EUROGRAPHICS 2017 / L. Barthe and B. Benes
(Guest Editors)

Volume 36 (2017), Number 2

General Point Sampling with Adaptive Density and Correlations

Riccardo Roveri[1], A. Cengiz Öztireli[1], Markus Gross[1]

[1]ETH Zurich

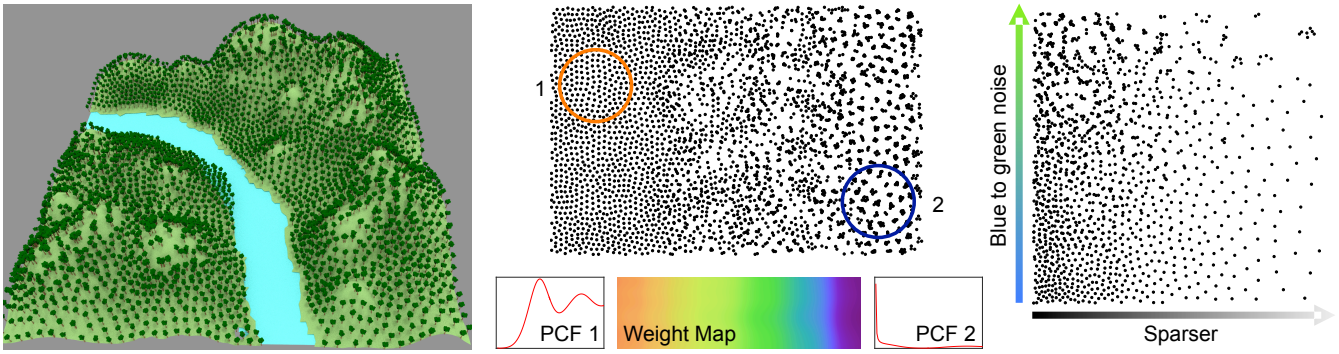

Figure 1: *Understanding natural or synthetic complex distributions such as the natural distribution of trees based on the altitude on the left is a difficult problem if the arrangement of the entities, resulting from pair-wise interactions, is spatially-adaptive, as shown for a canonical example in the middle. Our analysis technique provides an informative and comprehensive summary of such distributions. The correlations in our framework are represented with a set of extracted basis pair correlation functions (PCF 1 and 2, from local patches 1 and 2 in the example in the middle), and the corresponding weight maps illustrating how they are interpolated in space. Our synthesis algorithm utilizes these measures to synthesize distributions with adaptive density and correlations on Euclidean domains (right) or surfaces (left).*

Abstract

Analyzing and generating sampling patterns are fundamental problems for many applications in computer graphics. Ideally, point patterns should conform to the problem at hand with spatially adaptive density and correlations. Although there exist excellent algorithms that can generate point distributions with spatially adaptive density or anisotropy, the pair-wise correlation model, blue noise being the most common, is assumed to be constant throughout the space. Analogously, by relying on possibly modulated pair-wise difference vectors, the analysis methods are designed to study only such spatially constant correlations. In this paper, we present the first techniques to analyze and synthesize point patterns with adaptive density and correlations. This provides a comprehensive framework for understanding and utilizing general point sampling. Starting from fundamental measures from stochastic point processes, we propose an analysis framework for general distributions, and a novel synthesis algorithm that can generate point distributions with spatio-temporally adaptive density and correlations based on a locally stationary point process model. Our techniques also extend to general metric spaces. We illustrate the utility of the new techniques on the analysis and synthesis of real-world distributions, image reconstruction, spatio-temporal stippling, and geometry sampling.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Antialiasing I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Sampling

1. Introduction

Sampling patterns lie at the heart of many important applications in computer graphics such as representing and integrating functions, anti-aliasing, image and geometry sampling, physically-based simulation, non-photorealistic rendering, object and texture placement, and modeling natural distributions. It is thus essential to understand and control characteristics of sampling patterns.

Point distributions are generally characterized by the *density* and the arrangement of the sampling points given by the *pair-wise correlations* among point locations, as higher order correlations are not needed to study most patterns [IPSS08, OG12]. The most studied correlation model is based on variations of blue noise patterns where the points are randomly distributed with a minimum distance between pairs [Uli88]. Adaptive density, locally anisotropic distributions, or patterns on surfaces can be obtained by adapting the metrics used accordingly, while keeping the correlation model fixed [LWSF10, dGBOD12, JZW*15]. Recent works further explore matching a given target correlation model [ZHWW12, OG12, WPC*14, AHD15], resulting in an explicitly controlled arrangement of points. However, for all the synthesis methods, the correlation model is assumed to be constant throughout the space, resulting in translation invariant patterns up to local density variations. The analysis methods are thus also designed for studying such translation invariant correlations via statistics based on distributions of modulated difference vectors or distances between pairs of points and associated spectral measures [BWWM10, WW11, OG12, HSD13]. Hence, point patterns where the correlations are spatially varying cannot be handled with the current common analysis and synthesis techniques. In Figure 1, we show an example point distribution with spatially-adaptive correlations, where the pattern transitions from blue to green noise from left to right. Utilizing such patterns can lead to significantly improved results for many applications in computer graphics.

In this work, we present novel analysis and synthesis techniques for point patterns with adaptive density and *adaptive correlations*. Starting from the theory of stochastic point processes, we propose a novel analysis method for general point patterns, and a synthesis algorithm capable of generating distributions with spatially adaptive density and correlations. The proposed statistics converge to provably discriminative measures from point processes, and provide a comprehensive framework for point patterns. We illustrate how such patterns can improve image and geometry sampling with various examples. In summary, our main contributions are the following:

- The notion of adaptive correlations, and a comprehensive analysis framework for general point patterns. The proposed measures are based on well-known statistics form stochastic point processes, and reduce to previous analysis tools for the special case of translation invariant patterns.
- A synthesis algorithm for point patterns with adaptive density and correlations on general domains. We apply the algorithm to generate distributions on Euclidean domains as well as surfaces. In contrast to the previous works, the algorithm offers full control over the distributional characteristics.

2. Related Work

Analysis of point patterns. Determining characteristics of point patterns is essential for many applications in computer graphics such as stippling and halftoning [SGBW10, Fat11], anti-aliasing [Mit87, LD08, HSD13], object placement [Wei10], integration [PH10, SK13, PSC*15], or geometry sampling [OAG10]. A widely used analysis tool is the power spectrum, a 2D diagram computed by averaging the periodograms of point distributions that are instances of a certain point pattern. When the point pattern is translation invariant, many important characteristics such as anti-aliasing properties or anisotropy of the generated distributions can be inferred from the power spectrum [LD08, Uli88]. Other analysis methods rely on spatial measures such as minimum distance between points [LD08], discrepancy [Shi91], distributions of difference vectors [WW11], or distances [OG12, HSD13] between sample point locations. These methods can also be extended to non-Euclidean domains or point patterns with adaptive density or anisotropy [BWWM10, LWSF10, WW11]. For all cases, the underlying correlation model is assumed to be constant and translation invariant, and the adaptivity in density or space the points reside on is buried into the difference or distance measures used to compute the statistics. In contrast, we present a general analysis method that can handle point patterns with adaptive density and correlations. We prove that the proposed measures converge to provably discriminative statistics from stochastic point processes.

