General Point Sampling with Adaptive Density and Correlations
Supplementary Material - Algorithms

paper 1192

Abstract

In this supplementary material, we detail an alternative synthesis algorithm that we use to generate results with locally regular structures. We start from an existing discrete texture synthesis method and show how it can be extended to support locally stationary non-ergodic point processes. In the second Section, we explain how strong local anisotropy can be handled in our analysis and synthesis framework.

1. Synthesis of Locally Non-Ergodic Sampling Patterns

Our synthesis algorithm assumes a locally stationary point process and synthesizes a single instance, i.e. distribution, of such a point process. However, for non-ergodic processes, statistics cannot be reliably extracted from a single distribution. In other words, expected values over different distributions generated by the same point process cannot be estimated from a single larger distribution. An example of a non-ergodic stationary process can be obtained by taking a regular grid, and randomly translating or rotating it, which is also called uniform or isotropic jittering [RAMN12, Ö16] in the literature. For these processes, difference vectors between sample points always stay the same as the points always form a regular grid. However, each instance is actually different due to the random global translation of the grid. This cannot be inferred by observing a single distribution.

This is problematic for our synthesis algorihm as the output statistics are extracted and matched for a single distribution, the synthesized distribution. Thus, for locally non-ergodic stationary processes, we need a different synthesis approach. For these cases, we thus extend a recent texture synthesis based approach for synthesizing distributions of discrete entities [ROM*15].

Below we first explain the basics of this algorithm, and then elaborate on how it can be adapted to synthesize distributions from locally stationary ergodic or non-ergodic processes.

1.1. The Synthesis Algorithm

We start from the formulation of a recent patch based texture synthesis algorithm that can handle discrete repeated structures in general domains [ROM*15]. The algorithm takes a single example distribution as the input, and generates a distribution with similar local structures. Due to the point-based representation of the textures, and smooth formulation of the resulting optimization prob-

lem, this formulation forms a suitable foundation for our synthesis algorithm.

The original algorithm is based on a local similarity measure between a continuous output function $\mathbf{f}(\mathbf{x})$ and the continuous example function $\mathbf{e}(\mathbf{x})$. The *similarity error density* is defined, for the location $\mathbf{x}$, as

$$S(\mathbf{f}(\mathbf{x}), \mathbf{e}(\mathbf{m}(\mathbf{x}))) = \int_{\mathbb{R}^n} |\mathbf{f}(\mathbf{x}+\mathbf{s}) - \mathbf{e}(\mathbf{m}(\mathbf{x})+\mathbf{s})|^2 \, w(\mathbf{s}) \mathbf{ds} \quad (1)$$

where $w(.)$ is a window function which delimits a local neighborhood $\mathcal{N}$, and $\mathbf{m}(\mathbf{x})$ is a discontinuous mapping that matches each output domain point $\mathbf{x}$ to a point $\mathbf{m}(\mathbf{x})$ within the domain of the example function. Intuitively, S measures how well a local patch in f matches some other local patch in e. The extent of the patches is given by the window function w, and the center locations of the matching patches are $\mathbf{x}$ and $\mathbf{m}(\mathbf{x})$ for f and m, respectively.

The *total similarity error* between the output and input functions is then defined as

$$T = \int_{\mathcal{D}} S(\mathbf{x}, \mathbf{e}(\mathbf{m}(\mathbf{x}))) \mathbf{dx}, \quad (2)$$

where $\mathcal{D}$ is a subspace in the output domain. The T sums up all matching scores and thus measures the overall matching accuracy of each patch in f to some patch in e. For minimum T, we thus need to have that for each patch in the output f, we have a matching patch in the example e.

As in the common neighborhood matching based texture synthesis methods, the T is minimized alternatively for $\mathbf{m}(\mathbf{x})$ and $\mathbf{f}(\mathbf{x})$, with a matching step (where the best matching neighborhoods for the current $\mathbf{f}(\mathbf{x})$ are found, i.e. $\mathbf{m}(\mathbf{x})$ is computed), and a merging step (where the best output function $\mathbf{f}(\mathbf{x})$ is computed for the current $\mathbf{m}(\mathbf{x})$).

