

Differential Domain Analysis for Non-uniform Sampling

Li-Yi Wei

Microsoft Research

Rui Wang

University of Massachusetts Amherst

Abstract

Sampling is a core component for many graphics applications including rendering, imaging, animation, and geometry processing. The efficacy of these applications often crucially depends upon the distribution quality of the underlying samples. While uniform sampling can be analyzed by using existing spatial and spectral methods, these cannot be easily extended to general non-uniform settings, such as adaptive, anisotropic, or non-Euclidean domains.

We present new methods for analyzing non-uniform sample distributions. Our key insight is that standard Fourier analysis, which depends on samples' spatial locations, can be reformulated into an equivalent form that depends only on the distribution of their location *differentials*. We call this differential domain analysis. The main benefit of this reformulation is that it bridges the fundamental connection between the samples' spatial statistics and their spectral properties. In addition, it allows us to generalize our method with different computation kernels and differential measurements. Using this analysis, we can quantitatively measure the spatial and spectral properties of various non-uniform sample distributions, including adaptive, anisotropic, and non-Euclidean domains.

Keywords: differential domain, analysis, non-uniform, sampling, spectrum, noise

Links: [DL](#) [PDF](#)

1 Introduction

Sampling is a fundamental component for a variety of graphics algorithms, with applications ranging from rendering, imaging, animation, to geometry processing [Lloyd 1983; Dippé and Wold 1985; Cook 1986; Mitchell 1987; Turk 1992; Glassner 1994; Alliez et al. 2002; Dutre et al. 2002; Pharr and Humphreys 2004; Ostromoukhov et al. 2004; Kopf et al. 2006; Ostromoukhov 2007; Fu and Zhou 2008; Balzer et al. 2009; Wei 2010; Öztireli et al. 2010].

Despite the diverse algorithm characteristics and application domains, two common methodologies exist for evaluating the quality of samples: (1) spatial uniformity, including measures such as discrepancy [Shirley 1991] and ρ – the normalized minimum spacing between pairs of samples [Lagae and Dutré 2008]; (2) power spectrum analysis, including radial mean and anisotropy [Lagae and Dutré 2008]. However, existing methods are primarily designed for uniform Euclidean domains and can not be easily extended to general non-uniform scenarios, such as adaptive, anisotropic, or surface sampling (see Figure 1). To our knowledge, even though a few techniques exist for limited situations (e.g. warpable anisotropic domains [Li et al. 2010] or uniform surface domains [Bowers et al. 2010]), direct analysis of general non-uniform sampling patterns remains an important open problem. Specifically, many applications require certain forms of non-uniform sampling, and for a given




Figure 1: Differential domain analysis. Here we demonstrate uniform (top) and non-uniform (bottom) sampling patterns analyzed by traditional Fourier spectrum (left) and our method (right). Each group is produced by 10 sets of Poisson disk sampling with $r_{min} = 0.03$ and 628 samples per set. Within each group are the spectrum image, the corresponding radial mean profile (red curve), and the spatial sample pattern. The non-uniform sampling follows the importance function from [Ostromoukhov 2007]. As shown, traditional Fourier method fails to produce meaningful results for non-uniform sampling: note the excessive low frequency energy and the lack of typical blue noise characteristic as compared to the uniform sampling result. Our method analyzes the sample set in differential domain, and thus can well capture the blue noise characteristic: note the existence of a peak value around $r_{min} = 0.03$ in our radial mean profiles, and their consistent appearance across both uniform and non-uniform cases.

non-uniform pattern the underlying generation algorithm may be unknown and thus the analysis must be based on the samples only. Even when the sampling algorithm is known, its property in general non-uniform settings may not be reliably inferred from its behavior in the uniform domain (e.g. the hierarchical warping method in [Clarberg et al. 2005] that may introduce anisotropic stretch).

In this paper, we present new methods for analyzing non-uniform sample distributions, including adaptive, anisotropic, and surface domain samplings. Our key insight is that standard Fourier spectrum analysis, which depends on sample locations, can be reformulated into an equivalent form that only depends on the distribution of sample location *differentials*. We call this differential

domain analysis. The main benefit of the reformulation is that it bridges the fundamental connection between samples' spatial statistics and their spectral properties. This makes it possible to extend our method to non-uniform domains, by using the corresponding differential measurements that we present in the paper. Moreover, we can generalize our method with different computation kernels. For example, in this paper we focus on using the Gaussian kernel, which estimates a smoothed histogram of the location differentials. Finally, we present spatial and spectral analysis results using our method for different sampling methods in various non-uniform settings. In sum, our paper has the following contributions:

- A reformulation of standard Fourier spectrum analysis into a form that depends on sample location differentials.
- A generalization of this basic formulation, including different distance transformations for various domains, and range selection for better control of quality and speed.
- Applications in spectral and spatial analysis for non-uniform sample distributions.

2 Previous Work

Spatial analysis Various spatial quantities have been proposed to measure the spatial distribution properties of samples [Dale et al. 2002]. One choice is the discrepancy [Shirley 1991], which is a scalar measure for sample equidistribution. Another is the relative radius $\rho = \frac{r_{min}}{r_{max}}$ defined in [Lagae and Dutré 2008], where r_{min} is the minimum spacing between any pair of samples, and r_{max} is the average inter-sample distance computed from the maximum packing of a given number of samples. Lagae and Dutré [2008] applied ρ to measure the spatial uniformity of blue noise samples and proposed that the ideal ρ should be between [0.65 0.85]: too small a ρ indicates lack of uniformity (e.g. white noise for which ρ is close to 0) and too large a ρ indicates potential regular patterns (e.g. hexagonal lattice as a result of unconstrained Lloyd relaxation). To our knowledge ρ has been applied only to uniform sampling so far, and we show in this paper how it can be extended to non-uniform sampling with our method.

Other spatial measures, especially in the context of meshing, include the minimum and histogram of triangle areas and angles [Fu and Zhou 2008]. Some of these measures can be applied for restricted non-uniform settings (e.g. angles for isotropic domains), while others cannot (e.g. angles for anisotropic domains or areas for adaptive domains).

Spectral analysis Fourier spectrum analysis is a common method for evaluating sample distributions and has been shown to be effective in detecting sampling artifacts [Ulichney 1987; Cook 1986; Mitchell 1987; Kopf et al. 2006; Ostromoukhov 2007; Lagae and Dutré 2008; Wei 2008; Wei 2010; Li et al. 2010; Bowers et al. 2010]. However, except for a few rare cases discussed below, Fourier analysis so far has been applied primarily to uniform Euclidean domain sampling. Li et al. [2010] analyzed anisotropic distributions by warping the samples back into a uniform domain followed by standard Fourier spectrum analysis. Thus, the method is applicable only to anisotropic distributions with analytically invertible global warps. Bowers et al. [2010] applied spectral mesh basis [Karni and Gotsman 2000] to analyze uniform surface samples, but their method has several practical limitations caused by numerical computation (such as it can only measure a few hundred samples). In addition, it is not applicable to adaptive or anisotropic sampling. To our knowledge, this work proposes the first method that can perform spectral analysis for general non-uniform sampling.

A possible alternative solution is to derive generalized Fourier basis to extend standard Fourier analysis to non-uniform domains. For example, under isotropic surface sampling, one could apply the adaptive spectral mesh basis presented in [Huang et al. 2008] to an-

alyze adaptive surface samples. This can be treated as an extension to [Bowers et al. 2010]. We are not aware of any prior publication using this approach. But even if the idea works, it is likely to suffer from similar practical limitations as [Bowers et al. 2010], due to the large number of numerical basis functions that must be computed. Moreover, it remains unclear how to extend such an approach for anisotropic sampling. In contrast, our differential domain method is general and does not rely on the existence of any Fourier basis set. It is also fast to compute and can easily handle a large number of samples. Furthermore, as shown in our paper, even for uniform sampling our method can outperform traditional Fourier analysis methods in quality or speed.

Other analysis methods In spatial statistics, several methods exist [Dale et al. 2002; Bonetti and Pagano 2005] to deduce the distribution property of samples via autocorrelation, inter-point distance, variogram, or the Ripley's K function [Ripley 1977]. These methods are closely related to our differential measurement approach, but the main differences are: 1) they are typically used to verify that a particular distribution conforms with uniform white noise; 2) they are only applicable to uniform domains while we focus on non-uniform domains. Another related work is the second moment measure discussed in [Lau et al. 2003]. However, the connection of this measure to the Fourier spectrum has not been revealed, and the method has not been applied for non-uniform sampling. Other related methods include non-uniform Fourier transform [Dutt and Rokhlin 1993; Potts et al. 2000] and local image analysis/statistics. To our knowledge these methods mainly focus on analyzing *range* (e.g. color) properties. This differs from our goal of analyzing *domain* (i.e. spatial location) properties.

3 Core Ideas

This section presents our core ideas. We first show that given a sample set, its Fourier power spectrum is completely determined by the histogram (i.e. distribution) $p(\mathbf{d})$ of the sample location differentials \mathbf{d} through a cos transform (Equation 5). Then, by replacing the cos kernel with a Gaussian kernel, we can directly evaluate a smoothed histogram of $p(\mathbf{d})$ (Equation 6). Finally, we show how the use of $p(\mathbf{d})$ allows us to extend the analysis to general non-uniform sample domains (Equation 9).

Fourier power spectrum Let $\{\mathbf{s}^k\}_{k=0}^{N-1}$ be a set of N samples in an n dimensional space. Its Fourier transform $F(\mathbf{f})$, with \mathbf{f} being the frequency vector, is defined as follows:

$$F(\mathbf{f}) = \sum_{k=0}^{N-1} e^{-2\pi i(\mathbf{f} \cdot \mathbf{s}^k)} \quad (1)$$

where \cdot denotes the vector inner product. For the purpose of analyzing sample patterns, we are interested in the power spectrum $P(\mathbf{f})$ [Ulichney 1987; Bracewell 1999; Lagae and Dutré 2008], which measures the squared magnitude of $F(\mathbf{f})$:

$$\begin{aligned} P(\mathbf{f}) &= |F(\mathbf{f})|^2 = P_r(\mathbf{f}) + P_i(\mathbf{f}) \\ P_r(\mathbf{f}) &= \frac{1}{N} \left(\sum_{k=0}^{N-1} \cos(2\pi \mathbf{f} \cdot \mathbf{s}^k) \right)^2 \\ P_i(\mathbf{f}) &= \frac{1}{N} \left(\sum_{k=0}^{N-1} \sin(2\pi \mathbf{f} \cdot \mathbf{s}^k) \right)^2 \end{aligned} \quad (2)$$

Differential representation Using the product to sum rule in trigonometry, we can rewrite Equation 2 into the following form:

$$P(\mathbf{f}) = \frac{1}{N} \sum_{k=0}^{N-1} \sum_{j=0}^{N-1} \cos(2\pi \mathbf{f} \cdot (\mathbf{s}^k - \mathbf{s}^j)) \quad (3)$$

An alternative way to derive this equation is by using discrete auto-correlation and the Wiener-Khinchin theorem [Couch 2001]. Note that unlike Equation 2 which depends on the absolute coordinates of samples, Equation 3 only depends on the relative coordinates between every two samples, which we call pair-wise sample location *differentials*. This is a key observation that allows us to generalize the analysis to non-uniform sampling. By representing $(\mathbf{s}^k - \mathbf{s}^j)$ as $\mathbf{d}^{k,j}$, we can express Equation 3 as:

$$P(\mathbf{f}) = \frac{1}{N} \sum_{j,k=0}^{N-1} \cos(2\pi \mathbf{f} \cdot \mathbf{d}^{j,k}) \quad (4)$$

Integral form As Equation 4 sums over all pairs of samples, given sufficient samples, we can collect the statistics of $\mathbf{d}^{k,j}$, and rewrite the equation as an integral form:

$$P(\mathbf{f}) = N \int_{\Omega_d} \cos(2\pi \mathbf{f} \cdot \mathbf{d}) p(\mathbf{d}) \delta \mathbf{d} \quad (5)$$

where $p(\mathbf{d})$ is the probability density function of \mathbf{d} , and Ω_d is the differential domain for \mathbf{d} derived from the underlying sample space Ω . Note that this integral is simply a cosine transform of $p(\mathbf{d})$. This equation means that given any sample set, its Fourier power spectrum is completely determined by $p(\mathbf{d})$ via a cosine transform. Thus it bridges the fundamental connection between the samples' spatial statistics and their spectral properties.

