Eurographics Symposium on Geometry Processing 2010
Olga Sorkine and Bruno Lévy
(Guest Editors)

Volume 29 (2010), Number 5

Signing the Unsigned: Robust Surface Reconstruction from Raw Pointsets

Patrick Mullen[1] Fernando de Goes[1] Mathieu Desbrun[1] David Cohen-Steiner[2] Pierre Alliez[2]

[1] Caltech [2]INRIA Sophia Antipolis - Méditerranée

Abstract

We propose a modular framework for robust 3D reconstruction from unorganized, unoriented, noisy, and outlier-ridden geometric data. We gain robustness and scalability over previous methods through an unsigned distance approximation to the input data followed by a global stochastic signing of the function. An isosurface reconstruction is finally deduced via a sparse linear solve. We show with experiments on large, raw, geometric datasets that this approach is scalable while robust to noise, outliers, and holes. The modularity of our approach facilitates customization of the pipeline components to exploit specific idiosyncracies of datasets, while the simplicity of each component leads to a straightforward implementation.

1 Introduction

Surface reconstruction from measurements remains one of the most important concerns in geometry processing. While technological advances on sensors and scanners have greatly increased the availability of very detailed geometric measurements, current datasets are increasingly defect-ridden for several reasons: sensors are evolving from contact to contact-free and from short to long range, commodity scanners become cheaper but with higher levels of uncertainty, 3D information is increasingly inferred directly from photo collections, and practitioners often resort to large series of low-cost acquisitions instead of one accurate but expensive acquisition. Consequently, the need for reconstruction methods that are robust to noise and outliers is growing steadily. At the same time datasets are getting larger, so methods requiring advanced solvers or large amounts of memory do not scale to the level needed for iron-clad industrial applications.

1.1 Previous Work

While a large number of reconstruction methods have been proposed over the years, few are robust to noise and outliers. We now review the most notable exceptions, distinguishing the methods assuming attributes (such as oriented normals) from the methods able to process raw datasets.

Reconstruction of pointsets with attributes. Recent progress in scanning technology has led to some sensors providing not only points on a 3D object, but also attributes such as oriented normals, lines of sight, or confidence in the measurements. If reliable oriented normals are available, there now exist a number of mature approaches which can deal gracefully with noise, variable sampling, and holes. A common approach involves computing an approximate signed distance function to the inferred surface, either through a global variational formulation [CBC*01] or through local fitting of, for example, low degree implicit functions [OBA05]. Variational formulations have been proven more robust to noise, but recent attempts at fitting smooth local implicit functions [NOS09] show promise. Another popular approach computes an approximate indicator function of the unknown shape [KBH06], offering an efficient solution that is robust to both variable sampling and holes in the data. However, none of the aforementioned methods claim to handle outliers reliably. This issue can, however, be remedied for data coming from (stereo)photogrammetry where reliable lines of sight are available, as it makes outliers easier to classify as such [LPK09]. More generally, the more attributes the practitioner has the more robust the reconstruction can be.

Reconstruction from raw pointsets. Surface reconstruction from raw geometric data has received increasing attention due to the ever broadening range of geometric sensors and vision algorithms that provide little to no reliable attributes. The unavoidable presence of noise and outliers makes the scientific challenge even greater, and any progress in this direction can also directly benefit reconstruction from pointsets with attributes. A common approach to this prob-

lem involves filtering out outliers and inferring attributes before resorting to a reconstruction method as mentioned above. Typically, the data is first oriented (normals are computed) or, equivalently, signed (an inside/outside function is constructed based on the pointset). However, outlier removal often requires an interactive adjustment of parameters. Similarly, finding and orienting normals from raw geometric data can be as hard as reconstructing the whole surface itself: while there are several options to reliably estimate normal directions [MNG04, ACSTD07], robust normal *orientation* is considerable harder as recently reminded in Huang et al. [HLZ*09]. Spectral methods [KSO04] or graph cut approaches [HK06] can disambiguate the inside from the outside of a low-genus sampled object even in the presence of a significant number of outliers. As no smoothness prior is included, such reconstructions usually require significant post-treatment. Variational methods based on generalized eigenvalue problems [WCS05, ACSTD07] are better able to extract smooth isosurfaces from the data; alas, robustness to outliers is lost and their current computational complexity does not scale well with dataset size. Consequently there is a need for a reconstruction algorithm that promises scalability as well as robustness to noise, outliers, and undersampling.

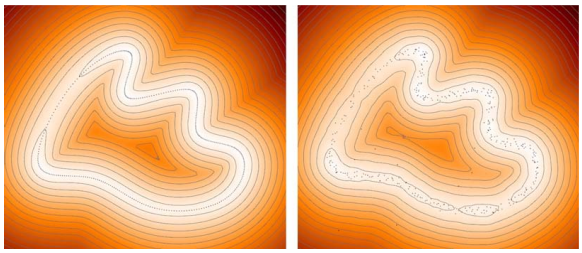

Figure 1: ***Contouring Unsigned Functions.*** *From defect-laden pointsets (here in 2D) , contouring a robust unsigned distance to the data does not lead to a proper surface reconstruction. However, thicker bands succeed in capturing the correct topology.*

1.2 Rationale and Contributions

In parallel to the reconstruction literature, the design of approximate unsigned distance functions which are robust to noise and outliers has recently made significant advances [CCSM09]. However, it has not yet benefited reconstruction as the resulting robust function does not lend itself to reliable surface reconstruction: contouring of an unsigned distance is unreliable, creating numerous geometric and topological artifacts (see Figure 1). In this paper, we propose to *sign* the unsigned distance in order to obtain an implicit function suitable for contouring. We leverage the fact that signing a *function*, rather than the data, can be made more robust by exploiting the property that signing a distance function makes the function smoother (see inset).

