3D Scattered Data Approximation with Adaptive Compactly Supported Radial Basis Functions

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 approach for surface reconstruction from a point set, which uses compactly supported RBFs whose centers are a subset of the initial point set. Each RBF has an associated support size, which is chosen adaptively according to the density and surface geometry at that RBF center. This method produces high quality noise-robust approximation of the surface.
- **[DS]**
  The paper constructs an implicit function whose zero level set approximates a surface constructed from a set of points scattered over a piecewise smooth surface with oriented normals.
  The function is constructed by using compactly supported RBFs centered at randomly chosen points from the set of input sample points. The support size of each RBF is adapted based on the surface geometry around that point as well as the desired smoothness of the resulting reconstruction.
  The algorithm produces an approximation that is robust to noise.
  The approach takes a partition of unity (PU) approximation of the surfaces and adds a normalized RBF function to improve the surface detail. This is done through a global error minimization.
- **[FP]**
  The authors propose surface reconstruction from an implicit function that involves a local polynomial fitting term (PU) and a local correction term (RBF). The authors define the global implicit function, by adding local approximations centered at adaptively chosen points from the input set. May be, the main contribution of the paper is the methodology proposed for the selection of the basis function support (scale) at each center. They use an energy function that finds an optimal support by trading off accuracy and large area coverage. The local correction term (RBF) improves the sharpness of the reconstruction.
- **[JD]**
  The paper presents another method for surface reconstruction using radial basis functions. The algorithm is robust to noise and gives better local detail than previous algorithms. It also accounts for confidence values from the data acquisition method.
- **[LF]**
- **[MK]**
  The authors propose a hybrid approach to surface reconstruction that combines partition of unity reconstruction with RBF fitting. The PU provides a solution that only requires performing a local blend of solutions. The RBF acts as a high-frequency refinement of the PU solution by adjusting the implicit function obtained from the PU approach to get the implicit surface to pass closer to the samples.
  In the course of implementation, the authors employ robust methods to ensure that the obtained solution does not over-fit noise. This includes both the careful estimation of the size of a point's neighborhood and the RBF fitting. (The former defines an energy that trades off fit quality for number of neighbors. The latter, uses the fact that the RBF solution provides an offsetting correction term, whose local frequency is determined by neighborhood size – so that penalizing the size of the offset, as a function of neighborhood size, reduces the high-frequency over-fitting.)

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

- **[AS]**
  The key insight of this paper is that although a compactly supported RBF modulated by a function, equivalent to the partition of unity (PU) approximation, produces high quality reconstructions, some of the high frequency detail is lost, giving the appearance of over smoothing in areas containing sharp features. In order to reconstruct fine features, a weighted normalized RBF is added to the PU approximation term, where the weights are computed by a minimization problem similar to ridge regression since it is able to fit the original data better in the presence of noise, bringing out local detail.
- **[DS]**
  The key insight is the use of the normalized RBFs whose support size is adapted based on the complexity of the geometry around their center points. These are used to improve the detail on the partition of unity reconstruction, through a global error computation.
- **[FP]**
  Description of the “base geometry” of the model is done with a minimal set of local components. Further detail is attained with a second correction term that fit to the sample points.
- **[JD]**
  The key insight of the paper is to use compactly supported RBFs. This yields a solution with a sparse matrix and allows the RBFs to adapt to local detail.
- **[LF]**
- **[MK]**
  In a sense, it seems that the approach the authors take is premised on the observation that while the reconstruction problem is inherently global, it may be able to get large parts of the solution using strictly local information (PU), and then incorporating a simple but global approach (RBF) to tie things together. (This provides a hierarchical approach that is not unlike multigrid – though this approach only has two levels.)

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

- **[AS]**
  This method does not guarantee watertight reconstructions -it can fill holes if the support size of the RBFs near the hole is big enough to cover the hole, but this is not always the case. The value for T_{sa} is lower for models with holes, which assumes some knowledge of the distribution of points in the point set, which is also not intuitive
- **[DS]**
  The method won't reconstruct parts of the surface that contain big holes. In order to fill in holes, that part of the surface must be within the support size of one or more of the chosen RBF centers.
- **[FP]**
  The first limitation is the requirement of normals. Normals are needed to define a global coherent sign of the local functions.
  The paper introduce several parameters (TSA, Treg, σmin,Toverlap) that were manually tunned by the authors. They present some set of values for these parameters that worked fined on the tested models, but may be, additional tunning could be required particular reconstructions.
  Finally, computation of the optimal support σ of a basis function seems to be a very expensive task. Evaluating the energy function ESA for a specific σ requires of fitting a complete model. Finding an optimal σ by doing this iteratively is a high cost.
- **[JD]**
  The algorithm can only fill holes within the range of its set of RBFs. The input is

assumed to be oriented. It is sensitive to point sets with sharply varying density resulting in poor performance compared to another methods using RBFs.

- **[LF]**
- **[MK]**
  The problem with this approach is the partition of unity is defined in terms of locally supported functions. As a result, the PU is only defined in the union of the supports of these functions, so the implicit function is undefined away from the samples. (In practice, this means that the method could have trouble filling in holes if they are larger than the support radii of the RBFs.)
  On a smaller note, the approach is careful to make energy terms dimensionless so that they can be combined (e.g. normalizing in terms of bounding box diagonal). However, this only ensures that the cumulative energy is scale invariant, it does not imply that a weighted sum of the two approaches is the right way to go. (For example, given two dimensionless energies E_1 andE_2, one can define a cumulative energy $E = f(E_1) + g(E_2)$ where $f, g: \mathbb{R} \to \mathbb{R}$ are any monotonically increasing functions.)