General kernel Since Equation 5 is essentially a weighted integral of $\cos(2\pi \mathbf{f} \cdot \mathbf{d})$ by $p(\mathbf{d})$, we can generalize it with an arbitrary kernel κ as follows:

$$P(\mathbf{q}) = N \int_{\Omega_d} \kappa(\mathbf{q}, \mathbf{d}) p(\mathbf{d}) \delta \mathbf{d} \quad (6)$$

where \mathbf{q} represents a general parameter set for κ . For example, in Fourier transform \mathbf{q} stands for the frequency. By picking a different kernel κ , we can choose to focus on analyzing different aspects of the sample distribution. When κ is a cos function, Equation 6 becomes traditional Fourier spectrum. In our analysis, we will choose a Gaussian kernel that essentially performs a kernel density estimation [Scott 1992] of $p(\mathbf{d})$. For notational simplicity, we will stick to the P symbol even though it might not actually measure "power" when a non-cos kernel is used.

Range selection Instead of taking into account all \mathbf{d} pairs as in Equation 6, for computational reasons we can also look at selected ranges of \mathbf{d} values. In general, we can incorporate a range selection function ξ into Equation 6 as follows:

$$P(\mathbf{q}) = N \int_{\Omega_d} \kappa(\mathbf{q}, \mathbf{d}) p(\mathbf{d}) \xi(\mathbf{d}) \delta \mathbf{d} \quad (7)$$

One common approach is to look at only $\mathbf{d} = \mathbf{s} - \mathbf{s}'$ with both \mathbf{s} and \mathbf{s}' within each other's local neighborhood:

$$\xi(\mathbf{d} = \mathbf{s} - \mathbf{s}') = \begin{cases} 1 & \text{if } \max(\mu(\mathbf{s}, \mathbf{s}'), \mu(\mathbf{s}', \mathbf{s})) \leq \epsilon \\ 0 & \text{else} \end{cases} \quad (8)$$

where $\mu(\mathbf{s}, \mathbf{s}')$ is a local distance measure of \mathbf{s}' with respect to the local frame centered at \mathbf{s} , and ϵ is a global constant controlling the local neighborhood size. In uniform Euclidean domains, $\mu(\mathbf{s}, \mathbf{s}') = |\mathbf{s} - \mathbf{s}'|$. For normalization, we usually divide μ by a constant which depends on the desired sampled density. For example, the constant may be r_{max} from the maximum packing of a give

number of samples. We will provide the specific formulations of μ for other scenarios in Section 5.1, including adaptive, anisotropic, and surface domains. The value of ϵ will depend on particular application needs, including computation speed (it is faster to consider only a shorter range \mathbf{d} values) and quality (e.g. detecting anomalies at a certain \mathbf{d} range). It is also possible to define ξ in other ways, e.g. with a smooth instead of binary profile.

Non-uniform domain For non-uniform domains, we can further extend Equation 7 as follows:

$$P(\mathbf{q}) = N \int_{\Omega_d} \kappa(\mathbf{q}, \chi(\mathbf{s}, \mathbf{s}', \mathbf{d})) p(\mathbf{d}) \xi(\mathbf{d}) \delta \mathbf{d} \quad (9)$$

where $\chi : \Omega_d \rightarrow \Omega_d$ is a differential domain transformation function that locally warps each \mathbf{d} from a non-uniform Ω_d to a (hypothetical) uniform Ω_d . By choosing the proper χ , κ , and ξ , we can analyze transformed domain properties for non-uniform sample distributions. Details will be presented in Section 5.1.

4 Analysis

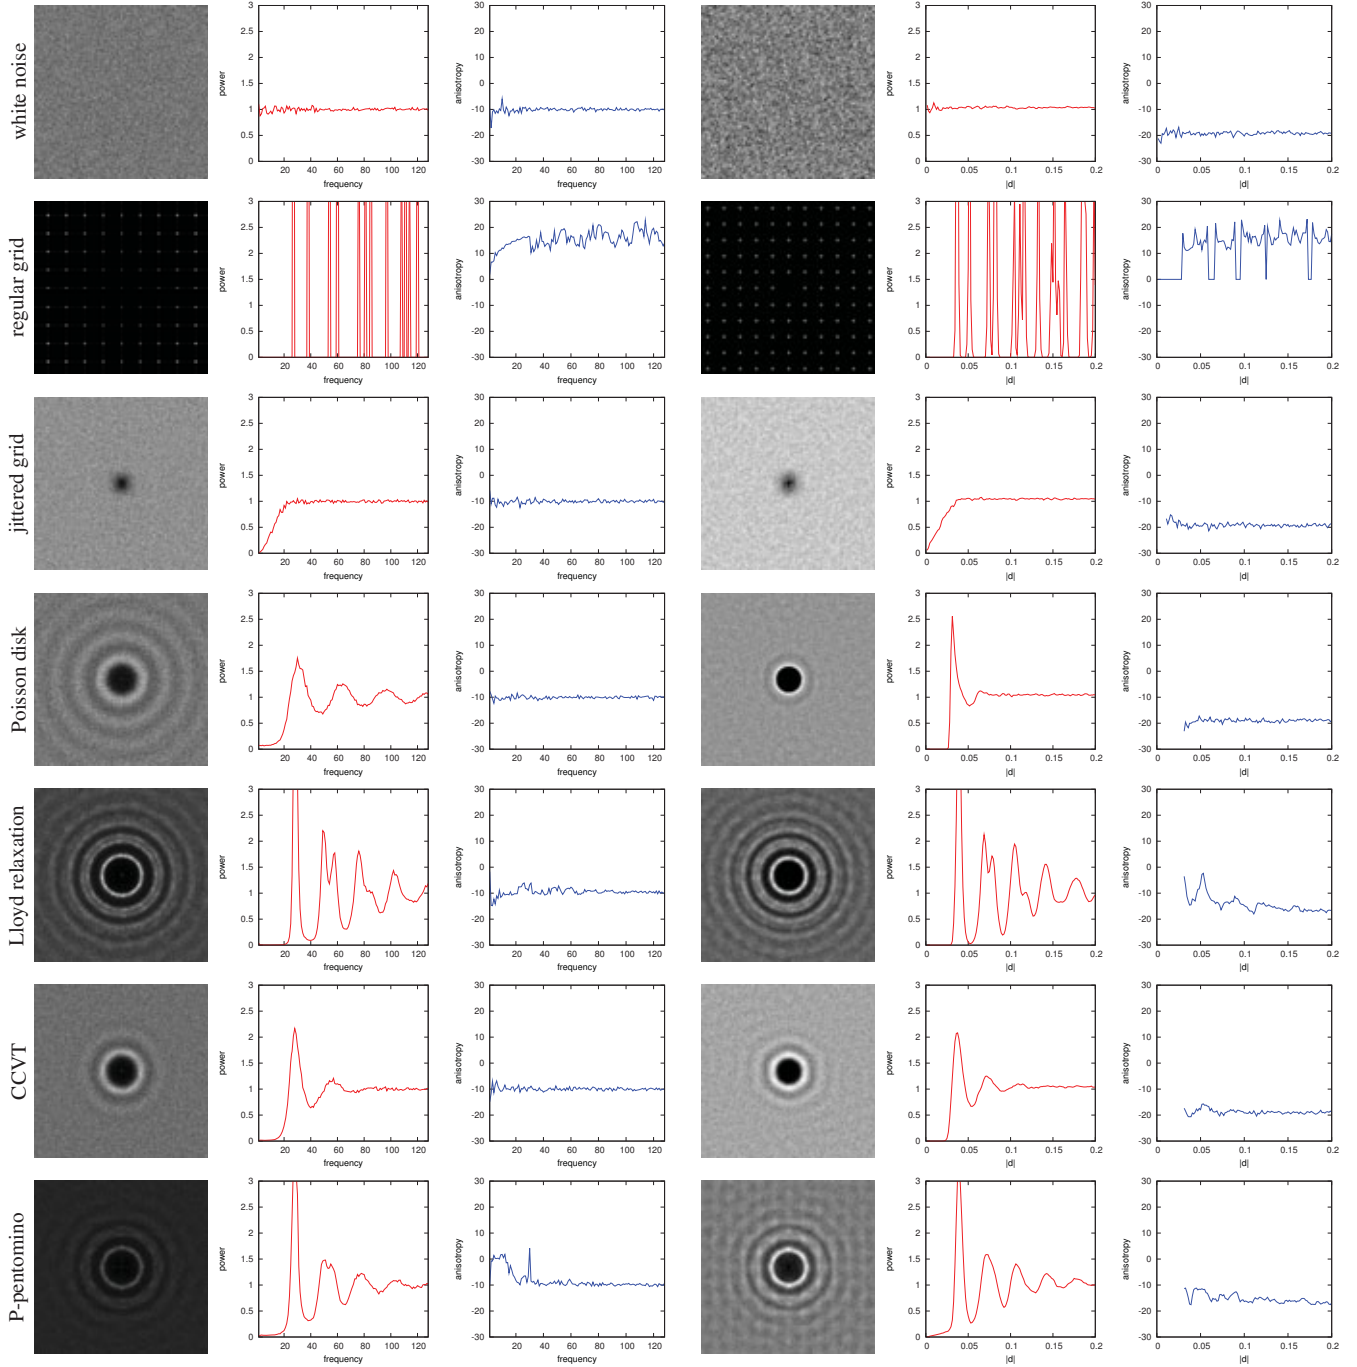
Comparisons with Fourier spectrum analysis The first question we need to answer is: how our method relates and compares to traditional Fourier spectrum analysis? As shown in Section 3, Fourier spectrum analysis can be considered as a special case of our method with a cos kernel in Equation 6. In Figure 2, we compare our method with Fourier spectrum analysis under a variety of different sampling methods, including white noise (uniform random), regular grid, jittered grid [Cook 1986], Poisson disk [Cook 1986], Lloyd relaxation [Lloyd 1983], CCVT [Balzer et al. 2009], and Polyominoes [Ostromoukhov 2007]. We choose these methods because they are common and possess a diverse repertoire of properties that are well known in the sampling literature. Since Fourier analysis is not applicable to general non-uniform domains, we use uniform domain sampling for this comparison to obtain ground truth information. As shown, our method relates well with Fourier spectrum analysis, capturing the key properties of all the methods. These include: the uniform and flat distribution of white noise; regular spikes in regular grid sampling; blue noise with small and large central "clear zones" for jittered and Poisson-disk samplings, respectively; potential hexagonal bias in Lloyd relaxation; a better relaxation method by CCVT; and the potential structured patterns in Polyomino sampling.