The key idea of the synthesis algorithm is forming these functions f and e in terms of the sampling points. Denoting the sample

points in the synthesized output as $\{(\mathbf{x}_i), i = 1..n\}$, and those in the given example distribution with $\{(\mathbf{e}_i), i = 1..m\}$, the continuous f and e can be formed by placing Gaussians $g(\mathbf{x}, \sigma) = e^{-|\mathbf{x}|^2/\sigma^2}$ at all point locations and summing them together, that is

$$
\begin{aligned}
\mathbf{f}(\mathbf{x}) &= \sum_i g(\mathbf{x} - \mathbf{x}_i, \sigma) \\
\mathbf{e}(\mathbf{x}) &= \sum_i g(\mathbf{x} - \mathbf{e}_i, \sigma).
\end{aligned} \tag{3}
$$

Then, the *similarity error density S* can be analytically evaluated (see the work by Roveri et al. [ROM*15] for details). For the discretization of the integral in Equation 2, the output domain $\mathcal{D}$ is sampled with a regular grid consisting of the background integration points $\mathbf{c}_k$. For each $\mathbf{c}_k$, there is a best matching location $\mathbf{m}_k$ in the example domain. Hence, each $\mathbf{c}_k$ defines a neighborhood $\mathcal{N}_k$ in the output domain, that will match to some neighborhood in the example domain. The *discrete similarity error* is then obtained as

$$
T = \sum_k S(\mathbf{q}_k, \mathbf{e}(\mathbf{m}_k)). \tag{4}
$$

1.2. Synthesis with Adaptive Density and Correlations

The original algorithm, as described in the last section, is designed for a single example.

Although the original algorithm from [ROM*15] can synthesize output point sets with general repeated patterns, these patterns are determined by a single example. Hence, it can be used to synthesize a stationary distribution with fixed pair-wise correlations, while the density can vary in the output space by scaling the example [ROM*15], resulting in the same limitations as the previous sampling papers. Having an example distribution with spatially-adaptive correlations will not generate adaptive correlations in the output since the neighborhoods from the example are assumed to be repeated throughout the output space. Hence, the main challenge of adaptive synthesis is revising the technique to allow for multiple examples and their combinations.

Adaptive Correlations In order to synthesize distributions with adaptive correlations, we thus extend the *similarity error density* to consider multiple input functions and weight their importance depending on the location in the output domain. In this case, instead of having a single example function e, we have multiple functions e_p.

For an arbitrary point $\mathbf{x}$ in the output domain $\mathcal{D}$, there is a weight $\mathbf{w}_p(\mathbf{x})$ that gives how much example e_p is influencing that point. Our *weighted similarity error density* then computes a weighted average of the local similarities between the output function and multiple input functions, as

$$
S_w(\mathbf{x}) = \sum_p \mathbf{w}_p(\mathbf{x}) S(\mathbf{f}(\mathbf{x}), \mathbf{e}_p(\mathbf{m}_p(\mathbf{x}))) \tag{5}
$$

where $S(\mathbf{f}(\mathbf{x}), \mathbf{e}_p(\mathbf{m}_p(\mathbf{x})))$ is the similarity between the output function f, and the p-th input function e_p, and $\mathbf{m}_p$ is the matching function from the output domain to the domain of e_p.

In the synthesis algorithm in the paper, we define a pair correlation function (PCF) at each point $\mathbf{x}$ as a linear combination of basis PCF-s. The example distributions that define e_p give these

basis PCF-s, and the weights w_p correspond to the computed or given weights for linearly combining the basis PCF-s. Hence, instead of having a PCF at each point in the output domain as a linear combination of the basis PCF-s, we have the corresponding basis example distributions and their linear combinations explicitly given in Equation 5.

Adaptive Density and Orientation The original method accounts for adaptive density and orientation by introducing a scale factor $\mathbf{s}(\mathbf{x})$ and a rotation matrix $\mathbf{R}(\mathbf{x})$ defined at every location in the output domain, which scales and rotates the example domain before matching. After computing the sizes of the neighborhoods $\mathcal{N}_k$ according to the given local density λ_k as described in the paper, we adjust the scaling field s such that at each neighborhood, we have the required number of points αn (please see the paper) in the neighborhoods of the same size in the example domain on average. The rotation for each neighborhood is used to rotate the corresponding examples for that neighborhood before synthesis, as done by Roveri et al. [ROM*15].

