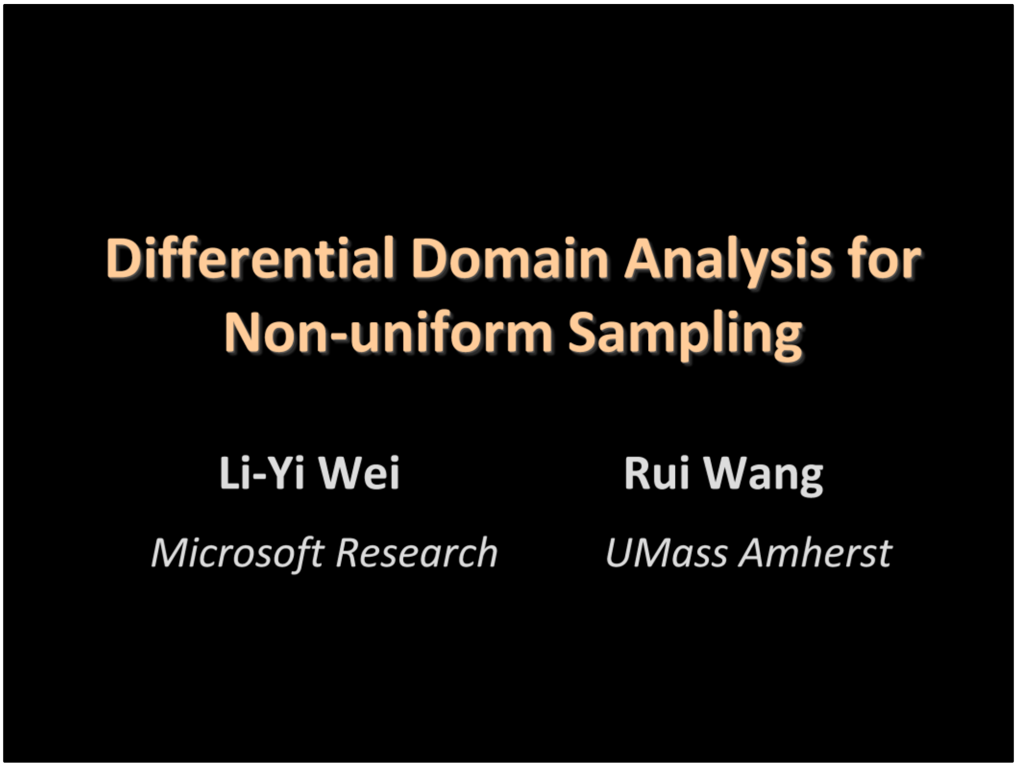

This paper is about a new method for analyzing the distribution property of non-uniform stochastic samples.

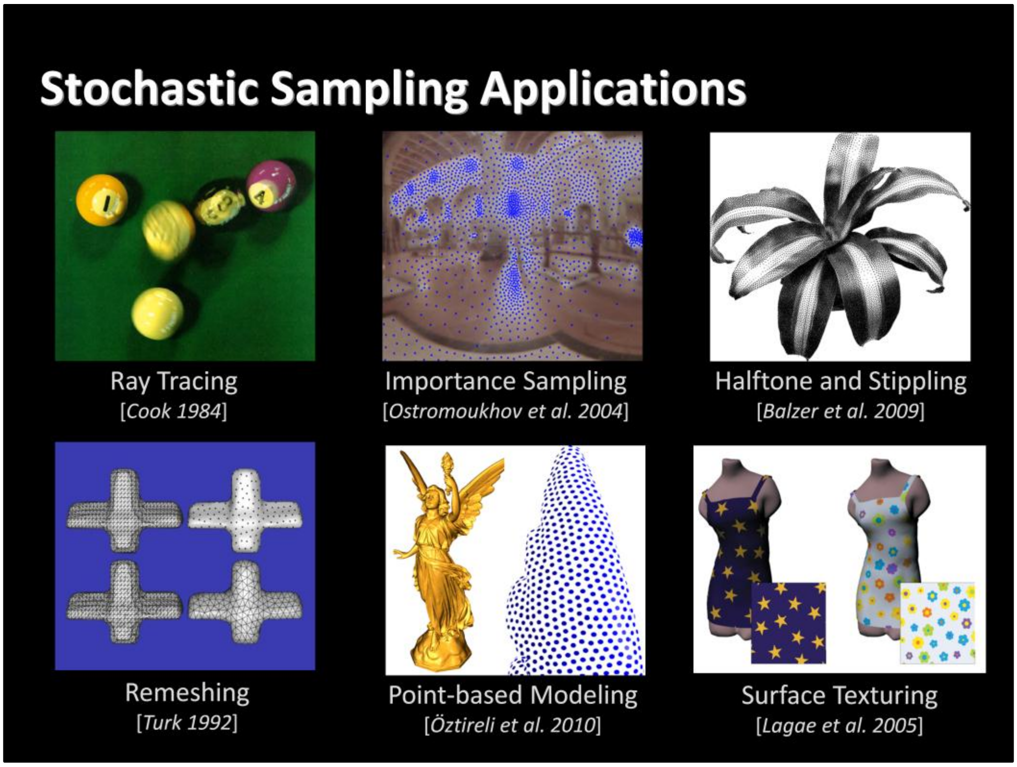

Stochastic sampling is a fundamental component in many graphics applications. Examples include rendering and imaging applications, such as ray tracing, importance sampling, half tone and stippling, as well as ■ geometry applications, such as remeshing, point-based modeling, and surface texturing. In most of these applications, the distribution property of samples can directly influence the result quality and accuracy. So it must be carefully evaluated.

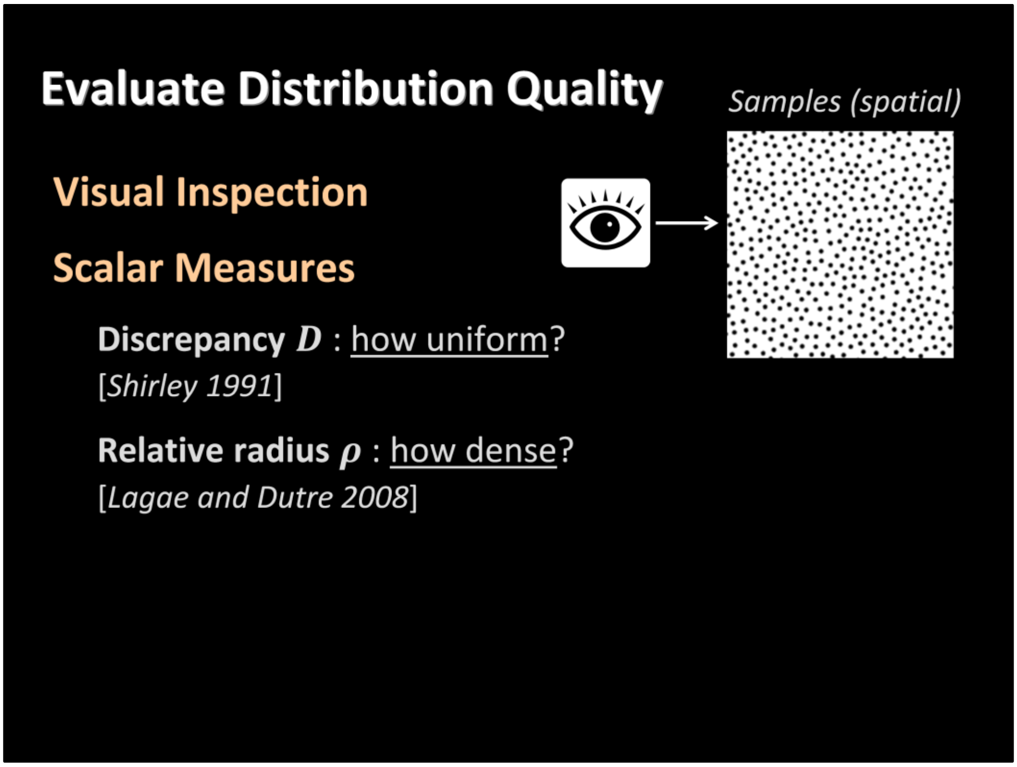

To start, we can visually inspect the samples, but this is too subjective and only suitable for qualitative inspection.

For ■ quantitative inspection, we can compute the discrepancy and the relative radius of the samples. These provide scalar measures to help us understand how uniform or dense the samples are globally.

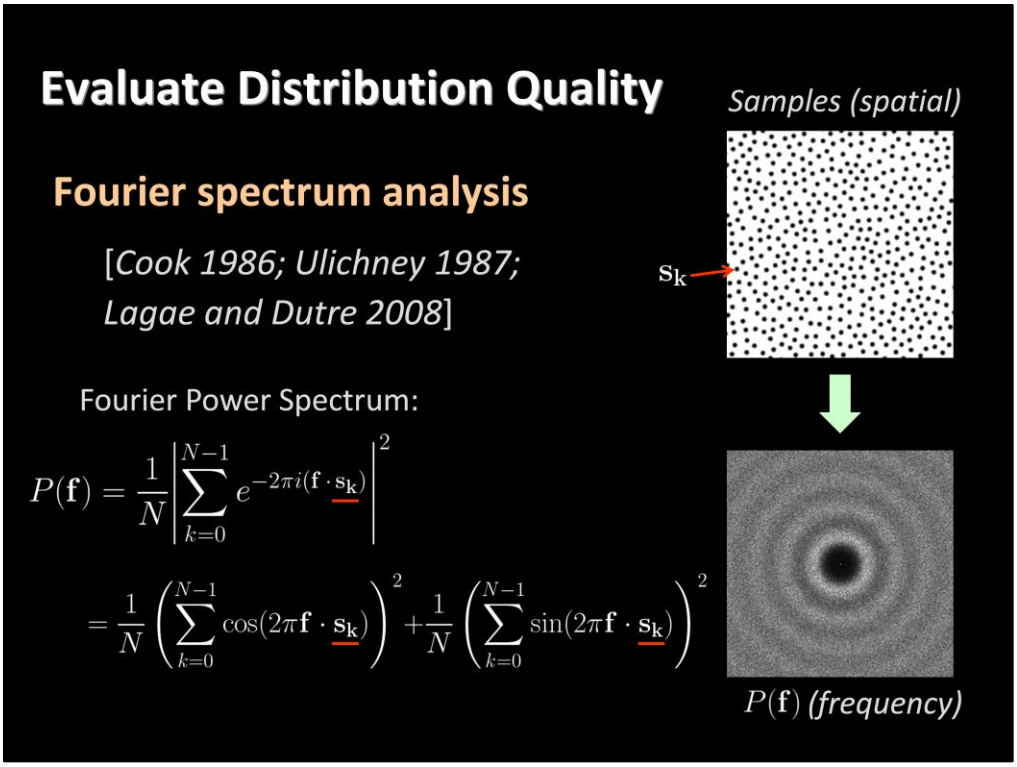

However, a more detailed analysis is typically done by inspecting the Fourier power spectrum, ■ which transforms the samples from their spatial domain to the frequency domain. This can be ■ computed by plugging in the sample locations s_k into this equation which measures the squared magnitude of the Fourier transform coefficients. The result can be plotted as a 2D spectrum image shown on the right.

Latex:

P(\mathbf{f})=\frac{1}{N}{\left| \sum_{k=0}^{N-1} e^{-2 \pi i (\mathbf{f} \, \cdot \,\mathbf{s_k})} \right|}^2

=\frac{1}{N}\left(\sum_{k=0}^{N-1} \cos(2 \pi \mathbf{f}\cdot\mathbf{s_k}) \right)^2 + \frac{1}{N}\left(\sum_{k=0}^{N-1} \sin(2 \pi \mathbf{f}\cdot\mathbf{s_k}) \right)^2

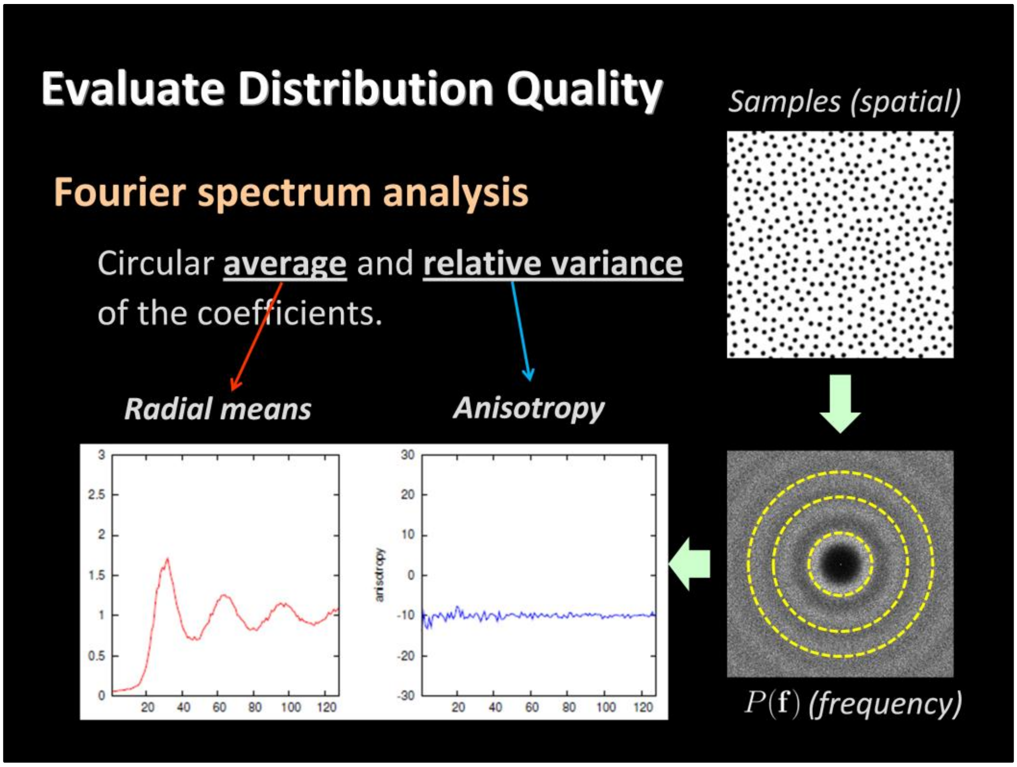

From this spectrum image, we can deduce a lot of useful information. For example, we can compute ■ the average and relative variance of the coefficients on each concentric circle. This results in two 1D graphs, which are called the radial means and the anisotropy. Previous work has shown that these graphs are very useful at predicting samples' behavior for suppressing aliasing artifacts, for improving numerical integration accuracy, and for producing visually pleasing spatial patterns. As a result, Fourier power spectrum has been a widely adopted standard for stochastic sample analysis.

