HHS Public Access

Author manuscript

Proc SPIE Int Soc Opt Eng. Author manuscript; available in PMC 2017 May 30.

Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2017 February 11; 10135: . doi:10.1117/12.2256025.

Fundamental limits of image registration performance: Effects of image noise and resolution in CT-guided interventions

M. D. Ketcha[a], **T. de Silva**[a], **R. Han**[a], **A. Uneri**[b], **J. Goerres**[a], **M. Jacobson**[a], **S. Vogt**[d], **G. Kleinszig**[d], and **J. H. Siewerdsen**[*,a,b,c]

[a]Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD

[b]Department of Computer Science, Johns Hopkins University, Baltimore, MD

[c]Department of Neurosurgery, Johns Hopkins University, Baltimore, MD

[d]Siemens Healthcare XP Division, Erlangen, Germany

Abstract

Purpose—In image-guided procedures, image acquisition is often performed primarily for the task of geometrically registering information from another image dataset, rather than detection / visualization of a particular feature. While the ability to detect a particular feature in an image has been studied extensively with respect to image quality characteristics (noise, resolution) and is an ongoing, active area of research, comparatively little has been accomplished to relate such image quality characteristics to registration performance.

Methods—To establish such a framework, we derived Cramer-Rao lower bounds (CRLB) for registration accuracy, revealing the underlying dependencies on image variance and gradient strength. The CRLB was analyzed as a function of image quality factors (in particular, dose) for various similarity metrics and compared to registration accuracy using CT images of an anthropomorphic head phantom at various simulated dose levels. Performance was evaluated in terms of root mean square error (RMSE) of the registration parameters.

Results—Analysis of the CRLB shows two primary dependencies: 1) noise variance (related to dose); and 2) sum of squared image gradients (related to spatial resolution and image content). Comparison of the measured RMSE to the CRLB showed that the best registration method, RMSE achieved the CRLB to within an efficiency factor of 0.21, and optimal estimators followed the predicted inverse proportionality between registration performance and radiation dose.

Conclusions—Analysis of the CRLB for image registration is an important step toward understanding and evaluating an intraoperative imaging system with respect to a registration task. While the CRLB is optimistic in absolute performance, it reveals a basis for relating the performance of registration estimators as a function of noise content and may be used to guide acquisition parameter selection (e.g., dose) for purposes of intraoperative registration.

[*] jeff.siewerdsen@jhu.edu; phone (443) 287-6269; fax (410) 955-9826; www.istar.jhu.edu.

Keywords

CT/CBCT; Fisher Information; Image Guided Procedures; Image Quality; Image Registration; Performance Limits

1. Description of Purpose

Many applications in the field of image-guided procedures use image-based registration to improve intraoperative navigation and guidance. Often, this registration is performed to map data (e.g., planning information) from a preoperative image to the coordinate system of an intraoperative image. The error associated with this registration process is an important factor that determines the overall accuracy of the image-guided procedure and clinical outcomes. The general understanding of the relationship between image quality and registration performance is such that “high-quality” images would likely result in improved registration performance. While such qualitative characterizations are often evaluated empirically, there is a potential benefit in developing an analytical form to understand how changes in quantitative image quality metrics [such as the modulation transfer function (MTF) and noise-power spectrum (NPS)] will affect registration performance.

In the case of point-based registration, a well-known analytical framework was developed by Fitzpatrick and West[1] to relate the error distribution in localizing the fiducial points to the distribution in target registration error (TRE), yielding a valuable method for understanding registration system performance and limitations. On the other hand, for medical imaging systems, analytical performance evaluation has primarily focused on relating image quality metrics to detection and visualization tasks[2,3], where metrics such as the detectability index provide a quantitative framework for understanding the effect of NPS and MTF on a feature identification task in the image. However, in image-guided procedures the task of interest is often image-based registration rather than visualization/detection. In this work, we developed an analytical framework to relate image signal and noise characteristics to registration performance.

