

LQG-Based Planning, Sensing, and Control of Steerable Needles

Jur van den Berg¹, Sachin Patil¹, Ron Alterovitz¹, Pieter Abbeel², Ken Goldberg²

Abstract This paper presents a technique for planning and controlling bevel-tip steerable needles towards a target location in 3-D anatomy under the guidance of partial, noisy sensor feedback. Our approach minimizes the probability that the needle intersects obstacles such as bones and sensitive organs by (1) explicitly taking into account motion uncertainty and sensor types, and (2) allowing for efficient optimization of sensor placement. We allow for needle trajectories of arbitrary curvature through duty-cycled spinning of the needle, which is believed to make a needle path small-time locally “trackable” [13]. This enables us to use LQG control to guide the needle along the path. For a given path and sensor placement, we show that a priori probability distributions of the needle state can be estimated in advance. Our approach then plans a set of candidate paths and sensor placements and selects the pair for which the estimated uncertainty is least likely to cause intersections with obstacles. We demonstrate the performance of our approach in a modeled prostate cancer treatment environment.

1 Introduction

We consider the problem of planning, sensing, and controlling a bevel-tip steerable needle towards a target in 3-D anatomy with obstacles, such as sensitive and impenetrable tissue. Needles are used in many forms of medical diagnosis and treatment, and accurately reaching a specific target is required in many procedures such as tissue biopsies and placement of radioactive seeds for cancer treatment. Bevel-tip steerable needles are asymmetric-tip, flexible needles that move along trajectories of constant curvature κ_0 when pushed forward [27]. The direction of motion can be changed by reorienting the bevel tip through twisting of the needle at its base. This allows for steering the needle around anatomical obstacles towards previously in-

¹University of North Carolina at Chapel Hill. {berg, sachin, ron}@cs.unc.edu.

²University of California at Berkeley. {pabbeel, goldberg}@berkeley.edu

accessible targets, and allows for significantly reducing patient trauma by avoiding the puncturing of sensitive tissues.

Planning and controlling the motion of a steerable needle is a challenging problem. A steerable needle is controlled from its base through only insertion and twisting, and we do not allow retractions and re-insertions as that results in excessive tissue damage. As such, a steerable needle is a highly underactuated non-holonomic system. In fact, the needle is not small-time locally controllable, and a natural needle path (with curvature κ_0) is not small-time locally “trackable” (i.e., the deviation with respect to the path is not small-time locally controllable) [13]. In addition, the motion of the needle is subject to uncertainty due to tissue inhomogeneity, tissue deformation, needle torsion, etc. [7, 23]. To address this issue, we follow the suggestion of Kallam [13] that a path with a smaller curvature $0 \leq \kappa < \kappa_0$ is small-time locally trackable and that this can be achieved using duty-cycled spinning of the needle during insertion. This would enable us to use feedback control to guide the needle along a pre-planned path. Our experiments suggest that this is indeed the case. The sensor feedback, however, may be noisy and partial as current medical imaging technology does not allow for measuring the full state of the needle tip (the imaging resolution is often too low to infer its orientation, for instance [12], and often only provides planar views in real-time feedback situations).

Our objective is to compute a sensor placement and a needle path to the target location, such that the path’s execution using LQG control has a minimal probability of intersecting obstacles in the anatomy, given a stochastic model of the motion and sensing uncertainty. Our approach is as follows. First, we build on work of [16] to encapsulate the (high-frequency) duty cycled spinning of the needle in a higher-level kinematic model that allows direct control of the curvature of the needle motion. We then derive an LQG controller (consisting of a Kalman filter for state estimation and an LQR control policy) for the extended kinematic model to optimally guide the needle along a given path. Based on the sensor placement, we can compute *in advance* the a priori probability distributions of the state of the needle along the path [25]. From these distributions and a geometric model of the anatomical obstacles, we can quickly compute the probability that the needle will intersect obstacles. Our method then plans a set of candidate paths using a variant of the RRT algorithm [14] and samples a set of feasible sensor placements, and then selects the pair for which this probability of obstacle intersection is minimal.

The type and placement of sensor(s) can have a big influence on which path is optimal (see Fig. 1). For example, if a sensor obtains a 2-D projection along the x -axis of the 3-D position of the needle tip, there will be more uncertainty in the x -coordinate of the needle state than in the y - and z -coordinates. If we then have to steer the needle through a passage that is narrow in either the y - or the x -direction, our path planner will prefer the passage for which the probability of intersecting obstacles is less, which is the one that is narrower in the y -direction. Given a set of candidate paths and the space of possible sensor placements, our approach will choose a needle path and find an axis along which a projection is obtained that minimizes the probability of collisions for the chosen needle path. We demonstrate




Fig. 1 Two examples of sensor placement, in which (left) only the x - and z -coordinate and (right) only the y - and z -coordinate of the needle-tip are measured by the imaging device (blue). Different paths will be optimal even as the obstacles (grey) and target location (cross) are the same.

the performance of our approach in modeled prostate cancer treatment environments with simulated uncertainty for different examples of sensor models.

The remainder of this paper is organized as follows. We discuss previous work in Section 2, and review the kinematic model of a steerable needle in Section 3. In Section 4 we derive an LQG controller to optimally guide the needle along a given path. In Section 5, we show how to estimate the probability of obstacle intersection for a given path and sensor placement, and present our path and sensor planner. We present simulation results in Section 6 and conclude in Section 7.


2 Related Work

A significant body of previous work exists on planning and/or controlling level-tip steerable needles. A kinematic model for a steerable needle generalizing a unicycle model was introduced in [27], and has been used by most subsequent work on needle steering, including this paper. In [16], it was shown that in addition to the insertion and rotation speed, the curvature of the needle path can be controlled through duty-cycled spinning of the needle during insertion.

2-D planners that address motion uncertainty have been presented in [2, 3], which optimize a Markov decision process (MDP) over a discretized state space to provide feedback control assuming full state observation. The approach was extended in [22] and integrated with imaging feedback. In [12], a feedback controller is presented to stabilize the needle in a plane. Tissue deformation is taken into account in the planner of [1], which optimizes a path using 2-D FEM simulation of soft tissue.