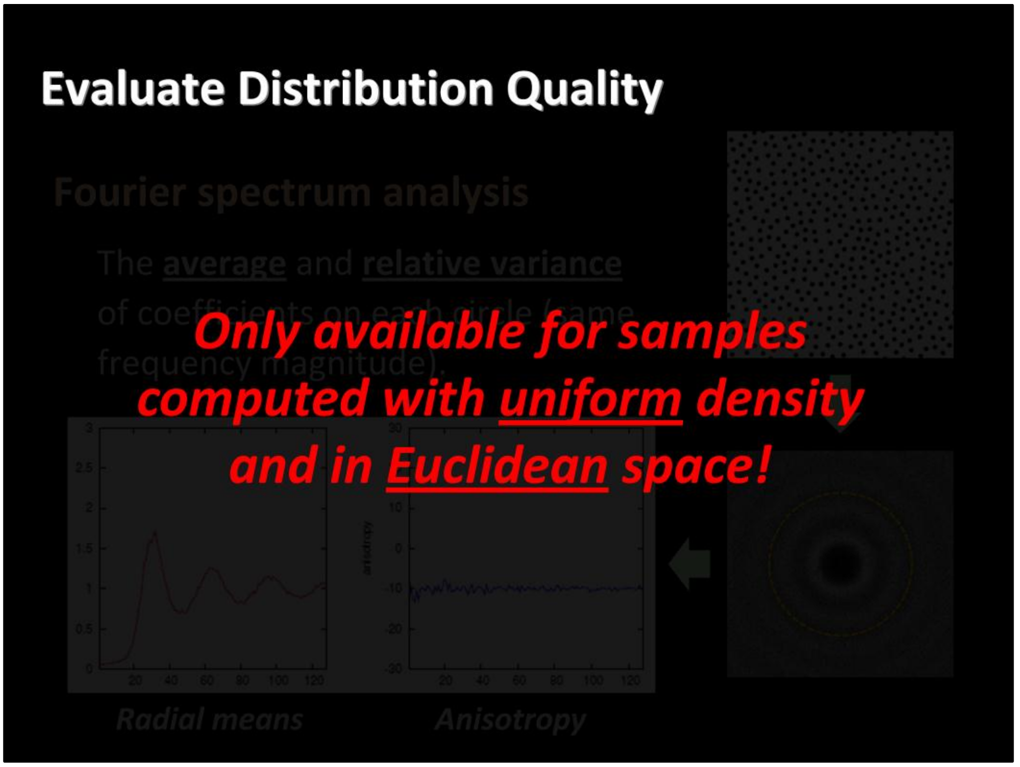

Unfortunately, this method is only available for samples computed with uniform density and in Euclidean space, because this is where the classic Fourier transform is defined.

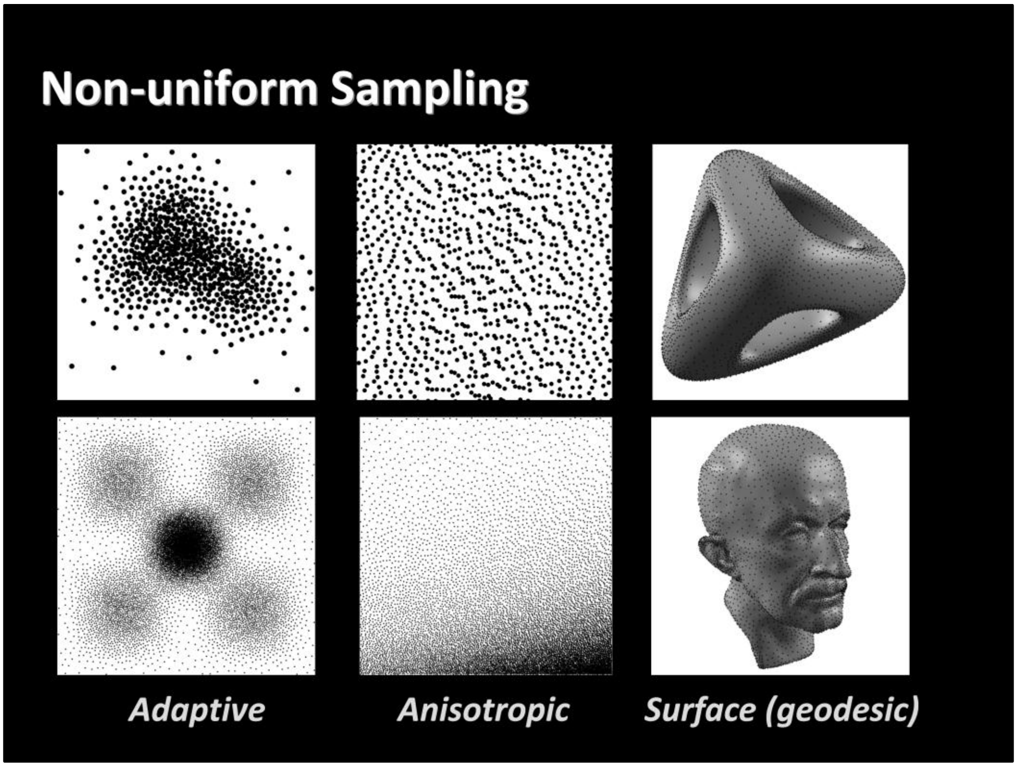

But there are lots of applications, such as adaptive, anisotropic, and surface samplings, where the samples must be computed with non-uniform density or in non-Euclidean spaces. In order to extend classic Fourier analysis to these cases, we would have to define and numerically compute a special Fourier basis for each case. This is a highly non-trivial task, and even when it could be done, it would be expensive and numerically unstable.

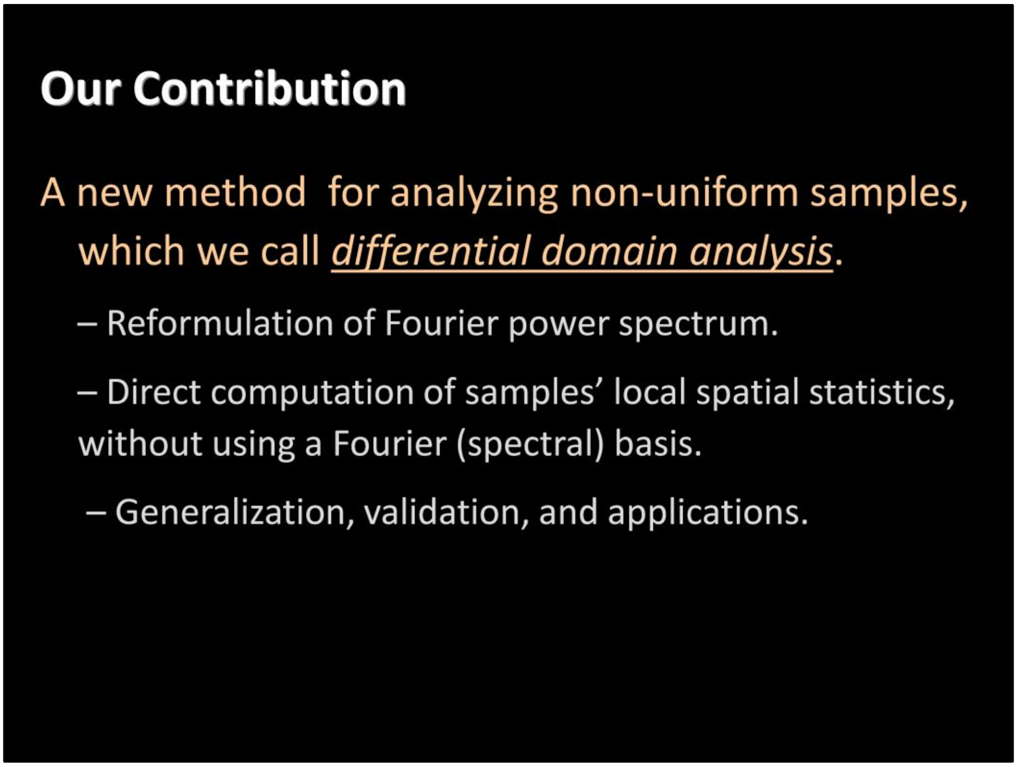

Our main contribution in this paper is a new method for analyzing non-uniform samples. We call this method the differential domain analysis. It is based on a reformulation of the Fourier power spectrum, and involves only computing the samples' local spatial statistics. It produces results equivalent to Fourier power spectrum, but without using a Fourier basis. Therefore we can easily generalize this method to non-uniform domains, as well as providing high computation efficiency.

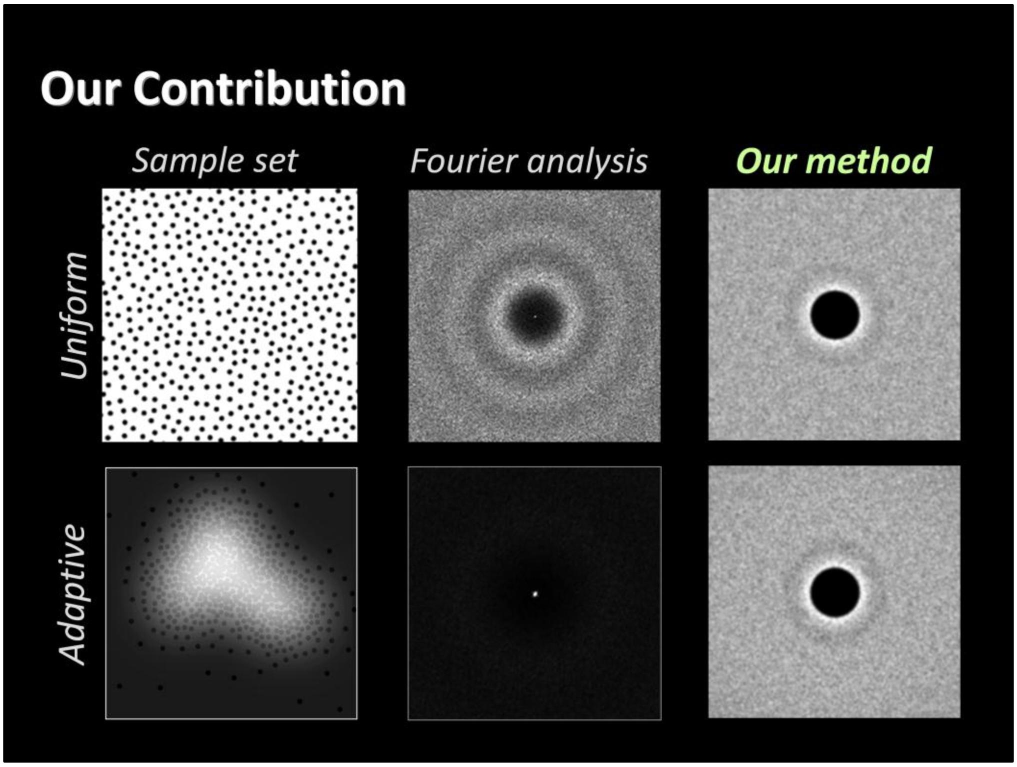

Here is an example that shows the ability of our method. First, given a uniform sample set, we can apply classic Fourier analysis, and the spectrum image tells us that the samples have a blue noise distribution. However, if ■ the samples are non-uniform, such as computed from a spatially varying density function, the Fourier analysis will produce nothing meaningful, because each sample's local distance metric is different. In contrast, ■ our method can successfully account for the spatially varying density, and reveals that the two sample sets have similar blue noise distribution.

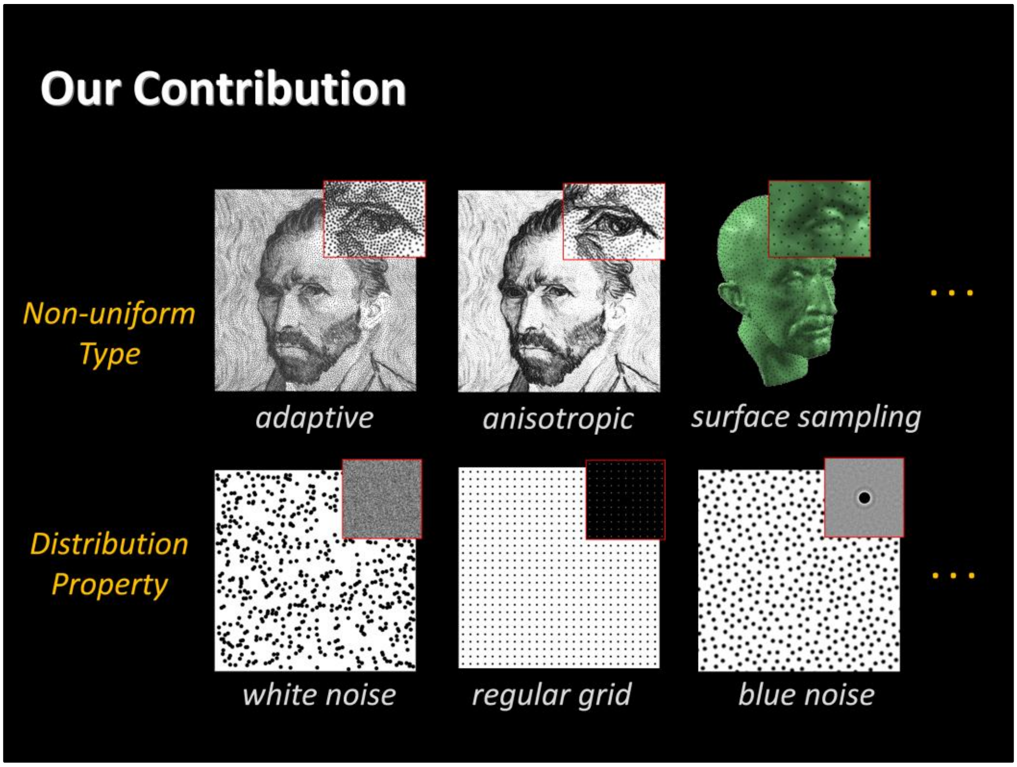

Using this method, we are able to analyze a variety of different non-uniform sampling types, including adaptive, anisotropic, surface, or a combination of them. ■ Our results can be used to capture and infer the samples' intrinsic distribution properties, whether they are white noise, regular grid, blue noise, or any other type.

