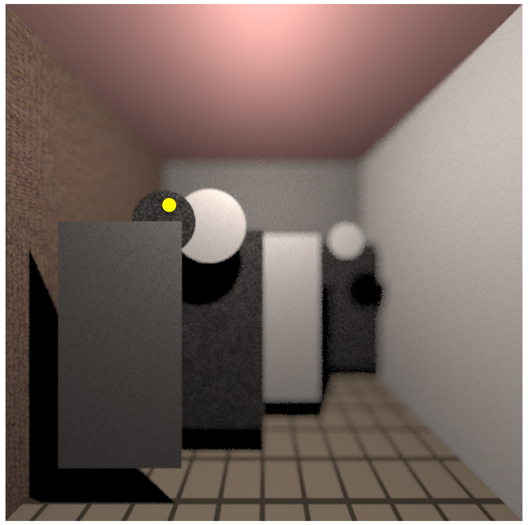

Depth of Field Analysis

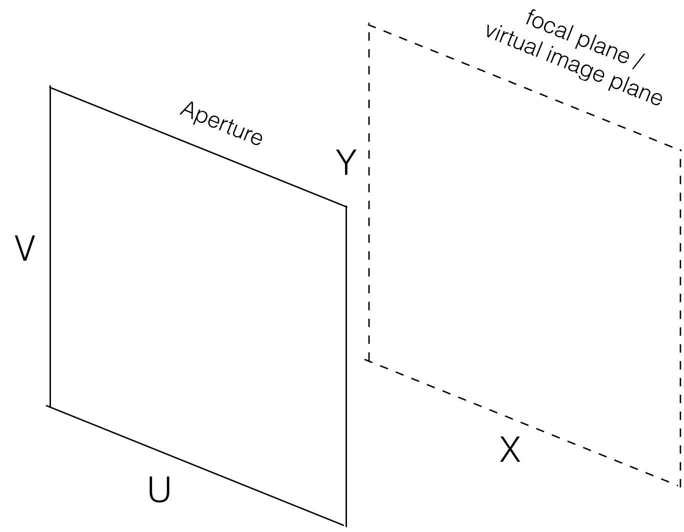

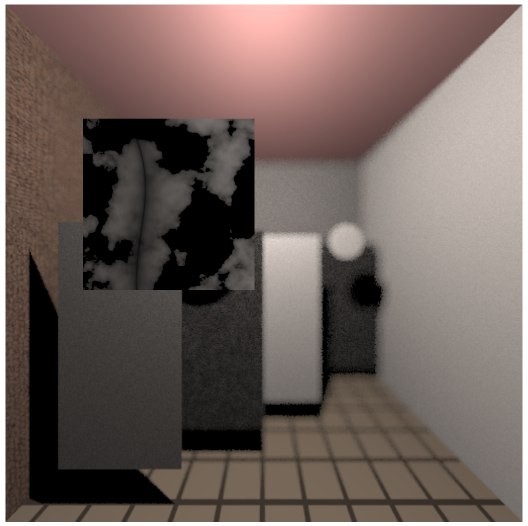

Depth of Field Analysis

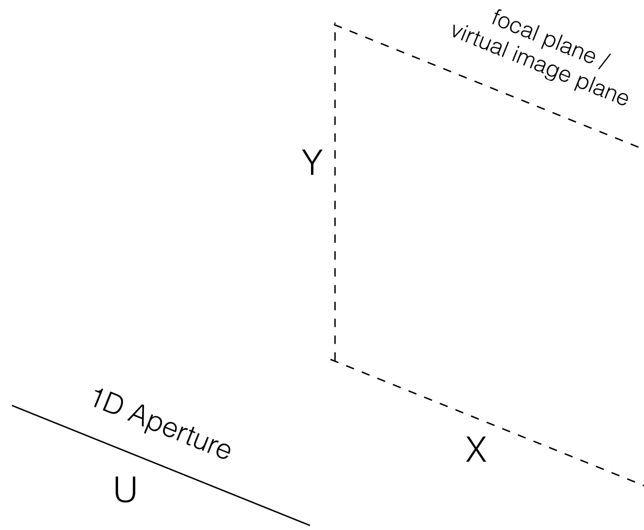

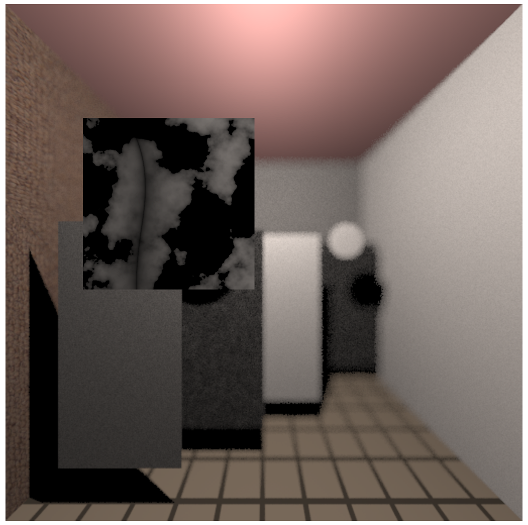

71

… and visualize this pixel…

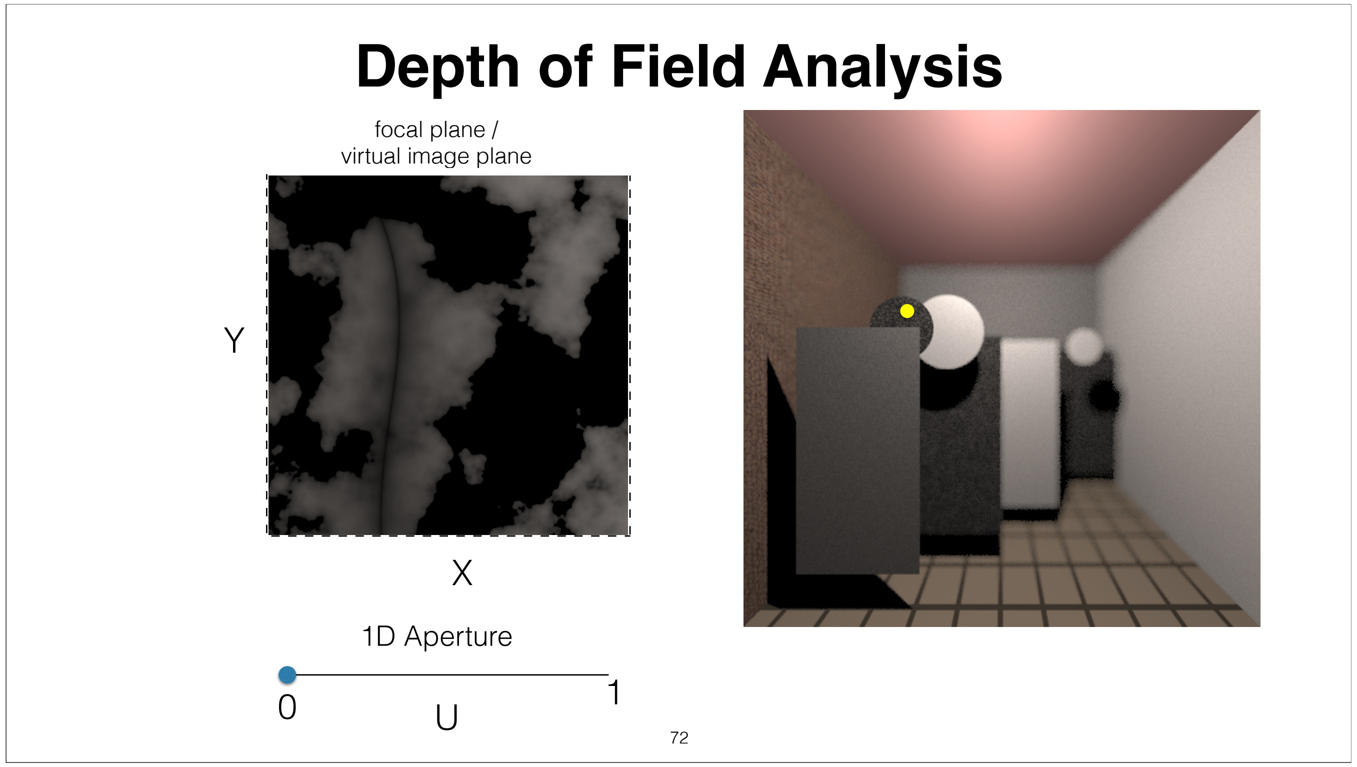

… on this virtual image plane as we move along the aperture. The underlying textures shifts…

Depth of Field Analysis

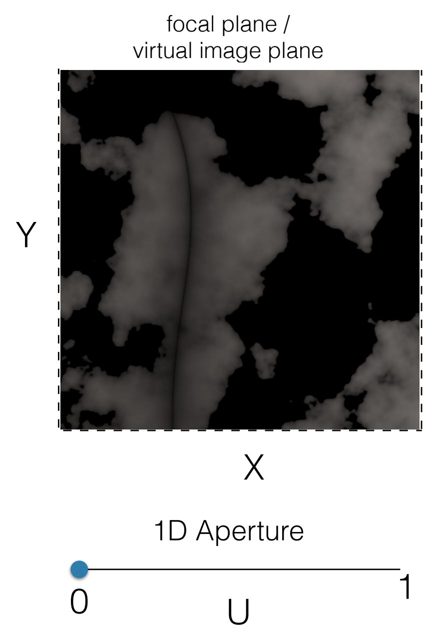

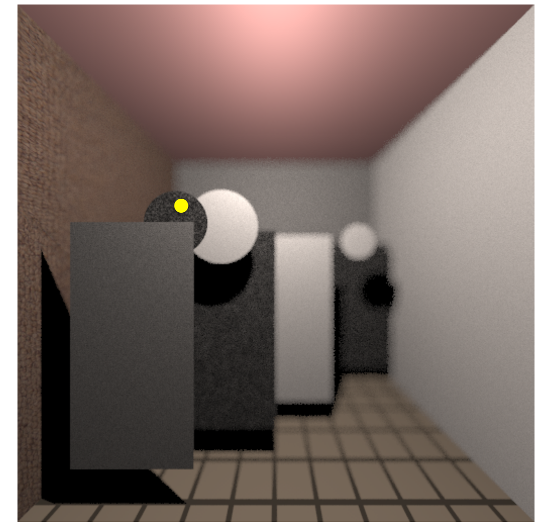

… as we go from left to right on this 1D aperture. This shifting in the XY plane results in…

