

# Performance Feedback for Closed-loop Particle Filters

Sami Siddique, Eugene Fiume, and David A. Jaffray

**Abstract**— Particle filters have emerged as a popular technique for solving nonlinear and non-Gaussian sequential estimation problems as they are in some sense "optimal" with sufficient sampling density. This can be brittle, however, for if the sampling threshold is not met, the estimator may be very poor. In this paper, a lower bound is developed that expresses the certainty in the estimate given by a particle filter in terms of the determinant of the state covariance matrix, the number of particles used, and the dimensionality of the state vector. A modified particle filter is presented in which feedback of a performance metric is actively used to modulate the number of particles used by the estimator. The utility of this feedback principle is first demonstrated by examining a fundamentally challenging case of tracking an object in a simulated 2D image sequence with varying levels of noise. The proposed algorithm is then evaluated for the task of tracking a fiducial marker in X-ray fluoroscopic images. Empirical results show that the proposed technique offers an improvement in the robustness and efficiency of particle filters while satisfying the performance expectations of demanding applications.

**Index Terms**— medicine, motion, tracking.

## 1 INTRODUCTION

In this paper, we examine the use of feedback of performance metrics to improve the robustness of tracking algorithms while satisfying the geometric performance requirements of demanding applications. Our work stems from the demand of clinical applications that require guarantees on geometric performance. In diagnostic and therapeutic interventions, intraoperative motion estimation is needed to guide interventional tools with high precision. In image-guided radiation therapy (IGRT), motion estimation is essential to improve the precision with which therapeutic radiation is delivered. Tracking objects of interest in IGRT can be challenging. For example, tracking a fiducial marker in the lung under X-ray fluoroscopy can present a heterogeneous background with nonstationary noise due to the pattern presented by ribs and breathing motion. Lack of sufficient geometric information leads to imprecision in intervention placement, an ineffective treatment, and can also cause irreversible toxicity.

If the tracking problem involves linear dynamics with additive Gaussian noise, then it can be handled well with a Kalman filter [9] as the entire density can be represented and propagated using only the mean and the variance. However, in practice there are many sources of nonlinear-

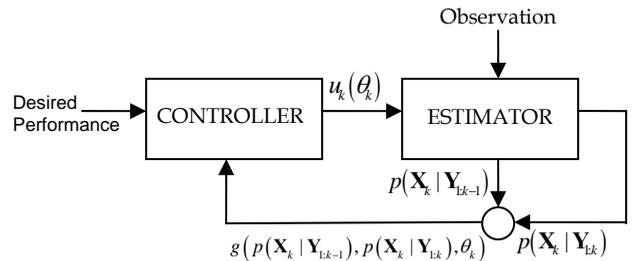


Fig. 1. System diagram illustrating feedback of performance in sequential estimation.  $u_k(\theta_k)$ ,  $p(\mathbf{X}_k | \mathbf{Y}_{t,k-1})$ , and  $p(\mathbf{X}_k | \mathbf{Y}_{tk})$  represent the controller input, and the prior and posterior probability densities of the state of the system, respectively.

ity and non-Gaussian noise. Particle filters and their variants (Condensation, Markov Chain Monte Carlo (MCMC) filter, Metropolis-Hastings importance sampling filter, Bayesian bootstrap filter, sequential importance sampling (SIS), sampling importance resampling (SIR) filter) [3,9,10,11] have emerged as promising solutions to tracking under such situations. These algorithms, however, are approximate, which converge in expected value with an asymptotically large number of particles. The number of particles,  $N$ , to use in a particle filter is an open question in computer vision and is an important decision as it affects the performance and computational expense of this estimator. A tradeoff exists between the reliability and precision of an estimate and the number of particles used, which affects the computation time. We note that while particle filters can operate in linear time in the number of particles asymptotically [22], the computational burden of the likelihood step requires in practice that we minimize the number of particles used. In this paper, we introduce a framework to control this tradeoff. Fig. 1 illustrates the proposed framework. Given a desired

- S. Siddique is with the Department of Computer Science, University of Toronto and the Radiation Medicine Program, Princess Margaret Hospital, Toronto, ON M5G 2M9. E-mail: sami.siddique@rmp.uhn.on.ca.
- E. Fiume is with the Department of Computer Science, University of Toronto, Toronto, ON M5S 2E4. E-mail: elf@dgp.toronto.edu.
- D. A. Jaffray is with the Departments of Medical Biophysics and Radiation Oncology, University of Toronto, and with the Ontario Cancer Institute / Princess Margaret Hospital, Toronto, ON M5G 2M9. E-mail: david.jaffray@rmp.uhn.on.ca.

tracking performance, as specified by a user, a controller uses estimates of the geometric performance of the tracking algorithm to dynamically update the parameters of the tracking algorithm in an attempt to maintain the desired performance within constraints.

A few attempts have been made to characterize the performance of particle filters [12, 20, 21, 22]. Carpenter et al. [22] propose an *effective sample size* (ESS) that requires running the particle filter independently over several replicates or runs. They also present an SIR filter based on stratified sampling that assigns a different proportion of the total number of particles  $N$  to each stratum. Kong et al. present an earlier definition of ESS [24]. MacCormick and M. Isard [20] introduce notions of *survival diagnostic* (similar to ESS in [24]) and *survival rate* in the context of partitioned sampling. They show how the number of particles assigned to different partitions can be varied when tracking objects in high-dimensional configuration spaces. In particular, they suggest through the task of tracking an articulated chain in simulation that partitions corresponding to parts that exhibit higher variance in dynamics be assigned more particles. In order to optimize the use of particles for high dimensional configuration spaces, Sminchisescu and Triggs [21] propose a covariance-scaled sampler that eigen-decomposes the covariance matrix from the previous time step's posterior, reconstitutes it to retain only a select few most uncertain eigendirections, and uses it to form the proposal density. For certain classes of state-space models, Rao-Blackwellization, a variance reduction method, can be applied to reduce the number of particles used [28]. This involves partitioning the state so that part of the computation can be performed analytically.

Approaches have been proposed for adapting the number of particles dynamically during tracking. In [13], use of the KL divergence between the estimated and true posterior densities of the state is made to express the minimum number of particles as a function of the number of "bins" in the true distribution. It is argued [14] that this expression is independent of the mismatch between the estimated and true distributions. An alternative expression for the minimum number of particles has been proposed in [14] that is also based on the KL divergence. However, the analysis has been based on a one dimensional state and the expression derived for the minimum number of particles is a function of the true mean value of the state. More recently, Guo and Qian [18] propose an adaptive particle filter in which they compare the ratio of ESS (as defined in [20] and [24]) to a threshold to determine if more samples are needed.

Feedback of performance in tracking applications has been proposed in computer vision literature. Erdem et al. demonstrated a framework for tracking contours of non-rigid objects that is robust to occlusion and background clutter [3, 4]. Their method employs feedback of image

cues including differences in color, edge information, motion between the object that is tracked and the background.

In this paper, we first derive an explicit bound on the performance of a particle filter based on sequential importance sampling. This bound provides valuable insight and forms the basis of a closed-loop particle filter in which feedback of geometric performance is used to control error. After establishing notation and providing a brief background of particle filters (Section 2), we present a bound on the performance of particle filters (Section 3), introduce a closed loop particle filter and evaluate it (Section 4).