4:30 here

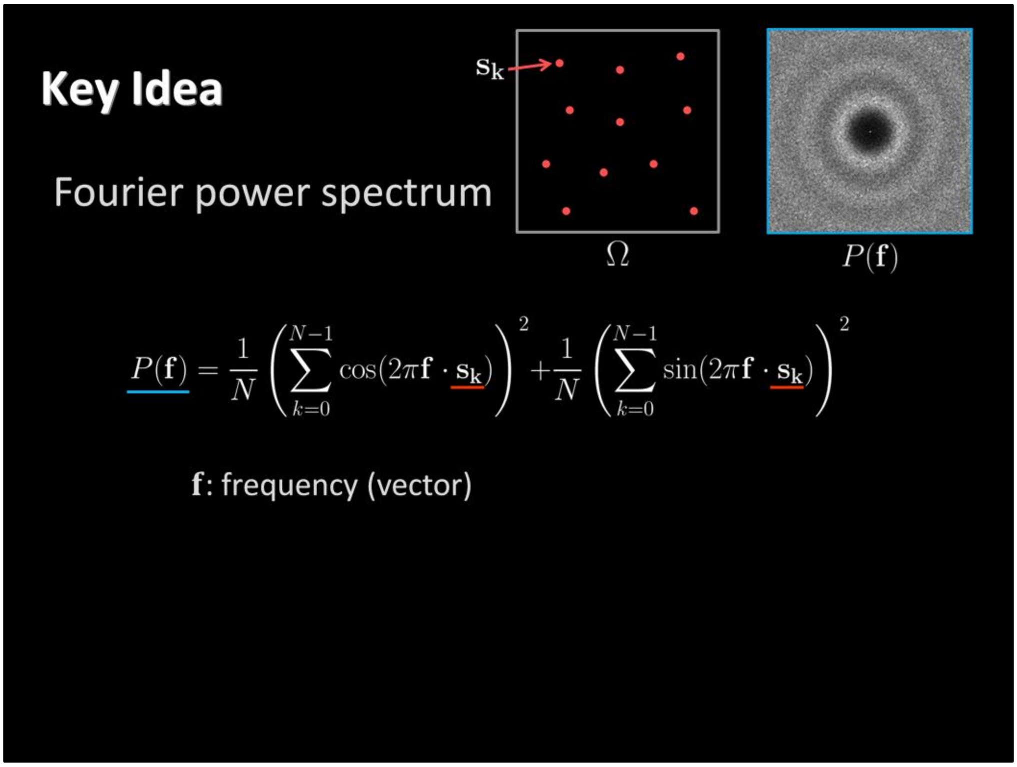

Now let me briefly explain the key idea behind our method. To start, let's assume a set of samples s_k are computed on the 2D plane. Their Fourier power spectrum ■ is defined by this equation where f is a vector frequency.

Latex:

P(\mathbf{f})=\frac{1}{N}\left(\sum_{k=0}^{N-1} \cos(2 \pi \mathbf{f}\cdot\mathbf{s_k}) \right)^2 + \frac{1}{N}\left(\sum_{k=0}^{N-1} \sin(2 \pi \mathbf{f}\cdot\mathbf{s_k}) \right)^2

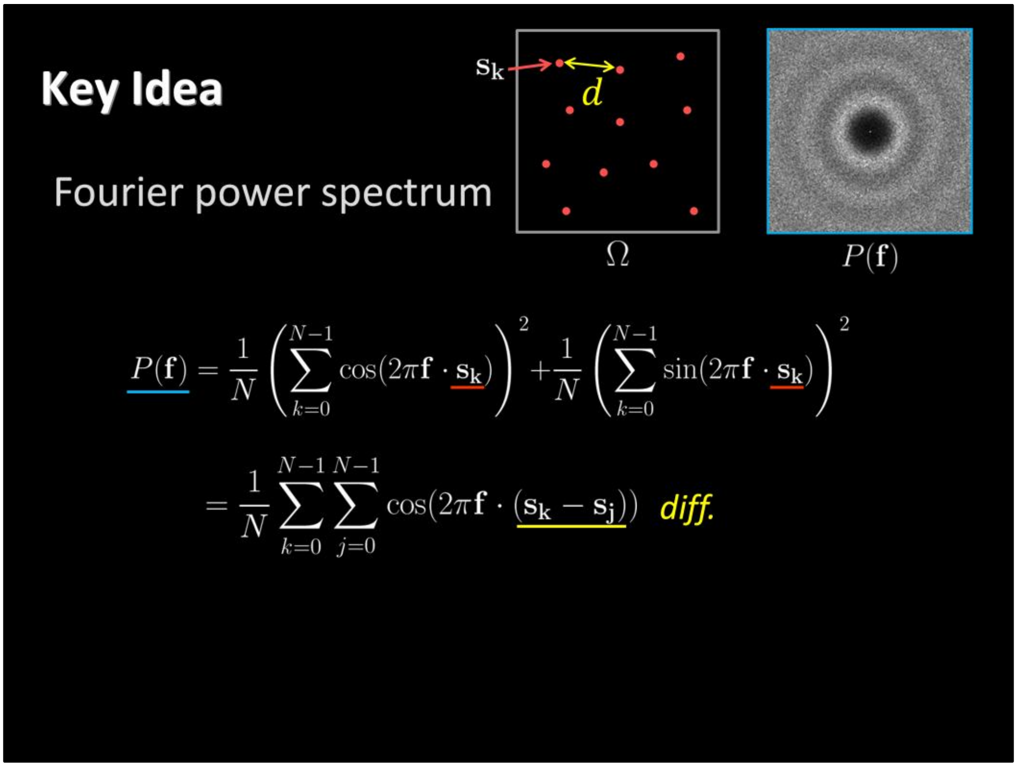

By using basic trigonometry, we can convert this equation into a simpler form that only relies on the difference between every two samples. We call this the differential vector d.

Latex:

P(\mathbf{f})=\left(\sum_{k=0}^{N-1} \cos(2 \pi \mathbf{f}\cdot\mathbf{x_k}) \right)^2 + \left(\sum_{k=0}^{N-1} \sin(2 \pi \mathbf{f}\cdot\mathbf{x_k}) \right)^2

=\frac{1}{N}\sum_{k=0}^{N-1} \sum_{j=0}^{N-1} \cos (2 \pi \mathbf{f} \cdot (\mathbf{s_k}-\mathbf{s_j}))

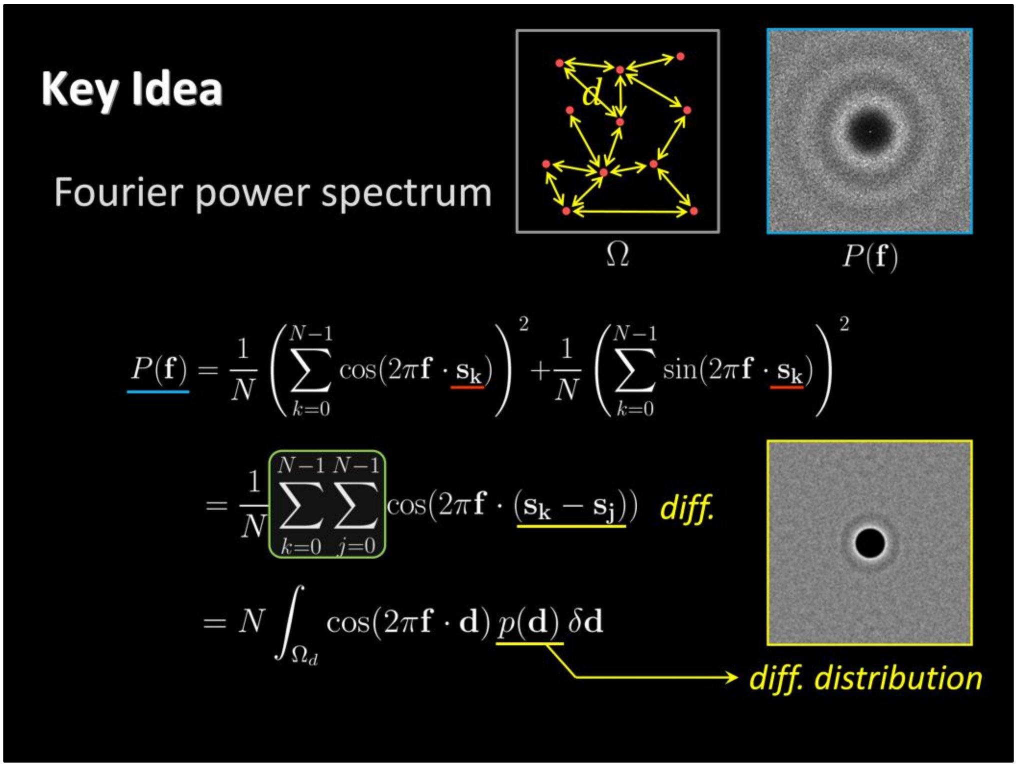

Since the equation sums over all such differential vectors, instead of going through every pair of samples individually to compute the summation, we can first construct a histogram of these differential vectors, and then ■ convert the summation into an integration form, where p(d) represents the histogram. Simply speaking, it counts how many pairs of samples are separated by any specific differential vector d, and ■ can be plotted out as a distribution function shown here.

=N \int_{{\Omega}_d}\cos(2\pi\mathbf{f} \cdot \mathbf{d}) \, p(\mathbf{d}) \, \delta \mathbf{d}

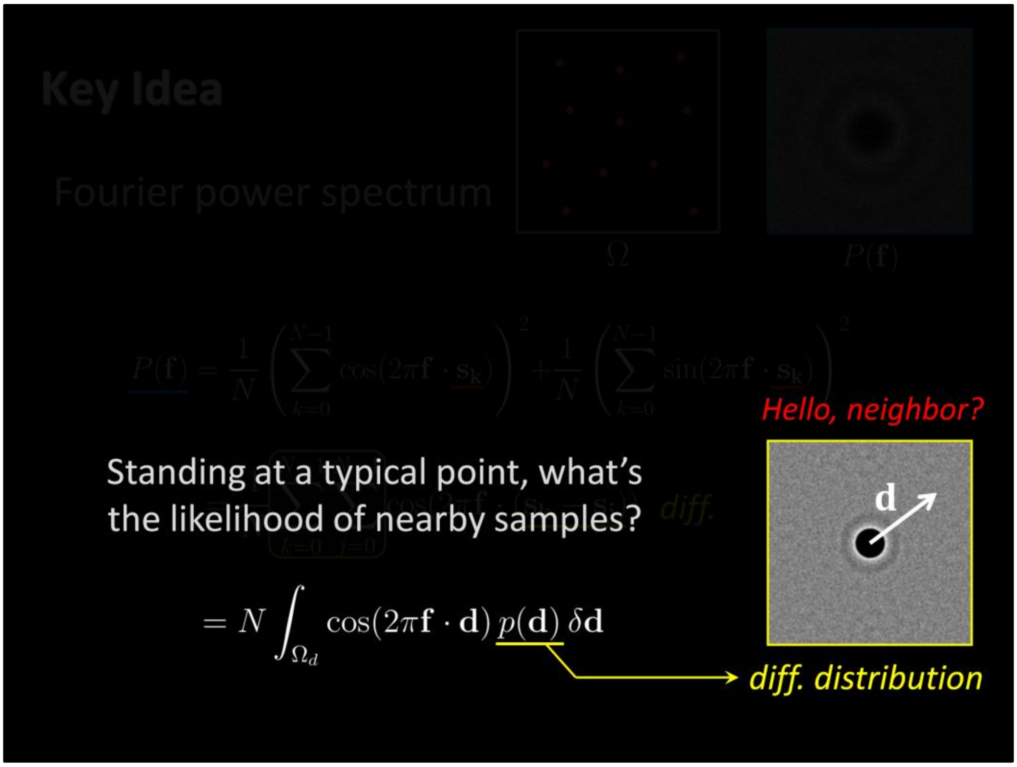

We can also understand this function as a quantity that tells us by standing at a typical point in this sample set, what's the likelihood of its nearby samples; in other words, ■ what's the probability that a sample can be found at any location in its nearby space.

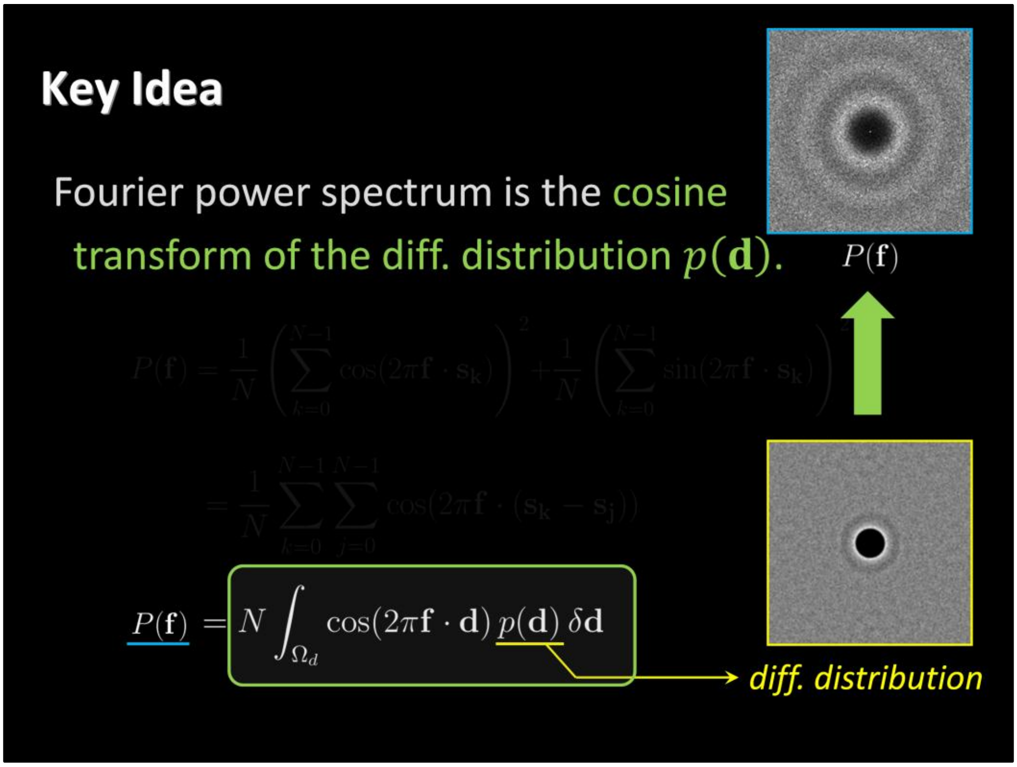

The significance of this function p(d) is that it completely determines the Fourier power spectrum. Now, from the equation that the Fourier power spectrum is simply a cosine transform of p(d). Therefore, a direct inspection of p(d) is ■ sufficient to infer everything about the power spectrum. This is the key to extend our method to non-uniform samples, without relying on a Fourier basis.

