Kalman Filtering for Frame-by-Frame CT to Ultrasound Rigid Registration

Haydar Talib[1], Matthias Peterhans[1], Jaime García[1],
Martin Styner[2], and Miguel A. González Ballester[1]

[1] MEM Research Center, Institute for Surgical Technology and Biomechanics
University of Bern, Switzerland
[2] Departments of Computer Science and Psychiatry, Neurodevelopmental Disorders
Research Center, University of North Carolina at Chapel Hill, USA
{Haydar.Talib, Matthias.Peterhans}@MEMcenter.unibe.ch*

Abstract. This paper presents a method for CT-US rigid registration in minimally-invasive computer-assisted orthopaedic surgery, whereby the registration procedure is reformulated to enable effectively real-time registrations. A linear Kalman filter based algorithm is compared to an Unscented Kalman filter based method in simulated and experimental scenarios. The validation schemes demonstrate that the linear Kalman filter is more accurate, more robust, and converges quicker than the UKF, yielding an effectively real-time method for rigid registration applications, circumventing surgeons' waiting times.

1 Introduction

In computer-assisted orthopaedic surgery (CAOS), surgeons often benefit from enhanced visualization by registering a pre-operatively acquired medical image, such as from CT or MRI, to the patient's anatomy during surgery. Registration is usually achieved by digitizing bone surface points from the patient using a navigated pointer, and determining the optimal transformation between the pre-operative data and the points. The use of navigated ultrasound (US) imaging for acquiring bone surface points to be used in registration is an ongoing area of research, and one of the main advantages of using US is that points can be acquired non-invasively [1,2,3].

The use of the Unscented Kalman Filter (UKF) was recently proposed for rigid registration in CAOS, and was shown to have improved performance compared to the Iterative Closest Point method [4]. The advantage of the Kalman filter is that it is a computationally efficient least-squares solver, which is an ideal feature for intra-operative registration applications. Furthermore, the UKF was originally formulated to avoid some of the linearizations that occur in the classic formulation of the Kalman filter.

* The authors thank the NCCR Co-Me and BrainLab AG for funding this project, as well as colleagues Urs Rohrer, Erland Mühlheim, Sebastian Marti, and Ronald Ramseier for the experimental set-up.

T. Dohi, I. Sakuma, and H. Liao (Eds.): MIAR 2008, LNCS 5128, pp. 185–192, 2008.

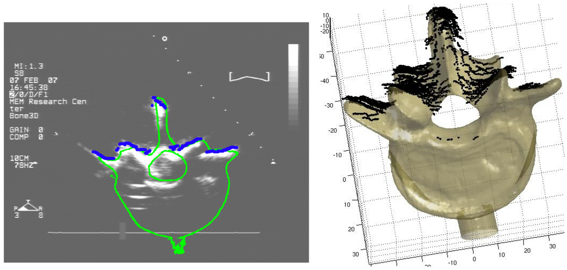

(a) (b)

Fig. 1. (a) Typical result of US segmentation (dark points) overlaid onto an US image and a contour generated from the ground truth (light grey) (b) segmented US points overlaid onto ground truth

We consider a US-based registration application in this paper, and reformulate the registration procedure such that the registration is effectively real-time. US images are considered as sequential signals, where each frame provides surface points of the anatomy. The registration is updated as frames are acquired, implying that iterations are no longer dependent on the number of points [4]. The registration itself begins during data acquisition, which implies that there is no need to wait until a full set of points is obtained before starting the algorithm.

The linear Kalman filter formulation used in this paper is akin to the perturbation Kalman filter [5]. The Extended Kalman Filter (EKF) could have been an alternative formulation, but the linearizations employed for this study were much simpler than the computation of Jacobian matrices [6]. We will demonstrate that the linear Kalman filter converges more quickly than the UKF, and is more robust with respect to starting positions. To evaluate the performance of these methods, we consider minimally-invasive interventions around the spine as potential target applications. The methods are compared using synthetic data and also in an experimental set-up using navigated 2D US, whereby the accuracy of the proposed registration method shows promise for use in clinical applications.

