TVL_1 Planarity Regularization for 3D Shape Approximation

Book Section

TVL_1 Planarity Regularization for 3D Shape Approximation

Eugen Funk[1,2], Laurence S. Dooley[1], and Anko Börner[2]

[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 (DLR), Berlin, Germany
eugen.funk@dlr.de, l.s.dooley@open.ac.uk, anko.boerner@dlr.de

Abstract. The modern emergence of automation in many industries has given impetus to extensive research into mobile robotics. Novel perception technologies now enable cars to drive autonomously, tractors to till a field automatically and underwater robots to construct pipelines. An essential requirement to facilitate both perception and autonomous navigation is the analysis of the 3D environment using sensors like laser scanners or stereo cameras. 3D sensors generate a very large number of 3D data points when sampling object shapes within an environment, but crucially do not provide any intrinsic information about the environment which the robots operate within.
This work focuses on the fundamental task of 3D shape reconstruction and modelling from 3D point clouds. The novelty lies in the representation of surfaces by algebraic functions having limited support, which enables the extraction of smooth consistent implicit shapes from noisy samples with a heterogeneous density. The minimization of total variation of second differential degree makes it possible to enforce planar surfaces which often occur in man-made environments.
Applying the new technique means that less accurate, low-cost 3D sensors can be employed without sacrificing the 3D shape reconstruction accuracy.

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 [25], spatial perception plays a pivotal role in robotics, having an impact onmany 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 [13], *inspection* in remote sensing [26] and *object recognition* [12].

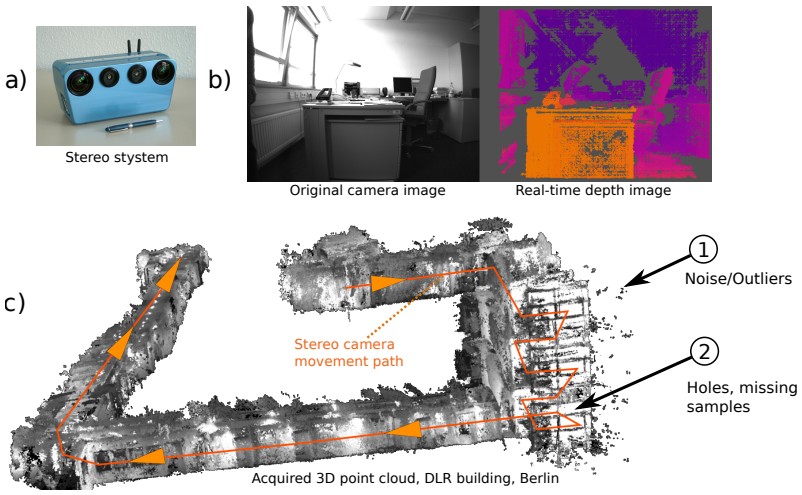

Fig. 1: a) Stereo System, b) sample RGB and depth image, c) acquired 3D point cloud.

Applying 3D sensors in uncontrolled practical environments, however, leads to strong noise and many data outliers. Homogeneous surface colours and dynamic illumination conditions lead to outliers and reduce drastically the quality of computed 3D samples. 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 or over-exposure frequently occur in realistic scenes (2) and make automated shape reconstruction even more challenging.

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 [36]. Since shape interpolation and extrapolation are required when dealing with error-prone 3D samples, the application of kernel-based 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 can be represented by a small set of planar shapes. This information can be exploited to help to achieve higher approximation accuracy. Integrating piecewise smoothness into the approximation task has attracted a lot of interest in the image processing community. Several research groups applied a regularization technique, also known as *Total Variation* (TV) minimization, to penalize strong variations in the colour values [34, 21]. Bredies [9] extended the traditional TV approach to second derivatives of the filtered image pixels. Figure 2 shows the comparison of Bredies's TVL_1 approach with state-of-the-art statistical filtering techniques. The extension of the

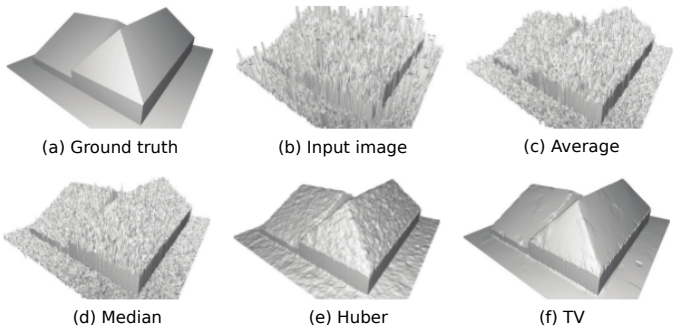

Fig. 2: Comparison of total variation minimization with standard statistical techniques for height maps filtering. Image courtesy: [9].

