

# Surface-Based Registration with a Particle Filter

Burton Ma and Randy E. Ellis

School of Computing, Queen's University at Kingston, Canada K7L 3N6

**Abstract.** We propose the use of a particle filter as a solution to the rigid shape-based registration problem commonly found in computer-assisted surgery. This approach is especially useful where there are only a few registration points corresponding to only a fraction of the surface model. Tests performed on patient models, with registration points collected during surgery, suggest that particle filters perform well and also provide novel quality measures to the surgeon.

## 1 Introduction

Preoperative 3D medical images, such as CT and MRI scans, can be registered intraoperatively to a patient's anatomy by estimating a transformation from surfaces in image coordinates to anatomical points in patient coordinates for use in image-guided surgery. Two notable limitations of current algorithms are: (a) most algorithms are non-incremental and process an additional anatomical point by reconsidering the entire set of anatomical points gathered during surgery; and (b) most algorithms report errors such as root-mean-square (RMS) or target registration (TRE) but do not report the probable distribution of errors. Particle filters offer an incremental method for computing probability distributions of the rotational and translational components of a rigid registration.

Most registration algorithms attempt to find the registration parameters that maximize the likelihood of the registration given the anatomical points, surface model, and measurement noise. This results in an *a posteriori* problem that admits a solution by expectation maximization, which has been solved by Chui and Rangarajan [2], Granger and Pennec [6], Dellaert [4], and even by Besl and McKay [1] (the ICP algorithm).

We propose to solve the registration problem by estimating a *filter distribution*. For a state vector  $\mathbf{x}_t$  at time  $t$  that represents the state of rotation and translation of the registration transformation, and given the observations  $\mathbf{y}_1, \dots, \mathbf{y}_t \equiv \mathbf{y}_{1:t}$ , the filter distribution  $p(\mathbf{x}_t | \mathbf{y}_{1:t})$  is the conditional probability distribution over the state vector given the observations. We use a particle filter (PF), specifically the unscented particle filter (UPF) described by van der Merwe et al. [13], to estimate the filter distribution. A PF represents the posterior distribution of the states as a set of weighted *particles* (or samples), making it easy to calculate statistical estimates of the state. This means that we can estimate the registration parameters using the mean, or any other statistical measure of location, and the precision of the parameters in terms of their standard deviations, confidence intervals, or any other statistical measure of spread. All of these estimates are updated *incrementally* as new registration points become available, which gives the surgeon an indication of whether or not more anatomical points are needed to achieve the desired level of precision. Multiple peaks in the distribution indicate

that there are multiple plausible registrations and that some action should be taken to resolve the ambiguity; broad peaks indicate that the anatomical points do not provide good localization and that the surgeon should take this into account while operating. Our framework can incorporate any prior knowledge, such as the location and localization uncertainty of landmarks, into the estimation process.

## 2 Previous Work

Particle filtering is a Monte Carlo method for estimating the posterior distribution of a state space model. The technique has been independently described by many different authors; a survey of sequential sampling methods for Bayesian filtering has been written by Doucet et al. [5]. An example of more recent work is the mixture particle filter of Vermaak et al. [15] and its application to multi-target tracking by Okuma et al. [11].

The state space model at time  $t$ , is described by a state transition model,  $p(\mathbf{x}_t|\mathbf{x}_{t-1})$ , and an observation model,  $p(\mathbf{y}_t|\mathbf{x}_t)$ . The state is modeled as a first-order Markov process, and the observations are conditionally independent given the states. We can express the state transition model (often called the process model),  $\mathbf{F}$ , and the observation model,  $\mathbf{H}$ , as

$$\begin{aligned}\mathbf{x}_t &= [x_{1_t} \ x_{2_t} \ x_{3_t} \ x_{4_t} \ x_{5_t} \ x_{6_t}]^T \\ \mathbf{x}_{t+1} &= \mathbf{F}(\mathbf{x}_t, \mathbf{u}_t, \mathbf{v}_t) \\ \mathbf{y}_t &= \mathbf{H}(\mathbf{x}_t, \mathbf{u}_t, \mathbf{n}_t)\end{aligned}\tag{1}$$