Synthesis of point patterns. Synthesis methods can generate point distributions with certain characteristics controlled by the construction of the synthesis algorithm or via explicitly provided statistics. Most techniques in computer graphics focus on blue noise distributions, where there is a minimum distance between pairs of points, and they are distributed randomly otherwise [Uli88, HSD13]. Variations can be generated by altering the distances between points and introducing randomness via adding, removing, or moving the sampling points [Llo82, MF92, BSD09, Fat11, dGBOD12, JZW*15], or tiling [Ost07] methods. Such distributions are very important for their anti-aliasing properties [Uli88, HSD13], and can be combined with adaptive density methods [LWSF10, dGBOD12, CGW*13], or generated on surfaces [JZW*15] for further applications. However, they cannot be utilized to model more complex patterns where this correlation model does not hold. To synthesize more general patterns with controlled characteristics, recent techniques rely on matching the statistics of output distributions with given target statistics [ZHWW12, OG12, WPC*14, AHD15]. These methods can also handle adaptive density, but are not trivial to extend to non-Euclidean domains such as surfaces [JZW*15]. Paralleling analysis methods, all synthesis algorithms so far assume a given constant pair-wise correlation model, and locally alter the density or anisotropy of points distributions. A notable exception is the work by Ju et al. [JCP*10], which, however, only models group motions and crowd behaviour based on a qualitative analysis. We present the first technique that can synthesize general point distributions with adaptive pair-wise correlations and density. It has been observed in rendering [Dur11, SK13, SNJ*14] that adapting both simultaneously can reduce the error in numerical integration. We show that such adaptivity can also significantly improve image and geometry sampling. The algorithm offers full control over the spectrum

of points on surfaces, in contrast to the previous surface sampling methods.

Stochastic point processes. The discipline of stochastic point processes [MW04, IPSS08] provides a principled mathematical treatment of general point patterns by characterizing generating processes that underlie point distributions. Hence, each distribution is considered a realization of a stochastic point process. A point process can be defined by setting a random variable at each point in space, and analyzing the correlations among these random variables. Equivalently, we can consider correlations among point locations over different realizations of a point process. Intuitively, first order correlations describe density, and second order pair-wise correlations determine the arrangement of points. Recent works in computer graphics explore utilizing statistics from point processes to analyze and synthesize point distributions [WW11, OG12, HSD13]. The main assumption of these works is that the underlying point process is stationary, i.e. the generated distributions are translation invariant up to density differences. Adaptive density can then be obtained by altering the distance or difference metric utilized. However, many important distributions from classical jittering patterns [Mit96] to complex distributions found in nature [IPSS08] cannot be modelled with these assumptions. We abandon the assumption of an underlying stationary correlation model, propose a comprehensive analysis framework for understanding a more general set of point processes, and develop the associated synthesis algorithms.

3. Analysis of General Sampling Patterns

In this section, we introduce a theoretical framework to study general point patterns, and a set of tools and diagrams that allow for qualitative and quantitative understanding of point distributions exhibiting spatially adaptive density and correlations. We will illustrate that applying existing analysis techniques does not yield meaningful statistics for this general case. We start with the most general case where first and second order correlations are considered, and move on to a locally stationary model that describes a wide range of natural and synthetic point patterns.

3.1. Stochastic Point Processes

The field of stochastic point processes provides a general mathematical framework to study point patterns. Intuitively, a point process is a generating algorithm or mechanism for a set of distributions that share common characteristics. We utilize this theory as a basis to understand and analyze general point patterns with adaptive correlations. We present a brief introduction to point processes, we refer the readers to the excellent books [MW04, IPSS08] for a more in-depth discussion.

The main construct to define a point process is assigning a random variable $X(B)$ to every Borel set $B \in \mathcal{D}$ for a given domain $\mathcal{D}$. Hence, a point process is described by infinitely many random variables. If we fix some sets B_i, we can stack all random variables for these sets to have the random vector $X = [X(B_1), \cdots, X(B_n)]^T$. The point process can then be fully defined by the joint probability $\mathbb{P}(X(B_1) \leqslant b_1, \cdots, X(B_n) \leqslant b_n)$ of the random variables at B_i for all n and all different sets B_i. A familiar example of such a random variable is the random number of points $N(B)$ in B. We can then study the joint probability of the random variables $N(B_1), \cdots, N(B_n)$ to characterize a point process. For the simplicity of the exposition, we will consider point processes with $\mathcal{D} \in \mathbb{R}^d$ in the discussion below, but the concepts also extend to general measure spaces.

General Gaussian processes. Gaussian processes cover almost all types of point processes encountered in applications [IPSS08]. These processes are characterized by having a Gaussian distribution for the random vectors X. Hence, the mean and covariance of the vectors are sufficient to describe Gaussian processes. The importance of such processes is highlighted by the term *second order dogma* in physics [IPSS08], as they very accurately model all distributions found in nature. For these processes, the first and second order moment measures and the associated product densities are sufficient for a complete specification.

Product densities. All Gaussian processes can be described by the first and second order product densities. The first order product density $\lambda(\mathbf{x})$ is called the intensity of the point process, and intuitively measures the average density of points at $\mathbf{x}$ over different distributions generated by the point process. It is proportional to the probability of finding a point in an infinitesimal volume $d\mathbf{x}$ around $\mathbf{x}$ such that $\lambda(\mathbf{x}) = p(\mathbf{x})d\mathbf{x}$. The expected number of points in a set B is given by the integral $\mathbb{E}_{\mathbf{X}} N(B) = \int_B \lambda(\mathbf{x})$, where the expectation is over different distributions $\mathbf{X} = [\mathbf{x}_i, \cdots]$, $\mathbf{x}_i \in \mathcal{D}$ generated by the point process. The second order product density $\rho(\mathbf{x}, \mathbf{y})$ is proportional to the joint probability of finding a pair of points in $d\mathbf{x}$ and $d\mathbf{y}$, $\rho(\mathbf{x}, \mathbf{y})d\mathbf{x}d\mathbf{y} = p(\mathbf{x}, \mathbf{y})$. This statistic determines the pair-wise correlation model, which can be spatially varying for general point processes.

Stationary and isotropic point processes. Stationarity and isotropy are common intrinsic assumptions in the literature. For stationary point processes, the generated distributions are translation invariant. Hence, $\lambda(\mathbf{x})$ is a constant, and $\rho(\mathbf{x}, \mathbf{y})$ turns into $\rho(\mathbf{x} - \mathbf{y})$, a function of the difference vector between two points in space. If we further assume rotation invariance, then $\rho(\|\mathbf{x} - \mathbf{y}\|)$ is a 1D function. For stationary distributions, the second order product density is often expressed in terms of the normalized pair correlation function (PCF) $g(\mathbf{h}) = \rho(\mathbf{h})/\lambda^2$, where we defined $\mathbf{h} = \mathbf{x} - \mathbf{y}$. Estimators of PCF are used for analysis and synthesis of stationary distributions in recent works [WW11, OG12, HSD13]. It is also closely related to the more commonly used periodograms with a Fourier transform [HSD13]. Utilizing this simplified form of pair-wise correlations lies at the heart of the main limitation of previous analysis and synthesis methods. Instead, we will utilize the general ρ, and a model with local stationarity for our techniques.

3.2. Locally Stationary Processes

Analysis with general first and second order statistics is in general difficult as the resulting measures are not intuitive, and hard to estimate unless many instances of the same pattern are available, due to the expectations $\mathbb{E}_{\mathbf{X}}$ involved (we present a formal derivation of the estimators of product densities, and related measures in the supplementary material). For many cases, however, it can be assumed that these measures exhibit a certain degree of smoothness in space.

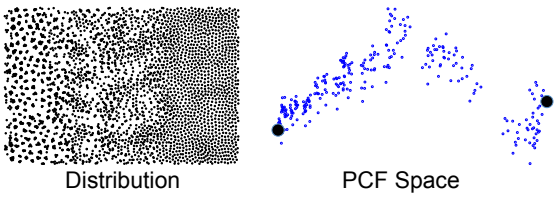

Figure 2: *A distribution with adaptive correlations (left), and the space of extracted PCF's (right) projected into a two dimensional subspace via PCA. Our algorithm detects the two indicated PCF's as the dominant ones, and represents all others as linear combinations.*

We can then decompose the pattern into *locally stationary* patterns around each point in space.

Within the neighborhood $\mathcal{N}_\mathbf{x}$ of a point $\mathbf{x}$, the pattern is then fully described by a constant intensity λ_x, and a PCF $g_x(\mathbf{h})$. For a stationary distribution, there can only be a global anisotropy due to the translation invariance of the generated point distributions [IPSS08], such that the PCF can also be written as $g_x(\|\mathbf{h}\|)$ once the neighborhood is reshaped with a matrix $\mathbf{M}_x$ by applying it to all $\mathbf{h}_{ij} = \mathbf{x}_i - \mathbf{x}_j$ for $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{N}_\mathbf{x}$ to cancel this global anisotropy (we elaborate on how $\mathbf{M}_x$ can be computed in the supplementary material).