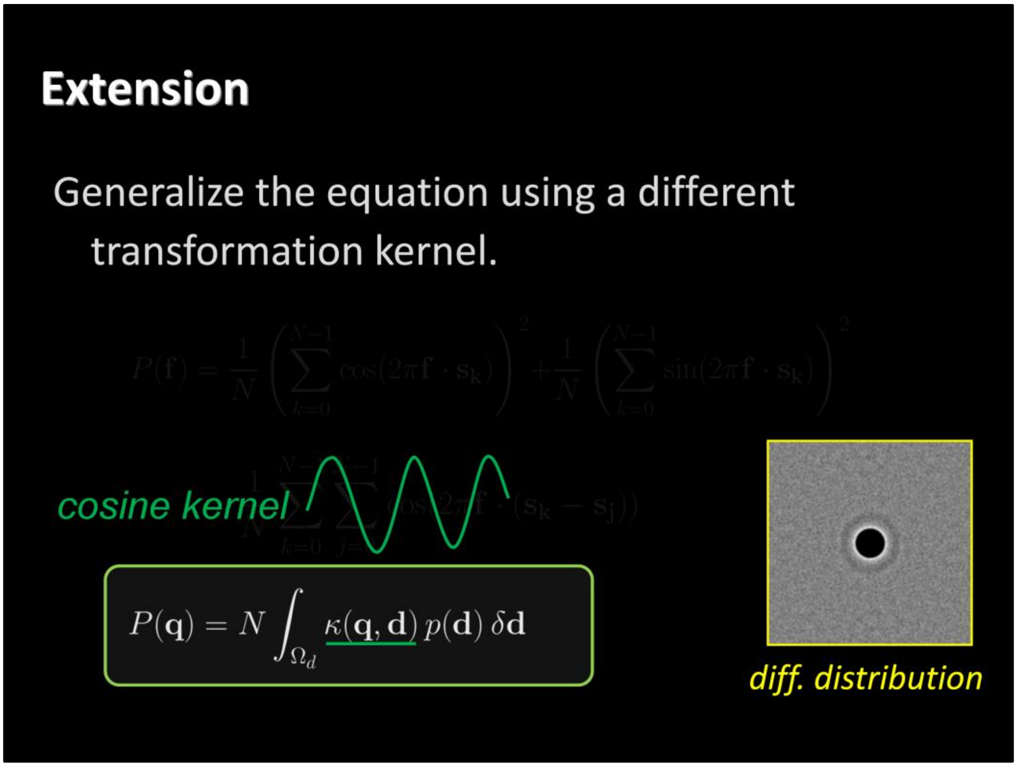

In addition, we can generalize this equation by using a different transformation kernel. This leads to a more general concept of power spectrum. For example, we can replace the cosine kernel, which has infinite support

$$P(\mathbf{q})=N \int_{{\Omega}_d} \kappa(\mathbf{q},\mathbf{d}) \, p(\mathbf{d}) \, \delta \mathbf{d}$$

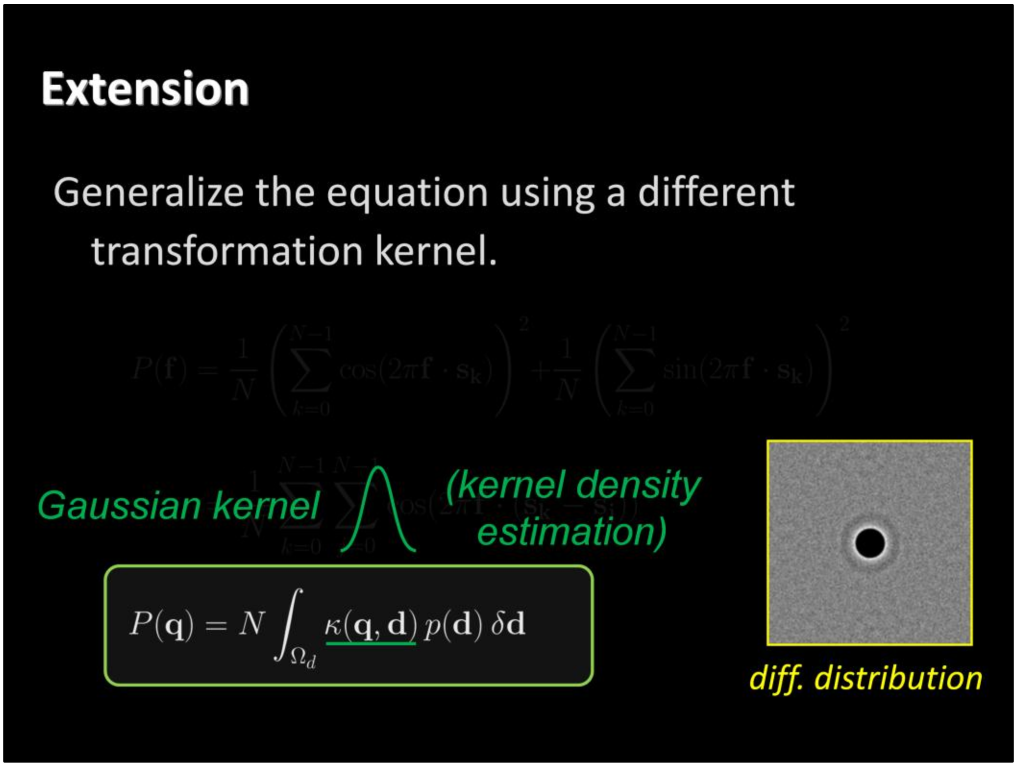

with a compact Gaussian kernel. This can be viewed as a kernel density estimation of p(d), which is what we have applied for generating all results in the paper.

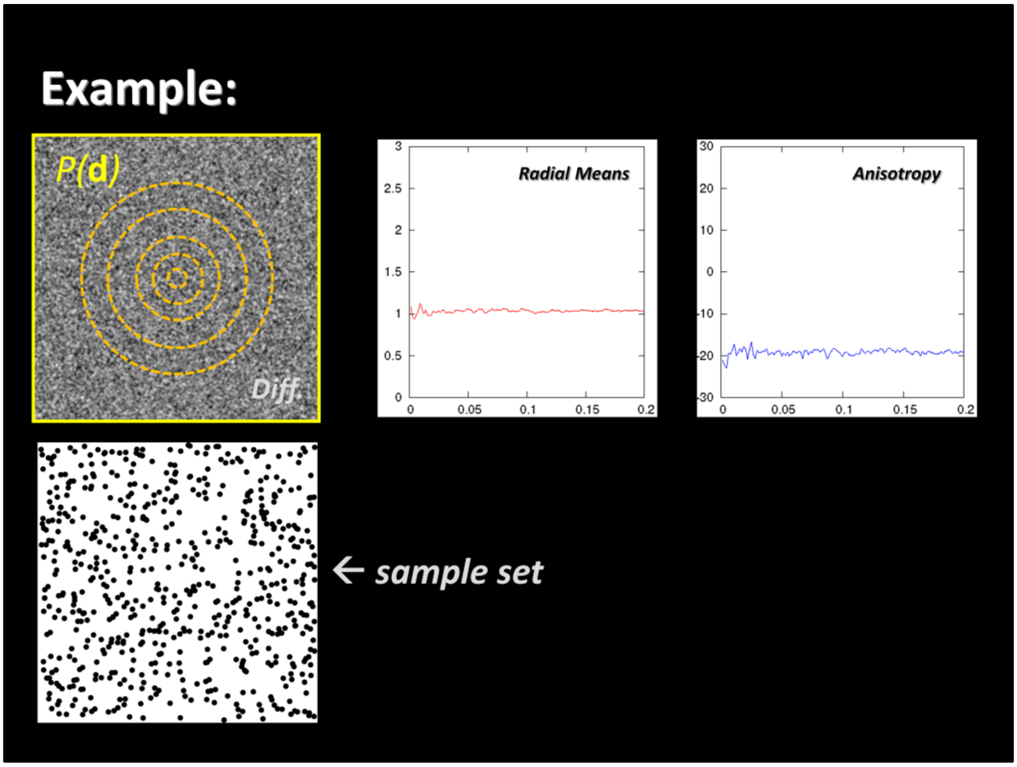

Now let's see how our method works. Given a sample set with unspecified property, we first compute its differential distribution p(d); then as before, we compute the circular average and relative variance, and this results in the radial means and anisotropy graphs. In this case, both graphs are flat, indicating that the samples are uncorrelated and equally distributed everywhere and in every direction, so we conclude that this a ■ white noise sample set.

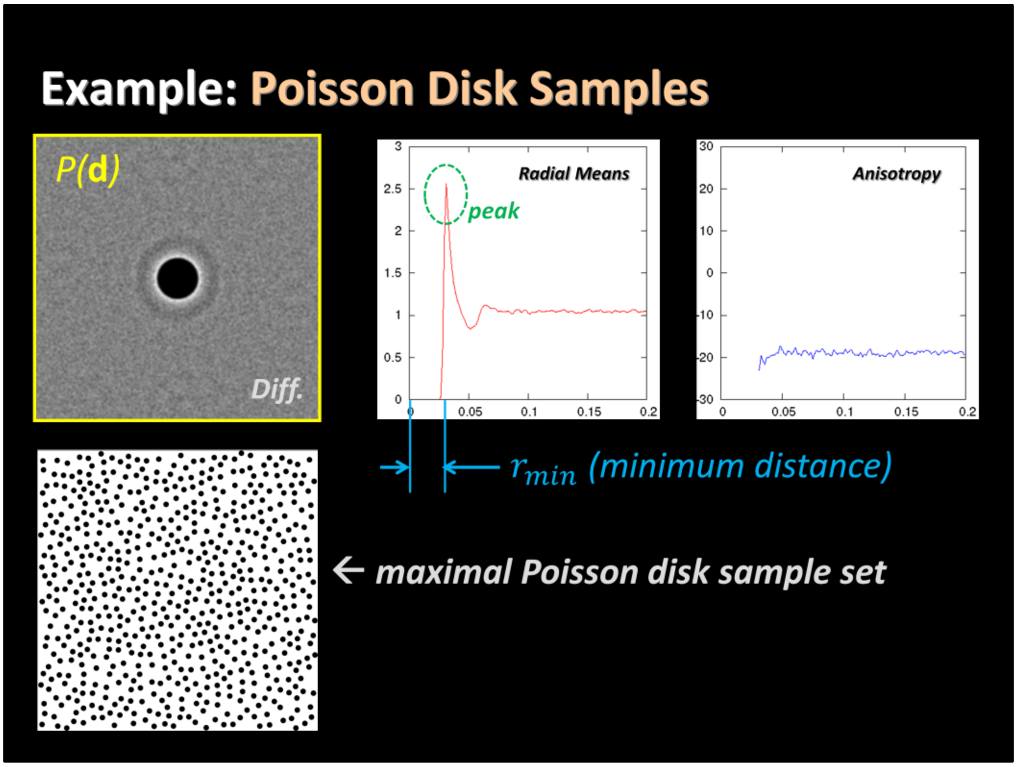

In this second example, we computed p(d) for a maximal Poisson disk sample set. Note that in the radial means graph, there is a very sharp peak right around the minimum sample distance rmin. This indicates that the samples are densely packed under the minimum distance constraint.

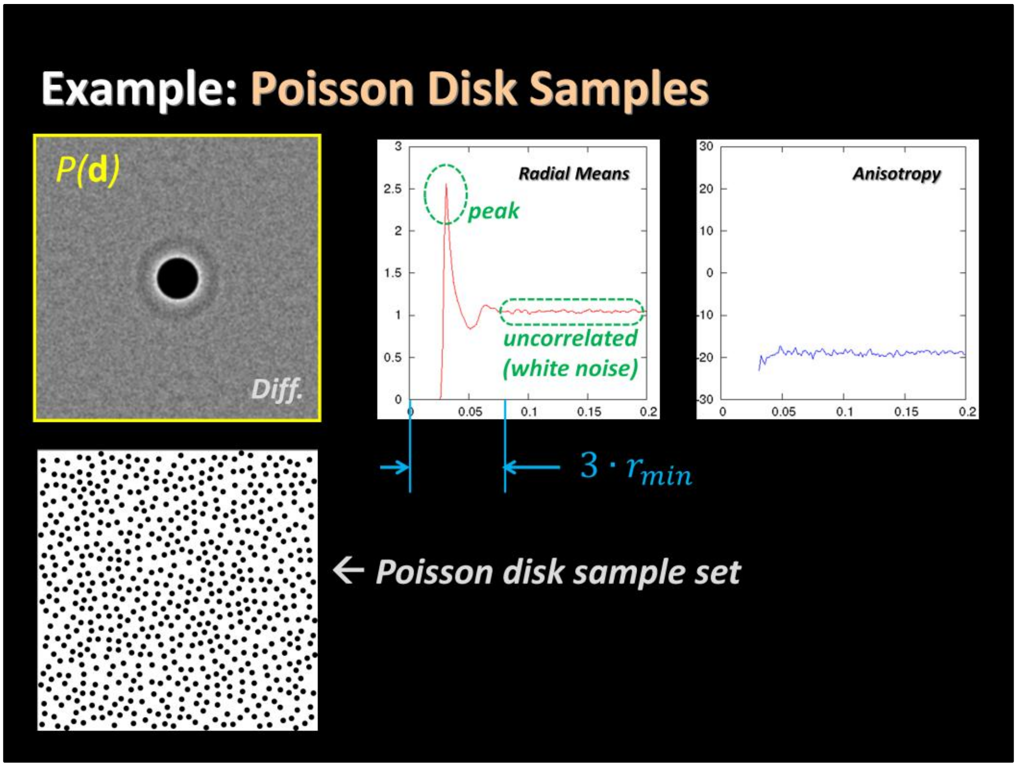

Also, notice that beyond 3 times rmin, the graph is essentially flat, indicating that samples beyond this distance are uncorrelated and appear as white noise to each other. Therefore, the interesting features about this sample set are all localized within a small range from 0 to 3rmin.