Kernel selection Our basic formulation in Equation 6 allows different choices for the computation kernel κ . We usually choose a Gaussian kernel defined as:

$$\kappa(\mathbf{q}, \mathbf{d}) = e^{-\frac{(\mathbf{q}-\mathbf{d})^2}{\sigma^2}} \quad (10)$$

where σ is the standard deviation and \mathbf{q} the query location differential. For a Gaussian kernel, we choose the output resolution so that on average each grid cell receives at least one sample, and we set σ to be 1 cell wide to ensure enough coverage and smoothness.

In contrast, a cos kernel may amplify some information while obscure others. As an example of the amplification, let's look at the Poisson disk samples in Figure 2. In this example, notice that $p(\mathbf{d})$ contains a peak value at $|\mathbf{d}| = r_{min}$, indicating a high probability density that the distance between two samples is around r_{min} . Next, the Fourier spectrum shows the well-known undulation profile, which is in fact caused by the peak value of $p(\mathbf{d})$. Specifically, from Equation 5 we can see that a cos kernel will produces peaks around $\mathbf{f} = \frac{m}{r_{min}}$ and troughs around $\mathbf{f} = \frac{m+0.5}{r_{min}}$ for any positive integer m . A Gaussian kernel, in contrast, displays the main peak clearly without undulations. Therefore it can manifest the distribution properties more clearly, e.g. more apparent characteristic



Fourier spectrum analysis

our method with Gaussian kernel

Figure 2: Comparisons between traditional Fourier spectrum analysis and our method under uniform domain with different sampling methods. Within each group are the spectrum image, radial mean, and anisotropy. The left group shows results generated by Fourier spectrum analysis while the right group by our method with a Gaussian kernel and $\epsilon = 10$. From top to bottom: white noise, regular grid, jittered grid, Poisson disk with $r_{\min} = 0.03$, Lloyd relaxation, CCVT [Balzer et al. 2009], and P-pentomino [Ostromoukhov 2007]. All cases are averaged over 10 runs each with ~ 720 samples. For several blue noise patterns the anisotropy values of our method within the range of r_{\min} are not shown as these regions are empty (i.e. no sample pairs fall within this range).

structures for CCVT as in Figure 2. Another advantage of using a Gaussian kernel is that it always returns a non-negative result for Equation 6 regardless of the \mathbf{d} distribution $p(\mathbf{d})$ and range selection. A cos kernel, in contrast, might produce negative values if not all \mathbf{d} pairs are considered.

Range selection The range selection function ξ (and the neighborhood size ϵ) plays an important role in deciding the speed and

quality of our method. Regarding speed, collecting all \mathbf{d} pairs can be expensive as it incurs an $O(N^2)$ time cost for N samples. Regarding quality, for general non-uniform sampling there might not exist an accurate χ for long range \mathbf{d} values. To resolve both issues, we have found it beneficial to consider only local \mathbf{d} pairs as in Equation 8. As shown in Figure 2, $\epsilon = 10$ already suffices to capture the major characteristics of common sample patterns. In our experience ϵ in the range of 10~12 works well in practice.

In theory a local range selection may miss long range structural artifacts. In practice, the impact of such artifacts is usually small, as they could contain no more than $\epsilon^{-n}N$ samples in an n dimensional space. For example, with $\epsilon = 10$ and $n = 2$, they can occupy no more than 1% of the total samples. To illustrate this, consider the pathological case where a regular grid is embedded within a Poisson disk sample set. This can be done by using a regular grid (with spacing d) as the initial samples, then performing Poisson disk sampling (with minimum spacing $r_{min} < d$) as usual. To detect such artifacts, our ϵ must satisfy $\epsilon r_{min} > d$. Thus, increasing d requires increasing ϵ as well. However, as d increases, the total number of offending samples and their contributions to $p(\mathbf{d})$ will decrease, reducing the total artifact energy. The theoretical maximum ratio ϵ^{-n} above is derived from assuming a maximal packing of d -spaced samples. Empirically, we have also found that our method is not worse than traditional Fourier analysis in detecting hidden structures; see the supplementary material for more details.

Computational complexity With N samples and F spectrum resolution in any given dimension, traditional Fourier power spectrum analysis has computational complexity $O(NF^n)$ in an n dimensional space Ω . This is assuming brute force computation in the continuous domain without any discrete approximation [Schlomer and Deussen 2010]. Although accelerations exist for computing non-uniform FFT (see e.g. [Potts et al. 2000]), they are mostly approximate methods. Using a Gaussian kernel, our method has complexity $O(N\epsilon^n\sigma^n)$, where σ is the Gaussian standard deviation and ϵ the range selection constant. Since $\epsilon\sigma$ is usually smaller than F , our method is computationally faster than traditional Fourier analysis. For example, in Figure 2 our method takes 1.4 seconds to accumulate each sample set while Fourier analysis takes 2.1 seconds, implying a 33% speedup using our method. The timing is measured on a PC with 2.5GHz CPU and 4GB memory. As a further acceleration, we can also select a uniform random subset of the \mathbf{d} pairs, which can still result in the same distribution without bias.

Radial measures Similar to standard Fourier spectrum analysis, we can compute the circular average and variance of $p(\mathbf{d})$. The former gives the radial mean, indicating the overall distance-based property of $p(\mathbf{d})$; the latter gives the anisotropy, which reveals if there is any directional bias/structure in the distribution.

Since $p(\mathbf{d})$ is evaluated based on a large number of \mathbf{d} pairs, we need to properly normalize it in order to compare different sample sets. Assume that we have computed $p(\mathbf{d})$ with M histogram bins from Q number of sample pairs. We normalize $p(\mathbf{d})$ with M/Q so that the expected value of each bin is 1. For white noise samples, this normalization will result in an expected value of 1 for the radial mean, as shown in the first sample of Figure 2.

To derive the expected anisotropy, we need to analyze the circular variance of $p(\mathbf{d})$. Let's again use white noise as an example: in this case, each bin of $p(\mathbf{d})$ has an independent value that can be modeled as a binomial distribution where the success probability of each trial is $\rho = 1/M$. Therefore the expected anisotropy can be calculated as the variance of this binomial distribution, which turns out to be $(M-1)/Q$, or $10 \log_{10}[(M-1)/Q]$ dB equivalently. Details of the derivation can be found in the supplementary material. Note that unlike traditional Fourier spectrum analysis, the anisotropy here is not directly related to the number of test runs.

If a general kernel is employed, such analysis may not be available. In that case, we can empirically compute reference radial mean and anisotropy from a white noise sampling of the same domain with identical parameters.

5 Applications

Now we present several applications of our core ideas in Section 3.

5.1 Spectral Analysis

Here, we apply Equation 9 to analyze non-uniform sample distributions, including adaptive, anisotropic, and non-Euclidean domains. Our key idea is to derive a proper χ and μ considering only $\mathbf{d} = s - s'$ that satisfies ξ (see discussions around Equation 7).

Exact χ For special circumstances where the transform function χ can be derived exactly, we can apply it directly to Equation 9. For example, if the output domain is warped from a uniform domain via a (forward) warp function φ [Li et al. 2010], we have

$$\chi(s, s', s - s') = \varphi^{-1}(s) - \varphi^{-1}(s') \quad (11)$$

However, such exact χ is usually not available for general non-uniform domains, for which we discuss solutions below.

Isotropic Euclidean domain For isotropic but adaptive sample domains with a scalar importance field $I(\cdot)$, we design χ as follows:

$$\chi(s, s', s - s') = \frac{2E(r)}{r(s) + r(s')}(s - s') \quad (12)$$

where $r(\cdot)$ is the distance field derived from $I(\cdot)$ with $r(\cdot) \propto I(\cdot)^{\frac{1}{n}}$ in a n -dimensional sample space Ω , and $E(r)$ is the mean of $r(\cdot)$ over Ω (derived through $I(\cdot)$). Equation 12 essentially performs a local scaling of $\mathbf{d} = s - s'$ via $r(\cdot)$.

Since $r(\cdot)$ is defined locally for every point, we can define the local distance measure μ as follows:

$$\mu(s, s') = \frac{|s - s'|}{r(s)} \quad (13)$$

Combining this equation with Equation 8, we have

$$\xi(\mathbf{d} = s - s') = \begin{cases} 1 & \text{if } |s - s'| \leq \epsilon \times \min(r(s), r(s')) \\ 0 & \text{else} \end{cases} \quad (14)$$

Now that χ , μ , and ξ are properly defined, we can enable our method for isotropic adaptive sampling. Figure 1 and Figure 3 show our results for a variety of isotropic sample domains, including Gaussian ramp, Gaussian blob, the density functions from [Ostromoukhov 2007] and [Balzer et al. 2009], and a complex bitmap image of van Gogh's portrait. For sampling method we choose white noise and Poisson disk because they are both common, have well-known characteristics, and have simple algorithms to compute adaptive samples. As shown, our method not only captures the known characteristics of the sampling methods (e.g. flat and peaky radial mean profiles for white-noise and Poisson-disk respectively) but also maintains consistent behavior across different sample domains, including: 1) the Gaussian ramp that has distinct distributions in the vertical and horizontal directions; 2) the Gaussian blob with radially symmetric \mathbf{d} distribution; 3) the density functions in [Balzer et al. 2009; Ostromoukhov 2007] with complex structures; 4) the van Gogh portrait with significant complexity.

Anisotropic Euclidean domain For anisotropic domains, χ will involve a general local Jacobian J [Li et al. 2010] as follows:

$$\chi(s, s', s - s') = \frac{1}{E(\lambda)} \left(\frac{J^{-1}(s) + J^{-1}(s')}{2} \right)^{-1} (s - s')^T \quad (15)$$

where $E(\lambda)$ is the mean of the eigenvalues of $J(\cdot)$ over Ω . Notice that we use the mean of J^{-1} instead of J so that Equation 15 will reduce to Equation 12 for isotropic cases, i.e.

