Smooth surface reconstruction from noisy range data

J. C. Carr[1] R. K. Beatson[2] B. C. McCallum[1] W. R. Fright[1] T. J. McLennan[1] T. J. Mitchell[1]

[1]Applied Research Associates NZ Ltd*
[2] University of Canterbury†

Abstract

This paper shows that scattered range data can be smoothed at low cost by fitting a Radial Basis Function (RBF) to the data and convolving with a smoothing kernel (low pass filtering). The RBF exactly describes the range data and interpolates across holes and gaps. The data is smoothed during evaluation of the RBF by simply changing the basic function. The amount of smoothing can be varied as required without having to fit a new RBF to the data. The key feature of our approach is that it avoids resampling the RBF on a fine grid or performing a numerical convolution. Furthermore, the computation required is independent of the extent of the smoothing kernel, i.e., the amount of smoothing. We show that particular smoothing kernels result in the applicability of fast numerical methods. We also discuss an alternative approach in which a discrete approximation to the smoothing kernel achieves similar results by adding new centres to the original RBF during evaluation. This approach allows arbitrary filter kernels, including anisotropic and spatially varying filters, to be applied while also using established fast evaluation methods. We illustrate both techniques with LIDAR laser scan data and noisy synthetic data.

CR Categories: I.3.5 [Curve, surface, solid, and object representations]: Splines;

Keywords: Surface reconstruction, point-cloud surfacing, noise reduction, spline smoothing, low pass filtering, anti-aliasing, Radial Basis Functions.

1 Introduction

This paper is concerned with smoothing large scattered surface data sets that contain noise or unwanted detail. The surfaces are typically irregularly sampled, exhibit varying sampling densities and may contain large regions with few or no measurements, e.g. due to occlusion. Such characteristics commonly occur when scanning complicated objects. Figure 1 illustrates a typical example. Few techniques are available for automatically modelling this type of data. Implicit modelling of surfaces with Radial Basis Functions (RBFs) is a relatively recent technique in computer graphics which offers the ability to interpolate across large, irregular holes in incomplete surface data without constraining the topology of an object or relying on *a priori* knowledge of its shape [Savchenko et al. 1995; Turk and O'Brien 1999; Morse et al. 2001; Carr et al. 2001]. While the results of RBF interpolation are encouraging for clean data, problems remain with modelling noisy real-world data. Furthermore, the issue of avoiding aliasing artefacts when evaluating an RBF at low resolutions has not been discussed in the literature. In this paper we show that an RBF representation of surface data set can be smoothed by simply substituting the appropriate basic function when evaluating the RBF. Furthermore, the applicability of existing fast evaluation methods for certain types of smoothing kernel make this approach computationally very effective.

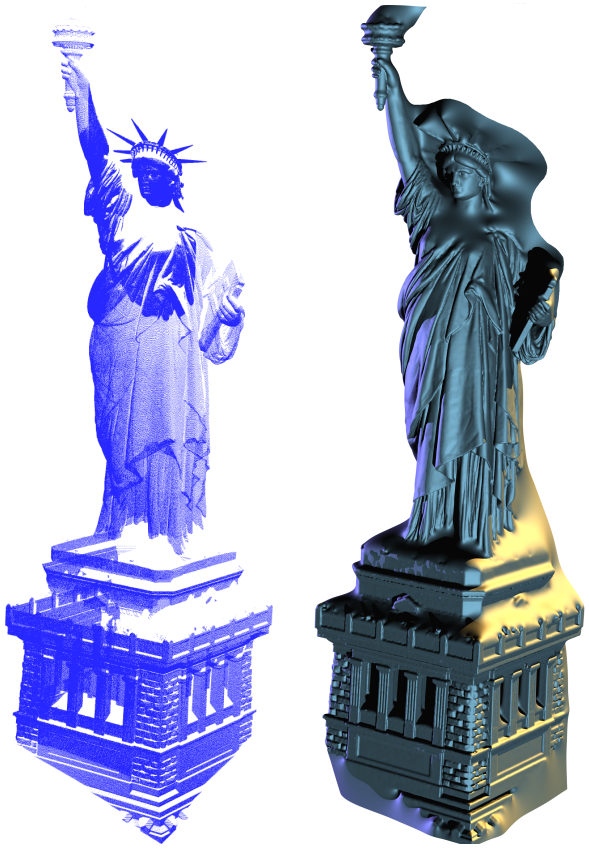

Figure 1: Example of a typical surface fitting problem. (Left) Raw point-cloud from a LIDAR scan of the Statue of Liberty. (Right) RBF surface. The few viewpoints, oblique scanning angles and large gaps make this a difficult problem for other techniques.

1.1 Implicit surface modelling with RBFs

[Savchenko et al. 1995] and [Turk and O'Brien 1999] first introduced the concept of modelling surfaces implicitly as the zero-set

*Applied Research Associates NZ Ltd, PO Box 3894, Christchurch, New Zealand. Email: [j.carr,r.fright,b.mccallum]@aranz.com Web: www.aranz.com

†Dept. Mathematics and Statistics, University of Canterbury, Christchurch, New Zealand, Email: r.beatson@math.canterbury.ac.nz

of an RBF for surface point-cloud data and mesh data. In that context the RBF is a continuous 3D function which is positive on one side of the surface and negative on the other (See Figure 2). The zero-set of the function defines a surface that smoothly interpolates between the input surface points or mesh vertices, depending on the case. RBFs are popular for interpolating scattered data as the associated system of linear equations is guaranteed to be invertible under very mild conditions on the locations of the data points [Cheney and Light 1999; Micchelli 1986]. For example, the pseudo-cubic spline used by Turk *et al.* only requires that the data points are not co-planar. The process of fitting and evaluating an RBF, however, ordinarily involves intensive computation and is infeasible for more than a few thousand surface points. Consequently early RBF surface models were constrained to small data sets. More recently [Morse et al. 2001] showed that large data sets can be modelled efficiently by using compactly supported basic functions. However, the nature of a compact support inherently limits the ability to interpolate across large holes and introduces a scale parameter.

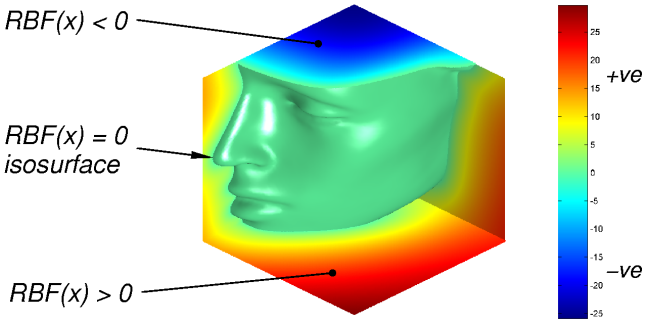

Figure 2: Implicit model of a surface using a biharmonic Radial Basis Function (RBF). The RBF describes the distance to the surface.