2 CT to US Rigid Registration

In the usual US-based registration scenario, navigated US images provide a set of bone surface points in the coordinate space of the patient's anatomy, which is referenced by the navigation system. To enhance the limited information provided by US, a surface model representation of the anatomy, as obtained from a CT for instance, is then registered to the US-acquired bone surface points.

The problem at hand is in finding the transformation that would place the CT-generated surface model in the coordinate space of the anatomy. For the Kalman filters, the measurement, $\mathbf{z}_k$, is a set of points in world coordinates acquired by segmenting the US images. At a given time step k, the state $\mathbf{x}_k$

is comprised of the true parameters that represent the transformation between the US points and the corresponding points on the surface model, $\hat{\mathbf{z}}_k^-$, which are found by searching for closest points using Euclidean distances. The estimate of this state is denoted $\hat{\mathbf{x}}_k$, and is a 6×1 vector $[\alpha, \beta, \gamma, t_x, t_y, t_z]^T$.

2.1 Kalman Filtering for Registration

For the UKF and the EKF, the state $\mathbf{x}_k$ represents the transformation from the CT surface model to the set of US-acquired points [4,6], and can be regarded as time-invariant. For the linear Kalman filter used here, however, we consider that $\mathbf{x}_k$ is time-variant, meaning that the state represents incremental transformations during the registration procedure [5]. For a given measurement $\mathbf{z}_k$, the corresponding points $\hat{\mathbf{z}}_k^-$ are found from the surface at the prior position, S_{k-1}.

To illustrate the registration method, consider the prediction and correction steps of the linear Kalman filter equations [7], with the unused terms omitted:

$$\hat{\mathbf{x}}_k^- = \hat{\mathbf{x}}_{k-1} \tag{1}$$

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_k^- + \mathbf{K}_k \left(\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_k^-\right) \tag{2}$$

At a given time step k, the correspondence is found between the new measurement $\mathbf{z}_k$ and S_{k-1}, yielding $\hat{\mathbf{z}}_k^-$. The predicted state $\hat{\mathbf{x}}_k^-$ is then updated (1) and applied to $\hat{\mathbf{z}}_k^-$, yielding the estimated measurement $\hat{\mathbf{z}}_k = \mathbf{H}\hat{\mathbf{x}}_k^-$. The predicted state is then corrected (2) and applied to the surface S_{k-1}, yielding S_k.

No model was used for propagating the state $\hat{\mathbf{x}}_k$. $\mathbf{z}_k$ and corresponding points $\hat{\mathbf{z}}_k^-$ are sets of vertically concatenated points such that they each become $3N \times 1$ column vectors. For instance, $\hat{\mathbf{z}}_k^-$ is written as $[\hat{z}^-_{x1}, \hat{z}^-_{y1}, \hat{z}^-_{z1}, \ldots, \hat{z}^-_{xN}, \hat{z}^-_{yN}, \hat{z}^-_{zN}]_k^T$, where N is the number of points segmented from US [4].

The matrix $\mathbf{H}$ relates the measurement $\mathbf{z}_k$ to the state $\hat{\mathbf{x}}_k$, and is used in all the correction steps of the Kalman filter. In the registration scenario, $\mathbf{H}$ should relate points in Cartesian coordinates to the parameters of a rigid transformation matrix, which is denoted $\mathbf{T}$, such that:

$$\hat{\mathbf{z}}_k = \mathbf{T}\hat{\mathbf{z}}_k^- = \mathbf{H}\hat{\mathbf{x}}_k^- \tag{3}$$

