

Be Water my Friend: Mesh Assimilation

Dennis R. Bukenberger, Hendrik P. A. Lensch

Department of Computer Graphics, Eberhard Karls University, Tübingen, Germany




Figure 1: Progress of our proposed method: Starting with a low-poly icosphere as initialization mesh, the surface grows within the target point cloud. The mesh first approximates the rough shape and further progresses by assimilating and refining to reconstruct fine detail.

Abstract

Inspired by the ability of water to assimilate any shape, if being poured into it, regardless if flat, round, sharp, or pointy we present a novel, high-quality meshing method. Our algorithm creates a triangulated mesh, which automatically refines where necessary and accurately aligns to any target, given as mesh, point cloud, or volumetric function. Our core optimization iterates over steps for mesh uniformity, point cloud projection, and mesh topology corrections, always guaranteeing mesh integrity and ϵ -close surface reconstructions. In contrast to similar approaches, our simple algorithm operates on an individual vertex basis. This allows for automated and seamless transitions between the optimization phases for rough shape approximation and fine detail reconstruction. Therefore, our proposed algorithm equals established techniques in terms of accuracy and robustness but supersedes them in terms of simplicity and better feature reconstruction, all controlled by a single parameter, the intended edge-length. Due to the overall increased versatility of input scenarios and robustness of the assimilation, our technique furthermore generalizes multiple established approaches such as ballooning or shrink wrapping.

CCS Concepts

• **Computing methodologies** → **Mesh models; Point-based models; Mesh geometry models;**

1. Introduction

Without the need for mappings, parameterizations, or specifically trained machine-learning models, our approach presents a simple yet powerful extension in the field of geometric surface reconstruction techniques. We demonstrate, how our algorithm can top similar state-of-the-art approaches; e.g. when reconstructing fine details and sharp feature edges, in terms of mesh uniformity or versatility of supported input data. Our method is based on the concept of water being quite formless and able to assimilate any shape if being poured into it. Starting from a small initial triangle mesh, the object can grow, shrink and locally refine to assimilate a target shape with controllable precision. Vertices of the initial mesh extend rather autonomously along surface normals until they reach the target hull space. The expanding surface mesh is supplemented with suitably interpolated vertices where it is stretched the most. In contrast to the surface tension of water, the inter-vertex energy optimization

promotes mesh uniformity but without enforcing smoothness priors, often found in other ballooning concepts. Once the vertices are close enough to the target hull, they individually transition from mesh growth mode to a projection scheme. Due to the bilateral weighting concept of the projection, vertices are effectively pulled into intersecting tangent planes of the hull, thus allowing for accurate reconstructions of small detail and sharp edges. The optimization of vertices is an iteration process, but individual for each vertex and can therefore be executed on the GPU as a massively parallelized operation. Per design, the algorithm assimilates to shape rather than a strict type of input data, which can be given as mesh, point cloud, signed distance, or other volumetric function. With its frugal input requirements and a minimal set of parameters to be adjusted, our method simplifies the rather complex nature of other established techniques; it is situated somewhere between the exhaustively studied fields of shape approximation [CSAD04], surface remeshing [AUGA08] and point cloud reconstruction [BTS*17].

1.1. Contributions

Our approach fulfills all criteria which one would ask of a state-of-the-art reconstruction technique: Guarantee for closed, watertight manifold meshes, the ability to cope with varying input density, robustness to noise and outliers as well as to missing data and control over higher genus topology. Moreover, we can summarize our main contributions: • **Mesh uniformity** is achieved as our mesh unravels with an inter-vertex energy, striving to equalize distances between vertices. • **Adaptive resolution** can be realized by locally adapting the specified target-edge-length to the provided sample density. • **Sharp features** are reconstructed, as the second vertex optimization step is designed to minimize projected distances to surface tangents, thus the vertices snap on close-by edges and corners. This is where most state-of-the-art methods still lack precision. • **Arbitrary initialization** meshes are possible, as demonstrated with various examples. • **Versatile input** data is supported, as our method is not bound to point clouds and can easily handle meshes as well as volume data and signed distance fields, currently explored in neural surface representations [MON*19, PFS*19].

1.2. Related Work

For decades, shape reconstruction has been, and is still an actively researched topic. While the ideas of ballooning or shrink-wrapping are not exactly new [DQ01], we are yet to see its full potential. In some aspects, our method follows similar principles like the concept of *Competing Fronts* [SLS*06], which is also a coarse-to-fine point cloud reconstruction technique, using growing mesh geometry. But instead of whole mesh fronts, which get frozen at some point, our method operates on an individual vertex basis so that the mesh remains flexible during the optimization. This allows for smoothing out irregularities, without the need for remeshing in every other iteration. Our hull projection scheme furthermore resolves the issue of reconstructing creases and sharp feature edges, an open problem for *Competing Fronts*. Unfortunately, a direct comparison was not possible as there is no reference implementation or result data available. Nevertheless, the concept gained some recent attention with the concept of *Cooperative Evolutions* [LL20], combining two mesh fronts enclosing on the point cloud from the in- and outside in parallel. A first draft utilizing this idea in a learning based approach was proposed with *Point2Mesh* [HMGCO20]. A thorough analysis of the actual distinctions of our approach to other ballooning concepts is featured in our discussion in Section 3.2.

Conceptually different approaches include the popular Screened Poisson Surface Reconstruction (SPSR) [KBH06, KH13, KH19, KCRH20], which is probably one of the most commonly used point cloud reconstruction methods. In direct comparison, however, SPSR tends to produce rather smoothed out results, whereas ours generally appear sharper and capture finer details, even at lower resolution. Due to the octree-nature of the SPSR approach, the resulting triangulation is quite inhomogeneous, while ours approaches both a uniform distribution and Delaunay-like connectivity. Scale Space Meshing (SSM) [DMSL11] is an approach that really takes advantage of modern high precision 3D scanners and is able to faithfully reconstruct even very fine details. However, as our comparison shows, sharp edges are not captured very well. The marching cubes

