

Anisotropic Noise Samples

Louis Feng, Ingrid Hotz, Bernd Hamann, and Kenneth I. Joy

Abstract—We present a practical approach to generate stochastic anisotropic samples with Poisson-disk characteristic over a two-dimensional domain. In contrast to isotropic samples, we understand anisotropic samples as non-overlapping ellipses whose size and density match a given anisotropic metric. Anisotropic noise samples are useful for many visualization and graphics applications. The spot samples can be used as input for texture generation, e.g., line integral convolution (LIC), but can also be used directly for visualization. The definition of the spot samples using a metric tensor makes them especially suitable for the visualization of tensor fields that can be translated into a metric. Our work combines ideas from sampling theory and mesh generation to approximate generalized blue noise properties.

To generate these samples with the desired properties we first construct a set of non-overlapping ellipses whose distribution closely matches the underlying metric. This set of samples is used as input for a generalized anisotropic *Lloyd relaxation* to distribute noise samples more evenly. Instead of computing the Voronoi tessellation explicitly, we introduce a discrete approach which combines the Voronoi cell and centroid computation in one step. Our method supports automatic packing of the elliptical samples, resulting in textures similar to those generated by anisotropic reaction-diffusion methods. We use Fourier analysis tools for quality measurement of uniformly distributed samples. The resulting samples have nice sampling properties, e.g., they satisfy a *blue noise property* where low frequencies in the power spectrum are reduced to a minimum.

Index Terms—tensor field visualization, glyph packing, anisotropic Voronoi diagram, blue noise

I. INTRODUCTION

NOISE samples with certain characteristics, such as spatially varying density and size, have many applications in computer graphics and visualization. The applications range from digital halftoning [1]–[4], mesh generation [5]–[8], Monte Carlo ray tracing [9], [10], to texture generation for visualization purposes [11]–[16]. While some of the desirable properties of the noise are similar across applications, the goals and appropriate sampling strategies are problem-dependent. For digital halftoning, the goal is to produce aperiodic isotropic patterns to avoid visually disturbing regular structures. For applications such as ray tracing, it might be important to have exact control over the number of samples. For mesh generation purposes, a primary goal may not be avoiding structural patterns but rather to produce an “optimal” triangulation to reflect a given anisotropic metric. Our interest in anisotropic spots with good sampling properties arose from a tensor field visualization method, where spot samples with spatially varying shape and density are used as input for line integral convolution (LIC) [17]. The shape and density of the spots determine the thickness and density of the resulting fiber structure. Any low-frequency patterns or holes in the spot samples will be visible as artifacts in the resulting texture, see Figure 13(a). Therefore, our method is designed to generate




Fig. 1. (a) LIC texture generated with a poorly sampled input texture. Holes in the spot texture result in holes in the LIC texture. (b) The same data set using our spot sampling method to generate the input.

anisotropic samples over a two-dimensional domain with the following properties:

- The samples vary in size and shape over a domain, represented by a Riemannian metric.
- The samples are mostly non-overlapping to approximate a generalized Poisson-disk characteristic.
- The samples cover the domain densely, without generating regular patterns.

The metric can be user-defined or derived from a scalar field, vector field, or tensor field. In general, it is spatially varying and anisotropic.

To achieve the objectives listed above we have designed a method which generates an anisotropic sample distribution in two main steps. First, we construct a set of non-overlapping ellipses whose distribution closely matches the underlying metric. This first sample set already exhibits most of the desired properties. Next, we use a generalized anisotropic Lloyd relaxation to distribute the noise samples more evenly. Our anisotropic Lloyd relaxation is a straight-forward generalization of the isotropic version, and coupled with the preprocessing steps, can produce both isotropic and anisotropic blue-noise. We use a localized anisotropic Voronoi cell and centroid definition based on the work of Labelle and Shewchuk [7] and Du et al. [6]. To speed up the relaxation we introduce a discrete approach which combines the Voronoi cell and centroid computation in one step, instead of computing the Voronoi tessellation explicitly. The total computation time is dominated by the relaxation step; the initial sampling time grows linearly with the required number of samples.

Anisotropic noise samples are appropriate for visualizing data sets with anisotropic non-uniform properties, and the most obvious application is the visualization of tensor fields. We have applied our method to stress tensor fields and diffusion tensor magnet resonance imaging (DT-MRI) data. Thereby we have used it as a visualization method in its own right

as well as as input for a continuous texture such as LIC. But anisotropic samples are also interesting for other areas in visualization where glyphs with spatially varying size and orientation are used. Our method automatically fills the entire domain with glyphs without clustering problems and reduces structural patterns to a minimum. A very different application is the use of anisotropic noise samples for artistic non-photorealistic rendering, where it opens a wide range of possibilities of which we discuss only some examples.

II. RELATED WORK

The generation of point or spot distributions with certain properties is the subject of research in different fields. Depending on the specific needs, many algorithms have been developed.

Generating uniformly distributed points with constant or varying density without large scale patterns has a long tradition in the area of noise generation, sampling or halftoning. These fields are closely related, many sampling algorithms are directly used to generate noise textures. Some techniques such as *dart-throwing* and others use a form of stochastic sampling, where random points are added or rejected according to certain criteria [18], [19]. Such methods often suffer from low convergence rates. Other approaches use relaxation techniques, in particular Lloyd’s relaxation [20], [21] and their variants, and they result in high quality blue-noise samples. To reduce the number of relaxation steps, a general strategy is to use a more efficient step to produce a set of samples which approximates blue-noise properties, followed by a few iterations of relaxation on this set to improve overall quality. To improve efficiency of the sampling algorithm further, several approaches have been suggested using tile sets which then are repeatedly tiled across the plane. Using this strategy, for example, Ostromoukhov et al. introduced a very efficient isotropic blue-noise sampling method based on *Penrose tiling* [9]. Most of these methods assume that the samples are isotropic. For a survey on sampling techniques, we refer the reader to [3], [18].

Anisotropic settings can be found in the area of stippling or automatic mosaic generation, where objects of different size and shape are distributed on a plane [22]–[24]. Different from our definition, the orientation of the distributed objects is not precisely predefined by the metric but can change during relaxation. Most of the proposed methods use a Lloyd relaxation based on a generalized Voronoi cell definition, where the Euclidean distance of the objects is approximated, using a hardware approach for computation of the Voronoi diagram [25]. The centroids used for relaxation are the geometric centers without considering the changed metric.