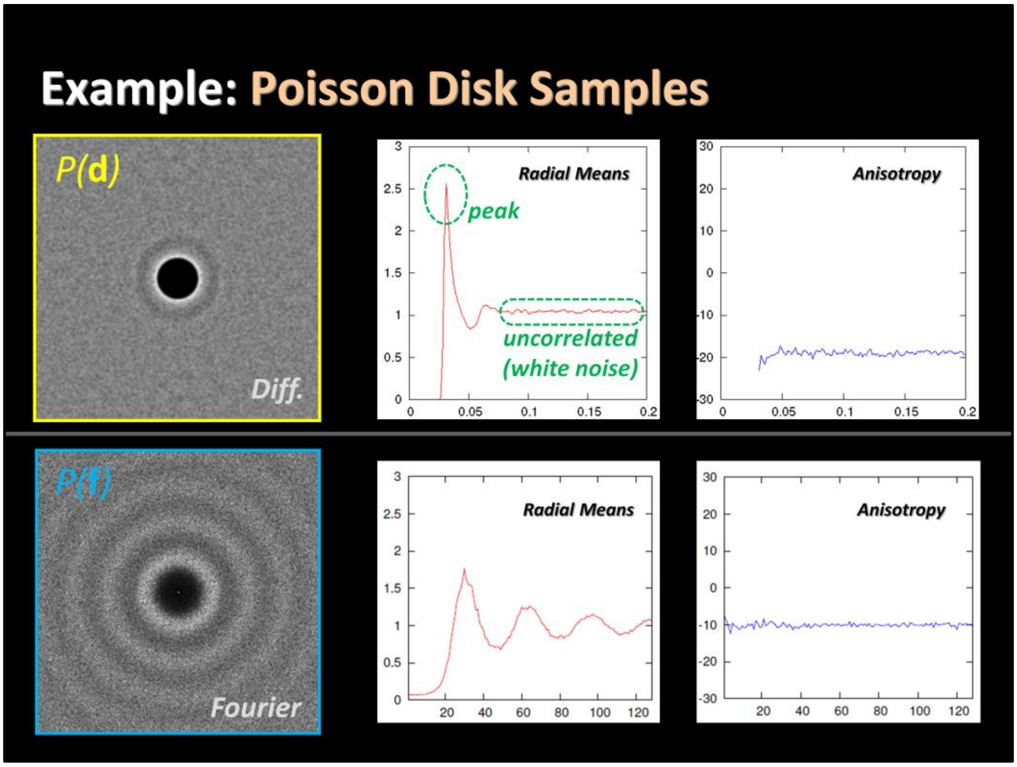

For the purpose of comparison, I also provided the Fourier analysis results here. As I said before, the Fourier power spectrum is simply a cosine transform of P(d). So you can the features in p(d), such as the high peak, directly lead to the features in the Fourier analysis results.

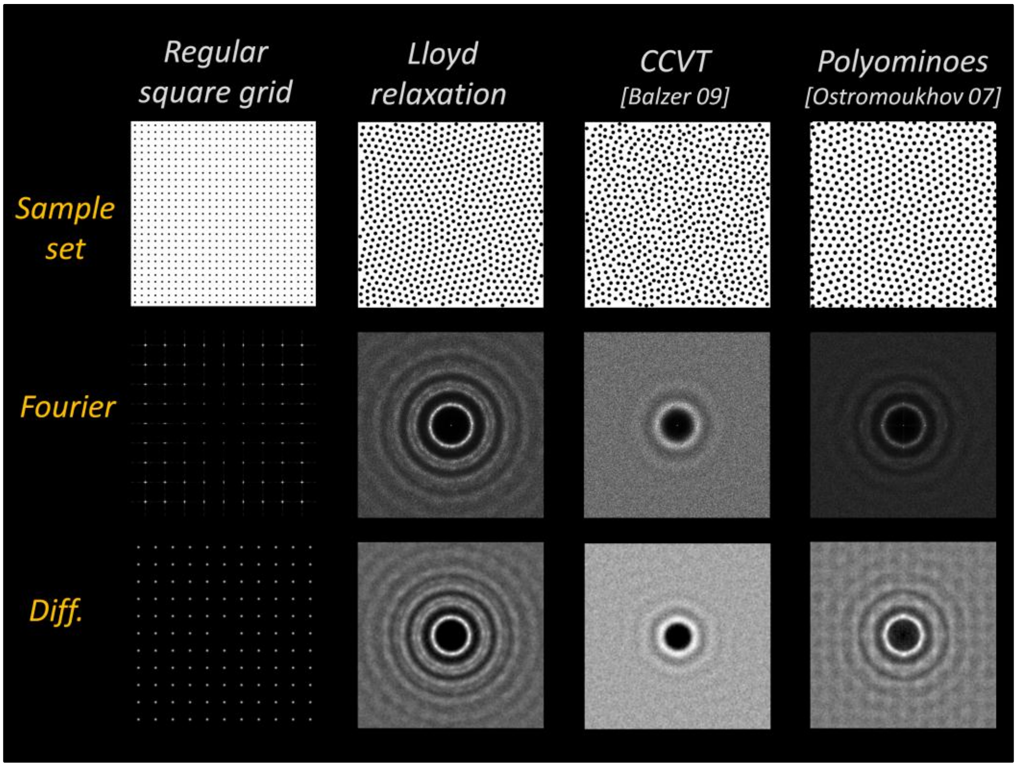

We have also analyzed a variety of other common sampling methods. These results can be found in the paper.

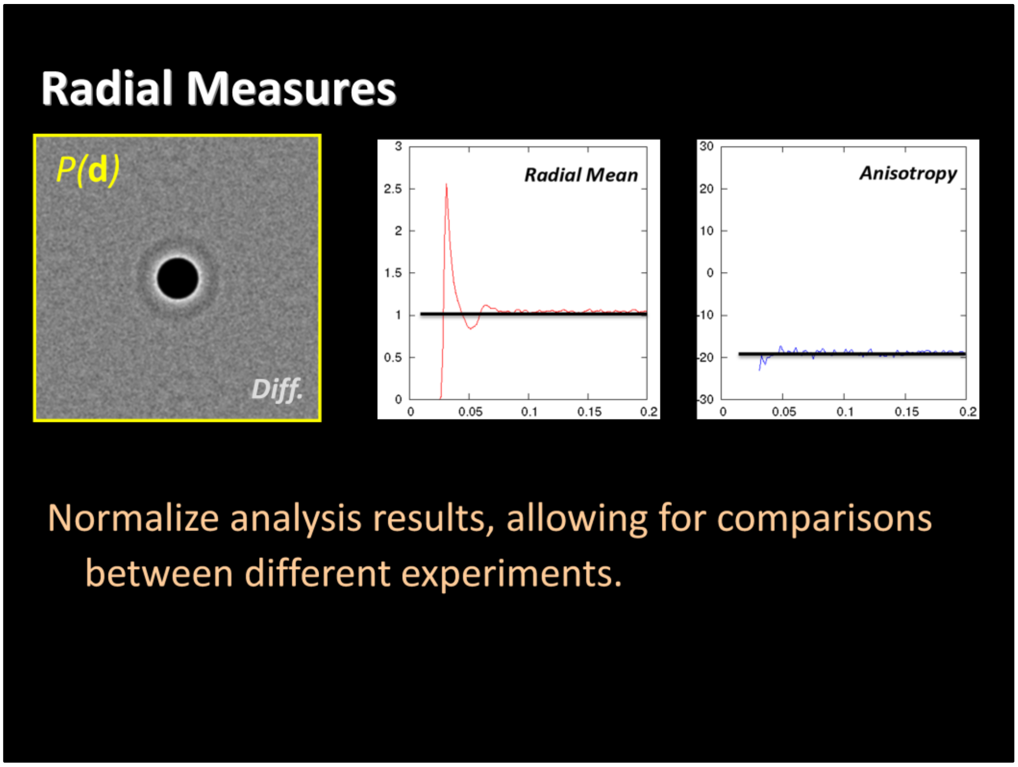

Finally, we have derived analytic radial measures to normalize our analysis results. This makes it possible to draw direct comparisons between experiments running with different settings and parameters. Details can be found in the paper.

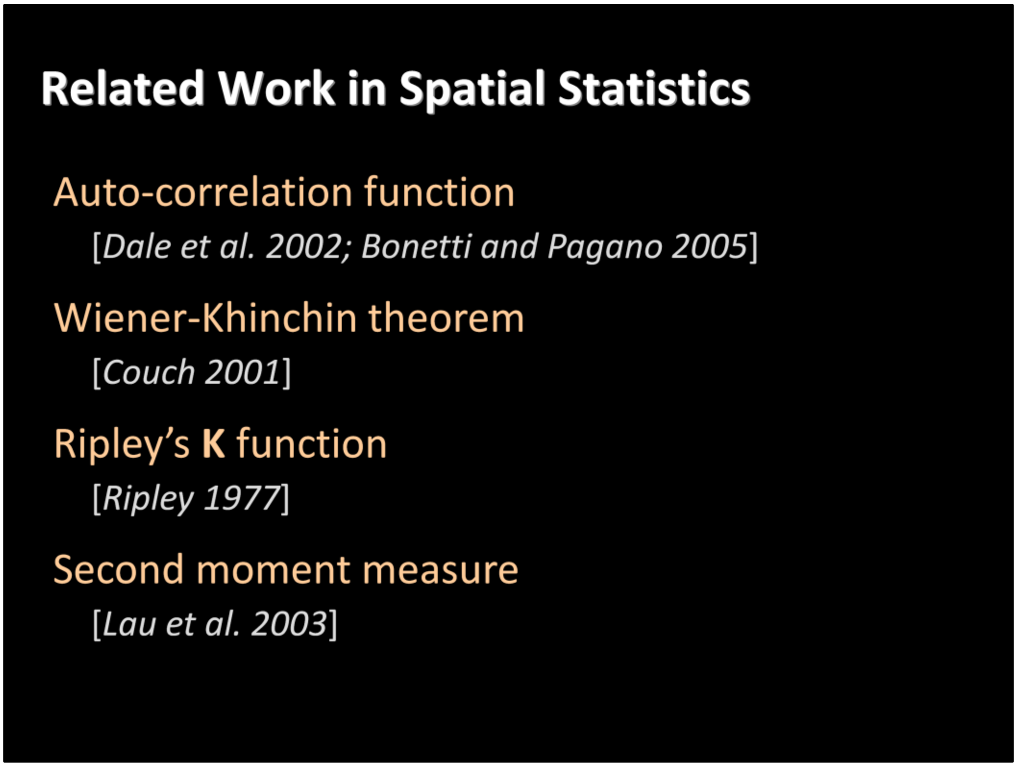

I should note here that our basic idea is fundamentally connected to several well-known concepts in the spatial statistics literature, particularly the autocorrelation function. Despite these connections, our novelty is in generalizing the basic idea beyond uniform sampling domains, and presenting a convenient tool for analyzing non-uniform samples. The key idea is to make use of the local nature of the differential distribution, and unwarp the local neighborhood of each sample to account for the spatially varying density. With this, I will give the podium to liyi, and he will continue with the rest of the talk.

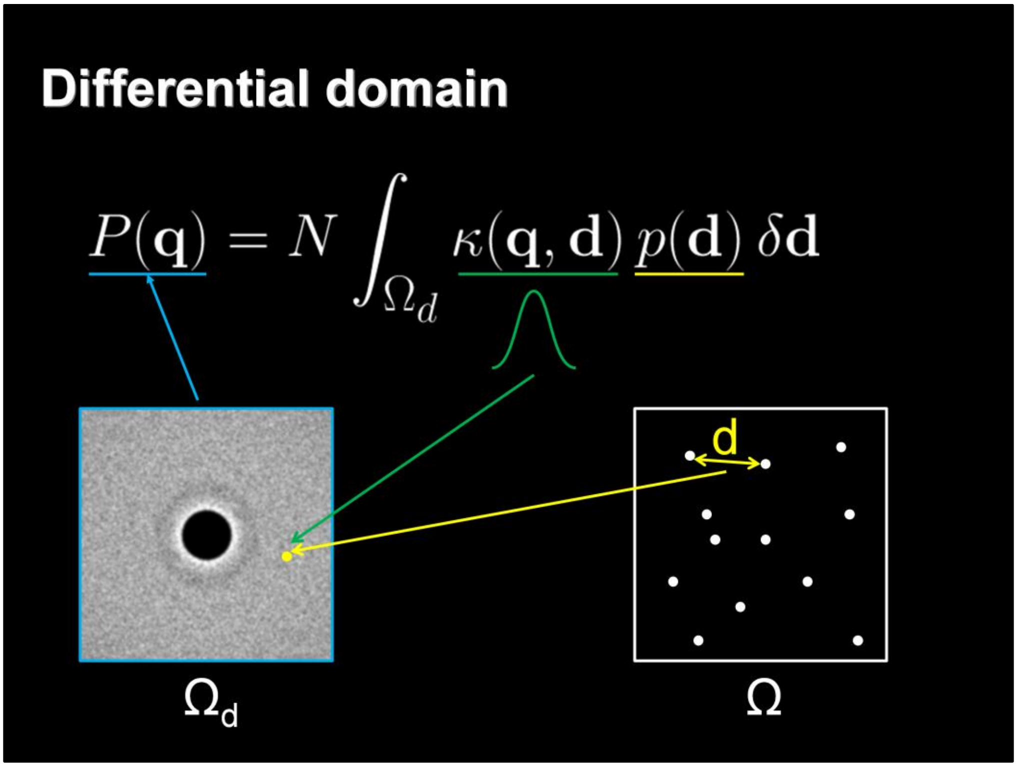
Differential domain
$P(\mathbf{q}) = N \int_{\Omega_d} \kappa(\mathbf{q}, \mathbf{d})\, p(\mathbf{d})\, \delta\mathbf{d}$
d
Ω_d
Ω