[Carr et al. 2001] have demonstrated that RBFs can be used to reconstruct surfaces from scattered point-cloud data consisting of millions of points by exploiting fast algorithms from recent work in numerical approximation [Beatson and Light 1997; Beatson et al. 2001]. In particular, they showed that the non-compactly supported family of RBFs, known as polyharmonic splines, exhibit remarkable extrapolation and interpolation (hole-filling) properties due to their least energy characterisation. In that work a polyharmonic spline was used to model a signed-distance metric computed near an object's surface. Off-surface points were projected along estimates of the surface normal. Points on the surface were assigned a zero value while points off the surface were assigned negative distances, if inside the object, and positive distances if outside. The line-of-sight to the scanner helped to determine which was the case. An RBF was then fitted to the scattered distance data, and the surface was reconstructed by iso-surfacing the RBF. The use of the biharmonic (thin-plate) spline interpolant gave a visually pleasing and physically plausible result. Figure 2 is an example of fitting a biharmonic spline to incomplete mesh data following the method of [Carr et al. 2001]. Near the surface the RBF approximates a signed-distance function. In this paper we extend this work by developing a Low Pass Filtering (LPF) strategy for noisy data which we call *implicit smoothing*.

1.2 RBF interpolation

Before commencing a review of relevant techniques for fitting RBFs to noisy data, it is helpful to make a few definitions. In general, an RBF is a function of the form

$$s(\boldsymbol{x}) = p(\boldsymbol{x}) + \sum_{i=1}^{N} \lambda_i \phi(|\boldsymbol{x} - \boldsymbol{x}_i|), \tag{1.1}$$

where p is a polynomial of low degree and the *basic function* ϕ is a real valued function on $[0, \infty)$, usually unbounded and of non-compact support. We call the λ_i's and the $\boldsymbol{x}_i$'s the *weights* and *centres* of the RBF. We refer to the process of solving for the weights as *fitting* the RBF, and computing $s(\boldsymbol{x})$ as *evaluating* the RBF. Popular choices for ϕ when modelling surface data include the non-compactly supported biharmonic spline $\phi(r) = r$, the triharmonic spline $\phi(r) = r^3$ with quadratic polynomial, and the pseudo-cubic spline $\phi(r) = r^3$ with linear polynomial [Turk et al. 2001].

The biharmonic and triharmonic splines are known as "smoothest" interpolators in the sense that they minimise certain energy functionals and interpolate the data. For example, given a set of nodes $\{\boldsymbol{x}_i\}_{i=1}^N \subset \mathbb{R}^3$ and a set of function values $\{f_i\}_{i=1}^N \subset \mathbb{R}$, the biharmonic RBF $s(\boldsymbol{x})$ satisfies the interpolation conditions $s(\boldsymbol{x_i}) = f_i$ and minimises

$$\|s\|^2 = \int_{\mathbb{R}^3} \left(\frac{\partial^2 s(\boldsymbol{x})}{\partial x^2}\right)^2 + \left(\frac{\partial^2 s(\boldsymbol{x})}{\partial y^2}\right)^2 + \left(\frac{\partial^2 s(\boldsymbol{x})}{\partial z^2}\right)^2 + 2\left(\frac{\partial^2 s(\boldsymbol{x})}{\partial x \partial y}\right)^2 + 2\left(\frac{\partial^2 s(\boldsymbol{x})}{\partial x \partial z}\right)^2 + 2\left(\frac{\partial^2 s(\boldsymbol{x})}{\partial y \partial z}\right)^2 d\boldsymbol{x}. \tag{1.2}$$

$\|s\|^2$ is a measure of the energy in the second derivative of s. The triharmonic spline minimises the corresponding energy in the 3rd derivative of s.

1.3 RBF approximation and smoothing

In many practical problems exact interpolation is not required nor, in the case of noisy data, desired. One way to acknowledge this is to solve the interpolation system only to within a certain accuracy. Thus, one seeks a spline of the form specified by Equation (1.1) with low energy (high smoothness), such that

$$|f_i - s(\boldsymbol{x}_i)| \leq \epsilon, \qquad i = 1, \ldots, N. \tag{1.3}$$

Such a model can be applied to fitting signed distance data from a laser scanner. The f_i would then be an estimate of the signed distance to the surface, and ϵ could be chosen as an estimate of the scanner error. This inexact fitting approach is followed by the greedy algorithm fits of [Carr et al. 2001] where they seek to add centres until the fit is sufficiently accurate. There is an inherent tendency for such greedy algorithms to favour absorption of the low frequencies first. Thus, while the criteria (1.3) does not enforce any optimality, it usually leads to satisfactorily smooth fits with considerable data reduction (i.e., fewer than N centres).

In many situations it will be reasonable to assume instead that each measurement is the combination of a smooth underlying signal and random noise. It is therefore sensible to trade smoothness of fitted function against fidelity to the data as in problem 1.1.

Problem 1.1. (Spline smoothing) *Given the nodes* $\{\boldsymbol{x}_i\}_{i=1}^N \subset \mathbb{R}^d$

$$\text{minimize} \qquad \rho\|s\|^2 + \frac{1}{N}\sum_{i=1}^{N} \big(s(\boldsymbol{x}_i) - f_i\big)^2, \tag{1.4a}$$

where $\rho \geq 0$ *and* $\|\cdot\|$ *denotes the smoothness penalty defined by Equation* (1.2).

This approach is known as *spline smoothing*. The parameter ρ balances smoothness against fidelity to the data [Wahba 1990]. The solution to Problem 1.1 is also an RBF of the form (1.1) but it approximates the given values f_i at the nodes $\boldsymbol{x}_i$. [Dinh et al. 2001] and [Carr et al. 2001] have used spline smoothing to smooth noisy surface data. Rather than choosing a global value for ρ, [Dinh et al. 2001] define a value at each data point, ρ_i, and adaptively vary smoothing to preserve sharp corners and edges in the surface.

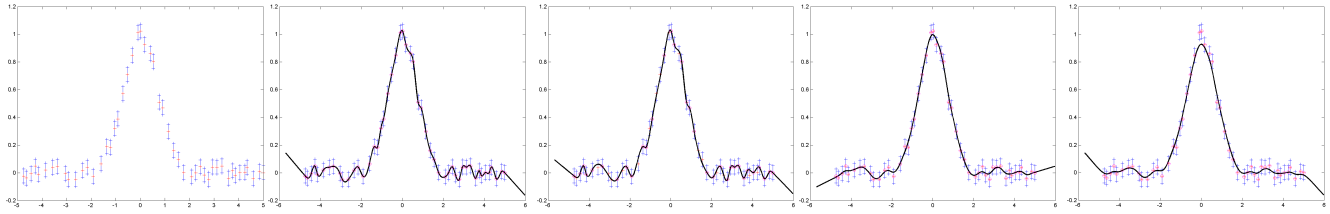

(a) Input data (b) Exact fit (c) "Greedy" fit (d) Spline smoothing (e) Implicit smoothing

Figure 3: Smoothing techniques for fitting an RBF in 1D.

The difficulty with spline smoothing lies with interpreting ρ in a physically meaningful way and choosing a sensible value. In the context of modelling a surface, we often want to attenuate high frequencies (random variations) in the surface data that are attributed to noise or to fit a smooth surface within some distance of the actual surface measurements. The latter requires an estimate of the *magnitude* of the noise present and this is related to the accuracy of the scanner. In the case of spline smoothing, there is no obvious relationship between ρ and the frequencies attenuated by fitting the approximating RBF. The deviation from the surface nodes is related to ρ but only by the product of ρ and the RBF weights λ_i which we are trying to solve for, i.e.,

$$s(\boldsymbol{x}_i) = \boldsymbol{f}_i + \rho\boldsymbol{\lambda}_i. \tag{1.5}$$

Consequently, estimating ρ while maintaining a given fitting tolerance is difficult. It is computationally expensive to iteratively choose ρ since this involves fitting and evaluating an RBF multiple times. Furthermore, in Equation (1.5) the λ_i are particular to a given data set and so the corresponding value for ρ will not generally have a similar effect on different data sets.