The goal of generating an anisotropic distribution following a given metric also appears in the area of mesh generation. Shimada et al. [8] based the mesh generation on a close packing of ellipsoidal bubbles. The packing is performed using a particle system, where particles move according to repulsive and attractive forces. The defined forces are not linear. The equations of motion are solved numerically to yield a force balancing configuration. A geometric approach for




Fig. 2. Generalized Poisson disk property. The minimum distance of two sample points is defined by the local ellipses which are not allowed to overlap.

anisotropic mesh generation was chosen by Du et al. [6] and Labelle et al. [7]. Labelle introduced an anisotropic Voronoi refinement of a set of sites which are required to appear in the mesh. In contrast to our goals, these sites are not allowed to move. The method by Du et al. uses a centroidal Voronoi relaxation to compute a triangulation. Both methods define a generalized Voronoi tessellation based on a non-Euclidean metric using different distance approximations as basis for the final triangulation. Our work builds on the ideas introduced in these methods. We have adapted the distance measures introduced by Labelle et al. since they match our demands better than the idea of centroidal relaxation of Du et al.

The use of glyphs for visualization of local field properties is common in visualization. The question of placing these glyphs has been subject of discussion in several contexts. The most common strategies are regular sampling, random sampling with or without Poisson property [14], [26] or procedural texture generation, e.g. using reaction diffusion. In vector field visualization, Turk and Banks proposed a method to place arrows along streamlines generated by streamline optimization [27]. Kindlmann introduced reaction-diffusion into the visualization community by applying it to diffusion tensor MRI data [15]. Sanderson et al. used a reaction-diffusion model to generate spot noise based on the underlying vector field placing glyphs at the spot center [28]. Reaction diffusion simulates a chemical process between two morphogenes, which react and diffuse. This changes their concentration and under certain conditions converge toward stable pattern representing the dynamic equilibrium of the system. The method provides automatic control of density, size and placement of patterns but the specification of appropriate parameters is not trivial. Stable patterns only form for a very narrow band of values for the parameters. In addition, it is computationally expensive. Recently Kindlmann and Westin proposed a glyph packing algorithm in the context of diffusion tensor visualization [16]. Their work is built on a particle approach simulating attractive and repulsive forces.

III. ASSUMPTIONS AND GOALS

The starting point for the generation of the elliptic noise samples is a metric g given over a domain $D \subset R^2$ which defines the sample properties. The metric can be user-defined or derived from scalar fields, vector fields, or tensor fields, see Section VI. The metric is given as a two-by-two symmetric, positive definite matrix depending on the location

$$P = (x, y) \in R^2,$$

$$g(x, y) = \begin{pmatrix} g_{11}(x, y) & g_{12}(x, y) \\ g_{12}(x, y) & g_{22}(x, y) \end{pmatrix}. \quad (1)$$

We assume that the metric is non-degenerate everywhere. In general, it is spatially varying and anisotropic. The goal is to generate an unstructured distribution of sampled ellipses, whose size and density matches the given metric. The ellipses are specified by their center $P_0 = (x_0, y_0)$ and their shape, which is defined as the unit circle with respect to the metric g_0 in P_0 , i.e.,

$$g_{011}(x-x_0)^2 + 2g_{012}(x-x_0)(y-y_0) + g_{022}(y-y_0)^2 = 1. \quad (2)$$

Their half-axes are aligned to the eigenvectors and their squared principal radii $a^2(x_0, y_0)$ and $b^2(x_0, y_0)$ are scaled according to the reciprocal eigenvalues

$$a^2(x_0, y_0) = \frac{1}{\lambda_1(x_0, y_0)} \quad \text{and} \quad b^2(x_0, y_0) = \frac{1}{\lambda_2(x_0, y_0)}, \quad (3)$$

where $\lambda_1(x_0, y_0)$ and $\lambda_2(x_0, y_0)$ are the eigenvalues of $g(x_0, y_0)$. The sample density is implicitly defined by the size of the ellipses. In order to make a glyph based visualization reasonable we further assume that the frequency of the generated spots is higher than the frequency of the change of the underlying metric. This means that density and eigendirections do not vary much from one sample to its neighbors. In summary, we have designed our algorithm to generate noise samples with the following properties:

- 1) The size and shape of the spots are determined by the local metric. By choosing the right scaling we can define the spots as unit circles, see Equation 2.
- 2) The spots are closely packed and non-intersecting having a minimum distance, defined by a generalized Poisson disk property, see Figure 1.
- 3) The spots are unstructured, i.e., inherent patterns are reduced to a minimum. In terms of the power spectrum of the samples this means absence of small frequencies. In contrast to the general blue-noise property we allow an angular dependency of the principal frequencies.

IV. ALGORITHM

Texture generation can be divided into two independent steps:

- A. Computation of a reasonable starting distribution of ellipses, where we generate a set of spot candidates based on a dense set of uniformly sampled jittered points, and then traverse the candidate set to select ellipses such that the resulting distribution fulfills a generalized Poisson disk property. This start distribution provides the basis for most of the properties of the resulting sample set.
- B. Optimization of the starting distribution using an anisotropic Lloyd relaxation.

Both steps are important. The first step determines the basic properties of the sample set, including the number of samples, the density and the Poisson disk property. The second step makes the samples more uniform and approaching stable configurations. It is important to use a sufficiently complex

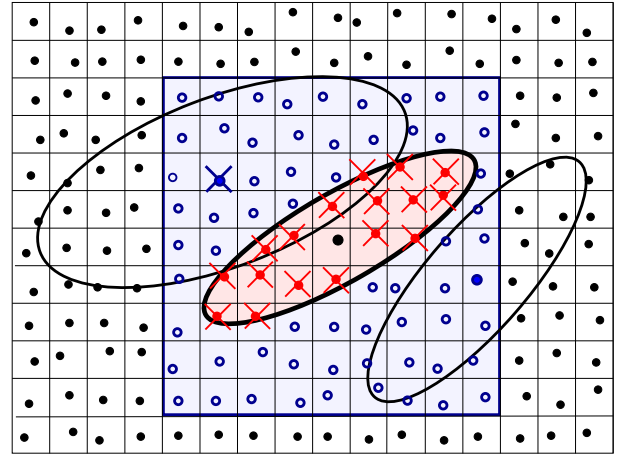


Fig. 3. First step of noise generation. Generation of a reasonable starting distribution of ellipses. A random point set defines candidate points, shown as black dots. Ellipses around these points still overlap each other. A point is accepted as sample if its ellipse does not intersect with any other previously accepted ellipse. In this example, the red ellipse is the first ellipse drawn in the considered region and is accepted. All points lying inside the ellipse, shown as red points, can immediately be removed from the candidate list. For simplicity, all points in the blue square containing the red ellipse are checked for intersection and are possibly removed. In this example, one of the two ellipses drawn would be removed.

first sample set as input for the relaxation and only a few iterations, since the relaxation tends to converge to stable pattern configurations, especially for uniform settings. We explain these steps in more detail in the following sections.