algorithm [LC87] is often used in a medical context to reconstruct surfaces from volumetric data or to remesh existing surfaces. Our algorithm is able to handle both as it can be used on signed distance fields (SDF) as well as on existing meshes with the advantage of producing more evenly distributed vertices and triangle faces. Both Poisson and Marching Cubes are able to recover surfaces but fall short in terms of triangle mesh homogeneity, a key contribution of our approach. Similar deficiencies can be attributed to α -shape- [EM94], Voronoi-based- [ABK98, AB99, ACK01, BL17] or the ball-pivoting-algorithm (BPA) [BMR*99]. For these algorithms, the achievable mesh geometry-, and sometimes topology-, quality crucially depends on the density and uniformity of the given input point set and sometimes requires prior outlier-filtering or noise-smoothing steps. Our approach is quite invariant to these artifacts, as it is able to cope with noise and can either ignore or incorporate varying sample densities. Impressive developments in quad-mesh reconstructions such as Instant Field Aligned Meshes (IFAM) [JTSSH15] or Online Surface Reconstruction (OSR) [STJ*17] produce nice feature-aligned mesh structures, but often struggle with low or varying sample density as demonstrated in comparisons to our results. Most recently, learning-based approaches also joined the reconstruction game [MON*19, PFS*19] but often require specifically trained models on an explicit object class. This is no longer a requirement with P2M, which reconstructs point clouds via shrink-wrapping and basically learns from the input itself. But their results also lack precision in corners and sharp object features.

2. Method

As illustrated in Figure 1, the goal of our algorithm is to assimilate to a target shape, starting from a simple initialization mesh. Therefore, the basis mesh first approximates the rough shape of the input, finer details are recovered as the optimization continues. This assimilation process follows a simple optimization scheme, iterating over the following five steps, performed in one iteration cycle:

- I. Subdivide the mesh where necessary
- II. Equalize inter-vertex distances
- III. Mesh expansion and hull projection
- IV. Fix mesh irregularities
- V. Add tunnels to increase genus (optional)

Termination criteria for the optimization can be specified, either as a target-edge-length l_t for the final mesh or an ϵ -closeness threshold to the target shape. Each operation is self-contained, meaning the surface remains a closed manifold and valid mesh at all times. Related coarse-to-fine approaches [SLS*06, LL20, HMGCO20] separate the *mesh development* and *final fitting* phase in their procedure or interleave the assimilation progress with remeshing operations. However, the strength of our method lies in the mixed design of the assimilation (III): Simple mesh growth seamlessly transitions into a projection scheme on an individual vertex basis.

Terminology Our method is first introduced with the focus on point cloud reconstruction, the adaptation to other input scenarios is later addressed in Section 2.1. In the following we associate *vertices* v_i with the growing mesh structure and *samples* s

with the input point cloud. Variables related to samples are denoted with an over-set dot. Further, we observe local 1-ring neighborhoods around vertices: For a vertex v_i , adjacent mesh neighbors are grouped in N_i where $k_i = |N_i|$ gives the number of neighbors. The average edge length around vertex v_i is therefore given as $l_i = \frac{1}{k_i} \sum_{j \in N_i} |v_i - v_j|$. Vertex normals n_i are derived from the face normals in this 1-ring neighborhood: The normalized portions of corner angles in the triangle fan are used to weigh their normals' contribution, respectively. Further, each vertex v_i is associated with a set S_i of 10 closest point cloud samples. This is cheap to initialize with a brute-forced search as the initial mesh only has very few vertices. When starting from a larger base mesh, a suitable data structure may speed up the first nearest neighbor associations. However, once established, this structure can be maintained cheaply using the mesh as a graph: To update the existing sets or create new ones for added vertices (**I.**) we then only query the sample sets in the vertices' 1-ring neighborhoods for potential new nearest neighbors. Further, each sample is interpreted as a splat-disc on the implicit hull, represented by a normal \hat{n} and a radius r . We compute the radius of a sample as half the median distance to its 10 closest point cloud neighbors. While all our results are generated using this simple approach, more advanced splat methods [WK04, CGAY13] may be suitable as well.




Figure 2: Mesh refinement: The 1-ring fan of a vertex is $\sqrt{3}$ -subdivided. New vertex positions are interpolated as PN triangle centers [VPBM01] to maintain local curvature.

I. Refinement The triangle mesh is constantly refined with targeted $\sqrt{3}$ -subdivision operations [Kob00]. A vertex qualifies for refinement when its average length l_i of incident edges is larger than the specified target-edge-length l . Or, if ϵ -closeness is specified as the optimization criteria, simply the vertices with the largest l_i get subsequently refined. Therefore, faces around a vertex are replaced as shown in Figure 2: For each triangle in this 1-ring fan, a new vertex is created along with three new triangles. The new vertices are the interpolated center points of the fan triangles, interpreted as curved PN triangles [VPBM01] using vertex normals. This allows us to maintain local curvature and the subdivided geometry smoothly integrates with the surrounding mesh as the outer 1-ring loop edges are not split. This operation increases the valence of the 1-ring neighbor vertices by 1 and may create very obtuse or pointy triangles. Nevertheless, both issues are fixed within the same cycle by upcoming steps **II.** and **IV.**

II. Equalization In order to obtain the most regular mesh structures with evenly distributed vertices, we employ a simple geometric optimization: Every vertex v_i strives to individually equalize the distances to its direct 1-ring neighbors $v_j \in N_i$, eventually resembling a Delaunay-like triangulation. Equation 1 formulates the update vectors, where $\vec{v}_{ij} = v_j - v_i$ and l_i is the average length of edges incident on v_i .

$$\vec{e}_i = \left(\frac{1}{k_i} \sum_{v_j \in N_i} v_j - l_i \frac{\vec{v}_{ij}}{|\vec{v}_{ij}|} \right) - v_i \quad (1)$$

These geometric increments mainly smooth out dense clusters created by step **I.** and can be applied with a weighted update on vertex positions as $v_i := v_i + \lambda_e \vec{e}_i$. Thus, λ_e allows adjusting the progression speed for mesh homogenization. We achieved the best results with a soft update, using $\lambda_e = 0.1$.

