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

Yutaka Ohtake [1,2] Alexander Belyaev [1] Hans-Peter Seidel [1]

[1] Computer Graphics Group, Max-Planck-Institut für Informatik,Germany

[2] Integrated V-CAD System Research Program, RIKEN, Japan

E-mails: {ohtake,belyaev,hpseidel}@mpi-sb.mpg.de

Abstract

In this paper, we develop an adaptive RBF fitting procedure for a high quality approximation of a set of points scattered over a piecewise smooth surface. We use compactly supported RBFs whose centers are randomly chosen from the points. The randomness is controlled by the point density and surface geometry. For each RBF, its support size is chosen adaptively according to surface geometry at a vicinity of the RBF center. All these lead to a noise-robust high quality approximation of the set. We also adapt our basic technique for shape reconstruction from registered range scans by taking into account measurement confidences. Finally, an interesting link between our RBF fitting procedure and partition of unity approximations is established and discussed.

Keywords: Adaptive RBF, surface reconstruction from scattered data

1 Introduction and Algorithm Overview

The *Radial Basis Functions* (RBF) approach has proven to be very useful in shape reconstruction from scattered, noisy, incomplete data, see, for example, [5, 6, 13, 19, 23, 31, 33, 34] and references therein. Recent studies on RBF-based shape reconstruction [11, 17, 7, 27] are focused on efficient shape reconstruction from large and/or noisy scattered point datasets produced by modern range finder devices. In this paper, we extend our adaptive RBF fitting technique [26] for processing noisy scattered point data. For the sake of completeness and since [26] is not a refereed paper, we include in this paper a description of a randomized procedure developed in [26] for selecting RBF centers.

Consider a set of points $\mathcal{P} = \{\mathbf{p}_1, \cdots, \mathbf{p}_N\}$ sampled from a surface and equipped with unit normals $\mathcal{N} = \{\mathbf{n}_1, \cdots, \mathbf{n}_N\}$ that indicate the surface orientation. In practice, points $\mathcal{P}$ are generated from range scans and normals $\mathcal{N}$ are usually estimated from range data during the shape acquisition phase or by local least-square fitting. Some points of a range scan are usually measured more accurately than others. This advocates assigning a certain confidence measure to each point of the scan [32, 10]. Thus we assume that each point $\mathbf{p}_i \in \mathcal{P}$ is attributed with a real number $c_i \in [0, 1]$ indicating the confidence of $\mathbf{p}_i$.

Our aim is to construct a function $y = f(\mathbf{x})$ such that its zero level-set $f(\mathbf{x}) = 0$ approximates $\mathcal{P}$.

Implicit surface $f(\mathbf{x}) = 0$ separates the space into two parts: $f(\mathbf{x}) > 0$ and $f(\mathbf{x}) < 0$. Let us assume that the normals $\mathcal{N}$ are pointing into the part of space where $f(\mathbf{x}) > 0$. Thus we can say that $f(\mathbf{x})$ has negative values outside the surface and positive values inside the surface.

Given a set of approximation centers $\mathcal{C} = \{\mathbf{c}_1, \cdots, \mathbf{c}_M\}$, $M < N$, we construct $f(\mathbf{x})$ approximating $\mathcal{P}$ in the following form suggested in [27]

$$\sum_{\mathbf{c}_i \in \mathcal{C}} [g_i(\mathbf{x}) + \lambda_i] \, \phi_{\sigma_i}(\|\mathbf{x} - \mathbf{c}_i\|), \tag{1}$$

where $\phi_\sigma(r) = \phi(r/\sigma)$, $\phi(r)$ is an RBF function, and $g_i(\mathbf{x})$ and λ_i are unknown functions and coefficients to be determined. We assume that $\mathcal{C}$ contains much less points than $\mathcal{P}$. Thus (1) delivers an economical approximation of $\mathcal{P}$. For each approximation center $\mathbf{c}_i$, we construct $g_i(\mathbf{x})$ as a local quadratic approximation of $\mathcal{P}$ in $\{\|\mathbf{x} - \mathbf{c}_i\| < \sigma_i\}$, the region of influence of $\mathbf{c}_i$. Then we determine $\{\lambda_i\}$ from M interpolation conditions

$$f(\mathbf{c}_i) = 0, \quad i = 1, \ldots, M. \tag{2}$$

Notice that (2) is a system of linear equations and $\{\phi_{\sigma_i}(r)\}$ form a set of *generalized* radial basis functions [20, 28, 16]. We choose $\phi(r)$ as a Gaussian-like function with compact support. Namely we set $\phi(r) = (1 - r)_+^4 (4r + 1)$, Wendland's compactly supported RBF [36].

