

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

Sachin Patil, Jia Pan, Pieter Abbeel, Ken Goldberg

University of California, Berkeley

Abstract. A “ribbon” is a surface traced out by sweeping a constant width line segment along a spatial curve. We consider the problem of planning multiple disjoint and collision-free ribbons of finite thickness along curvature and torsion constrained curves in 3D space. This problem is motivated by the need to route multiple smooth channels through a 3D printed structure for a healthcare application and is relevant to other applications such as defining cooling channels inside turbine blades, routing wires and cables, and planning trajectories for formations of aerial vehicles. We show that this problem is equivalent to planning motions for a rigid body, the cross-section of the ribbon, 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 present a two stage approach. In the first stage, we use sampling-based rapidly exploring random trees (RRTs) to generate feasible curvature and torsion constrained ribbons. In the second stage, we locally optimize the curvature and torsion along each ribbon using sequential quadratic programming (SQP). We evaluate this approach for a clinically motivated application: planning multiple channels inside 3D printed implants to temporarily insert high-dose radioactive sources to reach and cover tumors for intracavitary brachytherapy treatment. Constraints on the curvature and torsion avoid discontinuities (kinks) in the ribbons which would prevent insertion. In our experiments, our approach achieves an improvement of 46% in coverage of tumor volumes as compared to an earlier approach that generates each channel in isolation.

1 Introduction

Our work is motivated by applications where contiguous pathways or channels have to be routed through 3D environments while avoiding collisions with obstacles. In particular, we consider a clinical application of intracavitary brachytherapy where radioactive doses have to be delivered to cancerous tumors occurring in body cavities such as oral, rectal, gynecological, auditory, and nasal. Garg et al. [14] demonstrated that 3D printing can be used to design customized implants that conform to the patient anatomy. These implants have hollow internal channels that provide passage to radioactive sources. 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. They constructed implants with mutually collision-free channels that provide passage to catheters carrying a radioactive source (Fig. 1(a)).

The effectiveness of these implants for radiation treatment depends on how effectively candidate locations for placing the radioactive seed can cover the tumor surface




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. [14]. (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.

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, which looks like a ribbon (Fig. 1(b)), can increase coverage and hence achieve sufficient dose distribution to large tumor volumes. These ribbon-like arrangements are also spatially efficient, and allow for a larger number of channels to be embedded within the implant. Such arrangements are also better localized in space, thus improving the structural integrity of the 3D printed implant during the printing process. We note that such ribbon-like channel arrangements also commonly occur in other applications such as for routing wires inside electronic equipment such as computers and routing cooling channels through turbine blades [17].

The channels in these ribbons have to allow a radioactive source of finite dimensions to pass through [14], which imposes a constraint on the maximum instantaneous curvature of the ribbon. The spatial curve that generates the ribbon should also be continuous and differentiable (at least C^1 -continuous) since kinks in the curve 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. 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 [21]. However, unnecessary turns (curvature) and twists (torsion) in the channels can prevent the support material from being completely dissolved. In addition, each customized implant would have multi-

ple 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. This could potentially cause undesirable ambiguity in the motion of a catheter when pushed through such a channel.

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 line segment 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 [4] 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 [4].

We adopt a two stage approach for generating ribbons with the aforementioned specifications. We use a rapidly-exploring random trees (RRT) planner [20] that generates feasible curvature and torsion constrained candidate ribbons. 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) [10] that is modified for our application. 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. As shown in Fig. 1(b), ribbon-like arrangements lead to improved coverage of the tumor volume (46% improvement in our experiments), which allows for more effective treatment. We also show that ribbon-like arrangements are spatially efficient, allowing us to embed a larger number of channels within the implant as compared to prior work that embeds single channels [14, 10].