A valuable step toward this goal was previously proposed by Robinson and Milanfar[4], who suggested analyzing the relationship between registration performance and noise by examining the Cramer-Rao Lower Bound (CRLB) for 2D translation-only image registration. They considered a registration model where the two images to be registered contained the same underlying true image function, $g(x, y)$; however, both are contaminated by additive Gaussian white noise (AGWN) of equal magnitude.

$$I_1(x,y)=g(x,y)+n_1(x,y) \quad (1)$$

$$I_2(x,y)=g(x-u,y-v)+n_2(x,y) \quad (2)$$

As can be seen from Eqs. (1,2), the registration problem lies in estimating the unknown translational shift between the true underlying image functions represented by parameter vector $\boldsymbol{\theta} = [u, v]^T$. Their formulation had also been extended by Yetik and Nehorai[5] for the case of rigid (rotation and scale) and 3D registration. To relate this work to image-guided interventions, the method requires modification to account for disparate image quality characteristics (e.g., noise) between the preoperative and intraoperative image and a connection to image quality models that describe how these characteristics relate to imaging parameters such as radiation dose, acquisition protocol, and reconstruction method. In this work, we investigate theoretical lower bounds for image registration and the relationship to dose and noise for intraoperative cone-beam computed tomography (CBCT). We build a framework to relate established models of CBCT image quality[6,7,8] to registration performance.

2. Methods

2.1. The Cramer-Rao lower bound in image registration

The CRLB is commonly examined for many parameter estimation problems, as it provides a straightforward approach to derive lower bound for the mean squared error for any estimator. In particular, for the unbiased estimator, the CRLB covariance matrix (C_{LB}) is simply the inverse of the Fisher Information Matrix (FIM), a metric derived from the log-likelihood function [log $L(\boldsymbol{I}|\boldsymbol{\theta})$], i.e., the log-likelihood of the data ($\boldsymbol{I}$) conditioned on the parameter vector $\boldsymbol{\theta}$. By definition[9], we have:

$$[\mathrm{FIM}]_{ij} = -E\left\{\frac{\partial^2 \log L(\boldsymbol{I}|\boldsymbol{\theta})}{\partial\theta_i \partial\theta_j}\right\} \quad (3)$$

For image registration, $\boldsymbol{I}$ is the image data, and $\boldsymbol{\theta}$ is the transformation parameter vector ($[u, v]^T$ in this work). By noting that the subtraction of the two images at the true $\boldsymbol{\theta}$ shift leaves only the subtraction of the noise terms (which is itself a zero-mean Gaussian process), we can write the likelihood function as a jointly Gaussian distribution of the image difference conditioned on $\boldsymbol{\theta}$.

2.1.1. Signal-known-exactly (SKE) with white noise—While presented as having noise in both images, the derivation in Ref. [4] applies to the simplified case of a noiseless I_1 [simplified from Eqs. (1-2)] giving the signal-known-exactly (SKE) scenario [i.e., exactly known $g(x,y)$] described by:

$$I_1(x,y) = g(x,y) \quad (4)$$

$$I_2(x,y) = g(x-u, y-v) + n_2(x,y) \quad (5)$$

in which only the shifted image is considered to be contaminated by AGWN with variance σ^2. Following a conceptually similar derivation presented by Kay[9] (which treated 1D time delay estimation), we define the log-likelihood function for Eqs. (4-5) to be:

$$\begin{aligned}\log L(\boldsymbol{I}|\boldsymbol{\theta}) &= \log\left(\prod_{x,y} c\exp\left(\frac{-(I_2(x-u,y-v)-I_1(x,y))^2}{2\sigma^2}\right)\right)\\ &= \sum_{x,y}\frac{-1}{2\sigma^2}[I_2(x-u,y-v)-g(x,y)]^2+\text{const.}\end{aligned} \tag{6}$$

By noting that the expectation of I_2 evaluated at $\boldsymbol{\theta}$ is g and evaluating the second derivatives, we can compute the *FIM* from Eq. (3) as:

