
Dense Glyph Sampling for Visualization

Louis Feng¹, Ingrid Hotz², Bernd Hamann¹, and Kenneth Joy¹

¹ Institute for Data Analysis and Visualization (IDAV), Department of Computer Science, University of California, Davis CA 95616, USA

² Zuse Institute Berlin, Division for Visualization and Data Analysis, Berlin Germany

Summary. We present a simple and efficient approach to generate a dense set of anisotropic, spatially varying glyphs over a two-dimensional domain. Such glyph samples are useful for many visualization and graphics applications. The glyphs are embedded in a set of non-overlapping ellipses whose size and density match a given anisotropic metric. An additional parameter controls the arrangement of the ellipses on lines, which can be favorable for some applications, e.g., vector fields, and distracting for others. To generate samples with the desired properties we combine ideas from sampling theory and mesh generation. We start with constructing a first 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 samples more evenly.

Key words: tensor field visualization, glyph packing, anisotropic Voronoi diagram

1 Introduction

Anisotropic spot samples with certain characteristics, such as spatially varying density and size, have many applications in visualization and computer graphics ranging from glyph rendering and texture generation for visualization purposes [14, 8, 13, 11, 12, 3], digital halftoning [19, 15, 25] to mesh generation [2, 16, 22]. While some of the desirable properties of the samples are similar across applications, the goals and appropriate sampling strategies are problem-dependent. When using the samples as input for texture generation as, e.g., line integral convolution (LIC), it is important to avoid structural patterns. An ellipse alignment leads to distracting artifacts in the LIC texture. But an alignment of glyphs is desirable for other applications, e.g., vector field visualization, where it supports the impression of flow.

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. This first sample set already exhibits

most of the desired properties. Next, we use a generalized anisotropic Lloyd relaxation to distribute the spot samples more evenly. Our anisotropic Lloyd relaxation is a straight-forward generalization of the isotropic version and is based on the work of Labelle and Shewchuk [16] and Du et al. [2]. A parameter in the Voronoi cell definition controls the alignment of the ellipses. We have applied our method to several test data sets and various vector and tensor fields.

2 Related Work

The generation of point or spot distributions with certain properties is the subject of research in different fields. Dependent 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 use a form of stochastic sampling, where random points are added or rejected according to certain criterions. Such methods often suffer from low convergence rates. Other approaches use relaxation techniques, in particular Lloyd's relaxation [18, 1] and its variants resulting in high-quality blue-noise samples. To improve efficiency of the sampling algorithm several approaches have been suggested using tile sets, which then are repeatedly tiled across the plane. Using this strategy, e.g., Ostromoukhov et al. introduced a very efficient isotropic blue-noise sampling method based on Penrose tiling [20]. Most of these methods assume that the samples are isotropic. For a survey on sampling techniques, we refer the reader to [5, 24]. 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 [7, 6, 4]. Different from our definition, here the orientation of the distributed objects is not 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.

The goal of generating an anisotropic distribution following a given metric also appears in the area of mesh generation. Shimada et al. [22] introduced a mesh generation approach using a close packing of ellipsoidal bubbles. The packing is performed using a particle system, where particles move according to repulsive and attractive forces. The equations of motion are solved numerically to yield a force-balancing configuration. A geometric approach for anisotropic mesh generation was chosen by Du et al. [2] and Labelle et al. [16]. 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.




Fig. 1. Generalized Poisson disk property. The minimum distance of two sample points is defined by the local ellipses, which are not allowed to overlap.

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 [17, 13] 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 [23]. Kindlmann introduced reaction-diffusion into the visualization community applying it to diffusion tensor MRI data [11]. Sanderson et al. used a reaction-diffusion model to generate spot noise based on the underlying vector field placing glyphs at the spot center [21]. Reaction diffusion 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 [12]. Their work is built on a particle approach simulating attractive and repulsive forces. This work is an extension to our recent work for the generation of anisotropic noise samples [3] by adding a control over the alignment of the ellipses.

3 Assumptions and Goals

The starting point for the generation of the elliptic noise samples is a metric $g = (g_{ij})$ given over a domain $D \subset \mathbb{R}^2$, which defines the sample properties. The metric can be user-defined or derived from scalar fields, vector fields, or tensor fields, see Section 6. The metric is given as a two-by-two symmetric, positive definite matrix depending on the location $P = (x, y) \in \mathbb{R}^2$. We assume that the metric is non-degenerate everywhere. In general, it is spatially varying and anisotropic. Size and density of the ellipses are specified by the metric in their center $P_0 = (x_0, y_0)$. Their shape is defined as unit circle with respect to the metric g_0 in P_0 , i.e.,

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

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

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

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:

- 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 1.
- The spots are closely packed without holes resulting in an uniform density, defined as covered pixels per unit area.
- The spots are non-intersecting having a minimum distance, defined by a generalized Poisson disk property, see Figure 1.
- The degree of alignment of the spots can be controlled by a parameter.