2 Related Work

The geometry of swept surfaces and swept volumes has been extensively studied in the literature [12]. A ribbon is a particular example of a swept surface where a constant width rigid body is swept along a spatial curve [18, 13]. Ribbons find applications in a number of domains including computer aided geometric design [31], geometry modeling of DNA strands [16], for planning layout arrangements of roads [36] and for path planning for rigid formations of nonholonomic vehicles [2, 19]. Sampling-based planning was used to generate variable width ribbon-like paths between a camera and moving object in the environment [15]. 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, snow removal, and planning search and rescue operations [6].

Prior work has also explored the topic of modeling and finding minimum energy configurations of linear deformable objects such as suture threads [22, 5]. 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 these methods could be used to 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 [22, 5]. In our work, the ribbons are not physically-based and are carved out of the 3D implant volume, thus allowing us to directly generate ribbons by planning motions for a rigid body that describes the cross-section of the ribbon.

Planning smooth motions of rigid bodies in 3D has also been studied [38, 2, 3, 8]. This requires planning in the space of 3D positions and orientations, also commonly referred to 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.

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 [35]. 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 [1, 11, 25, 37, 33, 26, 10]. Our approach can be used to explicitly limit the torsional rotation of the needle, thus potentially improving planning and control of steerable needles [32].

The motivation behind using 3D printed implants for intracavitary brachytherapy and characterization of the effectiveness of treatment in terms of a *coverage* metric is presented by Garg et al. [14]. They used a sampling-based RRT planner to generate individual, curvature constrained channels. Duan et al. [10] used an optimization-based method to compute these channel arrangements for individual channels. In contrast, we consider the problem of generating ribbon-like arrangements of contiguous channels that is more spatially efficient. 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 length of 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.

3 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 continuous, 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 the spatial curve parameterized by parameter t . 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. 2). There are infinitely many choices for such a frame [4]. 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




Fig. 2. A ribbon is a swept surface traced out by sweeping a rigid body along a spatial curve. 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.

of interest in this work. We note that the following analysis also applies to other choices such as the Bishop's frame, also known as the rotation minimizing frame [4, 34].

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 \mathbf{t}_t , normal \mathbf{n}_t , and binormal \mathbf{b}_t 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

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

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 [4] [18] as

$$\dot{\mathbf{t}}_t = v_t \kappa_t \mathbf{n}_t, \quad (2a)$$

$$\dot{\mathbf{n}}_t = -v_t \kappa_t \mathbf{t}_t + v_t \tau_t \mathbf{b}_t, \quad (2b)$$

$$\dot{\mathbf{b}}_t = -v_t \tau_t \mathbf{n}_t, \quad (2c)$$

$$\dot{\mathbf{p}}_t = v_t \mathbf{t}_t, \quad (2d)$$

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 . These equations of motion can be rewritten as [29]

$$\begin{bmatrix} \dot{\mathbf{t}}_t & \dot{\mathbf{n}}_t & \dot{\mathbf{b}}_t & \dot{\mathbf{p}}_t \\ 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} \mathbf{t}_t & \mathbf{n}_t & \mathbf{b}_t & \mathbf{p}_t \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -v_t \kappa_t & 0 & v_t \\ v_t \kappa_t & 0 & -v_t \tau_t & 0 \\ 0 & v_t \tau_t & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (3)$$

Using Eqs. (1) and (3), we get the following relation

$$\dot{X}_t = X_t \begin{bmatrix} 0 & -v_t \kappa_t & 0 & v_t \\ v_t \kappa_t & 0 & -v_t \tau_t & 0 \\ 0 & v_t \tau_t & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} = X_t U_t, \quad (4)$$

where $U_t \in \mathfrak{se}(3)$ is the velocity twist of the rigid body in its local coordinate frame [23] and is completely described in terms of three parameters v_t , κ_t , and τ_t . In Sec. 5, 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 [23]

$$U_t = \begin{bmatrix} 0 & -v_t \kappa_t & 0 & v_t \\ v_t \kappa_t & 0 & -v_t \tau_t & 0 \\ 0 & v_t \tau_t & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} [\mathbf{w}_t]_{\times} & \mathbf{v}_t \\ \mathbf{0} & 0 \end{bmatrix}, \text{ where } \mathbf{w}_t = [v_t \tau_t, 0, v_t \kappa_t]^T, \mathbf{v}_t = [v_t, 0, 0]^T, \quad (5)$$

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 does not undergo any motion in the direction of the binormal vector \mathbf{b}_t , the curvature and torsion stay constant at points along \mathbf{b}_t . 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. We illustrate this phenomenon in Fig. 3. 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.