While simply scaling the input function with $\mathbf{s}(\mathbf{x})$, as proposed in [ROM*15], works for slight changes of scales, we found that, in order to properly handle adaptive densities, it is necessary to accordingly scale some additional parameters as well which are proportional to the input scale. The standard deviation σ of the Gaussians used to define the input and output functions in Equation 3 needs to be adapted according to the size of the neighborhoods and hence the scaling function s. If both σ and the local density of the sample points is small, there is little overlap between the Gaussians, and the minimization routine will not be able to compute a reliable gradient. Conversely, having Gaussians with a large overlap blurs out the distributional details and makes the synthesis not reliable. For a Gaussian placed at a sample position $\mathbf{x}$ in the output domain, we thus scale its initial σ with the scale factor $\mathbf{s}(\mathbf{x})$. We use the same scheme as in the paper for placing the neighborhood center points $\mathbf{c}_k$.

We show an example case where we compare to the original algorithm with non-adaptive neighborhood sizes and regular sampling of neighborhood centers in Figure 1. The example input distribution is shown in (a). Figure 1 (b) shows the output of our synthesis algorithm with our modifications, i.e. setting a proper spacing between the neighborhood centers, while (c) and (d) are obtained with the original algorithm where there are many overlapping large neighborhoods, and neighnorhoods with small overlap, respectively. Unlike the PCF based synthesis case, our adjustments for the neighborhood sizes and the corresponding neighborhood sampling are essential to get correct statistics for point distributions.

Initialization and Control Sampling We start by randomly initializing samples in the whole output domain as explained in the paper. In addition, the optimization procedure described in [ROM*15] includes a control sampling strategy to dynamically adapt the number of samples in the output domain and achieve better convergence. At each iteration, samples at random positions are added and existing samples are deleted if this would improve the similarity error density measure. We refer to the original paper for more details.

**Geometry Sampling and Local Anisotropy** The original

Figure 1: *Different distributions generated for the blue noise input example shown in (a). We obtain distributions with correct statistics when using our adjusted distribution of neighborhood centers (b), while having large (c) or small (d) overlaps leads to skewed characteristics.*

method allows to define guiding fields in 3D to accordingly orientate the input example and synthesize a distribution on a curved suface.

For small anisotropy factors, the original method allows to synthesize anisotropic distribution by simply scaling the input example. For strong anisotropy, the method can be extended to utilize anisotropic Gaussians.

Parameters The parameter σ determines the smoothness of the matched functions. In our implementation, it is set to the average spacing between the samples in the example. We refer to the [ROM*15] for other parameters. The neighborhoods should be large enough to capture the PCF of the distribution within them. For the distributions used in the paper, the neighborhoods sizes were set to include at least 50 samples.

1.3. Limitation

The main limitation of this discrete texture synthesis based algorithm results from using actual example distributions rather than extracted statistics as we do in the paper. A well-known limitation of such neighborhood based texture syntehsis approaches is that they cannot handle repetitions at multiple scales [ROM*15]. For the case of synthesizing point distributions, this means that clustered distributions, where the points in each cluster follow a certain distribution, and the clusters themselves follow another distribution, may not be reliably synthesized. In practice, we have observed that we still get visually accurate results for these cases. However, the statistics deviate from those of the examples slightly.

2. Analysis and Synthesis with Local Anisotropy

2.1. Analysis

The matrix $\mathbf{M}_x$ is a measure of anisotropy. We utilize 2D anisotropy measures from the spatial point processes literature [IPSS08] to compute $\mathbf{M}_x$. This method performs a simple density estimation on the difference vectors $\mathbf{h}_{ij}$. First, a radial histogram of the vectors $\mathbf{h}_{ij}$ is computed. Each bin thus corresponds to a direction on the unit hypersphere and a distance from its center. The dominant anisotropy is then extracted by picking the directions with the smallest and largest bin sums. We then form a covariance matrix $\mathbf{C}_x$ with two eigenvectors set as these directions, and the eigenvalues as the corresponding bin values. Finally, the matrix $\mathbf{M}_x$ is set as the whitening transform $\mathbf{M}_x = \mathbf{C}_x^{-1/2}$ such that the resulting difference vectors $\mathbf{h}_{ij}$ are distributed isotropically. Note that when we apply the transform $\mathbf{M}_x$, we assume that the volume $|\mathcal{N}_x|$ also scales so as to keep the intensity at λ_x. This step is designed for point distributions on two dimensional domains such as the 2D plane or two-manifold surfaces in 3D, or with limited number of anisotropy directions in higher dimensional domains.

