Motion Planning for Continuum Reconfigurable Incisionless Surgical Parallel Robots

Alan Kuntz, Arthur W. Mahoney, Nicolas E. Peckman, Patrick L. Anderson,
Fabien Maldonado, Robert J. Webster III, and Ron Alterovitz

Abstract— **Continuum Reconfigurable Incisionless Surgical Parallel (CRISP) robots consist of multiple needle-diameter flexible instruments that are assembled into a parallel structure inside the human body. With a camera placed at the tip of one of the instruments, the CRISP robot can be used to inspect anatomical sites in constrained body cavities in a minimally invasive manner. We introduce a motion planner for CRISP robots that computes manipulations of the flexible instruments outside the body such that the camera can visually inspect a user-specified site of clinical interest inside the body. Our sampling-based motion planner ensures avoidance of collisions with anatomical obstacles inside the body, enforces remote-center-of-motion constraints on the instrument's entry points into the body, and efficiently handles the expensive computation of CRISP robot kinematics. We also extend the motion planner to estimate the set of points inside a body cavity that can be visually inspected by the camera of a CRISP robot for a given setup. We demonstrate our method in a simulated endoscopic medical procedure in the pleural space around a lung.**

I. INTRODUCTION

The Continuum Reconfigurable Incisionless Surgical Parallel (CRISP) robot [1] is a new type of continuum robot [2], [3] that consists of multiple needle-diameter flexible instruments that are assembled together inside a body cavity to perform minimally invasive medical procedures. The CRISP robot typically includes (1) a flexible instrument with a working channel through which a tool (e.g., chip-tip camera, ablation probe, etc.) is passed, and (2) one or more additional flexible instruments that are inserted into the body cavity and attach to the first instrument via snares, creating a strong, parallel kinematic structure (see Fig. 1). The parallel nature of the structure provides strength to the robot, enabling the device to apply larger forces during medical procedures when required. The tool's tip can be repositioned and reoriented inside the body by robotically moving the instruments outside the body in concert.

The CRISP robot is an ideal platform for inspecting and manipulating tissues on the surface of a pleural effusion, which is a collection of excess fluid in the pleural space around the lungs. Pleural effusions can be caused by over 50 different diseases [4]. Accurate diagnosis of the disease

A. Kuntz, N. E. Peckman, and R. Alterovitz are with the Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA. {adkuntz,peckman,ron}@cs.unc.edu

A. W. Mahoney, P. L. Anderson, and R. J. Webster III are with the Department of Mechanical Engineering, Vanderbilt University, Nashville, TN 37235, USA. F. Maldonado is with the Vanderbilt University Medical Center, Nashville, TN 37232, USA.
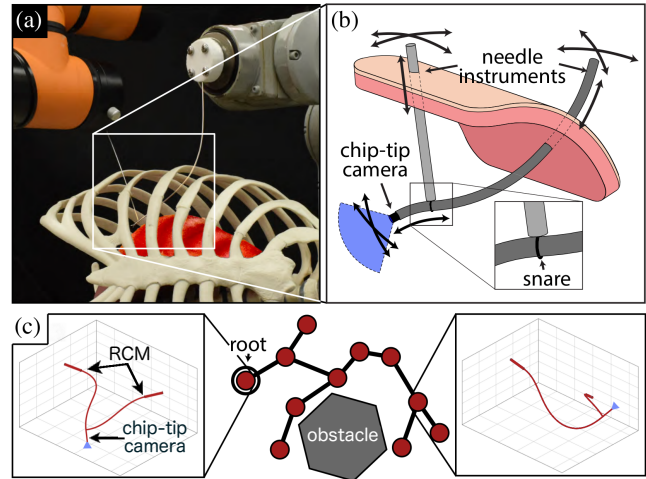
Fig. 1. An overview of the CRISP robotic system and motion planning framework. (a) The needle-diameter flexible instruments form a parallel structure inside the body whose shape is modified by actuating the instruments outside the body. (b) The instruments can be inserted and rotated to change the view of the camera at the tip. (c) The motion planner incrementally computes a tree data structure of collision-free robot configurations, which can be used to manipulate the instruments (shown in red) outside the body to reposition and reorient the camera tip while ensuring the instruments avoid anatomical obstacles inside the body.

is critical, as the underlying cause can be deadly and may have drastically different treatment paths. Thoracoscopy is the gold standard and involves insertion of endoscopic tools through the ribs [5]. Endoscopic tools give clinicians direct visualization of the pleural space. However, thoracoscopy is invasive: it requires incisions, and major complications are reported to be as high as 15% [6]. A CRISP robot, with a chip-tip camera deployed through the tool working channel, has the potential to combine the minimal invasiveness of needles with the ability of endoscopic tools to systematically inspect the interior surface of a patient's pleural effusion.