where the registration state has three rotation parameters  $x_{1_t}, x_{2_t}, x_{3_t}$  and three translation parameters  $x_{4_t}, x_{5_t}, x_{6_t}$ ;  $\mathbf{u}_t$  is a known control input; and  $\mathbf{n}_t$  is the observation noise. The process noise,  $\mathbf{v}_t$ , influences the system dynamics. The distribution of the initial state,  $p(\mathbf{x}_0)$ , is the prior distribution and must also be specified. The equations  $\mathbf{F}$  and  $\mathbf{H}$  need not be linear functions of their parameters.

The basic idea underlying the PF is that the posterior distribution  $p(\mathbf{x}_{0:t}|\mathbf{y}_{0:t})$  can be approximated by a weighted sum of particles from the posterior distribution. We cannot in general sample directly from the posterior, so instead we sample from a *proposal* distribution  $q(\mathbf{x}_{0:t}|\mathbf{y}_{0:t})$ . The weights are called the importance weights; particles with high weights correspond to regions of high density in the posterior. Under certain assumptions, the weights at time  $t$  can be computed sequentially using the weights at time  $t - 1$  to yield an incremental registration algorithm. It is inefficient to keep particles with low importance weights, so a resampling step is performed to remove particles with low weights and multiply particles with high weights. An optional Markov-chain Monte-Carlo step can be introduced after the resampling step to smooth the distribution of particles. The output of the PF is a set of equally weighted samples that approximate the posterior  $p(\mathbf{x}_{0:t}|\mathbf{y}_{0:t})$ . The filter distribution is simply the marginal distribution  $p(\mathbf{x}_t|\mathbf{y}_{0:t})$ .

The UPF uses the unscented Kalman filter (UKF) of Julier and Uhlmann [10] (see also [16] and the Gaussian filters of Ito and Xiong [8]) to compute a Gaussian approximation of the posterior for each particle; the proposal is drawn from the resulting Gaussian distribution. The UKF is a filter that assumes a state space model given by Equation 1.

- Initialize the filter by drawing  $i = 1, \dots, N$  particles  $\mathbf{x}_0^i$  from the prior distribution of the states  $p(\mathbf{x}_0)$ . Assign to each particle a covariance matrix  $\mathbf{P}_0^i$  (which will be propagated through the UKF).
- For time  $t = 1, 2, 3, \dots$ 
  - For  $i = 1, \dots, N$ 
    - \* Filter particle  $\mathbf{x}_{t-1}^i$  with covariance  $\mathbf{P}_{t-1}^i$  using the UKF to get an updated estimate  $\bar{\mathbf{x}}_t^i$  and  $\hat{\mathbf{P}}_t^i$
    - \* Sample  $\hat{\mathbf{x}}_t^i$  from the Gaussian proposal  $q(\mathbf{x}_t^i | \mathbf{x}_{0:t-1}^i, \mathbf{y}_{0:t}) = \mathcal{N}(\bar{\mathbf{x}}_t^i, \hat{\mathbf{P}}_t^i)$
    - \* Compute the observation likelihood  $p(\mathbf{y}_t | \hat{\mathbf{x}}_t^i)$
    - \* Compute the state transition prior  $p(\hat{\mathbf{x}}_t^i | \mathbf{x}_{t-1}^i)$
    - \* Compute the importance weight  $w_t^i \propto \frac{p(\mathbf{y}_t | \hat{\mathbf{x}}_t^i) p(\hat{\mathbf{x}}_t^i | \mathbf{x}_{t-1}^i)}{q(\hat{\mathbf{x}}_t^i | \mathbf{x}_{0:t-1}^i, \mathbf{y}_{0:t})}$
  - Normalize the importance weights  $w_t^i = w_t^i / \sum_{i=1}^N w_t^i$
  - Apply a resampling scheme to remove particles with low weights and multiply particles with high weights if necessary.
  - Output the  $N$  particles  $\mathbf{x}_t^i$ , or more generally the particle trajectories  $\mathbf{x}_{0:t}^i$

**Fig. 1.** The unscented particle filter algorithm.

It propagates the state mean and covariance through the (possibly non-linear) process and observation models by using a set of deterministically sampled points called sigma points. The previously cited references provide evidence of the superior performance of the UKF compared to the well known extended Kalman filter.