2.2. Synthesis

Before feeding into the estimator, the difference vectors $\mathbf{h}_{ij}$ are first whitened with $\mathbf{M}_k \mathbf{h}_{ij}$. The gradient $\frac{\partial}{\partial \mathbf{x}_i} E_k$ is then computed, multiplied with $\mathbf{M}_k^{-1}$, and summed over all neighborhoods to get the final gradient $\Delta_i = \sum_k \mathbf{M}_k \frac{\partial}{\partial \mathbf{x}_i} E_k$ for point $\mathbf{x}_i$.

References

[IPSS08] ILLIAN J., PENTTINEN A., STOYAN H., STOYAN D. (Eds.): *Statistical Analysis and Modelling of Spatial Point Patterns*. John Wiley and Sons, Ltd., 2008. 3

[Ö16] ÖZTIRELI A. C.: Integration with stochastic point processes. *ACM Trans. Graph. 35*, 5 (Aug. 2016), 160:1–160:16. 1

[RAMN12] RAMAMOORTHI R., ANDERSON J., MEYER M., NOWROUZEZAHRAI D.: A theory of monte carlo visibility sampling. *ACM Trans. Graph. 31*, 5 (Sept. 2012), 121:1–121:16. 1

[ROM*15] ROVERI R., ÖZTIRELI A. C., MARTIN S., SOLENTHALER B., GROSS M.: Example Based Repetitive Structure Synthesis. *Computer Graphics Forum* (2015). 1, 2, 3

General Point Sampling with Adaptive Density and Correlations Supplementary Material - Theory

paper 1192

Abstract

In this supplementary material, we present a theory of general point distributions in terms of the product densities of underlying point processes. We show how these can be estimated from example distributions and interpreted. We establish how this theory relates to and extends existing techniques that assume a stationary point process, and propose a measure of stationarity.

1. Analysis of General Sampling Patterns

1.1. Campbell's theorem

Campbell's theorem relates sums of values of functions at the point locations generated by a point process, to the integrals of those functions and the product densities. For our methods, the following two special cases of the theorem will be important,

$$\mathbb{E}_{\mathbf{X}} \sum_{\mathbf{x}_i \in \mathbf{X}} f(\mathbf{x}_i) = \int_{\mathbb{R}^d} f(\mathbf{x})\lambda(\mathbf{x})d\mathbf{x}, \qquad (1)$$

$$\mathbb{E}_{\mathbf{X}} \sum_{\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}}^{\neq} f(\mathbf{x}_i, \mathbf{x}_j) = \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} f(\mathbf{x}, \mathbf{y})\rho(\mathbf{x}, \mathbf{y})d\mathbf{x}d\mathbf{y}, \qquad (2)$$

for any positive function f. We will start from these expressions to derive our estimators in the next section. For brevity, we drop the integration domains where clear.

1.2. Estimating Product Densities

Our analysis relies on estimating the statistics $\lambda(\mathbf{x})$ and $\rho(\mathbf{x}, \mathbf{y})$ as defined in the previous section. These two statistics define *all* point patterns under the common assumption of Gaussianity. The density measure $\lambda(\mathbf{x})$ can be easily estimated using standard techniques such as kernel density estimation. The estimation of $\rho(\mathbf{x}, \mathbf{y})$ has so far not been considered except for the case of stationary or isotropic correlation models. We derive a reliable estimator that converges to ρ, and show that it can be interpreted as a density estimator in the higher dimensional space of aggregated point coordinates.

Given multiple distributions generated by a point process, we would like to design unbiased and low variance estimators for the product densities. We start with deriving an estimator for the intensity based on the Campbell's theorem and show how it naturally extends to second order product density ρ. We assume a monotonically decreasing positive kernel function k for the estimators. In this work, we assume a Gaussian kernel with $k_i(\mathbf{x}) := k(\mathbf{x}, \mathbf{x}_i) :=$

$\frac{1}{(\sqrt{\pi}\sigma)^d} e^{-||\mathbf{x} - \mathbf{x}_i||^2 / \sigma^2}$. If we plug this into the expression into Equation 1, we get

$$\mathbb{E}_{\mathbf{X}} \sum_{\mathbf{x}_i \in \mathbf{X}} k(\mathbf{x}, \mathbf{x}_i) = \int k(\mathbf{x}, \mathbf{y})\lambda(\mathbf{y})d\mathbf{y}. \qquad (3)$$

The convolution on the right-hand side converges to $\lambda(\mathbf{x})$ as $\sigma \to 0$. The left-hand side is the classical non-parametric density estimator. This estimator will get better as we get more instances of the point pattern.