$$\begin{aligned}[\text{FIM}]_{ij} &= -E\left\{\frac{\partial^2 \log L(\boldsymbol{I}|\boldsymbol{\theta})}{\partial\theta_i\partial\theta_j}\right\}\\ &= -E\left\{\frac{\partial}{\partial\theta_i}\left(\frac{-1}{2\sigma^2}\sum_{x,y}2[I_2(x-u,y-v)-g(x,y)]\left(\frac{\partial I_2(x-u,y-v)}{\partial\theta_i}\right)\right)\right\}\\ -E\left\{\frac{-1}{\sigma^2}\sum_{x,y}[I_2(x-u,y-v)-g(x,y)]\left(\frac{\partial^2 I_2(x-u,y-v)}{\partial\theta_j\partial\theta_i}\right)+\left(\frac{\partial I_2(x-u,y-v)}{\partial\theta_j}\right)\left(\frac{\partial I_2(x-u,y-v)}{\partial\theta_i}\right)\right\}\\ &= \frac{1}{\sigma^2}\sum_{x,y}\left(\frac{\partial g(x-u,y-v)}{\partial\theta_i}\right)\left(\frac{\partial g(x-u,y-v)}{\partial\theta_j}\right)\end{aligned} \tag{7}$$

where for the case of translation $\boldsymbol{\theta}$ [*u, v*], and $\frac{\partial g}{\partial\theta_i}$ can be shown to simply be the image derivative with respect to the translation direction, giving

$$\frac{\partial g(x-u,y-v)}{\partial u}=\left[\frac{-\partial g(m,n)}{\partial m}\right]_{\substack{m=x-u\\ n=y-v}} = -g_x(m,n)$$

, (and similarly for *v, y*), where $g_x(m, n)$ is the partial derivative image with respect to *x*. Therefore the FIM is:

$$\text{FIM}_{\text{SKE,AWGN}}=C_{LB}{}^{-1}=\frac{1}{\sigma^2}\begin{bmatrix}\sum_{m,n} g_x{}^2 & \sum_{m,n} g_x g_y\\ \sum_{m,n} g_x g_y & \sum_{m,n} g_y{}^2\end{bmatrix} \tag{8}$$

This agrees with the lower bound presented by Robinson and Milanfar[4] in the case that noise is only included in the moving image, I_2. The *FIM* suggests that registration performance (in this simple case) is dependent on two components: (1) the image noise magnitude; and (2) the sum of squared image gradients.

2.1.2. Noise-push approximation—Generally for image registration, noise is present in both images, thus the SKE case does not directly apply. In order to account for noise in both

images, as well as be consistent with the derivation of Eq. (8), we posit an analytical "push" of the noise from I_1 into I_2:

$$I_1(x,y)=g(x,y) \quad (9)$$

$$I_2(x,y)=g(x-u,y-v)+n_2(x,y)-n_1(x,y) \quad (10)$$

which for the AGWN case, provides us with the form given in Eqs. (4-5) of section 2.1.1. The FIM under such approximation is:

$$\mathrm{FIM}_{\mathrm{push,AWGN}}={C_{LB}}^{-1}=\frac{1}{{\sigma_1}^2+{\sigma_2}^2}\begin{bmatrix}\sum_{m,n}{g_x}^2 & \sum_{m,n}g_xg_y\\ \sum_{m,n}g_xg_y & \sum_{m,n}{g_y}^2\end{bmatrix} \quad (11)$$

Therefore, under the push approximation, the sum of the noise variance terms now enter the denominator of the FIM. We note then, that if the noise magnitudes between the images are equal, then denominator is increased by a factor of 2 compared to that of Eq. (8).

2.2. Application to intraoperative CBCT

From Eq. (11), we can see that the CRLB is dependent on both the image content (g) and the noise characteristics of the image (e.g., σ^2 in the AGWN case). For the setting of CBCT image registration, we can use established models of CBCT noise and resolution characteristics[6,7,8] to arrive at an approximation for the CRLB. By ignoring correlations in the noise, and only examining noise magnitude, we can compute σ^2 relating system and dose factors to the CRLB:

$$\sigma^2=\iint_{Nyq}\mathrm{NPS}_{15}(f_x,f_y)df_xdf_y=\frac{1}{m\bar{q}_0}\frac{1}{\mathrm{DQE}}\frac{1}{a_{xy}^3}\iint_{Nyq}\mathrm{MTF}^2(f_x,f_y)df_xdf_y \quad (12)$$

where dose is proportional the number of projections (m) and incident x-ray fluence ($\bar{q}_0$), and DQE is the detective quantum efficiency of the imaging system. Spatial resolution terms include the voxel size a_{xy} (and slice thickness, a_z, ignored in the current 2D analysis) and the "bandwidth integral" over system MTF squared. Considering the registration model of Eq. (11) and image quality model of Eq. (12) suggests an immediate result: since variance is inversely proportional to dose (via the $m\bar{q}_0$), and the CRLB (FIM^{-1}) is proportional to σ^2, then the lower bound on registration accuracy scales as (1/dose). When considering resolution effects, it is known that reduced MTF or increased voxel size may reduce the noise magnitude, however these effects will also lead to a reduction in the image gradient strength, suggesting that there may be an optimal resolution for registration performance.