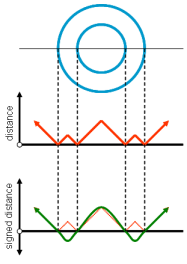

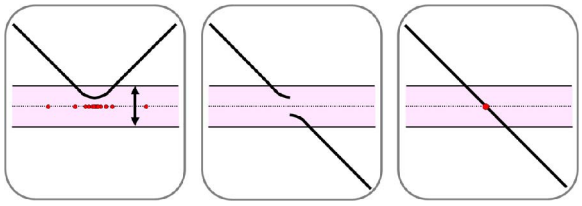

Figure 2: ***Reconstruction pipeline in 1D.*** *Stage 1 (left): We construct an unsigned distance function (robust to noise and outliers, shown in red), and a threshold for the width of the ε-band (in pink). Stage 2 (center): We estimate the sign (±1) of the function, along with a confidence of each estimate. Stage 3 (right): We construct the final signed function through smoothing, taking into account the unsigned function, estimated sign, and confidence.*

Our main contribution is a practical method for efficient reconstruction of closed surfaces from pointsets that are potentially noisy, outlier-ridden, and undersampled. Its simplicity and modularity facilitate customization and extensibility.

1.3 Overview of the Reconstruction Pipeline

Starting with a pointset possibly containing both noise and outliers, our surface reconstruction pipeline involves three main stages (see Figure 2):

- Computing an **unsigned distance** function to the input data. Robustness to outliers and noise is obtained by leveraging the advances in the design of Wasserstein-like metrics, allowing us to reliably identify an ε-band containing the densely sampled areas.
- Computing a global, stochastic **sign estimation** of the distance, first outside the ε-band where rays are traced against the ε-band to infer inside vs. outside, then inside the ε-band by propagating the sign estimates inward. The output of this step is a sign guess for the unsigned distance, along with its confidence ranging from zero to one.
- Smoothing the estimate to compute the **signed distance** through a linear solve to reconstruct a smooth, closed surface. This last step also serves to repair holes.

The remainder of this paper details the reconstruction pipeline stage by stage, before discussing results, limitations, and future work.

2 Robust Unsigned Distance to Data

The first stage in the pipeline involves simultaneously discretizing the domain while computing a robust unsigned distance function to the inferred reconstruction. This distance is then analyzed to determine the thickness ε of a volumetric band that best captures the geometric information available in the data. The ε-band is further refined and a more accurate distance to the pointset is computed inside.

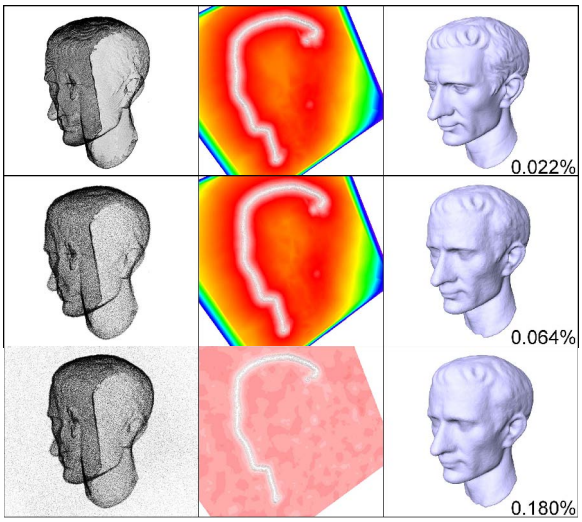

Figure 3: ***Robustness of Unsigned Distance.*** *Top: A* 380*K pointset of the Caesar model, its robust unsigned function displayed in a slice in false colors (K*=15*), and the final reconstruction along with the symmetric Hausdorff distance to the original model given as percentage of the bounding box diagonal. Middle: Same, with significant uniform noise added (K*=30*). Note the stability of the distance function in denser areas, and the change in the function range. Bottom: Same, with* 200*K outliers uniformly added (K*=70*). The stability of the function starts degrading.*

2.1 Outlier and Noise Robust Distance

In order to be resilient to outliers and noise we first compute the unsigned distance to the pointset based on a fast evaluation proposed in [CCSM09]. This approach leverages the notion of distances between *measures* to gain robustness while retaining the usual properties of distance functions including stability and semiconcavity. The distance function from a position x to a pointset is formulated as a minimization of the Wasserstein distance (a variant of the earth-mover's distance) between a delta function at x and a set of measures defined from the pointset. While one can define this distance for arbitrary sample types such as triangle soups, in the particular case of point samples the unsigned distance $d_U(x)$ is simply computed at any location x as

$$d_U(x) = \sqrt{\frac{1}{K} \sum_{p \in N_K(x)} ||x \quad p||^2}$$

where $N_K(x)$ is the set of K nearest neighbors to x, and K acts as a tradeoff between robustness and accuracy which may be adapted based on the data. We found choosing K in the 12 to 30 range was sufficient for all but the most extreme examples, and that the final results were not very sensitive to this choice. Figure 3 illustrates the robustness of this distance to noise and outliers. Note that this robustness is exactly what makes this unsigned function too inaccurate to precisely locate the surface; however, we use it to reliably identify the regions of space *near* the sampled surface within which we will use a more accurate distance.

