

TVL₁ Shape Approximation from Scattered 3D Data

E. Funk^{1,2}, L. S. Dooley¹ and A. Boerner²

¹ Department of Computing and Communications, The Open University, Milton Keynes, United Kingdom

² Department of Information Processing for Optical Systems, Institute of Optical Sensor Systems, German Aerospace Center, Berlin, Germany
l.s.dooley@open.ac.uk, {eugen.funk, anko.boerner}@dlr.de

Keywords: Shape Reconstruction, Radial Basis Function Interpolation, L_1 Total Variation Minimization, Iterative Large Scale Optimization

Abstract: With the emergence in 3D sensors such as laser scanners and 3D reconstruction from cameras, large 3D point clouds can now be sampled from physical objects within a scene. The raw 3D samples delivered by these sensors however, contain only a limited degree of information about the environment the objects exist in, which means that further geometrical high-level modelling is essential. In addition, issues like sparse data measurements, noise, missing samples due to occlusion, and the inherently huge datasets involved in such representations makes this task extremely challenging. This paper addresses these issues by presenting a new 3D shape modelling framework for samples acquired from 3D sensor. Motivated by the success of nonlinear kernel-based approximation techniques in the statistics domain, existing methods using radial basis functions are applied to 3D object shape approximation. The task is framed as an optimization problem and is extended using non-smooth L_1 total variation regularization. Appropriate convex energy functionals are constructed and solved by applying the Alternating Direction Method of Multipliers approach, which is then extended using Gauss-Seidel iterations. This significantly lowers the computational complexity involved in generating 3D shape from 3D samples, while both numerical and qualitative analysis confirms the superior shape modelling performance of this new framework compared with existing 3D shape reconstruction techniques.

1 INTRODUCTION

The analysis and perception of environments from static or mobile 3D sensors is widely envisioned as a major technological breakthrough and is expected to herald a significant impact upon both society and the economy in the future. As identified by the German Federal Ministry of Education and Research (Hägele, 2011), spatial perception plays a pivotal role in robotics, having an impact in many vital technologies in the fields of navigation, automotive, safety, security and human-robot-interaction. The key task in spatial perception is the reconstruction of the shape of the observed environment. Improvements in shape reconstruction have direct impact on three fundamental research disciplines: *self localization* from camera images (Canelhas et al., 2013), *inspection* in remote sensing (Hirschmüller, 2011) and *object recognition* (Canelhas, 2012). Applying 3D sensors in uncontrolled practical environments, however, leads to strong noise and many data outliers. Homogeneous surface colours and variable illumination conditions lead to outliers and make it a difficult task to obtain

3D samples from over-exposed areas. Figure 1 shows an example 3D point cloud obtained from a stereo camera traversing a building. Many 3D points such as marked by ① suffer from strong noise. Occlusions frequently occur in realistic scenes ② and make automated shape reconstruction even more challenging.

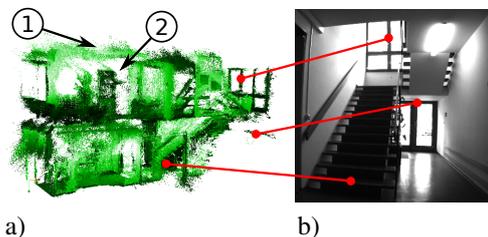


Figure 1: a) The acquired 3D samples shown as a point cloud. b) The scanned staircase as a photograph.

Dealing with noise and outliers inevitably involves applying statistical techniques. In the last decade, so-called kernel-based methods have become well-accepted in statistical processing. Successful

techniques like *deep learning* or *support vector machines* exploit kernel-based methods in the fields of machine learning and robotics for interpolation and extrapolation (Schölkopf and Smola, 2001). Since interpolation and extrapolation are required when dealing with error-prone 3D samples, the application of kernel-based approximation techniques for shape approximation is especially relevant to this domain. The initial aim was the investigation and development of a suitable kernel for geometrical shape modelling from noisy 3D samples.

Many indoor and urban outdoor environments contain a rather small set of large planar shapes. This information can be exploited to help reduce the noise sensitivity leading to higher approximation accuracy. Integrating piecewise smoothness into the approximation task has attracted a lot of interest in the image processing community. A regularization technique, also known as *Total Variation* (TV) minimization, is applied to penalize strong variations in the colour values (Rudin et al., 1992; Getreuer, 2012). Pock et al. (2011) extended the traditional TV approach to second derivatives of the filtered image pixels. Figure 2 shows the comparison of the TVL₁ approach with state of the art statistical filtering techniques. The extension of the TV technique to 3D shapes, is still a fertile area of research which is considered as the second aim in this paper.

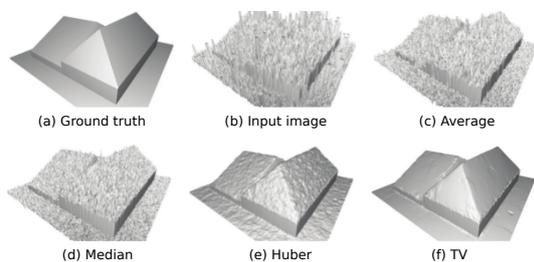


Figure 2: Comparing the total variation minimization with common statistical techniques for height maps filtering. Image courtesy: Bredies et al., (2011).

A further challenge in automated shape approximation is the processing of large datasets. A realistic dataset usually contains several millions of 3D points. However, kernel-based and total variation techniques suffer from high computational complexity prohibiting their application to datasets which contain more than a few thousand points (Bach et al., 2012). Such methods normally require a set of linear equations to be solved, which incurs $O(N^3)$ complexity. Even if such a method could feasibly process $N = 1,000$ points in 10 ms, it would still take 115 days to process 1,000,000 points. This major complexity issue motivated the third aim of this paper which is to develop

efficient strategies for handling non-smooth (L_1) total variation regularization on large datasets.

The remainder of this paper is organised as follows: A short overview of 3D shape reconstruction approaches is provided in Section 2, including issues such as approximation quality and stability. Section 3 discusses the three contributions of this work: i) application of RBFs for implicit 3D shape modelling, ii) integration of non-smooth TVL₁ regularization for noise suppression, and iii) efficient optimization reducing the computation complexity from $O(N^3)$ to $O(N)$. A critical quantitative analysis is then presented in Section 4, and concluding comments are provided in Section 5.