Needle path planners for 3-D environments with obstacles have been proposed in [8, 9], based on optimization and inverse kinematics, respectively. Rapidly-exploring random trees (RRTs) have been used in [29, 30, 20] to explore the entire 3D space of feasible paths. These approaches do not address issues such as uncertainty in motion and sensing. A diffusion-based planner was introduced in [18], but it does not take into account obstacles or sensor types. A feedback controller for 3-D needle steering was presented in [11], and proved robust against motion devia-

Fig. 2 Local coordinate frame attached to the needle tip. Its kinematics are shown by yellow arrows. Figure reproduced with permission from [9].



tion and sensing noise, even for a greedy control policy. The approach is not able to guide the needle along a prescribed path and does not address obstacle avoidance.

The approach we present in this paper extends these previous works and applies to 3-D environments with obstacles, takes into account motion and sensing uncertainty, and does not require discretization of the state-space. Our approach to needle steering is the first that specifically addresses the sensing capabilities and its effect on optimizing the path.

3 Needle Kinematics

We base our motion model of a bevel-tip steerable needle on the idealized kinematic model of [27], using the nomenclature of [8], in which the needle state is represented by a rigid body transformation. This model assumes that the motion of the needle is fully determined by the motion of the tip, which is beveled such that it follows a perfect arc of curvature κ_0 when pushed forward, independent of insertion speed and tissue properties. The model further assumes that the needle is flexurally flexible (it bends to follow the needle tip), but axially and torsionally stiff, such that the insertion and twisting of the needle at its base is directly transmitted to its tip.

The state of the needle tip can be described by a 3-D rigid body transformation relative to a world coordinate frame, which is compactly represented by a 4×4 transformation matrix $X \in SE(3)$;

$$X = \begin{bmatrix} R & \mathbf{p} \\ 0 & 1 \end{bmatrix}, \quad (1)$$

where $R \in SO(3)$ is a 3×3 rotation matrix describing the rigid body's orientation, and $\mathbf{p} \in \mathbb{R}^3$ a vector describing the rigid body's position.

The kinematics of a rigid body, i.e. the evolution of its state over time, can generally be described as follows. Let $\mathbf{v} \in \mathbb{R}^3$ and $\mathbf{w} \in \mathbb{R}^3$ be the vectors of instantaneous linear and angular velocities, respectively, expressed in the local coordinate frame attached to the rigid body. Then (using notation $'$ to refer to the time derivative):

$$X' = XU, \quad U = \begin{bmatrix} [\mathbf{w}] & \mathbf{v} \\ 0 & 0 \end{bmatrix}, \quad (2)$$

where 4×4 matrix $U \in se(3)$ is the twist of the rigid body. The notation $[\mathbf{a}]$ for a vector $\mathbf{a} = [a_x \ a_y \ a_z]^T \in \mathbb{R}^3$ refers to the following 3×3 skew-symmetric matrix:

$$[\mathbf{a}] = \begin{bmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{bmatrix}. \quad (3)$$

When the twist U is constant, the state of the rigid body at time t given the initial state at time 0 is computed explicitly by integrating Eq. (2), for which a closed form expression exists [8]:

$$X(t) = X(0)\exp(tU). \quad (4)$$

For the steerable needle, the local coordinate frame is rigidly attached to the tip of the needle such that the z -axis points along the forward direction of the needle, and the y -axis points along the bevel direction (see Fig. 2). The motion of the needle is determined by two control inputs: the linear forward speed (i.e. the speed with which the needle is inserted), denoted v , and the bevel orientation speed (i.e. the speed with which the needle is twisted at its base), denoted ω . Hence, the linear and angular velocities of the needle tip given the control inputs $v \geq 0$ and ω are $\mathbf{v} = [0 \ 0 \ v]^T$ and $\mathbf{w} = [v\kappa_0 \ 0 \ \omega]^T$, respectively, where κ_0 is the (fixed) curvature of the arc the needle follows through the tissue.

This needle model is constrained by the fact that the curvature κ_0 of the needle paths is fixed. In recent work, however, Minhas et al. [16] show that by performing duty cycled spinning of the needle during insertion, any curvature κ of the needle motion between 0 and κ_0 can be approximated. The greatest degree of curvature ($\kappa = \kappa_0$) is achieved by no spin at all, while a straight trajectory ($\kappa = 0$) is created by constantly spinning the needle at a high (infinite) rate during insertion. Any trajectory in between these two extrema can be approximated by duty cycling the spinning in a spin-stop-spin-stop fashion. Longer stop intervals create steeper curvature of the needle, and shorter stop intervals create straighter trajectories. To be precise, the proportion $0 \leq \alpha \leq 1$ of the time spent in spin intervals to approximate a curvature of $0 \leq \kappa \leq \kappa_0$ is given by:

$$\alpha = 1 - \kappa/\kappa_0. \quad (5)$$

Let the needle perform a 2π rotation each spin interval, and let the spin intervals be of a small and constant duration δ . Then, the period of one spin-stop cycle is δ/α . In order to incorporate the duty cycling into the kinematic model of the needle, the control input $\omega(t)$ (parameterized by time t) is adjusted to:

$$\omega(t) = \begin{cases} 2\pi/\delta + w & \text{if } j \leq t/(\delta/\alpha) < j + \alpha & \text{(spin)} \\ w & \text{if } j + \alpha \leq t/(\delta/\alpha) < j + 1, & \text{(stop)} \end{cases} \quad (6)$$

for any $j \in \mathbb{Z}$, where w is the higher-level control of the speed with which the needle is rotated at its base (at the lower level, it is augmented with rapid 2π rotations).