Given a distribution with the set of points $\{\mathbf{x}_i, \cdots\}$ generated by an underlying point process, the local intensity λ_x and PCF g_x can be estimated. The intensity has a natural estimator as we derive in the supplementary material with $\hat{\lambda}_x = \sum_{\mathbf{x}_i \in \mathcal{N}_\mathbf{x}} k(\mathbf{x}, \mathbf{x}_i)$ for a normalized kernel such as the Gaussian $k(\mathbf{x}, \mathbf{x}_i) = e^{-||\mathbf{x}-\mathbf{x}_i||^2/\sigma^2}/(\sqrt{\pi}\sigma)^d$. The smoothness of the estimation controlled by σ can be set to reflect the assumed density variation in the distribution.

The estimated PCF can be computed by utilizing an existing smooth estimator for isotropic processes [IPSS08, OG12]:

$$\hat{g}_x(\|\mathbf{h}\|) = \frac{1}{\hat{\lambda}^2 |\partial V_d| \|\mathbf{h}\|^{d-1}} \sum_{\mathbf{x}_i \neq \mathbf{x}_j \in \mathcal{N}_\mathbf{x}} \frac{k(\|\mathbf{h}\| - \|\mathbf{h}_{ij}\|)}{a_{\Pi_{\mathcal{N}_\mathbf{x}}}(\mathbf{h}_{ij})}, \qquad (1)$$

where $|\partial V_d|$ is the volume of a unit sphere in d dimensions, $a_{\Pi_{\mathcal{N}_\mathbf{x}}}$ is the autocorrelation function of the indicator function for $\mathcal{N}_\mathbf{x}$, i.e. $\Pi_{\mathcal{N}_\mathbf{x}}(\mathbf{y}) = 1$ for $\mathbf{y} \in \mathcal{N}_\mathbf{x}$ and zero otherwise, and k is a 1D normalized kernel (we use the Gaussian).

The intensity in the whole domain can be simply set as $\hat{\lambda}(\mathbf{x}) = \hat{\lambda}_x$. Similarly, the tensor field given by $\mathbf{M}_x$ can be interpolated or visualized as part of the analysis. However, setting a different PCF for each point in space adds many degrees of freedom, and hence makes qualitative and quantitative analysis, as well as synthesis difficult. Hence, we would like to compress the space of PCF's present in a given distribution.

3.3. Spatially Varying Correlations

It has been observed in previous works [OG12, HSD13] that the space of possible PCF's is rather limited, as valid PCF's lie on a subspace of low dimensionality (effectively 2 or 3 dimension). Due to this low dimensionality, all PCF's can be represented as linear combinations of a few basis PCF's. Although Öztireli and Gross [OG12] have shown that the subspace of the PCF's is approximately linear, our only assumption for analysis is its low dimensionality such that PCF's can be represented as linear combinations of a few basis PCF's (for synthesis, we will need the stronger linearity assumption, as elaborated on how we do synthesis). We thus would like to summarize the variability in $\mathcal{N}_\mathbf{x}$ with a PCF dictionary, and express the rest of the PCF's as linear combinations of the dictionary elements.

In practice, we will have a finite number of neighborhoods $\mathcal{N}_\mathbf{x}$ around certain points $\mathbf{c}_k$ in space. These points can be regarded as measurement points for the computed statistics. The corresponding PCF's $\hat{g}_k(\|\mathbf{h}\|)$ at $\mathbf{c}_k$'s are sampled and stored as vectors $\mathbf{g}_k$.

Given these $\mathbf{g}_k$, we would like to compute L basis PCF vectors $\mathbf{g}_l$ such that all others can be represented as a linear combination of these PCF's. Note that we would like to have a compact representation with an as small as possible L, while tolerating a certain error. The $\mathbf{g}_l$'s should also provide a meaningful summary of the adaptive correlations. We thus do not utilize sparseness based dictionary learning algorithms that generate an over-complete representation. Since there can be a variety of PCF's appearing in a given distribution, a simple clustering algorithm such as k-means clustering will also not give meaningful summaries, as can be observed in Figure 2.

Instead, we adopt an approach that is motivated by the structure of the PCF space. Öztireli and Gross [OG12] have observed that the main variation in the PCF space is due to the degree of irregularity in the generated distributions, and most variance can be captured with a few components. We thus first perform a PCA on the vectors $\mathbf{g}_k$ and reduce the number of dimensions such that we retain 99 percent of the variance. This typically results in three components, the first capturing most of the variance. We then choose the PCF's that are at the two ends of the line segment formed by the first component as shown in Figure 2. If these are already very close, we can conclude that the distribution is stationary. The weights w_{kl} are then computed as described below such that $\mathbf{g}_k = \sum_l w_{kl}\mathbf{g}_l + \mathbf{e}_k$. If $\max_k \|\mathbf{e}_k\| > e_g$ for a threshold e_g, the PCF that is furthest away from the already added PCF's is added. These two steps are performed alternately till the maximum error becomes sufficiently small. For each $\mathbf{g}_k$, the weights are computed by solving the following optimization problem for w_{kl} with quadratic programming:

$$\min \left\| \mathbf{g}_k - \sum_{l=1}^{L} w_{kl}\mathbf{g}_l \right\|^2 \qquad \sum_{l=1}^{L} w_{kl} = 1, \quad w_{kl} \geqslant 0. \qquad (2)$$

We illustrate an example of the chosen PCF's, corresponding patches, and weight maps for those PCF's in Figure 3. This is a synthetic example generated by our synthesis method we describe in the next section and thus comes with known ground truth PCF's and weight maps. Our analysis accurately figures out the two main PCF's and their interpolation weights.

3.4. The Analysis Framework

Our framework thus estimates spatially adaptive density, the basis PCF's $\mathbf{g}_l$, and the corresponding interpolation weights in space. We plot some of these diagrams for synthetic examples in Figures 1 and 3, and for a real example with distributions of locations of

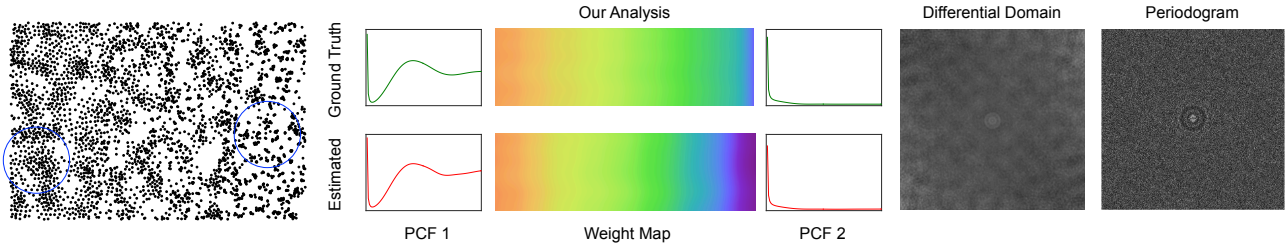

Figure 3: *An input distribution (left) and its analysis via our method, differential domain analysis [WW11], and periodogram. The patches that correspond to PCF 1 and 2 are marked on the point distribution. The weight map shown is w_1 for PCF 1. We do not show w_2 as it is equal to $1-w_1$. This is a synthetic example, with ground truth PCF's and weight maps provided. Our analysis successfully recovers the dominant PCF's and weights. Since the previous analysis methods mix different correlations, they cannot provide an informative summary.*

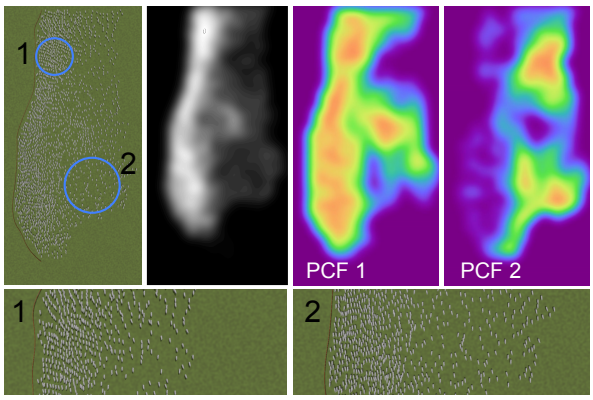

Figure 4: *(Top) From left to right: a real-world distribution of sheep, intensity map, and weight maps for the two marked patches. (Bottom) Zoomed regions around the marked patches.*

sheep acquired from a real scene in Figure 4. Note that before computing the PCF's, local anisotropy can be canceled as explained in the supplementary material, and the PCF's $\mathbf{g}_k$ are normalized with respect to density (Equation 1). Thus, we get the same PCF regardless of these degrees of freedom.