2 LITERATURE REVIEW

The problem of reconstructing a surface of an object from a set of scattered 3D points attracted a lot of attention in the literature (Alexa et al., 2001; Ohtake et al., 2003; Gomes et al., 2009) thus this section will review existing techniques relating to the aims of this paper, namely: shape representation using radial basis functions, statistical regularization model for noise suppression, and efficient optimization.

2.1 Shape Reconstruction

Two general shape representation approaches for 3D data currently exist: explicit and implicit representations.

Explicit models are polygon meshes, *non uniform B-Splines* (NURBS) or Bezier curves (Hughes et al., 2014). The research in computer graphics lead to large number of software frameworks such as OpenGL (Wolff, 2013) enabling the visualization of parametric polygon meshes with the help of parallel graphics hardware. For this reason initial research on automated shape reconstruction from 3D scattered points focused on the direct construction of triangle meshes, also known as Delaunay-Triangulation. Methods such as α shapes (Edelsbrunner and Mücke, 1994; Bernardini et al., 1999; Bodenmüller, 2009) aim at create a polygonal mesh by connecting the input samples by triangle edges. This however, leads to inaccurate results when error-prone samples are provided. Other family of parametric shapes are NURBS (Piegl and Tiller, 1997; Rogers, 2001) and Bezier curves (Agoston, 2005), which are commonly used in 3D modelling. These methods are able to create smooth surfaces for non-uniform control point sets. In order to apply these methods to automated shape reconstruction from scattered 3D points, the surface is

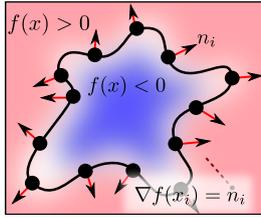


Figure 3: Smooth shape representation from scattered points and surface orientations (arrows) via an implicit function $f(\mathbf{x})$.

defined as a graph in the parameter space. This makes the problem *non polynomial* (NP) hard so its application to larger datasets is prohibited (Zhao et al., 2001).

Implicit models: Several state of the art techniques represent a shape *implicitly* by an indicator function $f(\mathbf{x})$ to indicate *inside* $f(\mathbf{x}) < 0$ or *outside* $f(\mathbf{x}) > 0$ of the object with $\mathbf{x} \in \mathbb{R}^3$ as the location in the 3D space. The surface of the object is the set of all \mathbf{x} , where f gives zero. Figure 3 illustrates an implicit shape where the dots indicate the samples on the surface ($f(\mathbf{x}) = 0$) and the point orientations the normal of the shape ($\nabla f(\mathbf{x}_i) = \mathbf{n}_i$). This representation allows to extract smooth surfaces from irregularly sampled, noisy and incomplete datasets (Gomes et al., 2009).

Facing the noise sensibility issues of Delaunay-Triangulation techniques, Alexa et al. proposed to apply *moving least squares* (MLS) for smoothing (averaging) the point samples prior to reconstructing a mesh via a Delaunay-Triangulation technique (Alexa et al., 2001). A simple implicit shape is for instance a plane, defined by its four parameters $\mathbf{n}^T \mathbf{x} + d = 0$ with $\mathbf{n} \in \mathbb{R}^3$ as the plane normal vector and d as offset to the origin along \mathbf{n} . Defining a shape function as $f(\mathbf{x}) = \mathbf{b}(\mathbf{x})^T \mathbf{u}$ with $\mathbf{b}(\mathbf{x}) = (x_1, x_2, x_3, 1)$ and \mathbf{u} as the plane coefficients $\mathbf{u} = (n_1, n_2, n_3, d)$ enables to find \mathbf{u} via a regression task (Alexa et al., 2003). Similarly, Guennebaud extended the shape model to spheres and proposed the popular *Algebraic Point Set Surfaces* (APSS) method (Guennebaud and Gross, 2007). Ohtake et al. and Oztireli et al. addressed the over smoothing issues by applying non linear regression for shape approximation (Ohtake et al., 2003; Oztireli et al., 2009). The MLS techniques are well capable of filtering data sets with moderate or small noise. However, it is still not feasible for realistic datasets, as introduced in Figure 1.

Implicit models with basis functions: Motivated by the drawbacks of MLS approaches, Calakli and Taubin proposed to apply a global optimization process (Calakli and Taubin, 2011). Acquired 3D samples are structured with an octree and the implicit values of $f(x)$ are distributed on the corners of the octree nodes (voxels). This approach enables large holes to

be closed and allows to handle sparse spatial samples which lead to isolated fragments when MLS is applied. A similar approach is proposed by Kazhdan and Hoppe, where the voxel corners are the B-Splines control points (Kazhdan and Hoppe, 2013). Both approaches suffer from the fundamental drawback that a priori information is required from an expert user to define the depth of the octree structure, which makes using it in automated applications very difficult.

Another family of implicit surface reconstruction algorithms uses smooth *radial basis functions* (RBF). The main difference between RBF based approximation and discrete octree models (Calakli and Taubin, 2011; Kazhdan and Hoppe, 2013) is that RBFs are not necessarily centred on the octree leaves but directly on the samples. This reduces the risk of applying inappropriate discretisation and to lose shape details (Gomes et al., 2009; Carr et al., 2001).

Novel approaches (Zach et al., 2007) propose to create a dense grid of a user specified resolution and to use the L_1 norm to penalize the changes between the implicit grid corner values. Accurate results are achieved when a fine grid is applied, although the approach does not consider the smoothness of the second derivative of the shape leading to non smooth reconstruction. Another drawback of the method is that it is restricted to small and compact objects since the computation time and memory consumption for the dense grid quickly become prohibitive.

Bredies et al. proposed to apply so-called generalized total variation minimization on depth images to penalize the variance in the second derivatives leading to piecewise smooth shapes (Figure 2). The accuracy of the method motivates its extension to 3D shapes, which has not been reported in the literature. Bredies et al. state that the stability of the approach heavily depends on the smoothness of the data, which is feasible when smooth RBFs are applied (Bredies et al., 2010). Thus, when developing an RBF based approximation model with a TV regularization, the choice of an appropriate RBFs type is crucial.

With a popular RBF example being Gaussian, which is of infinite differential degree, but tends to smooth out fine detail, Wahba studied the application of Duchon's *Thin Plate Splines* (Duchon, 1977) that facilitate control of the smoothness degree (Wahba, 1990). Due to its global definition domain, Thin Plate Splines do not result in sparse systems and lead to complex computations. Even more adverse, a change of a single RBF centre affects the complete shape model in the full approximation domain, which is not the case for RBF using compact support such as Gaussians. Later, Wendland proposed several RBF types with compact support of minimal smoothness degree

(Wendland, 1995). Wendland’s RBFs also control the smoothness of the approximated function and still lead to sparse and efficient linear regression systems. Moreover, as presented in Section 3, the smaller the smoothness degree the more stable is the regression process. The Thin Plate RBFs, however, have been shown to achieve superior approximation quality in the presence of noise (Tennakoon et al., 2013). Important aspects when selecting an appropriate RBF type are presented in Section 3.1.

