2018 IEEE International Conference on Robotics and Automation (ICRA)
May 21-25, 2018, Brisbane, Australia

Kinematic Design Optimization of a Parallel Surgical Robot to Maximize Anatomical Visibility via Motion Planning

Alan Kuntz[1], Chris Bowen[1], Cenk Baykal[2], Arthur W. Mahoney[3], Patrick L. Anderson[3], Fabien Maldonado[4], Robert J. Webster III[3], and Ron Alterovitz[1]

***Abstract*— We introduce a method to optimize on a patient-specific basis the kinematic design of the Continuum Reconfigurable Incisionless Surgical Parallel (CRISP) robot, a needle-diameter medical robot based on a parallel structure that is capable of performing minimally invasive procedures. Our objective is to maximize the ability of the robot's tip camera to view tissue surfaces in constrained spaces. The kinematic design of the CRISP robot, which greatly influences its ability to perform a task, includes parameters that are fixed before the procedure begins, such as entry points into the body and parallel structure connection points. We combine a global stochastic optimization algorithm, Adaptive Simulated Annealing (ASA), with a motion planner designed specifically for the CRISP robot. ASA facilitates exploration of the robot's design space while the motion planner enables evaluation of candidate designs based on their ability to successfully view target regions on a tissue surface. By leveraging motion planning, we ensure that the evaluation of a design only considers motions which do not collide with the patient's anatomy. We analytically show that the method asymptotically converges to a globally optimal solution and demonstrate our algorithm's ability to optimize kinematic designs of the CRISP robot on a patient-specific basis.**

I. Introduction

A robot's kinematic design, i.e., physical parameters fixed prior to the robot's use that affect a robot's kinematics, can greatly impact the robot's ability to perform a task. The quality of a specific kinematic design can vary as the robot's task and environment changes. In medical robotics, the quality of a robot's kinematic design is influenced by the specific surgical or interventional procedure it will perform as well as the specific patient anatomy in which it will operate. A suboptimal kinematic design may negatively impact patient outcomes. In this work, we introduce a method to optimize on a patient-specific basis the kinematic design of the Continuum Reconfigurable Incisionless Surgical Parallel (CRISP) robot [1], [2] to maximize the ability of the robot's tip camera to view tissues in constrained spaces.

The CRISP robot is a minimally invasive surgical robot composed of needle-diameter tubes which are inserted into the patient and assembled into a parallel structure (see Fig. 1). This assembly is performed using snares, which the tubes use to grip one another. In addition to snares, other tools can be passed through or placed at the end of the tubes, such as chip-tip cameras, biopsy needles, ablation probes, etc. As the tubes are then manipulated outside the body by robotic manipulators, the shape is changed inside the body to perform the surgical task. Due to the needle-like diameter of the tubes, the CRISP robot has the potential to further reduce invasiveness of minimally invasive surgical procedures such as abdominal, fetal, and neonatal surgery [1], [3]. When a chip-tip camera is mounted at the tip of one of the tubes, the robot can operate as an endoscope, enabling the physician to view the interior surface of an anatomical cavity, such

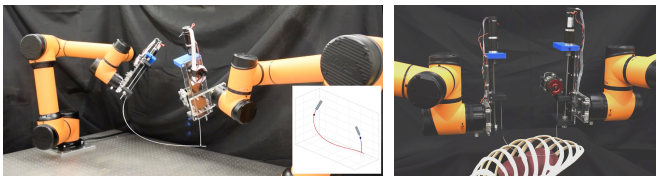

Fig. 1. Two images of the CRISP robot, where 2 manipulators adjust the base poses of needle-diameter tubes that are inserted into the body and assembled into a parallel structure.

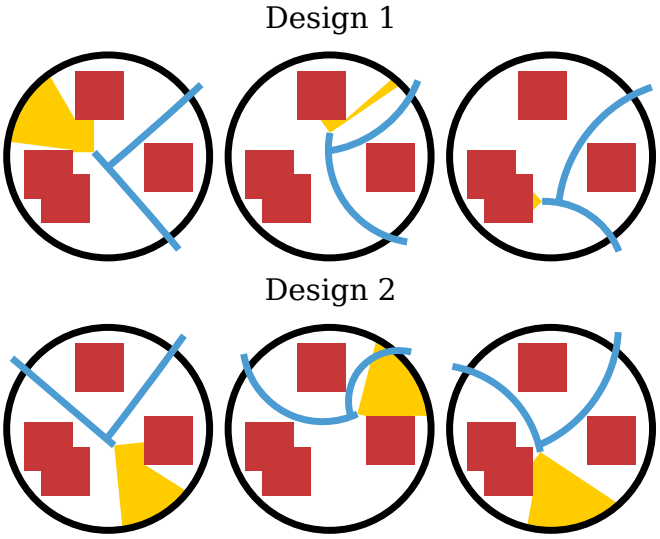

Fig. 2. Different designs can greatly impact the robot's ability to view the interior surface of a volume via its tip camera. Here we show 2D examples of a robot (blue) and its field of view (yellow) after being occluded by obstacles (red). In Design 1 (top), the robot cannot view much of the interior surface of the volume (the black circle) due to obstacles as it is manipulated through 3 collision-free configurations. However a different design (bottom), which has different entry points into the volume and a different snare grasping location, can view a much larger percentage of the interior surface.

This research was supported in part by the U.S. National Science Foundation (NSF) under Award IIS-1149965, a Vanderbilt Institute for Surgery and Engineering Seed Grant, and a Vanderbilt Discovery Grant.

[1]Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA. {adkuntz,cbbowen,ron}@cs.unc.edu

[2]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

[3]Department of Mechanical Engineering, Vanderbilt University, Nashville, TN 37235, USA.

[4]Vanderbilt University Medical Center, Nashville, TN 37235, USA.

as the abdomen or pleural space (between a lung and the chest wall) [4]. Whereas typical pleural endoscopes have diameters of 7mm and may require an incision as large as 10mm, the CRISP robot requires only needle-size (1-3mm diameter) entry points in the skin.

In this paper we investigate kinematic design optimization for the CRISP robot with a chip-tip camera to maximize the surface area of tissues visible to the physician in a constrained space. We consider as part of the optimization the anatomical entry points of each tube into the patient's body, which may be constrained by anatomical factors, as well as the snare grasping locations which define the parallel structure of the robot. Appropriately choosing these kinematic design parameters significantly impacts the robot's ability to view the anatomical sites of interest to the physician (see Fig. 2).

