DOI: 10.1111/cgf.70373

EUROGRAPHICS 2026 / B. Masia and J. Thies
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 45 (2026), Number 2

Contouring Signed Distance Fields by Approximating Gradients

M. Kohlbrenner [1] and M. Alexa[1]

[1]TU Berlin, Computer Graphics Group

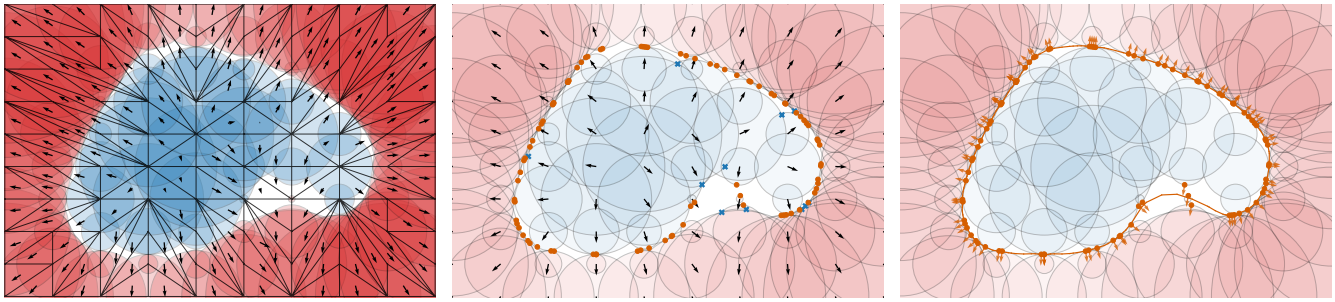

Figure 1: *Given a discrete set of signed distance values (visualized as circles with radius according to distance and color distinguishing sign), we first compute a triangulation of sample locations (left). The triangulation gives rise to a piecewise linear approximation of the signed distance function, with constant gradients in each triangle. Averaging the triangle gradients to the vertices allows approximating gradients for the distance samples (center). The gradients point towards the contour and, together with the distance value, yield points on the surface with normals. To avoid erroneous samples from the approximation, surface points contained in any of the circles are discarded. The contour is then reconstructed using Poisson Surface Reconstruction (right).*

Abstract

Signed distance fields are often represented by discrete samples (e.g., on a grid). Recovering the contour implicitly represented by the distance samples requires an approximation algorithm. Several recent approaches have shown that exploiting the information carried in each distance sample by explicitly constructing a surface point gives better results than classical contouring algorithms. We explore the idea of generating surface points by simply approximating the gradient of the signed distance function from a tesselation of the sample locations. The distance value together with gradient yields a potential surface point. To avoid problems resulting from bad approximation, surface points are removed if they are too close to any of the distance samples. Using the regular triangulation as tesselation facilitates this filtering. The resulting approximation algorithm is conceptually simple, easy to implement, and significantly faster than existing alternatives, yielding reconstructions that are on par.

CCS Concepts

• ***Methods/Applications*** → *Geometry;* • ***Modeling/Geometry*** → *Geometry Processing; Implicit Curves/Surfaces;*

1. Introduction

Implicit surfaces are a common representations of solids, useful in many geometry processing tasks. The solid is encoded by the sign of a scalar function, so that the zero-level set represents the surface. A signed distance field (SDF) is a particular implicit representation, where the scalar value represents the distance to the closest point on the surface. While the value of an arbitrary implicit function carries little information (apart from the sign), every value of the SDF provides the exact distance to the surface – and the gradient of the SDF provides the direction.

These properties have been exploited in algorithms for interrogating the surface, such as sphere tracing [Har96]. Recently, it has been observed that also for the task of recovering an approximation of the surface from a discrete set of samples the information carried in each sample is useful [SBS23]. In fact, the underlying idea of the most recent methods exploiting this information [SRBS24, KA25b] directly uses the information carried in the samples: find a point at the given distance from the sample location that is 'feasible', i.e., not closer to any other sample than their distance value. Such point is likely close to the true surface, as for

every distance sample there has to be a corresponding point on the surface. The estimated surface points are used for recovering the surface, using Poisson Surface Reconstruction [KBH06, KH13], although any method for this task could be used.

Sellan et al. [SRBS24] have designed an intricate optimization process for finding plausible locations of surface points, requiring multiple instances of PSR. Kohlbrenner and Alexa [KA25b] suggest a higher dimensional polyhedral construction of vertices of the Apollonius diagram to generate surface points. We argue that the information required for estimating a surface point from the distance values in the samples is readily contained in the given data, namely as the local variation of the SDF values, providing an estimate of the gradient in the samples.

We show that approximating the gradient from a tessellation of the samples is 'good enough' to estimate surface points. To avoid problems resulting from bad approximations we suggest filtering the candidates for surface points: similar to other methods we required that any surface point is not closer to any surface sample than its (absolute) distance value. Using the regular triangulation of the distance samples greatly facilitates this filtering [KA25a], because only a small constant set of samples need to be checked. Putting these two ideas together results in very simple, deterministic contouring approach for discrete SDFs.

In the following we first briefly provide some more background for this idea (Secs 2 and 3). We then explain the details of generating the surface points (Sec. 4, including several ways to approximate the gradients in the sample locations from the neighborhood provided by the regular triangulation and their accuracy compared to ground truth as well as the filtering stage. We compare the results of reconstructing the surface from the surface points to other methods and, as a baseline, surface points corresponding to the distance samples based on ground truth data (Sec. 5), showing that the surface points approximated from gradients are very fast, yet provide similar or better results. We conclude with a brief discussion in Sec. 6.

2. Related Work

This paper directly follows a recent line of work [SBS23, SRBS24, KA25a, KA25b]. We keep this section short and refer readers to these references.

2.1. Contouring

The classic primal construction for contouring is marching cubes [LC87], which is still amongst the most used today. The SurfaceNets construction of Gibson [Gib98] uses a simple dual construction, i.e. generates a vertex per volume element and connects the vertices based on sign information. The extended marching cubes algorithm [KBSS01] introduces a dual construction near sharp features. Ju introduces a more general dual contouring algorithm [JLSW02] and merges the vertex construction of of Kobbelt et al. [KBSS01] with the connectivity of the Gibson method [Gib98]. The dual marching cubes algorithm of Schaefer and Warren [SW04] applies primal contouring to dual grids and provides a manifold surface [SJW07]. Other work includes [Nie04] and [ZGG22, GZ22]. For simplicial domains, the marching tetrahedra algorithm [Blo88] is commonly used.

The contouring setup can also be extended to non-manifold geometry [BF95], some of which can be represented via vector-valued functions [DZCJ22, JDZ*24]. More recently, data-driven variants of classic contouring techniques have been proposed [CZ21, CTFZ22] that can reproduce sharp features more accurately.