We introduce a motion planner for CRISP robots that manipulates the flexible instruments outside the body such that the tool tip camera can see a user-specified site of clinical interest inside the body. The motion planner computes motions such that all the flexible instruments inside the body avoid collision with anatomical obstacles, including the chest wall, the lung surface, and potential connections between the lung and chest wall. The configuration of the CRISP robot is the position and orientation of the instruments outside the body, and the motion planner computes a sequence of configurations that avoids collisions with

anatomical obstacles inside the body, enforces remote-center-of-motion constraints on the tube's entry points into the body, and enables visibility of the desired clinical site inside the body with the tool's camera. Motion planning for CRISP robots is challenging because evaluating their kinematics for each configuration requires modeling the elastic and torsional interactions of the robot's constituent tubes, which is computationally expensive. We introduce a sampling-based motion planner that efficiently propagates presolved state information for the kinematic model through a tree data structure in configuration space to accelerate motion plan computation.

The set of sites that can be inspected via the camera at the tool tip of a CRISP robot is heavily influenced by the robot's setup, i.e., where on the skin surface the tubes enter the body and where the snares grasp the tool instrument. We demonstrate how our motion planner can be used to both estimate the set of points on the pleural effusion surface that can be seen by the tool tip camera of a CRISP robot for a given setup, as well as provide collision free motion plans for the robot to view the points on the pleural effusion surface. This analysis can provide physicians with insights into CRISP robot setups that are appropriate for specific clinical tasks that require pointing the camera at specific sites in the pleural effusion.

We demonstrate the speed and effectiveness of our new motion planner for CRISP robots in simulation using a pleural effusion segmented from a patient CT scan. We demonstrate both the method's ability to plan motions for the robot to view specific clinically relevant sites as well as the ability to estimate the set of points that can be seen by the CRISP robot's tool tip camera.

II. RELATED WORK

Motion planning for the CRISP robot [1] is influenced by the way the shape of the robot is calculated. This influence is not unique to the CRISP robot and is a consideration present in motion planning for other continuum surgical robots.

One continuum surgical robot with related mechanics is the concentric tube robot [7]. Concentric tube robots are needle-like surgical manipulators composed of thin, nested, pre-curved nitinol tubes. Similar to the mechanics of the CRISP robot, the tubes of the concentric tube robot elastically interact in different configurations to influence the robot's shape. Using various control methods, prior work has achieved position control of concentric tube robot tips [8], [9], [10]. Sampling based motion planning has also been used to control concentric tube robots. Torres et. al. use a combination of a precomputed roadmap and an inverse kinematics controller to achieve interactive rate planning for concentric tube robots [11]. Lyons et. al. apply optimization-based motion planning using a simplified kinematics model [12].

Another related medical robotic device for interventional medical procedures is steerable needles, which are composed of a highly flexible tube and employ an asymmetric tip to steer through soft tissue [13]. Motion planning for steerable needles has been achieved in a variety of ways [14], [15], [16], including sampling-based motion planning [17], [18], [19].

Automatically controlling the motion of cameras to view specific sites can be challenging, as discussed by Christie et al. [22]. Rosell et al. plan motions of a virtual bronchoscope to view lesions in the lung through the airway [20] but are restricted in their motion to the structure of the bronchial tree. In non-medical applications, probabilistic roadmaps have been applied to plan camera paths in virtual environments when given a specified goal position and orientation for the camera [21]. There has also been work in computer vision on how to plan new viewpoints for a camera such that object recognition is optimized [23], [24]. These works primarily consider how to plan the viewing angles of the camera, while we primarily focus on motion planning for a medical robot that contains a camera for purposes of viewing specific sites in cluttered and constrained spaces.

III. PROBLEM FORMULATION

A. CRISP Robot

We consider a CRISP robot composed of N needle-diameter tubes. One of these tubes has a chip-tip camera affixed to its tip which we refer to as the camera tube, T_c. The remaining $N-1$ tubes are deployed with snares, and will be referred to as snare tubes, T_s^k, where k is an integer uniquely identifying a specific snare tube. In order to perform accurate mechanical modeling, we require as input each tube's inner diameter (ID) and outer diameter (OD). We also require a description of the chip-tip camera, in the form of its angular field of view, θ_v.

A CRISP robot's set of tubes can be assembled into parallel structures inside the patient's body in an infinite number of ways. We require as input a description of the CRISP robot state (illustrated in Fig. 2). We make a distinction between the CRISP robot's *setup* state and the robot's *actuatable* state. We define the CRISP robot's *setup* state as:

$$\{\mathbf{r}_c, \mathbf{r}_1, \ldots, \mathbf{r}_{N-1}, s_1, \ldots, s_{N-1}\}, \tag{1}$$

where $\mathbf{r} \in \mathbb{R}^3$ denotes the tubes' entry points into the patient's body, expressed in a global coordinate system, and the scalars s_k denote the arc length along the camera tube at which the k^{th} snare tube attaches to the camera tube. The subscript c denotes a value corresponding to the camera tube, and an integer k denotes the value corresponding to the k^{th} snare tube. This *setup* state is set prior to the surgical procedure and is not varied during motion execution. The entry points $\mathbf{r}$ remain fixed as remote center of motion (RCM) constraints, preventing the system's tubes from pulling laterally on the patient's body during maneuvers, and the snare grasping locations s_k remain fixed throughout the surgical procedure. We then define the CRISP robot's *actuatable* state as:

$$\{R_c, R_1, \ldots, R_{N-1}, \ell_c, \ell_1, \ldots, \ell_{N-1}\}, \tag{2}$$

where $R \in SO(3)$ denotes the tubes' orientations at their entry points as expressed in a global coordinate system, and