2.2 Adaptive Domain Discretization

To discretize the unsigned distance we hierarchically triangulate the domain with a 3D Delaunay triangulation [CGA10], refining regions where the distance is small, i.e., likely near the inferred surface. The initial refinement is done in an octree-like manner. First, the eight corners of a loose bounding cube are assigned a depth level ℓ of 0 and added to a queue. We also compute the length of the diagonal H of the box for scaling purposes. Then, until the queue is empty, the next point p is popped out of the queue, and we test whether or not to insert the point in the domain discretization by checking if

$$d_U(p) < \frac{H}{2^\ell}.$$

If the point is inserted, all of its immediate neighbors in the next level of the octree are added into the queue with a level of $\ell+1$. We continue this procedure until the queue is empty or we have reached a predetermined maximum depth level L. Finally, a Delaunay triangulation is computed from the octree points; the mesh produced will be referred to as the coarse mesh $\mathbb{M}$, see Figure 12 (top, middle).

2.3 Automatic ε-Band Width Selection

For reliable sign estimation, we need to find an ε-band that captures the shape of the surface to be reconstructed. To find a good choice for the bandwidth ε we simultaneously analyze the topology of the band and density of input points inside it as a function of ε. We empirically found it sufficient to use the following function:

$$M(\varepsilon) = \frac{C(\varepsilon) + H(\varepsilon) + G(\varepsilon)}{D(\varepsilon)},$$

where C, H and G are the number of components, cavities, and tunnels in the ε-band, and D is the density of input points in the ε-band. To compute the function M, we first sort the nodes of the coarse mesh by distance value and use this to split the range of ε into 200 intervals each containing an equal number of coarse mesh nodes. We next bucket-sort the input points, along with the edges, faces, and tetrahedra of the coarse mesh into these intervals. We use a union-find algorithm to compute the evolution of the number of connected components as we increase the distance threshold ε. For each interval we also compute the volume of the ε-band along with its Euler characteristic $\chi = |V| \quad |E| + |F| \quad |T|$ where $|V|$, $|E|$, $|F|$ and $|T|$ are the number of vertices, edges, faces, and tetrahedra within the band respectively. Additionally, we compute the density of input points inside the band, given by the number of input points with distance less than ε divided by the volume of the band. We finally use the union-find algorithm in *reverse order of distance value* to get now the number of connected components of the complement of the band, from which we deduce the number of cavities in the band. This last calculation allows us to compute the

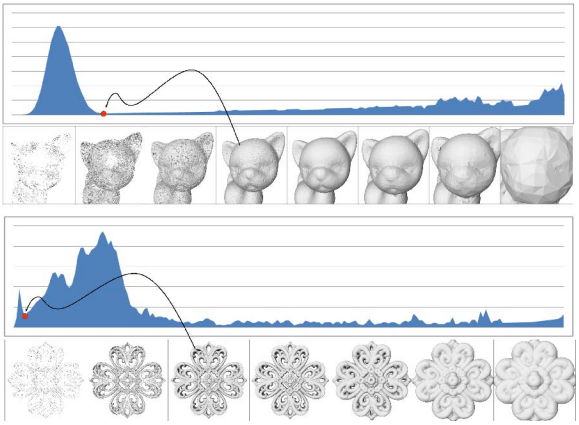

Figure 4: ***Automatic ε-Band width selection.*** *$M(\varepsilon)$ is plotted as a function of ε for two models with very different genus. The red dot indicates the first local minimum after the first local maximum value. As seen in the isocontours of varying ε's, the band corresponding to this chosen ε does well at capturing the correct topology while remaining as thin as possible to best preserve details.*

genus G of the band as:

$$G(\varepsilon) = C(\varepsilon) + H(\varepsilon) - \chi(\varepsilon).$$

We found that regardless of input topology, noise, and outliers, plotting M as a function of ε consistently resulted in a plot similar to those seen in Figure 4. The important feature of these curves is a bump at the start followed by a steady increase for trivial topology, and other bumps for complex topology. This first bump corresponds to values of ε that are slightly too small to enclose the surface we wish to reconstruct, resulting in a band pierced by a large number of cavities and tunnels. As we further increase ε the spurious topology disappears, after which the function slowly rises as the density decreases; hence, we are seeking the smallest ε immediately after this initial bump, as it will be the thinnest band that captures the best guess at the topology. In practice we subsample this curve, smooth it slightly to remove the noise, find the first local minimum after the first local maximum, and choose the ε corresponding to this point.

2.4 Improving Distance Inside the ε-Band

Once the ε-band has been identified, we can improve the estimation of distance to the ideal reconstruction inside the band: this thin band is now safely devoid of outliers, so a more accurate, noise-robust distance can be devised. We use stochastic sampling and PCA to bring robustness to noise: for any location in the band we find the K nearest neighbors in the band, but we now take m random subsets of size β and fit the best plane to each; from the m planes, we pick the one that achieves the best fit of its subset and use the distance between this plane and the location. We found that in our examples choosing $\beta = \frac{3}{4}K$ and $m = \frac{1}{2}K$ was sufficient for handling a variety of noise levels. If prior knowledge on the noise level of the data is known, other parameters or even other distance evaluations can of course be used to improve accuracy. This distance function is computed and stored on the nodes of a finer mesh that will be referred to as $\mathcal{M}$, obtained by simply refining $\mathbb{M}$ inside the ε-band through Delaunay refinement until it contains twice as many vertices as $\mathbb{M}$, to allow for denser sampling in these crucial regions.

