The 2009 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 11-15, 2009 St. Louis, USA

Robust Sensor Planning for a Partially Known Moving Target: Application to a Dynamic X-ray Imaging System

John D. Yamokoski, *Student Member, IEEE*, and Scott A. Banks

Abstract— **This paper discusses the extension of sample-based planning methods to the application of sensor planning for a partially unknown, moving target. This is achieved by modeling the target as a time-varying stochastic process with known mean and variance. Using the unscented transform, the target's statistics are then propagated through a user-supplied sensing effectiveness metric. The planner is demonstrated for a dynamic X-ray imaging platform tasked with viewing the human knee-joint center during normal, over-ground walking.**

I. Introduction

In the last two decades, the use of fluoroscopy or X-ray images as a means to measure in-vivo bone and implant motions has matured to a reliable biomechanical technique. Beginning with work by Banks et al. [1, 2], programs worldwide have since been established to pursue technical development and clinical applications of registration-based image measurements of in-vivo bone and implant motions. Despite these positive developments, clinical application and broadened research utilization of these measurement techniques are fundamentally limited by available imaging hardware. Existing fluoroscopy equipment provides superb images of anesthetized patients resting on an operating table, but these systems were not designed for large volume dynamic motion capture on moving subjects. Our group, along with others across the globe, has been developing novel dynamic imaging platforms which can move with patients to capture diagnostically insightful views of their joints during normal activity (Fig 1).

Our solution to the hardware problem is to mount light-weight X-ray equipment onto commercially available robotic manipulators. However, the coordination of two six DOF robot arms for the purpose of X-ray photography is not trivial. Even before permitting a patient within the reachable workspace of the robots, we need to 1) prove the feasibility of the measurement activity considering the geometric and dynamic constraints of the robots and 2) quantify the relative safety of the patient and efficiency of the sensing activity. Methods to answer these preliminary questions are found in the realm of motion planning and specifically in the specialized domain of view or sensor planning.

Even within the domain of sensor planning there are many different research thrusts of which only a few are applicable to our problem statement. Excellent surveys of automated 3D modeling and object inspection exist [3, 4]. These techniques focus on computing either the minimal number of views required to reconstruct a static object or, given the required views, the optimal path to visit these in minimal time or effort. An interesting tangent to this problem is the combined exploration/navigation problems of unknown environments (e.g. [5, 6]). While both research thrusts are interesting, our object of interest is moving (possibly with limited *a priori* knowledge of its motion) in a known environment.

A more relevant application area can be found in pursuit-evasion games. A short survey by Cheng summarizes a variety of techniques that have been developed to address different facets of this problem [7]. While recent work has focused on incorporating more realistic sensor models (e.g. [8]) or limitations on the kinematics or dynamics of the robotic agents (e.g. [9]), these methods often rely on the geometric properties of planar problem domains for their solutions.

Randomized or probabilistic motion planning algorithms are well-suited to solving high-dimensional planning problems [10]. A sample-based algorithm has been developed by Kehoe for sensor planning of stationary ground targets in cluttered environments by aerial vehicles [11]. By combining the rigorous development of a sensing metric with the adaptation of randomized dense tree (RDT) methods [12, 13], Kehoe was able to generate useful solutions for real-world sensor planning problems.

This paper discusses a novel extension to the method previously reported by Kehoe. We will present an algorithm for sensor planning of a moving target with partially known state information. Section II formalizes our problem definition. Section III gives a general overview of the algorithm and then details our novel contributions. Section IV details the application of this work to our dynamic X-ray imaging platform. Finally, sections V and VI provide experimental results and concluding remarks.

Manuscript received March 1, 2009. This work was supported in part by the University of Florida.

J. Yamokoski is with the Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, FL 32611 USA (phone: 352-392-9638; fax: 352-392-9638; e-mail: yamokosk@ufl.edu).

S. Banks is with the Department of Mechanical and Aerospace Engineering and the Department of Orthopaedics and Rehabilitation, University of Florida, Gainesville, FL 32611 USA (e-mail: banks@ufl.edu).

