DOI: 10.1111/cgf.70037
EUROGRAPHICS 2025 / A. Bousseau and A. Dai
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 44 (2025), Number 2

Isosurface Extraction for Signed Distance Functions using Power Diagrams

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

[1]TU Berlin, Computer Graphics Group

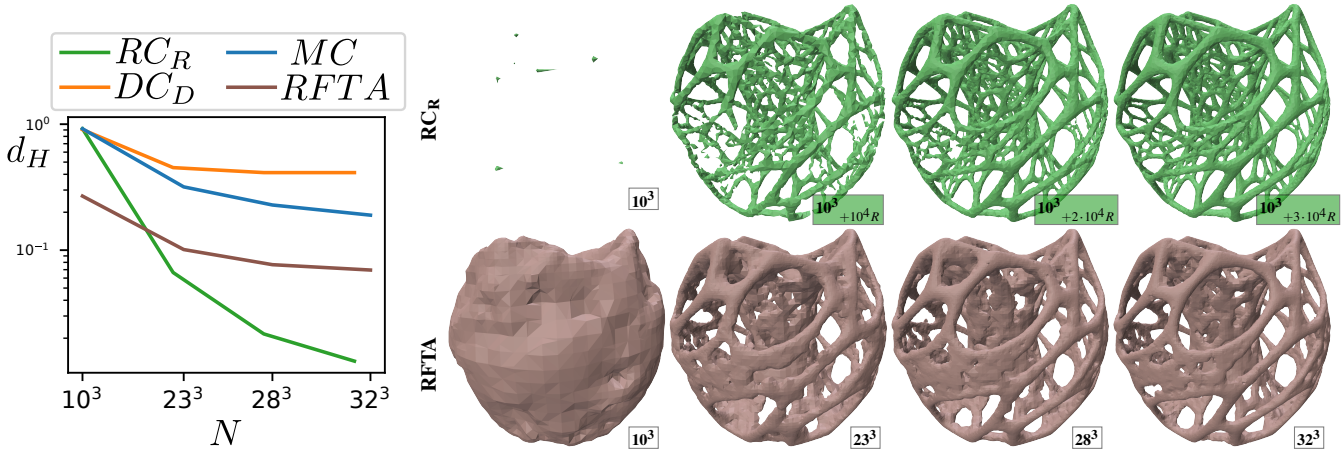

Figure 1: *Reconstructions using the regular contour with regular refinement of samples on a regular grid (RC_R, right, top row) and "Reach for the Arcs" (RFTA) (right, bottom row) on a comparable amount of samples on a grid as well as Hausdorff distances to the ground truth for different approaches (left). Boxes specify the amount of grid samples and additional steps of regular (R) refinement.*

Abstract

Contouring an implicit function typically considers function values in the vicinity of the desired level set, only. In a recent string of works, Sellán at al. have demonstrated that signed distance values contain useful information also if they are further away from the surface. This can be exploited to increase the resolution and amount of detail in surface reconstruction from signed distance values. We argue that the right tool for this analysis is a regular triangulation of the distance samples, with the weights chosen based on the distance values. The resulting triangulation is better suited for reconstructing the surface than a standard Delaunay triangulation of the samples. Moreover, the dual power diagram encodes the envelope enclosing the surface, consisting of spherical caps. We discuss how this information can be exploited for reconstructing the surface. In particular, the approach based on regular triangulations lends itself well to refining the sample set. Refining the sample set based on the power diagram outperforms other reconstruction methods relative to the sample count.

Keywords: Signed Distance Function, Contouring, Power Diagram, Weighted Delaunay Triangulation

CCS Concepts

• **Computing methodologies** → **Shape analysis;** • **Mathematics of computing** → *Mesh generation;*

1. Introduction

A common task in geometry processing is the conversion of a surface represented by its *signed distance function* (SDF) into an explicit representation such as a triangle mesh – a process referred to as *contouring*. In procedural modeling the SDF may be given as a function, allowing to query the signed distance to the surface from arbitrary locations. Sometimes only the result of such evaluation is given, for example distances sample on a discrete set of samples.

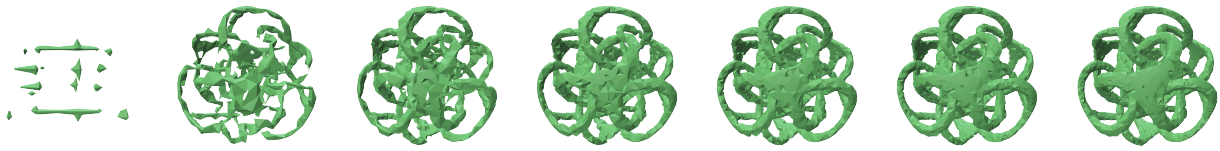

Figure 2: *Regular refinement from* 1*k to* 7*k SDF samples of the* METATRON *mesh.*

In either case, one may ask how to efficiently generate a mesh that closely approximates the true surface.

The 'classic' way to generate a surface approximation are contouring approaches based on a cell decomposition of the space based on the samples. If samples are located on a regular grid, one may create a surface approximation in each grid cell whose vertices are not entirely on one side of the surface [LC87]. More generally, the samples could be considered the vertices of a complex, and the surface is locally approximated for each edge that is incident on samples with opposing signs.

What is common among these approaches is that the reconstructed surface is governed by the samples that are close to the surface, i.e. whose distance values are small (in absolute value). Sellán et al. [SBS23] have recently exploited the observation that, in fact, all distance samples contain information. Every distance sample represents a ball that is either entirely inside or outside of the volume enclosed by the surface. Moreover, each such ball touches the surface and, generically, the region around this tangent point is covered by no other ball associated to the samples.

Our main contribution is the observation that information is lost because of an inappropriate choice of grid or triangulation, typically based on the proximity among the samples (for example a regular grid or the Delaunay triangulations). We suggest to use the available distance values to generate a *regular triangulation*, with the weights derived from the distance values. The dual of this triangulation, the *power diagram*, encodes the information about the spherical regions that are not covered by any other ball (see Section 4.4). For example, the faces separating power cells corresponding to inside samples from outside samples is a discrete approximation of the contour that is guaranteed to be not intersecting any of the empty balls – readily providing a surface reconstruction that can be extracted from the surface samples without the need for any optimization.

Beyond discussing the general properties of the regular triangulation and power diagram in relation to the surface represented by the SDF, we also exploit the properties for an improved surface reconstruction algorithm. Two unique features of the representation based on the power diagram are that (1) it is easy to incorporate information coming from additional surface samples and (2) it is clear where the reconstruction would benefit from more information. This leads to a refinement approach (see Section 6). The resulting surface reconstructions systematically outperform optimization of surface samples for Poisson reconstruction [SRBS24] as well as local refinement approaches based on sign changes of the SDF. We briefly discuss other possible ways to exploit the power diagram in Section 7 but leave their implementation and evaluation for future work.

