

Towards Optimizing a Convex Cover of Collision-Free Space for Trajectory Generation

Yuwei Wu, Igor Spasojevic, Pratik Chaudhari, Vijay Kumar

Abstract—We propose an online iterative algorithm to optimize a convex cover to under-approximate the free space for autonomous navigation to delineate Safe Flight Corridors (SFC). The convex cover consists of a set of polytopes such that the union of the polytopes represents obstacle-free space, allowing us to find trajectories for robots that lie within the convex cover. In order to find the SFC that facilitates trajectory optimization, we iteratively find overlapping polytopes of maximum volumes that include specified waypoints initialized by a geometric or kinematic planner. Constraints at waypoints appear in two alternating stages of a joint optimization problem, which is solved by a novel heuristic-based iterative algorithm with partially distributed variables. We validate the effectiveness of our proposed algorithm using a range of parameterized environments and show its applications for two-stage motion planning.

Index Terms—Collision Avoidance; Computational Geometry; Aerial Systems; Applications

I. INTRODUCTION

ALL motion and path planning algorithms require the representation of the environment for collision avoidance. Different algorithms choose different representations. Exact methods map obstacles from the workspace into the configuration space, this gives guarantees of completeness and optimality but may be computationally intractable [1]. Sampling-based methods do not need to represent the environment explicitly but they only offer probabilistic guarantees and may still require an excessively large number of samples in cluttered high-dimensional spaces [2]. A middle ground has become popular over the past few years, including convex approximation of the collision-free space. This is computationally attractive because it allows us to convert motion planning into an efficient optimization problem. This problem is also tractable because collision-free space can be described by a set of linear constraints corresponding to each of the polytopes in the cover.

One of the key issues in this approach pertains to computing a good approximation of the environment. One could find a *convex cover* [3] by a collection of convex sets such that the union of the convex sets can cover a subset of the collision-free region or include a given guess for trajectories. Solving

Manuscript received: November 6, 2024; Revised: January 29, 2025; Accepted: March 3, 2025. This paper was recommended for publication by Editor Giuseppe Loianno upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by TILOS funded by NSF grant CCR-2112665, USDA/NIFA 2022-67021-36856, IoT4Ag ERC funded through NSF grant EEC-1941529, and ONR grant N00014-20-S-B001. (Corresponding author: Yuwei Wu.)

The authors are with the GRASP Laboratory, University of Pennsylvania, Philadelphia, PA 19104 USA {yuweiwu, igorspas, pratikac, kumar}@seas.upenn.edu.

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE




Fig. 1. The influence of convex cover on trajectory generation. (a) Seed points (blue) determine the volume of the largest polytope (red) within a map (obstacles in grey). The largest-volume polytope can be identified by selecting the seed point at the orange star. (b) Polytopes may not always cover a large part of the trajectory, even if they have a large volume. (c, d) Given an initial path (black), the quality of the final trajectory (green) depends strongly upon which convex cover was used to optimize it.

the subproblem of extracting a *single convex region* from a seed point or convex body is usually involved in generating the convex cover [4]–[8]. Considering the example in Fig. 1(a), an ideal single convex region should capture most of the free space. There should also be enough overlaps between different polytopes (ellipsoids are another popular choice) in the convex cover, at least successive polytopes as a “Safe Flight Corridor” (SFC). The properties of the SFC will influence the ultimate geometric feasible region for optimization. Many existing approaches generate an SFC by selecting seed points around the initial path [9]–[11]. However, such covers are not ideal in practice because the selection of SFC is greedy and does not account for trajectory planning. It is difficult to determine how the geometric shapes and volumes of the polytopes interplay with the dynamical feasibility of the trajectories computed by the lower-level optimization procedure. Fig. 1(b)-(d) shows some examples of potential factors like polytope volumes and overlapping areas that can potentially change the final trajectory. As a result, the existing pipelines (illustrated at the top part of Fig. 2) cannot give any guarantees on the quality of the SFC or the trajectory. This suggests that coupled program: front-end path planning, free space approximation using a convex cover, and trajectory generation within this cover, cannot be solved independently.




Fig. 2. SFC generation pipeline of previous works and our method. Given waypoints from the path planning, the SFC generator selects seed points or lines from waypoints and outputs overlapping polytopes. The proposed method novelly formulates this process as a joint optimization problem to generate a better back-end trajectory downstream.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

To address the problem of SFC generation, we introduce a novel algorithm that optimizes the convex cover of the environment to aid in trajectory planning for quadrotors (and other mobile robots with dynamic constraints). It updates the polytopes that form the convex cover and maximizes the volumes of ellipsoids inscribed within these polytopes given the waypoints updated by using heuristic functions to approximate trajectory cost. These include two alternating stages of a joint optimization problem, which is solved by a heuristic iterative approach with partially distributed variables, inspired by the Alternating Direction Method of Multipliers (ADMM) [12] with partially distributed variables. The pipeline is shown in Fig. 2 in comparison to previous works. This approach is validated using a large variety of synthetic environments to demonstrate its flexibility and robustness.

Contributions: (1) We formulate an SFC generation problem to optimize a convex cover of a subset of the collision-free space. This subset is approximated by the boundaries of trajectories that belong to the same homology class as a given path within the collision-free space for trajectory optimization. (2) We design a novel iterative algorithm based on heuristics with partially distributed variables to maximize the efficiency of the computation. (3) We extensively evaluate our method, benchmark it in environments with varying complexities, and present the first thorough analysis of the influence of the front-end plan path and SFC on the trajectory cost.

II. RELATED WORKS

In real-world representations of geometric environments, obstacles or objects can be decomposed into a series of convex shapes to assist obstacle avoidance strategies. There are traditional methods to solve exact and approximate convex region [13]–[15] or neural networks to learn the convex primitives [16]. However, decomposing obstacles would result in a large number of small pieces that can lead to significant computational costs. Instead, extracting the convex representation of nonconvex free space is more efficient for motion planning problem formulation [11, 17]–[22]. Collision-free corridors or graphs can be generated using boxes [17, 18], sphere or ellipsoids [19, 20] and polytopes [11, 21] within environments represented by point clouds or obstacle maps.