2.2 Efficient L_1 Optimization

Extending the shape approximation with a L_1 penalty requires more advanced techniques to solve the optimization task. This issue has been discussed for some time in the statistics and numerical optimization community. However, efficient techniques being capable of dealing with thousands or millions of data samples are in the current research focus.

Tibshirani proposed the *Least Absolute Shrinkage and Selection* (Lasso) technique to minimize cost functions such as

$$\| \mathbf{y} - K\alpha \|_2^2 + \| \alpha \|_1 \quad (1)$$

with $\| \alpha \|_1 = \sum_j^N |\alpha_j|$ enabling its application on images with several hundreds of thousands entries in α (Tibshirani, 1994). This form is common for regression problems, where the signal \mathbf{y} is approximated linearly by the model matrix K . The additional $\| \cdot \|_1$ penalty term enforces only a small amount of non zeros entries in α . This behaviour is suitable for problems where the vector α is expected to have many zero entries. A common application is for example signal approximation by only a small set of frequencies represented by α .

When representing a shape with N RBFs

$$f(\mathbf{x}) = \sum_i^N \varphi(\mathbf{x}, \mathbf{x}_i) \alpha_i = k^T \alpha \quad (2)$$

with $k_i = \varphi(\mathbf{x}, \mathbf{x}_i)$ its second derivatives are penalized by $\| D\alpha \|_1$ with $D_{j,i} = \partial_{xx}^2 \varphi(\mathbf{x}_j, \mathbf{x}_i)$. This way $\| D\alpha \|_1$ penalizes the amount of non smooth regions in the extracted model. However, since the entries in $D\alpha$ are not separated as it is the case in (1), such problems are more difficult to solve and using the Lasso technique is not possible. Initially, interior active sets methods have been applied to solve the TVL_1 objective (Alizadeh et al., 2003). Chen et al. additionally demonstrated that the efficiency of primal-dual methods is of magnitudes higher than of the interior methods (Chen et al., 2010). Also Goldstein et al. proposed a primal-dual approach known as the *Bregman Split* (Bregman, 1967) to separate the smooth data

term $f_d = \| \mathbf{y} - K\alpha \|_2^2$ from the non smooth regularization term $f_r = \| D\alpha \|_1$ and to optimize each of them independently (Goldstein and Osher, 2009). Boyd et al. extended the Bregman Split approach by Dykstra’s alternating projections technique (Dykstra, 1982) and proposed the *Alternating Direction Method of Multipliers* (ADMM) (Boyd et al., 2011), which further improves the convergence. Discussions related to applications of ADMM are reported by Parikh and Boyd (2014).

The bottleneck of ADMM is the minimization of the smooth part $f_d = \| \mathbf{y} - K\alpha \|_2^2$. Solving this for α with efficient Cholesky factorization suffers from a complexity of $O(N^3)$. However, an iterative linear solver such as Jacobian or Gauss-Seidel may reduce the complexity to $O(N)$ as discussed by Saad or Friedman et al. relating to L_1 regularization (Saad, 2003; Friedman et al., 2010). Nevertheless, further investigations on the applicability of iterative linear solvers and ADMM on 3D shape modelling do not exist.

2.3 Summary

The presented state of the art in robust shape approximation and optimization methods covers several appropriate options for investigation. Table 1 shows the seminal methods summarizing the benefits and drawbacks of each technique. The TVL_1 approach (Bredies et al., 2010) delivers high quality with artefacts such as missing data, noise, outliers, or sharp edges in the image domain. This technique, however, suffers from high computational complexity and needs to be extended to 3D shape approximation. Section 2.2 states that the ADMM technique is expected to outperform the efficiency of existing TVL_1 algorithms when extended with an iterative solver.

The next Section investigates the impact of different RBFs applied for signal and shape approximation from scattered 3D points before the new ADMM technique for TVL_1 optimization on large datasets is presented.

3 THE METHOD

The first part of this section pursues the first research objective and discusses the fundamentals of RBF based approximation and compares different types of RBFs with respect to quality and stability when least squares optimization is performed. Section 3.2 applies the proposed RBFs and defines the convex optimization task to perform shape reconstruction from

Table 1: Shape approximation comparison. Here + indicates that a method is moderately successful in a particular aspect, and ++ indicates that a method is very successful.

Method	Missing Data	Noise	Outliers	Comput. Speed	Sharp Edges
α shapes (Edelsbrunner and Mücke, 1994)				++	+
Adaptive α shapes (Bernardini et al., 1999)	+			++	+
APSS (Guennebaud and Gross, 2007)		+	+	+	+
Smooth Signed Distance (Calakli and Taubin, 2011)	++	+	+	+	
Poisson (Kazhdan and Hoppe, 2013)	++	+	+	+	
TVL ₁ Depth Fusion (Bredies et al., 2010)	++	++	++		++

scattered 3D samples augmented with a TV regularization term. The last part of this section presents the developed optimization technique that allows to reduce the computational complexity while still being able to solve TVL₁ regularized approximation tasks.

3.1 Interpolation with Radial Basis Functions

When approximating any signal from a set of measurements, the general aim is to determine a function $f: \mathbb{R}^d \mapsto \mathbb{R}$ from a set of N sample values at $x_i \in \mathbb{R}^d$. The core idea of RBF based approximation is that the function $f(x)$ may be represented by a linear combination of M weighted functions such as

$$f(x) = \sum_j^M \varphi(x, x_j) \alpha_j. \quad (3)$$

Each of the basis functions $\varphi(x, x_j)$ is centered at each measurement x_j , and basically computes the similarity between x and $x_j \in \mathbb{R}^d$. One possible form for φ is a Gaussian $\varphi(x, x_j) = e^{-\|x-x_j\|/\sigma}$ with σ influencing the width of the support.

