Algorithmic Motion Planning Meets Minimally-Invasive Robotic Surgery

Oren Salzman
Technion—Israel Institute of Technology

Abstract

Robots for minimally-invasive surgery such as steerable needles and concentric-tube robots have the potential to dramatically alter the way common medical procedures are performed. They can decrease patient-recovery time, speed healing and reduce scarring. However, manually controlling such devices is highly un-intuitive and automatic planning methods are in need. For the automation of such medical procedures to be clinically accepted, it is critical from a patient care, safety, and regulatory perspective to certify the correctness and effectiveness of the motion-planning algorithms involved in procedure automation. In this paper, I survey recent and ongoing work where we develop efficient and effective planning capabilities for medical robots that provide provable guarantees on various planner attributes as well as introduce new and exciting research opportunities in the field.

1 Introduction

Continuum robots are biologically inspired systems composed of flexible, elastic, or soft materials that allow for complex bending motions. Examples include steerable needles (needles that have the ability to change their direction while inside tissue) [Reed *et al.*, 2011; Alterovitz *et al.*, 2003], and concentric tube robots (robotic systems consisting of multiple flexible tubes of different lengths and curvatures arranged concentrically that are controlled by actuating the relative movements of the tubes) [Gilbert *et al.*, 2013]. Their ability to follow 3D curvilinear trajectories make them excellent candidates for medical applications that require reaching clinically significant targets while safely avoiding critical anatomical structures [Alterovitz and Goldberg, 2008; Burgner-Kahrs *et al.*, 2015]. In contrast to conventional rigid medical instruments, medical continuum robots can minimize a patient's discomfort and enhance safety, providing minimally invasive access to previously unreachable targets. They have shown potential for various diagnostic and treatment procedures such as biopsy, drug delivery, and radioactive seed implantation for cancer therapy [Abolhassani *et al.*, 2007].

Manual control of such devices is highly unintuitive for human operators because of kinematic constraints on their 3D motion and the fact that the anatomical environment is highly cluttered. Therefore, automation is crucial to fully exploit their capabilities. Indeed, many planners have been proposed for different tasks and different continuum robots. However for the automation of such medical procedures to be clinically accepted, it is critical from a patient care, safety, and regulatory perspective to certify the correctness and effectiveness of the planning algorithms involved in procedure automation, a guarantee that is typically lacking from existing planners. To this end, in this paper, I survey recent and ongoing work where we develop efficient and effective planning capabilities for medical robots that provide provable guarantees on various planner attributes as well as introduce new and exciting research opportunities in the field.

2 Motion Planning for Steerable Needles

Steerable needles can be used to reach clinical targets for biopsy purposes while safely avoiding obstacles such as blood vessels. This clinical application can be modelled as a *motion-planning* problem which is the problem of determining a collision-free path or trajectory for a robot to move from its initial position to a desired goal position while avoiding obstacles in its environment. Ideally, a motion-planning algorithm should first guarantee that it will compute a solution, when one exists, in finite time, or notify the user that no solution exists. Moreover, the computed solution should strive to maximize some objective which in our setting is patient safety. This can be quantified using metrics such as minimizing trajectory length [Favaro *et al.*, 2018], maximizing clearance from obstacles [Kuntz *et al.*, 2015], and minimizing damage to sensitive tissue [Fu *et al.*, 2018; Bentley *et al.*, 2021].

Thus, a motion planner should be *complete*, i.e., find a solution plan in a finite number of steps, if one exists, and ideally should be *optimal*, i.e., ensure that the returned plan has a cost (for a given cost metric) that equals the global optimum. Unfortunately, no previously developed planner for steerable needles (see, e.g., [Xu *et al.*, 2008; Patil *et al.*, 2014; Pinzi *et al.*, 2021]) offers a formal guarantee on completeness, let alone optimality.