Fig. 2. A stiff outer sheath introduces the tubes into the body entry points $\mathbf{r}_c$, $\mathbf{r}_1$, and $\mathbf{r}_2$. The snares grasp the camera tube at arc lengths s_1 and s_2. A mechanics-based model predicts the states of the camera tube $\mathbf{x}_c$ and the snare tubes, $\mathbf{x}_1$ and $\mathbf{x}_2$, in arc length.

the scalars ℓ describe how far each of the tubes extend beyond their respective entry points into the body. This *actuatable* state represents the robot's state which will be varied during the execution of motions during the surgical procedure. The CRISP robot's total state then becomes the union of the setup and actuatable states.

B. Motion Planning

We consider the problem of planning motions for a CRISP robot. To produce motion, each tube can be actuated via changing its orientation at the entry point into the pleural effusion space and translating the tube into and out of the space.

We define an instance of the robot's actuatable state as a configuration $\mathbf{q} = \{R_c, R_1, \ldots R_{N-1}, l_c, l_1, \ldots, l_{N-1}\}$. The space of all configurations the robot can assume is then $Q \in \mathcal{SO}(3)^N \times \mathbb{R}^N$.

For a given configuration $\mathbf{q} \in Q$ we define the shape of the CRISP robot as a function

$$\mathbf{P}(\mathbf{q}, T, s) : \mathcal{SO}(3)^N \times \mathbb{R}^N \times \mathcal{N} \times \mathbb{R} \mapsto \mathbb{R}^3.$$

Function $\mathbf{P}$ is a 3D space curve representing the backbone of tube T, at arc length s in the domain $[0, l_{T\max}]$. Function $\mathbf{P}$, combined with knowledge of the cross-sectional OD of each tube allows us to calculate the shape of the entire CRISP robot. Also note, that as a special case, $\mathbf{p}_{\mathrm{camera}}$ is the 3D location of the camera on the tip of T_c, and it has a direction of view defined by the vector $\mathbf{v}_{\mathrm{camera}}$ which is tangent to the space curve at $\mathbf{p}_{\mathrm{camera}}$.

We then define a motion plan $\Pi = (\mathbf{q}_0, \mathbf{q}_1, \ldots, \mathbf{q}_n)$ as an ordered sequence of robot configurations. We define a collision free plan as a plan for which the shape of the robot at every configuration in the plan does not collide with obstacles in the environment, and an interpolation between adjacent configurations does not collide with obstacles in the environment. We then define a valid plan as a plan that is collision free and achievable given the robot hardware.

When computing a motion plan, our method takes as input a CRISP setup, an initial configuration $\mathbf{q}_0$, a Computed Tomography (CT) scan from which we will define the environment, and a goal point $\mathbf{p}_{\mathrm{goal}}$, the location in the anatomy that the physician is attempting to view with the

camera. Our method then produces as output a plan Π, which is a collision-free sequence of configurations that will result in the robot being able to view $\mathbf{p}_{\mathrm{goal}}$.

When evaluating the quality of a candidate setup, we require as input the setup, an initial configuration $\mathbf{q}_0$, and the CT scan. Then, instead of outputting a plan to a specific goal point, we instead output a set of cells on the interior surface of the pleural effusion which can be seen by manipulation of the CRISP robot with that specific setup, which we define as a visibility set.

IV. METHOD

A. Motion Planning

To model the environment for our motion planner, we use an occupancy grid. The grid has free cells, S_{free}, in which the robot is allowed to freely move, and occupied cells, S_{obs}, which the motion planner treats as an obstacle and restricts the robot from moving into. To generate these sets, we segment the pleural effusion from the CT scan using a semi-automatic region-growing method [25]. We set S_{free} to be the cells in the CT scan consistent with the pleural effusion and S_{obs} to be the inverse segmentation. We also define a third set, S_{bound} to be the cells in S_{obs} which are adjacent to S_{free} which will contain the goal point of interest $\mathbf{p}_{\mathrm{goal}}$ and which will be the set we attempt to visualize in the evaluation of a setup.

We solve the motion planning problem formulated in Sec. III-B using a sampling-based approach. We implement a planner based on the Rapidly-Exploring Random Trees (RRT) method [26]. The RRT method begins at a root node, the initial configuration, and iteratively and randomly constructs a tree structure where each node in the tree is a valid configuration, and an edge linking two nodes is a valid, collision free motion between them. As the tree grows, it expands and explores the obstacle free configuration space of the robot. Once a node is found which has a camera pose with a clear view of the goal point, $\mathbf{p}_{\mathrm{goal}}$ in the lung, the tree can be traced back to the root node, and a valid plan from the initial configuration to a goal configuration has been found.