## 2 BACKGROUND

Sequential Bayesian filtering provides a probabilistic framework for estimating the probability density of the state of a dynamic system based on noisy or indirect observations of the state. Let  $\mathbf{X}_k \in \mathbb{R}^d$  denote the state vector at time instant  $t_k = kT$ , where  $T \in \mathbb{R}$  is time period between successive observations with  $k \in \mathbb{N}$ , and let  $\mathbf{Y}_k \in \mathbb{R}^{d'}$  denote the observation at time  $t_k$ . Let  $\mathbf{Y}_{l:m}$  with  $l < m$  denote the sequence  $\{\mathbf{Y}_l, \dots, \mathbf{Y}_m\}$  and let  $\mathbf{X}_{l:m}$  be defined similarly. Each stage of this filtering process involves a *prediction step* at which the density  $p(\mathbf{X}_k | \mathbf{Y}_{1:k-1})$  is computed using a dynamic model, and an *update step* at which  $p(\mathbf{X}_k | \mathbf{Y}_{1:k})$  is computed using an observation model. Imposing a Markov assumption and an independence assumption,

$$\begin{aligned} \text{(i)} \quad & p(\mathbf{X}_i | \mathbf{X}_{1:i-1}) = p(\mathbf{X}_i | \mathbf{X}_{i-1}) \\ \text{(ii)} \quad & p(\mathbf{Y}_{i:k} | \mathbf{X}_i) = p(\mathbf{Y}_i | \mathbf{X}_i) p(\mathbf{Y}_{j:k} | \mathbf{X}_i) \end{aligned} \quad (1)$$

where  $i < j < k$ , the densities at the prediction and update steps can be shown to be given by [9]:

$$p(\mathbf{X}_k | \mathbf{Y}_{1:k-1}) = \int p(\mathbf{X}_k | \mathbf{X}_{k-1}) p(\mathbf{X}_{k-1} | \mathbf{Y}_{1:k-1}) d\mathbf{X}_{k-1} \quad (2)$$

and

$$p(\mathbf{X}_k | \mathbf{Y}_{1:k}) = c_k p(\mathbf{Y}_k | \mathbf{X}_k) p(\mathbf{X}_k | \mathbf{Y}_{1:k-1}) \quad (3)$$

where  $c_k$  is a normalizing constant and  $p(\mathbf{Y}_k | \mathbf{X}_k)$  is the likelihood.

Due to the inherent non-linearity, non-stationarity, and the non-Gaussian nature of many problems, the above equations are difficult to compute in closed form. Particle filters offer a solution to this problem by performing sequential Monte Carlo estimation [10, 11] based on a sampled representation of the probability density given by a set of weighted samples,

$$V_k = \left\{ \left\langle \mathbf{X}_k^{(i)}, w_k^{(i)} \right\rangle \mid i = 1, \dots, N \right\} \quad (4)$$

$$\text{with } \sum_{i=1}^N w_k^{(i)} = 1, 0 \leq w_k^{(i)} \leq 1, \forall k$$

where  $\mathbf{X}_k^{(i)}$  is a sample (or particle) drawn from the corresponding probability density,  $w_k^{(i)} \equiv w_k(\mathbf{X}_k^{(i)})$  is the weight associated with the  $i^{\text{th}}$  particle, and  $N$  is the number of particles used in the approximation. Under this representation, the posterior can be approximated using importance sampling [11] as

$$p(\mathbf{X}_k | \mathbf{Y}_{1:k}) \approx \sum_{i=1}^N w_k^{(i)} \delta(\mathbf{X}_k - \mathbf{X}_k^{(i)}) \quad (5)$$

where  $\delta(\cdot)$  is the Dirac delta function. Moreover, any moment of the posterior density can be approximated using a summation over the samples  $\mathbf{X}_k^{(i)}$ . Consider the first moment,

$$E(\mathbf{X}_k) = \int \mathbf{X}_k p(\mathbf{X}_k | \mathbf{Y}_{1:k}) d\mathbf{X}_k \quad (6)$$

which can be approximated as

$$\hat{E}_N(\mathbf{X}_k) = \sum_{i=1}^N \mathbf{X}_k^{(i)} w_k^{(i)} \quad (7)$$

where  $\mathbf{X}_k^{(i)}$  are samples drawn from the density  $p(\mathbf{X}_k | \mathbf{Y}_{1:k})$ . If the samples are independent, then it can be shown that under weak assumptions  $\hat{E}_N(\mathbf{X}_k)$  is an asymptotically unbiased estimator[12] of  $E(\mathbf{X}_k)$ , and according to the law of large numbers,  $\hat{E}_N(\mathbf{X}_k)$  will almost surely converge to  $E(\mathbf{X}_k)$  i.e.

$$\hat{E}_N(\mathbf{X}_k) \xrightarrow[N \rightarrow +\infty]{a.s.} E(\mathbf{X}_k) \quad (8)$$

Eq. 8 is an asymptotic bound with a specific inverse-root convergence as  $N$  grows arbitrarily. In practice we are interested in relatively small, finite  $N$ . From the central limit theorem, we expect the rate of convergence of a basic Monte Carlo estimate to be of the order  $O(N^{-1/2})$  [29]. However, this requires that the  $N$  particles used are statistically independent and that the function being integrated is known exactly. This is not necessarily the case in sequential Monte Carlo estimation [26, 27, 29]. According to Crisan and Doucet [27], an upper bound on the variance of the estimation error has the form  $cO(N^{-1})$ , where  $c$  is a constant. Daum and Huang [26],

however, argue that the constant in this upper bound is actually a function of the dimensionality of the state vector (approximately linear in state dimension). We address the problem of selecting an appropriate number of particles from a mathematical and practical standpoint and

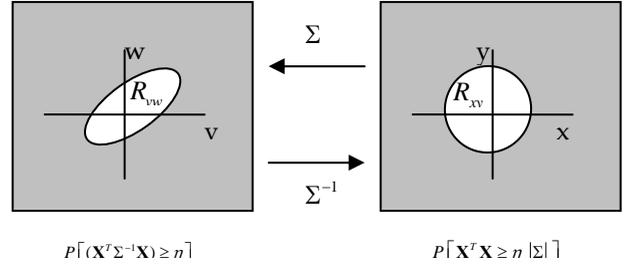


Fig. 2. Illustration of equiprobable events in  $\mathbb{R}^2$ . Regions defined by  $\mathbf{X}^T \Sigma^{-1} \mathbf{X} < \eta$  on the left and  $\mathbf{X}^T \mathbf{X} < \eta |\Sigma|$  on the right have the same area. The center of each figure represents the origin.

explore an online approach.

In the next section, we develop a probabilistic upper bound on the performance of a particle filter. We seek a practical value of  $N$  that meets the tracking requirements of a given application. As we show in the next section, this may require modulating the value of  $N$  while tracking. Throughout this paper we make use of the terms error and uncertainty. Intuitively, the term *error* refers to any deviation of an estimate from the true mean while the term *uncertainty* characterizes the dispersion of the estimate about the true mean. A formal definition of these terms is presented below.

### 3 FEEDBACK OF GEOMETRIC PERFORMANCE

We first develop a probabilistic upper bound on first order statistics that relates the uncertainty in an importance sampling based particle filter to the covariance and dimensionality of the state, and the number of particles used. This bound is then used to develop an algorithm for balancing tradeoffs between geometric precision and the number of particles used, which dictates the computational performance. The derivation of this bound is based on a generalization of the Chebyshev inequality to higher dimensions. A number of generalizations to the Chebyshev inequality have been presented in literature (e.g., [16] and [17]) that vary in the degree of tightness. Unlike these generalizations, here, a generalization is developed that is of the same form as the original inequality. The isotropic formulation makes it practical in a number of applications.