Range selection

$$\underline{P(\mathbf{q})} = N \int_{\Omega_d} \underline{\kappa(\mathbf{q}, \mathbf{d})}\ \underline{p(\mathbf{d})}\ \underline{\xi(\mathbf{d})}\ \delta\mathbf{d}$$

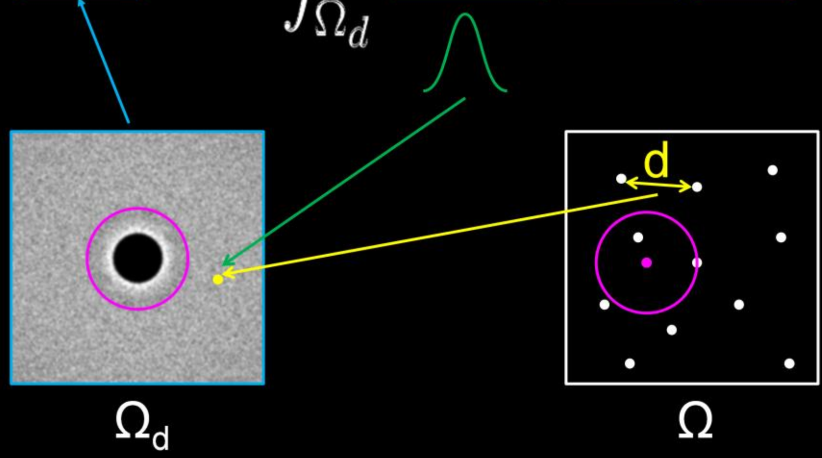

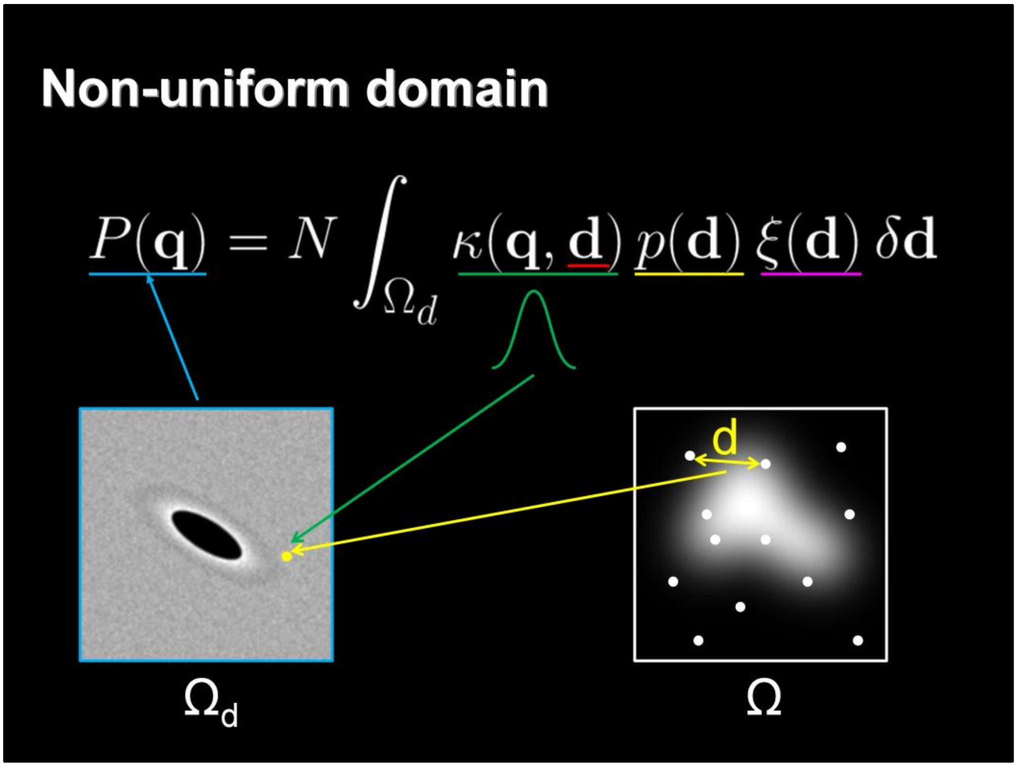
Non-uniform domain
$P(\mathbf{q}) = N \int_{\Omega_d} \kappa(\mathbf{q}, \mathbf{d})\, p(\mathbf{d})\, \xi(\mathbf{d})\, \delta\mathbf{d}$
d
Ω_d
Ω

Non-uniform domain

$$P(\mathbf{q}) = N \int_{\Omega_d} \kappa(\mathbf{q}, \chi(s, s', \mathbf{d}))\, p(\mathbf{d})\, \xi(\mathbf{d})\, \delta\mathbf{d}$$

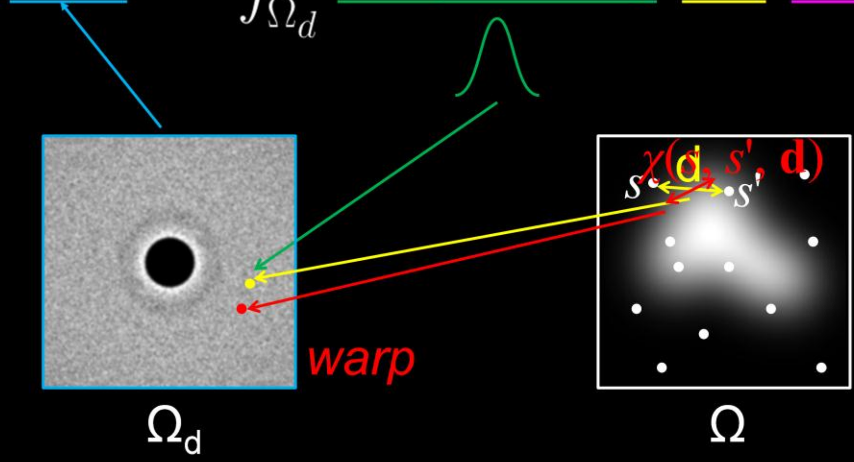

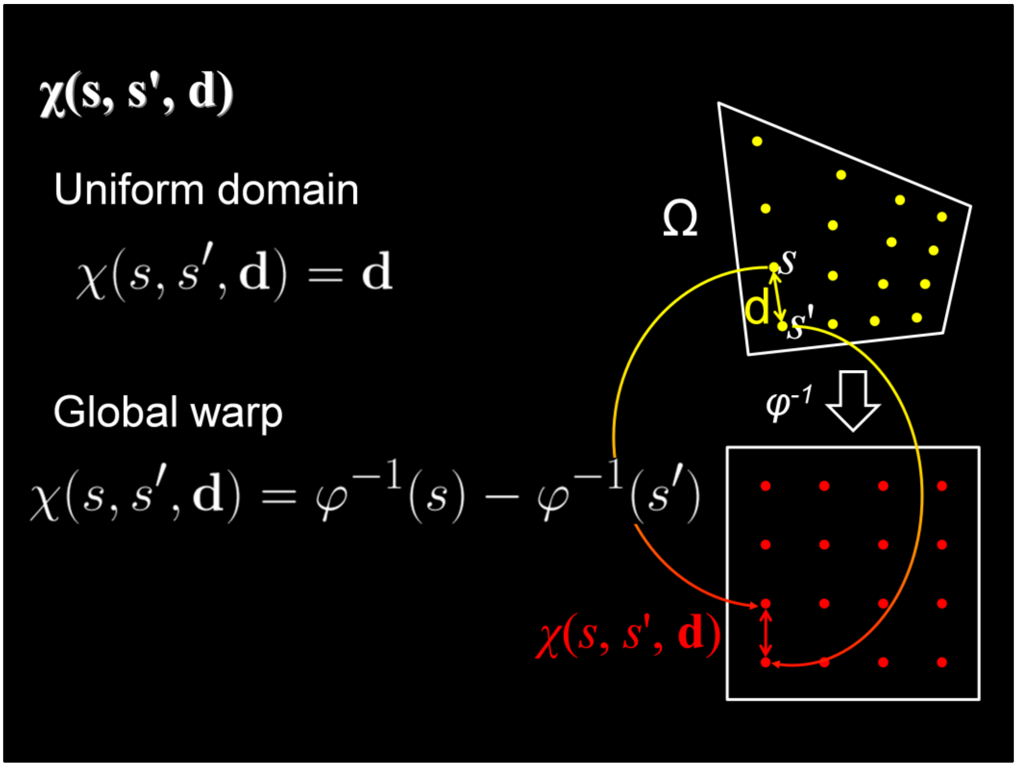

$\chi(s, s', d)$
Uniform domain
$\chi(s, s', \mathbf{d}) = \mathbf{d}$
Global warp
$\chi(s, s', \mathbf{d}) = \varphi^{-1}(s) - \varphi^{-1}(s')$
Ω
s
d
s'
φ^{-1}
$\chi(s, s', \mathbf{d})$

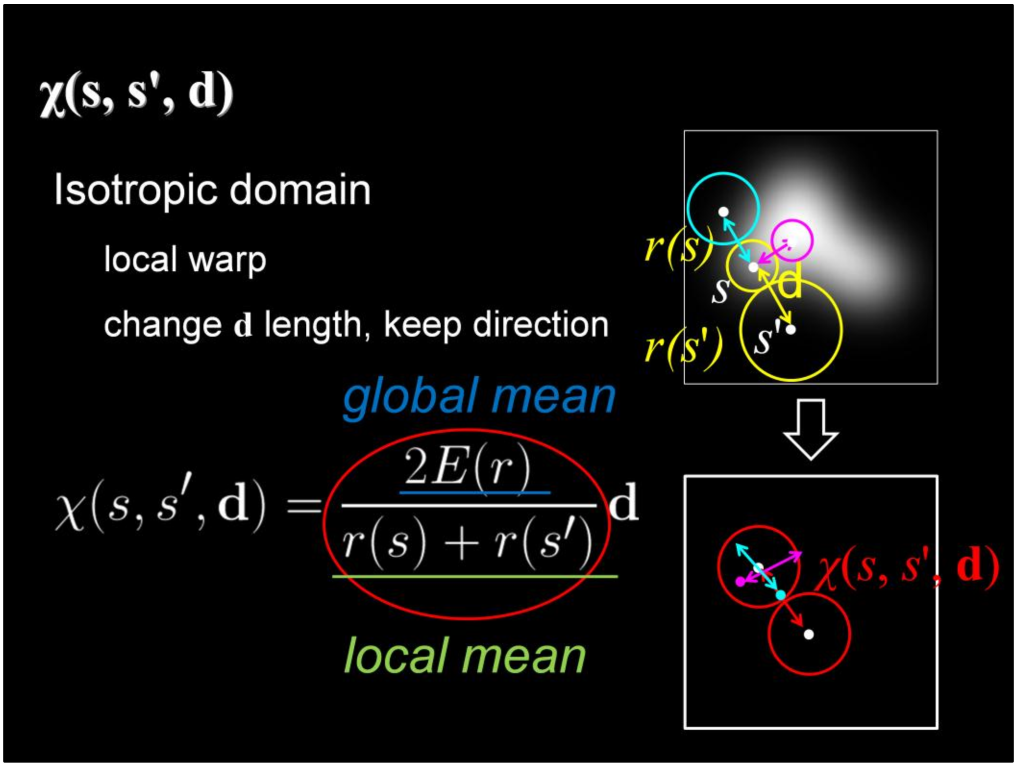
χ(s, s', d)
Isotropic domain
local warp
change d length, keep direction
global mean
$\chi(s, s', \mathbf{d}) = \frac{2E(r)}{r(s) + r(s')}\mathbf{d}$
local mean
r(s)
s
d
r(s')
s'
χ(s, s', d)

χ(s, s', d)