Our design optimization approach combines a global numerical optimization method with a motion planner to evaluate the ability of candidate kinematic designs to view the anatomical sites of interest. Specifically, we use the global optimization algorithm Adaptive Simulated Annealing (ASA) [5], [6] to generate candidate designs, which we evaluate using a sampling-based motion planner specifically designed for viewing anatomical sites using the CRISP robot [3]. By evaluating candidate designs with a motion planner, we ensure that a target is considered viewable by the robot only if viewing the target can be achieved by a robot motion that avoids collision with obstacles in the patient anatomy. This guarantees that the evaluation of the design takes into account the *motions* required by the robot to view the target regions, not just the robot's theoretical ability to do so in the absence of constrained anatomy.

Whereas other previous ASA based approaches [6] optimize a design's ability to reach goal regions, we focus on optimizing a design's ability to *view* target regions with its camera, an application which introduces unique challenges in providing theoretical guarantees. We prove that for this application the optimization is asymptotically optimal, i.e., almost surely converges to a globally optimal kinematic design. We then demonstrate the ability of the design optimization algorithm to produce high quality designs in two simulated scenarios, including a pleuroscopic scenario based on patient anatomy. We show that the algorithm generates designs which over time significantly improve the robot's ability to view target regions.

II. Related Work

Numerical optimization methods have been developed to optimize the design of various medical robots. Another continuum surgical robotic system for which design optimization methods have been developed is the concentric tube robot. Concentric tube robots are similar to the CRISP robot in that they are continuum systems with complex and computationally costly kinematics due to the elastic and torsional interactions between their tubes. The complex kinematics of these devices have motivated unique approaches to their design optimization.

Bergeles et al. provide a method that computationally optimizes concentric tube robot designs to reach a set of goal points while avoiding collisions with anatomy [7]. This method, however, does not provide a guarantee of global optimality. Burgner et al. leverage a grid-based search of the robot's configuration space with a nonlinear optimization method to optimize designs, maximizing reachable workspace while satisfying anatomical constraints [8]. Ha et al. optimize designs to maximize the concentric tube robot's elastic stability, a problem which is of particular concern for concentric tube robots [9]. Morimoto et al. take a unique approach to concentric tube robot design by providing a human with an interface to interactively design the tubes of the robot [10]. These works focus on computing goal configurations of the robot for a given design, and as such do not provide a guarantee that there exists a collision free *path*, or sequence of configurations, which allows the computed design to achieve a goal configuration.

To accomplish optimization with start to goal path guarantees, Torres et al. integrate a motion planner into the design optimization process, ensuring that valid paths exist to each reachable target for a given design [11], but this method suffers from slow performance. Baykal et al. build upon this method to compute a minimal set of designs that reach multiple targets [12]. Most recently, Baykal et al. provide asymptotic optimality in the design of piecewise cylindrical robots, which includes concentric tube robots, for reaching workspace targets [6]. In this paper, we build upon this prior work to optimize the kinematic design of the CRISP robot for anatomical visibility, rather than reachability, and show asymptotic optimality under this new optimization metric.

The optimization of surgical port placement, which has parallels to our optimization of needle entry points, has received previous study as well. Liu et al. optimize port placement and needle grasping locations for autonomous suturing [13], but they require as input a discrete set of candidate grasping and port placement locations. Hayashi et al. investigate abdominal port placement by examining the angular relationship between the port location and anatomical sites in the abdomen [14]. They do not, however consider the requirement of more complex motions during the procedure. Feng et al. optimize laparoscopic port placement for robotic assisted surgery by evaluating the robot's reachable workspace in the patient [15], but do not consider complex motions or obstacle avoidance during the procedure.

Design optimization has been studied outside of the surgical robotics domain as well. For instance, discrete parameter design optimization has been studied for a variety of applications, including multi-modal robots [16], [17], jumping robots [18], modular robots [19], [20], and protein chains [21]. These methods focus on optimizing over a finite, discrete set of features, whereas in this paper we optimize over continuous design parameters.

Methods for optimizing the design of serial manipulators has received a large amount of study and include grid-based approaches [22], geometric approaches [23], interval analysis [24], and genetic algorithms [25], [26], [27], [28]. Frequently

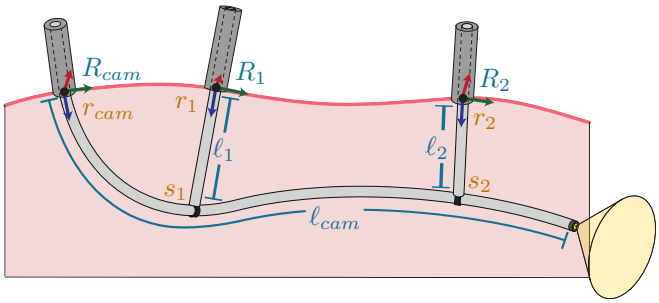

Fig. 3. The CRISP robot's kinematic design parameters (orange), which must be set before a procedure, include the entry points into the body (r_{cam} or r_k) and the snare grasping locations (s_k). The robot's configuration variables (blue), which can be continuously modified during a procedure, include the tubes' insertion lengths into the body (ℓ_{cam} or ℓ_k) and the tubes' orientations at the entry point (R_{cam} or R_k).

these methods require simplified kinematic models or restrictive assumptions to achieve computational feasibility, and typically provide no theoretical performance guarantees.

Taylor et al. present a non-linear constrained optimization approach to the simultaneous design of shape and motion for dynamic planar manipulation tasks [29]. Ha et al. explore the relationship between design and motion by leveraging the implicit function theorem to jointly optimize robot joint and motion parameters for manipulators and quadruped robots [30]. These methods, however, do not provide global guarantees and may be subject to local minima.

III. Problem Definition

We consider a CRISP robot composed of n needle-diameter tubes. One of the tubes, the camera tube ρ_{cam}, has a chip-tip camera attached to its tip. The rest of the $n-1$ tubes have snares with which they grip ρ_{cam} and are denoted ρ_k where k is a unique integer label.