2.3. Experimental methods

Experiments involved a digital simulation in which axial CT images of an anthropomorphic head were computed over a broad range in dose by forward-projecting a high-quality scan of the head (with soft tissue set to a constant value of 40 HU)[10], scaling the fluence down in proportion to dose, and adding Poisson noise in proportion to 1/sqrt((1+SPR)·dose) where the scatter to primary ratio (SPR)[11,12] was taken to be 9. Each reconstruction use $m = 720$ forward projections over 360° and the fluence was determined based on a total mAs at 100 kVp beam energy using Spektr[13]. A central 2D axial slice (390 × 485 at 0.5× 0.5 mm^2 voxel size) was extracted from the 3D filtered backprojection reconstructed image (shown in Fig. 1). Registration was performed among these images.

2.3.1. Methods and metrics—Empirical evaluation was performed using 3 methods of sub-pixel image registration. The three methods involved performing an optimization technique over the registration parameters where a cubic-spline interpolated image was computed to evaluate a similarity metric at each iteration. SimpleITK[14] was used to perform this L-BFGS optimization routine using 3 different similarity metrics: (1) Mean-squared difference (MSD), (2) Matte's mutual information[15] (MMI) (100 bins), and (3) joint-histogram mutual information[16] (JMI) (100 bins, 1.5 σ).

2.3.2. Registration of a High-Dose Image to a Low-Dose Image—To mimic the scenario of an image-guided procedure, we performed registration of a high-dose image to a low-dose image. Realizations of 100 mAs images were used as the high-dose images, where the lower-dose images had dose values ranging from 0.5 mAs to 100 mAs. A known shift of θ = [1.2 pix, 1.2 pix] was introduced in the high-dose image prior to registration, and registration was computed using the high-dose image as the moving image.

2.3.3 Registration Performance Evaluation—The root mean squared error (RMSE) of the estimated translation vector was used to evaluate performance. RMSE was examined as a function of dose and compared to the CRLB-Push estimate of Eq. (11). If we assume the registration estimators to be unbiased, we may set a lower bound for the RMSE as:

$$\mathrm{RMSE} \geq \sqrt{\mathrm{trace}(C_{LB})} = \sqrt{\mathrm{trace}(\mathrm{FIM}^{-1})} \quad (13)$$

The CRLB terms were computed by determining the mean image (to yield g) and NPS terms from a set of 20 instances of simulated noisy images. To compute g we took the mean over 20 realizations of the high-dose image. Noise power spectra were computed by subtracting out g from 20 realizations at each noise level and then computing the average periodogram among these images, from which $\sigma_1^2 + \sigma_2^2$ [the denominator in Eq. (11)] was computed by integrating $N_1 + N_2$.

The lower bound presented in Eq. (13) also leads to an interesting question of how efficient an empirical estimator performs with respect to this theoretical bound. Therefore we define the statistical registration efficiency (SRE):

$$\mathrm{SRE}=\frac{\mathrm{trace}(C_{LB})}{(\mathrm{RMSE})^2} \quad (14)$$

SRE provides a value (bounded between 0 and 1) describing the efficiency of an estimator, whereby an SRE of 1 implies that the estimator is achieving the lower bound performance. The efficiency of each registration method was examined as a function of dose.

3. Results and Breakthrough Work

3.1. Registration of a High-Dose Image to a Low-Dose Image

Figure 2 shows the registration performance among the 3 similarity metrics. Figure 2A depicts the RMSE performance as a function of the low-dose mAs value with a 100 mAs high-dose image. As expected, registration performance generally improved with higher dose for all metrics. While it is seen in this study that JHMI always performed comparatively poorly and MSD nearly always had the highest performance across all dose levels, we may also note that MMI has a dose-dependent optimality, with the observation of relatively improved performance in the low-dose region. Interestingly, the performance in the high-dose region for MSD and JHMI tends to follow the trend established by the CRLB (though the bound appears optimistic). However, in the low-dose range the trend with the CRLB falls off drastically for JHMI, indicating a threshold after which registration failures occur (i.e., a false optimum far from the main lobe of the true solution is found) caused by high noise levels that lead to arbitrarily large errors.

