

# Path Planning for Insertion of a Bevel-Tip Needle

Michael Lin  
Mechanical Engineering  
Stanford University  
mlinyang@stanford.edu

Xingchen Fan  
Mechanical Engineering  
Stanford University  
xfan@stanford.edu

Alex Gruebele  
Mechanical Engineering  
Stanford University  
agruebe2@stanford.edu

**Abstract**—Needle insertion is a common task in medical practice. Bevel-tip needles allow a physician to steer the needle to a target, however the actual path of the needle can be difficult to predict. In this paper, we demonstrate multiple methods to generate optimal policies for bevel-tip needle to reach a target while avoiding obstacles and minimizing insertion length. With a coarsely discretized tissue environment, we observe that the generated policy is robust but solving for it is computationally expensive, preventing us from further refining the discretization for potential performance improvement. Policies generated from a continuous POMDP are computationally less expensive when solved with Partially Observable Monte Carlo Planning (POMCP). Our results show that the continuous formulation of the problem yields policies that achieve a higher success rate in complicated problems. Also, it is able to achieve solutions with shorter paths as well as a wider variety of solutions compared to the discrete case.

## I. INTRODUCTION

Needle insertion is a critical aspect of many medical procedures, such as tissue sample acquisition, drug delivery, thermal ablation, and others. In all of these procedures the objective is to place the needle tip at a target location while minimizing damage to the surrounding tissues. Traditional needles are constrained to a straight insertion path due to their symmetric tip, however, recent innovations in needle design have enabled higher control over path through the use of bevel or kinked needle tips (Figure 1) [1]. We aim to predict the trajectory of the needle based on tip location and orientation, and provides advice for rotation actions based on some optimal policy, so the physician is able to make more intelligent rotational adjustments about the needle axis to reach the target while avoiding obstacles.

### A. Prior Work

Alterovitz et al. formulated a similar problem for two-dimensional needle steering using a discretized grid to simulate tissue, assuming full observability. They use previously obtained MRI images to assign transition probabilities for needle insertion model. Needle motion was simulated as a "Dubins Car" MDP problem [2].

In another approach by Patil et al., the path planning of needle insertion is done online to take advantage of real-time feedback from magnetic resonance imaging (MRI) or ultrasonic imaging. Physicians can update the plan based on where the needle has gone, as well as any unexpected physiological changes. They use a rapidly-exploring random tree (RRT) algorithm to speed up computation [3].

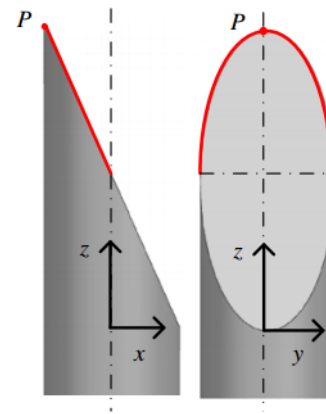


Fig. 1. Bevel-tip needle (picture taken from [1])

Patil, van den Berg and Alterovitz have approached the problem through modeling the deformable tissue environment through finite element method (FEM). In this work RRT was used to generate multiple path candidates for the needle to take. Then LQG control was used to select the path with the highest probability of success. [4].

Our work will tackle the same problem of path planning for needle insertion, but we will formulate the problem in two ways. One formulation will use a discrete environment similar to how was done in [2] and this will be compared to a continuous formulation of the problem. We will also include tip angle sensing in our POMDP problem formulation, which in practice can be achieved through shape sensing methods as done in [5].

## II. DISCRETE POMDP

### A. State Space ( $\mathcal{S}$ )

The 2D tissue is discretized into a  $10 \times 10$  grid. The nodes use the standard Cartesian coordinates. At any node, the needle tip can take 8 orientations, which are angles in  $45^\circ$  buckets. The first orientation bucket is aligned with the X-axis and it increments in counterclockwise direction. Thus, there are a total of 800 states. The target and obstacles are located at desired nodes. The start state is always at the boundary with the tip facing towards the tissue. The terminal states include the target node with any tip orientation, obstacle states with any tip orientation and any boundary nodes with the tip facing outward.

### B. Belief Space ( $\mathcal{B}$ )

The size of belief space is the same as that of state space, however because of our chosen observation model, only 27 elements in the belief vector are nonzero. They correspond to the states shown in Figure 3.