Specifically, we begin with our root node. We then sample a randomly generated point in configuration space. We linearly interpolate between the two configurations, using spherical linear interpolation (SLERP) to interpolate between the rotational degrees of freedom. We then propagate along the line segment starting at the root node for a random percentage of the line segment, computing the shape of the robot and checking it against S_{obs} at a fine discretization. If the robot collides with obstacles or reaches a configuration that the forward kinematics solver is unable to solve, we stop at the prior step of the discretization. We then add the last valid configuration and edge to the tree. This process is then repeated, but the node from which to start the propagation is chosen as the nearest neighbor in the tree to the newly sampled point.

This process continues until a time limit has passed or until a configuration has an unobstructed view of $\mathbf{p}_{\mathrm{goal}}$, whichever comes first.

To perform collision detection, we need an accurate estimate of the robot's shape at a given configuration. To calculate the robot's shape at a given configuration, we use a modification of the mechanics-based model developed by Mahoney et al. [1] and described in greater detail in Sec. IV-B. Having calculated the backbone shape of each tube, and knowing each tube's radius, we are able to efficiently check for collisions between the robot's geometry and the occupancy grid. This is done by interpolating along the shape of each flexible instrument, identifying which cells the shape will occupy, and doing an index lookup into the CT-derived occupancy grid for those cells.

To identify whether at a sampled configuration $\mathbf{p}_{\text{goal}}$ is visible from the camera on the tip of T_c we implement a ray trace. First, the camera position and direction of view, $\mathbf{p}_{\text{camera}}$ and $\mathbf{v}_{\text{camera}}$ are inferred from the tip of the shape computed for T_c. The vector between $\mathbf{p}_{\text{goal}}$ and $\mathbf{p}_{\text{camera}}$ is computed, and it is compared with $\mathbf{v}_{\text{camera}}$. To identify if $\mathbf{p}_{\text{goal}}$ lies within the field of view of the camera, we examine the angle between the two vectors. If the angle is larger than $\theta_v/2$, then $\mathbf{p}_{\text{goal}}$ does not lie within the field of view of the camera. If, however, the angle is less than $\theta_v/2$, then $\mathbf{p}_{\text{goal}}$ does lie in the field of view of the camera. This is not enough, however, because there must exist line of sight between the camera and $\mathbf{p}_{\text{goal}}$—the view of $\mathbf{p}_{\text{goal}}$ may be occluded by another part of the patient anatomy. To identify if there exists clear line of site, a ray is traced from $\mathbf{p}_{\text{camera}}$ to $\mathbf{p}_{\text{goal}}$. If the ray strikes an occupied cell in S_{obs} before it reaches $\mathbf{p}_{\text{goal}}$, there is not clear line of site and the motion planning continues. However, if there exists clear line of site then the plan is traced back to the root initial configuration and is returned.

B. Mechanics & Solution Seeding

One of the most computationally intensive aspects of the method is computing the forward kinematics of the CRISP robot that determines its shape. This is done for every node in the tree, and at every finely discretized point along each edge in the tree. The forward kinematics is calculated both to ensure the configuration is collision free everywhere on the CRISP robot's body and to identify the camera pose.

The forward kinematics of the CRISP robot results from its mechanics, which were initially presented in [1]. In this paper, we assume that the flexible instruments of a CRISP robot can be physically held by robot manipulators at the point where the tubes enter the patient's body. This reduces the dimensionality of the CRISP robot's actuation space to only include orientation of each tube at the body entry point and each tubes' insertion length into the body. This assumption also simplifies the system mechanics and can be physically implemented in practice using stiff introducer sheaths through which the flexible instruments can be deployed. What follows is a summary of the simplified CRISP robot's mechanics and forward kinematics.

We model the CRISP robot using the Cosserat rod equations that govern the backbone position $\mathbf{p} \in \mathbb{R}^3$, orientation expressed as a rotation matrix $R \in SO(3)$, internal moment

$\mathbf{m} \in \mathbb{R}^3$, and internal force $\mathbf{n} \in \mathbb{R}^3$ of each tube. The full system state $\mathbf{x}$ consists of the Cosserat-rod states for each tube packed as:

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_c & \mathbf{x}_1 & \dots & \mathbf{x}_{N-1} \end{bmatrix} \tag{3}$$

where $\mathbf{x}_c$ are the Cosserat-rod states of the camera tube (the tube grasped by the others), and $\mathbf{x}_k$ are the Cosserat-rod states of the k^{th} snare tube. Note that $\mathbf{x}$ is a column-vector but we express it in the form of (3) for compactness.