Depth of Field Analysis

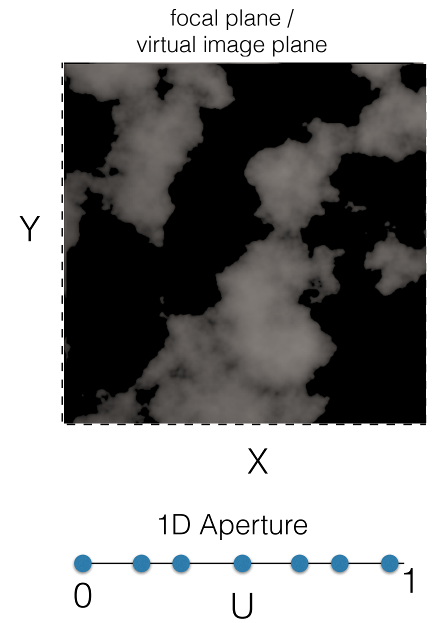

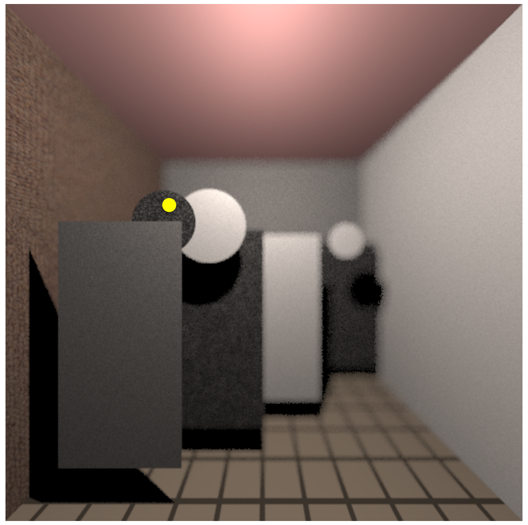

Depth of Field Analysis

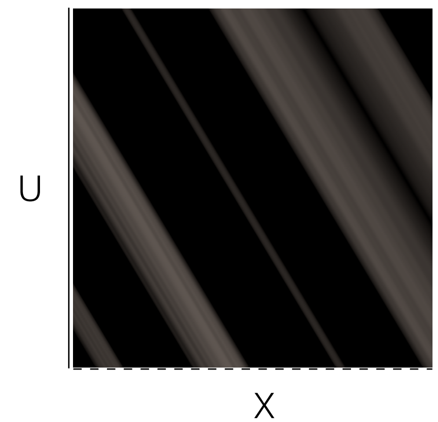

XU Slices

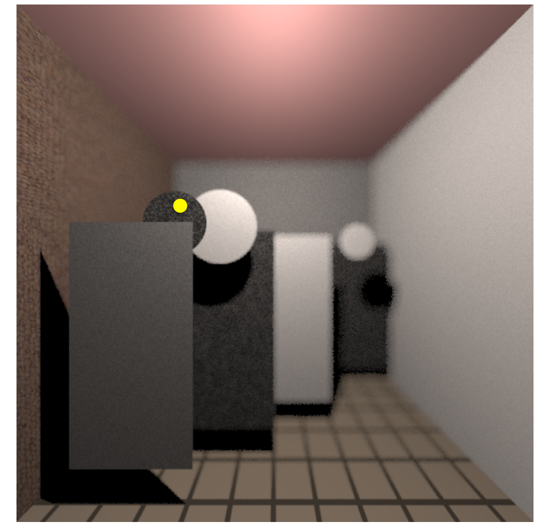

Durand et al. [2005]

…a shear in the XU projection. Note that, if we have an in-focus pixel…

Depth of Field Analysis

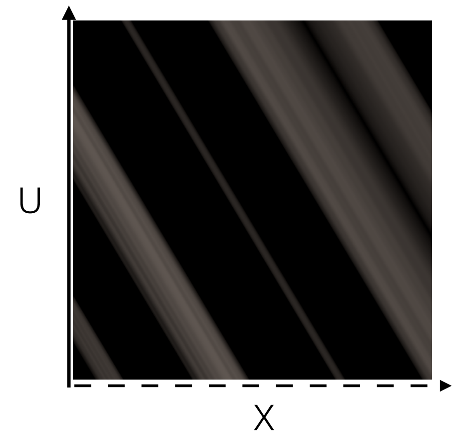

XU Slices

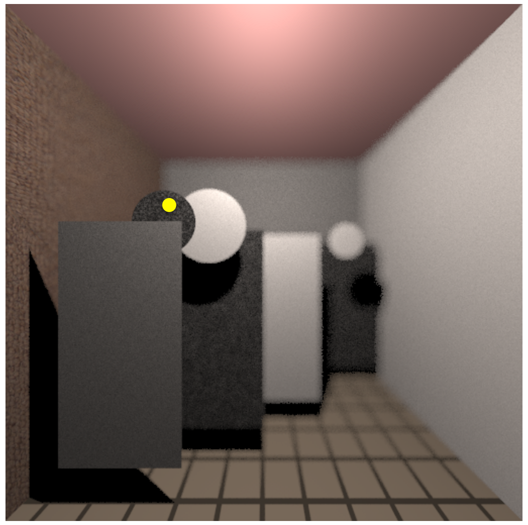

Durand et al. [2005]

…a shear in the XU projection. Note that, if we have an in-focus pixel…

Depth of Field Analysis

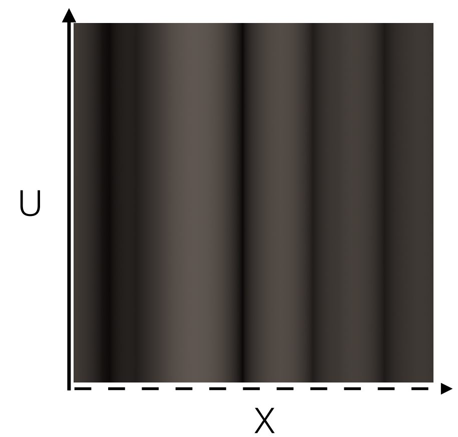

XU Slices

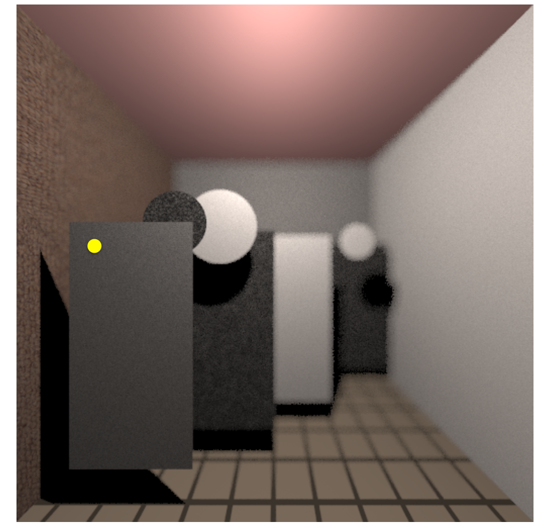

Durand et al. [2005]

…the XU projection will not show any variation along the U-axis. As a result, the corresponding Fourier power spectrum…

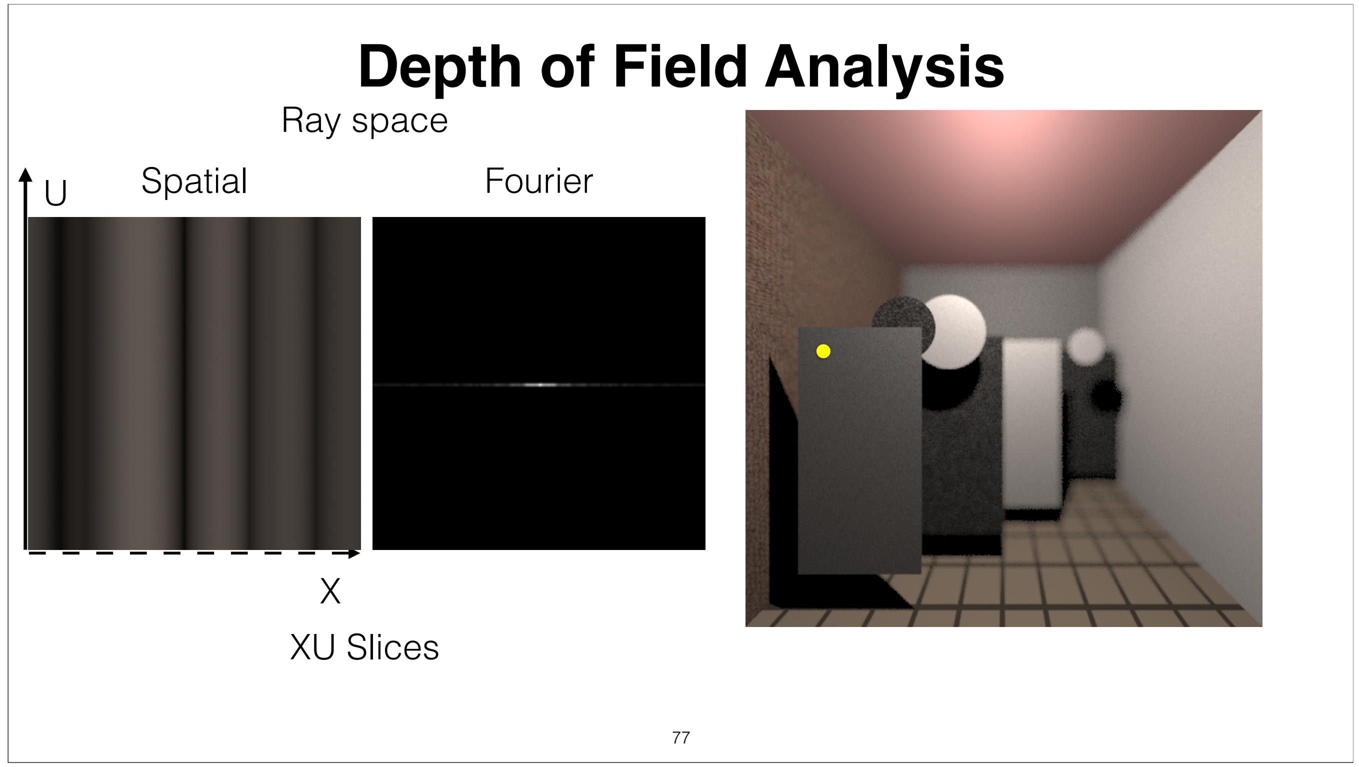

… would have all its energy only along the horizontal axis. But as we go far from the focal plane…