3 Estimating the Sign of an Unsigned Distance

We would now like to convert the robust (but unsigned) distance stored on the vertices of our adaptive triangulation of space into a signed distance. One could devise a variational approach to sign d_U into $\lambda(x)d_U(x)$ by solving for a sign field $\lambda(x)$ minimizing the energy

$$E(\lambda(x)) = \int_\Omega |\lambda(x)d_U(x)|_S^2 + \alpha(\lambda^2(x)-1)^2\,dx \qquad (1)$$

where Ω is the domain, $|.|_S$ is some smoothness norm (such as Sobolev) and α is a parameter enforcing how well the signed function should fit the unsigned one. However, this energy is quartic in λ and results in a nonconvex optimization unsuitable for the scalability we desire. Moreover, we would rather rely more on the unsigned distance in clearly sampled regions, while allowing new zeros to appear in undersampled regions in order to fill holes in the data. Hence, finding a suitable spatially constant α may be impossible. As we explain next, we reach robustness and scalability by *geometrically estimating the sign* while *identifying regions of uncertainty*. As in the previous stage, we compute an estimate outside of the ε band on the coarse mesh $\mathbb{M}$ first, and use this guess to propagate the information into the band on the fine mesh $\mathcal{M}$.

3.1 Coarse Estimation through Ray Shooting

Testing whether a point is inside or outside of a *closed* surface can be accomplished by shooting a ray from the point and counting the number of times it intersects the surface: if it intersects an odd number of times the point is inside, while an even number implies that the point is outside. However, one can not apply this procedure directly in our case: we do not have a surface but a pointset that is noisy and contains numerous outliers and holes, making the intersection count of a ray against the input set unreliable at best. Instead, we perform intersection tests against the ε band of the unsigned function stored on $\mathbb{M}$ rather than the input samples, and count the number of times rays intersect this band. While the ε band thickness was chosen to best capture the expected topology, it may still contain holes in undersampled regions. Hence, rather than shooting a single ray for each point, we take a stochastic approach and try several different rays for each point. The added benefit of this stochasticity is that the agreement between rays from the same point provides a confidence of the inside/outside guess at that point. A final necessary detail arises from the fact that we are testing for intersections with a band rather than a surface, and therefore rays that pass near or through the surface almost tangentially (a case referred to as shallow hits) may enter and exit the band

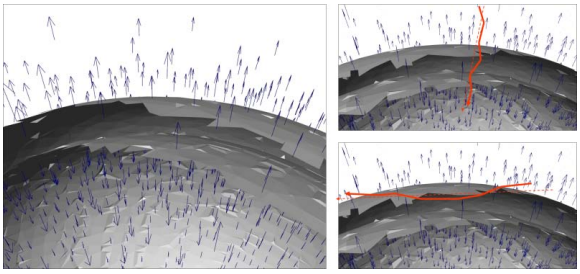

Figure 5: ***Ray shooting.*** *Left: Boundary of the ε-band depicted in gray through isosurfacing; the arrows show the gradient direction of the unsigned function. Right, top: A non-shallow ray (direction in dotted red, edge-based polyline in solid red) is shot from outside the ε-band; it intersects the band boundary twice, with entry and exit vertices having gradients of opposite orientations. Right, bottom: A shallow ray can also intersect the ε-band boundary twice, but its entry and exit vertices have gradients of similar orientations.*

on the same side of the surface, leading to an incorrect intersection count (see Figure 5). We detect these shallow hits using the gradient of the unsigned function at each pair of entrance/exit events along the ray: if the dot product of the gradient at the entrance point with that at the subsequent exit point is ever positive we consider the hit shallow and discard the two intersections. This procedure greatly reduces noise in the sign estimation near the band boundary. We show an example output of this sign estimation in Figure 6.

Sign and Confidence. In practice we visit each vertex of $\mathbb{M}$ outside the ε-band (i.e., with unsigned function greater than ε) and pick a random direction (with uniform distribution) in which to shoot the ray. We then walk along the edges from vertex to vertex for efficiency, always picking the next closest vertex in the chosen direction until we have reached the boundary of the domain (see Figure 5). As we visit each vertex v_i, we count the number of pairs of times we entered and exited sets of vertices with unsigned function less than the chosen ε. We store on the original vertex whether the count was even or odd and repeat the process, shooting a total of r rays for each vertex. Once this is complete, we assign an initial guess of the sign $\bar{\lambda}_i$ to each vertex equal to 1 or -1 if more rays returned an even or odd count respectively. In addition to the sign, we also assign a confidence $c_i = 2\max(e,o)/r - 1$ ranging from 0 to 1 indicating the level of agreement among the rays, where e/o are the number of rays giving an even/odd count respectively. While we obtain very reliable guesses and confidence both away from and near the densely sampled areas, it is an important feature that the confidence drops to 0 near regions missing data: it will allow high-quality reconstruction in dense sampling areas, and hole-filling through smoothing in sparse regions.

Smoothing the Sign Estimate. Notice that outside the band the sign should be mostly constant except in regions missing data (holes), and hence we may aggressively smooth it to help avoid artifacts due to incorrect sign guesses near the band boundary. While potentially many approaches may be used for this step, the smoothing term we will use in the final stage (Section 4) is perfectly appropriate for this task. Hence, we simply perform a few steps of the final stage restricted to the outside of the band and extract the sign from the smoothed function. Note that smoothing only occurs outside of the ε-band to help with signing, but does not smooth features of the reconstructed surface.