By taking the limit $\delta \rightarrow 0$, the resulting high-level kinematic model is similar to the low-level kinematic model described above, with the difference that ω is replaced by w , and the curvature κ is added to the set of control inputs. The high-level twist U given high-level control inputs v , w and κ is then given by:

$$U = \begin{bmatrix} [\mathbf{w}] & \mathbf{v} \\ 0 & 1 \end{bmatrix}, \quad \mathbf{v} = [0 \ 0 \ v]^T, \quad \mathbf{w} = [v\kappa \ 0 \ w]^T. \quad (7)$$

To account for the uncertainty the motion of the needle is subject to due to tissue inhomogeneity, tissue deformation, needle torsion, etc., we augment the model by assuming that the twist is corrupted by additive noise \tilde{U} drawn from a zero-mean Gaussian distribution with variance M :

$$\tilde{U} = \begin{bmatrix} [\tilde{\mathbf{w}}] & \tilde{\mathbf{v}} \\ 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} \tilde{\mathbf{v}} \\ \tilde{\mathbf{w}} \end{bmatrix} = \mathbf{m} \sim \mathcal{N}(\mathbf{0}, M), \quad (8)$$

The stochastic kinematics of the needle state X is then given by:

$$X' = X(U + \tilde{U}). \quad (9)$$

4 LQG Control for Steerable Needles

Let us be given a needle path Π consisting of states \hat{X} and control input twists \hat{U} formed from control inputs $\hat{v} > 0$, \hat{w} and $0 \leq \hat{\kappa} < \kappa_0$ (see Eq. (7)), such that

$$\hat{X}' = \hat{X}\hat{U}, \quad \hat{X} = \begin{bmatrix} \hat{R} & \hat{\mathbf{p}} \\ 0 & 1 \end{bmatrix}, \quad \hat{U} = \begin{bmatrix} [\hat{\mathbf{w}}] & \hat{\mathbf{v}} \\ 0 & 0 \end{bmatrix}. \quad (10)$$

That is, the path is consistent with the needle kinematics without noise, and as conjectured in [13], the path is small-time locally trackable since $0 \leq \hat{\kappa} < \kappa_0$.

During control of the needle along the path Π , we can assume that we obtain potentially noisy and partial observations of the state as feedback from sensors, in order to compensate for unexpected needle motion. We assume that this feedback will be according to the following (general) sensor model:

$$\mathbf{z} = h(X, \mathbf{q}), \quad \mathbf{q} \sim \mathcal{N}(\mathbf{0}, Q), \quad (11)$$

where \mathbf{z} is a vector of measurements that relates to the state X through function h , and \mathbf{q} is the measurement noise drawn from a zero-mean Gaussian with variance Q .

To control the needle along the needle path Π , we use the LQG-controller, since it provides optimal control for linear Gaussian motion and sensor models with a quadratic cost function penalizing deviation from the path. The LQG controller uses a Kalman filter for state estimation in parallel with an LQR control policy. Since our motion and sensor model are non-linear, we approximate them with local linearizations around the path. This is reasonable as the needle is controlled to stay close to the path during execution.

For purposes of linearization, we will not directly control the state itself, but rather the *deviation* of the state with respect to the path. This is also convenient for dealing with the 3-D orientation in the needle state, which either has a redundant but internally constrained representation (e.g. a quaternion or rotation matrix) or a minimal but singularity-prone representation (e.g. Euler angles). Assuming that the deviation is small, the orientation deviation is “far away” from any singularities, and can hence safely be represented by three mutually unconstrained parameters. Whereas many other works use quaternions to represent rotations [10, 15, 24], we

will linearize for the 4×4 matrix $X \in SE(3)$ (roughly following [4]), and derive a Kalman filter and LQR controller for the needle.

4.1 Model Linearization and Discretization

We define the state deviation \bar{X} as the transformation between the path state \hat{X} and the (unknown) true state X , and the twist deviation \bar{U} as the difference between the true control input twist U and the control input twist \hat{U} along the path:

$$\bar{X} = \begin{bmatrix} \bar{R} & \bar{\mathbf{p}} \\ 0 & 1 \end{bmatrix} = \hat{X}^{-1}X = \begin{bmatrix} \hat{R}^T R & \hat{R}^T(\mathbf{p} - \hat{\mathbf{p}}) \\ 0 & 1 \end{bmatrix}, \quad \bar{U} = \begin{bmatrix} [\hat{\mathbf{w}}] \hat{\mathbf{v}} \\ 0 & 0 \end{bmatrix} = U - \hat{U}, \quad (12)$$

where the last equality follows from the fact that $\hat{X}^{-1} = \begin{bmatrix} \hat{R}^T & -\hat{R}^T \hat{\mathbf{p}} \\ 0 & 1 \end{bmatrix}$. The kinematics of the state deviation, i.e. its evolution over time, is given by

$$\begin{aligned} \bar{X}' &= \hat{X}^{-1}X' + (\hat{X}^{-1})'X = \hat{X}^{-1}X(U + \bar{U}) - \hat{U}\hat{X}^{-1}X = \\ &= \bar{X}(U + \bar{U}) - \hat{U}\bar{X} = \bar{X}(\hat{U} + \bar{U} + \bar{U}) - \hat{U}\bar{X} = \\ &= \begin{bmatrix} \bar{R}[\hat{\mathbf{w}} + \bar{\mathbf{w}} + \bar{\mathbf{w}}] - [\hat{\mathbf{w}}]\bar{R} & \bar{R}(\hat{\mathbf{v}} + \bar{\mathbf{v}} + \bar{\mathbf{v}}) - [\hat{\mathbf{w}}]\bar{\mathbf{p}} - \hat{\mathbf{v}} \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} \bar{R}' & \bar{\mathbf{p}}' \\ 0 & 0 \end{bmatrix}, \end{aligned} \quad (13)$$

where the equalities follow from Eqs. (9) and (12), and the fact that $(\hat{X}^{-1})' = -\hat{U}\hat{X}^{-1}$ [21].

To get a non-redundant state vector, we represent the orientation deviation \bar{R} as a rotation of angle $\|\bar{\mathbf{r}}\|$ about axis $\bar{\mathbf{r}} \in \mathbb{R}^3$. Assuming this deviation is sufficiently small, it is approximated well by the following first-order Taylor expansion:

$$\bar{R} = I + [\bar{\mathbf{r}}]. \quad (14)$$

By substituting Eq. (14) into Eq. (13), and ignoring all second-order error terms, we get to first order (using the identities $[\mathbf{a}]\mathbf{b} = \mathbf{a} \times \mathbf{b} = -[\mathbf{b}]\mathbf{a}$ and $[\mathbf{a}][\mathbf{b}] - [\mathbf{b}][\mathbf{a}] = [\mathbf{a} \times \mathbf{b}]$):

$$\bar{\mathbf{p}}' = [\bar{\mathbf{r}}]\hat{\mathbf{v}} + \bar{\mathbf{v}} + \bar{\mathbf{v}} - [\hat{\mathbf{w}}]\bar{\mathbf{p}} = -[\hat{\mathbf{w}}]\bar{\mathbf{p}} - [\hat{\mathbf{v}}]\bar{\mathbf{r}} + \bar{\mathbf{v}} + \bar{\mathbf{v}}, \quad (15)$$

$$[\bar{\mathbf{r}}]' = [\bar{\mathbf{r}}][\hat{\mathbf{w}}] + [\bar{\mathbf{w}} + \bar{\mathbf{w}}] - [\hat{\mathbf{w}}][\bar{\mathbf{r}}] = -[\hat{\mathbf{w}}]\bar{\mathbf{r}} + [\bar{\mathbf{w}} + \bar{\mathbf{w}}]. \quad (16)$$