2.2. Contouring from Sparse SDF samples

Sellán et al. [SBS23] were the first to propose the geometric view of SDF samples as signed spheres in the context of contouring from sparse SDF samples. They define a SDF energy and use it to optimize an initial mesh to adhere more to tangency constraints. The explicit representation of the surface as a mesh limits the applicability to surfaces of known genus. In their follow up work, the authors express the surface via a set of explicit tangent points and optimize these locations [SRBS24]. Kohlbrenner and Alexa [KA25b] show that even a simple, heuristic choice of such tangent points based on largest, not intersecting spheres, can lead to very similar results without the need for optimisation. The same authors also propose the use of the weighted Delaunay / power diagram duality for this setting and present an efficient re-sampling strategy for the case that the SDF can be queried at novel positions [KA25a]. In a preprint [WZW*25], Wang et al. explore a very similar method that differs in the choice of tetrahedra for the resampling and in the additional insertion of contact points into the triangulation. The 'PoNQ' representation of Maruani et al. [MOAD24] can be effectively predicted by a neural network from discrete SDF samples.

3. Background

3.1. Signed distance fields

The boundary surface $S = d\Omega \subset \mathbb{R}^d$ of a volume $\Omega \subset \mathbb{R}^d$ defines its signed distance field (SDF) $\Phi : \mathbb{R}^d \to \mathbb{R}$

$$\Phi(\mathbf{x}) = \chi_\Omega(\mathbf{x}) \cdot \min_{\mathbf{y} \in S} \|\mathbf{x} - \mathbf{y}\|. \tag{1}$$

where

$$\chi_\Omega(\mathbf{x}) = \begin{cases} -1, & \mathbf{x} \in \Omega \\ 1, & \text{else.} \end{cases}$$

is the indicator function of Ω.

The surface S corresponds to the zero levelset of the SDF $S = \{\mathbf{x} \in \mathbb{R}^d : \Phi(\mathbf{x}) = 0\}$. Where the SDF is differentiable, its gradient is of unit length $\|\nabla\Phi\| = 1$ (Eikonal equation). Its singularities correspond to the medial axis: the positions in space where the closest point on the surface is not unique.

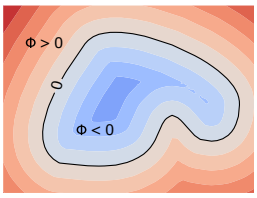

3.2. SDF contouring

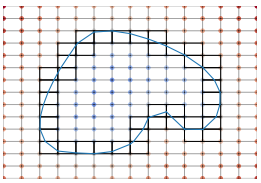

Standard contouring techniques such as marching cubes or marching tetrahedra approximate the function near the level set based on a discrete set of samples and extract the contour. Only the SDF values at the corners of volume elements in which a sign change occurs (highlighted squares in the inset) directly influence the contour .

In the special case of a SDF, each sample actually contains information on the zero level set (surface). Geometrically, each sample $(\mathbf{p}_i, \Phi(\mathbf{p}_i))$ can be interpreted as a sphere of radius $|\Phi(\mathbf{p}_i))|$ centered at $\mathbf{p}_i$ (see inset). The area outside of all spheres is called the "feasible region" by Sellán et al. [SRBS24]. The unknown surface separates positive from negative sphere centers, lies in the feasible region and is tangential to each sphere at least once [SBS23].

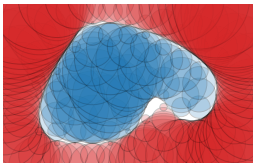

"Reach for the Arcs" (RFTA) [SRBS24] uses a set of oriented points located on the parts of the spheres that lie outside all other spheres as variables and uses sPSR to reconstruct a surface from them. This is a very efficient way to enforce tangency constraints. The location of these points is then used in an optimisation approach. Kohlbrenner and Alexa follow the same surface representation [KA25b] but show that even a single point per sphere, placed via a simple heuristic, can lead to very similar results. The difficulty in this approach lies in a consistent placement of such contact points.

3.3. Connectivity

While marching cubes assumes the samples $\mathbf{p}_i$ to lie on a regular grid, in general the samples could be placed arbitrarily. Any simplicial complex suffices to provide volumetric cells for the application of marching tetrahedra. Kohlbrenner and Alexa [KA25a] suggest to compute the cells for each sample from the squared distances to the spheres, leading to a power diagram and the simplicial complex being the weighted Delaunay triangulation of the weighted points $(\mathbf{p}_i, \Phi(\mathbf{p}_i)^2)$. It is shown that the connectivity encodes the parts of the spheres outside all other spheres. Moreover it is claimed that this weighted Delaunay triangulation leads to slightly better reconstructions from Marching Tetrahedra compared to a standard Delaunay triangulation.

4. Method

The purely geometric interpretation of a set of samples $\{(\mathbf{p}_i, \Phi(\mathbf{p}_i))\}$ as spheres obscures one crucial detail: Each sample corresponds to at least one contact point $\mathbf{c}_i = \min_{\mathbf{q}\in S} \|\mathbf{p} - \mathbf{q}\| \in S$, at which the closest distance is attained. If $\mathbf{p}$ is part of the medial axis, the contact point is not unique. For a sample not on the medial axis, the direction towards the contact point also is consistent with the gradient of the signed distance function $\nabla\Phi(\mathbf{p}_i)$. This means knowing the gradient at the sample location directly yields the contact point.

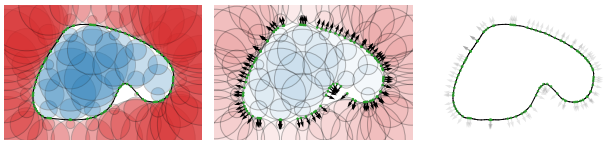

Figure 2: *Surface S, discrete samples of its SDF Φ visualized as red/blue circles and their green contact points with the surface (left). Contact points and the corresponding gradients $\nabla\Phi$ yield a sampling of the surface and its normal space (middle). sPSR surface reconstruction based on the contact points (right).*

$$\mathbf{c}_i = \mathbf{p}_i - \Phi(\mathbf{p}_i) \cdot \nabla\Phi(\mathbf{p}_i) \tag{2}$$

The contact points constitute an oriented sampling $\{(\mathbf{c}_i, \mathbf{n}_i)\}$ of the Surface S. The normal direction of the surface in each contact point is given by the gradient $\mathbf{n}_i = \nabla\Phi(\mathbf{p}_i)$. The contouring problem could then be formulated as a surface reconstruction from oriented points. Established methods such as Screened Poisson Reconstruction (sPSR) [KH13] yield good results (Figure 2).

