L-Tangent Norm: A Low Computational Cost Criterion for Choosing Regularization Weights and its Use for Range Surface Reconstruction

F. Brunet[1,2,3] A. Bartoli[1] R. Malgouyres[3] N. Navab[2]
[1]LASMEA (CNRS/UBP) [2]CAMPAR (TU München) [3]LAIC (UdA)
Clermont-Ferrand (France) Munich (Germany) Clermont-Ferrand (France)

Abstract

We are interested in fitting a surface model such as a tensor-product spline to range image data. This is commonly done by finding control points which minimize a compound cost including the goodness of fit and a regularizer, balanced by a regularization parameter. Many approaches choose this parameter as the minimizer of, for example, the cross-validation score or the L-curve criterion. Most of these criteria are expensive to compute and difficult to minimize.

We propose a novel criterion, the L-tangent norm, which overcomes these drawbacks. Even though it is empirical, it gives sensible results with a much lower computational cost. This new criterion has been successfully tested with synthetic and real range image data, and shows a behavior similar to cross-validation.

1. Introduction

This paper is concerned with the reconstruction of a surface from range image data (also known as 2.5D data). Such data are obtained by range sensors such as Time-of-Flight cameras or stereo imaging. Reconstructing a surface from scattered data points is important for computing geodesics, texture-mapping, global shape editing, etc.

A surface reconstructed from a range image is usually described as a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ called a surface model. A set of parameters controls the surface shape. Many models have been proposed including the Thin-Plate Spline [2, 6, 14, 15], Radial Basis Functions [8, 12, 13], Bézier surfaces [7] and tensor-product splines over the B-spline basis [3, 4].

The classical approach to reconstruct a surface is by finding the set of parameters which minimizes a cost function. This cost function has two terms: a term reflecting the goodness of fit and a regularization term. These two terms are related by a parameter, the so-called *regularization parameter*, which controls the importance given to the regularization. A small value for the regularization parameter results in a surface passing closely to the data points but prone to overfitting. On the contrary, a large regularization parameter results in a smooth surface which may not approximate the data very well.

One of the challenges in surface reconstruction (and in many other data fitting problems) is selecting the regularization parameter automatically. Many approaches have been proposed. Two of the mostly used methods are the cross-validation score minimization [16] and the L-curve criterion maximization [9, 11]. Cross-validation intends to maximize the ability of the reconstructed surface to generalize. The L-curve criterion selects the regularization parameter by connecting the residual (*i.e.* the closeness of the surface to the data points) and the solution norm (*i.e.* the smoothness of the surface). Unfortunately, these criteria are generally expensive to compute. Moreover, they require to solve optimization problems that are generally difficult. We propose a novel heuristic to select the regularization parameter: the *L-tangent norm*. This new approach is interesting for two main reasons. First, it is cheap to compute. Second, its 'shape' and numerical behavior make the selection of the regularization parameter easy.

Paper organization. Section 2 is dedicated to the problem of reconstructing the surface given a regularization parameter. Section 3 is a short review of existing methods to automatically select the regularization parameter. Then, our new approach, the L-tangent norm, is presented in section 4. Finally, experimental results are shown in section 5.

Notations. Scalars are in italics, *e.g.* x, vectors in bold right fonts, *e.g.* $\mathbf{p}$, and matrices in capitals, *e.g.* M. The vector and matrix transpose is denoted with the symbol $^\mathsf{T}$, *e.g.* M^T. Intervals are denoted with square brackets, *e.g.* $[a, b]$, $]a, b[$ and $]a, b]$ for respectively a closed, an open and a half-open interval.

2. Background

2.1. Reconstruction Given the Regularization Parameter

Assume one is given a set of n range data points. Such a set is composed of n twodimensional points $(x_i, y_i) \in \mathbb{R}^2$, associated to depth information $z_i \in \mathbb{R}$. This set of points is denoted:

$$\{(x_i, y_i) \leftrightarrow z_i \mid i = 1, \ldots, n\}. \tag{1}$$

The surface model is a function of h unknown parameters $\mathbf{p} = \begin{bmatrix} p_1, \ldots, p_h \end{bmatrix}^\mathsf{T} \in \mathbb{R}^h$ that control the shape of the surface:

$$\begin{array}{rccc} f(\cdot;\mathbf{p}): & \Omega \subset \mathbb{R}^2 & \longrightarrow & \mathbb{R} \\ & (x,y) & \longmapsto & f(x,y;\mathbf{p}). \end{array} \tag{2}$$

