Surface Reconstruction from Unorganized Points

1. Briefly summarize the paper's contributions. Does it address a new problem? Does it present a new approach? Does it show new types of results?

- **[AS]**
  This paper presents an algorithm that produces a simplicial surface, approximating a manifold surface M, from an unorganized set of points containing little to no information about M. It presents a new approach to the problem of recovering surfaces from range scans and other practical situations.
- **[DS]**
  The method takes as input an unorganized set of points and outputs a simplicial surface that approximates a manifold. Makes no assumptions on the topology, presence of boundaries or the geometry of the manifold.
- **[FP]**
  The authors propose a method to reconstruct a surface as the zero set of a signed distance function. They locally approximate the surface by fitting tangent planes at point neighborhoods and defining a coherent normal orientation. Their approach is valid for any raw set of points (no normal required), independent of the acquisition source. They present results in point sets obtained from meshes, ray tracing, range images and contours (horizontal slices).
- **[JD]**
  The paper introduces a method for general surface reconstruction from unorganized, unoriented points. The algorithm assumes nothing about the data and reconstructs surfaces of arbitrary topology with possible boundaries. It is the first use of moving least squares for surface reconstruction.
- **[LF]**
  The paper presents a method that can reconstruct surfaces of arbitrary topology unlike parametric methods, only requires point cloud data (no surface normals needed), and handles boundaries more naturally than implicit methods as specified by the authors.
- **[MK]**
  This paper proposes one of the first methods for implicit surface reconstruction from point-sets. The paper makes several contributions: (1) it describes an approach for consistently estimating normals for unoriented points and (2) it shows how to use the oriented points to construct a piecewise linear (though discontinuous along boundaries) approximation to the distance function.

2. What is the key insight of the paper? (in 1-2 sentences)

- **[AS]**
  The key insight of this paper is that a reasonable surface reconstruction can be obtained from data that contains no information about the topology or geometry of the manifold from which the data was sampled.
- **[DS]**
  The algorithm defines signed distance functions for regions near the surface by associating every data point with an oriented plane. For orientation consistency, the plane directions are propagated across all points by computing a minimum spanning tree on the Riemannian Graph.
- **[FP]**
  A surface is given by the zero set of its normal distance. Distance is computed by projecting points on near tangent planes. Reconstruction from the signed function is done by marching cubes.

- **[JD]**
  The key insight is to construct a best-fit plane from the k-nearest neighbors of a sample point to create a discontinuous piece-wise linear signed distance function for estimating the surface.
- **[LF]**
  The key insight of the paper perhaps is to generalize the surface reconstruction problem and perhaps identify portions of that are present, no matter how the point cloud data for that surface was collected. By contrast, in the other paper this week, Curless and Levoy have success in working to optimize reconstruction specifically from range images.
- **[MK]**
  The key observation of this paper is that, given a point on the surface of the model, the Euclidean Distance function restricted to the Voronoi region of that point is a linear function with a value of zero at the point and gradient equal to the point's normal. Thus, given a set of points with estimated normals, an implicit function can be constructed by computing the Voronoi region of each point and using the position and the normal to define a linear function.

3. What are the limitations of the method? Assumptions on the input? Lack of robustness? Demonstrated on practical data?

- **[AS]**
  One limitation of this method is that the parameter *k*used to compute the *k*-nearest neighbors is a user specified value. Since most practical data is noisy, *k*is a critical parameter in ensuring a good reconstruction. Hence the quality of the output depends on the user. It would be desirable to select and adapt *k*automatically.
- **[DS]**
  To handle boundaries, the algorithm assumes the point cloud is rho-dense and delta-noisy.
  No formal guarantees on the correctness of the reconstruction had been developed yet.
- **[FP]**
  Due to the locality of the method it seems not robust to noise, undersampling and outliers. In fact, a dense and uniform sampling (the rho-dense and delta-noisy conditions) is assumed by the authors.
  The reconstruction method tightly depends of a correct normal orientation. Finding a global optimal orientation is a NP hard problem so the authors prefer a more local (propagation based) approach. This reduction produced satisfactory results in their experiments but may not be robust enough on less uniform point inputs.
  The discontinuity of the signed distance function may introduce geometric artifacts on the marching cube's surface reconstruction.
- **[JD]**
  The method is a general algorithm that does not take into account regular sampling error. Therefore, the algorithm has difficulty with range-scanned data points. The k-neighborhood is fixed throughout the graph and does not adjust to areas of over and under-sampling.
- **[LF]**
  Assumes uniform sample density. The method is tested on sample sets from various sources, including meshes, ray traced points, range images and contours.
- **[MK]**
  I believe that the method implicitly assumes that the data is noise free, as this could corrupt the estimation of the normals and, independently, could result in far-reaching effects on the implicit function away from the point samples.

4. Are there any guarantees on the output? (Is it manifold? does it have boundaries?)

- **[AS]**
  The algorithm guarantees a manifold simplicial surface with or without boundary, approximating the sampled manifold M, from a set of unorganized points scattered on or near the surface of M.
- **[DS]**
  Produces a simplicial surface (piecewise linear surface with triangular faces) that approximates the manifold. Can handle boundaries.
- **[FP]**
  The marching cube method guarantees that the reconstruction is a manifold. For a densely and uniformly sampled surface, the generated triangle mesh is homeomorphic to the input surface (since the zero set will converge to the surface).
  In undersampled or empty regions the signed distance is not defined. This clamps the surface extension through neighbor cubes in the marching process and produces boundary edges.
- **[JD]**
  The output is guaranteed to be manifold, possibly with boundaries but not water-tight.
- **[LF]**
  The output should be topologically correct, but no formal guarentees of correctness were provided. The output should be manifold according to the implicit function theorem as stated by the authors.
- **[MK]**
  The reconstructed surface is obtained through a variant of marching cubes which should generate a manifold surface. Due to the clipping (when \delta is finite) the surface can have boundaries.