### C. Actions ( $\mathcal{A}$ )

The policy is presented to the physician during insertion in the form of which rotation is necessary to best proceed to the target. In our 2D simulation, there are only two orientations considered, so the actions available are clockwise (CW) and counterclockwise (CCW). They indicate, in the 2D field, the directions that the needle orientation tends to drive the tip at the next time step. By keeping the actions as an absolute orientation in global reference frame (rather than relative ones like "rotate" and "do not rotate"), we do not have to keep track of the previous orientation as part of the state. The rotation instructions for physicians can be outputted by separately processing the policy itself.

### D. Rewards ( $R$ )

Several types of rewards will be collected during needle insertion. Walls and obstacles have large negative rewards, while the target has a large positive reward. The time taken to reach the target is also penalized by small negative reward at every time step. Only relative magnitudes of the rewards matter for the problem.

### E. Transition Model ( $T$ )

The transition model assigns probabilities to three neighboring states based on beam bending mechanics. The needle can either move to the node in the front, facing forward or 45° sideways in the direction to which the tip force causes it to bend, or move to a diagonal node facing 45° sideways. These are shown in Figure 2.

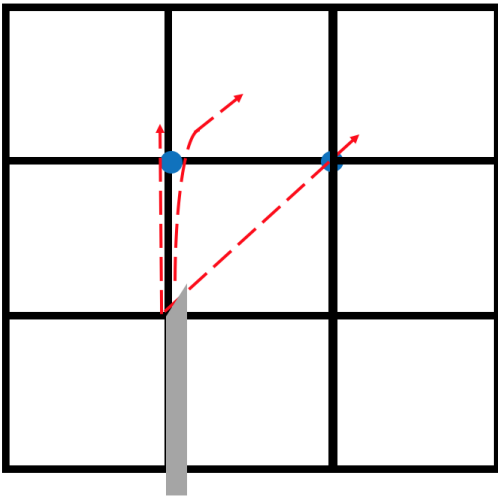


Fig. 2. The three states with non-zero probabilities in the discretized POMDP transition model

### F. Observation Model ( $O$ )

Uncertainty in this problem arises from the fact that sensor readings exhibit noise and artifacts. For the real problem, we would use ultrasonic imaging to locate the needle tip and needle shape sensing to determine the tip orientation. We use an observation model that assigns probabilities to 27 combinations of location and orientation at neighboring nodes to the actual state as shown in Figure 3.

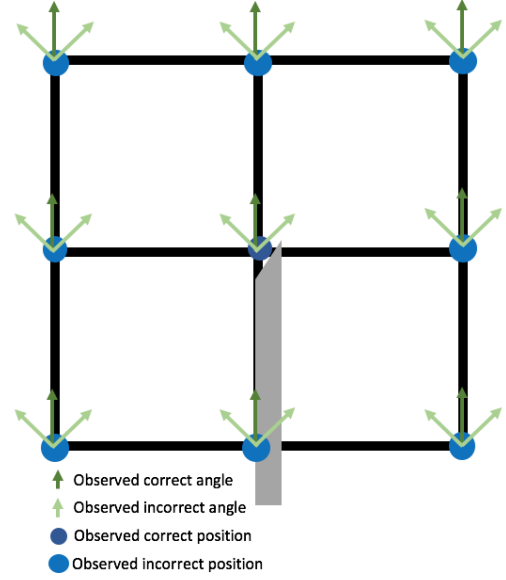


Fig. 3. The possible states that can be observed with the combination of ultrasound and needle-shape sensing

### G. Solver: SARSOP

We used the Julia SARSOP.jl implementation of Successive Approximations of the Reachable Space under Optimal Policies (SARSOP) to solve the discrete POMDP problem. This anytime solver samples only optimal points within the reachable belief state. It only calculates backup beliefs on states that are visited by the optimal policy, thus cutting computational costs associated with the larger set of calculations. It does this by performing value iteration on sampled points, followed by pruning away trees that are not explored by the optimal policy [6] [7].