Figure 2B shows the same registration data as Fig. 2A now depicting SRE. Similarly, the high-dose performance generally follows the trend established by the CRLB as seen by the flat SRE performance for MSD and JHMI. The behavior of MMI is interesting in that the efficiency seems to improve as dose decreases. This is likely due to a small bias associated with the estimator (and possibly dependent on the number of bins) that leads to a performance limit shown by the plateau region of MMI in Fig. 2A. Further, we can compare the estimators in terms of SRE value at each dose, and also in terms of robustness by examining the dose-level at which the registration failure threshold occurs (e.g. around 1 mAs for JHMI).

4. Conclusion

Understanding how image quality metrics relate to the purpose of image acquisition is an important step for evaluating the performance of any imaging system. In the field of image-guided procedures, an image is often acquired with the intent of relating pre-operative and intra-operative coordinate frames; therefore, it is important to understand the image quality characteristics with respect to registration error. The framework presented in this paper provides a means to relate image quality characteristics such as the NPS and MTF to registration performance via the CRLB, thereby creating a framework from which we can begin optimizing a system with respect to considerations on both dose and expected registration error. The work further provides a means to compare registrations methods in

terms of SRE, examining both the stability under low-noise conditions and the robustness against a low-dose registration failure threshold. Future work includes extension of this model beyond the white noise assumption and accounting for the effect of noise in both images beyond the current noise-push approximation.

Acknowledgments

The research was supported by NIH grant R01-EB-017226 and research collaboration with Siemens Healthineers (XP and AT groups, Forcheim Germany). Thanks to Jennifer Xu (Biomedical Engineering, Johns Hopkins University, Baltimore MD) for discussion of image quality models and Michael Mow and Hao Dang (Biomedical Engineering, Johns Hopkins University, Baltimore MD) for assistance with CBCT noise simulation.

References

1. Fitzpatrick JM, West JB. The distribution of target registration error in rigid-body point-based registration. IEEE Trans Med Imaging. 2001; 20(9):917–927. [PubMed: 11585208]
2. Barrett HH, Denny JL, Wagner RF, Myers KJ. Objective assessment of image quality. II. Fisher information, Fourier crosstalk, and figures of merit for task performance. J Opt Soc Am. 1995; 12(5):834–852.
3. Vennart, W. ICRU Report 54: Medical imaging—the assessment of image quality. Radiography; Maryland, USA: 1997.
4. Robinson D, Milanfar P. Fundamental performance limits in image registration. IEEE Trans Image Process. 2004; 13(9):1185–1199. [PubMed: 15449581]
5. Yetik IS, Nehorai A. Performance bounds on image registration. IEEE Trans Signal Process. 2006; 54(5):1737–1749.
6. Tward DJ, Siewerdsen JH. Cascaded systems analysis of the 3D noise transfer characteristics of flat-panel cone-beam CT. Med Phys. 2008; 35(12):5510–5529. [PubMed: 19175110]
7. Siewerdsen JH, Jaffray DA. Unified iso-SNR approach to task-directed imaging in flat-panel cone-beam CT. Proc SPIE. 2002; 4682:245–254.
8. Gang GJ, Lee J, Stayman JW, Tward DJ, Zbijewski W, Prince JL, Siewerdsen JH. Analysis of Fourier-domain task-based detectability index in tomosynthesis and cone-beam CT in relation to human observer performance. Med Phys. 2011; 38(4):1754–1768. [PubMed: 21626910]
9. Kay, SM. Fundamentals of Statistical Signal Processing: Estimation Theory. Vol. 3. Prentice-Hall; Englewood Cliffs, Ch: 1993.
10. Dang H, Stayman JW, Xu J, Zbijewski W, Sisniega A, Mow M, Wang X, Foos DH, Aygun N, Koliatsos VE, Siewerdsen JH. Task-based regularization design for high-quality cone-beam CT of the head. Phys Med Bio. Submitted.
11. Siewerdsen JH, Jaffray DA. Optimization of x-ray imaging geometry (with specific application to flat-panel cone-beam computed tomography). Med Phys. 2000; 27(8):1903–1914. [PubMed: 10984236]
12. Siewerdsen JH, Jaffray DA. Cone-beam computed tomography with a flat-panel imager: magnitude and effects of x-ray scatter. Med Phys. 2001; 28(2):220–231. [PubMed: 11243347]
13. Punnoose J, Xu J, Sisniega A, Zbijewski W, Siewerdsen JH. Technical Note: spektr 3.0—A computational tool for x-ray spectrum modeling and analysis. Med Phys. 2016; 43(8):4711–4717. [PubMed: 27487888]
14. Lowekamp BC, Chen DT, Ibáñez L, Blezek D. The design of SimpleITK. Front Neuroinform. 2013; 7:45. [PubMed: 24416015]
15. Mattes D, Haynor DR, Vesselle H, Lewellyn TK, Eubank W. Nonrigid multimodality image registration. Proc SPIE. 2001; 4322:1609–1620.
16. Thévenaz P, Unser M. Optimization of mutual information for multiresolution image registration. IEEE Trans Image Process. 2000; 9(12):2083–2099. [PubMed: 18262946]