III. Assimilation The goal of this step is to pull vertices to the surface and snap vertices on edges and corners. Once our mesh is close enough to the point cloud hull, a suitable projection technique is applied. However, when starting from a generic base mesh, e.g. a simple sphere inside of, or enclosing, the point cloud, it first has to grow or shrink into a rough approximation of the target shape, so that the projection may take over later. Whereas this is how it is usually done, we propose a mixed procedure, adaptive to each individual vertex. This step modifies the mesh's geometry by moving vertices with individual update vectors. These vectors \vec{a}_i may embody surface projection (**III.A**) or mesh growth (**III.B**). Equation 2 anticipates which one is to choose, dependent on an individual vertex v_i and its weights \hat{w}_{ij} , summed up in $\omega_i = \sum_{j \in S_i} \hat{w}_{ij}$. Other variables are elaborated in the following.

$$\vec{a}_i = \begin{cases} \frac{1}{\omega_i} \sum_{j \in S_i} \hat{w}_{ij} \vec{v}_{ij} & \text{III.A) if } \omega_i > \epsilon \\ \hat{n}_i \cdot \frac{\rho_i l_i}{4} & \text{III.B) else} \end{cases} \quad (2)$$

III.A) Projection For valid weights \hat{w}_{ij} , gathered from the point cloud samples associated with vertex v_i , the update vector \vec{a}_i is the weighted sum of v_i 's projection vectors \vec{v}_{ij} into the samples' tangent planes (Figure 3, green). Projected points compute as $\hat{v}_{ij} = v_i + \vec{v}_{ij}$ with vectors $\vec{v}_{ij} = \hat{n}_j \cdot (\hat{s}_j - v_i) \hat{n}_j$ where \hat{n}_j is the normal of point cloud sample \hat{s}_j . The bilateral weights \hat{w}_{ij} combine a distance and an angular component. This gives us the advantage of only pulling qualified vertices into edges and corners, while their direct neighbors reside in flat regions nearby. As meaningful distance measures we use the closest points (Figure 3, blue) on the individual splat-discs respectively, as formulated in Equation 3, where \hat{r}_j is the radius of sample \hat{s}_j . Further, the actual distance weight is the result of a Gaussian, centered on $\mu = 0$ with $\sigma = \frac{l_i}{3}$. The angular component computes as the dot-product of vertex and sample normal, clipped to values ≥ 0 . As formulated in Equation 4, \hat{w}_{ij} is simply the product of both components.

$$\hat{d}_{ij} = \begin{cases} |v_i - \hat{s}_j + \frac{\hat{v}_{ij} - \hat{s}_j}{|\hat{v}_{ij} - \hat{s}_j}| \cdot \hat{r}_j| & \text{if } |\hat{v}_{ij} - \hat{s}_j| > \hat{r}_j \\ |\hat{v}_{ij}| & \text{else} \end{cases} \quad (3)$$

$$\hat{w}_{ij} = \exp\left(-\frac{9\hat{d}_{ij}^2}{2l_i^2}\right) \cdot \max(0, \hat{n}_i \cdot \hat{n}_j) \quad (4)$$

Figure 3 visualizes the strength of our proposed concept. The contribution of a force, that pulls vertices into tangent planes is scaled down with increasing distance from its originator. However, the multiplication with the angular component is, what really

makes the difference when it comes to recovering more surface detail: Vertex v_0 lies on a flat part of the mesh, the contribution of samples *just around the corner* (\hat{s}_D, \hat{s}_E) is decreased significantly or even canceled out completely. For a vertex residing in a curved neighborhood or on an edge like v_1 , the attraction of samples from both sides of the implicit edge (\hat{s}_B, \hat{s}_C & \hat{s}_D, \hat{s}_E) effectively pull it evenly into the intersection of their tangent planes. Even if the splat-discs themselves do not intersect, this is guaranteed because only the projection directions contribute to the update vectors.




Figure 3: Point cloud samples \hat{s}_i are shown with their splat discs and normals \hat{n}_i . The example shows projections for vertices v_0 (left) and v_1 (right) as well as closest points on splat discs, respectively. Vertices move along the projection direction, weighted by distance and the angle between vertex normal and sample normal, and are pulled into the intersection of tangent planes, thus allow for better edge preservation.

III.B) Growth If the weights summed up in ω_i are too small, the vertex v_i is not yet sufficiently close to the point cloud hull. Therefore, the displacement vector \vec{v}_i is derived from its normal n_i , scaled by $\frac{l_i}{4}$. This very robustly ties the geometric increment of a vertex to the scale of its local mesh neighborhood. Further, the normal is oriented with a value $o_i = \pm 1$, which computes as the combined majority sign of all dot-products between normal n_i and projection vectors \vec{v}_{ij} . This allows for mesh expansion as shown in Figure 1, when the mesh is within the target shape, but also for shrinkage as in Figure 6 (top), with the base mesh enclosing the target shape. Further robustness is exemplified in Figure 23 (top) with an out-of-target initialization mesh, which first gets drawn into the target and then starts to grow.

Again, vertices are displaced using a differential update $v_i := v_i + \lambda_a \vec{a}_i$, where λ_a allows to adjust the mesh’s assimilation speed. However, as the differential increment is by design also linked to dimensions of the local mesh neighborhood, a large λ_a could provoke single vertices to grow out too fast. A moderate growth rate of $\lambda_a = 0.1$, allows step II. to level out length discrepancies before the next assimilation update, thus promotes a more even and homogeneous mesh expansion.




Figure 4: Mesh improvement operations: The concave (blue) and convex (green) edges (standing out in contrast to their neighbors) are fixed by rotation. A short (red) edge is collapsed to one vertex.

IV. Improvements To retain a high mesh quality, unfavorable edge constellations are detected and fixed [DQ01] on-the-fly: Concave or convex edges, which stand out in contrast to their direct triangle opposites are fixed with a simple rotation. Figure 4 shows examples of these cases and how they are resolved. Edges shorter than a certain length (e.g., $\frac{l_i}{10}$) are collapsed, merging two vertices into one. This might occur after the edge was rotated or if two neighboring vertices are drawn onto the same implicit edge and get too close to each other. As this step proactively prevents local foldovers, our results do not feature manifold violation artifacts.




Figure 5: Adapting the topology. Tunnels are invariant of the surface orientation and automatically apply for ballooning (top) as well as shrink-wrapping (bottom).

V. Object Topology Adaptation In the most basic scenario, the mesh assimilation is initialized on a simple spherical mesh with a target of the same genus 0. However, to support reconstructions of objects with higher genus, our algorithm supports a common ballooning technique: Tunnels are inserted between opposing mesh fronts to automatically adapt the mesh’s topology. The approach of *Cooperative Evolutions* [LL20] is to detect self-intersections after they have occurred, delete the affected regions, stitch them together, refine the topology and smooth the geometry. Our solution is more similar to the pro-active approach of *Competing Fronts* [SLS*06], using proximity tests on moving vertices to trigger the insertion of tunnels *before* the mesh self-intersects. Therefore, the 1-ring neighborhoods around close vertices or triangles are removed and replaced by tunnels, connecting the opposing edge rings. Tunnels are triangulated analogously to *Freestyle* [SCC11] using the Bresenham algorithm [Bre65] and can therefore be ensured to be watertight. Dedicated topology refinement or geometry smoothing is not required as both are automatically taken care of within the next iterations cycles.