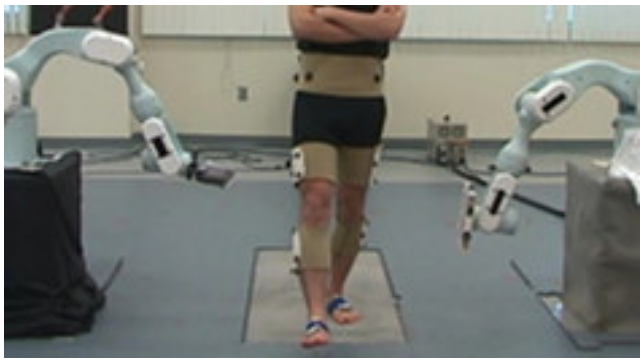

Figure 1. Prototype imaging platform that will be used for dynamic X-ray imaging as well as cone-beam computed tomography and traditional imaging tasks.

II. Problem Definition

A. Problem domain

To assist with further discussions, we will formally define the problem domain. Let $V = R^3 \times SO(3) \times T \subset R^7$ represent a standard, six dimensional Cartesian space augmented by time. A point in V represents the location and orientation of a generic 3D frustum at a particular instance in time. Let $F \subset R^p$ represent the parameter space of the frustum so that a point in F uniquely defines the size and shape of the frustum. Then the sensor space, $\mathbf{X}$ is simply $\mathbf{X} = V \times F \subset R^{7p}$.

The planning algorithm operates directly in $\mathbf{X}$, however functional mappings for a point $x \in \mathbf{X}$ into corresponding points in both the Cartesian workspace $W \subset R^6$ of the robots as well as the joint spaces $C_i \subset R^6$ for each robot $\Re_i$ are assumed to exist.

B. Target model

To account for a non-stationary target, we will define the target as a family of trajectories which belong to a specific trajectory population relevant to the sensing application. More specifically, we will treat the state time-history of the target as a time-varying stochastic process, $Y(t)$, where

$$\bar{y}_j = E\left[Y(t)\right]_{t=t_j} \tag{1}$$

$$P_{y_j} = E\left[\left(Y(t) - \bar{y}_j\right)\left(Y(t) - \bar{y}_j\right)^T\right]_{t=t_j} \tag{2}$$

where $\bar{y}_j$ and P_{y_j} are the mean and covariance of $Y(t_j)$ respectively.

C. Sensor effectiveness

Next we assume there exists a scalar function taking as input a realization of the target's time-history and a configuration of the sensing system. This function generates a measure of the sensing effectiveness of the configuration for the particular target realization. Mathematically, we assume there exists a real-valued function $f : (x, y) \rightarrow [0, 1]$.

D. Problem statement

The objective of the sensor planning algorithm is to find an open-loop control strategy which maximizes the expected sensing efficiency across the entire path.

III. Algorithm Description

A. Overview

Sample-based motion planning algorithms are recognized by the greater robotics community as viable methods to solve complex, high-dimensional planning problems. All sample-based methods share one common feature – they only implicitly model the problem space by sampling it at discrete points. This feature allows them to efficiently model very complex and high dimensional problem spaces using a finite amount of computer memory.

In general, RDT methods take as input a set S of randomly sampled points from the problem space. The algorithm then incrementally builds a tree data structure from an initial sample and returns a simply connected graph $G(N, E)$ of S where the nodes N are the sample points of S and the edges E represent collision free paths between the nodes. The graph G approximates the collision free topology of the problem space.

There are four major components to the algorithm: node selection, vantage point generation, extension, and solution check [11]:

1. **Node selection:** A node, N_i, from the existing search tree is selected probabilistically based on the node's fitness or weight.
2. **Vantage point generation:** Vantage points, x_j, are pseudo-randomly generated for a target (if stationary) or target segment (if moving) which has not yet been sensed.
3. **Extension:** Candidate paths are generated using a user-specified local planning method to connect the selected node, N_i, and the generated vantage points, x_j. Each candidate path is evaluated for its sensing effectiveness. The best candidate is chosen and then added to the search tree as a new branch.
4. **Solution check:** The new branch is checked to determine if it completes the sensing objective. If so, a trajectory is planned to the terminal configuration and the solution upper bound is updated.

