Minimizing dose during fluoroscopic tracking through geometric performance feedback

S. Siddique^{a)}

Princess Margaret Hospital/Ontario Cancer Research Institute, Toronto, Ontario M5G 2M9, Canada and
Department of Computer Science, University of Toronto, Toronto, Ontario M5S 3G4, Canada

E. Fiume

Department of Computer Science, University of Toronto, Toronto, Ontario M5S 3G4, Canada

D. A. Jaffray^{b)}

Princess Margaret Hospital/Ontario Cancer Research Institute, Toronto, Ontario M5G 2M9, Canada and
Department of Radiation Oncology and Department of Medical Biophysics, University of Toronto,
Toronto, Ontario M5S 3E2, Canada

(Received 16 June 2010; revised 3 November 2010; accepted for publication 10 February 2011;
published 29 April 2011)

Purpose: There is a growing concern regarding the dose delivered during x-ray fluoroscopy guided procedures, particularly in interventional cardiology and neuroradiology, and in real-time tumor tracking radiotherapy and radiosurgery. Many of these procedures involve long treatment times, and as such, there is cause for concern regarding the dose delivered and the associated radiation related risks. An insufficient dose, however, may convey less geometric information, which may lead to inaccuracy and imprecision in intervention placement. The purpose of this study is to investigate a method for achieving the required tracking uncertainty for a given interventional procedure using minimal dose.

Methods: A simple model is used to demonstrate that a relationship exists between imaging dose and tracking uncertainty. A feedback framework is introduced that exploits this relationship to modulate the tube current (and hence the dose) in order to maintain the required uncertainty for a given interventional procedure. This framework is evaluated in the context of a fiducial tracking problem associated with image-guided radiotherapy in the lung. A particle filter algorithm is used to robustly track the fiducial as it traverses through regions of high and low quantum noise. Published motion models are incorporated in a tracking test suite to evaluate the dose-localization performance trade-offs.

Results: It is shown that using this framework, the entrance surface exposure can be reduced by up to 28.6% when feedback is employed to operate at a geometric tracking uncertainty of 0.3 mm.

Conclusions: The analysis reveals a potentially powerful technique for dynamic optimization of fluoroscopic imaging parameters to control the applied dose by exploiting the trade-off between tracking uncertainty and x-ray exposure per frame. © 2011 American Association of Physicists in Medicine. [DOI: 10.1118/1.3560888]

Key words: image-guided intervention, tracking, uncertainty, geometric performance, intrafraction motion, feedback

I. INTRODUCTION

Imaging systems involve controllable acquisition parameters that affect image quality and therefore the ability to localize objects of interest. In current clinical practice, selection of an optimal set of imaging parameters is often made manually and is generally a function of factors, such as the size of the patient, the anatomy that needs to be visualized, and the uncertainty acceptable to the clinical procedure.

Consider an x-ray fluoroscopy system that consists of an x-ray source and a flat panel x-ray detector with a well defined geometry between the two. There are several imaging parameters involved, including the generator tube current (mA), beam energy (kVp), pulse width, frame rate, detector resolution, detector gain, detector field of view, imaging geometry, collimation, and filtration. In order to produce a sen-

sible image, a careful selection of these parameters needs to be made. Oosterkamp¹ was among the first to quantitatively show that the dose to the patient could be reduced while improving image quality by a careful selection of x-ray beam energy. Villagran *et al.*² demonstrated that a correct choice of filtration depending on the beam energy (kVp) and thickness of the patient can reduce the entrance surface exposure (ESE) by half. Theoretical models for an optimal beam quality have long been available.³ The effect of radiographic techniques (tube current and beam energy) on the image quality and the corresponding patient dose levels in the context of a number of clinical applications has recently gained attention (e.g., Refs. 4 and 5). It is desirable to have a high signal-to-noise ratio (SNR). Increasing the tube current improves the SNR; however, more dose is delivered to the patient, which increases radiation related risks.⁴ A good

selection of parameters would offer just enough image quality to carry out the treatment using minimal imaging dose.

X-ray image guidance is indispensable for a number of image-guided clinical procedures. The imaging dose delivered varies depending on the nature of the procedure. Examples of x-ray image-guided interventions with a potentially high imaging dose include real-time tracking in image-guided radiation therapy (RT), neuroradiologic vascular embolisations, percutaneous transluminal coronary angioplasty, transjugular interhepatic portosystemic shunt placement, and radiofrequency cardiac catheter ablations. These interventions involve long procedure times and there is a growing concern regarding the imaging dose delivered to the patient.⁶⁻¹⁷

Kemerink *et al.*¹⁰ reported an average fluoroscopic exposure time of 34.8 min with a maximum of 66 min for a set of neurointerventional procedures. The median patient surface dose according to a study by Gkanatsios *et al.* was found to be 2800 mGy, with a maximum of 5000 mGy.⁹ According to another study by O'Dea *et al.* on 522 cases, 40% of embolization procedures (PA view) deliver an entrance skin dose of greater than 2000 mGy and 6% of these procedures deliver greater than 6000 mGy. Johnson *et al.*⁷ provided a rigorous comparison of published patient dose surveys in interventional radiology between 1993 and 1998. Depending on the complexity of the procedure, the mean skin dose can be as high as 2520 mGy. It is recognized that these dose levels are of a concern.¹⁸

Real-time tumor-tracking radiation therapy (TRRT), particularly of the lung, also involves long periods of imaging-related exposure. Intrafraction motion owing to involuntary motion and movement due to patient discomfort as well as motion due to physiological dynamics such as bowel gas motion, peristalsis, cardiac motion, and respiration leads to geometric imprecision in intervention placement. Recent advances in imaging techniques and treatment delivery methods make it possible to guide the treatment beam based on the location of such moving targets. Imaging and treatment technologies are now available to track and treat tumors in real-time, both in radiosurgery¹⁹⁻²¹ and radiotherapy.²²⁻²⁷ However, there is also a concern regarding the magnitude of dose applied during these tracking procedures. Shirato *et al.*^{27,28} drew concern to the fact that due to the duration of treatment involved in TRRT (up to 30 min/fraction), the dose that is delivered to the patient is unacceptably high (having a surface dose rate as high as 28-980 mGy/h). The purpose of using fluoroscopy in TRRT is to quantitatively track the target. There is a trade-off between tracking uncertainty and the dose corresponding to the imaging technique used. A principled approach would allow for the selection of imaging parameters to balance this trade-off.

Radiation exposure to the operating staff is also a concern. Occupational radiation risk arises mainly from x-ray scatter arising from the patient. In fact, fluoroscopic procedures are the largest source of occupational exposure in medicine.¹⁵ Prolonged exposure during interventional radiology procedures can result in greater risk to staff than other radiological examinations. This is particularly of concern in interventional cardiology and neuroradiology as the nature of the interven-

tion dictates higher dose and longer operating times.^{10,14} Vano *et al.*¹⁵ showed that staff radiation dose rates in interventional cardiology are correlated with patient exposure.

Padovani *et al.*⁸ underlined that the selection of imaging parameters in fluoroscopy guided interventional cardiology procedures is currently patient, operator, and equipment dependent. They emphasized how interventional cardiology procedures are driven by the pathology of the patient and, for this reason, it is not possible to define a standard procedure. Instead, they proposed assigning procedures to classes of complexity and using an imaging technique based on the class into which the procedure falls. The selection of the class, in this case, is a function of several patient dependent factors. The variation in the entrance dose rates for a given procedure and a given homogeneous phantom of fixed thickness from one x-ray system to another⁶ and also across patient habitus²⁹ is also of concern. Moreover, these systems evolve over time and the image quality may change for a given entrance surface dose. An adaptive approach to the selection of imaging parameters may be possible and could prove beneficial.