#### Lemma 1

Let  $\mathbf{X} \in \mathbb{R}^d$  be a random vector and let  $\Sigma$  denote the covariance of  $\mathbf{X}$ . Assume that  $|\Sigma| \neq 0$ , where  $|\cdot|$  denotes the determinant. For any scalar  $\eta > 0$ ,

$$P[\mathbf{X}^T \Sigma^{-1} \mathbf{X} \geq \eta] = P[\mathbf{X}^T \mathbf{X} \geq \eta |\Sigma|]$$

#### Proof:

Consider a one-to-one mapping,  $\mathbf{W} = \Sigma \mathbf{X}$ , from the re-

gion  $R_{xy}$  enclosed by  $\mathbf{X}^T \mathbf{X} < \eta$  to the region  $R_{wy}$  enclosed by  $\mathbf{X}^T \Sigma^{-1} \mathbf{X} < \eta$ . Fig. 2 illustrates a two-dimensional case. The probability of the event corresponding to  $R_{wy}$  is given by

$$P[\mathbf{X}^T \Sigma^{-1} \mathbf{X} < \eta] = \int_{R_w} f_w(\mathbf{W}) d\mathbf{W} = |\Sigma^T| \int_{R_x} f_x(\mathbf{X}) d\mathbf{X}$$

where  $|\Sigma^T|$  is the Jacobian of the mapping,  $f_w$  and  $f_x$  are characteristic functions that are equal to 1 in the unshaded regions of the left and right plots in Fig. 2, respectively, and zero otherwise. Thus, this expression is equivalent to

$$|\Sigma| \int_{R_x} f_x(\mathbf{X}) d\mathbf{X} = |\Sigma| P[\mathbf{X}^T \mathbf{X} < \eta] = P[\mathbf{X}^T \mathbf{X} < \eta | \Sigma]$$

For example, the event described by right-hand side above defines an equivalent region (disc in  $\mathbb{R}^2$ , sphere in  $\mathbb{R}^3$ , etc.) of radius  $\sqrt{\eta |\Sigma|}$ . The complements of each set of events are also equal. Thus,

$$\begin{aligned} P[\mathbf{X}^T \Sigma^{-1} \mathbf{X} \geq \eta] &= 1 - P[\mathbf{X}^T \Sigma^{-1} \mathbf{X} < \eta] \\ &= 1 - P[\mathbf{X}^T \mathbf{X} < \eta | \Sigma] = P[\mathbf{X}^T \mathbf{X} \geq \eta | \Sigma] \quad \square \end{aligned}$$

### Theorem 1

Let  $\mathbf{X} \in \mathbb{R}^d$  be a random vector and let  $\Sigma$  denote the covariance of  $\mathbf{X}$ . Assume that  $|\Sigma| \neq 0$ , for any scalar  $\varepsilon > 0$

$$P[\|\mathbf{X} - E[\mathbf{X}]\| \geq \varepsilon] \leq \frac{|\Sigma| d}{\varepsilon^2}$$

where  $\|\cdot\|$  denotes the  $L_2$  norm.

### Proof:

Let  $\mathbf{Z} \equiv \mathbf{X} - E[\mathbf{X}]$  and define  $Y$  as

$$Y \equiv \begin{cases} \eta & \text{if } \mathbf{Z}^T \Sigma^{-1} \mathbf{Z} \geq \eta \\ 0 & \text{otherwise} \end{cases}$$

Clearly,

$$Y \leq \mathbf{Z}^T \Sigma^{-1} \mathbf{Z}.$$

The expected value of both sides is

$$\eta P[\mathbf{Z}^T \Sigma^{-1} \mathbf{Z} \geq \eta] \leq E[\mathbf{Z}^T \Sigma^{-1} \mathbf{Z}] \quad (9)$$

Since  $\mathbf{Z}^T \Sigma^{-1} \mathbf{Z}$  is a scalar,

$$\begin{aligned} E[\mathbf{Z}^T \Sigma^{-1} \mathbf{Z}] &= E[\text{tr}(\mathbf{Z}^T \Sigma^{-1} \mathbf{Z})] = E[\text{tr}(\mathbf{Z} \mathbf{Z}^T \Sigma^{-1})] \\ &= \text{tr}(E[\mathbf{Z} \mathbf{Z}^T] \Sigma^{-1}) = \text{tr}(\Sigma \Sigma^{-1}) = \text{tr}(\mathbf{I}) = d. \end{aligned}$$

Thus, (9) becomes

$$P[\mathbf{Z}^T \Sigma^{-1} \mathbf{Z} \geq \eta] \leq \frac{d}{\eta}.$$

From Lemma 1, it follows that

$$P[\mathbf{Z}^T \Sigma^{-1} \mathbf{Z} \geq \eta] = P[\mathbf{Z}^T \mathbf{Z} \geq \eta |\Sigma|] \leq \frac{d}{\eta}. \quad (10)$$

Defining  $\varepsilon^2 \equiv \eta |\Sigma|$ , and substituting into (10),

Initialize  $\hat{p}(\mathbf{X}_k | \mathbf{Y}_{1:k}) = \sum_{i=1}^N w_k^{(i)} \delta(\mathbf{X}_k - \mathbf{X}_k^{(i)})$

For each time step,  $k$ :

for  $n = 1:N_k$

draw  $\mathbf{X}_k^{(n)} \sim \hat{p}(\mathbf{X}_{k-1} | \mathbf{Y}_{1:k-1})$

predict by sampling from  $\hat{p}(\mathbf{X}_k | \mathbf{X}_{k-1}^{(n)})$

update  $w_k^{(n)} = \hat{p}(\mathbf{Y}_k | \mathbf{X}_k^{(n)})$

end for

normalize  $w_k$  so that  $\sum_{i=1}^N w_k^{(i)} = 1$

set  $\xi_k \equiv |\hat{\Sigma}_N|_k - |\hat{\Sigma}_N|_{ref}$

set  $N_{k+1} = \max(\min((\kappa_p \xi_k + \kappa_l \sum_{i=1}^k \xi_i + N_{ref}), N_{max}), N_{min})$

Fig. 3. Particle filtering algorithm with feedback of  $|\hat{\Sigma}|$  to modulate  $N$  using a reference uncertainty,  $|\hat{\Sigma}_N|_{ref}$ .

$$P[\mathbf{Z}^T \mathbf{Z} \geq \varepsilon^2] \leq \frac{d |\Sigma|}{\varepsilon^2}$$

$$P[\|\mathbf{Z}\|^2 \geq \varepsilon^2] \leq \frac{d |\Sigma|}{\varepsilon^2}$$

$$P[\|\mathbf{X} - E[\mathbf{X}]\| \geq \varepsilon] \leq \frac{d |\Sigma|}{\varepsilon^2}.$$

□

### Theorem 2

Let  $\mathbf{X} \in \mathbb{R}^d$  be a random vector with mean  $\boldsymbol{\mu} \equiv E[\mathbf{X}_k]$  and covariance  $\Sigma \equiv \text{Cov}[\mathbf{X}_k]$  with  $|\Sigma| \neq 0$ . Consider the estimated mean of  $\mathbf{X}$  as given by a particle filter based on sequential importance sampling using  $N$  particles, for sufficiently large  $N$ ,

$$\hat{E}_N(\mathbf{X}_k) = \sum_{i=1}^N \mathbf{X}_k^{(i)} w_k^{(i)}$$

Assume that the mean,  $\bar{\boldsymbol{\mu}}$ , and the variance,  $\hat{\Sigma}$ , of this estimate exist and are finite:

$$\bar{\boldsymbol{\mu}} \equiv E[\hat{E}_N[\mathbf{X}_k]]$$

$$\hat{\Sigma}_N \equiv E\left[\left(\hat{E}_N(\mathbf{X}_k) - \bar{\boldsymbol{\mu}}\right)\left(\hat{E}_N(\mathbf{X}_k) - \bar{\boldsymbol{\mu}}\right)^T\right]$$

Then for any scalar  $\varepsilon > 0$ ,

$$P\left[\left\|\hat{E}_N(\mathbf{X}_k) - \bar{\boldsymbol{\mu}}\right\| < \varepsilon\right] \geq 1 - \frac{|\hat{\Sigma}_N| d}{\varepsilon^2} = 1 - \frac{|\Sigma| d}{N^d \varepsilon^2}.$$

### Proof:

Under weak assumptions on  $\boldsymbol{\mu}$  and  $\Sigma$  discussed extensively in [12], particularly Theorems 1 and 2,

$$\bar{\boldsymbol{\mu}} \rightarrow \boldsymbol{\mu} \quad (11)$$

$$\hat{\Sigma}_N \rightarrow \frac{\Sigma}{N} \quad (12)$$

From Theorem 1,

$$P \left[ \left\| \hat{E}_N(\mathbf{X}_k) - \bar{\boldsymbol{\mu}} \right\| \geq \varepsilon \right] \leq \frac{|\hat{\Sigma}_N| d}{\varepsilon^2}.$$

