

Planning Curvature and Torsion Constrained Ribbons in 3D with Application to Intracavitary Brachytherapy

Sachin Patil *Member, IEEE*, Jia Pan *Member, IEEE*,
Pieter Abbeel *Senior Member, IEEE*, Ken Goldberg *Fellow, IEEE*

Abstract—We present an approach for planning ensembles of channels, *ribbons*, within 3D printed implants for facilitating radiation therapy treatment of cancer. The ribbons are traced out by sweeping a constant width rigid body (cuboid) along spatial curves. We propose a method for planning multiple disjoint and mutually collision-free ribbons of finite thickness along curvature and torsion constrained curves in 3D space. This is equivalent to planning motions for the cross-section of the ribbon along a spatial curve such that the cross-section is oriented along the unit binormal to the curve defined according to the *Frenet-Serret* frame. We propose a two-stage planning approach. In the first stage, a customized sampling-based planner uses rapidly exploring random trees (RRTs) to generate feasible curvature and torsion constrained ribbons. In the second stage, the curvature and torsion along each ribbon is locally optimized using sequential quadratic programming (SQP). We use this approach to design curved radiation delivery channels inside custom 3D printed implants that allow temporary insertion of a high-dose radioactive source that is threaded through the channels using a wire and allowed to dwell for specified times to expose cancerous tumors for intracavitary brachytherapy treatment. Constraints on the curvature and torsion are required for 3D printing (to allow flushing of sacrificial material) and for smooth insertion of radioactive sources. In simulation experiments, this approach achieves an improvement of 46% in tumor coverage compared with a greedy approach that generates channels sequentially.

Note to Practitioners—*Ribbons* are used in the design of efficient wiring, plumbing, transportation, architecture, and many other fields. This paper addresses the challenge of computing a set of smooth, disjoint ribbons within a constrained space to reach a specified set of target zones, motivated by the need to route multiple smooth channels through a custom 3D printed implant to deliver radiation to cancerous tumors. This paper proposes combining sampling-based motion planning with trajectory optimization to compute arrangements of ribbons. Experiments suggest that the resulting arrangements are superior to arrangements that consist of each channel in isolation.

Index Terms—Nonholonomic Motion Planning, Underactuated Systems, Trajectory Optimization, Multi-Robot Motion Planning, Intracavitary Brachytherapy

I. INTRODUCTION

Our work is motivated by applications where contiguous pathways or channels have to be routed through 3D environments that do not collide with each other or obstacles in the environment. In particular, we consider a clinical application of intracavitary brachytherapy where radioactive doses have to be delivered to cancerous tumors occurring

Sachin Patil, Pieter Abbeel, and Ken Goldberg are with the University of California, Berkeley, CA, USA, the Center for Automation and Learning for Medical Robotics (Cal-MR), and the People and Robots initiative of the UC Berkeley Center for Information Technology in the Interest of Society (CITRIS). (e-mail: {sachinpatil, pabbeel, goldberg}@berkeley.edu). Jia Pan is with the University of Hong Kong, HK, (e-mail: jpan@cs.hku.hk).






Fig. 1. **3D printed implants for intracavitary brachytherapy:** Hollow internal channels (mutually collision-free) embedded within a 3D printed implant for delivering radiation to tumors (red). The channels provide passage to radioactive sources. The candidate dose dwell segments (shown in dark blue) are aligned tangentially and placed proximal to the tumors to achieve sufficient dose distribution to the tumor volume while minimizing radiation exposure to healthy tissue. (a) Single channels computed using the approach of Garg et al. [24]. (b) Compared to single channels, multiple channels in ribbon-like arrangements can increase coverage of the tumor volume that is directly irradiated, leading to improved treatment outcomes. The challenge is to generate such mutually collision-free, curvature and torsion constrained ribbons from the candidate dwell segments to the entry region while staying within the implant volume.

in body cavities such as oral, rectal, gynecological, auditory, and nasal. Garg et al. [24] demonstrated that 3D printing can be used to design customized implants that conform to the patient anatomy. These implants allow precise positioning of radioactive sources that sufficiently irradiate the tumors while minimizing radiation to healthy tissues, which can potentially improve treatment outcomes. These implants have hollow internal channels that provide passage to radioactive sources. Garg et al. constructed implants with mutually collision-free channels with curvature constraints that provide passage to



(a) Radioactive source used for brachytherapy



(b) 3D printed implants with catheters inserted through hollow channels

Fig. 2. (a) Schematic of a typical ^{192}Ir source used in GYN Brachytherapy [10]. The source, typically 5 mm, imposes a constraint on the maximum instantaneous curvature of the channel. (b) A prototype 3D printed implant with multiple catheters inserted through the channels. The catheters have limited flexibility and have to be pushed/pulled through the channels during treatment, which imposes a constraint on the cumulative curvature and torsion (or twist) along the channel. This is also important from a fabrication perspective since the channels are printed with soluble support material that is later dissolved. Unnecessary turns and twists in the channels prevent the support material from being completely dissolved, thus resulting in blockages.

catheters carrying a radioactive source (Figs. 1(a) and 2(b)).

The effectiveness of these implants for radiation treatment depends on how effectively candidate locations for placing the radioactive seed can cover the tumor surface proximal to the implant. We refer to these locations as candidate dose dwell segments (Fig. 1(a)). The use of single channels is not sufficient for large tumors. In this work, we address this issue by generating contiguous channels within the 3D printed implant. The resulting arrangement of channels looks like a ribbon (Fig. 1(b)). These ribbon-like arrangements are compact, thus allowing for a larger number of channels to be embedded within the implant. This results in increased coverage and hence sufficient dose distributions can be applied to large tumor volumes. One could also consider alternate arrangements such as bundles. However, each channel in the bundle would have a different curvature and torsion, which is difficult from the perspective of planning to generate such arrangements.

The channels in these ribbons have to allow a radioactive source, typically 5 mm (Fig. 2(a)) in length, to pass through. This imposes a constraint on the maximum instantaneous curvature of the channels within the ribbon-like arrangement [24]. The ribbons should also be continuous and differentiable (at least C^1 -continuous) since kinks in the ribbon would not allow catheters to pass through the channels. The catheters carrying the radioactive source also have limited flexibility and have to be pushed/pulled through the channels during treatment, which imposes a constraint on the cumulative curvature and torsion (or twist) along the ribbon. In addition, each customized implant would have multiple channels that provide

access to different tumors and these channels would have to be mutually collision-free since any intersection between channels would lead to forks within the channels, potentially leading to undesirable ambiguity in the motion of a catheter when pushed through a channel.

Imposing constraints on the curvature and torsion is also important from the perspective of fabricating these implants. During the printing process, the spatial volume corresponding to channels is printed with a soluble support material that is later dissolved to create hollow channels [36]. This is a difficult problem in fabrication in general. However, unnecessary turns (curvature) and twists (torsion) in the channels further complicate this process by preventing the support material from being completely dissolved, thus resulting in blockages. Such implants are rendered unusable because the catheters cannot reach the desired dwell segments.

In this work, we consider the problem of generating such curvature and torsion constrained ribbons in 3D spaces that avoid collisions with obstacles and other ribbons. In geometry, a ribbon is a swept surface traced out by sweeping a constant width rigid body (cuboid) along a spatial curve. In our application, we consider a rigid body that describes the cross-section of the channels in a ribbon. The rigid body being swept out is oriented along the unit binormal to the curve. There are infinitely many choices of orthonormal frames [9] along a spatial curve for orienting the rigid body. In our work, we choose the *Frenet-Serret* frame which can be explicitly described in terms of the curvature and torsion along the curve [9]. However, our approach would also generalize to other choices of frames such as Bishop's frame, also known as rotation minimizing frames [9], [60].

We adopt a two stage approach for generating ribbons with the aforementioned specifications. We use a rapidly-exploring random trees (RRT) planner [34] that sequentially generates feasible curvature and torsion constrained candidate ribbons in a greedy fashion. These planners explore the configuration space by random sampling. However, the randomized nature of these planners can cause unnecessary changes in curvature and torsion along the ribbon. We locally minimize the curvature and torsion by using an optimization method based on sequential quadratic programming (SQP) [19] that simultaneously optimizes over all the ribbons. In doing so, we combine the benefits of a global exploration strategy using a randomized planner and a local optimization strategy to generate high quality ribbon trajectories.

We study the effectiveness of our approach for designing multiple ribbon-like arrangements of channels within 3D implants. We show that ribbon-like arrangements allow us to embed a larger number of channels within the implant as compared to prior work that embeds single channels [19], [24]. As a result, ribbon-like arrangements like the one shown in Fig. 1(b), lead to improved coverage of the tumor volume (46% improvement in our experiments), which allows for more effective treatment.

II. RELATED WORK

The geometry of swept surfaces and swept volumes has been extensively studied in the literature [21], [49]. A ribbon is

a particular example of a swept surface where a constant width rigid body is swept along a spatial curve [22], [28]. Ribbons find applications in a number of domains including computer aided geometric design [55], deformation of 3D shapes in solid modeling [37], modeling geometry of DNA strands [26] and polymers [12], [14], planning layout arrangements of roads [62], routing cables and wires in assembly tasks, and path planning of rigid formations of nonholonomic vehicles [3], [33]. Sampling-based planning was used to generate variable width ribbon-like paths between a camera and moving object in the environment [25]. The computed paths were smoothed to generate pleasing camera motions. The notion of grouping parallel paths, similar to ribbons, has also been studied for coverage applications in robotics such as machine milling, lawn mowing, and planning search and rescue operations [13].

Planning smooth motions of rigid bodies in 3D has also been studied [3], [8], [16], [66]. This requires planning in the space of 3D positions and orientations, also commonly referred to as the $SE(3)$ group. However, prior work has addressed the issue of generating minimum cost trajectories that inherently minimize curvature and torsion along the trajectory in environments without obstacles. Prior work on trajectory optimization on Lie groups has proposed Newton-like optimization methods [1], direct (collocation) methods for trajectory optimization for continuous time optimal control problems [47], a discrete variational formulation [32], and primitive-based motion planning [23]. However, these approaches do not address the issue of avoiding collisions with obstacles in the environment.