2. Related Work

Implicit Representations and Contouring One of the fundamental representations of geometry, implicit surfaces play a key role in many tasks in computer graphics and geometry processing. Implicit representations model geometry as the level set of a function and are the basis for many state-of-the-art algorithms in surface reconstruction. A wide range of applications reaches from classic PDE-based methods [KBH06, KH13], operating on point clouds, to recent volumetric scene representations [MST*21, KKLD23]. An important class of implicit functions are signed distance fields (SDFs) that are often employed in physically inspired settings such as collision detection [MEM*20] and geometric modeling [MBWB02]. Robust estimates of signed distance fields are still an ongoing research direction [FC24]. Physical constraints also find their way into recent neural methods via SDFs and the Eikonal equation [PFS*19, SMB*20]. Efforts are made to combine the physical restrictions of signed distance fields with intuitive editing techniques [MSLJ23].

While some use cases directly render the implicit function, e.g. using ray-tracing [Har96, VSJ22], others require a transformation of representation. The problem of contouring, i.e., transferring an implicit representation to an explicit one, most commonly a polygon mesh, is widely studied. For a detailed overview, we refer to the survey by de Araújo [dALJ*15]. Many well-known approaches are based on partitioning the ambient space into volumetric primitives and constructing the surface locally per primitive. Primitives could be cubes/squares [LC87] or simplices. The local construction can be based on the primal primitives [LC87] or a dual complex [JLSW02, KBSS01]. Recently, several Neural contouring approaches have emerged [CZ21, CTFZ22]. The interest in specifically contouring signed distance fields has lately been renewed in a line of work by Sellán et al. [SBS23, SRBS24]. The authors note that most contouring approaches fail to use all of the available data, ignoring information encoded by samples further away from the surface. We will review the geometric problem formulation of their work in detail in the next section.

Simplicial Meshes Besides regular or semi-regular structures like grids or octrees, irregular structures such as simplicial meshes are commonly used in order to discretize a domain. The flexibility that is gained by this representation comes at the price of irregular combinatorics. Different geometric properties can be desirable depending on the use case [She02]. Many methods exist to optimize the quality of mesh elements, some are based on flipping simplices or moving existing vertices [CX04, Che04, FAB*18, Ale19]. An important class of such methods relies on refining an existing complex by insertion of new vertices [She98]. A useful class of triangulations that contains and extends the common Delaunay triangulation

are *regular* or *weighted* Delaunay triangulations [AK00, AKL13]. A common application is the well-known sliver-exudation algorithm [CDE*00]. We will discuss the theory of regular triangulation more in detail in the next section.

3. Background

3.1. Regular Triangulations / Power Diagrams

Given a point set, Voronoi diagrams are a tessellation of the domain into regions that are closest to each of the points. They are dual to the Delaunay triangulation. This concept can be generalized to other notions of 'closest'. While additively or multiplicatively weighted distance functions lead to more complex diagrams with curved boundaries among the cells [AB86], a convenient generalization is based on *power distances*, leading to power diagrams and their dual *regular* triangulation (also called *weighted* Delaunay triangulation). We briefly review this idea below, but have to refer the reader to the literature for more details [AK00, AKL13, DLRS10] and the rich connection to the theory of polytopes [Zie95, Grü03].

Consider a set $P = \{\mathbf{p_1}, ..., \mathbf{p_n}\}$ of points carrying weights $\mathbf{p}_i = (p_i, w_i) \in \mathbb{R}^d \times \mathbb{R}$. The *power distance* between two weighted points is defined as:

$$\Pi(\mathbf{p_i}, \mathbf{p_j}) = |p_i - p_j|^2 - w_i - w_j \tag{1}$$

Note that the power distance is not affected by adding a constant weight to all points. This means we might as well consider all points to have positive weights. In this case we can think of w_i as the squared radius $r_i = \sqrt{w_i}$ of a sphere associated with points. Moreover, the weights could be increased so that the spheres are intersecting. The intersection of the spheres is orthogonal if the power distance of the points is zero.

For each d-simplex $t = (\mathbf{p}_1, ..., \mathbf{p}_{d+1})$, there exists a weighted point $\mathbf{c}_t = (c_t, w_t)$, commonly referred to as its *power sphere*, intersecting the spheres of the weighted points in t orthogonally, i.e. $\forall \mathbf{p}_i \in t : \Pi(\mathbf{c}_t, \mathbf{p_i}) = 0$. If the weights for all points in t are identical (or, equivalently, zero), c_t corresponds to the center of the circumscribed sphere of the t (and w_t is the square of its radius). A weighted point $\mathbf{q}$ is called *regular* if $\forall \mathbf{p} : \Pi(\mathbf{p}, \mathbf{q}) \geq 0$. A triangulation is said to be regular if the power spheres of all its simplices are regular. Points without weight can be treated as having a weight of 0. If all points have zero weight (more generally, the same weight), the regular triangulation is identical to the Delaunay triangulation.

The power diagram of a set of weighted points is dual to its regular triangulation and consists of power cells $\mathcal{C}_i = \{q \in \mathbb{R}^d : \Pi(\mathbf{p}_i, (q,0)) \leq \Pi(\mathbf{p}_j, (q,0)), \mathbf{p}_j \in \mathcal{P}, i \neq j\}$. Again, in the case of points with zero weight $w_i = 0, \forall i$, the power distance is the standard squared distance and the power diagram is equivalent to the Voronoi diagram, the dual of the Delaunay triangulation. A point p_i is part of the regular triangulation only if its power cell is not empty. Otherwise the point is *hidden*.

The practical relevance of regular triangulations and power diagrams, apart from their mathematical properties, stems from the fact that they can be computed in the exact same way as standard Delaunay triangulations and Voronoi diagrams [ES92, CDS12]. The central ingredient for the construction is a *predicate* deciding if a point is inside or outside of the sphere defined by $d+1$ points. This is done robustly by evaluating the determinant of a $(d+2) \times (d+2)$ matrix [She96] – and the matrix can be set up for weighted points in the same way as for unweighted points. Moreover, regular triangulations facilitate *point location* by walking the triangulation [Ede90, DPT01], in the same way as Delaunay triangulations.

The power diagram can be extracted from the regular triangulation using the bisector planes dual to edges of the triangulation as face planes. The vertices of the (convex) power cells are the power spheres of the simplices in the regular triangulations. Most implementations for the computation of the regular triangulations readily provide ways to extract the elements of the dual power cells.