3.2 Sign Propagation Inside the ε-Band

Once the sign and confidence have been estimated *outside* of the band, we would now like to "fill in" the guess (and confidence) for vertices inside the band as much as possible. This final step makes the initial signing guess cross zero close the the center of the band, resulting in less unwanted smoothing in the final stage of the pipeline. Conceptually, we would like to start at the boundary of the band and propagate the sign and confidence by copying known values to adjacent vertices along the gradient of the distance, i.e., to neighbors with a smaller distance. The implementation is simplified by sorting all vertices in the band by their distance. Starting with the vertex with the largest distance, we look at all of its neighbors with a larger distance (or equivalently, all of its neighbors with assigned signs and confidences). If all such neighbors have the same sign estimation, we assign this estimate to the vertex along with a confidence equal to the maximum confidence of its neighbors. If, on the other hand, some neighbors disagree on the sign, then the confidence is set to 0 to signal that the guess is unreliable. Note that if any neighbors had already been assigned zero confidence, this is considered a disagreement and the vertex is also assigned 0 confidence. Also, if the vertex is at a local maximum of the distance function, we assign it a new priority/distance equal

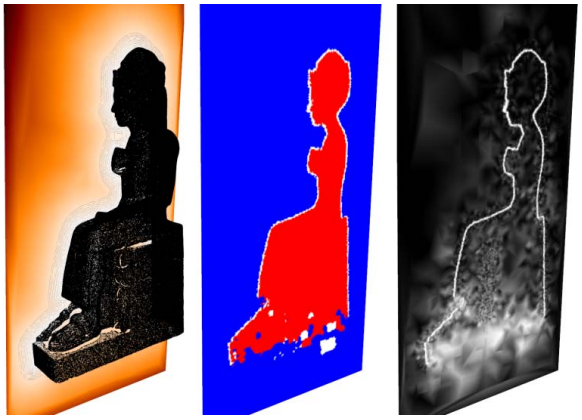

Figure 6: ***Sign guess and confidence.*** *Left: an input pointset (Ramses statue) and its color-coded unsigned distance in a slice. Middle: sign guess (white represents the ε-band or zero confidence where we cannot stochastically decide on a sign). Right: confidence function (black for high confidence, white for low confidence).*

to the average of its neighbors before reinserting it into the sorted list. We found this simple procedure to extend the sign guess to the inside of the ε-band to be reliable and efficient.

4 Solving for a Signed Implicit Function

Given the sign guess $\bar{\lambda}(x)$ and confidence $c(x)$ we solve for $\lambda(x)$ minimizing the energy

$$E(\lambda(x)) = \int_D \left[|\lambda(x) d_U(x)|_S^2 + W(c(x))(\lambda(x) \quad \bar{\lambda}(x))^2 \right] dx$$

where W is a positive weighting function mapping the local confidence of $\bar{\lambda}(x)$ to a weight for the fitting term. This energy can be seen as a quadraticization of the previous quartic energy in Eq. (1), additionally leveraging the confidence obtained in the second stage. The weighting function will yield faithful reconstruction in confident regions, while allowing the smoothing term to take over when necessary, such as in undersampled regions to fill holes. Modifying the weighting function W to incorporate a $d_U(x)^{\ 2}$ and writing $f(x) = \lambda(x) d_U(x)$ the energy becomes

$$E(f(x)) = \int_\Omega \left[|f(x)|_S^2 + W(c(x))(f(x) \quad \bar{\lambda}(x) d_U(x))^2 \right] dx$$

This quadratic energy in f is easily discretized on the refined mesh $\mathcal{M}$, and its minimizer satisfies the linear equation

$$(S + W_c)\, F = W_c\, \bar{\Lambda}\, \bar{F}$$

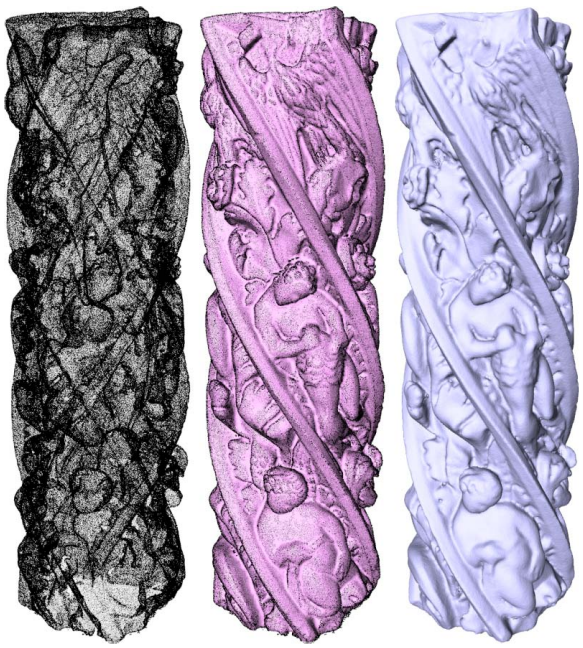

Figure 7: ***Column Scan.*** *Our robust signing directly reconstructs a surface for this challenging dataset which presents difficulties for conventional normal orienters due to the many folds and details. The reconstructed surface is shown with the original points (center) and alone (right).*

Figure 8: ***Stone Elephant.*** *Pointset from a Minolta laser scanner (left) and reconstruction, seen from two angles.*

where now F (resp., $\bar{F}$) is the vector of all the values of f (resp., d_U), $\bar{\Lambda}$ is a diagonal matrix containing the values of $\bar{\lambda}$, W_c is a diagonal matrix of weights computed from c, and S is a sparse matrix representing a linear operator. We found choosing S to be the standard Laplacian matrix yields high-quality results, but other choices, such as the bi-Laplacian, may also be used. We also found that choosing the diagonal entries of W_c to be αc_i was sufficient for good results, leaving the scalar α representing the strength of the fitting as a global user-defined parameter. For all of our examples we chose $\alpha = .1$ (and $\alpha = .01$ for the smoothing performed in Section 3.1).