However, as the initial mesh may be quite arbitrary, tunnels can also resolve inter-mesh situations as shown in Figure 19 (center) and Figure 23 (bottom), when initialized with separated geometry. As shown in Figure 5, the operation is not bound to surface orientation and may also form inverted tunnels, resulting in the correct genus nevertheless. This step is optional and may be omitted if the target genus is known to be 0 or the initialization mesh already has the correct genus.

2.1. Generalized Input

To be able to mesh more than just point clouds by assimilation, we actually only have to adapt step **III.** of the optimization. This is demonstrated with meshed input, signed distance fields, and volumetric data.




Figure 6: Top: Our algorithm assimilates the coarse icosphere to the target shape; given as quad mesh [Cra20]. Bottom: Swapped input and target lead to a coarsening; with a larger target-edge-length, short edges are incrementally collapsed. Progress is left to right.

Remeshing of existing meshes, e.g. to ensure hole-free reconstruction or uniform mesh quality is quite trivial to handle, as we can simply use actual surface triangles instead of splat discs. Distances, directions, and weighting follow straightforward. Two examples using a meshed input as the target are shown in Figure 6. On top, the sphere shrink-wraps around the target shape and refines for a smooth and curved surface. For the bottom case, the target-edge-length is increased, which leads to an incremental collapse of short edges and the mesh assimilates the coarse icosphere again.




Figure 7: Mesh construction from a signed distance field, visualized as a color coded cross-section slice: **outside**, **close to 0**, **inside**. As there are no splat discs (point cloud tangents), the hull projection step is omitted.

Signed distance fields are supported by our mesh construction method as well. Therefore, one only has to replace the orientation o_i in Equation 2 with a signed-distance. It is sufficient to sample the field once for each vertex position, acquiring single scalar values respectively, as the vertex normal is used as movement direction. The update vectors \vec{a}_i are always in mesh-growth mode (**III.B**), since

there is no trivial tangent space given, as with splat discs. This basically comes down to an analogous concept as *Cooperative Evolutions* [LL20] and, therefore, has the same limitations: As there is no surface projection, vertices are not automatically drawn towards edges or corners, thus sharp features are not recovered very well. Nevertheless, smooth and curved surfaces are still reconstructed accurately: Update vectors (Eq. 2) are automatically scaled down to 0 length as the (absolute) signed-distance decreases and, therefore, the mesh closely approaches the implicit target hull. Figure 7 illustrates surface reconstructions of a target shape, given as a signed distance field.




Figure 8: Tomography data [HBT*13]: A neuron, given as 1325 binary slices (not all shown) with 661×595 pixels each, reconstructed as surface mesh.

General volumetric 3D data, e.g. as acquired from computational tomography, require a bit more effort but can be reconstructed as well. Figure 8 shows the reconstruction of a CT scanned neuron, given as sliced binary volume. Analogous to the signed distance field, there is again no trivial tangent space for hull projection. Orientations are not so easily determined as with the SDF, but nevertheless, can be easily sampled and trilinearly interpolated from the volume at mesh vertex positions, with binary values giving the sign. Update vectors are scaled down to zero length, once the mesh vertex is situated between both in- and outside labeled voxels, indicating that the mesh reached a surface. By construction, the resulting mesh is watertight. Segments of dendrites separated by noise in the scan data are however not connected by our approach.

3. Experiments and Discussion

In this section, we experiment with the achievable mesh quality of our method, challenged with the five most common point cloud artifacts, as described by Berger et al. [BTS*17]. A numerical comparison follows in Section 3.1 as well as a qualitative discussion in Section 3.3.

Noise and outliers are prominent artifacts in 3D scan data and can be robustly dealt with by our algorithm. Figure 9 shows results of the challenge designed for *robust moving least-squares fitting* [FCOS05]. A point cloud, sampled reasonably dense but randomly, is perturbed with increasing magnitudes of random spherical jitter noise, scaled by the fandsk's bounding box diagonal. Without




Figure 9: Mesh construction on noisy input. Besides some kinks, strong edges remain sharp even with increasing noise of randomly jittered samples. The noise magnitude is given in percent of the bounding box diagonal.




Figure 10: Reconstruction on 3D scan data [TL14]. When aligned, the 10 individual partial scans (colored) create noisy overlap regions. Since every mesh vertex is attracted by multiple sample tangents, the noise is leveled out and the reconstruction is smooth.

noise, ground-truth-like results are achieved. At 0.2% noise scale [FCOS05], flat areas are still flat, curves smooth, edges sharp and only a few kinks become noticeable on very obtuse edges. With noise even further increased by factor 5, our reconstruction still robustly captures all major features of the object. Figure 10 shows details of the noisy overlap regions in a real 3D scan, along with the smoothly reconstructed surface. When not extracted in a pre-filtering step, outliers do not cause any difficulties in our approach either: Mesh vertices are associated with a fixed number of nearest point cloud neighbors. Therefore, far-out outliers are usually not part of any neighborhood, but if so anyway, their influence is still minimal as their contribution is weighted inverse to their distance.

Missing data in the input geometry is recovered by the reconstruction. The first example in Figure 11 shows the *twistcube* with cutout geometry on a curved side area, an edge, and on a corner. The reconstructed mesh smoothly interpolates the missing regions, stays within the implicit hull, and does not grow through the holes. The second example shows Igea, where 75% of the source mesh was cut out and removed in randomly shaped patches. Our algorithm still reconstructs a closed mesh and recovers fine details.

Sample density is not a decisive factor for successful reconstructions. While modern 3D scanners usually provide quite dense point clouds and more information always allows for more precise results, our algorithms splat discs also automatically adapt to fewer and sparser samples. Figure 12 shows the same object with different sampling but accurate results with smooth curves and sharp edges in both cases.




Figure 11: Filling up holes: Results on defected, then sampled input. The twistcube is missing portions on a curved side region, an edge and a corner. 75% of Igea's surface was removed by cutting out random patches.




Figure 12: Sample density has little influence on the reconstruction as the disc radii automatically adapt.




