

Biharmonic Distance Driven Voronoi Diagrams Restricted on Curved Surfaces

ARTICLE INFO

Article history:

Surface Voronoi diagram, biharmonic embedding, intrinsic distance, spectral geometry processing

ABSTRACT

Computing surface-based Voronoi diagrams is a fundamental operation in geometry processing, typically relying on either geodesic or straight-line distances as solvers. However, when the input is a 3D model containing thin-plate structures, geodesic distances incur significant computational overhead, whereas straight-line distances can result in ownerless regions. To address this issue, we propose integrating biharmonic embedding distances into the SurfaceVoronoi framework. Specifically, mesh vertices can be embedded into a high-dimensional spectral space, ensuring that the embedding distance closely approximates the straight-line distance between sufficiently close points. In contrast, when points reside on opposite sides of a thin plate, the embedding distance significantly exceeds the straight-line distance, effectively preventing dominance from penetrating through to the opposite side. Our proposed framework offers several advantages: 1) It operates efficiently, as the embedding distance can be rapidly evaluated as a straight-line distance in high-dimensional space. 2) It guarantees the “one site, one region” property, even for models consisting of thin-plate structures. 3) It enables high-quality triangulation through iterative repositioning of each site to the centroid of its dominant region. Extensive experiments conducted on benchmark meshes demonstrate these benefits.

© 2025 Elsevier B.V. All rights reserved.

1. Introduction

Partitioning a 2-manifold surface into regions based on proximity to a set of sites is a fundamental problem in geometry processing. The resulting structure, known as a surface-based Voronoi diagram (VD), has been widely applied in tasks such as surface sampling, remeshing, parameterization, and shape abstraction [1].

Over the past decades, two primary types of distance metrics have been commonly adopted for constructing surface-based Voronoi diagrams: geodesic distances and Euclidean distances. Geodesic distances are intrinsic and faithfully reflect the shortest path along the surface, enabling accurate and natural Voronoi partitioning [2, 3]. However, they are computationally expensive and often tightly coupled with the diagram construction process [4], limiting scalability to large meshes. In contrast, Euclidean-based Voronoi diagrams (RVDs) [5] provide a more efficient alternative by utilizing straightforward Euclidean

distances in three-dimensional space. While faster, these approaches are extrinsic and prone to topological inconsistencies such as fragmented or ownerless regions, particularly on highly curved or thin-sheet surfaces. Although the propagation-based SurfaceVoronoi framework [6] partially mitigates the ownerless region problem inherent to Euclidean distances, it still produces geometrically unreasonable Voronoi regions in certain critical areas, such as the boundaries of thin structures. As shown in Fig. 1 with EDBVD, these limitations significantly restrict the practical application scope of this approach.

In this work, we extend the SurfaceVoronoi framework by introducing a distance metric based on biharmonic embedding. Our key insight is recognizing that this distance metric inherently reflects the intrinsic geometric properties of the surface, making it an ideal candidate for integration with the SurfaceVoronoi framework. By performing a one-time embedding of each mesh vertex into a high-dimensional spectral space constructed from the eigenvectors of the biharmonic operator [7],