TV technique to 3D shapes is still a fertile area of research which is considered as the second aim in this work.

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 [5]. Methods that require a set of linear equations to be solved incur $\mathcal{O}(N^3)$ 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 work which is to develop efficient strategies for handling non-smooth (L_1) total variation regularization on large datasets.

The remainder of this paper is organized 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 main contributions of this work: i) application of smooth kernels for implicit 3D shape modelling, ii) integration of non-smooth TVL_1 regularization for noise suppression, and iii) efficient optimization reducing the computation complexity from $\mathcal{O}(N^3)$ to $\mathcal{O}(N)$. A critical quantitative analysis is 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 [2, 29, 23]. This section will review existing techniques relating to the aims of this paper, namely: shape representation using radial basis functions, statistical planarity-aware regularization model, and efficient optimization.

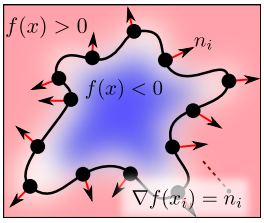

Fig. 3: Smooth shape representation from scattered points and surface orientations (arrows) via an implicit function $f(\mathbf{x})$.

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 [27]. Research in computer graphics leads to a large number of software frameworks such as OpenGL [42] that enables 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 [18, 6, 7] aim at creating a polygonal mesh by connecting the input samples with triangle edges. This, however, leads to inaccurate results when error-prone samples are provided. Another family of parametric shapes are NURBS [32, 33] and Bezier curves [1], 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 defined as a graph in the parameter space. This makes the problem *non-polynomial* (NP) hard so its application to larger datasets is prohibited [44].

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 [23].

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 [2]. 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 $b(\mathbf{x}) = (x_1, x_2, x_3, 1)$ and $\mathbf{u}$ as the plane coefficients $\mathbf{u} = (n_1, n_2, n_3, d)$ allows to find $\mathbf{u}$ via a regression task [3]. Similarly, Guennebaud extended the shape

model to spheres and proposed the popular *Algebraic Point Set Surfaces* (APSS) method [24]. Ohtake et al. and Oztireli et al. addressed the over-smoothing issues by applying non-linear regression for shape approximation [29, 31]. The MLS techniques are well capable of filtering datasets 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 applying a global optimization process [11]. 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 [28]. 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 [11, 28] is that RBFs are not necessarily centred on the octree leaves but directly on the samples. This reduces the risk of applying inappropriate discretization and to lose shape details [23, 14].

Novel approaches [43] propose creating 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 [9]. Thus, when developing an RBF-based approximation model with a TV regularization, the choice of an appropriate RBF 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* [16] that facilitate control of the smoothness degree [39]. Due to their 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 [40]. 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 [37]. Important aspects when selecting an appropriate RBF type are presented in Section 3.1.

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 [18]				++	+
Adaptive α shapes [6]	+			++	+
APSS [24]		+	+	+	+
SSD [11]	++	+	+	+	
Poisson [28]	++	+	+	+	
TVL_1 Depth Fusion [9]	++	++	++		++

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 focussed in current research.

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

$$\| \mathbf{y} - K\boldsymbol{\alpha} \|_2^2 + \| \boldsymbol{\alpha} \|_1 \tag{1}$$

with $\| \boldsymbol{\alpha} \|_1 = \sum_j^N |\alpha_i|$ enabling its application on images with several hundreds of thousands of entries in $\boldsymbol{\alpha}$ [38]. 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 $\boldsymbol{\alpha}$. This behaviour is suitable for problems where the vector $\boldsymbol{\alpha}$ is expected to have many zero entries. A common application is for example signal approximation by only a small set of frequencies represented by $\boldsymbol{\alpha}$.

When representing a shape with N RBFs

$$f(\mathbf{x}) = \sum_{i}^{N} \varphi(\mathbf{x}, \mathbf{x}_i)\alpha_i = k^T \boldsymbol{\alpha} \tag{2}$$