Using (11) and (12),

$$P \left[ \left\| \hat{E}_N(\mathbf{X}_k) - \boldsymbol{\mu} \right\| \geq \varepsilon \right] \leq \frac{\frac{1}{N} |\Sigma| d}{\varepsilon^2}.$$

The complement of this event is given by

$$P \left[ \left\| \hat{E}_N(\mathbf{X}_k) - \boldsymbol{\mu} \right\| < \varepsilon \right] \geq 1 - \frac{|\Sigma| d}{N^d \varepsilon^2} \quad \square$$

According to Theorem 2, for any given acceptable tracking precision,  $\varepsilon$ , the probability that the estimated mean is within  $\varepsilon$  of the true mean can be expressed as a function of the determinant of the covariance of the state, the dimensionality of the state, and the number of particles used. Many signals of interest exhibit nonstationarity. For instance, when tracking humans as the subject walks through a shadow or gets partially occluded by an object, the uncertainty in the location of the subject may increase. Accordingly,  $|\Sigma|$  increases and

$P \left[ \left\| \hat{E}_N(\mathbf{X}_k) - \boldsymbol{\mu} \right\| < \varepsilon \right]$  decreases. In practice, by modulating

$N$  to match the change in  $|\Sigma|$ , the probability bound,

$P \left[ \left\| \hat{E}_N(\mathbf{X}_k) - \boldsymbol{\mu} \right\| < \varepsilon \right]$ , can be maintained.

The quantity  $|\Sigma|$  is inaccessible but its estimate  $|\hat{\Sigma}_N|$  can be used as feedback (of uncertainty). An algorithm is presented in Fig. 3 that uses a proportional integral (PI) controller to modulate  $N$ . The proportional ( $\kappa_p$ ) and integral ( $\kappa_i$ ) gains of this PI controller were hand tuned.

When implementing the proposed algorithm, cases in which  $\hat{\Sigma}_N$  becomes singular must be handled. In these cases the determinant of this matrix will underestimate uncertainty. To avoid this situation,  $\hat{\Sigma}_N$  can be regularized by mixing it with an identity matrix. It should be noted that Theorem 2 presents a sufficient condition and requires that adequate observation and dynamic models are specified. An analysis of the controllability and stability of the controller is beyond the scope of this paper.

The proposed method uses just the number of particles required to maintain the error at the desired level. It offers a designer a more robust practical algorithm.

We apply the proposed algorithm and show some empirical results in Section 4. The analysis confirms a poten-

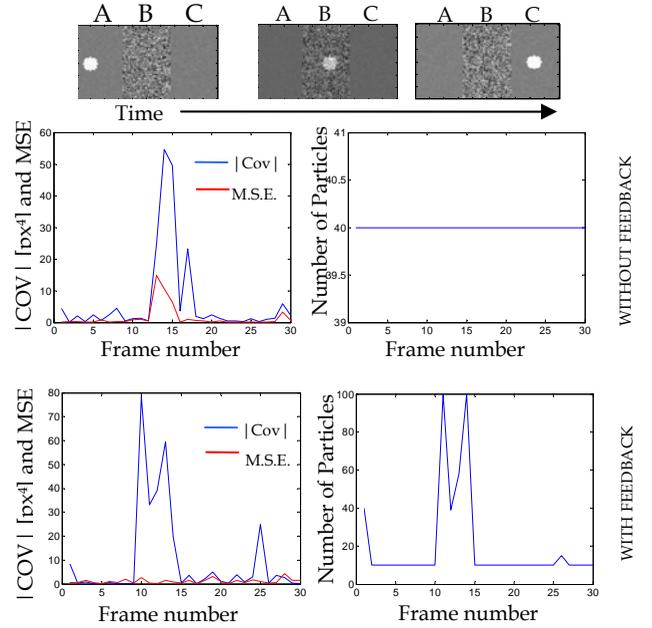


Fig. 4. Comparison of the performance of a particle filter with and without feedback of  $|\hat{\Sigma}|$ .

tially powerful technique for dynamic optimization of the number of particles to use in order to regulate the tradeoff between tracking uncertainty and computation time.

## 4 EXPERIMENTAL RESULTS AND DISCUSSION

Two situational examples are used to demonstrate the proposed particle filtering algorithm. The cases examined are clinically significant as the images used in the simulations are representative of tracking fiducials under X-ray fluoroscopy in an X-ray quantum noise limited system [30, 31].

The first example uses toy but representative data and involves the task of tracking a circular disc (of radius 5 pixels) in 2D image sequences (100x60 pixels; 30 frames). The state of interest is the location of the centre of the disc ( $x$  and  $y$  coordinates of the image). Feedback of the determinant of the covariance matrix,  $|\hat{\Sigma}|$ , as defined in Section 3.1, was used to modulate the number of particles,  $N$ , in an attempt to balance tradeoffs between computation time and certainty. Fig. 4 shows a few frames from the sequence used. The disc starts in region A in which an additive white Gaussian noise (AWGN) with an SNR of -10dB is present. It then moves into region B (-30dB; frames 9 to 20) followed by region C (-10 dB). An edge-based observation model similar to that used in [3] was used. The dynamics were assumed to be explained by a Brownian motion model with drift. The performance of the particle filter as the object moves in and out of the noisy region B, both with and without feedback of  $|\hat{\Sigma}|$ , is compared in Fig. 4. (The particle filter algorithms used here do not employ loss of diversity correction [19]). Both

estimators were initialized with 40 particles. In the case of the conventional particle filter without feedback,  $|\hat{\Sigma}|$  increases in the noisy region B and correspondingly the mean square error (M.S.E.) also increases to a maximum of 15 px<sup>2</sup> (pixels<sup>2</sup>). The number of particles is fixed at 40 in this conventional case. On the other hand, in the case of the particle filter with feedback,  $|\hat{\Sigma}|$  increases as the disc enters region B. Correspondingly, the number of particles increases in an attempt to maintain the required level of certainty. Thus, there is no significant increase in the M.S.E. error while the disc is in region B. When the disc exits region B, negative feedback of the fall in  $|\hat{\Sigma}|$  causes the number of particles to fall. Hence, the proposed approach modulates  $N$  in order to maintain the required certainty. Note that the value of  $|\hat{\Sigma}|$  as the disc enters region B is higher in the case of feedback because fewer particles (10 as opposed to 40) are used in this case prior to entering region B. Feedback thus provides a convenient means to select between targeting precision and computation speed.

The second example for demonstrating the proposed particle filtering algorithm uses real data. An experiment similar to the simulation described above was performed using an X-ray fluoroscopic imaging sequence. Compared to the additive Gaussian noise model used in the simulation above, the X-ray fluoroscopic images used here correspond to a more complex noise model governed by Poisson statistics [25].

A discrete-time sequence of 2D X-ray images of a fiducial marker (metallic sphere 2.0 mm in diameter - Beekley Y-Spot 102) was acquired using the setup shown in Fig. 5. (The X-ray projection of the marker is a sphere that spans 8 pixels in diameter). The marker was attached to a Styrofoam block that was moved using a programmable high precision linear motion stage. The images were acquired at 1 fps on a Varian Paxscan 4030A (2048x1536) under the kV imaging geometry typically found on kV enabled medical linear accelerator systems (1000 mm source-to-imaging-axis distance, 1550 mm source-to-detector distance). The images were acquired at 100 kVp and 0.5 mAs. A total of 50 frames were acquired.