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

- **[AS]**
  This method guarantees a smooth manifold surface, but cannot guarantee watertight surfaces -the surface might contain holes and boundaries.
- **[DS]**
  The output is an implicit function from which a manifold surface can be obtained. Not all holes will be filled, so a watertight surface is not guaranteed.
- **[FP]**
  The implicit function is not globally defined (its support is given by the union of the local functions supports), so the zero level set may not form a closed surface (the reconstruction does not fill holes in non-sampled regions). Additionally, if the support of the local functions is taken too small and data is noisy, extra zero level sets may form (authors set a threshold σmin to solve this). The mesh generation is done with Bloomenthal’s polygonazer which may guarantee manifold properties.
- **[JD]**
  The output is manifold but not water-tight.
- **[LF]**
- **[MK]**
  Since the solution is only defined within the support of the PU functions, the reconstructed surface can have boundaries. (Assuming non-vanishing gradients, it will be manifold.

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

- **[AS]**
  The time complexity and the number of RBF centers, and hence the space complexity, of this method depends on the size and the resolution of the reconstruction.
- **[DS]**
  Depends on the size of the input data set as well as the geometry of the surface. More complex surfaces will require a larger number of RBF centers to achieve the desired reconstruction detail.
  Finding function g is done through a minimization which reduces to solving a 6x6 linear problem by using SVD decomposition.
  Finding sigma is a one-dimensional minimization problem.

These two are done for every RBF center. The number of centers depends on the complexity of the surface.

- **[FP]**
  As the authors claim, the complexity on the reconstruction depends on the input complexity. Intricate geometry requires more centers to attain the desired accuracy on the local patches. The authors use Brent's method to find the optimal σ for each local function. They claim that this method converge to an accurate result (L/105) in few iterations (less than 10).
  The local correction weights (RBF) λ's are the optima of a regularized quadratic problem. This problem reduces to solving a sparse linear symmetric problem which is solved by preconditioned biconjugate gradient. Still, solution to this linear problem represents 40% of the time in the construction implicit function.
- **[JD]**
  The complexity is dependent on the number of RBF centers. The storage is linear.
- **[LF]**
- **[MK]**
  Ooof. The PU part should be linear (using good nearest-neighbor data structures). The RBF part requires a solve of a sparse symmetric system, for which the authors use a (bi-) conjugate-gradient solver. In principal, one would expect $O(N^{0.5})$ iterations for convergence, but it may be the case that the number of iterations is closer to $O(1)$ since the zero solution is already a good guess (since the PU surface should already come close to passing through the samples).
  There is also the issue of estimating neighborhood sizes, which requires fitting quadrics to differently sized neighborhoods about each point, but I suspect that this is can be done in near constant time per point (assuming that the maximum number of neighbors one ever considers is bounded by a constant independent of the size of the point set).

6. How could the approach be generalized?

- **[AS]**
  This approach can be generalized to include RBF centers that do not necessarily belong to the initial point set by using a point clustering approach. They hope that this will allow them to treat datasets with different levels of noise uniformly. This approach can also be extended to use a multi-scale approach, in order to be more robust against point density variations.
- **[DS]**
  Can be extended so that the choice of RBF centers is not restricted to the sample points. Can be combined with a multi-scale approach.
- **[FP]**
  Most hierarchical approaches to surface reconstruction from implicit function (Ohtake 2003, Kazhdan 2005, Calakli 2011) use regular structures such a octrees. The approach presented in this paper, could be extended to a non-regular hierarchical reconstruction using local functions. Local functions at each level refine the approximation provided by coarser levels. Following the idea that a small number of centers is desirable at each level.
  The author propose as a topic for further research to provide a method that do not restrict center position to sample points, but find the optimally position in 3D space.
- **[JD]**
  The algorithm estimates a function from a set of points for any input that has a distance metric and a first derivative.
- **[LF]**

- **[MK]**
  An extension of this approach could be to replace the parameters λ_i, which define a locally constant offset function with the parameters of a higher order function (e.g. using quadrics, just like for the PU part of the fitting.)

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 an extension to your method that might guarantee watertight reconstructions?
- **[DS]**
  How can this approach be extended so that it will fill in large holes?
- **[FP]**
  The second term of the implicit function (the local correction or RBF) is just a constant multiple (λi) of the weight function(φi). Therefore this term can be regrouped with the constant term of gi (the polynomial fitting) which also multiplies the weight function (φi). Why computing λi is not redundant (i.e., why it is not already considered in gi?), Could this be solved by using, on a local region to fit, the functions from overlapping regions already fitted?.
- **[JD]**
  Have you considered making sigma non-spherical?
- **[LF]**
- **[MK]**
  Why use a bi-conjugate-gradient solver? Since the matrix is SPD, a conjugate-gradient solver should do. (And it should converge more quickly.)
  In Equation (12), why set $E_{local}(\sigma) = L$ for $\sigma < \sigma_{min}$? This gives a discontinuity. Why not use $E_{local}(\sigma) = E_{local}(\sigma_{min})$ instead?