$$J(s) = I \cdot \frac{1}{r(s)} \quad (16)$$

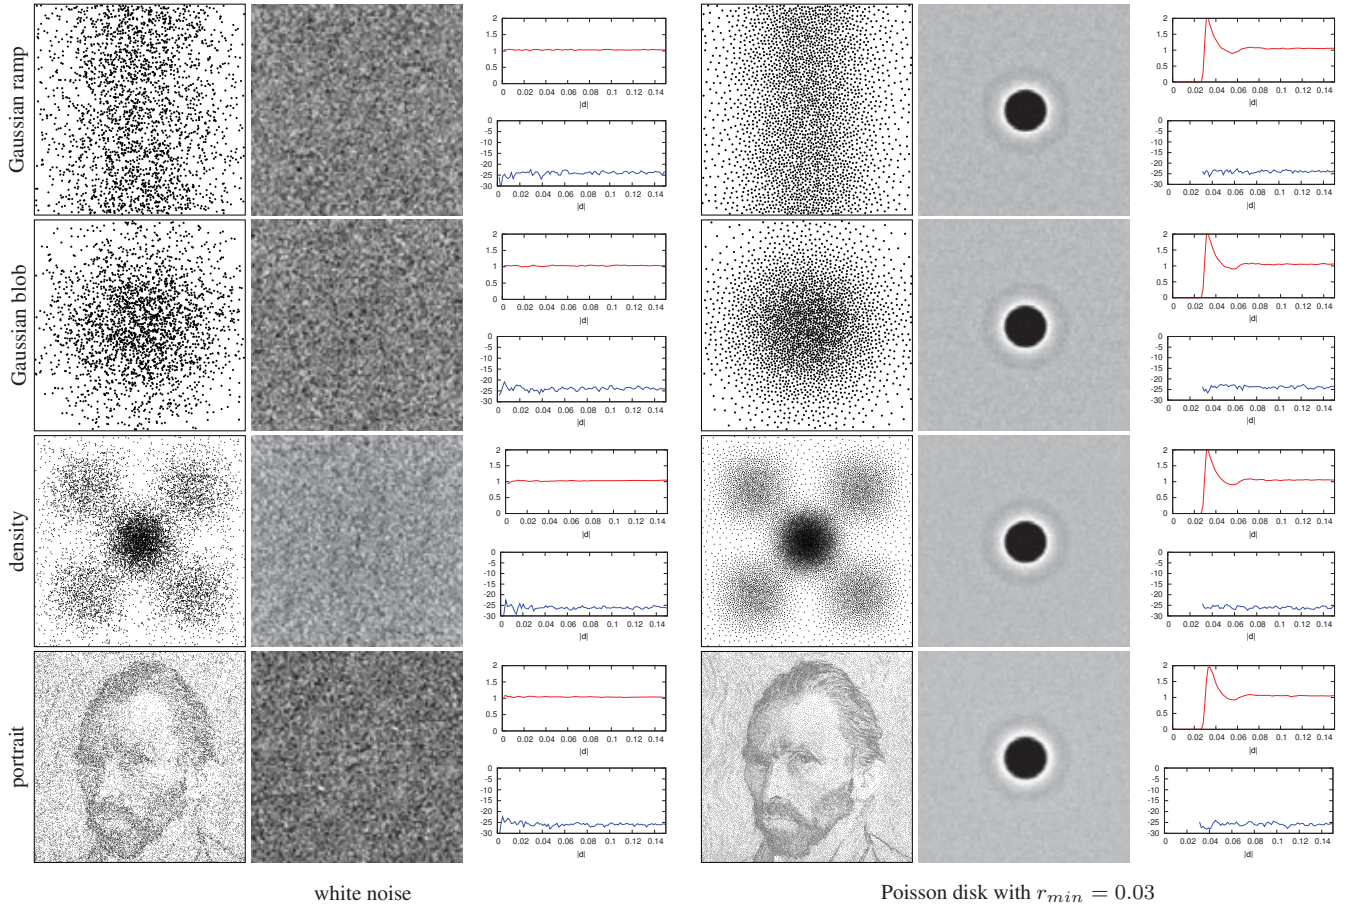


Figure 3: Isotropic adaptive sampling analysis via our method. The left group is generated by white noise while the right group by Poisson disk sampling. Within each group are the spatial sample set, spectrum image, radial mean, and radial anisotropy. From top to bottom: 1D Gaussian ramp, 2D Gaussian blob, the density function $e^{-20(x^2+y^2)} + 0.2 \sin^2(\pi x) \sin^2(\pi y)$ in [Balzer et al. 2009], and the van Gogh portrait. The ramp and blob cases are produced with 10 sets with ~ 2500 samples per set in a $[0, 2]^2$ domain, the density case with 4 sets with ~ 10000 samples per set in a $[0, 4]^2$ domain, and the portrait case with 1 set with ~ 40000 samples in a $[0, 8]^2$ domain. $\epsilon = 12$.

where I is the identity (Jacobian) matrix. Following the methodology in [Li et al. 2010], we can extend Equation 8 for anisotropic sampling as follows:

$$\mu(s, s') = |J(s)(s - s')^T| \quad (17)$$

$$\xi(\mathbf{d} = s - s') = \begin{cases} 1 & \text{if } \max(|J(s)(s - s')^T|, |J(s')(s - s')^T|) \leq \epsilon \\ 0 & \text{else} \end{cases} \quad (18)$$

It can be easily verified that Equation 18 reduces to Equation 14 with isotropic Jacobians (Equation 16).

Figure 4 demonstrates our results for anisotropic sampling, including four basic warps: scale, shear, perspective, and terrain [Wolberg 1994] as analyzed in [Li et al. 2010], and a more complex anisotropic field derived from a bitmap image. Similar to the isotropic case, we use white noise and Poisson disk distributions because they both have well-known statistical properties and can be produced over general anisotropic domains via methods in [Li et al. 2010]. As shown in Figure 4, our method well captures the properties of both white noise and Poisson disk sampling over a variety of anisotropic domains. Note that the approach by [Li et al. 2010] requires warping the anisotropic samples to a uniform domain and perform traditional Fourier analysis there. Our method, in contrast, can analyze samples directly in the anisotropic domain, without requiring analytic invertible warps. Thus our method is more general

than [Li et al. 2010] and can be applied to general anisotropic domains such as Jacobian fields derived from bitmap images.

Sampling rate As described in [Li et al. 2010], their anisotropic sampling methods approximate geodesic distances via local Jacobians and thus require sufficient sampling rates to work well. In Figure 5, we apply our methods to situations for which this requirement is not observed. We use the sinusoidal terrain warp from Figure 4 as it is both non-trivial to sample and has analytic inverse warps so we can apply Fourier analysis to generate comparisons. As shown in Figure 5, both Fourier analysis and our method can detect nuanced artifacts when the sampling rate is insufficient: note the anisotropic ripples in the Fourier spectrum image and anisotropic boundaries of the inner hole in the result of our method. However, at 2500 samples the anisotropy profile of Fourier analysis starts to become flat, whereas our method can still detect a high amount of anisotropy near $|\mathbf{d}| = 0.03$. At sufficiently high sampling rate such as used in Figure 4, both methods will result in flat anisotropy. This echoes our observation in Section 4 that our method may be able to manifest artifacts more clearly than Fourier spectrum analysis.

Figure 5 also illustrates the potential inaccuracy of our χ for handling long range \mathbf{d} values. As shown in the 625 sample case, the combination of $r_{min} = 0.03$ and $\epsilon = 12$ results in a total neighborhood size of 0.36. This is large enough to straddle over regions with significant changes in the terrain sinusoidal, causing inaccuracy in the long range χ values manifested by the darker corners in our

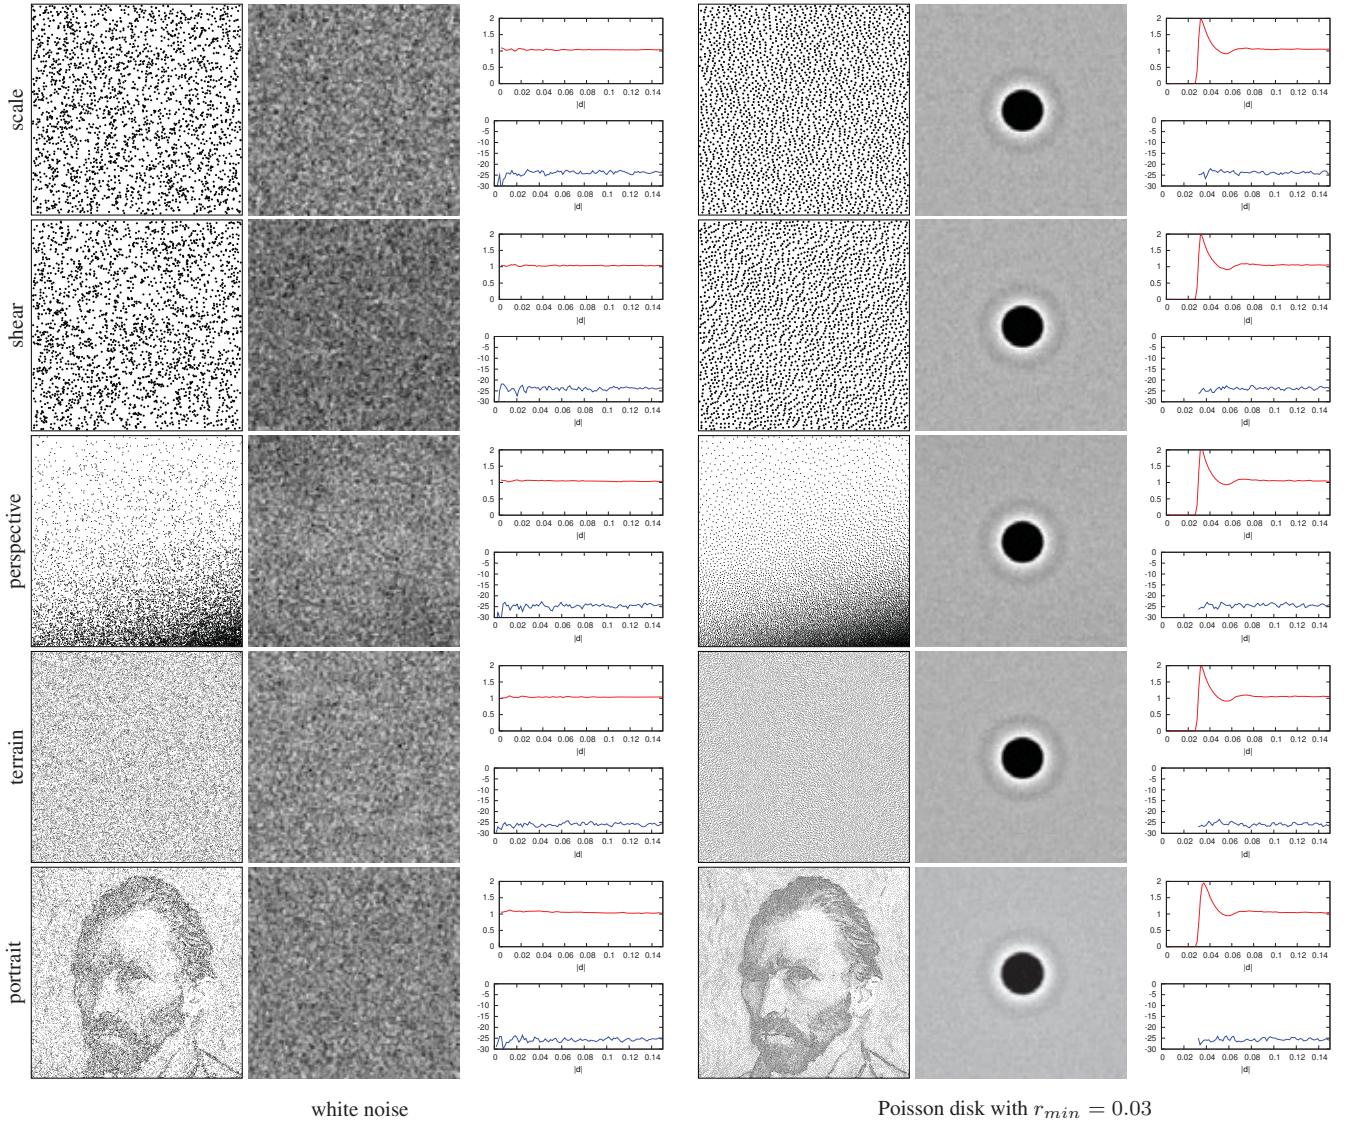


Figure 4: Anisotropic sampling analysis via our method. The left group is generated by white noise while the right group by Poisson disk sampling. Within each group are the spatial sample set, spectrum image, radial mean, and radial anisotropy. From top to bottom are the 4 warps from [Li et al. 2010]: scale, shear, perspective, and terrain, as well as the van Gogh portrait image. The scale and shear cases are produced with 10 sets with ~ 2500 samples per set in a $[0, 2]^2$ domain, the perspective case with 4 sets with ~ 10000 samples in a $[0, 4]^2$ domain, the terrain and portrait cases with 1 set with ~ 40000 samples in a $[0, 8]^2$ domain. $\epsilon = 12$.

image result. However, such long range artifacts only happen at insufficient sampling rate, which already causes short range artifacts that can be successfully detected by our method (see Figure 5).

