

UC Berkeley

UC Berkeley Previously Published Works

Title

PolyFly: Polytopic Optimal Planning for Collision-Free Cable-Suspended Aerial Payload Transportation

Permalink

<https://escholarship.org/uc/item/7d64d1xc>

Journal

IEEE Robotics and Automation Letters, 11(3)

ISSN

2377-3766

Authors

Sarvaiya, Mrunal

Li, Guanrui

Loianno, Giuseppe

Publication Date

2026

DOI

10.1109/lra.2026.3655278

Peer reviewed

PolyFly: Polytopic Optimal Planning for Collision-Free Cable-Suspended Aerial Payload Transportation

Mrunal Sarvaiya¹, Guanrui Li², and Giuseppe Loianno¹




Fig. 1: A quadrotor with a suspended payload navigating through cluttered environments. The left and middle images show different perspectives of Env. 3, where the system maneuvers around multiple obstacles. The right image depicts Env. 7, where the system successfully passes through a narrow gap - just wide enough for the cable - demonstrating the planner's ability to exploit the geometry of individual components.

Abstract—Aerial transportation robots using suspended cables have emerged as versatile platforms for disaster response and rescue operations. To maximize the capabilities of these systems, robots need to aggressively fly through tightly constrained environments, such as dense forests and structurally unsafe buildings, while minimizing flight time and avoiding obstacles. Existing methods geometrically over-approximate the vehicle and obstacles, leading to conservative maneuvers and increased flight times. We eliminate these restrictions by proposing PolyFly, an optimal global planner which considers a non-conservative representation for aerial transportation by modeling each physical component of the environment, and the robot (quadrotor, cable and payload), as independent polytopes. We further increase the model accuracy by incorporating the attitude of the physical components by constructing orientation-aware polytopes. The resulting optimal control problem is efficiently solved by converting the polytope constraints into smooth differentiable constraints via duality theory. We compare our method against the existing state-of-the-art approach in eight maze-like environments and show that PolyFly produces faster trajectories in each scenario. We also experimentally validate our proposed approach on a real quadrotor with a suspended payload, demonstrating the practical reliability and accuracy of our method.

Index Terms—Aerial Systems: Applications; Aerial Systems: Mechanics and Control

SUPPLEMENTARY MATERIAL

<https://mrunaljsarvaiya.github.io/polyfly.github.io/>

I. INTRODUCTION

Manuscript received: August 13, 2025; Revised: November 12, 2025; Accepted: December 29, 2025.

This paper was recommended for publication by Editor Soon-Jo Chung upon evaluation of the Associate Editor and Reviewers' comments.

This work was supported by the NSF CPS Grant CNS-2603416, the NSF CAREER Award 2546659, and the DARPA YFA Grant D22AP00156-00.

¹The authors are with the University of California Berkeley, Department of Electrical Engineering and Computer Sciences, Berkeley, CA 94720, USA. email: {mrunaljsarvaiya, loiannog}@eecs.berkeley.edu.

²The author is with the Worcester Polytechnic Institute, Robotics Engineering, Worcester, MA 01609, USA. email: gli7@wpi.edu.

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

MICRO Aerial Vehicles (MAVs), particularly quadrotors equipped with cable-suspended payload systems, have gained significant attention due to their versatility and effectiveness in diverse applications such as inspection [1], search and rescue missions [2], and package transportation [3]. Their ability to deliver essential materials in regions inaccessible by ground robots makes them invaluable during disaster relief operations [4]. Furthermore, aerial transportation systems offer considerable advantages in energy efficiency over ground vehicles for last-mile logistics, significantly reducing environmental footprints in urban delivery contexts [5].

In time critical scenarios, such as rescue missions, it is crucial for aerial vehicles to optimize their flight path to reduce the mission time, especially in challenging environments like dense forests or earthquake-damaged and structurally unsafe buildings. This requires the robot planners to avoid obstacles while maneuvering through narrow gaps and around tight corners, which need non-conservative physical approximations of the system components.