To simulate a simple form of a typical motion exhibited by fiducial markers and interventional tools, the marker attached to the linear stage was moved according to Brownian motion dynamics as given by

$$s_{k+1} = s_k + U_k,$$

where  $k$  denotes the frame number,  $s_k$  is the position of the marker in world coordinates and  $U_k$  is Gaussian noise,  $N(\mu, \sigma^2)$ , with  $\mu = \frac{1}{2}$  cm and  $\sigma^2 = \frac{1}{8}$  cm. The proposed algorithm's ability to handle changes in ob-

served images was tested by moving the marker behind a set of acrylic slabs (each 9.2 mm thick) as shown in Fig. 5. Increasing the number of slabs used increases the attenuation of X-rays and decreases the contrast-to-noise ratio (CNR) of the marker in the region behind the slabs. CNR is defined here as:

$$CNR = \frac{|m_{marker} - m_{bg}|}{\sigma_{bg}}$$

where  $m_{marker}$  is the average value of the pixels in the region corresponding to the marker and  $m_{bg}$  is the average value of the pixels in a region immediately outside the marker but within the slab region. The CNR corresponding to the different number of slabs used is shown in Table 1.

The setup described above was used to generate several image sequences, each corresponding to a different number of slabs. The particle filter described earlier was used to track the marker as it traversed through regions

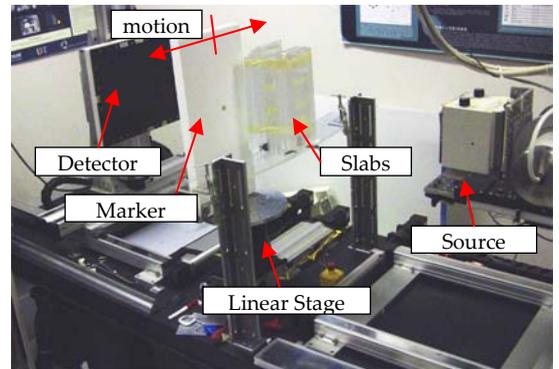
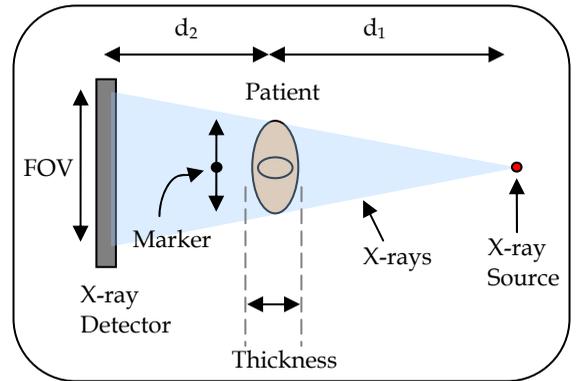


Fig. 5. (a): Diagram showing aerial view of setup. Geometry simulated in this study was taken from radiation treatment systems that are proposed for tracking fiducial markers ( $d_1 = 1000$  mm ;  $d_2 = 1550$  mm; thickness = 0 to 26 acrylic slabs of 9.2 mm thickness each; Field-of-view, FOV = 1024 pixels). The marker was posterior to the slabs to allow motion. (b): Experimental setup used to evaluate the performance of the proposed algorithm that incorporates feedback to modulate the number of particles used. Shown here is the configuration of the X-ray source, X-ray detector, fiducial marker, acrylic slabs, and the linear motion stage used to move the marker. Right: A close up view showing the configuration of the slabs used.

of low noise (corresponding to air) and high noise (corresponding to the region behind the slabs). For each image sequence, tracking was executed multiple times with different random seeds for the particle filter (each of these referred to as a *run*). Since a particle filter is a stochastic estimation technique, each run was repeated 20 times to evaluate the reproducibility of the estimate. The particle filter was initialized by specifying a uniform prior over a region containing the centre of the sphere. As in the simulation above, a Brownian motion model with drift and an edge-based observation model were used.

The marker was first tracked without using feedback of uncertainty. The metrics used for analysis were:

(i) Average  $L^1$  error, defined here as the absolute value of the difference between the estimated location of the marker along the horizontal axis and ground truth:

$$\frac{1}{KR} \sum_r \sum_k |\hat{u}_{k,r} - \tilde{u}_k|$$

where  $\hat{u}$  is the estimated location of the object of interest along the horizontal axis corresponding to frame  $k$  and run  $r$ ,  $\tilde{u}$  is the ground truth,  $K$  and  $R$  are the total number of frames and the total number of runs, respectively.

(ii) Reproducibility, defined as the expected value of the standard deviation in the estimated location over all runs:

$$\frac{1}{K} \sum_k \left( \sqrt{\frac{1}{R} \sum_r (\hat{u}_{k,r} - E_r[\hat{u}_{k,r}])^2} \right)$$

where  $E_r[\hat{u}_{k,r}]$  is the expected value of the estimated location along the horizontal axis over all  $R$  runs.

The metrics defined above were computed for all combinations of number of particles and number of slabs over all 20 runs and are presented in Table 1. Also shown in Table 1 is the average estimated uncertainty for each combination calculated as follows:

$$\frac{1}{KR} \sum_r \sum_k |\hat{\Sigma}_{N|k,r}|$$

where  $|\hat{\Sigma}_{N|k,r}|$  is the determinant of the covariance matrix corresponding to the posterior density of the  $k^{\text{th}}$  frame and  $r^{\text{th}}$  run. The analysis presented in Table 1 was performed on frames in which the marker was behind the slabs. Robust linear least squares fit of the average estimated uncertainty and the average  $L^1$  error to a linear polynomial reveals an  $R^2$  (coefficient of determination) value of 0.95.

The proposed framework employing feedback (Fig. 3) was then evaluated at operating precisions of 1.0 pixel and 3.0 pixels. To present a considerable challenge, the image sequence with 23 slabs was used. As in the case without feedback, each run was repeated 20 times. A P-

controller was used to modulate the number of particles used. The performance of the proposed algorithm was compared with a particle filter that did not use feedback. Runs from sequences with feedback corresponding to operating precisions of 1.0 pixel and 3.0 pixels are shown in Fig. 6(b) and Fig. 6(d), respectively. Shown in each quadrant of Fig. 6 are the determinant of the covariance matrix,  $|\hat{\Sigma}_{N|k}|$ , calculated from the updated probability density of the state, the  $L^1$  error per frame, and the number of particles used plotted as a function of frames over an exemplary run. Shaded regions in the figure indicate frames in which the marker falls behind the slabs and correspond to lower CNR. The feedback cases shown in Fig. 6(b) and Fig. 6(d) were compared with similar cases of a conventional particle filter that did not employ feedback. Fig. 6(a) and Fig. 6(c) show an exemplary run of a particle filter operating using a fixed number of particles: 20 and 100, respectively. These cases were chosen for comparison because the feedback cases operate using 20 and 75 particles when outside the noisy region behind the slabs. Also shown in Fig. 6(e) are the average  $L^1$  error, the  $L^\infty$  error and the average number of particles used per frame for each case. Analysis was performed on frames 14-50 and averaged over all 20 runs. These frames include frames in which the marker is in the noisy region behind the slabs and those in which the marker is in air. Tracking was initialized at frame 7 with 100 particles. The marker was behind the slabs in frames 8-37.

It can be seen from Table 1 that as the number of particles used increases, the accuracy improves as expected. What is more important to note is that as the noise increases with the addition of more slabs, the accuracy decreases in general. (The cases of 11 slabs and 17 slabs show a slightly higher error. This could be due to a slight shift in the marker away from ground truth.) As noise increases while tracking an object, the number of particles used can be increased to achieve the required operating accuracy, provided the object is still "visible". For the case of 26 slabs, due to the low SNR the marker is lost. The strong correlation between the average estimated uncertainty and the average  $L^1$  error makes the estimated uncertainty a good candidate for a metric of feedback for the given application.

Comparing the performance of the particle filter for the cases presented in Fig. 6(a-b), it can be observed that for the case without feedback, when the marker enters the noisy slab region, the  $L^1$  error rapidly increases and the estimator loses track of the marker. The feedback case, however, is able to use feedback of the increased  $|\hat{\Sigma}_{N|k}|$  to increase the number of particles in order to control the error. Similarly, for the cases shown in Fig. 6(c-d), while the error increases in the slab region for the case without feedback, the proposed algorithm is able to maintain the error. When the marker exits the slab region, the pro-

posed algorithm is able to use feedback of the lowered uncertainty estimate to reduce the number of particles used. From Fig. 6(e), it can be seen that the average  $L^1$  error and the  $L^\infty$  are lower for the proposed approach. Instead of specifying the number of particles to use, this approach lets the designer specify the error tolerance of the system leaving the algorithm to choose the number of particles. It is expected that the greater the disparity between regions, the more useful the proposed algorithm would be. Optimal controller designs could give better results.

