PERFORMANCE BOUNDS ON IMAGE REGISTRATION

İmam Şamil Yetik and Arye Nehorai

Dept. of ECE, Un. of Illinois at Chicago, USA

ABSTRACT

Although a vast number of publications have appeared on image registration, performance analysis is usually performed visually and little attention has been given to statistical performance bounds. Such bounds can be useful in evaluating image registration techniques, determining parameter regions where more successful registration is possible, and choosing features to be used for the registration. We derive Cramér-Rao bounds for a wide variety of geometric deformation models including translation, rotation, shearing, rigid, more general affine and non-linear transformations, and for intensity distortion and geometric deformation simultaneously. Numerical examples for illustrating the analytical performance bounds are presented.

1. INTRODUCTION

Image registration is the process of matching two images that differ in certain aspects but essentially represent the same object. The images to be registered may be obtained using different viewpoints, sensors, or time instants. Many engineering problems involve such cases, including motion detection where images taken at different time instants need to be registered, target recognition where images taken from different viewpoints should be combined, video processing where different frames should be registered for more efficient compression. Other application areas are 3-D modeling where 2-D images must be integrated to construct a 3-D model; medical imaging where combining information from different modalities is useful since each of them may have certain advantages, as well as registering images taken at different time instants to analyze tumor growth or regular child growth. In some of these applications the deformation itself is of interest, for example in motion estimation or tumor growth, whereas in others deformation analysis is required only to correct or align the images, such as combining medical images obtained using different modalities.

There are excellent tutorials on image registration including [1] focusing on geometric registration, [2] discussing intensity matching, and [3], which presents a well-organized classification of image registration algorithms.

In this paper we formulate the image registration problem as a statistical parameter estimation problem and derive Cramér-Rao bounds (CRB's) as performance measures. The CRB is a lower bound on the covariance of any unbiased estimator and is asymptotically achieved by the maximum likelihood (ML) estimator. It is an important benchmark performance measure that can be used to evaluate the efficiency of registration algorithms, to determine the parameter regions where good and poor estimates are expected, and to optimize image registration by selecting the features to be used. The CRB is widely studied in statistical signal processing areas, such as communications, radar, sonar, and biomedicine. Image registration literature lacks the study of this important bound except for a very limited deformation model (only translation) in the context of motion estimation [4]. This paper aims to fill this gap by deriving CRB expressions for a wide class of geometric deformation models including translation, rotation, scaling, shearing, rigid, more general affine and non-linear transformations shown in Figure 1, and for the case where the intensity distortion and geometric deformation occur simultaneously. In Section 2 we present

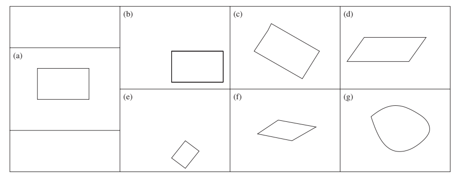

Fig. 1. Illustrations of geometric deformations. Part (a) shows the original image, (b) translation, (c) rotation, (d) skew, (e) rigid, (f) affine, and (g) non-linear transformations.

the image registration problem as a parameter estimation problem and give two basic frameworks: registration using (i) isolated points in Section 2.1 and (ii) the images as the features to be used in registration in Section 2.2 for geometric registration, and 2.3 for simulateneous intensity and geometric registration. We refer to these isolated points as simply "points" in the rest of the paper. We also derive CRB's using these two frameworks in Section 2. The maximum likelihood estimates (MLE's), their variances, and extensions to unknown real-world objects are given for some of the cases in [5]. We give numerical examples in Section 3 for easier visualisation of the analytically derived bounds. Section 4 has conclusion and discussion.

2. PROBLEM FORMULATION AND PERFORMANCE ANALYSIS

Let p be a real-world object of interest, and f and g two functions that represent the images to be registered of this object. The coordinates x', y', z' of the image f are geometrically deformed versions of the coordinates x, y, z of the image g:

$$(x', y', z') = (D_x(x, y, z), D_y(x, y, z), D_z(x, y, z)), \tag{2.1}$$

where D_x, D_y, D_z represents operators that map the coordinates (i.e. operate on coordinates and result in transformed coordinates). The intensity of the image f is also distorted resulting in the image g. Let H represent an operator which alters the intensity values of the image, then the overall relationship between the images f and g is

$$g(x, y, z) = H\{f[D_x(x, y, z), D_y(x, y, z), D_z(x, y, z)]\}, \tag{2.2}$$

Here, H operates on the geometrically distorted image changing its intensity values and resulting in the other image. The image registration problem is estimating D_x, D_y, D_z, and H according to some criterion using parametric models for the unknown geometric deformation and intensity distortion between f and g.