Some prior motion planners for steerable needles do aim to optimize motion plan cost but they lack *global* optimality guarantees [Liu *et al.*, 2016; Favaro *et al.*, 2018; Pinzi *et al.*, 2019; Favaro *et al.*, 2021]. Indeed, providing

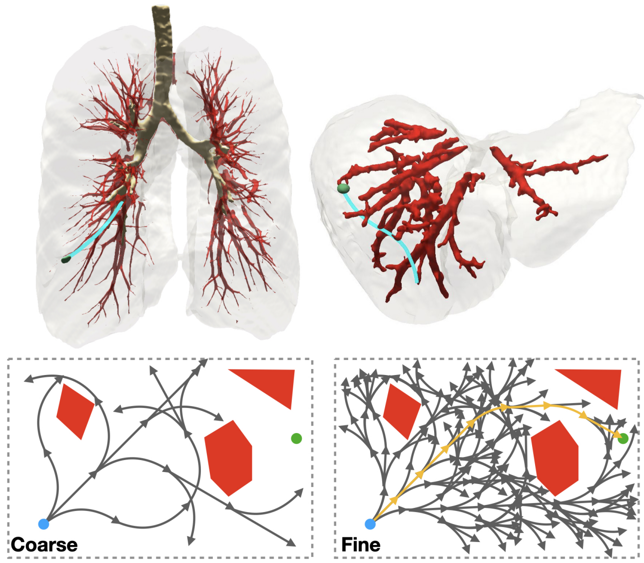

Figure 1: **Top:** A medical steerable needle (cyan) is used to reach a nodule (green) while avoiding major blood vessels (red) for biopsy or cancer treatment. The left and right figures depicts the lung parenchyma (where the bronchial tubes are in brown) and the liver, respectively. **Bottom:** Our resolution-complete motion planner uses search trees built using different resolutions, illustrated here in 2D. A valid motion plan goes from the start configuration (blue dot) to the goal point (green dot), while avoiding obstacles (red) and satisfying kinematic constraints. The left search tree uses a coarse resolution and fails to find a plan while the right one uses a finer resolution and successfully generates a motion plan (yellow).

completeness and optimality guarantees for a steerable needle motion planner is challenging in part because motion planning for steerable needles in 3D with curvature constraints is at least NP-hard [Solovey, 2020]. Some sampling-based planners are known to be both complete and optimal, albeit those properties are usually proven only for an asymptotic regime where the number of samples tends to infinity [LaValle, 2006; Salzman, 2019; Kleinbort *et al.*, 2020]. Interestingly, recent work has developed optimality guarantees for finite sampling, although those results cannot be currently applied to steerable needles as they deal with holonomic systems [Tsao *et al.*, 2020; Dayan *et al.*, 2021].

The aforementioned challenges inspired us to consider variants of completeness or optimality relevant to medical applications. Specifically, in a series of recent papers [Fu *et al.*, 2021b; Fu *et al.*, 2022], we focus on specific types of guarantees relevant to real-world medical applications: *resolution completeness* [LaValle, 2006] and *resolution optimality* [Barraquand and Latombe, 1993; Pivtoraiko *et al.*, 2009]. Generally speaking, a resolution characterizes the discretization of some space (e.g., state space, configuration space, action space, and time). An algorithm is resolution complete if there exists a fine-enough resolution with which the algorithm finds a plan in finite time when a qualified solution exists, and otherwise correctly returns that no such plan exists. An algorithm is resolution optimal if it is resolution complete and if, when it does return a motion plan, the plan's cost is guaranteed to be within a desired approximation factor of the cost of a globally optimal qualified motion plan. We illustrate at the bottom of Fig. 1 an example showing searches with different resolutions for needle steering.

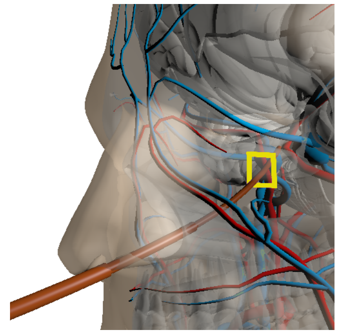

Figure 2: CTR deployed via the sinus to cut a window in the skull by following a reference path provided by a surgeon (yellow).

We first presented Resolution-Complete Search (RCS), an efficient, resolution-complete motion planner for steerable needles based on a novel adaptation of multi-resolution planning [Fu *et al.*, 2021b]. Under some mild assumptions on the system and the solution, the planner, in finite time, is guaranteed to find a plan as long as the problem admits a qualified solution. We then extended RCS to develop RCS*, that achieves resolution optimality [Fu *et al.*, 2022].