5. What is the space/time complexity of the approach?

- **[AS]**
  A naive implementation of this algorithm would take O(n2) time, since the EMST graph would take O(n2) time, where n is the number of points. However, if a hierarchical spatial partitioning scheme, such as octrees or k-D trees, is used, then the time complexity of the algorithm can be reduced to O(nk), which is the time taken to construct the Riemannian Graph, where k is the numberof neighbors considered in the k-nearest neighbors computation. Spatial complexity, using octrees or k-D trees, would be O(n).
- **[DS]**
  The Riemannian graph can be constructed in O(n*k) time (for k-nearest neighbors). Finding the best path is O(n * logn).
- **[FP]**
  The input samples are grouped in voxels of a regular spatial grid. Due to the uniform sampling assumption it is expected that each non empty voxel contain a constant number of samples. This data structure allow the computation of the Riemmanian graph in (nk) (n= number of samples, k= number of nearest neighbours) and the computation of the MST for normal propagation in $(n\log n)$. Signed distance function evaluation, which requires finding the closest center is done in constant time using this data structure.
  The cost of surface reconstruction is (h^2) on the resolution of the marching cubes (since doubling the resolution approximately quadruplicates the number of triangles in the output mesh).
- **[JD]**
  The authors break the complexity into parts: the EMST graph is O(n^2), finding the k-

nearest neighbors is O(n + klogn), and finding the nearest tangent plane is O(n). The space complexity is linear because the EMST has an edge for every point.

- **[LF]**
  The paper provides the time complexity of its main sub problems: (1) EMST graph => O(n^2), (2) k-nearest neighbors to a certain point => O(n + k*log(n)), and (3) nearest tangent plane origin to a given point => O(n). Spatial paritioning may be used to reduce these complexities by a factor of n as specified by the authors.
- **[MK]**
  The complexity of the algorithm is O(n log n), defined by the propagation of normal orientation, which requires maintaining an ordered queue.

6. How could the approach be generalized?

- **[AS]**
  The paper claims that their algorithm can be generalized to arbitrary dimensions -it is capable to reconstructing *d-1*dimensional manifolds in *d* dimensional space.
- **[DS]**
  This algorithm can be used to reconstruct d dimensional manifolds in d+1 dimensional space.
  The output of the algorithm can be used as the starting point for a subsequent spline surface fitting procedure.
- **[FP]**
  The algorithm is applicable for manifold reconstruction in arbitrary finite dimensional space. In particular, it is also applicable to planar curve reconstruction. The algorithm (as a whole) is formed by several modular methods that can be interchanged or extended using different approaches. As the authors claims, setting a global coherent normal orientation is a critic step for the reconstruction and some other (more global) techniques should be attempted.
  The defined signed distance function is discontinuous and this may produce artifacts in the reconstruction. Using a continuous distance function (for instance, by weighting projected distance according to center distance) could improve the results.
  Finally, other strategies (different to marching cubes) may be adapted to reconstruct the surface from the signed distance function. For instance, the authors propose a spline surface fitting or using Bezier triangles.
- **[JD]**
  Moving least squares can be applied to any domain where there exists a distance metric.
- **[LF]**
  The approach could be extended to reconstruct d-dimensional manifolds in d+1 dimensional space.
- **[MK]**
  In a sense, the approach can be thought of as a precursor to MLS (or perhaps more appropriately MPU) with the implicit function estimated independently in different regions of space. It seems like it would be straightforward to incorporate partition-of-unity “gluing” so that the independent pieces can be composited into a single, global, smooth function.
  Also, in the current description, the method assumes that the \epsilon and \delta values are constant over the point set. Can they be adapted (e.g. if points come with confidence values).

7. If you could ask the authors a question (e.g. “can you clarify” or “have you considered”) about the work, what would it be?

- **[AS]**
  Can you analytically show how this algorithm can be extended to arbitrary dimensions?
- **[DS]**
  How are you planning to prove correctness for your reconstruction algorithm?
  Why was the marching cubes algorithm the best choice for your method?
- **[FP]**
  The discontinuity of the signed distance function seems to be an undesired property of the method. How the discontinuity affects the reconstruction? What kind of artifacts does it produces? What is gained and lost by “smoothing” the distance function to a continuous function?
- **[JD]**
  Is this the first use of moving least squares in the literature? Are there any cases where the discontinuity between tangent planes causes irregularities? What about solving a system of tangent plane directions where the directions are "relaxed" so that they overlap to minimize/remove discontinuity?
- **[LF]**
- **[MK]**
  Can the method be extended to forego the \epsilon-dense assumption? For more recent approaches (e.g. power-crust, tight co-cone) the required sampling density adapts to the feature size. It would be interesting to see if this method could support a similarly laxer constraint on sampling density.
  If sample orientation is known in advance, why compute a Riemannian Graph at all? Since this is the computational bottle-neck of the system, it would be preferable to bypass the step altogether.
  The authors discuss allowing k to adapt locally. Why is this necessary? Doesn’t the use of k automatically adapt to sampling density? (As opposed, for example, to use a fixed \epsilon neighborhood.) What might be interesting, however, is to extend the method to adapt k not to the sampling density, but to the angular distribution around the point sample. (e.g. If the scanning is done in stripes, to ensure that k is large enough to include points in adjacent stripes – akin to computing a locally Delaunay triangulation and using a k-ring neighborhood.)