Theorem 2 relates the first order estimate of the true state to the determinant of a covariance matrix. It is expected that higher order estimates are related to the entropy of the state. It should be noted that a logarithmic relation exists between the determinant of a covariance matrix ( $C$ ) and entropy ( $S$ ),

$$S \propto \log |C|. \quad (13)$$

For example, for a Gaussian distribution, this relation is given by,

$$S = \frac{1}{2} \ln |C| + \frac{d}{2} \ln(2\pi e)$$

where  $d$  is the dimensionality of the state vector. In [15], through the use of the Legendre series expansion, it is shown that the entropy is related to the higher order moments of a distribution. It is not surprising to expect the minimum number of particles required for convergence to be a function of the entropy of the state. For instance, if a small image is being analyzed, fewer particles suffice as ‘‘asymptotic’’, while a larger image would require more particles. The maximum entropy of a state (achieved when all possible events are equiprobable) is given by the log of the cardinality of the state space. For example, for an image of size  $U$ -by- $V$  and a state  $(u_i, v_i) \in (1, 2, \dots, U) \times (1, 2, \dots, V)$ , a maximum of  $UV$  appropriately placed particles would be required: From Theorem 2 which suggests that  $N$  be chosen such that,  $N \propto |C|$  and from (13) above,  $N \propto e^S$ . Thus,  $\max N \propto e^{\max S} = e^{\log UV} = UV$ . The optimal value of  $N$  is affected by the discretization of the state space and is the subject of ongoing research.

In the algorithm presented in Fig. 3, the controller may introduce a lag; when the object being tracked transitions from the region of low noise to the region of high noise, there is a delay before  $N$  is increased to maintain the tracking uncertainty. This creates jumps in the RMSE at transition points. However, it is expected that if the target is in a neighborhood with certain noise characteristics at a given frame, it will be in neighborhoods with similar noise characteristics in the immediately following frames. Moreover, unlike the simulations in this study in which compositing of regions of low and high noise creates an artificial boundary that would otherwise not be observed, in practice changes in noise characteristics are expected to be less drastic. In a clinical context, when tracking objects

in the lung, if such boundaries are present, the algorithm could learn where such boundaries exist over a few breathing cycles and use this information along with the predicted location of the target to increase  $N$  just prior to entering a noise field. By learning the background context over a few breathing cycles, the controller would be better able to adapt the value of  $N$ . It is expected that the greater the variation in noise environments, the more useful the more useful the proposed algorithm would be. Better control algorithms with faster response times can be designed to address latency concerns, but we leave this problem to future work.

Note that a changing value of  $N$  implies a variable computation time. This may be an issue in certain applications when the estimator’s computation time exceeds the limited computational resources of a system. In radiotherapy, however, a variable computation time may be acceptable; if the estimator is unable to meet the desired targeting uncertainty or is unable to process an estimation step in time, the firing of the treatment beam can be delayed to the next available time step with an acceptable level of estimated uncertainty.

The tightness of the inequality presented in Theorem 2 and its relation to the Cramer-Rao bound need to be explored.

## 5 CONCLUSION

In this paper, it has been shown that the optimal choice of the number of particles in a particle filter depends on the noise in the image and the performance requirements of the application. A framework has been proposed in which the designer specifies an accuracy tolerance and the algorithm modulates the number of particles to use in order to meet the performance criteria using as many particles as necessary. This approach allows tradeoffs between the various requirements of an application to be balanced. Theorem 2 presents a sufficient condition; if there are modeling deficiencies in observation or dynamic models, then these can potentially be corrected adaptively while tracking. This approach would allow robustness to modeling imperfections and is the subject of future work.

	Number of Slabs							
	0	8	11	14	17	20	23	26
<b>CNR</b>	86.04	18.63	19.34	14.80	8.14	5.57	3.03	2.14
<b>Thickness (mm)</b>	0	73.6	101.2	128.8	156.4	184	211.6	239.2

**(a) Average  $L^1$  error (pixels) (Accuracy)**

No. of Particles	Number of Slabs							
	0	8	11	14	17	20	23	26
<b>10</b>	12.57	27.47	24.05	20.62	29.08	25.07	28.67	19.05
<b>20</b>	4.79	12.43	12.33	12.99	8.34	12.02	25.94	21.63
<b>50</b>	0.81	2.79	2.04	1.09	1.33	3.81	15.46	22.20
<b>100</b>	0.40	0.51	0.71	0.72	0.85	0.82	4.26	20.98
<b>200</b>	0.32	0.45	0.60	0.59	0.79	0.63	1.72	18.37
<b>500</b>	0.28	0.38	0.53	0.49	0.65	0.54	1.63	17.26
<b>1000</b>	0.26	0.33	0.52	0.45	0.66	0.51	1.65	16.10

**(b) Reproducibility (pixels)**

No. of Particles	Number of Slabs							
	0	8	11	14	17	20	23	26
<b>10</b>	12.62	24.22	25.91	22.43	22.12	15.74	20.51	16.89
<b>20</b>	4.87	10.14	15.96	8.86	5.98	12.98	15.95	14.96
<b>50</b>	1.25	2.57	2.47	1.38	1.38	5.54	11.10	9.99
<b>100</b>	0.32	0.53	0.82	0.70	0.92	1.29	7.20	6.26
<b>200</b>	0.19	0.29	0.55	0.38	0.51	0.88	2.81	3.44
<b>500</b>	0.13	0.24	0.30	0.23	0.39	0.53	1.63	3.32
<b>1000</b>	0.09	0.13	0.25	0.18	0.35	0.52	1.39	2.08

**(c) Average Estimated Uncertainty (pixels)**

No. of Particles	Number of Slabs							
	0	8	11	14	17	20	23	26
<b>10</b>	355.26	166.61	68.15	81.95	78.11	105.43	73.47	109.56
<b>20</b>	98.25	53.13	65.44	76.62	50.83	52.59	128.14	155.43
<b>50</b>	4.08	25.07	11.24	5.36	6.77	30.13	142.05	304.79
<b>100</b>	0.56	1.98	2.30	3.16	3.57	5.83	45.09	460.25
<b>200</b>	0.25	1.74	2.13	2.25	3.39	4.86	23.16	479.53
<b>500</b>	0.26	1.56	2.05	2.23	3.06	4.90	25.48	1182.61
<b>1000</b>	0.23	1.41	2.21	2.20	3.75	5.25	26.75	935.62

Table 1. Results of tracking a fiducial marker without feedback of uncertainty over a range of slab thicknesses using different number of particles. The average  $L^1$  error (a) and the average estimated uncertainty (c) over frames in which the marker falls behind the slabs, averaged over all 20 runs. (b) Reproducibility across all 20 runs.