A. Generating the Initial Sample Set

The generation of the initial sample set is done in two steps. In the first step, a set of jittered grid points is generated as locations for the candidate spots, see Figure 2. The initial set must have higher density than the target density. Four times the target density leads to good results. When using grid points without jittering the required initial density will be higher and may still introduce structures. The candidate spots in each location are defined by the local metric as “unit-circle.” For a general metric g , these are ellipses defined by Equation 2. At this stage, the generated candidate spots generally overlap. Once the initial set of points is generated, the algorithm traverses the set of points and determines whether each candidate spot should be accepted or not. A candidate is selected, when it satisfies the following conditions:

- The point has not been checked previously, and
- it fulfills the generalized Poisson disk property, meaning its ellipse does not overlap with the ellipses at any other already selected spots in S .

While traversing we check whether a point overlaps with other spots. The underlying regular grid structure of random points has the nice property that it supports efficient spatial search of neighboring points, therefore, simplifying the checking process. After selecting an ellipse, all candidate points in the immediate neighborhood of the ellipse are checked for intersection and are possibly removed from the point list. For simplicity we use a squared neighborhood with edge length

of twice the size of the larger ellipse half axis, see Figure 2, instead of a the ellipses themselves.

A simple way to traverse the sample set is visiting the points sequentially in a scan-line fashion. Unfortunately, traversing the points sequentially introduces artifacts in the resulting spot texture. Using a Hilbert curve, i.e., a special space filling curve, to traverse the points as proposed by Gilbert [29] leads to better results but still generates undesired patterns. We have achieved best results with a random traversal of the spots. Differences when using a random traversal or Hilbert curve traversal are rarely visible after the relaxation step.

B. Anisotropic Voronoi relaxation

After generating spot samples and selecting ellipses the sample set of ellipses has the following properties:

- The ellipses in the current sites represent the local metric.
- The ellipses do not overlap.
- The coverage of the domain with ellipses is dense.
- The sites do not follow a regular pattern.

By eliminating overlapping samples, holes can result in certain areas. To remove these artifacts we use a method similar to *Lloyd relaxation*.

Lloyd relaxation, also known as Voronoi iteration, is a method to generate evenly distributed samples. It is an iteration of constructing Voronoi tessellation and its centroids. In each iteration the sample points are moved into the cell centroid which corresponds to the center of mass of the cell. The process converges to a centroidal Voronoi diagram, where each sample point lies in its cell centroid. This diagram minimizes the energy given as

$$E = \sum_{i \in I} \int_{Vor_i} \rho(x) \|r - r_i\|^2 dr \quad (4)$$

where I is the index set for the samples, Vor_i the Voronoi cell of the i th sample, r_i its position and ρ a local scalar density.

Due to the anisotropy of the metric, we use an anisotropic Voronoi diagram and an anisotropic centroid computation for the relaxation step. Since the orientation of the ellipses is determined by its location it changes when moving its position. In contrast to Fritzsche et al.'s mosaic generation technique using Voronoi structures [24], our method does not perform explicit rotation to achieve a denser packing because the ellipses must reflect the local metric. The relaxation step does not change the density resulting from the first step much; it merely is a smoothing of the sample distribution. Therefore, the quality of the result after the relaxation process depends strongly on the result of the first sample set generation.

For the definition of the anisotropic Voronoi diagram and the centroid computation we built on the works of Labelle and Shewchuk [7] and Du et al. [6]. Our method is a combination of these two methods, satisfying our demands.

a) *Definition of the Voronoi regions:* Let $\{P_i \in D, i \in I\}$ be the set of sample points resulting from our previous step, where I is an index set for the points. The most natural way of generalizing the Voronoi tessellation to other more general metrics would be to define a Voronoi cell $Vor_g(P_i)$ of a point P_i as the set of all points $P \in D$ that are at least as close to




Fig. 4. (a) The anisotropic bisector of two points is a quadratic curve. (b) In the uniform case the bisector is a straight line passing through the midpoint of the two points. In general, it is not orthogonal to the connecting line segment of the two points.

P_i as to any other point $P_j, j \neq i$, using the geodesic distance as distance measure, i.e.,

$$Vor_g(P_i) = \{P \in D | d_g(P_i, P) \leq d_g(P_j, P) \text{ for all } j \in I \text{ with } i \neq j\}, \quad (5)$$

where $d_g(P, Q)$ is the geodesic distance of the two points P and Q defined by the metric g . This is the length of the shortest path connecting the two points P and Q . However, since the computation of this shortest path is difficult and computationally expensive we use an approximate distance function for two points.

Different distance approximations in the context of the definition of anisotropic Voronoi cells have been introduced. We follow the definitions proposed by Labelle and Shewchuk, because it matches our conditions well, i.e.,

$$d^2(P, Q) = (P - Q)^T g(P) (P - Q) \quad (6)$$

In this case, the distance is not symmetric $d(P, Q) \neq d(Q, P)$. Also, the triangle inequality is not necessarily satisfied. The bisector of two points P_1 and P_2 is defined as

$$B(P_1, P_2) = \{P | d(P_1, P) = d(P_2, P)\}, \quad (7)$$

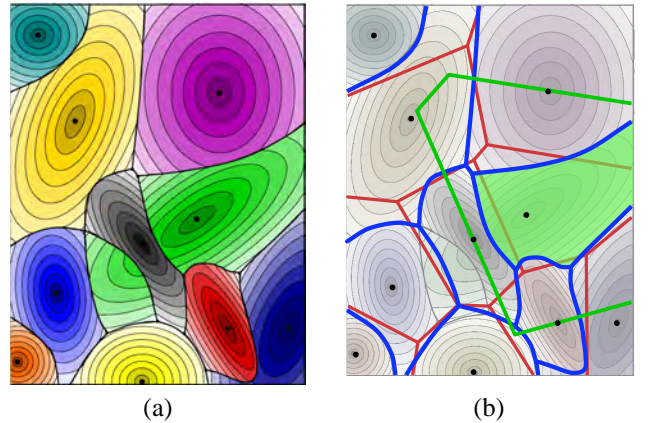


Fig. 5. (a) Voronoi cells resulting from the approximate distance function defined by the metric in the sample points. The green Voronoi cell is an example for a non-connected cell, having one orphan. Such orphans are undesirable for the relaxation process. (b) Comparison of the various Voronoi cells. The red lines show the boundaries of the standard Euclidean Voronoi cells. The gray lines show the boundaries of the anisotropic Voronoi cells with orphans as defined by Equation 8. Most of the lines are covered by the blue lines, which represent the localized anisotropic Voronoi cells as defined in Equation 9. The green line indicates the area to which the green Voronoi cell is restricted.