In order to determine $\mathbf{H}$, it becomes necessary to employ approximations for the terms in $\mathbf{T}$. We employ a zeroth order truncation of the Taylor Series representations of cosine and sine, such that for a given angle θ, $\cos\theta \approx 1$ and $\sin\theta \approx \theta$, which should be valid for small angles. We denote the approximation of $\mathbf{T}$ by $\mathbf{T}_{\mathrm{A}}$ (using the $\alpha\beta\gamma$ convention), and (3) for a single point can be written as:

$$\mathbf{T}_{\mathrm{A}}\hat{\mathbf{z}}_k^- = \begin{pmatrix} 1 & \gamma + \alpha\beta & \alpha\gamma - \beta & t_x \\ -\gamma & 1 - \alpha\beta\gamma & \alpha + \beta\gamma & t_y \\ \beta & -\alpha & 1 & t_z \end{pmatrix} \begin{pmatrix} \hat{z}^-_{xk} \\ \hat{z}^-_{yk} \\ \hat{z}^-_{zk} \\ 1 \end{pmatrix} \tag{4}$$

The residual in the Kalman filter update step (2), $\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_k^-$, can be reinterpreted as $\mathbf{z}_k - \mathbf{T}_{\mathrm{A}}\hat{\mathbf{z}}_k^-$, and $\mathbf{T}_{\mathrm{A}}\hat{\mathbf{z}}_k^-$ can finally be expressed as:

$$\mathbf{T}_{\mathrm{A}}\hat{\mathbf{z}}_k^- = \begin{pmatrix} \hat{z}_{xk}^- \\ \hat{z}_{yk}^- \\ \hat{z}_{zk}^- \end{pmatrix} + \mathbf{H}\hat{\mathbf{x}}_k^- + \begin{pmatrix} \alpha\beta\hat{z}_{yk}^- + \alpha\gamma\hat{z}_{zk}^- \\ \beta\gamma\hat{z}_{zk}^- - \alpha\beta\gamma\hat{z}_{yk}^- \\ 0 \end{pmatrix} \tag{5}$$

Where $\mathbf{H}$, to be used in the Kalman filter equations, becomes:

$$\mathbf{H} = \begin{pmatrix} 0 & -\hat{z}_{zk}^- & \hat{z}_{yk}^- & 1 & 0 & 0 \\ \hat{z}_{zk}^- & 0 & -\hat{z}_{xk}^- & 0 & 1 & 0 \\ -\hat{z}_{yk}^- & \hat{z}_{xk}^- & 0 & 0 & 0 & 1 \end{pmatrix} \tag{6}$$

For a single point, $\mathbf{H}$ is a 3×6 matrix that comprises the cross product matrix on the left side, and $\mathbf{I}_3$ on the right side. For N points, where $\hat{\mathbf{z}}_k^-$ is a $3N \times 1$ vector, $\mathbf{H}$ is a $3N \times 6$ matrix.

Moghari et al. [4] applied the UKF to estimate rigid transformation parameters in US-to-CT registration. By their implementation, the UKF iterates N times, where N is the number of points in their $\mathbf{z}_k$, and for each iteration, the number of points is gradually increasing. While this approach usually provides smooth filter behavior, it does not take advantage of the Kalman filter's sequential real-time nature.

With the UKF, no explicit relationship between the transformation parameters $\hat{\mathbf{x}}_k$ and the points $\mathbf{z}_k$ needs to be defined, but rather the UKF operates on the principle of an implicit linearization of the relationship. At each prediction step, the estimated state $\hat{\mathbf{x}}_k$ is applied to the surface at the initial position, S_1, producing S_k. Correspondence is determined between the measurement and S_k, and when the registration is completed, the resulting state is applied to S_1.

2.2 Frame-by-Frame Registration Procedure

In the usual approach to registration, the registration procedure does not begin until a full set of points is acquired from the anatomy. For the work presented here, we propose using the Kalman filter to update the registration as navigated 2D US images are acquired.

For both Kalman filters, each US frame is treated as one signal, which yields a coplanar cloud of points, $\mathbf{z}_k$. The registration is updated with each newly acquired frame, and so the time step k is now related to the number of US frames, as they are acquired, rather than the number of points. During image acquisition, the surgeon can then receive visual (fig. 1(b)) and numerical feedback in terms of how well the Kalman filter fits the surface to the US-acquired points. The result of this approach is that fewer iterations are needed for the registrations. To stabilize the frame-by-frame processing of the filters, a small subset of points from prior US frames is used to complement the point sets of new frames. In §3, the original UKF registration algorithm is compared to the frame-by-frame UKF and linear Kalman filter registration methods.