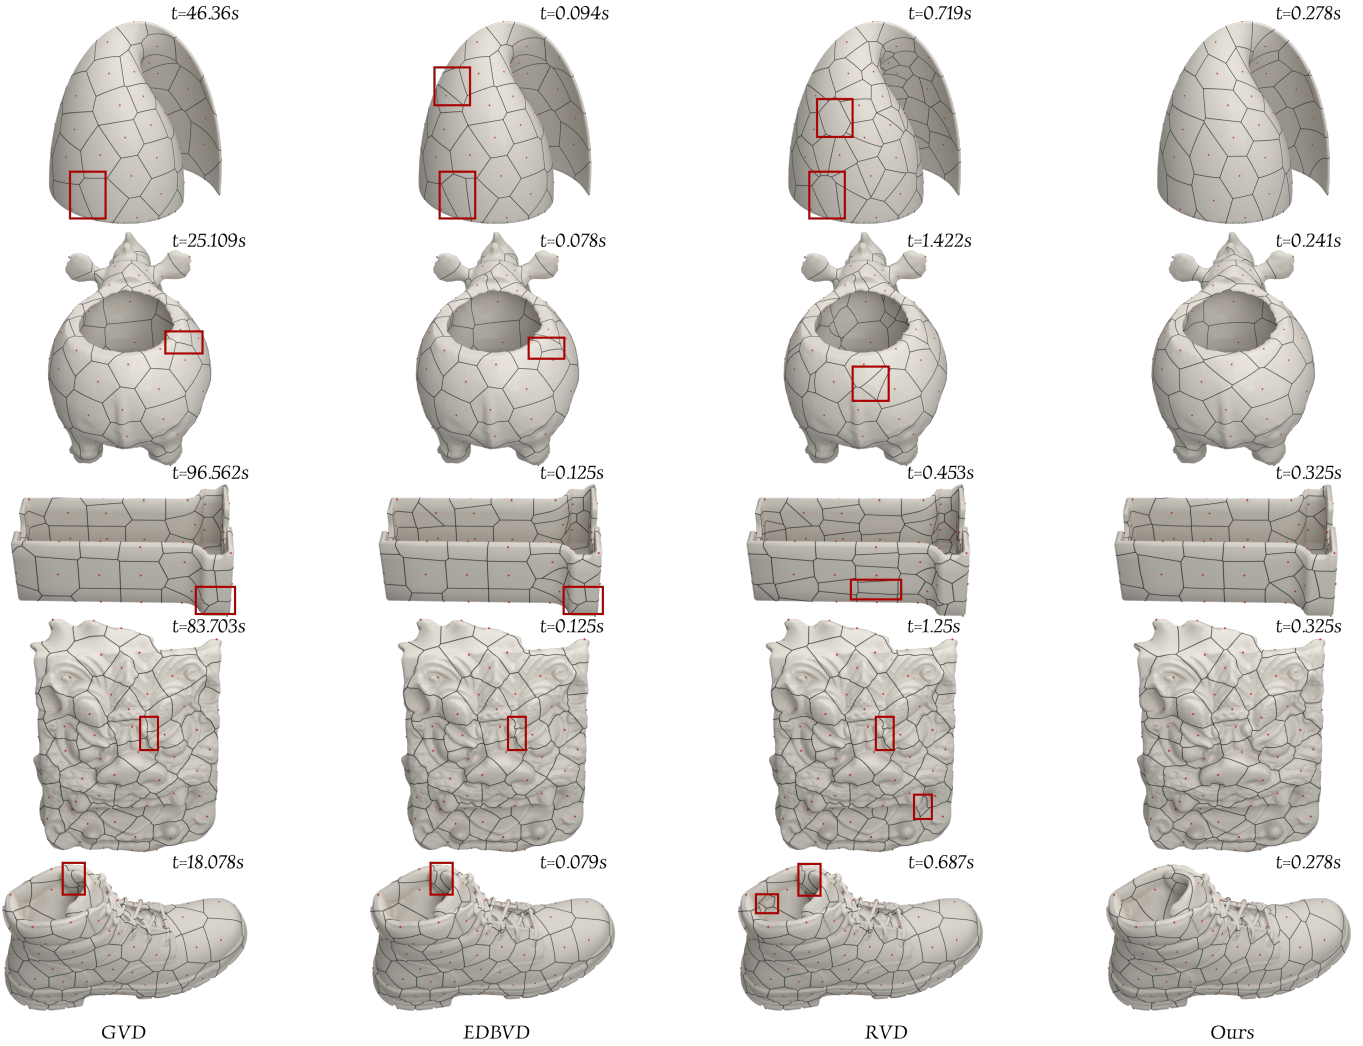


Fig. 1. Visual comparison of Voronoi diagrams generated by GVD, EDBVD, RVD, and our method on four models: Lampshade (a thin-shell model, top row), Pot (second row), Micro Block (third row), Shield (fourth row), and Boot (bottom row). Red dots indicate sampled sites; red boxes highlight noticeable differences among methods. Runtime (in seconds) is shown at the top-right of each subfigure. GVD yields the most accurate boundaries but is computationally expensive. RVD runs faster but often produces fragmented or ownerless regions. EDBVD improves efficiency and connectivity, but may suffer from jagged transitions due to non-metric propagation. In contrast, our method generates clean, well-connected partitions, avoiding the flaws of RVD and EDBVD, and offering a good balance between quality and efficiency. The Lampshade, Pot, Micro Block, Shield, and Boot models contain 20k, 20k, 13k, 35k vertices, and 19k vertices, respectively.

we enable distance calculations that are both computationally efficient and geometry-aware. Distances are simply evaluated as Euclidean norms between embedding vectors, avoiding the need for expensive geodesic computations while maintaining sensitivity to the surface’s intrinsic structure. This innovative combination yields a distance field that is smooth, globally coherent, and fully compatible with the existing SurfaceVoronoi pipeline without requiring any structural modifications, thereby enhancing partition quality while preserving computational efficiency.

Our embedding-driven SurfaceVoronoi algorithm offers a simple yet effective approach that overcomes key limitations of traditional distance metrics. By leveraging biharmonic embedding, our method elegantly handles challenging cases such as thin structures that conventional approaches struggle with, while maintaining computational efficiency and geometric fi-

delity.

2. Related Work

2.1. Surface-based Voronoi Diagram

The construction of Voronoi diagrams on curved surfaces, particularly non-differentiable polyhedral surfaces, plays a critical role in digital geometry processing applications. Surface-based Voronoi diagram computation methods can be broadly classified into three categories: (1) mesh-based approaches utilizing geodesic distances [8, 9], (2) parametrization-based methods operating on planar domains [10], and (3) operating in the 3D Euclidean space [11, 12, 13] through 3D Voronoi/Delaunay solvers.

Geodesic distance approaches offer theoretical precision but typically incur substantial computational costs [14]. Liu et

al. [8] analyzed the structure of iso-contours, bisectors, and Voronoi diagrams on triangular meshes, developing a geodesic Voronoi diagram (GVD) algorithm built upon the MMP algorithm's window propagation scheme [15]. Similar window propagation mechanisms have been employed in subsequent GVD approaches [16, 17, 18], but remain tightly coupled with geodesic computation frameworks.