Utilizing Campbell's theorem for higher order statistics, we can generalize this result. Specifically, for any two distinct point locations $\mathbf{x}_i$ and $\mathbf{x}_j$, if we plug $k(\mathbf{x}, \mathbf{x}_i)k(\mathbf{y}, \mathbf{x}_j)$ into Equation 2, we obtain

$$\mathbb{E}_{\mathbf{X}} \sum_{\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}}^{\neq} k_i(\mathbf{x})k_j(\mathbf{y}) = \int \int k(\mathbf{x}, \mathbf{z})k(\mathbf{y}, \mathbf{t})\rho(\mathbf{z}, \mathbf{t})d\mathbf{z}d\mathbf{t}. \qquad (4)$$

Analogous to the estimator for $\lambda(\mathbf{x})$, the right-hand side converges to $\rho(\mathbf{x}, \mathbf{y})$ as $\sigma \to 0$, and the left-hand side gives us the estimator for ρ. Hence, we define the estimator for ρ as

$$\hat{\rho}(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}}^{\neq} k_i(\mathbf{x})k_j(\mathbf{y}). \qquad (5)$$

This estimator can be computed for each $\mathbf{x}$ and $\mathbf{y}$ by averaging over different distributions. Note that this is a $2d$-dimensional statistic, to fully capture the rich correlation structure of the underlying point process.

Interpretation as a density estimator Since we assume Gaussian kernels, we can combine the two kernels in Equation 5 to obtain a single Gaussian that depends on the distance between the $2d$-dimensional vectors $[\mathbf{x}^T \mathbf{y}^T]^T$, and $[\mathbf{x}_i^T \mathbf{x}_j^T]^T$. Hence, this estimator can be considered as a non-parametric density estimator in the product space $\mathbb{R} \times \mathbb{R}$ for the points generated by aggregating all combinations of different points in the original point pattern.
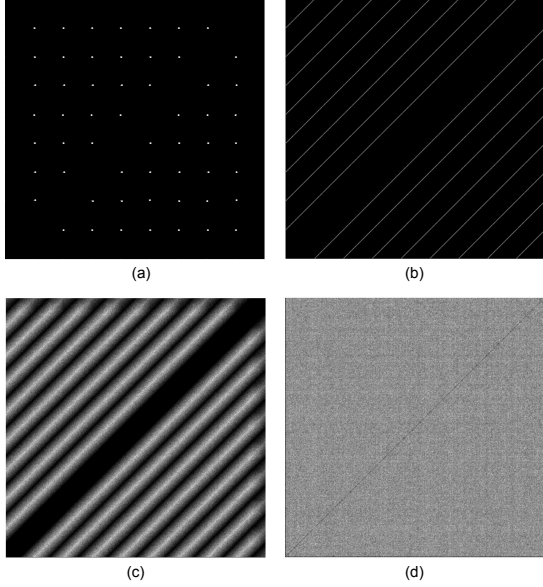
Figure 1: *Plot of ρ for different 1D distributions: regular grid (a), infinite regular grid with random translation (b), jittered distribution (c), random distribution (d).*

Figure 2: *Variance graph for some 2D distributions: regular grid (a), infinite regular grid with random translation (b), jittered distribution (c), random distribution (d).*

Figure 1 shows plots of ρ for different 1D distributions. For the regular grid, we get a regular grid again since the density in the product space is regular. For an infinite regular grid with random translation, we get lines parallel to the $x = y$ line as shown in Figure 1 (b). This is due to the correlated translation of sampling points, i.e. for each realization of this point process, the point locations are $x + t$ and $y + t$ for a random translation t. If this translation is independent for each sampling point, we get a jittered distribution [Mit96, Owe13], and the lines start to become blurred due to the loss of correlations (Figure 1 (c). Finally, for a completely random distribution, we get a constant ρ (Figure 1, d).