## III. CONTINUOUS POMDP

### A. State Space ( $\mathcal{S}$ )

The state space is continuous but conceptually it is similar to the discrete case. The tissue environment is set to be a  $100 \times 100$  mm bounding box. The target and obstacles are patches centered at certain nodes. Target nodes, obstacle nodes and all boundary nodes with any tip orientation are terminal states.

### B. Belief Space ( $\mathcal{B}$ )

Size is same as of states. Gaussian. Finite non-zero value everywhere.

### C. Actions ( $\mathcal{A}$ )

The actions will be the same as the ones used in discrete POMDP: counterclockwise and clockwise.

### D. Rewards ( $R$ )

The rewards will be similar to the ones used in discrete POMDP: the area centered at the target has positive reward while the areas around the obstacles and at the boundary have negative rewards. These areas are defined as a  $10mm \times 10mm$  square centered at the specified position. We also penalized the time taken to reach the target by giving a small negative reward at every state with the goal of minimizing the length of the path (which minimizes the damage to surrounding tissues).

### E. Transition Model ( $T$ )

The transition model is based on beam theory. The schematic is shown in Figure 4. The spatial and angular deflections of the needle tip are calculated at each time step using Equation 1 and 2:

$$u = \frac{PL^3}{3EI} \quad (1)$$

$$\theta = \frac{PL^2}{2EI} \quad (2)$$

The tip load  $P$  is modeled as a force with a mean of 40 N and a standard deviation of 5 N to account for tissue stiffness uncertainty. The needle is inserted by 2 mm for every time step and has a circular cross-section with an outer diameter of 0.9 mm and inner diameter of 0.6 mm. The Young's modulus of needle material is 3 GPa. A rotation matrix is used to keep track of the needle tip orientation in global reference frame.




Fig. 4. The beam bending model used to for the continuous POMDP transition model

### F. Observation Model ( $O$ )

Similar to the discrete case, uncertainty in this problem arises from the fact that sensor readings exhibit noise and artifacts. We simulated the tip location measurement through ultrasound and tip orientation measurement through shape sensing by taking the actual current state and using it as the mean of a Gaussian distribution to sample from. These observations were obtained from sampling a multivariate Gaussian

distribution as shown in Equation 3. In our case the distribution was a three dimensional Gaussian, where the mean values are  $\mu = [x_{actual}, y_{actual}, \theta_{actual}]$  and the covariance matrix is  $diagm([2.0mm, 2.0mm, 0.5rad])$ .

$$f(x|\mu, \Sigma) = \frac{1}{(2\pi)^{\frac{d}{2}} |\Sigma|^{1/2}} e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)} \quad (3)$$

### G. Solver: POMCP

We used the Julia POMCP.jl implementation of Partially Observable Monte Carlo Planning (POMCP) to solve the continuous problem online. In POMCP, each simulation starts from a state sampled from the belief state using a particle filter, and chooses the action based on the optimal value at the current step and some adjustable exploration bonus. Then it performs the selected action through Monte-Carlo simulation. It recursively evaluates the next state if it is already in the search tree, or inserts it into the tree and executes a rollout policy through simulation. It also updates the statistics of tree nodes by back-propagating the results up to the root. POMCP repeats this process and reports the best action based on the returned values of the tree search [8].

## IV. RESULTS

We set up a simple layout with one obstacle and one target for our first experiment. For the discrete POMDP case the obstacle was placed at  $[40mm, 60mm]$  and the target at  $[80mm, 40mm]$ . The rewards were set to 100 for reaching the target, -20 for running into obstacles and -1 for every step taken. We used a discount factor of 0.9. We set the insertion point of needle at  $[40mm, 100mm, 3\pi/2rad]$  with the tip facing downward. Figure 5 shows a sample simulation of our policy where the needle successfully reached the target. Big red dots show the obstacle positions in the grid and green dots show the target. Along the path of the needle we recorded the action taken at every step. Yellow dots correspond to a clockwise action and blue dots correspond to a counterclockwise action. We ran simulation with our policy 20 times which yielded a success rate of 70%. For the continuous POMDP case the obstacles and target were placed at the same locations as in the discrete case with the same rewards. In this case we used a discount factor of 0.95. Figure 6 shows a sample simulation of our policy with an insertion point at  $[40mm, 100mm, 3\pi/2rad]$ . We ran simulation with our policy 20 times which yielded a success rate of 70%.