The concept of feedback has been applied to guide the selection of imaging parameters in order to achieve a desired image quality. Automatic exposure control (AEC) methods select the tube current and beam energy based on image quality.³⁰⁻³² Initially, the only objective of AEC was to produce images of specific image intensity.³³ More recently, emphasis has also been placed on achieving lower dose with AEC.³⁴ Since the selection of imaging parameters in AEC is based on image quality, image quality needs to be defined. Several approaches use metrics such as SNR and contrast-to-noise performance to define image quality and there are other objective assessments of image quality.³⁵⁻³⁷ These approaches require three elements to be specified: (a) The task for which the images are being produced, typically classification or localization; (b) the observer performing this task, i.e., human observer or computer algorithm that will use the information present in the image; and (c) the patient population being imaged. In image-guided interventional procedures where the emphasis is on localization rather than diagnosis, image quality may be less important. In many cases, there are mobile anatomical targets and simple AEC approaches *per se* would not account for the dynamics of the target's motion. Additional challenges may be present when tracking; a mobile object (e.g., fiducial marker in a lung lesion) that is to be localized during an image-guided intervention may traverse regions with varying levels of noise and background intensities, some of which may cause occlusion or distract the tracking algorithm.

The concerns raised above, namely, the need to reduce imaging dose in complex image guidance and tracking procedures; the dependence of imaging parameters on patient pathology, procedure complexity, and imaging setup; and the nonstationarity of noise due to the heterogeneous nature of patient anatomy demand a more principled approach to the selection of imaging parameters. By modulating imaging parameters to achieve the desired localization uncertainty, dose savings can potentially be achieved. In this paper, a

(a)

(b)

FIG. 1. (a) A system diagram comparing a conventional open-loop image-guided procedure with the proposed closed-loop framework for image-guided therapy. In a conventional image-guided procedure, an experienced operator makes decisions regarding the choice of imaging parameters. In the proposed system, the operator specifies an operating uncertainty to a control unit that modulates the input parameters to an image acquisition system based on feedback from a state estimator (that provides, for example, uncertainty estimates). (b) Flow chart showing details of the state estimation module.

framework of dynamic tube current modulation is developed and its benefit is evaluated in a fiducial tracking problem associated with image-guided radiation therapy. Published motion models are incorporated into a test suite to evaluate the model's performance. Metrics of evaluation include the tracking error achieved with the method and the potential for dose reduction under a specified level of uncertainty.

II. METHODS

To produce a reliable tracking algorithm, ideally one could use the discrepancy between the estimated location and ground truth as feedback. However, in the absence of ground truth (i.e., in a real situation), a framework is proposed in which the system's tracking uncertainty is employed as a representative estimate of the accuracy for any single observation. In the proposed framework, the estimator is assumed to be unbiased (i.e., on average representing the true mean). 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.

The proposed framework for dynamic image parameter selection during image-guided interventions is illustrated in Fig. 1(a). In a conventional image-guided procedure, an experienced operator specifies the imaging parameters required to achieve the required tracking uncertainty for an intervention. In the proposed "closed loop" framework, the operator specifies a desired tracking uncertainty for the intervention and the system adapts itself in an attempt to maintain this level of uncertainty. The images that are produced by the imaging device are passed on to a state estimation module. This module produces a probability density function (pdf) for the state of the system. The state defines the parameters of interest in localization such as the location, velocity, and acceleration of the object of interest. The system's tracking uncertainty is computed from the state's pdf and is used as feedback by a controller along with the desired tracking uncertainty. This controller drives the imaging parameters within constraints to achieve the desired tracking uncertainty required for the therapy. The proposed framework is

FIG. 2. (a) Imaging apparatus and setup used for acquiring fluoroscopic images. Shown in the image is a flat panel x-ray detector on the left and an x-ray source on the right. (b) Photograph of the Teflon sphere used to simulate the object of interest. (c) Foam block containing Teflon sphere. (d) A frame from the fluoroscopic sequence acquired on the x-ray system.

illustrated through the problem of tracking a fiducial object in a 2D x-ray fluoroscopic sequence subject to respiratory motion. Here, the framework is adopted to determine the tracking error that can be achieved given the penalty associated with acquiring images in x-ray fluoroscopy, namely, the ionizing dose to the patient. In our feedback model, we *specify* a tracking uncertainty (in mm) and *measure* the performance in terms of an error (in mm).

II.A. Experimental setup

The experimental setup for this study is shown in Fig. 2. A Teflon sphere (1.5 mm in diameter; 1000 HU) was chosen as the fiducial for tracking as it presents relatively low contrast, and hence makes the task of tracking more challenging. This sphere was imaged under x-ray fluoroscopy in a low density uniform background (Styrofoam block) at tube currents ranging from 0.5 to 0.9 mA with an exposure time of 1/30 s and a fixed energy of 100 kVp. The images were acquired at 1 fps on a Varian Paxscan 4030A (2048 × 1536) under the kV imaging geometry typically found on kV enabled medical linear accelerator systems (1000 mm source to axis distance; 1550 mm source to detector distance). A total of 10 frames was acquired at each value of the tube current used. These images were first processed by applying dark and flood field corrections, and gain corrections.

The controller in the proposed framework requires a model that relates the imaging parameters to the resulting tracking uncertainty. One way to define the uncertainty in a location estimate of an object in a 2D image is by the determinant of the covariance matrix as computed using the state density. Let u and v denote the axes of a 2D image obtained from the imaging system and assume that $f[u, v]$ is a pdf given by the state estimation module that defines the proba-

bility of observing the center of the Teflon sphere at location (u, v) in the image. This density function can be used to define a covariance matrix,

$$C = \begin{bmatrix} \sigma_u^2 & \sigma_{uv}^2 \\ \sigma_{uv}^2 & \sigma_v^2 \end{bmatrix}.$$

The determinant of this matrix, $|C|$, is sometimes referred to as the generalized variance. In this paper, the fourth root of this determinant, $|C|^{1/4}$, is taken as a scalar metric of the uncertainty in the estimated location of the sphere's projection onto a 2D image. The fourth root is used in order to work in units of pixels. The determinant of the covariance matrix is proportional to the square of the area formed by an ellipse in 2D with major and minor semi-axes given by $|C|^{1/4} \sqrt{\sigma_u/\sigma_v}$ and $|C|^{1/4} \sqrt{\sigma_v/\sigma_u}$. If the covariance matrix is singular (the pdf is nonzero only on a line in the $u-v$ space), $|C|$ collapses to 0. For such cases, this metric will underestimate the uncertainty. Although this is highly unlikely given the nature of correlation and noise in 2D x-ray image formation, this singularity was avoided by regularizing the pdf as discussed with reference to the observation process below. Alternatively, the singularity can also be avoided by explicitly regularizing the covariance matrix to produce a new covariance matrix D ,

$$D = (1 - \alpha)C + \alpha \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

where $0 < \alpha < 1$ is a sufficiently small scalar. However, this approach was not used in the results presented here.

Using the observation process described below, the mean location of the center of the sphere in the image and the associated uncertainty was computed for each of the ten processed images individually at each of the tube current levels used. Given an image, the observation process produces a pdf for the state of the system (location of the center of the sphere). This was used to compute an uncertainty estimate in the localization. The average uncertainty across the 10 frames at each tube current level was then plotted as a function of the tube current to give the relation shown in Fig. 3(b).