Anisotropic domain

- local Jacobian J [Li et al. 2010]
- change $\mathbf{d}$ length and direction

$$\chi(s,s',\mathbf{d}) = \frac{1}{E(\lambda)}\left(\frac{J^{-1}(s)+J^{-1}(s')}{2}\right)^{-1}\mathbf{d}$$

global mean *local mean*

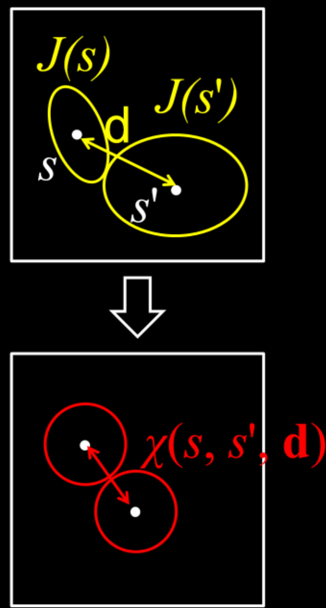

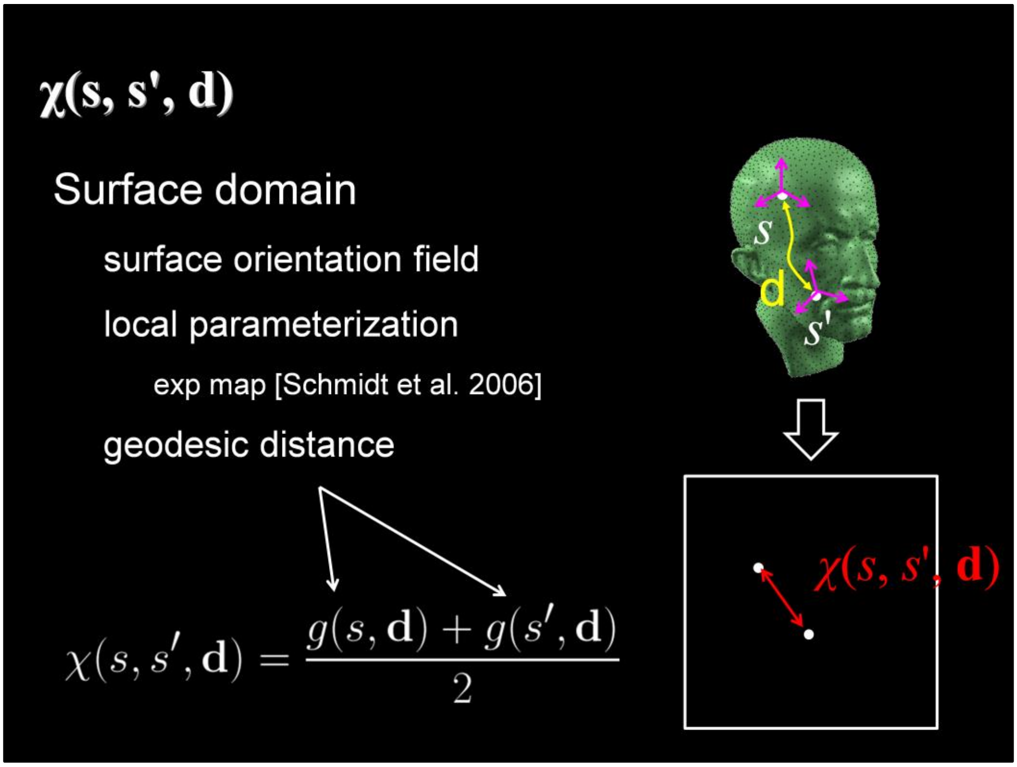

Clarify this brandon

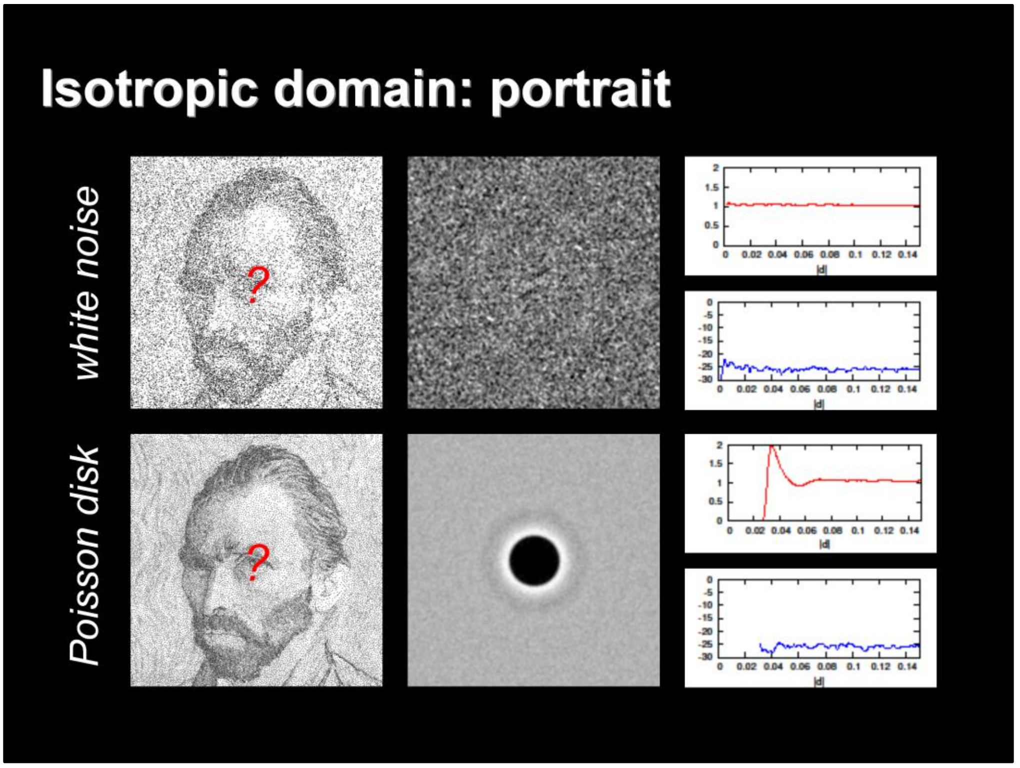

convert -density 10 portrait_0p03_white_0.pdf portrait_0p03_white_0.png

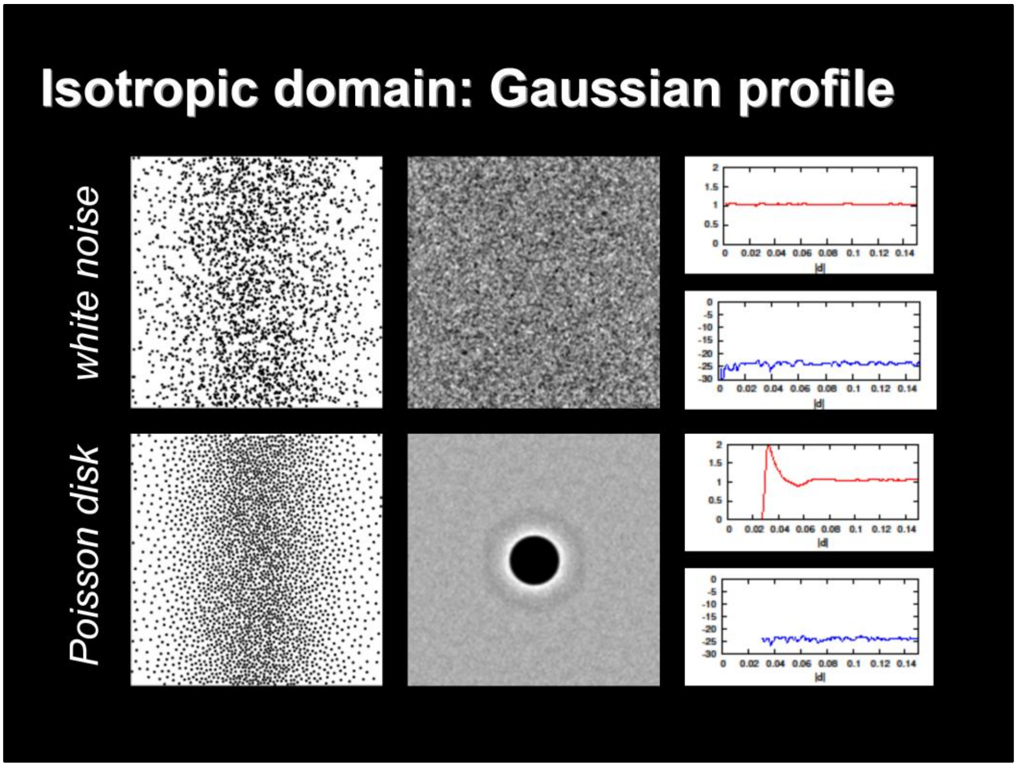
Isotropic domain: Gaussian profile
white noise
Poisson disk