We also note that the motion model that describes the evolution of the ribbon surface also applies to other domains, including modeling the motion of an airplane, a roller coaster, or bevel-tipped steerable needles [61]. In the case of steerable needles, the rigid body is assumed to be a point while in the case of the continuum robots, the rigid body is assumed to have a circular cross-section. In particular, motion planning for steerable needles has been extensively studied [2], [19], [44], [59], [64]. Our approach can be used to explicitly limit the torsional rotation of the needle, thus potentially improving planning and control of steerable needles [57]. In recent years, extensions to planning curvature-constrained trajectories in 3D have been proposed for unmanned aerial vehicles (UAVs) in environments without obstacles [20], [52], and with obstacles [30], [65].

Prior work has also explored the topic of modeling and finding minimum energy configurations of linear deformable objects such as suture threads [11], [41] and elastic ribbons or strips [5], [12]. The equilibrium or minimum energy configurations of these elastic objects inherently minimize curvature and torsion along the length of the object. These methods could be used to compute a ribbon configuration that minimizes curvature and torsion. However, it is not clear if they could impose constraints on the curvature and torsion along the length of the ribbon. Furthermore, these methods do not consider generation of collision-free configurations directly, and instead generate a sequence of deformations in the object that avoid collisions with obstacles [11], [41]. In our work, the ribbons are not physically-based and the fabrication process allows us to embed these ribbons within the 3D volume. We

can thus generate ribbons by planning motions for a rigid body that describes the ribbon cross-section.

In this work, we consider a clinical application involving cancer treatment using brachytherapy. Automation science has been applied to a number of healthcare applications to better treatment quality by improving repeatability and reliability. Huang et al. [29] studied planning of robotic therapy and task-oriented functions for hand rehabilitation, Solis et al. [54] explored the use of automation for studying human motor skills for medical task training. Mendez et al. [40] studied automation of drug delivery during anesthesia, and Subburaj et al. [56] studied computer assisted joint reconstruction surgery.

There are a number of commercially available implants (also known as applicators) for treating cervical and vaginal cancers, for e.g. the Fletcher applicator [17] and the Utrecht applicator [6]. Incorrect placement of these applicators and patient movement can cause shifts in the applicator position and hence result in undesired doses. Magne et al. [38] proposed the use of a customized silicone implant created using a plaster mold of the cavity and two linear catheter are embedded within the implant.

Garg et al. [24] proposed a novel approach for designing customized 3D printed implants that use MRI/CT scans to reconstruct a precise 3D model of patient anatomy. This model is provided as input to a sampling-based RRT planner to generate individual, curvature constrained channels that can be embedded within a 3D printed implant of the same shape as the cavity. Recent advances in 3D printing are poised to have major impact on many fields including healthcare [36]. Non-toxic, FDA approved materials are allowing 3D printed parts to be used for medical applications [39] such as bone replacement and oral surgery implants. Garg et al. also proposed a metric for effectiveness of treatment in terms of *coverage* of the tumor volume. Duan et al. [19] used an optimization-based method to compute arrangements for individual channels.

In contrast to Garg et al. [24] and Duan et al. [19], we consider the problem of generating ribbon-like arrangements of contiguous channels. Also, both approaches use a *stop-and-turn* strategy for modeling the kinematics, which is inconsequential for a single channel but would induce an instantaneous torsional twist in the ribbon. We consider a different kinematic model of the evolution of the ribbon that explicitly considers torsion to model the twist along the ribbon. We also combine the benefits of sampling-based and optimization-based planning methods instead of using them in isolation and show that a combination of the two methods is important for generating high-quality ribbons.

This paper is an extended version of a conference paper presented by the authors at the WAFR 2014 [45]. In this revised version, we provide additional details about design considerations of 3D printed implants for intracavitary brachytherapy, a thorough survey of related work, and additional details about the sampling-based RRT planner, optimization formulation, and the coverage metric for evaluation. We also present additional experimental results on simulated scenarios and analyze the effect of factors such as sequential versus simultaneous planning and possible scenarios for failure.




Fig. 3. A ribbon is a swept surface traced out by sweeping a rigid body (cuboid) along a spatial curve. The ribbon-like arrangement is used for routing multiple channels. A cross-section of the ribbon is shown. At each point along the curve, the rigid body is oriented according to the *Frenet-Serret* frame, where the \mathbf{x}_t , \mathbf{y}_t , and \mathbf{z}_t axes in the local coordinate frame of the rigid body are oriented along the tangent \mathbf{t}_t (red), normal \mathbf{n}_t (green), and binormal \mathbf{b}_t (blue) vectors of the curve. We consider the problem of generating such curvature and torsion constrained ribbons through 3D space that avoid collisions with obstacles in the environment.

III. KINEMATIC MODEL OF THE RIBBON

In this section, we define the kinematic model of how a ribbon can be constructed by sweeping a rigid body corresponding to the cross-section of the ribbon along a given continuous and differentiable spatial curve in 3D space.

Let $\mathbf{p}_t = [x_t, y_t, z_t]^T \in \mathbb{R}^3$ be a 3D point on a spatial curve parameterized by parameter t . At each point \mathbf{p}_t , we can define an orthonormal frame specified by the tangent \mathbf{t}_t , normal \mathbf{n}_t , and binormal vectors \mathbf{b}_t given by [9]

$$\mathbf{t}_t = \frac{\dot{\mathbf{p}}_t}{\|\dot{\mathbf{p}}_t\|}, \quad (1a)$$

$$\mathbf{n}_t = \frac{\dot{\mathbf{t}}_t}{\|\dot{\mathbf{t}}_t\|}, \quad (1b)$$

$$\mathbf{b}_t = \mathbf{t}_t \times \mathbf{n}_t. \quad (1c)$$

Intuitively, the tangent vector \mathbf{t}_t is a unit vector tangent to the curve that points in the direction of motion, \mathbf{n}_t is a unit vector that is normal to both the tangent vector and the curve itself, and the binormal vector \mathbf{b}_t is defined as the cross product of the tangent and normal vectors.

To define the local configuration of the ribbon cross-section at a given parameter value, we need to define an orthonormal frame at the given point (Fig. 3). There are infinitely many choices for such a frame [9]. In this work, we choose the Frenet-Serret frame, the evolution of which can be described in terms of the curvature and torsion of the spatial curve, which are quantities of interest in this work. We note that our algorithmic approach also applies to other choices of the orthonormal frame such as the Bishop's frame, which is also known as the rotation minimizing frame (RMF) [9], [60].

The Frenet-Serret frame is oriented so that the \mathbf{x}_t , \mathbf{y}_t , and \mathbf{z}_t axes in the local coordinate frame of the rigid body are oriented along the tangent, normal, and binormal vectors of the curve. The pose of a rigid body oriented according to the Frenet-Serret frame at point \mathbf{p}_t can be written as a 4×4 transformation matrix $X_t \in SE(3)$ given by

$$X_t = \begin{bmatrix} R_t & \mathbf{p}_t \\ \mathbf{0} & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{t}_t & \mathbf{n}_t & \mathbf{b}_t & \mathbf{p}_t \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (2)$$

where the 3×3 rotation matrix $R_t = [\mathbf{t}_t | \mathbf{n}_t | \mathbf{b}_t] \in SO(3)$ describes the orthonormal frame in terms of \mathbf{t}_t , \mathbf{n}_t , and \mathbf{b}_t .

For a given parameter t , the Frenet-Serret frame evolves according to the following differential equations of motion as

given by [9] [28] as

$$\dot{\mathbf{t}}_t = v_t \kappa_t \mathbf{n}_t, \quad (3a)$$

$$\dot{\mathbf{n}}_t = -v_t \kappa_t \mathbf{t}_t + v_t \tau_t \mathbf{b}_t, \quad (3b)$$

$$\dot{\mathbf{b}}_t = -v_t \tau_t \mathbf{n}_t, \quad (3c)$$

$$\dot{\mathbf{p}}_t = v_t \mathbf{t}_t, \quad (3d)$$

where v_t is the speed of the spatial curve, κ_t is the curvature, τ_t is the torsion of the spatial curve for a given parameter t , and the dot operator indicates the derivative with respect to the parameter t . We note that the Frenet-Serret frame does not exist at points where the speed or curvature vanishes. Since we consider spatial curves, the speed is never zero. Hence, for the remainder of this manuscript, we will assume that the degenerate case of curvature $\kappa_t = 0$ is never encountered. During planning, we mitigate this problem by adding a small non-zero perturbation factor (0.01 in our experiments) whenever the degenerate case of $\kappa_t = 0$ is encountered.

We follow the analysis of Selig et al. [51] to rewrite these equations of motion as

$$\begin{bmatrix} \dot{\mathbf{t}}_t & \dot{\mathbf{n}}_t & \dot{\mathbf{b}}_t & \dot{\mathbf{p}}_t \\ 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} \mathbf{t}_t & \mathbf{n}_t & \mathbf{b}_t & \mathbf{p}_t \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -v_t \kappa_t & 0 & v_t \\ v_t \kappa_t & 0 & -v_t \tau_t & 0 \\ 0 & v_t \tau_t & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (4)$$

Using Eqs. (2) and (4), we get the following relation

$$\dot{X}_t = X_t \begin{bmatrix} 0 & -v_t \kappa_t & 0 & v_t \\ v_t \kappa_t & 0 & -v_t \tau_t & 0 \\ 0 & v_t \tau_t & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} = X_t U_t, \quad (5)$$

where $U_t \in \mathfrak{se}(3)$ is the velocity twist of the rigid body in its local coordinate frame [42], where $\mathfrak{se}(3)$ is the Lie algebra associated with the Lie group $SE(3)$. U_t is completely described in terms of three parameters v_t , κ_t , and τ_t . In Sec. V, we will use the motion parameters $\mathbf{u}_t = [v_t, \kappa_t, \tau_t]^T$ as the basis for generating the desired curvature and torsion constrained ribbons in 3D space.