Fig. 6. The set $\{P | (P_i - P_j) \cdot (P - P_j) \geq 0\}$ defines an extended neighborhood of P_i (green) which contains the Euclidean Voronoi cell (red). We use this neighborhood to define a localization of our Voronoi cells.

and is a quadratic curve. If the metrics in P_1 and P_2 are the same, the bisector is a straight line passing through the midpoint of the connection of P_1 and P_2 , see Figure 3. Based on this approximate distance, a Voronoi cell of point P_i is defined as

$$Vor(P_i) = \{P \in D | d(P_i, P) \leq d(P_j, P) \text{ for all } j \in I \text{ with } i \neq j\}. \quad (8)$$

In contrast to Du et al. [6], this definition uses the metric in the point samples to determine distance. We have chosen this distance function because it guarantees that ellipses drawn in the sample points lie entirely inside the Voronoi cells for our start configuration with non-intersecting ellipses. This results from the fact that the ellipses are unit circles according to the local metric, meaning all points inside the ellipses have a distance smaller than one, and all points outside have a distance larger than one.

The resulting Voronoi cells are in general not convex and may not even be connected, see Figure 4(a). The part of the cell not containing the sample point is called an *orphan*. Orphans are undesirable for the relaxation process for two reasons. First, the relaxation should be determined by the local neighborhood of the sample point and not influenced by regions far away. Second, the distance approximation for those regions is very inaccurate. Therefore, we define a localized version of the Voronoi cells by reducing the possible region of influence using the Euclidean Voronoi tessellation, see Figure 5. It is possible that some of the orphans survive the localization but this happens very rarely, see Figure 4(b). We define

$$Vor_r(P_i) = \{P \in D | i \in I_P \text{ and } d(P_i, P) \leq d(P_j, P) \text{ for all } j \in I_P \text{ with } i \neq j\}, \quad (9)$$

with $I_P = \{i \in I | (P_i - P_j) \cdot (P - P_j) \leq 0, \forall j \neq i\}$.

b) *Centroid definition:* For the definition of the centroid we follow the idea of Du et al. [6], which is a straight-forward generalization of centroid definition as the center of mass to an anisotropic setting. The center of mass c_i of a Voronoi cell $Vor(P_i)$ is defined as

$$c_i = \frac{\int_{Vor(P_i)} d(r) r \, dr}{\int_{Vor(P_i)} d(r) \, dr}, \quad (10)$$

where d is an isotropic scalar density and $r = (x, y)$. By replacing the density d by the metric tensor g the centroid c_i is defined as

$$c_i = \left(\int_{Vor(P_i)} g(r) \, dr \right)^{-1} \cdot \left(\int_{Vor(P_i)} g(r) \cdot r \, dr \right). \quad (11)$$

As an integral over positive definite matrices, the left matrix is always invertible. By using an isotropic metric with

$$g(r) = \begin{pmatrix} d(r) & 0 \\ 0 & d(r) \end{pmatrix}, \quad (12)$$

this definition reduces to the standard weighted centroid definition, see Equation 10. If the metric is uniform, i.e., it does not depend on r , the anisotropic centroid definition coincides with isotropic uniform case.

After moving the sample point to the centroid of its cell we cannot guarantee that the ellipse is still entirely contained in the Voronoi cell and thus does not intersect other ellipses. This can be avoided by performing an intersection test with the neighboring ellipses and shortening the translation vector in case of intersection. This happens mostly for uniform data sets and close to the boundary. In our experience this check is not necessary for real data sets. For a more detailed discussion, see Section V.

C. Implementation:

a) *Intersection test:* The initial sampling requires intersection tests between neighboring samples. In the isotropic case, this intersection test is simply the circle to circle intersection test and can be done efficiently. In the more general case, the samples are represented by ellipses. The algebraic method of ellipse to ellipse intersection test involve solving a quartic polynomial which is computationally expensive and numerically unstable. We use polylines to approximate the ellipses during intersection test to reduce complexity. In our implementation, we approximate each ellipse by a closed polyline consisting of eight line segments. This approximation produces good results without the issues involved in the algebraic method.




Fig. 7. Anisotropic Voronoi cells resulting from the simplified distance measure. The grid used for relaxation has a resolution of 512×512 .

b) *Relaxation*: For the computation of the Voronoi cells and the centroid we use a discrete approach. Considering the domain as a set of uniform cells represented by their center R , the discretized version of Equation 11 results in

$$\underbrace{c_i - P_i}_{\equiv T_i} = \left(\underbrace{\sum_{R \in \text{Vor}(P_i)} g(R)}_{\equiv M_i} \right)^{-1} \cdot \underbrace{\sum_{R \in \text{Vor}(P_i)} g(R) \cdot (R - P_i)}_{\equiv t_i} \quad (13)$$

T_i , t_i , and M_i are used in the implementation.

Instead of computing the Voronoi cell explicitly and using these cells for the centroid computation, we perform both computations in one step. We initialize all sample positions $P_i, i \in I$, with a zero vector t_i and zero matrix M_i . Next, we march through the discretized domain performing the following steps for each cell represented by the point R :

- Find the Voronoi Cell $\text{Vor}_R(P_i)$ containing point R , then compare its distances to sample points lying inside a local bin to specify i .
- Update the matrix M_i and the vector t_i in the following way.

$$\begin{aligned} M_i &\leftarrow M_i + g(R) \\ t_i &\leftarrow t_i + g(R) \cdot (R - P_i) \end{aligned} \quad (14)$$

After traversing the entire domain, the new position of the sample points P_i , given by Equation 13, is determined by the translation vector T_i , i.e.,

$$T_i = M^{-1} \cdot t_i \text{ and } P_i \leftarrow P_i + T_i. \quad (15)$$

An example of the computed Voronoi cells of a synthetic dataset is shown in Figure 6. As the resolution of the discretization increases, the centroid can be computed more accurately. In our implementation, we have used grids of resolutions between 256×256 and 1024×1024 depending on the desired output resolution of the samples.

V. QUALITY MEASUREMENT

The evaluation of our algorithm is guided by the goals described in Section III. We use spatial (density, uniformness) as well as frequency (power spectrum) properties for quality measurements. In the following, we discuss the points defined as goals and the methods used to verify their quality. We discuss examples for a simple isotropic and anisotropic metric definition. The tensor used for the uniform example is aligned to the coordinate axes with eigenvalues 2.0 and 8.0. For the non-uniform example the eigenvector field is rotated about $\pi/2$ from left to right. The tensor field changes from an almost isotropic field with eigenvalues 2, to an anisotropic field with eigenvalues 2.0 and 8.0, respectively.

a) *Representation of the metric by sampling shape and density*: The size and shape of the spots are determined by the local metric. Thus, each spot reflects the metric values in the sample point exactly. The density of the spots is roughly determined by spot size. For an evaluation of density we consider two different density measures. The first is the standard scalar density d which is defined as covered pixels per unit area. This density can be measured by using a Gaussian




Fig. 8. Relaxation for uniform anisotropic samples. (a) first sample set, (b) after one iteration, (c) after six iterations.