Surface domain For surface sampling, an existing technique for analysis is via the spectral mesh basis [Bowers et al. 2010]. However, such a method is restricted to uniform surface sampling and can suffer from numerical computation issues that limit the maximum number of allowable samples to only a few hundred. Our method can be easily applied for analyzing surface samples without these restrictions. The basic idea is to use surface geodesics to measure $\mathbf{d} = s - s'$:

$$s - s' = \frac{g(s, s - s') + g(s', s - s')}{2} \quad (19)$$

where $g(s, s - s')$ is the geodesic differential $s - s'$ with respect to the local frame centered at s . For uniform surface sampling, this

allows us to rewrite the previous equations for χ , μ , ξ as:

$$\chi(s, s', s - s') = \frac{g(s, s - s') + g(s', s - s')}{2} \quad (20)$$

$$\mu(s, s') = \frac{|g(s, s - s')|}{r_{max}} \quad (21)$$

$$\xi(\mathbf{d} = s - s') = \begin{cases} 1 & \text{if } \max(|g(s, s - s')|, |g(s', s - s')|) \leq \epsilon r_{max} \\ 0 & \text{else} \end{cases} \quad (22)$$

For other non-uniform settings, such as adaptive and anisotropic surface sampling, we can similarly modify their corresponding equations defined for χ , μ , and ξ , by changing the distance metric from Euclidean to geodesic.

To fully define \mathbf{d} , we also need to provide a surface orientation field and a local parameterization. For the former, there is a rich literature in previous work on designing surface orientation fields with

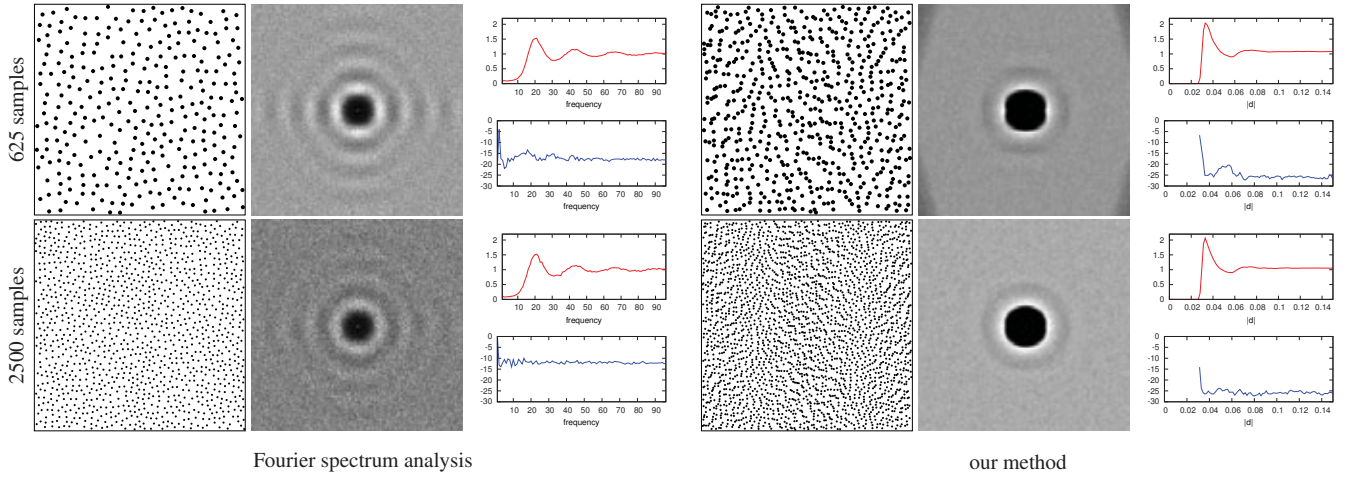


Figure 5: Analysis for different anisotropic sampling rates. Here, we apply the anisotropic dart throwing algorithm in [Li et al. 2010] to the terrain warp domain with different sampling rates: the top case is computed from 64 sets each with ~ 625 samples over a $[0, 1]^2$ domain whereas the bottom case from 16 sets each with ~ 2500 samples over a $[0, 2]^2$ domain. We set $r_{min} = 0.03$, and analyze the sample sets through both traditional Fourier spectrum analysis (on the left, by warping the samples back into a uniform domain as in [Li et al. 2010]) as well as our method (on the right, with $\epsilon = 12$). Within each group are: spatial samples (warped back into uniform domain for Fourier analysis), spectrum image, and radial mean/anisotropy.

desired properties [Palacios and Zhang 2007; Fisher et al. 2007]. Our method simply requires a sufficiently smooth orientation field that neither introduces nor obscures sampling artifacts. For example, a random orientation field may completely obscure any directional structure in the underlying samples, due to inconsistent local frames. In our current implementation, we follow the standard method of taking the cross product of the surface normal with a pre-defined global vector. This works well in practice. For defining local parameterization, our method requires a proper parameterization method that preserves local angles and geodesics. For this purpose, the discrete exponential map [Schmidt et al. 2006] is an excellent choice, as it is both fast and accurate. Specifically, we set the decal radius at each sample s according to ξ , then run the discrete exponential map algorithm to compute the geodesic coordinates (u, v) of every neighbor sample s' with respect to the center samples s . The vector (u, v) directly corresponds to $\mathbf{d} = s - s'$. The magnitude of the vector gives the geodesic distance between s and s' . Note that our approach does not require any global parameterization, therefore is feasible for arbitrary, unparametrized shapes.

In Figure 6 we show the results of our approach applied to a variety of models. These models are selected due to their diverse shapes and topological properties. We normalize the area of each model to be 1, and tested them with three different sampling methods: white noise, uniform, and non-uniform Poisson Disk sampling. White noise samples are generated with uniform random distribution on the surface. Since such sampling does not require estimating surface geodesics, it can be used to verify our approach. The top-left three examples in Figure 6 show the analysis of white noise samples. As expected, the radial mean and anisotropy plots are both flat, indicating a white noise distribution. Note that the use of discrete exponential map (or other proper local parametrization method) is critical: a naive approach, such as directly projecting the Euclidean vector, would cause the estimated $p(\mathbf{d})$ to be distorted either radially or angularly, leading to non-flat radial means or anisotropy.

For uniform Poisson Disk sampling, we set $r_{min} = 0.015$ and use the method presented in [Bowers et al. 2010] to compute samples for all models. This results in ~ 3000 samples per model. The analysis results are shown on the top-right and bottom-left of Figure 6. Despite each model’s different shape and topology, the radial mean and anisotropy plots are consistent and agree with the expectation.

Lastly, the bottom-right three examples in Figure 6 show non-uniform Poisson Disk sampling. We set $r(x) \propto \frac{1}{1+4\cdot\kappa^2(x)}$, where κ is the mean curvature, and we normalize all the $r(\cdot)$ values so that

their average (via the density space r^{-n}) is 0.015. As shown in the analysis results, our approach can successfully verify the distribution property of the samples, despite the spatial non-uniformity.

The method described here performs analysis directly on surfaces. Another possibility is to parameterize the surface into a 2D anisotropic domain [Li et al. 2010] and analyze sample distributions via our anisotropic method. However, this requires surface parametrization and the handling of chart boundaries, incurring significantly more complexity than a direct surface analysis approach.

5.2 Spatial Analysis

We can also easily apply our methods to extend the spatial uniformity measure to non-uniform domains. Specifically, the uniformity measure is defined as $\rho = \frac{r_{min}}{r_{max}}$ [Lagae and Dutré 2008], where r_{min} is the minimum spacing between any pair of samples and r_{max} is the average inter-sample distance computed from the maximum packing of a given number of samples. To handle non-uniform distributions, we simply compute r_{min} from the χ -transformed \mathbf{d} values as described in Section 5.1. Table 1 compares our method with directly measured ρ values without any χ correction. We use Poisson disk sampling as the benchmark as its r_{min} is known. As shown, our method accurately measures ρ across different non-uniform domains and distributions.

6 Limitations and Future Work

The main limitation of our method is that the requirement of a sufficiently accurate differential domain distance transform $\chi(s, s')$. Although it is sometimes possible to define such a transform for the entire domain Ω_d under special cases (e.g. analytical warps), for general non-uniform settings, obtaining an accurate χ is currently only possible for local sample pairs (s, s') . Thus, our method might not be able to capture the global \mathbf{d} properties, potentially missing the detection of certain long-distance biases. One potential future direction is to investigate the possibility of an accurate estimation of χ for arbitrary sample pairs under general non-uniform settings.

However, as shown in the paper, both theoretical and empirical evidences suggest that our method is no less discriminative than traditional Fourier analysis. In fact, for most sampling algorithms we are interested in, sample pairs over long distances tend to be uncorrelated, thus the local \mathbf{d} statistics is usually sufficient to capture the characteristics of these sampling patterns. Moreover, var-




Figure 6: Surface sampling analysis via our method. Here we apply our analysis on samples generated on different surfaces. Each case is produced from 8 sets with ~ 3000 samples each. We set $r_{min} = 0.015$ for Poisson disk sampling, and use $\epsilon = 10$ in all examples.

2D case		direct ρ	our ρ	surface case (Figure 6)		
Figure 1	density	0.32	0.70	uniform	eight	0.76
Figure 3	ramp	0.51	0.70		bunny	0.69
	blob	0.37	0.70		hand	0.72
	density	0.22	0.70		venus	0.74
	portrait	0.44	0.70		genus3	0.70
					head	0.66
Figure 4	scale	0.50	0.70	adaptive	venus	0.15
	shear	0.43	0.70		genus3	0.16
	perspective	0.09	0.70		head	0.16
	terrain	0.31	0.70			
	portrait	0.17	0.70			

Table 1: Spatial measure ρ for non-uniform sampling. We compare ρ computed directly vs. using our method. We use Poisson disk sampling as its r_{min} is known. The sample distributions are produced with a target $\rho = 0.70$ for the 2D cases and $\rho = 0.76$ for the surface cases; thus the measured ρ values cannot exceed these. As shown, our method accurately measures ρ across a variety of sample domains and distributions, while direct measurements (including both Euclidean and uncorrected geodesic distances for the surface case) usually under-estimate the ρ values.

ious prior sampling techniques have exploited local information to successfully produce high quality samples. Examples include dart throwing, relaxation, tiling [Kopf et al. 2006], and spectral sampling [Öztireli et al. 2010]. We believe our differential domain analysis will shed light on why these methods work well, and the intrinsic nature of these methods.

Another direction for future research is to investigate other computation kernels κ . For example, there may be scenarios that prefer kernels with multiresolution support, such as wavelets.

Finally, our idea can potentially benefit other applications, in addition to the ones shown in this paper. Examples include analyzing higher dimensional sample patterns, and designing novel sampling algorithms that can synthesize samples with desired distribution properties at fast computation speed.

Acknowledgements We would like to thank Hongwei Li for clarifying details in [Li et al. 2010], and SIGGRAPH anonymous reviewers for their suggestions. This work is supported in part by NSF grant CCF-0746577.