A. Design Space

Let $\mathbf{r}_{\text{cam}}, \mathbf{r}_1, \ldots, \mathbf{r}_{n-1}, \in \mathbb{R}^3$ represent the points at which the tubes of the CRISP robot enter the patient's body expressed in a global coordinate system. Let the set of all valid such entry points be $\mathcal{R} \subseteq \mathbb{R}^3$. The entry points into the patient's body act as remote center of motion (RCM) points, preventing the tubes from applying lateral load to the patient's skin during the procedure, and are fixed during the operation of the robot.

The CRISP robot is assembled into a parallel structure where each snare tube ρ_k grips the camera tube at $s_k \in \mathbb{R}$, the scalar valued arc length along ρ_{cam}. The snare grasping locations are fixed during operation.

We define a kinematic design $\mathbf{d}$ of the CRISP robot as a vector describing each tube's entry point into the body and each snare tube's grasping location. The open set of all kinematic designs is $\mathcal{D} \subseteq \mathbb{R}^{n-1} \times \mathbb{R}^{3n}$ (see Fig. 3).

B. Configuration Space

Once the kinematic design $\mathbf{d}$ is fixed and the CRISP robot is assembled inside the patient, the operation of the robot may then be performed by rotating the tubes about their RCM entry points and inserting and withdrawing the tubes from the patient. We define a configuration of the robot as

$$\mathbf{q} = (R_{\text{cam}}, R_1, \ldots R_{n-1}, \ell_{\text{cam}}, \ell_1, \ldots, \ell_{n-1}),$$

where $R \in \mathcal{SO}(3)$ represents a tube's rotation at its entry point, and $\ell \in \mathbb{R}$ represents a tube's insertion length into the patient's body. The open set of all configurations then becomes $\mathcal{Q} \subseteq \mathcal{SO}(3)^n \times \mathbb{R}^n$. We define a path as a continuous function $\sigma : [0,1] \to \mathcal{Q}$, where $\sigma(0) = \mathbf{q}_0$, with $\mathbf{q}_0$ defined as the starting configuration of the robot, which may vary for different designs (see Fig. 3).

C. Workspace

We define the robot's workspace as $\mathcal{W} \subseteq \mathbb{R}^3$. We define $\text{Shape} : \mathcal{D} \times \mathcal{Q} \times \mathcal{K} \to \mathcal{W}$ as the continuous shape function of the robot, where $\mathcal{K}$ is the compact set of points representing the robot's geometry.

Given a compact obstacle set $\mathcal{O} \subseteq \mathcal{W}$, the set of points which must be avoided by the robot's geometry, we say a design-configuration pair $(\mathbf{d}, \mathbf{q}) \in \mathcal{D} \times \mathcal{Q}$ is in collision if there exists $\mathbf{k} \in \mathcal{K}$ such that $\text{Shape}(\mathbf{d}, \mathbf{q}, \mathbf{k}) \in \mathcal{O}$.

A design-configuration pair $(\mathbf{d}, \mathbf{q})$ is said to be *reachable* if there exists a path $\sigma : [0,1] \to \mathcal{Q}$ with $\sigma(0) = \mathbf{q}_0$ and $\sigma(1) = \mathbf{q}$ such that for all $0 \leq s \leq 1$, $(\mathbf{d}, \sigma(s))$ is not in collision.

Let $\text{View} : \mathcal{D} \times \mathcal{Q} \times \mathcal{V} \times [0,1] \to \mathcal{W}$ be a continuous function defining the view of the robot, where $\mathcal{V}$ represents the set of rays along which the robot can see as determined by the camera's field of view, and $[0,1]$ is the domain of a distance along the ray. A target region $T \subseteq \mathcal{W}$ is said to be *visible* from a design-configuration pair if there exists a ray in $\mathcal{V}$ that reaches any point in the target region, unoccluded by any obstacle. More formally, a target region $T \subseteq \mathcal{W}$ is visible from a design-configuration pair $(\mathbf{d}, \mathbf{q})$ if $\exists \mathbf{v} \in \mathcal{V}, r \in [0,1] : \text{View}(\mathbf{d}, \mathbf{q}, \mathbf{v}, r) \in T \wedge \forall s \in [0,r] : \text{View}(\mathbf{d}, \mathbf{q}, \mathbf{v}, s) \notin \mathcal{O}$.

D. Maximizing Viewable Anatomy

Let $\mathcal{T} = \{T_1, \ldots, T_m\}$ denote the set of m physician-specified target regions that we seek to view, where each target region denotes an open set of points in the workspace, $T_i \subseteq \mathcal{W}, \forall i \in [m]$. We define a target region as *viewable* under design $\mathbf{d}$ if the target region can be viewed from any configuration that can be reached by following a collision-free path. We then define the number of viewable target regions from a given $\mathbf{d}$ as $\Pi(\mathbf{d}) : \mathcal{D} \to [0, m]$, where

$$\Pi(\mathbf{d}) := |\text{TargetRegionsViewable}(\mathbf{d})|. \tag{1}$$

$\text{TargetRegionsViewable}(\mathbf{d})$ denotes the set of viewable target regions when using kinematic design $\mathbf{d}$. Our goal then becomes finding an optimal design $\mathbf{d}^*$ such that $\Pi(\mathbf{d})$ is maximized.

IV. Method

We present a method which optimizes the kinematic design of a CRISP robot to maximize the robot's viewable target regions in a specific anatomy while avoiding collision with obstacles, and which does so in an asymptotically

optimal way (see Sec. V). The method, detailed in Alg. 1, is based on the global stochastic optimization algorithm Adaptive Simulated Annealing (ASA) [5], [31], which iteratively samples and evaluates candidate designs. We evaluate candidate kinematic designs with a motion planner designed specifically for the CRISP robot, and which can determine the viewable target regions for a given design [3].