Despite the difficulties with choosing parameters, spline smoothing is very effective (Figure 7) and many authors make their parameter choices based on experience with particular types of data.

1.4 A low pass filter approach to smoothing

In this paper we consider traditional low pass filtering (LPF) of noisy data. The advantages of this approach in the context of interpolating surface data with RBFs are

- There is an intuitive control parameter (noise frequency)
- Efficient implementation — computation is independent of kernel size
- Smoothing can be varied *a posteriori* during RBF evaluation

The method involves fitting an RBF to the data and then convolving with a smoothing kernel (a low pass filter). The linear nature of the RBF means that this convolution can be carried out *implicitly* through a change of basic function during the evaluation of the RBF. The computation otherwise associated with smoothing the raw data through convolution, by spatial or Fourier means, is avoided. Furthermore, there is no requirement for the RBF to be evaluated on a fine grid, and the degree of smoothing can be varied. The computation required is independent of the extent of the corresponding kernel (i.e., the amount of smoothing). Fast evaluation methods can be used to evaluate the smoothed RBF for certain special choices of ϕ and corresponding smoothing kernels.

Figure 3 illustrates in 1D the different approaches considered here. Figure 3(a) is a noisy test data set, Figure 3(b) an exact fit RBF and Figure 3(c) is a 'greedy' fit as described in [Carr et al. 2001]. Figure 3(d) is an example of spline smoothing while Figure 3(e) illustrates the LPF approach achieved by implicit smoothing.

The remainder of the paper is organised as follows. Section 2 introduces the LPF approach to smoothing. Section 3 describes how filtering is realised implicitly by a change of basic function and gives the smoothing kernels that allow the use of fast evaluation methods. In Section 4 we develop a discrete approximation that leads to the possibility of alternative kernels, including anisotropic and spatially varying kernels. Section 5 discusses results, and conclusions are presented in Section 6.

2 Low pass filtering

Low pass filtering is well-understood in signal processing. Although such techniques have been applied to the manipulation of geometric data in computer graphics before [Taubin 1995; Pauly and Gross 2001], the use of Fourier-based smoothing techniques has been hindered by the irregular nature of geometric data, both point cloud and mesh data. [Taubin 1995; Pauly and Gross 2001] and others have developed techniques for estimating the geometric spectrum locally over 2D surface patches and consequently have applied 2D linear filtering to smooth meshes or enhance edges. In this section we consider filtering the 3D distance field described by an RBF rather than the surface which it implicitly defines. A very similar surface smoothing effect is achieved. This indirect approach avoids the former methods' problem of deriving a tessellation of 2D surface patches that are topologically equivalent to discs over the object's surface. The linear nature of the RBF representation means that adapting Fourier filtering techniques is relatively straightforward, despite the highly irregular sampling of the underlying surface data.

2.1 The RBF as a convolution

We begin by noting that the RBF $s(\boldsymbol{x})$ defined in Equation (1.1) can be written as

$$s(\boldsymbol{x}) = p(\boldsymbol{x}) + (W \star \Phi)(\boldsymbol{x}) \tag{2.1}$$

where $\star$ is the convolution operator and $\Phi(\boldsymbol{x}) = \phi(|\boldsymbol{x}|)$. The term $W(\boldsymbol{x})$ depends solely on the weights λ_i and centres $\boldsymbol{x}_i$, as defined by

$$W(\boldsymbol{x}) = \sum_{i=1}^{N} \lambda_i \delta(\boldsymbol{x} - \boldsymbol{x}_i) \tag{2.2}$$

where $\delta(\boldsymbol{x})$ is the Dirac delta function.

We now consider forming a smoothed function $s_s(\boldsymbol{x})$ by applying a low pass filter to the RBF $s(\boldsymbol{x})$. Mathematically this is represented by a convolution,

$$s_s(\boldsymbol{x}) = (s \star h)(\boldsymbol{x}) \tag{2.3}$$

where $h(\boldsymbol{x})$ is the impulse response of the filter. In an image processing context $h(\boldsymbol{x})$ is also known as the *point-spread function*

(PSF) of a blurring (smoothing) process. In this paper we refer to $h(\boldsymbol{x})$ as the smoothing kernel which we convolve with an RBF. Since convolution is commutative, Equation (2.3) can be written as either

$$s_s(\boldsymbol{x}) = (p \star h)(\boldsymbol{x}) + [W \star (\Phi \star h)]\,(\boldsymbol{x}) \tag{2.4}$$

or

$$s_s(\boldsymbol{x}) = (p \star h)(\boldsymbol{x}) + [(W \star h) \star \Phi]\,(\boldsymbol{x}) \tag{2.5}$$

The two forms of the convolution suggest two different smoothing schemes, the first is substitution of a smoother function for the original basic function (Section 3), the second is a discrete approximation achieved by adding new centres to the RBF (Section 4).

3 Basic function substitution

Equation (2.4) suggests that smoothing, or indeed any type of linear filtering, can be achieved by replacing the basic function Φ with a new, *smoother* basic function Ψ where

$$\Psi(\boldsymbol{x}) = (\Phi \star h)(\boldsymbol{x}). \tag{3.1}$$

Unlike spline smoothing, where the smoothing is achieved during fitting, this approach gives us the ability to smooth data *after* an RBF has been fitted, during evaluation. The computationally intensive task of solving the RBF system is done once. The weights λ_i are determined for an 'exact' fit to the data using Φ as the basic function. During evaluation, Φ is replaced by the smoother basic function Ψ. The degree of smoothing is varied by choosing Ψ appropriately. The full detail of the raw data is retained in the λ_i coefficients. This approach is useful when an *a priori* estimate of noise is not available, or when the RBF is being sub-sampled at different resolutions, in which case $h(\boldsymbol{x})$ is interpreted as an anti-aliasing filter. Filtering is achieved *implicitly* through the change of basic function, thus avoiding the computationally intensive alternatives of convolution on a regular grid by spatial or Fourier methods.

The observation that a smoother basic function Ψ can be substituted for Φ during RBF evaluation is of limited practical benefit unless corresponding fast methods exist for the efficient evaluation of an RBF containing Ψ. In the special case where Φ is an odd power of the modulus, which includes the biharmonic ($|\boldsymbol{x}|$) and triharmonic ($|\boldsymbol{x}|^3$) basic functions, smoothing kernels exist that, when convolved with Φ, result in multiquadric basic functions (as defined in Section 3.1). The availability of fast fitters for powers of the modulus [Beatson et al. 2000], and fast evaluators for the generalised multiquadric ($\mathcal{O}$(N log N)) [Cherrie et al. 2002], allows the method to be applied to large data sets involving millions of scattered data points.

This implicit smoothing point of view has been used to analyse the error in quasi-interpolation by 1D multiquadrics [Beatson and Dyn 1996]. Here we realise the advantages of applying the approach to smoothing 3D data.