Constructing surface points from a sample of the SDF and its gradient is straightforward [SSSC22]. In the context of SDF contouring, Sellán et al. discuss a contact point 'oracle' [SRBS24] and Wang et al. [WZW*25] use contact points as locations for re-sampling. Sellán et al. dismiss a purely contact point-based approach compared to their more elaborate optimization-based approach (see Fig. 9 [SRBS24]). While the RFTA optimisation can achieve results that are superior to the ones generated from ground truth contact points, we notice that this is not generally the case and depends on the parameters of the surface reconstruction technique (see Figure 8, in which "GTCP" corresponds to the ground truth contact points).

Previous approaches that construct the contour from a fixed set of samples assume that the gradient/contour point information is not available [SBS23, SRBS24, KA25b] . We propose to frame the contouring problem as an approximation of the contact points in two steps: (1) we make use of the local variation of the function over a simplicial complex in order to calculate vertex-wise approximate gradients and then (2) use the regular triangulation of the domain described in [KA25a] in order to efficiently filter out contact points that lie within other spheres.

4.1. Approximate gradients

We consider the piecewise-linear function f defined by the function values $\Phi(\mathbf{p}_i)$ at the vertices of a simplicial complex (V, F). The gradient of this function is constant in each simplex. It is common to assign a gradient to the vertex by averaging the gradients of the incident simplices t_j (cf. vertex surface normals in triangle meshes):

$$\nabla\tilde{f}(\mathbf{p}_i) \propto \sum_{i\ni\partial t_j} w_j \nabla f_j. \tag{3}$$

```
def GCP_contouring({p_i, Φ_i}):
    T = weighted_delaunay({p_i, Φ_i^2})
    contact_points = []
    for i in range(1,N+1):
        g̃_i = approximate_gradient(T,i)
        n_i = g̃_i / |g̃_i|
        c_i = p_i − Φ_i · n_i
        insert = True
        for (p_n, Φ_n) in T.neighbors(i):
            if ‖p_n − p_i‖^2 < Ω_n^2:
                insert=False
                break
        if insert:
            contact_points.append((c_i, n_i))

    return poisson_reconstruct(contact_points)
```

Listing 1: *Overview: "Gradient Contact Point" algorithm (GCP).*

For the weights we consider three choices:

1. Uniform weights: $w_j = 1$
2. Weights based on the volume of the simplices, motivated by the concept of *area vectors* for triangles and the observation that this construction at least preserves integrated gradients: $w_j = vol(t_j)$
3. The weights suggested by James [Jam20]) for consistent polynomial interpolation of the gradients.

4.2. Approximate contact points and filtering

We use the approximate gradient $\nabla\tilde{f}(\mathbf{p}_i)$ and construct an approximation of the contact point as in Equation (2) as $\tilde{\mathbf{c}}_i = \mathbf{p}_i - \Phi(\mathbf{p}_i) \cdot \nabla\tilde{\Phi}(\mathbf{p}_i)$. While many such $\tilde{\mathbf{c}}_i$ are samples on the surface, errors in the gradient approximation, e.g. stemming from tetrahedra that are near the medial axis, lead to erroneous contact points. We use a simple filtering heuristic and remove all points that lie within any other sphere. We make use of the fact that the part of a sphere that lies outside all other spheres can be efficiently calculated based on the neighborhood in the regular triangulation (V,F) [KA25a]. For each $\mathbf{p_i} \in V$, we iterate over its neighbors in the regular triangulation and discard all contact points that lie within any of the adjacent spheres. The remaining samples consistently lie near the surface.

In practice, many samples get discarded. In order to increase the amount of samples contributing with a contact point, we also evaluate using additional contact point directions based on the surrounding face normals at a vertex. If the contact point $\tilde{\mathbf{c}}_i$ has been discarded when generated with the approximate vertex gradient $\nabla\tilde{\Phi}_i$, we repeat the construction with each of the adjacent face gradients ∇f_j until all have been tested or a valid contact point is found.

4.3. Algorithm

Our algorithm takes as input a set of sample position and SDF values $\{(\mathbf{p}_i, \Phi_i)\}$ and outputs a triangle mesh. An intermediate output is a set of contact points. We summarize the algorithm in Listing 1 and visualize the individual steps in Figure 1. Time-wise, the limiting factor is the construction of the regular triangulation. We make use of the optimized algorithms of the CGAL library [JPT25], in practice the construction is linear.

5. Results

The real set of contact points $\mathbf{c}_i$ (see Equation (2)) lies at the core of a discrete SDF, but has not been discussed in previous publications. Even though not available in the common reconstruction setting, this set of actual surface samples yields a good upper bound of the limit of what we can expect of a contact point based approach. We refer to the these contact points $\mathbf{c}_i$ and reconstructions based on them as "GTCP" or "Ground truth contact points" In contrast to this, we call our algorithm based on the approximate gradients and contact points $\tilde{\mathbf{c}}_i$ "GCP"or "Gradient contact points".

We first evaluate different choices of weights for the gradient approximation in Section 5.1. When talking about generated contact points in Section 5.2, we also compare to the "Reach for the arcs" (RFTA) method of Sellán et al. [SRBS24]. In Section 5.3, when comparing contouring solutions, we also work with the "Maximal empty spheres" (MES) method of Kohlbrenner and Alexa. As baselines, we use marching cubes (MC), and marching tetrahedra (MT) based on a Delaunay triangulation (DT) as well as the regular triangulation (RT) for SDF samples introduced in [KA25a].

A factor that influences the evaluation is the location of the surface in the bounding box, in which the SDF samples are drawn. The effective resolution of a N^3 regular grid depends on its placement with respect to the bounding box. We use regular grids or random samples over the $[-1,1]^3$ cube and scale the ground truth meshes to lie within the $[-0.9,0.9]^3$ cube if not noted otherwise.

We use a couple of common meshes, as well as the THINGI10K dataset [ZJ16], remeshed using TETWILD [HZG*18].

5.1. Gradient approximation

We evaluate the three sets of weights for vertex-based gradient approximation described in Section 4.1. In addition to the obvious choice of the regular triangulation as the simplicial complex, we also evaluate the Delaunay triangulation. In addition to this, we also evaluate a k nearest-neighbors (KNN) approach: the gradient of the linear function that interpolates the SDF function at the vertex and has the least squared error to the function values at its k nearest neighbors. The values of k are chosen to lead to symmetric neighborhoods on samples distributed on a regular grid.
The better the approximated gradient, the fewer points are discarded for lying in adjacent spheres, we therefore evaluate the total number of valid points $|\{\tilde{\mathbf{c}}_i\}|$ (more is better). A perfect approximation of the vertex gradient leads to a contact point that lies on the surface, where the SDF gradient aligns with the vertex normal. We measure the quality of valid contact points and normals $(\tilde{\mathbf{c}}_i, \tilde{\mathbf{n}}_i)$ using its SDF value $|\Phi(\tilde{\mathbf{c}}_\mathbf{i})|$ (smaller is better) and the alignment of its normal with the SDF gradient $\tilde{\mathbf{n}}_i \cdot \nabla\Phi(\tilde{\mathbf{c}}_\mathbf{i})$ (larger is better). We show results averaged over a subset of THINGI10K in Figure 3. We observe only minimal differences between weights based on volumes and uniform weights. In comparison, the method described by James (DJ) leads to considerably more contact points, that generally have smaller SDF values and better aligned gradients.