We demonstrated the performance of our planners in two clinically realistic scenarios where the needle should reach a target while safely avoiding obstacles (e.g., blood vessels). In the setting of (i) lung biopsy, where the needle is deployed through a bronchoscope and must steer through the lung parenchyma (the tissue of the lung outside the bronchial tubes) and in the setting of (ii) liver biopsy, where the needle is deployed into the liver through its anterior surface and must steer through the liver tissue. We compared in simulation our planner with several other steerable needle planners and demonstrated experimentally that RCS and RCS* outperform the state-of-the-art in terms of computation time, success rate, and plan quality.

3 Path Following for Concentric Tube Robots

The dexterity and small diameter of Concentric tube robots, or CTRs [Gilbert *et al.*, 2013] enable minimally-invasive surgery in constrained areas, such as accessing the pituitary gland via the sinuses. For example, they can be used to cut away a window of tissue during a procedure where the tissue to be cut is expected to be provided by a surgeon as a path that the CTR's end effector should trace or follow. This can be modelled as a *path-following problem* where we are given an entire *path* R for the robot's tip and are tasked with computing a robot plan such that when following this plan, it's end effector traces R "as closely as possible" (this is formally defined by Holladay et al. [2019]).

Although existing planners [Torres *et al.*, 2015] enable CTRs to reach specified points in task-space (space of a robot's end-effector positions), this is often insufficient—existing path-following algorithms [Berenson *et al.*, 2009;

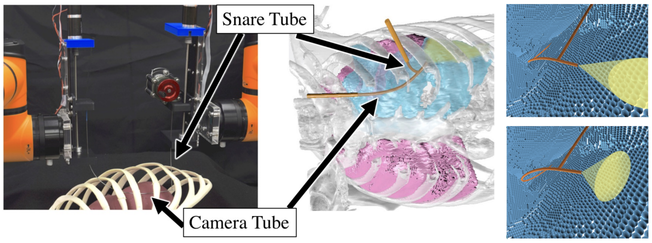

Figure 3: Example of Human anatomy inspection. **Left:** The CRISP robot is composed of needle-diameter tubes assembled into a parallel structure inside the patient's body (in which a tube uses a snare system to grip a tube with a camera affixed to its tip) and then manipulated outside the body, allowing for smaller incisions and faster recovery times compared to traditional endoscopic tools (which have larger diameters). **Middle:** The CRISP robot in simulation inspecting a collapsed lung, a scenario segmented from a CT scan of a real patient with this condition. Visualization shows the robot (orange), the lungs (pink), and the pleural surface visible (green) and not visible (blue) by the robot's camera sensor in its current configuration. **Right:** Two example configurations with inspected POI. The CRISP robot (orange) inspects POI (blue) on the organ surface with visible points covered by the cone shape (yellow).

Holladay *et al.*, 2019] have difficulty planning robustly in the cluttered environments that arise in our medical domains (e.g., cutting a window in the skull during brain surgery as depicted in Fig. 2). To this end, we introduce the NNF (Nearest-Neighbor Fréchet) motion planner—the first planner in this domain able to compute a trajectory that closely follows such a task-space path [Niyaz *et al.*, 2018].

However, our key insight is that truly performing our task optimally requires not only robust planning but also optimizing the *robot placement* (namely, where the fixed based of the robot is placed prior to the surgery). This is a critical choice in a highly-constrained environment, given that certain placements will allow the robot a greater range of motion within the already-limited space, and thus better enable it to follow the path with its end-effector.

Our final approach [Niyaz *et al.*, 2019] therefore uses the output of the NNF planner to inform a gradient-free optimizer over the robot's placement, and thus simultaneously plans both the motions to execute and the placement to execute them from. In order to make this approach as efficient as possible, we also enable tight integration between the motion planner and optimizer to bound unnecessary computation. Our experiments confirm both that (i) incorporating the choice of placement dramatically improves the ability to follow a given path and (ii) tightly integrating the planner and optimizer in this way affords our approach a significant reduction in compute time.

4 Inspection Planning for CRISP Robot

In certain medical applications, physicians may want to inspect some region of interest for diagnostic purposes completing the procedure as fast as is safely possible to reduce costs and improve patient outcomes, especially if the patient is under anesthesia during the procedure. For example, the needle-diameter Continuum Reconfigurable Incisionless Surgical Parallel (CRISP) robot [Anderson *et al.*, 2017; Mahoney *et al.*, 2016] was suggested to assist in the diagnosis of the cause of a pleural effusion (a serious medical condition that can cause the collapse of a patient's lung) by visually inspecting the surface of the collapsed lung and chest wall inside the body (see Fig. 3) [Kuntz *et al.*, 2018]. This can modelled as an *inspection-planning problem* [Almadhoun *et al.*, 2016; Galceran and Carreras, 2013] where we are given a robot equipped with a sensor and a set of points of interest (POI) in the environment to be inspected by the sensor. The problem calls for computing a minimal-length motion plan for the robot that maximizes the number of POI inspected.

