Cramer-Rao Bounds for Nonparametric Surface Reconstruction from Range Data

Tolga Tasdizen   Ross Whitaker

University of Utah, School of Computing
Technical Report UUCS-03-006

Also submitted to the 4th Int. Conference on 3-D Digital Imaging and Modeling for review

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

April 18, 2003

Abstract

The Cramer-Rao error bound provides a fundamental limit on the expected performance of a statistical estimator. The error bound depends on the general properties of the system, but not on the specific properties of the estimator or the solution. The Cramer-Rao error bound has been applied to scalar- and vector-valued estimators and recently to parametric shape estimators. However, nonparametric, low-level surface representations are an important important tool in 3D reconstruction, and are particularly useful for representing complex scenes with arbitrary shapes and topologies. This paper presents a generalization of the Cramer-Rao error bound to nonparametric shape estimators. Specifically, we derive the error bound for the full 3D reconstruction of scenes from multiple range images.

Chapter 1

Introduction

A confluence of several technologies has created new opportunities for reconstructing 3D models of complex objects and scenes. More precise and less expensive range measurement systems combined with better computing capabilities enable us to build, visualize, and analyze 3D models of the world. The difficulty of reconstructing surfaces from range images stems from inadequacies in the data. Range measurements presents several significant problems, such as measurement noise, variations in measurement density, occlusions, and errors in the registration of multiple range images. Hence, the reconstructed surfaces are not perfect, they are merely *estimates* of the true surfaces. As the use of measured 3D models becomes more commonplace, there will be a greater need for quantifying the errors associated with these models. For instance, the use of 3D models in forensics, to model crime scenes [1], will invariably raise the question, "How much can we trust these models?"[1]

Signal processing, and estimation theory in particular, provides a tool, the Cramer-Rao error bound (CRB), for quantifying the performance of statistical estimators. However, the CRB has traditionally been applied to *parameter estimation* problems. That is, problems in which the number of parameters and their relationship to the physical measurements is fixed and known. In order apply these tools to surface reconstruction, we must first define a notion of error for surfaces and then adapt these tools to a 3D geometric setting.

The analysis of reconstruction errors depends on the surface representation. For this discussion we divide the space of surface models into two classes: parametric and nonparametric. Parametric models are those that represent shapes indirectly via a finite set of variables that

[1]http://www.3rdtech.com/DeltaSphere_at_AAFS_Conf.htm

control the local or global position of the surface. Parametric models range from simple primitives that have a few parameters to more complicated algebraic polynomial surfaces and piecewise smooth models, such as splines. Parametric approaches are particularly well suited to higher-level tasks such as object recognition. In the context of estimation, the number of parameters and their relationship to the shape is not usually considered as a random variable. Therefore, parametric models restrict the solution to the space of shapes that are spanned by the associated parameters.

The alternative is a *nonparametric* model, which, for the purposes of this paper, refers to those representations in which the position of any point on the surface is controlled directly and is independent (to within a finite resolution) from the positions of other points on the surface. According to this definition, surface meshes, volumes, and level sets are examples of nonparametric shape representations. Nonparametric models typically have many more free parameters (e.g. each surface point, their number, and their configuration) and they represent a much broader class of shapes. However, nonparametric models impose other limitations such as finite resolution and, in the case of implicit models, closed boundaries. Nevertheless, the literature has shown that nonparametric models are preferred when reconstructing surfaces of complex objects or scenes with arbitrarily topology and very little a-priori knowledge about shape [2, 3, 4, 5]. This paper introduces a novel formulation for computing expected errors of nonparametric surface estimates using point-wise Cramer-Rao bounds.

The rest of this report is organized as follows. Chapter 2 discusses related work, Chapter 3 summarizes the maximum likelihood nonparametric surface estimation process, and Chapter 4 derives a CRB for nonparametric surface estimators and gives results for synthetic data. Chapter 5 presents results for real data. Chapter 6 summarizes the contributions of this paper and discusses possibilities for future research directions.

Chapter 2

Related Work

The CRB states the minimum achievable error for an estimator, and therefore, provides fundamental limits on the performance of any estimation process. The expression for the CRB is independent of the specific form of the estimator; it depends only on the statistics of the input measurements and the bias of the estimator. Moreover, for asymptotically efficient estimators, such as the maximum likelihood estimator (MLE), the CRB is a tight lower bound, i.e. for MLEs the CRB is achievable. Thus, the CRB quantifies the expected error of the output of an estimation process in the absence of ground truth. In the context of surface reconstruction, it provides a well-founded, systematic mechanism for computing the error of a reconstructed surface.