References

- ALLIEZ, P., MEYER, M., AND DESBRUN, M. 2002. Interactive geometry remeshing. In *SIGGRAPH '02*, 347–354.
- BALZER, M., SCHLOMER, T., AND DEUSSEN, O. 2009. Capacity-constrained point distributions: A variant of Lloyd's method. In *SIGGRAPH '09*, 86:1–8.
- BONETTI, M., AND PAGANO, M. 2005. The interpoint distance distribution as a descriptor of point patterns, with an application to spatial disease clustering. *Statistics in Medicine* 24, 5, 753–73.
- BOWERS, J., WANG, R., WEI, L.-Y., AND MALETZ, D. 2010. Parallel Poisson disk sampling with spectrum analysis on surfaces. In *SIGGRAPH Asia '10*, 166:1–10.
- BRACEWELL, R. 1999. *The Fourier Transform and Its Applications*. McGraw-Hill.
- CLARBERG, P., JAROSZ, W., AKENINE-MÖLLER, T., AND JENSEN, H. W. 2005. Wavelet importance sampling: efficiently evaluating products of complex functions. In *SIGGRAPH '05*, 1166–1175.
- COOK, R. L. 1986. Stochastic sampling in computer graphics. *ACM Trans. Graph.* 5, 1, 51–72.
- COUCH, II, L. W. 2001. *Digital and Analog Communication Systems*, 6th ed. Prentice Hall, New Jersey.
- DALE, M. R. T., DIXON, P., FORTIN, M.-J., LEGENDRE, P., MYERS, D. E., AND ROSENBERG, M. S. 2002. Conceptual and mathematical relationships among methods for spatial analysis. *ECOGRAPHY*, 25, 558–577.
- DIPPÉ, M. A. Z., AND WOLD, E. H. 1985. Antialiasing through stochastic sampling. In *SIGGRAPH '85*, 69–78.
- DUTRE, P., BALA, K., AND BEKAERT, P. 2002. *Advanced Global Illumination*. A. K. Peters, Ltd., Natick, MA, USA.
- DUTT, A., AND ROKHLIN, V. 1993. Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Comput.* 14, 6, 1368–1393.
- FISHER, M., SCHRÖDER, P., DESBRUN, M., AND HOPPE, H. 2007. Design of tangent vector fields. In *SIGGRAPH '07*, 56.
- FU, Y., AND ZHOU, B. 2008. Direct sampling on surfaces for high quality remeshing. In *SPM '08*, 115–124.
- GLASSNER, A. S. 1994. *Principles of Digital Image Synthesis*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- HUANG, J., ZHANG, M., MA, J., LIU, X., KOBELT, L., AND BAO, H. 2008. Spectral quadrangulation with orientation and alignment control. In *SIGGRAPH Asia '08*, 147:1–9.
- KARNI, Z., AND GOTSMAN, C. 2000. Spectral compression of mesh geometry. In *SIGGRAPH '00*, 279–286.
- KOPF, J., COHEN-OR, D., DEUSSEN, O., AND LISCHINSKI, D. 2006. Recursive Wang tiles for real-time blue noise. In *SIGGRAPH '06*, 509–518.
- LAGAE, A., AND DUTRÉ, P. 2008. A comparison of methods for generating Poisson disk distributions. *Computer Graphics Forum* 21, 1, 114–129.
- LAU, D., ULICHNEY, R., AND ARCE, G. 2003. Fundamental characteristics of halftone textures: blue-noise and green-noise. *IEEE Signal Processing Magazine* 20, 4 (July), 28–38.
- LI, H., WEI, L.-Y., SANDER, P., AND FU, C.-W. 2010. Anisotropic blue noise sampling. In *SIGGRAPH Asia '10*, 167:1–12.
- LLOYD, S. 1983. An optimization approach to relaxation labeling algorithms. *Image and Vision Computing* 1, 2.
- MITCHELL, D. P. 1987. Generating antialiased images at low sampling densities. In *SIGGRAPH '87*, 65–72.
- OSTROMOUKHOV, V., DONOHUE, C., AND JODOIN, P.-M. 2004. Fast hierarchical importance sampling with blue noise properties. In *SIGGRAPH '04*, 488–495.
- OSTROMOUKHOV, V. 2007. Sampling with polyominoes. In *SIGGRAPH '07*, 78.
- ÖZTIRELI, A. C., ALEXA, M., AND GROSS, M. 2010. Spectral sampling of manifolds. In *SIGGRAPH ASIA '10*, 168:1–8.
- PALACIOS, J., AND ZHANG, E. 2007. Rotational symmetry field design on surfaces. In *SIGGRAPH '07*, 55.
- PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann Publishers Inc.
- POTTS, D., STEIDL, G., AND TASCHE, M., 2000. Fast Fourier transforms for nonequispaced data: A tutorial.
- RIPLEY, B. D. 1977. Modeling spatial patterns. *J. R. Stat. Soc.* 39, 2, 172–212.
- SCHLOMER, T., AND DEUSSEN, O. 2010. Towards a standardized spectral analysis of point sets with applications in graphics. Tech. rep., University of Konstanz.
- SCHMIDT, R., GRIMM, C., AND WYVILL, B. 2006. Interactive decal compositing with discrete exponential maps. In *SIGGRAPH '06*, 605–613.
- SCOTT, D. W. 1992. *Multivariate density estimation: theory, practice, and visualization*. Wiley.
- SHIRLEY, P. 1991. Discrepancy as a quality measure for sample distributions. In *Eurographics '91*, 183–194.
- TURK, G. 1992. Re-tiling polygonal surfaces. In *SIGGRAPH '92*, 55–64.
- ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press, Cambridge, MA.
- WEI, L.-Y. 2008. Parallel Poisson disk sampling. In *SIGGRAPH '08*, 20:1–9.
- WEI, L.-Y. 2010. Multi-class blue noise sampling. In *SIGGRAPH '10*, 79:1–8.
- WOLBERG, G. 1994. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA, USA.

Supplementary Material

(For the electronic version of the paper.)

A Hierarchical Warping

Here we provide evidence where a sampling method that behaves well in the uniform domain might not remain so in a non-uniform domain. For this, we use the hierarchical warping method in [Clarberg et al. 2005] which produces adaptive sample patterns by warping an initial set of uniform samples. The method performs well in uniform sampling because in that case it would not alter the initial samples. However, for adaptive sampling (under a given importance field), the hierarchical warping can easily introduce significant anisotropy. Such artifacts cannot be detected by existing methods as they only work in uniform domains. Our method, in contrast, can easily detect such artifacts, as demonstrated in Figure 7. Note that [Clarberg et al. 2005] is not specifically designed for blue noise sampling: we employ it here as a simple example to demonstrate that analysis performed in the uniform domain does not necessarily transfer to the behavior of an algorithm in the non-uniform domain.




Figure 7: Analysis for hierarchical warping. We perform our differential domain analysis to both adaptive Poisson disk sampling (left) and hierarchical warping [Clarberg et al. 2005] (right) using a very simple importance field as in Figure 5(b) of [Clarberg et al. 2005] (with relative importance 0.6, 0.2, 0.05, 0.15 for each quarter cell). Notice the high amount of anisotropy in our analysis result for [Clarberg et al. 2005] even though it starts with a reference initial sample pattern produced by uniform Poisson disk sampling (right).

B Radial Measures

Here we give a theoretical analysis of the expected anisotropy in the radial measures of our approach. We use white noise as an example as it is easy to analyze. First, note that the anisotropy measures the circular variance, and it should relate to the number of sample pairs Q . In particular, as Q approaches infinity, the distribution function $p(\mathbf{d})$ becomes flat, and hence the variance on any concentric circle will approach zero. This is the similar to averaging periodograms in Fourier spectrum analysis, as doing so will reduce the variance of values. Next, the anisotropy should also relate to the number of histogram bins M that we use to represent $p(\mathbf{d})$, because using less bins is equivalent to averaging sample values.

To derive the anisotropy of $p(\mathbf{d})$, consider that at any neighborhood, the distribution of nearby samples is white noise, hence the distances between them are also white noise. Thus collecting Q pair values and place them in M disjoint bins can be modeled as a binomial process: at any particular bin, a random pair value has a probability of $\rho = \frac{1}{M}$ landing in that bin, and all bins are independent. Thus for each bin, there is a probability of

$C(Q, i)\rho^i(1-\rho)^{M-i}$ that the total number of samples that fall into it is i , where $C(Q, i)$ denotes Q select i . The expectation of this binomial process is $Q \times \rho = Q/M$, hence we normalize the value of each bin by M/Q such that the expected value of each bin is 1. Therefore with probability of $C(Q, i)\rho^i(1-\rho)^{M-i}$ a bin has value $M/Q \cdot i$. It turns out that the variance of such a binomial variable is $(M^2/Q^2)Q\rho(1-\rho) = (M-1)/Q$. As each bin is independent, if we look at any set of bins (such as a circular set of bins), the variance of them is still the same. Thus $(M-1)/Q$ is the expected anisotropy for a white noise sample set.

The above analysis applies to the unmodified $p(\mathbf{d})$. To account for the smoothing caused by applying the Gaussian kernel, we multiply Q by a factor of 2.37 before using the above formula. Here is a brief justification. Since we set the standard deviation σ of the Gaussian kernel to be 1 cell wide, it will average primarily 3 adjacent bins: the center bin x_1 with a weight w , and the two neighbor bins x_0 and x_2 with weight $(1-w)/2$ each. Thus their average is $w \cdot x_1 + (1-w)/2 \cdot x_0 + (1-w)/2 \cdot x_2$. Assume the variance of each bin variable is v , then due to independence, the variance of their average is $(w^2 + 2[(1-w)/2]^2) \cdot v$. By plugging in the normalized Gaussian weight $w = \frac{1}{1+2e^{-1}}$ we have the factor 2.37.

C Explanation for Equation 12




Figure 8: Isotropic χ .

Here we provide an intuitive explanation on our formulation of Equation 12. As shown in Figure 8, assuming we have two samples s and s' for which we wish to compute $\chi(s, s', s - s')$. Let m be the mid-way point that partitions $s - s'$ in proportion to $r(s)$ and $r(s')$. Then intuitively we should adjust the two segments $m - s'$ and $s - m$ according to the ratios $\frac{E(r)}{r(s')}$ and $\frac{E(r)}{r(s)}$, respectively:

$$\begin{aligned} \chi(s, s', s - s') &= \frac{E(r)}{r(s')} (m - s') + \frac{E(r)}{r(s)} (s - m) \\ &= \frac{E(r)}{r(s')} \frac{r(s')}{r(s) + r(s')} (s - s') + \\ &\quad \frac{E(r)}{r(s)} \frac{r(s)}{r(s) + r(s')} (s - s') \\ &= \frac{2E(r)}{r(s) + r(s')} (s - s') \end{aligned} \quad (23)$$

D Explanation for Equation 15

Equation 15 is formulated by considering the fact that J provides a “local warp” from the anisotropic to a uniform domain [Li et al. 2010]; and we ensure that Equation 15 will reduce to Equation 12 under isotropic settings.