By definition, U_t can be decomposed into the instantaneous linear \mathbf{v}_t and angular \mathbf{w}_t velocities in the local coordinate frame of the body as [42]

$$U_t = \begin{bmatrix} [\mathbf{w}_t]_{\times} & \mathbf{v}_t \\ \mathbf{0} & 0 \end{bmatrix}, \quad (6)$$




Fig. 4. Ribbons generated by sweeping a rigid body (cuboid) along a circular arc shown as a dashed line. (Left) The ribbon is oriented along the binormal vector \mathbf{b} pointing into the page. If the rigid body moves only along the tangent vector, each channel along the ribbon has a constant curvature κ , torsion, and length, which equals the curvature, torsion (zero in this case), and length of the circular arc. (Right) The ribbon lies in the plane of the page. If the rigid body moves along the normal vector \mathbf{n} , each channel along the ribbon has a different curvature, torsion value, and length, which makes the planning problem much harder.

where $\mathbf{w}_t = [v_t \tau_t, 0, v_t \kappa_t]^T$, $\mathbf{v}_t = [v_t, 0, 0]^T$, and the notation $[\cdot]_{\times}$ stands for a 3×3 skew-symmetric matrix.

Note that the above kinematic model is subject to nonholonomic constraints. Informally, this implies that since the rigid body only moves along the tangent vector \mathbf{t}_t and does not undergo any instantaneous rotation about the normal vector \mathbf{n}_t , the curvature and torsion stay constant along the width of the ribbon. Hence, to generate a curvature and torsion constrained ribbon, it suffices to plan motions of a rigid body describing the cross-section of the ribbon along a single spatial curve, such that the cross-section is oriented along the binormal vector to the curve, as shown in Fig. 4. If we consider alternate kinematic models where the rigid body might undergo motion along the normal vector \mathbf{n}_t , different channels along the ribbon would have different curvatures, torsional values, and lengths. This makes the planning problem harder because it is difficult to impose separate constraints on the curvature and torsion of individual channels in the ribbon.

When the velocity twist U_t is held constant over time interval of duration δ , the differential motion model given by Eq. (5) can be explicitly integrated as

$$X_{t+\delta} = X_t \exp(\delta U_t) \quad (7)$$

where $\exp : \mathfrak{se}(3) \rightarrow SE(3)$ is the exponential operator, for which an analytical expression exists and can be evaluated in closed-form [42].

IV. PROBLEM DEFINITION

We consider the problem of generating a set $\mathcal{R} = \{r^1, \dots, r^n\}$ of n collision-free ribbons within an implant that reach candidate dose dwell segments proximal to tumors and are curvature and torsion constrained (Fig. 1). We are provided the following inputs:

- Description of the 3D external geometry of the implant I as a triangle mesh.
- Description of the geometry of the entry region E at the base of the implant.
- Description of the geometry of external cancerous tumors $\mathcal{C} = \{C^1, \dots, C^m\}$.
- Set of n poses $\mathcal{D} = \{X^1, \dots, X^n\}$ that describe the 3D positions and orientations of groups of dose dwell segments $\{d^1, \dots, d^n\}$ corresponding to each ribbon. The poses of

the dwell segments are computed such that sufficient coverage of the tumors can be achieved.

- Set \mathcal{O} of other obstacles (voids or forbidden regions) within the implant.
- Channel width w .
- Maximum limits on speed \bar{v} , instantaneous curvature $\bar{\kappa}$, cumulative curvature $\bar{\kappa}^c$, instantaneous torsion $\bar{\tau}$, and cumulative torsion $\bar{\tau}^c$. These values are ascertained based on the flexibility of the catheters used and the dimensions of the radioactive source that is carried by the catheters.

For planning purposes, we assume that each ribbon r^i , $1 \leq i \leq n$, is discretized into a set of T^i time steps, each of constant duration δ . From here on, we assume that time is discretized and the configuration of the i^{th} ribbon at time step t is specified by the pose $X_t^i = \begin{bmatrix} R_t^i & \mathbf{p}_t^i \\ \mathbf{0} & 1 \end{bmatrix}$ of a rigid body describing the cross-section of the ribbon for $0 \leq t \leq T^i$. We further assume that the twist U_t^i , described in terms of motion parameters $\mathbf{u}_t^i = [v_t^i, \kappa_t^i, \tau_t^i]^T$ (Eq. (6)), remains constant for the duration of the step t for $0 \leq t < T^i$.

For sake of conciseness, we introduce the notation $\mathcal{X}^i = \{X_t^i : 0 \leq t \leq T^i\}$ to denote the set of all poses, and $\mathcal{U}^i = \{\mathbf{u}_t^i : 0 \leq t < T^i\}$ to denote the set of all control inputs for ribbon r^i . The entire ribbon r^i can be parameterized as $[\mathcal{X}^i, \mathcal{U}^i]$, and can be generated by integrating the constant twist between subsequent time steps.

The planning objective can be formally stated as: Generate the set $\mathcal{R} = \{r^1, \dots, r^n\}$ of ribbons, such that $\forall r^i = [\mathcal{X}^i, \mathcal{U}^i]$, the following constraints are satisfied:

- $X_0^i = X^i$: The initial pose is constrained to be the pose of the i^{th} dose target.
- $X_{T^i}^i \in E$: The cross-section of the ribbon at final time step T^i lies within the entry region to permit insertion of catheters.
- $X_{t+1}^i = X_t^i \exp(\delta U_t^i)$: The poses at consecutive time steps are related according to the kinematics model given by Eq. (7).
- $(r^i \cap I = \emptyset) \wedge (r^i \cap \mathcal{O} = \emptyset)$: Ribbon r^i does not collide with the implant boundary I and does not collide with other obstacles \mathcal{O} in the environment.
- $r^i \cap r^j = \emptyset$, $1 \leq j \leq n, j \neq i$: All ribbons are mutually collision-free.
- $(|\kappa_t^i| < \bar{\kappa}) \wedge (|\tau_t^i| < \bar{\tau})$ for $0 \leq t < T^i$: The instantaneous curvature and torsion values are within their respective bounds.
- $(\sum_{t=0}^{T^i-1} |\delta v_t^i \kappa_t^i| < \bar{\kappa}^c) \wedge (\sum_{t=0}^{T^i-1} |\delta v_t^i \tau_t^i| < \bar{\tau}^c)$: The cumulative curvature and torsion along the ribbon is respectively constrained.

In addition, for practical applications, it is desirable to minimize the cumulative curvature and torsion along the length of each ribbon. Formally, given user supplied weights α_{κ} and α_{τ} , we wish to minimize the following objective for each ribbon r^i :

$$C(\mathcal{U}^i) = \alpha_{\kappa} \sum_{t=0}^{T^i-1} (\delta v_t^i \kappa_t^i)^2 + \alpha_{\tau} \sum_{t=0}^{T^i-1} (\delta v_t^i \tau_t^i)^2, \quad (8)$$

which is equivalent to minimizing the energy or rotational strain along a curve [41].

Algorithm 1 $\mathcal{R} \leftarrow \text{ribbon_generation}(I, E, D, \mathcal{O}, w, \bar{v}, \bar{\kappa}, \bar{\kappa}^c, \bar{\tau}, \bar{\tau}^c)$

```

1:  $\mathcal{R} = \emptyset$ 
2: for all  $d^i \in \mathcal{D}$  do                                     ▷ Process dose dwell segments  $\{d^1, \dots, d^n\}$  sequentially
3:    $\mathcal{T} \leftarrow \emptyset$                                        ▷ Initialize RRT tree
4:    $\mathcal{T} \leftarrow \text{add\_vertex}(X^i)$                              ▷ Add pose  $X^i$  as root node
5:   repeat
6:      $\mathbf{p} \leftarrow \text{random\_point}(I, \mathcal{O}, \mathcal{R}, E)$                  ▷ Sample collision-free point in  $\mathbb{R}^3$ 
7:      $X \leftarrow \text{nearest\_neighbor}(\mathbf{p}, \mathcal{T})$                    ▷ Reachability-guided neighbor search [53]
8:      $\mathbf{u} \leftarrow \text{control\_sampling}(X, \mathbf{p}, \bar{v}, \bar{\kappa}, \bar{\tau})$        ▷ Sampling for best control  $[v, \kappa, \tau]^T$ 
9:      $X' \leftarrow \text{integrate\_twist}(X, \mathbf{u}, \delta)$                  ▷ Integrate constant twist using Eq. (7)
10:    if  $\text{feasible}(X', \bar{\kappa}^c, \bar{\tau}^c) \wedge$ 
         $\text{collision\_free}(X, X', I, \mathcal{O}, \mathcal{R})$  then                 ▷ Check feasibility of the edge
11:       $\mathcal{T} \leftarrow \text{add\_vertex}(X')$                              ▷ Add pose to tree
12:       $\mathcal{T} \leftarrow \text{add\_edge}(X, X')$                          ▷ Add edge to tree
13:    end if
14:    until  $(X' \in E) \vee \text{max\_iterations\_reached}$            ▷ Repeat till entry region reached
15:     $r^i \leftarrow \text{generate\_ribbon}(\mathcal{T}, X', w)$                  ▷ Generate ribbon from  $X^i$  to  $X'$ 
16:     $\mathcal{R} \leftarrow \mathcal{R} \cup r^i$                                    ▷ Add ribbon to set  $\mathcal{R}$ 
17:  end for
18: return  $\mathcal{R}$ 

```

V. APPROACH

In this section, we describe our two stage planning approach. In the first stage, we use a sampling-based rapidly-exploring random trees (RRT) planner that sequentially explores the free space in the environment to compute feasible candidate ribbons. These candidate ribbons are then simultaneously locally optimized using sequential quadratic programming (SQP) to minimize the cumulative curvature and torsion along the length of each ribbon.

A. Rapidly-exploring Random Trees (RRT)

In the first stage, we use a customized sampling-based RRT planner [34] to plan motions of the rigid body according to the nonholonomic kinematic model described in Sec. III. We sequentially generate a feasible ribbon corresponding to each dwell segment group d^i , $1 \leq i \leq n$ in \mathcal{D} using the RRT planner. We note that it is also possible to simultaneously plan for all the dose segments but in our experiments, we found that doing so failed to find feasible solutions in the kind of constrained environments considered in this work. In the second stage, however, we jointly optimize over all ribbons.