Parametrization-based approaches [10, 19] require precomputing global parametrizations, making them susceptible to numerical instabilities and topological challenges when handling complex geometries.

Restricted Voronoi diagrams (RVD) [11, 12, 13] offer significant performance improvements by substituting geodesic distances with Euclidean distances. However, for models with thin structural elements, RVDs may violate the critical "one site, one region" property unless site distributions are sufficiently dense. Yan et al. [12] introduced localized RVD (LRVD) to address thin-plate models by first establishing neighborhood relationships, though triangulation quality issues remained unresolved. Wang et al. [13] later developed a fast post-processing technique to remediate problematic RVD cells by enforcing the "one site, one region" property.

Alternative approaches include the diffusion diagrams proposed by Herholz et al. [20], which leverage heat diffusion principles, though these are computationally equivalent to back-substitution only when system matrix factorization is precomputed.

The most recent advancement in computing Voronoi diagrams on surfaces is the SurfaceVoronoi framework [6], which innovatively decouples Voronoi diagram computation from distance metrics, presenting a generalized framework for Voronoi diagram construction. Wang et al. [21] extended this approach to 3D space. While its propagation strategy partially addresses the ownerless region problem encountered when using Euclidean distances, it remains constrained by the inherent properties of Euclidean metrics, leading to unreasonable behaviors in certain areas, particularly at the thin-plate edge regions illustrated in Fig. 1.

2.2. Biharmonic Distance

Lipman et al. [7] proposed the *biharmonic distance* (BH), obtained by weighting Laplacian eigenmodes with the inverse squared eigenvalues. This mathematical formulation bridges the gap between previously established distance measures while addressing their limitations. Unlike geodesic distances, which lack shape awareness and exhibit derivative discontinuities, or diffusion distances, which require careful parameter tuning, BH achieves an elegant balance of local and global properties. Their theoretical analysis demonstrated that biharmonic distance possesses three key advantages: (i) it is smooth and locally isotropic, mirroring geodesic behaviour close to a point; (ii) it remains sensitive to large-scale structure without requiring a user-chosen diffusion time; and (iii) it is numerically robust to mesh irregularities and mild topological noise. These properties make biharmonic distance particularly suitable for mesh processing tasks requiring both local precision and global shape awareness.

In this paper, we propose integrating biharmonic distances into the SurfaceVoronoi framework, effectively addressing the challenges of thin-plate structures while maintaining computational efficiency and topological correctness. Our approach ensures the "one site, one region" property even for complex models and enables high-quality triangulation through iterative site repositioning.

3. Method

Our approach leverages the flexibility of the SurfaceVoronoi framework, which supports arbitrary distance metrics for Voronoi diagram computation. We employ high-dimensional biharmonic embedding to transform the mesh vertices into a spectral space where Euclidean distances between embedded vectors effectively represent biharmonic distances on the original surface. By feeding these biharmonic distance measurements into the SurfaceVoronoi framework, we obtain partitioning results that significantly outperform conventional Euclidean distance-based approaches while preserving computational efficiency.

The proposed algorithm comprises three principal stages: (1) embedding of input sites into the triangle mesh, a necessary preprocessing step as the biharmonic distance formulation only supports calculations between mesh vertices; (2) biharmonic embedding of mesh vertices into a high-dimensional spectral space; and (3) Voronoi diagram computation on the original surface using distances derived from this embedding. This methodology maintains computational efficiency while substantially enhancing geometric fidelity, particularly in challenging surface configurations such as thin structures and regions exhibiting high curvature, where traditional approaches often produce topological inconsistencies or fragmented partitions.

Furthermore, we distinguish our approach from prior work. Although SurfaceVoronoi [6] briefly includes an example using biharmonic distance, it implements this approach inefficiently by computing a separate global distance field for each site, requiring pairwise calculations between each site and all mesh vertices. In contrast, our method first embeds all vertices into a high-dimensional spectral space, enabling direct computation of the Voronoi diagram through local triangle-wise operations. This approach significantly improves efficiency and scalability while preserving geometric fidelity.

3.1. Site Integration for Mesh Preprocessing

A critical constraint of our biharmonic embedding approach is that it exclusively operates on mesh vertices, whereas Voronoi sites rarely coincide with these predefined vertex locations. Given our fundamental assumption that all sites lie on the surface, we must integrate these sites into the mesh structure prior to computing both the biharmonic embedding and the resulting Voronoi diagram.

We implement a straightforward yet effective approach for site integration. For sites s_1, \dots, s_n located on a mesh triangle $\Delta v_1 v_2 v_3$, all of which inherently occupy the same plane, we perform a Delaunay triangulation on the combined point set $\{s_1, \dots, s_n, v_1, v_2, v_3\}$. This triangulation necessarily preserves

the three edges of the original triangle $\Delta v_1 v_2 v_3$ (in cases where a site lies precisely on a triangle edge, we apply a small perturbation to relocate it to the triangle interior, avoiding degenerate configurations). Through this procedure, we effectively embed sites s_1, \dots, s_n into the original triangular face while maintaining the mesh's topological integrity.