Combining Eqs. (15) and (16), we get in matrix form:

$$\begin{bmatrix} \bar{\mathbf{p}}' \\ [\bar{\mathbf{r}}]' \end{bmatrix} = \begin{bmatrix} -[\hat{\mathbf{w}}] & -[\hat{\mathbf{v}}] \\ 0 & -[\hat{\mathbf{w}}] \end{bmatrix} \begin{bmatrix} \bar{\mathbf{p}} \\ [\bar{\mathbf{r}}] \end{bmatrix} + \begin{bmatrix} \bar{\mathbf{v}} \\ [\bar{\mathbf{w}}] \end{bmatrix} + \begin{bmatrix} \bar{\mathbf{v}} \\ [\bar{\mathbf{w}}] \end{bmatrix}, \quad (17)$$

which we can write as

$$\bar{\mathbf{x}}' = F\bar{\mathbf{x}} + G\bar{\mathbf{u}} + \mathbf{m}, \quad \mathbf{m} \sim \mathcal{N}(\mathbf{0}, M), \quad (18)$$

where $\bar{\mathbf{x}} = \begin{bmatrix} \bar{\mathbf{p}} \\ [\bar{\mathbf{r}}] \end{bmatrix}$, $F = \begin{bmatrix} -[\hat{\mathbf{w}}] & -[\hat{\mathbf{v}}] \\ 0 & -[\hat{\mathbf{w}}] \end{bmatrix}$, \mathbf{m} is as defined in Eq. (8), and $\bar{\mathbf{u}}$ and G are defined as follows:

$$\bar{\mathbf{u}} = \begin{bmatrix} v - \hat{v} \\ w - \hat{w} \\ v\kappa - \hat{v}\hat{\kappa} \end{bmatrix}, \quad G = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}^T. \quad (19)$$

Let us discretize time into stages of duration τ and assume that the control inputs v , w , κ and \hat{v} , \hat{w} , $\hat{\kappa}$ and variance M are constant for the duration of each stage k . The path Π then consists of a series of states and control input twists $(\hat{X}_0, \hat{U}_0, \dots, \hat{X}_\ell, \hat{U}_\ell)$, where ℓ is the number of stages of the path, such that

$$\hat{X}_{k+1} = \hat{X}_k \exp(\tau \hat{U}_k). \quad (20)$$

We can then integrate Eq. (18) to get [15]:

$$\bar{\mathbf{x}}_{k+1} = A_k \bar{\mathbf{x}}_k + B_k \bar{\mathbf{u}}_k + \mathbf{n}_k, \quad \mathbf{n}_k \sim \mathcal{N}(\mathbf{0}, N_k), \quad (21)$$

where

$$A_k = e^{\tau F_k}, \quad B_k = \int_0^\tau e^{(\tau-t)F_k} G dt, \quad N_k = \int_0^\tau e^{(\tau-t)F_k} M_k e^{(\tau-t)F_k^T} dt. \quad (22)$$

Eq. (21) is the linearized and discretized motion model of the deviation of the needle state from the path.

The sensor model (see Eq. (11)) is discretized by assuming that in each stage k we obtain a measurement \mathbf{z}_k . To relate the state deviation vector $\bar{\mathbf{x}}$ (as opposed to the state matrix X as in Eq. (11)) to a measurement \mathbf{z} , we define (note that $\bar{\mathbf{p}}$ and $\bar{\mathbf{r}}$ are part of $\bar{\mathbf{x}}$):

$$\bar{h}_k(\bar{\mathbf{x}}, \mathbf{q}) = h\left(\hat{X}_k \begin{bmatrix} I + [\bar{\mathbf{r}}] & \bar{\mathbf{p}} \\ 0 & 1 \end{bmatrix}, \mathbf{q}\right), \quad (23)$$

where we use the fact that $X = \hat{X}\bar{X}$ (see Eq. (12)) to reconstruct the state matrix X from the state deviation vector $\bar{\mathbf{x}}$. Linearizing \bar{h}_k around the path Π gives

$$\bar{\mathbf{z}}_k = H_k \bar{\mathbf{x}}_k + W_k \mathbf{q}_k, \quad \mathbf{q}_k \sim \mathcal{N}(\mathbf{0}, Q_k), \quad (24)$$

where

$$\bar{\mathbf{z}}_k = \mathbf{z}_k - \bar{h}_k(\mathbf{0}, \mathbf{0}), \quad H_k = \frac{\partial \bar{h}_k}{\partial \bar{\mathbf{x}}}(\mathbf{0}, \mathbf{0}), \quad W_k = \frac{\partial \bar{h}_k}{\partial \mathbf{q}}(\mathbf{0}, \mathbf{0}). \quad (25)$$

Eq. (24) is the linearized and discretized sensor model of the deviation of the needle state from the path.

4.2 Kalman Filter and LQR Controller

Eqs. (21) and (24) form a standard linear Gaussian model, of which $\bar{\mathbf{x}}$ is the state, $\bar{\mathbf{u}}$ the control input and $\bar{\mathbf{z}}$ the measurement. The Kalman filter keeps track of the estimate $\hat{\mathbf{x}}$ and variance P of the true state $\bar{\mathbf{x}}$ during control. It continually performs two steps; a process update to propagate the applied control input $\bar{\mathbf{u}}$, and a measurement update to incorporate the obtained measurement $\bar{\mathbf{z}}$:

Process update step:

$$\hat{\mathbf{x}}_k^- = A_{k-1} \hat{\mathbf{x}}_{k-1}, \quad (26)$$

$$P_k^- = A_{k-1} P_{k-1} A_{k-1}^T + N_{k-1}. \quad (27)$$

Measurement update step:

$$K_k = P_k^- H_k^T (H_k P_k^- H_k^T + W_k Q_k W_k^T)^{-1}, \quad (28)$$

$$\hat{\mathbf{x}}_k = K_k (\bar{\mathbf{z}}_k - H_k \hat{\mathbf{x}}_k^-), \quad (29)$$

$$P_k = (I - K_k H_k) P_k^-. \quad (30)$$

The Kalman-gain matrices K_k can be computed in advance (i.e. before execution) given the initial variance P_0 , without knowledge of the actual control inputs $\bar{\mathbf{u}}$ and measurements $\bar{\mathbf{z}}$. We refer the reader to [28] for additional details.