Fig. 3. Ribbons generated by sweeping a line segment along a circular arc (shown as dashed line). (Left) The ribbon is oriented along the binormal vector \mathbf{b} pointing into the page. If the line segment does not undergo any motion along the binormal 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 line segment undergoes motion along the normal vector \mathbf{n} , each channel along the ribbon has different curvatures, torsional values, and lengths, which makes the planning problem much harder.

When the velocity twist U_t is held constant over time interval of duration δ , the differential motion model given by Eq. (4) can be explicitly integrated as

$$X_{t+1} = X_t \exp(\delta U_t) \quad (6)$$

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 [23].

4 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.
- 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.
- 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 $\bar{\tau}^c$.

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 Sec. 3, the local configuration of the i^{th} ribbon at time step t is specified by the pose $X_t^i = \begin{bmatrix} R_t^i & p_t^i \\ \mathbf{0} & 1 \end{bmatrix}$ 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. (5)), 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. (6).
- $(r^i \cap I = \emptyset) \wedge (r^i \cap \mathcal{O} = \emptyset)$: 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 :

$$C(\mathcal{U}^i) = \alpha_\kappa \sum_{t=0}^{T^i-1} (\delta v_t^i \kappa_t^i)^2 + \alpha_\tau \sum_{t=0}^{T^i-1} (\delta v_t^i \tau_t^i)^2, \quad (7)$$

which is equivalent to minimizing the energy or rotational strain along a curve [22].

5 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.

5.1 Rapidly-exploring Random Trees (RRT)

In the first stage, we use a customized sampling-based RRT planner [20] to plan motions of the rigid body according to the nonholonomic kinematic model described in Sec. 3. 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 5.1. 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. 4). 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} . We also bias the sampling towards the entry region to ensure progress [20] 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 [26]. 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 has been shown to work well in practice [30, 26].

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 [37]. Let X' be the new pose obtained by integrating the controls starting from pose X (Eq. (6)). 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