By applying this operation to each triangle containing sites, we seamlessly integrate all Voronoi sites into the mesh structure, enabling subsequent biharmonic embedding across the augmented vertex set.

3.2. Biharmonic Embedding

After integrating Voronoi sites into the mesh, we compute the biharmonic embedding of the augmented mesh vertices. This embedding maps each vertex into a high-dimensional spectral space derived from the eigenvectors of the Laplace–Beltrami operator [22, 23].

The key insight of our approach is that in this embedding space, Euclidean distances between vectors closely approximate biharmonic distances on the original surface [24, 25]. Formally, the squared biharmonic distance between two vertices x and y is defined as

$$d_B(x, y)^2 = \sum_{k=1}^{\infty} \frac{(\phi_k(x) - \phi_k(y))^2}{\lambda_k^2}, \quad (1)$$

where ϕ_k and λ_k are the k -th eigenfunction and eigenvalue of the Laplace–Beltrami operator, respectively [7].

By representing each mesh vertex as a point in this spectral space, we transform the complex problem of computing intrinsic surface distances into simple Euclidean distance calculations.

We construct this embedding by first computing the cotangent Laplacian matrix of the mesh and solving for its eigenvectors and eigenvalues. An equivalent form of the biharmonic distance uses the Green's function of the biharmonic operator:

$$d_B(x, y)^2 = g_B(x, x) + g_B(y, y) - 2g_B(x, y), \quad (2)$$

$$g_B(x, y) = \sum_{k=1}^{\infty} \frac{\phi_k(x)\phi_k(y)}{\lambda_k^2}. \quad (3)$$

The embedding coordinates are then formed by weighting these eigenvectors according to the inverse squared eigenvalues, following the formulation proposed by Lipman et al. [7].

This embedding approach offers significant computational advantages while preserving the desirable properties of biharmonic distances, including their balance between local isotropy and global shape awareness. The resulting distance metric is both efficient to compute and geometrically meaningful, making it particularly well-suited for integration with the SurfaceVoronoi framework in the next stage of our algorithm.

3.3. Voronoi Diagram Computation

After obtaining the high-dimensional biharmonic embedding, we utilize it to compute the Voronoi diagram on the original surface by integrating with the SurfaceVoronoi framework.




Fig. 2. Illustration of the triangle-level cutting process used to trace Voronoi boundaries. Following [6], we perform incremental half-plane cutting to compute the lower envelope of distance fields inside each triangle, enabling precise Voronoi region extraction.

SurfaceVoronoi consists of two main stages: propagation and cutting. The propagation stage identifies contributing distance fields for each triangle on the mesh. In our implementation, Voronoi sites propagate outward from their containing triangles following the SurfaceVoronoi framework, but we replace the original distance calculation with measurements derived from our biharmonic embedding. Specifically, the distance between two points is computed as the Euclidean distance between their corresponding embedding vectors.

In the cutting stage, we adopt the squared distance linear approximation assumption from SurfaceVoronoi to improve computational accuracy. Specifically, we adopt the local linear approximation of distance fields from the SurfaceVoronoi method [6], where the distance at any point (x, y) inside a triangle is modeled as a linear function:

$$d(x, y) = ax + by + c, \quad (4)$$

with coefficients obtained by solving the linear system

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{bmatrix}^{-1} \begin{bmatrix} d_1 \\ d_2 \\ d_3 \end{bmatrix}, \quad (5)$$

where (x_j, y_j) are the coordinates of the triangle's three vertices, and d_j are their respective distances to a given Voronoi site (measured in the embedding space). The cutting process used to determine Voronoi boundaries inside each triangle is illustrated in Fig. 2.

This approach allows us to determine the Voronoi partition within each triangle with higher precision while utilizing the geometric information captured by the biharmonic embedding.

This integration creates an efficient pipeline that maintains the intrinsic geometric properties of the biharmonic distance while leveraging the established computational framework of the SurfaceVoronoi algorithm [26].

4. Experiments

4.1. Experimental Setting

We implemented our method in C++. All experiments were conducted on a Windows 11 machine with an Intel Core i9-13900HX CPU (2.20 GHz) and 32GB of RAM. All reported




Fig. 3. Surface Voronoi diagrams generated by our method on 15 diverse models. Red dots indicate sampled sites. These results highlight the robustness and adaptability of our method across different surface geometries and complexities.