Discussion. The density estimate we compute is closely related to previous works that utilize such local modulations of difference vectors or distances [LWSF10, WW11]. The mentioned methods compute these directly from the spaces or functions to be estimated, e.g. replace Euclidean distance with geodesic distances, or compute the anisotropy tensors from an image to warp the difference vectors. After local normalization with density and anisotropy, previous analysis methods compute global statistics such as PCF's or periodograms on the whole dataset, leading to blending of different correlations, as discussed above and shown in Figure 3. In contrast to these approaches, our analysis technique can separate the important components of correlation models and density apart, and have spatially-adaptive correlations explicitly built in for handling complex point distributions.

Especially for clustering distributions, there is an inherent ambiguity on whether the intensity or PCF is causing the fluctuations in the density of points [IPSS08]. Hence, given a distribution such as in Figure 3, it is hard to disentangle these two different statistics. Our strategy is letting the user assume a certain degree of smoothness for intensity, such as the one we show in Figure 4.

In practice, in this work we assume local isotropy, as we work with small neighborhoods and we did not encounter locally strongly anisotropic distributions in practical cases. In the supplementary material, we describe how strong local anisotropy can be handled with standard methods from point processes.

Parameters. For computing the PCF's we use the same parameters as in a previous work [OG12]. All parameters are relative to r_{max}, the minimum distance between pairs of points for the maximum packing of points in a domain [LD08]. We then set $\sigma = 0.25$ for the Gaussian kernel in Equation 1, the lower and upper limits for the PCF to $r_a = 0.01\sigma$ and $r_b = 2.5$, respectively, and use a regular sampling of $\|\mathbf{h}\|$ with 100 samples to convert $g_k(\|\mathbf{h}\|)$ to the vectors $\mathbf{g}_k$. The smoothness of the intensity $\lambda(\mathbf{x})$ is a user given parameter, to disambiguate the intensity-correlation duality as described above.

4. Synthesis of General Sampling Patterns

Our synthesis algorithm follows the same model of local stationarity as elaborated on in the last section. We assume that the intensity λ_k and PCF $g_k(r) = g_k(\|\mathbf{h}\|)$ for each neighborhood $\mathcal{N}_k$ are provided or estimated from one or more example distributions with the proposed analysis framework. Then, the main idea of the synthesis algorithm is to generate a new distribution with statistics matching these target statistics.

4.1. The Synthesis Algorithm

We assume that we are given a domain and n points to be distributed to match the target characteristics. The intensity function $\lambda(\mathbf{x})$ can be scaled with a constant factor such that its integral is equal to n over the domain. For each neighborhood $\mathcal{N}_k$, the fitting error can then be computed as:

$$E_k(\mathcal{X}_k) = \int_{r_a}^{r_b} \left(g_k^E(r) - g_k(r) \right)^2 . \tag{3}$$

Here, $\mathcal{X}_k$ denotes the set of sampling points within $\mathcal{N}_k$, $g_k^E(r)$ is the estimated PCF from these points with the estimator in Equation 1, and $g_k(r)$ is the target PCF to be matched to. The gradient $\frac{\partial}{\partial \mathbf{x}_i} E_k$ is then computed and summed over all neighborhoods to get the final

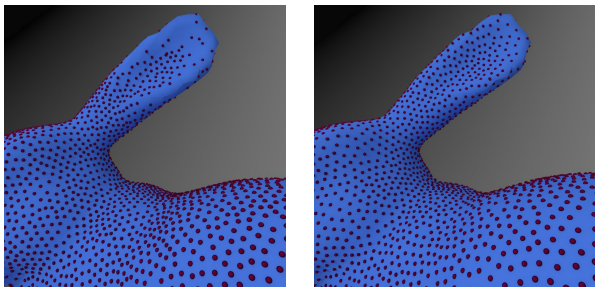

Figure 5: *Adaptive neighborhoods (right) can provide more accurate matching of target sample statistics than isotropic ones (left).*

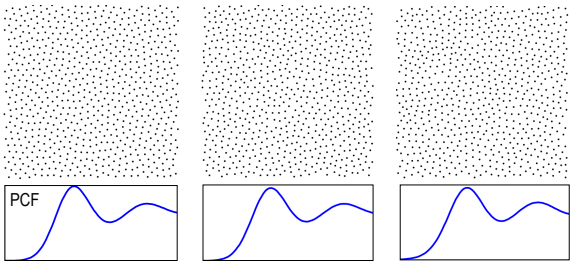

Figure 6: *A blue noise pattern [BSD09] (left) is used as the input for our texture synthesis approach (center) and for our PCF based synthesis algorithm (right).*

gradient $\Delta_i = \sum_k \frac{\partial}{\partial \mathbf{x}_i} E_k$ for point $\mathbf{x}_i$. We then perform a gradient descent $\mathbf{x}_i = \mathbf{x}_i - \lambda\Delta_i$, where we choose λ with a line search at each iteration. Similar to the previous section, we discretize the PCF's such that the integral in Equation 3 turns into a sum (we explain how this can be adapted to handle strong local anisotropy in the supplementary material).

Size and distribution of neighborhoods. We assume that the domain is divided into overlapping spherical neighborhoods $\mathcal{N}_k$ around center points $\mathbf{c}_k$. The volume $|\mathcal{N}_k|$ of each neighborhood is computed such that there are the same number of points in each $\mathcal{N}_k$, by setting $|\mathcal{N}_k| = \alpha n/\lambda_k$ for a constant factor α and the total number of points n. Since we assume that the point process is locally stationary in $\mathcal{N}_k$, the expected total number of points in $\mathcal{N}_k$ is then given by [IPSS08] $\int_{|\mathcal{N}_k|} \lambda(\mathbf{x}) = |\mathcal{N}_k|\lambda_k = \alpha n$. Note that by fixing the number of points in each neighborhood, we also fix the r_{max} (Section 3.4, parameters) when estimating the PCF's.

To get correctly blended characteristics, each neighborhood should also see points belonging to the others. For a neighborhood of radius R_k, we thus retrieve all sample points that fall into a hypersphere of radius $R_k + r_b$ when computing the gradients $\frac{\partial}{\partial \mathbf{x}_i} E_k$ for all $\mathbf{x}_i \in \mathcal{N}_k$. For constant $R_k = R$, the neighborhood centers $\mathbf{c}_k$ lie on a regular grid. The spacing of the grid is set as $T = R/2$.

Adaptive neighborhoods. In the case of adaptive λ_k, the neighborhood size and thus R_k changes. Decreasing the spacing, i.e. having neighborhoods such that $T \leqslant R/2$, does not alter the synthesized distributions, but degrades the performance. Hence, we utilize a conservative greedy non-uniform sampling of $\mathbf{c}_k$ for efficiency. We start from a sparse grid such that $T = \min_k R_k$. Each grid point $\mathbf{c}_k$ is then subdivided, starting from $\mathbf{c}_k$ with the largest λ_k, such that for each one-ring neighbor $T \leqslant R_k/2$.

Discontinuities in the density function violate our assumption of local smoothness in isotropic neighborhoods. For these cases, it is important to have an adaptive neighborhood that aligns itself along the discontinuity and thus avoids it. Typical examples of such cases are stippling images when the intensity changes abruptly, or geometry sampling when there are sharp features. For such discontinuities, we utilize adaptive neighborhoods computed by confining the neighborhood to one part of the discontinuity. For images, a bilateral filter on intensities, and for geometry on surface normals is first applied to cluster similar pixels/geometry points. Then, any neighborhood that contains different clusters is subdivided along the discontinuity. An example where we apply isotropic and adaptive neighborhoods for geometry sampling with blue noise is shown in Figure 5.