Naïvely computed inspection plans may enable inspection of only a subset of the POI and may require motion plans significantly longer than an optimal plan, and hence may be undesirable or infeasible due to time constraints. Our goal is to compute kinematically feasible collision-free inspection plans that maximize the number of POI inspected, and of the plans that inspect those POI we compute a shortest one.

Inspection planning is computationally challenging because we need to simultaneously reason both about motion planning in a high-dimensional configuration space $\mathcal{X}$ (the space of all parameters that determine the shape of the robot) [LaValle, 2006] as well as about maximizing the number of POI inspected.

There are multiple approaches to computing inspection plans. Optimization-based methods locally search over the space of all inspection plans [Bircher *et al.*, 2015; Bogaerts *et al.*, 2018]. Decoupled approaches first independently select suitable viewpoints and then determine a visiting sequence, i.e., a motion plan for the robot that realizes this sequence [Englot and Hover, 2011]. Finally, recent progress in motion planning [Karaman and Frazzoli, 2011] has enabled methods to exhaustively search over the space of all motion plans [Bircher *et al.*, 2017; Kafka *et al.*, 2016] thus guaranteeing asymptotic optimality. Roughly speaking, asymptotic optimality for inspection planning means these methods produce inspection plans whose length and the number of points inspected will asymptotically converge to those of an optimal inspection plan, given enough planning time.

Of all the aforementioned methods, only algorithms in the latter group provide any formal guarantees on the quality of the solution. This guarantee is achieved by attempting to exhaustively compute the set of Pareto-optimal inspection plans embedded in $\mathcal{X}$ for which full coverage has not been obtained. Namely, for every configuration $x \in \mathcal{X}$, they (asymptotically) compute the set of paths Π_x starting from a given start configuration x_s such that $\forall \pi_1, \pi_2 \in \Pi_x$, either π_1 is shorter than π_2 and π_2 covers POI not covered by π_1 or vice versa. Once Π_x contains a path π^*_x that covers all POI, this path is considered as a candidate solution. In our setting, the set of Pareto-optimal inspection plans is the minimal set of inspection plans such that each plan is either shorter or has better coverage of the POI than any other inspection plan.[1]

[1]More formally, an inspection plan P connecting two configurations $\mathbf{q}, \mathbf{q}' \in \mathcal{X}$ is said to be Pareto optimal in our setting if any other plan connecting $\mathbf{q}$ to $\mathbf{q}'$ is either longer or does not inspect a point visible to P.

Unfortunately, this comes at the price of very long computation times as the size of the search space is exponential in the number of POI.

To this end, we introduce Incremental Random Inspection-roadmap Search (IRIS), an asymptotically optimal inspection-planning algorithm [Fu *et al.*, 2019; Fu *et al.*, 2021a]. IRIS incrementally constructs a sequence of increasingly dense roadmaps—graphs embedded in $\mathcal{X}$ wherein each vertex represents a collision-free configuration and each edge a collision-free transition between configurations—and computes an inspection plan on the roadmaps as they are constructed.

Unfortunately, even the problem of computing an optimal inspection plan on a graph (and not in the configuration space) is computationally hard. To this end, our key insight is to compute a *near-optimal* inspection plan on each roadmap. Namely, we compute an inspection plan that is at most $1 + \varepsilon$ the length of an optimal plan while covering at least p-percent of the number of POI (for any $\varepsilon \geq 0$ and $p \in (0, 1]$). This additional flexibility allows us to improve the quality of our inspection plan in an *anytime* manner, i.e., the algorithm can be stopped at any time and return the best inspection plan found up until that point. We achieve this by incrementally densifying the roadmap and then searching over the densified roadmap for a near-optimal inspection plan—a process that is repeated as time allows. By reducing the approximation factor between iterations, we ensure that our method is asymptotically optimal.