The surface model we use in this paper is given in section 2.2. The range surface reconstruction problem consists in finding the best set of parameters $\mathbf{p}^*_\lambda \in \mathbb{R}^h$ (which depends on λ) such that:

$$\mathbf{p}^*_\lambda = \arg\min_{\mathbf{p}\in\mathbb{R}^h} \mathcal{E}_d(\mathbf{p}) + \frac{\lambda}{1-\lambda}\mathcal{E}_r(\mathbf{p}) \qquad \lambda \in]0,1[, \tag{3}$$

where $\mathcal{E}_d$ and $\mathcal{E}_r$ are respectively called the *data term* and the *regularization term*. The data term is a function that measures the closeness of the surface to the whole set of data points. The regularization term is a measure of the surface regularity (or smoothness). These two terms are related by the *regularization parameter* λ which controls the trade-off between the goodness of fit and the regularity. In the limit $\lambda \to 0$, the surface is likely to overfit the data. If λ is large, the surface becomes very smooth but may not reflect the data very well ; for instance, with many models, the reconstructed surface is almost a plane when λ is close to 1.

The data term and the regularization term can be chosen as respectively the Mean Squared Residual (MSR) and the bending energy:

$$\mathcal{E}_d(\mathbf{p}) = \frac{1}{n}\sum_{i=1}^{n} \left(f(x_i, y_i; \mathbf{p}) - z_i\right)^2 \tag{4}$$

$$\mathcal{E}_r(\mathbf{p}) = \iint_\Omega \sum_{d=0}^{2} \binom{2}{d} \left(\frac{\partial^2 f(\cdot;\mathbf{p})}{\partial x^{2-d}\partial y^d}(x,y)\right)^2 \mathrm{d}x\mathrm{d}y. \tag{5}$$

We consider that the surface model is linear with respect to its parameter:

$$\forall (x,y) \in \mathbb{R}^2, \exists\, \mathbf{v}_{x,y} \in \mathbb{R}^h : f(x,y;\mathbf{p}) = \mathbf{v}^\mathsf{T}_{x,y}\mathbf{p},$$

The MSR can be written as:

$$\mathcal{E}_d(\mathbf{p}) = \|M\mathbf{p} - \mathbf{z}\|_2^2 \tag{6}$$

where M is the *collocation matrix* and $\mathbf{z}$ is the vector containing all the depths:

$$M = \begin{bmatrix} \mathbf{v}^\mathsf{T}_{x_1,y_1} & \ldots & \mathbf{v}^\mathsf{T}_{x_n,y_n} \end{bmatrix}^\mathsf{T} \in \mathbb{R}^{n\times h},$$
$$\mathbf{z} = \begin{bmatrix} z_1 & \ldots & z_n \end{bmatrix}^\mathsf{T} \in \mathbb{R}^n.$$

The bending energy can be approximated by discretizing the integral sum of equation (5) over a regular grid:

$$\mathcal{E}_r(\mathbf{p}) \approx \frac{1}{ab}\sum_{i=0}^{a-1}\sum_{j=0}^{b-1}\sum_{d=0}^{2} \binom{2}{d} \left(\frac{\partial^2 f\left(\frac{i}{a}, \frac{j}{b}; \mathbf{p}\right)}{\partial x^{2-d}\partial y^d}\right)^2 \tag{7}$$

$$\iff \quad \mathcal{E}_r(\mathbf{p}) \approx \|R\mathbf{p}\|_2^2 \tag{8}$$

where R is the *regularization matrix*. Note that the partial derivatives of a linear model are also linear with respect to $\mathbf{p}$, *i.e.*, for all $d \in \{0,1,2\}$ and for all $(x,y) \in \mathbb{R}^2$, there exists $\mathbf{w}_{x,y,d} \in \mathbb{R}^h$ such that:

$$\frac{\partial^2 f(x,y;\mathbf{p})}{\partial x^{2-d}\partial y^d} = \mathbf{w}^\mathsf{T}_{x,y,d}\mathbf{p}. \tag{9}$$

If $\mathbf{r}^\mathsf{T}_k$ is the kth row of R, then we have that:

$$\mathbf{r}^\mathsf{T}_{dab+ib+j+1} = \binom{2}{d}\mathbf{w}^\mathsf{T}_{\frac{i}{a},\frac{j}{b},d}. \tag{10}$$

Finally, equation (3) is equivalent to a Linear Least Squares (LLS) minimization problem [1]:

$$\mathbf{p}^*_\lambda = \arg\min_{\mathbf{p}\in\mathbb{R}^h} \left\| \begin{bmatrix} M \\ \frac{\lambda}{1-\lambda}R \end{bmatrix} \mathbf{p} - \begin{bmatrix} \mathbf{z} \\ 0 \end{bmatrix} \right\|_2^2. \tag{11}$$

The solution of this problem is given by:

$$\mathbf{p}^*_\lambda = \left(M^\mathsf{T}M + \left(\tfrac{\lambda}{1-\lambda}\right)^2 R^\mathsf{T}R\right)^{-1} M^\mathsf{T}\mathbf{z}. \tag{12}$$

2.2. The Surface Model

We have chosen to use the tensor-product splines over the B-spline basis (TPBS) model. A reason is that the influence of a control point is bounded to its neighborhood due to the local support of the B-splines basis [3, 4, 7]. This property leads to sparse collocation and regularization matrices. This makes the computations fast.