Relation to previous analysis techniques The ρ reduces to previously used statistics for the special case of stationary distributions. For these distributions, $\rho(\mathbf{x}, \mathbf{y}) = \rho(\mathbf{x} - \mathbf{y})$ is constant for all $\mathbf{x}$ and $\mathbf{y}$ with $\mathbf{y} = \mathbf{x} + \mathbf{h}$ for a constant $\mathbf{h}$. Then, we can integrate over all such $\mathbf{x}$ and $\mathbf{y}$ pairs to get an estimator of $\rho(\mathbf{h})$ for stationary distributions as follows

$$\int \mathbb{E}_{\mathbf{X}} \sum_{\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}}^{\neq} k_i(\mathbf{x}) k_j(\mathbf{x} + \mathbf{h}) d\mathbf{x} d\mathbf{y}$$

$$= \mathbb{E}_{\mathbf{X}} \sum_{\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}}^{\neq} \int k(\mathbf{x} - \mathbf{x}_i) k(\mathbf{x} - \mathbf{x}_j + \mathbf{h}) d\mathbf{x} d\mathbf{y} \qquad (6)$$

$$= \mathbb{E}_{\mathbf{X}} \sum_{\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}}^{\neq} k'(\mathbf{h} - \mathbf{h}_{ij}),$$

with $\mathbf{h}_{ij} = \mathbf{x}_i - \mathbf{x}_j$ and k' is a Gaussian with standard deviation 2σ. This results in the estimator for stationary distributions proposed by Wei et al. [WW11]. Similarly, if we integrate for constant $||\mathbf{h}||$, we can recover the estimator for isotropic point processes in the work

by Öztireli et al. [OG12]. Since the ρ is related to periodogram with a Fourier transform [HSD13], spectral analysis can be similarly treated in the same framework for stationary distributions. Note that for the distributions with adaptive density, previous works alter the difference or distance metric used, but still assume that the underlying correlation model is fixed by a single translation invariant ρ. The only method that does not operate under the invariance assumption is proposed by Subr and Kautz [SK13], where they utilize the Fourier spectrum instead of the phaseless power spectrum. However, spatially adaptive correlations can be mixed into different frequencies in the spectral domain, and thus it is not clear how to interpret and utilize the resulting diagrams for adaptive correlations.

A measure of stationarity The above observation can be utilized to devise a test of stationarity for general distributions. To normalize with respect to density, we work with the pair correlation function (PCF) $g(\mathbf{x}, \mathbf{y}) := \rho(\mathbf{x}, \mathbf{y})/\lambda(\mathbf{x})\lambda(\mathbf{y})$. We expect $g(\mathbf{x}, \mathbf{x} + \mathbf{h})$ to be constant for a given $\mathbf{h}$. Hence, we can define the variance of g at this $\mathbf{h}$ as a measure of stationarity.

$$\hat{s}(\mathbf{h}) = \frac{1}{|V|} \int_V \hat{g}^2(\mathbf{x}, \mathbf{x} + \mathbf{h}) d\mathbf{x} - \left[\frac{1}{|V|} \int_V \hat{g}(\mathbf{x}, \mathbf{x} + \mathbf{h}) d\mathbf{x} \right]^2, \qquad (7)$$

for a given domain V and its volume $|V|$. In the product space, this means that we are measuring variance of the values of g along lines parallel to the $\mathbf{x} = \mathbf{y}$ line. For stationary distributions such as the uniformly translated grid in Figure 1 (b) and random distribution in Figure 1 (d), we get low variance along these lines. We further plot variance graphs for some 2D distributions in Figure 2. As expected, stationary distributions result in much less variance. We will utilize this measure to understand the local behaviour of our synthesis algorithm in the next section.

References

[HSD13] HECK D., SCHLÖMER T., DEUSSEN O.: Blue noise sampling with controlled aliasing. *ACM Trans. Graph. 32*, 3 (July 2013), 25:1–25:12. 2

[Mit96] MITCHELL D. P.: Consequences of stratified sampling in graphics. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1996), SIGGRAPH '96, ACM, pp. 277–280. 2

[OG12] ÖZTIRELI A. C., GROSS M.: Analysis and synthesis of point distributions based on pair correlation. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 170:1–170:10. 2

[Owe13] OWEN A. B.: *Monte Carlo Theory, Methods and Examples*. To be published, 2013. 2

[SK13] SUBR K., KAUTZ J.: Fourier analysis of stochastic sampling strategies for assessing bias and variance in integration. *ACM Trans. Graph. 32*, 4 (July 2013), 128:1–128:12. 2

[WW11] WEI L.-Y., WANG R.: Differential domain analysis for non-uniform sampling. *ACM Trans. Graph. 30*, 4 (July 2011), 50:1–50:10. 2

For each example, in order:

1. Blue Noise
2. Regular and Blue Noise with sharp transitions
3. Regular and Blue Noise with smooth transitions