4 Algorithm

Texture generation can be divided into two independent steps:

1. Computation of a reasonable starting distribution of ellipses, where we generate a set of spot candidate 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.
2. Optimization of the starting distribution using an anisotropic Lloyd relaxation. Dependent on a parameter controlling the anisotropy the relaxation more or less favors an alignment.

Both steps are important. The first step determines the number of samples, the density and the Poisson disk property. The second step leads to a more uniform sample distribution, approaching a stable configuration. In the following we shortly explain the single steps. For more details we refer to [3].

4.1 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.

The initial set must have higher density than the target density. target density leads to good results. 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 1. At this stage, the generated candidate spots generally overlap. Once the initial set of points is generated, the algorithm traverses the set of points. A candidate spot is accepted when its ellipse does not overlap with the ellipses at any other already selected samples. 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.

4.2 Anisotropic Voronoi Relaxation

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 against a centroidal Voronoi diagram, where each sample point lies in its cell centroid. This diagram minimizes the energy given as

$$E = \sum_{i \in I} \int_{Vor_i} \rho(x) \|r - r_i\|^2 dr \quad (3)$$

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. For the definition of the anisotropic Voronoi diagram and the centroid computation we built on the works of Labelle and Shewchuk [16] and Du et al. [2]. Our method is a combination of these two methods, satisfying our demands. Depending on the special choice of the metric used to define the Voronoi cells the alignment of ellipses is more or less supported.

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 $Vor(P_i)$ of a point P_i as the set of all points $P \in D$ that are at least as close to P_i as to any other point $P_j, j \neq i$, using the geodesic instead of the Euclidean distance. However, since the computation of this shortest path is difficult and computationally expensive we use an approximate distance function for two points proposed by Labelle and Shewchuk, because it matches our conditions well, i.e.,

$$d^2(P, Q) = (P - Q)^T g(P)(P - Q). \quad (4)$$

The distance measure is not symmetric $d(P, Q) \neq d(Q, P)$. Also, the triangle inequality is not necessarily satisfied. Based on this approximate distance, a Voronoi cell of point P_i is defined as

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

When using the metric defining the ellipses, this distance function guarantees that ellipses of our start configuration lie entirely inside the Voronoi cells. The resulting Voronoi cells are in general not convex and may not even be connected. Therefore, we define a localized version of the Voronoi cells considering only the part containing the sample point. For more details we refer to [3]. We define

$$\begin{aligned} Vor_r(P_i) &= \{P \in D | i \in I_P \text{ and } d(P_i, P) \leq d(P_j, P) \text{ for all } j \in I_P \text{ with } i \neq j, \\ &\text{with } I_P = \{i \in I | (P_i - P_j) \cdot (P - P_i) \leq 0, \forall j \neq i\}. \end{aligned} \quad (6)$$

Centroid definition

For the definition of the centroid we follow the idea of Du et al. [2], 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

$$c_i = \frac{\int_{Vor(P_i)} d(r) r \, dr}{\int_{Vor(P_i)} d(r) \, dr}, \quad (7)$$

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

$$c_i = \left(\int_{Vor(P_i)} g(r) \, dr \right)^{-1} \cdot \left(\int_{Vor(P_i)} g(r) \cdot r \, dr \right). \quad (8)$$

As an integral over positive definite matrices, the left matrix is always invertible. When using an isotropic metric this definition reduces to the standard weighted centroid definition. If the metric is uniform, i.e., it does not depend on r , the anisotropic centroid definition coincides with isotropic uniform case.

4.3 Implementation

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 quadratic polynomial, which is computationally expensive and numerically unstable. We use polylines to approximate the ellipses during intersection test to reduce complexity. This approximation produces good results without the issues involved in the algebraic method.