The underlying idea of RBF approximation is illustrated for a one-dimensional signal in Figure 4, where f is defined as a sum of all given Gaussians with their weights α_j respectively. Usually it is assumed that the widths σ of the basis functions are known a priori so only the weighting factors α_j are to be found, leading to $f(x)$. The task is therefore to perform regression over N samples and to find M weights via minimization of

$$\min_{\alpha} \sum_i^N (y_i - \sum_j^M \alpha_j \varphi(x_i, x_j))^2$$

where y_i is the i -th measured sample at position x_i . This can also be rewritten in matrix-vector form as:

$$\min_{\alpha} \| \mathbf{y} - \mathbf{K} \alpha \|^2. \quad (4)$$

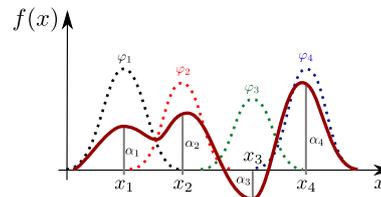


Figure 4: Illustrative example of a function f (red line) constructed by a weighted linear combination of Gaussian radial basis functions.

where K is often referred to as the design matrix or the *kernel* matrix with $K_{i,j} = \varphi(x_i, x_j)$. The solution is obtained via

$$\alpha = A^{-1} K^T \mathbf{y} \quad (5)$$

with $A = (K^T K)$. This is the well known linear *least squares regression*. Note that the function $f(x)$ itself is not restricted to be linear.

In the last two decades several types of RBFs have been proposed for different applications. For the application on shape approximation three RBF types are investigated. i) The Gaussian which is the state of the art, ii) Thin-Plate splines (Duchon, 1977) with global smooth properties and the iii) compactly supported RBFs (CSRBFs) (Wendland, 1995) which enable sparse regression systems to be created and to control the smoothness of the solution. Table 2 shows the three types of the investigated RBFs with the corresponding explicit forms for data dimension $d = 3$. Note that the scaling of each RBF type is achieved by scaling the argument

$$\varphi_s(r) = \varphi\left(\frac{r}{s}\right) \text{ with } r = \|x_i - x_j\|_2. \quad (6)$$

In order to make a systematic decision which RBF type is best suited for the underlying application, the stability and the approximation quality is considered. When solving for α in (5), the condition of K $\text{cond } K = \left| \frac{\lambda_{\max}}{\lambda_{\min}} \right|$ plays an important role. λ_{\min} and

Table 2: Investigated radial basis functions for data dimension $d = 3$.

Type	$\varphi(r)$	Cont. m
Gaussian	e^{-r^2}	C^∞
CSRBF	$(1-r)_+^2$	C^0
	$(1-r)_+^4(4r+1)$	C^2
	$(1-r)_+^6(35r^2+18r+1)$	C^4
Thin-Plate	r^{2m-3}	C^m

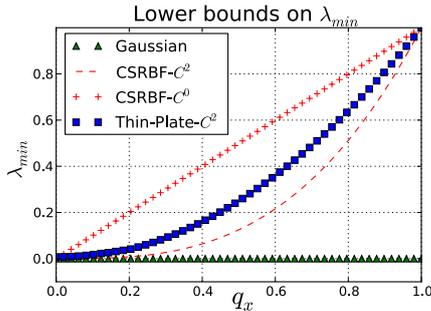


Figure 5: The lower bounds for λ_{min} (higher is better) of different RBFs depending on the normalized sampling radius q_x , $d = 3$ and different continuities C^m .

λ_{max} are the minimal and maximal eigenvalues of K respectively. In practice, it is not feasible to evaluate the condition number on large systems since the computation of the eigenvalues has a complexity of $O(N^3)$. Therefore, a generalized approach to assess the stability a priori is proposed.

Considering the minimal distance between two samples as $q_x := \frac{1}{2} \min_{j \neq i} \|x_i - x_j\|_2$ and interpreting $f(x) = \sum_j^M \varphi(x) \alpha_j$ as a transfer function, it is proposed to analyse the system behaviour in the frequency domain. The key to this is the Fourier-Bessel transform of $\varphi(r)$ (Wendland, 2004). Interpreting the frequency ω as the minimal distance q_x between the approximated samples provides the best-case stability of the regression model without having to perform experiments on data. More practically, the boundaries for the lowest eigenvalue are discussed and put in relation to the expected sample radius q_x . This enables qualitative assessment of a basis function without performing any numerical experiments. Regarding the stability evaluation presented in Figure 5, the Thin-Plate RBF is the only type, which remains stable for all q_x . Depending on the smoothness, CSRBF with C^0 and C^2 follow. The Gaussian RBF is the least stable basis function with λ_{min} slightly below zero.

Another important aspect when selecting an RBF type is the approximation quality. Similarly to the case of stability assessment, numerical experiments often only indicate the behaviour of the RBFs restricted to the given dataset. Thus, it is proposed to

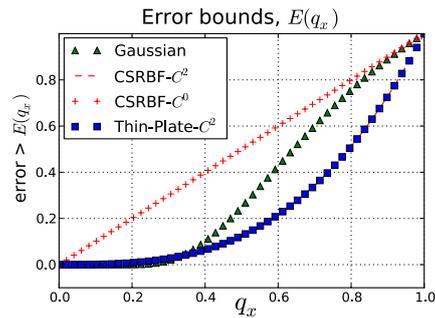


Figure 6: The lower bounds (lower is better) for the approximation error of each RBF type.

appraise the theoretical error bounds in a similar way, as has been shown with the generalized stability. The diagram in Figure 6 presents the best achievable error up to a positive scale factor for each RBF type given the sampling density q_x . It is clear that the higher the sampling density, the better the approximation quality. Notably, the CSRBF- C^2 achieves higher quality than other compact RBFs with lower sampling density q_x and is very similar (overlays) with the global Thin-Plate RBF.

This evaluation indicates the superior performance of the Thin-Plate RBF, though this is not applicable in most realistic applications because the support is not restricted to the neighbouring domain. Furthermore, the presented evaluations claim that applying the compactly supported RBFs with C^2 or C^0 achieves comparable properties. Table 3 shows the summarized investigation results, where a higher number of plus signs reflect better performance. According to the evaluation it is clear that CSRBF is more stable and more accurate than the Gaussian RBFs and provides comparable performance to the Thin-Plate splines without requiring global support. These key observations imply that using CSRBF for 3D data approximation is an attractive option which will now be examined in the next Section.