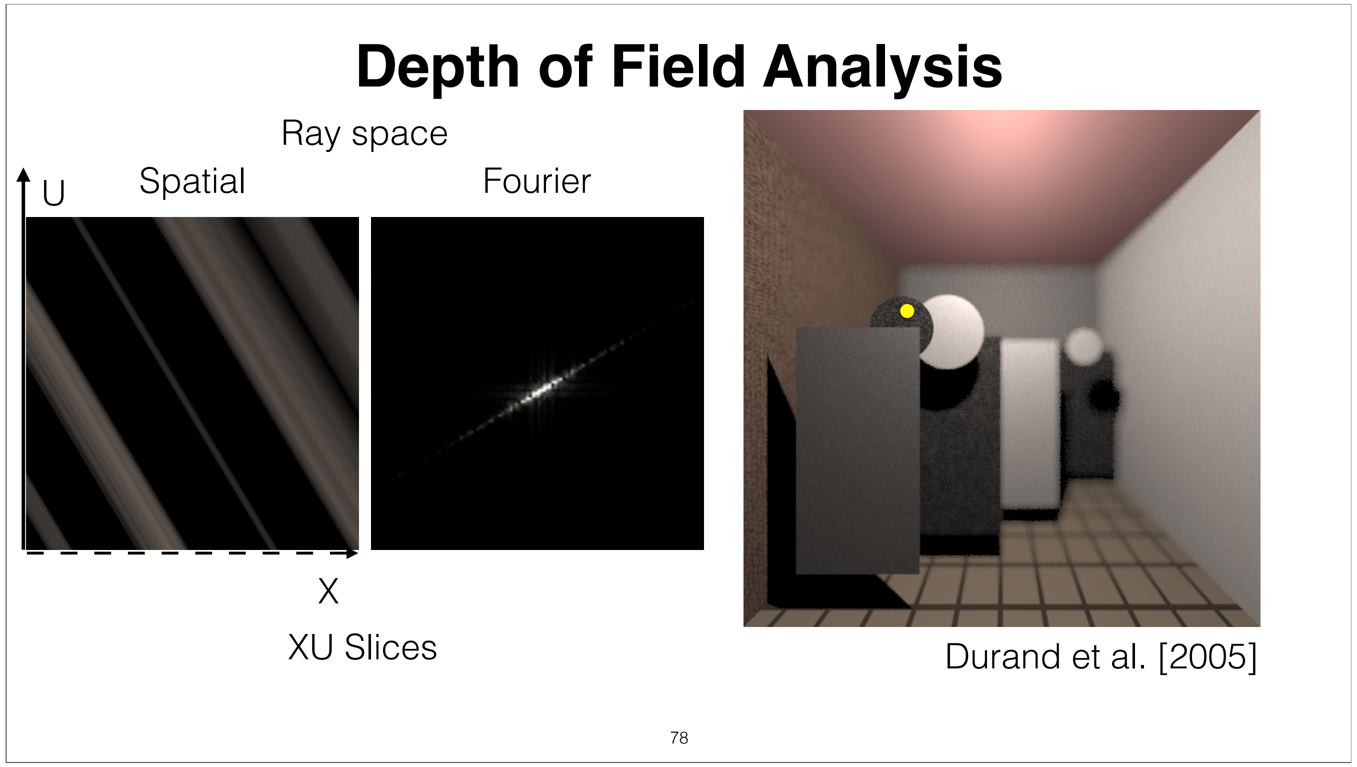

…we observe a shear in the light field.

Depth of Field Analysis

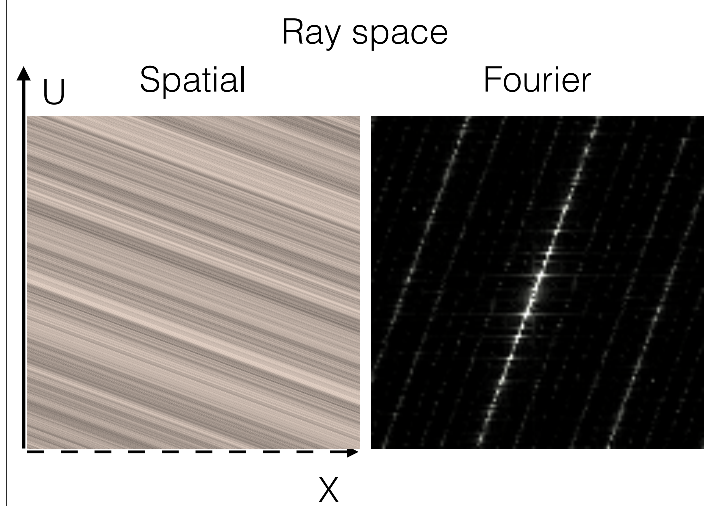

XU Slices

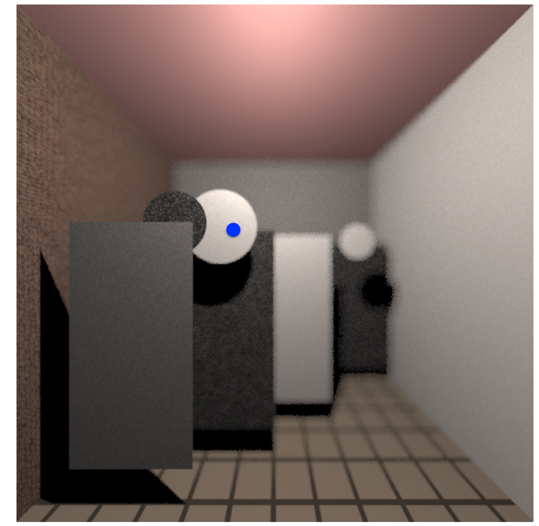

Durand et al. [2005]

…we observe a shear in the light field.

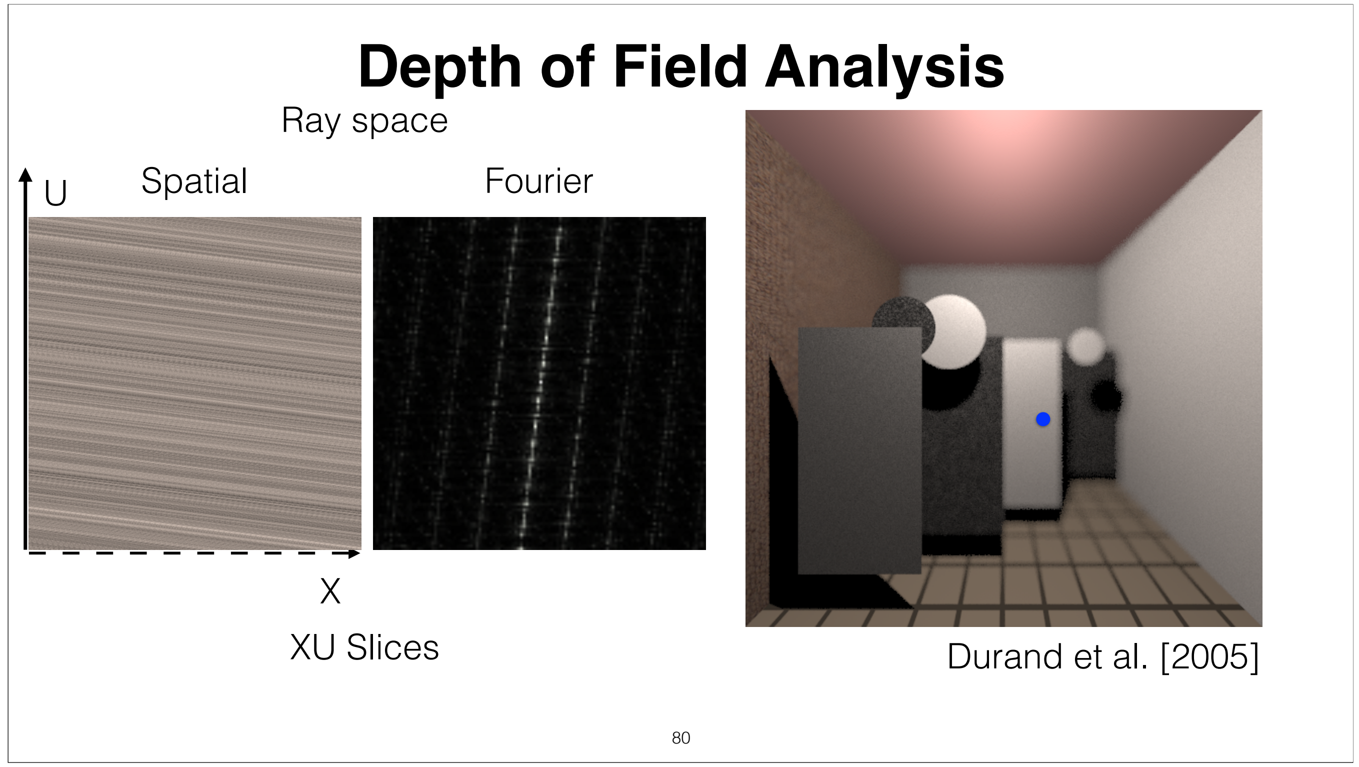

…we observe a shear in the light field.

Depth of Field Analysis

Ray space

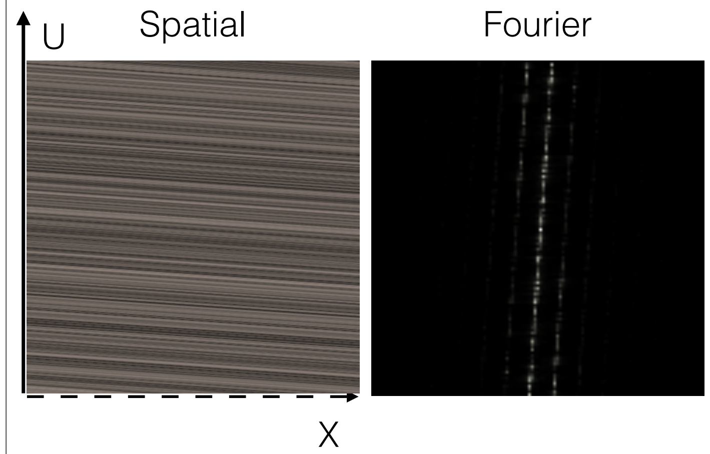

XU Slices

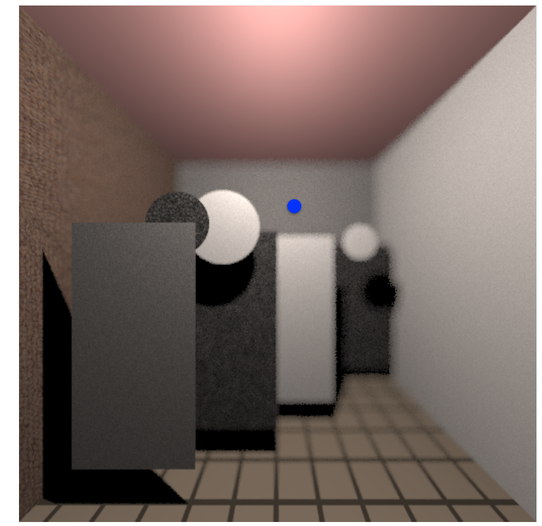

Durand et al. [2005]