3.1 Generalised multiquadrics

We define the generalised multiquadric basic function as

$$\Psi_{\beta,c}(\boldsymbol{x}) = (|\boldsymbol{x}|^2 + c^2)^{\beta/2}, \qquad \boldsymbol{x} \in \mathbb{R}^d. \tag{3.2}$$

where $c > 0$. These functions are most often considered in the case that β is a positive odd integer. $\Psi_{\beta,c}$ can be viewed as a smoothed out version of $|\boldsymbol{x}|^\beta$. This is readily apparent in Figure 4(a) which compares $|\boldsymbol{x}|$ with the multiquadric ($\beta = 1$ case).

By generalising the 1D results of [Beatson and Dyn 1996] the relationship between $|\boldsymbol{x}|^\beta$ and $(|\boldsymbol{x}|^2 + c^2)^{\beta/2}$ can be expressed as

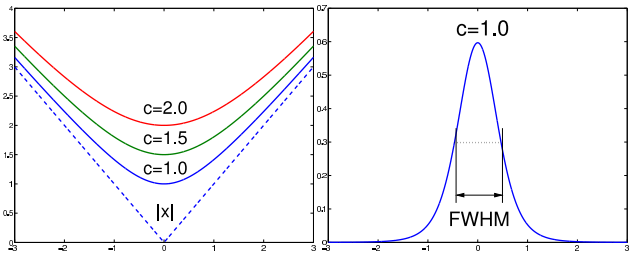

(a) The multiquadric (for various values of c) and 3D biharmonic basic functions.

(b) The $h_{3,1}(\boldsymbol{x})$ smoothing kernel plotted with c=1.0.

Figure 4: The multiquadric depicted in (a) can be written as the convolution of $|\boldsymbol{x}|$, the biharmonic basic function in 3D (also shown in (a)) with the $h_{3,1}(\boldsymbol{x})$ smoothing kernel plotted in (b).

smoothing by convolution. More precisely, in odd dimensions (odd d) and $\beta > 0$

$$\Psi_{\beta,c}(\boldsymbol{x}) = (\Phi_\beta \star h_{d,\beta})(\boldsymbol{x}), \tag{3.3}$$

where $\Phi_\beta = |\boldsymbol{x}|^\beta$ and the convolution kernel $h_{d,\beta}(\boldsymbol{x})$ is the generalised multiquadric with negative index $\Psi_{-\beta-2d,c}(\boldsymbol{x})$, normalised to have integral one. That is

$$h_{d,\beta}(\boldsymbol{x}) = a_{d,\beta}\Psi_{-\beta-2d,c}(\boldsymbol{x}), \tag{3.4}$$

where

$$a_{d,\beta} = \pi^{-d/2} c^{d+\beta} \frac{\Gamma((\beta+2d)/2)}{\Gamma((\beta+d)/2)}. \tag{3.5}$$

and Γ denotes the Gamma function [Abramowitz and Stegun 1965]. We are particularly interested in the biharmonic spline in 3D ($d = 3, \beta = 1$). In Table 1 we give the biharmonic and triharmonic basic functions that have been used to model surface data and their corresponding smoothed basic function and the associated smoothing kernel. The smoothing kernel $h_{3,1}$ corresponding to the 3D biharmonic basic function is plotted in Figure 4(b). The linear polynomial component of the 3D biharmonic RBF is preserved under convolution with the kernel $h_{3,1}$.

The Fourier transform of $h_{d,\beta}(\boldsymbol{x})$ is

$$H_{d,\beta}(\boldsymbol{u}) = \frac{2^{1-(d+\beta)/2}}{\Gamma((\beta+d)/2)} K_{\frac{\beta+d}{2}}(c|\boldsymbol{u}|)\,(c|\boldsymbol{u}|)^{(\beta+d)/2}, \tag{3.6}$$

where K_ν is the usual modified Bessel function [Abramowitz and Stegun 1965, pp.374–379]. Figure 5 illustrates the low pass nature of the power spectrum of the 3D biharmonic smoothing kernel for $c = 1$. The 3dB (half power) frequency is also shown.

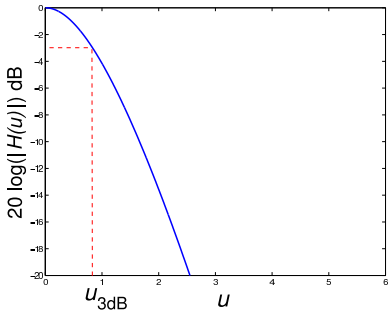

Figure 5: Power spectrum of the $h_{3,1}(\boldsymbol{x})$ smoothing kernel.

The relationship between the multiquadric parameter c and the degree of smoothing can be characterised by the effective width

β	Un-smoothed $\Phi(\boldsymbol{x}) = \lvert\boldsymbol{x}\rvert^{\beta}$	Smoothed $\Psi(\boldsymbol{x}) = \left(\lvert\boldsymbol{x}\rvert^2 + c^2\right)^{\frac{\beta}{2}}$	Smoothing kernel $h_{d,\beta}$	FWHM
1	3D biharmonic $\lvert\boldsymbol{x}\rvert$	$\sqrt{\lvert\boldsymbol{x}\rvert^2 + c^2}$	$\frac{15c^4}{8\pi}\left(\lvert\boldsymbol{x}\rvert^2 + c^2\right)^{-7/2}$	$2\sqrt{2^{2/7} - 1}\, c$ $\approx 0.9360\, c$
3	3D triharmonic $\lvert\boldsymbol{x}\rvert^3$	order 3 multiquadric $\left(\lvert\boldsymbol{x}\rvert^2 + c^2\right)^{(3/2)}$	$\frac{105c^6}{32\pi}\left(\lvert\boldsymbol{x}\rvert^2 + c^2\right)^{-9/2}$	$2\sqrt{2^{2/9} - 1}\, c$ $\approx 0.8162\, c$

Table 1: The relationship between basic functions often used to model surfaces, the multiquadric and corresponding smoothing kernel h.

of the smoothing kernel $h(\boldsymbol{x})$ at half its maximum. This is easily computed from the kernel formulae and is more convenient than computing the 3dB cutoff frequency shown in Figure 5 which is determined numerically from Equation (3.6). In Table 1 the Full Width at Half Maximum (FWHM) is given in terms of the parameter c. Since the FWHM $\approx c$, then c can be loosely interpreted as the scale below which finer detail is lost or severely attenuated.

Figure 6 shows the results of fitting a 3D biharmonic RBF 'exactly' to noisy LIDAR data (to an accuracy of 10^{-6}mm), then evaluating the RBF without smoothing to generate the fitted surface. A small sub-region, approximately 2m $\times$ 0.9m with a depth of 0.5m, taken from a larger scan, is shown. Figure 7 illustrates the results of implicit smoothing, i.e. evaluating the same RBF with the multiquadric basic function. Results are given for several values of the smoothing parameter c and compared with similar results using spline smoothing. Note that for each spline smoothing result a new RBF was computed and evaluated while a single RBF was fitted once and evaluated with three different values for c in the implicit smoothing case. The values of ρ were chosen using trial and error to find values giving similar amounts of smoothing.