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 of Equation 8 results in

$$\underbrace{c_i - P_i}_{\equiv T_i} = \left(\underbrace{\sum_{R \in \text{Vor}(P_i)} g(R)}_{\equiv M_i} \right)^{-1} \cdot \underbrace{\sum_{R \in \text{Vor}(P_i)} g(R) \cdot (R - P_i)}_{\equiv t_i} \quad (9)$$

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 to specify i , by comparing the distances to sample points lying inside a local bin.
- Update the matrix M_i and the vector t_i in the following way:

$$\begin{aligned} M_i &\rightarrow M_i + g(R) \\ t_i &\rightarrow t_i + g(R) \cdot (R - P_i) \end{aligned} \quad (10)$$

After traversing the entire domain, the new position of the sample points P_i , given by Equation 9, is determined by the translation vector T_i , i.e.,

$$T_i = M_i^{-1} \cdot t_i \text{ and } P_i \rightarrow P'_i = P_i + T_i. \quad (11)$$

5 Structural Behavior

The distance approximation makes general statements about the convergence behavior of the point set difficult. For the quality of the results the effect of a couple of relaxation steps is more important than convergence. We can identify fix-points of the relaxation process for the uniform case, as e.g., hexagonal structures or any other point symmetric configurations. In our examples we observe that after several iterations regions with hexagonal patterns are forming, see Figure 2. Inside these regions the patterns become soon relatively stable. Between these regions the structure still changes slightly even after many iterations.

Dependent on the orientation of the ellipses in relation to the orientation of the hexagonal structure, these configurations result in an aligned structure or not. A similar behavior can be observed for slowly varying fields. Whether an alignment is desirable depends on the specific application. While it generates




Fig. 2. After many iterations the sample positions converge and lead to stable patterns. This is especially visible for uniform data sets. After 50 iterations (b) we can already see the formation of basically two different structures (highlighted by the red box). These regions hardly change over the next 75 iterations (c). Looking at difference images between the single iterations we can see that there is still a fluctuation in the regions between the stable patterns. (a) shows the start configuration before relaxation.

artifacts when the samples are used as input for texture generation or in non-photorealistic rendering applications, it enhances perception of flow field data sets.

This behavior can be controlled by the anisotropy of the metric used for the Voronoi cells in the relaxation. Figure 3 shows Voronoi cells for different types of anisotropies: isotropic, given by the ellipse and with an exaggerated anisotropy. The shape of the Voronoi cell determines the movement of the sample point in the next iteration. Especially if the start configuration is not very dense, the anisotropy given by the ellipses is not sufficient to prevent the samples from aligning.

In our implementation we have adjusted the anisotropy by multiplying the larger eigenvalue λ_2 with a positive parameter p . Since only the ratio of the eigenvalues is important it is enough to manipulate one eigenvalue. A value of λ_2/λ_1 results in an isotropic metric. A value larger than one leads to a higher




Fig. 3. Example of the Voronoi diagram of six points using three different uniform metrics. (a) isotropic metric, (b) metric given by the ellipses, (c) metric with higher anisotropy.




Fig. 4. Relaxation of uniform anisotropic samples. (a) shows the samples after ten relaxations with intersection test. The other images show results after ten iterations without intersection test using three different metrics for the relaxation: (b) original metric, (c) scaling of the of the larger eigenvalue with two, (d) scaling of the larger eigenvalues with three.

degree of anisotropy. Since alignment often goes hand in hand with overlapping of the ellipses, we also implemented a relaxation with intersection test. In case of an intersection with neighboring ellipses the translation vector is shortened. Whereas this prevents the ellipses from overlapping and reduces the alignment it results in a less uniform distribution, see Figure 5(a). It is important to note that only the anisotropy of the relaxation process is influenced by the alignment parameter. Size and shape of the represented glyph are not changed and thus still represent the local field properties. It also does not change the general convergence properties but the characteristics of the resulting texture see Figure 5(b-d).




Fig. 5. Relaxation of nonuniform anisotropic samples. The definition of the Voronoi cells uses the original metric. Top row shows the sample set. (a) first sample set, (b) after one iteration, (c) after ten iterations. The bottom row shows the images after applying a Gaussian blur.

6 Results

The evaluation of our algorithm is guided by the goals described in Section 3. We first discuss examples for a simple isotropic and anisotropic metric definition, which already exhibit most characteristic behaviors of our method. Then we show some results for applications in different contexts. Our main applications are related to visualization, but we also considered "artistic" image rendering applications. The use of glyph sampling with varying density and size is appropriate for any glyph-based visualization, using glyphs that can be embedded into an elliptical shape. For visualization purposes the main step is the definition of the metric, ensuring that it incorporates the most important features of the data. A further analysis of the results in frequency space can be found in [3].

Representation of the metric by sampling shape and density

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 scalar density d , which is defined as covered pixels per unit area, can be measured by using a Gaussian filter. The local density is then given as gray value. Due to the discrete structure of the samples we cannot expect a constant, but almost uniform density. An example for an anisotropic data set is shown in Figure 6. The size of the Gaussian filter used for these examples is the same for both

examples. It can be observed that the density is fast approaching a uniform distribution. After six relaxation steps there are no holes visible anymore. In particular there is no dependence 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.