Motion sequences were generated by randomly drawing 100 frames uniformly from the 10 acquired frames and translating them to simulate the influence of respiratory motion using motion data from Seppenwoolde *et al.* (Table II).³⁸ Once a motion sequence for each tube current level was generated, the sequences were composited to include a region of low noise and a region of high noise. Four motion sequences were thus generated (2.0/1.0, 2.8/1.4, 3.6/1.8, and 4.0/2.0 mA) with the tube current corresponding to the region of high noise being half that of the region of low noise. The motivation for the use of the low and high noise regions comes from a clinical context: When tracking an object in fluoroscopic imaging, the object may travel into regions of varying background noise (and mean signal). For example, when tracking a fiducial marker in the thorax, the marker could become obscured by overlying ribs during respiration or be periodically occluded by a part of the patient support system. In this study, we simulate the effect by using a simple stepwise transition from low

FIG. 3. Observation operator. (a) Computation of the likelihood of the state (location) of the sphere. Left: A zoomed-in view of a frame from an x-ray fluoroscopic sequence of the Teflon sphere. Center: Image of the sphere in edge space with an overlay of the observation model (four profiles, $g_1 - g_4$). The bright pixels indicate locations on the overlaid profiles where an edge element is observed. Right: A plot of the likelihood of observing the sphere at a specific location as computed in edge space. (b) Top: Localization uncertainty as given by the observation model in (a) under varying tube current levels along with a polynomial fit. Here, $a = 2.0$, $b = -3/2$, and $c = 0.1$. Bottom: Sample images from fluoroscopic sequences of the Teflon sphere under varying tube currents illustrating reduction in CNR with increasing tube current. A 40×40 pixel region of interest is shown for illustration purposes.

to high noise. The mean signal between these regions has been normalized to allow a simplified observer model to be employed (i.e., no jump discontinuity in the mean signal).

The x-ray pulse width and the magnification factor were fixed at $1/30$ s and 1.55, respectively. The frame rate was fixed at 10 fps. A frame from a generated sequence and an overlay of a sample trajectory taken by the sphere is shown in Fig. 4(a). Sample trajectories corresponding to selected patients from the data of Seppenwoolde *et al.*³⁸ with varying dynamics are shown in Fig. 4(b).

An acceptable level of uncertainty was specified and based on feedback of the estimated uncertainty, the controller was allowed to modulate the tube current in order to maintain the specified level of uncertainty as the sphere was tracked through the region of low exposure. Let $X[k] = [x[k], y[k]]$ with $X[k] \in \chi$, $\chi \subset \mathbb{R}^2$ denoting the true state of the system at time step k , where $[x[k], y[k]]$ are the world coordinates of the center of the Teflon sphere. The motion of the sphere was restricted to a 2D plane (100 cm from source) in world coordinates. The task was to report the estimated 2D location of the center of the sphere in world coordinates given the observed images ($Y[1:k]$) up to time step k ,

$$\hat{X}[k] = [\hat{x}[k], \hat{y}[k]]^T \equiv E(X[k]|Y[1:k]).$$

The components of the proposed framework for this task are described next.

II.B. Framework

II.B.1. State estimation module

The task of the state estimation module is to produce a probability density estimate of the state of the system at each time step using the observed images and prior knowledge in the form of observation and dynamic models of the object of interest. X-ray fluoroscopic sequences involve varying levels of noise. Objects of interest may be occluded by overlapping structures, and anatomical structures present background clutter. For the proposed framework, an algorithm is needed that is robust to such challenging circumstances. In order to be robust to partial occlusion, a method that can propagate a multimodal probability distribution of the state of the system is needed. When the object of interest is occluded by structures such as a graticule, patient support apparatus, and bony structures, then using a multimodal belief representation can help provide a more faithful estimate of the object's state. The dynamic model of objects of interest in the context of interventional procedures and radiotherapy is not necessarily linear with additive Gaussian noise. Nonlinearities can arise in a wide variety of circumstances: The typical motion of a surgeon's hand in manipulating a tool is nonlinear; respiratory motion is nonlinear; tumor motion may also be subject to highly nonlinear motion, for example, due to coughing, peristalsis, or fluid movement within the patient.

FIG. 4. Model for evaluating the closed-loop framework. (a) A frame from the simulated “low-high-low” noise x-ray fluoroscopic sequence used to evaluate the proposed framework. Each frame is composed of a region of low noise and a region of high noise. Shown in this image is an example of a simulated trajectory taken by the sphere through regions of low and high noise. (b) Sample trajectories corresponding to selected patients from Seppenwoolde *et al.* (Ref. 38) illustrating the range of motion and form. Trajectories are shown to scale in pixel coordinates.

A brief review of filtering methods is provided to motivate the use of particle filters. 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 and lends itself well to the problem. This class of algorithms seeks to compute recursively some degree of belief in the state using all the observations available up to the current time step. This proceeds in two steps: (i) *Prediction* of the future distribution of the state of the object using a dynamic motion model and (ii) *correction* (or *update*) of the predicted probability density using an observation model to obtain a posterior density. Under challenging tracking circumstances involving noise, occlusion, and nonlinear dynamics, it may be difficult to compute the posterior density in closed form. Sequential importance sampling (also known as particle filters,^{39,40} condensation algorithm,⁴¹ bootstrap filtering,⁴² and interacting particle systems⁴³) is a class of sequential Bayesian estimators that overcomes this challenge by performing sequential Monte Carlo estimation of the distribution of the state of the system.⁴⁴ These estimators maintain a sampled representation of the distribution of the state, update it appropriately on receiving a new observation, and propagate it over time. Any statistic that is to be drawn from this sampled distribution is obtained using the Monte Carlo integration approximation. The assumptions imposed by other Bayesian filters can be too restrictive for the problem of robustly tracking

objects of interest under x-ray fluoroscopy. The Kalman filter (e.g., Ref. 45) requires linear observation and dynamic models and can only work under additive Gaussian noise. The extended Kalman filter (EKF) (e.g., Ref. 46) can handle nonlinear observations and dynamics with additive white Gaussian noise; however, it can only propagate unimodal densities. The unimodal belief limitation of Kalman filters and EKFs can be overcome by using a Gaussian mixture model based multi-hypothesis tracker (MHT). However, updating the density in this case is a computationally complex task. Moreover, if left unchecked, the number of Gaussian components can grow exponentially with time.⁴¹ Furthermore, the linearity assumptions of Kalman filters must still hold for each hypothesis of the MHT.

In this study, a particle filter is used for state estimation because of its ability to handle multimodal probability densities, nonlinear dynamics, and non-Gaussian noise. Specifically, the condensation algorithm⁴¹ is used here without any loss of diversity correction.⁴⁸ New samples are drawn from a dynamic model prior. The ensemble of particles (1000 in this analysis) is propagated using a dynamic model and with each new observation, the weights of the particles are updated with weights derived from an observation process applied at each particle location. The sampled representation of the probability density of the 2D location of the sphere thus generated is used to construct the covariance matrix and the location estimate ($\hat{X}$) using the Monte Carlo integration

approximation. The steps involved in state estimation module are summarized in Fig. 1(b). The dynamic model and the observation model used are described below. The particle filter was initialized with particles obtained by sampling uniformly from a rectangular region containing the sphere in the first frame.