The LQR controller provides optimal control for a motion model of the type given by Eq. (21) [5]. The optimal control inputs $\bar{\mathbf{u}}_k$ are found by minimizing the cost function

$$\min_{\bar{\mathbf{u}}} \left(\sum_{k=0}^{\ell} (\bar{\mathbf{x}}_k^T C \bar{\mathbf{x}}_k + \bar{\mathbf{u}}_k^T D \bar{\mathbf{u}}_k) \right), \quad (31)$$

which quadratically penalizes deviations from the path Π through given positive-definite weight matrices C and D . Matrix C specifies the cost for deviating from the planned path, while D specifies the cost for deviating from the planned control input. Penalizing the magnitude of $\bar{\mathbf{u}}$ is reasonable as the linearized motion model is only valid when $\bar{\mathbf{u}}$ is small.

Solving Eq. (31) gives the control policy $\bar{\mathbf{u}}_k = L_k \bar{\mathbf{x}}_k$, for feedback matrices L_k that are pre-computed using a standard recursive procedure (for $0 \leq k < \ell$) [5]. As the true state $\bar{\mathbf{x}}_k$ is unknown, the estimate $\hat{\mathbf{x}}_k$ of the state which is obtained from the Kalman filter is used to determine the control input $\bar{\mathbf{u}}_k$ at each stage k during execution of the path. Hence, the control policy is $\bar{\mathbf{u}}_k = L_k \hat{\mathbf{x}}_k$. We refer the reader to [5, 25] for additional details.

The actual control inputs v_k , w_k and κ_k that are applied to the needle are found using Eq. (19), given $\bar{\mathbf{u}}_k$ and the control inputs \hat{v}_k , \hat{w}_k and $\hat{\kappa}_k$ of path Π . After application of the control input, the Kalman filter produces the estimate of the next state from which in turn a new control input is determined. This cycle repeats until the execution of the path is complete.

5 Optimal Path and Sensor Planning

The first objective is to plan a needle path towards a target location \mathbf{g} inside a workspace $\mathcal{W} \subset \mathbb{R}^3$. For simplicity, we assume that the workspace is the rectangular region $[0, x_{\max}] \times [0, y_{\max}] \times [0, z_{\max}]$, and that the needle can enter the workspace at any point in the plane $z = 0$. Further, the workspace may contain obstacles defined by a region $O \subset \mathcal{W}$ that models impenetrable or sensitive tissue. The second objective is to select a placement of the sensor that will provide feedback during control.

The quality measure of a path Π and a sensor placement is the probability that the needle will intersect obstacles when the path is executed using LQG control. We will first discuss how to compute this probability for a given path Π and a

given sensor model, and then discuss how we plan a set of candidate paths and sensor placements from which we select the pair that minimizes the probability of intersecting obstacles during control.

5.1 Obstacle Intersection Probability Along a Path

Given the LQG controller for a path Π and sensor model h (see Section 4), we can analyze in advance how the true state $\bar{\mathbf{x}}_t$ and the estimated state $\hat{\mathbf{x}}_t$ will evolve during control as functions of each other. The evolution of the true state $\bar{\mathbf{x}}_t$ is dependent on the estimated state through the LQR control policy and the evolution of the estimated state $\hat{\mathbf{x}}_t$ is dependent on the true state through the measurement obtained in the Kalman filter. This gives the following equation in matrix form (see [25] for more details):

$$\begin{bmatrix} \bar{\mathbf{x}}_{k+1} \\ \hat{\mathbf{x}}_{k+1} \end{bmatrix} = \begin{bmatrix} A_k & B_k L_k \\ K_{k+1} H_{k+1} A_k & A_k + B_k L_k - K_{k+1} H_{k+1} A_k \end{bmatrix} \begin{bmatrix} \bar{\mathbf{x}}_k \\ \hat{\mathbf{x}}_k \end{bmatrix} + \begin{bmatrix} I & 0 \\ K_{k+1} H_{k+1} & K_{k+1} W_{k+1} \end{bmatrix} \begin{bmatrix} \mathbf{n}_k \\ \mathbf{q}_{k+1} \end{bmatrix}, \quad \begin{bmatrix} \mathbf{n}_k \\ \mathbf{q}_{k+1} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \begin{bmatrix} N_k & 0 \\ 0 & Q_{k+1} \end{bmatrix}), \quad (32)$$

which we write shorthand—for the appropriate definitions of \mathbf{y}_k , Y_k , V_k , \mathbf{s}_k and Z_k —as (note that Y_k , V_k and Z_k can all be computed in advance for a given a path Π):

$$\mathbf{y}_{k+1} = Y_k \mathbf{y}_k + V_k \mathbf{s}_k, \quad \mathbf{s}_k \sim \mathcal{N}(\mathbf{0}, Z_k). \quad (33)$$

From this, we can compute the mean $\hat{\mathbf{y}}_k$ and the variance Σ_k of $\mathbf{y}_k = \begin{bmatrix} \bar{\mathbf{x}}_k \\ \hat{\mathbf{x}}_k \end{bmatrix}$ for any stage k of the execution of the path:

$$\hat{\mathbf{y}}_{k+1} = Y_k \hat{\mathbf{y}}_k, \quad \hat{\mathbf{y}}_0 = \mathbf{0}, \quad (34)$$

$$\Sigma_{k+1} = Y_k \Sigma_k Y_k^T + V_k Z_k V_k^T, \quad \Sigma_0 = \begin{bmatrix} P_0 & 0 \\ 0 & 0 \end{bmatrix}. \quad (35)$$

Note that the mean $\hat{\mathbf{y}}_k$ is zero for all stages t . Hence, $\begin{bmatrix} \bar{\mathbf{x}}_k \\ \hat{\mathbf{x}}_k \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \Sigma_k)$.