To extract convex regions from nonconvex environments, [4] proposed an algorithm of Iterative Regional Inflation by Semidefinite Programming (IRIS) to generate maximum volume polytopes around convex obstacles given a seed point. The iterative optimization process involved sequentially quadratic programming to find separating hyperplanes toward obstacles, followed by semidefinite programming to determine the maximum-volume ellipsoid (inner *Löwner-John Ellipsoid* [23]) within this polytope. However, considering the expensive computation of this approach, most other works focused on improving the efficiency of the process, such as generating polytopes on line segments [9], or convex bodies [8]. With given voxel maps as inputs, [22] proposed parallel convex inflation of axis-aligned cubes on axis directions by checking the voxel points incrementally, which can achieve resolution near real-time. Moreover, in [24], the

authors introduced a stereographic projection technique to project obstacle points or point clouds and convert them into a convex hull. The generative region is over-conservative for application, which is improved in [7] using sphere flipping to generate star-convex shapes and further refined into convex. The work in [8] improves the polytope inflation and maximum volume inscribed ellipsoid processes in IRIS, providing a more efficient solution through a linear-complexity algorithm. Except for inflating a single convex region within naturally nonconvex environments, determining a convex cover of local or global space remains challenging. Several offline methods face challenges in selecting appropriate seed points to effectively generate a convex cover of the environment [4, 25]. To improve the sampling efficiency, the work [6] leveraged the topology property of the objects and employed skeleton structures to position convex primitives. Furthermore, [21] iteratively conducted sampling to update clique covers of visibility graphs and inflate polytopes. With the convex free-space maps, the trajectory generation can be addressed by mix integer programming [25, 26].

For real-world applications to trajectory optimization, computational efficiency is crucial and is typically achieved by decoupling the problem into subproblems or by narrowing the solution space. A two-stage planning framework, involving front-end path planning with geometric constraints followed by back-end trajectory optimization within the same homotopy class, has been extensively discussed in [27] and applied to SFC-decomposed environments. The authors in [9] proposed a jump point search (JPS) to find a geometric path composed of line segments and initialize ellipsoids on lines with one iteration polytope inflation. For general polytopes generation using seed points or lines, different paths from front-end pathfinding can be applied by iteratively checking the farthest point along the path within the polytope as the next sample points [10, 11]. However, the above methods either fix line segments or employ a greedy approach to find the next waypoint as seeds, potentially resulting in suboptimal trajectories, particularly in narrow hallways.

Unlike previous methods, we introduce a novel approach that optimizes the selection of seeds in polytope generation for better trajectory optimization. Our approach uniquely focuses on decomposing collision-free space to SFC, addressing critical factors such as convex approximation accuracy, trajectory segment division, and connectivity of convex regions between waypoints. This integrated approach ensures that the trajectory optimization process is both optimized and computationally efficient, setting our method apart from existing techniques.

III. PROBLEM STATEMENT

A. Trajectory Optimization Formulation

We consider the problem of finding an optimal energy minimization trajectory for a s -order integrator system, represented by states $x(t)$ and control inputs $u(t)$, with traversal time T that passes through collision-free space \mathcal{F}_{free} . Given a start state $x_0 \in \mathcal{F}_{free}$ and a feasible goal region $\mathcal{F}_{goal} \subseteq$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

\mathcal{F}_{free} , the trajectory optimization problem is formulated as

$$\min_{x(t), T} \int_0^T \|u(t)\|^2 dt + w_t \gamma(T), \quad (1a)$$

$$\text{s.t. } x^{(s)}(t) = u(t), \quad \forall t \in [0, T], \quad (1b)$$

$$x(0) = x_0, \quad x(T) \in \mathcal{F}_{goal}, \quad (1c)$$

$$x(t) \in \mathcal{F}_{free}, \quad \forall t \in [0, T]. \quad (1d)$$

The objective function minimizes a positive combination of the control effort and the total traversal time. $\gamma(\cdot)$ is a time regularization function with weight w_t , that can encode several different requirements such as minimizing the total time, reaching the goal region at a target time, etc. The constraint (1b) represents robot dynamics, and $x^{(s)}$ denotes the s -th derivative of x . The constraints (1c) refer to start and end conditions. This constrained optimization problem is nontrivial to solve due to the non-convexity of collision-free configuration space.

B. Convex Cover Generation

We consider n -dimensional geometric space represented as $\mathcal{X} = \mathcal{X}_{free} \cup \mathcal{X}_{obs} \subseteq \mathbb{R}^n$, where $n \in \{2, 3\}$. \mathcal{X} consists of free space \mathcal{X}_{free} and obstacles \mathcal{X}_{obs} , where \mathcal{X}_{obs} is defined as a union of convex bodies. Several different representations exist for the convex bodies themselves. Typically they are either convex polytopes (bounded intersections of closed half-spaces) or spheres which are commonly used while manipulating raw point-cloud observations of the environment.

In general, for navigation scenarios where there is enough free space in the environment, a conservative yet efficient strategy for preventing collisions involves uniformly inflating the grid map or obstacles and simplifying the robot to a point mass. Such pre-processing allows for trajectory planning directly within the geometric space to generate a collision-free geometrical path. We consider an initial geometric path represented as a concatenation of straight line segments, whose endpoints form the sequence of waypoints $p = (p_0, \dots, p_M)$,

$$\Gamma(p) = \{x \in \mathbb{R}^n \mid x = \theta p_{i-1} + (1 - \theta)p_i, \quad (2)$$

$$\forall \theta \in [0, 1], \quad \forall i = 1, \dots, M\}.$$

With such a collision-free path, for example, output by a path planning algorithm, the convex cover problem can be formulated as finding a sequence of overlapping convex sets in collision-free space that covers the given path. To obtain a parameterized family of convex sets suitable for optimization, we begin by defining a collision-free convex polytope using \mathcal{H} -representation $\mathcal{P}(A, b) = \{x \in \mathbb{R}^n \mid Ax \leq b\}$, with $A \in \mathbb{R}^{v \times n}$ and $b \in \mathbb{R}^v$. We employ a sequence of overlapping polytopes for covering a subset of the collision-free space that contains the class of paths homotopic to the original path. Hence, the convex cover of a given path is to generate an SFC such that

$$\min_{\mathcal{P}} \sum_{i=1}^M f(\mathcal{P}_i), \quad (3a)$$

$$\text{s.t. } \Gamma \in \bigcup_{i=1}^M \mathcal{P}_i, \quad (3b)$$

$$\mathcal{P}_i \cap \mathcal{P}_{i+1} \neq \emptyset, \quad \forall 1 \leq i \leq M - 1. \quad (3c)$$

