Voronoi-based Variational Reconstruction of Unoriented Point Sets

1. Briefly summarize the paper's contributions. Does it address a new problem? Does it present a new approach? Does it show new types of results?

- **[AS]**
  This paper presents a new approach to estimating unoriented normals from an unoriented set of input points, and using this tensor field to find an implicit function representing the surface defined by the input point set.
- **[DS]**
  The paper presents an algorithm for reconstructing (debatably) watertight surfaces from unoriented point sets. They use the Voronoi diagram of the point set to compute a tensor field that represents the most likely direction of the normal to the surface and the confidence of the estimate. An implicit function is found by solving a generalized eigenvalue problem so that the gradient is aligned with the normal direction estimates. The approach is resilient to noise, and provides parameters to adjust data fitting and smoothness of the reconstructed surface.
- **[FP]**
  Implicit surface methods (e.g., Othake 2003, Kazhdan 2006, Shen 2004) require of accurate surface normal at sample position. When this normal is not provided, traditional approaches to compute it are based in PCA (e.g. Hoppe 1992), or Voronoi Poles. Global orientation of normal is then attained using MST on a normal propagation graph. This paper looks for a simultaneous solution of the implicit surface and the normal orientation problems, by minimizing a global energy on a tensor field. This tensor field is given by the covariance matrix on expanded Voronoi regions. The optimization step compute an implicit function whose gradient align with the principal directions of the tensor field and is constrained both in smoothness of the gradient field and data fitting (i.e., sample positions at zero level).
- **[JD]**
- **[LF]**
  The paper presents a new technique for approximating unoriented normals that combines qualities of the PCA and Voronoi pole approaches for doing so and outperforms those two approaches.
- **[MK]**
  The paper proposes a new variational approach to surface reconstruction that does not require knowing, or even estimating, surface normals.

2. What is the key insight of the paper? (in 1-2 sentences)

- **[AS]**
  The key insight of this paper is that by combining the idea behind PCA and Voronoi poles, and estimating unoriented normals from covariance matrices of a union of Voronoi cells, we can obtain better estimates of the normal direction. Isosurfacing the implicit function whose gradient is best aligned with the principal axes of this covariant tensor field will provide the best fitting reconstruction of the surface defined by the input points set.
- **[DS]**
  The key insight is the use of the tensor field's principal axes and eccentricities as estimates of the normal direction of the surface and the confidence of this estimate. They also provide a novel Voronoi-PCA method to estimate normals when they are not available.

- **[FP]**
  Implicit surface description and global normal orientation can be simultaneously solved by fitting a smooth gradient field that minimizes energy on a tensor field which measures the confidence on the normal direction.
- **[JD]**
- **[LF]**
  According to the authors: The covariance matrix of a Voronoi cell can provide both an estimate of the unoriented normal direction (eigenvector of largest eigenvalue) and a measure of confidence in that estimate (anistropy).
- **[MK]**
  The paper makes two key insights. First, that the local covariance at a point can be estimated by integrating the covariance over the (possibly locally agglomerated) Voronoi regions. (The authors show that this is more stable than either using PCA or using the Vornoi poles.) Second, the authors show by using a constrained-variational/Lagrange-multiplier formulation that it is possible to fit a function whose gradient aligns with principal direction of variance (i.e. normal) without requiring knowledge of the sign.

3. What are the limitations of the method? Assumptions on the input? Lack of robustness? Demonstrated on practical data?

- **[AS]**
  One limitation of this method is that the Cholesky Factorization step is a bottleneck in both the speed and the memory requirements, making this method slow and hard to scale.
- **[DS]**
  The Choleski factorization affects the scalability of the system. Out-of-core factorization can be used but this results in increased computation times.
- **[FP]**
  In order to discretely define the implicit function and its gradient, it is required to set a space tetrahedral mesh. Based on the connectivity of this tetrahedral mesh, the authors set the operators (the Drichlet energy and the Lapalcian constraint) of the optimization problem. Therefore, the accuracy of the method seems to highly relie on the quality (regularity, density) of the tetrahedral mesh.
  The optimization problem is reduced to a Generalized Eignevalue Problem. The solution to this problem seems to be hard both in terms of conditioning and performance. To improve conditioning the authors add two regularizing terms that also contribute to data fitting and smoothness of the implicit function. On the other hand, the method has a bottleneck in terms of performance due to memory requirements on the Cholesky factorization used to solve the GEP.
- **[JD]**
- **[LF]**
  The normal estimation method seems to be more robust than its predecessors and the later reconstruction results on more complex point sets look great.
- **[MK]**
  If interpolation constraints are not introduced, a serious limitation of this method is there is no reason that the obtained solution should have similar values at the input point samples (so using the average value of the reconstructed function over the input samples becomes meaningless). As an example, consider the 2D case of points sampled from an annulus (circle within a circle). The smoothest function that aligns with covariance will likely be something like a distance function, with all gradients points to away from the origin, but then the values on the two circles will be quite different.

The incorporation of value constraints addresses this problem, but it would seem that this results in a “competition” between the energies, and it is not clear that this would be easy to balance. (e.g. How does one set the weights and how should the weights adapt to local distribution of information?)
More generally, this points at the fact that “just making the gradients do the right thing locally” does not guarantee global consistency.