The forward kinematics of a multi-tube system are formulated as a *multi-point boundary value differential equation*, where the Cosserat-rod states of the snare tubes propagate along their backbone in arc length $0 \le s \le \ell_k$, as

$$\mathbf{x}'_k(s) = \begin{bmatrix} \mathbf{p}'_k & R'_k & \mathbf{m}'_k & \mathbf{n}'_k \end{bmatrix}, \tag{4}$$

and the states of the camera tube propagate along its backbone in arc length $0 \le s \le \ell_c$, as

$$\mathbf{x}'_c(s) = \begin{bmatrix} \mathbf{p}'_c & R'_c & \mathbf{m}'_c + \alpha & \mathbf{n}'_c + \beta \end{bmatrix}, \tag{5}$$

where the arc length derivative of the Cosserat-rod states as well as the terms α and β can be found in [1]. The tube lengths (ℓ_c and ℓ_k) and the initial values of the tube position ($\mathbf{p}_k(0)$ and $\mathbf{p}_c(0)$) and orientation ($R_k(0)$ and $R_c(0)$) at the body entry points are given by the corresponding entry point position and orientation from the starting state, (1) and (2). The initial values of the Cosserat-rod internal moments ($\mathbf{m}_k(0)$ and $\mathbf{m}_c(0)$) and forces ($\mathbf{n}_k(0)$ and $\mathbf{n}_k(0)$) for both the snare and camera tubes are determined later to satisfy the constraints of the multi-point boundary value problem.

The constraints of the multi-point boundary value problem include a constraint at each of the grasp points s_k on the camera tube's body. The grasp constraints enforce the tip position of the snare tube to be coincident with the camera tube's position at arc length s_k and the tip pose of the snare tube is constrained so that there is a constant rigid body rotation that maps the snare tip orientation to the backbone pose of the camera tube's orientation at arc length s_k as

$$\mathbf{c}_k = \begin{bmatrix} \mathbf{p}_k(\ell_k) - \mathbf{p}_c(s_k) \\ R_c^T(s_k)R_k(\ell_k)R_{\mathbf{x}} - I \end{bmatrix} = \mathbf{0} \tag{6}$$

where $R_{\mathbf{x}} \in SO(3)$ is the rotation in the $[-1, 0, 0]^T$ direction by $90°$ and $I \in \mathbb{R}^{3 \times 3}$ is the identity matrix.

Under the assumption that the system is quasistatic and in the absence of applied forces and moments at the camera tube's tip, the force and moment at the camera tube's tip will be zero, leading to the additional constraint of

$$\mathbf{c}_c(\ell_c) = \begin{bmatrix} \mathbf{m}_c(\ell_c) \\ \mathbf{n}_c(\ell_c) \end{bmatrix} = \mathbf{0}. \tag{7}$$

The $N - 1$ grasp constraints $\mathbf{c}_i$ and the tip constraint $\mathbf{c}_c$ can be packed into the total constraint vector

$$\mathbf{c} = \begin{bmatrix} \mathbf{c}_c & \mathbf{c}_1 & \dots & \mathbf{c}_{N-1} \end{bmatrix} = \mathbf{0}. \tag{8}$$

The multi-point boundary value problem is solved by varying the snare and camera tube's initial conditions of their moments ($\mathbf{m}_k(0)$ and $\mathbf{m}_c(0)$) and forces ($\mathbf{n}_k(0)$ and $\mathbf{n}_c(0)$) so

that the total constraint equation (8) is satisfied. When solved, the multi-point boundary value problem yields the system's forward kinematics. We accomplish this using a numerical optimization routine known as a "shooting" method, where an initial seed of the snare and camera tube's initial internal moment are iteratively perturbed to minimize $\|\mathbf{c}\|$ [1].

The runtime of the forward kinematics computation is heavily dependent on the initial conditions seeded into the shooting method. If the initial conditions lie far from the true solution, not only will the shooting method converge more slowly, but it may not converge to a solution at all. However, if the initial conditions lie close to the true solution, then the shooting method will run much more quickly and the forward kinematics will be solved faster.

This insight is a significant motivation for our choice of an RRT based motion planner. RRT's incremental growth property means that we can always view tree expansion as a sequential small perturbation on an already solved forward kinematics problem. More specifically, as we expand the tree, we seed the initial conditions of each subsequent shape calculation with the true values found at the state from which it is propagating. Because each step is relatively small, we are always seeding the initial moments and forces with a solution that lies close to the true solution. We note a substantial computational speedup associated with this property compared to seeding the initial moments and forces with a generic set of initial conditions, as discussed in Sec. V-C.

C. Candidate CRISP Setup Evaluation

An adaptation of our method can be used to evaluate a specific candidate CRISP setup and initial configuration by generating a visibility set. Rather than attempting to find a plan from the initial configuration to a viewpoint for a specific $\mathbf{p}_{goal}$, one can ask the question "What is the total set of all points that can be viewed in the space, if we start at a specific $\mathbf{q}_0$ using a given candidate setup?" An answer to this question may be useful in evaluating how effective an initial configuration and setup is. To answer this question, we attempt to generate the set of all points in S_{bound} which can be viewed by the robot starting at $\mathbf{q}_0$ with the candidate setup. This can be viewed as the endoscopic equivalent of evaluating the reachable workspace of the robot.