2.1. Registration Using Isolated Points

Many registration algorithms use isolated *points* to find the geometric distortion [6], [7] and then determine the optimum intensity matching using geometrically aligned images. Let $\boldsymbol{p} = [p_1, p_2, \ldots, p_l]^{\mathrm{T}}$ represent the l points corresponding to l features of p, $\boldsymbol{f} = [f_1, g_2, \ldots, f_l]^{\mathrm{T}}$

This work was supported by the Air Force Office of Scientific Research Grant F49620-02-1-0339, the National Science Foundation Grants CCR-0105334 and CCR-0330342.

be the corresponding isolated points of f, and $\boldsymbol{g} = [g_1, g_2, \ldots, g_l]^{\mathrm{T}}$ the isolated points of g representing the same l features. In this framework $\boldsymbol{f}$ and $\boldsymbol{g}$ are the observations (hence random functions due to measurement noise) and $\boldsymbol{p}$ is the real-world coordinates set (assumed to be deterministic.) Assuming a global geometric deformation (i.e. the same deformation for all isolated points), the statistical model is

$$\begin{bmatrix} \mathcal{F}_{3\times l} \\ \mathcal{G}_{3\times l} \end{bmatrix} = \begin{bmatrix} I_{3\times 3} \\ D_{3\times 3} \end{bmatrix} \mathcal{P}_{3\times l} + \begin{bmatrix} \mathbf{0}_{3\times l} \\ \boldsymbol{t}_{3\times 1} \otimes \mathbf{1}_{1\times l} \end{bmatrix} + N, \tag{2.3}$$

$$\mathcal{F} = \begin{bmatrix} f_{1x} & & f_{lx} \\ f_{1y} & \cdots & f_{ly} \\ f_{1z} & & f_{lz} \end{bmatrix}, \mathcal{G} = \begin{bmatrix} g_{1x} & & g_{lx} \\ g_{1y} & \cdots & g_{ly} \\ g_{1z} & & g_{lz} \end{bmatrix}, \mathcal{P} = \begin{bmatrix} p_{1x} & & p_{lx} \\ p_{1y} & \cdots & p_{ly} \\ p_{1z} & & p_{lz} \end{bmatrix},$$

$$\boldsymbol{t} = [t_x, t_y, t_z]^{\mathrm{T}}, \tag{2.4}$$

where the subscripts of the block matrices denote their dimensions, the subscripts "$_{ij}$" of the points denote the jth component (x, y or z) of the ith point, I is the identity matrix, $\mathbf{0}$ a matrix with all zero entries, $\mathbf{1}$ a matrix with all one entries, "$\otimes$" the Kronecker product, D the geometric deformation, t_x, t_y, t_z the translation, and finally N is a $6 \times l$ matrix denoting additive noise. This noise is assumed to be white Gaussian and independent for different points and directions. It is allowed to have different variances along different directions with covariance. Let $\boldsymbol{\theta}$ be the unknown parameters that define the geometric deformation and intensity distortion and $P_{\theta}(f,g)$ the probability density function of the random functions f and g given $\boldsymbol{\theta}$. The CRB matrix $\mathrm{CRB}(\boldsymbol{\theta})$ for the unknown parameters is the inverse of the Fisher information matrix (FIM) denoted by $J(\boldsymbol{\theta})$ [8]:

$$\mathrm{CRB}(\boldsymbol{\theta}) = J^{-1}(\boldsymbol{\theta}), \tag{2.5}$$

$$J_{ij}(\boldsymbol{\theta}) = -\mathrm{E}\left[\frac{\partial^2 \log P_{\theta}(f,g)}{\partial\theta_i\,\partial\theta_j}\right], \tag{2.6}$$

where $J_{ij}(\boldsymbol{\theta})$ denotes the ijth entry of $J(\boldsymbol{\theta})$, θ_i the ith component of $\boldsymbol{\theta}$, and $\mathrm{E}[\cdot]$ denotes the expectation operator. We derive expressions for the FIM and CRB matrices for a wide class of geometric deformations. All of the derivations follow three main steps: (i) finding $\log P_{\theta}(f,g)$, (ii) taking the partial derivatives, (iii) calculating the expectation.