Algorithm 1: CRISP Robot Design Optimization

Input:
T: target regions
$\mathcal{O}$: obstacles
i_{init}: initial number of RRT iterations
i_Δ: RRT iteration increment
Output: $\mathbf{d}^*$: optimal CRISP design maximizing (1)

1 $i \leftarrow i_{\text{init}}$; $\text{Temp} \leftarrow \text{Temp}_{\text{initial}}$; $\hat{\Pi}_{\text{current}} \leftarrow 0$; $\hat{\Pi}^* \leftarrow 0$
2 $\mathbf{d}_{\text{current}}, \mathbf{q}_0 \leftarrow$ random initial design
3 $\mathbf{d}^* \leftarrow \mathbf{d}_{\text{current}}$
4 **while** time allows **do**
5 $\quad \mathbf{d}', \mathbf{q}_0' \leftarrow$ SAMPLEDESIGN$(\mathbf{d}_{\text{current}}, \text{Temp}, i)$;
6 $\quad$ targetsViewed $\leftarrow$ CRISPRRT$(\mathbf{d}', i, \mathcal{O}, \mathbf{q}_0')$;
7 $\quad \hat{\Pi}' \leftarrow |\text{targetsViewed}|$;
8 $\quad$ **if** ACCEPT$(\hat{\Pi}', \hat{\Pi}_{\text{current}}, \text{Temp})$ **then**
9 $\quad\quad \mathbf{d}_{\text{current}} \leftarrow \mathbf{d}'$;
10 $\quad\quad \hat{\Pi}_{\text{current}} \leftarrow \hat{\Pi}'$;
11 $\quad$ **end**
12 $\quad$ **if** $\hat{\Pi}' > \hat{\Pi}^*$ **then**
13 $\quad\quad \mathbf{d}^* \leftarrow \mathbf{d}'$;
14 $\quad\quad \hat{\Pi}^* \leftarrow \hat{\Pi}'$;
15 $\quad$ **end**
16 $\quad i \leftarrow i + i_\Delta$;
17 $\quad \text{Temp} \leftarrow$ UPDATETEMPERATURE(Temp);
18 **end**

A. Exploring Design Space

To explore the CRISP robot's design space, we leverage the ASA method. ASA is a global optimization algorithm, ensuring it does not become trapped in local optima. ASA works by iteratively improving upon the current best design. It first samples a candidate design some distance from the current design in design space (SAMPLEDESIGN at line 5 of Alg. 1). It then evaluates the quality of the candidate design (ACCEPT at line 8 of Alg. 1). If the candidate design is of higher quality than the current design, i.e., $\hat{\Pi}' > \hat{\Pi}_{\text{current}}$, ASA will accept the candidate design. However, with some probability ACCEPT will accept an inferior design. It is this mechanism that allows the algorithm to avoid local maxima. Both the distance at which it samples and the probability of allowing an inferior design to be accepted depend upon a "temperature" value (Temp in Alg. 1). The temperature is initially set to a large value, but is decreased over time according to a cooling schedule. In this way, early in the optimization process the algorithm is much more likely to explore more distant designs or accept inferior designs, but later in the process converges to a high quality design. Our algorithm keeps track of the best design found through the course of its execution and returns the best design at its conclusion.

B. Evaluating Candidate Designs

We evaluate the number of target regions viewable by a candidate design $\mathbf{d}$, denoted $\Pi(\mathbf{d})$, using a sampling-based motion planner for the CRISP robot [3]. This approach enables our method to estimate the total number of targets which are viewable by a given design, while only considering those configurations reachable by following collision-free paths. The motion planner, CRISPRRT, is based on the Rapidly-exploring Random Tree (RRT) algorithm [32] and has the property that the estimate of viewable target regions approaches the correct value as the number of motion planning iterations rises.

1) Sampling an Initial Configuration: The motion planner requires as input a collision-free starting configuration $\mathbf{q}_0$ from which to plan motions. In the case of the CRISP robot, $\mathbf{q}_0$ depends on $\mathbf{d}$ and is non-trivial to generate.

A key challenge with planning the motions for the CRISP robot is that evaluating the forward kinematics to solve for the shape of the robot is costly and requires an initial guess as to the forces and moments being applied to the tubes of the robot. Generally, if the state of the first configuration is known then subsequent configurations can be propagated out using the forces and moments of the previous configuration to seed that guess. This is a key insight behind the CRISP motion planner [3]. There is, however, no previous state with which to seed the initial configuration $\mathbf{q}_0$, as all subsequent configurations are grown from it, and it is unique to a given design. To solve this problem, we require $\mathbf{q}_0$ to be a load free configuration, i.e., one in which the tubes are applying no forces or moments to each other. This geometrically restricts the shape of the robot to a right triangle without any axial twist applied to the tubes. There are, however, an infinite number of right triangles that pass through the entry points defined by $\mathbf{d}$.

We address this challenge by selecting initial configurations uniformly at random from the set of possible initial configurations which satisfy the design constraints (note that the set of valid initial configurations from which to sample can be defined relatively simply geometrically). The sampled initial configuration is then collision checked against the environment. If it is found to be in collision, it is discarded and another starting configuration is sampled. This process repeats at most i times (the iteration parameter), or until a collision-free initial configuration is found, whichever comes first. If a collision-free initial configuration is not found, the design is assigned a value of 0, and the motion planner immediately returns.

2) Running the Motion Planner: If a valid initial configuration is found, the motion planner is run to determine the design's set of viewable targets. The motion planner builds a tree in configuration space with collision-free configurations as nodes and collision-free transitions between configurations

as edges. During CRISPRRT (line 6 of Alg. 1), the motion planner attempts to expand its tree i times, keeping track of the set of all target regions viewable by the robot configurations represented by nodes of the tree.

A sampling-based motion planner which runs for a *finite* duration may only return an approximation of the viewable set of targets for a given design (as noted in [6] for the case of reachability). To guarantee that the approximations increasingly approach the true value, the number of iterations for which the planner is run is increased by i_Δ at each iteration of Alg. 1 (line 16). This allows the algorithm to evaluate the quality of candidate designs with increasing accuracy as the algorithm executes.

V. ANALYSIS

In this section, we prove under mild assumptions the asymptotic optimality of the proposed method, which provides the user a guarantee that the method's solution will approach a globally optimal kinematic design as more computation time is allowed.

Specifically, we show that the method almost surely converges to a design under which the maximum number of target regions are viewable. This proof builds on ideas from [6] and adapts them to the case of maximizing viewable goal regions. Our approach reduces to showing that the set of optimal designs has non-zero measure and applying results from prior work on design optimization based on the motivating property of ASA. We prove this via open sets, which are either empty or have non-zero measure. In Lemma 1, we show openness of the set of design-configuration pairs for which the configuration is reachable under the design, a useful lemma in its own right. In Lemma 2, we show that the set of design-configuration pairs from which a given target region is visible is also open. We combine these results in Lemma 3 to show that the set of optimal designs is open. Straightforward application of prior work is then sufficient to prove asymptotic optimality in Theorem 1.