with $k_i = \varphi(\mathbf{x}, \mathbf{x}_i)$, its second derivatives are penalized by $\| D\boldsymbol{\alpha} \|_1$ with $D_{j,i} = \partial^2_{xx}\varphi(\mathbf{x}_j, \mathbf{x}_i)$. This way $\| D\boldsymbol{\alpha} \|_1$ penalizes the amount of non-smooth regions in the extracted model. However, since the entries in $D\boldsymbol{\alpha}$ are not separated as it is the case in (1), such problems are solved with more difficulty and using the Lasso technique is not possible. Initially, interior active sets methods have been applied to solve the TVL_1 objective [4]. Chen et al. additionally demonstrated that the efficiency of primal-dual methods is of magnitudes higher than that of the interior methods [15]. Also Goldstein et al. proposed a primal-dual approach known as the *Bregman Split* [10] to separate the smooth data term $f_d = \| \mathbf{y} - K\boldsymbol{\alpha} \|_2^2$ from the non-smooth regularization term $f_r = \| D\boldsymbol{\alpha} \|_1$ and to optimize each of them independently [22]. Boyd et al. extended the Bregman Split approach by Dykstra's alternating projections technique [17] and proposed the *Alternating Direction Method of Multipliers* (ADMM) [8], 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 $f_d = \| \mathbf{y} - K\boldsymbol{\alpha} \|_2^2$. Solving this for $\boldsymbol{\alpha}$ with efficient Cholesky factorization suffers from a complexity of $\mathcal{O}(N^3)$. However, an iterative linear solver such as Jacobian or Gauss-Seidel may reduce the complexity to $\mathcal{O}(N)$ as discussed by Saad or Friedman et al. relating to L_1 regularization [35, 20]. 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 [9] 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 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 TVL_1 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

$$f(x) = \sum_{j}^{M} \varphi(x, x_j)\alpha_j. \tag{3}$$

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

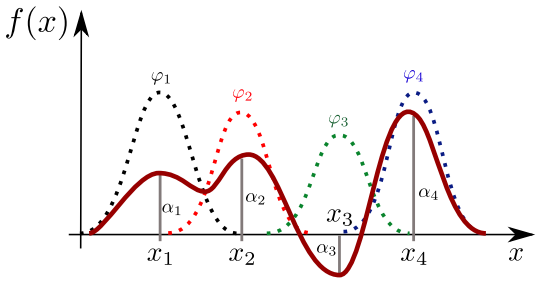

Fig. 4: Illustrative example of smooth $f(x)$ (red line) constructed by a weighted linear combination of Gaussian radial basis functions φ_i.

to find M weights via minimization of

$$\min_{\boldsymbol{\alpha}} \sum_{i}^{N} (y_i - \sum_{j}^{M} \alpha_j \varphi(x_i, x_j))^2,$$

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

$$\min_{\boldsymbol{\alpha}} \parallel \mathbf{y} - K\boldsymbol{\alpha} \parallel_2^2 \tag{4}$$

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

$$\boldsymbol{\alpha} = A^{-1}K^T\mathbf{y} \tag{5}$$

with $A = (K^TK)$. 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 [16] with global smooth properties and the iii) compactly supported RBFs (CSRBFS) [40] that 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

$$\varphi_s(r) = \varphi(\frac{r}{s}) \text{ with } r = \parallel x_i - x_j \parallel_2 . \tag{6}$$

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

Type	$\boldsymbol{\varphi(r)}$	**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

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 $\boldsymbol{\alpha}$ in (5), the condition of K cond $K = |\frac{\lambda_{max}}{\lambda_{min}}|$ plays an important role. λ_{min} and 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 $\mathcal{O}(N^3)$. Therefore, a generalized approach to assess the stability a priori is proposed.

Considering the minimal distance between two samples as $q_x := \frac{1}{2}\min_{j\neq i} \parallel x_i - x_j \parallel_2$ and interpreting $f(x) = \sum_j^M \varphi(x)\alpha_j$ as a transfer function, it is proposed to analyze the system behaviour in the frequency domain. The key to this is the Fourier-Bessel transform of $\varphi(r)$ [41]. Interpreting the frequency ω as

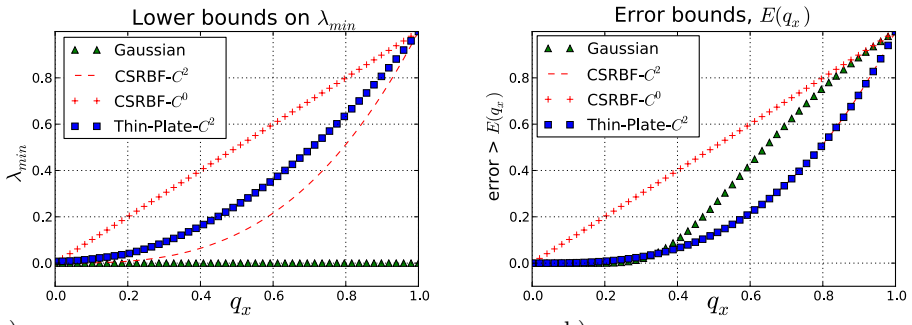

a) b)

Fig. 5: a) Lower bounds for λ_{min} (higher is better), b) lower bounds (lower is better) for the approximation error of each RBF type.