We present the general form of $\log P_{\theta}(f,g)$ here, omit the special forms and also refer the reader to [5] for the resulting FIM and CRB matrices. Considering the model in (2.3) we have

$$\begin{aligned} \log P_{\theta}(f,g) = & \sum_{i=1}^{l}\left[\frac{(f_{ix}-p_{ix})^2}{-2\sigma_x^2} + \frac{(f_{iy}-p_{iy})^2}{-2\sigma_y^2} + \frac{(f_{iz}-p_{iz})^2}{-2\sigma_z^2}\right] \\ & + \sum_{i=1}^{l}\left[\frac{(g_{ix}-s_{ix})^2}{-2\sigma_x^2} + \frac{(g_{iy}-s_{iy})^2}{-2\sigma_y^2} + \frac{(g_{iz}-s_{iz})^2}{-2\sigma_z^2}\right] \\ & + \text{const.}, \end{aligned} \tag{2.7}$$

where

$$\begin{bmatrix} s_{ix} \\ s_{iy} \\ s_{iz} \end{bmatrix} = D \begin{bmatrix} p_{ix} \\ p_{iy} \\ p_{iz} \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix}, \quad i = 1, \ldots, l. \tag{2.8}$$

Note that the first and last terms in (2.7) do not involve deformation parameters; hence, they become zero in computing the CRB because of the derivatives with respect to the parameters in step (ii). Taking the derivatives and expectation gives the ijth element of FIM

$$J_{ij} = \frac{1}{\sigma_x^2}\sum_{i'=1}^{l}\frac{\partial s_{i'x}}{\partial\theta_i}\frac{\partial s_{i'x}}{\partial\theta_j} + \frac{1}{\sigma_y^2}\sum_{i'=1}^{l}\frac{\partial s_{i'y}}{\partial\theta_i}\frac{\partial s_{i'y}}{\partial\theta_j} + \frac{1}{\sigma_z^2}\sum_{i'=1}^{l}\frac{\partial s_{i'z}}{\partial\theta_i}\frac{\partial s_{i'z}}{\partial\theta_j}. \tag{2.9}$$

We use equation (2.9) to calculate the FIM's and CRB's for a wide class of geometric deformations using the described framework. The MLE's for some of the deformation parameters and expressions of their variances are also given in [5]. We also extend to unknown $\boldsymbol{p}$ and observe how the bounds change for some of the cases.

2.2. Geometric Registration Using Images

In this case no pre-processing is made on the images to extract the features; hence, we do not know the corresponding points. We also note that the two images will not have the same domain as explained in detail in [2]. Here, we will consider the overlapping domain sets of the two images. It is not possible to obtain explicit expressions for the MLE's since the unknown parameters are hidden in the coordinates and a search is necessary to find the maximum of the likelihood function.

We assume that the intensity values are not altered except for the additive noise. The statistical model we use is

$$f(x,y,z) = p(x,y,z) + n(x,y,z), \tag{2.10}$$

$$g(x,y,z) = p(x',y',z') + n(x,y,z), \tag{2.11}$$

where x', y', and z' are the transformed coordinates and $n(x,y,z)$ is the additive noise. The unknown parameters are hidden in x', y', and z':

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = D \begin{bmatrix} x \\ y \\ z \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix}. \tag{2.12}$$

Using the short-hand notation for the image coordinates $\boldsymbol{r} = (x,y,z)$ and $\boldsymbol{r}' = (x',y',z')$ our observations are now $f(\boldsymbol{r})$, $g(\boldsymbol{r})$ and the term $\log P_{\theta}(f,g)$ is

$$\log P_{\theta}(f,g) = \sum_{\boldsymbol{r}} \frac{[f(\boldsymbol{r}) - p(\boldsymbol{r})]^2 + [g(\boldsymbol{r}) - p(\boldsymbol{r}')]^2}{-2\sigma^2} + \text{const.} \tag{2.13}$$