Researchers have extensively used CRBs for problems where the estimator is relatively simple, such as scalar or vector quantities. For instance, parameter estimation to determine the location, size and orientation of a target has been studied using CRB analysis [6]. More recently, several authors have derived CRB expressions for parametric shape estimators. Hero *et al.* [7] compute the CRB for B-spline parameters of star-shapes estimated from magnetic resonance imagery. Ye *et al.* [8] compute the CRB for more general parametric shape estimators. Confidence intervals for shape estimators can be computed using CRBs [9], which provides an important computational advantage over using a Monte-Carlo simulation [10]. However, these results apply only to parametric shape estimators. The goal of this paper is to fill an gap in 3D surface reconstruction by deriving the CRB for nonparametric shape estimators and expressing the error in terms of a statistical model of a scanning laser range finder.

Chapter 3

Maximum Likelihood Surface Reconstruction

This chapter describes a particular formulation for a nonparametric MLE surface estimator. The results in this paper establish a bound that applies to any nonparametric surface estimator. However, these results provide a tight bound for MLE estimators, and the formulation for the MLE estimator introduces some basic concepts that are important for the CRB.

We begin by describing a mathematical model of a range image. A *range finder* is a device that measures distances to the closest point on an object along a particular *line of sight*. A *range scanner* produces a 2D array (image) of range measurements $r_{i,j}$, through a scanning mechanism that aims the line of sight accordingly, see Figure 3.1. Therefore each element or pixel of a range image consists of two things: a line of sight and a range measurement, which together describe a 3D point. We denote a single range image $R^{(j)}$ and a collection of range images taken from different scanner locations as $\{R^{(1)}, ..., R^{(M)}\}$.

The object or scene also requires a precise specification. We define the surface $\mathcal{S}$ as the closure of a compact subset of 3D, Ω. Thus, $\mathcal{S}$ is the "skin" that covers the solid Ω.

The range measurements are random variables, but if we know the sensor model, we can compute the probability of a particular set of range image conditional on the scene as $P\left(R^{(1)}, ..., R^{(M)}|S\right)$. This is the *likelihood*. An MLE estimator is defined as

$$\hat{S} = \arg\sup_{S} P\left(R^{(1)}, ..., R^{(M)}|S\right). \tag{3.1}$$

Any estimator that minimizes the likelihood is asymtotically efficient and unbiased. That

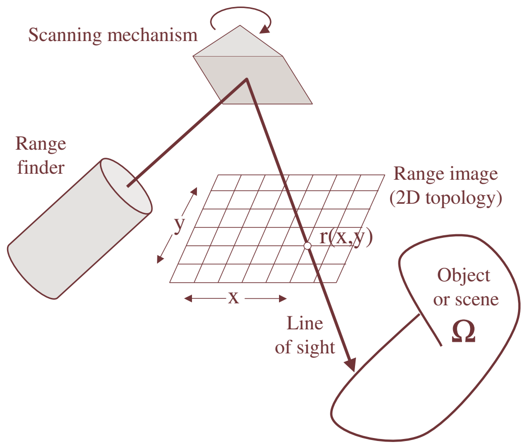

Figure 3.1: A range finder produces a dense range map of a scene.

is, as the number of measurements goes to infinity, the estimator is correct on the average and is as good as any other unbiased estimator.

Whitaker [5] shows that the maximum likelihood estimator of such a collection of range images can be computed as the set of zero crossings of scalar function $G(\boldsymbol{x})$. That is

$$\hat{S} = \left\{ \mathbf{x} \in \mathbb{R}^3 \mid G(\mathbf{x}) = 0 \right\}, \tag{3.2}$$

Curless [3] uses a similar implicit formulation to reconstruct surfaces from multiple range scans.

For the MLE formulation $G(\boldsymbol{x})$ is

$$G(\mathbf{x}) = \sum_{j=1}^{M} h\left(D_j\left(\mathbf{x}\right)\right), \tag{3.3}$$

where $D_j(\mathbf{x})$ is the distance from the j'th scanner location to the point $\mathbf{x} \in \mathbb{R}^3$, and h is the derivative of the logarithm of the pdf for the range measurement error model. The model assumes that the range measurements within a single scan are sufficiently close and are suitable for interpolation. Notice that if the sensor model is Gaussian h is linear. Certain classes of range measurements, such as ladar, have been shown to have noise characteristics that can be described as Gaussian with outliers [11].

Chapter 4

Cramer-Rao Error Bounds

The only description of the surface offered by nonparametric estimators, (3.2), is the set of points $\hat{\mathbf{x}}$ in 3D that lie on the surface. Therefore we formulate the error as a separate bound for each surface point. Errors on points are directional, but without any correspondence between the estimate and the true surface, the only important aspect of the error is how far each point is from the nearest point on the actual surface. Given a point $\hat{\mathbf{x}}$, we can compute the CRB as $E\left[||\hat{\mathbf{x}} - \mathcal{S}||^2/\mathcal{S}\right]$, where $||\hat{\mathbf{x}} - \mathcal{S}||$ denotes the shortest Euclidean distance between $\hat{\mathbf{x}}$ and S, the actual surface, and E denotes the expected value for all possible $\hat{\mathbf{x}}$. The local error bound gives us a map of errors over the entire surface estimate. This is a more useful and general result than a global error bound.