The approach described at the end of Section 4.2, in which we try further gradient directions for discarded samples is marked by +FN in Figure 12. While it leads to considerably more valid contact

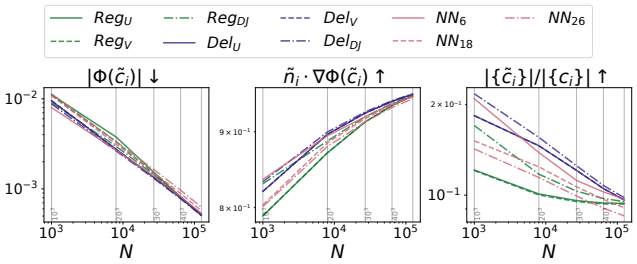

Figure 3: *Contact point quality metrics for different gradient approximation methods, averaged over* 490 *random meshes taken from the* THINGI10K *dataset. We compare different choices of normal estimation based on regular / Delaunay triangulations (Reg/Del) and different aggregation weights (U/V/DJ) with nearest neighbor approaches (*$k \in \{6, 18, 26\}$*).*

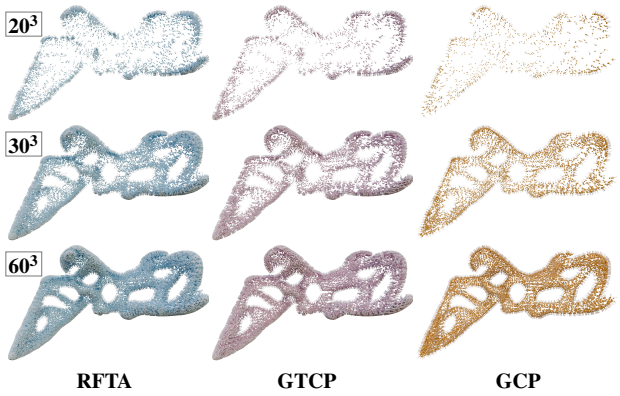

Figure 4: *Contact points for RFTA (left), GTCP (center) and GCP (right) on the Nightingale model at different resolutions.*

points, these points have larger SDF values and their directions are generally less aligned. We do not further pursue this direction in the rest of this work and all other (GCP) results in this section are using the James weights (DJ) only. If not otherwise noted, we use the gradient approximation based on the regular triangulation which we require for the filtering anyways.

5.2. Contact points

We show approximated contact points generated by our method, together with the true contact points and the ones generated by RFTA on the same set of samples in Figure 4. Hidden vertices in the regular triangulation and the discarding of contact points falling in neighboring spheres leads to our method producing fewer samples. However our samples yield a generally precise and even sampling of the surface. In convex regions of the shape, we observe a regularity in the GTCP normal directions, a result of the SDF samples lying on a regular grid.

5.3. Contouring

To measure the accuracy of a contour represented as a surface mesh $M = (V, F)$, we use the distance to the ground truth mesh $M_{GT} = (V_{GT}, F_{GT})$ as:

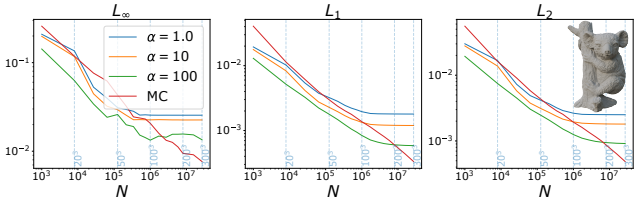

Figure 5: *Distance metrics for sPSR reconstructions from ground truth contact points (GTCP) using different choices of screening weight* α *as well as marching cubes (MC).*

$$L_\infty(M, M_{GT}) = \max\left(\max_{x\in M}\min_{y\in M_{GT}}\|x-y\|, \max_{x\in M_{GT}}\min_{y\in M}\|x-y\|\right)$$

$$L_1(M, M_{GT}) = \tfrac{1}{2}\left(\int_M \min_{y\in M_{GT}}\|x-y\|dx + \int_{M_{GT}}\min_{y\in M}\|x-y\|dx\right)$$

$$L_2(M, M_{GT}) = \tfrac{1}{2}\left(\sqrt{\int_M \min_{y\in M_{GT}}\|x-y\|^2 dx} + \sqrt{\int_{M_{GT}}\min_{y\in M}\|x-y\|^2 dx}\right)$$

In practice, we approximate the integrals based on finite samples on subdivided meshes.

We use the same implementation and parameters of poisson surface reconstruction (sPSR) for all contact point based approaches. While higher screening weight α can lead to smaller mesh distances, we observe that this can come at the cost of surface artifacts in areas where the resolution of the contact points is not high enough. Figure 5 shows the reconstruction accuracy of the baseline GTCP for different choices of α. If not otherwise noted, we choose a screening weight $\alpha = 1.$ for all methods.

5.3.1. Comparison: marching tetrahedra-based approaches

If connectivity of the sample points is unknown and possibly irregular, contouring commonly requires the computation of a triangulation anyways, the simplest being the (weighted) Delaunay triangulation. Kohlbrenner and Alexa have shown that computing a specific regular triangulation of the samples can generally be beneficial in such settings. Our approach can extend this with little overhead since it is also based on a regular triangulation. We show the timings together with reconstructions and error metrics on random (Figure 6) and regular samples (Figure 7) of the Armadillo model. GCP consistently reaches lower errors and visually more correct reconstructions, as seen on the ears of the Armadillo. We can follow the argument made in [KA25a] for the regular triangulation and state that if the data is unstructured, one might as well construct the regular triangulation and use it not only for marching tets but use GCP to reach higher detail at little additional cost.

5.3.2. Marching cubes

Like other contact point based methods, GCP consistently outperforms marching cubes at lower resolutions Figures 11 and 14 to 17. Its simplicity makes it the first contact point based approach that scales to higher resolutions, we show some results and metrics in Figure 9. The reconstruction seem to have a higher level of detail and GCP reaches superior metrics in an input range that is comparable with that of the real contact points in Figure 5.