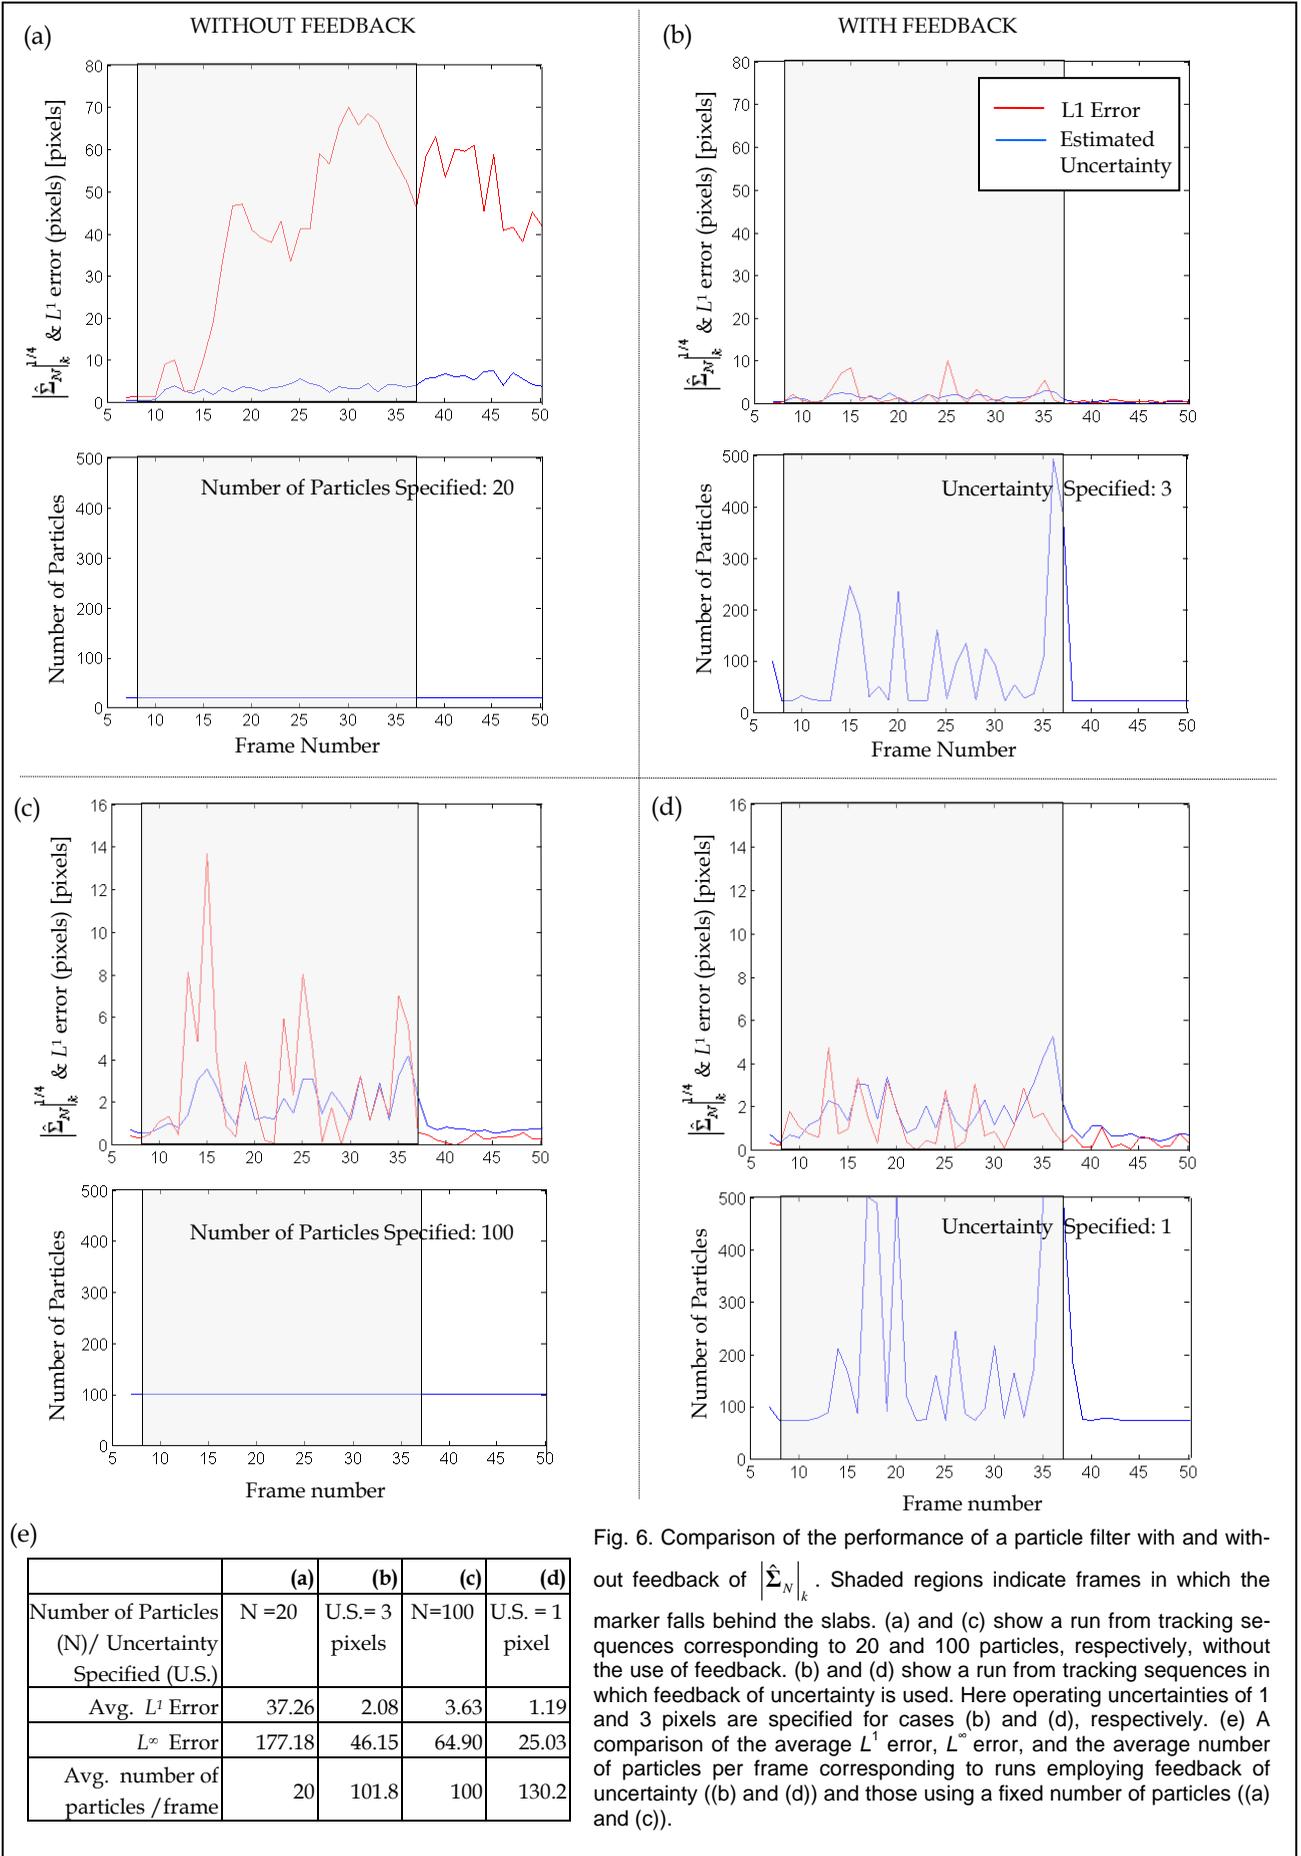


Fig. 6. Comparison of the performance of a particle filter with and without feedback of  $\left| \hat{\Sigma}_N \right|_k$ . Shaded regions indicate frames in which the marker falls behind the slabs. (a) and (c) show a run from tracking sequences corresponding to 20 and 100 particles, respectively, without the use of feedback. (b) and (d) show a run from tracking sequences in which feedback of uncertainty is used. Here operating uncertainties of 1 and 3 pixels are specified for cases (b) and (d), respectively. (e) A comparison of the average  $L^1$  error,  $L^\infty$  error, and the average number of particles per frame corresponding to runs employing feedback of uncertainty ((b) and (d)) and those using a fixed number of particles ((a) and (c)).

## ACKNOWLEDGMENT

This work was supported in part by the National Institute on Aging/National Institutes of Health under grant R33 AG019381, Elekta Oncology Systems, and through the Fidani Family Chair in Radiation Physics at Princess Margaret Hospital.