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 5a, the Thin Plate RBF is the only type that 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 appraise the theoretical error bounds in a similar way, as has been shown with the generalized stability. The diagram in Figure 5b 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	+ + +	+ + +	++
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:

$$\min_{\alpha} \sum_{i}^{N} \parallel f(\mathbf{x}_i) \parallel_2^2 + \parallel \mathbf{n}_i - \nabla f(\mathbf{x}_i) \parallel_2^2 . \tag{7}$$

To simplify the optimization problem the normalization term $\parallel \nabla f(\mathbf{x}_i) \parallel_2 = 1$ 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

$$\min_{\boldsymbol{\alpha}} \parallel K\boldsymbol{\alpha} \parallel_2^2 + \parallel \mathbf{n} - K_{\nabla}\boldsymbol{\alpha} \parallel_2^2 . \tag{8}$$

The matrix K contains the values of the RBFs $K_{n,m} = \varphi(\mathbf{x}_n, \mathbf{x}_m) \in \mathbb{R}$ and $K_{\nabla n,m} = \nabla\varphi(\mathbf{x}_n, \mathbf{x}_m) \in \mathbb{R}^3$ 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 zeros for distant points $\mathbf{x}_n, \mathbf{x}_m$ which leads to sparse matrices K and K_{∇} improving the storage and the computation efficiency. Figure 6 shows example matrices K and K_{∇} when a RBF with compact support is applied. Black dots illustrate the non-zero matrix values. Since most of the entries in K and K_{∇} are zero, they can be dismissed in the computation process.

Figure 7 shows an example of applying CSRBF-C^2 on a synthetic point set. The red line in Figure 7a) indicates the cut-plane at which Figure 7b has been rendered, while Figure 7c 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

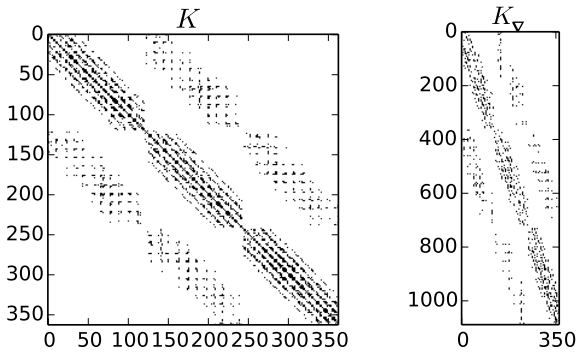

Fig. 6: Example of sparse matrices K_∇ and K for equation (10) when CSRBF is applied.

accepted practise [21] 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 $\varphi(r)$. Comparing the single computation $\partial^2_{rr}\varphi(r)$ to the evaluation of the 3×3 Hessian 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:

$$TV(\mathbf{x}) = \sum_{m}^{M} \partial^2_{rr}\varphi(r)\alpha_m \tag{9}$$

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

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

with $D_{n,m} = \partial^2_{rr}\varphi(\mathbf{x}_n, \mathbf{x}_m)$ 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 8 shows

a) b) c)

Fig. 7: a) Synthetic input points, b) the cut plane visualizing $f(\mathbf{x})$ as red ($f < 0$) and green ($f > 0$), c) reconstructed shape from a).

an example where the input samples have been perturbed by noise (Figure 9a 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 9. While the simple LSQ model does not achieve a smooth shape (Figure 9b the new regularized approach in Figure 9c shows considerable perceptual improvement in terms of the quality of the shape reconstruction.

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

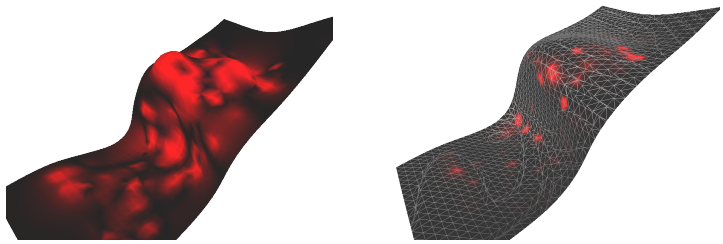

a) b)

Fig. 8: 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).

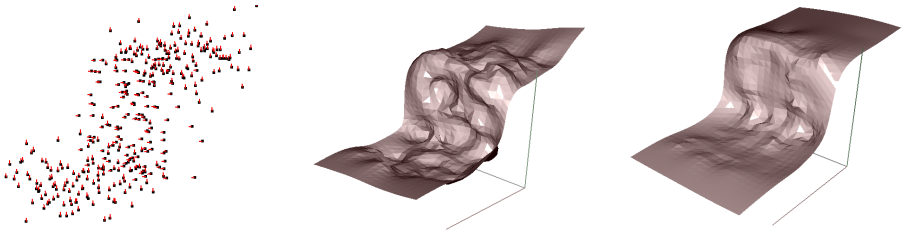

a) b) c)

Fig. 9: a) Noisy 3D samples of the step function. b) Direct LSQ. c) TVL_1 regularized approximation.