To solve this linear system we notice that since W_c is a positive diagonal matrix and the mesh is Delaunay (and assuming S is the Laplacian matrix), the matrix in the linear system is symmetric positive definite and diagonally dominant. We thus use Jacobi iterations, a trivially parallelizable process, to efficiently solve for F, generally using 30 iterations.

5 Results

Our implementation follows the pipeline we described in this paper using the CGAL library [CGA10]. Surface extraction was done using Delaunay refinement [BO05], but any isosurface extraction method may be used. Note that rather than contouring the surface with 0 isovalue, we get slightly improved results by taking the isovalue as the median signed distance value over the input points. All timings given were taken on an Intel P4 clocked at 3GHz. Our implementation did not take advantage of the very parallelizable nature of some of the stages, and doing so should increase the efficiency.

A typical example such as the hand in Figure 12, containing 1.8M points and run with a depth of 9, took a total of 420 seconds. Of that time the KNN data structure construction took 14s, the construction of $\mathbb{M}$ took 8.6s, the unsigned function computation took 3.7s, the analysis for finding ε took 31s, the sign guess took 181s, the smoothing outside the band took .3s, the refinement inside the band to get $\mathcal{M}$ took 127s, the sign propagation into the band took 37s, and the final solve including computing the Laplacian matrix and performing 30 Jacobi iterations took 80s. By far the longest

example was the elephant taking 4630s with a depth of 12 to capture the very fine details, while the Caesar model was the shortest taking 350s with a depth of 9. We found the only parameters we needed to play with were the depth of the mesh $\mathbb{M}$, generally taken between 8 and 12 based on the level of detail desired, and K generally ranging from 12 to 30 based on the number of outliers. For our ray shooting we never took more than $r = 20$ rays per point. Examples of parameters used for our figures are given in the following table.

Model	depth	K	#rays	Time (s)
Elephant	12	12	15	4,630
Column	11	15	15	2,350
Indonesian lady	8	30	15	1,270
Hand	9	20	15	420
Caesar	9	15	20	350

On dense, noise-free examples (Figure 8) our approach performs comparably to state-of-the-art reconstruction methods such as Poisson reconstruction [KBH06], although our current code cannot compare in terms of scalability with advanced implementations such as [BKBH07, MPS08]. For such cases when normals are given or can be properly estimated and oriented, the added value of our approach is the simplicity and efficiency of the Jacobi iterations and its trivial parallelization.

Our approach is thus primarily targeted at dealing with the cases where the point set comes with no attributes and the normal orientation and estimation algorithms (e.g., [HDD*92, AB99]) fail because of bad sampling conditions. On the hand model scanned with a Kreon laser scanner for example (Figure 9), the anisotropic and noisy nature of the sampling made it impossible to properly orient the normals, and these failures are reflected in the reconstruction produced by Poisson surface reconstruction. Our approach deals with this dataset without issue.

The flexible lady model (Figure 10) and the column model (Figure 7) are also shown because they challenge conventional normal orienters. More specifically, the MST-based normal orienter [HDD*92] proceeds by constructing a graph over the K nearest neighbors and propagating a seed normal orientation through a minimum spanning tree over the graph. When the graph is disconnected the normals of disconnected regions cannot be properly oriented. Choosing a larger parameter K is a solution but it leads to the connection of nearby surfaces of opposite orientations. Instead, our approach handles these datasets gracefully.

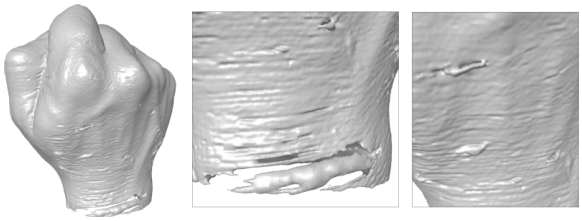

Figure 9: ***Brittleness of Normal Orientation.*** *On the highly anisotropic and noisy hand scan from Figure 12, normal orientation fails at several places which translates into artifacts when Poisson reconstruction [KBH06] is used.*

Figure 10: ***Flexible Lady.*** *Robust reconstruction from irregular sampling. Top: input point set, oriented normals (red segments are normals where the orienter fails) and points with unoriented normals. Bottom: our reconstructed isosurface with and without input points, and output of Poisson reconstruction with same depth.*

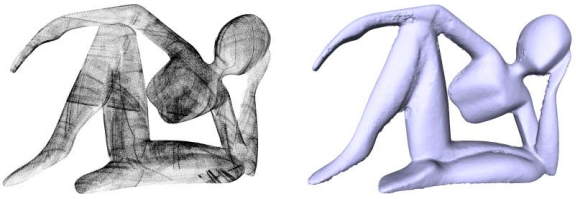

Figure 11: ***Indonesian Lady.*** *Robust reconstruction from a highly non-uniform pointset with no attributes.*

The robustness of our approach to outliers is illustrated on the Caesar head model in Figure 3: the bottom figure has additional noise and 200K outliers, but our outlier-robust method still captures the shape quite well. Also, while the robust unsigned function is not designed to properly deal with nonuniform sampling, our automatic band width computation combined with the use of a more accurate distance inside of the band helps to handle very irregular sampling. Figure 11 illustrates our approach on a model with a non-uniform sampling, taken from [HLZ*09] for comparison.