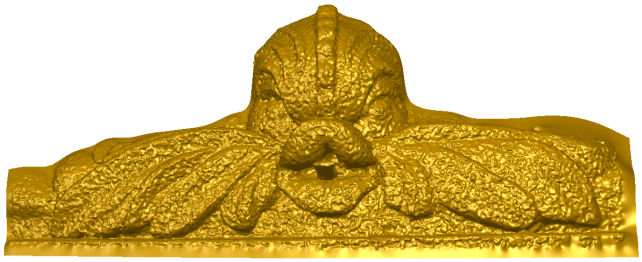

Figure 6: An 'exact' fit, to an accuracy of 10^{-6}mm, of an RBF to the noisy LIDAR data studied in Figure 7.

Figure 8 is an example using artificial data. Figure 8(a) is the pseudo 'distance' data. The samples lie within the unit cube and the values of the distance function lie in a range from -0.5 to 0.5. Figure 8(b) is the result of fitting an RBF to the data and evaluating an isosurface corresponding to $s(\boldsymbol{x}) = 0$. In Figure 8(c) random uniform noise of magnitude ± 0.1 was added to the data and an RBF was then fitted and isosurfaced at the same threshold. The addition of the noise has added a high frequency component to the isosurface which is related to the local sampling density of the raw data. Figure 8(d) is the result of evaluating the same RBF but with implicit smoothing and a kernel width corresponding to $c = 0.1$. Features occurring at a smaller scale than 0.1 are attenuated while the low frequency components of the data, which correspond to the genuine noise-free signal, are largely preserved.

4 Discrete smoothing

Recall that in Section 2.1 the smoothed RBF $s_s(\boldsymbol{x})$ was written in two different forms. The second form (Equation (2.5)) suggests that it would be useful to consider the case in which $h(\boldsymbol{x})$ consists of a number of delta functions, since then the term $W \star h$ would also consist of delta functions. Let us call such a kernel $\overline{h}(\boldsymbol{x})$ since we will show that one can think of $\overline{h}(\boldsymbol{x})$ as a sampled version of a continuous kernel $h(\boldsymbol{x})$. On writing

$$\overline{h}(\boldsymbol{x}) = \sum_{j=1}^{M} h_j \delta(\boldsymbol{x} - \boldsymbol{l}_j) \tag{4.1}$$

it follows from Equation (1.1) that

$$s_s(\boldsymbol{x}) = \sum_{j=1}^{M} h_j p(\boldsymbol{x} - \boldsymbol{l}_j) + \sum_{i=1}^{N} \lambda_i \sum_{j=1}^{M} h_j \Phi(\boldsymbol{x} - \boldsymbol{x}_i - \boldsymbol{l}_j) \tag{4.2}$$

Note that the first term is a polynomial of the same degree as $p(\boldsymbol{x})$. Now on comparing Equation (4.2) with Equation (1.1), it is evident that $s_s(\boldsymbol{x})$ is also an RBF, with M times as many centres as that of $s(\boldsymbol{x})$. Thus, the fast methods used in [Carr et al. 2001] can be used to evaluate $s_s(\boldsymbol{x})$ efficiently. The convolution can thereby be performed implicitly (as was the case for the method described in Section 3), rather than evaluating $s(\boldsymbol{x})$ on a regular grid and convolving explicitly.

It is convenient if the polynomial term of $s_s(\boldsymbol{x})$ remains unchanged; i.e. it equates to $p(\boldsymbol{x})$. For a 3D biharmonic RBF, the polynomial term is linear; i.e. $p(\boldsymbol{x}) = \boldsymbol{a} \cdot \boldsymbol{x} + b$, in which case

$$\sum_{j=1}^{M} h_j p(\boldsymbol{x} - \boldsymbol{l}_j) = \sum_{j=1}^{M} h_j \left[\boldsymbol{a} \cdot (\boldsymbol{x} - \boldsymbol{l}_j) + b\right] \tag{4.3}$$

which equals $p(\boldsymbol{x})$ if

$$\sum_j h_j = 1 \quad \text{and} \quad \sum_j h_j \boldsymbol{l}_j = \boldsymbol{0} \tag{4.4}$$

One can also think of this approach as being equivalent to replacing the basic function $\Phi(\boldsymbol{x})$ with a smoother basic function $\Psi(\boldsymbol{x})$, as in Equation (3.1), but where

$$\Psi(\boldsymbol{x}) = \sum_{j=1}^{M} h_j \Phi(\boldsymbol{x} - \boldsymbol{l}_j) \tag{4.5}$$

3x3x3 kernel weights

1	2	1
2	4	2
1	2	1

2	4	2
4	8	4
2	4	2

1	2	1
2	4	2
1	2	1

Normalisation constant 64

Figure 9: Discrete triangle filter weights for $3 \times 3 \times 3$ kernel.

Consider the 27 sample triangular kernel $\overline{h}(\boldsymbol{x})$ whose weights h_j are shown in Figure 9 and whose $\boldsymbol{l}_j$ lie on a 3x3x3 grid centred at $\boldsymbol{0}$ with spacing L. In Figure 10 we have plotted $\Psi(x, 0, 0)$,

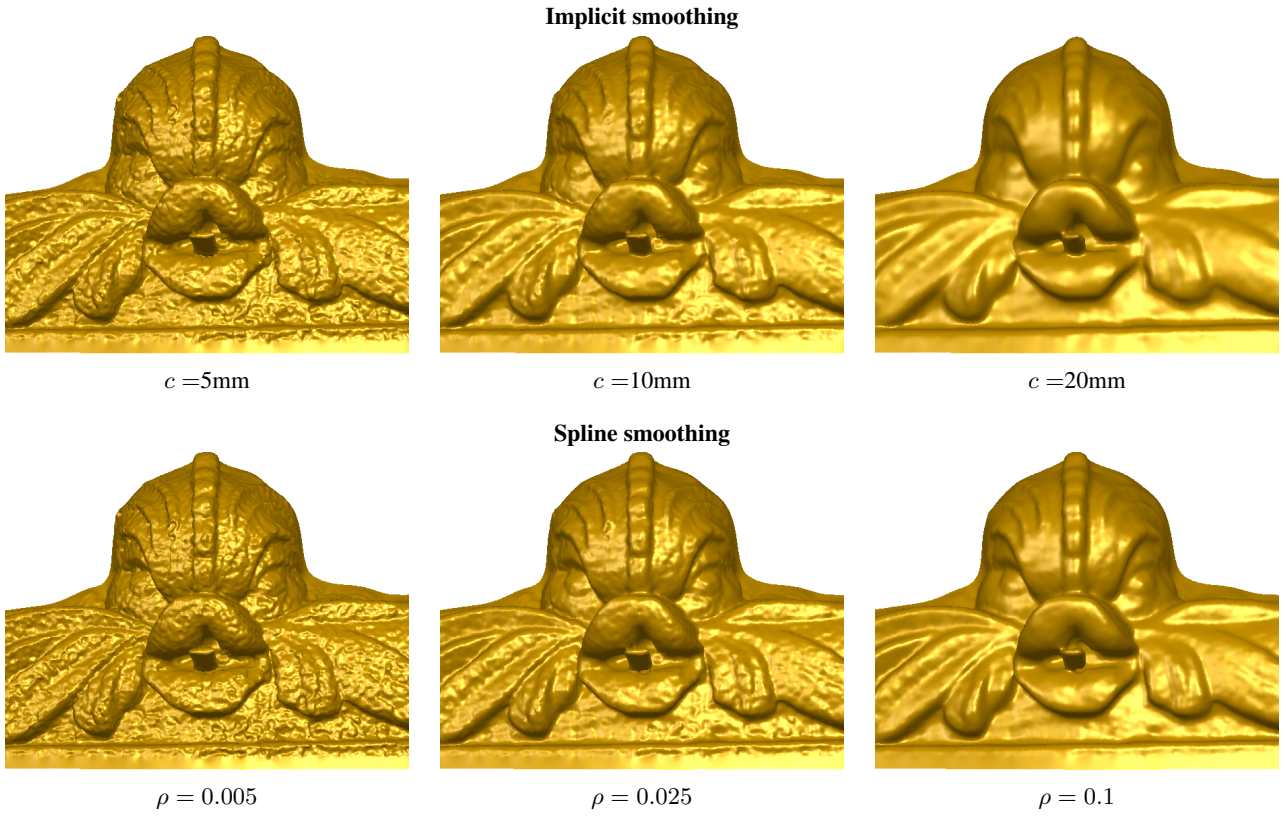

Figure 7: Comparison of implicit smoothing (low pass filtering) and spline smoothing techniques.