We represent sequential polytopes $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_M)$, with the i -th polytope \mathcal{P}_i parameterized by $A_i = [a_{i,1}^T, \dots, a_{i,v_i}^T]^T$ and $b_i = [b_{i,1}, \dots, b_{i,v_i}]^T$. The function $f(\cdot)$ defines an objective function to find the *best* SFC. Examples of possible criteria encoded by f include the volume of polytopes $f(\mathcal{P}_i) = -\text{vol}(\mathcal{P}_i)$, or the volume of the inner Löwner-John Ellipsoid of each polytope. Nevertheless, while previous classes of criteria are intuitively appealing, it is not obvious how they lead toward synthesizing better trajectories. In the rest of the paper, we will discuss the factors and properties of SFC that are essential for effective trajectory generation.

C. Trajectory Optimization with SFC

With different representations of the dynamics and trajectories, we can formulate different optimization problems while minimizing control efforts. We will discuss one type of back-end optimization method that is applied to provide a heuristic cost function in our algorithm and thus can be extended to any back-end planner. If a trajectory is represented in flat-output space [28], we can further characterize it by $2s - 1$ degree n -dimensional piecewise polynomials

$$\sigma(t) = \{\sigma_i(t) \mid t \in [\sum_{k=1}^{i-1} \Delta t_k, \sum_{k=1}^i \Delta t_k], \quad (4)$$

$$\tau = [\Delta t_1, \dots, \Delta t_M]^T\}.$$

The $\tau \in \mathbb{R}_{>0}^M$ represents a time interval vector with a total traversal time $\mathbf{1}^T \tau = T$, where $\mathbf{1}$ is a one-vector of size M . The time-optimal minimum control trajectory optimization with approximate convex cover can be formulated as a parametric nonlinear programming problem,

$$\min_{\sigma, \tau} J(\sigma, \tau) = \int_0^{\mathbf{1}^T \tau} \|\sigma^{(s)}(t)\|_2^2 dt + w_t \gamma(\tau), \quad (5a)$$

$$\text{s.t. } \sigma_1^{[s-1]}(0) = \bar{p}_0, \quad \sigma_M^{[s-1]}(\Delta t_M) = \bar{p}_f, \quad (5b)$$

$$\sigma_i^{[s-1]}(\Delta t_i) = \bar{p}_i, \quad \forall i = 1, \dots, M, \quad (5c)$$

$$A_i \sigma_i(t) \leq b_i, \quad \forall i = 1, \dots, M, \quad \forall t \in [0, \Delta t_i], \quad (5d)$$

where $\bar{p}_0, \bar{p}_i, \bar{p}_f \in \mathbb{R}^{n \times s}$ are initial, intermediate and end derivative vectors, $\sigma^{[s-1]} = (\sigma(t), \dots, \sigma(t)^{(s-1)})$. As discussed in [29], the piecewise polynomial trajectory can be parameterized as $\sigma(t) = \sigma(p, \tau)(t)$, where $p = (\bar{p}_0, \dots, \bar{p}_i, \dots, \bar{p}_f)$. For brevity, we will use p and τ to represent the parameters that determine the polynomial trajectory. Trajectory optimization with piece-wise polynomial reduces the complexity of dynamics and scaling issues of discretizations, whereas introducing nonlinearity in objectives, and constraints coupled with coefficients, times, and SFC variables. Previous research in piece-wise trajectory with corridors mainly focuses on time allocation [28, 30, 31], while the optimality and feasibility of convex approximation of environments are rarely discussed.

IV. PROPOSED ALGORITHMS

In a single convex decomposition problem, an ellipsoid is typically coupled with the polytope during optimization [4].

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

This is partly because the volume of the polytope is highly correlated with the volume of its inscribed ellipsoid, and optimizing the volume of the latter offers computational advantages. To find the trajectory-optimal SFC for covering paths in the same homotopy equivalence class, we introduce a sequence of intersecting ellipsoids to encode a set of valid deformations of the current guess of the optimal path. For computational efficiency, we use the Cholesky representation of the ellipsoid and define it as $\mathcal{E}(L, d) = \{x \in \mathbb{R}^n \mid \|L^{-T}(x - d)\|_2 \leq 1\}$, where $L \in \mathbb{R}^{n \times n}$ is lower triangular with strictly positive diagonal entries, and $d \in \mathbb{R}^n$ is the center of the ellipsoid. Our goal is to jointly find a geometric path together with a sequence of covering ellipsoids that simultaneously optimize the trajectory and the volume of the ellipsoids that certify it lies in free space. The problem is therefore formulated as

$$\min_{L_i, d_i, p_i, \forall i} w_v \tilde{f}(L_i) + w_c \tilde{J}(p, \tau), \quad (6a)$$

$$\text{s.t. } \|L_i^{-T}(p_i - d_i)\|_2 \leq 1, \quad \forall i = 1, \dots, M \quad (6b)$$

$$\|L_i^{-T}(p_{i-1} - d_i)\|_2 \leq 1, \quad \forall i = 1, \dots, M \quad (6c)$$