We remind the reader some basic facts about B-splines. An extensive review of splines can be found in [3, 4, 7].

The B-spline of degree $k > 0$ (order $k+1$) having the increasing knot sequence $\mu_0 \leq \ldots \leq \mu_{g+1}$ can be defined recursively by:

$$\begin{cases} N_{i,1}(x) = 1 & \text{if } x \in [\mu_i, \mu_{i+1}[\\ N_{i,1}(x) = 0 & \text{if } x \notin [\mu_i, \mu_{i+1}[\\ N_{i,k+1}(x) = \frac{x-\mu_i}{\mu_{i+k}-\mu_i}N_{i,k}(x) \\ \qquad\qquad + \frac{\mu_{i+k+1}-x}{\mu_{i+k+1}-\mu_{i+1}}N_{i+1,k}(x). \end{cases} \tag{13}$$

A TPBS is defined as a linear combination of the B-spline basis functions weighted by the *control points* $C_{i,j}$:

$$s(x,y) = \sum_{i=-k_1}^{g_1} \sum_{j=-k_2}^{g_2} C_{i,j} N_{i,k_1+1}(x) N_{j,k_2+1}(y). \quad (14)$$

Note that equation (14) is linear in the $C_{i,j}$. Note also that it is well-known [4] that the partial derivatives of a TPBS are also TPBS.

In the sequel, we always consider uniform knot sequences. Besides, we take as many knots as possible (regarding the computational complexity) so that the flexibility of the surface model is sufficient to approximate complex shapes.

3. Previous Work

3.1. Cross-Validation

The goal of Ordinary Cross-Validation (OCV) [15, 16] is to choose the regularization parameter so that the reconstructed surface generalizes well. In other words, the reconstructed surface must have a good behavior between the data points. The optimal regularization parameter is the minimizer of the so-called *OCV score*:

$$\lambda^* = \arg \min_{\lambda \in]0,1[} OCV(\lambda). \quad (15)$$

An example of the OCV criterion is given in figure 1a. The OCV score is defined by fitting the model without the ith data point, giving the parameter vector $\mathbf{p}_\lambda^{[i]}$. This is used to predict the ith measurement as $f(x_i, y_i; \mathbf{p}_\lambda^{[i]})$. This prediction is compared against the actual value z_i. This is averaged over the n data points, giving:

$$OCV(\lambda) = \frac{1}{n} \sum_{i=1}^{n} \left(f(x_i, y_i, \mathbf{p}_\lambda^{[i]}) - z_i \right)^2. \quad (16)$$

It is almost impossible to directly use this definition of the OCV as its evaluation for a single value of λ requires the reconstruction of n surfaces. It is well-known [15] that there exists a non-iterative formula that approximates closely the OCV score:

$$OCV(\lambda) = \frac{1}{n} \left\| \Delta \left(\frac{1}{\mathbf{1} - \Delta(H_\lambda)} \right) (H_\lambda - I)\mathbf{z} \right\|_2^2 \quad (17)$$

where I is the identity matrix, Δ the diagonal operator (*i.e.* $\Delta(\mathbf{u})$ is a square matrix having $\mathbf{u}$ as its main diagonal and $\Delta(A)$ extracts the diagonal entries of matrix A as a vector) and H_λ the *influence matrix*:

$$H_\lambda = M \left(M^\mathsf{T} M + \left(\frac{\lambda}{1-\lambda} \right)^2 R^\mathsf{T} R \right)^{-1} M^\mathsf{T}. \quad (18)$$

Even with the non-iterative equation (17), two problems remain. First, the amount of computation required to solve the minimization problem (15) is still too heavy for large datasets (say $n > 1000$). Second, minimizing the OCV score can be difficult. Indeed, this criterion is numerically unstable. This has the effect to introduce high frequency oscillations (see figure 1b) . It is thus difficult to estimate the criterion derivative which would be useful in an optimization process such as gradient descent.

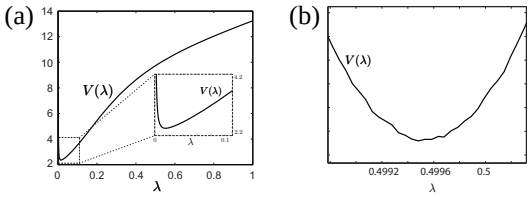

Figure 1. (a) A typical cross-validation score function. (b) High frequency oscillations resulting of numerical instability of the cross-validation score.