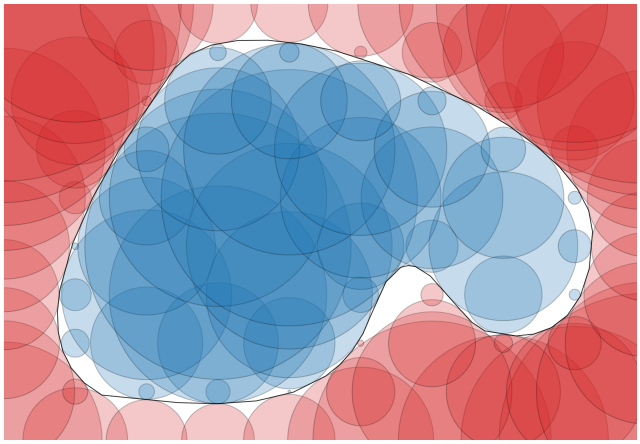

Figure 3: *Surface S and samples of its SDF on a regular grid. Distance values visualized as circles, shaded according to the sign.*

3.2. SDF Contouring

We employ the geometric view on SDF contouring by Sellán et al. [SBS23], who credit Batty [Bat11] and Kobbelt et al. [KBSS01] for mentioning it previously. Consider a surface $S = \partial\Omega$ bounding the volume Ω. At a point x, the signed distance function can be defined as $\phi_S(x) = \chi_\Omega(x) \min_{p \in S} |x - p|$ with χ_Ω the indicator function returning -1 for points in Ω and 1 elsewhere. We consider a finite set of samples $\mathcal{P} = \{(p_i, s_i) \in \mathbb{R}^d \times \mathbb{R} : s_i = \phi_S(p_i)), i = 1, ..., n\}$ from the SDF of an (unknown) surface S. Geometrically, each (p_i, s_i) defines a signed sphere (circle in $2D$) such that (cf. [SBS23]):

1. The sphere center lies on the side of S indicated by the sign.
2. S does not intersect the interior of the sphere.
3. S is tangent to the sphere in at least one point.

An example of a surface S and SDF samples $\mathcal{P}$ is shown in Fig. 3.

By (1) and (2), the interiors of two spheres with opposite sign are always separated by S and are not intersecting in their interiors. It follows that also the unions of all spheres with the same sign are not intersecting the interior of the union of the spheres with the opposite sign. All surfaces consistent with (1) and (2) are not intersecting the interiors of these unions, so they lie in the region

between the two unions, which is referred to as the 'feasible region' [SRBS24]. The boundary of this region consists of spherical caps (resp. arcs in 2D), which are referred to as 'feasible arcs'. A surface satisfying the tangency constraint (3) has to be tangent to each sphere. While this restricts the set of possible surfaces, there is generally more than one surface satisfying all constraints implied by a set of samples $\mathcal{P}$. Parameterizing the space of admissible surfaces is non-trivial and reconstructing a surface consequently still requires additional regularization.

General contouring algorithms [LC87, JLSW02] reconstruct level sets of implicit functions and satisfy (1) but generally violate (2) and (3), failing to exploit the complete information carried by the samples.

Reach for the Spheres Sellán et al. [SBS23] take the geometric view, recognizing the difficulty of parameterizing the possible solution. They suggest an energy

$$\mathcal{E}_\phi(\Omega) = \frac{1}{2}\sum_{i=1}^{n}(\phi_S(p_i) - s_i)^2 \tag{2}$$

that encourages the surface to satisfy constraints (2) and (3). Starting from an arbitrary initial mesh, the energy is minimized by following its negative gradient. We refer to this method as RFTS. If the initial mesh is computed using standard contouring techniques such as marching cubes, condition (1) is satisfied for the initial mesh, and the optimization is expected to preserve this property. While RFTS often leads to greatly improved results, no guarantees are available.

Reach for the arcs In follow-up work, Sellán et al. aim at exploiting that the tangency constraints (3) need to be satisfied in a point on the feasible arc (and not an arbitrary point on the sphere) [SRBS24]. The strategy for using this observation is to generate point samples on the feasible arcs and then use Poisson surface reconstruction [KBH06] to generate the surface. We refer to this approach as RFTA. The main problem is finding good points samples, as not *every* point in a feasible region is close to the tangent point. Clever choices yield a robust and visually impressive SDF reconstruction algorithm, without necessarily satisfying the conditions (1-3).

4. A Regular Triangulation for SDF Samples

We begin with a set of points and signed distance values $\mathcal{P} = \{(p_i, s_i) \in \mathbb{R}^d \times \mathbb{R} : s_i = \phi_S(p_i)), i = 1, ..., n\}$. For each point, we define the weighted point $\tilde{\mathbf{p}}_i = (p_i, s_i^2)$ and store its sidedness as $o_i = \pm 1$. We want to emphasize that all radii are squared and therefore non-negative. Whenever we are talking of positive/negative spheres, we are referring to the sign stored in o_i – it is not encoded in the weights or the resulting triangulation.

The main aim of this section is to discuss the properties of the regular triangulation and dual power diagram of this weighted point set (including the additional sign information). As we will see, these properties are very useful for the purpose of contouring the SDF.

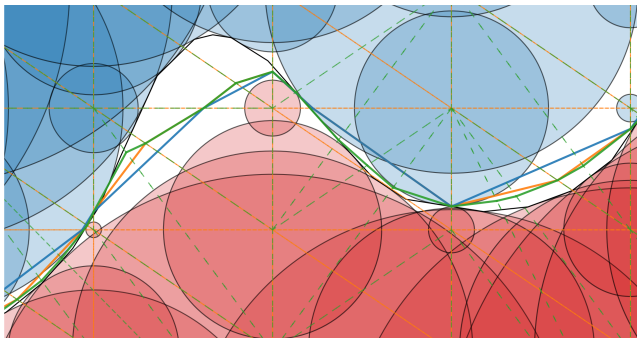

Figure 4: *Marching squares (blue) and the marching triangles contour on an arbitrary (Delaunay) triangulation (orange) and the regular triangulation of the weighted points (green). Contours are solid lines, edges of the triangulations dashed.*

4.1. Power distance

In contrast to the unweighted (Delaunay) case, the regular triangulation is based on the (power-)distances to spheres as opposed to the sample points. Similar to the (signed) L^1 and L^2 distance to the sphere, the power distance $\Pi_{\tilde{\mathbf{p}}}(x) = |p - x|^2 - s^2$ is equal to zero on the sphere, smaller than zero inside, and larger than zero outside. However, it behaves differently from a proper distance function and, in particular, its gradient at the zero level set depends on the radius, the power distance increases faster when moving away from a circle with larger radius.

4.2. Power Diagram

The power diagram partitions space into cells, each associated to a generator vertex $\tilde{\mathbf{p}}_i$. All points in a cell are closer to their generator than to any other in power distance. In contrast to Voronoi cells, a sample may not be contained in its own power cell. This is also true in our setting and, in fact, happens regularly. The boundary facets of the power cells are given by the set of points that share the same two closest spheres. The vertices of the power diagram are located where $d+1$ generators are the closest and their weight is defined such that they are orthogonal to all of them. The weight therefore corresponds to the closest distance value at the location. As discussed in the previous section, all generator weights are positive and the sign of the power distance encodes whether a point is inside or outside of a sphere. This has an important and useful consequence:

Observation 1 A power vertex is located in the feasible region if (and only if) its associated weight is positive.

Fig. 5 shows the power cells shaded according to the sign of their generator, and the boundary facets in green. As we can see, power cells may intersect many spheres, yet only of the same sign. We visualize the weights of a set of power vertices in the feasible region in Fig. 13.