Table 3: Comparative overview of the RBF models

	Gaussian	Thin Plate	CS-RBF
Stability		+++	+++
Approximation	++	+++	+++
Smoothness	+++	+++	+(+) C^2
Efficiency	+		++

3.2 Shape Reconstruction from Scattered Points

The principal idea of shape modelling with RBF is to extract an implicit function which represents the shape by its zero value as introduced in Figure 3. More formally, an algebraic function $f(\mathbf{x}), f: \mathbb{R}^3 \mapsto \mathbb{R}$ needs to be constructed by regression. Given a set of measured 3D points, the task is further to find a function $f(\mathbf{x})$ which returns zero on every i -th sample \mathbf{x}_i and interpolates well between the samples. Since the zero level alone does not provide information about the orientation of the surface, the surface normals \mathbf{n}_i at every sample are used as constraints for the gradient $\nabla f(\mathbf{x})$ wrt. \mathbf{x} . The task is now to find $f(\mathbf{x}_i) = 0$ giving zero at every sample position *and* $\nabla f(\mathbf{x}_i) = \mathbf{n}_i$. Integrating all this information, a convex cost functional is defined.

$$\min_{\alpha} \sum_i^N \| f(\mathbf{x}_i) \|_2^2 + \| \mathbf{n}_i - \nabla f(\mathbf{x}_i) \|_2^2 \quad (7)$$

To simplify the optimization problem the normalization term $\| \nabla f(\mathbf{x}_i) \|_2 = 1$ is omitted. In order to obtain the gradient ∇f , only the gradient of ϕ needs to be computed, which is precomputed analytically. Putting (7) into matrix notation leads to the short form of the cost function

$$\min_{\alpha} \| K\alpha \|_2^2 + \| \mathbf{n} - K_{\nabla}\alpha \|_2^2. \quad (8)$$

The matrix K contains the values of the RBFs $K_{n,m} = \phi(\mathbf{x}_n, \mathbf{x}_m) \in \mathbb{R}$ and $K_{\nabla n,m} = \nabla\phi(\mathbf{x}_n, \mathbf{x}_m) \in \mathbb{R}^3$ represents the derivatives of ϕ wrt. \mathbf{x}_n . The matrices are of sizes $K \in \mathbb{R}^{N \times M}$ and $K_{\nabla} \in \mathbb{R}^{3N \times M}$. At this point it becomes clear that radial basis functions with local support return for distant points $\mathbf{x}_n, \mathbf{x}_m$ zeros which leads to sparse matrices K and K_{∇} improving the storage and the computation efficiency. Figure 7 shows example matrices K and K_{∇} when a RBF with compact support is applied. Black dots illustrate the non-zero matrix values.

Figure 8 shows an example of applying CSRBF- C^2 on a synthetic point set. The red line in Figure 8a) indicates the cut-plane at which the Figure 8b) has been rendered, while Figure 8c) shows the 3D shape reconstruction.

Next, it is proposed to extend the cost term with an additional total variation regularization term enforcing piecewise smoothness. In computer vision it is accepted practise (Getreuer, 2012) to measure the total variation by computing the Frobenius norm of a Hessian matrix. In contrast, it is proposed to compute the second derivatives with respect to the radius r of the RBF $\phi(r)$. Comparing the single computation $\partial_{rr}^2\phi(r)$ to the evaluation of the 3×3 Hessian

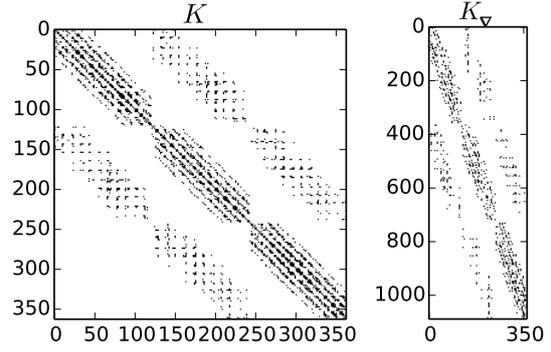


Figure 7: Example sparse matrices K_{∇} and K for (10) when CSRBF is applied.

matrix with nine elements, this reduces the computational cost by a factor of nine and is easier to compute analytically. Similar to the case when computing the gradients of f , the second derivative is also a sum of derived RBFs

$$TV(\mathbf{x}) = \sum_m^M \partial_{rr}^2\phi(r)\alpha_m \quad (9)$$

with $r = \| \mathbf{x}_m - \mathbf{x} \|_2$. Applying the TV regularization, the cost function becomes

$$\min_{\alpha} \| K\alpha \|_2^2 + \| \mathbf{n} - K_{\nabla}\alpha \|_2^2 + \lambda \| D\alpha \|_1 \quad (10)$$

with $D_{n,m} = \partial_{rr}^2\phi(\mathbf{x}_n, \mathbf{x}_m)$ and λ as the weighting of the regularization term. The factors α_m corresponding to the largest eigenvalue of D are attenuated the most while weights lying in the kernel of D , are not affected at all. Figure 9 shows an example when the input samples have been perturbed by noise (Figure 10a) and the shape is reconstructed via a) simple least squares (LSQ) and b) TVL_1 . In both images the red colour corresponds to the TV cost intensity (9). Clearly, when applying TV minimization the shape accuracy of the reconstruction is improved and the red TV intensity is reduced significantly.

Increasing the normally distributed sample noise up to $\sigma \approx 30\%$ of the bounding box, the effect of the regularization is demonstrated in Figure 10. While the simple LSQ model does not achieve a smooth shape

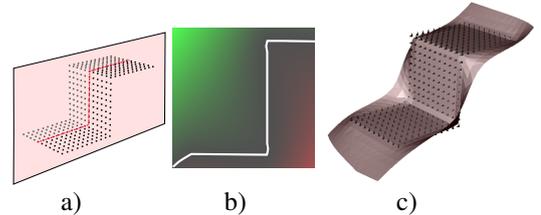


Figure 8: a) Synthetic input points, b) the cut plane visualizing $f(\mathbf{x})$ as red ($f < 0$) green ($f > 0$), c) reconstructed shape from a).

(Figure 10b) the new regularized approach in Figure 10c) shows considerable perceptual improvement in terms of the quality of the shape reconstruction.

In the next section, the proposed numerical technique to efficiently solve the TVL_1 task is presented.

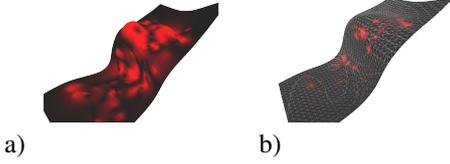


Figure 9: a) The TV cost (red) overlaid with the unregularized shape obtained via LSQ. b) The reduced TV cost (less red colour) after performing regularized approximation following (10).