II.B.1.a. Dynamic model. A dynamic model attempts to explain how the dynamics of the object of interest evolve over time. A recursive dynamic model (required for sequential estimation) was derived based on the respiratory motion model proposed by Lujan *et al.*,⁴⁹

$$f(t) = s_0 + S \cos^{2n} \left(\frac{\pi t}{\tau} - \phi \right), \quad (1)$$

where s_0 is the exhale position, S is the spatial amplitude of motion in one direction, τ is the breathing period, n is a shape parameter, and ϕ is the starting phase.⁴⁹ Here, $f(t)$ gives the location along one axis of the image and the same model holds for the other image axis except possibly for differences in the values of some parameters. For numerical stability reasons, a recursive formulation was developed by defining

$$s(t) \equiv \cos^2 \left(\frac{\pi t}{\tau} - \phi \right), \quad (2)$$

together with a simple mapping to image space,

$$\hat{f}(t) = s_0 + S[s(t)]^n.$$

Differentiating and reintegrating Eq. (2) using Beeman's algorithm⁴⁷ gives

$$s(t + \Delta t) = s(t) + \dot{s}(t)\Delta t + \frac{2}{3}\ddot{s}(t)\Delta t^2 - \frac{1}{6}\ddot{s}(t - \Delta t)\Delta t^2 + N(0, \sigma_s^2), \quad (3)$$

$$\dot{s}(t + \Delta t) = \dot{s}(t) + \frac{1}{3}\ddot{s}(t + \Delta t)\Delta t + \frac{5}{6}\ddot{s}(t)\Delta t - \frac{1}{6}\ddot{s}(t - \Delta t)\Delta t + N(0, \sigma_{\dot{s}}^2), \quad (4)$$

where $N(0, \sigma_s)$ and $N(0, \sigma_{\dot{s}})$ are zero mean Gaussian random variables with variances σ_s and $\sigma_{\dot{s}}$ that approximate the higher order terms in the expansion, and the dot operator represents the first derivative. Computing the second derivative of Eq. (2) and simplifying yields

$$\ddot{s}(t) = \frac{2\pi}{\tau} \left[1 - 2 \cos^2 \left(\frac{\pi t}{\tau} - \phi \right) \right] = \frac{2\pi}{\tau} [1 - 2s(t)], \quad (5)$$

This expression can be substituted for the acceleration term in Eqs. (3) and (4). Equations (3) and (4) thus allow a simple recursive formulation for Lujan's model. The parametrized position given by $s_u[k]$ and $s_v[k]$ evolve independently, each according to Eqs. (3) and (4), with the acceleration term substituted using Eq. (5). The particle filter's state carries distributions for $s_u[k]$, $s_v[k]$, $\dot{s}_u[k]$, and $\dot{s}_v[k]$, which map to image coordinates along the u and v axes [by Eq. (1)] at time step k by

$$u[k] = u_0 + A s_u^n[k], \\ v[k] = v_0 + B s_v^n[k],$$

where u_0 and v_0 are offsets along the u and v axes as defined in Eq. (1), and A and B represent the amplitude of motion along the anterior-posterior and cranial-caudal directions. $u[k]$ and $v[k]$ relate to world coordinates through scaling: $x[k] = c_u u[k]$ and $y[k] = c_v v[k]$.

Using a dynamic model alone is not enough for robustly tracking objects of interest as it may not be exact (thus accumulating residual errors) and does not account for disturbances. Using the observation process described below, the state distributions are updated when requested. In this study, the state distributions are updated at the point of image acquisition.

II.B.1.b. Observation process. The observation process, here, refers to the process of updating the weights of the particles using images and prior information (a geometric model of the object being tracked). In this analysis, the geometric model of the object of interest (i.e., the sphere) is a disk of a specified radius. The observation process used here proceeds as follows: On each acquired image, a Canny edge detector is applied to produce an edge map (indicating the regions where edges are observed) and an orientation map (likewise showing the orientation of these edge elements). For a given evaluation point, four profiles perpendicular to the expected locations of edges of a sphere centered at the evaluation point are examined based on the object's geometric model, as shown in Fig. 3(a). In this analysis, the profiles correspond to the top, bottom, left, and right edges. The likelihood of the sphere's center being at the evaluation point, given the model, is then computed as follows:

likelihood

$$= P(\text{Object is present at location } (u, v) \text{ in image} | \text{Object model}) \\ = P(g_1 = m_{g_1}, g_2 = m_{g_2}, g_3 = m_{g_3}, \\ g_4 = m_{g_4} | \text{Object model}), \quad (6)$$

where $g_j = m_{g_j}$ is the event that an edge of the correct orientation is found at location m_{g_j} on profile g_j . We make the following independence assumption:

$$P(g_i | g_j, \text{Object model}) \\ = P(g_i | \text{Object model}) P(g_j | \text{Object model}) \quad (7)$$

for $i, j = 1, \dots, 4$, $i \neq j$. This independence assumption simplifies Eq. (6) to

$$\text{likelihood} = \prod_{j=1}^4 P(g_j = m_{g_j} | \text{Object model}). \quad (8)$$

This likelihood is used to update the particle weights via the factored sampling algorithm.⁴¹ Let π_i denote the likelihood corresponding to the i th sample. The weights corresponding to the samples are then given by

$$w_i = \frac{\pi_i}{\sum_{j=1}^N \pi_j}.$$

The set of samples $(u_i[k], v_i[k], w_i[k])$ for $i = 1, \dots, N$ constitutes a sampled representation of the pdf of the location of the sphere in the image and was used to define the location of the sphere and the covariance matrix, C , using the Monte Carlo integration approximation.

The probability mass functions (pmf's), $P(g_j = m_{g_j} | \text{Object model})$, used in the likelihood computation above were learned empirically by building histograms using a set of images with the object's geometric model placed to

coincide with the known locations of the sphere. These histograms were then smoothed using a Gaussian filter to effect a distance transform, normalized to produce a probability mass function, and regularized by mixing the resulting pmf with that of a uniform pmf at a 95:5 ratio. The mixing with a uniform pmf was done to allow robustness to displaced or missing edges. This mixing also ensures that weights assigned to the particle filter's samples are nonzero and reduces the chances of C becoming singular.

II.B.2. Control module

A discrete-time proportional-integral (PI) controller was used to modulate the tube current to the generator. The inputs to this controller were the desired uncertainty and feedback of the uncertainty estimate produced by the state estimation module. Given a desired operating uncertainty, a reference tube current (i_0) was established using the relation developed in Fig. 3(b). The PI controller was then allowed to regulate the tube current about this reference value in response to changes in the estimated tracking uncertainty in order to maintain the specified operating uncertainty. The controller uses the difference between the desired operating uncertainty and the estimated uncertainty. Denote this difference by ε_k , where k is the current time step. The output of the controller is:

$$i_{k+1} = i_0 + G_P \varepsilon_k + \sum_{k'=k-q}^k G_I \varepsilon_{k'},$$

where G_P and G_I are the proportional and integral gains, respectively, and $q < k$. This is then quantized so it corresponds to one of the available tube current levels on the generator and is saturated so that it does not exceed a safe operational range. Four tube current levels were available as the output of the controller: 2.0, 2.8, 3.6, and 4.0 mA, corresponding to the motion sequences 2.0/1.0, 2.8/1.4, 3.6/1.8, and 4.0/2.0 mA, as described earlier. A saturator was applied to limit the tube current between 2.0 and 4.0 mA.

II.C. Validation

The framework described above was used to track a sphere through regions of low and high noise, as described with reference to Fig. 4(a). This framework was executed multiple times with different random seeds for the particle filter (each of these referred to as a run) for motion data corresponding to all 21 patients as described earlier. Since the particle filter is a stochastic estimation technique, each run (1000 particles) was repeated ten times to evaluate the reproducibility of the estimate. The framework was evaluated at operating uncertainties of 0.15, 0.2, and 0.3 mm. The proposed framework was compared to a control case in which feedback was not employed and the tube current was fixed. In this case, the Teflon sphere was tracked at tube currents of 2.0, 2.8, 3.7, and 4.0 mA. As in the case with feedback, each run was repeated ten times. The metrics listed below were computed for each patient and averaged over the ten runs. The analysis was performed on two cycles for all patients except for patients 6, 16, and 18, in which case only one cycle was

available for analysis. In all cases, the analysis was started at frame 10 of the generated sequence to ignore transient effects associated with initializing the state with a uniform prior distribution. The following metrics were used for evaluation.