3.3 $\mathbf{TVL_1}$ Solver

To minimize the task (10), it is proposed to apply the Lagrangian approach from the *Alternating Direction Method of Multipliers* (ADMM) [8]. Formally, (10) is restated to

$$\begin{aligned} \min_{\alpha,z} L(\boldsymbol{\alpha}, \mathbf{z}) = f_1(\boldsymbol{\alpha}) + f_2(\mathbf{z}) + \\ \mathbf{b}^T(D\boldsymbol{\alpha} - \mathbf{z}) + \frac{\rho}{2} \parallel D\boldsymbol{\alpha} - \mathbf{z} \parallel_2^2, \end{aligned} \tag{11}$$

where $f_1(\boldsymbol{\alpha})$ is the data part from (10) depending on $\boldsymbol{\alpha}$, and $f_2(z) = \lambda \parallel \mathbf{z} \parallel_1$ is the non-smooth regularization part weighted by λ. The basic approach is to minimize for $\boldsymbol{\alpha}$, then for $\mathbf{z}$ iteratively. The terms $\mathbf{b}^T(D\boldsymbol{\alpha} - \mathbf{z})$ and $\parallel D\boldsymbol{\alpha} - \mathbf{z} \parallel_2$ make sure that $D\boldsymbol{\alpha}$ 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 $\boldsymbol{\alpha}$ 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 [19] being applied on each element z_i independently:

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

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

Algorithm 1: ADMM for L_1 approximation

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

While steps 2 and 3 are direct evaluations and can be performed in parallel after $D\boldsymbol{\alpha}^{k+1}$ has been precomputed, step 1 incurs high computational complexity. It is proposed to solve $\boldsymbol{\alpha}^{k+1}$ via Gauss-Seidel iterations, which are well known from large scale linear system optimization [35]. 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 changes to

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

with $y_i = K_{\nabla i}^T \mathbf{n} + D_i^T(\mathbf{z}\rho - \mathbf{b})$ 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.

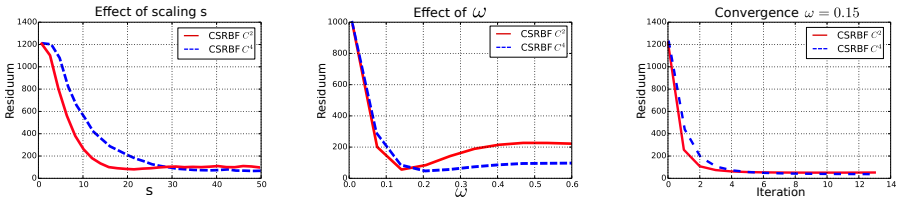

a) b) c)

Fig. 10: a) The impact of the scaling parameter s, b) the over relaxation weighting ω, c) convergence behaviour for $\omega = 0.15$ when CSRBF-C^2 or CSRBF-C^4 are applied.

Algorithm 2: Matrix free TVL$_1$

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

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 10 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 7. Figure 10a 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 5, where $s = 10$ is $q_x = 0.1$. Furthermore, the empirical impact analysis of the over relaxation parameter ω on the convergence concludes that $\omega \leq 0.15$ allows 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.

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 [28] and the *Smooth-Signed-Distance* (SSD) algorithm [11]. The selected methods have been identified as successful techniques

