Algebraic Point Set Surfaces

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's main contribution is a new Point Set Surface (PSS) definition based on Moving Least Squares (MLS) fitting of algebraic spheres which can be expressed either by a projection procedure or as an implicit form.
- **[DS]**
  The paper defines a point set surface based on an MLS fitting of algebraic spheres, whose representation can be expressed as a projection operator or an implicit surface. The method has a more stable projection operator under low sampling and high curvature. It naturally handles planar point clouds. It provides an estimate of the surface's mean curvature at no additional cost and allows robust handling of sharp features and boundaries. Also presents a novel normal estimation technique.
- **[FP]**
  The paper proposes a MLS algorithm which outperforms previous approaches on dealing with undersampled regions. Instead of locally fitting a base plane and a polynomial, the authors use an algebraic spherical fitting to do projection, implicit function evaluation, and normal estimation. The authors also present methods to compute and propagate normal on a raw point set, as well as a tagging method to preserve features. The computation of the algebraic fitting is done by solving a least square problem, which is much simpler than solving the optimal geometrical fitting, and provides a good approximation of the signed distance for points close to the sphere.
- **[JD]**
- **[LF]**
- **[MK]**
  The paper proposes several extensions to Alexa et al.'s PSS paper:
  1] The authors propose fitting an algebraic sphere (which can also characterize planes) to each point. Iteratively computing the best fit sphere in the vicinity of a point and then mapping the point on to the nearest point on the sphere provides an approximate projection approach, similar to the approach in the original work.
  2] In the case that normals are unknown, the solution to the best sphere problem is obtained by solving a generalized eigenvalue problem, so that the parameters of the best fit sphere are only defined up to sign (and hence the implicit function cannot distinguish between interior and exterior). This does not affect the algorithm since the projection only requires knowing the zero-set (which is invariant to negation of parameter values). However, if the normals are given, then the solution is obtained by solving a least-squares problem, and the locally computed implicit function can be used to define a (consistent) global explicit. [Though the zero set of this global set is not equal to the fixed points of the projection operation, since that operation is only approximate and requires iteration.]
  3] The authors propose a variation on Hoppe's normal estimation that uses the algebraic spheres to "transport" the normals to the neighbors, so that neighbors are not assumed to have parallel normals (making the method more robust near corners).
  4] The authors propose an approach for marking and consistent propagation that can be used to reconstruct surfaces with sharp creases and corners.

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

- **[AS]**
  The key insight of this paper is that fitting an algebraic sphere, rather than a plane,

provides significantly improved stability where planar MLS fails, for instance, in the case of undersampling, and in the presence of high curvature.

- **[DS]**
  The key insight is the procedures that can efficiently perform MLS fitting to algebraic spheres (instead of a plane) to a given point set, which provides increased stability and correct handling of sheet separation as well as undersampled regions and high curvature surfaces. The presented algorithm run on point sets with normals, reduces to linear least squares, so a normal fitting must be performed for point sets without normals.
- **[FP]**
  Spherical fitting provide a fair approximation for a wide range of curvature and sampling density conditions. Algebraic fitting is efficiently computed, and used in projection, normal estimation, etc.
- **[JD]**
- **[LF]**
- **[MK]**
  The key insight of the paper is that by fitting algebraic spheres rather than planes, the point-set surface can be made to be more robust in the presence of noise and approximate sheets.

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 it can be sensitive to unwanted stationary points away from the surface.
- **[DS]**
  The input is a point set, but normals are estimated beforehand to improve robustness, quality and performance. Normals are also necessary for proper rendering (in/out).
- **[FP]**
  In order to implicitly define the surface as a zero level set it is required a global coherent normal orientation. Therefore accurate normal must be initially available or being computed and propagated according to their proposal.
  The spherical fitting is done in a traditional least square way which may be not robust in the presence of noise or outliers.
- **[JD]**
- **[LF]**
- **[MK]**
  The method to work well on clean data, though not necessarily densely-sampled, data. It is less clear how well the method would work on real-world data with noise. In particular, how robust/stable/etc. is the zero level-set of the global implicit function (assuming that normals are either given or estimated using the proposed approach)? And, does it suffer from the same narrow band problem noted by Amenta and Kil, where noise in the input can result in a situation where the input samples themselves don’t project onto the correct surface?