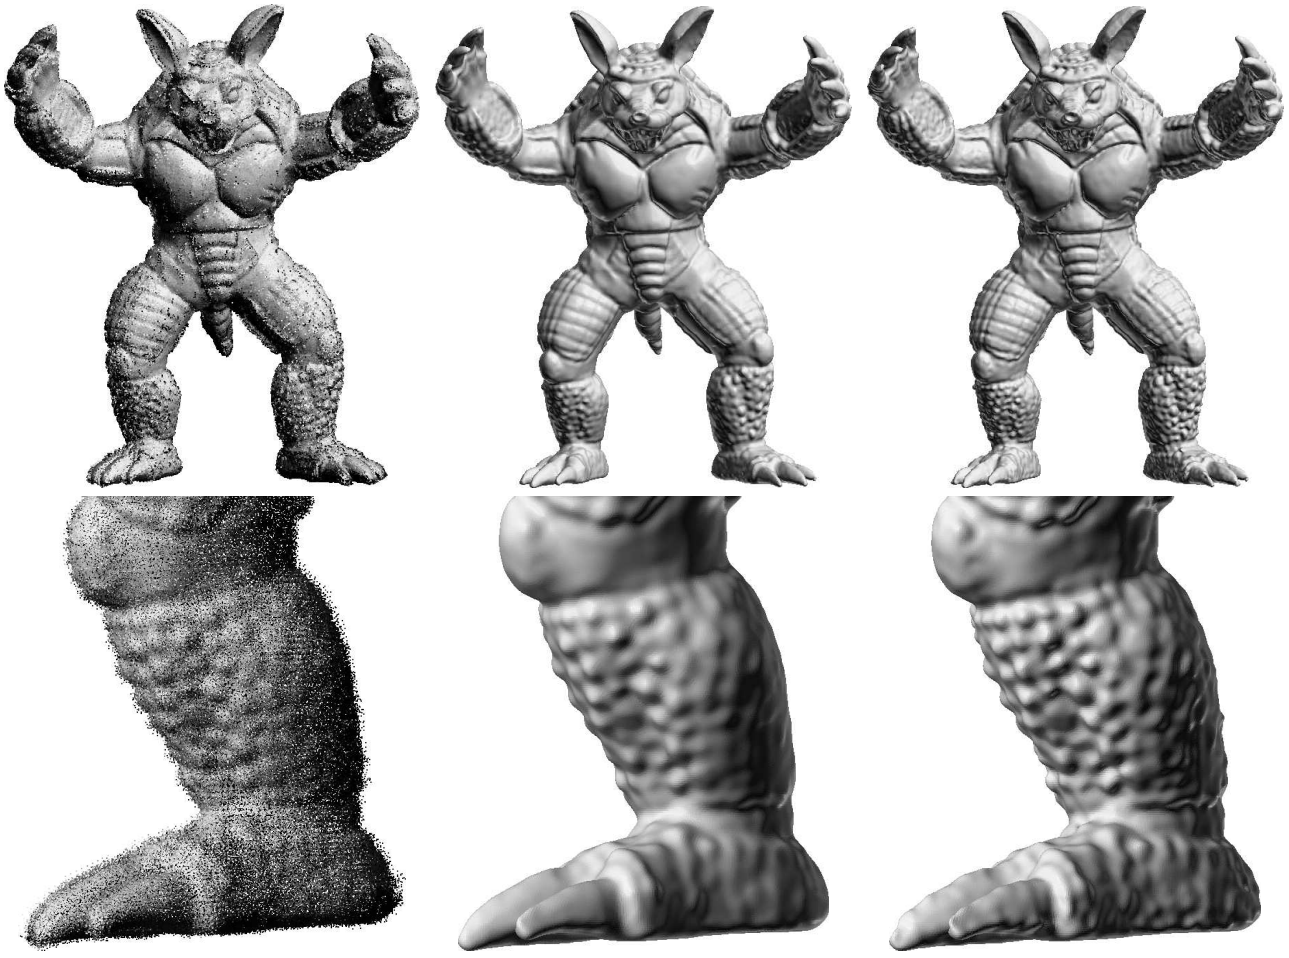

Initial noisy point dataset | PU approximation | PU + RBF approximation

Fig. 1. Reconstruction of the Stanford Armadillo dataset consisting of $2,366$K points 41K with approximation centers. L^2 approximation error equals 7.22×10^{-4} for PU approximation (middle) and 5.99×10^{-4} for PU + RBF approximation (right).

We can rewrite (1) in the form

$$\underbrace{\sum_{\mathbf{c}_i \in \mathcal{C}} g_i(\mathbf{x}) \phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|)}_{\text{base approximation}} + \underbrace{\sum_{\mathbf{c}_i \in \mathcal{C}} \lambda_i \phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|)}_{\text{local details}} \quad (3)$$

The first term of the right-hand side of (3) can be considered as a base approximation of $f(\mathbf{x})$ while the second term represents local details.

One can notice that the base approximation term in (3) has the same zero level-set as a partition of unity approximation (PU)

$$\sum_{\mathbf{c}_i \in \mathcal{C}} g_i(\mathbf{x}) \Phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|), \quad (4)$$

where functions

$$\Phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|) = \frac{\phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|)}{\sum_j \phi_{\sigma_j}(\|\mathbf{x}-\mathbf{c}_j\|)}$$

form a partition of unity (PU). PU approximations are now very popular in computational mechanics [2, 3] and can deliver fast-constructed high quality approximations of scattered data [37, 1, 25]. Functions $\{\Phi_{\sigma_i}(\cdot)\}$ belong to the class of the so-called normalized RBFs [22, 14] which often show a better performance than RBFs do in fitting non-uniform data [15, § 6.7].

So we will approximate points $\mathcal{P}$ attributed with normals $\mathcal{N}$ and confidences $\{c_i\}$ using normalized RBFs $\{\Phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|)\}$:

$$\underbrace{\sum_{\mathbf{c}_i \in \mathcal{C}} g_i(\mathbf{x}) \Phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|)}_{\text{adaptive PU}} + \underbrace{\sum_{\mathbf{c}_i \in \mathcal{C}} \lambda_i \Phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|)}_{\text{normalized RBF}} = 0 \quad (5)$$

Fig. 1 demonstrates the adaptive PU and PU+RBF approximations of the Stanford Armadillo dataset. Notice that the PU approximation alone delivers a high quality recon-

struction. Another interesting feature of our approach is building quite economical approximations: the dataset consisting of 2,366K points is accurately reconstructed using only 41K approximation centers.

In the following sections, we explain how to choose approximation centers $\{\mathbf{c}_i\}$ and determine their influences σ_i, construct local approximations $g_i(\mathbf{x})$ and use them to a build the PU approximation given by the first sum in (5), and finally refine the PU approximation by the RBF terms.

2 Constructing Adaptive PU Approximation

Typically $\mathcal{P}$ is generated from several overlapped range scans and the density of points is higher at the overlapped regions. In order to take into account density irregularities and the confidence measures attributed to the points, each point $\mathbf{p}_i \in \mathcal{P}$ is assigned a weight d_i given by

$$d_i = c_i \sum_{j=1}^{K} \|\mathbf{p}_i - \mathbf{p}_j\|^2, \tag{6}$$

where $\{\mathbf{p}_j\}_{j=1}^{K}$ are the K-nearest neighbors of $\mathbf{p}_i$. From our numerical experiments, we found that $K = 20$ is a good choice.

Constructing local approximations. For each approximation center $\mathbf{c}_i$ we construct a local quadratic approximation $g_i(\mathbf{x})$ of $\mathcal{P}$ in the σ-neighborhood of $\mathbf{c}_i$, $\{\mathbf{x} : \|\mathbf{x} - \mathbf{c}_i\| < \sigma\}$. Let us define local orthogonal coordinate system (u, v, w) at $\mathbf{c}_i$ such that the positive direction of w coincides with the direction defined by weighted averaging of normals

$$\sum_j d_j \, \phi_\sigma(\|\mathbf{p}_j - \mathbf{c}_i\|) \, \mathbf{n}_j. \tag{7}$$

Here $\mathbf{n}_j$ is the unit normal assigned to $\mathbf{p}_j$ and sum is taken over all points $\mathbf{p}_j \in \mathcal{P}$ from the σ-neighborhood of $\mathbf{c}_i$. The results of our numerical experiments show that weighted averaging (7) is sufficiently resistant to noise. The coefficients of the local fitting function

$$w = h(u, v) \equiv Au^2 + 2Buv + Cv^2 + Du + Ev + F$$

are determined by the following minimization procedure

$$\sum_j d_j \, \phi_\sigma(\|\mathbf{p}_j - \mathbf{c}_i\|) \, g_i(\mathbf{p}_j)^2 \to \min. \tag{8}$$

Given σ, the solution to this least-squares problem is obtained by using the normal equation approach [29]. The corresponding 6×6 linear system is solved by the matrix inverting via the singular value decomposition approach.