Figure 1 is an abbreviated version of the UPF algorithm; the algorithm is discussed in much greater detail by van der Merwe et al. [13] and the application of the filter to object tracking is given by Rui and Chen [12].

### 3 UPF Registration

We use the state space model

$$\mathbf{x}_{t+1} = \mathbf{F}(\mathbf{x}_t, \mathbf{v}_t) = \mathbf{x}_t + \mathcal{N}(\mathbf{0}, \mathbf{Q}_t) \quad (2)$$

$$\mathbf{y}_t = \mathbf{H}(\mathbf{x}_t, \mathbf{u}_t, \mathbf{n}_t) = \begin{bmatrix} \mathbf{r}(x_{1_t}, x_{2_t}, x_{3_t})(\mathbf{u}_1 + [x_{4_t} \ x_{5_t} \ x_{6_t}]^T) \\ \vdots \\ \mathbf{r}(x_{1_t}, x_{2_t}, x_{3_t})(\mathbf{u}_1 + [x_{4_t} \ x_{5_t} \ x_{6_t}]^T) \end{bmatrix} + \mathcal{N}(\mathbf{0}, \mathbf{R}_t) \quad (3)$$

Equation 2 is the process model for estimating the registration state of the three rotation parameters and the three translation parameters; the model has time-invariant state, except for the additive process noise, because we are estimating a constant. The zero-mean Gaussian process noise with covariance  $\mathbf{Q}_t$  allows us to move from possibly poor initial estimates of the state to successively better estimates. We anneal  $\mathbf{Q}_t$  towards  $\mathbf{0}$  over time as our estimates become better. For this article, we took  $\mathbf{Q}_t$  to be uncorrelated with initial variances of  $1.22 \times 10^{-3} \text{rad}^2$  and  $4 \text{mm}^2$  for the rotational and translational components; we annealed by a factor of 0.7 after each time step. We prefer to use rotation parameters that surgeons are most familiar with, measured as order-independent angles

of rotation in the coronal, sagittal, and horizontal planes. Iyun et al. [9]<sup>1</sup> presented a simplified version of this rotation matrix.

Our observation model is given by Equation 3 where  $\mathbf{r}(x_{1_t}, x_{2_t}, x_{3_t})$  is a rotation matrix,  $[x_{4_t} \ x_{5_t} \ x_{6_t}]^T$  is a translation vector,  $\mathbf{u}_i$  is the  $i$ th registration point in patient coordinates. The model is simply the estimated registration transformation applied to the registration points (which we supply as the control input) concatenated into a single vector; the length of the vector at time  $t$  is  $3t$ . We assume additive, zero-mean Gaussian noise with covariance  $\mathbf{R}_t$ ; the noise is the displacement of each transformed registration point to the surface of the model. For this article, we took  $\mathbf{R}_t$  to be uncorrelated with initial variances of  $4\text{mm}^2$  and annealed by a factor of 0.8 after each time step.

The actual observation,  $\mathbf{y}_t$ , should be the vector formed by the concatenation of the model points corresponding to the registration points; of course, these correspondences are unknown. We used the nearest-neighbor approximation made popular by Besl and McKay [1]; computationally more expensive alternatives such as the Mahalanobis distance (Granger and Pennec [6]) or Gaussian weighted distance (Chui and Rangarajan [2], Grimson et al. [7]) could also be used.