Existing methods often simplify these systems into spheres [6] or large prisms [7], [8] to address computational demands, trading off solution quality for computational speed. In contrast, our approach directly addresses these limitations by representing each component - quadrotor, payload, and cable - as distinct polytopes. We further enhance our planner by incorporating the quadrotor's orientation into the polytopic representation. This orientation-aware modeling enables more precise collision prediction during aggressive maneuvers in confined spaces such as the environment shown in Fig. 5. This representation provides planners with detailed geometric information, enhancing their capabilities to navigate maze-like and highly constrained environments.

Current aerial transportation planning methodologies can broadly be categorized into sampling-based and optimization-based approaches. Within optimization methods, trajectories are either parameterized as polynomials [6], [9] or discretized

and solved via Model Predictive Control (MPC) frameworks [7], [10]–[12]. Our work builds upon MPC-based global planning techniques and specifically addresses the limitations posed by existing approaches that simplify system geometry. Methods that employ Euclidean Signed Distance Fields (ESDFs) result in discontinuous optimization constraints [13] while cylindrical obstacle representations limit the complexity of environments that can be modeled. We tackle these issues by formulating obstacle avoidance explicitly through polytopic models and represent each robot component and every obstacle as distinct polytopes. As we show in this work, this representation allows us to both incorporate smooth differentiable constraints into our global planner, and generate aggressive trajectories in tightly constrained environments.

To summarize, we present the following contributions:

- A novel non-conservative representation for aerial transportation that models the quadrotor, cable, and payload as separate components. To enhance physical accuracy, we incorporate the attitude of both the robot and the obstacles into the environment representation, constructing orientation-aware polytopes that more accurately capture the system’s spatial footprint.
- An optimal global planning method for a robot composed of independent polytopes that produces collision-free trajectories in maze-like environments by leveraging duality theory to formulate non-linear polytopic collision constraints as smooth differentiable constraints.
- Experimental validation on hardware by performing real-world experiments using a quadrotor carrying a slung payload. We illustrate the advantages of our representation and planning method by tracking trajectories in tightly constrained environments. We also demonstrate that our approach generates faster trajectories than the existing state-of-the-art method in all ten test environments.

We release our implementation as an open-source, standalone package with minimal dependencies.

II. RELATED WORKS

In this section, we survey the main design choices in (i) how the environment and robots are spatially represented and (ii) how the states and inputs are mathematically modeled.

Environment Representations. Obstacle-aware planners typically encode the workspace either implicitly, through Euclidean Signed Distance Fields (ESDFs), or explicitly, by representing obstacles with geometric shapes. Recent aerial transportation planners [6], [9], voxelize the environment and store an ESDF, turning collision checks into fast table look-ups, but introduce non-smooth, discontinuous gradients when used in optimization-based solvers.

Other works express the environment via geometric objects. [10]–[12] employ cylindrical and spherical objects, while [7], [8], [14] rely on polytopes. While ellipsoids and cylinders limit the range of scenes that can be modeled, polytopes significantly increase the complexity of environments that can be captured. In fact, works that use ESDFs also primarily consider environments composed of multiple cuboids [9].




Fig. 2: Different polytopes used in PolyFly: **Quadrotor** N_Q , **Payload** N_L , **Cable** N_C , and **Obstacles** M_i . The quadrotor polytope enables accurate distance computations h_{N_Q, M_i} , whereas most methods use a sphere around the quadrotor resulting in conservative distance computations $\tilde{h}_{N_Q, M_i}$.

Furthermore, in the context of trajectory planning, adopting polytopes to represent the environment yields distance constraints with good optimization properties. Duality theory can be used to convert non-linear polytope-to-polytope distance inequalities into smooth, differentiable constraints. [15], [16] used polytopic collision constraints for short-horizon reactive MPC frameworks but are unable to optimize over global paths due to the short MPC horizon. [7] used this approach for global planning but did not consider complex environments and cannot navigate environments through narrow gaps, eg. Fig. 1, due to their conservative representation that models the robot as a single cuboid.