$$h_1(p, \tau, \mathcal{P}) \leq 0, \quad (6d)$$

$$h_2(\mathcal{E}(L, d), \mathcal{P}) \leq 0. \quad (6e)$$

The objectives are decoupled into two parts: the approximation of the polytope volume in Eq.(3a) as the volume of inner ellipsoid $f(L)$, and an approximation of trajectory cost in Eq.(5a) as $\tilde{J}(p, \tau)$. These terms are weighted by w_v and w_c respectively. The constraints (6b) and (6c) encode the requirement that every pair of consecutive waypoints lie in the corresponding ellipsoid. To rewrite the latter pair of constraints in a more streamlined way, we introduce the function $h_0(\mathcal{E}, p) = \|L^{-T}(p - d)\|_2 - 1$. The functions $h_1(\cdot), h_2(\cdot)$ in (6b) and (6c) represent separate conditions for ellipsoids and trajectory, which are further elaborated in section IV-B and section IV-C. We adopt the notation $\mathcal{E} = (\mathcal{E}_1, \dots, \mathcal{E}_M)$, where each \mathcal{E}_i is parameterized by L_i and d_i . We reformulate the problem as follows:

$$\min_{\mathcal{E}, p} w_v \sum_{i=1}^M \tilde{f}(\mathcal{E}_i) + w_c \tilde{J}(p, \tau), \quad (7a)$$

$$\text{s.t. } g(h_0(\mathcal{E}_i, p_i)) = 0, \quad \forall i = 1, \dots, M, \quad (7b)$$

$$g(h_0(\mathcal{E}_i, p_{i-1})) = 0, \quad \forall i = 1, \dots, M, \quad (7c)$$

$$h_1(p, \tau, \mathcal{P}) \leq 0, \quad (7d)$$

$$h_2(\mathcal{E}, \mathcal{P}) \leq 0, \quad (7e)$$

where $g(\cdot) = \max\{0, \cdot\}^2$ for coupled conditions of ellipsoid \mathcal{E} and waypoint p variables. The independent constraints (7d) and (7e) are complex and may introduce additional variables such as trajectory time intervals and polytopes. Therefore, we incorporate the coupled constraints and formulate the scaled-form augmented Lagrangian as

$$\begin{aligned} \mathcal{L}(\mathcal{E}, p, y) = & w_v \sum_{i=1}^M \tilde{f}(\mathcal{E}_i) + w_c \tilde{J}(p, \tau) - \frac{\rho}{2} \|y\|_2^2 \\ & + \sum_{i=1}^M \frac{\rho}{2} \|[g(h_0(\mathcal{E}_i, p_i)), g(h_0(\mathcal{E}_i, p_{i-1}))]\|^T + y_i \end{aligned} \quad (8)$$

where $y = (y_1, \dots, y_M)$ are Lagrange multipliers, $\rho > 0$ is a parameter. As our problem consists of decoupled subproblems, including convex region inflation, ellipsoid generation, and trajectory generation, multi-level decomposition [12, 32]–[34] provides a suitable approach for efficiently solving the original problem. To solve the problem, we propose an iterative algorithm, detailed in Alg. (1), that leverages a scaled-form ADMM [12] in a partially distributed manner. We relax the constraints on independent variables of subproblems for efficient updates. The input variables include an initial geometric path and obstacles, while the output is SFC represented by a sequence of polytopes. The constraints $h_i(\cdot)$ would be satisfied during each sub-problem. Additionally, the cost function and constraints for ellipsoid generation are separable and can be updated in parallel.

Algorithm 1: Given an initial geometric path and obstacles, find an SFC to cover the path while avoiding all the obstacles.

Input: path p^0 , obstacles \mathcal{O}

Output: SFC \mathcal{P}

$\mathcal{E} \leftarrow \text{InitEllipsoids}(p^0)$

repeat

$k \leftarrow 0, \mathcal{P} \leftarrow \text{InflatePolytopes}(\mathcal{E}, \mathcal{O})$

repeat

$p^{k+1} := \arg \min_p \{\mathcal{L}(\mathcal{E}^k, p, y^k) \mid h_1(p, \tau, \mathcal{P}) \leq 0\}$

$\mathcal{E}_i^{k+1} := \arg \min_{\mathcal{E}_i} \{\mathcal{L}(\mathcal{E}_i, p^k, y_i^k) \mid h_2(\mathcal{E}_i, \mathcal{P}) \leq 0\}, \forall i = 1, \dots, M,$

$y_i^{k+1} := y_i^k + [g(h_0(\mathcal{E}_i^{k+1}, p_i^{k+1})), g(h_0(\mathcal{E}_i^{k+1}, p_{i-1}^{k+1}))]^T, \forall i = 1, \dots, M,$

$k \leftarrow k + 1$

until $k = k_{max};$

$\mathcal{E} \leftarrow \mathcal{E}^k$

until end condition;

A. Initialization and Polytopes Generation

Given an initial path p^0 , we upsample waypoints by adding auxiliary intermediate waypoints between consecutive original waypoints to avoid creating excessively long ellipsoids that exceed a specified distance threshold α . We use the procedure **InitEllipsoids** that takes as input line segments of the initial path and outputs the ellipsoid centered at $d_i = (p_i + p_{i-1})/2$. Let $\mu_i = \frac{p_i - p_{i-1}}{\|p_i - p_{i-1}\|_2}$ be the unit vector along the displacement from waypoint p_{i-1} to waypoint p_i . Then L_i is chosen so that

$$L_i L_i^T = \mu_i \mu_i^T \frac{\|p_i - p_{i-1}\|_2^2}{4} + (I - \mu_i \mu_i^T) \epsilon^2. \quad (9)$$

We assume the initial path has ϵ -clearance from obstacles; this implies that the constructed ellipsoids are collision-free. In **InflatePolytopes**, to find a maximum volume of polytopes in \mathcal{X}_{free} , we reformulate the optimization problem (3) as

$$\min_{\mathcal{P}} \sum_{i=1}^M -\text{vol}(\mathcal{P}_i), \quad (10a)$$

$$\text{s.t. } \mathcal{E}_i \subset \mathcal{P}_i, \mathcal{P}_i \cap \mathcal{O} = \emptyset \quad \forall i = 1, \dots, M. \quad (10b)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

We adopt a method similar to [4] to solve this problem by finding tangent separating planes from the ellipsoid. To save computation, we define a local bounding box by adding a margin of $\pm l$ around the ellipsoid for each polytope generation. This margin ensures that only the *local* obstacles within this expanded range are considered. The overlap condition $\mathcal{P}_i \cap \mathcal{P}_{i+1} \neq \emptyset$ is satisfied by construction. Additionally, we store the obstacle points near each initialized ellipsoid for parallel updating of polytopes during iteration.