The prior distribution  $\mathbf{x}_0$  depends on how the initial estimate of the registration transformation was obtained, e.g., by performing a Monte Carlo simulation of the initial estimation process. If anatomic landmarks or fiducial markers are present then the prior should reflect this information.

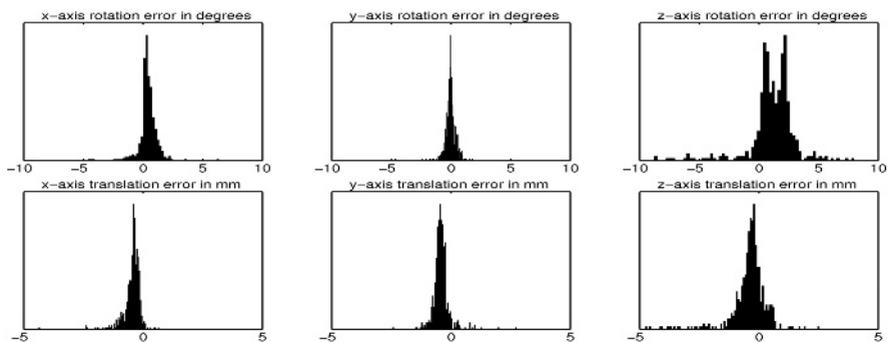
## 4 Experiments

We performed experiments to test for accuracy and precision using synthetic and intraoperatively digitized registration points for the proximal tibia, distal radius, and proximal femur. All surface models were derived from CT scans of patients who consented to surgery and research use of data in a study approved by the Research Ethics Board of Queen's University and Kingston General Hospital.

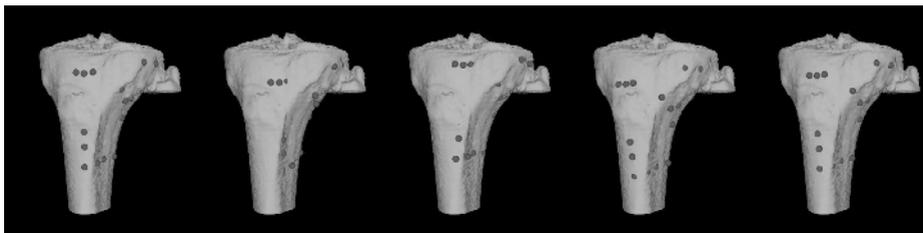
### 4.1 Synthetic Points and the Proximal Tibia

We tested for registration accuracy and precision by massively subsampling points from a realistic surgical exposure of the proximal tibia. Fifteen points were selected from the surface model generated from a CT scan of a patient that underwent high tibial osteotomy (HTO). We generated 500 random transformations by drawing from the uniform distribution  $\mathcal{U}(\pm 20^\circ, \pm 20^\circ, \pm 20^\circ, \pm 10\text{mm}, \pm 10\text{mm}, \pm 10\text{mm})$ ; each transformation was applied to the points to which we added random noise drawn from the uniform distribution  $\mathcal{U}(\pm 0.5\text{mm}, \pm 0.5\text{mm}, \pm 0.5\text{mm})$ . The range of angles and displacements are large compared to the error of the initial estimate that we typically obtain in surgery; the mean total angular displacement was  $19^\circ$  and mean total translation was 10mm. We used 5000 particles and a prior of  $\mathcal{U}(\pm 25^\circ, \pm 25^\circ, \pm 25^\circ, \pm 12\text{mm}, \pm 12\text{mm}, \pm 12\text{mm})$ . We used the UPF to register the transformed point sets, taking the mean of the posterior distribution for the registration parameters. The distribution of rotation and translation errors are shown in Figure 2.

<sup>1</sup> There is a typographical error in this reference; in Equation 1, page 235, the term  $dc - af$  should be  $af - dc$



**Fig. 2.** Proximal tibia and synthetic data results. Error distributions for 500 UPF registrations; standard deviations are  $0.98^\circ$ ,  $0.49^\circ$ ,  $2.02^\circ$  for the  $x$ ,  $y$ ,  $z$  axis rotation errors, and  $0.39\text{mm}$ ,  $0.36\text{mm}$ ,  $0.77\text{mm}$  for the  $x$ ,  $y$ ,  $z$  axis translation errors.