3.3 TVL_1 Solver

To minimize the task (10) it is proposed to apply the Lagrangian approach from the *Alternating Direction Method of Multipliers* (ADMM) (Boyd et al., 2011). Formally, (10) is restated to

$$\begin{aligned} \min_{\alpha, \mathbf{z}} L(\alpha, \mathbf{z}) &= f_1(\alpha) + f_2(\mathbf{z}) + \\ &\mathbf{b}^T (D\alpha - \mathbf{z}) + \frac{\rho}{2} \|D\alpha - \mathbf{z}\|_2^2 \end{aligned} \quad (11)$$

where $f_1(\alpha)$ is the data part from (10) depending on α , and $f_2(\mathbf{z}) = \lambda \|\mathbf{z}\|_1$ is the non-smooth regularization part weighted by λ . The basic approach is to minimize for α , then for \mathbf{z} iteratively. The terms $\mathbf{b}^T (D\alpha - \mathbf{z})$ and $\frac{\rho}{2} \|D\alpha - \mathbf{z}\|_2^2$ make sure that $D\alpha$ is close to \mathbf{z} after an iteration finishes reducing the duality gap. This restriction is controlled by ρ which is usually a large scalar. The iterative optimization process between α and \mathbf{z} is summarized in Algorithm 1. The minimization for \mathbf{z} involves a sub gradient over $\|\cdot\|_1$ and its solution is known as the shrinkage operator (Efron et al., 2004) being applied on each element z_i independently:

$$\begin{aligned} z_i &= \text{shrink}(a, b) \\ &= a - b \cdot \text{sign}(b - a)_+ \\ \text{where } a &= \frac{b_i}{\rho} + (D\alpha)_i, \quad b = \frac{\lambda}{\rho} \end{aligned} \quad (12)$$

with $(D\alpha)_i$ as the i -th element of the vector $D\alpha$ and $\text{sign}(b - a)_+$ gives 1 if $b > a$ and zero otherwise.

Algorithm 1: ADMM for L_1 approximation

1. Solve for α :
 $(K_{\nabla}^T K_{\nabla} + K^T K + \rho D^T D)\alpha = K_{\nabla}^T \mathbf{n} + D^T (\rho \mathbf{z} - \mathbf{b})$.
2. Evaluate: $z_i^{k+1} = \text{shrink}(\frac{b_i}{\rho} + (D\alpha)_i, \lambda/\rho)$
3. Evaluate: $\mathbf{b}^{k+1} := \mathbf{b}^k + (D\alpha^{k+1} - \mathbf{z}^{k+1})\rho$

While steps 2 and 3 are direct evaluations and can be performed in parallel after $D\alpha^{k+1}$ has been precomputed, step 1 incurs high computational complexity. It is proposed to solve α^{k+1} via Gauss-Seidel iterations well known from large scale linear system optimization (Saad, 2003). However, the standard Gauss-Seidel process suffers from difficult convergence conditions. Thus, *successive over relaxation* (SOR) is applied with a weight factor ω . By applying SOR in step 1 Algorithm 1 is solved via:

$$\alpha_i^{k+1} = \alpha_i^k + \omega \frac{y_i - (K_{\nabla}^T K_{\nabla} + K_i^T K + \rho D_i^T D)\alpha^k}{K_{\nabla}^T K_{\nabla} + K_i^T K + \rho D_i^T D} \quad (13)$$

with $y_i = K_{\nabla}^T \mathbf{n} + D_i^T (\rho \mathbf{z} - \mathbf{b})$ and i th columns of a matrix respectively. Considering that K_{∇} , K and D are sparse when CSRBF is applied, the computation is reduced to Algorithm 2.

Algorithm 2: Matrix free TVL_1

1. For each RBF centre α_i compute:
 - (a) Find all neighbouring centres and all neighbouring samples located in the support of α_i .
 - (b) Compute the equation (13) using only the collected neighbours.
 - (c) Precompute $D\alpha$ with the new α_i^{k+1} .
2. Evaluate: $z_i^{k+1} = \text{shrink}(\frac{b_i}{\rho} + (D\alpha)_i, \lambda/\rho)$
3. Evaluate: $\mathbf{b}^{k+1} := \mathbf{b}^k + (D\alpha^{k+1} - \mathbf{z}^{k+1})\rho$

The optimization is controlled by two important parameters: ω for the successive over relaxation and the RBF scaling s , as introduced in table 2 and (6). Figure 11 shows the effect of these parameters on the approximation quality and the achieved convergence rates. The experiments have been performed on the synthetic dataset from Figure 8. Figure 11a) illustrates the approximation quality over the scaling s , the quality attains its optimum when $s = 10$ is reached. This observation corresponds to the generalized investigations from Figure 6, where $s = 10$ is $q_x = 0.1$. Furthermore, the empirical impact analysis of the over relaxation parameter ω on the convergence concludes, that $\omega \leq 0.15$ enables to remove the instability issues for CSRBF- C^2 and CSRBF- C^4 when SOR is applied. Note that when applying the Gaussian RBF, ω is required to be very small ($\omega \approx 1e-3$), leading to an impractically high number of iterations. This fact is a consequence of the stability properties of the Gaussian, investigated in Section 3.1.

The next section evaluates the proposed TVL_1 shape approximation framework with respect to existing methods by applying the algorithms on a large dataset with an existing ground truth.

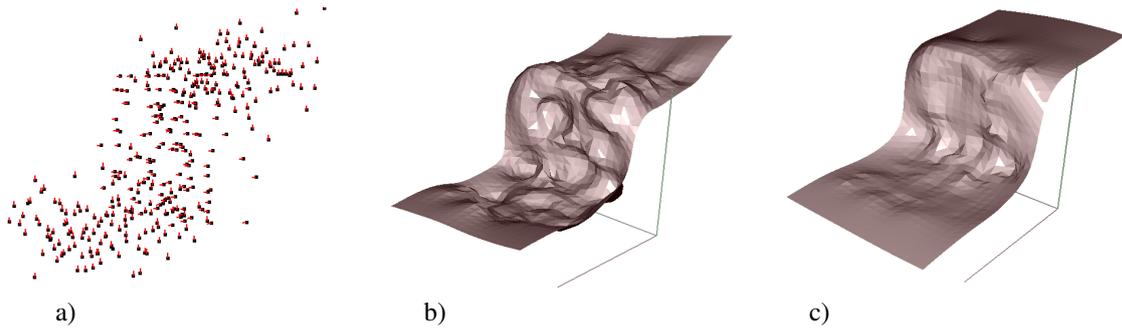


Figure 10: a) Noisy 3D samples of the step function. b) Direct LSQ. c) TVL_1 regularized approximation.