Spatial Robot Representation. There are two primary design decisions with respect to modeling the robot. First, how should the physical geometry of the robot be modeled? Existing methods typically use multiple spheres or ellipses to represent the quadrotor and its payload [6], [9]–[12] or a single prism or polytope [7], [8]. These representations over-approximate at least one component (quadrotor, payload, cable or environment obstacles), therefore limiting their applicability in highly constrained or cluttered environments. Spheres over-approximate the quadrotor, cable, and obstacles while ellipses over-approximate the quadrotors and obstacles used in this work. For example, ellipsoidal representations fail in the ceiling-wall environments in Fig. 7, while single polytope representations [7] fails in the ceiling-floor and narrow gap environments in Fig. 5 and 7. In contrast, our method addresses these limitations by leveraging orientation-aware polytopes for each robot component - quadrotor, cable, and payload. Combined with polytopic representation of the environments, we enable non-conservative distance evaluation, as shown in Fig. 2, facilitating aerial transportation in maze-like environments and narrow passages that are slightly larger than the width of the cable.

State and Input Modeling. The second design decision is how states and inputs are mathematically represented. A

TABLE I: Notation table

$\mathcal{I}, \mathcal{B}$	inertial, and robot frame
$m_L, m_Q \in \mathbb{R}$	mass of payload and robot
$\mathbf{x}_L, \mathbf{x}_Q \in \mathbb{R}^3$	position of payload and robot in $\mathcal{I}$
$\dot{\mathbf{x}}_L, \ddot{\mathbf{x}}_L, \ddot{\mathbf{x}}_L \in \mathbb{R}^3$	payload velocity, acceleration, jerk in $\mathcal{I}$
$\dot{\mathbf{x}}_Q, \ddot{\mathbf{x}}_Q \in \mathbb{R}^3$	robot velocity, acceleration in $\mathcal{I}$
$\mathbf{R}_Q \in SO(3)$	robot orientation with respect to $\mathcal{I}$
$\boldsymbol{\Omega} \in \mathbb{R}^3$	robot angular velocity in $\mathcal{B}$
$\mathbf{f} \in \mathbb{R}, \mathbf{M} \in \mathbb{R}^3$	collective thrust and moment on robot in $\mathcal{B}$
$\mathbf{J} \in \mathbb{R}^{3 \times 3}$	moment of inertia of robot in $\mathcal{B}$
$\boldsymbol{\xi} \in S^2$	unit vector from robot to payload in $\mathcal{I}$
$f_i, l, g \in \mathbb{R}$	i^{th} motor thrust, cable length, gravity
$N_O \in \mathbb{Z}$	total number of obstacles
M_i, N_j	i^{th} obstacle and j^{th} robot polytope,

popular strategy is to express the states and inputs as polynomials to facilitate trajectory optimization [6], [9]. This is often paired with MINCO [17] to reformulate the problem into an unconstrained optimization to boost computational efficiency. While this approach can be used to achieve real-time capabilities, our experiments in Section V-Q1 indicate that it may compromise both solution quality and convergence. In contrast, similar to [7], [10]–[12], PolyFly operates directly in the native state space. Unlike polynomial parameterizations, the native state space approach places no restrictions on the shape of the state or control trajectories, making them inherently more expressive.

The method presented in [6], [9], serves as our primary baseline, employs an ESDF to represent the environment, approximates the quadrotor and payload with spheres, and parameterizes the states and inputs via polynomials. [9] reports an ablation study comparing the computational times and success rates of their method with IPOPT-MP, a state-based global planning approach. However, they omit the comparison of the optimized trajectory times of their method with IPOPT-MP, and do not address navigation close to obstacles in maze-like environments. We demonstrate that by using the native state-based modeling and a non-conservative spatial workspace representation via polytopes, PolyFly outperforms this method in all test environments.