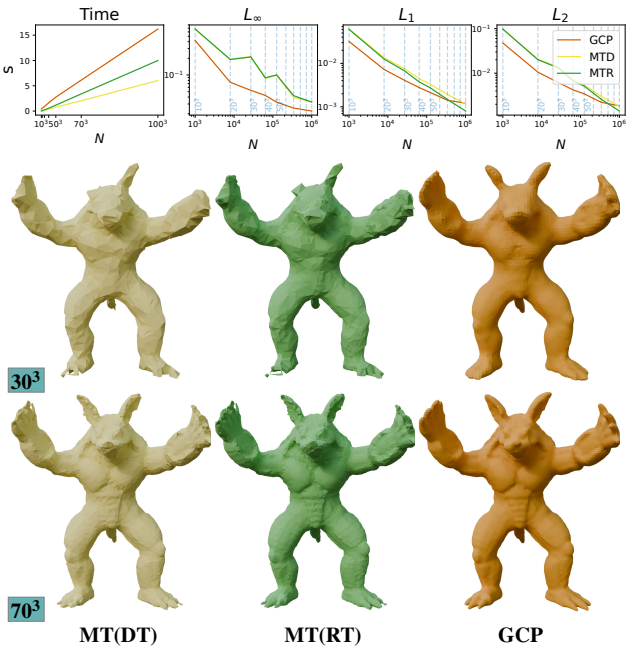

Figure 6: *Comparison to marching tetrahedra (MT) on the Delaunay (DT, left) and regular (RT, middle) triangulation, samples randomly distributed.*

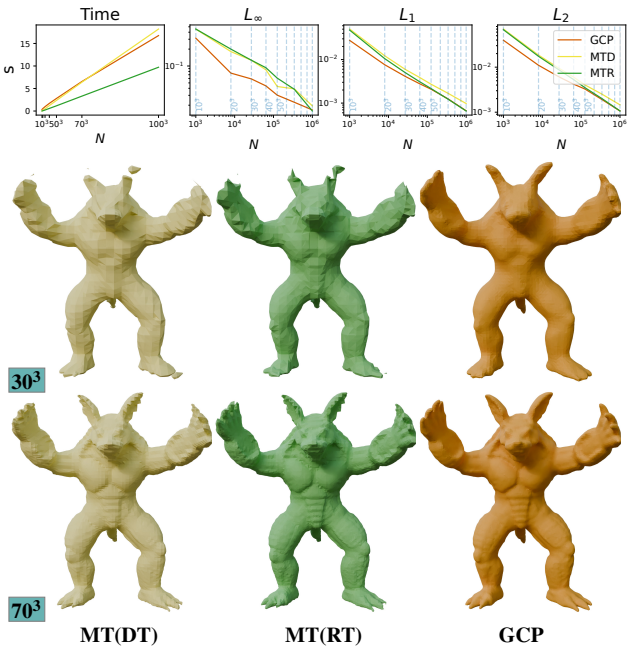

Figure 7: *Comparison to marching tetrahedra (MT) on the Delaunay (DT, left) and regular (RT, middle) triangulation, samples on a regular grid.*

5.3.3. Contact point-based approaches

We compare our method to the two state of the art methods RFTA and MES in Figures 11, 16 and 17. Our method is conceptually much simpler, which leads to a significantly faster runtime Figure 10. Even though the more complex methods generate significantly more contact points (Figure 4), we do not see a clear advantage from the additional computation. In particular, our method is able to capture the surfaces of complicated genus such as the difficult CellularThing model in Figure 11 at comparable resolutions at a fraction of the runtime. Thin structures such as exhibited in the model in Figure 15 can also be recovered at similar resolutions. The chair shown in Figure 17 contains many small details that are very hard to reconstruct at lower resolutions but no method seems clearly superior to its competitors. On surfaces with a lot of detail, MES and RFTA sometimes appear to generate more detail. This is visible e.g. on the koala for higher resolutions in Figure 16. However, this does not seem to translate to a lower error to the ground truth. If we take a closer look at the error distribution over the mesh in Figure 13, some of the more detailed regions even appear to feature a higher error.

We show the average Chamfer error over a subset of meshes from THINGI10K in Figure 8. For consistency with [SRBS24], the ground truth meshes are scaled to lie within the $[-0.5, 0.5]^3$ cube and SDF samples are drawn on a regular grid of $[-1, 1]^3$. While no method consistently outperforms its competitors, our approach is considerably faster.

6. Discussion

Our main contribution is the observation that surface (=contact) points can be generated from samples of a signed distance function in a much simpler fashion than suggested by previous work, without sacrificing quality of the reconstructed surface.

We have demonstrated that the ground truth contact points contain more useful information than previously assumed, and have thereby put emphasis on a class of methods that had previously been dismissed. While the ground truth contact points provide a useful upper limit of the reconstruction quality for this class of methods, we have also found that the results are quite sensitive to the parameters of PSR, in particular for higher resolutions and smaller errors. This not only shows the limits of this approach in general, it also suggests that it may be worthwhile to experiment with other methods for reconstructing a surface from boundary samples, or tailor a method to this specific use case.

Acknowledgements

Funded, in part, by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101055448, ERC Advanced Grand EMERGE). Open Access funding enabled and organized by Projekt DEAL. We thank the authors of the used models for their publication: KOALA [Thu19], ARMADILLO [Sta96], NIGHTINGALE [Emm21], as well as CELLULARTHING (id 61258), CHAIR (id 93366), and SPIRALCUP (id 37358) taken from THINGI10K [ZJ16]. Open Access funding enabled and organized by Projekt DEAL.

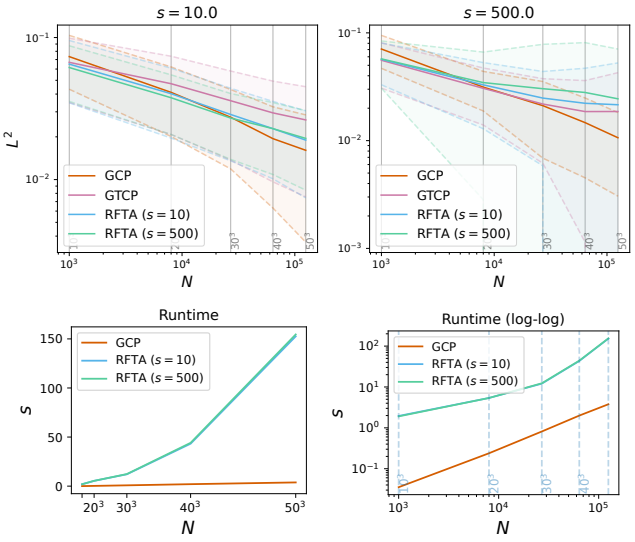

Figure 8: *Chamfer errors (top, log-log) and timings (bottom) on a subset of 100 random meshed shapes from the* THINGI10K *dataset. Solid lines denote meaned quantities, the shaded area between the dashed lines denotes mean ± one standard deviation. The two bottom plots show the same runtime data, using standard (left) and log-log axes (right).*