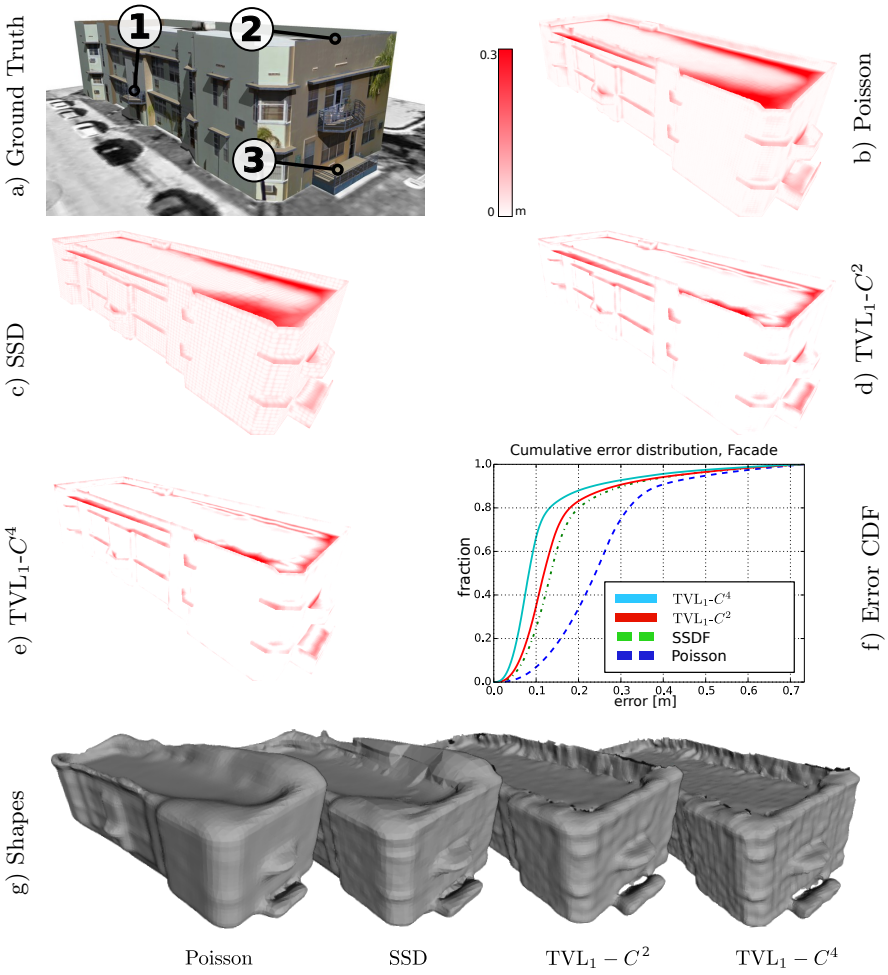

Fig. 11: Evaluation results for the proposed TVL_1 and the compared techniques. See text for details.

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 presented analysis uses a 3D point cloud as input. 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 [30]. This enabled the control of the noise level on the samples and an accurate model of the observed object. An outdoor scene was selected since similar environments are used in many robotic applications. Figure 11a shows the ground truth model used for the simulated measurements in assessing the quality of reconstructed 3D shapes in Figure 11g. Figures 11b - e show the model coloured according to the local reconstruction error. Red areas indicate larger errors. 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. At these locations TVL_1 interpolates a shape which is more similar to the ground truth than other 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 11f shows the cumulative error distribution of the reconstructed shapes. It states for example, that only 77% of all samples have a smaller error than $0.3m$ when Poisson is applied. The diagram is produced by re-sampling the ground truth model and the approximated shapes with $5M$ points and by measuring the distance between a reconstructed sample and its nearest neighbour from the ground truth set. The point-to-point (ptp) error is shown on the horizontal axis. The increased accuracy of TVL_1 techniques can also be observed in the coloured error models in Figures 11b-e.

Another evaluation scenario has been considered where real point clouds from a mobile stereo system have been processed to shapes. As previously illustrated in Figure 1, the small stereo system has been carried inside of a building at high speed, computing the 3D point clouds in real time. This data has been processed by the Poisson, SSD and TVL_1-C^2 techniques, and is shown in Figure 12. Figure 12a and below show larger overview, which has been selected because of the difficult conditions. Floor reflections, occlusions and blending from ceiling illumination lead to error prone data. Poisson and SSD are designed to reconstruct closed surfaces and thus generate wrong surfaces even in the open entries. TVL_1 however, extrapolates the measurements to some extent but does not close the entries correctly indicating open and traversable space.

Figures 12b show a part from the DLR building. It can be observed, that Poisson provides over-smoothed surfaces, SSD tends to interpolate noisy samples and TVL_1 manages smooth surface but also representing the edges between the walls and the floor.