## REFERENCES

- [1] K. P. Kording and D. M. Wolpert, "Bayesian decision theory in sensorimotor control," *Trends Cogn. Sci.*, vol. 10, pp. 319-326, Jul. 2006.
- [2] K. P. Kording and D. M. Wolpert, "Bayesian integration in sensorimotor learning," *Nature*, vol. 427, pp. 244-247, Jan 15. 2004.
- [3] Isard, M. and Blake, A., "Conditional Density Propagation for Visual Tracking," *Int. J. Computer Vision*, 29, 1, 5-28, 1998.
- [4] C. Erdem, A. Tekalp and B. Sankur, "Video object tracking with feedback of performance measures," *IEEE Transactions Circuits and Systems for Video Technology*, vol. 13, pp. 310, April 2003. 2003.
- [5] J. Y. Aloimonos, I. Weiss and A. Bandopadhyay, "Active Vision," *International Journal on Computer Vision*, vol. Vol.1, pp. 333-333- 356, 1987.
- [6] R. Bajcsy, "Active Perception," *Proceedings of the IEEE*, vol. 76, pp. 996-1005, 1988.
- [7] D. Ballard, "Animate Vision," *Artificial Intelligence*, vol. 48, pp. 57-86, 1991.
- [8] H. I. Christensen and P. L. Corke, "Visual Servoing," *I. J. Robotic Res.*, vol. 22, pp. 779-780, 2003.
- [9] D. A. Forsyth and J. Ponce, *Computer Vision: A Modern Approach*. {Prentice Hall}, 2002,
- [10] J. M. Hammersley and K. W. Morton, "Poor Man's Monte Carlo," *J. Royal Statistical Soc. Series B (Methodological)*, vol. 16, pp. 23--38, 1954.
- [11] A. Doucet, N. d. Freitas and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Springer, 2001, XIV, pp. 581.
- [12] J. Geweke, "Bayesian Inference in Econometric Models Using Monte Carlo Integration," *Econometrica*, vol. 57, pp. 1317-1339, 1989.
- [13] D. Fox, "Adapting the Sample Size in Particle Filters Through KLD-Sampling," *I. J. Robotic Res.*, vol. 22, pp. 985-1004, 2003.
- [14] A. Soto. Self adaptive particle filter. Presented at Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, Morgan Kaufmann, 2005, pp 1398-1403.
- [15] N. Ebrahimi, E. Maasoumi and E. S. Soofi, "Ordering univariate distributions by entropy and variance," *J. Econ.*, vol. 90, pp. 317-336, June. 1999.
- [16] X. Sun, "On a generalization of the Chebyshev inequality," *Acta Mathematica Hungarica*, September, 1996, vol.72, number 3, pp.269-281.
- [17] D. Monhor, "A Chebyshev Inequality For Multivariate Normal Distribution," *Probab.Eng.Inf.Sci.*, vol. 21, pp. 289-300, 2007.
- [18] F. Guo and G. Qian, "3D arm movement tracking using adaptive particle filter," in *Image Processing (ICIP)*, 16th IEEE International Conference on Image Processing (ICIP), 2009, pp. 4069-4072.
- [19] C. Musso, N. Oudjane and F. L. Gland, "Improving regularized particle filters. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, chapter 12, pages 247--271, Springer--Verlag, New York, 2001," pp. 247--271, 2001.
- [20] J. MacCormick and M. Isard, "Partitioned sampling, articulated objects, and interface-quality hand tracking," in *ECCV '00: Proceedings of the 6th European Conference on Computer Vision-Part II*, 2000, pp. 3-19.
- [21] C. Sminchisescu and B. Triggs, "Covariance Scaled Sampling for Monocular 3D Body Tracking," *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 1, pp. 447, 2001.
- [22] J. Carpenter, P. Clifford and P. Fearnhead, "Improved particle filter for nonlinear problems," *Radar, Sonar and Navigation, IEE Proceedings -*, vol. 146, pp. 2-7, 1999.
- [23] Y. Chen and Y. Rui, "Real-time speaker tracking using particle filter sensor fusion," *Proceedings of the IEEE*, vol. 92, pp. 485-494, 2004.
- [24] A. Kong, J. S. Liu, and W. H. Hong, "Sequential imputations and Bayesian missing data problems", *Journal of the American Statistical Association*, vol. 89, no. 425, pp. 278-288, 1994.
- [25] H. E. Johns and J. R. Cunningham, *Physics of Radiology*; 4th Ed. Hingham, MA: Charles River Media, 1983.
- [26] F. Daum and J. Huang, "Curse of dimensionality and particle filters," in *Aerospace Conference, 2003. Proceedings. 2003 IEEE*, 2003, pp. 4\_1979-4\_1993.
- [27] D. Crisan and A. Doucet, "A survey of convergence results on particle filtering methods for practitioners," *Signal Processing, IEEE Transactions on*, vol. 50, pp. 736-746, 2002.
- [28] Z. Khan, T. Balch and F. Dellaert, "A Rao-Blackwellized Particle Filter for EigenTracking," *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 2, pp. 980-986, 2004.
- [29] B. Ristic, S. Arulampalam and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2004.
- [30] J. Wong, D. A. Jaffray, and F. Yin, "Task Group -104: In-Room kV Imaging for Image-Guided Radiation Therapy", *AAPM*, July 2010.
- [31] S. Siddique, E., Fiume, and D. A. Jaffray, "Minimizing Dose during Fluoroscopic Tracking through Geometric Performance Feedback", *Med Phys*. May 2011.

**Sami Siddique** is a PhD student at the Department of Computer Science, University of Toronto and is affiliated with the Radiation Medicine Program at Princess Margaret Hospital, Toronto.

**Eugene Fiume** Eugene Fiume is Professor and past Chair of the Department of Computer Science at the University of Toronto, where he also co-directs the Dynamic Graphics Project. Following his B.Math. degree from the University of Waterloo and M.Sc. and Ph.D. degrees from the University of Toronto, he was an NSERC Postdoctoral Fellow and Maitre Assistant at the University of Geneva, Switzerland. He was awarded an NSERC University Research Fellowship in 1987 and returned to the University of Toronto to a faculty position. His past board memberships include the Scientific Advisory Board of GMD (Germany, now Fraunhofer); the Board of Directors of Truespectra Inc. (Toronto, acquired by Scene7 Adobe), and CITO (Ottawa); the Advisory Boards of CastleHill Ventures, PlateSpin, BitFlash, OctigaBay Systems (acquired by Cray Systems), the Max-Planck Center for Visual Computing and Communication (Germany), NGRain Corporation (Vancouver), View22 Technology Inc. (Richmond Hill), and the Department of CSE at the Hong Kong University of Science and Technology; and the Executive Advisory Board of the IBM Lab in Toronto. He currently sits on the Board of Directors of

Tucows Inc. (Toronto). He has published two books and (co-)authored over 115 papers on many aspects of computer graphics and simulation.

**David A. Jaffray** is the Head of Radiation Physics at Princess Margaret Hospital and a Senior Scientist of the Ontario Cancer Institute. Dr. David Jaffray graduated from the University of Alberta, Edmonton, Alberta, Canada with a BSc in Physics in 1988 and completed his PhD in the Department of Medical Biophysics at the University of Western Ontario, London, Ontario in 1994. The following eight years were spent as a Clinical Physicist in the Department of Radiation Oncology at William Beaumont Hospital, Royal Oak, Michigan. Dr. Jaffray became a Board Certified Medical Physicist (ABMP-Radiation Oncology) in 1999. In 2002, Dr. Jaffray joined the Princess Margaret Hospital in Toronto, Canada as Head of Radiation Physics and Senior Scientist within the Ontario Cancer Institute. David holds the Fidani Chair in Radiation Therapy Physics and is a principal in the Image-guided Therapy Group of the University Health Network. He has an appointment as Vice Chair and Professor in the Departments of Radiation Oncology, Medical Biophysics, Institute of Biomaterials and Biomedical Engineering (IBBME) at the University of Toronto. Dr. Jaffray holds numerous peer-reviewed and industry sponsored research grants in the field of image-guided radiation therapy. His current research interests focus on the development of novel approaches to directing and applying radiation therapy and advancing the standard of radiation therapy delivery.