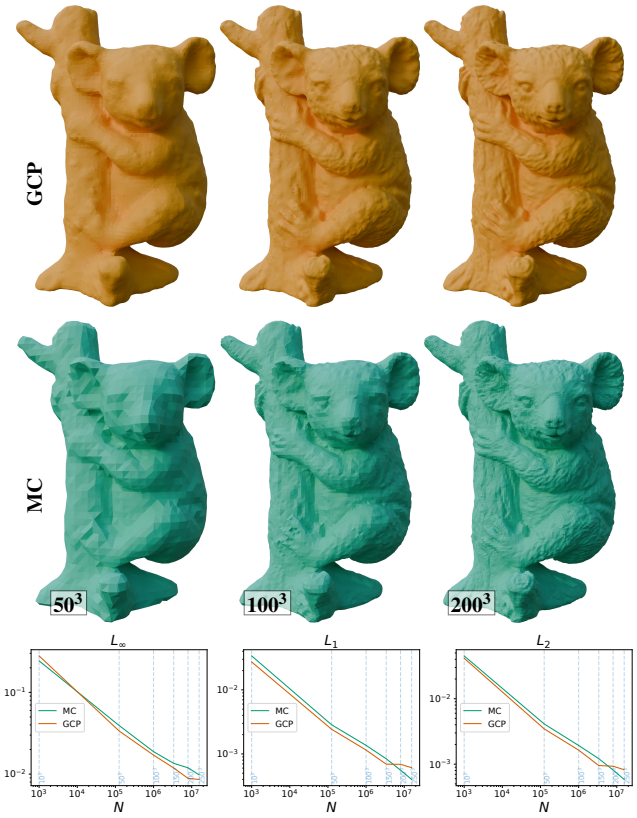

Figure 9: *GCP results up to higher resolutions, compared to the marching cubes reconstruction.*

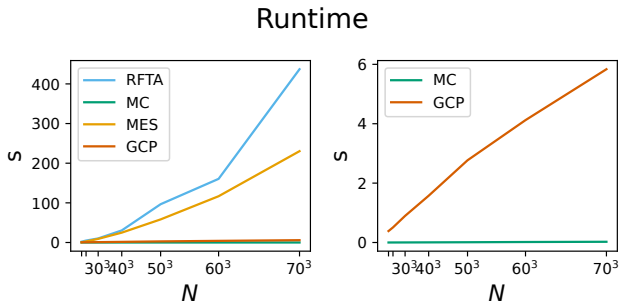

Figure 10: *Total runtime on the* KOALA *model for GCP in comparison to MC (right) and the same values also including the slower contact point based methods that are currently state of the art (left).*

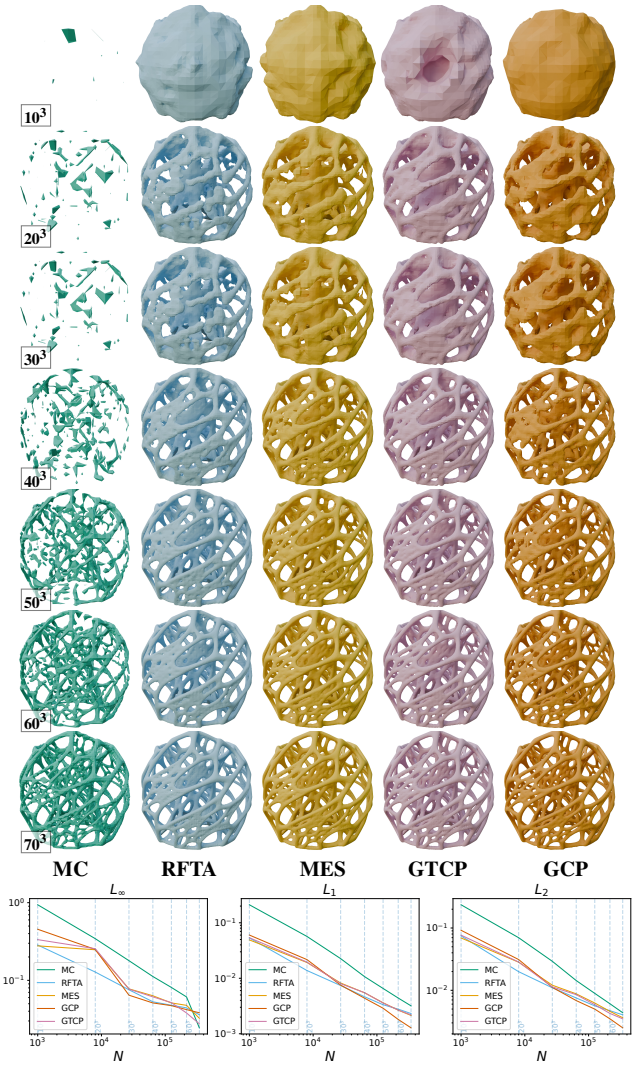

Figure 11: *(top) Contouring results on regular grids of SDF samples of different sizes (rows) and methods (columns) on the* CELLULARTHING *model. (bottom) distance metrics to the ground truth mesh (log-log).*

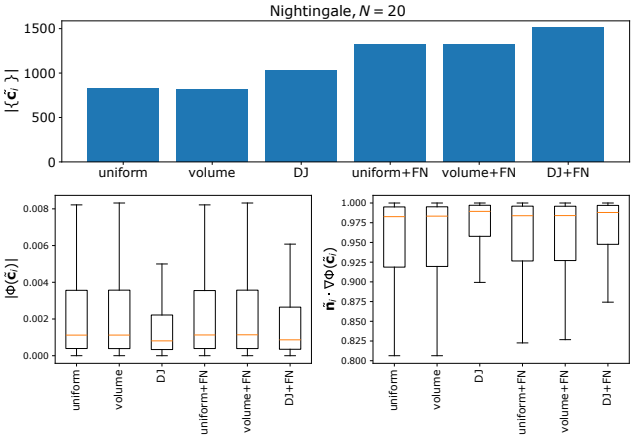

Figure 12: *Number of contact points (top), boxplots of the SDF value (left) and normal alignment with the sdf gradient (right) at the contact points on the* NIGHTINGALE *model ($N = 20$).*