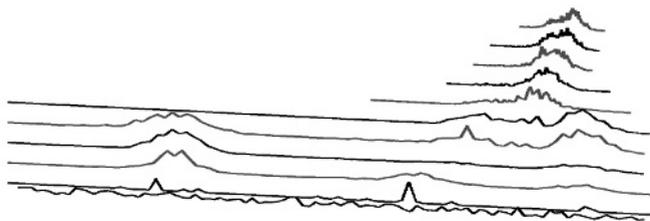
**Fig. 3.** Proximal tibia and synthetic data results. Left: the registration data points. The other four images are of registrations with large rotation and/or translation errors. The spheres indicating the registration points are  $4\text{mm}$  in diameter.

The consistency of the registrations was good, particularly in the rotations about and the translations along the  $x$  and  $y$  axes. Rotation errors about and translation errors along the  $z$  axis (the long axis of the bone) were worse, but this was expected given the shape of the proximal tibia and the location of the surgically accessible region. The registrations that were far from the correct registrations still produced small RMS errors between the registered points and the surface model; the filter simply found a different fit of the very small number of registration points to the model. Figure 3 illustrates the potential difficulty in finding the correct registration.

The evolution of the posterior distribution of the rotation about the  $z$  axis for one trial is shown in Figure 4. When there are only a few registration points, the UPF finds a few different possible values for the angle of rotation. As more points are added, the ambiguity is resolved and the posterior converges to narrow range of values.

## 4.2 Intraoperative Points and the Distal Radius

We performed a retrospective analysis of registration data from seven patients who underwent computer-assisted distal radius osteotomy (Croitoru et al. [3]). We used the registration data collected intraoperatively using a calibrated pointer and an Optotrak



**Fig. 4.** Proximal tibia and synthetic data results. The evolution of the posterior distribution of the  $z$ -axis rotation. The prior distribution, sampled from  $\mathcal{U}(\pm 25^\circ)$  is at the bottom. Moving towards the top of the page are the posteriors after 6, 7, ..., 15 registration points have been processed.

**Table 1.** Distal radius and intraoperative registration points experimental results. Standard deviations, over 100 trials, of the registration state using the UPF to register to the distal radius. The  $z$  axis is the long axis of the bone.

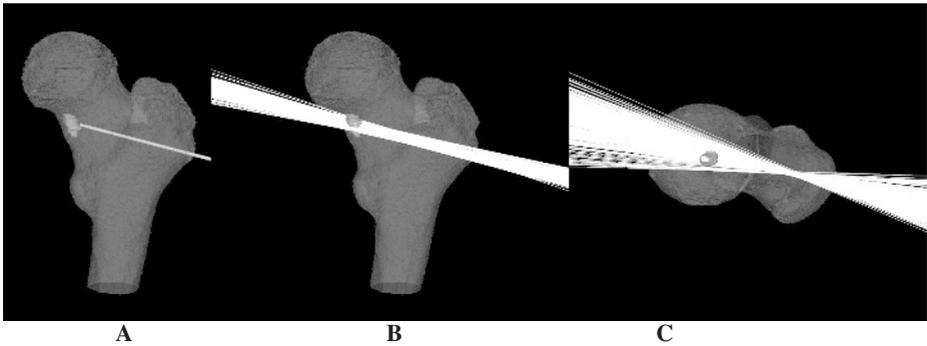
| Num. Reg. Pts. | $\theta_x$   | $\theta_y$   | $\theta_z$   | $t_x$  | $t_y$  | $t_z$  |
|----------------|--------------|--------------|--------------|--------|--------|--------|
| 11             | $0.74^\circ$ | $0.54^\circ$ | $2.48^\circ$ | 0.36mm | 0.20mm | 0.50mm |
| 10             | $0.51^\circ$ | $0.83^\circ$ | $1.26^\circ$ | 0.10mm | 0.06mm | 0.31mm |
| 11             | $0.54^\circ$ | $0.37^\circ$ | $0.69^\circ$ | 0.08mm | 0.10mm | 0.30mm |
| 11             | $0.36^\circ$ | $0.48^\circ$ | $0.42^\circ$ | 0.14mm | 0.12mm | 0.80mm |
| 10             | $0.38^\circ$ | $0.27^\circ$ | $0.66^\circ$ | 0.16mm | 0.07mm | 0.65mm |
| 11             | $0.26^\circ$ | $0.52^\circ$ | $1.47^\circ$ | 0.05mm | 0.06mm | 0.13mm |
| 12             | $0.14^\circ$ | $0.33^\circ$ | $0.69^\circ$ | 0.07mm | 0.03mm | 0.11mm |