The key contribution of this work is a computationally efficient algorithm to compute provably near-optimal inspection plans on graphs. Coupled with our method for generating this graph, this algorithmic building block enables us to dramatically outperform Rapidly-exploring Random Tree Of Trees (RRTOT) [Bircher *et al.*, 2017]—a state-of-the-art asymptotically optimal inspection planner. Specifically, we demonstrated the efficacy of our approach in simulation for several complex robotic systems (Fig. 3).

5 Summary and Future Work

In this paper we reviewed several algorithmic planning problems (motion planning, path following and inspection planning) and demonstrated their applicability and dedicated solutions for minimally invasive robot surgery that come with some guarantees. We continue to briefly introduce new and exciting research opportunities in the field.

Task Inspection Planning In semi-autonomous minimally-invasive robotic surgery we have one robot tele-operated by a surgeon who is tasked with suturing or removing a tumor while a second, fully autonomous robot, is in charge of autonomously providing the surgeon with visual feedback on the progress of the task. This can be modelled as a *task inspection planning* where we are given two robots: one, which we call the *task robot*, that is in charge of performing a given task, and the second, which we call the *inspection robot*, is in charge of inspecting the task as performed by the task robot using on-board sensors.

This problem is not new [Robin and Lacroix, 2016] however it is typically considered in the *adversarial setting* where the inspection robot is required to avoid being observed by the inspection robot. Our version bares resemblance to the path following and inspection-planning problems we mentioned and we anticipate that the algorithmic building blocks developed for these problems will serve as foundations to solving the task-inspection planning problem. However, a unique challenge we foresee in addressing this problem is robustness—computing a path that is robust to execution uncertainty of both robots while guaranteeing that the task is indeed inspected.

Therapy Planning As steerable medical devices become more capable of taking expressive paths, they unlock the door for therapy delivery in ways not previously possible. This may include drug-based therapies such as gene therapy compounds [Lonser *et al.*, 2020] and cellucidal compounds [Vogelbaum and Aghi, 2015], or energy-based therapy such as laser/radiofrequency ablation [Ashraf *et al.*, 2018; Shimamoto *et al.*, 2019; Voges *et al.*, 2018]. Automating the planning of such therapy for *straight-line* paths of medical devices is itself a challenging problem [Zhang *et al.*, 2019], and considering non-straight paths only adds complexity.

We believe that these are fundamentally coupled problems: Planning how to deliver therapy (e.g., pose and volume of liquid drug delivery, or pose, intensity, and duration of energy delivery) and planning how the steerable device moves to those locations must be considered in concert. We deem this problem the *therapy planning* problem.

Preliminary, we are abstracting the therapy planning problem as a multi-criteria optimization problem in which we must simultaneously consider objectives that incorporate the delivery of therapy to the intended tissue, the delivery of therapy to unintended tissue, some cost associated with moving the robot to the therapy delivery poses, and potentially other medical objectives (e.g., total time under anesthesia, etc.).

This is an extremely complex problem. Modeling the therapy involves potentially non-uniform dispersion of drug or diffusion of energy [Knavel and Brace, 2013]. Modeling the motion of the robot itself is complex for the reasons listed above, and multi-criteria optimization is inherently complex itself as well (see, e.g., [Salzman *et al.*, 2023]).

Acknowledgements

This research was supported by the United States-Israel Binational Science Foundation (BSF) grants no. 2019703 and 2021643 and by the Israeli Ministry of Science & Technology grants No. 3-16079 and 3-17385. I am greatly indebted to all my coauthors. Specifically, this research was and is carried out in collaboration with Ron Alterovitz, Mengyu Fu, Alan Kuntz, Michael Bentley, Caleb Rucker, Chakravarthy Reddy Siddhartha Srinivasa, Sherdil Niyaz, and Kiril Solovey.

References

[Abolhassani *et al.*, 2007] Niki Abolhassani, Rajni Patel, and Mehrdad Moallem. Needle insertion into soft tissue: A survey. *Medical Engineering & Physics*, 29(4):413–431, 2007.

[Almadhoun *et al.*, 2016] Randa Almadhoun, Tarek Taha, Lakmal Seneviratne, Jorge Dias, and Guowei Cai. A survey on inspecting structures using robotic systems. *International Journal of Advanced Robotic Systems*, 13(6), 2016.

[Alterovitz and Goldberg, 2008] Ron Alterovitz and Kenneth Y. Goldberg. *Motion Planning in Medicine: Optimization and Simulation Algorithms for Image-Guided Procedures*, volume 50 of *Springer Tracts in Advanced Robotics*. Springer, 2008.