To make the problem more challenging we added more obstacles around the target. For both the discrete and continuous POMDP, we placed obstacles at nodes  $[40mm, 60mm]$ ,  $[60mm, 60mm]$  and  $[70mm, 40mm]$ , and the target at node  $[50mm, 40mm]$ . In the discrete problem we set the discount factor to 0.92. Figure 7 shows a sample simulation of our policy with an insertion point at  $[30mm, 100mm, 3\pi/2rad]$ . We ran simulation with our policy 20 times which yielded a success rate of 70%. In the continuous problem we set the discount factor to 0.9. Figure 8 shows a sample simulation of our policy with an insertion point at  $[30mm, 100mm, 3\pi/2rad]$ .




Fig. 5. Discount factor = 0.9, target reward = 100, obstacle reward = -20

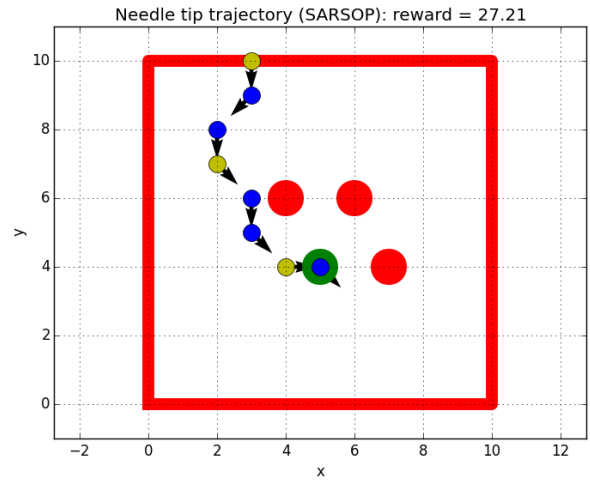


Fig. 7. Discount factor = 0.92, target reward = 100, obstacle reward = -20, boundary reward = -100




Fig. 6. Discount factor = 0.95, target reward = 100, obstacle reward = -20, boundary reward = -100




Fig. 8. Discount factor = 0.9, target reward = 100, obstacle reward = -20, boundary reward = -100

We ran simulation with our policy 20 times which yielded a success rate of 90%.

We performed another experiment where we increased the discount factor for both the continuous and discrete cases and ran their solvers with the more challenging layout. The policy returned by SARSOP did not vary much, however, POMCP returned a significantly different policy. Figure 9 and Figure 10 show a successful and an unsuccessful simulation with this policy respectively.

## V. DISCUSSION

One interesting result that we observed in the experiments is shown in Figure 5 and 6. In both discrete and continuous POMDP cases, the solvers learned to drive the needle along straight lines. By continuously switching the orientation, the needle is able to go straight to a target. This is a phenomenon

that was intentionally designed into a duty-cycle driven steerable needle designed by Majewicz et al. [9]. From these two figures we can also see an advantage of the continuous model. When attempting to take a direct route, the solution for discrete POMDP has no choice but to traverse four half circle arcs to approximate a straight path. Whereas, the solution for the continuous POMDP is able to approximate a straight path more closely.

In the more difficult problem with three obstacles around the target, both discrete and continuous POMDPs arrive at the same policy with a discount factor of 0.92 and 0.9 respectively as shown in Figures 7 and 8, respectively. When the discount factor increases, the policy obtained from discrete POMDP remains the same. Whereas the continuous POMDP is able to find a shorter path by navigating through a narrow path in




Fig. 9. Discount factor = 0.95, target reward = 100, obstacle reward = -20, boundary reward = -100

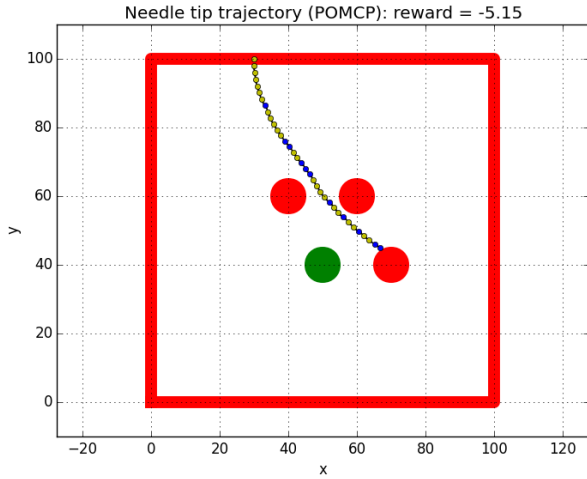


Fig. 10. Discount factor = 0.95, target reward = 100, obstacle reward = -20, boundary reward = -100