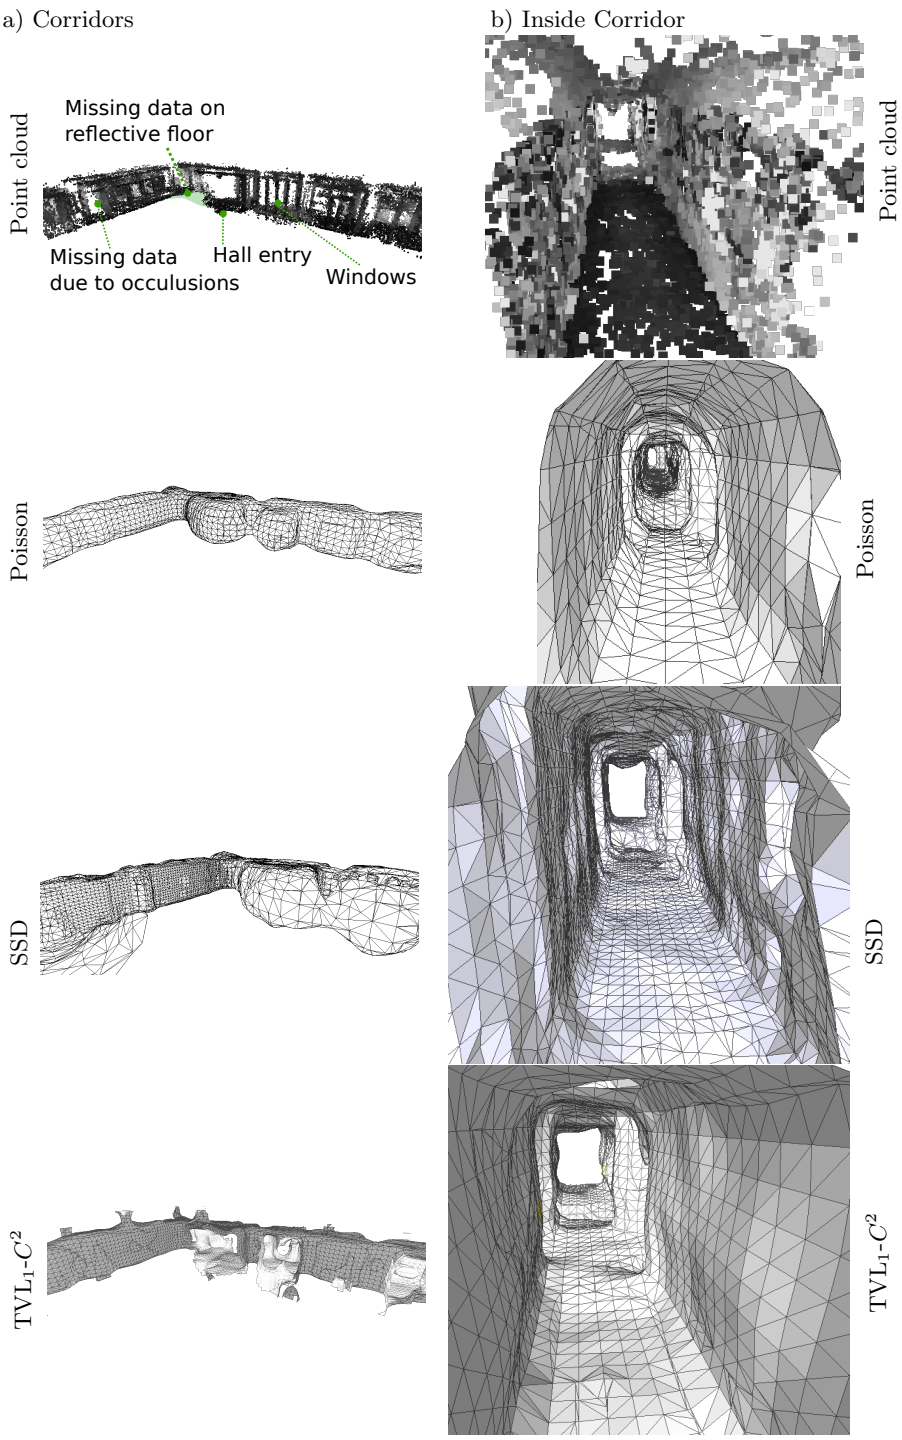

Fig. 12: Shape reconstruction from real time point clouds. See text for details.

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 TVL_1 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 million 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 then will be able to recursively modify its shape as new measurements become available.

References

1. M. Agoston. *Computer Graphics and Geometric Modelling: Implementation & Algorithms*. Computer Graphics and Geometric Modeling. Springer, 2005.
2. M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. T. Silva. Point set surfaces. In *Proceedings of the Conference on Visualization '01*, VIS '01, pages 21–28, Washington, DC, USA, 2001. IEEE Computer Society.
3. M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. T. Silva. Computing and rendering point set surfaces. *Visualization and Computer Graphics, IEEE Transactions on*, 9(1):3–15, Jan 2003.
4. F. Alizadeh, F. Alizadeh, D. Goldfarb, and D. Goldfarb. Second-order cone programming. *Mathematical Programming*, 95:3–51, 2003.
5. F. R. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
6. F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359, Oct. 1999.
7. T. Bodenmüller. *Streaming surface reconstruction from real time 3D-measurements*. PhD thesis, Technical University Munich, 2009.
8. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, Jan. 2011.
9. K. Bredies, K. Kunisch, and T. Pock. Total generalized variation. *SIAM J. Img. Sci.*, 3(3):492–526, Sept. 2010.
10. L. Bregman. 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, 1967.