III. SYSTEM DYNAMICS

In this section, we summarize the non-linear dynamics of a quadrotor carrying a suspended payload. We follow the derivation presented in [3] and assume that the cable remains taut. Applying the Lagrange–d’Alembert principle, we obtain

$$\frac{d\mathbf{x}_L}{dt} = \dot{\mathbf{x}}_L, \quad \frac{d\mathbf{x}_Q}{dt} = \dot{\mathbf{x}}_Q, \quad \dot{\mathbf{R}}_Q = \mathbf{R}_Q \hat{\boldsymbol{\Omega}}, \quad (1)$$

$$(m + m_L)(\ddot{\mathbf{x}}_L + \mathbf{g}) = (\boldsymbol{\xi} \cdot f \mathbf{R}_Q \mathbf{e}_3 - ml(\dot{\boldsymbol{\xi}} \cdot \dot{\boldsymbol{\xi}})) \boldsymbol{\xi}, \quad (2)$$

$$ml(\ddot{\boldsymbol{\xi}} + (\dot{\boldsymbol{\xi}} \cdot \dot{\boldsymbol{\xi}}) \boldsymbol{\xi}) = \boldsymbol{\xi} \times (\boldsymbol{\xi} \times f \mathbf{R}_Q \mathbf{e}_3), \quad (3)$$

$$\mathbf{M} = \mathbf{J} \dot{\boldsymbol{\Omega}} + \boldsymbol{\Omega} \times \mathbf{J} \boldsymbol{\Omega}, \quad (4)$$

$$[f, \mathbf{M}]^\top = \mathbf{D} [f_1, f_2, f_3, f_4]^\top. \quad (5)$$

where $\mathbf{g} = g\mathbf{e}_3$, $g = 9.81 \text{ m/s}^2$, $\mathbf{e}_3 = [0 \ 0 \ 1]^\top$, and $\hat{\boldsymbol{\Omega}}$ is the skew-symmetric matrix of the quadrotor angular velocity $\boldsymbol{\Omega}$. Table I defines the remaining variables. Qualitatively, eq. (2) describes the relationship between the payload acceleration and collective motor thrust, while eq. (3) shows that cable’s angular acceleration depends on the collective thrust and the

quadrotor’s orientation. The matrix $\mathbf{D}$ provides a linear mapping between collective thrust-moments and individual motor thrusts. It is defined in terms of the robot’s physical properties and propellers aerodynamics characteristics. Readers can refer to [18] for additional information.

By leveraging the differential flatness property of the system [3], we can model the state dynamics described in eqs. (1)–(5) in terms of the flat outputs $\{\mathbf{x}_L, \psi\}$, where ψ is the quadrotor’s yaw angle. In our method, we set $\psi = 0$ and assume negligible cable angular accelerations. This allows us to express the system through the state $\mathbf{x} = [\mathbf{x}_L^\top, \dot{\mathbf{x}}_L^\top, \ddot{\mathbf{x}}_L^\top]^\top \in \mathbb{R}^9$ and input $\mathbf{u} = [\ddot{\mathbf{x}}_L] \in \mathbb{R}^3$.

The induced system dynamics are therefore modeled as a third-order system. Given $\mathbf{x}$, $\mathbf{u}$, the robot’s physical parameters, and assuming that cable angular accelerations are zero at the states optimized by the planner, we can use the non-linear differential flatness mapping [3] to calculate the robot position $\mathbf{x}_Q$, robot velocity $\dot{\mathbf{x}}_Q$, and orientation $\mathbf{R}_Q$. This allows our trajectory generator to plan in payload space and add constraints to the quadrotor’s states. Since our planner operates in a discrete space, we convert the continuous time representation above to its discrete form. We obtain the Runge-Kutta approximation $\mathbf{x}(k+1) = F(\mathbf{x}(k), \mathbf{u}(k), \Delta t_k)$, where F is the discretized dynamics function and Δt_k the time step duration at stage k .