Control of alignment

To evaluate the influence of the shape of the Voronoi cells on the relaxation we started with a uniform anisotropic data set. We used the same data set as for Figure 2, where the initial sample set can be seen in (a). Figure 5 shows results using different relaxation methods after ten iterations. For the generation of Figure 5(a) and (b) the Voronoi cells are computed using the original metric as given by the ellipses. In (a) we performed an intersection test after each iteration. This enforces the maintenance of the Poisson disk property but also hinders the relaxation process. There are holes in the dataset even after ten iterations. (b) shows the result without intersection test. We can observe hexagonal structures with and without alignment of the ellipses. There almost no holes left, but there are a few regions where the ellipses start overlapping. For Figure 5(c) and (d) we used an exaggerated anisotropy for the Voronoi cell computation. Especially (d) shows a very uniform structure with almost no overlapping and alignment along the major eigenvector.

Vector field visualization

One of the most direct vector field visualization methods is the use of arrows or other icons. We applied our glyph sampling method to provide a dense placement of glyphs without clustering based on a synthetic vector field. The major eigendirection of the metric is determined by the direction of the vector field. The major eigenvalue is specified by vector magnitude $\lambda_1 = |v|$ the minor eigenvalue λ_2 is defined as a constant. The metric is given as

$$g = \lambda_1 e_v \cdot e_v^T + \lambda_2 e_v^\perp \cdot e_v^{\perp T}. \quad (12)$$

Figure 7 shows the results for two different degrees of anisotropy in the relaxation after five iterations. The left image uses the original metric g , in the right image the larger eigenvalue is scaled by a factor of three. Both images show a uniform sampling of the ellipses, but the perception of flow is much better in (a). To achieve similar results Turk and Banks proposed a method to place arrows along streamlines generated using streamline optimization [23]. Sanderson et al. [21] used a reaction-diffusion model to generate spot noise based on the underlying vector field, and places glyphs at spot centers.

Tensor field 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 [9]. 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 the respective eigenvectors. We define the metric for the sample generation as

$$g = \frac{1}{\sqrt{\lambda_1}} v_1 \cdot v_1^T + \frac{1}{\sqrt{\lambda_2}} v_2 \cdot v_2^T \quad (13)$$

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.




Fig. 6. Effect of manipulating the anisotropy value during the relaxation process: First row shows the 10th relaxation step of a uniform data set starting from the same sample configuration. Second row shows our method for vector field visualization. The left images (a,c) use the original metric for relaxation, for (b,d) the larger eigenvalue has been multiplied by three. It can be seen that the original metric favors an alignment along the field lines whereas the exaggerated anisotropy favors an alignment orthogonal to the field lines.

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 8(a) 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-axes indicate compression, large half-axis indicate expansion in the respective direction. Ellipses with high eccentricity mean strong shear forces.

We also 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 [10], is commonly done for visualizing such data sets. The glyphs are mostly placed in grid points or are randomly spread [17]. Figure 8(b) 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. [12] using a particle simulating approach with repulsive and attractive forces.

Non-photorealistic Rendering

Non-photorealistic rendering is often used to simulate painting or drawing styles an artist would use. There are many techniques to simulate these styles. Anisotropic noise samples can be used for generating “artistic images” where




Fig. 7. (a) Slice of a numerical simulation of a solid block with two forces acting on the block, one pushing and one pulling force. The image shows the tensor data as ellipses. The ellipses given an idea of directions of contraction or deletion inside the material. (b) shows a close-up view obtained after three relaxation steps of a diffusion MRI slice. The color code is the standard color map of encoding the major eigendirection. The projected tensors are represented by ellipses. Each ellipse is defined by the tensor value given at its center.




Fig. 8. Mosaic-like images generated by our technique. The metric used for ellipse generation results from the gradient field of the blurred original image. The top image shows the result before relaxation and the bottom image after three relaxation steps.

elements of the image have directional properties, such as paint brush direction or rectangular mosaic tiles. Our example images were 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 were blurred by applying a Gaussian filter. We 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. One example can be seen in Figure 9, the left image shows the ellipses before relaxation the right image after three relaxations using the original metric.

7 Conclusion

We have introduced a method to generate a dense set of uniformly spread glyphs. Besides the local control of size and density the methods provides a parameter to control the alignment of the glyphs. This 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 method as described is applicable for two-dimensional fields. A generalization to three-dimensional is principally possible but yields as well time as perceptually issues.