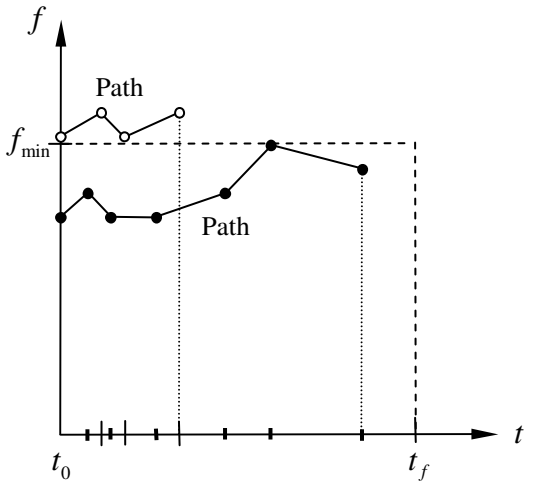

Figure 2. Illustrative example of two fictitious paths and their expected sensing efficiency.

B. Node selection and fitness evaluation

The probability of selecting any particular node for expansion is proportional to that node's fitness. More specifically, the stochastic universal sampling (SUS) algorithm, developed originally for use in genetic optimization routines [14], was chosen for node selection.

To aid discussion, Fig 2 illustrates the metric evaluations along two fictitious paths from a search tree. The first path has an average metric score above a user specified minimum. The second path has a lower average metric score; however it is close to a completed sensing solution. For our purposes, a node fitness rubric which penalized long, low-average metric paths and rewarded short, high-average metric paths was desired. Through a series of simple experiments we developed a node fitness equation,

$$w_i = \left[\alpha \left(\frac{F_i}{F_{\min}} \right)^2 + (1-\alpha) \left(\frac{d_{\max} - d_i}{d_{avg}} \right)^2 \right]^{1/2} \tag{3}$$

where,

$$F_i = \int_{t_i} f(x,y)dt \tag{4}$$

$$F_{\min} = (t_f - t_0) f_{\min} \tag{5}$$

And $d_{\max}$ and d_{avg} are the maximum and average path lengths in the search tree, respectively, and d_N is the path length up to node *i*. The free parameter, α, allows one to manipulate the overall behavior of the search tree. Setting α=1 removes the penalty on long paths and therefore results in a greedy or depth-first search type of behavior. Conversely, by setting α=0, node sensing effectiveness is ignored and only shorter paths will be favored for selection and subsequent expansion, resulting in breath-first search behavior.

C. Sensing effectiveness for a partially known target

As stated in the problem definition, the target time-history is modeled as a family of trajectories with a mean $\bar{y}_i$ and covariance P_{y_i}. The goal is to understand how this target information (or target uncertainty) is propagated through the sensing effectiveness metric. The unscented transformation (UT), a technique often found in non-linear Kalman filters, can be used for this purpose. By computation of a set of 2N+1 carefully chosen sigma points $S_j = \{W_j, \chi_j\}$, where N is the dimension of our target's state vector, we can propagate the statistics of the target data [15]. The sigma points are calculated as follows,

$$\begin{aligned}
&\chi_0 = \bar{y}_i && j = 0 \\
&W_0 = \kappa/(N+\kappa) && \\
&\chi_j = \bar{y}_i + \left(\sqrt{(N+\kappa)P_{y_i}}\right)_j && j = 1 \ldots N \\
&W_j = 1/(2(N+\kappa)) && \\
&\chi_j = \bar{y}_i - \left(\sqrt{(N+\kappa)P_{y_i}}\right)_j && j = N+1 \ldots 2N \\
&W_j = 1/(2(N+\kappa)) &&
\end{aligned}$$

where κ is a scaling parameter and $\left(\sqrt{(N+\kappa)P_{y_i}}\right)_j$ is the j[th] row of the matrix square root of $(N+\kappa)P_{y_i}$. These sigma points are then propagated through the sensing metric,

$$f_{j,i} = F(x, \chi_j) \tag{6}$$

When the propagated points are combined with the sigma weights, the estimated mean and covariance of f_i can be computed as,

$$\bar{f}_i = \sum_j W_j f_{j,i} \tag{7}$$

$$P_{f_i} = \sum_j W_j \left(f_{j,i} - \bar{f}_i\right)\left(f_{j,i} - \bar{f}_i\right)^T \tag{8}$$

Maximizing this expected sensing effectiveness is then the focus of sensor planning algorithm.