Here, the summation $\sum_{\boldsymbol{r}}$ is over all the coordinates of the overlapping domain of the two images. The first term does not depend on the unknown parameters and becomes zero during the derivative operation when calculating the FIM. We need to take the derivative of a function of three variables ($\boldsymbol{r}'$) which are all functions of the elements of $\boldsymbol{\theta}$. Hence, we use the chain rule

$$\frac{\partial p(\boldsymbol{r}')}{\partial\theta_i} = \frac{\partial p(\boldsymbol{r}')}{\partial x'}\frac{\partial x'}{\partial\theta_i} + \frac{\partial p(\boldsymbol{r}')}{\partial y'}\frac{\partial y'}{\partial\theta_i} + \frac{\partial p(\boldsymbol{r}')}{\partial z'}\frac{\partial z'}{\partial\theta_i}. \tag{2.14}$$

Computing the derivatives using (2.14) and then taking the expectations gives the ijth element of the FIM for 3-D:

$$\begin{aligned} J_{ij} = & \mathrm{E}\left[\frac{\partial^2 \log P_{\theta}(f,g)}{\partial\theta_i\partial\theta_j}\right] = \frac{-1}{\sigma^2}\sum_{\boldsymbol{r}}\Big[p_x^2(\boldsymbol{r}')\frac{\partial x'}{\partial\theta_i}\frac{\partial x'}{\partial\theta_j} \\ & + p_x(\boldsymbol{r}')p_y(\boldsymbol{r}')\frac{\partial x'}{\partial\theta_i}\frac{\partial y'}{\partial\theta_j} + p_x(\boldsymbol{r}')p_z(\boldsymbol{r}')\frac{\partial x'}{\partial\theta_i}\frac{\partial z'}{\partial\theta_j} \\ & + p_y^2(\boldsymbol{r}')\frac{\partial y'}{\partial\theta_i}\frac{\partial y'}{\partial\theta_j} + p_x(\boldsymbol{r}')p_y(\boldsymbol{r}')\frac{\partial y'}{\partial\theta_i}\frac{\partial x'}{\partial\theta_j} + p_y(\boldsymbol{r}')p_z(\boldsymbol{r}')\frac{\partial y'}{\partial\theta_i}\frac{\partial z'}{\partial\theta_j} \\ & + p_z^2(\boldsymbol{r}')\frac{\partial z'}{\partial\theta_i}\frac{\partial z'}{\partial\theta_j} + p_y(\boldsymbol{r}')p_z(\boldsymbol{r}')\frac{\partial z'}{\partial\theta_i}\frac{\partial y'}{\partial\theta_j} + p_x(\boldsymbol{r}')p_z(\boldsymbol{r}')\frac{\partial z'}{\partial\theta_i}\frac{\partial x'}{\partial\theta_j}\Big]. \end{aligned} \tag{2.15}$$

where the subscripts "x, y, z" denote partial derivatives with respect to x, y, z respectively. These derivatives of p can be calculated by first interpolating the discrete image and then using the resulting interpolated continuous function for direct derivative calculation. It is also possible to approximate the derivatives using the difference function. The resulting FIM and CRB matrices for translation, rotation, rigid, skew, affine and bi-variate polynomial transformations can be found in [5].

2.3. Simultaneous Intensity and Geometric Registration

Consider two images where one of them is both intensity distorted and geometrically deformed version of the other. The model becomes

$$f(\boldsymbol{r}) = p(\boldsymbol{r}) + n(\boldsymbol{r}), \tag{2.16}$$

$$g(\boldsymbol{r}) \quad = \quad \sum_{\boldsymbol{r}_i \epsilon S_h} H(\boldsymbol{r}_i) p((\boldsymbol{r}-\boldsymbol{r}_i)') + n(\boldsymbol{r}), \qquad (2.17)$$

where $(\boldsymbol{r}-\boldsymbol{r}_i)'$ denotes the geometrically transformed version of $(\boldsymbol{r}-\boldsymbol{r}_i)$:

$$\begin{bmatrix} (x-i_x)' \\ (y-i_y)' \\ (z-i_z)' \end{bmatrix} = D \begin{bmatrix} x-i_x \\ y-i_y \\ z-i_z \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \\ t_y \end{bmatrix}. \qquad (2.18)$$