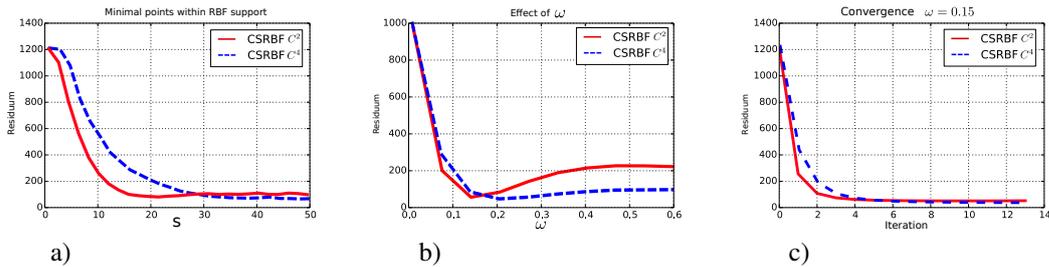


Figure 11: a) The impact of the scaling parameter s , b) the over relaxation weighting ω , c) convergence behavior for $\omega = 0.15$ when CSRBF- C^2 or CSRBF- C^4 are applied.

4 EVALUATION

This Section evaluates the proposed shape reconstruction framework on different datasets and compares it with two successful surface reconstruction techniques: The *Poisson* approximation (Kazhdan and Hoppe, 2013), and the *Smooth-Signed-Distance* (SSD) algorithm (Calakli and Taubin, 2011). The selected methods have been identified as successful techniques for shape reconstruction under strong noise. Both use the implicit model to represent the shapes, however, in contrast to the presented work the compared methods structure the data via an octree of predefined depth and apply discrete optimization via finite differences to extract the zero level of the surface.

The analysis uses a 3D point cloud as the input, and since 3D samples with a ground truth were not available, a virtual camera flight was simulated in order to generate error prone data with an established ground truth. The simulation of the moving camera and noisy 3D measurements were achieved by extending the CAD software Blender (Open Source Community, 2014). This enabled different noise levels on the measurements to be applied and provided an accurate model of the observed environment. Figure 12a) shows the ground truth model used for the simulated measurements in assessing the quality of reconstructed 3D shapes in Figure 12e). An outdoor scene was selected since similar environments are used in

many robotic applications. The facade consisted of large planar areas with a number of sharp edges as identified in point ①. The proposed TVL_1 technique performs significantly better than the Poisson approach and similarly well compared to SSD. During the simulation several areas ② have been occluded by the railing, thus have not been sampled. This increases the difficulty of the reconstruction task. Here, TVL_1 interpolates a shape, which corresponds more to the ground truth than the two compared techniques. The area marked by ③ is the balcony, where only a small part of the floor has been sampled. In such areas, both TVL_1 and SSD perform well, significantly outperforming the Poisson approach. The diagram in Figure 12f) shows the error distribution of the reconstructed shapes. It is produced by re-sampling the ground truth model and the approximated shapes with $5M$ points and by measuring the minimal distance between each reconstructed sample and a ground truth sample. The point-to-point (ptp) distance is shown on the horizontal axis. The histogram height is normalized resulting in a fraction of samples being inside a histogram bin. The further the peak moves to the left of the diagram, the better is the reconstruction accuracy. As also illustrated by the Figures 12 a)-d), TVL_1 slightly outperforms SSD and is significantly more accurate than Poisson.

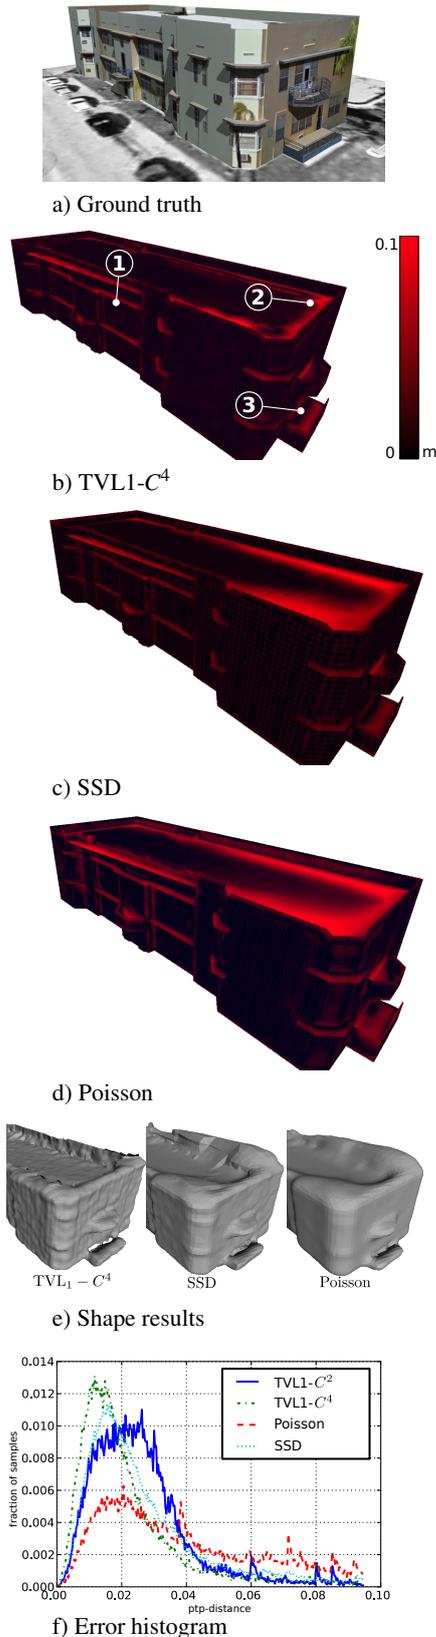


Figure 12: Evaluation results for the proposed TVL₁ and the compared techniques.

5 Conclusion

This paper has presented a new 3D shape modelling strategy for noisy error prone 3D data samples. Modelling 3D shapes with radial basis functions has been proposed with the choice of the most appropriate RBF corroborated using generalized stability and approximation quality assessments. The shape regression model has been extended by non-smooth L_1 regularization assuming planar areas to improve the accuracy of the reconstructed shape in indoor and urban environments. Since the TVL₁ optimization task is computationally expensive, a low complexity optimization technique has been developed. The optimization process exploits the Lagrangian form of the optimization task with an iterative over relaxation technique. This enables realistic datasets containing several millions points to be effectively processed. Quantitative analysis confirms that the proposed method achieves superior accuracy on the synthetic objects.