Initialization. We use a simple initialization strategy with random sampling. Around a randomly chosen neighborhood center $\mathbf{c}_k$, we iteratively pick a random point, and keep this point if all neighborhoods containing this point have not reached the desired number of points. We then discard the ones that already contain enough points, and continue with random sampling around the remaining neighborhood centers.

Non-ergodic processes. So far we have considered ergodic distributions where the statistics of the underlying point process can be estimated by observing a single distribution. There exist stationary point processes that are non-ergodic [IPSS08]. An important example that we encounter in practice is the locally regular distribution, where the points lie on a regular grid with fixed orientation but random global translation. Such a distribution is referred to as uniform or isotropic jittering in the literature [RAMN12, Ö16]. The statistics for this case cannot be extracted or matched to by considering just a single distribution, as it will have a constant global translation. In other words, expected values computed over many distributions are not equal to those over a single larger distribution.

For these cases, we propose to extend an existing discrete texture synthesis algorithm [ROM*15] as we elaborate in the supplementary material. For ergodic distributions, this method provides identical results to those generated by the synthesis algorithm presented in this section (Figure 6), as the same statistical measures are assumed for both cases. For all non-ergodic synthesis results in this paper, i.e. where locally regular distributions need to be generated, we employ this method.

5. Results and Discussion

We test our analysis and synthesis algorithms for a variety of distributions on Euclidean domains and curved spaces, and illustrate a series of applications for these generalized sets of distributions.

5.1. Analysis and Synthesis of Complex Distributions

We illustrate several examples with adaptive density and correlations in Figures 1, 3, 4. The distributions in Figures 1 and 3 are synthesized with known characteristics, i.e. PCF's and weight maps.

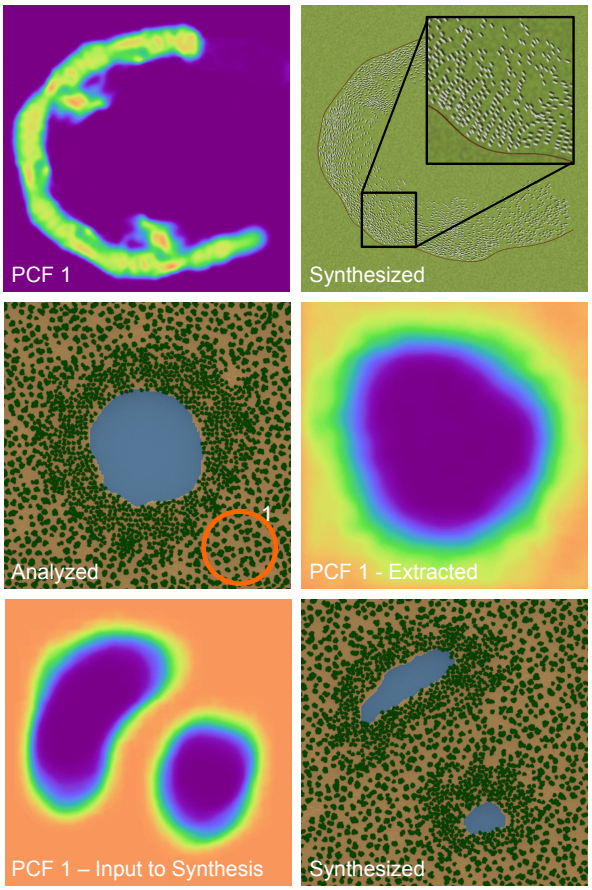

Figure 7: *(Top) The extracted PCF's and weight maps from a real-world distribution of sheep (Figure 4) are used to synthesize a new distribution with a new drawn weight map for PCF 1. (Bottom) Distribution of trees in a forest is analyzed, and a new forest with a custom weight map for PCF 1 is synthesized.*

Our analysis recovers the ground truth parameters, and our synthesis reproduces local PCF's very accurately. Notice that, while linear transitions in PCF's may not translate into visually linear transitions of distributions, our algorithm synthesizes new distributions with the same visual transitions of the given example distribution. In Figure 4, the characteristics of a sheep distribution are extracted from a real-world distribution. It can be observed that they form more regular structures near the fences on the left. We then take these weights and warp the one for PCF 1 to simulate a circular fence in Figure 7, and synthesize a distribution. The result accurately reproduces the distributional characteristics of the sheep in accordance with the environment.

A similar analysis and synthesis result is shown in Figure 7, second and third rows, for the distribution of trees in a forest. Instead of extracting from existing natural distributions, the characteristics can also be specified by pre-determined rules. In Figure 1, the trees are forced to form more regular distributions at lower altitudes, in accordance with real-world observations.

5.2. Image Sampling and Reconstruction

Being able to control local characteristics of point distributions opens new ways of improving irregular sampling for function reconstruction and representation. We illustrate an example to image plane sampling and reconstruction. Instead of using patterns with anti-aliasing properties such as blue noise, or merely adjusting density or metrics based on image content, we propose to steer both the density and correlations in order to obtain better image reconstructions from irregular samplings.

The main idea is that by distributing sampling points regularly along main image edges, and ensuring that we do not run into aliasing artifacts by a smooth transition to blue noise, we can get sharper and artifact-free reconstructions. We illustrate examples where we compare to a blue noise pattern and the bilateral blue noise sampling of Chen et al. [CGW*13] in Figure 8. We first distribute samples with different methods. For our synthesis algorithm, we extract the edges in an image with Canny edge detector, and smooth them to obtain the weighting map (Figure 8, left, insets) for the regular distribution. This is used to interpolate between the regular and blue noise distributions. We align the regular distributions with the edges by computing the local orientation of the closest edge.

Each sample carries a color value. These samples are used to reconstruct an image. For the reconstruction, we use isotropic Gaussian kernels (Figure 8, top) and iterative bilateral filtering (Figure 8, bottom), when comparing to blue noise and Chen et al. [CGW*13], respectively. Combination of regular and blue noise sampling leads to significantly improved results, especially around the edges, for both cases.

5.3. Image and Video Stippling

Such rules can also be defined for image stippling to generate alternative stippling styles. In Figure 9, we illustrate how combinations of blue noise and regular sampling can be utilized to generate stippled images with a novel style. For these images, the intensity at a pixel determines the density of the points as in previous works (e.g. [Fat11, ZHWW12]), and a smoothed edge map is used as the weight map for interpolating between a blue noise and regular distribution. Similarly to our image sampling results, we first detect the main edges of the image, and set the weight associated with the regular distribution for a neighborhood $\mathcal{N}_k$ to be inversely proportional to the distance between $\mathbf{c}_k$ and the closest edge, such that as we move away from the edges, we get a more blue noise type distribution. The smooth transition between the two correlation models result in artifact-free distributions. This is in contrast with distributions generated by simply defining separate regions for blue noise and regular sampling as illustrated in Figure 9, middle. For this case, additional structures appear around the edges, leading to visually unpleasant results.