B. Waypoints Update with Trajectory Heuristics

The SFC serves as the environmental approximation for trajectory optimization, therefore the trajectory cost is the key factor during this phase. For back-end optimization, minimum control effort (energy), distance, and traversal time are three main factors. The optimal control trajectory $\sigma(t)$ is therefore encoded through its waypoints and time intervals $\{p, \tau\}$, which formulates the problem

$$\begin{aligned} \min_{p, \tau} w_c \tilde{J}(p, \tau) + & \quad (11a) \\ \sum_{i=1}^M \frac{\rho}{2} \|[g(h_0(\mathcal{E}_i^k, p_i)), g(h_0(\mathcal{E}_i^k, p_{i-1}))]\]^T + y_i\|_2^2, & \\ \text{s.t. } h_1(p, \tau, \mathcal{P}) \leq 0. & \quad (11b) \end{aligned}$$

Incorporating the energy cost introduces the dynamic model of the robot, which is complex to accurately derive and expensive to optimize simultaneously during SFC generation. In the rest of the section, we will discuss the approximation of the objective \tilde{J} through minimum path length (denoted as \tilde{J}_1) and minimum control effort (\tilde{J}_2) for updating waypoints.

1) *Minimum Path Length*: Without the consideration of dynamic feasibility, distance serves as a common criterion for geometric pathfinding. We employ Euclidean distance between waypoints to find the shortest path inside the corridors and ellipsoids, as

$$\begin{aligned} \min_p w_c \sum_{i=1}^M \|p_i - p_{i-1}\|_2 + & \quad (12a) \\ \sum_{i=1}^M \frac{\rho}{2} \|[g(h_0(\mathcal{E}_i^k, p_i)), g(h_0(\mathcal{E}_i^k, p_{i-1}))]\]^T + y_i\|_2^2, & \\ \text{s.t. } p_i \in \mathcal{P}_i \cap \mathcal{P}_{i+1}, \forall i = 1, \dots, M-1, & \quad (12c) \end{aligned}$$

where constraint (12c) ensures waypoints are generated within the overlapping region of every two consecutive polytopes.




Fig. 3. Polynomial trajectories can take on various shapes by changing time allocations.

2) *Minimum Control Effort*: The time allocation significantly influences the control efforts and the shape of the polynomials, shown in Fig. 3. Including time allocation in optimization will also introduce nonlinear constraints. In addition, time allocation is also correlated to dynamic limits like maximum acceleration and velocities, which need to be specified. However, in rest-to-rest trajectory planning, one can always adjust a time factor to re-scale the time allocation to push the trajectory inside corridors. In this case, we assume the rest-to-rest scenarios and solve trajectory optimization with geometric constraints as

$$\begin{aligned} \min_p w_c \int_0^{1^T \tau} \|\sigma(p, \tau)^{(s)}(t)\|_2^2 dt + & \quad (13a) \\ \sum_{i=1}^M \frac{\rho}{2} \|[g(h_0(\mathcal{E}_i^k, p_i)), g(h_0(\mathcal{E}_i^k, p_{i-1}))]\]^T + y_i\|_2^2, & \\ \text{s.t. } M\sigma(p, \tau) = e, & \quad (13c) \\ A_i \sigma_i(p, \tau) \leq b_i, \forall i = 1, \dots, M, & \quad (13d) \end{aligned}$$

where σ_i is the i -th trajectory segment mapping from derivative vectors and time intervals, which has been discussed in [35]. The constraints (13c) include boundary and continuity constraints for polynomials, and constraints (13d) enforces trajectory inside corridors. Practically, we can use a fixed time allocation with a time factor to provide sub-optimal heuristics.

C. Maximum Volume Ellipsoids with Inner&Outer Constraints

We introduce a cost function related to the volume of the convex region that can inflate to increase the potential solution space for trajectory optimization. To maintain the connected constraints for every convex region, we have to ensure the overlapping conditions of polytopes or ellipsoids. For each sub-problem, the goal is to find the maximal volume inscribed ellipsoid of a convex body (polytope) and the volume circumscribed ellipsoid outside a convex body (a line segment or multiple convex bodies). As illustrated in Fig.4, an ellipsoid is constrained inside the polytope with different outer conditions. Given current waypoints p^k , we can formulate the sub-problem




Fig. 4. Demonstration of geometric constraints for an ellipsoid (blue). (a) The general formulation of inner and outer Löwner-John ellipsoid. (b) The Inner Löwner-John ellipsoid with outer line segment constraints.

of finding the i -th ellipsoid $\mathcal{E}_i(L_i, d_i)$ by

$$\min_{L_i, d_i} -w_v \log \det(L_i) + \quad (14a)$$

$$\frac{\rho}{2} \|[g(h_0(\mathcal{E}_i, p_i^k)), g(h_0(\mathcal{E}_i, p_{i-1}^k))]\]^T + y_i\|_2^2, \quad (14b)$$

$$\text{s.t. } \|(a_{i,j})^T L_i\|_2 + (a_{i,j})^T d_i \leq b_{i,j}, \quad j = 1, \dots, v_i \quad (14c)$$

$$L_i \text{ is lower triangular}, \quad (14d)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

where the first term of objective, $\log \det(L_i)$, is proportional to the volume of the ellipsoid, and (14c) represents corridor constraints. We employ a distributed style to update \mathcal{E}^{k+1} .