4.3. Regular Triangulation

A generator is a vertex of the regular triangulation (not hidden) if its associated power cell is not empty. We recall that each SDF

sphere shares at least one tangent point with the surface. This point is generically unique for smooth surfaces, i.e. not shared with another SDF sample. This implies that the corresponding power cell, which has to contain this point, is not empty and, consequently, the point cannot be hidden. Different implementations might handle the edge case of a power cell collapsing to a single point differently but this has no effect on the applications discussed here.

Two vertices are connected in the regular triangulation if their power cells are adjacent, meaning they share a point from which both are the closest vertices. As opposed to the Delaunay triangulation, this is measured in power distance and takes into account the radii of the spheres. For samples with opposing signs this favors pairs of vertices with spheres that are close to tangential. As shown in Fig. 4, this leads to triangles reliably connecting vertices of opposite sign that leave the least space between them in densely sample areas. While the power distance is reliable at the spheres, it increases slower for small radii, which can lead to small radii being favored in sparsely sampled areas.

4.4. Feasible region and arcs

[SRBS24] call the the space outside of all spheres "feasible region" and the part of a SDF sphere that bounds it a "feasible arc". On a feasible arc, the power distance is zero to its generating sphere and positive to all other spheres. The arc lies in the power cell of its sphere. Its boundary consists of points shared by two spheres, points on the facets bounding the power cell. We summarize:

Observation 2 The feasible arc of the sphere corresponding to an SDF sample is the intersection of the sphere with the power cell of the weighted sample.

It was believed that the construction of the feasible arcs was intractable at least in $3D$ [SRBS24]. However, the power diagram provides a direct method for explicit construction. While in general the intersection of a convex polytope with a sphere is NP-hard, the particular case encountered for SDF is simpler: the power cells corresponding to the samples of one sign can be traversed, explicitly tracing out the boundaries of the feasible arcs on the spheres. This will also reveal the possibly non-trivial topology of the feasible arcs, which can consist of disconnected arcs in 2D or, components that are not simply-connected in three or more dimensions (for example, an annulus on a sphere in 3D).

A simpler approach to construct samples on the feasible arcs is rejection sampling: generate samples on a sphere and check the half spaces of the corresponding power cell. Sampling directly on the arcs allows choosing the distribution of the final samples on the boundary of the feasible region.

4.5. Power contour

The power cells partition the space in two regions, according to the sign of their generators. The regions are separated by a piecewise-linear surface $\mathcal{S}_\mathcal{P}$ formed by the boundary facets separating the cells corresponding to generators with opposing sign. We will call this contour the *power contour*. A visualization is provided in in Fig. 5. It marks the points in space at which two vertices of opposite sign have the closest power distance to two samples of opposite sign. By design, it separates the set of positive and negative spheres without intersecting any. To the best of our knowledge, this is the first construction that is guaranteed to satisfy conditions (1) and (2), and can be constructed deterministically. Its vertices are dual to the primal simplices, hence their location is only directly influenced by vertices that are part of a simplex with vertices of different signs. All other vertices only influence the contour via the connectivity of the triangulation. The surface lies tangent to a circle if it directly touches a circle of opposite sign, so the contour generally fails to satisfy (3). Given the power diagram, $\mathcal{S}_\mathcal{P}$ can be efficiently extracted by iterating over the edges of the primal regular triangulation and inserting the facets dual to the edges connecting vertices of different sign.

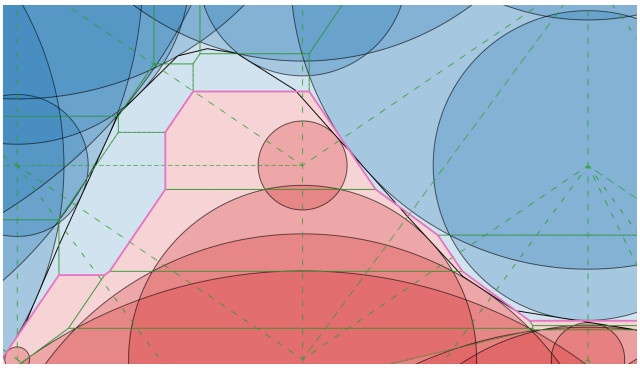

Figure 5: *Power cells: boundaries in solid green, interiors colored according to the orientation of their generator vertex. Power contour shown in pink, regular triangulation in dashed green.*

5. Regular Contouring Approaches

We first examine the effect of the regular triangulation for (simplicial) primal contouring and examine the power contour. In addition to the perceived level of detail, we evaluate the quality of the tessellations using three error metrics applied to the reconstructed mesh R and the ground truth surface S:

- **Hausdorff distance**:
  $d_H(R,S) = \max(\max_{x\in S}\min_{y\in R} d(x,y), \max_{y\in R}\min_{x\in S} d(x,y))$
- **Chamfer distance**:
  $d_C(R,S) = \frac{1}{2}(\int_S \min_{y\in R} d(x,y)dx + \int_R \min_{x\in S} d(x,y)dy)$
- **SDF normal error**, with $n_R(x)$ the unit normal of R at point x:
  $d_N(R,S) = \int_R (1 - \nabla\phi_S(x)\cdot n_R(x))dx$

All metrics are discretely approximated using subdivided meshes.

5.1. Marching Tetrahedra (Primal)

The standard methods to contour an implicit function on a regular grid is marching cubes (MC) [LC87]. Subdividing the cubes into tetrahedra introduces degrees of freedom in the choice of the diagonal that influence the contour. This may be interpreted as taking the the Delaunay triangulation of the samples, which is not unique for the vertices of a rectangular cuboid because they are co-spherical.

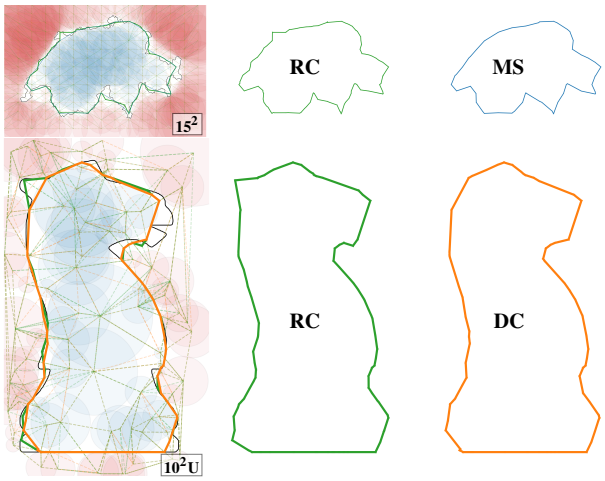

Figure 6: *The marching triangles contour depends on the triangulation. For regular samples (top), the regular triangulation (RC) leads to additional detail compared to "Marching Squares" (MS). For random samples (bottom), the choice of triangles increases detail when compared to the Delaunay triangulation (DC).*