between obstacles at [40mm, 60mm] and [60mm, 60mm] as shown in Figure 9. The continuous POMDP formulation has the advantage of discovering these more aggressive policies because it is able to change actions more frequently compared to the discrete POMDP formulation. The trade-off of taking these more aggressive policies is that the failure rate increases to 20%.

## VI. CONCLUSION

This work investigates two ways to model needle insertion in tissues using POMDP. Both the discrete method and the continuous method are able to generate policies that can drive needle to the target while dodging obstacles under transition and observation uncertainty. We have shown that the continuous POMDP problem is easier to solve than discrete POMDP problem, and the continuous policy has higher spatial

resolution hence better performance. In conclusion, continuous POMDP with an online solver like POMCP is more suitable for path planning of bevel-tip needle insertion.

## VII. FUTURE WORK

The most direct extension of this work is to go to a three-dimensional tissue model. For the discrete problem, the size of state space will increase due to the extra dimension. For example, there will be 27 neighboring grid nodes, each with 27 possible orientations of the needle tip. For the continuous problem, the same beam theory still holds and we will need to keep track of the normal vector of the bevel surface at the tip. Six variables, three for location and three for orientation, will be used to define the state.

Another interesting parameter that was not accounted for in our models is the physician action uncertainty. Humans do not necessarily act rationally. Through experience, physicians have learned to trust what they see and feel, and this intuition can overwrite what instructions they are given by our algorithm. This uncertainty could be included in the problem in two ways: i) introducing noise into the actions or ii) spreading the probabilities assigned in the transition model to states corresponding to the opposite action.

## ACKNOWLEDGMENT

The authors would like to thank Johannes Weickenmeier, Zach Sunberg, Ethan Chan, and Mykel Kochenderfer for their technical support to this work.

## REFERENCES

- [1] P. Han, D. Che, K. Pallav, and K. Ehmann, "Models of the cutting edge geometry of medical needles with applications to needle design," *International Journal of Mechanical Sciences*, vol. 65, no. 1, pp. 157–167, 2012.
- [2] R. Alterovitz, M. Branicky, and K. Goldberg, "Motion planning under uncertainty for image-guided medical needle steering," *The International journal of robotics research*, vol. 27, no. 11-12, pp. 1361–1374, 2008.
- [3] S. Patil and R. Alterovitz, "Interactive motion planning for steerable needles in 3d environments with obstacles," in *Biomedical Robotics and Biomechanics (BioRob), 2010 3rd IEEE RAS and EMBS International Conference on*. IEEE, 2010, pp. 893–899.
- [4] S. Patil *et al.*, "Motion planning under uncertainty in highly deformable environments," *Robotics science and systems: online proceedings*, 2011.
- [5] Y.-L. Park, S. Elayaperumal, B. Daniel, S. C. Ryu, M. Shin, J. Savall, R. J. Black, B. Moslehi, and M. R. Cutkosky, "Real-time estimation of 3-d needle shape and deflection for mri-guided interventions," *IEEE/ASME Transactions On Mechatronics*, vol. 15, no. 6, pp. 906–915, 2010.
- [6] H. Kurniawati, D. Hsu, and W. S. Lee, "Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces." in *Robotics: Science and Systems*, vol. 2008. Zurich, Switzerland, 2008.
- [7] G. Shani, J. Pineau, and R. Kaplow, "A survey of point-based pomdp solvers," *Autonomous Agents and Multi-Agent Systems*, vol. 27, no. 1, pp. 1–51, 2013.
- [8] D. Silver and J. Veness, "Monte-carlo planning in large pomdps," in *Advances in neural information processing systems*, 2010, pp. 2164–2172.
- [9] A. Majewicz, J. J. Siegel, A. A. Stanley, and A. M. Okamura, "Design and evaluation of duty-cycling steering algorithms for robotically-driven steerable needles," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 5883–5888.