3.2. The L-Curve

The L-curve was introduced in [11]. An extensive review of this approach can be found in [9, 10]. The idea of this criterion is to find the best compromise between the goodness of fit and the surface smoothness. To do so, these two quantities are plotted against each other as functions of the regularization parameter.

Let $\rho(\lambda) = \|M\mathbf{p}_\lambda^* - \mathbf{z}\|$ be the *residual norm* and $\eta(\lambda) = \|R\mathbf{p}_\lambda^*\|$ be the *solution norm*. The L-curve is a continuous curve parametrized by the regularization parameter λ and defined by:

$$\left\{ \left(\hat{\rho} = \log \rho(\lambda), \hat{\eta} = \log \eta(\lambda) \right) \in \mathbb{R}_+^2 \mid \lambda \in]0,1[\right\}. \quad (19)$$

The L-curve method chooses one of the maximizers of the L-curve curvature, leading to:

$$\lambda^* = \arg \max_{\lambda \in]0,1[} \kappa(\lambda), \quad (20)$$

where κ is the curvature of the L-curve:

$$\kappa(\lambda) = 2 \frac{\hat{\rho}' \hat{\eta}'' - \hat{\rho}'' \hat{\eta}'}{\left(\hat{\rho}'^2 + \hat{\eta}'^2 \right)^{3/2}}. \quad (21)$$

When the L-curve has the shape of the letter L (see figure 2a), the 'corner' of the curve is well defined: the curvature (figure 2b) has one maximum which corresponds to the regularization parameter λ^* we are searching for. Unfortunately, the curvature often exhibits multiple maxima (see figure 3). In such cases, it is not clear how to choose the regularization parameter.

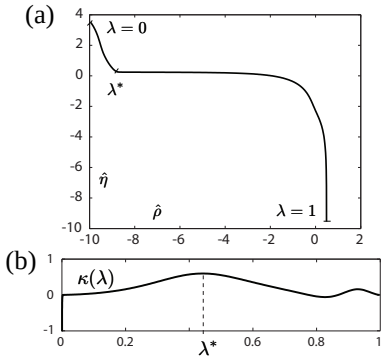

Figure 2. An L-curve (a) and its curvature (b).

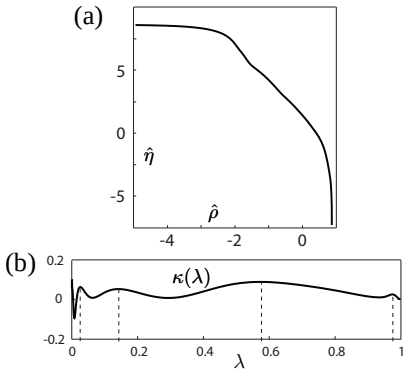

Figure 3. A pathological L-curve (a) and its curvature (b) obtained with the data used in figure 5.

4. The L-Tangent Norm

4.1. The Proposed Criterion

One thing that can easily be noticed when dealing with L-curves is that their parametrization is not uniform. In particular, one can observe that there exists a range of values for λ where the tangent vector norm is significantly smaller than elsewhere. Our new criterion is based on this observation. The regularization parameter is chosen as the one for which the L-curve tangent norm is minimal. Intuitively, such a regularization parameter is the one for which a small variation of the regularization parameter has the lowest impact in the trade-off between the goodness of fit and the surface smoothness.

The L-tangent norm criterion can be written as:

$$\lambda^* = \arg\min_{]0,1[} L(\lambda) \tag{22}$$

$$\text{with } L(\lambda) = \|(\overline{\eta}'_\lambda, \overline{\rho}'_\lambda)\|_2^2 . \tag{23}$$

$\overline{\rho}'_\lambda$ and $\overline{\eta}'_\lambda$ are the derivatives with respect to λ of the normalized residual and solution norms:

$$\overline{\rho}_\lambda = \frac{\rho_\lambda - \rho_\varepsilon}{\rho_{1-\varepsilon} - \rho_\varepsilon}, \qquad \overline{\eta}_\lambda = \frac{\eta_\lambda - \eta_{1-\varepsilon}}{\eta_\varepsilon - \eta_{1-\varepsilon}} \tag{24}$$

for ε a small positive constant (10^{-6}, for instance).

4.2. Properties of the L-Tangent Norm Criterion

A typical example of the L-tangent norm criterion is shown in figure 4c. Even if our criterion is not convex, it is continuous and smooth enough to make it interesting from the optimization point of view. Moreover, neglecting the values of λ very close to 1, our criterion often has a unique minimum, which is not the case of the L-curve criterion.