as defined in Equation (4.5). For comparison, we have also plotted the multiquadric basic function for $c = 1$. We have chosen L such that Ψ has the same value as the multiquadric at the origin, i.e., $\Psi(0) = c$. Evaluating Equation (4.5) at 0 for this kernel gives $\Psi(0) = \frac{1}{64}(8\sqrt{3} + 24\sqrt{2} + 24)L = 1.12L$. Setting $\Psi(0) = c$ gives $L = 0.9c$. The similarity suggests that the discrete triangle filter is reasonable approximation to the multiquadric $h_{3,1}$ smoothing kernel in Figure 4(b). Indeed, the result obtained for smoothing the fish data set in Figure 6 with $L = 9$mm (Figure 11) is very similar to that using the multiquadric in Figure 7 for $c = 10$mm.

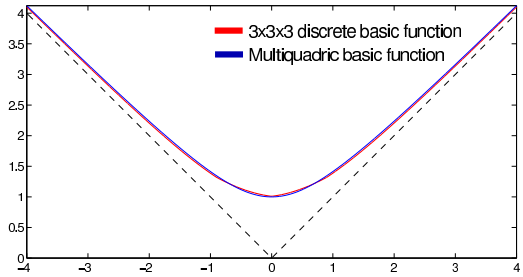

Figure 10: Plot of $\Psi(x, 0, 0)$, a cross-section through the discrete basic function (red) formed by summing 27 biharmonic basic functions (dashed) weighted by the coefficients of the $3 \times 3 \times 3$ kernel in Figure 9. The blue plot is the comparable multiquadric.

If $\overline{h}(\boldsymbol{x})$ consists of delta functions on a regular grid of spacing Δ, then $\overline{h}(\boldsymbol{x})$ can be viewed as a sampled version of a continuous kernel $h(\boldsymbol{x})$. Thus, a result equivalent to convolving an RBF with some desired continuous kernel $h(\boldsymbol{x})$ could be achieved with this method by using the sampled version $\overline{h}(\boldsymbol{x})$. However, in doing so

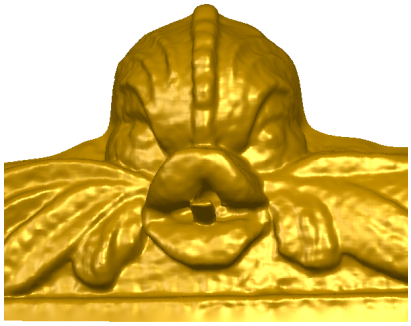

Figure 11: Smoothing using a discrete approximation to the multiquadric smoothing kernel.

the sample spacing Δ must be chosen to ensure that $s \star \overline{h}$ is an adequate approximation to $s \star h$. The effect of sampling $h(\boldsymbol{x})$ is the repetition of its spectrum $H(\boldsymbol{u})$ in the frequency domain. Assuming band-limited signals, Figure 12 graphically depicts the sampling criterion to avoid aliasing (overlap between the RBF spectrum and the repeated $H(\boldsymbol{u})$ spectrum). We generalise this idea by noting that although $S(\boldsymbol{u})$ (the spectrum of the RBF $s(\boldsymbol{x})$) and $H(\boldsymbol{u})$ may have infinite extent, in practice each will have some frequency beyond which the energy is negligible, and these limits are denoted u_S and u_H respectively. Consequently, if Δ is chosen such that $1/\Delta >> u_S + u_H$, then aliasing is reduced and so $s \star \overline{h}$ is a reasonable approximation of $s \star h$.

It is worthwhile noting the extra computational cost of evaluating the RBF $s_s(\boldsymbol{x})$ (Eq. 4.2) over that of evaluating $s(\boldsymbol{x})$. To evaluate an RBF comprising N centres by the fast methods described in [Carr et al. 2001] requires $\mathcal{O}(N \log N)$ flops for setup, then $\mathcal{O}(1)$ flops per evaluation. Therefore the setup phase incurs extra computational cost, but the subsequent evaluation at each point does not.

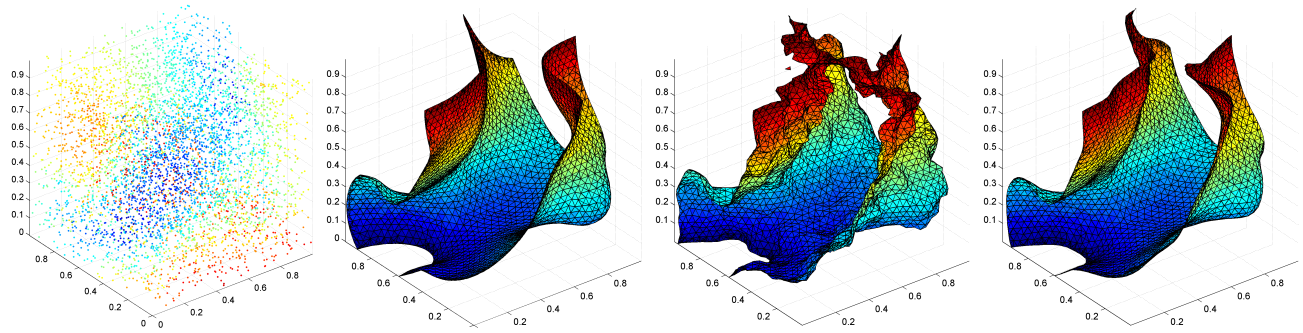

(a) Artificial 'distance' data

(b) Isosurface of RBF fitted to noise-free data

(c) Isosurface of RBF fitted to noisy data

(d) Isosurfacing noisy data with implicit smoothing

Figure 8: Illustration of implicit LPF smoothing with artificial data.

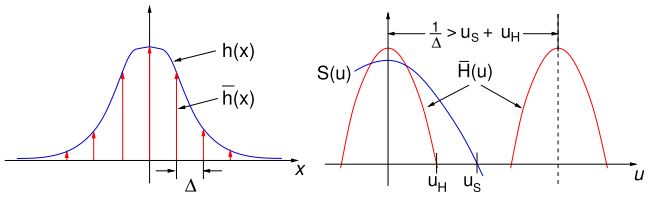

(a) Continuous kernel $h(\boldsymbol{x})$ and sampled kernel $\overline{h}(\boldsymbol{x})$

(b) Spectrum of RBF $S(\boldsymbol{u})$ and the sampled kernel $\overline{H}(\boldsymbol{u})$

Figure 12: Establishing the sampling criterion for a discretely sampled smoothing kernel $h(\boldsymbol{x})$.