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  $\{d^1, \dots, d^n\}$  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 3D collision-free point
7:      $X \leftarrow \text{nearest\_neighbor}(\mathbf{p}, \mathcal{T})$  ▷ Reachability-guided neighbor search
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. (6)
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
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:     $r^i \leftarrow \text{build\_ribbon}(\mathcal{T}, X', w)$  ▷ 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  $\mathcal{R}$ 

```

a plan by traversing the tree \mathcal{T} starting backwards from X' till the root node is reached. Given this plan, we 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} .

It is well known that randomized planners compute sub-optimal or non-smooth plans [20]. We illustrate this in a scenario with no obstacles in Fig. 4. The RRT planner computes a feasible solution to the entry region but the solution is clearly sub-optimal, with unnecessary twists and turns (Fig. 4(a)). We also note that the RRT planner does not optimize the objective stated in Eq. (7). As is standard practice, we then 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. 4(b)). However, we note that the nonholonomic kinematic model and planning in the pose space makes it difficult to employ standard smoothing heuristics suggested in the literature [20].

5.2 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. 4 as:

$$\min_{\substack{\mathcal{X}^i, \mathcal{U}^i \\ 1 \leq i \leq n}} \sum_{i=1}^n C(\mathcal{U}^i), \quad \text{subject to constraints.} \quad (8)$$

Objective and constraints: The optimization problem stated above has a quadratic objective (Eq. (7)), which is well suited for optimization. However, the constraints are




Fig. 4. 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.

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. [28] 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. 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 [24] 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. We refer the reader to [24] for additional details on the ℓ_1 -SQP method. This method has been successfully used for robot motion planning in a variety of contexts [28, 10].

The optimization problem outlined above is, however, described directly over the set of poses \mathcal{X} . Using a global parameterization of the rotation group, such as axis-angle coordinates or Euler angles, is not suitable for direct optimization [27]. Instead, we follow the approach of Saccon et al. [27] and Duan et al. [10]. We consider a local coordinate parameterization of the pose given by the Lie algebra $\mathfrak{se}(3)$, which is defined as the tangent vector space at the identity of $SE(3)$. 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 [23]. 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. We refer the reader to Duan et al. [10] for additional details.

Note on simultaneous versus sequential optimization: Duan et al. [10] propose a sequential optimization strategy for optimizing the generated channels. A sequential strategy involves a lesser number of variables and collision avoidance constraints in the optimization problem. However, in our experiments, we found that the sequential strategy failed to generate a feasible set of ribbons \mathcal{R} because it imposes an ordering on which ribbons to optimize for. We believe that this is especially true in spatially constrained environments. The simultaneous optimization strategy, although more expensive, is able to resolve conflicts by jointly optimizing over all the ribbons.

5.3 Discussion

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.

In the first scenario, as shown in Fig. 4, 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. 4(a)). The local optimization method is able to optimize the ribbon to generate one with zero curvature and torsion (Fig. 4(b)).

In the second scenario, as shown in Fig. 5, 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. 5(b)). The RRT planner is able to resolve collisions but generates a sub-optimal ribbon (Fig. 5(a)). However, applying local optimization to the RRT-generated ribbon generates a high quality ribbon that is able to avoid collisions with obstacles (Fig. 5(c)).




Fig. 5. 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. 6. 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.

In the third scenario, as shown in Fig. 6, 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. 6(a)). Local optimization in isolation is unable to find a feasible solution. However, applying local optimization to the RRT-generated ribbon is able to locally optimize the solution (Fig. 6(b)). However, 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 [20], 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.

6 Designing 3D Implants for Intracavitary Brachytherapy