For future research, the presented solution will be adapted and extended to recursive, real-time 3D mapping applications where environment measurements are received as a data stream. The corresponding 3D shape approximation model will be then able to recursively modify its shape as new measurements become available.

REFERENCES

- Agoston, M. (2005). *Computer Graphics and Geometric Modelling: Implementation & Algorithms*. Computer Graphics and Geometric Modeling. Springer.
- Alexa, M., Behr, J., Cohen-Or, D., Fleishman, S., Levin, D., and Silva, C. T. (2001). Point set surfaces. In *Proceedings of the Conference on Visualization '01, VIS '01*, pages 21–28, Washington, DC, USA. IEEE Computer Society.
- Alexa, M., Behr, J., Cohen-Or, D., Fleishman, S., Levin, D., and Silva, C. T. (2003). Computing and rendering point set surfaces. *Visualization and Computer Graphics, IEEE Transactions on*, 9(1):3–15.
- Alizadeh, F., Alizadeh, F., Goldfarb, D., and Goldfarb, D. (2003). Second-order cone programming. *Mathematical Programming*, 95:3–51.
- Bach, F. R., Jenatton, R., Mairal, J., and Obozinski, G. (2012). Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106.
- Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., and Taubin, G. (1999). The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359.

- Bodenmüller, T. (2009). *Streaming surface reconstruction from real time 3D-measurements*. PhD thesis, Technical University Munich.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122.
- Bredies, K., Kunisch, K., and Pock, T. (2010). Total generalized variation. *SIAM J. Img. Sci.*, 3(3):492–526.
- Bregman, L. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *{USSR} Computational Mathematics and Mathematical Physics*, 7(3):200 – 217.
- Calakli, F. and Taubin, G. (2011). SSD: Smooth signed distance surface reconstruction. *Computer Graphics Forum*, 30(7):1993–2002.
- Canelhas, D. R. (2012). *Scene Representation, Registration and Object Detection in a Truncated Signed Distance Function Representation of 3D Space*. PhD thesis, Örebro University.
- Canelhas, D. R., Stoyanov, T., and Lilienthal, A. J. (2013). Sdf tracker: A parallel algorithm for on-line pose estimation and scene reconstruction from depth images. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 3671–3676. IEEE.
- 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 the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01*, pages 67–76, New York, NY, USA. ACM.
- Chen, X., Lin, Q., Kim, S., Peña, J., Carbonell, J. G., and Xing, E. P. (2010). An efficient proximal-gradient method for single and multi-task regression with structured sparsity. *CoRR*, abs/1005.4717.
- Duchon, J. (1977). Splines minimizing rotation-invariant semi-norms in sobolev spaces. In Schempp, W. and Zeller, K., editors, *Constructive Theory of Functions of Several Variables*, volume 571 of *Lecture Notes in Mathematics*, pages 85–100. Springer Berlin Heidelberg.
- Dykstra, R. (1982). *An Algorithm for Restricted Least Squares Regression*. Technical report, mathematical sciences. University of Missouri-Columbia, Department of Statistics.
- Edelsbrunner, H. and Mücke, E. P. (1994). Three-dimensional alpha shapes. *ACM Trans. Graph.*, 13(1):43–72.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. *Annals of Statistics*, 32:407–499.
- Friedman, J. H., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22.
- Getreuer, P. (2012). Rudin-Osher-Fatemi Total Variation Denoising using Split Bregman. *Image Processing On Line*, 2:74–95.
- Goldstein, T. and Osher, S. (2009). The split bregman method for l1-regularized problems. *SIAM J. Img. Sci.*, 2(2):323–343.
- Gomes, A., Voiculescu, I., Jorge, J., Wyvill, B., and Galbraith, C. (2009). *Implicit Curves and Surfaces: Mathematics, Data Structures and Algorithms*. Springer Publishing Company, Incorporated, 1st edition.
- Guennebaud, G. and Gross, M. (2007). Algebraic point set surfaces. *ACM Trans. Graph.*, 26(3).
- Hägele, M. (2011). Wirtschaftlichkeitsanalysen neuartiger Servicerobotik-Anwendungen und ihre Bedeutung für die Robotik-Entwicklung.
- Hirschmüller, H. (2011). Semi-global matching - motivation, developments and applications. In Fritsch, D., editor, *Photogrammetric Week*, pages 173–184. Wichmann.
- Hughes, J., Foley, J., van Dam, A., and Feiner, S. (2014). *Computer Graphics: Principles and Practice*. The systems programming series. Addison-Wesley.
- Kazhdan, M. and Hoppe, H. (2013). Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3):29:1–29:13.
- Ohtake, Y., Belyaev, A., Alexa, M., Turk, G., and Seidel, H.-P. (2003). Multi-level partition of unity implicit. *ACM Trans. Graph.*, 22(3):463–470.
- Open Source Community (2014). Blender, open source film production software. <http://blender.org>. Accessed: 2014-06-6.
- Oztireli, C., Guennebaud, G., and Gross, M. (2009). Feature Preserving Point Set Surfaces based on Non-Linear Kernel Regression. *Computer Graphics Forum*, 28(2):493–501.
- Piegl, L. and Tiller, W. (1997). *The NURBS Book*. Monographs in Visual Communication. U.S. Government Printing Office.
- Rogers, D. F. (2001). Preface. In Rogers, D. F., editor, *An Introduction to NURBS*, The Morgan Kaufmann Series in Computer Graphics, pages xv – xvii. Morgan Kaufmann, San Francisco.
- Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268.
- Saad, Y. (2003). *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition.
- Schölkopf, B. and Smola, A. J. (2001). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA.
- Tennakoon, R., Bab-Hadiashar, A., Suter, D., and Cao, Z. (2013). Robust data modelling using thin plate splines. In *Digital Image Computing: Techniques and Applications (DICTA), 2013 International Conference on*, pages 1–8.
- Tibshirani, R. (1994). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288.

- Wahba, G. (1990). *Spline models for observational data*, volume 59 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA.
- Wendland, H. (1995). Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4(1):389–396.
- Wendland, H. (2004). *Scattered Data Approximation*. Cambridge University Press.
- Wolff, D. (2013). *OpenGL 4 Shading Language Cookbook, Second Edition*. EBL-Schweitzer. Packt Publishing.
- Zach, C., Pock, T., and Bischof, H. (2007). A globally optimal algorithm for robust tv-l1 range image integration. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8.
- Zhao, H., Oshery, S., and Fedkiwz, R. (2001). Fast surface reconstruction using the level set method. In *In VLSM 01: Proceedings of the IEEE Workshop on Variational and Level Set Methods*.