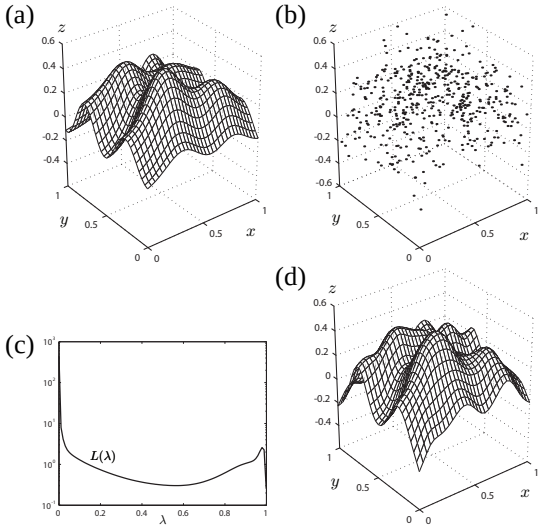

Figure 4. Example of the L-tangent norm criterion. (a) An initial surface. (b) The initial surface sampled on a set of 500 points and with added gaussian noise. (c) The L-tangent norm criterion. (d) The reconstructed surface using the optimal regularization parameter found with the L-tangent norm.

It sometimes happens that there are two minima. In such cases, it seems that these two local minima are both meaningful. The smaller one (*i.e.* the global minimum) corresponds to the regularization parameter giving the best of the two 'explanations' of the data. The second one seems to appear when the data contains, for instance, a lot of small oscillations. In this case, it is not clear (even for a human being) whether the surface must interpolate the data or approximate it, considering the oscillations as some kind of noise. This situation is illustrated in figure 5.

The evaluation of the L-tangent norm criterion requires only the computation of the residual and solution norm derivatives. This makes our new criterion faster to compute

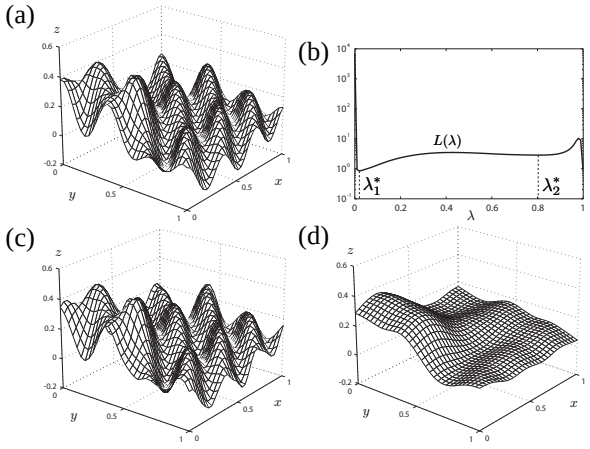

Figure 5. An example of the L-tangent norm criterion presenting two meaningful minima. (a) An initial surface containing a lot of small oscillations. (b) The L-tangent norm criterion presents two minima (excluding the one reached for λ close to 1). (c) The reconstructed surface using the first minimum ($\lambda_1^* = 0.0189$). (d) The reconstructed surface using the second minimum ($\lambda_2^* = 0.8073$).

than cross-validation. In particular, our criterion allows one to improve the computation time when the surface model leads to sparse collocation and regularization matrices (as it is the case with the TPBS model). This is not possible with the cross-validation because the influence matrix is generally not sparse.

Another advantage of the L-tangent norm criterion is that it would still be performant with a non-linear surface model. While cross-validation needs a non-iterative formula to achieve acceptable computational time (which does not necessarily exists for such surface models), our criterion just needs the computation of the residual and the solution norms.

4.3. The Optimization Process

In order to use the L-tangent norm criterion, the computation of the residual and the solution norms is needed. This can be done using finite differences:

$$\overline{\eta}'_\lambda \approx \frac{\overline{\eta}_{\lambda+\delta} - \overline{\eta}_\lambda}{\delta} \qquad \overline{\rho}'_\lambda \approx \frac{\overline{\rho}_{\lambda+\delta} - \overline{\rho}_\lambda}{\delta} \tag{25}$$

where δ is a small positive constant (say, for instance, 10^{-6}).

We have used a multistart Sequential Quadratic Programming (SQP) method in order to solve the optimization problem (23). The criterion is first evaluated over a small set of values (such as $\left\{\frac{1}{10}, \frac{3}{10}, \frac{5}{10}, \frac{7}{10}, \frac{9}{10}\right\}$) in order to determine the starting point of the SQP algorithm. Then the criterion is minimized using the *fmincon* function of Matlab.

5. Experimental Results

5.1. Data

Synthetic data. The first type of data we have used in these experiments are generated by taking sample points (with added noise) of surfaces defined by:

$$g(x,y) = \sum_{i=1}^{8} \frac{2(1-d)c_1}{5} g_1(x,y) + \frac{dc_2}{5} g_2(x,y)$$
$$g_1(x,y) = \exp\Big(-\frac{20\big(a_1(x-a_2)^2 + a_3(y-a_4)^2\big)}{a_5}\Big)$$
$$g_2(x,y) = \sin 4\pi \Big(b_1(x+2b_2)^{\frac{1}{2}+b_3} + b_4(y+2b_5)^{\frac{1}{2}+b_6}\Big)$$

where $a_1, \ldots, a_5, b_1, \ldots, b_6, c_1, c_2$ are randomly chosen in $[0,1]$ and where d is randomly chosen in $\{0,1\}$. Examples of generated surfaces are given in figures 4a and 5a. A noisy sample is shown in figure 4b.