We consider a scenario where a 3D printed implant is used for treatment of OB/GYN tumors, as shown in Fig. 7. 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. [14]. 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 [9] in diameter. In this scenario, we consider channels of width $w = 2.5$ mm. We impose a constraint on




Fig. 7. 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.

the instantaneous curvature of $\bar{\kappa} = 1\text{cm}^{-1}$ based on the maximum allowable curvature reported by Garg et al. [14] for radioactive sources that are 1 mm in diameter and 5 mm in length. We also impose 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 constrain the cross-section of each ribbon at the respective final time steps to lie within the specified entry region.

We implemented our planning approach in C++ and used the Bullet collision checking library [7] for collision checking queries. Fig. 7(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. 7(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 volume.

Table 1 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.

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 1. 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.




Fig. 8. Generating individual channels for groups of candidate dwell dose segments following the approach of Garg et al. [14]. Even with 3 dwell segments per group (18 channels in all), 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 and can also compromise the structural integrity of the implant.

We also compared the use of using ribbon-like arrangements versus generating single channels for groups of dose dwell segments. Fig. 8 shows the arrangement of individual channels for groups of 3 dwell segments, beyond which the planner is unable to compute feasible solutions. The use of single channels leads to inefficient utilization of space within the implant volume, as shown in Fig. 7(b).

The use of single channels also limits the number of reachable dose dwell segments. This affects the coverage of tumor volume that can receive radiation, thus limiting the treatment effectiveness. We used the coverage metric proposed by Garg et al. [14] to compute coverage of the tumor volume using reachable candidate dose dwell segments. Each dose dwell segment is partitioned into dwell positions for the radioactive source in intervals of 5 mm, which corresponds to the length of the source. Any point inside a tumor is then said to be *covered* if a dwell position lies within a ball of ϵ radius. A smaller coverage radius (ϵ) parameter results in lower dose to healthy organs while supplying sufficient radiation to the tumor volumes. Hence a higher tumor coverage with small ϵ is preferred for brachytherapy treatment.

In this scenario, we found that dwell positions generated with ribbon-like channel arrangements (Fig. 7(b)) achieved 100% coverage of the tumor volume for $\epsilon = 2.1$ cm. In contrast, the coverage achieved for single channels (Fig. 8) for $\epsilon = 2.1$ cm is only 54%. Achieving 100% coverage is important to prevent cancer recurrence. Using ribbon-like arrangements allows us to achieve this coverage while minimizing damage to surrounding healthy tissue. We refer the reader to Garg et al. [14] for additional details on the coverage metric.

7 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 can 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 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, such as bundles or tubes instead of ribbons will also be considered. 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 and formations of aerial vehicles.

Acknowledgments

This research has been funded in part by AFOSR-YIP Award #FA9550-12-1-0345, by NSF under award IIS-1227536, by a DARPA Young Faculty Award #D13AP00046, CITRIS Seed Grant, and by a Sloan Fellowship.

References

1. Alterovitz, R., Siméon, T., Goldberg, K.: The stochastic motion roadmap: A sampling framework for planning with markov motion uncertainty. In: Robotics: Science and Systems (RSS) (2007)
2. Belta, C., Kumar, V.: Optimal motion generation for groups of robots: a geometric approach. *Journal of Mechanical Design* 126(1), 63–70 (2004)
3. Biggs, J., Holderbaum, W.: Planning rigid body motions using elastic curves. *Mathematics of Control, Signals, and Systems* 20(4), 351–367 (2008)
4. Bishop, R.L.: There is more than one way to frame a curve. *American Mathematical Monthly* pp. 246–251 (1975)
5. Bretl, T., McCarthy, Z.: Quasi-static manipulation of a Kirchhoff elastic rod based on a geometric analysis of equilibrium configurations. *Int. Journal of Robotics Research* 33(1), 48–68 (2014)
6. Choset, H.: Coverage for robotics—A survey of recent results. *Annals of mathematics and artificial intelligence* 31(1-4), 113–126 (2001)
7. Coumans, E.: Bullet Collision Detection and Physics Library. Available: <http://bulletphysics.org> (2013)