We can develop and extend other stippling styles as well thanks to the generality of the distributions we can handle. In Figure 10, we show a different stippling style, which is closer to the one proposed by Kim et al. [KSL*08]. In this style, we extract stylistic smooth edge maps [KLC07] (Figure 10 top, middle), and compute the weight map for regular distribution (Figure 10 top, right) by dilating and smoothing them with a Gaussian. To achieve a

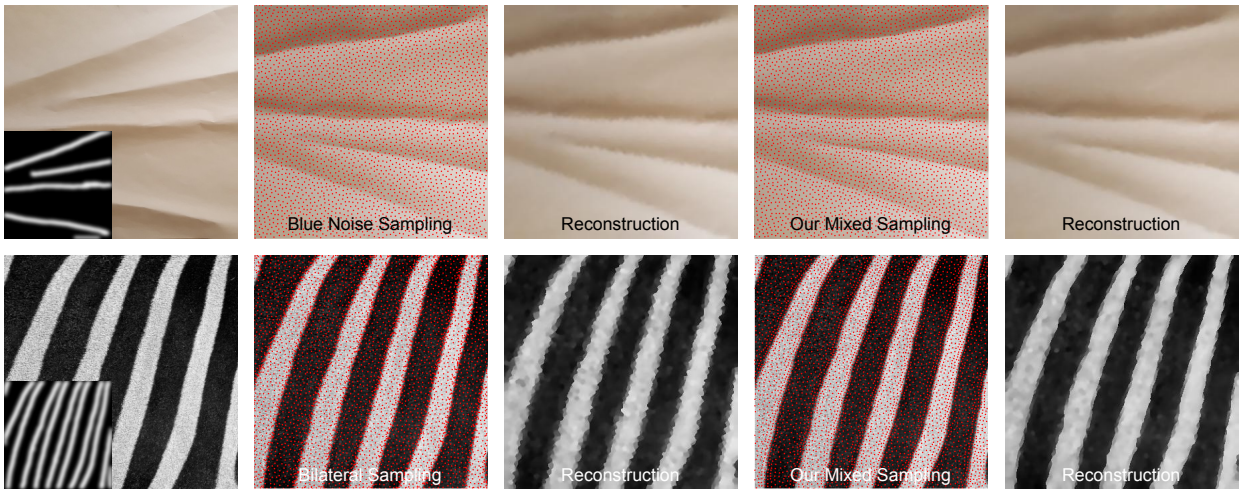

Figure 8: *Input images (with extracted and smoothed edges used for determining the weighting for regular sampling in our synthesis shown in the insets), sampling results with blue noise, bilateral blue noise [CGW*13], and our technique, and the corresponding recontructions with isotropic Gaussian kernels (top) and iterative bilateral filtering (bottom).*

more regular looking style, we generate larger regular regions compared to the previous results. We use the PCF of a blue noise pattern [BSD09] as the second PCF, and blend between this and the regular distribution using the computed weight map. The edge tangent flow [KLC07] (a smooth vector flow describing the salient edge tangent direction in the image) is used to orient the regular distribution along the extracted edges. The density is constant and the sizes of the dots are changed depending on the intensity of the image, as in the work by Kim et al. [KSL*08]. Finally, the extracted lines are shown together with the points.

Interpolating between oriented regular distribution close to the edges and blue noise elsewhere allows us to avoid artifacts in regions were multiple lines with different orientations intersect. Even though Kim et al. [KSL*08] propose additional controls to handle such cases, their method, purely based on regular distributions, can result in structured artifacts, as illustrated in Figure 11, left. By placing blue noise at the intersections, our method replaces these structures with blue noise instead (Figure 11, right).

Video stippling. The neighborhoods in our synthesis algorithm can also be extended in time to get spatio-temporally smooth sampling. We show an example of video stippling in the accompanying video, using the previously explained stippling style. For these results, in addition to blending a regular and a blue noise distribution, we set the PCF of a neighborhood at frame t to be a combination of its PCF and that of the same neighborhood in the previous frame. The weight assigned to the previous frame determines the trade-off between temporal smoothness and fidelity to the current frame. We found that a weighted average of PCF's where the PCF of the previous frame is given a weight of 0.25 works well in practice, i.e. $g = 0.25g_{t-1} + 0.75g_t$.

5.4. Geometry Sampling

Our synthesis algorithm extends to curved surfaces. Although recent works [JZW*15] explore adapting the spectrum of blue noise distributions on surfaces, to the best of our knowledge, there does not exist any algorithm to generate distributions with general characteristics on curved domains. We illustrate different noise patterns on surfaces in Figure 12. Similar to the planar case, our technique allows to generate smooth transitions between different patterns on surfaces as well, as shown for a combination of blue and green noise in Figures 12 and 1. We use a simple approximation of surface geodesics [BWWM10] in the synthesis algorithm. For every local patch, the distances are computed on the tangent plane of the surface, and the points are projected back to the surface after being moved, similar to previous techniques for blue noise sampling on surfaces [OAG10].

5.5. Performance

Most of the 2D distributions shown in the paper contain about 5000 samples and were generated with about 200 neighborhood centers $\mathbf{c}_k$. The stippling images with adaptive density (zebra) contain about 25000 samples and were generated with 10000 neighborhoods. The geometry sampling examples have up to 15000 samples and were generated with 2000 neighborhoods. To compute the gradient of a sample, the PCF's of all the neighborhoods in which the sample is included are taken into consideration. Thus, the performance is mostly influenced by the radii R_k and the spacing of neighborhood centers $\mathbf{c}_k$. On average, 25 iterations were needed until convergence. Our unoptimized, single core implementation takes up to one minute to complete one iteration on a PC with an Intel i7-3770K CPU. The bottleneck of the algorithm is the update of the PCF's of all the neighborhoods containing a point, after moving it, which could be optimized by computing some of the gaussians only once to update multiple PCF's. Furthermore, our synthesis algorithm can be significantly speeded up by moving multiple samples simultaneously, as they influence only their local neighborhood.

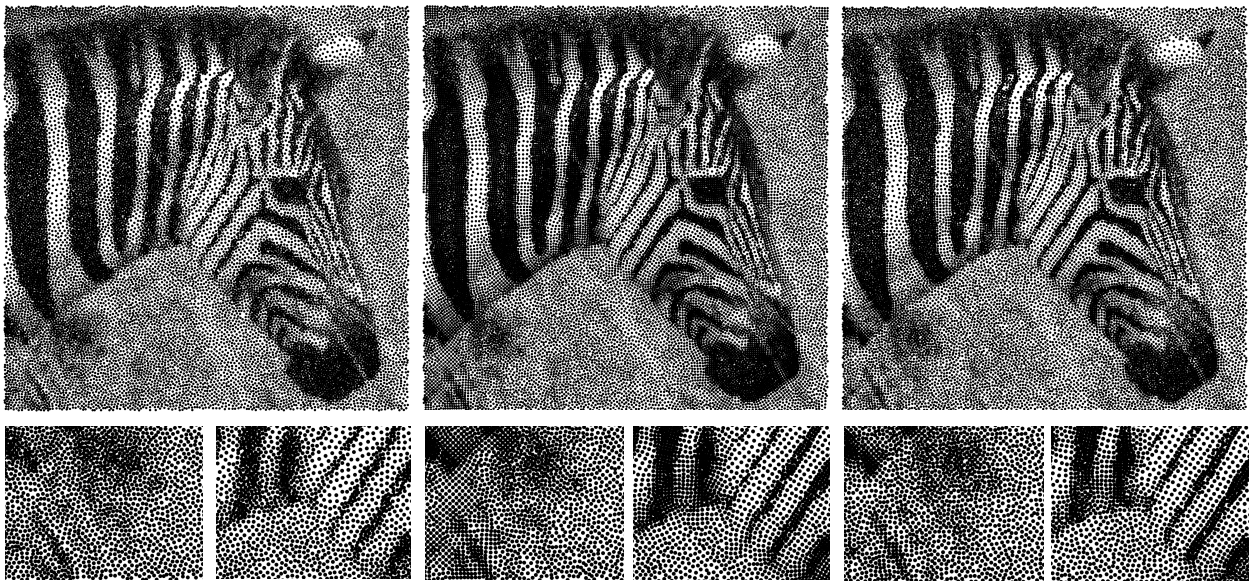

Figure 9: *From left to right: stippling with blue noise sampling, regular and blue noise sampling with a sharp transition, and regular and blue noise sampling with smooth transitions along the edges. The combination of blue noise and regular sampling with smooth transitions avoid artifacts, while providing a novel sharper stippling style, as illustrated in the insets below. This figure is best viewed on a computer screen, please zoom-in to see the details clearly. For vector graphics images we refer to the supplementary material.*

5.6. Limitations

Theoretically, it is possible that there exist point patterns that violate the local stationarity assumption, with dense discontinuities of the first and second order correlations throughout the space, although we have not encountered such distributions in practice. As we elaborate in the supplementary material, a more general analysis is possible, but will lead to major difficulties since it requires multiple instances of the same distribution.