Real data. The second type of data we have used are range images acquired by stereo imaging means. A range image is an image for which each pixel (x,y) is associated to a depth information $D(x,y)$ (and, possibly, a color information $C(x,y)$). Figure 6a is a representation of a range image as a textured surface for which a pixel (x,y) has color $C(x,y)$ and elevation $D(x,y)$. We call *depth map* the picture such that the pixels have a color proportional to their elevation. The depth map corresponding to the range image of figure 6a is shown in figure 6b.

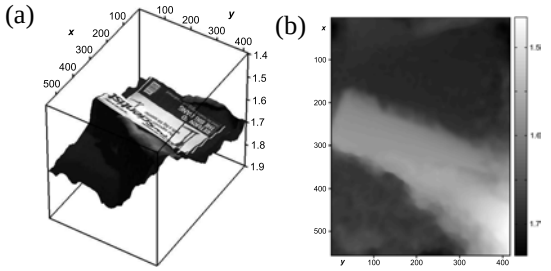

Figure 6. (a) A range image represented as a textured surface. (b) Depth map of the range image (a).

The range images we have used in these experiments are large: their size is approximately 400×600 pixels. It is thus difficult (even impossible) to reconstruct a surface from the original datasets. This is the reason why the range images have been subsampled over a regular grid of size 30×45.

However, the full resolution image is used when comparing a reconstructed surface to the initial dataset.

Three range images have been used in these experiments. The first one is represented in figure 6. The two others are shown in figure 7.

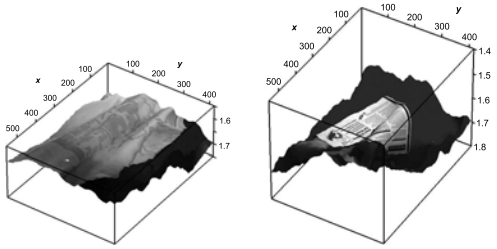

Figure 7. Range images number 2 and 3.

5.2. Computation Timings

Single point evaluation. We intend to compare the computation time of the evaluation for a single value of the regularization parameter of the cross-validation score and the L-tangent norm. To do so, we take a surface and we sample it for several number of points. The timings reported in figure 8 have been obtained with the *cputime* function of Matlab and for the (arbitrary) regularization parameter $\lambda = \frac{1}{2}$. Note that the timing for each distinct number of points has been repeated 5 times in order to get reliable results. Not surprinsingly, figure 8 tells us that the evaluation

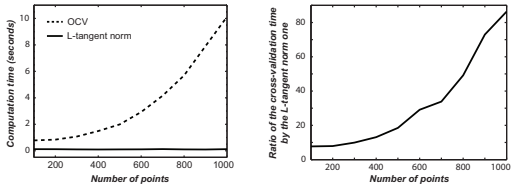

Figure 8. Comparison of the cross-validation versus the L-tangent norm computation time for a single point evaluation.

for a single point with the L-tangent norm is far much faster than with cross-validation. This comes from the fact that the inversion of a matrix is needed in the computation of the cross-validation score while only multiplications between sparse matrices and vectors are involved in the L-tangent norm computation.

Optimization of the criterion. In this experiment, we are interested in the computation time of the whole optimization process for both the L-tangent norm and cross-validation. We have taken 300 examples of randomly generated surfaces known through a noisy sampling. The optimization of the L-tangent norm is realized with the process described in section 4.3. The cross-validation optimization process is performed using a golden section search (implemented in the *fminbnd* function of Matlab). The results are shown in figure 9. As in the previous experiment, the optimization of the L-tangent norm is faster than for cross-validation.

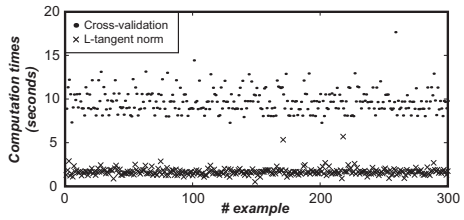

Figure 9. Computation times needed to optimize the L-tangent norm and the cross-validation.

Reconstruction of whole surfaces. Figure 10 shows the computation times needed to the whole surface reconstruction problem with the three range images presented in figure 6a and figure 7. Timings for both the L-tangent norm criterion and the cross-validation score are given in figure 10. As expected, using the L-tangent norm is faster than using cross-validation.