II.C.1. Tracking error

The location estimates given by the state estimation module were compared to the corresponding "ground truth" values for each of the trajectories in the test suite. For each patient, the root mean square error (RMSE) between the estimated location, $\hat{X}_r[k]$, of the sphere for run r and ground truth, $\tilde{X}[k]$, was computed using the Euclidean distance between the two and normalized by the number of frames used, K , in a run to give the average RMSE per frame,

$$\frac{1}{KR} \sum_{k=1}^K \sum_{r=1}^R \|\hat{X}_r[k] - \tilde{X}[k]\|_2.$$

Here, R is the total number of runs per patient and was fixed at 10. These values are reported in Table I. This statistic was computed over two cycles for all patients except patients 18, 16, and 6, in which case only one cycle was used. The ground truth values were computed by first manually identifying the location of the center of the Teflon sphere in the ten acquired images for each choice of imaging parameters. The mean location over these ten images was taken as the ground truth for the stationary image. This mean location along with the known translations applied to the acquired images when simulating respiratory motion trajectories were then used to define ground truth. Bias or systematic error in the estimated location was removed by determining the offset that minimized the total RMSE for all tube current levels and all runs when operating without feedback of uncertainty and adding this offset to the estimates.

II.C.2. Entrance Surface Exposure (ESE)

Using a measured exposure rate of 4.74 mR/mAs at the detector, the total ESE at the object for each run was computed and averaged over all ten repetitions. These values were then normalized by the number of frames in the corresponding run and reported as the average ESE per frame (Table I),

$$\frac{1}{10K} \sum_{k=1}^K \sum_{r=1}^{10} f_r[k] \left(\frac{\text{SDD}}{\text{SAD}} \right)^2 4.74,$$

where $f_r[k]$ is the exposure for the k th frame and r th run, SDD is the source to detector distance, and SAD is the source to axis distance.

II.C.3. Normalized area greater than threshold (A_{GT})

A metric to quantify error referred to as the normalized area (A_{GT}) greater than a chosen threshold, ε , was defined as follows:

$$\begin{aligned} & \frac{1}{KR} \sum_{r=1}^R \sum_{k=1}^K g[k], \text{ where } g[k] \\ &= \begin{cases} \|\hat{X}[k] - \tilde{X}[k]\|_2 - \varepsilon & \text{if } \|\hat{X}[k] - \tilde{X}[k]\|_2 > \varepsilon \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

TABLE I. Quantitative comparison of performance with and without feedback of uncertainty. Shown here are the parameters corresponding to Lujan's motion model for each of the 21 patients in Seppenwoolde *et al.* (Table II) (Ref. 38) and the associated performance in terms of the average RMSE per frame, the average ESE per frame, and the normalized area (A_{GT}) greater than the chosen threshold, ϵ , of 0.2 mm. Also shown are the dose reductions using feedback of uncertainty relative to the cases without feedback operating at tube currents of 2.0 and 2.8 mA. For comparison purposes, the equivalent average tube current values employed by the feedback case are shown. CC and AP indicate the extent of motion along the cranial-caudal and anterior-posterior axes. s is the breathing period in seconds, n is a shape parameter, and "cycles" indicates the number of cycles over which the estimation was run. The mean and standard deviation across the patient population for RMSE and ESE are also shown. CC, AP, and RMSE values are expressed in mm in world coordinates at the patient stage. (Average tube current for the feedback case is shown in mA and not 10^{-3} mA.)

Patient Parameters							Without Feedback ($\times 10^{-3}$)												With Feedback ($\times 10^{-3}$)				Dose Change Relative to 2.8 mA	Dose Change Relative to 2.0 mA
							2.0 mA			2.8 mA			3.6 mA			4.0 mA			Operating Uncertainty of 0.2 mm					
ID	CC [mm]	AP [mm]	s sec.	n	cycles	RMSE [mm/f]	ESE [mR/f]	A _{GT} > ε [mm/f]	RMSE [mm/f]	ESE [mR/f]	A _{GT} > ε [mm/f]	RMSE [mm/f]	ESE [mR/f]	A _{GT} > ε [mm/f]	RMSE [mm/f]	ESE [mR/f]	A _{GT} > ε [mm/f]	Avg. Tube Current [mA]	(%)	(%)				
20	24.6	1.9	3.8	1	2	107	↑	8	102	↑	5	71	↑	1	65	↑	0	96	878	5	2.3	-18.4	+14.2	
9a	13.1	3.6	3.6	2	2	152	↑	49	88	↑	1	67	↑	0	65	↑	0	99	944	14	2.5	-12.3	+22.7	
10	11.8	2.6	4	2	2	173	↑	70	90	↑	4	65	↑	0	62	↑	0	103	949	22	2.5	-11.8	+23.5	
3	10	1.7	3	2	2	167	↑	58	101	↑	7	80	↑	8	65	↑	2	103	1000	19	2.6	-7.1	+30.0	
7	7.7	0.9	3.3	2	2	168	↑	66	85	↑	4	78	↑	3	69	↑	0	98	963	12	2.5	-10.6	+25.2	
18	5.4	0.9	6.6	9	1	347	↑	236	125	↑	27	87	↑	2	83	↑	0	157	984	60	2.6	-8.6	+28.0	
17	1	1.7	3.8	1	2	123	↑	15	85	↑	1	82	↑	0	65	↑	0	121	852	17	2.2	-20.9	+10.8	
16	0.2	0.2	4.9	1	1	117	↑	10	88	↑	1	78	↑	0	66	↑	0	113	782	9	2.0	-27.4	+1.7	
19	11.1	5.8	3.2	2	2	140	↑	40	85	↑	1	68	↑	1	70	↑	0	96	949	11	2.5	-11.9	+23.4	
6	0.7	1.2	4.8	2	1	156	769	45	90	1077	2	63	1384	0	69	1538	0	105	1006	16	2.6	-6.5	+30.8	
8	8.7	8.2	3.8	3	2	297	↑	188	107	↑	15	76	↑	5	67	↑	0	103	956	16	2.5	-11.2	+24.3	
12	4.2	4.1	3.2	2	2	186	↑	83	96	↑	3	72	↑	1	63	↑	0	92	959	9	2.5	-11.0	+24.7	
9b	3	2.5	3.5	1	2	110	↑	13	84	↑	2	69	↑	0	65	↑	0	99	841	11	2.2	-21.9	+9.4	
1	2.8	1.5	2.8	2	2	162	↑	59	86	↑	1	77	↑	7	67	↑	0	93	959	8	2.5	-11.0	+24.6	
5	2.8	0.8	2.7	2	2	157	↑	62	80	↑	1	73	↑	0	69	↑	0	91	946	10	2.5	-12.1	+23.0	
11	2.5	1.4	3.2	1	2	104	↑	10	84	↑	1	80	↑	0	67	↑	0	98	784	8	2.0	-27.2	+1.9	
15	2	2.7	3.2	1	2	127	↑	18	90	↑	0	71	↑	0	60	↑	0	125	799	17	2.1	-25.8	+3.9	
14	1.5	1.4	3.7	2	2	200	↑	85	81	↑	2	71	↑	4	68	↑	0	104	974	12	2.5	-9.5	+26.6	
4	1.2	2	3.2	1	2	99	↑	7	85	↑	3	82	↑	1	59	↑	0	98	810	5	2.1	-24.7	+5.4	
2	1	1.2	3.3	1	2	114	↓	10	88	↓	1	79	↓	0	67	↓	0	106	826	10	2.1	-23.3	+7.4	
13	0.7	0.6	3.7	1	2	91	↑	2	90	↑	1	80	↑	0	62	↑	0	92	791	2	2.1	-26.5	+2.8	
Mean						157	769	Mean	91	1077	Mean	75	1384	Mean	66	1538	Mean	104	902					
Std						63	0	Std	10	0	Std	6	0	Std	5	0	Std	15	79					

This metric was evaluated for each patient and is reported in Table I in units of mm/frame.

III. RESULTS

As discussed in Sec. II A, a plot of the average estimated localization uncertainty, $|C|^{1/4}$, over 10 frames at each tube current level as a function of tube current is shown in Fig. 3(b). A polynomial fit to the measured data was computed giving the following relation:

$$|C|^{1/4} = \frac{2.0}{i\sqrt{i}} + 0.1,$$

where i is the tube current. The constant 0.1 is due to overfitting and the limited range of tube currents explored.