When analyzing a given distribution, it is in general an ill-posed problem to determine the intensity, anisotropy, and second order correlations without certain assumptions. Our choices of a smooth intensity, a simple anisotropy model, and interpolated PCF's provide a set of such priors. Note, however, that the forward problem of synthesis does not suffer from this limitation, and we do not assume smooth density in that case.

Finally, our synthesis algorithm shares some of the limitations of previous PCF based fitting algorithms [OG12], as it simplifies to those for the case of a globally constant PCF. In particular, we cannot guarantee a minimum distance between point locations when synthesizing blue noise distributions, as the contribution of a pair of points can be blurred out by many others in the PCF. This can lead to small fitting errors.

6. Conclusions and Future Work

We introduced novel analysis and synthesis techniques for point distributions with adaptive density and correlations. The analysis framework provides an informative view of complex distributions with extracted maps capturing different distributional characteristics. Based on the same characteristics, the synthesis algorithm combines adaptive density and correlations and extends to general domains.

Sampling and anti-aliasing. The proposed general framework offers new possibilities for accurate representation or anti-aliasing of images, instead of the generic patterns with blue noise properties. Exploiting the redundancy of image patches, measures based on the local content such as edges, textures, and other structures can be computed and utilized to steer the synthesis algorithm for general images.

Rendering. Rendering involves computing integrals of complex functions. Traditionally, density adaptation of the samples via importance sampling has been the norm to improve the rendering quality by reducing noise while avoiding artifacts due to aliasing. Recent works [Dur11, RAMN12, SK13, SNJ*14] have shown that adapting density and correlations simultaneously can significantly improve the rendering results. However, the researchers have been limited by the analysis techniques [SK13] and pattern generation algorithms. Our techniques can be instrumental in developing new sampling methods and understanding existing patterns utilized in rendering such as for distributed ray tracing. As an example, combining adaptive correlations with the recent works that explore correlation models for rendering, e.g. [Ö16], can be an exciting future direction.

Understanding natural phenomena. Distributions and patterns in nature are inherently spatially adaptive due to environmental factors. We have illustrated that (e.g. Figure 4), it might not be possible to explain natural distributions by simply adaptive density. Our framework can be utilized to understand a spectrum of distributions ranging from surface details [YHJ*14], facial features [BBN*12], fluid particles [OG12], to crowds [JCP*10].

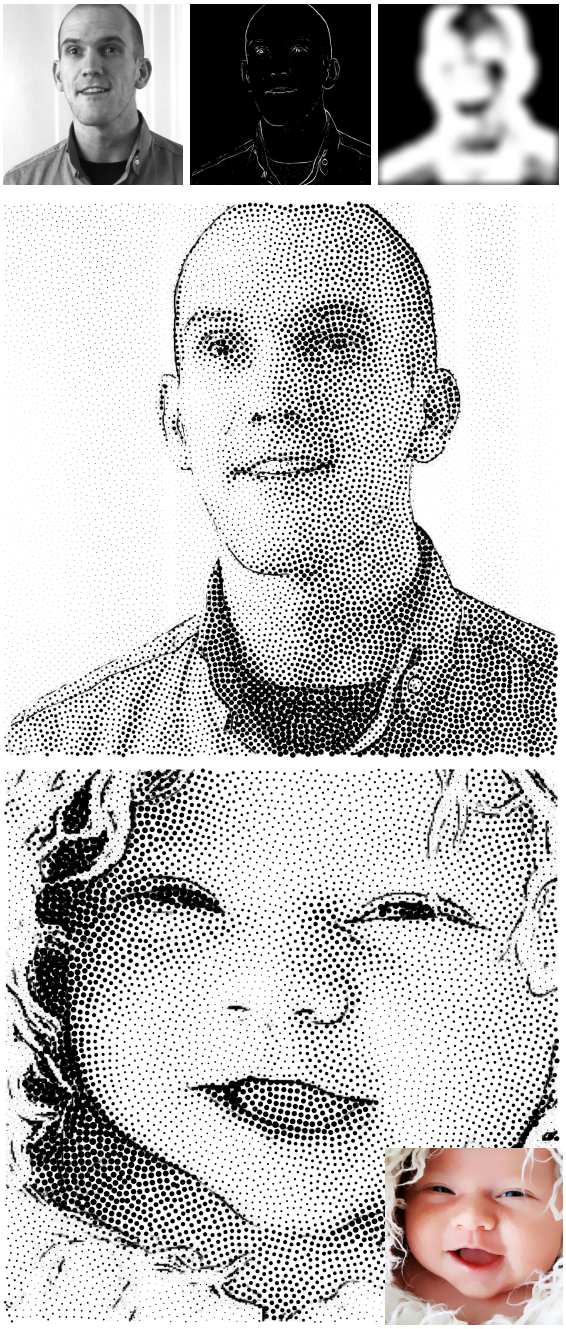

Figure 10: *Stippling with a different style. On the top row, we also show the extracted edges, and the corresponding smoothed map for the regular distribution.*

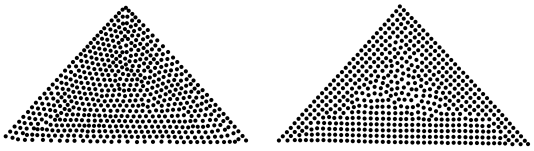

Figure 11: *Stippling a triangle using the method of Kim et al. [KSL*08] (left), and our style (right).*

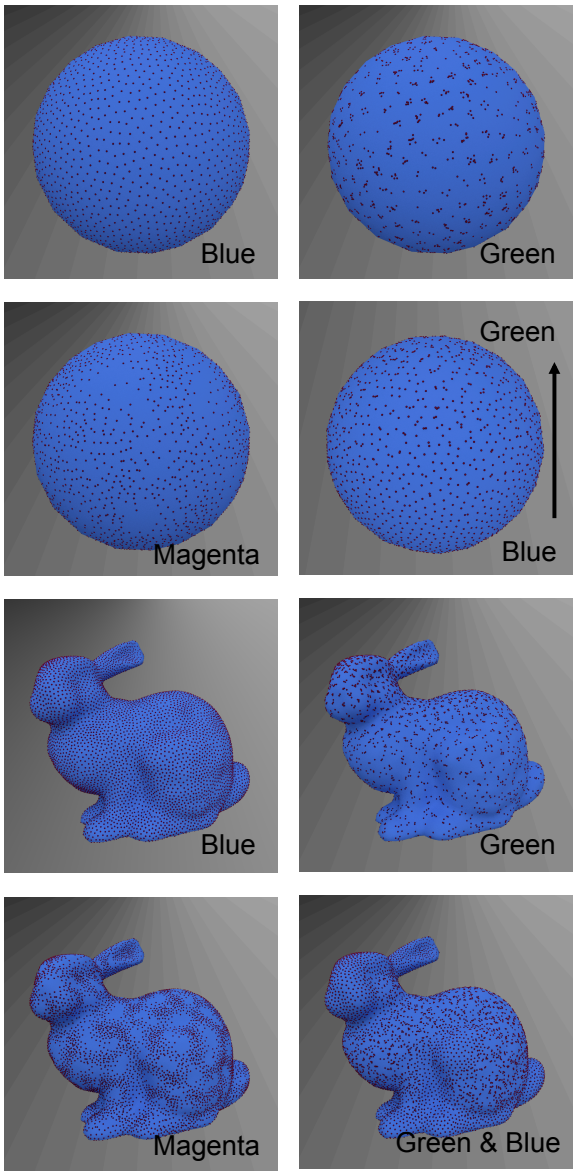

Figure 12: *Our synthesis algorithm can generate patterns with controllable characteristics and transitions on surfaces.*

Geometry reconstruction and remeshing. Reconstruction of geometry from point samples, and remeshing surfaces for rendering or simulations critically depend on the quality of the sampling patterns, given by both the density and the pair-wise correlations [OAG10, JZW*15]. Our synthesis algorithm extends full correlation control to surfaces, unlocking a significantly extended set of sampling patterns for geometry sampling. Similarly to our image reconstruction application, sharp features in surface reconstruction could be better preserved by aligning regular samples along them. For remeshing, a regular sampling could be ideal to generate quads, and a blue noise distribution could be adopted for transitions.