Let n be the number of scanners (range images) to which the point $\hat{\mathbf{x}}$ is visible. Each of the n scanners has one line of sight, $\mathbf{L}^{(i)}$, associated with $\hat{\mathbf{x}}$. This is the vector from the scanner location to $\hat{\mathbf{x}}$. Let $r^{(i)}$ denote the range measurement for the *i*'th scanner taken for the line of sight $\mathbf{L}^{(i)}$. In principle, $\hat{\mathbf{x}}$ is a function only of the set $r^{(1)}, ..., r^{(n)}$. Using the Cramer-Rao error bound formula for unbiased estimators [12], we find

$$E \geq \frac{1}{\sum_{i=1}^{n} \int_0^\infty \left(\frac{dP\left(r^{(i)}|\mathcal{S}\right)}{d\mathcal{S}} \right)^2 \frac{1}{P\left(r^{(i)}|\mathcal{S}\right)} \, dr^{(i)}}. \tag{4.1}$$

In the rest of this paper, we use E to denote $E\left[||\hat{\mathbf{x}} - \mathcal{S}||^2/\mathcal{S}\right]$. Equation (4.1) is a function of the derivative of the pdf with respect to $\mathcal{S}$. To compute this derivative, we use a local first-order approximation to the surface, i.e. the tangent plane. Let $\mathbf{N}$ denote the surface normal vector, which is perpendicular to the tangent plane. Then, perturbations of the surface can be locally approximated by moving the tangent plane along the normal direction $d\mathcal{S} = \epsilon\mathbf{N}$, see Figure 4.1(a). Let $\theta^{(i)}$ denote the angle between $\mathbf{N}$ and $-\mathbf{L}^{(i)}$. Also, let $r_t^{(i)}$ denote

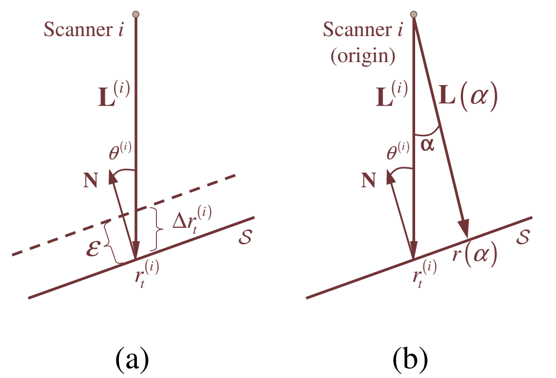

Figure 4.1: (a) The relationship between perturbations of the surface and its true distance from the scanner. (b) The 2D geometry of line of sight error.

the true distance from the scanner to the surface along $\mathbf{L}^{(i)}$. The geometric relationship between the surface perturbation ϵ and the change in the true distance from the scanner to the surface, see Figure 4.1(a), dictates

$$\Delta r_t^{(i)} = \frac{\epsilon}{\cos \theta^{(i)}}. \tag{4.2}$$

Using this relationship, we obtain

$$\begin{aligned}
\frac{dP\left(r^{(i)}|\mathcal{S}\right)}{d\mathcal{S}} &= \lim_{\epsilon \to 0} \frac{P\left(r^{(i)}|\mathcal{S} + \epsilon \mathbf{N}\right) - P\left(r^{(i)}|\mathcal{S}\right)}{\epsilon} \\
&= \lim_{\epsilon \to 0} \frac{P\left(r^{(i)}|r_t^{(i)} + \Delta r_t^{(i)}\right) - P\left(r^{(i)}|r_t^{(i)}\right)}{\epsilon} \\
&= \frac{\lim_{\Delta r_t^{(i)} \to 0} \frac{P\left(r^{(i)}|r_t^{(i)} + \Delta r_t^{(i)}\right) - P\left(r^{(i)}|r_t^{(i)}\right)}{\Delta r_t^{(i)}}}{\cos \theta^{(i)}} \\
&= \frac{1}{\cos \theta^{(i)}} \frac{dP\left(r^{(i)}|r_t^{(i)}\right)}{dr_t^{(i)}}.
\end{aligned} \tag{4.3}$$

Substituting this result into (4.1) yields

$$E \geq \frac{1}{\sum_{i=1}^{n} \sec \theta^{(i)} \int_0^{\infty} \left(\frac{dP\left(r^{(i)}|r_t^{(i)}\right)}{dr_t^{(i)}} \right)^2 \frac{1}{P\left(r^{(i)}|r_t^{(i)}\right)} \, dr^{(i)}}. \tag{4.4}$$

The next step is to formulate the conditional pdf for range measurements. There are two possible sources of error in each range measurement:

- the angular error in aiming the line of sight, and
- the error in the distance measurement along the actual line of sight.