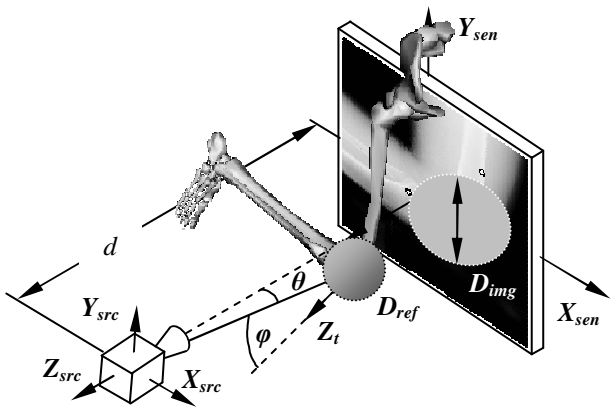

Figure 3. Illustration of the various parameters which comprise the sensing effectiveness metric.

IV. Application

A. Overview

We have begun early testing of a new imaging platform which, in addition to standard fluoroscopy and cone-beam computed tomography tasks, could provide the ability to move with patients to capture diagnostically insightful views of their joints during normal activity (Fig 1). With these abilities to acquire dynamic images of joints in motion and tomographic scans, this novel platform could provide all the required information needed to implement routine image-based measurement of dynamic 3D skeletal motion. Application of the sensor planning algorithm presented in this paper relies heavily on the existence of a meaningful sensing effectiveness metric. In what follows, we will develop a metric for this dual-arm robotic imaging system.

B. Metric description

The goal of this metric is to evaluate the utility of a particular robot configuration. The metric can be broken into three components. The first two components evaluate the alignment of a configuration's actual imaging axis, $P_{src\to sen}$, to the target's desired view vector, Z_t. The angles ϕ and θ, between the desired view vector and the actual viewing axis and the source to target vector, $P_{src\to tg}$, and actual viewing axis respectively, are parameters to Gaussian radial basis functions centered at zero degrees (Fig 3).

$$f_\phi = gaussian\_rbf(\phi,0) \tag{9}$$

$$f_\theta = gaussian\_rbf(\theta,0) \tag{10}$$

Note that no explicit constraint is set on the alignment of the X-ray source and detector (e.g. the alignment of Z_{src} and Z_{sen}). This equipment is mounted to robotic manipulators that have excellent spatial resolution. Therefore the location and orientation of this equipment is assumed to be measurable at the moment of image acquisition. If that assumption holds true and the angles θ and ϕ are minimized, then we will be able to deduce the 3D spatial kinematics of the target object using standard registration methods.

The third component attempts to evaluate the remaining imaging parameters by approximation of the image area projected by the object of interest. The objects we want to image range from the small bones in the human ankle joint system to total hip replacements. To capture such a diverse set of potential targets, we define a 3D sphere and associated reference diameter, D_{ref}, which is meant to bound our region of interest (Fig 3).

For a first-order approximation, we assume that the projection of our reference sphere is a perfect circle with diameter D_{img}. Utilizing the properties of similar triangles, the relationship between D_{ref} and D_{img} is,

$$\frac{D_{img}}{D_{ref}} = \frac{d}{r*\cos(\theta)} \tag{11}$$

where $d = \left\|P_{src\to sen}\right\|$ is the distance between the source and sensor robots, $r = \left\|P_{src\to tg}\right\|$ is the tracking distance, and θ is,

$$\theta = \arccos\left(\frac{P_{src\to tg}\cdot P_{src\to sen}}{r*d}\right) \tag{12}$$

Using (5) and normalizing by the total area of the detector (in units of pixels^2), we can write an approximation of the projected image area of the reference object as a function of the sensor's configuration:

$$\frac{A_{img}}{A_{total}} = \hat{A}_{img} = \frac{\pi}{4}\left(\frac{d}{S*r*\cos(\theta)}D_{ref}\right)^2 \tag{13}$$