[Alterovitz *et al.*, 2003] Ron Alterovitz, Kenneth Y. Goldberg, Jean Pouliot, Richard Tascherau, and I-Chow Hsu. Needle insertion and radioactive seed implantation in human tissues: simulation and sensitivity analysis. In *ICRA*, pages 1793–1799, 2003.

[Anderson *et al.*, 2017] Patrick L Anderson, Arthur W Mahoney, and Robert J Webster III. Continuum reconfigurable parallel robots for surgery: Shape sensing and state estimation with uncertainty. *IEEE Robot. Autom. Lett.*, 2(3):1617–1624, 2017.

[Ashraf *et al.*, 2018] Omar Ashraf, Nitesh V. Patel, Simon Hanft, and Shabbar F. Danish. Laser-Induced Thermal Therapy in Neuro-Oncology: A Review. *World Neurosurgery*, 112:166–177, April 2018.

[Barraquand and Latombe, 1993] Jérôme Barraquand and Jean-Claude Latombe. Nonholonomic multibody mobile robots: Controllability and motion planning in the presence of obstacles. *Algorithmica*, 10(2):121–155, 1993.

[Bentley *et al.*, 2021] Michael Bentley, D. Caleb Rucker, Chakravarthy Reddy, Oren Salzman, and Alan Kuntz. A novel shaft-to-tissue force model for safer motion planning of steerable needles. *CoRR*, abs/2101.02246, 2021.

[Berenson *et al.*, 2009] D. Berenson, S.S. Srinivasa, D. Ferguson, and J. Kuffner. Manipulation planning on constraint manifolds. In *ICRA*, pages 625–632, 2009.

[Bircher *et al.*, 2015] Andreas Bircher, Kostas Alexis, Michael Burri, Philipp Oettershagen, Sammy Omari, Thomas Mantel, and Roland Siegwart. Structural inspection path planning via iterative viewpoint resampling with application to aerial robotics. In *ICRA*, pages 6423–6430. IEEE, 2015.

[Bircher *et al.*, 2017] Andreas Bircher, Kostas Alexis, Ulrich Schwesinger, Sammy Omari, Michael Burri, and Roland Siegwart. An incremental sampling–based approach to inspection planning: The rapidly–exploring random tree of trees. *Robotica*, 35(6):1327–1340, 2017.

[Bogaerts *et al.*, 2018] Boris Bogaerts, Seppe Sels, Steve Vanlanduit, and Rudi Penne. A gradient-based inspection path optimization approach. *IEEE Robot. Autom. Lett.*, 3(3):2646–2653, 2018.

[Burgner-Kahrs *et al.*, 2015] Jessica Burgner-Kahrs, D. Caleb Rucker, and Howie Choset. Continuum robots for medical applications: A survey. *IEEE Trans. Rob.*, 31(6):1261–1280, 2015.

[Dayan *et al.*, 2021] Dror Dayan, Kiril Solovey, Marco Pavone, and Dan Halperin. Near-optimal multi-robot motion planning with finite sampling. In *ICRA*, pages 9190–9196, 2021.

[Englot and Hover, 2011] Brendan J Englot and Franz S Hover. Planning complex inspection tasks using redundant roadmaps. In *ISRR*, pages 327–343, 2011.

[Favaro *et al.*, 2018] Alberto Favaro, Leonardo Cerri, Stefano Galvan, Ferdinando Rodriguez Y Baena, and Elena De Momi. Automatic optimized 3D path planner for steerable catheters with heuristic search and uncertainty tolerance. In *ICRA*, pages 9–16. IEEE, 2018.

[Favaro *et al.*, 2021] Alberto Favaro, Alice Segato, Federico Muretti, and Elena De Momi. An evolutionary-optimized surgical path planner for a programmable bevel-tip needle. *IEEE Trans. Rob.*, 37(4):1039–1050, 2021.

[Fu *et al.*, 2018] Mengyu Fu, Alan Kuntz, Robert J Webster III, and Ron Alterovitz. Safe motion planning for steerable needles using cost maps automatically extracted from pulmonary images. In *IROS*, pages 4942–4949. IEEE, 2018.