A demonstration of the proposed framework's ability to predict uncertainty and its relation to error is shown in Fig. 5 for the control case without feedback. Plots of the uncertainty, tube current, and root mean square (RMS) error for a run corresponding to tube current levels 2.0, 2.8, 3.6, and 4.0 mA are shown. Regions of higher noise are shaded.

In Fig. 6(a), a sample run employing the proposed feedback framework operating at 0.2 mm is compared to a run from the control case operating at a fixed tube current of 2.0 mA without feedback of uncertainty. For all ten runs of each

of the 21 patients in which feedback was not used, the average values of the RMSE in regions of low and high noise were plotted as a function of the average estimated uncertainty in the corresponding regions to produce the scatterplot shown in Fig. 6(b). Also shown in this figure is a similar plot for the proposed framework using feedback, operating at an uncertainty of 0.2 mm.

A quantitative comparison of the performance of the simulations with and without feedback is presented in Table I. For each of the motion trajectories corresponding to patients mentioned in Seppenwoolde *et al.* (Table II),³⁸ the validation metrics described above were evaluated. CC and AP refer to the extent of motion in the cranial-caudal and anterior-posterior directions, s is the breathing period, and n is a shape parameter as described with respect to Eq. (1). The RMSE per frame, ESE per frame, and A_{GT} values for the case without feedback are compared to the feedback case operating at a targeting uncertainty of 0.2 mm. The mean and standard deviation of these metrics across the patient population are also provided. The last two columns of Table I provide the dose reductions for the feedback case as compared to the case without feedback at fixed tube currents of 2.0 and 2.8 mA, respectively. Note that the motion models with higher order shape parameter (n) demonstrated higher error and this is consistent with the high run-to-run variance seen in the

FIG. 5. Demonstration of the proposed method's ability to predict uncertainty and its relation to error. Shown in this figure are statistics corresponding to tracking a run of a sequence corresponding to patient 20. The plots here show how the predicted uncertainty in mm, the tube current in mA, and the RMS vary as a function of time as the sphere moves in and out of the noise field. The feedback method is not yet employed in these results. The four quadrants correspond to tube current levels of (a) 2.0 mA, (b) 2.8 mA, (c) 3.6 mA, and (d) 4.0 mA. Shaded regions indicate regions of higher noise or noise fields. Averages for each region are also shown.

FIG. 6. Comparison of performance with and without feedback of uncertainty. Shown from top to bottom are the predicted uncertainty, the tube current, and the RMS error over one simulated trajectory corresponding to patient 20. (a) Left: Operation without feedback of uncertainty as in Fig. 5. Right: Operation of the proposed framework for tube current modulation using feedback of uncertainty. In this case, an operating uncertainty of 0.2 mm is specified. (b) Scatterplots of the average estimated uncertainty and the average error for each of the regions in Figs. 5 and 6(a) for the case without feedback (shown on the left) and the case with feedback (shown on the right). These plots contain points for all ten runs corresponding to each of the 21 patient trajectories employed in this study. The inset on the right shows a close-up view of the region around 0.2 mm, which is the specified operating point.

FIG. 7. Average values of RMSE per frame (top) and the ESE per frame across all 21 patients plotted as a function of the operating uncertainty specified. The height of each error bar represents 2σ . The standard deviations at operating uncertainties of 0.15, 0.2, and 0.3 mm are 10×10^{-3} , 15×10^{-3} , and 16×10^{-3} , respectively, for the average RMSE per frame and 55×10^{-3} , 79×10^{-3} , and 69×10^{-3} for the average ESE per frame.

particle filter's estimate (ensemble of 10, Sec. II C). This may be due to a need for a higher number of particles to accurately track higher order motion ($n \geq 3$). The number of particles to use in a particle filter is an open question in computer vision.

The proposed framework employing feedback of uncertainty was also evaluated for operating uncertainties of 0.3 and 0.15 mm in which cases dose savings of up to 18.4% and 28.6% were observed, respectively. The average ESE per frame and the average RMSE per frame across all 21 patients in the test suite at operating uncertainties of 0.15, 0.2, and 0.3 mm are shown in Fig. 7.

In terms of implementation, the algorithms used here were implemented in MATLAB and no effort was made to accelerate the algorithm. For both the feedback and without feedback cases, it took 11 min for tracking in 100 frames on an Intel Core 2 Duo T9300 2.5 GHz processor with 3 GB of RAM. These speeds would have to be substantially increased to implement in a real-time context, but this is not expected to be a major challenge provided the observer model is not computationally heavy.

IV. DISCUSSION

The trade-off between dose and operating uncertainty is demonstrated in Fig. 3(b). With this type of dependence, it is possible to operate either at a lower uncertainty using a high dose rate (which is proportional to tube current) or at a higher uncertainty using a low dose rate. This relation

between dose and uncertainty allows the proposed framework to offer a method of modulating the tube current in order to maintain the desired level of uncertainty in response to changes in the observations.

It can be seen from the operation of the framework without feedback of uncertainty (as shown in Fig. 5) that in the regions of higher noise (shaded regions), the estimated uncertainty is higher and so is the corresponding RMSE. For the proposed framework to work, faithful estimates of the error are required. Figure 6(b) suggests that the estimated uncertainty is well correlated with the average RMSE in each region. This allows the controller to increase the tube current as necessary when the estimated uncertainty increases. As shown in Fig. 6(a) for the feedback case, when the sphere enters the region with higher noise, the controller increases the tube current in response to the increase in the estimated uncertainty, thereby maintaining the RMSE through the noise field. In the low noise region, feedback of a lower value of the estimated uncertainty causes the controller to lower the tube current. The framework thus attempts to maintain the desired uncertainty throughout a procedure with a minimum dose rate. In practice, there is a set of quantized tube current values that can be used when imaging and it is often difficult to select one level that works throughout a procedure. Continuously modulating the tube current would help overcome this limitation.

The RMSE values in Table I show that the framework is able to track the Teflon sphere as it traverses the noise field. The average RMSE per frame values across the patient population are 0.2, 0.09, 0.08, and 0.07 mm for the case without feedback operating at 2.0, 2.8, 3.6, and 4.0 mA, respectively. The corresponding value for the feedback case operating at an uncertainty of 0.2 mm is 0.10 mm/frame. For a specified operating uncertainty of 0.2 mm, the value of the average RMSE per frame for the feedback case falls between that of the cases without feedback operating at 2.0 and 2.8 mA. Likewise, the average ESE per frame for the feedback case also falls between that of the cases operating at 2.0 and 2.8 mA without feedback. For each patient, the feedback case uses less dose compared to the 2.8 mA case without feedback but more dose compared to the 2.0 mA case without feedback. As can be seen in Table I, the feedback case employs an average tube current value of between 2.0 and 2.6 mA, depending on what is necessary to achieve the specified operating uncertainty. The advantage of the proposed framework in this case is that it allows one to achieve an RMSE and ESE intermediate to those for the fixed 2.0 and 2.8 mA cases, while satisfying the uncertainty requirement in the context of changing noise conditions. The feedback case adopts a lower tube current under low noise conditions (e.g., less attenuating background object) and a higher tube current in high noise conditions (e.g., more attenuating background object). Depending on the proportion of time spent under these two conditions, the dose saving advantages will be moderated. For example, in the case of patient 20, the feedback case maintains the RMSE per frame (0.096 mm vs 0.102 mm in the nonfeedback case) yet does so using 18.4% less dose. In contrast, for patients 3 and

12, the feedback approach provides only 7.1% and 11.0% less dose, respectively.