IV. PLANNING

A. Preliminaries

We develop a global planner that computes a collision-free trajectory $\boldsymbol{\tau}$ for a quadrotor with a suspended payload where

$$\boldsymbol{\tau} = [\boldsymbol{\tau}_0, \dots, \boldsymbol{\tau}_N], \quad \boldsymbol{\tau}_k = [\mathbf{x}_L(k), \dot{\mathbf{x}}_L(k), \ddot{\mathbf{x}}_L(k)]^\top, \quad (6)$$

where $k \in (0, N)$. The robot starts at state $\mathbf{x}_O$ and must navigate to its goal state $\mathbf{x}_N$. The environment is described by the set $\mathbf{E}$, which contains N_O polytopic obstacles M_i for $i \in (0, N_O)$. The aerial system is modeled using three separate polytopes, N_j for $j \in \{Q, C, L\}$ that represent the quadrotor, cable, and load polytopes respectively, as shown in Fig. 2.

B. Polytopic Approximation

Let us consider two polytopes M_i and N_j , as defined in Table I. Let $\mathbf{y}_{M_i}$ be any point within the polytope M_i and $\mathbf{y}_{N_j}$ be any point within the polytope N_j . $\mathbf{A}_{M_i}, \mathbf{B}_{M_i}, \mathbf{A}_{N_j}$ and $\mathbf{B}_{N_j}$ are constants that define their geometric shape. These polytopes can be represented by

$$\mathbf{A}_{M_i} \mathbf{y}_{M_i} \leq \mathbf{B}_{M_i}, \quad \mathbf{A}_{N_j} \mathbf{y}_{N_j} \leq \mathbf{B}_{N_j}. \quad (7)$$

While not strictly necessary, we simplify computation by representing one polytope’s points in the other polytope’s frame. We represent the obstacle points $\mathbf{y}_{M_i}$ using the relative position vector $\mathbf{y}_{M_i}^{N_j}$, which gives

$$\mathbf{y}_{M_i} = \mathbf{R}_{N_j} \mathbf{y}_{M_i}^{N_j} + \mathbf{O}_{N_j}, \quad (8)$$

where $\mathbf{R}_{N_j}$ is N_j ’s rotation matrix, and $\mathbf{O}_{N_j}$ is its origin. Substituting this into the polytopic form expressed by eq. (7), we can represent polytope M_i as

$$(\mathbf{A}_{M_i} \mathbf{R}_{N_j}) \mathbf{y}_{M_i}^{N_j} \leq \mathbf{B}_{M_i} - \mathbf{A}_{M_i} \mathbf{O}_{N_j}. \quad (9)$$

Since $\mathbf{R}_{N_j}$ and $\mathbf{O}_{N_j}$ are functions of the chosen state vector $\mathbf{x}$, M_i is represented by

$$\begin{aligned} \mathbf{A}_{M_i}^{N_j}(\mathbf{x})\mathbf{y}_{M_i}^{N_j} &\leq \mathbf{B}_{M_i}^{N_j}(\mathbf{x}), \\ \mathbf{A}_{M_i}^{N_j}(\mathbf{x}) &= \mathbf{A}_{M_i}\mathbf{R}_{N_j}(\mathbf{x}), \mathbf{B}_{M_i}^{N_j}(\mathbf{x}) = \mathbf{B}_{M_i} - \mathbf{A}_{M_i}\mathbf{O}_{N_j}(\mathbf{x}). \end{aligned} \quad (10)$$

C. Distance Metric Between Polytopes

Our optimization problem seeks to lower bound h_{M_i, N_j} , the distance between polytopes M_i and N_j , with a positive margin β . Specifically,

$$h_{M_i, N_j}(\mathbf{x}) \geq \beta, \quad h_{M_i, N_j}(\mathbf{x}) := d(M_i, N_j), \quad (11)$$

where d is the distance function described in [19, Eq. (8a)]. The constraint defined by eq. (11) is non-differentiable and difficult to integrate into optimization problems when the obstacles and robots are represented as convex polytopes [19]. To handle this complexity, we follow the approach in [19, Proposition 1]. and reformulate eq. (11) using duality theory into the following set of smooth non-linear differentiable constraints. The equivalent dual problem is