[Fu *et al.*, 2019] Mengyu Fu, Alan Kuntz, Oren Salzman, and Ron Alterovitz. Toward asymptotically-optimal inspection planning via efficient near-optimal graph search. In *RSS*, 2019.

[Fu *et al.*, 2021a] Mengyu Fu, Oren Salzman, and Ron Alterovitz. Computationally-efficient roadmap-based inspection planning via incremental lazy search. In *ICRA*, pages 7449–7456, 2021.

[Fu *et al.*, 2021b] Mengyu Fu, Oren Salzman, and Ron Alterovitz. Toward certifiable motion planning for medical steerable needles. In *RSS*, 2021.

[Fu *et al.*, 2022] Mengyu Fu, Kiril Solovey, Oren Salzman, and Ron Alterovitz. Resolution-optimal motion planning for steerable needles. In *ICRA*, pages 9652–9659, 2022.

[Galceran and Carreras, 2013] Enric Galceran and Marc Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12):1258–1276, 2013.

[Gilbert *et al.*, 2013] H. Gilbert, D. Rucker, and R. J. Webster III. Concentric tube robots: The state of the art and future directions. In *ISRR*, pages 253–269, 2013.

[Holladay *et al.*, 2019] Rachel M. Holladay, Oren Salzman, and Siddhartha S. Srinivasa. Minimizing task-space fréchet error via efficient incremental graph search. *IEEE Robot. Autom. Lett.*, 4(2):1999–2006, 2019.

[Kafka *et al.*, 2016] Přemysl Kafka, Jan Faigl, and Petr Váňa. Random inspection tree algorithm in visual inspection with a realistic sensing model and differential constraints. In *ICRA*, pages 2782–2787, 2016.

[Karaman and Frazzoli, 2011] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *Int. J. of Rob. Res.* , 30(7):846–894, 2011.

[Kleinbort *et al.*, 2020] Michal Kleinbort, Edgar Granados, Kiril Solovey, Riccardo Bonalli, Kostas E Bekris, and Dan Halperin. Refined analysis of asymptotically-optimal kinodynamic planning in the state-cost space. In *ICRA*, pages 6344–6350. IEEE, 2020.

[Knavel and Brace, 2013] Erica M. Knavel and Christopher L. Brace. Tumor ablation: Common modalities and general practices. *Techniques in Vascular and Interventional Radiology*, 16(4):192–200, December 2013.

[Kuntz *et al.*, 2015] Alan Kuntz, Luis G Torres, Richard H Feins, Robert J Webster III, and Ron Alterovitz. Motion planning for a three-stage multilumen transoral lung access system. In *IROS*, pages 3255–3261. IEEE, 2015.

[Kuntz *et al.*, 2018] Alan Kuntz, Chris Bowen, Cenk Baykal, Arthur W Mahoney, Patrick L Anderson, Fabien Maldonado, Robert J Webster III, and Ron Alterovitz. Kinematic design optimization of a parallel surgical robot to maximize anatomical visibility via motion planning. In *ICRA*, pages 926–933, 2018.

[LaValle, 2006] Steven M LaValle. *Planning algorithms*. Cambridge university press, 2006.

[Liu *et al.*, 2016] Fangde Liu, Arnau Garriga-Casanovas, Riccardo Secoli, and Ferdinando Rodriguez y Baena. Fast and adaptive fractal tree-based path planning for programmable bevel tip steerable needles. *IEEE Robot. Autom. Lett.*, 1(2):601–608, 2016.

[Lonser *et al.*, 2020] Russell R. Lonser, Asad S. Akhter, Mirosław Zabek, J. Bradley Elder, and Krystof S. Bankiewicz. Direct convective delivery of adeno-associated virus gene therapy for treatment of neurological disorders. *Journal of Neurosurgery*, pages 1–13, July 2020.

[Mahoney *et al.*, 2016] Arthur W Mahoney, Patrick L Anderson, Philip J Swaney, Fabien Maldonado, and Robert J Webster III. Reconfigurable parallel continuum robots for incisionless surgery. In *IROS*, pages 4330–4336, 2016.

[Niyaz *et al.*, 2018] S. Niyaz, A. Kuntz, O. Salzman, R. Alterovitz, and S. S. Srinivasa. Following surgical trajectories with concentric tube robots. In *ISER*, 2018.