E Anisotropic Importance Jacobian

According to [Li et al. 2010], the Jacobian J derived from an importance image I may not be a square matrix. However, we need a square J for Equation 15 to preserve the dimension of \mathbf{d} . We overcome this issue by a simple observation: since the dart throwing algorithm in [Li et al. 2010] depends on only $J^T J$, not J itself, all

we need is to derive a square J' so that

$$J^T J = J'^T J' \quad (24)$$

This can be achieved by the standard matrix square root method:

$$J^T J = V^T D V \quad (25)$$

where V is an orthonormal matrix and D a diagonal matrix. Note that since $J^T J$ is positive definite, D will contain only non-negative diagonal elements. Thus

$$J' = V^T \sqrt{D} V \quad (26)$$

F Quantifying Blue Noise

Blue noise refers to a general class of sample distributions that exhibit the absence of low-frequency energy and the existence of broadband high-frequency energy in the power spectrum. As documented in previous work [Cook 1986], these properties generally lead to preferred visual effects in sampling. In [1987], Ulichney provided a qualitative description of blue noise, but no quantitative definition was given. To our knowledge there is no explicit quantitative measurement for blue noise in existing work. Here we provide such a definition using our method. Given the distribution $p(\mathbf{d})$ of a sample set, one simple way to quantify blue noise is to measure its difference with an ideal blue noise profile $b(\mathbf{d})$:

$$b(\mathbf{d}) = \begin{cases} 0 & \text{if } |\mathbf{d}| < r_{min} \\ 1 & \text{else} \end{cases} \quad (27)$$

where r_{min} is the desired minimum inter sample spacing. It can be shown that the Fourier power spectrum of this ideal profile b (by plugging b into Equation 5) satisfies the blue noise properties with additional undulations similar to a Poisson disk sample set (see the next section for more details). As discussed in Section 4, the undulation is caused by the nature of the \cos kernel instead of the inherent properties of the sample distribution. We thus prefer a flat profile for our definition of b as in Equation 27.

Given a $p(\mathbf{d})$ and our definition of $b(\mathbf{d})$ (both distributions are normalized), we can measure their discrepancy as follows:

$$\begin{aligned} \bar{\beta}(p, b) = & \frac{|\Omega_d|}{2|\Omega_c|} \int_{\Omega_c} |p(\mathbf{d}) - b(\mathbf{d})| \delta \mathbf{d} \\ & + \frac{|\Omega_d|}{2|\Omega_d - \Omega_c|} \int_{\Omega_d - \Omega_c} |p(\mathbf{d}) - b(\mathbf{d})| \delta \mathbf{d} \end{aligned} \quad (28)$$

where Ω_c is the inner core of Ω_d with radius r_{min} and $|\cdot|$ indicates the volume of the domain. Intuitively, $\bar{\beta}$ sums up the absolute differences between p and b , but we take care to normalize the terms within and outside Ω_c so that the result will not be impacted by the relative size of Ω_c and Ω_d . Thus, $\bar{\beta}$ will be in the range $[0, 2]$, where smaller values indicating closer resemblance to an ideal blue noise.

case	$\bar{\beta}$
CCVT	0.11
Poisson disk	0.12
Polymino	0.14
Lloyd relaxation	0.15
jittered grid	0.34
white noise	0.54
regular grid	0.90

Table 2 measures $\bar{\beta}$ for a variety of distributions from Figure 2. As shown, the measured $\bar{\beta}$ values reflect our intuitions about the *blueness* of various distributions quite well. By using other formulations of $b(\mathbf{d})$, we can quantify blue noise with different criteria

Table 2: Quantifying blue noise.
Here, we measure $\bar{\beta}$, the deviation of $p(\mathbf{d})$ from various distributions in Figure 2 with respect to an ideal blue noise distribution $b(\mathbf{d})$ with $r_{min} = 0.03$.

(e.g. having a peak around r_{min}), or noise with other spectral colors. Moreover, by obtaining an analytic representation of $p(\mathbf{d})$ and then applying Equation 5, it is possible to derive an analytic formula for the Fourier power spectrum of standard sample sets such as Poisson disk samples. This is because $p(\mathbf{d})$ is often much easier to derive (at least empirically) than the power spectrum itself, as it does not contain the undulation caused by the \cos kernel.

G Cosine Transform for b

Here we derive the cosine transform for our ideal blue noise profile b in Equation 27. First, to simplify the derivation, we define the complement function

$$b'(\mathbf{d}) = 1 - b(\mathbf{d}) = \begin{cases} 1 & \text{if } |\mathbf{d}| < r_{min} \\ 0 & \text{else} \end{cases} \quad (29)$$

Without loss of generality, let's assign $r_{min} = 1$ in subsequent derivations. Due to linearity, the cosine transform of b is simply a powered dirac delta function δ minus the cosine transform of b' . This is because the cosine transform of constant 1 is δ^n in n -dimensional space. Next, the cosine transform of b' is defined as

$$P(b') = \int_{\Omega} \cos(\omega \cdot \mathbf{d}) b'(\mathbf{d}) \delta \mathbf{d} = \int_{\mathbf{b}(0,1)} \cos(\omega \cdot \mathbf{d}) \delta \mathbf{d} \quad (30)$$

where $\mathbf{b}(0,1)$ is a unit n -dimensional ball centered at origin with radius 1. Due to radial symmetry/isotropy, it suffices to compute the result for any given frequency orientation. Without loss of generality, we pick the frequency vector $\omega = \omega(1, 0, 0, \dots)$, in other words, a vector that's along the first axis in n -dimensional space and has magnitude ω . Thus the integral we are computing is $\int_{\mathbf{b}(0,1)} \cos(\omega \cdot d_1) \delta \mathbf{d}$. Essentially, this is integrating function $\cos(\omega \cdot d_1)$ inside the unit n -dimensional ball. We can expand this multi-dimensional integral to a 1D integral as:

$$\int_{-1}^1 \cos(\omega \cdot d_1) V^{n-1}(\sqrt{1 - d_1 \cdot d_1}) \delta d_1 \quad (31)$$

where $V^{n-1}(\sqrt{1 - d_1 \cdot d_1})$ is the volume of an $(n-1)$ -dimensional ball with radius $\sqrt{1 - d_1 \cdot d_1}$. It is known that the volume of an n -dimensional ball is:

$$V^n(R) = \frac{\pi^{n/2}}{\Gamma(n/2 + 1)} R^n$$

where Γ is the gamma function, and R is the radius of the ball. Thus Equation 31 is further developed to:

$$\int_{-1}^1 \cos(\omega \cdot d_1) \frac{\pi^{(n-1)/2}}{\Gamma((n-1)/2 + 1)} (\sqrt{1 - d_1 \cdot d_1})^{(n-1)} \delta d_1$$

Using the property of the Γ function, it turns out that this 1D integral has analytic solutions. After normalization, the solution is:

$$\begin{cases} \frac{2^{n/2} (n/2)! \text{BesselJ}(n/2, \omega)}{\omega^{n/2}} & \text{if } n \text{ is even} \\ {}_0F_1\left[\frac{n}{2} + 1, -\frac{\omega^2}{4}\right] & \text{if } n \text{ is odd} \end{cases} \quad (32)$$

where BesselJ is the Bessel function of the first kind, $!$ denotes the factorial, and ${}_0F_1$ is a confluent hypergeometric function. Considering the complement relationship of b and b' , and expand ${}_0F_1$, we can express the cosine transform of b for the first 4 dimensions as:

$$\begin{cases} 1 - \frac{\sin \omega}{\omega} & n=1 \\ 1 - \frac{2 \text{BesselJ}(1, \omega)}{\omega} & n=2 \\ 1 - \frac{3(-\omega \cos \omega + \sin \omega)}{\omega^3} & n=3 \\ 1 - \frac{8 \text{BesselJ}(2, \omega)}{\omega^2} & n=4 \end{cases} \quad (33)$$

Figure 9 plots these four functions. Note that these curves directly correspond to the radial mean of the Fourier power spectrum of our ideal blue noise profile. Observe that they all exhibit the basic features of blue noise (with undulation), indicated by the lack of low-frequency energy and the existence of broadband high-frequency energy.




Figure 9: Cosine transform of b for the first 4 dimensions.

H Additional Measures

Here, we describe several additional measures on top of our core formulation (Section 3) that might complement our methods but have not been employed in our experimental results.

Measure similarity In addition to spatial proximity (Equation 8), it will also be good to consider local *measure* similarity. Specifically, we can define the local measure for s as follows:

$$\tau(s) = \frac{|s - s'|}{\mu(s, s')} \quad (34)$$

where s' is a nearby sample that is infinitely close to (but distinct from) s . Specifically, $\tau = r$ for isotropic sampling and $\tau = J^{-1}$ for anisotropic sampling. From this, we factor the following local measure similarity into Equation 9:

$$\nu(\mathbf{d} = s - s') = \begin{cases} 1 & \text{if } |\tau(s) - \tau(s')| \leq \varepsilon \times \min(\tau(s), \tau(s')) \\ 0 & \text{else} \end{cases} \quad (35)$$

Theoretically, ν would allow us to adaptively adjust the local neighborhood size for collecting \mathbf{d} pairs depending on the local variations of domain property manifested by τ , but throughout our experimental results across a variety of adaptive, anisotropic, and surface domains we have not found it to be necessary.

I Hidden Structure

Here, we provide experimental results for the potential pathological cases of our range selection as discussed in Section 4. The worst possible case for our method would be a maximum packing with inter-sample spacing d just beyond εr_{min} of our method. (In 2D, the maximum packing corresponds to hexagonal lattices.) However, as shown in Figure 10, our method is no less discriminative than Fourier spectrum analysis. A similar experiment using square lattice is shown in Figure 11.