V. RESULTS

A. Implementation Details

We employ the environments provided in [27] which include both simulated and real-world point clouds for cluttered obstacle maps, maze maps, and real-world maps. We follow their procedure which defines a quantity called the Environmental Complexity Signature (ECS) to generate simulated environments of varying complexities (obstacle structures, densities, etc). All environments are stored as voxel maps. We use the RRT* [2] algorithm from the OMPL library [36] to get the initial geometric path for all methods. We use the L-BFGS algorithm [37] to minimize the augmented Lagrangian for all the constrained optimization problems in our approach and replace the rest hard constraints with penalties. In practice, we run the inner iteration once and update the polytope through several outer iterations. Our approach can be deployed on our custom hardware platform [38] to provide an environmental representation for two-stage motion planning.¹

B. Numerical Analysis

We conduct several tests in a parameterized environment of size $20 \times 20 \times 5 \text{ m}^3$ with randomly chosen start and end points, at least 10 m apart. All obstacles are inflated by the radius of the robot. We use total volumes of ellipsoids ($\text{vol}(\mathcal{E})$) and polytopes ($\text{vol}(\mathcal{P})$), and overlapping volumes of each neighboring pair of polytopes ($\text{vol}(\cap \mathcal{P})$) to showcase the performance of the proposed algorithm during several iterations. For evaluation among different test cases, the values were normalized by the maximum values observed during the iterations of a trial.




Fig. 5. Visualization of ellipsoids (blue) and polytopes (red) during 4 iterations. The black line segments represent the path, with waypoints in magenta. The green curve is the trajectory generated within polytopes.

1) *Ellipsoids and Polytopes over Iterations.*: To validate the effectiveness of our proposed method, we show one example using heuristics \tilde{J}_1 and visualize the ellipsoids and polytopes of four iterations in Fig.5. The center and orientation (or axes) impact the maximum volume an ellipsoid can inflate, and our proposed method is able to adjust it with iterations. Fig.6 illustrates the actual volumes of additional tests. The




Fig. 6. The variance in normalized volumes of ellipsoids and polytopes after 20 iterations of our proposed algorithm (across ten trials).

normalized volumes of ellipsoids (left) and polytopes (right) consistently grow with each iteration and contain more than 90% of the maximum volumes within 5-10 iterations in most cases. In what follows, we will cap the number of iterations between 5 and 10; we empirically observed that this comes without sacrificing performance.




Fig. 7. Comparison of performance by trajectory heuristics based on minimum path length (\tilde{J}_1) and minimum control effort (\tilde{J}_2) over 100 test cases.

2) *The Effects of Cost Heuristics.*: We compare two trajectory heuristics with minimum path length and minimum control effort (using jerk) for waypoint updates and show the dependence of normalized volumes of ellipsoids, polytopes, and overlapping areas of the convex cover with the number of iterations. We use a uniform time allocation proportional to the path length in minimum control effort to ensure that the waypoints are distributed uniformly as they move and that the cost term is updated consistently. The performances using different trajectory heuristics costs are shown in Fig.7. The cost \tilde{J}_2 performs slightly better than \tilde{J}_1 , regarding the coverage rate of ellipsoids and polytopes volumes. The cost value of distance has less margin to decrease as the initial path from RRT* has already minimized the distance to a large extent.

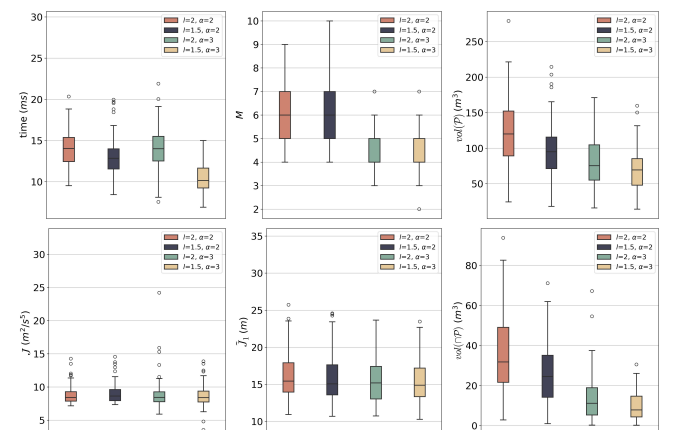


Fig. 8. Performance by different settings with $l \in \{1.5, 2\}$ and $\alpha \in \{2, 3\}$.

¹We will open-source at https://github.com/KumarRobotics/kr_opt_sfc.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

3) Configurable Performance in Different Environments:

We leverage parameters such as local range l and upsample threshold α to generate SFC with specified accuracy and efficiency. Our analysis considers the following four setups with a local range of 1.5 m and 2.0 m and an upsample threshold of 2.0 m and 3.0 m. The upsample threshold influences the number of segments, and the local range determines the number of obstacles considered for each polytope generation. These settings are evaluated in terms of computation time, trajectory cost (J), number of segments (M), path length (\tilde{J}_1), volumes of polytopes, and overlapping areas. Trajectory cost is evaluated by optimizing time intervals and waypoints in [39]. As shown in Fig.8, the computation time can be reduced by decreasing the local range and increasing the upsample threshold. The path length, polytope volumes, overlapping volumes, and trajectory cost also decrease accordingly. We conclude that within the same homotopic class, the optimal trajectory is indeed influenced by the convex cover.