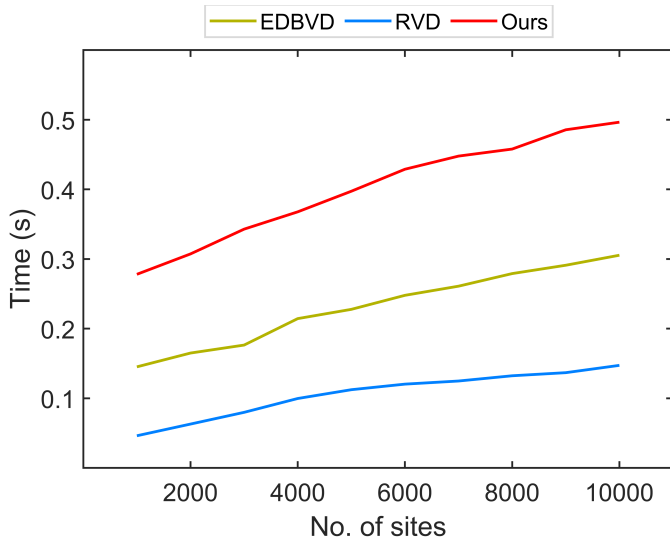


Fig. 4. Computational efficiency comparison of EDBVD, RVD, and ours as site count increases from 1,000 to 10,000. Total runtime averaged across 100 models with diverse geometries (biharmonic embedding dimension $d = 10$). GVD excluded due to prohibitive computational cost.

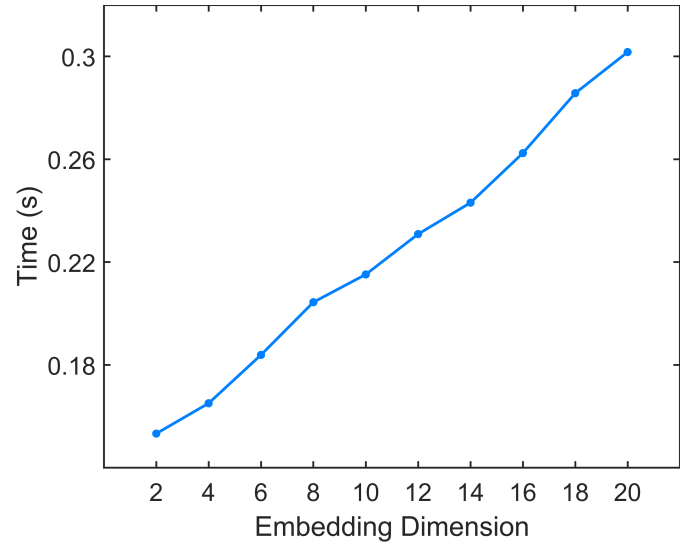


Fig. 5. Relationship between computational embedding time and embedding dimension across 100 models. The number of sampled sites is fixed at 100. Embedding time exhibits approximately linear growth with increasing dimension due to additional spectral computations, while maintaining efficiency for practical applications.

timings include both the distance propagation and bisector extraction stages.

Furthermore, our method relies on biharmonic embedding derived from Laplace–Beltrami eigenfunctions, which are mathematically defined on smooth Riemannian manifolds [7]. This formulation necessitates that input surfaces be well-formed 2-manifolds. Application of spectral methods to non-manifold geometries can introduce numerical instability due to ill-defined cotangent weights and singular Laplacian operators. Our implementation strictly adheres to these input assumptions. Consequently, all test models in our experimental evaluation are manifold meshes.

4.2. Visual Comparison

To qualitatively assess the partitioning quality of different methods, we conduct a visual comparison on three representative surface models: Lmpshade, Pot, Micro Block, Bunny and Shield. We evaluate Geodesic Voronoi Diagrams (GVD), Restricted Voronoi Diagrams (RVD), Euclidean Distance-Based Voronoi Diagrams (EDBVD) using the SurfaceVoronoi framework, and our proposed biharmonic embedding-based method. For each model, 100 sites are selected via blue noise sampling. The same set of sampled sites is used for all methods to ensure consistency and fairness.

Fig. 1 shows the results of different methods. From the visual results, we observe the following:




Fig. 6. Ablation study of embedding dimension. Voronoi partitions generated with different embedding dimensions ($d = 3, 4, 6, 10, 50$) on five models: Lampshade, Vase, Pig, Boot, each containing approximately 20K vertices, and Shield with around 35K vertices. Red dots denote sampled sites, and red boxes highlight areas with noticeable differences among embedding settings.

- **GVD** yields the most geometrically accurate partitions, with smooth and well-aligned boundaries. However, it suffers from extremely high computational cost, making it impractical for large-scale models or time-sensitive applications.
- **RVD** offers very fast computation but frequently generates fragmented or undefined regions, especially in thin plate or high-curvature areas. This occurs because Euclidean distance cannot effectively capture the intrinsic properties of the surface.
- **EDBVD** partially mitigates the ownerless region problem through its propagation mechanism, but it remains constrained by the limitations of Euclidean distance. As a result, it still produces geometrically unreasonable partitioning, particularly along the boundaries of thin structures.
- **Our method** consistently produces clean, connected, and well-shaped regions. Although not the fastest among all methods, it avoids the topological flaws of RVD and the boundary artifacts of EDBVD. This results in a favorable balance between computational efficiency and visual quality.

In addition, Fig. 3 shows results on 15 diverse models, further demonstrating that our method consistently produces robust and semantically meaningful partitions across a wide range of surface geometries. These results reinforce the adaptability and generalization capability of our embedding-driven approach.