3 Results

3.1 Simulated Data

In order to compare their performance on ideal data, the three Kalman filter registration methods were used with synthetic data. For the measurement, consider 100 randomly selected points from a 3D surface model of an L4 vertebra, in the area that would be accessible by US. $\mathcal{N}(0, 1[mm^2])$ noise was added to the coordinates of the points. In the original UKF registration method (point-by-point UKF), each point is treated 99/2 times on average. To achieve a comparison using

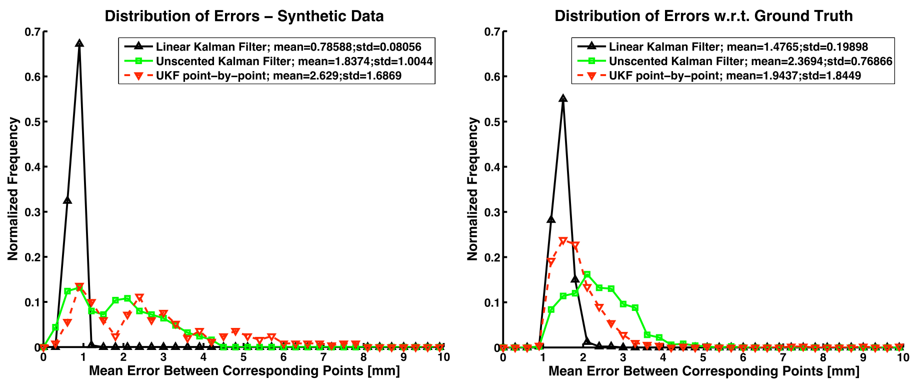

(a) (b)

Fig. 2. Distribution of RMS corresponding point-to-point errors, computed between surfaces at estimated positions and the ground truth, for (a) synthetic data and (b) experimental validation, comparing the three Kalman filter-based approaches

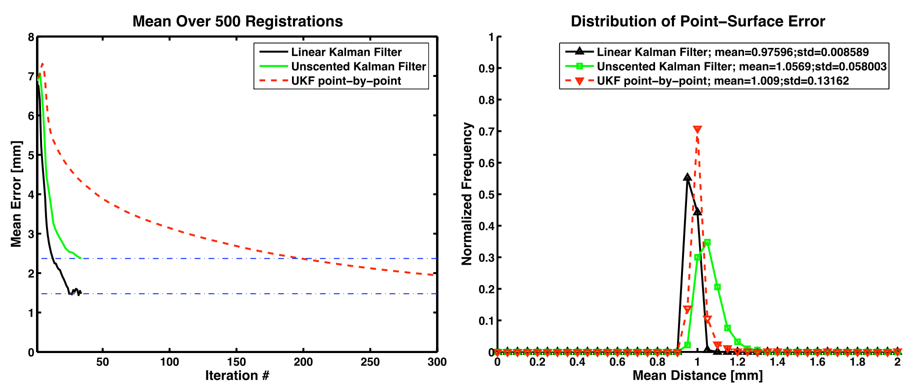

(a) (b)

Fig. 3. (a) Mean errors plotted with respect to number of iterations and (b) distribution of RMS point-to-surface errors, computed using nearest points between estimated surfaces and ground truth

the full-dataset approach, the frame-by-frame filters were iterated 50 times over the 100 points. A set of 500 randomly-generated rigid transformations (within a range of $\pm$ 5 [deg.] for the three angles, $\pm$ 10 [mm] for two translations and 0-20 [mm] in the US scanning direction) was applied to the surface model, and each Kalman filter was used to register the surface to the points for each transformation. Errors were computed by measuring the distances between corresponding points on the registered surface and the surface at the true position.