C. Comparison of Convex Cover

We compare our proposed algorithm with two state-of-the-art SFC generation methods, IRIS [4] and the method presented in [9], denoted as RILS. IRIS was originally designed for graph cover, and we use its applications for a convex cover of a path with incremental point collision checking. The iteration number of IRIS is set as the major iteration in their paper, the same as our proposed method. To get inner Löwner-John ellipsoids in IRIS, we directly employ the Cholesky representation for optimization in [8]. The local range for considering obstacles within the bounding box is set to $l = 2$ m to optimize computational efficiency. We first evaluate the performance of three methods across three types of randomly sampled maps generated from [27]. We conduct 1,000 trials for each type of environment, with start and end points randomly selected, ensuring a minimum distance of 10 meters between them. As shown in Table I, our method outperforms other benchmark methods for generating better trajectories, specifically in unstructured environments. We visualize one




Fig. 9. The illustration of SFC generated by three methods. More comparisons are available at <https://youtu.be/k7C17-fgSXE>

It is clear that the performance of any algorithm is influenced by the complexity of both the environment and the task, which makes it difficult to draw conclusions based on average performance. To further compare with benchmark methods, we conduct 10,000 additional tests in simulated environments characterized by the ECS (density index, cluster index, and structure index). The results are filtered to exclude all invalid cases encountered by any of the methods. We categorize test cases based on the range of the density index to represent environment complexity and show the results in Fig. 10.

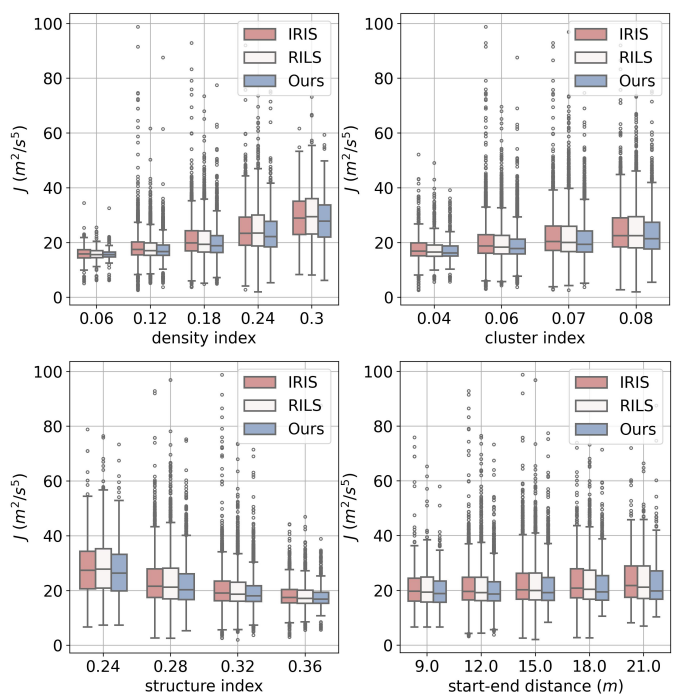


Fig. 10. Benchmark comparison with different environmental parameters, including density index, cluster index, structure index (as defined in [27]), and start-to-goal distance.

Overall, our method achieves lower trajectory costs across various environments. As density and cluster indices increase, environments become more complex, which leads to increased trajectory costs across all three methods. In contrast, more structured environments allow for trajectory generation with lower costs. The comparison results categorized by start-to-end distance (the Euclidean distance between the start and end points), show that our method outperforms others as the distance increases. More importantly, this is the *first systematic analysis of multiple factors that influence the trajectory cost*: the waypoint and number of segments in the front-end plan, the volumes of the polytopes in the convex cover, and the volumes of consecutive polytopes.

TABLE I
COMPARISON ON DIFFERENT TYPES OF MAPS.

Map	Method	J (m^2/s^5)	$vol(\mathcal{P})$ (m^3)	Time (ms)	\tilde{J}_1 (m)
Obstacle	IRIS*	35.39	107.82	45.93	20.17
	RILS	45.76	60.08	11.52	19.37
	Ours	23.83	84.79	24.76	19.66
Maze	IRIS*	33.82	118.38	33.85	20.31
	RILS	33.48	83.93	6.06	18.81
	Ours	33.02	121.79	30.46	19.56
Real	IRIS*	20.53	144.86	17.68	14.04
	RILS	28.28	158.45	3.11	13.80
	Ours	16.45	178.72	19.32	14.04

*IRIS search range is cropped.

trial as an instance, illustrated in Fig. 9. Our method iteratively refines the path for better trajectory generation, particularly in complex environments such as narrow mazes where RILS show poor performance around corners.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

VI. CONCLUSION AND FUTURE WORK

This paper introduces a novel algorithm for optimizing the convex cover of collision-free space as SFC for trajectory optimization. The convex cover is formulated as a joint optimization of ellipsoids, polytopes, and intermediate waypoints, which are updated iteratively to find optimal solutions. The proposed method can be applied to two-stage motion planning as an intermediate process, or optimizing with trajectory generation jointly. Our future work will address whole-body trajectory planning for $SE(2)$ and $SE(3)$.