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

- **[AS]**
  The output can handle both sharp features and boundaries.
- **[DS]**
  The output is a continuous approximation or estimation of the input point set.

- **[FP]**
  The authors do not claim any guarantee on the output properties, but it is sound that the image of the projection operator (i.e. the set of points which maps to themselves) should be an infinitely smooth manifold as in the case of Levin's MLS.
  Due to the MLS nature of the approach the result should be a closed surface (i.e., no boundaries). The authors guarantee sharp features reconstruction by using a tagging algorithm.
- **[JD]**
- **[LF]**
- **[MK]**
  A priori, the method generates a smooth manifold but without boundaries. However, using sharp feature and boundary labels, the authors show that the method can be extended to generate non-smooth surfaces with boundaries.

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

- **[AS]**
  The generalized eigenvalue problem is the most time consuming procedure in this method. Since, the matrices in the system are sparse, the time complexity of this method is roughly O(N*k*logN), where N is the number of points and *k*is the size of the BSP neighborhood. The space complexity of this approach is roughly O(N*m*), where *m*is the number of Eigen pairs to be calculated.
- **[DS]**
  The algorithm reduces to linear least squares for a point set with normals.
  Up to 45 million points per second, which can be used for real-time calculations.
- **[FP]**
- The algebraic fitting phase (computation of vector u) involves the solution of a minimal eigenvalue problem. This can be efficiently done, for instance, using SVD. Once the spherical fit is computed, point projection, implicit function evaluation and normal estimation are all (1). The final projection of a point into the surface is done by iterative sphere projection, requiring an algebraic fitting update.
  The normal propagation phase requires the computation of a BSP neighborhood (worst case (n^2)) , graph weights computation ($O(n)$) and finding the MST on the respective graph ($O(n)$ in a sparse graph).
- **[JD]**
- **[LF]**
- **[MK]**
  As with previous APSS methods, I believe that the method only requires being able to estimate neighbors (either k-nearest, or within an epsilon ball) which should be doable in linear time with a good datastructure (assuming points are not clustered too tightly into disjoint groups).

6. How could the approach be generalized?

- **[AS]**
  This work can be generalized to non-manifold geometry, as well as real-time ray-tracing of dynamic APSS.
- **[DS]**
  The method can be applied to higher order surfaces (not just algebraic spheres), for example, ellipsoids are more suitable for highly anisotropic objects.
  Extensions to non-manifold geometry as well as real-time ray tracing of dynamic algebraic point set surfaces.

- **[FP]**
  As the authors claim, a natural approach is to consider algebraic ellipsoidal fitting instead of spherical. However, they also claim that this possibility was disregarded since there is no natural degeneration of ellipsoids into planes. According to the authors spherical fitting still gives good results in non-umbilical regions, but probably, more accurate results can be obtained by locally fitting using a non-umbilical base shapes. Therefore I think trying to find extensions in this direction is worth.
- **[JD]**
- **[LF]**
- **[MK]**
  Natural generalizations include the use of either higher order algebraic surfaces, or even extending the proposed case to support ellipses.

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]**
  Have you considered how the performance of your method varies as more noise is added to the point set?
- **[DS]**
  How would we render the resulting surface? What would the algorithm look like for a ray tracing application?
- **[FP]**
  The authors propose algebraic fitting due to performance reasons. What is the quality trade-off? How close is the reconstruction using algebraic fitting to the one obtained using geometrical fitting?
  When the authors say there is no natural degeneration of a ellipsoid to a plane, Do they refer to geometrical degeneration or algebraic degeneration?
- **[JD]**
- **[LF]**
- **[MK]**
  When fitting an algebraic sphere to points with normals, how sensitive is the fit to normal flips?