The uncertainty in the line-of-sight can be used to describe several sources of error. First, the scanner measures range along a discrete grid of line-of-sights, and therefore, it introduces a sampling error. Moreover, given an intended line-of-sight on the discrete grid, there is an error in aiming the range finder. We will refer to this discrepancy between the intended and the actual line of sights as pointing error. Finally, when estimating surfaces from multiple range images, error is introduced by imperfections in the registration of the different range images to each other.

For most range scanners, such as *Ladar*, the pointing error is small compared to the error in the distance measurement. Hence, it is common to assume a perfectly aimed line of sight, to simplify the formulation of the conditional pdf. In this case, $P\left(r^{(i)}|\mathcal{S}\right)$ depends only on the true distance from the i'th scanner to $\mathcal{S}$ along the vector $\mathbf{L}^{(i)}$. We can assume a Gaussian distribution for the noise in the distance measurement [5], and therefore

$$P\left(r^{(i)}|\mathcal{S}\right) = P\left(r^{(i)}|r_t^{(i)}\right) \sim \mathcal{N}(r_t^{(i)}, \sigma_r^2). \tag{4.5}$$

Using results for Gaussian pdf's from [12], we find that

$$\int_{r^{(i)}} \left(\frac{dP(r^{(i)}|r_t^{(i)})}{dr_t^{(i)}}\right)^2 \frac{1}{P(r^{(i)}|r_t^{(i)})} dr^{(i)} = \sigma_r^2 \tag{4.6}$$

Substituting this result into (4.4), we get

$$E \geq \frac{1}{\sum_{i=1}^{n} 1/\sigma_r^2 \cos\theta^{(i)}} = \frac{\sigma_r^2}{\sum_{i=1}^{n} \sec\theta^{(i)}}. \tag{4.7}$$

This result states that if any of the n scanners have a line of sight that is perpendicular to the normal vector at $\hat{\mathbf{x}}$ ($\theta^{(i)} = 0$), the error bound for that point is zero. Figure 4.2 demonstrates this result with a sphere. We compute the CRB for estimating a surface from six noisy range images of a sphere with unit radius taken along the six cardinal directions. Figure 4.2 shows the CRB as a colormap on the surface; the units are the radius of the sphere. The scanners are located on the axis along the purple regions on the estimator. As predicted by (4.7) these are the regions of highest expected error. The red regions, where the CRB is 0, form six circles on the sphere. These circles are the silhouettes of the sphere as seen from the scanner locations. Therefore, according to this incomplete CRB derivation, it should be possible to determine the location of any desired point exactly by

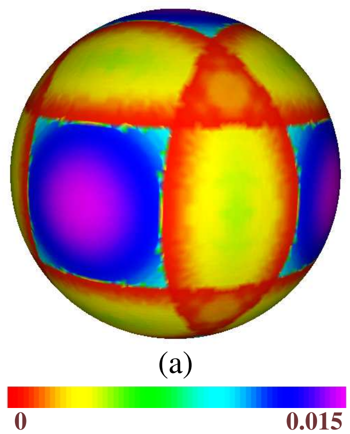

(a)

(b)

Figure 4.2: [Color]The incomplete CRB shown as a colormap on the sphere, and (b) the color map for the CRB. The radius of the sphere is 1 unit.

repositioning the scanner. This counter-intuitive result is due to ignoring the angular error in the line of sight. In practice, this error is non-zero, and we can not determine any point on an object error-free. We derive a complete conditional pdf and CRB in the rest of this section.

4.1 Error bound in 2D

We can derive an accurate conditional pdf for the range measurement if we take the pointing error in the line of sight into account. Let us first examine the simpler 2D case, where $\Omega \subset \mathbb{R}^2$ and $\mathcal{S}$ is a curve. The vector from the i'th scanner to $\hat{\mathbf{x}}$ on the surface estimator, $\mathbf{L}^{(i)}$, is now the intended line of sight. Figure 4.1(b) illustrates $\mathbf{L}(\alpha)$ and α which represent the actual line of sight (random variable A) and the angle it makes with the intended line of sight, respectively. We assume that the pdf for α is a Gaussian with zero mean (there is no constant offset error in aiming the scanner) and σ_a standard deviation

$$P(\alpha) \sim \mathcal{N}(0, \sigma_a^2). \tag{4.8}$$

Given this actual line of sight, we assume a Gaussian distribution for the distance measurement (random variable B)

$$P\left(r^{(i)} | \mathbf{L}(\alpha), \mathcal{S}\right) = P\left(r^{(i)} | r(\alpha)\right) \sim \mathcal{N}(r(\alpha), \sigma_r^2), \tag{4.9}$$

where $r(\alpha)$ is the actual distance to the surface $\mathcal{S}$ along $\mathbf{L}(\alpha)$. Random variables A and B are independent; therefore, their joint probability is the product of (4.8) and (4.9). Integrating this joint probability over the domain of α, we compute the marginal distribution

$$P\left(r^{(i)}|\mathcal{S}\right) = \int_{-\pi}^{\pi} P\left(r^{(i)}|\mathbf{L}(\alpha), \mathcal{S}\right) P(\alpha) d\alpha. \tag{4.10}$$

To evaluate this probability, we still need to determine the expression for $r(\alpha)$ in (4.9).