Now $g_i(\mathbf{x})$ is defined by

$$g_i(\mathbf{x}) = w - h(u, v),$$

where (u, v, w) are local coordinates of $\mathbf{x}$. Note that local approximation $g_i(\mathbf{x})$ depends on σ.

Local errors and influence parameters. To determine an optimal value for influence parameter σ_i associated with approximation center $\mathbf{c}_i$ we define an error function $E_{\text{local}}(\sigma)$ at $\mathbf{c}_i$ by

$$E_{\text{local}}(\sigma) = \frac{1}{L} \sqrt{\frac{\sum_j d_j \phi_\sigma(\|\mathbf{p}_j - \mathbf{c}_i\|) \left(\dfrac{g_i(\mathbf{p}_j)}{\|\nabla g_i(\mathbf{p}_j)\|} \right)^2}{\sum_j d_j \phi_\sigma(\|\mathbf{p}_j - \mathbf{c}_i\|)}}, \tag{9}$$

where L is the main diagonal of the bounding box of $\mathcal{P}$. The factor $1/L$ is used to make (9) scale-independent.

We can assume that $E_{\text{local}}(\sigma)$ is monotonically decreasing to zero as $\sigma \to 0$.

Now to determine an optimal support size σ_i we use a variant of Rissanen's minimum description length (MDL) principle [30] which is rooted in Occam's razor

"Entities should not be multiplied beyond necessity"

The MDL principle states that

"From several alternative models, the best one gives the minimum length of combined description of the model and the residuals"

The principle is closely connected with Sparse Approximation (SA) techniques [12, Section 8.1] which proved to be an excellent tool for an accurate reconstruction of natural signals from noisy data (see, for example, [8, 18] and references therein). Given a large collection of basis approximants (such a collection is usually called a dictionary), SA techniques seek an approximation of a noisy signal as a linear combination of the smallest number of approximants from the collection.

Let us assume that near approximation center $\mathbf{c}_i$ points of $\mathcal{P}_i \subset \mathcal{P}$ are generated by a stochastic process such that their mean positions are on local approximation $g_i(\mathbf{x}) = 0$ and their distances to $g_i(\mathbf{x}) = 0$ are normally distributed. Since the distance from $\mathbf{p}$ to $g_i(\mathbf{x}) = 0$ is accurately approximated by $g(\mathbf{p})/\|\nabla g(\mathbf{p})\|$, $E_{\text{local}}(\sigma)^2$ is proportional to the negative logarithm of the probability of observing points $\mathcal{P}_i$. Thus $E_{\text{local}}(\sigma)^2$ is proportional to the number of bits required to describe the points of $\mathcal{P}_i$ near $\mathbf{c}_i$.

We choose σ_i as the argmin for

$$E_{\text{SA}}(\sigma) = E_{\text{local}}(\sigma)^2 + \frac{C}{\sigma^2}, \tag{10}$$

where C is a positive constant. Note that the second term in (10) is proportional to the number of approximation centers used. Thus minimizing (10) estimates optimal influence parameter σ_i for approximation center $\mathbf{c}_i$ according to the MDL principle and leads to a sparse approximation of $\mathcal{P}$. Parameter C in (10) controls the trade off between sparsity

and approximation. As we will see later, C also controls reconstruction smoothness.

Let us describe the graph of $E_{\mathrm{SA}}(\sigma)$ qualitatively. If σ is large then the total number of approximation centers $\{\mathbf{c}_i\}$ is small and local approximation error $E_{\mathrm{local}}(\sigma)$ is large. We can conclude that $E_{\mathrm{SA}}(\sigma)$ grows drastically as $\sigma \to \infty$. If σ is small enough then local approximation error $E_{\mathrm{local}}(\sigma_i)$ is also small and the behavior of (10) determined by the second term equal to the number of approximation centers. The number of approximation centers grows sharply as $\sigma \to 0$ since for a sufficiently small σ the zero level-set of (1) must reproduce noisy behavior of $\mathcal{P}$.

In practice, we set C in (10) proportional to L^2, namely we use

$$C = (T_{\mathrm{SA}} L)^2. \tag{11}$$

We analyze a dependence between T_{SA} and the reconstruction quality in Section 4.

Fig. 2 presents typical graphs for $E_{\mathrm{local}}(\sigma)$ and $E_{\mathrm{SA}}(\sigma)$ measure at two different points of the Stanford bunny model.

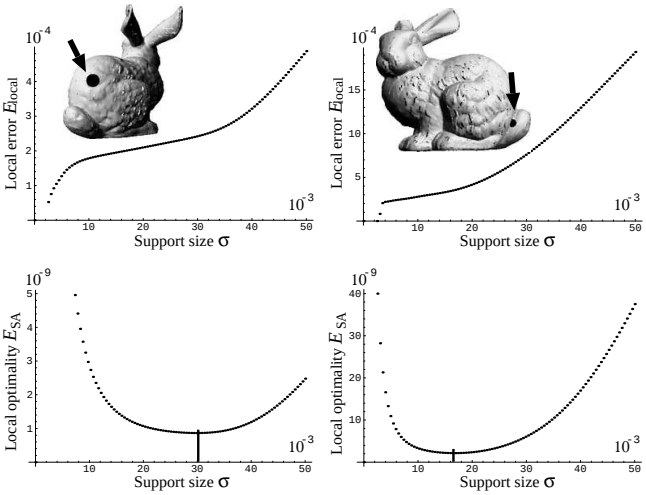

Fig. 2. Graphs of $E_{\mathrm{local}}(\sigma)$ and $E_{\mathrm{SA}}(\sigma)$ for two different points of the Stanford bunny model.

Fig. 3 visualize the spherical regions of influence (shown at 50% of their real size) centered at approximation centers $\mathbf{c}_i$. Notice that the value of σ_i reflects surface geometric complexity at $\mathbf{c}_i$: the bigger complexity, the smaller σ_i. Here the local geometric complexity depends on local quadratic approximation $g_i(\mathbf{x})$ at $\mathbf{c}_i$.