Figure 13: Adaptive meshing results: The bust was importance-sampled based on local surface curvature. The reconstructed mesh adapts accordingly. The sphere features a very steep density gradient with 10k samples located on the upper hemisphere and 1k samples on the lower hemisphere.

One might argue that equidistantly distributed vertices - the default outcome of our algorithm - are not always a favorable choice for meshed objects. Therefore, we can easily switch to a sample-density adaptive mode, as exemplified in Figure 13. To simulate varying density, the ground-truth mesh was importance-sampled, based on surface curvature. This leads to a higher sample density on the facial features and fewer, sparsely distributed samples on the head or neck region. Now we can locally scale the target-edge-length for each vertex individually, based on the radius of the closest splat discs. Figure 13 (right) shows a more challenging ex-

ample with a very steep density gradient: 10 times more samples are placed on the upper than on the lower hemisphere. The result adopts this change in density with coarser geometry on the bottom and fine geometry on the top.

3.1. Comparisons

Besides validation against common point cloud challenges or the qualitative discussion in the upcoming section, we provide a numerical comparison to a variety of state-of-the-art reconstruction techniques, evaluated on over 2000 randomly selected objects from the Thingi10K data set [ZJ16]. A comparison to the most closely related technique, *Competing Fronts* [SLS*06], is not included, since there is no implementation available and the authors could not provide reference results. Ground truth objects are scaled-down and centered to fit in a $[-1:1]$ cube. Oriented point clouds are sampled uniformly from the source objects, where the number of samples computes as object surface area $\times 1000$. Where possible, we specified $l_t \approx 0.025$, except for *IFAM* where we set the target vertex count to half the amount of the input point cloud samples, to actually produce a decent result. Other parameters were left on auto- or the default settings respectively. Total failure cases ($<5\%$), where the applications crashed and produced no output, could not be taken into account.

Results of our numerical comparison are listed in Table 1 with means and medians of different error measures, respectively. In the average HD, our method is only outranked by two other approaches but supersedes (or is close to) the competition in all other measures. This small drawback can be attributed to a limitation of our method, namely very, very thin geometry. Such a case is included in Figure 20 with object 1489590: With geometry, becoming thinner than the target-edge-length, the mesh growth comes to a halt, as there are no further refinements. Portions of the point cloud remain unexplored, thus result in rather high HD errors. A heuristic, indicating how many point cloud samples are actually in use for hull projections, could potentially improve on this matter. The other models in Figure 20 and 21 exemplify, why some procedures fall behind in our numerical comparison: While flat or curved surfaces are often reconstructed quite well, edges on feature lines or sharp geometry are smoothed out. This especially shows on CAD-like models with sharp edges, e.g. for object 200080, our result is the only one that accurately reconstructed the threads on the screw.

3.2. Other Ballooning Methods

Whereas the numerical comparisons are mostly focused on established state-of-the-art meshing procedures, we also want to point out clear conceptual distinctions between our and related ballooning approaches.

- **Competing Fronts (CF)** [SLS*06] also uses a deformable mesh complex, iteratively approaching the target shape. However, CF strictly grows from within the target as vertices only move along positive normal directions and get frozen once close enough to the point cloud. Our mesh remains flexible throughout the optimization

as vertices move and align individually. Further, our method generalizes ballooning and shrink-wrapping, as the mesh may grow from within the target, but also enclose on it from the outside. As shown in Figure 23, our mesh’s growth direction is robust enough to even overcome a misaligned initialization. In order to maintain high-quality mesh structures, CF employs remeshing operations [BK04] every 6-18 iteration cycles. In our iteration, corrective operations are performed in every cycle but targeted at the affected mesh regions. The increase of an object’s genus is approached in a similar defensive manner, as both methods replace opposing geometry with a triangulated tunnel before the mesh self-intersects. See step V. of Section 2 for more details. CF finalizes the optimization with a dedicated projection step to guarantee ϵ -close reconstruction results, which, however, is invariant of the point cloud tangent space, thus does not recover any sharp features. Our optimization is designed to automatically transition from mesh growth to a projection scheme on an individual vertex basis. Further, in our projection, the vertices are explicitly drawn towards intersections of tangent planes, thus faithfully reconstruct sharp corners and edges.

- **Cooperative Evolutions (CE)** [LL20] is based on *two* mesh structures, approaching the target shape from the interior *and* exterior simultaneously. Once both hulls are close enough, individual vertices from both hulls have to be matched to each other and the final geometry is extracted from the space between both meshes. Objects of a higher genus are also possible: Self-intersections are identified, the affected geometry is removed, again stitched together, smoothed, and refined. The mesh evolution in CE is based on a signed distance field, recomputed every iteration cycle, and mixed with a normal-distance field. In contrast, our mesh vertices are linked to the individual, fixed-size sets of nearest samples, easily updated with sets from their 1-ring neighbor vertices, respectively. Mesh integrity in CE is, as in CF, maintained with dedicated global remeshing [BK04]. Further, CE is invariant of the point cloud orientations, thus solely relies on the distance fields. As shown in our SDF experiments (Section 2.1), our method can also easily operate on distance fields. However, this also comes with the same drawbacks and limitations of CE, as our method is then also no longer able to recover sharp edges.




Figure 14: Comparing results of the learning based approach Point2Mesh (P2M) [HMGCO20] to ours.

		HD		RMSE		SAD		TVD		NAD	WT
[BMR*99, Dig14]	BPA	0.057 155	0.052 247	0.002 710	0.002 693	-0.093 504	-0.050 309	0.065 322	-0.014 808	0.151 215	0.000 000
[DMSL11, Dig15]	SSM	0.089 543	0.090 939	0.004 422	0.004 029	-0.470 502	-0.425 178	0.028 261	-0.447 994	0.220 108	0.001 329
[JTPSH15]	IFAM	0.139 415	0.129 359	0.004 021	0.004 036	-0.413 034	-0.371 409	0.127 780	-0.379 604	0.197 236	0.000 000
[STJ*17]	OSR	0.124 078	0.101 473	0.007 759	0.006 079	-0.133 871	-0.016 991	0.025 842	-0.001 302	0.287 676	0.517 084
[CT12]	SSDR	0.029 466	0.025 658	0.001 922	0.001 747	-0.031 290	-0.022 540	0.115 445	-0.003 310	0.092 074	0.931 677
[KBH06, KH13]	SPSR	0.034 983	0.033 962	0.002 623	0.002 557	-0.043 226	-0.031 528	0.010 974	0.000 109	0.128 020	0.998 227
	Ours	0.040 563	0.019 009	0.001 269	0.000 773	-0.015 890	0.000 910	0.003 979	0.000 570	0.040 147	1.000 000

Table 1: Results of multiple reconstruction techniques, compared to the ground truth of over 2000 objects from the Thing10K data set [ZJ16], as mean and median. The errors (best close to 0, except WT best close to 1) include the symmetric Hausdorff-Distance (HD) measured with 200k samples, the root-mean-squared-error (RMSE) measuring the closest-point distance of result faces to the ground truth, the surface-area-deviation (SAD) and total-volume-deviation (TVD) between reconstruction and ground truth, as well as the amount of watertight (WT) results, all given in percent, except for the normal-angle-deviation (NAD) which is given in radians. See Figures 20 and 21 for examples.