4. Are there any guarantees on the output? (Is it manifold? does it have boundaries?)

- **[AS]**
  The output is guaranteed to be a smooth manifold without boundary.
- **[DS]**
  An implicit function from which a manifold surface can be extracted. Watertightness is debatable.
- **[FP]**
  Due to the implicit approach, the final reconstruction is a watertight surface. Form the set of experiments presented in the paper, the authors claim the superiority of their approach to previous normal estimation techniques. Also claim that results attain sharper details than Poisson reconstruction.
- **[JD]**
- **[LF]**
  The authors say the result is water tight, but Misha mentioned in discussion that appears instead to be manifold without boundary.
- **[MK]**
  The output is a manifold without boundary, assuming non-vanishing gradients along the zero set (i.e. in the context of discrete iso-surface, as long as vertices don’t have zero value.)

5. What is the space/time complexity of the approach?

- **[AS]**
- The time complexity is driven by the time required for the Cholesky Factorization, which can be performed in $O(N^2)$ time.
- The data structures they use for Delaunay triangulation scale linearly in memory. However, the bottleneck in memory is caused by Cholesky Factorization, which in the worst case results in $O(N^2)$ non-zero elements.
- **[DS]**
  The space complexity is bottlenecked by the Choleski factorization.
  The same is probably true for the time complexity depending on the resulting factorization. Otherwise, the Arnoldi iteration procedure might be the bottleneck.
- **[FP]**
  The initial stage requires computing a Voronoi Diagram/Delanauy triangulation from the sample points and a subsequent refining. Since the refinement is required to be tight just close the sample points, the space and time required for setting the tetrahedral mesh is small.
  Instead, solving the GEP using Cholesky factorization is an expensive task. According to the authors an in-core factorization has high memory requirements, while an out-core factorization implies some increments in overall timing.
- **[JD]**
- **[LF]**
- **[MK]**
  This is debatable. The method bottle-necks on the computation of a Cholesky

factorization of a sparse NxN matrix. Since there is no guarantee that the factorization is sparse, I believe that the method can be as bad as O(N^1.5)-O(N^2) (I believe “1.5” is for planar and “2” for 3D) though in practice the complexity seems to be closer to linear.

6. How could the approach be generalized?

- **[AS]**
  This approach can be generalized to accommodate varied amounts of data fitting and separation between connected components, as well as to allow for controlled smoothness.
- **[DS]**
  Improve scalability, maybe by finding a better, more efficient way to do the Choleski factorization.
  A hybrid between their method and the Poisson reconstruction is proposed for large data.
- **[FP]**
  The method contains some regularization parameters that provide some control on the smoothness of the reconstruction and how close it fit the sample points. Selecting these parameters adaptively (instead of globally as done in the paper) should give reconstruction with distinct local attributes.
  As the authors claim, the idea of using a tensor field to define a vector field has many applications. They point as an interesting application using this idea to curvature aligned quad remeshing.
  Since a tensor field provides a measure on the confidence of the orientation of an implicit vector field, this kind of techniques should provide more robust results than gradient domain methods. It would be interesting to see if tensor fields can be used to improve gradient domain image processing techniques such as tone mapping.
  The authors also propose using the normal orientation provided by their method for the initialization of other implicit and faster methods such as Poisson reconstruction.
- **[JD]**
- **[LF]**
  Though the method does not appear to require any user defined parameters, it mentions that smoothness can be controlled by (1) changing a fitting factor which adjusts the amount of data fitting and separation between connected components and (2) constraining B by splines-over-tension energy to achieve a better tradeoff between smoothness and fitting of normal directions. Matrices A and B can be modified in other ways for application-specific purposes.
- **[MK]**
  I don’t see an obvious extension.

7. If you could ask the authors a question (e.g. “can you clarify” or “have you considered”) about the work, what would it be?

- **[AS]**
  Is there a way to show that the ratio between the smallest and the largest eigenvalue must indicate the isotropy of the Voronoi cell, as defined in this paper?
- **[DS]**
  How can the choice of points around the bounding sphere (for Voronoi diagram estimation) affect the results?
  How does the mu_fit parameter affect the results?
- **[FP]**

The Laplacian term in the energy function is fundamental to guarantee a continuous (smooth) transition of the surface normal. However, for surface with sharp features (say a cube) this condition may be undesired. How does the method perform in surface with sharp features? It produces an oversmooth result or does it introduce additional artifacts?

- **[JD]**
- **[LF]**
  Could you elaborate on the time/space performance of your method?
- **[MK]**
  In the formulation of the B matrix, it seems that one should incorporate the inverse of the Hodge star (diagonally-lumped mass matrix) to account for the fact that the output of the DDG Laplacian takes values on the dual matrix. The authors suggest that the simplified implementation works well in practice. Would incorporating the Hodge star change things? For the better or for the worse? Why or why not?
  In defining anisotropy, why compare the ratio of largest to smallest eigenvalues, wouldn't it make more sense to compare largest to second largest (i.e. to ensure that the points are locally on a plane, not a curve)?