[Niyaz *et al.*, 2019] S. Niyaz, A. Kuntz, O. Salzman, R. Alterovitz, and S. S. Srinivasa. Optimizing motion-planning problem setup via bounded evaluation with application to following surgical trajectories. In *IROS*, 2019.

[Patil *et al.*, 2014] Sachin Patil, Jessica Burgner, Robert J Webster III, and Ron Alterovitz. Needle steering in 3D via rapid replanning. *IEEE Trans. Rob.*, 30(4):853–864, 2014.

[Pinzi *et al.*, 2019] Marlene Pinzi, Stefano Galvan, and Ferdinando Rodriguez y Baena. The adaptive hermite fractal tree (AHFT): a novel surgical 3D path planning approach with curvature and heading constraints. *International Journal of Computer Assisted Radiology and Surgery*, 14(4):659–670, 2019.

[Pinzi *et al.*, 2021] Marlene Pinzi, Thomas Watts, Riccardo Secoli, Stefano Galvan, and Ferdinando Rodriguez y Baena. Path replanning for orientation-constrained needle steering. *IEEE Transactions on Biomedical Engineering*, 68(5):1459–1466, 2021.

[Pivtoraiko *et al.*, 2009] Mihail Pivtoraiko, Ross A Knepper, and Alonzo Kelly. Differentially constrained mobile robot motion planning in state lattices. *J. of Field Robotics*, 26(3):308–333, 2009.

[Reed *et al.*, 2011] Kyle B. Reed, Ann Majewicz, Vinutha Kallem, Ron Alterovitz, Kenneth Y. Goldberg, Noah J. Cowan, and Allison M. Okamura. Robot-assisted needle steering. *IEEE Robotics Autom. Mag.*, 18(4):35–46, 2011.

[Robin and Lacroix, 2016] Cyril Robin and Simon Lacroix. Multi-robot target detection and tracking: taxonomy and survey. *Autonomous Robots*, 40:729–760, 2016.

[Salzman *et al.*, 2023] Oren Salzman, Ariel Felner, Carlos Hernandez, Han Zhang, Shao-Hung Chan, and Sven Koenig. Heuristic-search approaches for the multi-objective shortest-path problem: Progress and research opportunities. In *IJCAI*, 2023.

[Salzman, 2019] Oren Salzman. Sampling-based robot motion planning. *Commun. ACM*, 62(10):54–63, 2019.

[Shimamoto *et al.*, 2019] Shoichi Shimamoto, Chengyuan Wu, and Michael R. Sperling. Laser interstitial thermal therapy in drug-resistant epilepsy. *Current Opinion in Neurology*, 32(2):237–245, April 2019.

[Solovey, 2020] Kiril Solovey. Complexity of planning. *arXiv preprint arXiv:2003.03632v2 [cs.RO]*, 2020.

[Torres *et al.*, 2015] L. Torres, A. Kuntz, H. Gilbert, P. Swaney, R. Hendrick, R. J. Webster III, and R. Alterovitz. A motion planning approach to automatic obstacle avoidance during concentric tube robot teleoperation. In *ICRA*, pages 2361–2367, 2015.

[Tsao *et al.*, 2020] Matthew Tsao, Kiril Solovey, and Marco Pavone. Sample complexity of probabilistic roadmaps via ϵ-nets. In *ICRA*, pages 2196–2202. IEEE, 2020.

[Vogelbaum and Aghi, 2015] Michael A. Vogelbaum and Manish K. Aghi. Convection-enhanced delivery for the treatment of glioblastoma. *Neuro-Oncology*, 17 Suppl 2:ii3–ii8, March 2015.

[Voges *et al.*, 2018] J. Voges, L. Büntjen, and F. C. Schmitt. Radiofrequency-thermoablation: General principle, historical overview and modern applications for epilepsy. *Epilepsy Research*, 142:113–116, May 2018.

[Xu *et al.*, 2008] Jijie Xu, Vincent Duindam, Ron Alterovitz, and Ken Goldberg. Motion planning for steerable needles in 3D environments with obstacles using rapidly-exploring random trees and backchaining. pages 41–46. IEEE, 2008.

[Zhang *et al.*, 2019] Rui Zhang, Shuicai Wu, Weiwei Wu, Hongjian Gao, and Zhuhuang Zhou. Computer-assisted needle trajectory planning and mathematical modeling for liver tumor thermal ablation: A review. *Mathematical Biosciences and Engineering*, 16(5):4846–4872, 2019.