The reduction in dose is dependent on the operating uncertainty chosen. Dose savings of up to 18.4%, 27.4%, and 28.6% were observed in the feedback cases operating at 0.15, 0.2, and 0.3 mm, respectively, relative to the case without feedback operating at 2.8 mA. Given the duration of treatment involved in complex interventional procedures described earlier, these savings are promising. The less stringent the uncertainty requirements are, the greater the dose savings. From Fig. 7, it can be seen that with increasing values of the operating uncertainty specified, the average RMSE per frame increases and the average ESE per frame decreases, as expected. Thus, in the proposed framework, operation is not dictated by the values selected for the imaging parameters but rather by the uncertainty requirements of the interventional procedure.

The A_{GT} values for the feedback case never exceed that of the case without feedback operating at 2.0 mA with the exception of patient 17 (Table I). For this patient, the A_{GT} value for the feedback case operating at 0.2 mm is slightly larger when compared to the case without feedback operating at 2.0 mA. This can be attributed to the fact that for this patient the extent of motion along the AP direction is greater than that along the CC direction. As a result, in the simulation, the sphere moves along an artificial edge between the region of high noise and the region of low noise. The worst performance in terms of the average RMSE per frame is observed for patient 18. This is due to the highly nonlinear nature of motion ($n = 9$) exhibited by the sphere. Nevertheless, the framework is able to track the sphere and employ feedback of uncertainty. For patients 3, 1, and 14, the A_{GT} values for the case without feedback operating at 3.6 mA are slightly greater than that for an operating uncertainty of 2.8 mA. This can be attributed to an artifact of structured noise present in a few frames of the 3.6 mA sequence.

Across the patient population, the trends (Fig. 7) show that RMSE and ESE correlate with operating uncertainty and inversely with operating uncertainty, respectively. In general, the lower the uncertainty specified, the more dose is used and consequently the lower the standard deviation in the average RMSE per frame. However, the range (see error bars in Fig. 7) reflects patient-specific variations in dose benefit. Since the results are also a function of the dynamic model besides the observation model, the nature of motion has an effect on dose savings, i.e., for slow and smooth motion (no sharp turns), dose savings are greater. For example, for patient 16, dose savings of 27.4% savings are observed when operating at an uncertainty of 0.20 mm with feedback as compared to the case without feedback operating at 2.8 mA.

In the implementation used here, the uncertainty estimated upon observing the image at the current time step is used to modulate the tube current for the next time step (the uncertainty is computed from the posterior distribution). This introduces a lag; when the sphere being tracked transitions from the region of low noise to the region of high noise, there is a delay before the tube current 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 in the simulations in this study where 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. Nevertheless, if such boundaries are present, the framework 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 the tube current just prior to entering a noise field. This is also useful when the object being tracked moves in a cluttered background. If the object's neighborhood changes very rapidly, there may be a delay before the controller can modify the tube current. By learning the background context over a few breathing cycles, the controller would be better able to adapt the tube current.

In this study, a simple case of tracking a sphere through two different backgrounds is examined. It is expected that the greater the disparity between the background regions, the higher the potential for dose savings. Clinical images tend to be more complex and may better reveal the potential of the proposed framework. The simulations in this study were restricted to tube current levels at which sphere remains detectable. For further robustness, the framework would need to be developed to handle tracking failures. From an implementation standpoint, the proposed framework involves relatively low cost modifications to existing x-ray based tracking systems.

In this study, only tube current modulation was considered. A number of other parameters can also be modulated. For example, if the object being tracked moves into a region with higher noise, the tube current may need to be increased to improve the SNR; if the object is occluded by bone, the energy (kVp) may need to be reduced to improve contrast-to-noise ratio (CNR); if the object makes a sharp turn, the frame rate may need to be increased. More elaborate observation models that account for quantum noise, anatomical noise (in the form of background clutter), and the object's dynamics would need to be explored.

The proposed framework can also be used to allocate varying uncertainty requirements to different regions, i.e., one region may require a higher precision during intervention than another. For example, during RTTRT, one-half of the tumor's trajectory may be less tolerant to tracking uncertainty than the other as the tumor may be treated only during part of the breathing cycle with the other half of the breathing cycle requiring sufficient performance to simply to avoid losing track of the tumor.

V. CONCLUSIONS

In this paper, it has been identified that relationships exist between tracking uncertainty and the parameters of an imaging modality. A framework has been introduced in which such relations can be exploited to control the parameters of

an image-guided system using metrics of geometric performance as feedback in order to meet the geometric tracking requirements. The efficacy of this framework has been illustrated through the problem of estimating the location of a fiducial object undergoing respiratory motion during x-ray fluoroscopy. This framework offers the potential for (i) reducing dose to the patient and the operating staff, (ii) automatically maintaining the geometric objectives of the therapy and reporting to the operator when these objectives cannot be achieved, (iii) optimizing the system parameters by dynamically assigning them based on feedback of geometric performance, and (iv) reducing the level of human intervention required to optimize the tracking performance. Finally, a detailed analysis of the accuracy, tracking uncertainty, and dose savings needs to be conducted in a more clinically relevant context that involves more complex noise models and dynamics. This is the subject of future work.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Douglas Moseley and Dr. Jeffrey H. Siewerdsen for providing valuable insight and for many stimulating discussions. This research was supported, in part, by the National Institute on Aging/National Institutes of Health under Grant No. R33 AG019381, Elekta Oncology Systems, and through the Fidani Family Chair in Radiation Physics at Princess Margaret Hospital.

^{a)}Electronic mail: sami.siddique@rmp.uhn.on.ca

^{b)}Author to whom correspondence should be addressed. Electronic mail: david.jaffray@rmp.uhn.on.ca; Telephone: 416-946-4501 (ext. 5387); Fax: 416-946-4578.

¹I. W. Oosterkamp, "Monochromatic x-rays for medical fluoroscopy and radiography," *Med. Mundi* **7**, 68–77 (1961).

²J. E. Villagran, B. B. Hobbs, and K. W. Taylor, "Reduction of patient exposure by use of heavy elements as radiation filters in diagnostic radiology," *Radiology* **127**(1), 249–254 (1978). Medline

³J. W. Motz and M. Danos, "Image information content and patient exposure," *Med. Phys.* **5**(1), 8–22 (1978).

⁴T. S. Curry, J. E. Dowdey, and R. C. Murry, *Christensen's Physics of Diagnostic Radiology*, 1990. 4th ed. Lea and Febiger, Philadelphia, PA.

⁵N. A. Gkanatsios, W. Huda, and K. R. Peters, "Effect of radiographic techniques (kVp and mAs) on image quality and patient doses in digital subtraction angiography," *Med. Phys.* **29**(8), 1643–1650 (2002).

⁶J. Geleijns *et al.*, "Reference dose rates for fluoroscopy guided interventions," *Radiat. Prot. Dosim.* **80**(1–3), 135–138 (1998).

⁷D. R. Johnson, J. Kyriou, E. J. Morton, A. Clifton, M. Fitzgerald, and E. Macsweeney, "Radiation protection in interventional radiology," *Clin Radiol.* **56**(2), 99–106 (2001).

⁸R. Padovani, G. Bernardi, M. R. Malisan, E. Vaño, G. Morocutti, and P. M. Fioretti, "Patient dose related to the complexity of interventional cardiology procedures," *Radiat. Prot. Dosim.* **94**(1–2), 189–192 (2001).

⁹N. A. Gkanatsios, W. Huda, and K. R. Peters, "Adult patient doses in interventional neuroradiology," *Med. Phys.* **29**(5), 717–723 (2002).

¹⁰G. J. Kemerink, M. J. Frantzen, K. Oei, M. Sluzewski, W. J. van Rooij, J. Wilmink, J. M. van Engelsen. Patient and occupational dose in neuro-interventional procedures. *Neuroradiology* **44**, 522–528 (2002).

¹¹G. Singer, "Occupational radiation exposure to the surgeon," *J. Am. Acad. Orthop. Surg.* **13**(1), 69–76 (2005). Medline

¹²E. Vano, L. Gonzalez, J. M. Fernandez, C. Prieto, and E. Guibelalde, "Influence of patient thickness and operation modes on occupational and patient radiation doses in interventional cardiology," *Radiat. Prot. Dosim.* **118**(3), 325–330 (2006).