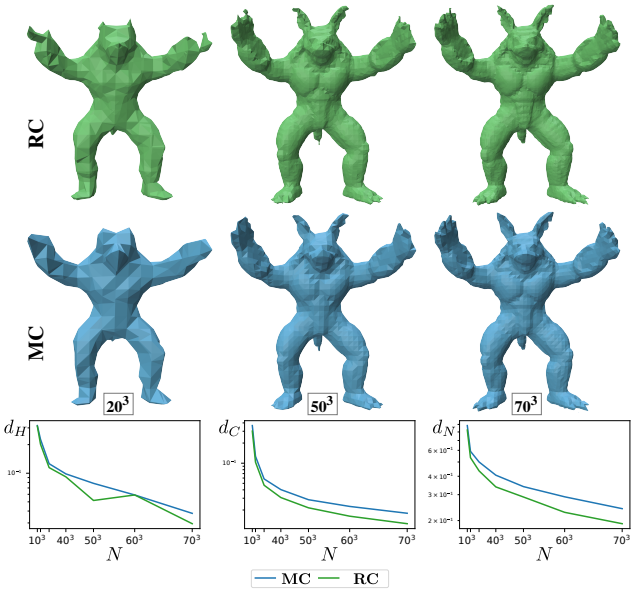

Figure 7: *Regular Marching Tets (RC) and Marching Cubes (MC) reconstructions of the* ARMADILLO *from regular grids of N SDF samples and distances to the ground truth.*

Applying marching tetrahedra to the regular triangulation of the SDF samples (RC) leads to more SDF values contributing to the contour. This results in additional detail clearly visible in the contour at the top of Fig. 6 as well as a higher surface detail on the ARMADILLO 7. For the 3*D* Armadillo, this improvement is also visible in the Hausdorff and Chamfer distance to the ground truth mesh as well as the average normal deviation to the SDF gradient (Figure 7, bottom). For samples in irregular locations, the power distance leads the regular triangulation to favor the contribution of larger circles as opposed to the Delaunay triangulation as shown in Figure 6 (bottom) and Figure 8. This again leads to more segments in the contour and additional detail.

All of the mentioned primal contouring approaches only directly take into account the information provided by a single cuboid / simplex. In areas where the feasible region is large, the additional information provided by the rest of the spheres is not used.

5.2. Power Contour (Dual)

Fig. 9 shows the power contour (PC) in pink together with a marching triangles contour of its dual regular triangulation (RC). Each vertex of the power contour is the power sphere of a triangle that connects vertices of different sign in the primal triangulation - the triangles in which the RC contour lies. The general structure of the reconstructions is therefore very similar. However, the power contour is defined through a lower convex hull, not a level set. Different than the primal reconstruction, the power contour guarantees not to intersect any of the spheres. The structure of the power diagram leads the facets to align with pairs of sample vertices, prominently visible at lower resolutions (Fig. 9).

5.3. Comparison: Reach for the Spheres/Arcs

While we have seen that the local choice of spheres taken by the regular triangulation can increase the perceived quality, it only takes a subset of spheres into consideration for the actual contour. Both methods of Sellán et al. [SBS23, SRBS24] explicitly take into account all available spheres. This can lead to an increased amount of surface detail as shown for RFTA in Fig. 10. For all results we used the code provided by the authors, without investigating the cause of apparent lower quality results. As shown in the follow-up work [SRBS24], the flow-based optimization [SBS23] may run into singularities. The qualitative differences are also reflected in the distance metrics, where RFTA clearly outperforms the other methods on a fixed set of samples. We show a comparison of runtime in Figure 11, the Delaunay / Regular triangulation based approaches require no optimization and are considerably faster than RFTA.

6. Regular Refinement

As we have seen, contouring based on simplices of the regular triangulation or the polygons of the power diagram cannot make use of all the information encoded in the distance spheres. In fact, not even all of the information encoded in the triangulation is used. As noted earlier, the dual spheres have positive radii if and only if they lie outside of all circles in the regular triangulation (**Observation 1**). All dual vertices with positive radius therefore lie in the feasible region and their radius indicates how far away they are from the closest feasible arcs. Fig. 13 shows all dual circles with positive radius in light green.

This naturally leads to a refinement strategy similar to the common Delaunay refinement that we visualize in Fig. 12. Out of all dual spheres with positive weight, we choose the one with the

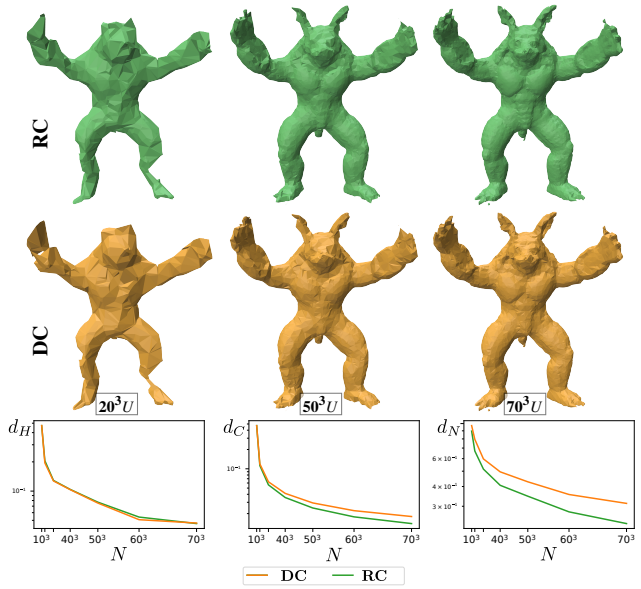

Figure 8: *Marching Tets reconstruction on the Regular (RC) and Delaunay (DC) Triangulation of a set of N random (U) SDF samples and distance metrics to the* ARMADILLO *ground truth averaged over* 5 *sets of samples.*

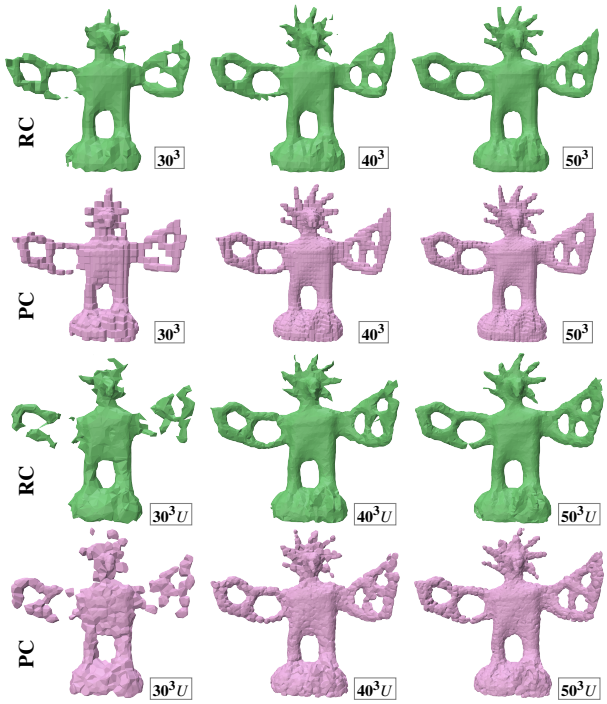

Figure 9: *Marching Tets on the regular triangulation (primal, RC) and the power contour (dual, PC) on regular (top two rows) and random (U, bottom two rows) samples of the* FIREBIRD *model.*

largest radius and sample the SDF at its center. We construct the respective circle (which generally has a different, smaller radius as seen in Fig. 12, second image) and insert it in the regular triangulation. The choice of location allows us to reduce the volume of the feasible region in a very effective way. We use the fact that regular triangulations admit inserting additional points at low computational cost and iterate, generating more and more detail in the reconstruction – this is illustrated in Fig. 2. The general procedure is almost identical to standard Delaunay refinement, but instead of inserting the circumcenters of simplices that contain the contour, we add the location of the weighted circumcenters in the feasible region. We show execution times on the NIGHTINGALE model in Fig. 11 (right), inserting into regular grids of increasing base resolution. Refining from a lower resolution is faster than building the regular triangulation for a comparable amount of samples.