$$\begin{aligned} g_{M_i, N_j}(\mathbf{x}) &= \max_{\lambda_{N_j}, \lambda_{M_i}} -\lambda_{N_j}^\top \mathbf{B}_{N_j} - \lambda_{M_i}^\top \mathbf{B}_{M_i}^{N_j}(\mathbf{x}) \\ \text{s.t. } \quad &\lambda_{N_j}^\top \mathbf{A}_{N_j} + \lambda_{M_i}^\top \mathbf{A}_{M_i}^{N_j}(\mathbf{x}) = 0 \\ &\lambda_{M_i} \geq 0, \quad \lambda_{N_j} \geq 0, \quad \|\lambda_{M_i} * \mathbf{A}_{M_i}\|_2 \leq 1, \end{aligned} \quad (12)$$

where $\lambda_{N_j} \in \mathbb{R}^{n_j}$ and $\lambda_{M_i} \in \mathbb{R}^{m_i}$ are the dual variables associated with the obstacle avoidance problem. For our setup, $m_i = n_j = 6$. Results in [19], [20] show that Strong Duality holds for the dual problem defined in eq. (12), which gives us the lower bound of the primal cost function $h_{M_i, N_j}(\mathbf{x})$

$$\bar{g}_{M, N}(\mathbf{x}, \lambda_{M_i}, \lambda_{N_j}) := -\lambda_{N_j}^\top \mathbf{B}_{N_j} - \lambda_{M_i}^\top \mathbf{B}_{M_i}^{N_j}(\mathbf{x}) \leq h_{M_i, N_j}(\mathbf{x}). \quad (13)$$

Therefore, to enforce a minimum distance of β between M_i and N_j , we constrain the lower bound $\bar{g}_{M, N} \geq \beta$, giving us a smooth non-conservative reformulation of eq. (11)

$$\begin{aligned} -\lambda_{N_j}^\top \mathbf{B}_{N_j} - \lambda_{M_i}^\top \mathbf{B}_{M_i}^{N_j}(\mathbf{x}) &\geq \beta, \\ \lambda_{N_j}^\top \mathbf{A}_{N_j} + \lambda_{M_i}^\top \mathbf{A}_{M_i}^{N_j}(\mathbf{x}) &= 0, \\ \lambda_{M_i} \geq 0, \quad \lambda_{N_j} \geq 0, \quad \|\lambda_{M_i} * \mathbf{A}_{M_i}\|_2 &\leq 1. \end{aligned} \quad (14)$$

We add the set of constraints defined in eq. (14) for collisions between each (M_i, N_j) pair, where $i \in (0, N_O)$. For notational convenience, we combine the dual variables into a vector $\Lambda \in \mathbb{R}^Q$, where $Q = N \times N_O \times N_R \times (m_i + n_j)$, and $N_R = 3$ is the number of robot components.

D. Optimization Problem Formulation

Our planner solves the following optimization problem

$$\begin{aligned} \min_{\mathbf{x}(k), \mathbf{u}(k), \Delta t_k, \Lambda} \quad &\sum_{k=0}^{N-1} L(\mathbf{x}(k), \mathbf{u}(k), \Delta t_k) \\ \text{subject to} \quad &\mathbf{x}(k+1) = F(\mathbf{x}(k), \mathbf{u}(k), \Delta t_k), \\ &\mathbf{x}(0) = \mathbf{x}_0, \quad \mathbf{x}(N) = \mathbf{x}_N, \\ &H_l(\mathbf{x}(k), \mathbf{u}(k), \Lambda) \leq 0, \quad \Delta t_{\min} \leq \Delta t_k \leq \Delta t_{\max}, \end{aligned} \quad (15)$$




Fig. 3: A* trajectories for Env. 6 (left) and Env. 8 (right)

where $\forall k \in (0, N-1)$, $\forall l \in (0, N_h)$, N is the horizon length, Δt_k is the time duration for stage k , and x_0 and x_N are the initial and terminal states respectively. $\Delta t_{\min}$ and $\Delta t_{\max}$ are the min. and max. time durations set by the user and are discussed in Section IV-D1. H_l comprises the N_h inequality constraints from the collision-avoidance formulation in eq. (14) together with the state and input bounds specified in Section IV-D7. Finally, $L = L_t + L_u + L_g + L_{td}$ is the loss function defined in the subsequent subsections.