Without loss of generality, define the scanner location and $\mathbf{L}^{(i)}$ to be the origin of the coordinate frame and the *y-axis*, respectively. Then, we have $\mathbf{L}(\alpha) = \begin{pmatrix} \sin\alpha \\ \cos\alpha \end{pmatrix}$ and $\mathbf{N} = \begin{pmatrix} \sin\theta^{(i)} \\ \cos\theta^{(i)} \end{pmatrix}$, see Figure 4.1(b). Using the equation for the tangent line

$$\left(r(\alpha)\mathbf{L}(\alpha) - r_t^{(i)}\mathbf{L}^{(i)}\right) \cdot \mathbf{N} = 0, \tag{4.11}$$

the distance $r(\alpha)$ can be found as

$$\begin{aligned} r(\alpha) &= r_t^{(i)} \frac{\mathbf{L}^{(i)} \cdot \mathbf{N}}{\mathbf{L}(\alpha) \cdot \mathbf{N}} \\ &= r_t^{(i)} \frac{\cos\theta^{(i)}}{\cos\alpha\cos\theta^{(i)} - \sin\alpha\sin\theta^{(i)}} \\ &\approx r_t^{(i)} \frac{\cos\theta^{(i)}}{\cos\theta^{(i)} - \alpha\sin\theta^{(i)}} \qquad (4.12) \\ &\approx r_t^{(i)} \left(1 + \alpha\tan\theta^{(i)}\right) \qquad (4.13) \end{aligned}$$

Equation (4.12) follows from the fact that α is very small, and therefore $\sin\alpha \approx \alpha$ and $\cos\alpha \approx 1$. Equation (4.13) is obtained using a first order Taylor series approximation to (4.12) around $\alpha = 0$.

Let us put together equations (4.8-4.10) and (4.13) to evaluate $P\left(r^{(i)}|\mathcal{S}\right)$

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_r} e^{\frac{\left(r^{(i)} - r_t^{(i)}\left(1+\alpha\tan\theta^{(i)}\right)\right)^2}{2\sigma_r^2}} \frac{1}{\sqrt{2\pi}\sigma_a} e^{-\frac{\alpha^2}{2\sigma_a^2}} d\alpha. \tag{4.14}$$

We can change the range of integration from $(-\pi, \pi)$ in (4.10) to $(-\infty, \infty)$ above because $\sigma_a \ll \pi$, and hence $P(\alpha) \approx 0$ for $|\alpha| \geq \pi$. Then, using the change of variables $y = r_t^{(i)} \tan\theta^{(i)}\alpha$, the above integral can be rewritten as

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_r} e^{\frac{\left(r^{(i)} - r_t^{(i)} - y\right)^2}{2\sigma_r^2}} \frac{1}{\sqrt{2\pi}\sigma_a'} e^{-\frac{y^2}{2(\sigma_a')^2}}\, dy, \tag{4.15}$$

where

$$\sigma_a' = r_t^{(i)} \tan\theta^{(i)} \sigma_a. \tag{4.16}$$

Equation (4.15) is in the form of convolution of two Gaussians. Consequently, probability theory states that the result is the Gaussian pdf

$$P\left(r^{(i)}|\mathcal{S}\right) = P\left(r^{(i)}|r_t^{(i)}, \theta^{(i)}\right) \sim \mathcal{N}\left(r_t^{(i)}, \sigma_r^2 + (\sigma_a')^2\right). \tag{4.17}$$

For the purposes of differentiating this pdf with respect to $r_t^{(i)}$, we ignore the dependence of σ_a' on $r_t^{(i)}$. Hence, the derivation (4.3) also applies here. Using results for Gaussian pdf's from [12], we find that

$$\int_{r^{(i)}} \left(\frac{dP(r^{(i)}|r_t^{(i)}, \theta^{(i)})}{dr_t^{(i)}}\right)^2 \frac{1}{P(r^{(i)}|r_t^{(i)}, \theta^{(i)})} dr^{(i)} = \sigma_r^2 + (\sigma_a')^2. \tag{4.18}$$

Finally, using these result in (4.1)

$$\begin{aligned} E &\geq \frac{1}{\sum_{i=1}^n \sec\theta^{(i)} / \left(\sigma_r^2 + \left(r_t^{(i)} \tan\theta^{(i)} \sigma_a'\right)^2\right)} \\ &= \frac{1}{\sum_{i=1}^n \frac{1}{\cos\theta^{(i)} \sigma_r^2 + \sin\theta^{(i)} \left(r_t^{(i)} \sigma_a'\right)^2}} \end{aligned} \tag{4.19}$$

Note that, provided $\sigma_r^2 \neq 0$ and $\sigma_a^2 \neq 0$, which is always the case for a real range finder, (4.19) is never0.