4.3. Time Comparison

We evaluate the runtime efficiency of our method from two perspectives: (1) the impact of the number of sites on the prop-




Fig. 7. Voronoi partitions generated with 20, 50, 100, 200, and 500 sampled sites on four models: Lampshade, Fish Pen Pot, Trinket Boxes, and Camera. The first three are thin-shell models, while the Camera is a relatively regular solid model. Red dots denote sampled sites. The experiment aims to evaluate the consistency and robustness of partitioning under varying sampling densities. Even under sparse sampling, our method produces well-formed regions without fragmentation or ambiguity.

agation and cutting stages, and (2) the effect of the embedding dimension on the embedding computation cost.

4.3.1. Runtime Performance with Increasing Site Density

We conducted performance analysis on computational scalability with respect to increasing Voronoi site density. Our evaluation utilized a dataset of 100 surface meshes exhibiting varied topological characteristics, with an average resolution of approximately 20,000 vertices per model. We fix the embedding dimension at $d = 10$ and vary sampled sites from 1,000 to 10,000, measuring the runtime of both propagation and cutting stages. All sites were generated using blue noise sampling to ensure consistent distribution.

Fig. 4 compares RVD, EDBVD, and our method. We omit GVD results due to its prohibitive runtime (exceeding 60 seconds for high site counts). Our method demonstrates linear scalability with respect to site count. While moderately slower than EDBVD due to the embedding computation overhead, our approach provides significantly better geometric consistency and

smoother Voronoi partitions. This moderate performance cost represents a reasonable trade-off for achieving high-quality surface Voronoi diagrams that avoid the artifacts present in faster methods.

4.3.2. Computational Efficiency of Biharmonic Embedding

We conducted a systematic assessment of biharmonic embedding computation time across dimensions from $d = 2$ to $d = 20$ on 100 models, while maintaining a constant sample density of 100 Voronoi sites per model. Fig. 5 illustrates that computational time increases approximately linearly with embedding dimension. This behavior corresponds to the additional computational requirements of calculating higher-order eigenfunctions in the Laplace–Beltrami spectral decomposition. Despite this linear growth, the embedding procedure maintains practical efficiency, with mean processing times consistently below one second even at the maximum tested dimension ($d = 20$). This performance ensures the method’s viability for integration into comprehensive surface analysis workflows.




Fig. 8. Voronoi diagram computation on meshes with poor triangulation. The input meshes (left) exhibit numerous elongated and ill-conditioned triangles. Our method produces well-structured Voronoi partitioning (right) despite the irregular triangulation, demonstrating algorithmic robustness to poor mesh quality.




Fig. 9. Surface power diagram with weighted sites. Sites with positive weights are shown as red dots, and those with zero weights are shown as dark green dots.

4.4. Ablation Study

We analyze how key parameters affect the quality of Voronoi partitions generated by our method by examining embedding dimension.

We evaluate how embedding dimensionality affects partitioning quality using the Lampshade, Vase, Pig, Shield, and Boot models with 100 sites each, varying dimension d among 3, 4, 6, 10, 50).

Fig. 6 shows that increasing the embedding dimension leads to progressively more accurate partitioning by better capturing the intrinsic geometry. Notably, the transition from $d = 3$ to $d = 6$ significantly improves boundary smoothness and region coherence. However, the results computed at $d = 10$ and $d = 50$ exhibit only slight differences in the resulting partitions. Based on these observations, we adopt $d = 10$ as our default configuration, as it offers a good balance between partition quality and computational cost.

4.5. Voronoi Diagrams at Different Site Densities

We conducted experimental evaluations on several models using a fixed embedding dimension $d = 10$, generating Voronoi




Fig. 10. Surface Voronoi diagram with user-defined breakline constraints. The blue curve represents a breakline, which acts as a barrier to prevent Voronoi cells from propagating across breakline.

diagrams at varying site densities: 20, 50, 100, 200, and 500 sites per model.

Fig. 7 shows the visual progression as site count increases. With fewer sites (20-50), the partitions naturally contain larger regions. As the number of sites increases to 100, 200, and 500, the diagram exhibits progressively finer granularity. The visualizations demonstrate our method's ability to handle different sampling densities while maintaining consistent partitioning behavior across all tested models.

4.6. Robustness

The ability to compute high-quality Voronoi diagrams on imperfect meshes is crucial for practical applications, as real-world models often contain triangulation defects. This presents a significant challenge since the cotangent Laplacian formulation underlying our method can be influenced by mesh quality. To assess robustness, we evaluated our approach on meshes with suboptimal triangulation characteristics. Fig. 8 presents results on test meshes containing elongated triangles and non-uniform vertex distribution. The resulting Voronoi partitions demonstrate acceptable quality without significant distortions or topological inconsistencies, suggesting that our computation exhibits a degree of resilience to imperfections in the underlying triangulation.

4.7. Extensions and Applications

Our method preserves the flexible extensibility of the SurfaceVoronoi framework. By leveraging this foundation, we can directly support various diagram types with only minor modifications to the distance computation.

4.7.1. Power Diagram

Our approach easily supports power diagrams by modifying the distance function for a site p with weight w to:

$$d(v, p) = \|v - p\|^2 - w. \quad (6)$$

Fig. 9 demonstrates this capability, where weighted sites (20 out of 1000 with nonzero weights) produce appropriately larger Voronoi cells.