Marked and space-time processes. We have illustrated a simple application of spatio-temporal sampling for stippling videos with a trivial extension. A more in-depth analysis with space-time pro-

cesses can be used to principally extend our techniques. Similarly, extensions of the framework to multi-class sampling can be developed with the theory of marked processes [IPSS08].

Acknowledgments. Riccardo Roveri is supported by the Swiss National Science Foundation (grant No. 200021_146227).

References

[AHD15] AHMED A. G. M., HUANG H., DEUSSEN O.: Aa patterns for point sets with controlled spectral properties. *ACM Trans. Graph. 34*, 6 (Oct. 2015), 212:1–212:8. 2

[BBN*12] BEELER T., BICKEL B., NORIS G., MARSCHNER S., BEARDSLEY P., SUMNER R. W., GROSS M.: Coupled 3d reconstruction of sparse facial hair and skin. *ACM Trans. Graph. 31* (August 2012), 117:1–117:10. 9

[BSD09] BALZER M., SCHLÖMER T., DEUSSEN O.: Capacity-constrained point distributions: A variant of lloyd's method. *ACM Trans. Graph. 28*, 3 (July 2009), 86:1–86:8. 2, 6, 8

[BWWM10] BOWERS J., WANG R., WEI L.-Y., MALETZ D.: Parallel poisson disk sampling with spectrum analysis on surfaces. *ACM Trans. Graph. 29* (December 2010), 166:1–166:10. 2, 8

[CGW*13] CHEN J., GE X., WEI L.-Y., WANG B., WANG Y., WANG H., FEI Y., QIAN K.-L., YONG J.-H., WANG W.: Bilateral blue noise sampling. *ACM Trans. Graph. 32*, 6 (Nov. 2013), 216:1–216:11. 2, 7, 8

[dGBOD12] DE GOES F., BREEDEN K., OSTROMOUKHOV V., DESBRUN M.: Blue noise through optimal transport. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 171:1–171:11. 2

[Dur11] DURAND F.: *A frequency analysis of Monte-Carlo and other numerical integration schemes*. Tech. Rep. MIT-CSAILTR-2011-052, CSAIL, MIT,, MA, February 2011. 2, 9

[Fat11] FATTAL R.: Blue-noise point sampling using kernel density model. *ACM Trans. Graph. 30*, 4 (July 2011), 48:1–48:12. 2, 7

[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, 3, 4

[IPSS08] ILLIAN J., PENTTINEN A., STOYAN H., STOYAN D. (Eds.): *Statistical Analysis and Modelling of Spatial Point Patterns*. John Wiley and Sons, Ltd., 2008. 2, 3, 4, 5, 6, 11

[JCP*10] JU E., CHOI M. G., PARK M., LEE J., LEE K. H., TAKAHASHI S.: Morphable crowds. *ACM Trans. Graph. 29*, 6 (Dec. 2010), 140:1–140:10. 2, 9

[JZW*15] JIANG M., ZHOU Y., WANG R., SOUTHERN R., ZHANG J. J.: Blue noise sampling using an sph-based method. *ACM Trans. Graph. 34*, 6 (Oct. 2015), 211:1–211:11. 2, 8, 10

[KLC07] KANG H., LEE S., CHUI C. K.: Coherent line drawing. In *ACM Symposium on Non-Photorealistic Animation and Rendering (NPAR)* (Aug. 2007), pp. 43–50. 7, 8

[KSL*08] KIM D., SON M., LEE Y., KANG H., LEE S.: Feature-guided image stippling. *Computer Graphics Forum 27*, 4 (2008), 1209–1216. 7, 8, 10

[LD08] LAGAE A., DUTRÉ P.: A comparison of methods for generating Poisson disk distributions. *Comput. Graph. Forum 27*, 1 (March 2008), 114–129. 2, 5

[Llo82] LLOYD S.: Least squares quantization in pcm. *Information Theory, IEEE Transactions on 28*, 2 (Mar 1982), 129–137. 2

[LWSF10] LI H., WEI L.-Y., SANDER P. V., FU C.-W.: Anisotropic blue noise sampling. *ACM Trans. Graph. 29*, 6 (Dec. 2010), 167:1–167:12. 2, 5

[MF92] MCCOOL M., FIUME E.: Hierarchical poisson disk sampling distributions. In *Proceedings of the Conference on Graphics Interface '92* (San Francisco, CA, USA, 1992), Morgan Kaufmann Publishers Inc., pp. 94–105. 2

[Mit87] MITCHELL D. P.: Generating antialiased images at low sampling densities. *SIGGRAPH Comput. Graph. 21*, 4 (Aug. 1987), 65–72. 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. 3

[MW04] MØLLER J., WAAGEPETERSEN R. P.: *Statistical inference and simulation for spatial point processes*. Chapman & Hall/CRC, 2003, Boca Raton (Fl.), London, New York, 2004. 3

[Ö16] ÖZTIRELI A. C.: Integration with stochastic point processes. *ACM Trans. Graph. 35*, 5 (Aug. 2016), 160:1–160:16. 6, 9

[OAG10] ÖZTIRELI A. C., ALEXA M., GROSS M.: Spectral sampling of manifolds. *ACM Trans. Graph. 29*, 6 (Dec. 2010), 168:1–168:8. 2, 8, 10

[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, 3, 4, 5, 9

[Ost07] OSTROMOUKHOV V.: Sampling with polyominoes. *ACM Trans. Graph. 26*, 3 (July 2007). 2

[PH10] PHARR M., HUMPHREYS G.: *Physically Based Rendering, Second Edition: From Theory To Implementation*, 2nd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2010. 2

[PSC*15] PILLEBOUE A., SINGH G., COEURJOLLY D., KAZHDAN M., OSTROMOUKHOV V.: Variance analysis for monte carlo integration. *ACM Trans. Graph. 34*, 4 (July 2015), 124:1–124:14. 2

[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. 6, 9

[ROM*15] ROVERI R., ÖZTIRELI A. C., MARTIN S., SOLENTHALER B., GROSS M.: Example Based Repetitive Structure Synthesis. *Computer Graphics Forum* (2015). 6

[SGBW10] SCHMALTZ C., GWOSDEK P., BRUHN A., WEICKERT J.: Electrostatic halftoning. *Comput. Graph. Forum 29*, 8 (2010), 2313–2327. 2

[Shi91] SHIRLEY P.: Discrepancy as a quality measure for sample distributions. In *In Eurographics '91* (1991), Elsevier Science Publishers, pp. 183–194. 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, 9

[SNJ*14] SUBR K., NOWROUZEZAHRAI D., JAROSZ W., KAUTZ J., MITCHELL K.: Error analysis of estimators that use combinations of stochastic sampling strategies for direct illumination. *Comput. Graph. Forum 33*, 4 (2014), 93–102. 2, 9

[Uli88] ULICHNEY R.: Dithering with blue noise. *Proceedings of the IEEE 76*, 1 (Jan 1988), 56–79. 2

[Wei10] WEI L.-Y.: Multi-class blue noise sampling. *ACM Trans. Graph. 29*, 4 (July 2010), 79:1–79:8. 2

[WPC*14] WACHTEL F., PILLEBOUE A., COEURJOLLY D., BREEDEN K., SINGH G., CATHELIN G., DE GOES F., DESBRUN M., OSTROMOUKHOV V.: Fast tile-based adaptive sampling with user-specified fourier spectra. *ACM Trans. Graph. 33*, 4 (July 2014), 56:1–56:11. 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, 3, 5

[YHJ*14] YAN L.-Q., HAŠAN M., JAKOB W., LAWRENCE J., MARSCHNER S., RAMAMOORTHI R.: Rendering glints on high-resolution normal-mapped specular surfaces. *ACM Trans. Graph. 33*, 4 (July 2014), 116:1–116:9. 9

[ZHWW12] ZHOU Y., HUANG H., WEI L.-Y., WANG R.: Point sampling with general noise spectrum. *ACM Trans. Graph. 31*, 4 (July 2012), 76:1–76:11. 2, 7