REFERENCES

- [1] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006. Available at <http://planning.cs.uiuc.edu/>.
- [2] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30:846–894, 2011.
- [3] A. Sarmientoy, R. Murrieta-Cidz, and S. Hutchinson. A sample-based convex cover for rapidly finding an object in a 3-d environment. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3486–3491, 2005.
- [4] Robin Deits and Russ Tedrake. Computing large convex regions of obstacle-free space through semi-definite programming. In *Algorithmic Foundations of Robotics XI: Selected Contributions of the Eleventh International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 109–124, 2015.
- [5] Fei Gao and Shaojie Shen. Online quadrotor trajectory generation and autonomous navigation on point clouds. In *2016 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 139–146, 2016.
- [6] Mukulika Ghosh, Shawna Thomas, and Nancy M. Amato. Fast collision detection for motion planning using shape primitive skeletons. In Marco Morales, Lydia Tapia, Gildardo Sánchez-Ante, and Seth Hutchinson, editors, *Algorithmic Foundations of Robotics XIII*, pages 36–51, Cham, 2020. Springer International Publishing.
- [7] Xingguang Zhong, Yuwei Wu, Dong Wang, Qianhao Wang, Chao Xu, and Fei Gao. Generating large convex polytopes directly on point clouds. *arXiv preprint arXiv:2010.08744*, 2020.
- [8] Qianhao Wang, Zhepei Wang, Mingyang Wang, Jialin Ji, Zhichao Han, Tianyue Wu, Rui Jin, Yuman Gao, Chao Xu, and Fei Gao. Fast iterative region inflation for computing large 2-d/3-d convex regions of obstacle-free space, 2024.
- [9] Sikang Liu, Michael Watterson, Kartik Mohta, Ke Sun, Subhrajit Bhattacharya, Camillo J. Taylor, and Vijay Kumar. Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments. *IEEE Robotics and Automation Letters*, 2(3):1688–1695, 2017.
- [10] Yuwei Wu, Ziming Ding, Chao Xu, and Fei Gao. External forces resilient safe motion planning for quadrotor. *IEEE Robotics and Automation Letters*, 6(4):8506–8513, 2021.
- [11] Zhepei Wang, Chao Xu, and Fei Gao. Robust trajectory planning for spatial-temporal multi-drone coordination in large scenes. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12182–12188, 2022.
- [12] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [13] Bernard M. Chazelle. Convex decompositions of polyhedra. In *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing, STOC '81*, page 70–79, New York, NY, USA, 1981. Association for Computing Machinery.
- [14] Jyh-Ming Lien and Nancy M. Amato. Approximate convex decomposition of polyhedra. In *Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling, SPM '07*, page 121–131, New York, NY, USA, 2007. Association for Computing Machinery.
- [15] Xinyue Wei, Minghua Liu, Zhan Ling, and Hao Su. Approximate convex decomposition for 3d meshes with collision-aware concavity and tree search. *ACM Transactions on Graphics (TOG)*, 41(4):1–18, 2022.
- [16] Boyang Deng, Kyle Genova, Soroosh Yazdani, Sofien Bouaziz, Geoffrey Hinton, and Andrea Tagliasacchi. Cvxnet: Learnable convex decomposition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 31–41, 2020.
- [17] Jing Chen, Tianbo Liu, and Shaojie Shen. Online generation of collision-free trajectories for quadrotor flight in unknown cluttered environments. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1476–1483, 2016.
- [18] Tobia Marcucci, Parth Nobel, Russ Tedrake, and Stephen Boyd. Fast path planning through large collections of safe boxes, 2024.
- [19] Aaron Ray, Alyssa Pierson, and Daniela Rus. Free-space ellipsoid graphs for multi-agent target monitoring. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 6860–6866, 2022.
- [20] Yunfan Ren, Fangcheng Zhu, Wenyi Liu, Zhepei Wang, Yi Lin, Fei Gao, and Fu Zhang. Bubble planner: Planning high-speed smooth quadrotor trajectories using receding corridors. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6332–6339, 2022.
- [21] Peter Werner, Alexandre Amice, Tobia Marcucci, Daniela Rus, and Russ Tedrake. Approximating robot configuration spaces with few convex sets using clique covers of visibility graphs, 2024.
- [22] Fei Gao, Luqi Wang, Boyu Zhou, Xin Zhou, Jie Pan, and Shaojie Shen. Teach-Repeat-Replan: A complete and robust system for aggressive flight in complex environments. *IEEE Transactions on Robotics*, 36(5):1526 – 1545, 2020.
- [23] Fritz John. *Extremum Problems with Inequalities as Subsidiary Conditions*, pages 197–215. Springer Basel, Basel, 2014.
- [24] Sergei Savin. An algorithm for generating convex obstacle-free regions based on stereographic projection. In *2017 International Siberian Conference on Control and Communications (SIBCON)*, pages 1–6, 2017.
- [25] Robin Deits and Russ Tedrake. Efficient mixed-integer planning for uavs in cluttered environments. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 42–49, 2015.
- [26] Tobia Marcucci, Jack Umenberger, Pablo Parrilo, and Russ Tedrake. Shortest paths in graphs of convex sets. *SIAM Journal on Optimization*, 34(1):507–532, 2024.
- [27] Yifei Simon Shao, Yuwei Wu, Laura Jarin-Lipschitz, Pratik Chaudhari, and Vijay Kumar. Design and evaluation of motion planners for quadrotors in environments with varying complexities. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10033–10039, 2024.
- [28] Daniel Mellinger and Vijay Kumar. Minimum snap trajectory generation and control for quadrotors. In *2011 IEEE International Conference on Robotics and Automation*, pages 2520–2525, 2011.
- [29] Zhepei Wang, Hongkai Ye, Chao Xu, and Fei Gao. Generating large-scale trajectories efficiently using double descriptions of polynomials. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7436–7442, 2021.
- [30] Weidong Sun, Gao Tang, and Kris Hauser. Fast uav trajectory optimization using bilevel optimization with analytical gradients. In *2020 American Control Conference (ACC)*, pages 82–87, 2020.
- [31] Yuwei Wu, Xiaotao Sun, Igor Spasojevic, and Vijay Kumar. Deep learning for optimization of trajectories for quadrotors. *IEEE Robotics and Automation Letters*, 9(3):2479–2486, 2024.
- [32] Sanjay Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on optimization*, 2(4):575–601, 1992.
- [33] Gnana Anandalingam and Terry L Friesz. Hierarchical optimization: An introduction. *Annals of Operations Research*, 34:1–11, 1992.
- [34] Mung Chiang, Steven H. Low, A. Robert Calderbank, and John C. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE*, 95(1):255–312, 2007.
- [35] Charles Richter, Adam Bry, and Nicholas Roy. Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments. In *Robotics Research: The 16th International Symposium ISRR*, pages 649–666. Springer, 2016.
- [36] Ioan A. Şucan, Mark Moll, and Lydia E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012. <https://ompl.kavrakilab.org>.
- [37] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- [38] Yuezhan Tao, Yuwei Wu, Beiming Li, Fernando Cladera, Alex Zhou, Dinesh Thakur, and Vijay Kumar. Seer: Safe efficient exploration for aerial robots using learning to predict information gain. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1235–1241, 2023.
- [39] Zhepei Wang, Xin Zhou, Chao Xu, and Fei Gao. Geometrically constrained trajectory optimization for multicopters. *IEEE Transactions on Robotics*, 38(5):3259–3278, 2022.