As can be inferred from the graph in fig. 2(a), the two UKF methods have similarly distributed mean errors, but the point-by-point UKF had much higher errors in several cases. This result indicates that the point-by-point UKF was more susceptible to initial positions than the full-dataset approach, and that 100 points were not enough for it to converge to a desirable result. The linear Kalman filter rapidly converged to the correct solution (overall mean error of 0.79 [mm]) and was robust with respect to starting positions. Nevertheless, it is to be expected that given a larger set of points, and more iterations during the registration, the UKF would eventually achieve an accurate registration.

3.2 Experimental Validation with Real US Data

The experimental validation was performed on a navigated plastic L4 vertebra immersed in a water bath. A ground truth registration was obtained using a navigated pointer to digitize surface points and then applying a surface matching algorithm (point-to-surface error was 0.5 [mm]). After calibrating the US probe [8], navigated 2D B-mode US images where then acquired using the Philips Sonos 7500 ultrasound system. 36 tracked US frames were acquired, focusing on the areas that would be visible in a real patient set-up (fig. 1(b)). Bone contour points were then automatically extracted by combining Otsu thresholding with a morphological opening, and a thinning of the resulting border along the US scan lines. This preliminary segmentation method is suitable for water bath scans, with an expected segmentation error of 1-2 [mm], and less than one second of computation per frame. For each frame, the segmentation was sampled to yield 100 bone contour points. Fig. 1(a) shows an ultrasound image with a typical outcome of the segmentation procedure and the contour from the ground truth.

As in §3.1, 500 transformations were randomly generated and applied to the ground truth. The Kalman filters were then used to register the transformed surface to the US points. For the frame-by-frame filters, 8 points were randomly selected from each prior US image as new frames were used for the registration. The point-by-point UKF was halted once 300 points were processed in order to keep this method computationally feasible.

Fig. 2(b) shows the distribution of mean errors over all registrations, which were computed by measuring distances of corresponding points between registered surfaces and the ground truth, providing a measure in the areas that were not scanned by ultrasound (i.e. the vertebral body). The linear Kalman filter presented in this paper showed superior performance in terms of accuracy (overall mean error of 1.48 [mm]) and robustness with respect to starting positions (standard deviation of 0.20 [mm]). Although the point-by-point UKF achieved a

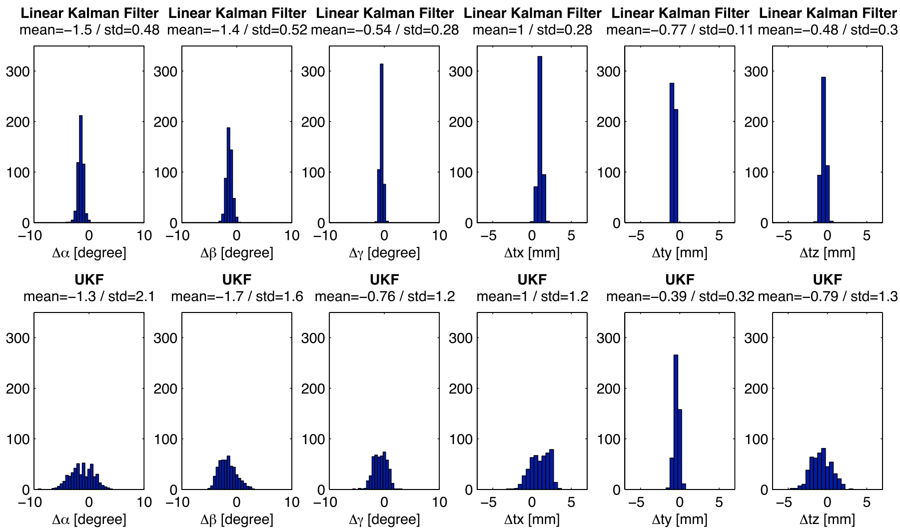

Fig. 4. Distributions of transformation parameter errors

lower mean error than the frame-by-frame UKF, it demonstrated higher sensitivity to starting positions in 6 cases, when it had mean errors higher than 5 [mm]. The point-by-point UKF results may have been improved if it was iterated over more points, but this would have led to an exorbitant amount of computation (e.g. for N points, computing matrix inverses for $3N \times 3N$ matrices).