Limitations. While the unsigned distance function has proven robust to noise and outliers, the measure-based approach does not deal well with structured outliers (i.e., outliers with regular patterns, making them look like features). We see one such example in Figure 13(top), where the automatic selection of the ε-band (which should enclose most

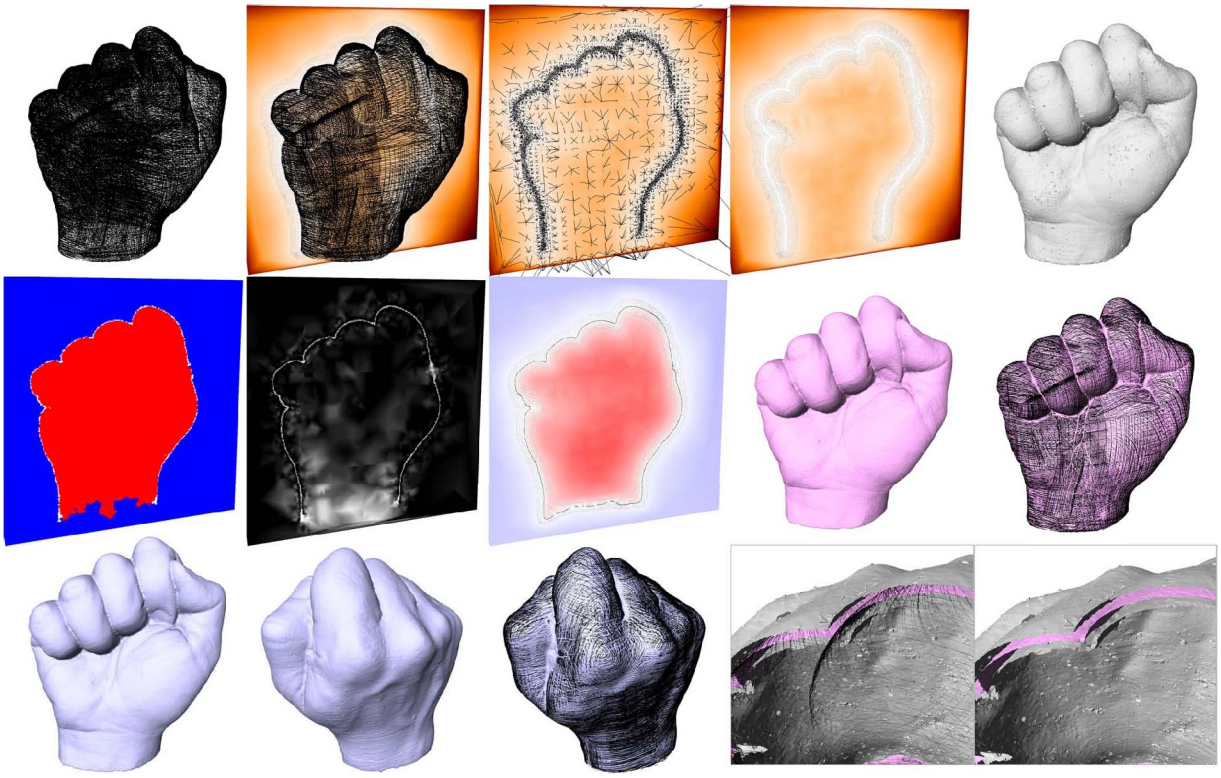

Figure 12: ***Plaster Hand.*** *Data scanned with a Kreon laser scanner mounted on an articulated arm; the* $1.8M$ *point sampling is very anisotropic as it was obtained by manual sweeping of a 1D contact sensor. Top: input point set (with a big hole at the bottom and others due to occlusions between the fingers), point set and 2D cut of unsigned function, same 2D cut with nearby edges of the coarse mesh* $\mathbb{M}$*, same 2D cut alone, and full* ε*-band. Middle: 2D cuts of sign guess (red for inside, blue for outside and white uncertain), confidence (which decreases in the holes), signed function after smoothing, isosurface of the robust unsigned function obtained by marching tetrahedra in the lattice mesh, and same isosurface superimposed with input points. Bottom: views of the reconstructed surface obtained by Delaunay refinement without and with points added, and cut view of the* ε*-band with the reconstructed isosurface of the signed function inside, with and without the input points.*

inliers but not too many outliers) fails. Finally, our approach is not targeted at reconstructing piecewise smooth surfaces as shown in Figure 13(bottom).

6 Conclusion and Future Work

We have presented a modular framework for robust 3D reconstruction from unorganized, unoriented, noisy, and outlier-ridden geometric data. Our approach departs from previous work as we reconstruct a closed surface from the data by signing an unsigned distance to the dataset. We have shown that using information *away* from the surface allows improved robustness in reconstruction. By signing a distance on an entire domain we can reliably handle noise, outliers, and missing data. Tackling the problem sequentially avoids resorting to complex solvers or algorithms, and results in a modular approach that may be customized to better serve more specific applications.

Each stage of the pipeline independently lends itself well to future research. Advances in unsigned distance functions robust to variable sampling density and structured outliers would allow reconstruction from increasingly challenging datasets. Also, our novel geometric and topological band width selection may be refined by better choices of the function M, as well as using a spatially varying width to better handle nonuniform sampling. Sign estimation in itself is an interesting problem, and it may be possible to reliably guess a consistent sign away from a surface without shooting rays across the entire domain, for example, through a multiresolution process. This may be particularly important if attempting to reconstruct very large and detailed surfaces requiring out of core storage. The final stage may be able to make better use of the information in the confidence to perform more advanced hole filling. Finally, we believe it is important to notice that the last two stages provide a general way of robustly contouring unsigned distance functions. While it allowed us to leverage the current advances in robust distance functions, our approach also permits future advances to be immediately incorporated into surface reconstruction.