Our customized RRT algorithm is summarized in Algorithm IV. Starting from each pose X^i , we iteratively grow a tree \mathcal{T} in the pose space of 3D positions and orientations that grows towards the entry region E subject to constraints (Sec. IV). A node is iteratively added to the tree as follows. First, a point $\mathbf{p} \in \mathbb{R}^3$ is randomly sampled from the implant volume such that it does not collide with obstacles \mathcal{O} or other ribbons \mathcal{R} . As opposed to directly sampling a pose, this strategy has been shown to work better for nonholonomic systems such as the ribbon and bevel-tipped steerable needles [44]. We also bias the sampling towards the entry region to ensure progress [34]. In our experiments, 5% of all random samples were drawn from the entry region to ensure progress.

Given \mathbf{p} , we search for a node X in the tree that is nearest to \mathbf{p} according to a reachability-guided distance measure, which accounts for the nonholonomic constraints on the kinematic motion model of the ribbon. This strategy has been shown to work well in practice [44], [53]. The ribbon has a maximum instantaneous curvature $\bar{\kappa}$ and not all points will be reachable from a given pose. The reachable set from a pose $X = \begin{bmatrix} R^X & \mathbf{p}^X \\ \mathbf{0} & 1 \end{bmatrix}$ consists of all points that can be connected to \mathbf{p}^X by a circular arc that has a radius $r \geq 1/\bar{\kappa}$ and is tangent to the tangent \mathbf{t}^X of the local coordinate frame. We then define the distance metric as the length of such a circular arc connecting \mathbf{p} and X if \mathbf{p} is in the reachable set of X , and infinity otherwise. This strategy restricts the search domain to only those nodes that are within the reachable set of the nearest node (X), thus increasing the likelihood of coverage of the state space. Since the rigid body motion is only constrained along the tangent vector, the reachability-guided distance measure only depends on the speed v , and curvature κ , of the curve (Eq. 3a), and is not dependent on the torsion τ .

The node X is then expanded as follows. We randomly sample control inputs $\mathbf{u} = [v, \kappa, \tau]^T$ in the ranges $[0, \bar{v}]$, $[-\bar{\kappa}, \bar{\kappa}]$, and $[-\bar{\tau}, \bar{\tau}]$, respectively. We select the best set of controls \mathbf{u} that gets closest to \mathbf{p} using the Euclidean distance metric [64]. Let X' be the new pose obtained by integrating the controls starting from pose X (Eq. (7)). We then check if X' is feasible, i.e., it does not violate the cumulative curvature $\bar{\kappa}^c$ and torsion $\bar{\tau}^c$ constraints. This is easy to check by storing the cumulative curvature and torsion values at each node in the tree as it is grown. If feasible, we then check if the constant twist trajectory between X and X' does not collide with the implant boundary I , obstacles \mathcal{O} , and the existing set of ribbons \mathcal{R} . If collision-free, we add X' and the edge from X to X' to the tree \mathcal{T} . If X' lies within the entry region, we stop growing the tree and compute a plan by traversing the tree \mathcal{T} starting backwards from X' till the root node is reached. Given this plan, we




Fig. 5. **No obstacle scenario:** Ribbons generated by (a) the RRT planner and (b) the local optimization method. The RRT planner uses random sampling of control inputs, which leads to unnecessary twists and turns in the generated ribbon. The local optimization is able to compute an optimal ribbon with zero curvature and torsion.

generate the entire ribbon r^i by integrating the sequence of controls along each edge of the plan. This ribbon r^i is added to the set of existing ribbons \mathcal{R} .

We note that any sampling-based motion planner, including our method, cannot guarantee that a globally optimal solution will be found in a finite-time interval. Planners like RRT* [31] can compute optimal motion plans as computation time is allowed to increase, but require solving a two-point boundary value problem, which is very challenging for a nonholonomic system such as one considered in this work. Recently, Li et al. [35] and Xie et al. [63] have proposed planners that are asymptotically optimal for kinodynamic systems but it is not clear if they will be able to extend to our application and this is a topic for future research.

It is well known that randomized planners compute sub-optimal, non-smooth plans [34]. We illustrate this in a scenario with no obstacles in Fig. 5. The RRT planner computes a feasible solution to the entry region but the solution is clearly sub-optimal, with unnecessary twists and turns (Fig. 5(a)). We also note that the RRT planner does not optimize the objective stated in Eq. (8)). As is standard practice, we further locally optimize the RRT solution using a local optimization procedure outlined below. In the scenario with no obstacles, the local optimization improves the RRT initialization to a straight ribbon with zero curvature and torsion, as is expected (Fig. 5(b)). However, we note that the nonholonomic kinematics and planning in the pose space makes it difficult to employ standard smoothing heuristics suggested in the literature [34].

B. Local Optimization

In this stage, we simultaneously optimize the cumulative curvature and torsion along all ribbons in the set \mathcal{R} of feasible ribbons generated by the RRT planner. Each ribbon $r^i \in \mathcal{R}$ is parameterized as a sequence of poses and controls as $[\mathcal{X}^i, \mathcal{U}^i]$

and has T^i time steps. We formulate the optimization problem using the objective and constraints described in Sec. IV as:

$$\min_{\mathcal{X}^i, \mathcal{U}^i} \sum_{i=1}^n C(\mathcal{U}^i), \quad \text{subject to constraints.} \quad (9)$$

Objective and constraints: The optimization problem stated above has a quadratic objective (Eq. (8)), which is well suited for optimization. However, the constraints are nonlinear. These constraints are converted into the standard equality and inequality constraints for optimization as described below:

- Constraints already expressed in standard form: $X_0^i = X^i$, $(|\kappa_t^i| < \bar{\kappa})$, and $(|\tau_t^i| < \bar{\tau})$.
- Non-convex constraints in standard form include: $X_{t+1}^i = X_t^i \exp(\delta U_t^i)$, $\sum_{t=0}^{T^i-1} |\delta v_t^i \kappa_t^i| < \bar{\kappa}^c$, and $\sum_{t=0}^{T^i-1} |\delta v_t^i \tau_t^i| < \bar{\tau}^c$
- In this work, the entry region E is defined as a convex region. We use a bounding circle or rectangle depending on the scenario. The constraint $X_{T^i}^i \in E$ is then formulated as a nonlinear inequality constraint based on whether the ribbon cross-section at time step T^i lies within the bounds of E .
- The collision avoidance constraints $(r^i \cap I = \emptyset)$ and $(r^i \cap \mathcal{O} = \emptyset)$ are encoded as nonlinear inequality constraints $\text{sd}(X_t^i, X_{t+1}^i, I) > 0$ and $\text{sd}(X_t^i, X_{t+1}^i, \mathcal{O}) > 0$, respectively, for $0 \leq t < T^i, 1 \leq i \leq n$. Here, sd denotes the signed distance. Similarly, the constraint $r^i \cap r^j = \emptyset$ is encoded as $\text{sd}(X_t^i, X_{t+1}^i, X_t^j, X_{t+1}^j) > 0$ for $1 \leq j \leq n, j \neq i$. We refer the reader to Schulman et al. [50] for details on how to efficiently compute the signed distance between convex objects and linearize such constraints.

Optimization method: We solve this constrained nonlinear optimization problem via sequential quadratic programming (SQP), where we repeatedly construct a quadratic program (QP) that locally approximates the original problem around the current solution $\mathbf{y} = \{\mathcal{X}^i, \mathcal{U}^i\}, 1 \leq i \leq n$. In particular, we use the ℓ_1 -SQP method proposed in [43]. This method has been successfully used for robot motion planning in a variety of contexts [19], [50] and we include a brief description below for completeness.

Two key ingredients of a sequential convex optimization algorithm are: (1) a method for constraining the step to be small, so the solution vector remains within the region where the approximations are valid; (2) a strategy for turning the infeasible constraints into penalties, which eventually drives all of the constraint violations to zero. For (1), we use a trust region modeled as a box constraint around the current iterate. For (2) we use ℓ_1 penalties: each inequality constraint $g_i(\mathbf{y}) \leq \mathbf{0}$ becomes the penalty $|g_i(\mathbf{y})|^+$, where $|\mathbf{y}|^+ = \max(\mathbf{y}, \mathbf{0})$; each equality constraint $h_i(\mathbf{y}) = \mathbf{0}$ becomes the absolute value penalty $|h_i(\mathbf{y})|$. In both cases, the penalty is multiplied by some coefficient μ , which is sequentially increased, usually by multiplying by a constant scaling factor at each step, during the optimization to drive constraint violations to zero. Note that ℓ_1 penalties are non-differentiable but convex, and convex optimization algorithms can efficiently minimize them.

In our formulation, the objective is directly expressed in the quadratic form. However, the constraints are nonlinear and

have to be linearized for inclusion in the QP. The QP is then solved and we compute a step based on a merit function [43] to ensure that progress is made on the original problem. To satisfy constraints up to a tolerance, we use ℓ_1 penalties that are progressively increased over the SQP iterations.

The optimization problem outlined above is, however, described directly over the set of poses \mathcal{X} . One could use a global parameterization of the rotation group, such as axis-angle coordinates or Euler angles. The drawback of those parameterizations is that they distort the geometry—for example, consider how a map of the world is distorted around the poles. This distortion can severely slow down an optimization algorithm, by reducing the neighborhood where local (first and second-order) approximations are good [47].