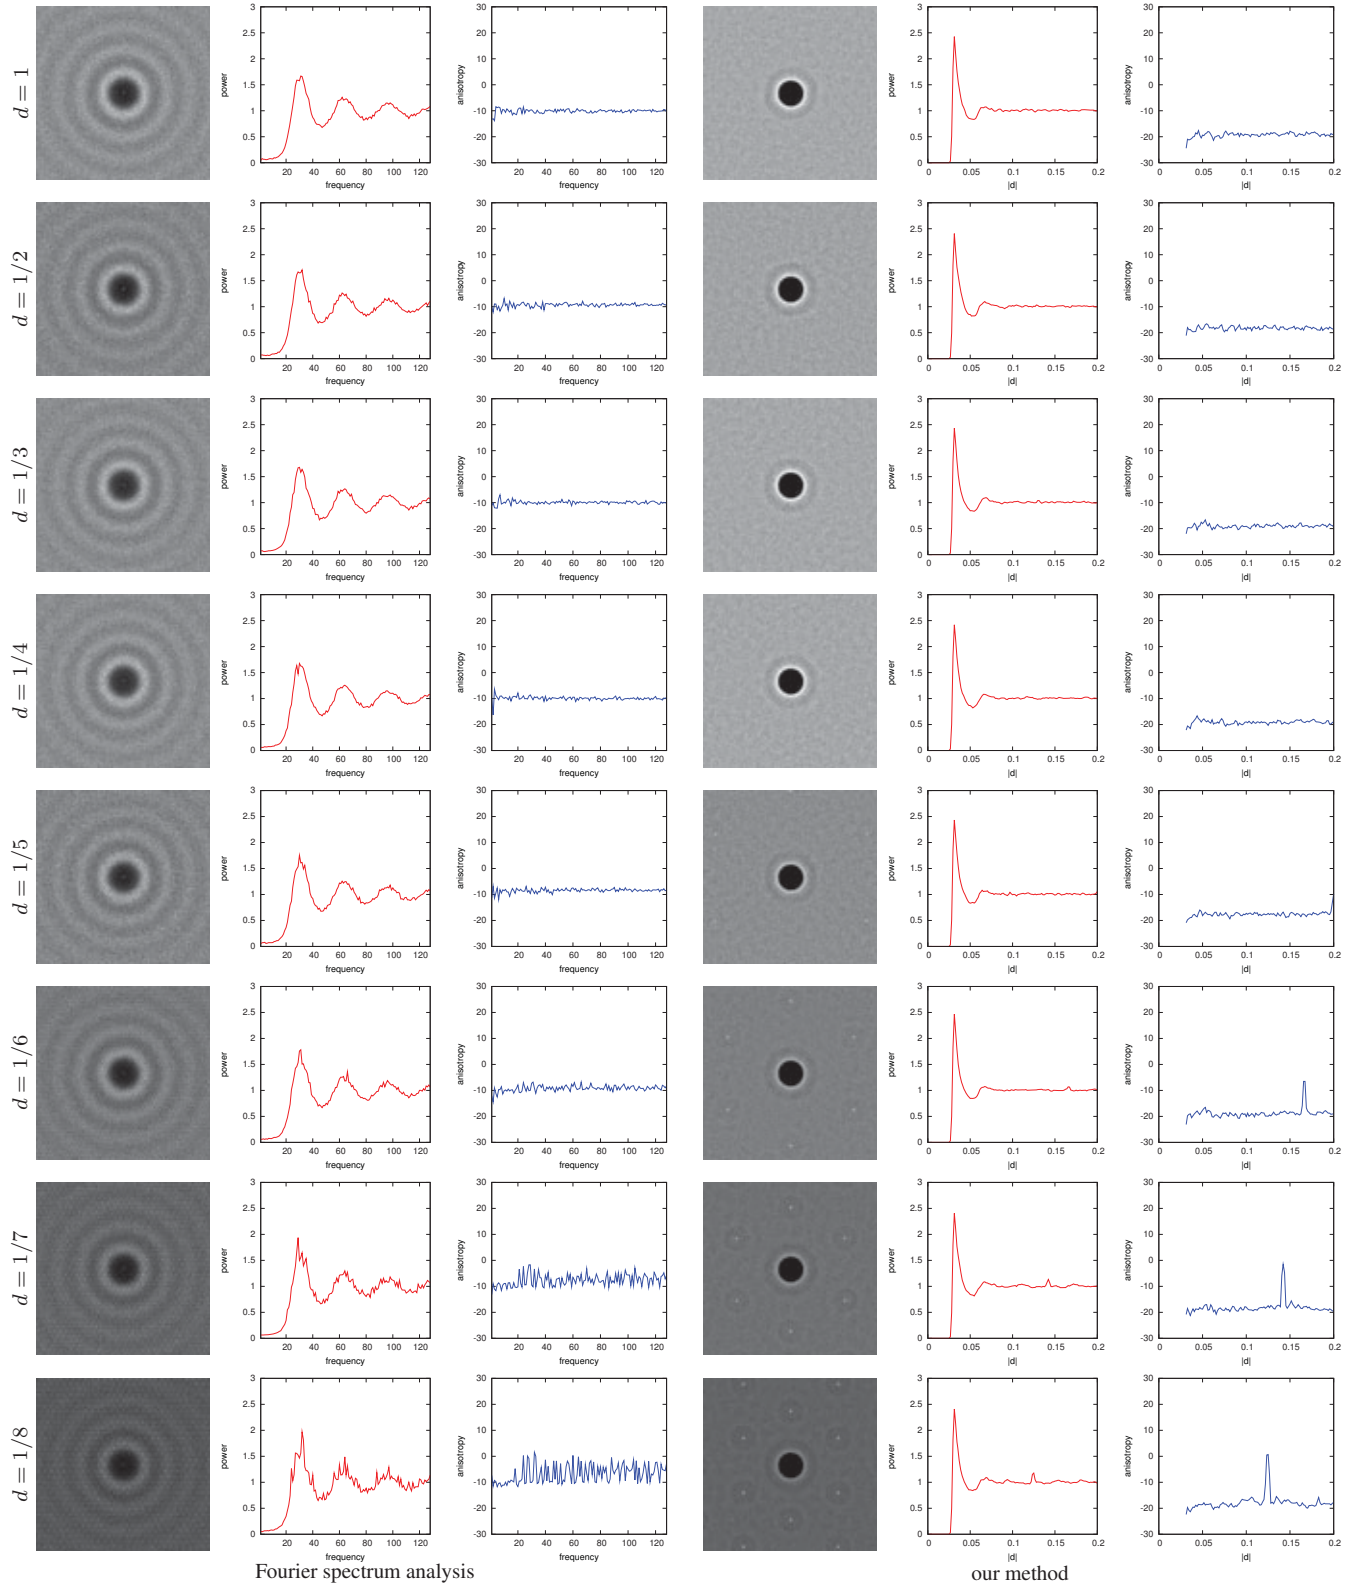


Figure 10: Detecting hidden hexagonal lattice. Here, we first produce a hexagonal pattern with different inter-sample spacing d followed by Poisson disk sampling with $r_{\min} = 0.03$. We then apply Fourier spectrum analysis and our method over the sample sets similar to Figure 2.

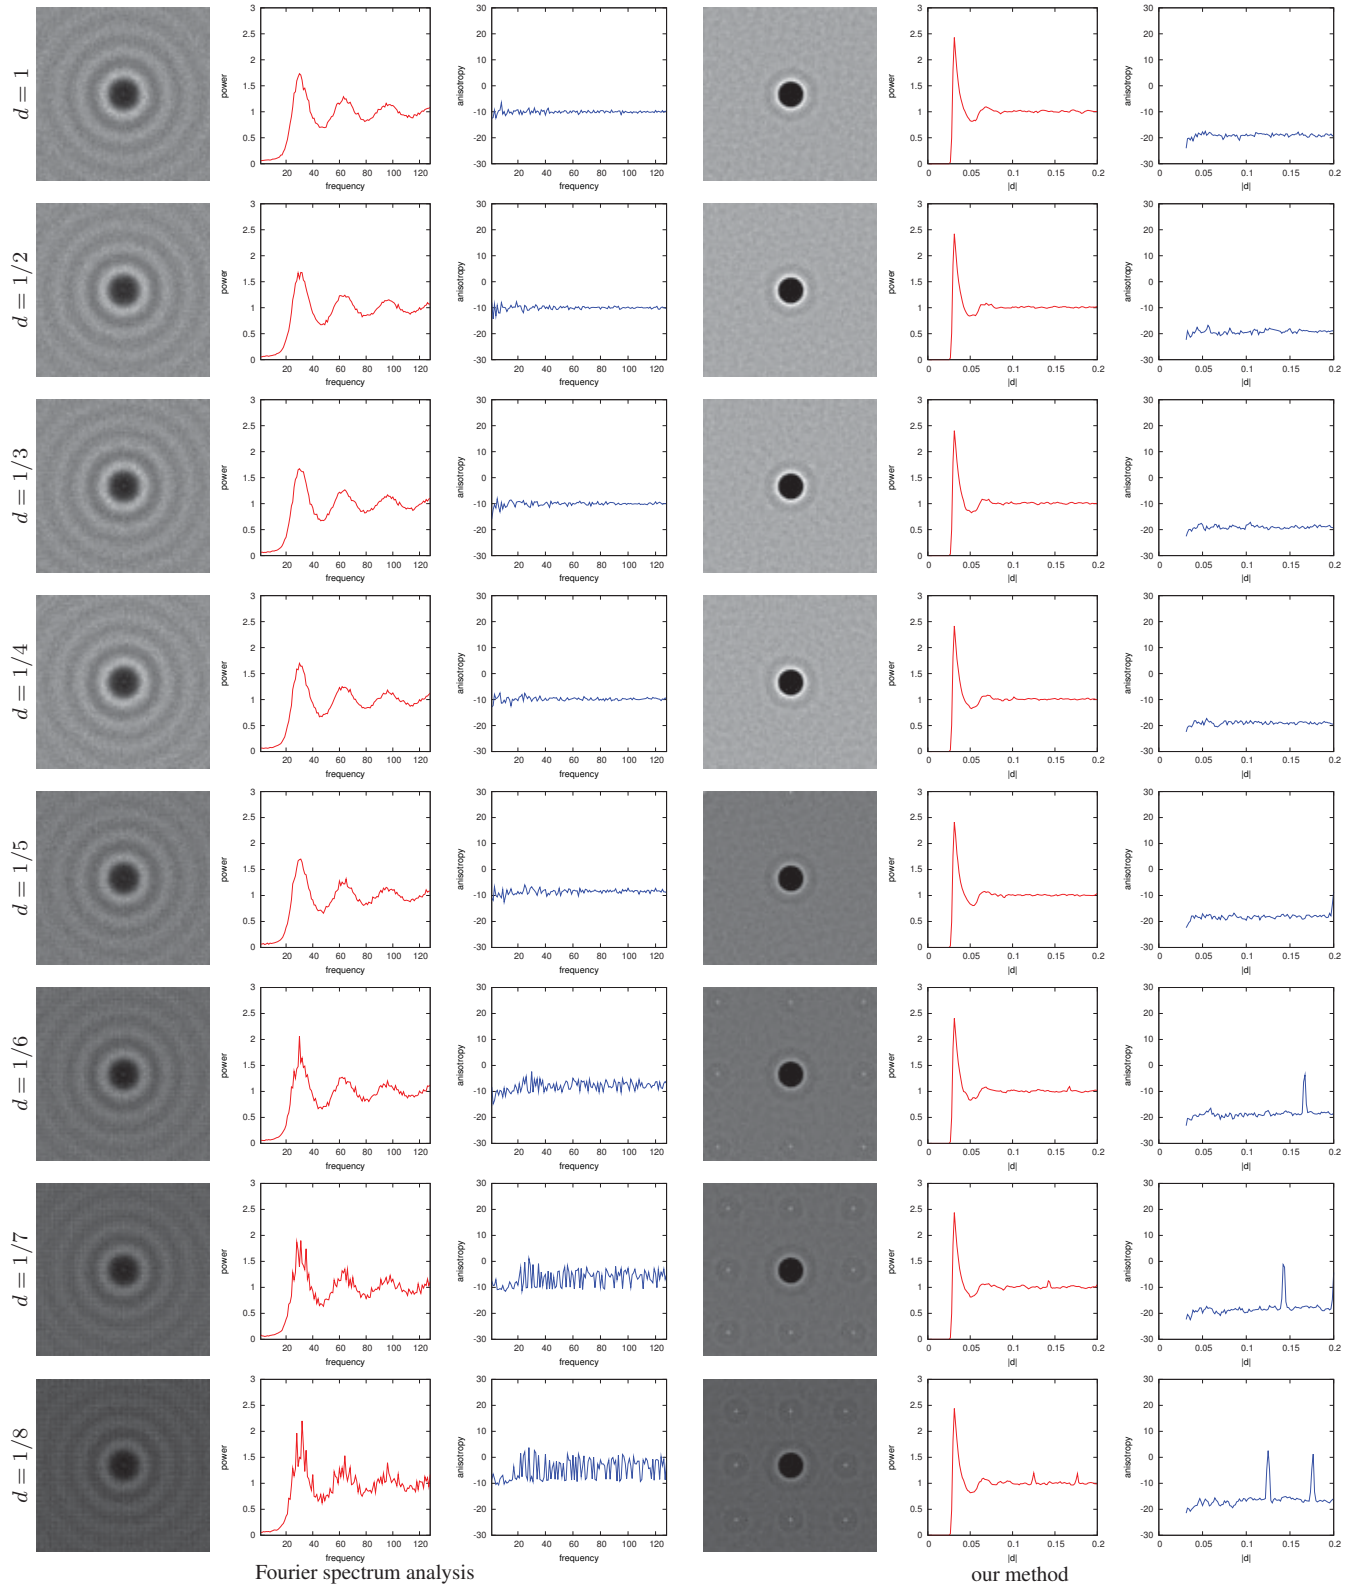


Figure 11: Detecting hidden square lattice. This is a similar experiment to Figure 10, except the use of square instead of hexagonal lattices.