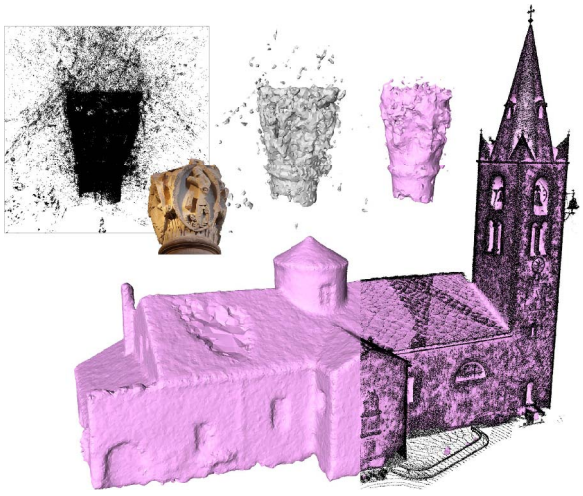

Figure 13: ***Failures.*** *(Top) Pointset of a column capital with a large number of structured outliers (i.e., with systematic errors typically appearing when data come from dense photogrammetry). The original data is shown on the left (along with a photo), followed by the ε-band (middle) and reconstruction (right). Other choices of* ε *similarly fail to capture the right shape, as our algorithm cannot separate the outliers from the data due to the misleading structure of the outliers. (Bottom) Our method also fails to capture sharp edges of piecewise smooth surfaces.*

We hope this encourages further research in unsigned functions with surface reconstruction now as a direct application.

Acknowledgments. The AIM@SHAPE consortium, ISTI-CNR's Visual Computing Lab, Hui Huang (UBC), and B. Vallet (Imagine) provided the 3D models used in this paper. We wish to thank Jean-Philippe Pons for early discussions on outliers. This research was partially funded by NSF grants (CCF-0811373, CMMI-0757106, and CCF-1011944) and by an INRIA associate teams programme with Caltech.

References

[AB99] AMENTA N., BERN M. W.: Surface reconstruction by Voronoi filtering. *Discrete & Computational Geometry 22* (1999). 7

[ACSTD07] ALLIEZ P., COHEN-STEINER D., TONG Y., DESBRUN M.: Voronoi-based variational reconstruction of unoriented point sets. In *EUROGRAPHICS Symposium on Geometry Processing* (2007), pp. 39–48. 2

[BKBH07] BOLITHO M., KAZHDAN M., BURNS R., HOPPE H.: Multilevel streaming for out-of-core surface reconstruction. In *EUROGRAPHICS Symposium on Geometry Processing* (2007), pp. 69–78. 7

[BO05] BOISSONNAT J.-D., OUDOT S.: Provably good sampling and meshing of surfaces. *Graphical Models 67*, 5 (2005), 405–451. 6

[CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3D objects with radial basis functions. In *SIGGRAPH* (2001), pp. 67–76. 1

[CCSM09] CHAZAL F., COHEN-STEINER D., MÉRIGOT Q.: *Geometric Inference for Measures based on Distance Functions*. Tech. Rep. 00383685, INRIA, 2009. 2, 3

[CGA10] Computational Geometry Algorithms Library CGAL-3.6. http://www.cgal.org/, 2010. 3, 6

[HDD*92] HOPPE H., DEROSE T., DUCHAMP T., MCDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *ACM SIGGRAPH* (1992), pp. 71–78. 7

[HK06] HORNUNG A., KOBBELT L.: Robust reconstruction of watertight 3D models from non-uniformly sampled point clouds without normal information. In *EUROGRAPHICS Symposium on Geometry Processing* (2006), pp. 41–50. 2

[HLZ*09] HUANG H., LI D., ZHANG H., ASCHER U., COHEN-OR D.: Consolidation of unorganized point clouds for surface reconstruction. *ACM Transactions on Graphics 28*, 5 (2009). 2, 7

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson Surface Reconstruction. In *EUROGRAPHICS Symposium on Geometry Processing* (2006), pp. 61–70. 1, 7

[KSO04] KOLLURI R., SHEWCHUK J. R., O'BRIEN J. F.: Spectral surface reconstruction from noisy point clouds. In *EUROGRAPHICS Symposium on Geometry Processing* (2004), pp. 11–21. 2

[LPK09] LABATUT P., PONS J.-P., KERIVEN R.: Robust and efficient surface reconstruction from range data. *Computer Graphics Forum 28*, 8 (2009), 2275–2290. 1

[MNG04] MITRA N., NGUYEN A., GUIBAS L.: Estimating surface normals in noisy point cloud data. *International Journal of Computational Geometry and Applications* (2004). 2

[MPS08] MANSON J., PETROVA G., SCHAEFER S.: Streaming surface reconstruction using wavelets. *Computer Graphics Forum 27*, 5 (2008), 1411–1420. 7

[NOS09] NAGAI Y., OHTAKE Y., SUZUKI H.: Smoothing of partition of unity implicit surfaces for noise robust surface reconstruction. *Computer Graphics Forum 28*, 5 (2009), 1339–1348. 1

[OBA05] OHTAKE Y., BELYAEV A., ALEXA M.: Sparse low-degree implicits with applications to high quality rendering, feature extraction, and smoothing. In *EUROGRAPHICS Symposium on Geometry Processing* (2005), pp. 149–158. 1

[WCS05] WALDER C., CHAPELLE O., SCHÖLKOPF B.: Implicit surface modelling as an eigenvalue problem. In *Machine Learning ICML 2005* (2005), pp. 936–939. 2