In this work, we follow the approach of Saccon et al. [47] and Duan et al. [19] to generalize SQP to the case where the domain is a differentiable manifold such as the $SE(3)$ Lie group rather than \mathbb{R}^n by considering a local coordinate parameterization of the manifold [47]. This parameterization is given by the Lie algebra $\mathfrak{se}(3)$, which is defined as the tangent vector space at the identity of $SE(3)$. The $SE(3)$ group and $\mathfrak{se}(3)$ algebra are related via the exponential and log maps, $\exp : \mathfrak{se}(3) \rightarrow SE(3)$ and $\log : SE(3) \rightarrow \mathfrak{se}(3)$ [42]. The local neighborhood of a nominal pose $X \in SE(3)$ is then defined in terms of the incremental twist $\bar{\mathbf{x}} \in \mathbb{R}^6$ using the exp map [42]. Given a vector $\bar{\mathbf{x}} = \begin{bmatrix} \bar{\mathbf{p}} \\ \bar{\mathbf{r}} \end{bmatrix} \in \mathbb{R}^6$ that represents the incremental twist, the corresponding Lie algebra element is given by the mapping $\wedge : \mathbb{R}^6 \rightarrow \mathfrak{se}(3)$ as $\bar{\mathbf{x}}^\wedge = \begin{bmatrix} [\bar{\mathbf{r}}] & \bar{\mathbf{p}} \\ \mathbf{0}_3^T & 0 \end{bmatrix}$, where the notation $[\bar{\mathbf{r}}]$ for the vector $\bar{\mathbf{r}} = [\bar{\mathbf{r}}_x \ \bar{\mathbf{r}}_y \ \bar{\mathbf{r}}_z]^T \in \mathbb{R}^3$ is the 3×3 skew-symmetric matrix. Intuitively, $\bar{\mathbf{r}}$ represents the incremental rotation and $\bar{\mathbf{p}}$ represents the incremental translation to be applied to a nominal pose. The inverse is defined by the operator $\vee : \mathfrak{se}(3) \rightarrow \mathbb{R}^6$ to recover $\bar{\mathbf{x}}$ given a Lie algebra element, i.e., $\begin{bmatrix} [\bar{\mathbf{r}}] & \bar{\mathbf{p}} \\ \mathbf{0}_3^T & 0 \end{bmatrix}^\vee = \bar{\mathbf{x}}$.

We then construct and solve each QP in terms of the increments to the previous solution. At each SQP iteration, the set of poses \mathcal{X}^i is updated based on the incremental twists computed by solving the QP approximation. As an example, consider the j^{th} iteration of SQP and let $\bar{\mathcal{X}}^{(j)} = \{\bar{\mathbf{x}}_0^{(j)}, \dots, \bar{\mathbf{x}}_T^{(j)}\}$ be the sequence of incremental twists (step) computed by solving the convex subproblem. Given a trajectory consisting of a sequence of nominal poses $\hat{\mathcal{X}}^{(j)} = \{\hat{\mathbf{x}}_0^{(j)}, \dots, \hat{\mathbf{x}}_T^{(j)}\}$, the subsequent sequence of poses is obtained by applying $\bar{\mathcal{X}}^{(j)}$ as $\hat{\mathcal{X}}^{(j+1)} = \{\exp(\bar{\mathbf{x}}_0^{(j)\wedge}) \cdot \hat{\mathbf{x}}_0^{(j)}, \dots, \exp(\bar{\mathbf{x}}_T^{(j)\wedge}) \cdot \hat{\mathbf{x}}_T^{(j)}\}$, where $\wedge : \mathbb{R}^6 \rightarrow \mathfrak{se}(3)$ is the wedge operator that maps a twist vector in \mathbb{R}^6 to the $\mathfrak{se}(3)$ Lie algebra [42]. We refer the reader to [19], [47] for additional details.

Simultaneous versus sequential optimization: Duan et al. [19] propose a sequential optimization strategy for optimizing the generated channels. A sequential strategy optimizes the ribbons one at a time while treating all the other ribbons as obstacles for formulating the collision avoidance constraints. The optimization problem involves a lesser number of variables and collision avoidance constraints in the optimization problem. However, in our experiments, we found that the sequential strategy is unable to optimize the objective when multiple ribbons are considered because the ribbons are processed in

a chosen order (Fig. 7(b)). In our experiments, the sequential optimization strategy also fails to generate a feasible set of ribbons when the environment is very spatially constrained, such as shown in Fig. 1(b). The simultaneous optimization strategy is more computationally expensive since it involves a larger number of variables in the optimization formulation and collision avoidance constraints. However, it is able to resolve conflicts by jointly optimizing over all the ribbons (Fig. 7(c)).

We note that a similar tradeoff occurs in multi-robot motion planning problems where the sequential strategy corresponds to prioritized motion planning for multiple robots [4], [58] and the simultaneous optimization strategy corresponds to centralized multi-robot motion planning [34], [48]. Planning for multiple ribbons is also an instance of the multi-robot motion planning problem with additional kinematic motion constraints and curvature and torsion constraints.

VI. EXPERIMENTS

We consider three planning scenarios to highlight the merits and demerits of the RRT planner and the local optimization method. In each of these scenarios, we consider a box-shaped implant volume. The objective is to generate a ribbon from a specified target configuration to the entry region, which is defined by one of the faces of the box.

No obstacle scenario: In the first scenario, as shown in Fig. 5, we do not consider any obstacles. The RRT planner generates a ribbon with unnecessary twists and turns as a result of the randomized nature of the algorithm (Fig. 5(a)). The local optimization method is able to optimize the ribbon to generate one with zero curvature and torsion (Fig. 5(b)).

Two walls scenario: In the second scenario, as shown in Fig. 6, we consider two box-shaped obstacles that block the straight line path from the target configuration to the entry region. If we use local optimization in isolation with a naïve straight line initialization for the ribbon, the optimization is unable to resolve collisions with the two obstacles (Fig. 6(b)). The RRT planner is able to resolve collisions but generates a sub-optimal ribbon (Fig. 6(a)). However, applying local optimization to the RRT-generated ribbon generates a high quality ribbon that is able to avoid collisions with obstacles within the environment (Fig. 6(c)).

Multiple ribbon scenario: In the third scenario, as shown in Fig. 7, we consider an empty box-shaped implant volume where the objective is to reach the entry region corresponding to the face of the box. The RRT planner is used to sequentially generate mutually collision-free ribbons, as shown in Fig. 7(a). In this sequential strategy, the planner treats previously generated ribbons as obstacles. The generated set of ribbons is sub-optimal in terms of the curvature and torsion along the length of the ribbons, which is further optimized using local optimization. If we use a sequential optimization strategy, we optimize each ribbon one at a time but the solution is dependent on the order in which the ribbons are optimized and we end up with a sub-optimal set of ribbons with unnecessary changes in curvature and torsion, as shown in Fig. 7(b). Our simultaneous optimization strategy jointly optimizes over all the ribbons at once. The joint optimization, although more




Fig. 6. **Two walls scenario:** Two walls are positioned such that a direct path to the entry region is blocked. (a) The RRT planner explores the free space to find a sub-optimal ribbon. (b) Local optimization fails to find a feasible solution starting from an infeasible, straight line initialization to the entry region. (c) The RRT generated ribbon is provided as initialization to the local optimization, to generate a ribbon that minimizes curvature and has zero torsion.




Fig. 7. **Multiple ribbon scenario:** (a) We use the RRT planner to sequentially generate mutually collision-free ribbons that reach the entry region. (b) Sequential optimization locally optimizes the curvature and torsion along each ribbon one at a time and considers all the other ribbons as potential obstacles during the optimization. The optimization problem involves a lesser number of variables but the solution is highly dependent on the order in which the ribbons are optimized. In this case, we end up with a sub-optimal set of ribbons with unnecessary changes in curvature and torsion. (c) The simultaneous optimization strategy jointly optimizes over all the ribbons at once. The optimization problem involves a larger number of variables and is computationally expensive to solve but the joint optimization is able to optimize the curvature and torsion along all the ribbons to zero.

computationally expensive, does not suffer from the problem of sequential ordering and successfully optimizes the curvature and torsion along all the ribbons to zero, as shown in Fig. 7(c).

VII. DESIGNING 3D IMPLANTS FOR INTRACAVITARY BRACHYTHERAPY

We considered a scenario where a 3D printed implant is used for treatment of OB/GYN tumors, as shown in Fig. 8. The implant was modeled as a cylinder of height 7 cm and radius 2.5 cm, with an attached hemisphere with radius 2.5 cm. The dimensions of the implant was designed based on dimensions reported by Garg et al. [24]. We considered 3 tumors that are targeted for intracavitary brachytherapy treatment. We placed 6 groups of candidate dose dwell segments that are placed proximal to the tumors within the implant volume and oriented tangentially to maximize dose distribution to the tumor volumes. The circular entry region is the base of the implant. The objective is to generate mutually collision-free, curvature and torsion constrained ribbons within the implant volume that reach the candidate dose dwell segments.

Standard catheters used for brachytherapy are 1.65 mm – 2 mm [18] in diameter. In this scenario, we consider channels of width $w = 2.5$ mm. We imposed a constraint on the instantaneous curvature of $\bar{\kappa} = 1\text{cm}^{-1}$ based on the maximum allowable curvature reported by Garg et al. [24] for radioactive sources that are 1 mm in diameter and 5 mm in length. We also imposed a constraint on the instantaneous torsion of

$\bar{\tau} = 0.1$ radians. We set maximum limits on the cumulative curvature and torsion for each ribbon as $\bar{\kappa}^c = \frac{\pi}{2}$ and $\bar{\tau}^c = \frac{\pi}{2}$ units, respectively. We also constrained the cross-section of each ribbon at the respective final time steps to lie within the specified entry region, which is the bottom face of the implant.

We implemented our planning approach in C++ and used the Bullet collision checking library [15] for collision checking queries. We chose Bullet because of the high-performance signed distance implementations for convex-convex collision checking. Fig. 8(a) shows the candidate ribbons computed using the RRT planner. The candidate ribbons are mutually collision-free but contain unnecessary changes in curvature and torsion along the ribbons since the RRT planner does not optimize the considered objective. Fig. 8(b) shows the results of jointly optimizing the ribbons using the local optimization approach. In practice, this optimized arrangement of channels would be printed with support material and later dissolved to compute the hollow internal channels within the implant.

Table I summarizes the computation time required to generate these ribbons using the RRT planner and local optimization as we vary the number of channels per ribbon. The RRT planner is faster since it generates these ribbons sequentially while the local optimization jointly optimizes over all ribbons and hence is computationally expensive. In this scenario, our approach was not able to find a solution for greater than 6 channels per ribbon due to the limited free space within the implant volume.