This increase…

Light Field gets Sheared

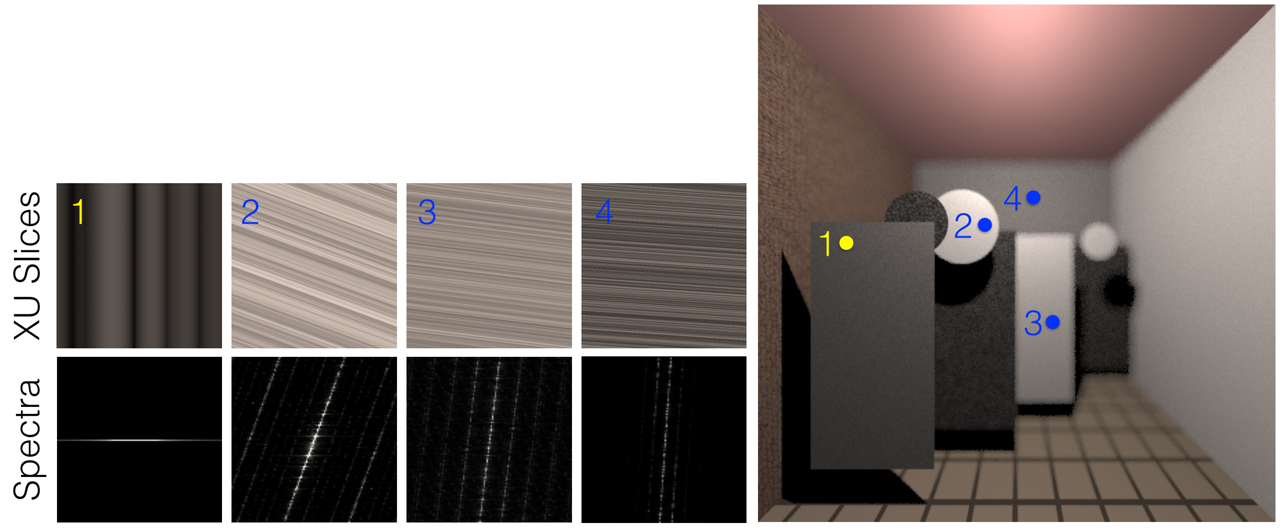

…in shear with objects depth can be easily represented in a mathematical form [CLICK] using the following equation, where [CLICK] F is the focal distance and [CLICK] d is the object depth.
This shows that it is enough to know the depth for each pixel to compute the shear, given we already know the focal distance. Our oracle gives us the depth per pixel which we then use to shear the samples. Note that, given the different projections in this 4D light field…