Minimizing $E_{\mathrm{SA}}(\sigma)$ is an one-dimensional problem. We use Brent's method [29] to minimize (10). In the most cases less than ten iterations are required to reach the minimum with $L/10^5$ accuracy.

It is interesting to compare two different strategies for selecting σ_i for the adaptive PU approximation given by the first sum in (5):

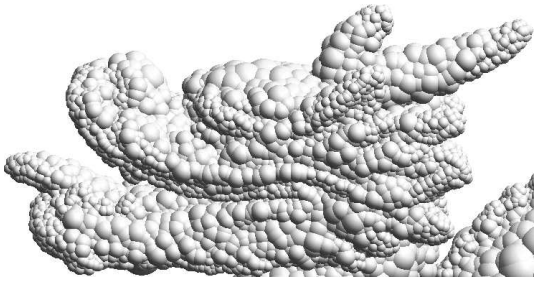

Fig. 3. Optimal σ_i for a part of the Stanford Dragon model are visualized as spheres of radius $\sigma_i/2$ centered at $\mathbf{c}_i$.

- minimizing SA energy (10) penalizing the number of local approximations;
- the selection of σ by solving the equation $E_{\mathrm{local}}(\sigma) = \epsilon_0$, where ϵ_0 is a user-specified accuracy [26].

Too small ϵ_0 leads reconstruction of noise and choosing too large ϵ_0 produces oversmoothing. As one can judge from the left images of Fig. 4, the optimal value of ϵ_0 may not exist: some regions of the Stanford bunny model are oversmoothed and noise is reconstructed at some other regions. In contrast, the MDL-based selection procedure leads to a very good reconstruction, as seen in the right images of Fig. 4.

Thus one can see that the sparse and adaptive PU approximation constructed via minimizing regularized energy (10) does a substantially better job than our previous technique [26] even for "clean" point datasets.

Selection of approximation centers. It is clear that a "good" cover is important for constructing a high quality partition of unity approximation. In our case, the cover consists of balls $\operatorname{supp}\phi_{\sigma_i}(\mathbf{c}_i)$ of radius σ_i centered at $\mathbf{c}_i$. We choose approximation centers $\{\mathbf{c}_i\}$ such that $\{\operatorname{supp}\phi_{\sigma_i}(\|\mathbf{x}-\mathbf{c}_i\|)\}$ covers all the points of $\mathcal{P}$. Further our aim is to generate a minimal cover with an amount of overlap greater than a certain threshold.

It is natural to measure the amount of overlap at $\mathbf{p}_j \in \mathcal{P}$ by

$$v_j = \sum_{i=1}^{M} \phi_{\sigma_i}(\|\mathbf{p}_j - \mathbf{c}_i\|).$$

We control overlapping by a user-specified parameter T_{overlap} and determine approximation centers $\{\mathbf{c}_i\}$ according to the following procedure proposed in [26].

Step 1. Assign $v_j = 0$ to each point $p_j \in \mathcal{P}$.

Step 2. Choose randomly m points (in our current implementation, we use $m = 15$) with $v < T_{\mathrm{overlap}}$.

Step 3. Select a point with minimum v among the points chosen at the previous step.

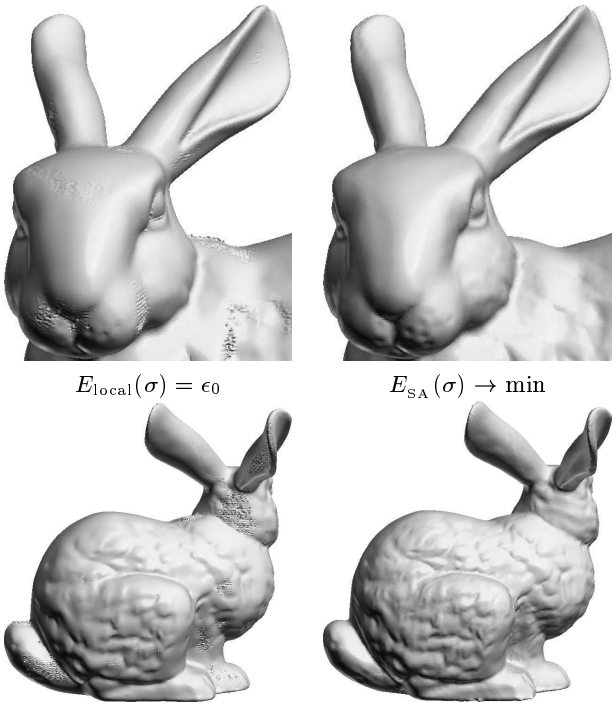

$E_{\rm local}(\sigma) = \epsilon_0$ $E_{\rm SA}(\sigma) \to \min$

Fig. 4. Reconstruction of the scanned bunny data consisting of 362K points using adaptive PU. Left: reconstruction using the measure in [26]. Some noisy parts are observed. Right: reconstruction using our new measure. No noise is observed while fine surface features (see, for example, the nose part of the bunny) are accurately reconstructed.

Step 4. Choose the point selected at Step 3 as an approximation center $\mathbf{c}_k \in \mathcal{C}$. Set $v_k = T_{\rm overlap}$ for $\mathbf{c}_k$.

Step 5. Find optimal support size σ_k and local polynomial approximation $g_k(\mathbf{x})$ at center $\mathbf{c}_k$ determined at the previous step.

Step 6. Update overlapping numbers v_j for all $\mathbf{p}_j \in \mathcal{P} \setminus \mathcal{C}$ by adding $\phi_{\sigma_k}(\|\mathbf{p}_j - \mathbf{c}_k\|)$.

Step 7. If there are points $\mathbf{p}_j \in \mathcal{P}$ with $v_j < T_{\rm overlap}$, go to Step 2.

Steps 2 and 3 implement a multiple choice technique, a powerful tool for randomized optimization [21] introduced recently in geometric modeling [38, 39].

According to our numerical experiments, the above procedure with $T_{\rm overlap} = 1.5$ produces a good cover. Selecting bigger values of overlapping parameter $T_{\rm overlap}$ does not improve the approximation quality and wastes computational time.

Fig. 5 demonstrates several intermediate steps in construction the approximation centers $\{\mathbf{c}_i\}$ and their corresponding influence radii $\{\sigma_i\}$ for the Stanford bunny model.