Finally to finish the image area metric, the normalized image area is passed as a parameter into a Gompertz function.

$$f_{area} = gompertz\left(\hat{A}_{img}\right) \tag{14}$$

This class of function traditionally is used to model time series data where growth is slowest at the start and end of a time period but rapid in the middle. This maximizes function sensitivity when the desired object fills much, but

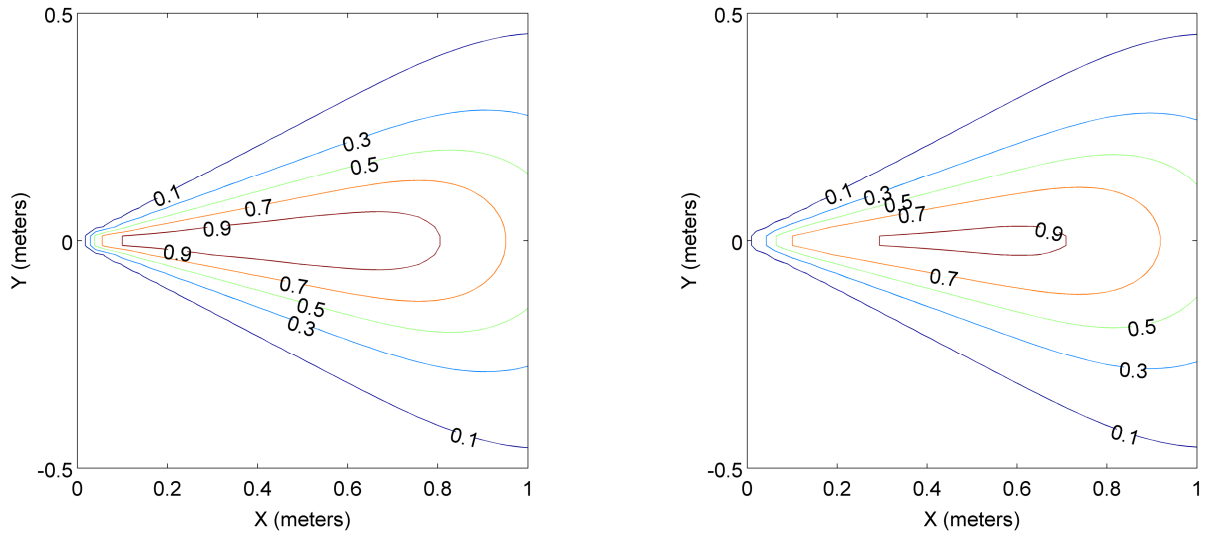

Figure 4. Contour map of the example sensing effectiveness metric for a target located at the origin with a desired view vector in the -Y direction. The contour was calculated for a target with (a) no uncertainty and (b) with uncertainty.

not all (e.g. $\hat{A}_{img} > 80\%$), of the sensor. The final metric is computed as the product of the three components,

$$f_{tg}(x, y) = f_{area} * f_{\phi} * f_{\theta} \tag{15}$$

To demonstrate the metric in operation, we selected the NDI-160-20 solid-state X-ray detector (Varian Medical Systems, Palo Alto, CA) where $S = 0.179$ meters. Choosing $D_{ref} = 0.09$ meters, which can bound the geometry of a standard total knee replacement, we generated a contour map of the metric (Fig 4a). The map was generated for a target with no uncertainty located at the origin with a desired view vector of (0,-1,0). The actual image axis was arbitrarily constrained to align with the global –y direction and the separation distance between the source and sensor robots was fixed at 1.0 meters. Metric evaluations were then performed over an X-Y grid with a uniform spacing of 0.02 meters.