Fig. 9. Sampling sets for an anisotropic spatially varying metric. The upper row shows how the sample set evolves over the computation. The row below shows the image after applying a Gaussian filter. (a) First sample set after the first step of our algorithm; (b) result after one relaxation step; and (c) result after six relaxation steps.

filter, where the local density is given as gray value. Due to the discrete structure of the samples we cannot expect a constant density but it should be almost uniform over the entire domain even for non-uniform samples. An example for a uniform anisotropic data set is shown in Figure 7, an example for a nonuniform data-set in Figure 8. The coverage is computed by averaging the pixel grey value of the output image. For a visual analysis, we used blurring to measure the uniformity of the samples. The size of the Gaussian filter used for these examples is the same for both examples. In both cases it can be observed that the density is fast approaching a uniform distributions. After six relaxation steps there are no holes visible anymore. In particular there is no dependency of the coverage on the size and shape of the samples. Close to the boundary, a slightly higher density can be seen for both data sets. The entire coverage of the space is approximately 40% for the uniform and the nonuniform case.

The second density we are interested in is an anisotropic density. We define the density $d_a(\theta)$ for samples with non-zero area as average number of sample hits when traversing a unit area along a straight line in direction θ , see Figure 9(a).

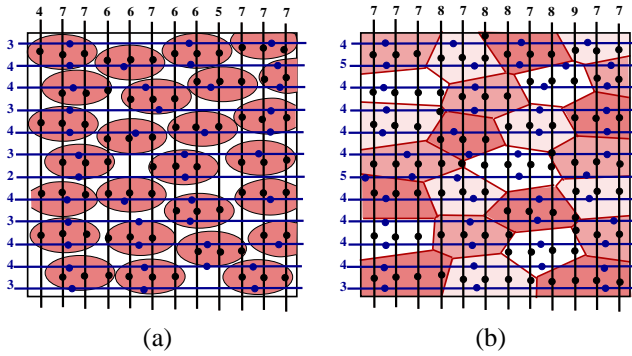


Fig. 10. Definition of an anisotropic density as average number of sample hits when traversing a unit cell in a certain direction. Here, sample hits for two different directions are shown. The average number of sample hits in horizontal direction (blue) is circa 3.4 and circa 6.25 in vertical direction (black).

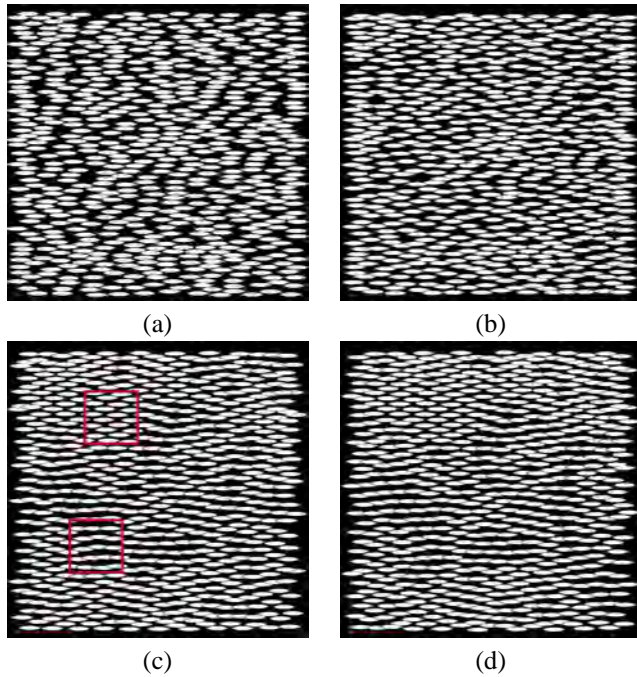


Fig. 11. For a uniform anisotropic data set, regions of regular patterns may form after running the relaxation for many iterations. (a) shows the start configuration before relaxations, (b) after one iteration, (c) after 50 iterations and (d) after 125 iterations. After 50 iterations we can see the formation of two different structures (highlighted by the red box). These regions hardly change over the next 75 iterations. Looking at difference images between the iterations we can see that there is still a fluctuation in the regions between the stable patterns.

Practically we count the number of Voronoi cells we hit for a given direction, see Figure 9(b). This value is especially interesting when we use the samples as input for LIC generation, because it determines the fiber density orthogonal to the convolution direction. The ratio of the anisotropic density in principal direction and the half-axis of the ellipses should be the same for both eigendirections. For the uniform case as shown in Figure 7(c) the two ratios vary about 1.5%.

b) Generalized Poisson disk property: By construction, the generalized Poisson disk condition is satisfied after the first sample set generation. The relaxation itself does not guarantee that this property is maintained. Since the samples

tend to move towards holes in the texture and the overall density is not changed significantly during the relaxation, in most cases the Poisson disk property is not violated after several iterations, except close to the boundary. In the case of uniform data-sets hexagonal patterns are fix-points for the iteration process. After several iterations regions with hexagonal patterns are forming, see Figure 10. Inside of these regions the pattern become relatively stable fairly quickly. In between these regions the structure still changes slightly after many iterations. This can lead to an unwanted overlapping of the ellipses. To avoid overlapping we have implemented an additional intersection test during the relaxation. Since for non-uniform data-sets there are no such regular stable configurations this test is in general not necessary, we mostly skip it for performance reasons.

TABLE I
PERFORMANCE

Grid size	# initial samples	Removing overlaps	# final samples	Relaxation
Anisotropic				
128x128	16384	0.188s	259	6.266s
256x256	65536	0.812s	1038	26.437s
512x512	262144	3.313s	4254	104.39s
Isotropic				
128x128	16384	0.047s	786	0.078s
256x256	65536	0.234s	3075	0.281s
512x512	262144	0.922s	12217	1.141s

c) Unstructured sampling: Half-toning methods are usually evaluated by their radially averaged power spectrum. This measures the frequency content of the spatial distribution of points, also called “blue-noise characteristic.” Low frequencies in this spectrum lead to large-scale artifacts, that will become visible. For the isotropic case the expected pattern in Fourier space shows no frequencies close to the origin, and it shows a maximum at a radius representing the principal frequency, representing the inverse average distance of samples. Since we use elliptical Poisson disks, this principal frequency (PF) is direction-dependent, varying from a minimum to a maximum PF. The PFs are visible in frequency space as light circle for the isotropic case and as light ellipse for the anisotropic case, see Figure 17 and 18. The higher frequencies representing the randomness are almost homogeneous and independent of the direction even for the anisotropic case.