11. F. Calakli and G. Taubin. SSD: Smooth signed distance surface reconstruction. *Computer Graphics Forum*, 30(7):1993–2002, 2011.
12. D. R. Canelhas. *Scene Representation, Registration and Object Detection in a Truncated Signed Distance Function Representation of 3D Space*. PhD thesis, Örebro University, 2012.
13. D. R. Canelhas, T. Stoyanov, and A. J. Lilienthal. 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, 2013.
14. J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, pages 67–76, New York, NY, USA, 2001. ACM.
15. X. Chen, Q. Lin, S. Kim, J. Peña, J. G. Carbonell, and E. P. Xing. An efficient proximal-gradient method for single and multi-task regression with structured sparsity. *CoRR*, abs/1005.4717, 2010.
16. J. Duchon. Splines minimizing rotation-invariant semi-norms in sobolev spaces. In W. Schempp and K. Zeller, editors, *Constructive Theory of Functions of Several Variables*, volume 571 of *Lecture Notes in Mathematics*, pages 85–100. Springer Berlin Heidelberg, 1977.
17. R. Dykstra. *An Algorithm for Restricted Least Squares Regression*. Technical report, mathematical sciences. University of Missouri-Columbia, Department of Statistics, 1982.
18. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graph.*, 13(1):43–72, Jan. 1994.
19. B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32:407–499, 2004.
20. J. H. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2 2010.
21. P. Getreuer. Rudin-Osher-Fatemi Total Variation Denoising using Split Bregman. *Image Processing On Line*, 2:74–95, 2012.
22. T. Goldstein and S. Osher. The split bregman method for l1-regularized problems. *SIAM J. Img. Sci.*, 2(2):323–343, Apr. 2009.
23. A. Gomes, I. Voiculescu, J. Jorge, B. Wyvill, and C. Galbraith. *Implicit Curves and Surfaces: Mathematics, Data Structures and Algorithms*. Springer Publishing Company, Incorporated, 1st edition, 2009.
24. G. Guennebaud and M. Gross. Algebraic point set surfaces. *ACM Trans. Graph.*, 26(3), July 2007.
25. M. Hägele. Wirtschaftlichkeitsanalysen neuartiger Servicerobotik-Anwendungen und ihre Bedeutung für die Robotik-Entwicklung. 2011.
26. H. Hirschmüller. Semi-global matching - motivation, developments and applications. In D. Fritsch, editor, *Photogrammetric Week*, pages 173–184. Wichmann, September 2011.
27. J. Hughes, J. Foley, A. van Dam, and S. Feiner. *Computer Graphics: Principles and Practice*. The systems programming series. Addison-Wesley, 2014.
28. M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3):29:1–29:13, July 2013.
29. Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, and H.-P. Seidel. Multi-level partition of unity implicits. *ACM Trans. Graph.*, 22(3):463–470, 2003.

30. Open Source Community. Blender, open source film production software. `http://blender.org`, 2014. Accessed: 2014-06-6.
31. C. Oztireli, G. Guennebaud, and M. Gross. Feature Preserving Point Set Surfaces based on Non-Linear Kernel Regression. *Computer Graphics Forum*, 28(2):493–501, 2009.
32. L. Piegl and W. Tiller. *The NURBS Book*. Monographs in Visual Communication. U.S. Government Printing Office, 1997.
33. D. F. Rogers. Preface. In D. F. Rogers, editor, *An Introduction to NURBS*, The Morgan Kaufmann Series in Computer Graphics, pages xv – xvii. Morgan Kaufmann, San Francisco, 2001.
34. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268, Nov. 1992.
35. Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition, 2003.
36. B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
37. R. Tennakoon, A. Bab-Hadiashar, D. Suter, and Z. Cao. Robust data modelling using thin plate splines. In *Digital Image Computing: Techniques and Applications (DICTA), 2013 International Conference on*, pages 1–8, Nov 2013.
38. R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1994.
39. G. Wahba. *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, 1990.
40. H. Wendland. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4(1):389–396, 1995.
41. H. Wendland. *Scattered Data Approximation*. Cambridge University Press, 2004.
42. D. Wolff. *OpenGL 4 Shading Language Cookbook, Second Edition*. EBL-Schweitzer. Packt Publishing, 2013.
43. C. Zach, T. Pock, and H. Bischof. 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, Oct 2007.
44. H. Zhao, S. Oshery, and R. Fedkiwz. Fast surface reconstruction using the level set method. In *In VLSM 01: Proceedings of the IEEE Workshop on Variational and Level Set Methods*, 2001.