A. Preliminaries

Our proofs use the fact that the inverse images of open and closed sets under continuous functions are themselves open and closed respectively where the inverse image of $A \subseteq B$ under $f : C \to B$ (denoted $f^{-1}[A]$) is the set $\{c \in C \mid f(c) \in A\}$. In the following proofs, we will also frequently refer to the topological projection (or simply projection) from a Cartesian product of topological spaces $X \times Y$ to X. The projection of a set $Z \subseteq X \times Y$ to X is the set $\{x \in X \mid \exists y \in Y : (x, y) \in Z\}$. Two useful properties of projections enable the proofs below. First, the projection of an open set is itself open. Second, if Y is compact, then the projection $X \times Y \to X$ of a closed set is itself closed. This latter property is referred to as the Tube Lemma below.

Assumption 1 (Target Regions as Open Sets): Each target region $T_i \in \mathcal{T}$, is defined as an open set.

Assumption 2 (Continuity of Shape): $\text{Shape} : \mathcal{D} \times \mathcal{Q} \times \mathcal{K} \to \mathcal{W}$ is continuous.

Assumption 3 (Continuity of View): $\text{View} : \mathcal{D} \times \mathcal{Q} \times \mathcal{V} \times [0, 1] \to \mathcal{W}$ is continuous.

B. Sampling Optimal Designs Infinitely Often

Lemma 1: The set of reachable design-configuration pairs, denoted R, is open.

Proof: Consider a reachable design-configuration pair $(\mathbf{d}, \mathbf{q})$ for which we wish to construct a reachable neighborhood. By definition of reachable, there must exist some path $\sigma \in [0, 1] \to \mathcal{Q}$ with $\sigma(0) = \mathbf{q}_0$ and $\sigma(1) = \mathbf{q}$ such that $\forall s \in [0, 1] : \forall \mathbf{k} \in \mathcal{K} : \text{Shape}(\mathbf{d}, \sigma(s), \mathbf{k}) \in (\mathcal{W} \setminus \mathcal{O})$. Let $\sigma_{\mathbf{q}'}(s) = \sigma(s) + s \cdot (\mathbf{q}' - \mathbf{q})$. $\sigma_{q'}$ is continuous by continuity of σ and $\sigma_{\mathbf{q}'}(1) = \mathbf{q}'$ by construction, so $\sigma_{\mathbf{q}'}$ is a path to $\mathbf{q}'$. We thus have only to show that $\sigma_{\mathbf{q}'}$ is collision-free under each design $\mathbf{d}'$ for all $(\mathbf{d}', \mathbf{q}')$ in a neighborhood of $(\mathbf{d}, \mathbf{q})$.

Observe that the mapping $L : \mathcal{D} \times \mathcal{Q} \times \mathcal{K} \times [0, 1] \to \mathcal{W}$ given by

$$\mathbf{d}', \mathbf{q}', \mathbf{k}, s \mapsto \text{Shape}(\mathbf{d}', \sigma_{\mathbf{q}'}(s), \mathbf{k}) \tag{2}$$

is continuous by continuity of $\sigma_{q'}$ and Shape. We then have that $B = L^{-1}[\mathcal{O}] \subseteq \mathcal{D} \times \mathcal{Q} \times \mathcal{K} \times [0, 1]$ is closed by closedness of $\mathcal{O}$. Let C be the projection of B to $\mathcal{D} \times \mathcal{Q}$. C is thus the set of all $(\mathbf{d}', \mathbf{q}')$ for which $\sigma_{\mathbf{q}'}(s)$ is in collision under design $\mathbf{d}'$ for some $s \in [0, 1]$, and is closed by compactness of $\mathcal{K} \times [0, 1]$ and the Tube Lemma.

Let $\overline{C} = (\mathcal{D} \times \mathcal{Q}) \setminus C$. Now observe that $(\mathbf{d}, \mathbf{q}) \in \overline{C}$ because $\sigma_q = \sigma$, and σ is collision-free for design $\mathbf{d}$ by definition. But $\overline{C}$ is open, so it covers some neighborhood N of $(\mathbf{d}, \mathbf{q})$. N is thus a neighborhood of $(\mathbf{d}, \mathbf{q})$ in which $\sigma_{\mathbf{q}'}$ is collision-free under design $\mathbf{d}'$ for all $(\mathbf{d}', \mathbf{q}') \in N$. ■

Lemma 2: The set of design-configuration pairs from which target region T is visible, denoted $G(T)$, is open.

Proof: Consider $A = \text{View}^{-1}[\mathcal{O}]$, the set of all design-configuration-ray-distance tuples which yield points inside obstacles, which is closed because $\mathcal{O}$ is closed and View is continuous. We next construct $\hat{A}$, the set of all design-configuration pairs from which $\mathcal{O}$ is visible. $\hat{A}$ is the projection of A to $\mathcal{D} \times \mathcal{Q}$, and because $\mathcal{V} \times [0, 1]$ is compact and A is closed, $\hat{A}$ is also closed by the Tube Lemma.

Let $L = \{(s, r) \mid s \leq r \wedge (s, r) \in [0, 1]^2\}$ denote the set of all (s, r) such that if an obstacle were at distance s, it would occlude a target region at distance r and observe that this set is closed by construction. From A, $\hat{A}$, and L, we construct

$$B = (A \times [0, 1]) \cap (\hat{A} \times \mathcal{V} \times L),$$

which as the finite intersection of finite products of closed sets is, itself, closed. This is the set of all tuples $(\mathbf{d}, \mathbf{q}, \mathbf{v}, s, r)$ such that under a design $\mathbf{d}$ in configuration $\mathbf{q}$ an obstacle at distance s would occlude a target region at distance r along ray $\mathbf{v}$.

Let C be the image of B under the projection

$$\mathbf{d}, \mathbf{q}, \mathbf{v}, s, r \mapsto \mathbf{d}, \mathbf{q}, \mathbf{v}, r.$$

Again by the Tube Lemma, C is closed because B is closed and the interval $[0, 1]$ is compact. C is the set of all $(\mathbf{d}, \mathbf{q}, \mathbf{v}, r)$ such that under a design $\mathbf{d}$ in configuration $\mathbf{q}$ ray $\mathbf{v}$ is occluded at or before a distance r.