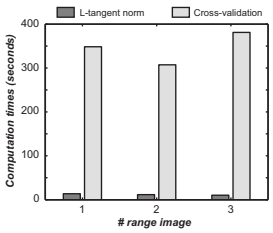

Figure 10. Computation time needed to reconstruct the whole surface from the range images of figures 6a and figure 7 using the L-tangent norm and cross-validation.

5.3. Is L-Tangent Norm an Approximation of Cross-Validation?

This experiment aims to compare the regularization parameter obtained with our L-tangent norm and with the cross-validation criterion. To do so, we have taken noisy samples of randomly generated surfaces. Then, the regularization parameters obtained with cross-validation (the λ_c^*)

for each dataset are plotted versus the regularization parameter determined with the L-tangent norm (the λ_l^*). The results are reported in figure 11.

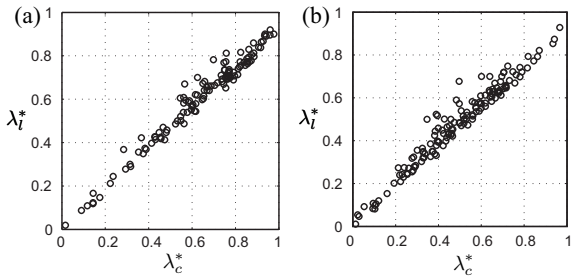

Figure 11. Comparison of the regularization parameters obtained with the L-tangent norm (λ_l^*) with the ones obtained with cross-validation (λ_c^*). (a) Gaussian noise. (b) Uniform noise.

We see from figure 11 that the regularization parameters obtained with the L-tangent norm are often close to the ones obtained with cross-validation. One can observe that the L-tangent norm tends to slightly under-estimate large regularization parameters. However, large regularization parameters are usually obtained for datasets with a lot of noise or badly constrained. In such cases, the accuracy of the regularization parameter does not matter so much.

5.4. Reconstructed Surfaces

Synthetic data. In this experiment, we compare the surfaces reconstructed from data obtained as noisy discretization of randomly generated surfaces. Let us denote f the original randomly generated surface, f_c, f_l and f_n the surfaces reconstructed using respectively cross-validation, the L-curve criterion and our L-tangent norm. The difference between the original surface and the reconstructed ones is measured with the *Integral Relative Error (IRE)*. If the functions f, f_c, f_l and f_n are all defined over the domain Ω, the IRE is given by:

$$e(f,g) = \frac{\iint_\Omega |g(x,y) - f(x,y)| \, \mathrm{d}x\mathrm{d}y}{\left| \max_{(x,y)\in\Omega} f(x,y) - \min_{(x,y)\in\Omega} f(x,y) \right|}, \qquad (26)$$

where the function g is f_c, f_l or f_n. The results of this experiment are reported in figure 12. This figure tells us that the reconstruction errors are small and similar for cross-validation and the L-tangent norm. The IRE for surfaces reconstructed using the L-curve criterion are much larger than with the two other criteria. Moreover, only the IRE less than 1 are reported for the L-curve criterion: the IRE was greater than 1 for 48 test surfaces. These large IRE are mainly due to a failure in the maximization of the L-curve criterion.

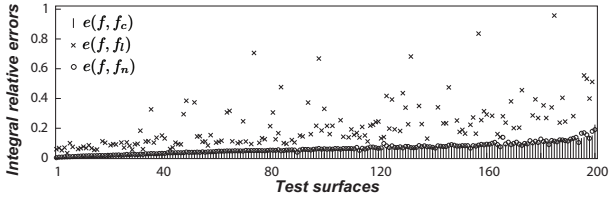

Figure 12. Integral relative errors for 200 randomly generated surfaces sampled over 500 points with added gaussian noise.

Range images. In this last experiment, we intend to compare the surfaces reconstructed from real range images. To do so, we take again the three range images of figure 6 and 7. Let f_i be the original range image (before subsampling). Let f_l and f_c be the reconstructed surfaces using respectively the L-tangent norm and cross-validation to choose the regularization parameter. The results of this experiment are presented in the form of *Relative Error Maps (REM)*. The REM for the surface reconstructed using the L-tangent norm is a picture such that each pixels (x, y) is associated to a color $C_l(x, y)$ proportional to the difference of depth between the reconstructed surface and the original one. This is written as:

$$C_l(x,y) = \frac{|f_i(x,y) - f_l(x,y)|}{\left| \max_{(u,v)} f_i(u,v) - \min_{(u,v)} f_i(u,v) \right|}. \qquad (27)$$

The REM for the surface reconstructed using cross-validation, C_c, is defined similarly to equation (27) except that f_l is replaced by f_c. We also define the *Difference Error Map (DEM)* by $C_{l,c}(x, y) = |C_c(x, y) - C_l(x, y)|$. The results of the comparison between surfaces reconstructed from range images using the L-tangent norm and the cross-validation are reported in figure 13. On this figure, only the error map for the L-tangent norm is reported. Indeed, as it is shown in figures 13(d-f), the two reconstructed surfaces are very similar (which is the point of main interest in this experiment). Even if the reconstruction errors are not negligible (figures 13(a-c)), they are still small. The main reason for these errors is the subsampling of the initial datasets.