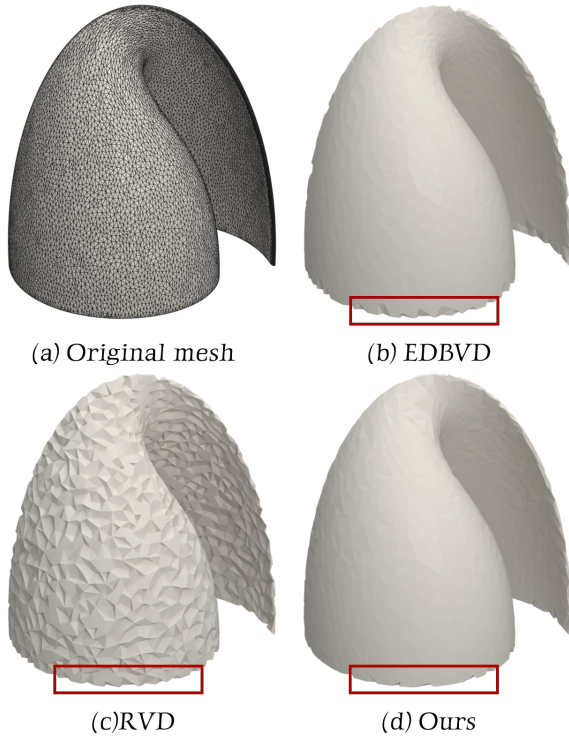


Table 1. Quantitative comparison between remeshed and original surfaces across 100 test models (3000 sites per model). All distance metrics are normalized relative to the bounding box diagonal.

Method	Hausdorff Distance	Chamfer Distance
EDBVD	0.0231	0.0064
RVD	0.0238	0.0066
Ours	0.0219	0.0061

ric fidelity across diverse models.

5. Conclusion

We presented an enhancement to the SurfaceVoronoi framework by replacing the original Euclidean embedding metric in EDBVD with a biharmonic embedding. By leveraging the smoothness and intrinsic distance-preserving properties of biharmonic embeddings, our method offers a more faithful approximation of surface geometry while retaining the computational efficiency of embedding-based distance evaluation.

Extensive experiments on a variety of mesh models demonstrate that our approach consistently produces high-quality Voronoi partitions with clearer boundaries and fewer topological artifacts. Compared to geodesic-based methods such as GVD, our method achieves substantial reductions in computation time. Compared to existing EDBVD and RVD methods, it significantly improves partition accuracy and region coherence.

Moreover, our method exhibits strong scalability with respect to both the number of sampled sites and the mesh complexity, making it suitable for large-scale surface processing tasks. These results confirm that biharmonic embeddings provide a promising direction for efficient and intrinsic-aware surface partitioning.

In future work, we plan to explore adaptive embedding dimensionality strategies, investigate faster spectral solvers for large-scale Laplacian eigenproblems, and extend the framework to support anisotropic or direction-aware distance metrics.

References

- [1] P. Giblin, B. B. Kimia, A formal classification of 3d medial axis points and their local geometry, *IEEE Trans. Pattern Anal. Mach. Intell.* 26 (2) (2004) 238–251.
- [2] V. Surazhsky, T. Surazhsky, D. Kirsanov, S. J. Gortler, H. Hoppe, Fast exact and approximate geodesics on meshes, *ACM Trans. Graph.* 24 (3) (2005) 553–560.
- [3] K. Crane, C. Weischedel, M. Wardetzky, Geodesics in heat: A new approach to computing distance based on heat flow, *ACM Trans. Graph.* 32 (5) (Oct. 2013).
- [4] R. Kimmel, J. A. Sethian, Computing geodesic paths on manifolds, *Proceedings of the National Academy of Sciences* 95 (15) (1998) 8431–8435.
- [5] D. Boltcheva, B. Lévy, Surface reconstruction by computing restricted voronoi cells in parallel, *Comput. Aided Des.* 90 (C) (2017) 123–134.
- [6] S. Xin, P. Wang, R. Xu, D. Yan, S. Chen, W. Wang, C. Zhang, C. Tu, SurfaceVoronoi: Efficiently computing voronoi diagrams over mesh surfaces with arbitrary distance solvers, *ACM Trans. Graph.* 41 (6) (Nov. 2022).
- [7] Y. Lipman, R. M. Rustamov, T. A. Funkhouser, Biharmonic distance, *ACM Trans. Graph.* 29 (3) (Jul. 2010).

Fig. 11. Remeshing comparison on the *Lampshade* model with 3000 sampled sites. (a) Original mesh. (b) EDBVD-based remeshing. (c) RVD-based remeshing. (d) Remeshing using our method. Our method produces more feature-aligned triangulations, demonstrating superior quality under identical sampling conditions.

4.7.2. Breakline Constraints

Our method also handles breakline constraints, where user-defined feature lines prevent Voronoi regions from crossing important surface features. As shown in Fig. 10, the diagram properly respects these constraints by limiting propagation across breaklines.