Given these a priori distributions of the state deviation, we can compute the probability that the needle will intersect an obstacle during the execution of path Π . Let $\Sigma_k^{\mathbf{p}}$ be the variance of the position deviation $\bar{\mathbf{p}}_k$, which is the upper-left 3×3 -submatrix of the 12×12 -matrix Σ_k (note that $\bar{\mathbf{x}} = \begin{bmatrix} \bar{\mathbf{p}} \\ \bar{\mathbf{r}} \end{bmatrix}$). As $\bar{\mathbf{p}} = \hat{R}^T (\mathbf{p} - \hat{\mathbf{p}})$ (see Eq. (12)), we have that $\mathbf{p} = \hat{\mathbf{p}} + \hat{R} \bar{\mathbf{p}}$, so the a priori distribution of the position \mathbf{p}_k of the needle tip at stage k along path Π is given by $\mathcal{N}(\hat{\mathbf{p}}_k, \hat{R}_k \Sigma_k^{\mathbf{p}} \hat{R}_k^T)$. Hence, the probability p_k that the needle intersects the obstacle region at stage k along path Π is computed as:

$$p_k = \int_O \frac{\exp(-\frac{1}{2}(\mathbf{b} - \hat{\mathbf{p}}_k)^T (\hat{R}_k P_k^{\mathbf{p}} \hat{R}_k^T)^{-1} (\mathbf{b} - \hat{\mathbf{p}}_k))}{\sqrt{8\pi^3 \det(\hat{R}_k P_k^{\mathbf{p}} \hat{R}_k^T)}} d\mathbf{b}, \quad (36)$$

which is the integral over the obstacle region O of the probability density function of Gaussian $\mathcal{N}(\hat{\mathbf{p}}_k, \hat{R}_k P_k^{\mathbf{p}} \hat{R}_k^T)$.

Instead of computing the probabilities p_k exactly, we will use an approximation that can be computed efficiently. To this end, we look at the maximum factor by which the ellipsoid of one standard deviation of the a priori probability distribution can be scaled such that it does not intersect an obstacle. Let this factor be c_k for stage k along the path. For a multivariate Gaussian distribution of dimension n , the probability that a sample is outside c_k standard deviations is given by

$$\hat{p}_k = 1 - \Gamma(n/2, c_k^2/2), \quad (37)$$

where Γ is the regularized Gamma function [17]. It provides a (conservative) upper bound of the true probability of intersecting obstacles at stage k .

The value of c_k for stage k is efficiently computed by using a collision-checker capable of performing distance calculations and linear transformations on the obstacle geometry, such as SOLID [26]. Let E_k be a matrix such that $E_k E_k^T = \Sigma_k^p$. The set of positions within one standard deviation is then an ellipsoid centered at the mean $\hat{\mathbf{p}}_k$ obtained by transforming a unit sphere by E_k . Transforming the environment by E_k^{-1} (such that the uncertainty ellipsoid becomes a unit sphere), and calculating the Euclidean distance between $\hat{\mathbf{p}}_k$ and the nearest obstacle in the transformed environment gives the value of c_k for stage k , from which the approximate probability \hat{p}_k of intersecting obstacles at stage k can be computed using Eq. (37).

Assuming (somewhat opportunistically) that the probabilities \hat{p}_k are independent, it follows that the probability $p(\Pi)$ that the needle intersects the obstacle region anywhere along path Π is given by $p(\Pi) = 1 - \prod_{k=0}^{\ell} (1 - \hat{p}_k)$.

5.2 Planning a Needle Path and Sensor Placement

To plan an optimal pair of a needle path and sensor placement, we (randomly) generate large sets of possible needle paths and sensor placements, and evaluate the probability of intersecting obstacles for each pair.

To generate a large set of (random) needle paths we use the RRT (random rapidly-exploring tree) algorithm [14], as it is known to create trees of paths that uniformly cover the space and handle kinematically constrained systems, such as steerable needles, well. As the target location is a specific point \mathbf{g} , and the entry location can be anywhere in a pre-defined entry zone, it is convenient to plan backwards from the target location. Backward kinematics are identical to forward kinematics, except that the (forward) control inputs \hat{v} , \hat{w} , and $\hat{\kappa}$ are integrated over a negative time-step. As the actual path traced out by the needle only depends on the ratio w/v and not on the values of the individual terms [8], we set $\hat{v} = 1$ cm/s and only vary $\hat{w} \in [-w_{\max}, w_{\max}]$ and $\hat{\kappa} \in [0, \kappa_0]$. We will not describe the RRT-planner in detail here, but refer to [20, 29, 30] for details on RRT implementations for steerable needles. All the paths in the resulting tree that reach the entry zone are valid.

Let \mathcal{S} be the space of sensor placements, and let the sensor model $h(X, \mathbf{q})$ corresponding to a placement $s \in \mathcal{S}$ be denoted $h_s(X, \mathbf{q})$. We generate a large set of placements by random sampling from \mathcal{S} . We iterate over all valid paths contained in the RRT-tree and all placements sampled from \mathcal{S} , and select the pair for which the needle has minimal probability of intersecting obstacles as computed above.




Fig. 3 (a) Given a candidate path (gray) from left to right, we illustrate samples (red spheres) from 100 simulated executions of the LQG controller and extended Kalman filter showing convergence to the path. The simulations included artificially generated process and measurement noise. The blue ellipsoids show the a priori distributions computed by the planner along the path. (b) A simulated example trajectory (shown in red) using LQG control.

6 Simulation Results

We experimentally evaluate our approach using simulations of the computed candidate paths using LQG control with simulated process and measurement noise. In our experiments, we use an anatomical model of the human male pelvic region to simulate needle insertion in tissue for delivering radioactive doses to targets within the prostate region for cancer treatments. We first show that the needle is controllable along a candidate path. We then show the effect the sensor can have on the optimal path, and how the sensor placement can be optimized for a given problem.

We implemented the system in C++ and tested it on a 3.33 Ghz 4-core Intel[®] i7[™] PC. All experiments utilized only a single core for computation, but our approach could be parallelized over multiple cores to yield significant speedups. In our experiments, we model the needle motion and noise using the following parameters: $w_{\max} = 2\pi$ rad/s, $\tau = 0.1$ s, $\kappa = 0.2$ cm⁻¹, M is a diagonal matrix with 0.01 (cm/s)² for the position components and 0.05 (rad/s)² for the rotational components, and Q is a diagonal matrix with 0.05 cm² along the diagonal.