The relaxation process degrades the anisotropic structure in the power spectrum which goes hand in hand with the pattern formation and overlapping of ellipses mentioned in the last section. Similar in isotropic settings, one can show that all hexagonal structures are stable under uniform anisotropic relaxation. The anisotropy is expressed in the shape of the cells but not based on the locations of the centroids. The frequency pattern approaches the standard blue noise pattern as we apply more relaxations but still maintains anisotropy after the samples stabilize at final position. This tendency can be diminished by a small margin when using the intersection test during the relaxation process to ensure non-overlapping samples. The resulting samples with intersection test produce some irregularities which can be seen in the frequency domain,

see Figure 19.

d) *Performance*: Table ?? shows the running times on an Intel P4 3.2GHz with 2GB memory. Important observation is the linear growth of the processing time for both isotropic and anisotropic tests. In the isotropic case the time for the generation of the first sample set and the relaxation steps is about the same. For the anisotropic case the computation time is dominated by the relaxation process.

VI. APPLICATIONS

We have used the anisotropic noise generation in different contexts. Our main applications are related to visualization, but we also considered artistic image rendering applications. For visualization purposes the main step is the definition of the metric, ensuring that it incorporates the most important features of the data. In the following sections, we explain the definition for specific applications.

A. Tensor Visualization

To be able to use anisotropic noise for the visualization of tensor data, we must define a metric based on the given tensors. Some of the tensor fields we are interested in are already positive definite, e.g., diffusion tensor fields. But other tensor fields, like stress or strain fields, also have negative eigenvalues. To be able to treat such tensor fields we interpret them as distortion of a flat metric [17].

Assume that we have a positive definite tensor field T defined over a domain D . Let λ_1 and λ_2 be its eigenvalues and v_1 and v_2 be the respective eigenvectors. We define the metric for the sample generation as

$$g = \frac{1}{\sqrt{\lambda_1}} v_1 \cdot v_1^T + \frac{1}{\sqrt{\lambda_2}} v_2 \cdot v_2^T \quad (16)$$




Fig. 12. Slice of a numerical simulation of a solid block with two forces acting on the block, one pushing and one pulling force. This image shows the tensor data as ellipses. The ellipses provide an idea about the directions of contraction or expansion inside the material.

the resulting samples are ellipses aligned to v_1 and v_2 and scaled according to the eigenvalues. Depending on the application it may be necessary to normalize the eigenvalues.

Our first example is a stress field of a solid block with two applied loads with opposite sign, resulting from a three dimensional numerical finite element simulation. Figure 11 shows a slice of this data set orthogonal to the applied forces. The displayed ellipses represent the shape of a unit sphere deformed according to the local stress field. Small ellipse half-axis indicate compression, large half-axis indicate expansion in the respective direction. Ellipses with high eccentricity mean strong shear forces.

As second example we applied our method to a slice of a diffusion tensor MRI dataset of a brain. The use of glyphs, ranging from simple ellipses to more advanced glyphs as superquadrics [30], is a common method for visualizing of such data sets. The glyphs are mostly placed in grid points or are randomly spread [26]. Figure 12 shows a result using our sample generation. We used a mask image representing the confidence values of the tensors as provided by Gordon Kindlmann together with the data set. The color is used to represent the principal diffusion direction. The result is a uniform and dense representation of the data independent from the grid points. Similar results were obtained by Kindlmann et al. [16] using a particle simulating approach with repulsive and attractive forces.

B. Other Applications

The use of spot samples with varying density and size is not limited to tensor field visualization. It is also appropriate for any other glyph based visualization methods, using glyphs that can be embedded into an elliptical shape. In this section we demonstrate manifold possibilities using two exemplary areas, vector field visualization and non-photorealistic rendering.

1) *Vector field visualization*: One of the most direct vector field visualization methods is the use of arrows or other icons. The placement of glyphs without clustering or structural pattern formation is a challenging problem. Turk and Banks proposed a method to place arrows along streamlines generated using streamline optimization [27]. Sanderson et al. [28] used a reaction-diffusion model to generate spot noise based on the underlying vector field, and places glyphs at spot centers. Anisotropic noise samples can be used as an alternative to these methods. To demonstrate this technique we have used anisotropic noise samples to place arrows based on a synthetic vector field.

The direction of our metric is determined by the direction of the vector field and the orthogonal direction. We tried different ways to assign the eigenvalues to the metric. In Figure 14(a)-(d), we use the vector magnitude $|v|$ to define the ellipse size in direction of the flow, $\lambda_1 = c|v|$, and a constant value for the orthogonal direction, $\lambda_2 = d$. c and d are constants. Thus the widths of the arrows remain the same for all arrows.

$$g = \lambda_1 e_v \cdot e_v^T + \lambda_2 e_v^\perp \cdot e_v^{\perp T} \quad (17)$$

2) *Non-photorealistic rendering*: Non-photorealistic rendering is often used to simulate painting or drawing styles an




Fig. 13. These images show a slice of diffusion MRI data. The color code is the standard color map of encoding the major eigendirections in three dimensions. The projected tensors are represented by ellipses. Each ellipse is defined by the tensor value given at its center. The left image (a) shows the sampling set before relaxation, the right image (b) a close-up after three relaxation steps with densely packed ellipses. (c) shows the color coding for the direction.



Fig. 16. Mosaic-like images generated by our technique. The metric used for ellipse generation results from the gradient field of the blurred original image. The left image shows the result before relaxation and the right image after three relaxation steps.

artist would use. There are many techniques to simulate these styles. Anisotropic noise samples can be used for generating “artistic images” where elements of the image have directional properties, such as paint brush direction or rectangular mosaic tiles. Our example images are generated by constructing a gradient vector field based on the intensity values of the images. To reduce noise in the vector field, the original images are blurred by applying a Gaussian filter. We have defined a tensor metric over the image using the gradient vector field and its orthogonal vector field. The orthogonal vector field essentially points in the direction tangent to the boundary features in the images. In Figure 15 and 16, the samples are circular in areas where there is little gradient change, where the elliptic samples in the image follow the boundaries and make them more visible.

VII. CONCLUSION AND FUTURE WORK

We have introduced a method to generate anisotropic spot noise on 2D domains. We have demonstrated our method for different applications. An automatic generation of uniformly spread glyphs with locally varying size is a desirable property not only in tensor field visualization. As for all glyph based methods, the resolution of the representation is limited by the size of the glyphs that are used. The same is valid for the resulting fabric when using the spots as input for LIC.