Contouring based on the regular triangulation using the regular refinement strategy (RC_R) clearly outperforms both, Delaunay refinement (DC_D) as well as MC and RFTA when evaluated at similar resolutions. We show reconstructions on the KOALA in Fig. 17. Regular refinement quickly improves in visual detail and proximity to the ground truth. Delaunay refinement also increases the quality measurably, regular refinement leads to faster improvement.

For the more complex models shown in Fig. 14 with thinner parts and holes, our method manages to quickly recover the correct topology. In contrast, Delaunay refinement always samples near the known level sets and potentially misses disconnected components not part of the initial approximation of the level set. Clearly visible on the difficult CELLULARTHING model in Fig. 15, Delaunay refinement eventually reconstructs details but needs to slowly grow the known level set. The regular sampling strategy explicitly encourages sampling in unobserved regions, and all thin structures emerge.

The adaptive domain tessellation together based on a sampling strategy that effectively explores the unknown areas of the feasible region enables the method to reconstruct very complex geometry from a comparatively low amount of SDF samples, clearly outperforming the RFTA approach (Fig. 1). The spheres added during the refinement are generally not as effective in improving the RFTA reconstruction. Applying RFTA on the sets of SDF samples obtained from the refinement methods ($RFTA_R$/$RFTA_D$) tends to produce artifacts as shown in Figure 17.

7. Discussion

The main aim of this work is to point out that regular triangulation and their dual power diagrams are a natural fit for samples of SDFs. We showed how the resulting triangulation can benefit reconstruction quality and in particular lead to a simple yet effective refinement scheme, yielding reconstructed surfaces of high complexity and high quality from far fewer samples than other methods.

We believe that this is only one of many possibilities of how the power diagram can be used for SDF tesselation and see the following directions for future work:

SDF estimation at the dual positions The refinement scheme currently requires the possibility to sample new SDF values. We be-

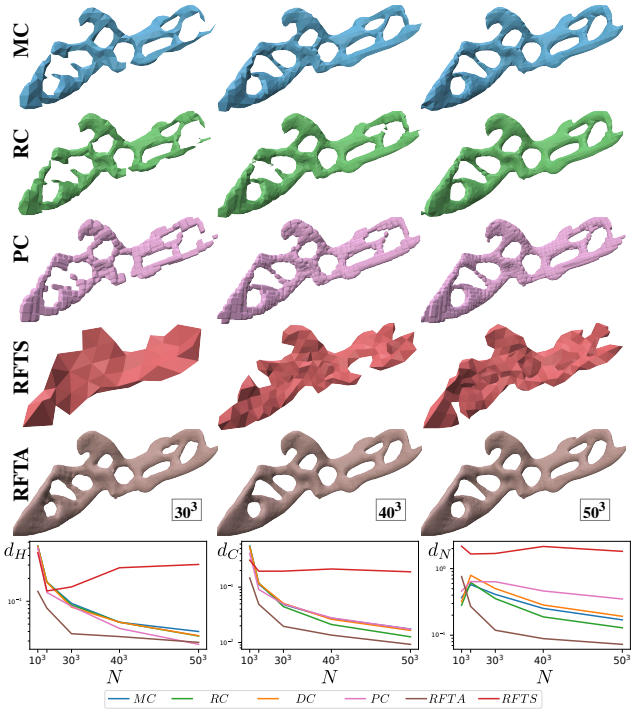

Figure 10: *Contourings of the* NIGHTINGALE *model from regular SDF samples and distance metrics to the ground truth.*

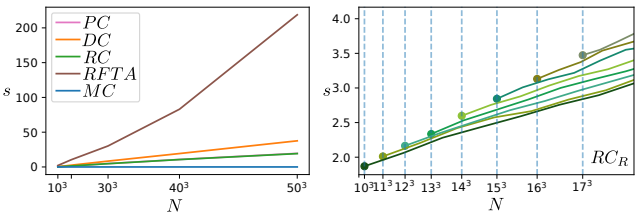

Figure 11: *Runtime comparison on the* NIGHTINGALE *model between methods (left) and for for the incremental approach RC_R (right) starting at different resolutions.*

lieve that estimating the values based on the data derived from the given samples might also help improving the reconstruction.

Adapting RFTS/RFTA Both of the methods by Sellán et al. could profit from the information encoded in the regular triangulation. The power contour as a surface with certain guarantees might be a good starting point for the RFTS optimization. First experiments suggest that additional explicit regularization would be needed to deal with the high initial resolution and the influence of the power distance on the choice of tangents. The RFTA approach relies on finding a point on each feasible arc for an initial tangency set. As we explained earlier, the arcs can explicitly be recovered as the intersection of a sphere with its power cell. In addition to this, the neighborhood information in the regular triangulation can be used for an efficient, direct rejection sampling on the spheres. This can be used as a means to get a dense, uniform sampling directly on the

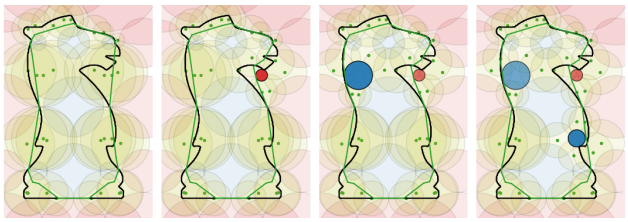

Figure 12: *Regular refinement: we incrementally insert new SDF samples at the location of the dual spheres with largest radius (left to right). We highlight the newly inserted circles and show the marching triangles contour in green.*

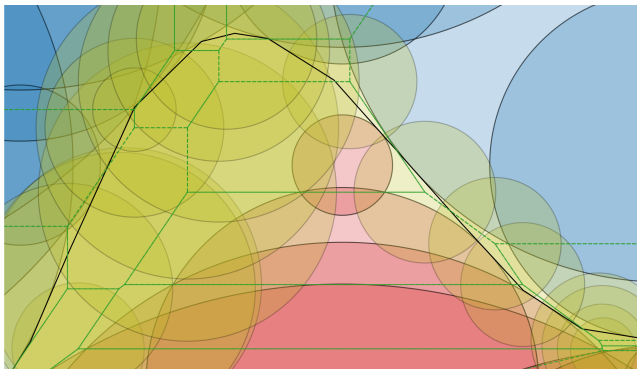

Figure 13: *Dual circles with positive radius lie inside the feasible region. Large spheres are further away from the feasible arcs.*

spheres and also help to accelerate the optimization by enabling a cheap back projection to the feasible arcs.