6.1 Needle Controllability

We first demonstrate the controllability of the needle along a candidate path using an LQG controller with artificially generated process and measurement noise. We assume that the sensor can only measure the position \mathbf{p} of the needle tip and not the orientation. Fig. 3a shows the samples obtained from 100 simulations of executions of the path using the LQG controller, demonstrating that the a priori probability distributions computed by the planner are close approximations of the true distribution of the states along the path and that the needle follows the candidate path closely. For an insertion length of 11.25 cm, the standard deviation of the distance to the target was found to be 0.17 cm.

To emulate uncertainties arising due to tissue heterogeneities, we applied spatially varying process noise sampled from Gaussian distributions with zero mean and variances up to three times the variance M . Following the same candidate path of length 11.25 cm, the standard deviation of the distance of the final needle-tip position to the target across 1000 simulation runs was 0.24 cm.




Fig. 4 Optimal needle paths in an anatomical environment modeling the human prostate and surrounding tissues. Two examples of a sensor placement, in which (left) only the x - and y -coordinate and (right) only the x - and z -coordinate of the needle-tip are measured by an imaging device by projecting the anatomy on the imaging plane (shown in the two views). The optimal path predominantly lies in the imaging plane to minimize uncertainty in the viewing direction.

6.2 Effect of Sensor Placement on Optimal Path

We demonstrate the use of LQG to select plans that minimize the probability of failure and also show the effect of sensor placement on the optimal path. We performed these experiments using an anatomical model of the human pelvic region (see Fig. 4). As an example sensor model, we consider a 2-D image of the anatomy (for example, an x-ray or 2-D ultrasound image). The image is projected along the z -axis using an imaging device, from which we can only measure the x - and y -coordinate of the position $\mathbf{p} = [x, y, z]$ of the needle tip. This gives the following observation function h (note that \mathbf{p} is part of X):

$$h(X, \mathbf{q}) = \begin{bmatrix} x \\ y \end{bmatrix} + \mathbf{q}, \quad \mathbf{q} \sim \mathcal{N}(\mathbf{0}, Q) \quad (38)$$

Fig. 4 (left) shows the optimal needle path within the RRT-tree for this sensor. The optimal path predominantly lies in the x - y plane to minimize the uncertainty along the viewing direction (z -axis). It took 41 seconds to generate a set of 1000 candidate paths and selecting an optimal path from this set required 4.6 seconds. Similarly, if we instead place the sensor such that it obtains 2-D images projected along the y -axis, the optimal path (shown in Fig. 4 (right)) is predominantly in the x - z plane.

To quantify the effect the sensor location has on the probability of success of a path, we compare the results of two paths assuming sensing along the z axis: the optimal path and a path that lies predominantly in the x - z plane. For each path, we estimated the “ground truth” probability of success by performing 1000 simulated executions using the LQG controller with artificially generated process and measurement noise. The path that is optimal according to our method, which lies predominantly in the x - y plane, has a probability of success of 89%. In contrast, the path that primarily lies in the x - z plane had a probability of success of only 44%,




Fig. 5 Using an x-ray imager mounted on a rotating C-arm, it is possible to rotate the sensor about the horizontal axis along which the patient is positioned (left). The anatomy as viewed from the computed optimal sensor placement (right). The optimal path predominantly lies in the imaging plane to minimize uncertainty in the viewing direction.

which is to be expected since there is greater uncertainty in needle pose along the viewing direction (z -axis).

6.3 Optimizing the Sensor Placement

In many clinical procedures involving bevel-tip steerable needles, the physician can select where to place an intra-operative sensing device for real-time feedback. We consider the case of an x-ray imaging sensor attached to a C-arm, a commonly used setup in operating rooms that allows the physician to rotate the sensor in a circle about the patient as shown in Fig. 5. The placement of the imaging device can be parameterized by angle θ relative to the horizontal axis of the patient. A 2-D image of the anatomy can be obtained by parallel projection along the viewing direction (along the radius of the arm). This gives the following observation function h :

$$h(X, \mathbf{q}) = \begin{bmatrix} x \\ y \cos \theta + z \sin \theta \end{bmatrix} + \mathbf{q}, \quad \mathbf{q} \sim \mathcal{N}(\mathbf{0}, \mathcal{Q}) \quad (39)$$

To optimize sensor placement, we iterate over all paths computed by the RRT-based planner and all sampled sensor placements (obtained by varying the angle of rotation about the horizontal axis θ in regular increments), and select the pair for which the needle has minimal probability of intersecting obstacles. Fig. 5 shows the optimal sensor placement for a given set of candidate paths and a 2-D view of the anatomy as visible from the imaging device. For a set of 1000 candidate paths and 36 possible sensor placements (obtained by discretizing over the interval $[0, \pi]$ in intervals of 5 degrees), our implementation took 185 seconds to compute the optimal pair over the set of candidate paths and sensor placements. It should be noted that with modern multi-core processors, this computation can be trivially parallelized to bring down the computation time within clinically acceptable limits.

7 Conclusion and Future Work

In this paper, we presented a technique for planning and controlling flexible steerable needles towards a target location in 3-D anatomy with obstacles such as bones and sensitive organs. Our approach minimizes the probability that the needle intersects obstacles by explicitly taking into account both needle motion uncertainty and the sensors used to obtain (noisy, partial) feedback on the needle state. We demonstrated how the sensor influences the optimal path and optimize over the set of candidate paths and feasible sensor placements to select the pair for which the estimated uncertainty is least likely to cause intersections with obstacles.

In our current implementation, the LQR controller does not bound the control inputs within permissible limits during feedback. This can be a problem when the control for needle curvature exceeds the attainable curvature of the needle. We plan to address this issue in future work. We also plan to incorporate uncertainty introduced by tissue deformation using a physically accurate FEM simulator [6].

We also plan to evaluate our new planning and control approach using a robotic device that steers a needle in phantom tissue as in prior 2-D experiments [22]. We will utilize the ability of this planner and controller to account for unexpected events that might occur due to tissue inhomogeneity, tissue deformation, and estimation errors in motion model parameters, as well as to optimize sensor placement to minimize the probability of intersecting an obstacle.