Consider $U = \text{View}^{-1}[T]$, the set of all design-configuration-ray-distance tuples via which T is visible,

which is open by openness of T and continuity of View. $V = U \cap \overline{C}$ is then the set of *unoccluded* design-configuration-ray-distance tuples via which target region T is visible. Furthermore, as the finite intersection of open sets, V is itself open. But $G(T)$ is simply the projection of V to $\mathcal{D} \times \mathcal{Q}$, and projections preserve openness. ■

Lemma 3: Given a set of target regions $T_1, \ldots, T_m$, the set of designs under which a maximum number of these target regions are viewable, denoted $\mathcal{D}^*$, is open.

Proof: Let $R_i = G(T_i) \cap R$, the set of reachable design-configuration pairs from which T_i is visible. By Lemma 1, R is open, and by Lemma 2, $G(T_i)$ is open, so their intersection R_i is also open. Let $\mathcal{D}_i$ denote the projection of R_i to its designs. Projection is an open mapping, so each $\mathcal{D}_i$ is open. Let m^* be the number of target regions viewable by an optimal design. Observe that the union of all m^*-wise intersections of $\{\mathcal{D}_1, \ldots, \mathcal{D}_m\}$ is the set of optimal designs. This is a finite union of finite intersections of open sets, and is thus open itself. ■

Corollary 1: $\mathcal{D}^*$ has non-zero measure.

Proof: By Lemma 3, $\mathcal{D}^*$ is open, and by definition, it contains an optimal design. Every non-empty open set has non-zero measure. ■

Corollary 2: Designs from $\mathcal{D}^*$ will be sampled and evaluated infinitely often by Alg. 1.

Proof: This result follows readily from the fact that ASA samples from non-zero measure sets infinitely often [5], [31]. ■

C. Asymptotic Optimality

We conclude the asymptotic optimality of our algorithm by invoking Corollary 2, established above, and Theorem 5 of [6] which extends directly to the objective of visibility.

Let $(\mathcal{Y})_{k \in \mathbb{N}}$ denote the sequence of random variables such that for each $k \in \mathbb{N}$, $\mathcal{Y}_k$ denotes the maximum number of viewable target regions attained over all the designs sampled in optimization iterations $1, \ldots, k$. Let m^* be the number of target regions viewable by an optimal design, as in the proof of Lemma 3.

Theorem 1 (Asymptotic Optimality): The solution generated by Alg. 1 almost surely converges to a globally optimal design $\mathbf{d}^* \in \mathcal{D}^*$, i.e.,

$$\mathbb{P}\left(\lim_{k \to \infty} \mathcal{Y}_k = m^*\right) = 1.$$

For completeness, we note that as an implementation detail of our method, the robot's initial configuration depends upon the design. The theorem above continues to hold in this context because the initial configuration may conceptually be incorporated into the robot's design space by introducing $\mathcal{D}' = \mathcal{D} \times \mathcal{Q}$ and letting $\text{Shape}'(\mathbf{d}', \mathbf{q}, \mathbf{k}) = \text{Shape}(\mathbf{d}, \mathbf{q} + \mathbf{q}_0, \mathbf{k})$ and $\text{View}'(\mathbf{d}', \mathbf{q}, \mathbf{v}, r) = \text{View}(\mathbf{d}, \mathbf{q} + \mathbf{q}_0, \mathbf{v}, r)$ where $\mathbf{d}' = (\mathbf{d}, \mathbf{q}_0)$.

VI. Results

We evaluate the performance of our algorithm, denoted ASA+MP, in two scenarios. Scenario 1 is an anatomically inspired but generic volume defined by an ellipsoid with a flattened side and with cylindrical obstacles in the interior (see Fig. 4). Scenario 2 is based on a pleuroscopic scenario using patient anatomy (see Fig. 5). In Scenario 2, the CRISP robot enters the volume of a pleural effusion, a serious medical condition which causes the collapse of a patient's lung [4]. The robot enters the effusion space between the patient's ribs and maneuvers inside the space to enable a physician to view the internal surface of the volume.

For both scenarios we consider a CRISP robot with a camera tube grasped by one snare tube. In the experiments we define the CRISP robot's set of valid entry points into the volume as a sequence of line segments, on the top surface of the volume for Scenario 1 and between the patient's ribs for Scenario 2. In both scenarios, the pink spheres shown in Figs. 4 and 5 were used as both the set of target regions to be viewed and as obstacles for the motion planner. All results were generated on a 3.40GHz Intel Xeon E5-1680 CPU with 64GB of RAM.

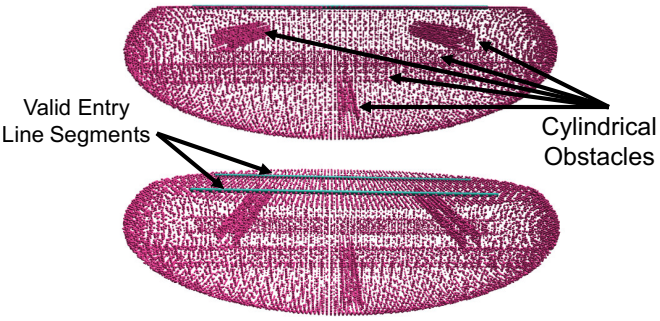

Fig. 4. Two views of the environment for Scenario 1. The environment is a flat-topped ellipsoidal volume with 5 internal cylindrical obstacles. The line segments which define the valid points of entry are shown as green lines on the top of the volume. The pink spheres (down-sampled for ease of visualization) represent the target regions for the motion planner to view under a sampled design, as well as the obstacles which must be avoided by the robot's geometry during planning.

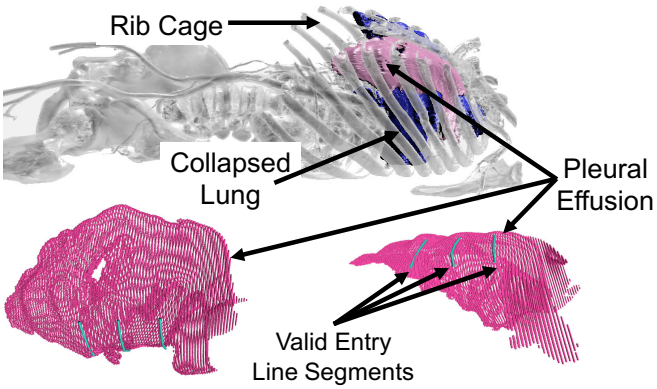

Fig. 5. Scenario 2 is a segmented pleural effusion volume from a real patient CT scan. The pleural effusion is shown in pink, displacing the collapsed lung (blue). The valid entry line segments (green) are placed such that the robot can enter the pleural space between the ribs. The target regions are rendered as pink spheres in the bottom two images, and which serve both as target regions to be viewed under a design and as obstacles for the motion planner.