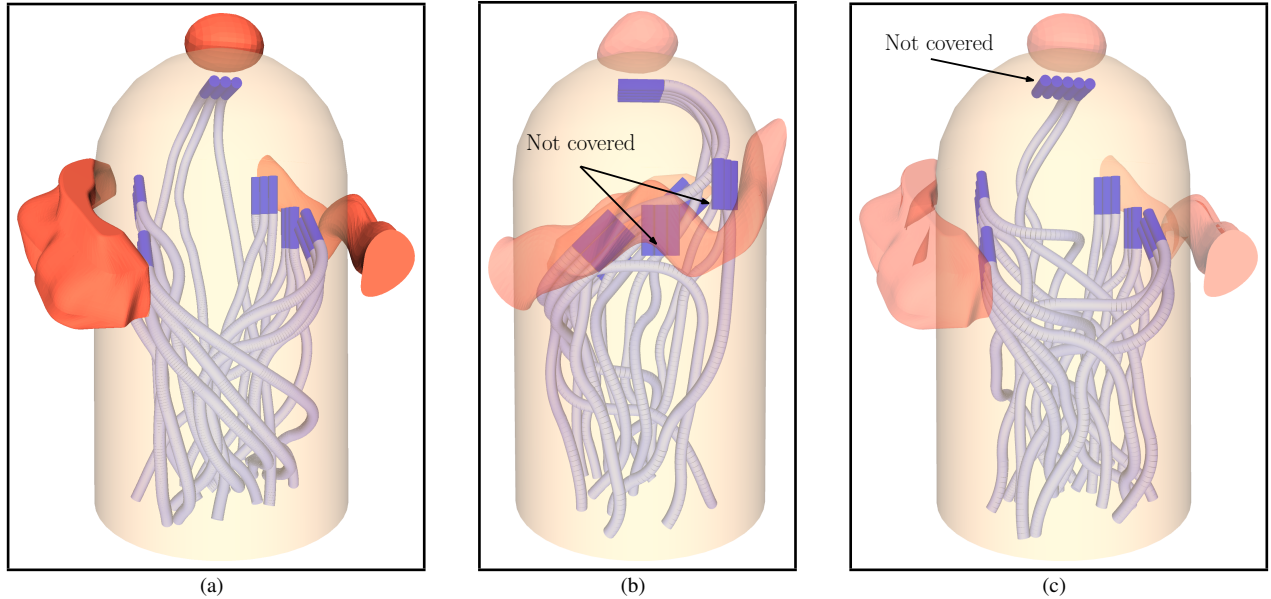


Fig. 9. **Comparison with single channel arrangements:** We compare our approach to generating individual channels for groups of candidate dwell dose segments following the approach of Garg et al. [24]. (a) Even with 3 dwell segments per group (18 channels in all) in an effort to cover all tumors (shown in red), the implant volume is occupied with internal channels. The planner is unable to generate single channels for greater than 3 channels per group. This leads to less effective coverage of the tumor volume in terms of dose distribution. (b,c) The tumors are rendered as translucent to visualize dwell segments that are not covered by the planner. (b) The quality of the sequential RRT planning process is dependent on the order in which dwell positions are processed by the planner. In this situation, we processed dwell segments in decreasing distance to the entry region. The planner is unable to find feasible channels that visit dwell segments closer to the entry region since the channels computed earlier in the planning process impede the generation of subsequent channels. (c) In this situation, we processed dwell segments in increasing distance to the entry region. The planner is able to successfully generate channels for the dwell segments that are closer but is unable to compute feasible channels for dwell segments that are farthest from the entry region as there is insufficient free space to generate subsequent channels.

Num. channels	1	2	3	4	5	6
RRT time(s)	0.6	1.1	3.5	9.6	16.7	38.4
Opt. time(s)	53.7	81.9	143.5	247.3	313.9	397.1

TABLE I

PERFORMANCE OF OUR PLANNING APPROACH WITH DIFFERENT NUMBER OF CHANNELS PER RIBBON. THE CUMULATIVE TIME IS THE SUM OF THE RRT AND OPTIMIZATION TIMES. THE REPORTED TIMES FOR THE RRT PLANNER ARE AVERAGED OVER 10 RUNS. ALL EXECUTION TIMES ARE BASED ON EXPERIMENTS RUN ON A SINGLE 3.5 GHZ INTEL I7 PROCESSOR CORE.

A. Comparison with Single Channel Arrangements

We also compared the use of ribbon-like arrangements versus generating single channels for groups of dose dwell segments following the approach of Garg et al. [24] that uses an RRT planner to generate curvature constrained channels by processing the dwell segments sequentially. The channels also have to be mutually collision-free and this leads to ineffective utilization of space within the implant volume as the number of channels increases. The quality of the solution is also highly dependent on the order in which the dwell segments are processed by the planner.

The ability to deliver radiation doses depends on the arrangement of dwell segments and their proximity to tumors. Fig. 9(a) shows an example arrangement of individual channels in groups of 3 dwell segments. Here, the objective was to compute channels that would attempt to reach all the tumors. However, the RRT planner is unable to compute feasible solutions since the channels occupy the entire volume within the implant, thus preventing the planner from adding additional channels to the implant. This results in less effective coverage

of the tumors, as described in Sec. VII-B.

The solution quality is dependent on the order in which the dwell segments are processed during planning. Fig. 9(b) shows an example arrangement of individual channels applied to the original placement of dwell segments considered in Fig. 8(b). In this situation, we processed dwell segments in the order of decreasing distance to the entry region, i.e., we first processed dwell segments that are farthest from the entry region. The planner is unable to find feasible channels that visit dwell segments closer to the entry region since the channels computed earlier in the planning process act as obstacles and impede the generation of subsequent channels.

Fig. 9(c) shows an example arrangement of individual channels where we processed dwell segments in the order of increasing distance to the entry region, i.e., we first processed dwell segments that are closest to the entry region. In this situation, the planner is able to successfully generate channels for the dwell segments that are closer but is unable to compute feasible channels for dwell segments that are farthest from the entry region as there is insufficient free space to generate subsequent channels.

B. Coverage Quality Metric

Each dose dwell segment d^i , $i = 1, \dots, n$ is divided into dwell points $\{dP\}$ that are spaced approximately 5 mm apart, which corresponds to the dimensions of the source. Let the set of reachable dwell positions be denoted as $\mathcal{P} = \{dP\}$. The radiation dosage from the source follows an inverse square




Fig. 8. We consider an implant volume modeled as a cylinder and attached hemisphere. The objective is to generate 6 mutually collision-free ribbons that reach groups of candidate dwell segment groups proximally located and oriented tangentially to the tumors. (a) The RRT planner is able to sequentially plan for each dwell segment group to generate collision-free ribbons that have unnecessary twists. (b) The candidate ribbons are jointly optimized locally to minimize the curvature and torsion along the length of each ribbon.

law i.e., the radiation decay is inversely proportional to the squared of the distance to the source [10], [27].

The ability to deliver radiation doses depends on the arrangement of potential dwell points and their proximity to tumors. We measure the quality of an implant by the percentage of tumor volume that is covered by the set of dwell points, where coverage is a function of the distance between a dwell point (source) and a tumor point (target). Higher quality reduces the maximum dwell time needed to treat tumors and the potential for hot spots that can harm healthy tissue.

To compare implants and channels for a given set of tumors \mathcal{C} , we considered the set of reachable dwell positions $\mathcal{P} = \{dP\}$ and how thoroughly they cover the set of tumors. We discretized the set of tumors into a set of evenly spaced points $\mathcal{C} = \{dC\}$ defined on a regular grid. We quantified the proximity of a dwell position dP from a tumor point dC with the *coverage radius* ε such that: if dP lies within a ball of radius ε centered at dC , then dP is said to cover dC . Hence the cover of dC is the set as given by Garg et al. [24] as

$$\text{coverage}(dC, \varepsilon) = \{dP : \|dP - dC\|_2 \leq \varepsilon, dP \in \mathcal{P}\}. \quad (10)$$

Reaching 100% coverage with a smaller radiation radius and more dwell positions can reduce occurrence of radiation hot spots and increase dose conformation to the tumor geometry to spare healthy tissue. A smaller coverage radius parameter results in lower dose to healthy organs while supplying suf-




Fig. 10. **Limitation: Completeness:** We consider a narrow passage scenario. The target is oriented so that a feasible ribbon would have to simultaneously twist and turn and make its way through a narrow passage. (a) Feasible solution found by the RRT planner. (b) The RRT solution is further locally optimized to minimize curvature and torsion. Without the feasible initialization, the local optimization is unable to find a feasible solution. The narrow passage is a known problem for sampling-based planners like RRT. For narrower passages than the one considered here, our approach is unable to find a solution even though a feasible solution exists for a passage that is just wider than the ribbon width.


ficient radiation to the tumor volumes. Hence a higher tumor coverage with small ε is preferred for brachytherapy treatment.

The use of single channels limits the number of reachable dose dwell segments. This affects the coverage of tumor volume that can receive radiation, thus limiting the treatment effectiveness. In the scenario considered above, we found that dwell positions generated with ribbon-like channel arrangements (Fig. 8(b)) achieved 100% coverage of the tumor volume for $\varepsilon = 2.1$ cm. In contrast, the coverage achieved for single channels (Fig. 9(a)) for $\varepsilon = 2.1$ cm is only 54%. It is important to achieve 100% coverage to destroy all cancerous tissue to prevent cancer recurrence. Using ribbon-like arrangements allows us to achieve this coverage while minimizing damage to surrounding healthy tissue.

VIII. LIMITATIONS

In this section, we consider the limitations of our planning approach through illustrative examples. We consider the following two scenarios:

Limitation: Completeness: In this scenario, as shown in Fig. 10, we consider two box-shaped obstacles that only permit passage to the entry region via a narrow passage. The target orientation is specified such that the ribbon would have to simultaneously twist and turn to be able to traverse the narrow passage. Narrow passages are a known problem for sampling-based planners like RRT because it is difficult to generate and connect to collision-free samples in the narrow passage. This problem becomes especially difficult when the kinematic model is nonholonomic. The RRT planner is able to resolve collisions after repeated attempts (8 in our case) to generate a sub-optimal ribbon (Fig. 10(a)). Local optimization in isolation



(a) Initialization in different homotopy classes (Left) RRT initialization. (Right) Local optimization. (b) Initialization in same homotopy class (Left) RRT initialization. (Right) Local optimization.

Fig. 11. Limitation: Two stage planning: We consider a scenario with a cylindrical obstacle that spans the extent of the bounding box, thus dividing the free space into two distinct homotopy classes. (a) RRT solution in which the two ribbons are in different homotopy classes. Local optimization succeeds in optimizing the curvature and torsion in the two ribbons. (b) RRT solution in which the two plans are in the same homotopy class. The solution obtained using local optimization also lies in the same homotopy class, and is sub-optimal because of larger cumulative curvature in the ribbons.

is unable to find a feasible solution. However, applying local optimization to the RRT-generated ribbon is able to locally optimize the solution (Fig. 10(b)).