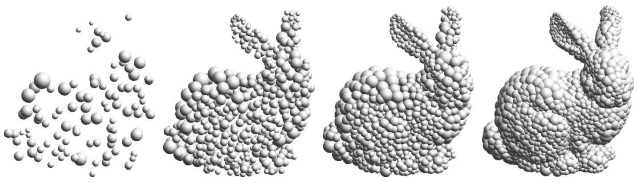

Fig. 5. Four intermediate stages of the approximation center selection procedure. The number of approximation centers increases from left to right and is equal to 100, 500, 1000, and 2000, respectively.

Reconstruction from noisy data. Small extra zero level-sets may appear when we apply our PU approximation procedure to datasets with a high level of noise. One possible remedy to avoid such unwanted artifacts consists of preventing influence parameter σ from being too small. So we introduce one more parameter $\sigma_{\rm min}$ and redefine $E_{\rm local}(\sigma)$ in interval $[0, \sigma_{\rm min}]$ by setting

$$E_{\rm local}(\sigma) = L \quad \text{if} \quad \sigma < \sigma_{\rm min}. \tag{12}$$

Then optimal σ is found from minimizing (10), (11) with redefined $E_{\rm local}(\sigma)$.

Fig. 6 compares reconstruction of the noisy Armadillo dataset without (the left image) and with (the right image) condition (12). Notice that the reconstruction with (12) leads to oversmoothing and small features are lost. However the RBF term in (5) allows us to reconstruct the lost details.

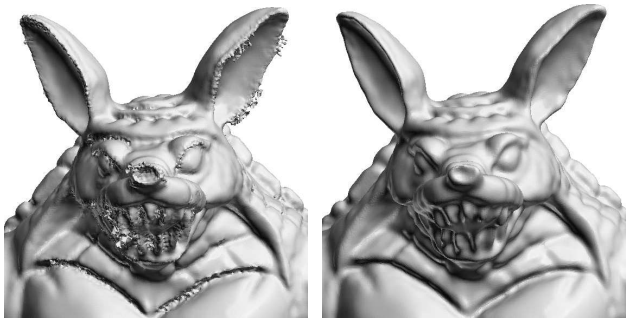

$\sigma_{\rm min} = 0, M = 134K$ $\sigma_{\rm min} = L/100, M = 42K$

Fig. 6. Adaptive PU reconstruction of a part of the Stanford Armadillo dataset without (left) and with (right) condition (12).

3 Least-Squares RBF Approximation

After the PU approximation is constructed we use a least square fitting procedure to determine unknown RBF weights $\{\lambda_i\}$ in (5). Let us define a global L^2-error metric by

$$E_{\rm global}(\boldsymbol{\lambda}) = \frac{1}{L}\sqrt{\frac{\sum_{j=1}^{N} d_j\, f(\mathbf{p}_j)^2}{\sum_{j=1}^{N} d_j}},$$

where $\boldsymbol{\lambda} = (\lambda_1, ..., \lambda_M)$ and the weights $\{d_j\}$ are defined in (6). Now the RBF weight can be found by minimizing $E_{\text{global}}(\boldsymbol{\lambda})$. However such least-squares fitting noisy data with RBFs often produces undesirable results, as demonstrated in Fig. 7.

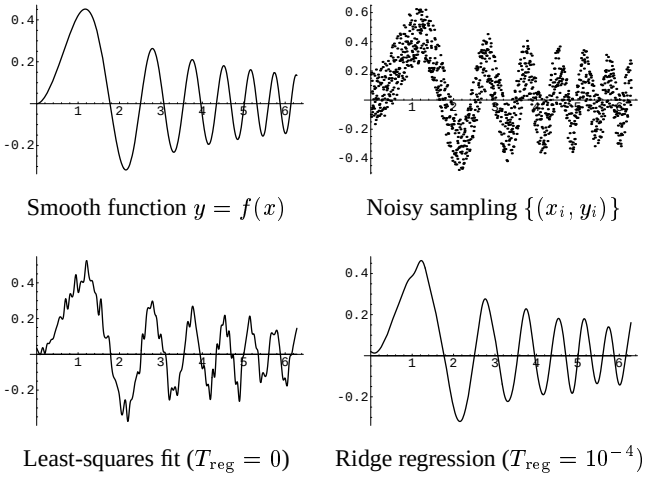

Fig. 7. Top-left: the graph of a smooth function. Top-right: noisy sampling data $\{(x_i, y_i)\}$, $y_i = f(x_i) + \epsilon_i$, is created by adding uniformly random vertical displacements ϵ_i. Bottom-left: Reconstruction from noisy data via least-squares RBF fitting; L^2-error between reconstructed graph and noisy data is 1.12×10^{-1}; L^2-error between reconstructed graph and $f(x)$ is 3.36×10^{-2}. Bottom-right: Reconstruction from noisy data via ridge regression: L^2-error between reconstructed graph and noisy data is 1.16×10^{-1}; L^2-error between reconstructed graph and $f(x)$ is 1.54×10^{-2}. Thus least squares give a better fit to the noisy data while ridge regression delivers a better approximation of the original data.

The overfitting problem exposed by the bottom-left image can be eliminated using the so-called regularization approach. In order to suppress oscillations let us determine $\boldsymbol{\lambda}$ from the following minimization problem

$$E_{\text{reg}}(\boldsymbol{\lambda}) = E_{\text{global}}(\boldsymbol{\lambda})^2 + T_{\text{reg}} \, \|\boldsymbol{\lambda}\|^2 \to \min, \qquad (13)$$

where $\|\boldsymbol{\lambda}\|$ is weighted norm given by

$$\|\boldsymbol{\lambda}\| = \sqrt{\frac{1}{M} \sum_{i=1}^{M} \left(\frac{\lambda_i}{\sigma_i}\right)^2}.$$

Minimization problem (13) is similar to ridge regression, a popular statistical robust estimation technique [15]. We have found that $T_{\text{reg}} = 10^{-5}$ works well for all our models. For example, the bottom-right image of Fig. 7 demonstrates how (13) suppresses overfitting.