Although not quite so convenient as substituting a multiquadric for the basic function during an RBF evaluation (Section 3.1), this discrete approach offers the possibility of specifying an arbitrary kernel. For example, an asymmetric kernel can be used, which may be desirable when reconstructing surfaces from laser range data where range noise and pointing (azimuth and elevation) errors differ. Furthermore, a spatially varying kernel may be implemented by varying the weights and/or the positions of the delta functions throughout space. This could be useful for preserving sharp edges and corners. Figure 13 shows the result of using the kernel in Figure 9, but spatially varying L between 0 and 10mm.

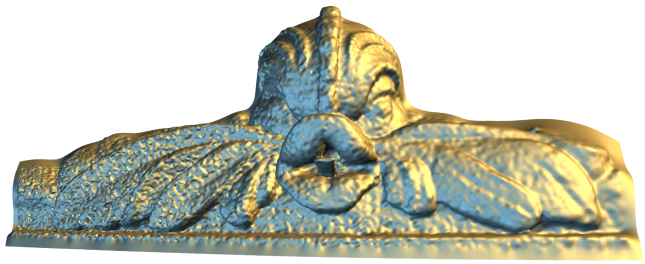

Figure 13: Example of a spatially varying smoothing kernel. The spacing L of the 3x3x3 kernel increases across the figure from 0 at left (no smoothing) to 10mm at right (maximum smoothing).

5 Discussion

Implicit smoothing offers a very convenient way to filter irregularly sampled 3D data where Fourier techniques have otherwise been restricted to regularly sampled (gridded) data. The concept of linear filtering is simpler than spline smoothing and the controlling parameter, the cut-off frequency or filter width, has a clear physical meaning. Table 2 compares timings for smoothed and non-smoothed mesh generation from RBF models fitted to LIDAR scans. The cost of performing a smooth evaluation with basic function substitution is approximately two to three times that of a standard RBF evaluation. However, this cost is relatively invariant to the amount of smoothing, i.e., the value of c (the approximate width of the smoothing kernel).

High frequencies, irrespective of their origin (whether due to noise or genuine structure in the data) are globally attenuated by a low pass filter. Consequently, sharp edges and corners can be expected to be chamfered (smoothed) just as high frequency features due to noise are. However, with our implicit smoothing methods it is possible to locally vary the amount of smoothing when evaluating an RBF and so genuine edges could be preserved by reducing smoothing in their vicinity. [Dinh et al. 2001] demonstrate such a strategy for preserving edges in their work.

5.1 Data reduction

Unlike methods which apply smoothing during fitting, no reduction in the number of RBF centres occurs when using implicit smoothing. The ability to vary smoothing when evaluating an RBF implies that the RBF retains all the detail of the raw data. However, data reduction can be achieved for any particular amount of smoothing by fitting a second RBF to the implicitly-smoothed data using a 'greedy' algorithm as in [Carr et al. 2001]. First an 'exact-fit' RBF is evaluated at the original data points using implicit smoothing. A second RBF is then fitted to these smoothed values using a 'greedy' algorithm resulting in fewer centres. Fewer centres means subsequent evaluations are much faster. Figure 14 illustrates low pass filtering of an RBF fitted to a LIDAR scan of a statue. Fitting a new RBF reduces the 457,586 centres of the original fit to 45,988 and 28,950 respectively for the smoothed data.

5.2 Aliasing

An advantage of a mathematical model, such as an RBF, is that we can evaluate (i.e., sample) it at any resolution. However, to avoid aliasing artifacts, the RBF should first be bandlimited to half the sampling frequency (the Nyquist limit) as with any other signal. Employing implicit smoothing during mesh evaluation, with the smoothing parameter c (the kernel width) set approximately to the mesh resolution, is an elegant solution to this problem.

5.3 Discrete approximation

The interesting possibilities raised by the discrete approximation discussed in Section 4 warrant further exploration. This would ben-

Figure	Number of centres	Number of faces	Time to isosurface		
			No smoothing (s)	10mm smoothing (s)	20mm smoothing (s)
Fish (Fig.6)	89,996	604,382	456	1367	1359
Stone King (Fig.14)	457,586	121,348	193	391	400

Table 2: Timings for smooth surface evaluations compared to timings without smoothing conducted on a 550MHz PIII.

efit from a customised evaluator that can efficiently manage the number of extra centres arising from sampling large kernels. Presumably, the symmetry within a kernel and the local repetition of the added centres can be exploited in this respect.

5.4 Outliers

Real-world data may exhibit occasional outliers or *shot noise*. This can be acknowledged in the RBF fitting process by introducing a *confidence* parameter which allows a small percentage of the strongest (i.e., most difficult to achieve) fitting constraints to be ignored. We have found that enforcing, say 95-99% of the constraints generally results in faster fitting times with little, if any, visible difference to the resulting surface. This approach is relevant to both low pass filtering and spline smoothing. All figures in this paper, however, were generated with 100% of constraints being enforced.

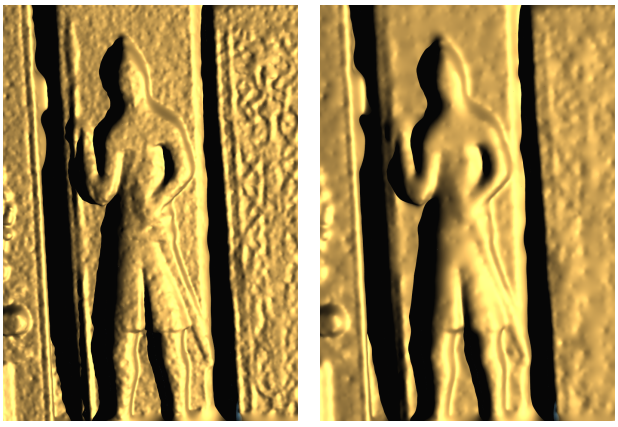

(a) c=10mm, 45,988 centres (b) c=20mm, 28,950 centres

Figure 14: Removing unwanted detail from a surface reconstructed from LIDAR data with low pass filtering and then fitting an RBF with fewer centres to the smoothed data.

6 Conclusions

The LPF approach presented in this paper addresses the difficulties of traditional spline-smoothing when fitting RBFs to noisy data, namely, the computational requirements and the lack of an intuitive control parameter. We have shown that the substitution of a smoother basic function during RBF evaluation is, in special cases, equivalent to convolution with a LPF. This results in an elegant and computationally efficient approach to smoothing noisy surface data due to the applicability of fast numerical methods. The success of LPF depends on the degree to which noise can be discriminated from genuine signal (the actual surface shape) based on frequency content alone. Unlike spline smoothing and approximation techniques which apply when *fitting* an RBF to noisy data, our implicit implementation allows *a posteriori* smoothing of the RBF to suit the desired level of detail required when generating a mesh from the RBF.

Finally, it should be noted that low pass filtering is not just relevant to reconstructing surfaces, but applicable to any scattered 3D data distribution modelled by an RBF. Indirectly, we have shown that Fourier signal processing methods are applicable to irregularly sampled data.