We note that for passages narrower than the one considered here, our approach is unable to find a feasible solution even though one exists for passages that are just wider than the ribbon width. Even though the RRT planner is theoretically probabilistically complete [34], i.e., it is guaranteed to find a solution if one exists, we cannot guarantee that this would be true in practice for our application.

Limitation: Two stage planning: In this scenario, we consider a scenario with a cylindrical obstacle that spans the extent of the bounding box, thus dividing the free space into two distinct homotopy classes. We also note that the ribbons themselves also partition the free space into distinct homotopy classes since they are required to be mutually collision-free.

If the RRT planner finds a solution in which the two ribbons are in different homotopy classes, the local optimization is successful in minimizing the curvature and torsion along the ribbons to obtain the solution that would be expected in such a scenario (as shown in Fig. 11(a)). However, if the RRT planner finds a solution in which the ribbons are in the same homotopy class, the local optimization subsequently finds a sub-optimal solution with larger cumulative curvatures along both ribbons (as shown in Fig. 11(b)). This is a well known limitation of gradient-based trajectory optimization methods [34] and is also a limitation of our approach. This shows that the topology of the RRT solution that is used as initialization has an effect on the quality of the final solution. Recently, Bhattacharya et al. [7] and Pokorný et al. [46] have shown that it is possible to classify and compute paths in distinct homotopy classes. Extending these concepts to nonholonomic motion planning and our approach is a promising avenue for future research.

IX. CONCLUSION

In this work, we posed the problem of planning curvature and torsion constrained ribbons that avoid collisions with obstacles in 3D environments. We showed that this problem is equivalent to planning motions for a rigid body along a spatial curve, such that the rigid body is oriented along the unit binormal to the curve defined according to the Frenet-Serret frame. We used a combination of sequential sampling-based (RRT) planning and simultaneous local optimization (SQP) to compute high quality ribbons for designing channels within 3D printed implants for intracavitary brachytherapy.

This opens up several avenues for future work. We plan to extend this work to automatically compute dose dwell segments that are proximally located with respect to the external tumors. We plan to test our planning approach to design and print 3D implants and evaluate them in clinical brachytherapy trials. Alternate arrangements of channels will also be considered. Given an set of dose dwell segments, one could also first plan individual channels starting from the dose dwell segments into a ribbon shape, which is a challenging planning problem. One can then plan a path of the ribbon to the entry region (similar to routing of wires inside electrical equipment). We envision that our approach will also be useful for other applications such as routing ribbon-like arrangements of cooling channels, wires and cables, and planning motions for bevel-tipped steerable needles.

ACKNOWLEDGMENTS

This research has been funded in part by by U.S. National Science Foundation under Award IIS-1227536: Multilateral Manipulation by Human-Robot Collaborative Systems, AFOSR-YIP Award #FA9550-12-1-0345, by a DARPA Young Faculty Award #D13AP00046, a seed grant from the UC Berkeley Center for Information Technology in the Interest of Society (CITRIS), and by a Sloan Fellowship. This work has been supported in part by funding from Google and Cisco. The authors thank Jean Pouliot, J. Adam M. Cunha, I-Chow Hsu, Timmy Siau, and Animesh Garg for ongoing insights and advice on this topic and the staff at UCSF Medical Center at Mt. Zion, San Francisco, CA for input on the clinical procedures for GYN HDR brachytherapy.

REFERENCES

- [1] P. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.
- [2] R. Alterovitz, T. Siméon, and K. Goldberg, "The Stochastic Motion Roadmap: A Sampling Framework for Planning with Markov Motion Uncertainty," in *Robotics: Science and Systems (RSS)*, 2007.
- [3] C. Belta and V. Kumar, "Optimal Motion Generation for Groups of Robots: A Geometric Approach," *Journal of Mechanical Design*, vol. 126, no. 1, pp. 63–70, 2004.
- [4] M. Bennewitz, W. Burgard, and S. Thrun, "Optimizing Schedules for Prioritized Path Planning of Multi-Robot Systems," in *Proc. Int. Conf. Robotics and Automation (ICRA)*, vol. 1, 2001, pp. 271–276.
- [5] M. Bergou, M. Wardetzky, S. Robinson, B. Audoly, and E. Grinspun, "Discrete Elastic Rods," *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3, p. 63, 2008.

- [6] M. Bernstein, K. J. Mehta, R. Yaparparvi, H. Kuo, and S. Kalnicki, "Results of the Hybrid Interstitial-Intracavitary Utrecht Applicator for Cervical Cancer in an Outpatient Setting," *Radiotherapy and Oncology*, vol. 103, p. S116, 2012.
- [7] S. Bhattacharya, M. Likhachev, and V. Kumar, "Topological Constraints in Search-based Robot Path Planning," *Autonomous Robots*, vol. 33, no. 3, pp. 273–290, 2012.
- [8] J. Biggs and W. Holderbaum, "Planning Rigid Body Motions using Elastic Curves," *Mathematics of Control, Signals, and Systems*, vol. 20, no. 4, pp. 351–367, 2008.
- [9] R. L. Bishop, "There is more than one way to Frame a Curve," *American Mathematical Monthly*, pp. 246–251, 1975.
- [10] J. Borg and D. Rogers, "Monte Carlo Calculations of Photon Spectra in Air from 192Ir Sources," *National Research Council Report PIRS-629; Ontario, Canada*, 1999.
- [11] T. Bretl and Z. McCarthy, "Quasi-Static Manipulation of a Kirchhoff Elastic Rod based on a Geometric Analysis of Equilibrium Configurations," *Int. Journal of Robotics Research*, vol. 33, no. 1, pp. 48–68, 2014.
- [12] J. Chopin and A. Kudrolli, "Helicoids, Wrinkles, and Loops in Twisted Ribbons," *Physical Review Letters*, vol. 111, no. 17, p. 174302, 2013.
- [13] H. Choset, "Coverage for Robotics—A Survey of Recent Results," *Annals of mathematics and artificial intelligence*, vol. 31, no. 1-4, pp. 113–126, 2001.
- [14] D. Chubelaschwili and U. Pinkall, "Elastic Strips," *manuscripta mathematica*, vol. 133, no. 3-4, pp. 307–326, 2010.
- [15] E. Coumans, "Bullet Collision Detection and Physics Library," Available: <http://bulletphysics.org>, 2013.
- [16] R. Cripps and G. Mullineux, "Constructing 3D Motions from Curvature and Torsion Profiles," *Computer-Aided Design*, vol. 44, no. 5, pp. 379–387, 2012.
- [17] L. Delclos, G. H. Fletcher, E. Bailey Moore, and V. A. Sampiere, "Minicolpostats, Dome Cylinders, Other Additions and Improvements of the Fletcher-Suit Afterloadable System: Indications and Limitations of their Use," *International Journal of Radiation Oncology, Biology, Physics*, vol. 6, no. 9, pp. 1195–1206, 1980.
- [18] P. Devlin, *Brachytherapy: Applications and Techniques*. Lippincott Williams & Wilkins, 2007.
- [19] Y. Duan, S. Patil, J. Schulman, K. Goldberg, and P. Abbeel, "Planning Locally Optimal, Curvature-Constrained Trajectories in 3D using Sequential Convex Optimization," in *Proc. Int. Conf. Robotics and Automation (ICRA)*, 2014, pp. 5889–5895.
- [20] A. Elnagar and A. Hussein, "On Optimal Constrained Trajectory Planning in 3D Environments," *Robotics and Autonomous systems*, vol. 33, no. 4, pp. 195–206, 2000.
- [21] G. E. Farin, *Curves and Surfaces for Computer-Aided Geometric Design: A Practical Code*. Academic Press, Inc., 1996.
- [22] T. Farmer, "A New Model for Ribbons in \mathbb{R}^3 ," *Mathematics Magazine*, vol. 79, no. 1, p. 31, 2006.
- [23] E. Frazzoli, M. A. Dahleh, and E. Feron, "Maneuver-based Motion Planning for Nonlinear Systems with Symmetries," *IEEE Trans. on Robotics*, vol. 21, no. 6, pp. 1077–1091, 2005.
- [24] A. Garg, S. Patil, T. Siau, J. A. M. Cunha, I.-C. Hsu, P. Abbeel, J. Pouliot, and K. Goldberg, "An Algorithm for Computing Customized 3D Printed Implants with Curvature Constrained Channels for Enhancing Intracavitary Brachytherapy Radiation Delivery," in *IEEE Int. Conf. on Automation Science and Engg. (CASE)*, vol. 4, 2013, pp. 3306–3312.
- [25] O. Goemans and M. Overmars, "Automatic Generation of Camera Motion to Track a Moving Guide," in *Algorithmic Foundations of Robotics VI*. Springer, 2005, pp. 187–202.
- [26] A. Goriely and P. Shipman, "Dynamics of Helical Strips," *Physical Review E*, vol. 61, no. 4, p. 4508, 2000.
- [27] E. C. Halperin, L. W. Brady, D. E. Wazer, and C. A. Perez, *Perez & Brady's Principles and Practice of Radiation Oncology*. Lippincott Williams & Wilkins, 2013.
- [28] A. J. Hanson, "Quaternion Frenet Frames: Making Optimal Tubes and Ribbons from Curves," Tech. Rep. 407, Indiana University Computer Science Department, Tech. Rep., 1994.
- [29] Y. Y. Huang and K. H. Low, "Comprehensive Planning of Robotic Therapy and Assessment of Task-Oriented Functions via Improved QFD Applicable to Hand Rehabilitation," in *IEEE Int. Conf. on Automation Science and Engg. (CASE)*, 2010, pp. 252–257.
- [30] M. Hwangbo, J. Kuffner, and T. Kanade, "Efficient Two-Phase 3D Motion Planning for Small Fixed-Wing UAVs," in *Proc. Int. Conf. Robotics and Automation (ICRA)*, 2007, pp. 1035–1041.
- [31] S. Karaman and E. Frazzoli, "Sampling-based Algorithms for Optimal Motion Planning," *Int. Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [32] M. B. Kobilarov and J. E. Marsden, "Discrete Geometric Optimal Control on Lie Groups," *IEEE Trans. on Robotics*, vol. 27, no. 4, pp. 641–655, 2011.
- [33] A. Krontiris, S. Louis, and K. E. Bekris, "Simulating Formations of Non-holonomic Systems with Control Limits along Curvilinear Coordinates," in *Motion in Games*, 2010, pp. 121–133.
- [34] S. LaValle, *Planning Algorithms*. Cambridge University press, 2006.
- [35] Y. Li, Z. Littlefield, and K. Bekris, "Asymptotically Optimal Sampling-based Kinodynamic Planning," *arXiv preprint arXiv:1407.2896*, 2014.
- [36] H. Lipson and M. Kurman, *Fabricated: The New World of 3D Printing*. John Wiley & Sons, 2013.
- [37] I. Llamas, A. Powell, J. Rossignac, and C. D. Shaw, "Bender: A Virtual Ribbon for Deforming 3D Shapes in Biomedical and Styling Applications," in *Proc. ACM symposium on Solid and physical modeling*, 2005, pp. 89–99.
- [38] N. Magné, C. Chargari, N. SanFilippo, T. Messai, A. Gerbaulet, and C. Haic-Meder, "Technical Aspects and Perspectives of the Vaginal Mold Applicator for Brachytherapy of Gynecologic Malignancies," *Brachytherapy*, vol. 9, no. 3, pp. 274–277, 2010.
- [39] F. P. W. Melchels, J. Feijen, and D. W. Grijpma, "A Review on Stereolithography and its Applications in Biomedical Engineering," *Biomaterials*, vol. 31, no. 24, pp. 6121–30, 2010.
- [40] J. A. Mendez, S. Torres, J. A. Reboso, and H. Rebozo, "Model-based Controller for Anesthesia Automation," in *IEEE Int. Conf. on Automation Science and Engg. (CASE)*, 2009, pp. 379–384.
- [41] M. Moll and L. E. Kavraki, "Path Planning for Deformable Linear Objects," *IEEE Trans. on Robotics*, vol. 22, no. 4, pp. 625–636, 2006.
- [42] R. M. Murray and S. S. Shankar, *A Mathematical Introduction to Robotic Manipulation*. CRC press, 1994.
- [43] J. Nocedal and S. Wright, *Numerical Optimization*. Springer Verlag, 2006.
- [44] S. Patil, J. Burgner, R. J. Webster III, and R. Alterovitz, "Needle Steering in 3D via Rapid Replanning," *IEEE Trans. on Robotics*, p. to appear, 2014.
- [45] S. Patil, J. Pan, P. Abbeel, and K. Goldberg, "Planning Curvature and Torsion Constrained Ribbons in 3D with Application to Intracavitary Brachytherapy," 2014.
- [46] F. Pokorny, M. Hawasly, and S. Ramamoorthy, "Multiscale Topological Trajectory Classification with Persistent Homology," in *Robotics: Science and Systems (RSS)*, 2014.
- [47] A. Saccon, J. Hauser, and A. P. Aguiar, "Optimal Control on Lie Groups: The Projection Operator Approach," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2230–2245, 2013.
- [48] G. Sanchez and J.-C. Latombe, "Using a PRM Planner to Compare Centralized and Decoupled Planning for Multi-Robot Systems," in *Proc. Int. Conf. Robotics and Automation (ICRA)*, vol. 2, 2002, pp. 2112–2119.
- [49] W. J. Schroeder, W. E. Lorensen, and S. Linthicum, "Implicit Modeling of Swept Surfaces and Volumes," in *Proc. of the Conf. on Visualization*, 1994, pp. 40–45.
- [50] J. Schulman, J. Ho, A. Lee, H. Bradlow, I. Awwal, and P. Abbeel, "Finding Locally Optimal, Collision-Free Trajectories with Sequential Convex Optimization," in *Robotics: Science and Systems (RSS)*, 2013.
- [51] J. Selig, "Characterisation of Frenet-Serret and Bishop Motions with Applications to Needle Steering," *Robotica*, vol. 31, no. 06, pp. 981–992, 2013.
- [52] M. Shanmugavel, A. Tsourdos, R. Zbikowski, and B. White, "3D Path Planning for Multiple UAVs using Pythagorean Hodograph Curves," in *AIAA Guidance, Navigation, and Control Conference*, 2007, pp. 20–23.
- [53] A. Shkolnik, M. Walter, and R. Tedrake, "Reachability-Guided Sampling for Planning under Differential Constraints," in *Proc. Int. Conf. Robotics and Automation (ICRA)*, 2009, pp. 2859–2865.
- [54] J. Solis and A. Takamishi, "Towards Enhancing the Understanding of Human Motor Learning," in *IEEE Int. Conf. on Automation Science and Engg. (CASE)*, 2009, pp. 591–596.
- [55] K. Sprott and B. Ravani, "Kinematic Generation of Ruled Surfaces," *Advances in Computational Mathematics*, vol. 17, no. 1-2, pp. 115–133, 2002.
- [56] K. Subburaj, B. Ravi, and M. G. Agarwal, "Automated 3D Geometric Reasoning in Computer Assisted Joint Reconstructive Surgery," in *IEEE Int. Conf. on Automation Science and Engg. (CASE)*, 2009, pp. 367–372.
- [57] J. P. Swensen and N. J. Cowan, "Torsional Dynamics Compensation Enhances Robotic Control of Tip-Steerable Needles," in *Proc. Int. Conf. Robotics and Automation (ICRA)*, 2012, pp. 1601–1606.