¹³C. T. Mehlman and T. G. DiPasquale, "Radiation exposure to the orthopaedic surgical team during fluoroscopy: 'How far away is far enough?'," *J. Orthop. Trauma* **11**(6), 392–398 (1997).

¹⁴E. Vano, "Radiation exposure to cardiologists: How it could be reduced," *Heart* **89**(10), 1123–1124 (2003).

¹⁵E. Vano, C. Ubeda, F. Leyton, P. Miranda, and L. Gonzalez, "Staff radiation doses in interventional cardiology: Correlation with patient exposure," *Pediatr. Cardiol.* **30**(4), 409–413 (2009).

¹⁶V. Tsapaki *et al.*, "Correlation of patient and staff doses in interventional cardiology," *Radiat. Prot. Dosim.* **117**(1–3), 26–29 (2005).

¹⁷T. M. Bashore, "Radiation safety in the cardiac catheterization laboratory," *Am Heart J.* **147**(3), 375–378 (2004).

¹⁸E. J. Hall and D. J. Brenner, "Cancer risks from diagnostic radiology," *Br. J. Radiol.* **81**(965), 362–378 (2008).

¹⁹A. Schweikard, H. Shiomi, and J. Adler, "Respiration tracking in radio-surgery," *Med. Phys.* **31**(10), 2738–2741 (2004).

²⁰A. Schweikard, G. Glosser, M. Bodduluri, M. J. Murphy, and J. R. Adler, "Robotic motion compensation for respiratory movement during radio-surgery," *Comput. Aided Surg.* **5**(4), 263–277 (2000).

²¹J. J. Kresl, J. D. Luketich, and H. C. Urschel, in *Robotic Radiosurgery: Treating Tumors that Move with Respiration*, edited by J. J. Kresl, J. D. Luketich, and H. C. Urschel (Springer, Berlin, 2007).

²²P. J. Keall, V. R. Kini, S. S. Vedam, and R. Mohan, "Motion adaptive x-ray therapy: A feasibility study," *Phys. Med. Biol.* **46**(1), 1–10 (2001).

²³T. Neicu, H. Shirato, Y. Seppenwoolde, and S. B. Jiang, "Synchronized moving aperture radiation therapy (SMART): Average tumour trajectory for lung patients," *Phys. Med. Biol.* **48**(5), 587–598 (2003).

²⁴W. D. D'Souza, S. A. Naqvi, and C. X. Yu, "Real-time intra-fraction-motion tracking using the treatment couch: A feasibility study," *Phys. Med. Biol.* **50**(17), 4021–4033 (2005).

²⁵H. Shirato *et al.*, "Four-dimensional treatment planning and fluoroscopic real-time tumor tracking radiotherapy for moving tumor," *Int. J. Radiat. Oncol., Biol., Phys.* **48**(2), 435–442 (2000).

²⁶S. Shimizu *et al.*, "Detection of lung tumor movement in real-time tumor-tracking radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **51**(2), 304–310 (2001).

²⁷C. Ozhasoglu and M. J. Murphy, "Issues in respiratory motion compensation during external-beam radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **52**(5), 1389–1399 (2002).

²⁸H. Shirato, M. Oita, K. Fujita, Y. Watanabe, and K. Miyasaka, "Feasibility of synchronization of realtime tumor-tracking radiotherapy and intensity-modulated radiotherapy from viewpoint of excessive dose from fluoroscopy," *Int. J. Radiat. Oncol., Biol., Phys.* **60**(1), 335–341 (2004).

²⁹W. K. Laskey, M. Wondrow, and D. R. Holmes, Jr., "Variability in fluoroscopic x-ray exposure in contemporary cardiac catheterization laboratories," *J. Am. Coll. Cardiol.* **48**(7), 1361–1364 (2006).

³⁰M. K. Kalra, S. M. Rizzo, and R. A. Novelline, "Reducing radiation dose in emergency computed tomography with automatic exposure control techniques," *Emerg. Radiol.* **11**(5), 267–274 (2005).

³¹I. A. Elbakri, A. V. Lakshminarayanan, and M. M. Tesic, "Automatic exposure control for a slot scanning full field digital mammography system," *Med. Phys.* **32**(9), 2763–2770 (2005).

³²P. Doyle, D. Gentle, and C. J. Martin, "Optimising automatic exposure control in computed radiography and the impact on patient dose," *Radiat. Prot. Dosim.* **114**(1–3), 236–239 (2005).

³³P. Cooney, D. M. Marsh, and J. F. Malone, "Automatic exposure control in fluoroscopic imaging," *Radiat. Prot. Dosim.* **57**(1–4), 269–272 (1995).

³⁴S. M. R. Rizzo, M. K. Kalra, B. Schmidt, J. Paul, A. Sigal-Cinqualbre, and H. Abada, "Automatic exposure control techniques for individual dose adaptation. Dr. Paul and colleagues respond," *Radiology* **235**(1), 335–336 (2005).

³⁵H. H. Barrett, "Objective assessment of image quality: Effects of quantum noise and object variability," *J. Opt. Soc. Am. A* **7**(7), 1266–1278 (1990).

³⁶H. H. Barrett, J. L. Denny, R. F. Wagner, and K. J. Myers, "Objective assessment of image quality. I. Fisher information, Fourier crosstalk, and figures of merit for task performance," *J. Opt. Soc. Am. A* **12**(5), 834–852 (1995).

³⁷H. H. Barrett, C. K. Abbey, and E. Clarkson, "Objective assessment of image quality. II. ROC metrics, ideal observers, and likelihood-generating functions," *J. Opt. Soc. Am. A* **15**(6), 1520–1535 (1998).

³⁸Y. Seppenwoolde *et al.*, "Precise and real-time measurement of 3D tumor motion in lung due to breathing and heartbeat, measured during radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **53**(4), 822–834 (2002).

³⁹P. M. Djuric *et al.*, "Particle filtering," *IEEE Signal Process. Mag.* **20**(5), 19–38 (2003).

⁴⁰J. Carpenter, P. Clifford, P. Fearnhead, "An improved particle filter for nonlinear problems," *IEEE Proc Radar Sonar Navig* **146**, 2–7 (1999).

⁴¹M. Isard and A. Blake, "Condensation—Conditional density propagation for visual tracking," *Int. J. Comput. Vis.* **29**, 5–28 (1998).

- ⁴²N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proc. F, Radar Signal Process.* **140**(2), 107–113 (1993).
- ⁴³P. Del Moral, (1998). Measure valued processes and interacting particle systems. Application to nonlinear filtering problems. *Ann. Appl. Probab.* **8** 438–495.
- ⁴⁴B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*, Artech House, Boston, MA (2004).
- ⁴⁵T. K. Moon and W. C. Stirling, *Mathematical Methods and Algorithms for Signal Processing*, Englewood Cliffs, Prentice hall, NJ, (2000).
- ⁴⁶B. D. O. Anderson and J. B. Moore, *Optimal Filtering* (Prentice-Hall, Englewood Cliffs, NJ, 1979).
- ⁴⁷D. Beeman, "Some multistep methods for use in molecular dynamics calculations," *J. Comput. Phys.* **20**(2), 130–139 (1976).
- ⁴⁸C. Musso, N. Oudjane, and F. L. Gland, "Improving regularized particle filters," in *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, edited by A. Doucet, N. De Freitas, and N. Gordon (Springer-Verlag, New York, 2001), Chap. 12, pp. 247–271.
- ⁴⁹A. E. Lujan, E. W. Larsen, J. M. Balter, and R. K. Ten Haken, "A method for incorporating organ motion due to breathing into 3D dose calculations," *Med. Phys.* **26**(5), 715–720 (1999).