V. Results

A twelve-camera passive-marker motion camera system (Motion Analysis Corporation, Santa Rosa, CA) was used to measure one subject's body motion during normal-speed, over-ground walking. Data from several trials were averaged together and fit with third-order, piece-wise cubic splines to yield a mean position and velocity of the knee-joint center. Desired orientation of the target frame, which determines the desired view vector of the object, was chosen to be normal to the sagittal or X-Z plane. The covariance of the target state information was chosen arbitrarily. The target's state variables were assumed to be independent and normally distributed with the following standard deviations,

$$\sigma_x = \sigma_y = \sigma_z = 0.01 \text{ meters}$$

$$\sigma_\theta = \sigma_\phi = \sigma_\varphi = 5 \text{ degrees}$$

$$\sigma_{\dot{x}} = \sigma_{\dot{y}} = \sigma_{\dot{z}} = 0.1 \text{ m/s}$$

$$\sigma_{\dot{\theta}} = \sigma_{\dot{\phi}} = \sigma_{\dot{\varphi}} = 0 \text{ degrees/s.}$$

Figure 4b demonstrates the effect this uncertainty has on the metric. The contour map was generated under the same conditions as Figure 4a except with the target deviations defined above.

To limit the computational requirements of the planner's execution, connection of states during the extension step simply used a straight line in the joint space of the manipulators and assumed every path was collision free.

The planner was run three times (n=3) for 500 iterations. The free parameter, α, was chosen to be 0.5 and the minimum sensing effectiveness from (5) was chosen to be 0.85. Figure 5a shows the maximum and mean node weight at the completion of each iteration. The maximum node weights exhibited monotonically increasing values with greater slope than the mean node weights. Figure 5b shows the best normalized sensing efficiency over all potential paths at the completion of each iteration. The error bars in Figure 5b represent the 95% confidence interval. Sensing effectiveness also increased monotonically.

VI. Conclusion

We have demonstrated the extension of sample-based planners to the application of sensor planning for partially-known moving targets. The node weighting scheme proposed in (3) yields nice node weight population behavior as demonstrated in Figure 5a. The ability of the algorithm to maximize the expected sensing efficiency is demonstrated in Figure 5b.

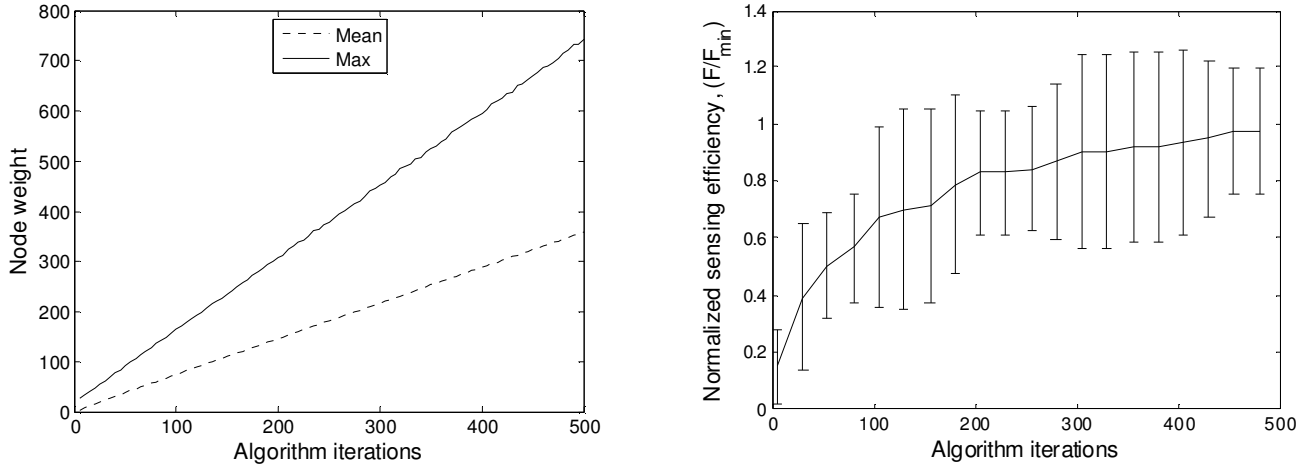

(a) **(b)**

Figure 5. n=3 runs of the sensor planning algorithm for 500 iterations. (a) The average and maximum node weights at the completion of each iteration. (b) The best normalized sensing effectiveness over all potential paths at the completion of each planning iteration. Error bars represent the 95% confidence interval.