3020 localizer (NDI, Waterloo, Canada) and examined the consistency of the UPF registrations. Lister's tubercle is a prominent anatomic landmark on the distal radius. We can use the tubercle to directly estimate the translation component of the registration, thus reducing the range of the translation component of the prior.

We generated 100 random transformations, starting from the surgical registration, for each bone by drawing from the uniform distribution  $\mathcal{U}(\pm 10^\circ, \pm 10^\circ, \pm 10^\circ, \pm 1\text{mm}, \pm 1\text{mm}, \pm 3\text{mm})$ . We applied these transformations to the registration points and computed the UPF registrations using 2000 particles and a prior of  $\mathcal{U}(\pm 12^\circ, \pm 12^\circ, \pm 12^\circ, \pm 3\text{mm}, \pm 3\text{mm}, \pm 5\text{mm})$ ; standard deviations over the 100 trials are shown in Table 1.

### 4.3 Intraoperative Points and the Proximal Femur

We examined the utility of obtaining the filter distribution by retrospectively calculating the distribution of probable drill paths for a computer-assisted minimally invasive removal of a deep bone tumor. The registration points for this procedure were collected through a small incision on the lateral side of the hip and percutaneously using a needle-tipped probe. The cylindrical shape of the proximal femur makes it difficult to precisely estimate the axial rotation component of the registration transformation; the standard deviation for this component was  $3.34^\circ$ . Figure 5 illustrates the osteoma, planned drill



**Fig. 5.** Proximal femur and intraoperative registration points experimental results. **A** an osteoid osteoma located in the neck of the femur and the planned drill path. **B** predicted drill paths after registration viewed in the frontal plane. **C** predicted drill paths in the axial plane; the axial rotation component has the highest uncertainty due to the almost cylindrical shape of the proximal femur.

path, and the drill path predicted by each of the 2000 particles. It is easy to appreciate the consequences of uncertainty in the axial rotation by direct visualization of the drill paths; target registration error (TRE) is a single number that is less informative.

## 5 Discussion

UPF registration offers notable advantages to most alternative approaches. Because it is iterative, a surgeon can collect only as many points as needed, which can reduce operative time. More importantly, the surgeon can directly visualize the *quality* of the registration, by examining either the posterior distribution (which gives an overall view of the error space) or of the transformation of a planned trajectory (which gives a task-specific view of error effects). For example, had UPF registration been available when the osteoid osteoma case was performed, the surgeons would have been vividly aware that they were well on target but that the drill bit was more likely to track posteriorly.<sup>2</sup>

Limitations of our study include running times, an unoptimized implementation, and unknown robustness to statistical outliers. The UPF runs in an acceptable period of time when using a small number of registration points. For example, using an Intel Pentium 4 PC (2.8 GHz) we updated the state estimate with 2000 particles and 15 registration points in 1.5 seconds; much of the time was spent performing nearest-neighbor searches. We have not thoroughly investigated the relationship between the number of particles and the estimation process; it may be possible to safely reduce the number of particles. It may even be possible to avoid particle filtering altogether by estimating the posterior density using a Gaussian mixture filter (van der Merwe and Wan [14] or Ito and Xiong [8]). We also have not addressed the issue of robustness to outliers, but we believe that the PF framework will allow us to improve our algorithm in this respect.