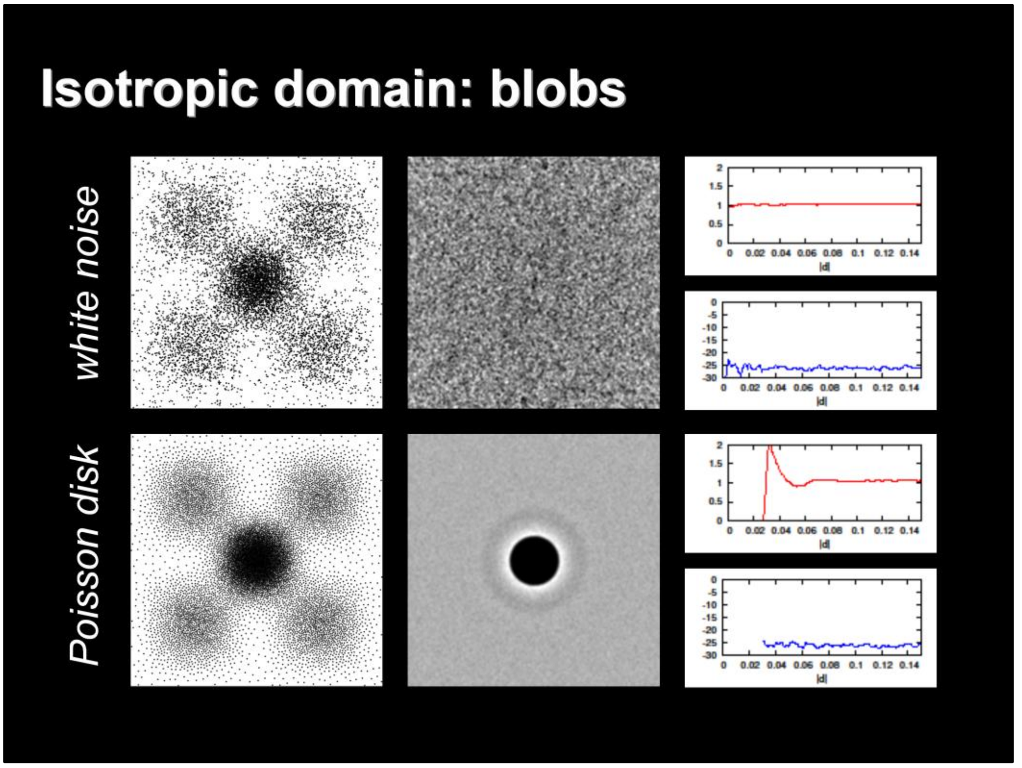

convert -density 10

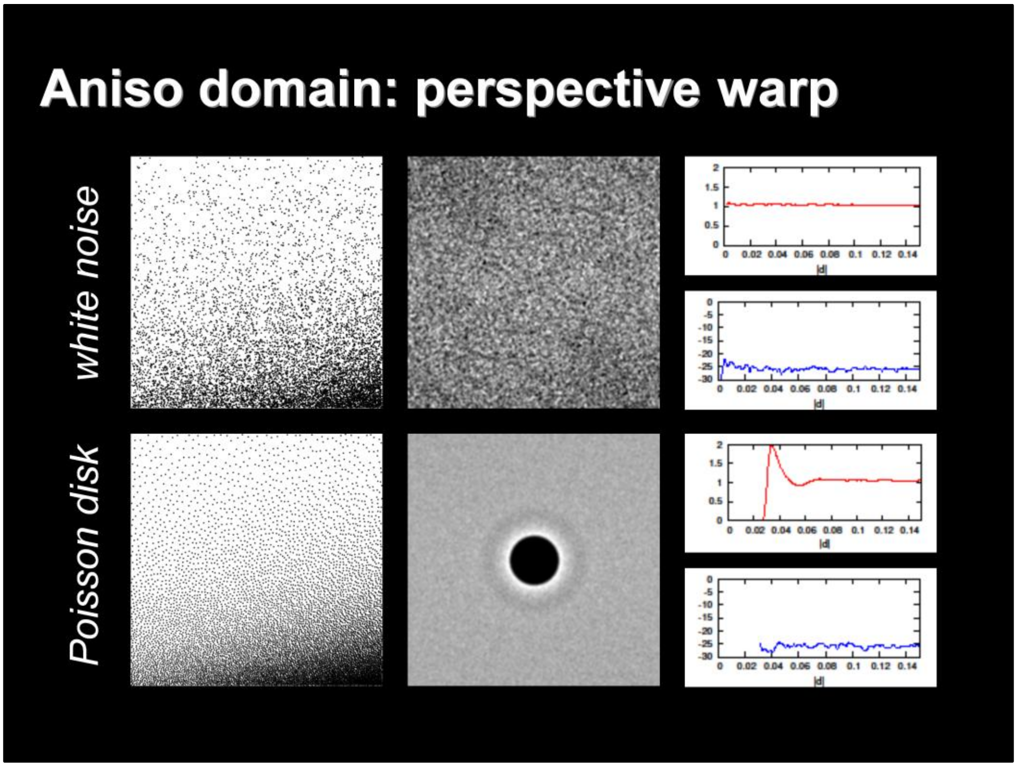

convert -density 10 portrait_0p03_white_0.pdf portrait_0p03_white_0.png

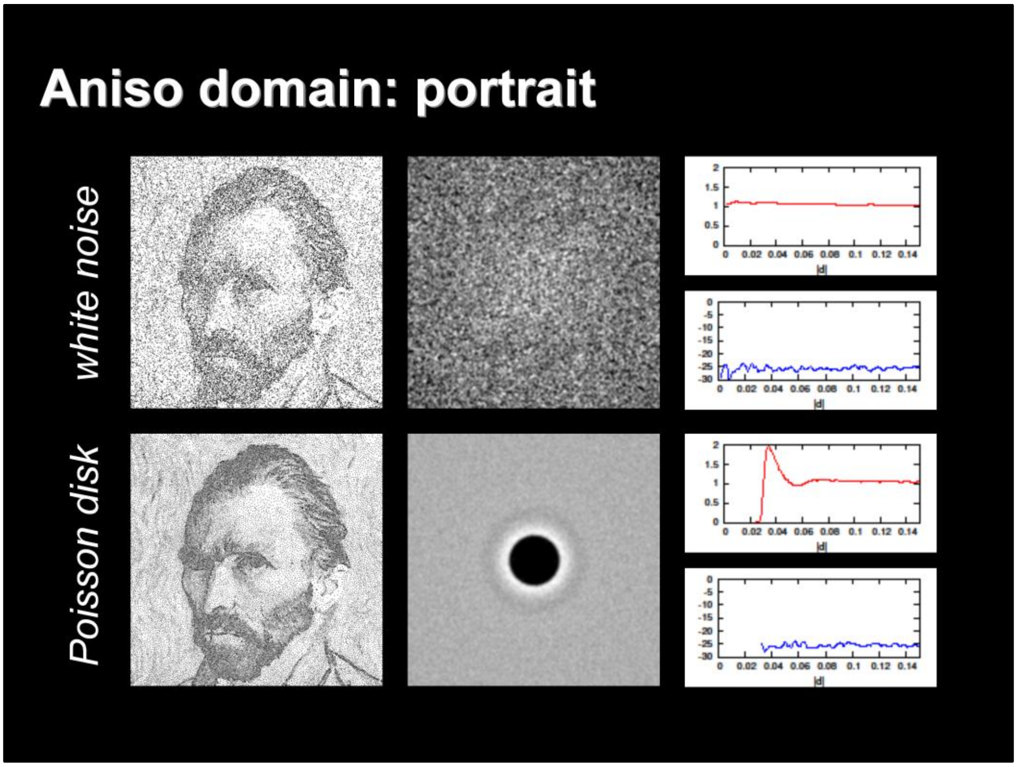

convert -density 10 portrait_0p03_white_0.pdf portrait_0p03_white_0.png

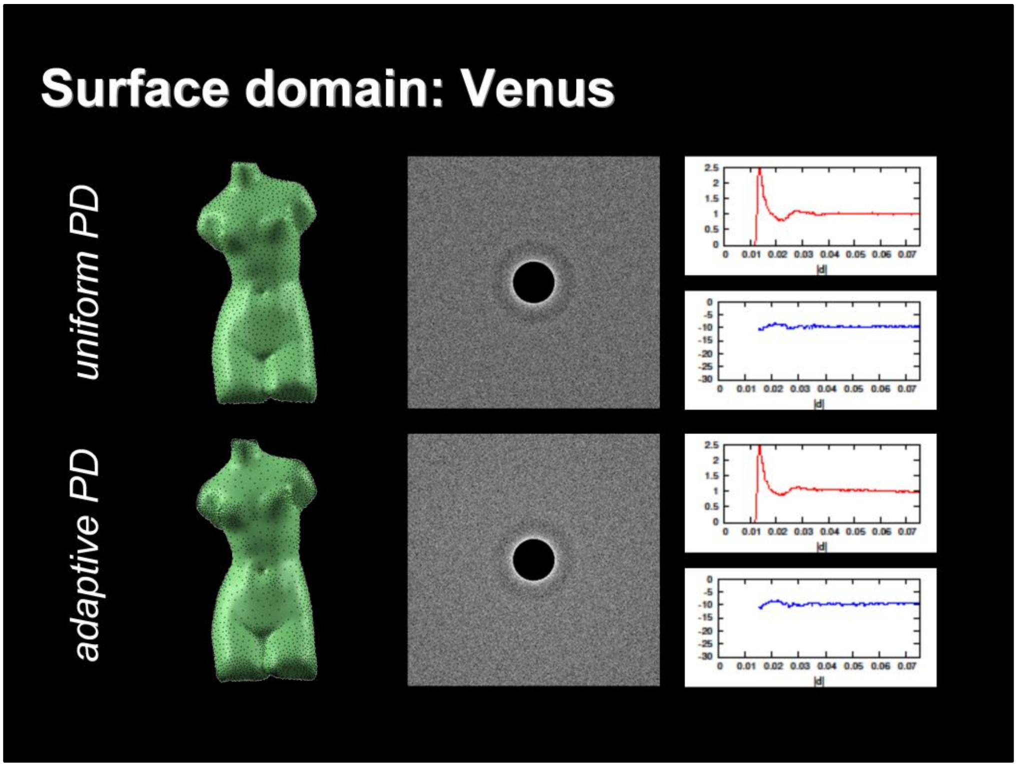

convert -density 10 portrait_0p03_white_0.pdf portrait_0p03_white_0.png

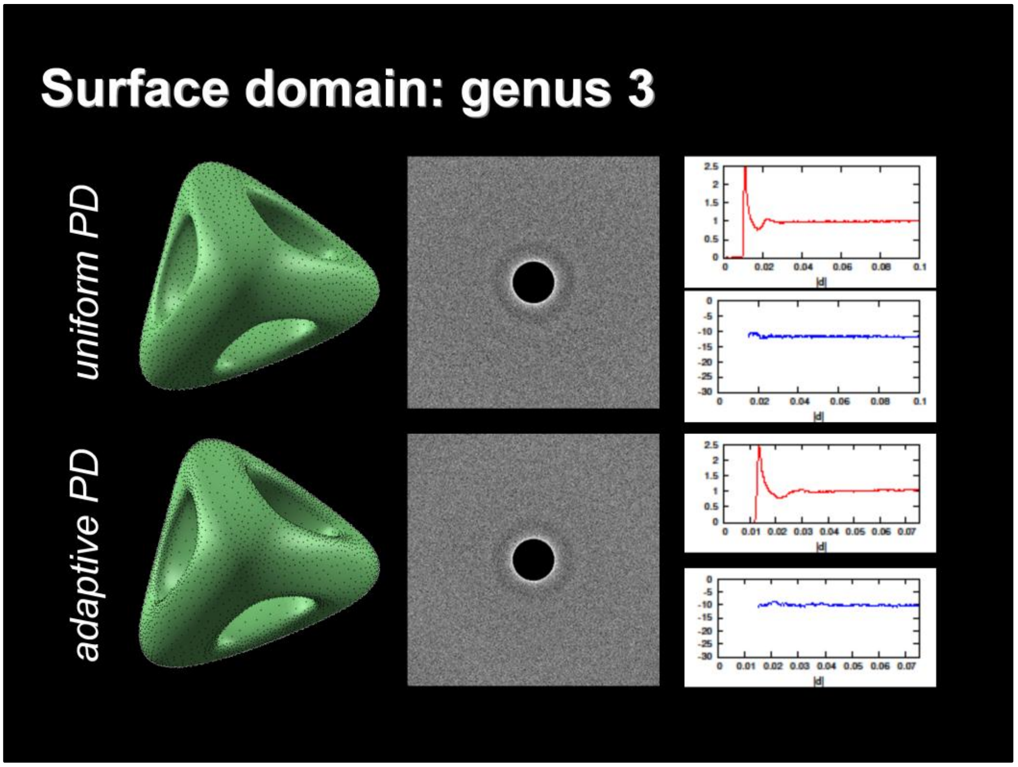

convert -density 10 portrait_0p03_white_0.pdf portrait_0p03_white_0.png

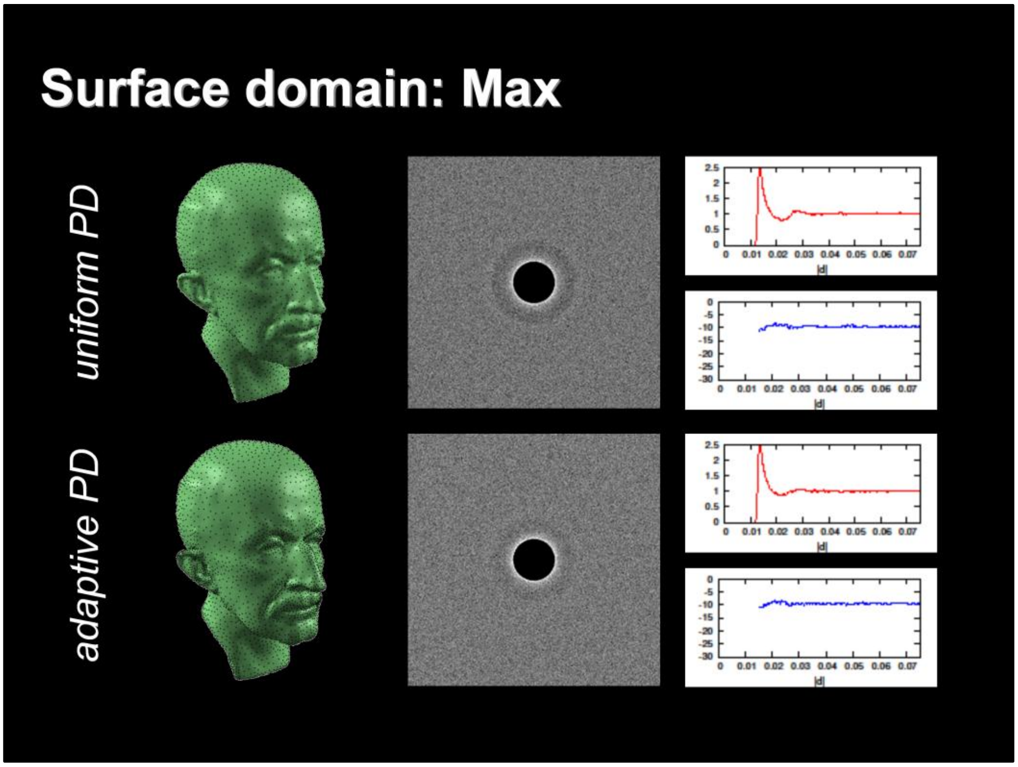

convert -density 10 portrait_0p03_white_0.pdf portrait_0p03_white_0.png

Other applications

Spatial analysis

- ρ: normalized min spacing [Lagae and Dutre 2008]
- Use local warp $\chi(s, s', \mathbf{d})$

Quantifying blue noise

- See extended paper

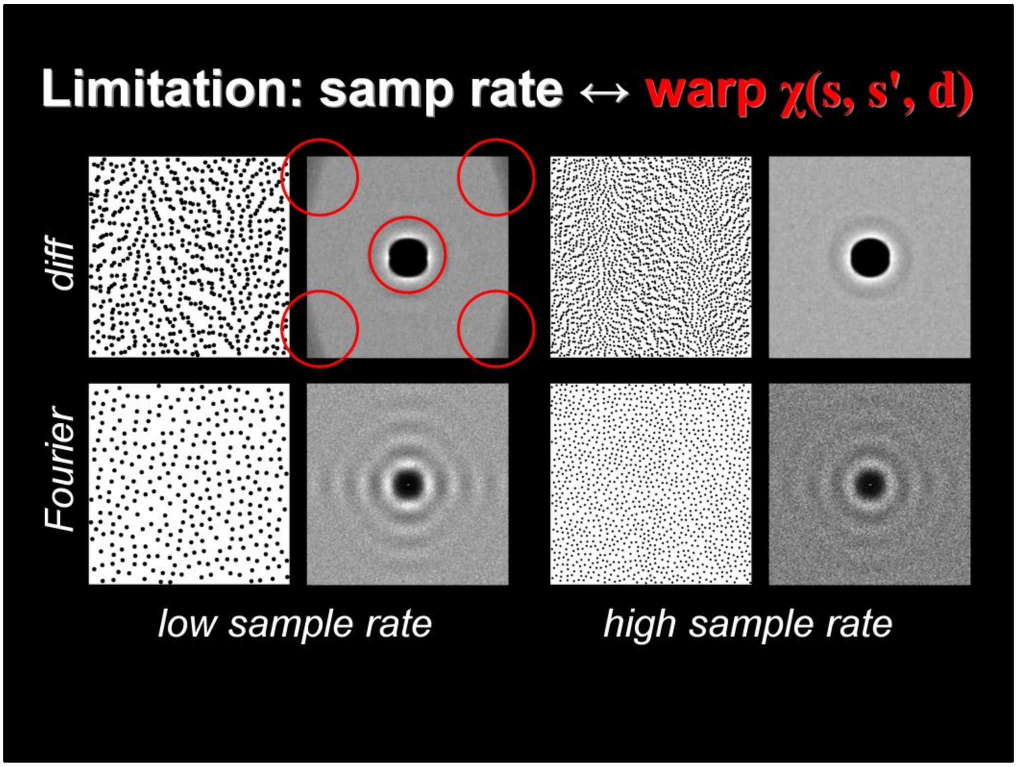
Limitation: samp rate ↔ warp χ(s, s', d)
diff
Fourier
low sample rate
high sample rate

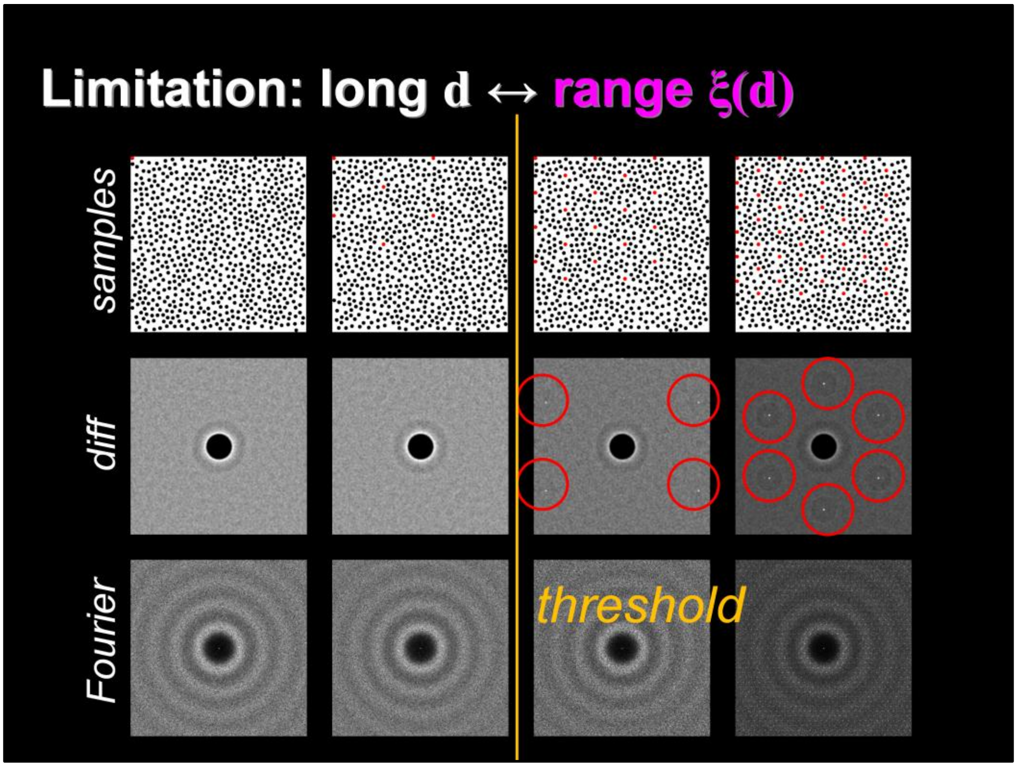
Limitation: long d ↔ range ξ(d)
samples
diff
Fourier
threshold

Future work

Higher dimensional space

Analyzing other sampling methods

e.g. tiling [Kopf et al. 2006; Ostromoukhov 2007]

New sampling algorithms

Acknowledgements

Paper

- Hongwei Li
- SIGGRAPH reviewers
- NSF grant CCF-0746577

Talk

- Hugues Hoppe
- Xin Tong
- Brandon Lloyd
- Yahan Zhou
- Steve Giguere
- Will Stumpf