4.7.3. Remeshing

Surface remeshing is a fundamental application in geometry processing with significant practical implications. We conducted experimental evaluation using this application to assess our method’s effectiveness. The remeshing process involves computing the Voronoi diagram on the surface and extracting its dual to generate a triangulated mesh. We compared our approach against EDBVD [6] and RVD on the *Lampshade* model (20K vertices), using identical configurations of 3000 sampled sites. Fig. 11 presents comparative results: the original mesh (a), followed by remeshing results from EDBVD (b), RVD (c) and ours (d). The results demonstrate that our approach yields enhanced triangulation quality with superior feature preservation.

For quantitative validation, we conducted experiments on 100 manifold models with 3000 blue-noise distributed sites per model. Table 1 presents the average bidirectional Hausdorff Distance (HD) and bidirectional Chamfer Distance (CD) between remeshed and original surfaces. Our method achieves the lowest values in both metrics, confirming superior geomet-

- [8] Y.-J. Liu, Z. Chen, K. Tang, Construction of iso-contours, bisectors, and voronoi diagrams on triangulated surfaces, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33 (8) (2011) 1502–1517.
- [9] X. Wang, X. Ying, Y.-J. Liu, S.-Q. Xin, W. Wang, X. Gu, W. Mueller-Wittig, Y. He, Intrinsic computation of centroidal Voronoi tessellation (CVT) on meshes, *Computer-Aided Design* 58 (2015) 51–61.
- [10] G. Rong, M. Jin, X. Guo, Hyperbolic centroidal voronoi tessellation, in: *Proceedings of the 14th ACM Symposium on Solid and Physical Modeling*, ACM, NY, USA, 2010, p. 117–126.
- [11] D.-M. Yan, B. Lévy, Y. Liu, F. Sun, W. Wang, Isotropic remeshing with fast and exact computation of restricted Voronoi diagram, Vol. 28, Wiley Online Library, 2009, pp. 1445–1454.
- [12] D.-M. Yan, G. Bao, X. Zhang, P. Wonka, Low-resolution remeshing using the localized restricted Voronoi diagram, *IEEE transactions on visualization and computer graphics* 20 (10) (2014) 1418–1427.
- [13] P. Wang, S. Xin, C. Tu, D. Yan, Y. Zhou, C. Zhang, Robustly computing restricted Voronoi diagrams (RVD) on thin-plate models, *Computer Aided Geometric Design* (2020) 101848.
- [14] H.-T. D. Liu, M. Gillespie, B. Chislett, N. Sharp, A. Jacobson, K. Crane, Surface simplification using intrinsic error metrics, *ACM Trans. Graph.* 42 (4) (Jul. 2023).
- [15] J. S. Mitchell, D. M. Mount, C. H. Papadimitriou, The discrete geodesic problem, *SIAM Journal on Computing* 16 (4) (1987) 647–668.
- [16] C. Xu, Y.-J. Liu, Q. Sun, J. Li, Y. He, Polyline-sourced geodesic voronoi diagrams on triangle meshes, Vol. 33, Wiley Online Library, 2014, pp. 161–170.
- [17] Y. Qin, H. Yu, J. Zhang, Fast and memory-efficient voronoi diagram construction on triangle meshes, Vol. 36, 2017, pp. 93–104.
- [18] Y. Qin, X. Han, H. Yu, Y. Yu, J. Zhang, Fast and exact discrete geodesic computation based on triangle-oriented wavefront propagation, *ACM Trans. Graph.* 35 (4) (Jul. 2016).
- [19] P. Alliez, É. C. de Verdière, O. Devillers, M. Isenburg, Centroidal Voronoi diagrams for isotropic surface remeshing, *Graph. Model.* 67 (2005) 204–231.
- [20] P. Herholz, F. Haase, M. Alexa, Diffusion diagrams: Voronoi cells and centroids from diffusion, *Computer Graphics Forum* 36 (2) (2017) 163–175.
- [21] P. Wang, J. Song, L. Wang, S. Xin, D.-M. Yan, S. Chen, C. Tu, W. Wang, Towards voronoi diagrams of surface patches, *IEEE Transactions on Visualization and Computer Graphics* (2025) 1–15.
- [22] R. M. Rustamov, Laplace-beltrami eigenfunctions for deformation invariant shape representation, in: *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, SGP '07, Eurographics Association, Goslar, DEU, 2007, p. 225–233.
- [23] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, L. Guibas, Functional maps: a flexible representation of maps between shapes, *ACM Trans. Graph.* 31 (4) (Jul. 2012).
- [24] A. G. Belyaev, P.-A. Fayolle, On variational and pde-based distance function approximations, *Computer Graphics Forum* 34 (8) (2015) 104–118.
- [25] F. Fouss, A. Pirotte, J.-m. Renders, M. Saerens, Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation, *IEEE Transactions on Knowledge and Data Engineering* 19 (3) (2007) 355–369.
- [26] N. Maruani, R. Klokov, M. Ovsjanikov, P. Alliez, M. Desbrun, Voromesh: Learning watertight surface meshes with voronoi diagrams (2023). [arXiv:2308.14616](https://arxiv.org/abs/2308.14616).