We now describe the optimization constraints, costs and initialization methods. The variables $\Delta t_{\min}$, $\Delta t_{\max}$, α_{to} , α_u , α_g , α_{td} and N , are user defined constants whose values are defined in Section V.

1) *Time Minimization*: We incentivize the planner to minimize the total trajectory time by adding a large cost to each time step duration. Specifically,

$$L_t = T \frac{\alpha_{to}}{N}, T = \sum_{k=0}^{N-1} \Delta t_k, \Delta t_{\min} \leq \Delta t_k \leq \Delta t_{\max}, \quad (16)$$

where $\Delta t_{\min}, \Delta t_{\max}, \alpha_{to} > 0$.

2) *Input Smoothness Regularization*: We add a small cost that penalizes large changes between consecutive inputs, thereby regularizing their rate of change and improving the smoothness of the resulting robot states.

$$L_u = \frac{\alpha_u}{N} \sum_{k=1}^N \|\mathbf{u}_k - \mathbf{u}_{k-1}\|^2, \quad \alpha_u > 0. \quad (17)$$

3) *Regularizing Proximity to the Initial Guess*: Even though we use a simple A* initialization, we expect the optimal trajectory's positions to lie in the vicinity of the initialized positions. We achieve this by adding a small penalty to deviations of the optimized positions from the initial guess. The study supporting the benefits of this term is discussed in Section V-Q4.

$$L_g = \frac{\alpha_g}{N-1} \sum_{k=1}^{N-1} \|\mathbf{x}(k) - \mathbf{x}_g(k)\|^2, \quad (18)$$

where $\mathbf{x}_g(k)$ is the initial guess at stage index k and $\alpha_g > 0$.

4) *Regularizing Subsequent Time Step Durations*: We add a small cost to the difference between subsequent time step durations to help our solver converge. Specifically, we add

$$L_{td} = \frac{\alpha_{td}}{N} \sum_{k=0}^{N-1} \|\Delta t_{k+1} - \Delta t_k\|^2, \quad \alpha_{td} > 0. \quad (19)$$

5) *Solver Initialization*: Since the optimization problem combines non-linear dynamics with non-convex collision constraints, the resulting solution is sensitive to its initial guess. Whereas many prior planners [6], [9] adopt a kino-dynamic initialization strategy, our approach begins from a simple A^* path that ignores the non-linear dynamics and assumes the cable is always aligned with the global z -axis, i.e., $\ddot{\mathbf{x}}_L = 0$. Fig. 3 visualizes this initialization for two representative environments. We then use a lightweight initialization method for the load velocities using the equation

$$\dot{\mathbf{x}}_{L,k} = \frac{\mathbf{x}_{L,k+1} - \mathbf{x}_{L,k}}{\Delta t_{max} - \Delta t_{min}}, \quad (20)$$

where $0 < \Delta t_{min} < \Delta t_{max}$. Even with this relatively simple initialization method, our planner produces high quality trajectories that are consistently better than the baseline.

6) *Collision Constraints*: We account for the collision between the quadrotor, cable, and payload polytopes against each obstacle. For each time stage k and each polytope pair, we add the constraints derived in Section IV-C. As an example, to handle collisions between the payload polytope N_L and an obstacle M_i , we add the following constraints for all stages k and obstacles i

$$\begin{aligned} \forall k, \forall i \quad & -\lambda_{N_L,k}^\top \mathbf{B}_{N_L} - \lambda_{M_i,k}^\top \mathbf{B}_{M_i}^{N_L}(\mathbf{x}_k) \geq \beta, \\ & \lambda_{N_L,k}^\top \mathbf{A}_{N_L} + \lambda_{M_i,k}^\top \mathbf{A}_{M_i}^{N_L}(\mathbf{x}_k) = 0, \\ & \lambda_{M_i,k} \geq 0, \quad \lambda_{N_L,k} \geq 0, \quad \|\lambda_{M_i,k} * \mathbf{A}_{M_i}\|_2 \leq 1. \end{aligned} \quad (21)$$