Problem (13) is a quadratic minimization problem:

$$\frac{\partial E_{\text{reg}}(\boldsymbol{\lambda})}{\partial \boldsymbol{\lambda}} = \mathbf{0} \quad \Longleftrightarrow \quad (\mathbf{A} + T_{\text{reg}} \, \mathbf{D}) \, \boldsymbol{\lambda} = \mathbf{b}. \qquad (14)$$

Here $\mathbf{A}$ and $\mathbf{D}$ are $M \times M$ matrices and $\mathbf{b}$ is a column vector defined by

$$\begin{cases} A_{ij} = \dfrac{\sum_{k=1}^{N} d_k \, \Phi_{\sigma_i}(\|\mathbf{p}_k - \mathbf{c}_i\|) \, \Phi_{\sigma_j}(\|\mathbf{p}_k - \mathbf{c}_j\|)}{L^2 \, \sum_{k=1}^{N} d_k}, \\ D_{ii} = \dfrac{1}{M} \left(\dfrac{1}{\sigma_i}\right)^2, \\ b_i = \dfrac{\sum_{k=1}^{N} d_k \, \Phi_{\sigma_i}(\|\mathbf{p}_k - \mathbf{c}_i\|) \, \left(-f(\mathbf{p}_k)|_{\boldsymbol{\lambda}=\mathbf{0}}\right)}{L^2 \, \sum_{k=1}^{N} d_k} \end{cases}$$

Notice that $\mathbf{D}$ is a diagonal matrix and A_{ij} is equal to zero if there is no point $\mathbf{p}_k \in \mathcal{P}$ in the intersection of $\operatorname{supp} \phi_{\sigma_i}(\|\mathbf{x} - \mathbf{c}_i\|)$ and $\operatorname{supp} \phi_{\sigma_j}(\|\mathbf{x} - \mathbf{c}_j\|)$. Thus $\mathbf{A}$ is a sparse matrix in which the number of the non-zero elements in each row grows as T_{overlap} increases. When $T_{\text{overlap}} = 1.5$, the average number of non-zero elements in each row of $\mathbf{A}$ is about 100.

To solve (14), a sparse system of linear equations, we use the preconditioned biconjugate gradient method [29]. According to our numerical experiments, a few hundred iterations are required to reach a sufficient accuracy for M in the range of 10^3–10^6.

Fig. 8 presents adaptive PU + RBF reconstruction of the Armadillo part shown in Fig. 6 where only adaptive PU reconstruction used. Notice how well Armadillo's fine features are reconstructed.

4 Results and Discussion

Parameter selection. The smoothness of a reconstructed model depends mainly on parameter T_{SA} in (11) and (10). Three dragons of Fig. 9 are reconstructed via our adaptive PU+RBF approximation defined by (5). One can see how effectively T_{SA} controls the smoothness of the reconstructed models. If no smoothing or hole filling is required, we set $T_{\text{SA}} = 2 \times 10^{-6}$.

Moreover we found out that the L^2-error depends linearly on T_{SA} as demonstrated in the left image of Fig 10. We have no theoretical explanation of such an almost perfect linear dependence.

Another interesting observation concerns a dependence between reconstruction accuracy E_{global} and number of approximation centers M. A typical dependence graph is shown in the right image of Fig. 10. The number of approximation centers is reduced drastically if the reconstruction accuracy decreases.

Parameters T_{overlap} and T_{reg} are fixed in all our experiments. We set $\sigma_{\min} = 0$ for low-noise models and use $\sigma_{\min} = L/100$ for noisy datasets like the Stanford Armadillo model.

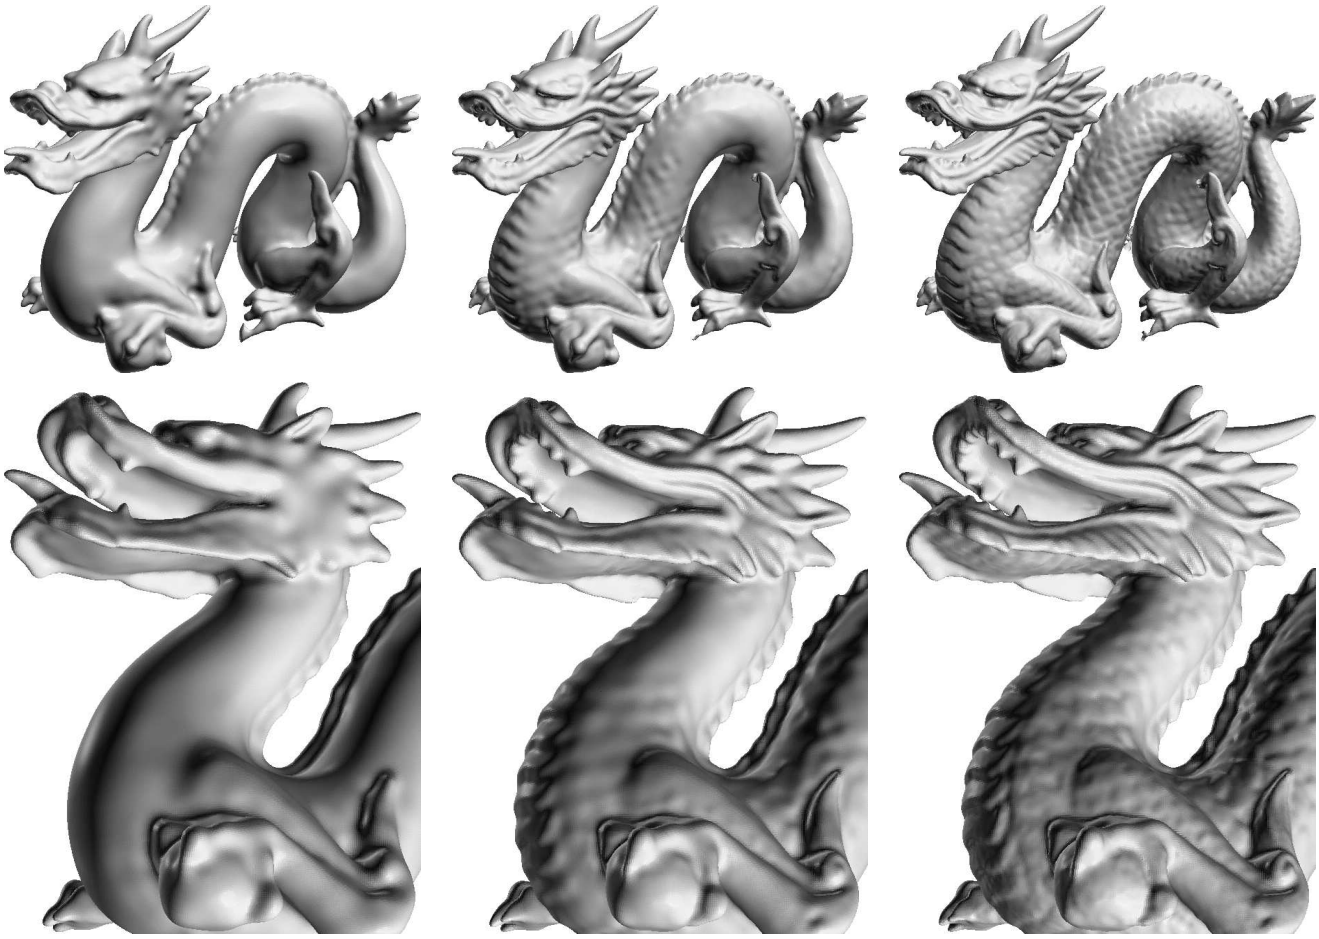

Fig. 9. Parameter T_{SA} in (11) and (10) controls reconstruction smoothness.