We extend our method to construct a visibility set by allowing the tree to expand for a fixed and relatively long duration of time, while observing which cells in S_{bound} can be seen from each configuration in the tree. Rather than checking whether a specific $\mathbf{p}_{goal}$ can be seen, we instead ask what the set of cells is which is visible from the configuration. This is done in a similar fashion as above, but all cells in S_{bound} are evaluated and filtered by their relative angle to $\mathbf{v}_{camera}$. For each cell that lies in the field of view of the camera, a ray is then traced to a point $\mathbf{p}$ in the cell, and the cell in S_{bound} at which the ray terminates is added to the visibility set. The union of the sets for each configuration in the tree then becomes the total visibility set and is returned.
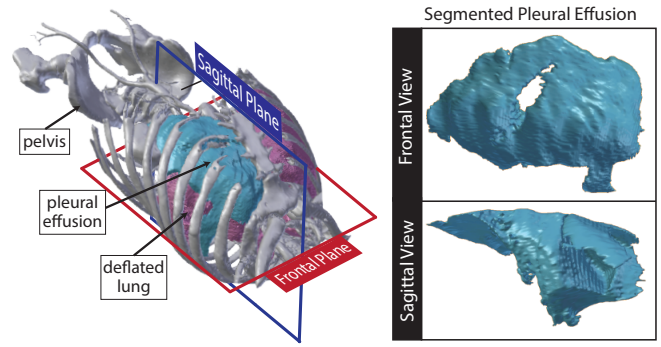

Fig. 3. An isometric view of the segmented open pleural effusion volume, bone structure, and lung is shown on the left. Frontal and sagittal views of the pleural effusion volume are shown on the right. Note that the collapsed lung lies on the posterior side of the pleural effusion.

V. RESULTS

We demonstrate the speed and effectiveness of our new motion planner for CRISP robots in a simulated scenario based on a pleural effusion segmented from a patient CT scan. We ran the motion planner on an Intel Xeon E5-1680 CPU with 8 cores running at 3.40GHz and 64GB of RAM.

A. Generating Visibility Sets

We first evaluated the ability of the method to generate visibility sets for two CRISP robot setups. Each setup included one snare tube and the camera tube, where the snare tube is affixed to the camera tube 1 cm from its tip. Both tubes have an OD of 1.02 mm and an ID of 0.84 mm. The two starting setups are shown in Fig. 4 in the pleural effusion space. The setups initially point the camera in different directions and have differing entry points and orientations into the effusion.

For each setup, we ran the motion planner for 1 hour to explore the space and recorded cells on the pleural effusion surface that could be seen by the camera. As can be seen in Fig. 5, from each setup the robot is able to visualize a large portion of the interior surface of the pleural effusion. After 1 hour, Setup 1 has visualized 38% percent of the pleural effusion surface and Setup 2 has visualized 57% percent of the pleural effusion surface. The difference in the visibility sets between the two setups implies that the ability of CRISP to visually inspect a particular set of goal points is dependent on choosing a high quality setup. The computed visibility sets can provide feedback to a physician on the usefulness of each specific setup. The union of the two visibility sets covers 80% of the total effusion surface, illustrating that using multiple setups increases the size of the visibility set and can enable the physician to see most of the pleural effusion.

B. Motion Planning to View a Specific Goal Point

We next evaluate the motion planner for computing a motion to view a specific goal point using the chip-tip camera. Using $\approx 49,500$ goal points on the surface of the pleural effusion, we show in Fig. 6 the percentage of the goal points which have been visualized as a function of time. The
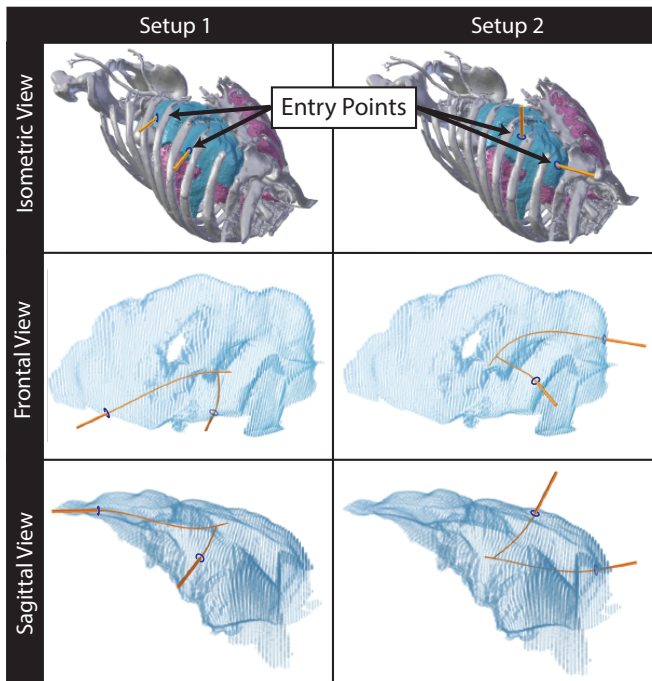
Fig. 4. The setups and initial configurations for both Setup 1 (left column) and Setup 2 (right column). The pleural effusion is rendered transparent so the initial shape of the robot can be visualized.
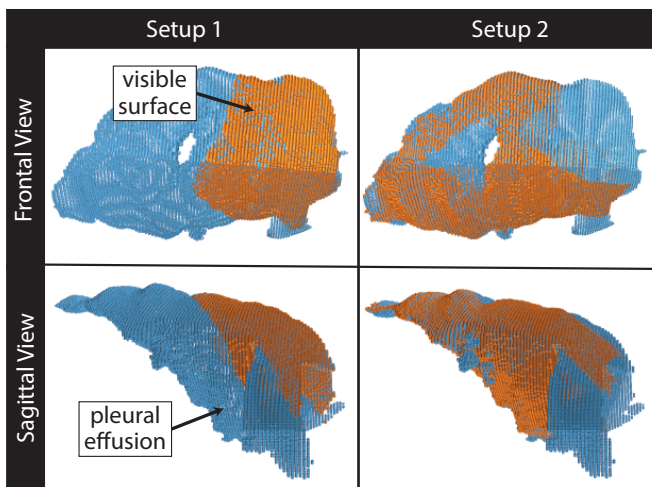