The function $\log P_\theta(f,g)$ becomes

$$\log P_\theta(f,g) \quad = \quad \frac{-1}{2\sigma^2}\left\{\sum_{\boldsymbol{r}}[f(\boldsymbol{r})-p(\boldsymbol{r}]^2 \right.$$
$$\left. + \sum_{\boldsymbol{r}}[g(\boldsymbol{r}) - \sum_{\boldsymbol{r}_i \epsilon S_h} H(\boldsymbol{r}_i)p((\boldsymbol{r}-\boldsymbol{r}_i)')]^2\right\} + \text{ const.} \qquad (2.19)$$

We calculate expressions of the FIM components in three groups: (i) components of geometrical deformation parameters, (ii) components of intensity distortion parameters, (iii) cross-components between geometrical deformation and intensity distortion parameters.
(i) Components of geometrical deformation parameters: These will have the same form as in the geometric deformations (Section 2.2) except $p(\boldsymbol{r}')$ will be replaced by

$$\sum_{\boldsymbol{r}_i \epsilon S_h} H(\boldsymbol{r}_i)p((\boldsymbol{r}-\boldsymbol{r}_i)').$$

The FIM expressions will then be similar to those of the geometric deformations, except the image will be filtered by H before the derivative operations. This result implies that the estimation accuracy of the geometric distortion parameters suffers from smaller values of the filter coefficients. This is eventually a result of decreased SNR.
(ii) Components of intensity distortion parameters: The FIM components of the intensity distortion parameters are

$$J_{h_a,h_b} = \frac{1}{\sigma^2}\sum_{\boldsymbol{r}} p((\boldsymbol{r}-\boldsymbol{r}_a)')p((\boldsymbol{r}-\boldsymbol{r}_b)'). \qquad (2.20)$$

This FIM shows that more correlated image pixels result in better estimates of the filter coefficients.
(iii) Cross-components between geometrical deformation and intensity distortion parameters: In this case we will have the first derivative with respect to the geometrical deformation parameters and the second derivative with respect to the intensity distortion parameters; hence, the FIM expressions will not be similar to any of the previous cases. These can be obtained using equation (2.5).

3. NUMERICAL EXAMPLES

We use basic examples to calculate and plot the CRB's and MLE variances that are analytically derived in Section 2 for better visualisation. These plots illustrate the analytical results of Section 2.

3.1. Registration Using Isolated Points - Geometric Deformations

We use isolated points that are uniformly distributed on the image in both directions (x and y) and several parameters are varied as explained in the following.
2-D Rotation: There is only one unknown parameter α resulting in a scalar CRB. We plot this CRB as a function of the rotation angle for the case of equal and different variances along x and y directions. Figure 2 shows the CRB on the rotation angle for $\sigma_x^2 = 1, \sigma_y^2 = 3$; $\sigma_x^2 = 1, \sigma_y^2 = 2$; and $\sigma_x^2 = 1, \sigma_y^2 = 1$. In this figure we have used points that form a 5×5 grid resulting in twenty-five points. Observe that the CRB approaches a constant as the variances along x and y directions become closer. We also plot the variance of the MLE as a function of number of points and compare it with the CRB in Figure 3(d) for $\alpha = 60$deg. The MLE variance asymptotically achieves the CRB as expected. 2-D Affine Transformations: The unknown parameters are $\boldsymbol{\theta} = [d_1, \ldots, d_6]^{\mathrm{T}}$. We have 12 CRB components; however, none of these depends on the deformation parameters. Therefore, we plot them as a function of the number of points. We present the plots for only d_1, d_2, and d_3 since the others have similar forms. Figures 3a-c show the diagonal CRB components and corresponding MLE variances, and part (d) shows three cross components $\mathrm{CRB}_{d_1,d_2}, \mathrm{CRB}_{d_1,d_3}, \mathrm{CRB}_{d_2,d_3}$. Parts (a)-(c) show how the MLE variances approach CRB values as the number of points gets larger. Observe also that the CRB components for d_3 (which is the translation parameter) are smaller than those of d_1 and d_2.