Our method is a purely geometric process. The centroid can be computed explicitly without involving numerics. 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 a uniform distribution. Thus the method is reasonable fast.

Due to the lack of repulsive forces the Voronoi relaxation does not necessarily preserve the Poisson disk property. We have shown that we can reduce the violation of the Poisson disk property by manipulating the shape of the Voronoi cell appropriately. The key entity thereby is the anisotropy of the metric used for the Voronoi cell definition. In contrast to expensive intersection tests, this approach does not hinder the relaxation process and does not introduce any additional computational costs.

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 partially supported by the German Research Foundation DFG (Emmy-Noether Research group) and 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. Further we thank the members of the Visualization and Computer Graphics Research Groups at the Zuse Institute Berlin and the Institute for Data Analysis and Visualization (IDAV) at the University of California, Davis.

References

1. Q. Du, V. Faber, and M. Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999.
2. Q. Du and D. Wang. Anisotropic centroidal voronoi tessellations and their applications. *SIAM Journal on Scientific Computing*, 26(3):737–761, 2005.
3. L. Feng, I. Hotz, B. Hamann, and K. Joy. Anisotropic noise samples. *IEEE Transactions on Visualization and Computer Graphics*, 14(2):342–354, 2008.
4. 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.
5. A. S. Glassner. *Principles of Digital Image Synthesis: Volume Two*. The Morgan Kaufmann Series in Computer Graphics and Geometric Modeling. Morgan Kaufmann, 1995. GLA a 95:2 1.Ex.
6. A. Hausner. Simulating decorative mosaics. In E. Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 573–578, 2001.
7. S. Hiller, H. Hellwig, and O. Deussen. Beyond stippling - methods for distributing objects on the plane. *Comput. Graph. Forum*, 22(3):515–522, 2003.

8. 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*, pages 46–53, 1993.
9. I. Hotz, L. Feng, H. Hagen, B. Hamann, B. Jeremic, and K. I. Joy. Physically based methods for tensor field visualization. In *VIS '04: Proceedings of IEEE Visualization 2004*, pages 123–130. IEEE Computer Society Press, oct 2004.
10. G. Kindlmann. Superquadric tensor glyphs. In *Proceeding of The Joint Eurographics - IEEE TCVG Symposium on Visualization*, pages 147–154, May 2004.
11. G. Kindlmann, D. Weinstein, and D. Hart. Strategies for direct volume rendering of diffusion tensor fields. *IEEE Transactions on Visualization and Computer Graphics*, 6(2):124–138, /2000.
12. G. Kindlmann and C.-F. Westin. Diffusion tensor visualization with glyph packing. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization / Information Visualization 2006)*, 12(5):1329–1336, September–October 2006.
13. R. M. Kirby, H. Marmanis, and D. H. Laidlaw. Visualizing multivalued data from 2D incompressible flows using concepts from painting. In D. Ebert, M. Gross, and B. Hamann, editors, *IEEE Visualization '99*, pages 333–340, San Francisco, 1999.
14. M.-H. Kiu and D. C. Banks. Multi-frequency noise for lic. In *VIS '96: Proceedings of the 7th conference on Visualization '96*, pages 121–126, Los Alamitos, CA, USA, 1996. IEEE Computer Society Press.
15. D. E. Knuth. Digital halftones by dot diffusion. *ACM Transactions on Graphics*, 6(4):245–273, 1987.
16. 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*, pages 191–200, New York, NY, USA, 2003. ACM Press.
17. 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 D. Ebert, H. Hagen, and H. Rushmeier, editors, *IEEE Visualization '98 (VIS '98)*, pages 127–134, Washington - Brussels - Tokyo, Oct. 1998. IEEE, IEEE Computer Society Press.
18. S. P. Lloyd. Least square quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, Mar. 1982.
19. V. Ostromoukhov. A simple and efficient error-diffusion algorithm. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 567–572. ACM Press, 2001.
20. V. Ostromoukhov, C. Donohue, and P.-M. Jodoin. Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph.*, 23(3):488–495, 2004.
21. 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*, pages 115–122, Washington, DC, USA, 2004. IEEE Computer Society.
22. K. Shimada, A. Yamada, and T. Itoh. Anisotropic triangulation of parametric surfaces via close packing of ellipsoids. *Int. J. Comput. Geometry Appl.*, 10(4):417–440, 2000.
23. G. Turk and D. Banks. Image-guided streamline placement. *Computer Graphics*, 30(Annual Conference Series):453–460, 1996.

24. R. Ulichney. *Digital halftoning*. MIT Press, Cambridge, MA, USA, 1987.
25. 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*, pages 81–90. ACM Press, 1991.