We compare our method against an implementation which uses the Nelder-Mead algorithm to explore the design space, and which has been used successfully in concentric tube robot design optimization before [8]. Even when using Nelder-Mead to explore the design space, we still use the CRISP motion planner to evaluate the candidate designs, keeping in the spirit of our desire to consider only collision-free motions. We represent this as NM+MP.

For Scenario 1 we allow both algorithms to optimize designs for 16 hours and average the results over 20 different runs. For Scenario 2 we run each algorithm for 8 hours and average over 5 runs. As can be seen in Fig. 6, ASA+MP performs well in both scenarios. In Scenario 1, ASA+MP goes from designs that are able to view approximately 20% of the target regions early in the optimization to designs which are able to view approximately 64% of the target regions late in the optimization. In Scenario 2, ASA+MP goes from designs that can view approximately 14% to designs that can view approximately 46% of the target regions. ASA+MP also outperforms NM+MP in both Scenarios, finding kinematic designs that enable more target regions to be viewed in less computation time.

Fig. 7 shows a comparison between a design early in the ASA+MP optimization process and a design late in the optimization process for Scenario 2. The ability of the later design to visualize a significantly larger percentage of the target regions demonstrates the efficacy of design optimization in this scenario.

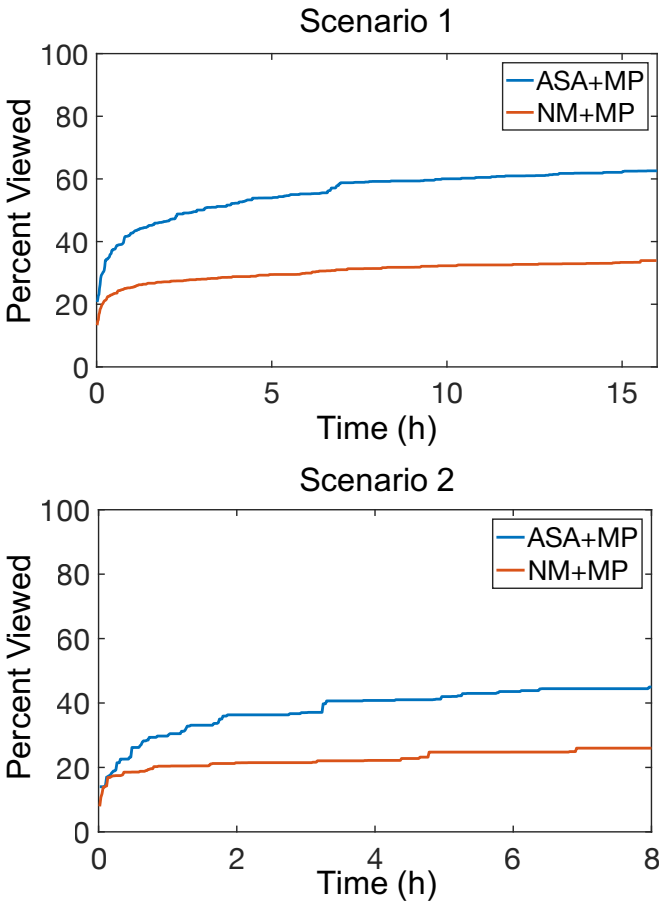

Fig. 6. The percent of target regions viewed by the best design found over time for Scenarios 1 (top) and 2 (bottom). The blue line represents the results for the ASA algorithm with motion planning, and the red line represents results for the Nelder-Mead algorithm with motion planning. The results are averaged over multiple runs, 20 for Scenario 1 and 5 for Scenario 2.

VII. Conclusion

Design optimization can have a large impact on a robot's ability to successfully perform a task. This is especially important in surgical robotics where positive patient outcomes are so vital. In this work, we demonstrated a method for optimizing the kinematic design of the CRISP robot for endoscopic purposes on a patient-specific basis. Our method leverages Adaptive Simulated Annealing (ASA) combined with sampling-based motion planning to ensure that candidate designs are evaluated in such a way that only valid motions, i.e., motions which do not cause the robot to collide with the patient's anatomy, are considered. The results show that the method is able to significantly improve the performance of the robot.

In future work, we plan to expand upon our results to help bring the CRISP robot closer to clinical use. We plan to implement and evaluate this design optimization method and the motion planner on a physical prototype of the robot. Because two designs may be near optimal but view different regions of the anatomy, we also plan to extend the work to optimize *sets* of designs to maximize the visibility of target regions, as multiple designs could potentially be employed during a single procedure. We also believe that our method and analysis can be extended for purposes of asymptotically optimal design optimization with respect to other continuous objectives and intend to pursue this idea as well. Due to the hours-long scale of the optimization, the current work would require time between imaging and the use of the robot in the clinical setting. The current time-scale would limit the use of this method to procedures for which that delay is clinically feasible, such as the pleuroscopic procedure described above. We would like to explore ways to make the optimization faster, so that it has the potential to be used in more emergent clinical procedures as well.

Acknowledgment

The authors would like to thank Nicolas E. Peckman for his valuable contributions to the visualization software.

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 *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2016, pp. 4330–4336.

[2] P. L. Anderson, A. W. Mahoney, and R. J. Webster, "Continuum reconfigurable parallel robots for surgery: Shape sensing and state estimation with uncertainty," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1617–1624, July 2017.

[3] A. Kuntz, A. W. Mahoney, N. E. Peckman, P. L. Anderson, F. Maldonado, R. J. Webster III, and R. Alterovitz, "Motion planning for continuum reconfigurable incisionless surgical parallel robots," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, Sept 2017, pp. 6463–6469.