Fig. 5. The visibility sets found by the motion planner for Setup 1 (left column) and Setup 2 (right column) after 1 hour of computation. The portion of the surface visualized is rendered in orange.


Fig. 6. The percent of goal points seen by the camera as a function of the motion planner's exploration time for each setup. Solid lines show the percent of goal points seen by the camera with respect to each setup's visibility set. Dashed lines represent the percent of goal points seen with respect to the total number of points on the pleural effusion surface.


Fig. 7. A motion plan viewed from inside the pleural effusion. Potential points of interest on the interior surface of the pleural effusion are rendered as blue spheres, with the specific goal point rendered in pink. The plan goes from the initial configuration (a), through collision free intermediate configurations (b) and (c), to a configuration in which the tip of the camera tube can view the goal point in (d).

percentage of points found can be viewed as the probability that the motion planner finds a plan to visualize a goal point if it were sampled uniformly from the set of visible points found by that specific setup (solid lines) or if it were sampled from all points on the interior surface of the pleural effusion (dashed lines). An example of a motion being planned to view a specific goal point can be seen in Fig. 7.

C. Mechanics Solution Seeding

To evaluate the efficacy of seeding the forward kinematics computation with the solution of an already known near-

by shape, we repeated the experiments but without the the seeding. Instead, we initialize each shape calculation with the solved solution of only the initial configuration, corresponding to the root node in the tree.

At the end of the hour-long planning time, the motion planner for Setup 1 without the mechanics solution seeding had only planned visualizations for 30% of the total cells in S_{bounds} compared to the 38% found by the planner with the mechanics solution seeding. For Setup 2, the planner without the seeding had only computed motion plans to visualize 38% of the cells compared to the 57% found by the motion planner with the mechanics solution seeding. A stark difference was also found in the number of configurations being added to the tree. For example, from Setup 1, after 1

hour the motion planner without the seeding had only added 2,177 configurations to the tree, while the motion planner with seeding had added 37,855 configurations. For Setup 2, the motion planner without the seeding had added 2,627 configurations, while the motion planner with the seeding had added 36,610 configurations. The ≈ 14 times increase in the number of configurations when using the seeding suggests a significantly more expansive motion planning tree capable of viewing more user-specified points of interest.

VI. CONCLUSION

CRISP robots provide new avenues for performing minimally invasive, incisionless medical procedures. Our motion planner for CRISP robots computes manipulations of the needle-diameter flexible instruments outside the body such that the camera can visually inspect a user-specified site of clinical interest inside the body. Our sampling-based motion planner ensures avoidance of collisions with anatomical obstacles inside the body, enforces remote-center-of-motion constraints on the flexible instrument's entry points into the body, and efficiently handles the expensive computation of CRISP robot kinematics. We also extended the motion planner to estimate the set of points inside a body cavity that can be visually inspected by the camera of a CRISP robot for a given setup.

In future work, we plan to build upon this new motion planner to bring CRISP robots closer to clinical use. We plan to further accelerate the motion planner using sampling heuristics, precomputation, and parallelization. Additionally, we plan to extend the motion planner to account for uncertainty such as patient motion during the procedure. We also plan to develop methods that utilize the motion planner to automatically optimize the setup of a CRISP robot on a patient-specific basis to ensure the camera can see sites of clinical interest. We also plan to integrate our motion planner with a physical robot prototype and evaluate the performance of the integrated physical system.

REFERENCES

[1] A. W. Mahoney, P. L. Anderson, P. J. Swaney, F. Maldonaldo, and R. J. Webster III, "Reconfigurable parallel continuum robots for incisionless surgery," in *IEEE Int. Conf. Intelligent Robots and Systems (IROS)*, 2016, pp. 4330–4336.

[2] J. Burgner-Kahrs, D. C. Rucker, and H. Choset, "Continuum robots for medical applications: a survey," *IEEE Trans. Robotics*, vol. 31, no. 6, pp. 1261–1280, 2015.

[3] G. Chirikjian, "Conformational modeling of continuum structures in robotics and structural biology: a review," *Adv. Robot.*, vol. 29, no. 13, pp. 817–829, 2015.

[4] R. W. Light, *Pleural diseases*. Lippincott Williams & Wilkins, 2007.

[5] M. Noppen, "The utility of thoracoscopy in the diagnosis and management of pleural disease," in *Seminars in respiratory and critical care medicine*, vol. 31, no. 06. Thieme Medical Publishers, 2010, pp. 751–759.