Filling holes. Our adaptive PU and PU + RBF reconstruction schemes have reasonable abilities in reconstructing missed data and filling holes. Fig. 11 demonstrates reconstruction of a squirrel model from an incomplete set of range scans.

Performance. To evaluate function $f(\mathbf{x})$ defined by the right hand-side of (5) at point $\mathbf{x}$ we need to find all the approximation centers $\mathbf{c}_i$ such that $\mathbf{x}$ belongs the intersection of their regions of influence. In order to do it quickly we use a range searching octree-based data structure proposed in [35] for illuminance storage.

For visualization of $f(\mathbf{x}) = 0$ we use Bloomenthal's polygonizer [4]. We found out that only one linear interpolation pass is required to find roots of $f(\mathbf{x}) = 0$ with a sufficient accuracy because near its zero level-set $f(\mathbf{x})$ mimics the distance function to the zero level-set. During the polygonization, $f(\mathbf{x})$ can be evaluated about $10,000$ - $20,000$ times per second. It means that only a few minutes are requited to generate a mesh with more than a million triangles.

Table 1 presents performance of our adaptive PU and PU + RBF reconstruction schemes. Computations were performed on a 1.6 GHz Mobile Pentium 4 PC with 1GB RAM, and timings are listed as minutes:seconds. Note that for a given point dataset the computational time and number of approximation centers depend not only on the size of the dataset but also on the geometric complexity. For the incomplete squirrel dataset choose a slightly lower value of T_{SA} than for the other models. Since M, the number of the approximation centers, is much smaller than N, the total number of points, the RBF reconstruction step is also fast.

Model	N	PU	RBF	T_{SA}	M
Squirrel	46K	00 : 06	00 : 04	5.0×10^{-5}	1.6K
Bunny	362K	01 : 00	00 : 42	2.0×10^{-6}	23K
Dragon	2.11M	12 : 58	06 : 20	2.0×10^{-6}	36K
Armadillo	2.37M	14 : 21	09 : 50	2.0×10^{-6}	41K

Table 1. Performance of our adaptive PU and PU + RBF reconstruction schemes.

Fig. 8. Adaptive PU + RBF reconstruction of the Armadillo part used also in Fig. 6.

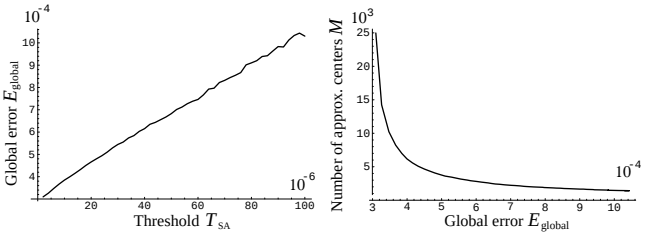

Fig. 10. Left: almost perfect linear dependence between T_{SA} and fitting error E_{global}. Right: typical dependence of reconstruction accuracy E_{global} from M, the number of approximation centers.

Statistical tools. An interesting feature of our approach consists of using statistical tools for an accurate fitting and automatic estimation of parameters. In this paper, we have employed a simple variant of the minimal description length principle, the ridge regression technique, and a multiple-choice procedure. We feel that statistical approaches have a huge potential in geometric modeling. It is also worth to note here that so far the main field of applications for RBF-based techniques is statistical learning [16, 15].

Directions for future research. In this study, we have selected the approximation centers $\mathcal{C}$ as a subset of the scattered points $\mathcal{P}$. In future we hope to drop this restriction and develop an method for optimal selecting $\mathcal{C}$. We think that a point clustering approach similar to that proposed in [9] may be useful. In particular, we hope that it will help us to treat uniformly datasets with small and large levels of noise.

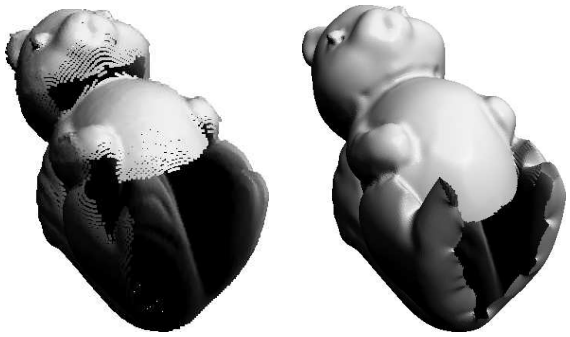

Fig. 11. PU+RBF reconstruction of a squirrel model from an incomplete set of range scans. Small holes are very well filled.

The method developed in this paper is more sensitive to density variations in scattered point datasets than, for example, a multi-scale approach proposed in [27] (see Fig. 12 for an extreme example). In future we plan to combine both the approaches.

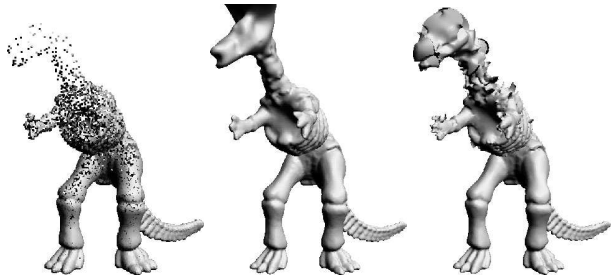

Fig. 12. Left: a scattered point dataset with sharply varying density. Middle: the points are interpolated by the method proposed in [27]. Right: the points are approximated via the technique developed in this paper.

Finally we are intrigued with deep relations between sparse approximations, MDL and similar statistical principles, and the general regularization theory [24].

Acknowledgments

We are grateful to the anonymous reviewers of this paper for their constructive comments and useful suggestions.

The Armadillo, Bunny and Dragon datasets are courtesy of the Stanford 3D scanning repository. The Dinosaur dataset is due to Cyberware.

References

[1] M. Alexa. Hierarchical partition of unity approximation. Technical report, TU Darmstadt, August 2002.