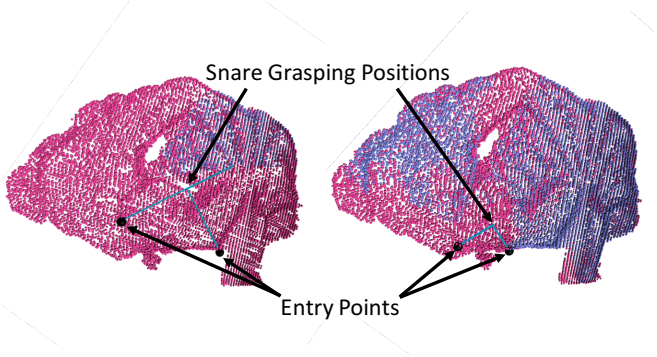

Fig. 7. A comparison between a design generated after 2 minutes of optimization (left) and a design generated after 8 hours of optimization (right). Starting configurations for each design are overlaid in light blue, the target regions viewed by the design during motion planning are shown in purple, and target regions not viewed are shown in pink. The design generated after 2 minutes is only able to view $< 7\%$ of the target regions, while the design generated after 8 hours of optimization is capable of viewing $> 50\%$ of the target regions.

[4] 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.

[5] L. Ingber, "Very fast simulated re-annealing," *Mathematical and Computer Modelling*, vol. 12, no. 8, pp. 967–973, 1989.

[6] C. Baykal and R. Alterovitz, "Asymptotically optimal design of piecewise cylindrical robots using motion planning," in *Robotics: Science and Systems (RSS)*, Jul 2017.

[7] C. Bergeles, A. H. Gosline, N. V. Vasilyev, P. Codd, P. J. del Nido, and P. E. Dupont, "Concentric tube robot design and optimization based on task and anatomical constraints," *IEEE Trans. Robotics*, vol. 31, no. 1, pp. 67–84, Feb. 2015.

[8] J. Burgner, H. B. Gilbert, and R. J. Webster III, "On the computational design of concentric tube robots: Incorporating volume-based objectives," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, May 2013, pp. 1185–1190.

[9] J. Ha, F. C. Park, and P. E. Dupont, "Achieving elastic stability of concentric tube robots through optimization of tube precurvature," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, Sept. 2014, pp. 864–870.

[10] T. K. Morimoto, J. D. Greer, M. H. Hsieh, and A. M. Okamura, "Surgeon design interface for patient-specific concentric tube robots," in *Proc. IEEE Int. Conf. Biomedical Robotics and Biomechatronics (BioRob)*, 2016, pp. 41–48.

[11] L. G. Torres, R. J. Webster III, and R. Alterovitz, "Task-oriented design of concentric tube robots using mechanics-based models," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, Oct. 2012, pp. 4449–4455.

[12] C. Baykal, L. G. Torres, and R. Alterovitz, "Optimizing design parameters for sets of concentric tube robots using sampling-based motion planning," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, Sept. 2015, pp. 4381–4387.

[13] T. Liu and M. C. Cavusoglu, "Needle grasp and entry port selection for automatic execution of suturing tasks in robotic minimally invasive surgery," *IEEE Transactions on Automation Science and Engineering*, vol. 13, no. 2, pp. 552–563, April 2016.

[14] Y. Hayashi, K. Misawa, and K. Mori, "Optimal port placement planning method for laparoscopic gastrectomy," *International Journal of Computer Assisted Radiology and Surgery*, pp. 1677–1684, Mar 2017.

[15] M. Feng, X. Jin, W. Tong, X. Guo, J. Zhao, and Y. Fu, "Pose optimization and port placement for robot-assisted minimally invasive surgery in cholecystectomy," *The International Journal of Medical Robotics and Computer Assisted Surgery*, 2017.

[16] Y. Mulgaonkar, B. Araki, J.-s. Koh, L. Guerrero-Bonilla, D. M. Aukes, A. Makineni, M. T. Tolley, D. Rus, R. J. Wood, and V. Kumar, "The flying monkey: a mesoscale robot that can run, fly, and grasp," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2016, pp. 4672–4679.

[17] C. Wright, A. Buchan, B. Brown, J. Geist, M. Schwerin, D. Rollinson, M. Tesch, and H. Choset, "Design and architecture of the unified modular snake robot," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2012, pp. 4347–4354.

[18] M. M. Plecnik, D. W. Haldane, J. K. Yim, and R. S. Fearing, "Design exploration and kinematic tuning of a power modulating jumping monopod," *Journal of Mechanisms and Robotics*, vol. 9, no. 1, p. 011009, 2017.

[19] G. Jing, T. Tosun, M. Yim, and H. Kress-Gazit, "An end-to-end system for accomplishing tasks with modular robots." in *Robotics: Science and Systems*, 2016.

[20] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, pp. 467–475, 2015.

[21] L. Denarie, K. Molloy, M. Vaisset, T. Siméon, and J. Cortés, "Combining system design and path planning," in *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.

[22] J.-Y. Park, P.-H. Chang, and J.-Y. Yang, "Task-oriented design of robot kinematics using the grid method," *Advanced robotics*, vol. 17, no. 9, pp. 879–907, 2003.

[23] R. Vijaykumar, K. J. Waldron, and M. J. Tsai, "Geometric optimization of serial chain manipulator structures for working volume and dexterity," *Int. J. Robotics Research*, vol. 5, no. 2, pp. 91–103, 1986.

[24] J.-P. Merlet, "Optimal design of robots," in *Robotics: Science and Systems (RSS)*, 2005.

[25] O. Chocron, "Evolutionary design of modular robotic arms," *Robotica*, vol. 26, no. 3, pp. 323–330, 2008.

[26] L. K. Katragadda, "A language and framework for robotic design," Ph.D. dissertation, Carnegie Mellon University, 1997.

[27] C. Leger, "Automated synthesis and optimization of robot configurations: An evolutionary approach," Ph.D. dissertation, Carnegie Mellon University, 1999.

[28] D. Salle, P. Bidaud, and G. Morel, "Optimal design of high dexterity modular MIS instrument for coronary artery bypass grafting," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, Apr. 2004, pp. 1276–1281.

[29] O. Taylor and A. Rodriguez, "Optimal shape and motion planning for dynamic planar manipulation," in *Robotics: Science and Systems (RSS)*, Jul 2017.

[30] S. Ha, S. Coros, A. Alspach, J. Kim, and K. Yamane, "Joint optimization of robot design and motion parameters using the implicit function theorem," in *Robotics: Science and Systems (RSS)*, Jul 2017.

[31] M. Locatelli, "Simulated annealing algorithms for continuous global optimization," in *Handbook of Global Optimization*. Springer, 2002, pp. 179–229.

[32] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006.