7 Acknowledgements

We are particularly grateful to Allen Instruments and Supplies, USA, for the LIDAR data in Figure 7, and to Cyra Technologies for the LIDAR data in Figure 1 and Figure 14.

References

ABRAMOWITZ, M., AND STEGUN, I. A., Eds. 1965. *Handbook of Mathematical Functions*. Dover, New York.

BEATSON, R. K., AND DYN, N. 1996. Multiquadric B-splines. *J. Approximation Theory 87*, 1–24.

BEATSON, R. K., AND LIGHT, W. A. 1997. Fast evaluation of radial basis functions: Methods for two-dimensional polyharmonic splines. *IMA Journal of Numerical Analysis 17*, 343–372.

BEATSON, R. K., LIGHT, W. A., AND BILLINGS, S. 2000. Fast solution of the radial basis function interpolation equations: Domain decomposition methods. *SIAM J. Sci. Comput. 22*, 1717–1740.

BEATSON, R. K., CHERRIE, J. B., AND RAGOZIN, D. L. 2001. Fast evaluation of radial basis functions: Methods for four-dimensional polyharmonic splines. *SIAM J. Math. Anal. 32*, 6, 1272–1310.

CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of SIGGRAPH 2001*, ACM Press/ACM SIGGRAPH, E. Fiume, Ed., 67–76.

CHENEY, E. W., AND LIGHT, W. A. 1999. *A Course in Approximation Theory*. Brooks Cole, Pacific Grove.

CHERRIE, J. B., BEATSON, R. K., AND NEWSAM, G. N. 2002. Fast evaluation of radial basis functions: Methods for generalized multiquadrics in $\mathbb{R}^n$. *SIAM J. Sci. Comput. 23*, 1549–1571.

DINH, H. Q., SLABAUGH, G., AND TURK, G. 2001. Reconstructing surfaces using anisotropic basis functions. In *International Conference on Computer Vision (ICCV) 2001, Vancouver, Canada,*, 606–613.

MICCHELLI, C. A. 1986. Interpolation of scattered data: Distance matrices and conditionally positive definite functions. *Constr. Approx. 2*, 11–22.

MORSE, B. S., S., T., RHEINGANS, Y. P., CHEN, D. T., AND SUBRAMANIAN, K. R. 2001. Interpolating implicit surfaces from scattered surface data using compactly supported radial basis functions. In *Proceedings of International Conference on Shape Modeling and Applications '01*, IEEE Computer Society Press, 89–98.

PAULY, M., AND GROSS, M. 2001. Spectral-processing of point-sampled geometry. In *Proceedings of ACM SIGGRAPH 2001*, ACM Press/ACM SIGGRAPH, E. Fiume, Ed., 379–386.

SAVCHENKO, V. V., PASKO, A. A., OKUNEV, O. G., AND KUNII, T. L. 1995. Function representation of solids reconstructed from scattered surface points and contours. *Computer Graphics Forum 14*, 4, 181–188.

TAUBIN, G. 1995. A signal processing approach to fair surface design. In *Proceedings of ACM SIGGRAPH 95*, ACM Press/ACM SIGGRAPH, 351–358.

TURK, G., AND O'BRIEN, J. F. 1999. Variational implicit surfaces. Tech. Rep. GIT-GVU-99-15, Georgia Institute of Technology, May.

TURK, G., DINH, H. Q., AND O'BRIEN, J. F. 2001. Implicit surfaces that interpolate. In *Shape Modelling International 2001*, 62–74.

WAHBA, G. 1990. *Spline Models for Observational Data*. No. 59 in CBMS-NSF Regional Conference Series in Applied Math. SIAM.

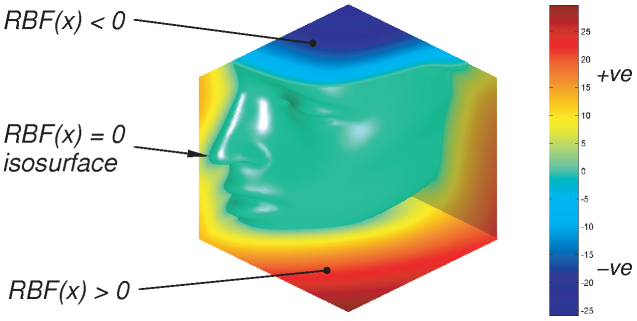

Figure 2

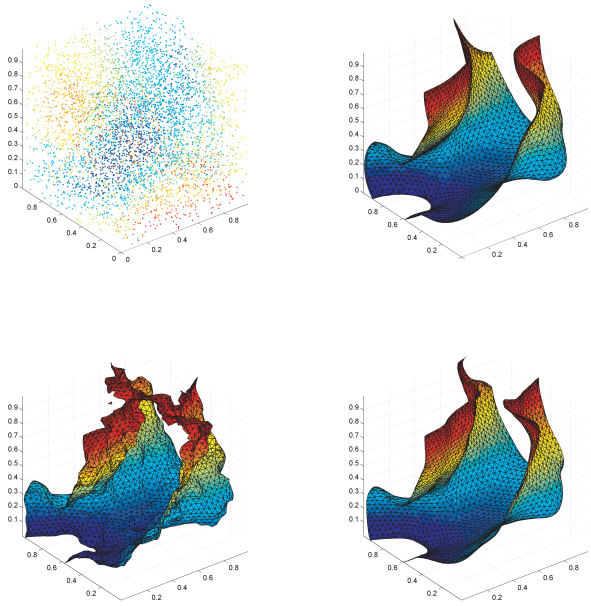

Figure 8

Smooth surface reconstruction from noisy range data
J. C. Carr, R. K. Beatson, B. C. McCallum, W. R. Fright, T. J. McLennan, T. J. Mitchell

Figure 15

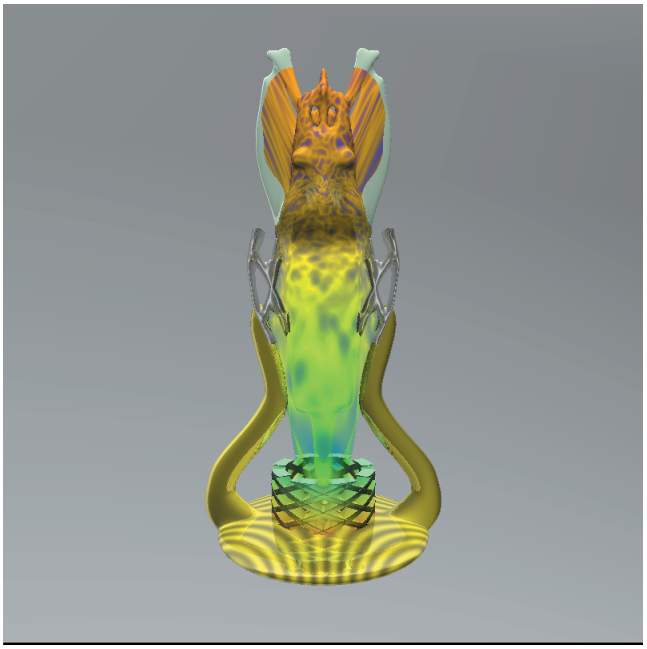

Figure 16

Shape-driven deformations of functionally defined heterogeneous volumetric objects
Benjamin Schmitt, Alexander Pasko, Christophe Schlick