The sampling characteristics of the resulting spot distribution is of high quality. For the uniform case, this can be seen by considering the frequency behavior of the samples. In contrast to models using repulsive forces the Voronoi cell based relaxation is very stable. Using a good start configuration of the samples only a few relaxation steps are needed to achieve good results. Thus the method is reasonable fast. Due to the lack of repulsive forces the Voronoi relaxation does not preserve the Poisson disk property. It turns out that in most cases this is not




Fig. 14. Our method used for vector field visualization. We can generate a dense sample of vector glyphs. These images show two simple datasets defined by blending simple critical-point configurations. The principal directions of the metric are defined in parallel and orthogonal to the vector field. The corresponding eigenvalues are given by the velocity. It can be seen that already a small number of relaxations improves the result. (a) and (c) Glyphs before relaxation, (b) and (d) after two relaxation steps.

a problem if only a few relaxation steps are performed. But in uniform regions unpleasing patterns with overlapping ellipses may form. This can be avoided introducing an additional intersection test, which is not entirely satisfactory because



Fig. 15. Mosaic-like image generated by our technique. The metric used for ellipse generation results from the gradient field of the blurred original image. The ellipses automatically align with edges with high gradients and thus emphasize image structure.

it also hinders the relaxation. As an extension of this work we want to look for possibilities to change the relaxation process in order to maintain the Poisson disk distribution without giving up the simplicity of the method. Further we plan on extending our work to anisotropic textures on arbitrary surfaces and volumes.

ACKNOWLEDGMENTS

The brain dataset is courtesy of Gordon Kindlmann, Scientific Computing and Imaging Institute, University of Utah, and Andrew Alexander, W. M. Keck Laboratory for Functional Brain Imaging and Behavior, University of Wisconsin-Madison. This work was supported by the National Science Foundation under contracts ACI 9624034 (CAREER Award), through the Large Scientific and Software Data Set Visualization (LSSDSV) program under contract ACI 9982251, and a large Information Technology Research (ITR) grant; the National Institutes of Health under contract P20 MH60975-06A2, funded by the National Institute of Mental Health and the National Science Foundation; and the U.S. Bureau of Reclamation. We gratefully acknowledge the support of the W.M. Keck Foundation provided to the UC Davis Center for Active Visualization in the Earth Sciences (CAVES). We thank the members of the Visualization and Computer Graphics Research Group at the Institute for Data Analysis and Visualization (IDAV) at the University of California, Davis.

REFERENCES

- [1] V. Ostromoukhov, "A simple and efficient error-diffusion algorithm," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM Press, 2001, pp. 567–572.
- [2] D. E. Knuth, "Digital halftones by dot diffusion," *ACM Transactions on Graphics*, vol. 6, no. 4, pp. 245–273, 1987.
- [3] R. Ulichney, *Digital halftoning*. Cambridge, MA, USA: MIT Press, 1987.
- [4] L. Velho and J. de Miranda Gomes, "Digital halftoning with space filling curves," in *Proceedings of the 18th annual conference on Computer graphics and interactive techniques*. ACM Press, 1991, pp. 81–90.
- [5] P. Alliez, Éric Colin de Verdière, O. Devillers, and M. Isenburg, "Isotropic surface remeshing," in *SMI '03: Proceedings of the Shape Modeling International 2003*. Washington, DC, USA: IEEE Computer Society, 2003, p. 49.
- [6] Q. Du and D. Wang, "Anisotropic centroidal voronoi tessellations and their applications," *SIAM Journal on Scientific Computing*, vol. 26, no. 3, pp. 737–761, 2005.
- [7] F. Labelle and J. R. Shewchuk, "Anisotropic voronoi diagrams and guaranteed-quality anisotropic mesh generation," in *SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry*. New York, NY, USA: ACM Press, 2003, pp. 191–200.
- [8] K. Shimada, A. Yamada, and T. Itoh, "Anisotropic triangulation of parametric surfaces via close packing of ellipsoids," *Int. J. Comput. Geometry Appl.*, vol. 10, no. 4, pp. 417–440, 2000.
- [9] V. Ostromoukhov, C. Donohue, and P.-M. Jodoin, "Fast hierarchical importance sampling with blue noise properties," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 488–495, 2004.
- [10] D. P. Mitchell, "Spectrally optimal sampling for distribution ray tracing," in *Proceedings of the 18th annual conference on Computer graphics and interactive techniques*. ACM Press, 1991, pp. 157–164.
- [11] B. Cabral and L. Leedom, "Imaging vector fields using line integral convolution," in *Proceedings of SIGGRAPH 1993*, J. T. Kajiya, Ed., vol. 27, August 1993 1993, pp. 263–272.
- [12] M.-H. Kiu and D. C. Banks, "Multi-frequency noise for lic," in *VIS '96: Proceedings of the 7th conference on Visualization '96*. Los Alamitos, CA, USA: IEEE Computer Society Press, 1996, pp. 121–126.
- [13] A. J. S. Hin and F. H. Post, "Visualization of turbulent flow with particles," in *VIS '93: Proceedings of the 4th conference on Visualization '93*, 1993, pp. 46–53.