From the standpoint of the algorithm's power to prove the existence or non-existence of a solution, the advantages of probabilistically representing the target state information is clear. In one single pass, we are able to replicate what would require a Monte Carlo experiment for a non-probabilistic target framework to achieve the same sensitivity analysis. Performing such an experiment could be expensive computationally since most real-world motion planning problems require significant computational time, especially when using sample-based methods.

Another less obvious advantage of the probabilistic target definition is that it allows the application designer more freedom in defining their planning query. For instance, the probabilistic framework can be used to model not only the variability of human motion from one subject or population group to the next; but the same framework also allows freedom in the definition of the desired view vector. By choosing appropriate variances for the target's orientation, one can simultaneously check for the existence of any path which might, for instance, use a desired view cone instead of a single view vector.

Future efforts will be directed towards continual characterization of the algorithm for different levels of target uncertainty. We also plan to use a more sophisticated local planner which accounts for the robot's kinematic and dynamic constraints. Finally, for our application, we intend to build a database of common activities and desired viewing constraints to assist clinicians in planning use the dynamic imaging platform.

ACKNOWLEDGMENT

J.D. Yamokoski gratefully acknowledges members of the Orthopaedic Biomechanics laboratory for their insightful comments concerning this work.

REFERENCES

[1] S. Banks, "Model based 3d kinematic estimation from 2d perspective silhouettes: Application with total knee prostheses," Ph.D. dissertation, Massachusetts Institute of Technology, 1992.

[2] S. Banks and W. Hodge, "Accurate measurement of three-dimensional knee replacement kinematics using single-plane fluoroscopy," *IEEE Transactions on Biomedical Engineering*, vol. 43, pp. 638–649, 1996.

[3] W. Scott, G. Roth, and J. Rivest, "View planning for automated three-dimensional object reconstruction and inspection," *ACM Computing Surveys (CSUR)*, vol. 35, no. 1, pp. 64–96, 2003.

[4] S. Dutta Roy, S. Chaudhury, and S. Banerjee, "Active recognition through next view planning: a survey," *Pattern Recognition*, vol. 37, no. 3, pp. 429–446, 2004.

[5] P. Wang and K. Gupta, "View planning for exploration via maximal C-space entropy reduction for robot mounted range sensors," *Advanced Robotics*, vol. 21, no. 7, pp. 771–792, 2007.

[6] H. Gonzalez-Banos and J. Latombe, "Robot Navigation for Automatic Model Construction Using Safe Regions," Lecture Notes in Control and Information Sciences, pp. 405–416, 2001.

[7] P. Cheng, "A Short Survey on Pursuit-Evasion Games," *Department of Computer Science, University of Illinois at Urbana-Champaign*, 2003.

[8] B. Gerkey, S. Thrun, and G. Gordon, "Visibility-based pursuit-evasion with limited field of view," *The International Journal of Robotics Research*, vol. 25, no. 4, p. 299, 2006.

[9] B. Tovar and S. LaValle, "Visibility-Based Pursuit-Evasion with Bounded Speed," in *Proceedings Workshop on Algorithmic Foundations of Robotics*. Springer, 2006.

[10] S. LaValle, *Planning algorithms*. Cambridge University Press, pp. 185-248, 2006.

[11] J. J. Kehoe, "Trajectory generation for effective sensing of a close proximity environment," Ph.D. dissertation, University of Florida, 2007.

[12] S. LaValle and J. Kuffner Jr, "Randomized kinodynamic planning," *The International Journal of Robotics Research*, vol. 20, no. 5, p. 378, 2001.

[13] D. Hsu, R. Kindel, J. Latombe, and S. Rock, "Randomized kinodynamic motion planning with moving obstacles," *The International Journal of Robotics Research*, vol. 21, no. 3, p. 233, 2002.

[14] J. Baker, "Reducing bias and inefficiency in the selection algorithm," in *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application table of contents*. L. Erlbaum Associates Inc. Hillsdale, NJ, USA, 1987, pp. 14–21.

[15] S. Julier and J. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," in *Int. Symp. Aerospace/Defense Sensing, Simul. and Controls*, vol. 3, 1997.