• **Point2Mesh (P2M)** [HMGCO20] is a state-of-the-art learning-based approach, iteratively shrink-wrapping a convex hull mesh around a target point cloud. The individual vertex displacements are the results of a learning self-prior network, which does not penalize manifold violations. Therefore, P2M requires heavy manifold reconstructions [HSG18] every few iteration cycles to actually maintain a valid surface. Adapting the genus within the optimization is not possible. However, P2M may also start on a given mesh with the correct genus, i.e., a coarse alpha shape [EM94] or Poisson [KH13] mesh. This is also natively supported by our method, moreover, our algorithm may even start on separated geometry as shown in Figure 19 and 23. Comparisons featured in Figure 14 again exemplify the strengths of our approach, i.e., generating a uniform mesh structure and reconstructing sharp features and fine details. This can be seen, especially on the edges of the **G** as well as on the toes and epoccipitals on the triceratops’ frill. The target-edge-lengths for our results were set to the average edge length of the P2M meshes, respectively. On the same hardware (GTX 1080 Ti) and starting from the same convex hull mesh, P2M finished the given 6000 iterations for the triceratops within 70 minutes, whereas ours performed 1500 iterations and terminated after 108 seconds.




Figure 15: Error comparisons with related ballooning approaches, visualizing distance errors (top), normal deviations (bottom) and listing the number of faces (#f).

Ballooning Results Figure 15 provides a qualitative comparison against results of CF, CE, and P2M. The visualization shows distance errors and normal angle deviations between reconstruction and input point cloud. Although the CF and CE results have a higher mesh resolution than ours, they still struggle to capture the rounded ridges and valleys of the bunny. With the original point cloud (362k samples) and a memory-limited size of 65.5k faces, the official P2M code totally maxed out the 24 GB of memory on a Titan RTX card, performed 30k iterations, ran for 237 h but the result is still riddled with artifacts. This is the best result we were able to generate. Even without special attention to sharp features, our method still provides the most accurate reconstruction of round and organic shapes.

3.3. Discussion

With our method, we achieve the same ϵ -closeness to the target shape as *Competing Fronts* [SLS*06], which is a guaranteed result of our projection scheme. Figure 17 and 20 visualize magnified deviation errors from the sampled ground truth meshes. An effect that becomes visible here is the mesh’s tendency to exaggerate curved regions. This is caused by the tangent projection, which, on concave regions intersect inside and for convex regions intersect outside of the mesh. By design, mesh vertices are drawn towards intersections of tangent planes, thus resulting in vertices slightly in or outside of the actual hull. However, this is a small trade-off in our bilateral sample-weighting scheme, which generally allows for a robust and accurate reconstruction of smooth regions and sharp edge features respectively, even under the presence of noise. The listed numerical values in Figure 17 and Table 1 for HD and RMSE play in the same order of magnitude as compared state-of-the-art methods. Figure 18 compares reconstructions of the Tanagra high precision 3D scan [DAL*11]. While the PSR [KBH06] result is again quite smoothed out, the MLS [FCOS05] filtered BPA [BMR*99] result features better details. Nevertheless, our mesh equals or even surpasses the sharpness and fine-detail quality of the SSM [DMSL11] result. The zoom-ins shown in Figure 16 demonstrate the achievable mesh quality. Whereas other methods’ mesh structure either directly depends on the input points (BPA, SSM), smooth out details for the sake of regularity (IFAM, OSR), or triangulate octree-cells (SSDR, SPSR), our results natively approximate Delaunay-like triangulations.

Section 2.1 introduced the ability to adapt our method with little effort to various other input scenarios, giving examples for high-quality point cloud reconstruction, remeshing, volume- and CT-scan meshing, or surface reconstruction from signed distance fields which are currently popular in neural surface representations. Further examples for the versatility of our algorithm are shown in Figure 6, where *Spot* is given as a quad-meshed target, enclosed by a coarse initial mesh, which automatically assimilates its shape by *shrink wrapping*. Swapping target and init mesh leads to the inverse operation: The only parameter to be adapted is a larger target-edge-length, automatically coarsening the mesh as small edges incrementally collapse.

Examples in Figure 19 show reconstructions of the Bunny given as point cloud but starting from different initialization meshes. The initial mesh on the left is a coarse voxelization of the point cloud volume. As this initial hull already fills out most parts of the target




Figure 16: Comparison on mesh uniformity with a zoom-in on one of the nobs of object 101634 from Figure 20.




Figure 17: Errors analogous to Table 1 but multiplied by 100 for better readability. Hull projections on tangent planes cause vertices to exaggerate on concave (vertex **inside**) and convex (vertex **outside**) curved regions. This effect is not as prominent on the horse as it is reconstructed from an SDF without hull projections.