Explicit Tangent-Point Construction [SRBS24] optimize a set of tangent points on the feasible arcs of the spheres. We observed that attributing each sphere to its closest point on the power contour can

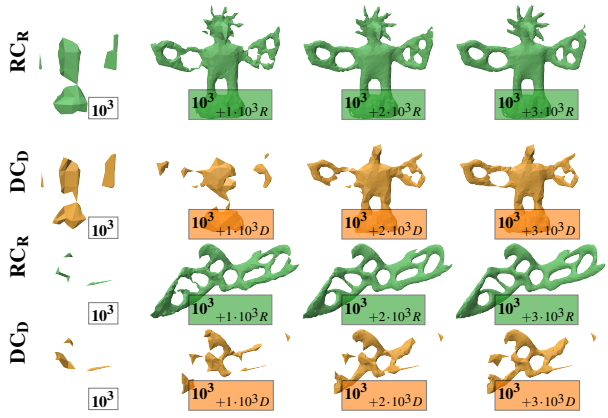

Figure 14: *Starting with a 10^3 grid (left column), we refine 1000 steps per column using the regular (R,green) and Delaunay (D,yellow) triangulation and refinement.*

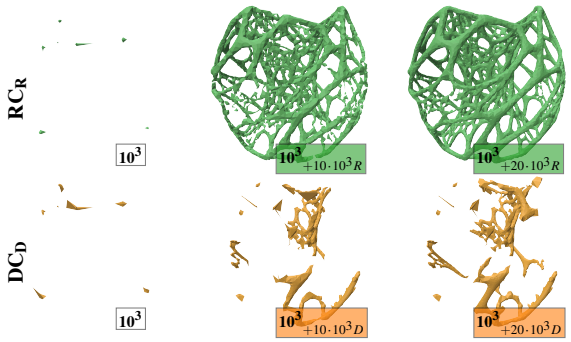

Figure 15: *Regular (R,top) and Delaunay (D,bottom) refinement from* 1*k to* 21*k samples on the difficult* CELLULARTHING *model.*

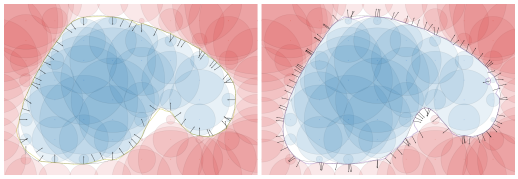

Figure 16: *Contact points in direction of the closest point on the regular contour in the own power cell (left) and RFTA (right)*

yield reasonable tangent directions – we show a simple example in Fig. 16. The quality of the surface is lower compared to RFTA, but the power diagram might contain information to derive a simple heuristic for the choice of tangent points, without optimization.

Connectivity Optimization The weights of the regular triangulation enable optimizing the connectivity of the triangulation. This could be used to further optimize the triangulation, considering e.g. the amount of spheres that contribute to the power contour.

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.

References

[AB86] ASH P. F., BOLKER E. D.: Generalized dirichlet tessellations. *Geometriae Dedicata 20*, 2 (1986), 209–243. doi:10.1007/BF00164401. 3

[AK00] AURENHAMMER F., KLEIN R.: Chapter 5 - voronoi diagrams. In *Handbook of Computational Geometry*, Sack J.-R., Urrutia J., (Eds.). North-Holland, Amsterdam, 2000, pp. 201–290. doi:https://doi.org/10.1016/B978-044482537-7/50006-1. 2, 3

[AKL13] AURENHAMMER F., KLEIN R., LEE D.-T.: *Voronoi Diagrams and Delaunay Triangulations*. WORLD SCIENTIFIC, 2013. arXiv:https://www.worldscientific.com/doi/pdf/10.1142/8685, doi:10.1142/8685. 2, 3

[Ale19] ALEXA M.: Harmonic triangulations. *ACM Transactions on Graphics (TOG) 38*, 4 (2019), 1–14. 2

[Bat11] BATTY C.: Problem: Reconstructing meshes with sharp features from signed distance data, 2011. URL: https://web.archive.org/web/20111112202136/http://www.cs.columbia.edu/~batty/misc/levelset_meshing/level_set_reconstruction.html. 3

[CDE*00] CHENG S.-W., DEY T. K., EDELSBRUNNER H., FACELLO M. A., TENG S.-H.: Sliver exudation. *J. ACM 47*, 5 (Sept. 2000), 883–904. doi:10.1145/355483.355487. 2

[CDS12] CHENG S.-W., DEY T. K., SHEWCHUK J.: *Delaunay Mesh Generation*, 1st ed. Chapman & Hall/CRC, 2012. 3

[Che04] CHEN L.: Mesh smoothing schemes based on optimal delaunay triangulations. In *Proceedings of the 13th International Meshing Roundtable, IMR 2004, Williamsburg, Virginia, USA, September 19-22, 2004* (2004), Üngör A., (Ed.), pp. 109–120. URL: http://imr.sandia.gov/papers/abstracts/Ch317.html. 2

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

[CX04] CHEN L., XU J.-C.: Optimal delaunay triangulations. *Journal of Computational Mathematics* (2004), 299–308. 2

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

[dALJ*15] DE ARAÚJO B. R., LOPES D. S., JEPP P., JORGE J. A., WYVILL B.: A survey on implicit surface polygonization. *ACM Comput. Surv. 47*, 4 (May 2015). doi:10.1145/2732197. 2

[DLRS10] DE LOERA J. A., RAMBAU J., SANTOS F.: *Triangulations: Structures for Algorithms and Applications*, 1st ed. Springer, Berlin,Heidelberg, 2010. 3

[DPT01] DEVILLERS O., PION S., TEILLAUD M.: *Walking in a triangulation*. Tech. Rep. RR-4120, INRIA, 2001. URL: https://hal.inria.fr/inria-00072509. 3

[Ede90] EDELSBRUNNER H.: An acyclicity theorem for cell complexes ind dimension. *Combinatorica 10*, 3 (Sep 1990), 251–260. doi:10.1007/BF02122779. 3

[ES92] EDELSBRUNNER H., SHAH N. R.: Incremental topological flipping works for regular triangulations. In *Proceedings of the Eighth Annual Symposium on Computational Geometry* (New York, NY, USA, 1992), ACM, pp. 43–52. doi:10.1145/142675.142688. 3

[FAB*18] FENG L., ALLIEZ P., BUSÉ L., DELINGETTE H., DESBRUN M.: Curved optimal delaunay triangulation. *ACM Transactions on Graphics 37*, 4 (2018), 16. 2

[FC24] FENG N., CRANE K.: A heat method for generalized signed distance. *ACM Trans. Graph. 43*, 4 (July 2024). doi:10.1145/3658220. 2