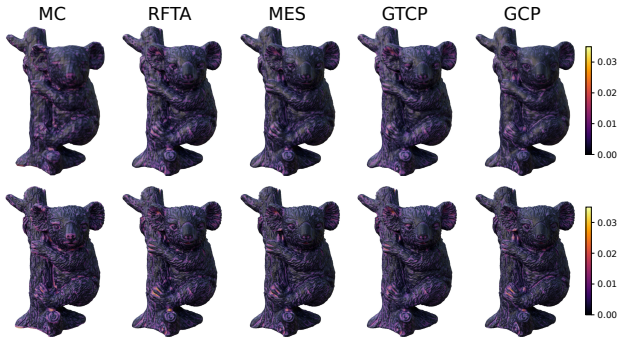

Figure 13: *Distance to the closest point on the ground truth (top row) and to the respective reconstruction from the ground truth (bottom row) for different reconstruction techniques. The SDF samples are on a regular grid of resolution $N = 60^3$. The mean of these values integrated over both meshes is shown as L_1 in Figure 16.*

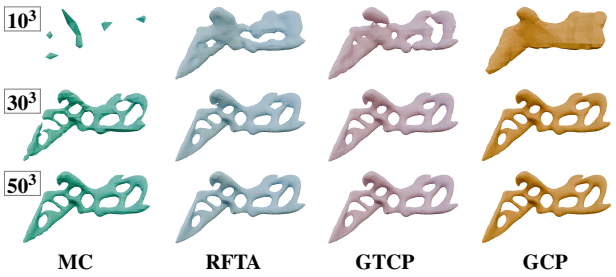

Figure 14: *Contact-point based approaches reconstruct the topology of the* NIGHTINGALE *model faster than marching cubes.*

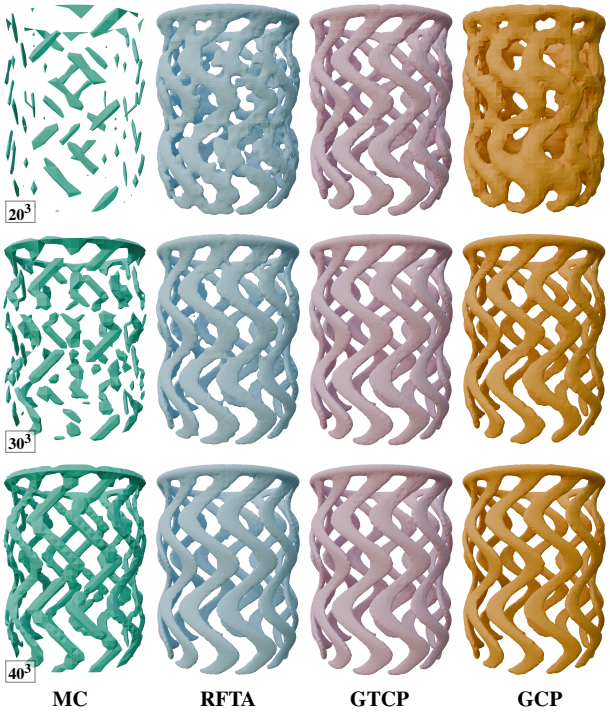

Figure 15: *Thin structures are challenging to reconstruct at lower resolutions, visible on the* SPIRALCUP *model.*

Figure 16: *(top) Contouring results on regular grids of SDF samples of the* KOALA *model. (bottom) distance metrics to the ground truth mesh (log-log).*

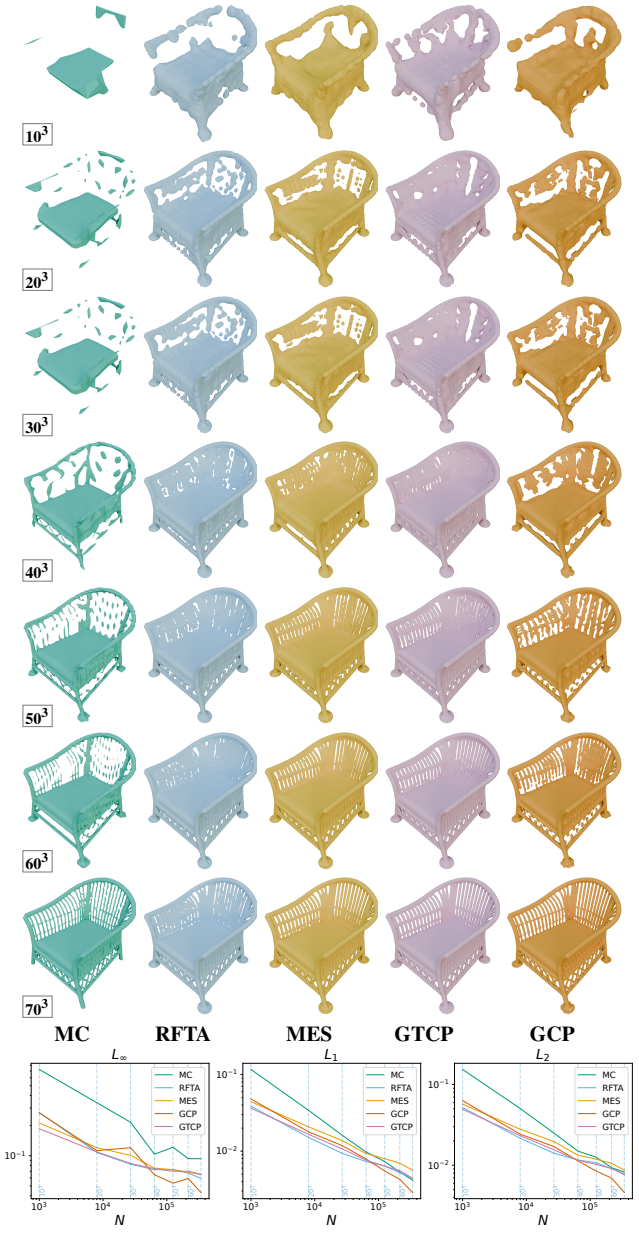

Figure 17: *(top) Contouring results on regular grids of SDF samples of the* CHAIR *model. (bottom) distance metrics to the ground truth mesh (log-log).*

References

[BF95] BLOOMENTHAL J., FERGUSON K.: Polygonization of non-manifold implicit surfaces. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1995), SIGGRAPH '95, Association for Computing Machinery, p. 309–316. doi:10.1145/218380.218462. 2

[Blo88] BLOOMENTHAL J.: Polygonization of implicit surfaces. *Computer Aided Geometric Design 5*, 4 (1988), 341–355. doi:10.1016/0167-8396(88)90013-1. 2

[CTFZ22] CHEN Z., TAGLIASACCHI A., FUNKHOUSER T., ZHANG H.: Neural dual contouring. *ACM Trans. Graph. 41*, 4 (July 2022). doi:10.1145/3528223.3530108. 2

[CZ21] CHEN Z., ZHANG H.: Neural marching cubes. *ACM Trans. Graph. 40*, 6 (Dec. 2021). doi:10.1145/3478513.3480518. 2