<sup>2</sup> The tumor was successfully excised using robust registration for guidance.

We have illustrated an application of estimating the filter distribution for surgical removal of an osteoid osteoma. Any quantities that depend on having a registration transformation, such as instrument trajectories, can also be estimated and the uncertainty of the registration can be propagated into the estimations. We believe that UPF registration will allow us to estimate the stability of surface-based registration, which is an important but poorly understood part of computer-assisted surgical systems.

**Acknowledgments.** This research was supported in part by the Institute for Robotics and Intelligent Systems, the Ontario Research and Development Challenge Fund, and the Natural Sciences and Engineering Research Council of Canada.

## References

- [1] P. Besl and N. McKay. A method for registration of 3-D shapes. *IEEE Trans Pattern Anal*, 14(2):239–256, 1992.
- [2] H. Chui and A. Rangarajan. A feature registration framework using mixture models. In *IEEE Workshop on Math Methods in Biomed Im Anal*, pages 190–197, 2000.
- [3] H. Croitoru, R. E. Ellis, C. F. Small, R. Prihar, and D. R. Pichora. Fixation-based surgery: A new technique for distal radius osteotomy. *J Comp Aided Surg*, 6:160–169, 2001.
- [4] F. Dellaert. *Monte Carlo EM for Data-Association and its Applications in Computer Vision*. PhD thesis, Carnegie Mellon University, 2001.
- [5] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat Comp*, 10:197–208, 2000.
- [6] S. Granger and X. Pennec. Multi-scale EM-ICP: A fast and robust approach for surface registration. In *ECCV 2002*, pages 418–432, 2002.
- [7] W. E. L. Grimson, G. J. Ettinger, S. J. White, T. Lozano-Pérez, W. M. Wells III, and R. Kikinis. An automatic registration method for frameless stereotaxy, image guided surgery, and enhanced reality visualization. *IEEE Trans Med Imaging*, 15(2):129–140, 1996.
- [8] K. Ito and K. Xiong. Gaussian filters for nonlinear filtering problems. *IEEE Trans Automat Con*, 45(5):910–927, 2000.
- [9] O. Iyun, D. P. Borschneck, and R. E. Ellis. Computer-assisted correction of bone deformities. In T. Dohi and R. Kikinis, editors, *MICCAI 2002*, pages 232–240, Tokyo, 2002.
- [10] S. Julier and J. Uhlmann. A new extension of the Kalman filter to nonlinear systems. In *Proc. of AeroSense*, Orlando, Florida, 1997.
- [11] K. Okuma, A. Taleghani, N. de Freitas, J. Little, and D. Lowe. A boosted particle filter: Multitarget detection and tracking. In *ECCV 2004*, 2004.
- [12] Y. Rui and Y. Chen. Better proposal distributions: Object tracking using unscented particle filter. In *Proc. of IEEE Conf on Comp Vis and Patt Recog*, pages 786–793, 2001.
- [13] R. van der Merwe, A. Doucet, N. de Freitas, and E. Wan. The unscented particle filter. Technical Report CUED/F-INFENG/TR 380, Cambridge Univ. Dept. of Eng., 2000.
- [14] R. van der Merwe and E. Wan. Gaussian mixture sigma-point particle filters for sequential probabilistic inference in dynamic state-space models. In *Proc of the Int Conf on Acous, Speech, and Sig Proc*, 2003.
- [15] J. Vermaak, A. Doucet, and P. Pérez. Maintaining multi-modality through mixture tracking. In *Int. Conf. on Computer Vision*, Nice, France, 2003.
- [16] E. Wan and R. van der Merwe. The unscented kalman filter for nonlinear estimation. In *Proc of Sym on Adapt Sys for Sig Proc, Comm and Control*, Lake Louise, Canada, 2000.