[Grü03] GRÜNBAUM B.: *Convex Polytopes*. Springer, 2003. doi:10.1007/978-1-4613-0019-9. 3

[Har96] HART J. C.: Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer 12*, 10 (Dec. 1996), 527–545. doi:10.1007/s003710050084. 2

[JLSW02] JU T., LOSASSO F., SCHAEFER S., WARREN J.: Dual contouring of hermite data. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2002), SIGGRAPH '02, Association for Computing Machinery, p. 339–346. doi:10.1145/566570.566586. 2, 4

[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. 2, 4

[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, 3

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

[KKLD23] KERBL B., KOPANAS G., LEIMKUEHLER T., DRETTAKIS G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph. 42*, 4 (July 2023). doi:10.1145/3592433. 2

[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, 4, 5

[MBWB02] MUSETH K., BREEN D. E., WHITAKER R. T., BARR A. H.: Level set surface editing operators. *ACM Trans. Graph. 21*, 3 (July 2002), 330–338. doi:10.1145/566654.566585. 2

[MEM*20] MACKLIN M., ERLEBEN K., MÜLLER M., CHENTANEZ N., JESCHKE S., CORSE Z.: Local optimization for robust signed distance field collision. *Proc. ACM Comput. Graph. Interact. Tech. 3*, 1 (May 2020). doi:10.1145/3384538. 2

[MSLJ23] MARSCHNER Z., SELLÁN S., LIU H.-T. D., JACOBSON A.: Constructive solid geometry on neural signed distance fields. In *SIGGRAPH Asia 2023 Conference Papers* (New York, NY, USA, 2023), SA '23, Association for Computing Machinery. doi:10.1145/3610548.3618170. 2

[MST*21] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHI R., NG R.: Nerf: representing scenes as neural radiance fields for view synthesis. *Commun. ACM 65*, 1 (Dec. 2021), 99–106. doi:10.1145/3503250. 2

[PFS*19] PARK J. J., FLORENCE P. R., STRAUB J., NEWCOMBE R. A., LOVEGROVE S.: Deepsdf: Learning continuous signed distance functions for shape representation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019* (2019), Computer Vision Foundation / IEEE, pp. 165–174. doi:10.1109/CVPR.2019.00025. 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), SA '23, Association for Computing Machinery. doi:10.1145/3610548.3618196. 2, 3, 4, 6

[She96] SHEWCHUK J. R.: Robust adaptive floating-point geometric predicates. In *Proceedings of the Twelfth Annual Symposium on Computational Geometry* (New York, NY, USA, 1996), SCG '96, Association for Computing Machinery, p. 141–150. doi:10.1145/237218.237337. 3

[She98] SHEWCHUK J. R.: Tetrahedral mesh generation by delaunay refinement. In *Proceedings of the Fourteenth Annual Symposium on Computational Geometry* (New York, NY, USA, 1998), SCG '98, Association for Computing Machinery, p. 86–95. doi:10.1145/276884.276894. 2

[She02] SHEWCHUK J. R.: What is a good linear element? interpolation, conditioning, and quality measures. In *IMR* (2002), pp. 115–126. 2

[SMB*20] SITZMANN V., MARTEL J. N. P., BERGMAN A. W., LINDELL D. B., WETZSTEIN G.: Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual* (2020), Larochelle H., Ranzato M., Hadsell R., Balcan M., Lin H., (Eds.). 2

[SRBS24] SELLÁN S., REN Y., BATTY C., STEIN O.: Reach for the arcs: Reconstructing surfaces from sdfs via tangent points. In *SIGGRAPH 2024 Conference Papers* (2024). 2, 3, 4, 5, 6, 9

[VSJ22] VICINI D., SPEIERER S., JAKOB W.: Differentiable signed distance function rendering. *ACM Trans. Graph. 41*, 4 (July 2022). doi:10.1145/3528223.3530139. 2

[Zie95] ZIEGLER G. M.: *Lectures on Polytopes*. Springer, 1995. doi:10.1007/978-1-4613-8431-1. 3

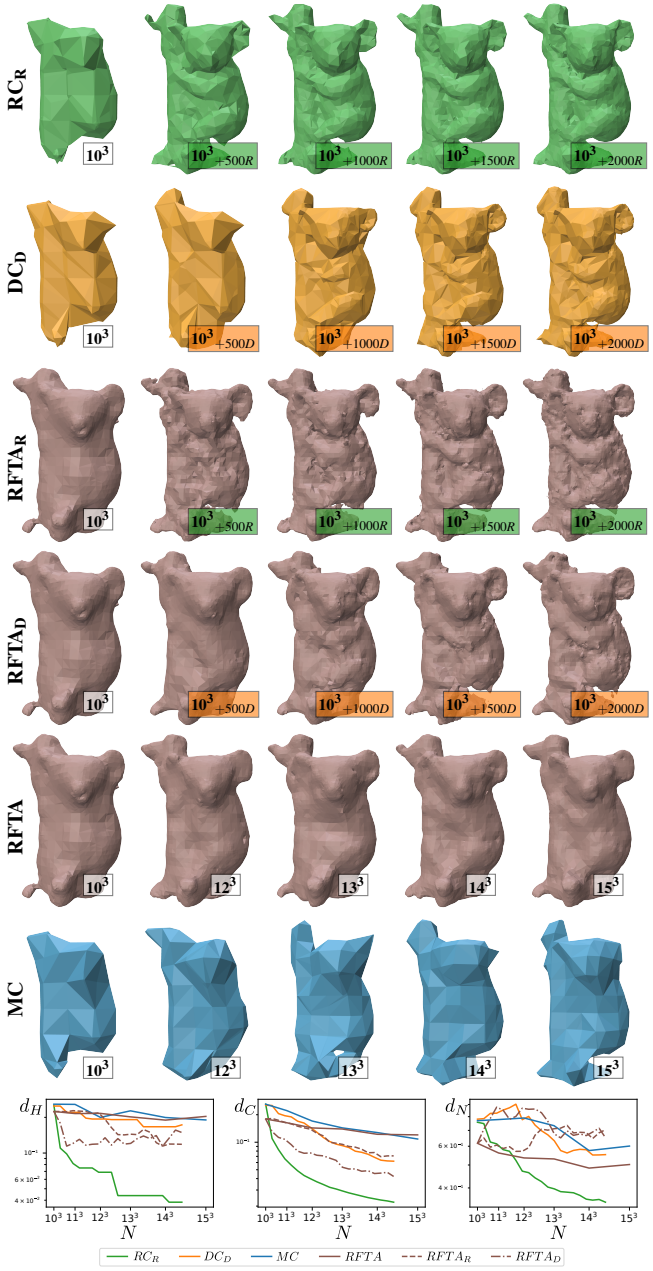

Figure 17: *Reconstructions from SDF samples on regular grids with and without additional incremental samples and difference metrics to the ground truth. Grid resolutions and number and type of incremental samples are given in the boxes on the bottom right of each reconstruction. In the subscripts and boxes, D/R indicate the use of Delaunay / Regular refinement respectively.*