4.2 Error bound in 3D

The 3D case differs from the 2D because there are two angles of deviation in the line of sight that we need to consider; however, we show that only one of these angles matter and the 3D case can be reduced to the 2D case. Consider the geometry illustrated in Figure 4.3(a). Without loss of generality, define the scanner location and $\mathbf{L}^{(i)}$ (the intended line of sight) to be the origin of the coordinate frame and the *z-axis*, respectively. The surface can be locally approximated by its tangent plane defined by the surface normal vector $\mathbf{N}$. Using $\mathbf{N}$, define the following unit vectors:

$$\mathbf{v} = \mathbf{N} \times \mathbf{z} \text{ and } \mathbf{u} = \mathbf{v} \times \mathbf{z}, \tag{4.20}$$

where $\mathbf{z}$ is the unit vector in the direction of the *z-axis*, and "$\times$" denotes the vector cross product. The unit vectors $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{z}$ define an orthonormal coordinate frame.

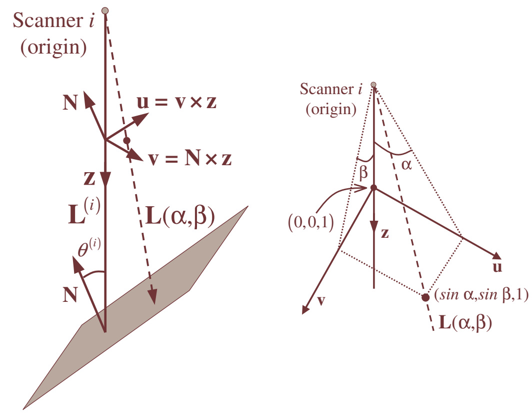

(a) (b)

Figure 4.3: (a) The 3D geometry of line of sight error, (b) the *u*-*v*-*z* coordinate frame.

The two angles of deviation for the line of sight can be defined as rotations of the intended line of sight around any pair of orthogonal pair of vectors in the plane perpendicular to the *z-axis*. Without loss of generality, we choose the angles of deviation α and β to be rotations around the $\mathbf{v}$ and $\mathbf{u}$ axes, respectively, as shown in Figure 4.3(b). The random variables α and β are independent and identically distributed with the pdf $\mathcal{N}\left(0, \sigma_a^2\right)$. The intended line of sight is $\mathbf{L}^{(i)} = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}^T$ in the *u*-*v*-*z* coordinate frame. The actual lines of sight can be expressed in the *u*-*v*-*z* coordinate frame as

$$\begin{aligned} \mathbf{L}\left(\alpha, \beta\right) &= \frac{1}{\sqrt{\sin^2\alpha + \sin^2\beta + 1}} \begin{pmatrix} \sin\alpha \\ \sin\beta \\ 1 \end{pmatrix} \\ &\approx \frac{1}{\sqrt{\alpha^2 + \beta^2 + 1}} \begin{pmatrix} \alpha \\ \beta \\ 1 \end{pmatrix}, \end{aligned} \tag{4.21}$$

where we have used the small angle approximation $\sin\alpha \approx \alpha$ and $\sin\beta \approx \beta$. We can also express the surface normal in the same coordinate frame. By definition, $\mathbf{N} \cdot \mathbf{L}^{(i)} = \mathbf{N} \cdot \mathbf{z} = -\cos\theta^{(i)}$, $\mathbf{N} \cdot \mathbf{v} = \mathbf{N} \cdot (\mathbf{N} \times \mathbf{z}) = 0$ and $||\mathbf{N}|| = 1$; therefore,

$$\mathbf{N} = \begin{pmatrix} \sin\theta^{(i)} \\ 0 \\ \cos\theta^{(i)} \end{pmatrix}.$$

We can use the equation for the tangent plane, which is the same as equation (4.11) for the tangent in 2D, to find the distance $r(\alpha, \beta)$

$$\begin{aligned} r(\alpha, \beta) &= r_t^{(i)} \frac{\mathbf{L}^{(i)} \cdot \mathbf{N}}{\mathbf{L}(\alpha, \beta) \cdot \mathbf{N}} \\ &\approx r_t^{(i)} \frac{\cos \theta^{(i)} \sqrt{\alpha^2 + \beta^2 + 1}}{\cos \theta^{(i)} - \alpha \sin \theta^{(i)}} \\ &\approx r_t^{(i)} \left(1 + \alpha \tan \theta^{(i)}\right) \end{aligned} \tag{4.22}$$

Equation (4.22), which is exactly the same result as (4.13) for 2D, follows from a first order Taylor series approximation around $\alpha = \beta = 0$. Hence, the Cramer-Rao bound (4.19) also applies to the 3D case.