[DZCJ22] DU X., ZHOU Q., CARR N., JU T.: Robust computation of implicit surface networks for piecewise linear functions. *ACM Trans. Graph. 41*, 4 (July 2022). doi:10.1145/3528223.3530176. 2

[Emm21] EMM: Stravinsky fountain: the nightingale. https://sketchfab.com/3d-models/stravinsky-fountain-the-nightingale-66589dd1f2b04c2f9d2e828be5241d77, 2021. 6

[Gib98] GIBSON S. F. F.: Using distance maps for accurate surface representation in sampled volumes. In *Proceedings of the 1998 IEEE Symposium on Volume Visualization* (New York, NY, USA, 1998), VVS '98, Association for Computing Machinery, p. 23–30. doi:10.1145/288126.288142. 2

[GZ22] GROSSO R., ZINT D.: A parallel dual marching cubes approach to quad only surface reconstruction. *Vis. Comput. 38*, 4 (Apr. 2022), 1301–1316. doi:10.1007/s00371-021-02139-w. 2

[Har96] HART J. C.: Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer 12* (1996), 527–545. URL: https://api.semanticscholar.org/CorpusID:20278. 1

[HZG*18] HU Y., ZHOU Q., GAO X., JACOBSON A., ZORIN D., PANOZZO D.: Tetrahedral meshing in the wild. *ACM Trans. Graph. 37*, 4 (July 2018), 60:1–60:14. doi:10.1145/3197517.3201353. 4

[Jam20] JAMES D. L.: Phong deformation: a better c0 interpolant for embedded deformation. *ACM Trans. Graph. 39*, 4 (Aug. 2020). doi:10.1145/3386569.3392371. 4

[JDZ*24] JU Y., DU X., ZHOU Q., CARR N., JU T.: Adaptive grid generation for discretizing implicit complexes. *ACM Trans. Graph. 43*, 4 (July 2024). doi:10.1145/3658215. 2

[JLSW02] JU T., LOSASSO F., SCHAEFER S., WARREN J.: Dual contouring of hermite data. *ACM Trans. Graph. 21*, 3 (July 2002), 339–346. URL: https://doi.org/10.1145/566654.566586, doi:10.1145/566654.566586. 2

[JPT25] JAMIN C., PION S., TEILLAUD M.: 3D triangulations. In *CGAL User and Reference Manual*, 6.0.2 ed. CGAL Editorial Board, 2025. URL: https://doc.cgal.org/6.0.2/Manual/packages.html#PkgTriangulation3. 4

[KA25a] KOHLBRENNER M., ALEXA M.: Isosurface Extraction for Signed Distance Functions using Power Diagrams. *Computer Graphics Forum* (2025). 2, 3, 4, 5

[KA25b] KOHLBRENNER M., ALEXA M.: A polyhedral construction of empty spheres in discrete distance fields. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (New York, NY, USA, 2025), SIGGRAPH Conference Papers '25, Association for Computing Machinery. doi:10.1145/3721238.3730748. 1, 2, 3

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2006), SGP '06, Eurographics Association, pp. 61–70. URL: http://dl.acm.org/citation.cfm?id=1281957.1281965. 2

[KBSS01] KOBBELT L. P., BOTSCH M., SCHWANECKE U., SEIDEL H.-P.: Feature sensitive surface extraction from volume data. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, Association for Computing Machinery, p. 57–66. doi:10.1145/383259.383265. 2

[KH13] KAZHDAN M., HOPPE H.: Screened poisson surface reconstruction. *ACM Trans. Graph. 32*, 3 (July 2013). doi:10.1145/2487228.2487237. 2, 3

[LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph. 21*, 4 (Aug. 1987), 163–169. doi:10.1145/37402.37422. 2

[MOAD24] MARUANI N., OVSJANIKOV M., ALLIEZ P., DESBRUN M.: Ponq: a neural qem-based mesh representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2024), pp. 3647–3657. 2

[Nie04] NIELSON G. M.: Dual marching cubes. In *Proceedings of the Conference on Visualization '04* (USA, 2004), VIS '04, IEEE Computer Society, p. 489–496. doi:10.1109/VISUAL.2004.28. 2

[SBS23] SELLÁN S., BATTY C., STEIN O.: Reach for the spheres: Tangency-aware surface reconstruction of sdfs. In *SIGGRAPH Asia 2023 Conference Papers* (New York, NY, USA, 2023), Association for Computing Machinery. doi:10.1145/3610548.3618196. 1, 2, 3

[SJW07] SCHAEFER S., JU T., WARREN J.: Manifold dual contouring. *IEEE Transactions on Visualization and Computer Graphics 13*, 3 (May 2007), 610–619. doi:10.1109/TVCG.2007.1012. 2

[SRBS24] SELLÁN S., REN Y., BATTY C., STEIN O.: Reach for the arcs: Reconstructing surfaces from sdfs via tangent points. In *ACM SIGGRAPH 2024 Conference Papers* (New York, NY, USA, 2024), Association for Computing Machinery. doi:10.1145/3641519.3657419. 1, 2, 3, 4, 6

[SSSC22] SOMMER C., SANG L., SCHUBERT D., CREMERS D.: Gradient-sdf: A semi-implicit surface representation for 3d reconstruction. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), pp. 6270–6279. doi:10.1109/CVPR52688.2022.00618. 3

[Sta96] STANFORD: The stanford 3d scanning repository: Armadillo man. https://graphics.stanford.edu/data/3Dscanrep/, 1996. 6

[SW04] SCHAEFER S., WARREN J.: Dual marching cubes: Primal contouring of dual grids. In *Proceedings of the Computer Graphics and Applications, 12th Pacific Conference* (USA, 2004), PG '04, IEEE Computer Society, p. 70–76. 2

[Thu19] THUNK3D: koala bear. https://sketchfab.com/3d-models/koala-bear-221d8d6519944a65b473ea56fc032570, 2019. 6

[WZW*25] WANG P., ZHANG Z., WANG W., CHEN S., LU L., XIN S., TU C.: Power diagram enhanced adaptive isosurface extraction from signed distance fields, 2025. URL: https://arxiv.org/abs/2506.09579, arXiv:2506.09579. 2, 3

[ZGG22] ZINT D., GROSSO R., GÜRTLER P.: Resolving Non-Manifoldness on Meshes from Dual Marching Cubes . In *Eurographics 2022 - Short Papers* (2022), Pelechano N., Vanderhaeghe D., (Eds.), The Eurographics Association. doi:10.2312/egs.20221029. 2

[ZJ16] ZHOU Q., JACOBSON A.: Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797* (2016). URL: https://arxiv.org/abs/1605.04797. 4, 6