Acknowledgments

This research was supported in part by the National Science Foundation (NSF) under grants IIS-0905344 and IIS-0904672 and by the National Institute of Health (NIH) under grants R01-EB006435 and R21-EB011628. The authors also thank members of the needle steering community for their valuable input.

References

1. R. Alterovitz, K. Goldberg, A. Okamura. Planning for steerable bevel-tip needle insertion through 2D soft tissue with obstacles. *IEEE Int. Conf. on Robotics and Automation*, 2005.
2. R. Alterovitz, M. Branicky, K. Goldberg. Motion planning under uncertainty for image-guided medical needle steering. *Int. J. Robotics Research*, 27(11-12): 1361-1374, 2008.
3. R. Alterovitz, T. Siméon, K. Goldberg. The stochastic motion roadmap: a sampling framework for planning with Markov motion uncertainty. *Robotics: Science and Systems*, 2007.
4. G. Baldwin, R. Mahony, J. Trumph, T. Hamel, T. Chervon. Complementary filter design on the Special Euclidean group $SE(3)$. *European Control Conference*, 2007.
5. D. Bertsekas. *Dynamic programming and optimal control*. Athena Scientific, 2001.
6. N. Chentanez, R. Alterovitz, D. Ritchie, L. Cho, K. Hauser, K. Goldberg, J. Shewchuk, J. F. O'Brien. Interactive Simulation of Surgical Needle Insertion and Steering. *ACM Trans. on Graphics* 28(3):88.1-88.10, 2009.
7. N. J. Cowan, K. Goldberg, G. S. Chirikjian, G. Fichtinger, R. Alterovitz, K. B. Reed, V. Kallem, W. Park, S. Misra, A. M. Okamura. Robotic Needle Steering: Design, Modeling, Planning, and Image Guidance. *Surgical Robotics - Systems, Applications, and Visions*, Springer, 2010.

8. V. Duindam, R. Alterovitz, S. Sastry, K. Goldberg. Screw-based motion planning for bevel-tip flexible needles in 3D environments with obstacles. *IEEE Int. Conf. on Robotics and Automation*, 2008.
9. V. Duindam, J. Xu, R. Alterovitz, S. Sastry, K. Goldberg. Three-dimensional motion planning algorithms for steerable needles using inverse kinematics. *Int. J. Robotics Research* 29(7):789–800, 2010.
10. J. Hall, N. Knoebel, T. McLain. Quaternion attitude estimation for miniature air vehicles using a multiplicative extended Kalman filter. *IEEE/ION Position, Location and Navigation Symp.*, 2008.
11. K. Hauser, R. Alterovitz, N. Chentanez, A. Okamura, K. Goldberg. Feedback control for steering needles through 3D deformable tissue using helical paths. *Robotics: Science and Systems*, 2009.
12. V. Kallem, N. Cowan. Image-guided control of flexible bevel-tip needles. *IEEE Int. Conf. on Robotics and Automation*, 2007.
13. V. Kallem. *Vision-based control on lie groups with application to needle steering*. PhD Thesis, Johns Hopkins University, 2008.
14. S. LaValle, J. Kuffner. Randomized kinodynamic planning. *Int. J. Robotics Research* 20(5):378–400, 2001.
15. E. Lefferts, F. Markley, M. Shuster. Kalman filtering for spacecraft attitude estimation. *Journal of Guidance, Control and Dynamics* 5(5):417–429, 1982.
16. D. Minhas, J. Engh, M. Fenske, C. Riviere. Modeling of needle steering via duty-cycled spinning. *Int. Conf. IEEE Engineering in Medicine and Biology Society*, 2007.
17. F. W. Olver, D. W. Lozier, R. F. Boisvert, C. W. Clark. *NIST Handbook of Mathematical Functions*. Cambridge University Press, 2010.
18. W. Park, J. Kim, Y. Zhou, N. Cowan, A. Okamura, G. Chirikjian. Diffusion-based motion planning for a nonholonomic flexible needle model. *IEEE Int. Conf. on Robotics and Automation*, 2005.
19. W. Park, Y. Wang, G. Chirikjian. Path planning for flexible needles using second order error propagation. *Workshop on Algorithmic Foundations of Robotics*, 2008.
20. S. Patil, R. Alterovitz. Interactive motion planning for steerable needles in 3D environments with obstacles. *IEEE Int. Conf. on Biomedical Robotics and Biomechanics*, 2010.
21. K. Petersen, M. Pedersen. *The Matrix Cookbook*. Tech. Univ. of Denmark, 2008.
22. K. Reed, V. Kallem, R. Alterovitz, K. Goldberg, A. Okamura, N. Cowan. Integrated planning and image-guided control for planar needle steering. *IEEE/RAS-EMBS Int. Conf. Biomedical Robotics and Biomechanics*, 2008.
23. K. Reed, A. Okamura, N. Cowan. Controlling robotically steered needle in the presence of torsional friction. *IEEE Int. Conf. on Robotics and Automation*, 2009.
24. N. Trawny, S. Roumeliotis. Indirect Kalman filter for 3D attitude estimation; a tutorial for quaternion algebra. *Tech. report TR-2005-002*, University of Minnesota, 2005.
25. J. van den Berg, P. Abbeel, K. Goldberg. LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. *Robotics: Science and Systems*, 2010.
26. G. van den Bergen. *Collision detection in interactive 3D environments*. Morgan Kaufmann Publishers, 2004.
27. R. Webster, J. Kim, N. Cowan, G. Chirikjian, A. Okamura. Nonholonomic modeling of needle steering. *Int. J. Robotics Research* 25(5-6):509–525, 2006.
28. G. Welch, G. Bishop. An introduction to the Kalman filter. *Tech. Report TR 95-041*, University of North Carolina at Chapel Hill, 2006.
29. J. Xu, V. Duindam, R. Alterovitz, K. Goldberg. Motion planning for steerable needles in 3D environments with obstacles using rapidly-exploring random trees and backchaining. *IEEE Int. Conf. on Automation Science and Engineering*, 2008.
30. J. Xu, V. Duindam, R. Alterovitz, J. Pouliot, A. Cunha, I. Hsu, K. Goldberg. Planning fire-works trajectories for steerable medical needles to reduce patient trauma. *IEEE Int. Conf. on Intelligent Robots and Systems*, 2009.