Light Field gets Sheared

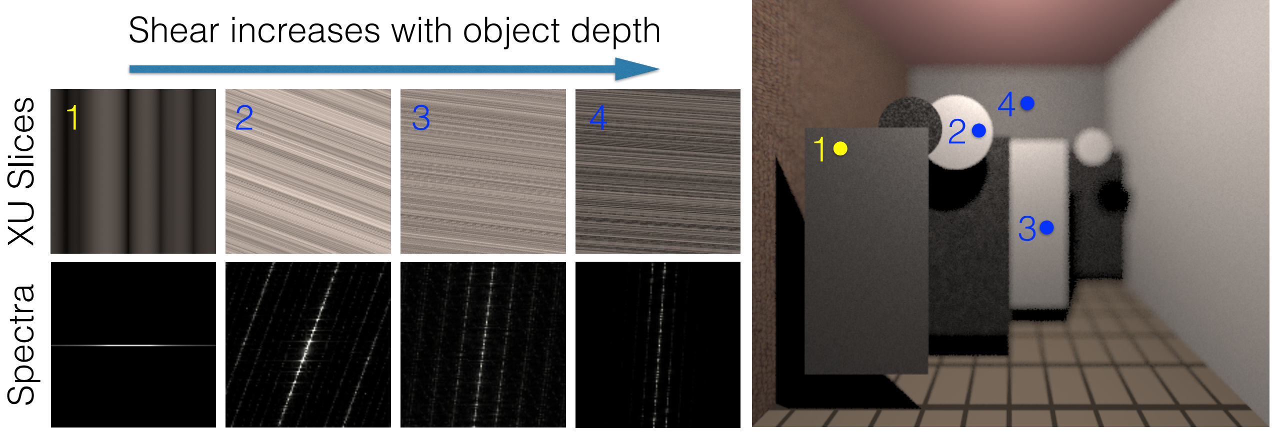

Light Field gets Sheared

$$x = x + u\frac{F-d}{d}\,,$$

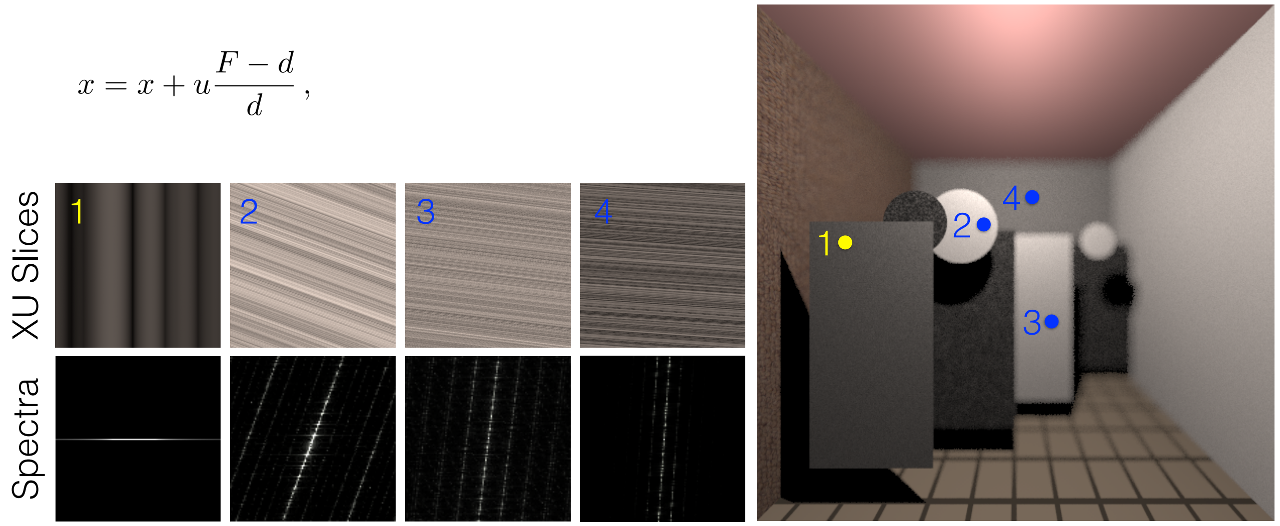

Light Field gets Sheared

$x = x + u\frac{F-d}{d}$, $F:$ focal distance

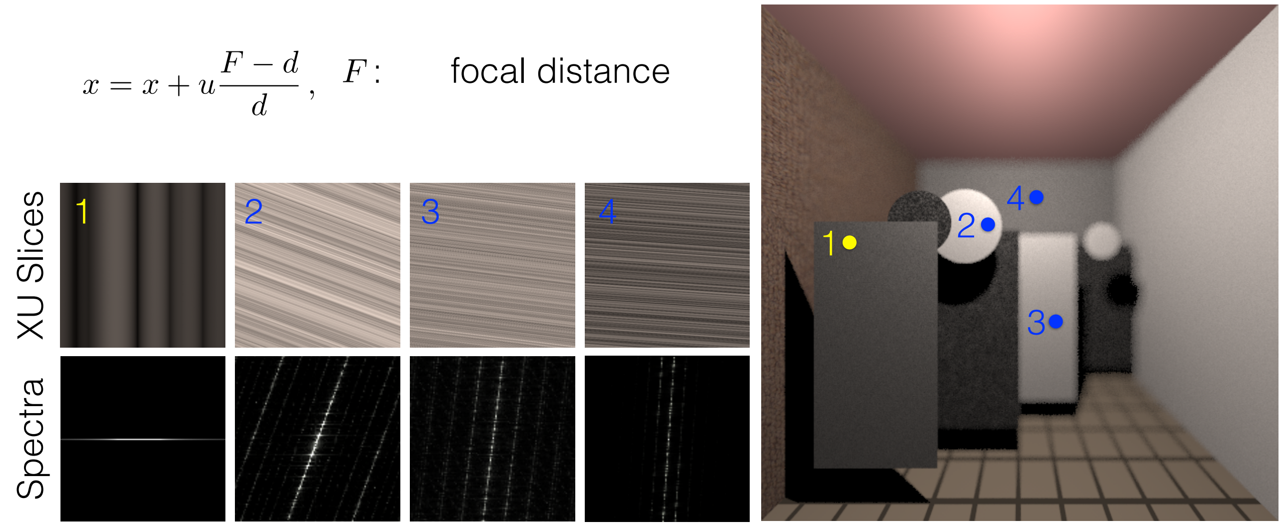

Light Field gets Sheared

$$x = x + u\frac{F-d}{d}\,,$$

$F:$ focal distance

$d:$ depth of the hit object

XU Slices

Spectra

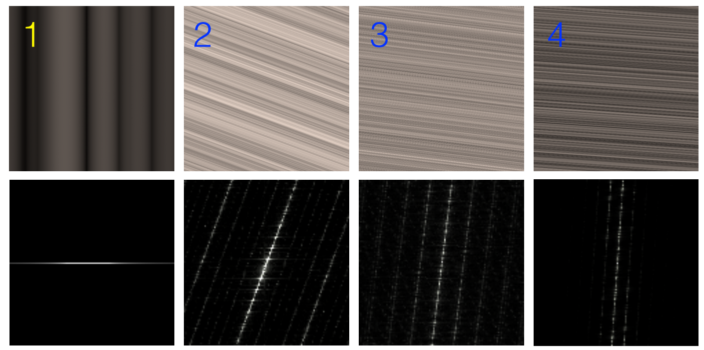

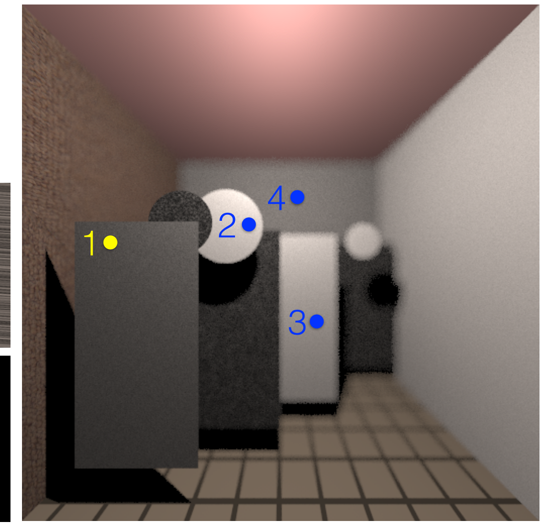

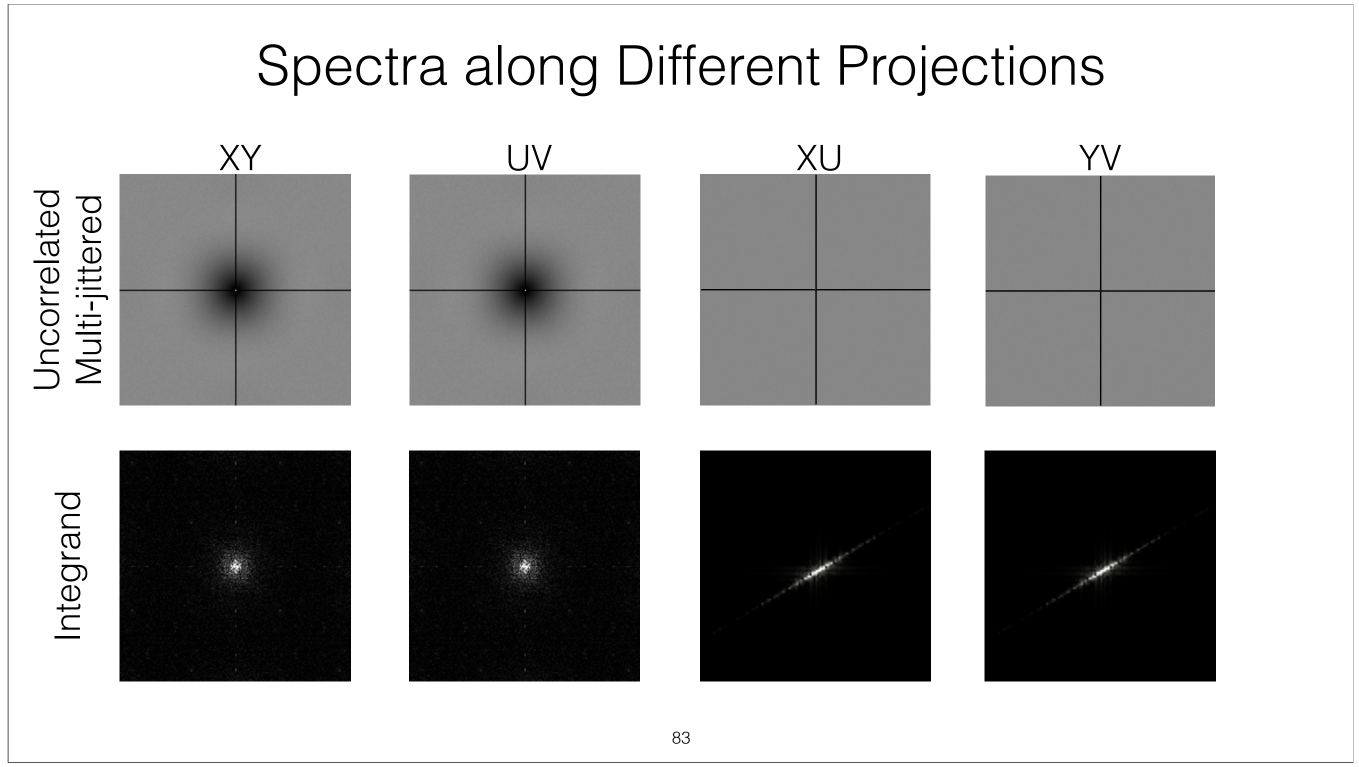

…shearing happens only in the [CLICK] XU and YV projections, as shown in the bottom row. This means, we only need to shear the samples…

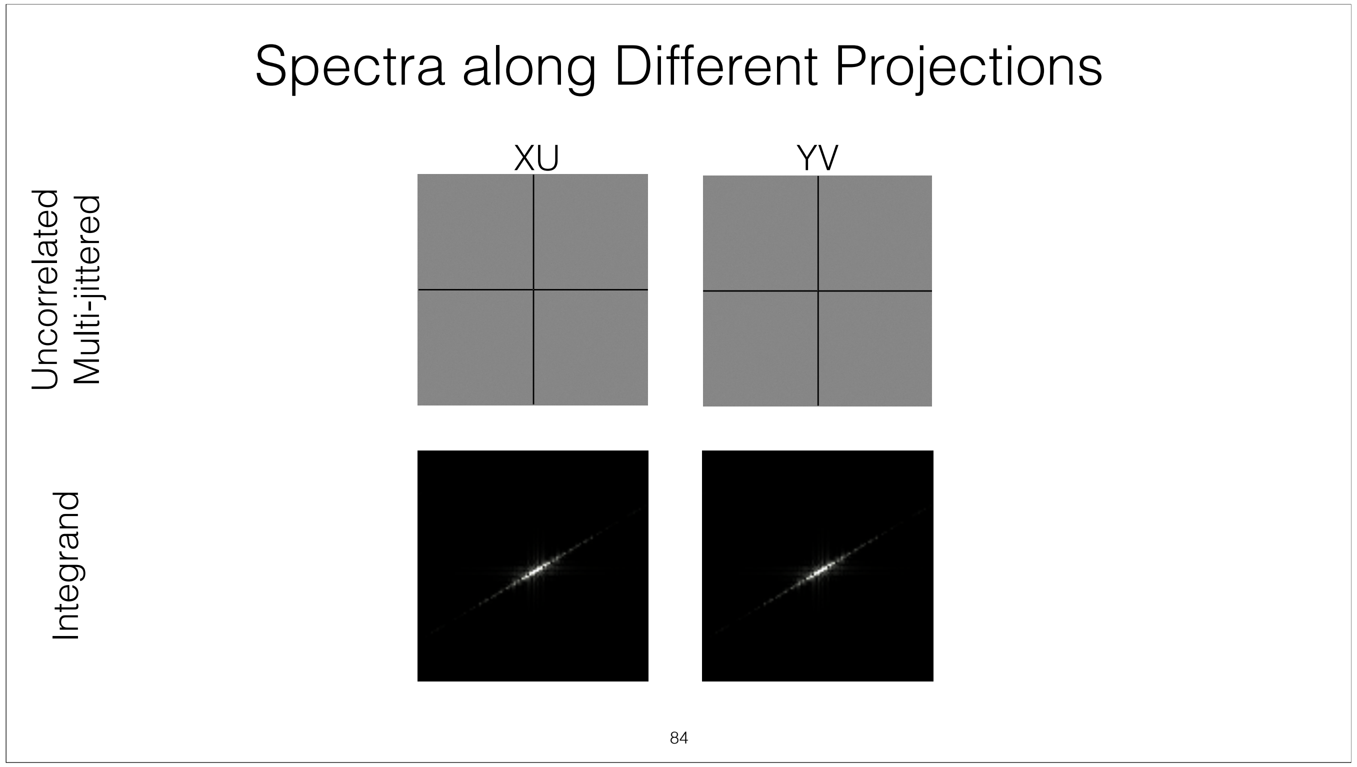

…shearing happens only in the [CLICK] XU and YV projections, as shown in the bottom row. This means, we only need to shear the samples…