3.2. Registration Using Images - Geometric

We use two images where one of them has an image derivative with larger intensity values so that we can observe the effect of the image derivatives on the CRB's (see Section 2.2.) We approximate the image derivatives using the difference between the pixels. In Figure 4a, we plot $\mathrm{CRB}_{\alpha,\alpha}$, $\mathrm{CRB}_{s,s}$, CRB_{t_x,t_x} using a 2-D rigid transformation. The first row shows the two images used, second row the CRB on the rotation angle, third row for the scaling parameter, and fourth row for the translation parameter along x direction. First column shows the bounds for the first image and second column for the other image. We observe that the CRB values for the first image are larger. This can be explained partly by the fact that the first image has smaller derivative values compared with the second image (the second image changes more rapidly).

3.3. Registration Using Images - Simultaneous

We use a low-pass filter to model the intensity distortion and observe the effect of simultaneous intensity and geometric registration on the CRB on the geometric deformation parameters (see Section 2.3 for details.) The low-pass filter we choose is

$$H = \frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}. \qquad (3.21)$$

We plot the same set of CRB components as in Section 3.2 in Figure 4b. Observe that the simultaneous estimation of the filter coefficients increases the CRB values for the geometric deformation parameters.

4. CONCLUSION

We have derived statistical performance bounds for image registration algorithms. For some of the cases we have derived also the expressions for the MLE's and their finite-sample variances. We have considered a wide class of geometric deformations, intensity distortion using an MA model, as well as simultaneous geometric deformation and intensity distortion. We summarize the results below, see [5] for a more detailed discussion.

For registration using isolated points, The bounds on the variances of the translation parameter estimates (when the deformation model is translation only) depend only on the number of points but not on the amount of translation of the locations of the points. The CRB's are twice as large when we include the estimation of the locations of the points. The variances of the MLE's are equal to the CRB values for any number of points. Better estimates can be obtained for the rotation angle (if the deformation model is rotation only) when the points are further from the rotation center. The bounds are independent of the rotation amount for the case of equal noise variances for the

x and y directions. For rigid transformations (scaling, rotation, and translation), better estimates are obtained for the rotation angle when points further from the rotation center are used for the case of equal noise variances for all directions. Larger scaling values degrade the performance of the estimation of the rotation angle. The bound on the variance of the scaling parameter is independent of the amount of scaling. For 2-D skew, the bounds for the shear parameters along the two directions are decoupled. The bounds depend only on the number and locations of the points used but not on the shear amount. The MLE variances are equal to the CRB components. The bounds for the parameters of affine and bi-variate polynomial transformations (2-D and 3-D) are independent of the parameter values, but depend only on the points and noise variances.

For registration using images, the bounds for all geometric deformation parameters depend on the values of the deformation parameters in contrast to some of the cases of registration using isolated points. For the case of simultaneous registration considering both geometric deformation and intensity distortion, the bounds for the estimates of the geometric deformation parameters suffer from smaller filter coefficients. This is related to decreased SNR as a result of smaller filter coefficients.

As a future work it is possible to compare the CRB's to the performances of standard methods, such as information-based or intensity-similarity-based methods. It would also be useful to derive the confidence intervals for easier visualization. Another direction is to consider the registration of video frames, and exploit the fact that the deformation (e.g. motion of a certain object in the video) is correlated in time. In this case, iterative estimation can be applied and the performance analysis for this case is of interest. It would also be interesting to search for distance metrics other than mean-squared error. For instance, the error for the rotation angle may be represented using covariance of vector angular error [9]. The performance bounds will then be on these metrics rather than conventional mean-squared error.

5. REFERENCES

[1] L. G. Brown, "A survey of image registration techniques," *AMC Computing Surveys*, vol. 24, pp. 325–376, 1992.

[2] D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, "Medical image registration," *Physics in Medicine and Biology*, vol. 46, pp. R41–R45, 2001.

[3] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, pp. 1–36, 1998.

[4] D. Robinson and P. Milanfar, "Fundamental performance limits in image registration," *IEEE Trans. Image Proc.*, submitted.

[5] I. S. Yetik and A. Nehorai, "Performance bounds on image registration," *IEEE Transactions on Signal Processing*, submitted.