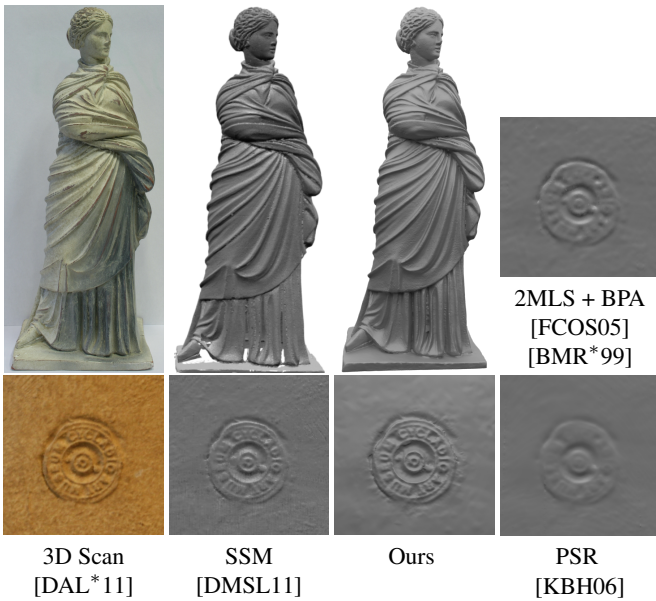


Figure 18: Reconstructions of the Tanagra 3D scan with PSR, MSL & BPA, SSM and ours. The reconstructed mesh is significantly smoother in appearance along the folds of the toga while the inscription (on the back of the statue) is much better readable. Image Source [DMSL11]

shape, it is quite easy for our algorithm to project vertices into the point cloud while equalizing their distances. In the second case, the algorithm starts from separate objects which individually grow and connect via inter-mesh tunnels as they approach each other. The rightmost example is the default mesh, which automatically grows into the desired shape from within the target point cloud. How-

ever, while the same initialization meshes lead to identical results, the algorithm will not produce canonical results from different init meshes. As shown with the Igea example in Figure 23 (top), the init. mesh does not necessarily have to be placed fully within the target shape. The normal-based orientation for mesh growth direction is robust enough, to cope with miss-aligned starting conditions so that the initial mesh will get *sucked in* and then starts to grow.




Figure 19: Reconstructions of the Bunny [TL14] from three different initialization meshes: Starting from a coarse octree-voxelization, separate geometric objects and the default icosphere.

Performance comparisons to other related work might yet not be very meaningful, as they are often single-core CPU implementations and ours a mixture of Python and Cuda code. All vertex-based operations in our pipeline are easily parallelized, thus executed on the GPU. Single-threaded operations on the mesh data structure only interleave the fast optimization steps. Therefore, our algorithm is able to successfully finish hundreds of iterations within seconds. Full reconstructions of larger results, however, play in the same order of magnitude as competing techniques, ranging from seconds to a few minutes. A speedup can be gained, e.g., from tailored input like coarse voxelizations, as shown in Figure 19 (left), or well placed initial geometry, as shown in Figure 23 (bottom), where separate initial meshes grow in parallel until fused to one data structure via inter-mesh tunnels.

4. Conclusion

In this work we present a mesh construction/remeshing algorithm that is easy to implement but still checks all the important boxes of the state-of-the-art in this well-researched field: ϵ -close surface reconstructions of point clouds and other (re)meshing tasks, e.g., volumetric functions or CT-scans, robustness to noise or outliers, and support for higher genus object topology with a guaranteed closed, watertight manifold surface. Moreover, we can note several improvements over similar approaches: Instead of global subdivisions our targeted mesh refinements avoid unnecessary increases of mesh resolution or repeated remeshing steps. The mesh resolution may adapt to the given sample density or distribute vertices uniformly. With a minimal parameter space (target-edge-length or closeness-threshold) to be specified, the process robustly completes fully automated, without the need for different reconstruction modes, finalizing projection, or smoothing passes. This can be attributed to our

concept of a seamless transition between mesh growth and surface assimilation on an individual vertex basis. Our bilateral weighing scheme for surface projections allows for perfectly captured sharp object features as well as smooth areas. This is substantiated with a bulk error-benchmark on a large dataset comparing state-of-the-art reconstruction methods as well as various qualitative comparisons, including other ballooning concepts. Furthermore, we demonstrate the flexibility of our approach to cope with arbitrary input and initialization situations. As our method is still able to adapt to any given target shape, it exemplifies a powerful advancement and generalization of related *ballooning* and *shrink wrapping* approaches.