[2] I. Babuška and J. M. Melenk. The partition of unity method. *International Journal of Numerical Methods in Engineering*, 40:727–758, 1997.

[3] T. Belytschko, Y. Krongauz, D. Organ, M. Fleming, and P. Krysl. Meshless methods: An overview and recent developments. *Computer Methods in Applied Mechanics and Engineering*, 139:3–47, 1996.

[4] J. Bloomenthal. An implicit surface polygonizer. *Graphics Gems IV*, pages 324–349, 1994.

[5] D. Buhmann, M. *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003.

[6] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans. Reconstruction and representation of 3D objects with radial basis functions. In *Proceedings of ACM SIGGRAPH 2001*, pages 67–76, August 2001.

[7] J. C. Carr, R. K. Beatson, B. C. McCallum, W. R. Fright, T. J. McLennan, and T. J. Mitchell. Reconstruction and representation of 3D objects with radial basis functions. In *Proceedings of ACM GRAPHITE 2003*, pages 119–126, Melbourne, Australia, February 2003.

[8] D. L. Chen, S. S.and Donoho and Saunders M. A. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.

[9] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Analysis Machine Intell.*, 24(5):603–619, 2002.

[10] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of ACM SIGGRAPH 1996*, pages 303–312, 1996.

[11] H. Q. Dinh, G. Slabaugh, and G. Turk. Reconstructing surfaces using anisotropic basis functions. In *International Conference on Computer Vision (ICCV) 2001*, pages 606–613, Vancouver, Canada, July 2001.

[12] T. Evgeniou, M. Pontil, and T. Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13(1):1–50, 2002.

[13] M. S. Floater and A. A. Iske. Multistep scattered data interpolation using compactly supported radial basis functions. *Journal of Comp. Appl. Math.*, 73:65–78, 1996.

[14] F. Girosi, M. Jones, and T. Poggio. Regularization theory and neural networks architectures. *Neural Computation*, 7:219–269, 1995.

[15] H. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning*. Springer, 2001.

[16] S. Haykin. *Neural Networks: A Comprehensive Foundation*. Macmillan College Publishing Company, Inc., 1994.

[17] A. Iske and J. Levesley. Multilevel scattered data approximation by adaptive domain decomposition. Technical Report TUM-M0208, Technische Universität München, Juli 2002.

[18] E. LePennec and S. Mallat. Sparse geometric image representation with bandelets. *IEEE Trans. on Image Processing*.

[19] S. K. Lodha and R. Franke. Scattered data techniques for surfaces. In H. Hagen, G. Nielson, and F. Post, editors, *Proceedings of Dagstuhl Conference on Scientific Visualization*, pages 182–222. IEEE Computer Society Press, 1999.

[20] D. Lowe. Adaptive radial basis function nonlinearities, and the problem of generalization. In *1st IEE International Conference on Artificial Neural Networks*, pages 29–33, Bournemouth, UK, 1999.

[21] M. Mitzenmacher, A. Richa, and R. Sitaraman. The power of two random choices: A survey of techniques and results. In *Handbook of Randomized Computing, Chapter 9*. Kluwer, 2001.

[22] J. Moody and C. Darken. Fast learning in networks of locally-tuned processing units. *Neural Computation*, 1(2):281–294, 1989.

[23] B. S. Morse, T. S. Yoo, P. Rheingans, D. T. Chen, and K. R. Subramanian. Interpolating implicit surfaces from scattered surface data using compactly supported radial basis functions. In *Shape Modeling International 2001*, pages 89–98, Genova, Italy, May 2001.

[24] Tikhonov A. N. and V. Y. Arsenin. *Solution of Ill-Posed Problems*. Winston, Washington D. C., 1977.

[25] Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, and H.-P. Seidel. Multi-level partition of unity implicits. *ACM Transactions on Graphics*, 22(3):463–470, July 2003. Proceedings of SIGGRAPH 2003.

[26] Y. Ohtake, A. G. Belyaev, and H.-P. Seidel. Multi-scale and adaptive CS-RBFs for shape reconstruction from cloud of points. In *MINGLE workshop on Multiresolution in Geometric Modelling*, pages 337–348, Cambridge, UK, September 2003.

[27] Y. Ohtake, A. G. Belyaev, and H.-P. Seidel. A multi-scale approach to 3D scattered data interpolation with compactly supported basis functions. In *Shape Modeling International 2003*, pages 153–161, Seoul, Korea, May 2003.

[28] T. Poggio and F. Girosi. Networks for approximation and learning. *Proceedings of the IEEE*, 78:1481–1497, 1990.

[29] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1993.

[30] J. Rissanen. A universal prior for integers and estimation by Minimal Description Length. *The Annals of Statistics*, 11:131–138, 1983.

[31] V. V. Savchenko, A. A. Pasko, O. G. Okunev, and T. L. Kunii. Function representation of solids reconstructed from scattered surface points and contours. *Computer Graphics Forum*, 14(4):181–188, 1995.

[32] G. Turk and M. Levoy. Zippered polygon meshes from range images. In *Proceedings of ACM SIGGRAPH 1994*, pages 311–318, July 1994.

[33] G. Turk and J. O'Brien. Shape transformation using variational implicit surfaces. In *Proceedings of SIGGRAPH 1999*, pages 335–342, August 1999.

[34] G. Turk and J. O'Brien. Modelling with implicit surfaces that interpolate. *ACM Transactions on Graphics*, 21(4):855–873, October 2002.

[35] G. J. Ward, F. M. Rubinstein, and R. D. Clear. A ray tracing solution for diffuse interreflection. In *Proceedings of ACM SIGGRAPH 1988*, pages 85 – 92, 1988.

[36] H. Wendland. Piecewise polynomial, positive definite and compactly supported radial basis functions of minimal degree. *Advances in Computational Mathematics*, 4:389–396, 1995.

[37] H. Wendland. Fast evaluation of radial basis functions: Methods based on partition of unity. In L. Schumaker and J. Stöckler, editors, *Approximation Theory X: Wavelets, Splines, and Applications*, pages 473–483. Vanderbilt University Press, Nashville, 2002.

[38] J. Wu and L. P. Kobbelt. Fast mesh decimation by multiple-choice techniques. In *Vision, Modeling, Visualization 2002 Proceedings*, pages 241–248, Erlangen, Germany, November 2002.

[39] J. Wu and L. P. Kobbelt. A stream algorithm for the decimation of massive meshes. In *Graphics Interface 2003 Proceedings*, pages 185–192, Halifax, Canada, June 2003.