[6] J. Fitzpatrick, J. West, and C. M. Jr, "Predicting error in rigid-body, point-based registration," *IEEE Trans. Medical Imaging*, vol. 17, pp. 694–702, 1998.

[7] D. L. G. Hill, D. J. Hawkes, M. J. G. J. E.C̃rossman, T. C. S. Cox, A. J. S. E. E. C. M. L.B̃racey, and P. Graves, "Registration of mr and ct images for skull base surgery using point-like anatomical features," *Br. J. Radiol.*, vol. 64, pp. 1030–1035, 1991.

[8] V. Poor, *An introduction to signal detection and estimation*, 2nd ed. Springer, 1994.

[9] A. Nehorai and E. Paldi, "Vector-sensor array processing for electromagnetic source localization," *IEEE Trans. Sig. Proc.*, vol. 42, pp. 376–398, 1994.

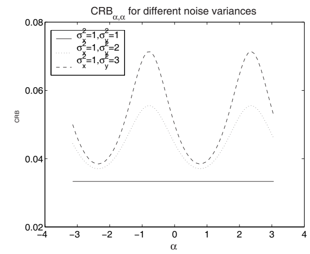

Fig. 2. $\mathrm{CRB}_{\alpha,\alpha}$ as a function of the rotation angle α assuming the geometric deformation is only rotation. The dashed line shows the CRB for $\sigma_x^2 = 1, \sigma_y^2 = 3$, dotted line for $\sigma_x^2 = 1, \sigma_y^2 = 2$, and solid line for $\sigma_x^2 = 1, \sigma_y^2 = 1$. The CRB approaches a constant as the variances along x and y directions become closer.

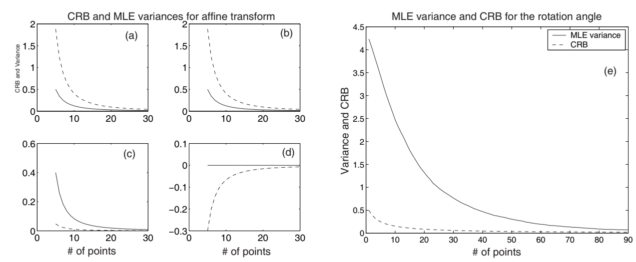

Fig. 3. Parts (a)-(c) compare MLE variances (dashed line) with CRB's (solid line) for the affine transformation parameters as a function of number of points; (a) is for the parameter d_1, (b) d_2, and (c) d_3. Part (d) shows CRB_{d_1,d_2} (solid line), CRB_{d_1,d_3} (dashed line), and CRB_{d_2,d_3} (dotted line). Part (e) shows MLE variance (solid line) and CRB on the rotation angle (dashed line) as a function of number of registration points assuming the geometric deformation is only rotation for $\alpha = \pi/3$. The MLE achieves the CRB as the number of points gets larger.

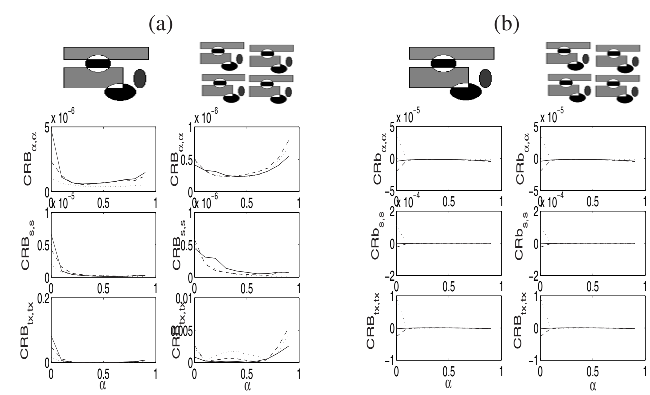

Fig. 4. (a) CRB components for two different images assuming rigid transformations and no intensity distortion. The first row shows the two images used, second row $\mathrm{CRB}_{\alpha,\alpha}$, third row $\mathrm{CRB}_{s,s}$, and last row CRB_{t_x,t_x} all as a function of α. The plots in the first column correspond to the first image on the left and the second column to the second image on the right. (b) same as part (a) except there is a low-pass intensity distortion.