Spectra along Different Projections

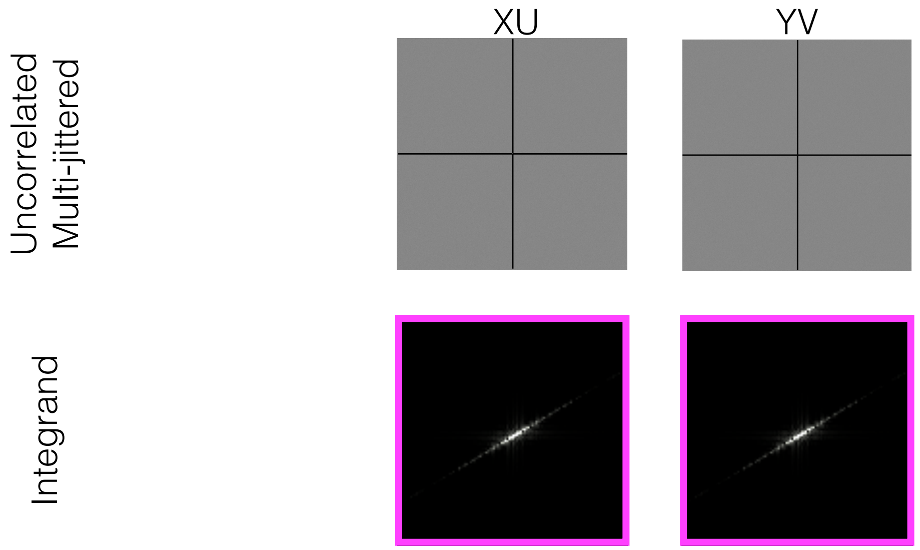

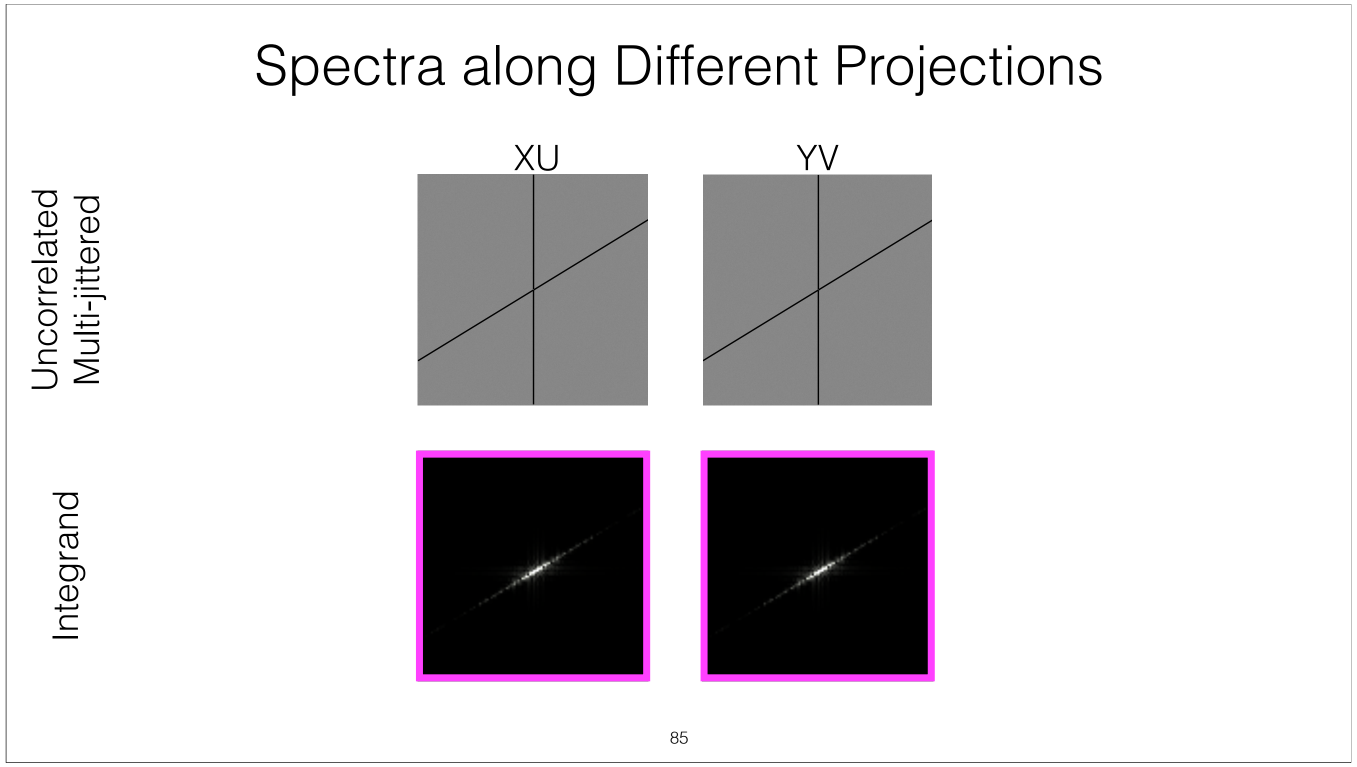

…in the XU and YV projections. Now that we know the shear parameters, we can apply this shear to a real rendering to analyze the benefits.

Variance & Convergence Analysis with Sheared Samples

…variance and convergence rates. We consider the same cornell box scene…

Cornell Box Scene

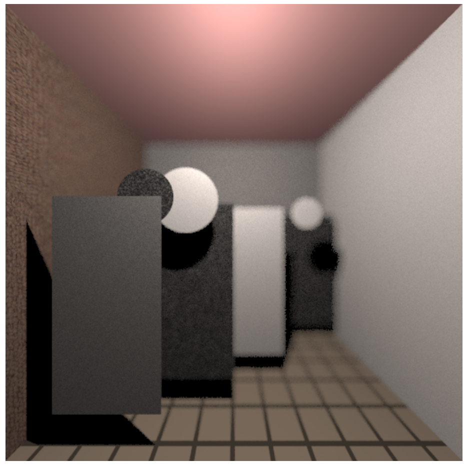

87

…with defocus blur. Since the scene is lit with a point light source, the out of focus pixels, for example, the [CLICK] one on the back wall, will have an underlying 4D function which has the following [CLICK] sheared integrand spectrum in the XU projection. If we sample the underlying 4D function with an uncorrelated-multijittered samples…

Cornell Box Scene

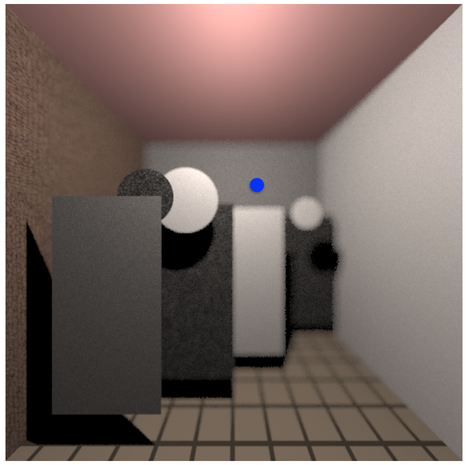

$$\int_x \int_y \int_u \int_v f(x, y, u, v)\; dv\; du\; dy\; dx$$

Cornell Box Scene

XU Projection

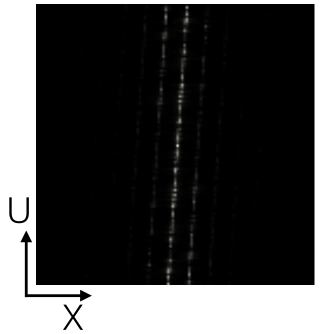

Integrand Spectrum

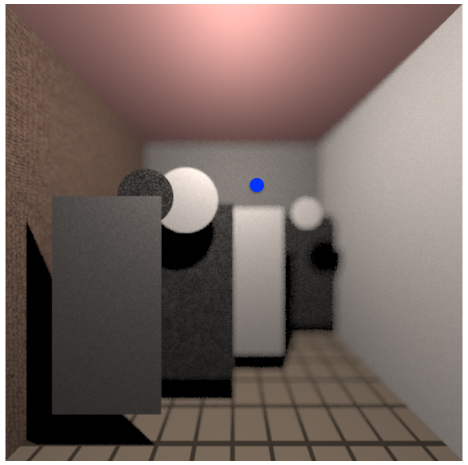

$$\int_x \int_y \int_u \int_v f(x, y, u, v)\, dv\, du\, dy\, dx$$

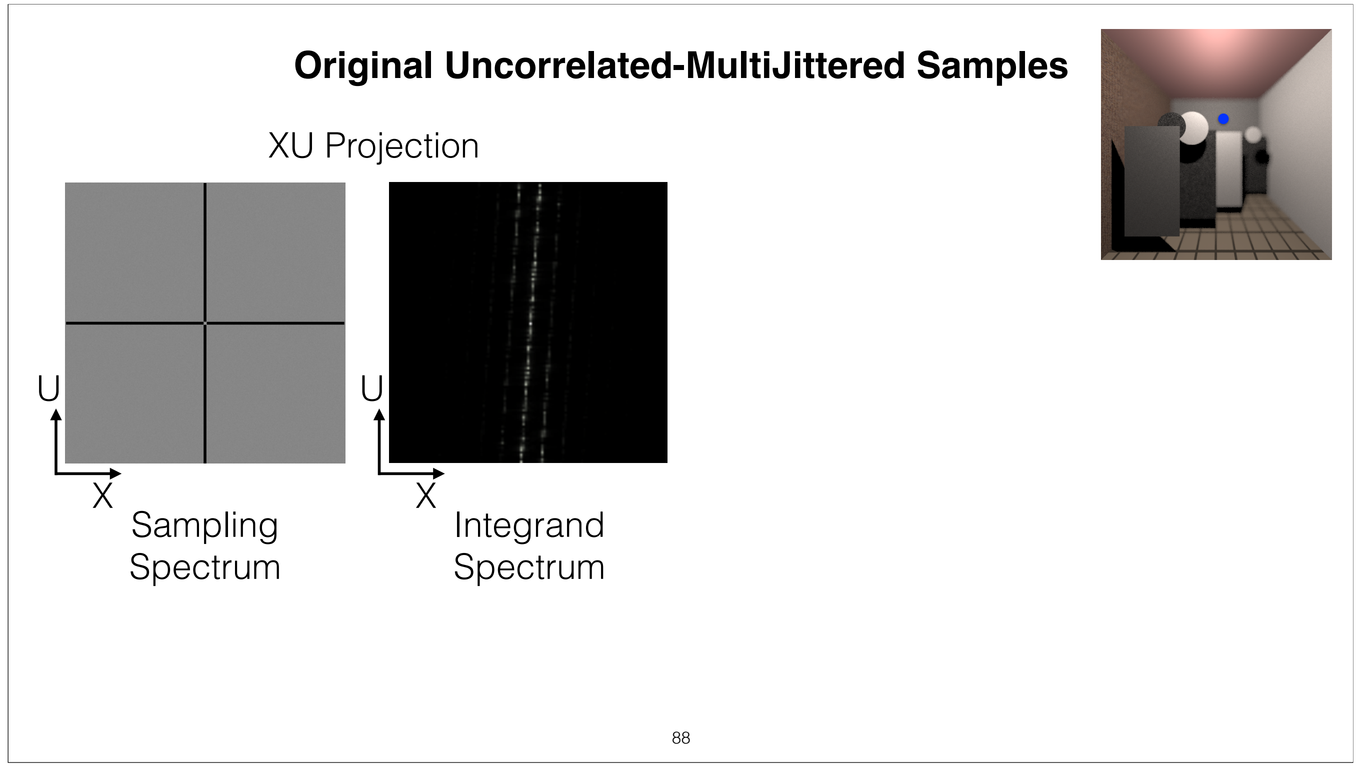

…we get a cross in the XU projection of the sampling spectrum. Since this cross is not aligned with the integrand spectrum, it won't help in variance reduction. Consequently, we see a convergence rate of (N^-1) for this pixel.
However, if we shear the samples…