6. Conclusion

We proposed a novel approach to automatically select the regularization weight in the problem of surface reconstruction from range data. Experimental results show that it compares well to one of the most renowned criterion: the

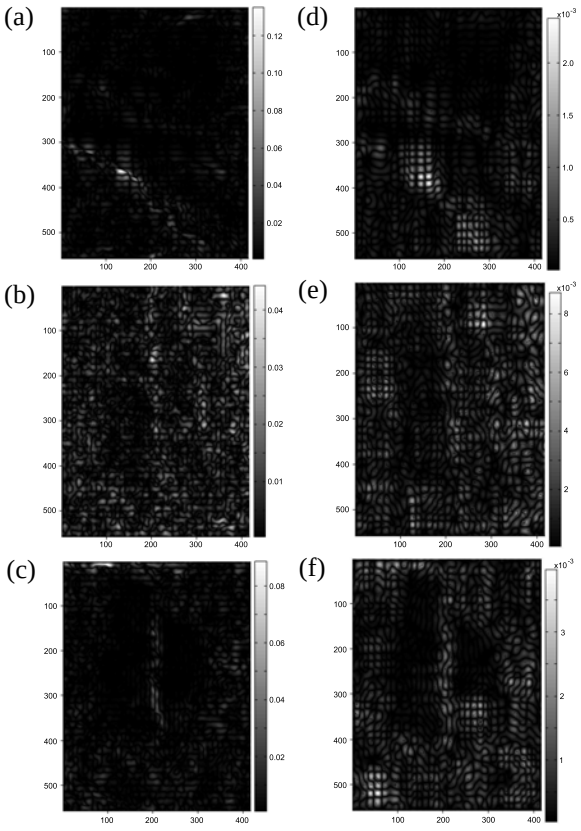

Figure 13. (a-c) REM for the surfaces reconstructed using the L-tangent norm for the three range images. (d-f) DEM between the surfaces reconstructed using the L-tangent norm and the cross-validation.

cross-validation score. Our L-tangent norm criterion is interesting for two main reasons. First, its shape and numerical behavior makes the selection of the regularization parameter easy. Second, it is computationnally cheap. This is especially important for real time processing required by recently proposed range sensors.

Future ways of research include finding a good theoretical explanation to our criterion and its extension to multiple regularization terms.

Acknowledgments The authors thank Toby Collins for providing the range images. This work has been supported by the regional council of Auvergne.

References

[1] Å. Björck. *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, 1996.

[2] F. L. Bookstein. Principal warps: Thin-Plate Splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(6):567–585, 1989.

[3] C. de Boor. *A Practical Guide to Splines, Revised Edition*, volume 27 of *Applied Mathematical Sciences*. Springer Verlag, 2001.

[4] P. Dierckx. *Curve and surface fitting with splines*. Oxford University Press, Inc., New York, NY, USA, 1993.

[5] G. Donato and S. Belongie. Approximate Thin-Plate Spline mappings. *Lecture Notes In Computer Science*, 2352:21–31, 2002.

[6] J. Duchon. Splines minimizing rotation-invariant seminorms in sobolev spaces. *Constructive Theory of Functions of Several Variables*, 1:85–100, 1977.

[7] G. E. Farin. *Curves and Surfaces for Computer Aided Geometric Design*. Academic Press, 1997.

[8] M. Fornefett, K. Rohr, and H. Stiehl. Radial basis functions with compact support for elastic registration of medical images. *Image and Vision Computing*, 19:87–96, 2001.

[9] P. Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Rev.*, 34(4):561–580, 1992.

[10] P. Hansen. The L-curve and its use in the numerical treatment of inverse problems. Technical report, Technical University of Denmark, 2005.

[11] C. L. Lawson and R. J. Hanson. *Solving Least Squares Problem*. Prentice Hall, Englewood Cliffs, 1974.

[12] M. J. Powell. Five lectures on radial basis functions. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, 2005.

[13] R. M. Rifkin and R. A. Lippert. Notes on regularized least squares. Technical report, Computer Science and Artificial Intelligence Laboratory, May 2007.

[14] S. Roberts and L. Stals. Discrete thin plate spline smoothing in 3D. *ANZIAM*, 45:646–659, 2004.

[15] G. Wahba. *Spline models for observational data*. SIAM [Society for Industrial and Applied Mathematics], 1990.

[16] G. Wahba and S. Wold. A completely automatic French curve: Fitting spline functions by cross validation. *Communications in Statistics*, 4:1–18, 1975.