Fig. 3(a) shows the trend of mean errors with respect to iteration number for each of the three methods. The point-by-point UKF was iterated 300 times, and showed a generally smooth behavior in its estimation. The frame-by-frame methods, which were iterated 36 times, demonstrated much quicker convergence in their estimation, with the linear Kalman filter converging the quickest. Fig. 3(b) shows the distribution of mean distances between the US points and the surface. This result is analogous to the result in [4], and consistent with the the more reliable measure in fig. 2(b).

The histograms in fig. 4 reflect the differences between transformation parameters obtained by the frame-by-frame Kalman filters and those of the 500 randomly-generated transformations, whereby the linear Kalman filter showed a more consistent estimation of the parameters than the UKF.

4 Conclusions

The experimental results suggest that the linear Kalman filter produces more accurate results compared to the Unscented Kalman Filter in the frame-by-frame registration framework presented in this paper. The greatest reason for this unexpected outcome stems from the approach to the design of the filters. Both in this work and in [4], batch processing of points for each iteration was performed by concatenating the points into higher-dimensional vectors. In this way, it is assumed that noise on the points is independent, which is the same assumption as would be made if the points were treated sequentially [6]. In concatenating the points, however, their covariance as calculated from the state, a 3×3 matrix for each point, is also concatenated to yield a $3N \times 3N$ matrix. This would then lead to an ill-conditioned system, but this effect is implicitly mitigated by the Kalman filter due to the additive measurement noise, which regularizes the ill-conditioned covariance matrix. Further discussion on the additive measurement noise can be found in [4], but it would be interesting to investigate more optimal approaches to the sequential batch processing employed here [9].

Nevertheless, the current design demonstrated rapid convergence, reliability under different initial conditions, as well as an accuracy high enough to motivate more thorough investigation, nearer to clinical situations. Although the validation schemes focused on potential applications in the spine, the methods were formulated generally enough that they may be applicable to various minimally-invasive orthopedic procedures.

References

1. Barratt, D., Penney, G., Chan, C., Slomczykowski, M., Carter, T., Edwards, P., Hawkes, D.: Self-calibrating 3d-ultrasound-based bone registration for minimally invasive orthopedic surgery. IEEE Trans. Med. Imaging 25(3), 312–323 (2006)
2. Muratore, D.M., Russ, J.H., Dawant, B.M., Galloway, R.L.: Three-dimensional image registration of phantom vertebrae for image-guided surgery: a preliminary study. Comput. Aided Surg. 7(6), 342–352 (2002)
3. Talib, H., Rajamani, K., Kowal, J., Nolte, L.P., Styner, M., Ballester, M.A.G.: A comparison study assessing the feasibility of ultrasound-initialized deformable bone models. Comput Aided Surg 10(5-6), 293–299 (2005)
4. Moghari, M.H., Abolmaesumi, P.: Point-based rigid-body registration using an unscented kalman filter. IEEE Transactions on Medical Imaging 26(12), 1708–1728 (2007)
5. Grewal, M.S., Andrews, A.P.: Kalman Filtering Theory and Practice. Information and System Sciences Series. Prentice-Hall, Englewood Cliffs (1993)
6. Pennec, X., Thirion, J.: A framework for uncertainty and validation of 3d registration methods based on points and frames. International Journal of Computer Vision 25(3), 203–229 (1997)
7. Welch, G., Bishop, G.: An introduction to the kalman filter. Technical Report TR 95-041, University of North Carolina, Chapel Hill (1995)
8. Kowal, J., Amstutz, C., Nolte, L.P.: On the development and comparative evaluation of an ultrasound b-mode probe calibration unit. Comput. Aided Surg. 8(3), 107–119 (2003)
9. Maybeck, P.S.: Stochastic Models, Estimation, and Control. 1 of Mathematics in Science and Engineering, vol. 141. Academic Press, London (1979)