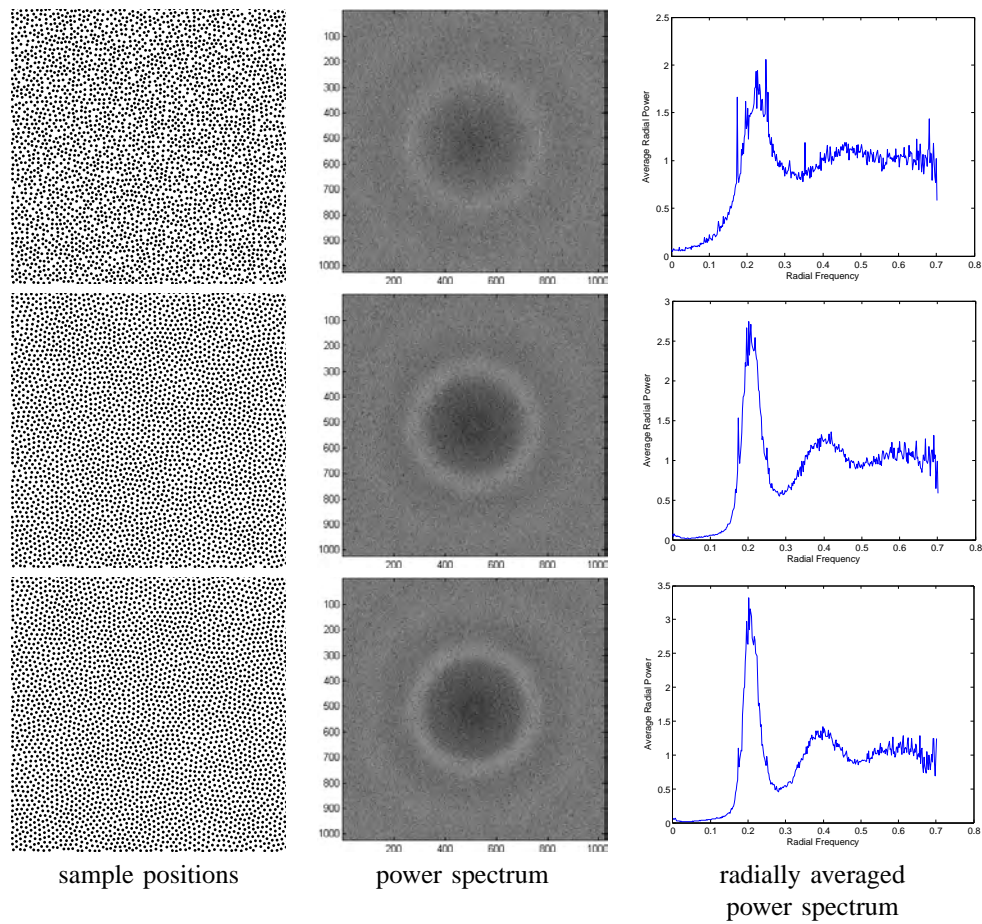


Fig. 17. The frequency analysis of isotropic, uniform samples generated with our method shows the typical blue-noise characteristics - no low frequencies (dark spot around the origin) and a peak (light circle) for the principal frequency, representing the average distance between nearest neighbors. These characteristics become more expressed after the first iterations. After the third iterations the changes are hardly visible anymore. First row: initial point set, second row: after five iterations, third row: after ten iterations.

- [14] R. M. Kirby, H. Marmanis, and D. H. Laidlaw, "Visualizing multivalued data from 2D incompressible flows using concepts from painting," in *IEEE Visualization '99*, D. Ebert, M. Gross, and B. Hamann, Eds., San Francisco, 1999, pp. 333–340. [Online]. Available: citeseer.ist.psu.edu/kirby99visualizing.html
- [15] G. Kindlmann, D. Weinstein, and D. Hart, "Strategies for direct volume rendering of diffusion tensor fields," *IEEE Transactions on Visualization and Computer Graphics*, vol. 6, no. 2, pp. 124–138, /2000. [Online]. Available: citeseer.ist.psu.edu/kindlmann00strategies.html
- [16] G. Kindlmann and C.-F. Westin, "Diffusion tensor visualization with glyph packing," *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization / Information Visualization 2006)*, vol. 12, no. 5, p. (to appear), September-October 2006.
- [17] I. Hotz, L. Feng, H. Hagen, B. Hamann, B. Jeremic, and K. I. Joy, "Physically based methods for tensor field visualization," in *Proceedings of IEEE Visualization 2004*. IEEE Computer Society Press, oct 2004, pp. 123–130.
- [18] A. S. Glassner, *Principles of Digital Image Synthesis: Volume Two*, ser. The Morgan Kaufmann Series in Computer Graphics and Geometric Modeling. Morgan Kaufmann, 1995, gLA a 95:2 1.Ex.
- [19] R. L. Cook, "Stochastic sampling in computer graphics," *ACM Trans. Graph.*, vol. 5, no. 1, pp. 51–72, 1986.
- [20] S. P. Lloyd, "Least square quantization in pcm," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, Mar. 1982.
- [21] Q. Du, V. Faber, and M. Gunzburger, "Centroidal voronoi tessellations: Applications and algorithms," *SIAM Review*, vol. 41, no. 4, pp. 637–676, 1999.
- [22] S. Hiller, H. Hellwig, and O. Deussen, "Beyond stippling - methods for distributing objects on the plane." *Comput. Graph. Forum*, vol. 22, no. 3, pp. 515–522, 2003.
- [23] A. Hausner, "Simulating decorative mosaics," in *SIGGRAPH 2001, Computer Graphics Proceedings*, E. Fiume, Ed., 2001, pp. 573–578. [Online]. Available: <http://citeseer.ist.psu.edu/hausner01simulating.html>
- [24] L.-P. Fritzsche, H. Hellwig, S. Hiller, and O. Deussen, "Interactive design of authentic looking mosaics using voronoi structures," in *Proceedings of 2nd International Symposium on Voronoi Diagrams in Science and Engineering*, Seoul, Korea, 2005. [Online]. Available: <http://citeseer.ist.psu.edu/744345.html>
- [25] K. E. Hoff, J. Keyser, M. Lin, D. Manocha, and T. Culver, "Fast computation of generalized Voronoi diagrams using graphics hardware," *Computer Graphics*, vol. 33, no. Annual Conference Series, pp. 277–286, 1999. [Online]. Available: citeseer.ist.psu.edu/hoff99fast.html
- [26] D. H. Laidlaw, E. T. Ahrens, D. Kremers, M. J. Avalos, R. E. Jacobs, and C. Readhead, "Visualizing diffusion tensor images of the mouse spinal cord," in *IEEE Visualization '98 (VIS '98)*, D. Ebert, H. Hagen, and H. Rushmeier, Eds., IEEE. Washington - Brussels - Tokyo: IEEE Computer Society Press, Oct. 1998, pp. 127–134. [Online]. Available: <http://visinfo.zib.de/EVlib/Show?EVL-1998-331>
- [27] G. Turk and D. Banks, "Image-guided streamline placement," *Computer Graphics*, vol. 30, no. Annual Conference Series, pp. 453–460, 1996. [Online]. Available: citeseer.csail.mit.edu/turk96imageguided.html
- [28] A. R. Sanderson, C. R. Johnson, and R. M. Kirby, "Display of vector fields using a reaction-diffusion model," in *VIS '04: Proceedings of the conference on Visualization '04*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 115–122.
- [29] E. N. Gilbert, "Gray codes and paths on the n-cube," *Bell System Technical Journal*, vol. 37, no. 3, pp. 815–826, 1958.
- [30] G. Kindlmann, "Superquadric tensor glyphs," in *Proceeding of The Joint Eurographics - IEEE TCVG Symposium on Visualization*, May 2004, pp. 147–154.




Fig. 18. The frequency analysis of an anisotropic, uniform sample set for the initial point set (column (a)) clearly shows an anisotropic characteristic. As we apply relaxation to the samples, the anisotropy property start to diminish. When the samples converge to a stable position, the power spectrum still shows fairly strong anisotropy. Similar to the blue-noise samples there are almost no low frequencies but here the principal frequency is direction-dependent. This is a result of the generalized Poisson disk property.

Column (a): initial point set, column (b): after 25 iterations, column (c): after 50 iterations.




Fig. 19. When an intersection test is enforced during relaxation, the samples converge to stable locations faster because their movements are more restricted. The power spectrum shows a slightly improved anisotropic behavior for non-overlap samples but with irregularities due to the restricted relaxation process.

(a) Initial sample set, (b) after 25 iterations, (c) after 35 iterations.



Louis Feng Biography text here.

Bernd Hamann Biography text here.

Ken Joy Biography text here.

Ingrid Hotz Biography text here.