Original Uncorrelated-MultiJittered Samples

XU Projection

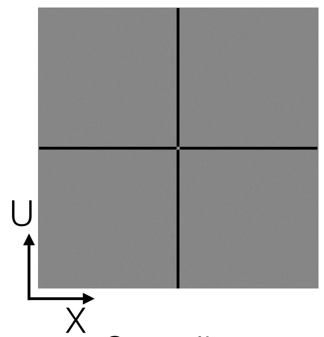

Sampling Spectrum

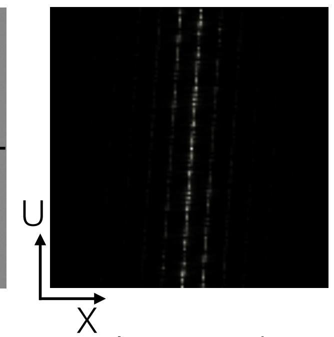

Integrand Spectrum

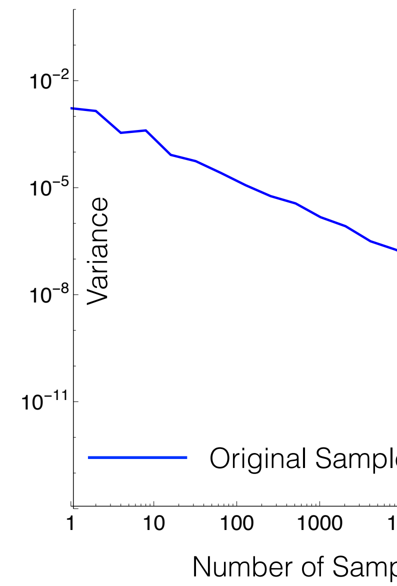

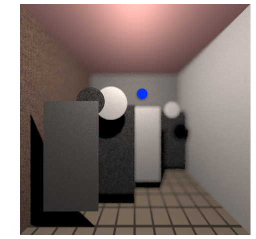

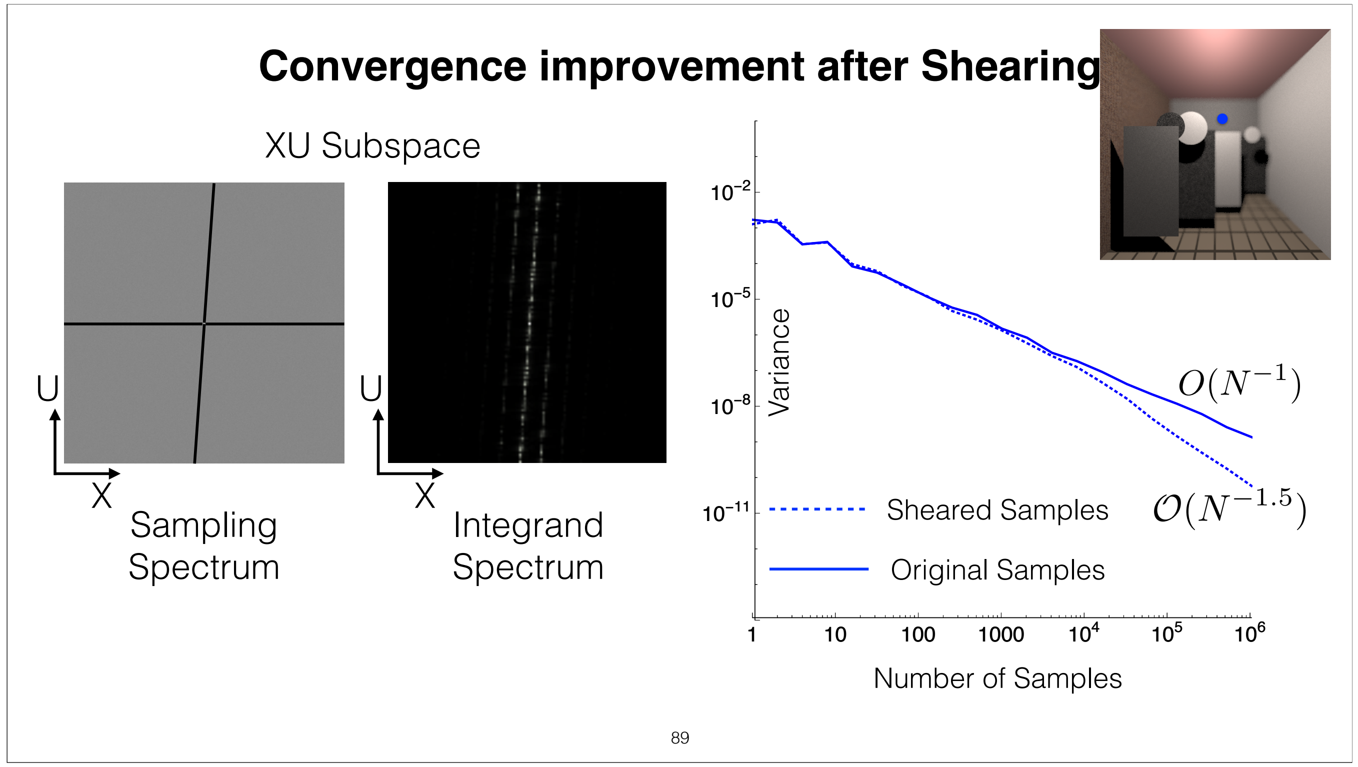

…to align the sampling spectrum with the integrand spectrum, we observe a 2D convergence of (N^-1.5) for this 4D integral. However, this is not true for all pixels. There are some pixels which show…

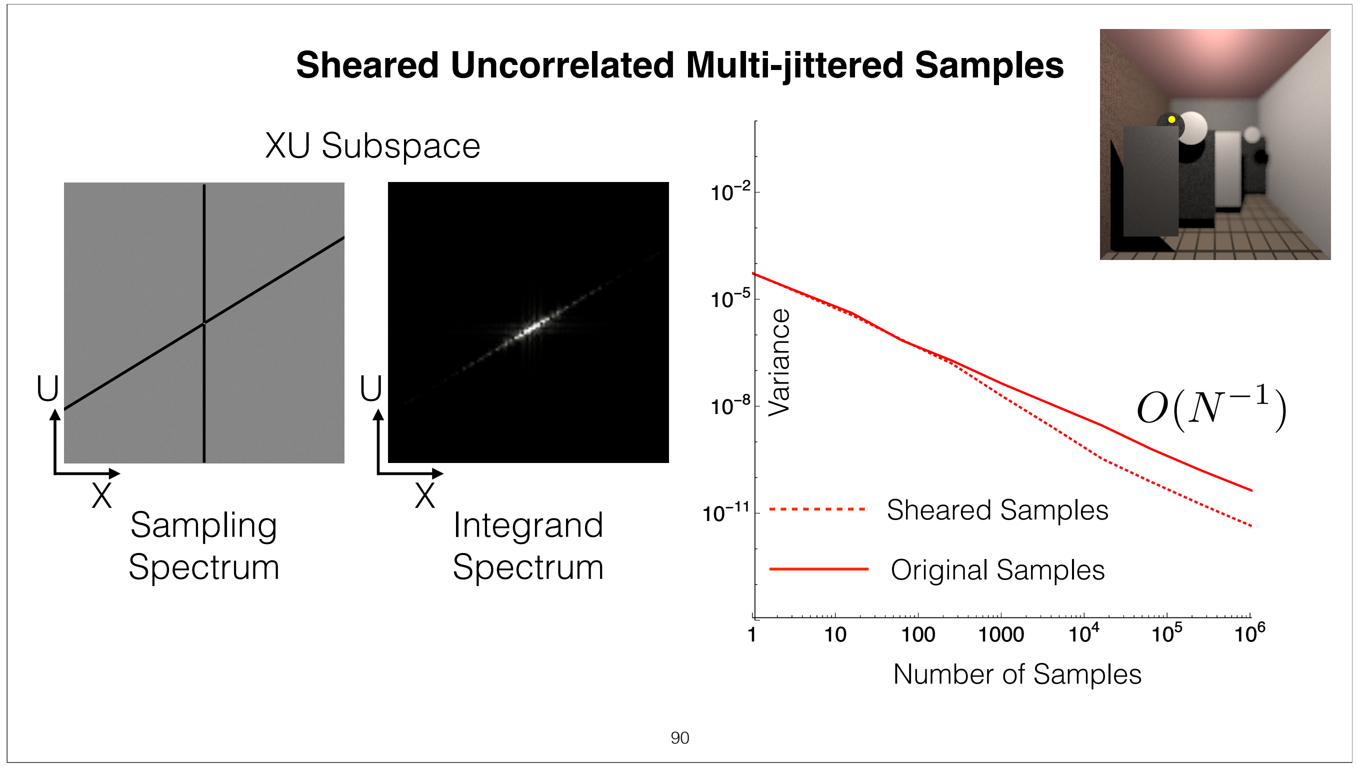

…improvement only in variance after shearing. For example, the following pixel shows 10X improvement in variance but only at a high sample count. Note that, this idea of shearing is not limited to stochastic samples. We apply this idea to deterministic samplers like…

Variance improvement after Shearing

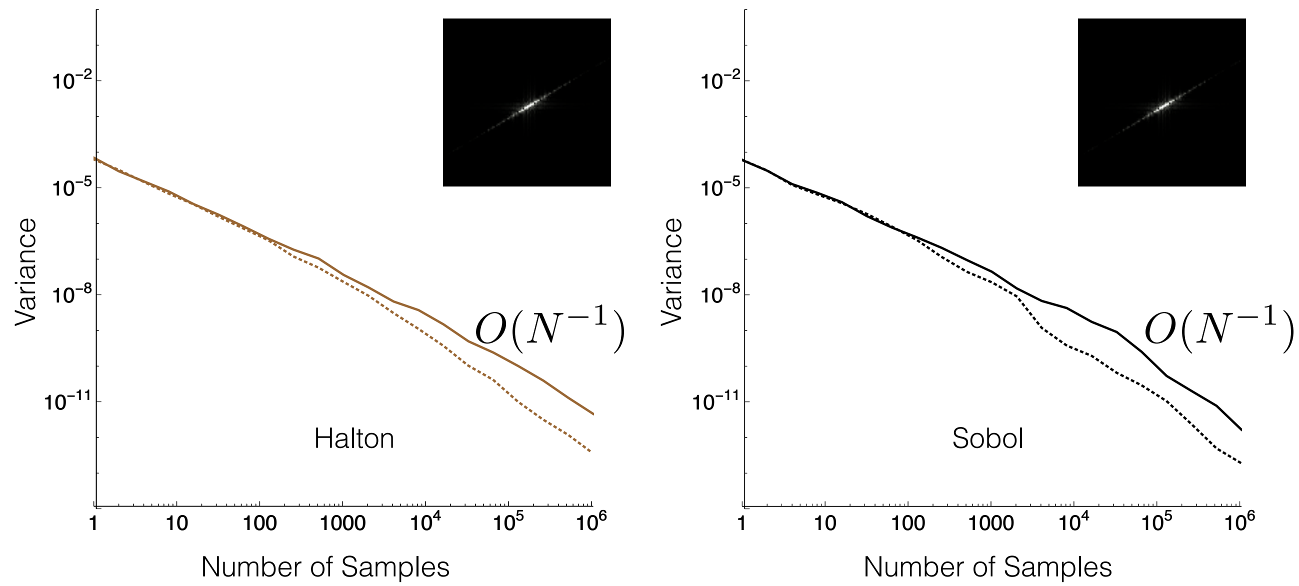

…halton and sobol and observe similar improvements in variance and convergence rates. One thing you may have noticed here is that the improvements are visible after a large sample count. This happens due to the…

Challenging Cases: XU & YV Projections

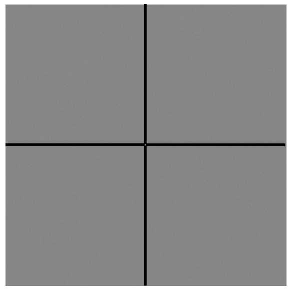

Sampling
XU Spectrum

…these very thin hairline anisotropic structures in the sampling spectra. These hairline anisotropic structures are only useful [CLICK] when we have hairline structures in the integrand spectrum, for which we can shear…