Chapter 5

Results and Discussion

Figure 5.1 shows the results of the experiment from Figure 4.2 with the complete CRB formulation (4.19), which includes angular error. As expected the CRB does not go to zero along occlusion boundaries. This is in contrast to the incomplete CRB formulation, which predicts a zero error bound for the silhouette of a sphere.

As an alternative to the CRB results in this paper, consider the method of computing error measures on simple scalar estimators, which averages all measurements. For such an estimator to be the error $E \geq \sigma^2/n$, where n is the number of measurements and σ^2 is the variance of any one measurement. An naive application of this scheme to surface reconstruction would produce an error measure that depended only on the number of scanners that are visible from any particular surface point. Unlike the CRB derived in this paper, this trivial result does not take the sensor-model geometry into account and is not correct. In Figure 5.1, the CRB is highest (purple-blue) in regions seen by a single scanner, and lowest (yellow) in regions seen by three scanners. However, the CRB varies significantly within these regions which can not be predicted by the trivial approach that discounts the sensor-model geometry.

Next, we compute an actual estimator using a level-set surface representation [5]. Figure 5.1(b) shows a close-up view of the CRB colormapped on to this actual estimator. If we consider the *roughness* of the estimated surface as a subjective indicator of error, we observe that the actual estimation errors are approximately proportional to the error predicted by the CRB. In other words, the estimator is indeed more noisy in blue-purple regions of the CRB compared to the yellow regions.

Finally, we demonstrate the importance of error bounds in a real surface reconstruction

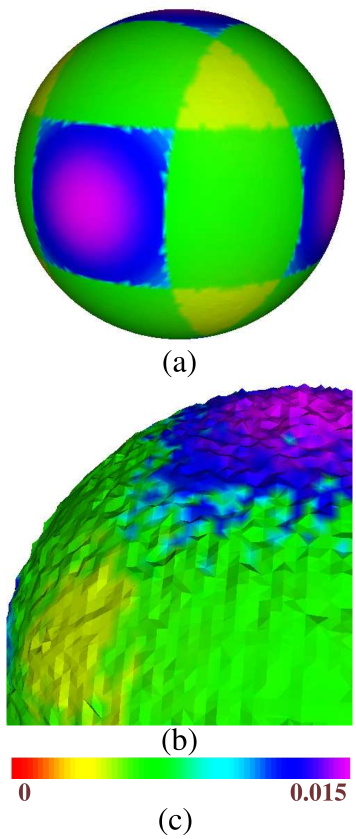

Figure 5.1: [Color](a) The CRB colormapped on the sphere, (b) a close-up view of the CRB shown on an instance of the estimator, and (c) the color map for the CRB. The radius of the sphere is 1 unit.

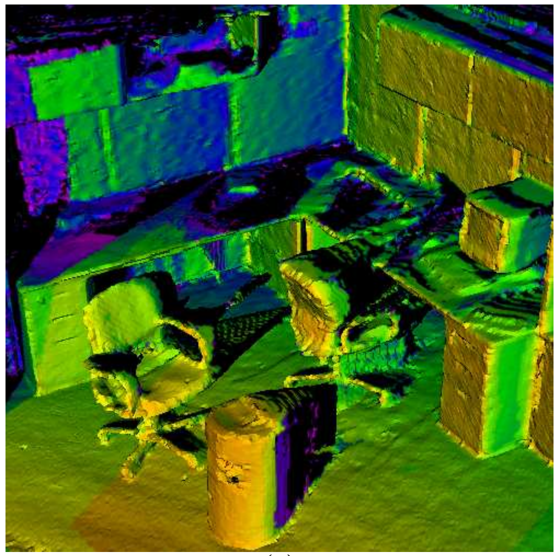

(a)

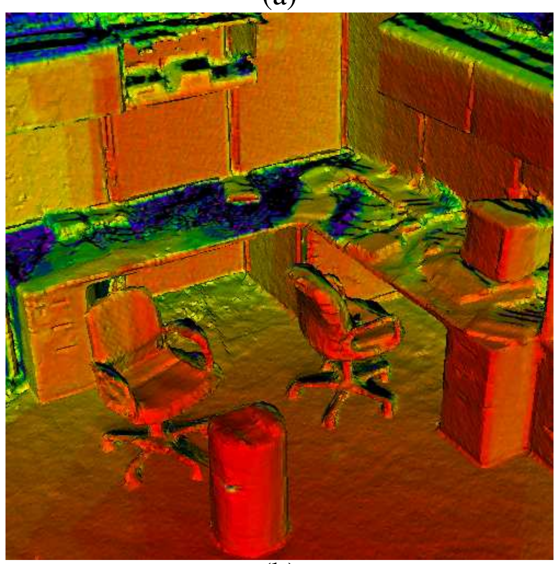

(b)

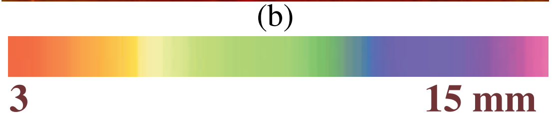

(c)