8. Cripps, R., Mullineux, G.: Constructing 3D motions from curvature and torsion profiles. *Computer-Aided Design* 44(5), 379–387 (2012)
9. Devlin, P.: Brachytherapy: Applications and techniques. Lippincott Williams & Wilkins (2007)
10. Duan, Y., Patil, S., Schulman, J., Goldberg, K., Abbeel, P.: Planning locally optimal, curvature-constrained trajectories in 3D using sequential convex optimization. In: Proc. Int. Conf. Robotics and Automation (ICRA). p. to appear (2014)
11. Duindam, V., Xu, J., Alterovitz, R., SastrySastry, S., Goldberg, K.: Three-dimensional motion planning algorithms for steerable needles using inverse kinematics. *Int. Journal of Robotics Research* 29(7), 789–800 (2010)
12. Farin, G.E.: Curves and Surfaces for Computer-Aided Geometric Design: A Practical Code. Academic Press, Inc. (1996)
13. Farmer, T.: A new model for ribbons in \mathbb{R}^3 . *Mathematics Magazine* 79(1), 31 (2006)
14. Garg, A., Patil, S., Siau, T., Cunha, J.A.M., Hsu, I.C., Abbeel, P., Pouliot, J., Goldberg, K.: 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, pp. 3306–3312 (2013)
15. Goemans, O., Overmars, M.: Automatic generation of camera motion to track a moving guide. In: Algorithmic Foundations of Robotics VI, pp. 187–202. Springer (2005)

16. Goriely, A., Shipman, P.: Dynamics of helical strips. *Physical Review E* 61(4), 4508 (2000)
17. Han, J.C., Datta, S., Ekkad, S.: Gas turbine heat transfer and cooling technology. CRC Press (2013)
18. Hanson, A.J.: Quaternion Frenet frames: Making optimal tubes and ribbons from curves. Tech. rep., Tech. Rep. 407, Indiana University Computer Science Department (1994)
19. Krontiris, A., Louis, S., Bekris, K.E.: Simulating formations of non-holonomic systems with control limits along curvilinear coordinates. In: *Motion in Games*. pp. 121–133 (2010)
20. LaValle, S.: Planning algorithms. Cambridge University press (2006)
21. Lipson, H., Kurman, M.: Fabricated: The new world of 3D printing. John Wiley & Sons (2013)
22. Moll, M., Kavraki, L.E.: Path planning for deformable linear objects. *IEEE Trans. on Robotics* 22(4), 625–636 (2006)
23. Murray, R.M., Shankar, S.S.: A mathematical introduction to robotic manipulation. CRC press (1994)
24. Nocedal, J., Wright, S.: Numerical Optimization. Springer Verlag (2006)
25. Park, W., Wang, Y., Chirikjian, G.: The path-of-probability algorithm for steering and feedback control of flexible needles. *Int. Journal of Robotics Research* 29(7), 813–830 (2010)
26. Patil, S., Burgner, J., Webster III, R.J., Alterovitz, R.: Needle steering in 3D via rapid replanning. *IEEE Trans. on Robotics* p. to appear (2014)
27. Saccon, A., Hauser, J., Aguiar, A.P.: Optimal control on Lie groups: The projection operator approach. *IEEE Transactions on Automatic Control* 58(9), 2230–2245 (2013)
28. Schulman, J., Ho, J., Lee, A., Bradlow, H., Awwal, I., Abbeel, P.: Finding locally optimal, collision-free trajectories with sequential convex optimization. In: *Robotics: Science and Systems (RSS)* (2013)
29. Selig, J.: Characterisation of frenet–serret and bishop motions with applications to needle steering. *Robotica* 31(06), 981–992 (2013)
30. Shkolnik, A., Walter, M., Tedrake, R.: Reachability-guided sampling for planning under differential constraints. In: *Proc. Int. Conf. Robotics and Automation (ICRA)*. pp. 2859–2865 (2009)
31. Sprott, K., Ravani, B.: Kinematic generation of ruled surfaces. *Advances in Computational Mathematics* 17(1-2), 115–133 (2002)
32. Swensen, J.P., Cowan, N.J.: Torsional dynamics compensation enhances robotic control of tip-steerable needles. In: *Proc. Int. Conf. Robotics and Automation (ICRA)*. pp. 1601–1606 (2012)
33. van den Berg, J., Patil, S., Alterovitz, R., Abbeel, P., Goldberg, K.: LQG-based planning, sensing, and control of steerable needles. In: *Proc. Workshop Algorithmic Foundations of Robotics (WAFR)*. pp. 373–389 (2010)
34. Wang, W., Jüttler, B., Zheng, D., Liu, Y.: Computation of rotation minimizing frames. *ACM Transactions on Graphics (TOG)* 27(1), 2 (2008)
35. Webster III, R.J., Kim, J.S., Cowan, N.J., Chirikjian, G.S., Okamura, A.M.: Nonholonomic modeling of needle steering. *Int. Journal of Robotics Research* 25(5-6), 509–525 (2006)
36. Willemsen, P., Kearney, J.K., Wang, H.: Ribbon networks for modeling navigable paths of autonomous agents in virtual environments. *IEEE Transactions on Visualization and Computer Graphics* 12(3), 331–342 (2006)
37. Xu, J., Duindam, V., Alterovitz, R., Pouliot, J., Cunha, J.A., Hsu, I., Goldberg, K.: Planning fireworks trajectories for steerable medical needles to reduce patient trauma. In: *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*. pp. 4517–4522 (2009)
38. Zefran, M., Kumar, V., Croke, C.B.: On the generation of smooth three-dimensional rigid body motions. *IEEE Transactions on Robotics and Automation* 14(4), 576–589 (1998)