- [58] J. Van Den Berg and M. H. Overmars, "Prioritized Motion Planning for Multiple Robots," in *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005, pp. 430–435.
- [59] J. van den Berg, S. Patil, R. Alterovitz, P. Abbeel, and K. Goldberg, "LQG-Based Planning, Sensing, and Control of Steerable Needles," in *Proc. Workshop Algorithmic Foundations of Robotics (WAFR)*, 2010, pp. 373–389.
- [60] W. Wang, B. Jüttler, D. Zheng, and Y. Liu, "Computation of Rotation Minimizing Frames," *ACM Transactions on Graphics (TOG)*, vol. 27, no. 1, p. 2, 2008.
- [61] R. J. Webster III, J. S. Kim, N. J. Cowan, G. S. Chirikjian, and A. M. Okamura, "Nonholonomic Modeling of Needle Steering," *Int. Journal of Robotics Research*, vol. 25, no. 5-6, pp. 509–525, 2006.
- [62] P. Willemsen, J. K. Kearney, and H. Wang, "Ribbon Networks for Modeling Navigable Paths of Autonomous Agents in Virtual Environments," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 3, pp. 331–342, 2006.
- [63] C. Xie, J. van den Berg, S. Patil, and P. Abbeel, "Toward Asymptotically Optimal Motion Planning for Kinodynamic Systems using a Two-Point Boundary Value Problem Solver," in *Proc. Int. Conf. Robotics and Automation (ICRA)*, 2015.
- [64] J. Xu, V. Duindam, R. Alterovitz, J. Pouliot, J. A. Cunha, I. Hsu, and K. Goldberg, "Planning Fireworks Trajectories for Steerable Medical Needles to Reduce Patient Trauma," in *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, 2009, pp. 4517–4522.
- [65] K. Yang and S. Sukkarieh, "An Analytical Continuous-Curvature Path-Smoothing Algorithm," *IEEE Trans. on Robotics*, vol. 26, no. 3, pp. 561–568, 2010.
- [66] M. Zefran, V. Kumar, and C. B. Croke, "On the Generation of Smooth Three-Dimensional Rigid Body Motions," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 4, pp. 576–589, 1998.



Ken Goldberg is Professor of Industrial Engineering and Operations Research at UC Berkeley, with appointments in Electrical Engineering, Computer Science, Art Practice, and the School of Information. He was appointed Editor-in-Chief of the IEEE Transactions on Automation Science and Engineering (T-ASE) in 2011 and served two terms (2006–2009) as Vice-President of Technical Activities for the IEEE Robotics and Automation Society. Goldberg earned his PhD in Computer Science from Carnegie Mellon University in 1990. Goldberg is Founding Co-Chair of the IEEE Technical Committee on Networked Robots and Founding Chair of the (T-ASE) Advisory Board. Goldberg has published over 200 refereed papers and awarded eight US patents, the NSF Presidential Faculty Fellowship (1995), the Joseph Engelberger Award (2000), the IEEE Major Educational Innovation Award (2001) and in 2005 was named IEEE Fellow.



Sachin Patil received his B.Tech degree in Computer Science and Engineering from the Indian Institute of Technology, Bombay, in 2006 and his Ph.D. degree in Computer Science from the University of North Carolina at Chapel Hill, NC in 2012. He is now a postdoctoral researcher in the EECS department at University of California, Berkeley. His research interests include motion planning, planning under uncertainty, and medical robotics.



interests include motion planning, GPGPU, and machine learning for robotics.

Jia Pan received his BE degree from the Department of Automation, Tsinghua University in 2005, a MS degree from the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences in 2008, and his Ph.D. degree in Computer Science from the University of North Carolina at Chapel Hill, NC in 2013. He was a postdoctoral researcher in the EECS department at University of California, Berkeley. He joined the faculty in the Department of Computer Science at the University of Hong Kong in 2014. His research



Pieter Abbeel received a BS/MS in Electrical Engineering from KU Leuven (Belgium) and received his Ph.D. degree in Computer Science from Stanford University in 2008. He joined the faculty at UC Berkeley in Fall 2008, with an appointment in the Department of Electrical Engineering and Computer Sciences. His current research focuses on robotics and machine learning with a particular focus on challenges in personal robotics, surgical robotics and connectomics. He has won numerous awards, including best paper awards at ICML and ICRA,

the Sloan Fellowship, the Air Force Office of Scientific Research Young Investigator Program (AFOSR-YIP) award, the Office of Naval Research Young Investigator Program (ONR-YIP) award, the Darpa Young Faculty Award (Darpa-YFA), the Okawa Foundation award, the TR35, the IEEE Robotics and Automation Society (RAS) Early Career Award, and the Dick Volz Best U.S. Ph.D. Thesis in Robotics and Automation Award.