Challenging Cases: XU & YV Projections

Hairline Anisotropy

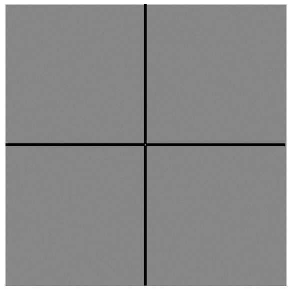

Sampling
XU Spectrum

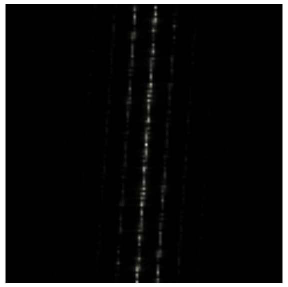

Pixel A
XU Spectrum

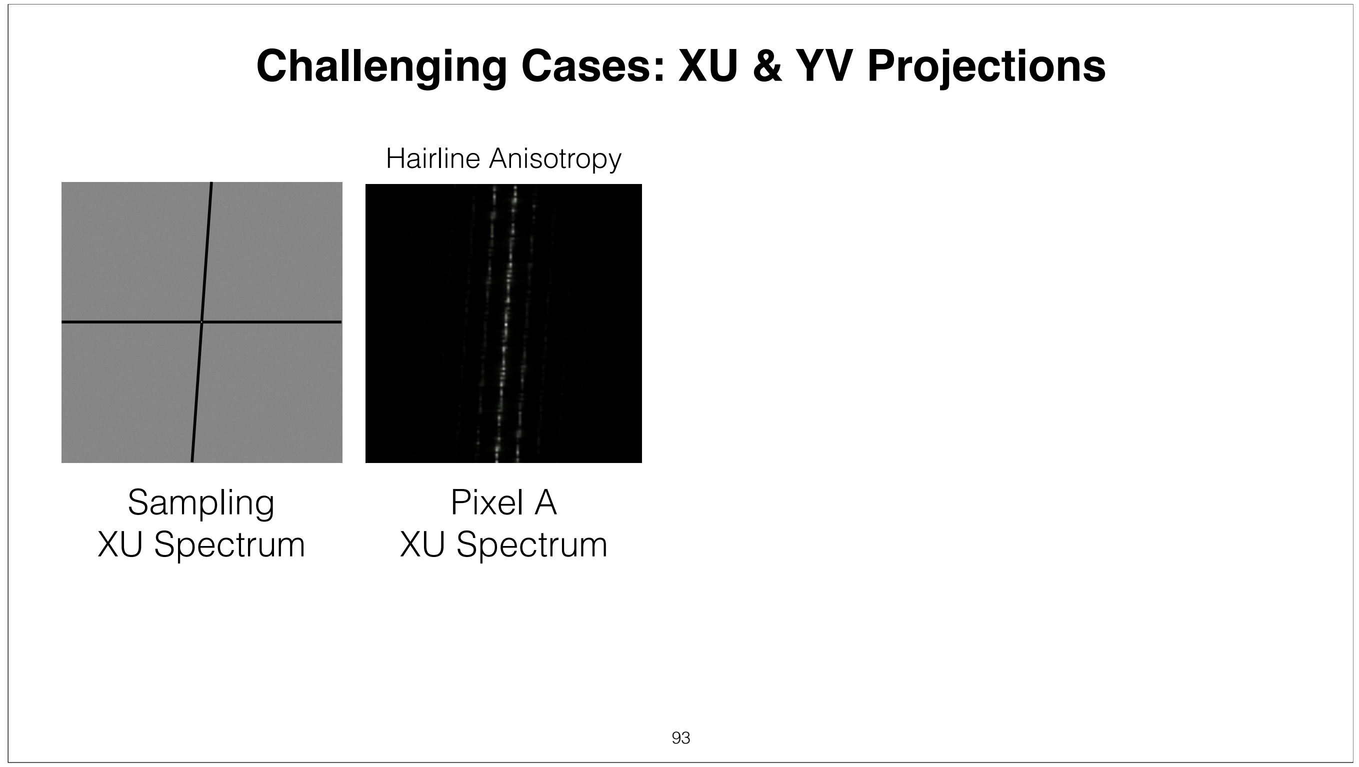

…the sampling spectrum to get benefits. However, even this alignment heavily depends [CLICK] on the accuracy of the oracle that provides the shear parameters. On the other hand, it is also very common to have pixels with a lot of occluders which results in a frequency footprint of a [CLICK] double wedge spectrum. The existing samplers are not even close to handle these double wedge structures [CLICK] due to the presence of only hairline anisotropies.
To handle these cases, our analysis suggests that…

Challenging Cases: XU & YV Projections

Hairline Anisotropy

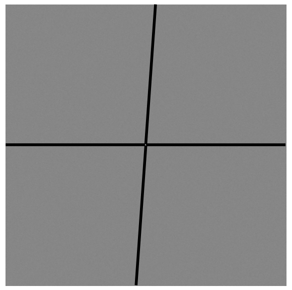

Sampling
XU Spectrum

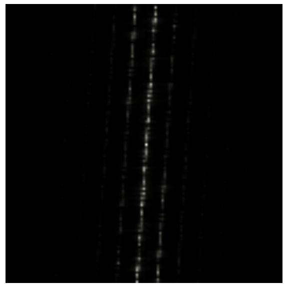

Pixel A
XU Spectrum

Oracle Accuracy

Challenging Cases: XU & YV Projections

Hairline Anisotropy

Double-wedge Anisotropy

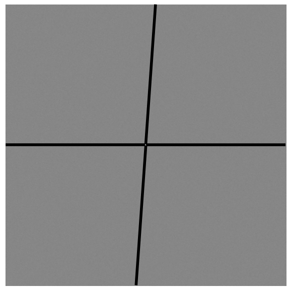

Sampling
XU Spectrum

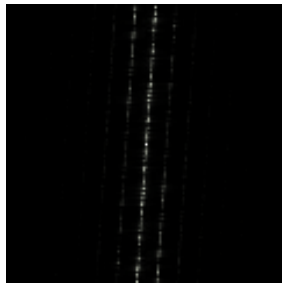

Pixel A
XU Spectrum

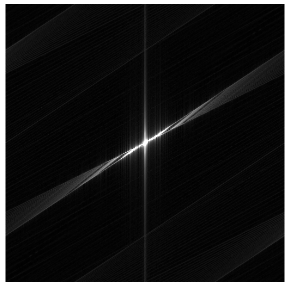

Pixel B
XU Spectrum

Oracle Accuracy

Challenging Cases: XU & YV Projections

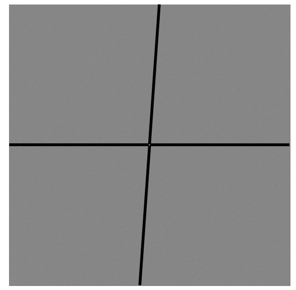

Sampling
XU Spectrum

Hairline Anisotropy

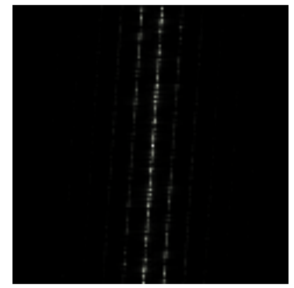

Pixel A
XU Spectrum

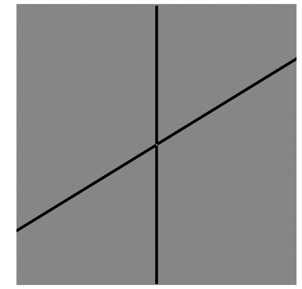

Sampling
XU Spectrum

Double-wedge Anisotropy

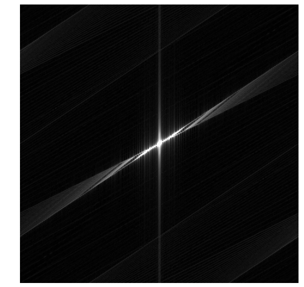

Pixel B
XU Spectrum

Oracle Accuracy

Double-wedge Spectrum

Design Principles for New Sampling Patterns

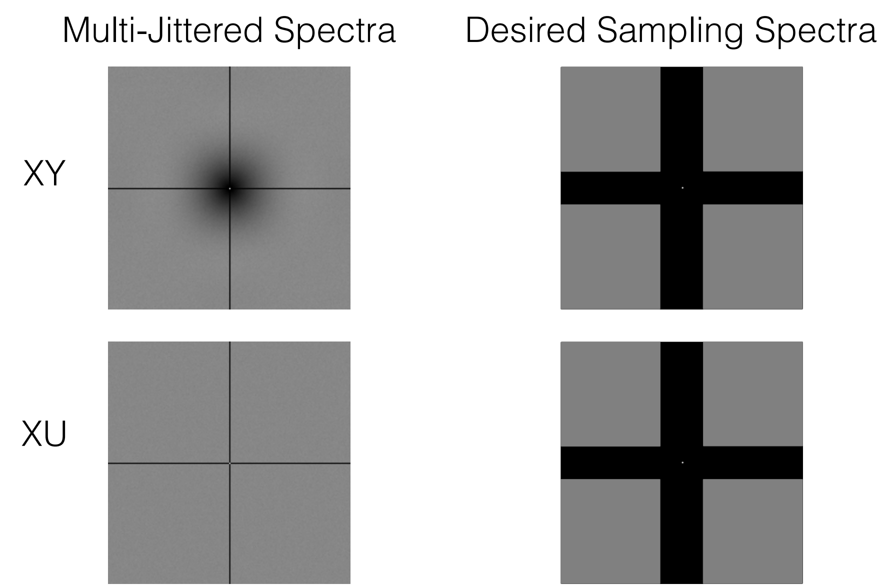

…where a desired sampling spectrum can have wider anisotropic structures in all the projections, that can greatly improve convergence rates and can also promise huge variance reduction. For pixels with a double-wedge…

Design Principles for New Sampling Patterns

Integrand Spectrum

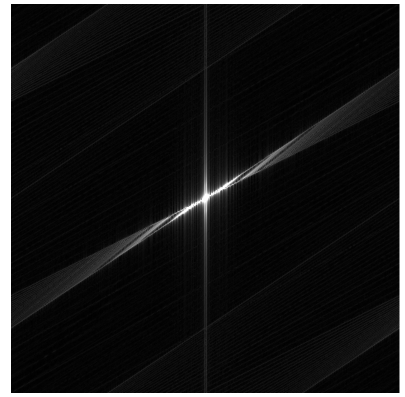

Desired Sampling Spectra

In both XU and YV Projections

95

… shaped spectrum, it could also be interesting to create sampling patterns that matches the wedge shaped target spectrum. With this I would like to conclude my talk and I will be happy to take any questions you may have. Thank you..

Thank you for your attention!