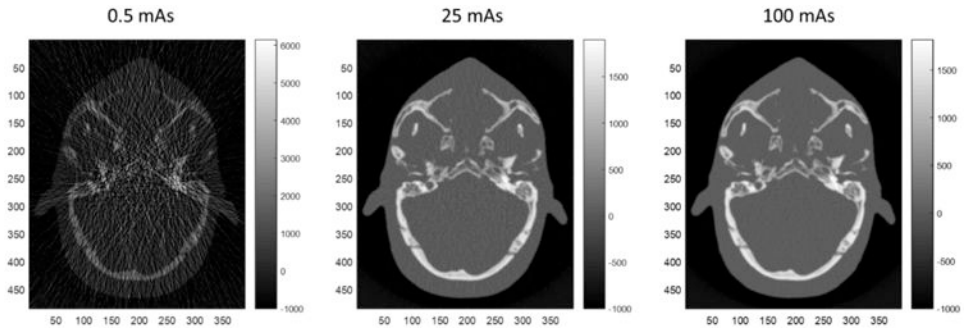

Figure 1.
Example images of the anthropomorphic head phantom at various dose values.

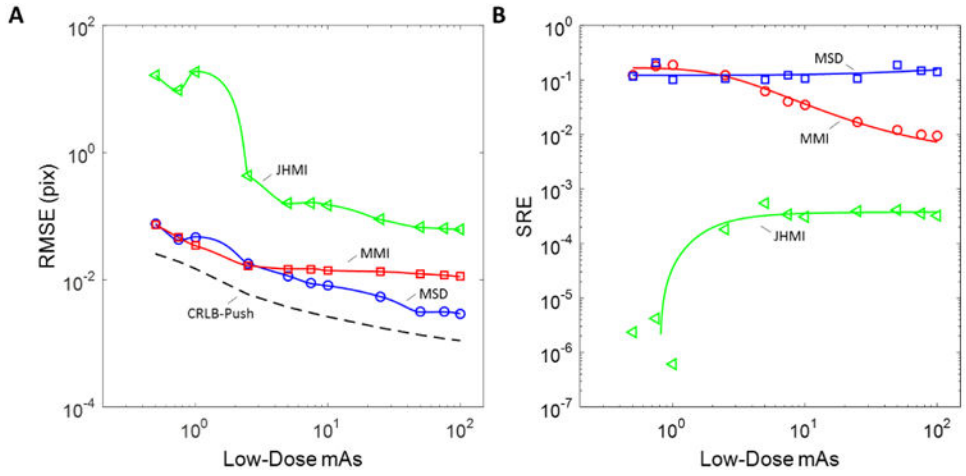

Figure 2.

Registration performance as a function of the low-dose image mAs value for a 100 mAs high-dose image. (A) Performance in terms of RMSE for three similarity metrics (MSD, MMI, and JHMI) with the CRLB-Push approximation of Eq. (11) shown as the dashed line. (B) The same experimental data plotted in terms of the SRE.