[6] R. J. Harris, M. S. Kavuru, A. C. Mehta, S. V. Medendorp, H. P. Wiedemann, T. J. Kirby, and T. W. Bice, "The impact of thoracoscopy on the management of pleural disease," *Chest*, vol. 107, no. 3, pp. 845–852, 1995.

[7] A. W. Mahoney, H. B. Gilbert, and R. J. Webster III, "A review of concentric tube robots: modeling, control, design, planning, and sensing," *Encyclopedia of Medical Robotics, Minimally Invasive Surgical Robotics*, 2016.

[8] P. E. Dupont, J. Lock, B. Itkowitz, and E. Butler, "Design and control of concentric-tube robots," *IEEE Trans. Robotics*, vol. 26, no. 2, pp. 209–225, Apr. 2010.

[9] D. C. Rucker, "The mechanics of continuum robots: model-based sensing and control," Ph.D. dissertation, Vanderbilt University, 2011.

[10] R. Xu, A. Asadian, A. S. Naidu, and R. V. Patel, "Position control of concentric-tube continuum robots using a modified Jacobian-based approach," in *IEEE Int. Conf. Robotics and Automation (ICRA)*, May 2013, pp. 5793–5798.

[11] L. G. Torres, A. Kuntz, H. B. Gilbert, P. J. Swaney, R. J. Hendrick, R. J. Webster III, and R. Alterovitz, "A motion planning approach to automatic obstacle avoidance during concentric tube robot teleoperation," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, May 2015, pp. 2361–2367.

[12] L. A. Lyons, R. J. Webster III, and R. Alterovitz, "Motion planning for active cannulas," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, Oct. 2009, pp. 801–806.

[13] N. J. Cowan, K. Goldberg, G. S. Chirikjian, G. Fichtinger, R. Alterovitz, K. B. Reed, V. Kallem, W. Park, S. Misra, and A. M. Okamura, "Robotic needle steering: design, modeling, planning, and image guidance," in *Surgical Robotics: System Applications and Visions*, J. Rosen, B. Hannaford, and R. M. Satava, Eds. Springer, 2011, ch. 23, pp. 557–582.

[14] K. Hauser, R. Alterovitz, N. Chentanez, A. M. Okamura, and K. Goldberg, "Feedback control for steering needles through 3D deformable tissue using helical paths," in *Robotics: Science and Systems (RSS)*, Jun. 2009.

[15] M. C. Bernardes, B. V. Adorno, P. Poignet, and G. A. Borges, "Robot-assisted automatic insertion of steerable needles with closed-loop imaging feedback and intraoperative trajectory replanning," *Mechatronics*, vol. 23, pp. 630–645, Jul. 2013.

[16] K. M. Seiler, S. P. Singh, S. Sukkarieh, and H. Durrant-Whyte, "Using Lie group symmetries for fast corrective motion planning," *Int. J. Robotics Research*, vol. 31, no. 2, pp. 151–166, Dec. 2011.

[17] S. Patil, J. Burgner, R. J. Webster III, and R. Alterovitz, "Needle steering in 3-D via rapid replanning," *IEEE Trans. Robotics*, vol. 30, no. 4, pp. 853–864, Aug. 2014.

[18] W. Sun, S. Patil, and R. Alterovitz, "High-frequency replanning under uncertainty using parallel sampling-based motion planning," *IEEE Trans. Robotics*, vol. 31, no. 1, pp. 104–116, Feb. 2015.

[19] A. Kuntz, L. G. Torres, R. H. Feins, R. J. Webster III, and R. Alterovitz, "Motion planning for a three-stage multilumen transoral lung access system," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, Sep. 2015, pp. 3255–3261.

[20] J. Rosell, A. Perez, P. Cabras, and A. Rosell, "Motion planning for the virtual bronchoscopy," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, May 2012, pp. 2932–2937.

[21] D. Nieuwenhuisen and M. H. Overmars, "Motion planning for camera movements," in *IEEE Int. Conf. Robotics and Automation (ICRA)*, 2004, pp. 3870–3876.

[22] M. Christie, P. Olivier, and J.-M. Normand, "Camera control in computer graphics," *Computer Graphics Forum*, vol. 27, no. 8, pp. 2197–2218, 2008.

[23] A. Doumanoglou, R. Kouskouridas, S. Malassiotis, and T.-K. Kim, "Recovering 6D object pose and predicting next-best-view in the crowd," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 3583–3592.

[24] P. Ammirato, P. Poirson, E. Park, J. Kosecka, and A. C. Berg, "A dataset for developing and benchmarking active vision," in *IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 1378–1385.

[25] H. J. Johnson, M. McCormick, L. Ibáñez, and Insight Software Consortium, "The ITK Software Guide," Available: http://www.itk.org/ItkSoftwareGuide.pdf, Dec. 2013.

[26] S. M. LaValle and J. J. Kuffner, "Rapidly-exploring random trees: Progress and prospects," in *Algorithmic and Computational Robotics: New Directions*, B. R. Donald and Others, Eds. Natick, MA: AK Peters, 2001, pp. 293–308.