References

- [AB99] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry* 22, 4 (1999), 481–504. 2
- [ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998), pp. 415–421. 2
- [ACK01] AMENTA N., CHOI S., KOLLURI R. K.: The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications* (2001), pp. 249–266. 2
- [AUGA08] ALLIEZ P., UCELLI G., GOTSMAN C., ATTENE M.: Recent advances in remeshing of surfaces. In *Shape analysis and structuring*. Springer, 2008, pp. 53–82. 1
- [BK04] BOTSCH M., KOBBELT L.: A remeshing approach to multiresolution modeling. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2004), pp. 185–192. 7
- [BL17] BOLTCHIEVA D., LÉVY B.: Surface reconstruction by computing restricted voronoi cells in parallel. *Computer-Aided Design* 90 (2017), 123–134. 2
- [BMR*99] BERNARDINI F., MITTLEMAN J., RUSHMEIER H., SILVA C., TAUBIN G.: The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics* 5, 4 (1999), 349–359. 2, 8, 9, 11, 12
- [Bre65] BRESENHAM J. E.: Algorithm for computer control of a digital plotter. *IBM Systems journal* 4, 1 (1965), 25–30. 4
- [BTS*17] BERGER M., TAGLIASACCHI A., SEVERSKY L. M., ALLIEZ P., GUENNEBAUD G., LEVINE J. A., SHARF A., SILVA C. T.: A survey of surface reconstruction from point clouds. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 301–329. 1, 5
- [CGAY13] CAMPOS R., GARCIA R., ALLIEZ P., YVINEC M.: Splat-based surface reconstruction from defect-laden point sets. *Graphical Models* 75, 6 (2013), 346–361. 3
- [Cra20] CRANE K.: Keenan’s 3D Model Repository, 2020. <http://www.cs.cmu.edu/~kmcraane/Projects/ModelRepository/>. 5
- [CSAD04] COHEN-STEINER D., ALLIEZ P., DESBRUN M.: Variational shape approximation. In *ACM SIGGRAPH 2004 Papers* (2004), pp. 905–914. 1
- [CT12] CALAKLI F., TAUBIN G.: Ssd-c: Smooth signed distance colored surface reconstruction. In *Expanding the Frontiers of Visual Analytics and Visualization*, Dill J., Earnshaw R., Kasik D., Vince J., Wong P. C., (Eds.). Springer London, 2012, pp. 323–338. URL: http://dx.doi.org/10.1007/978-1-4471-2804-5_18. 8, 11, 12
- [DAL*11] DIGNE J., AUDFRAY N., LARTIGUE C., MEHDI-SOUZANI C., MOREL J.-M.: Farman Institute 3D Point Sets - High Precision 3D Data Sets. *Image Processing On Line* 1 (2011), 281–291. doi: [10.5201/ipol.2011.dalnm_ps](https://doi.org/10.5201/ipol.2011.dalnm_ps). 8, 9
- [Dig14] DIGNE J.: An Analysis and Implementation of a Parallel Ball Pivoting Algorithm. *Image Processing On Line* 4 (2014), 149–168. <https://doi.org/10.5201/ipol.2014.81>. 8, 11, 12
- [Dig15] DIGNE J.: An Implementation and Parallelization of the Scale Space Meshing Algorithm. *Image Processing On Line* 5 (2015), 282–295. <https://doi.org/10.5201/ipol.2015.102>. 8, 11, 12
- [DMSL11] DIGNE J., MOREL J.-M., SOUZANI C.-M., LARTIGUE C.: Scale space meshing of raw data point sets. In *Computer Graphics Forum* (2011), vol. 30.6, Wiley Online Library, pp. 1630–1642. 2, 8, 9, 11, 12
- [DQ01] DUAN Y., QIN H.: Intelligent balloon: a subdivision-based deformable model for surface reconstruction of arbitrary topology. In *Proceedings of the sixth ACM symposium on Solid modeling and applications* (2001), pp. 47–58. 2, 4
- [EM94] EDELSBRUNNER H., MÜCKE E. P.: Three-dimensional alpha shapes. *ACM Transactions on Graphics (TOG)* 13, 1 (1994), 43–72. 2, 8
- [FCOS05] FLEISHMAN S., COHEN-OR D., SILVA C. T.: Robust moving least-squares fitting with sharp features. *ACM transactions on graphics (TOG)* 24, 3 (2005), 544–552. 5, 6, 8, 9
- [HBT*13] HELMSTAEDTER M., BRIGGMAN K. L., TURAGA S. C., JAIN V., SEUNG H. S., DENK W.: Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature* 500, 7461 (2013), 168–174. 5
- [HMGCO20] HANOCKA R., METZER G., GIRYES R., COHEN-OR D.: Point2mesh: A self-prior for deformable meshes. *ACM Trans. Graph.* 39, 4 (July 2020). URL: <https://doi.org/10.1145/3386569.3392415>, doi: [10.1145/3386569.3392415](https://doi.org/10.1145/3386569.3392415). 2, 7, 8
- [HSG18] HUANG J., SU H., GUIBAS L.: Robust watertight manifold surface generation method for shapenet models. *arXiv preprint arXiv:1802.01698* (2018). 8
- [JTPSH15] JAKOB W., TARINI M., PANOZZO D., SORKINE-HORNUNG O.: Instant field-aligned meshes. *ACM Trans. Graph.* 34, 6 (2015), 189–1. 2, 8, 11, 12
- [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, 8, 9, 11, 12
- [KCRH20] KAZHDAN M., CHUANG M., RUSINKIEWICZ S., HOPPE H.: Poisson surface reconstruction with envelope constraints. In *Computer Graphics Forum* (2020), vol. 39.5, Wiley Online Library, pp. 173–182. 2, 11, 12
- [KH13] KAZHDAN M., HOPPE H.: Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)* 32, 3 (2013), 1–13. 2, 8, 11, 12
- [KH19] KAZHDAN M., HOPPE H.: An adaptive multi-grid solver for applications in computer graphics. In *Computer Graphics Forum* (2019), vol. 38.1, Wiley Online Library, pp. 138–150. 2, 11, 12
- [Kob00] KOBBELT L.: $\sqrt{3}$ -subdivision. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), pp. 103–112. 3
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics* 21, 4 (1987), 163–169. 2
- [LL20] LU W., LIU L.: Surface reconstruction via cooperative evolutions. *Computer Aided Geometric Design* 77 (2020), 101831. 2, 4, 5, 7
- [MON*19] MESCHEDER L., OECHSLE M., NIEMEYER M., NOWOZIN S., GEIGER A.: Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (2019). 2
- [PFS*19] PARK J. J., FLORENCE P., STRAUB J., NEWCOMBE R., LOVEGROVE S.: DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 165–174. 2
- [SCC11] STANCULESCU L., CHAINE R., CANI M.-P.: Freestyle: Sculpting meshes with self-adaptive topology. *Computers & Graphics* 35, 3 (2011), 614–622. 4




Figure 20: Randomly selected objects (IDs on the right) from the Thing10K [ZJ16] data set serve as ground truth. All objects are scaled down to fit in a $[-1:1]$ cube and are uniformly sampled. The number of samples computes as surface area \times 1000. 89914 is a thin hollow shell and therefore not as simple as it may appear. 1489590 is a failure case, discussed in Section 3.1, of our approach where the geometry is too thin for the mesh to grow.




Figure 21: The same results as from Figure 20 with visualized surface normal errors. Many methods perform well on curved and flat regions, but sharp edges as on the CAD-like objects are often not recovered very well.




Figure 22: Objects of higher genus are reconstructed, using intra-mesh tunnels. Black arrows indicate normal growth process. Proximity checks trigger tunneling operations on close (but not directly adjacent) geometry, shown in blue. The iteration continues normally afterwards.




Figure 23: Top: The mesh growth orientation is robust enough to overcome miss-aligned initialization meshes. The Hilbert curve: Assimilated from one (mid row) starting position and from eight (bottom) individual positions, connected via inter-mesh tunnels.

[SLS*06] SHARF A., LEWINER T., SHAMIR A., KOBBELT L., COHEN-OR D.: Competing fronts for coarse-to-fine surface reconstruction. In *Computer Graphics Forum* (2006), vol. 25.3, Wiley Online Library, pp. 389–398. 2, 4, 7, 8

[STJ*17] SCHERTLER N., TARINI M., JAKOB W., KAZHDAN M., GUMHOLD S., PANOZZO D.: Field-aligned online surface reconstruction. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 77. 2, 8, 11, 12

[TL14] TURK G., LEVOY M.: The Stanford 3D Scanning Repository, 2014. <http://graphics.stanford.edu/data/3Dscanrep/>. 6, 9

[VPBM01] VLACHOS A., PETERS J., BOYD C., MITCHELL J. L.: Curved pn triangles. In *Proceedings of the 2001 symposium on Interactive 3D graphics* (2001), ACM, pp. 159–166. 3

[WK04] WU J., KOBBELT L.: Optimized sub-sampling of point sets for surface splatting. In *Computer Graphics Forum* (2004), vol. 23.3, Wiley Online Library, pp. 643–652. 3

[ZJ16] ZHOU Q., JACOBSON A.: Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797* (2016). 7, 8, 11, 12