Figure 5.2: [Color] MLE reconstruction from (a) 4 range images, (b) 12 range images, and (c) CRB colorbar. The units of the CRB is in millimeters (mm). The diameter of the barrel in the scene is approximately 500 mm. The black regions have infinite CRB; these are the points not seen by any scanner.

problem. Figure 5.2 illustrates the CRBs computed for reconstructions of an office scene. Twelve range images were taken and registered with the methods described in [14]. Then using a level set representation, we reconstruct a surface model [5]. In the first reconstruction, we use only 4 out of the 12 range images. The occlusion shadows of the barrel and the chair are observed as the black regions on the reconstructed surface in Figure 5.2(a). Very high CRB values (purple) are also observed at various locations including the top of the desk, and on the bookshelves due to the occlusions of objects placed on it. Unlike the occlusion shadows of the chair and the barrel, these artifacts are not immediately observable from the reconstructed surface. Hence, the CRB image brings out useful information that can be used to choose further scanning locations. In the second reconstruction, we use all 12 range images. Overall, the average CRB is lower as expected and there are much fewer occluded regions. However, notice that certain parts of the desk and bookshelves still have infinite CRB values (black), indicating that these parts are occluded in all 12 range images. This result can be used to add another range image from a scanner location that can see these parts. Or alternatively, it can inform users (or some subsequent processing) not to trust the surface estimate in these locations.

Chapter 6

Conclusion

This paper shows the derivation of a systematic error measure for nonparametric surface reconstruction that uses the Cramer-Rao bound. The CRB is a tight lower error bound for unbiased estimators such as the maximum likelihood. However, there are some limitations in this formulation. We have assumed no knowledge of surface shape other than that given by the measurements. However, in practice shape reconstruction often includes some *a-priori* knowledge about surface shape, such as smoothness. The inclusion of such priors corresponds to a maximum posteriori estimation process. The current formulation still gives meaningful results—it tells us to what extent a particular estimate is warranted by the data. That is, it gives us some idea of the relative weighting of the data and the prior *at each point* on the surface. Future work will include a study of how to incorporate priors and estimator bias into these error bounds.

Acknowledgements

This work is supported by the Office of Naval Research under grant #N00014-01-10033 and the National Science Foundation under grant #CCR0092065.

Bibliography

[1] L. Nyland, A. Lastra, D. McAllister, V. Popescu, and C. McCue, "Capturing,Processing and Rendering Real-World Scenes", Videometrics and Optical Methods for 3D Shape Measurement, Electronic Imaging 2001.

[2] G. Turk and M. Levoy, "Zippered polygon meshes from range images", Proc. SIGGRAPH'94, pp. 311–318, 1994.

[3] B. Curless and M. Levoy, "A volumetric method for building complex models from range images", Proc. SIGGRAPH'96, pp. 303–312, 1996.

[4] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface Reconstruction from Unorganized Points", Computer Graphics, 26(2), pp. 71–78, 1992.

[5] R. T. Whitaker, "A Level-Set Approach to 3D Reconstruction From Range Data", IJCV, 29(3), pp. 203–231, 1998.

[6] D. J. Rossi and A. S. Willsky, "Reconstruction from projections based on detection and estimation of objects - Part I: Performance Analysis", IEEE Trans. Acoustic Speech and Signal Processing, pp. 886–897, 1984.

[7] A. O. Hero, R. Piramuthu, J. A. Fessler, and S. R. Titus, "Minimax emission computed tomography using high resolution anatomical side information and B-spline models", IEEE Trans. Information Theory, pp. 920–938, 1999.

[8] J. C. Ye, Y. Bresler and P. Moulin, "Cramer-Rao bounds for 2D target shape estimation in nonlinear inverse scattering problems with applications to passive radar", IEEE Trans. Antennas and Propagation, pp. 771–783, 2001.

[9] J. C. Ye, Y. Bresler and P. Moulin, "Asymptotic Global Confidence Regions in Parametric Shape Estimation Problems", IEEE Trans. Information Theory, 46(5), pp. 1881–1895, 2000.

[10] K. M. Hanson, G. S. Gunningham, and R. J. McKee, "Uncertainty assesement for reconstructions based on deformable geometry", I. J. Imaging Systems & Technology, pp. 506–512, 1997.

[11] J. Gregor and R. Whitaker, "Reconstructing Indoor Scene Models From Sets of Noisy Range Images", Graphical Models, 63(5), pp. 304–332, 2002.

[12] N. E. Nahi, “Estimation Theory and Applications”, John Wiley & Sons Inc., 1969.

[13] S. Osher and J. Sethian, ”Fronts Propogating with Curvature-Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations”, J. Comp. Physics, 79, pp. 12–49, 1988.

[14] R. Whitaker and J. Gregor, “A Maximum Likelihood Surface Estimator For Dense Range Data”, IEEE Trans. Pattern Analysis and Machine Intelligence, 24(10), 2002.