7) *State and Input Constraints*: The system's initial and final positions are specified based on the environment and are enforced via equality constraints. We also constrain the start and end velocities, accelerations and jerks to zero to account for hardware limitations. Specifically,

$$\begin{aligned} \mathbf{x}(0) &= \mathbf{x}_0, \quad \mathbf{x}(N) = \mathbf{x}_N, \quad \mathbf{u}(0) = \mathbf{u}(N) = 0, \\ \dot{\mathbf{x}}_L(0) &= \dot{\mathbf{x}}_L(N) = 0, \quad \ddot{\mathbf{x}}_L(0) = \ddot{\mathbf{x}}_L(N) = 0, \\ \dot{\mathbf{x}}_Q(0) &= \dot{\mathbf{x}}_Q(N) = 0, \quad \ddot{\mathbf{x}}_Q(0) = \ddot{\mathbf{x}}_Q(N) = 0. \end{aligned} \quad (22)$$

Finally, we bound the state and inputs based on the robot's limitations and the environment dimensions.

$$\mathbf{x}_l \leq \mathbf{x} \leq \mathbf{x}_h, \quad \mathbf{u}_l \leq \mathbf{u} \leq \mathbf{u}_h, \quad (23)$$

where $\mathbf{x}_l, \mathbf{u}_l$ and $\mathbf{x}_h, \mathbf{u}_h$ are the lower and upper bounds.

8) *Incorporating the Quadrotor's Rotation into the Polytopic Constraints*: To improve the accuracy of the polytopic representation, we rotate the quadrotor polytope by the quadrotor's rotation. As discussed in Section III, we can derive the robot's rotation $\mathbf{R}_Q$ using $\mathbf{x}$ and $\mathbf{u}$, and use it to compute the $\mathbf{A}_{M_i}^{N_Q}$ matrix in eq. (10).

V. EXPERIMENTAL RESULTS

We conduct various experiments to evaluate the capabilities of this method. We seek to answer the following:

- Q1** Can this planning method produce trajectories that navigate tightly around corners of polytopic obstacles in cluttered maze environments?
- Q2** What are the benefits of modeling each component of the robot separately?

Q3 What are the benefits of adding the quadrotor's orientation into the polytopic representation?

Q4 Is adding a cost that regularizes the proximity to the initial guess important?

Q5 Does our velocity initialization method help the solver converge?

Q6 Can we run these aggressive trajectories on a real quadrotor with a suspended payload?

For all experiments, we set $\Delta t_{min} = 0.01s$, $\Delta t_{max} = 0.20s$, $\alpha_u = 5.0$, $\alpha_g = 5.0$, $\alpha_{td} = 600.0$ and $N = 100$. As described in Section V-Q1, α_{to} is tuned to maximize performance in each environment.

Q1 Navigating Cluttered Environments

We create 10 maze environments to evaluate the performance of our planning method. The optimization problem is constructed in CasADi [21] and solved with IPOPT [22].

We select AutoTrans, the MINCO approach [17] adopted in [6], [9] as our primary baseline, which is an optimization-based trajectory generator that parametrizes trajectories as a sequence of 7th order polynomials, initialized with kinodynamic hybrid A^* . They model the quadrotor and payload as independent spheres, the cable as a set of discrete points, and encode the environment using an ESDF. To solve the resulting problem, they reformulate the problem into an unconstrained one by applying L_1 penalties on the constraints.

To ensure a fair comparison, we adjusted AutoTrans's safety margins, ESDF map resolution, kinodynamic search resolution, and quadrotor parameters to be consistent with our environment setup and removed AutoTrans' hardware limitations constraints on the quadrotor thrust and tilt angle. Furthermore, for each environment, we tune AutoTrans' and PolyFly's time minimization weights and report the mean results based on the planning output over 10 trials in Table II. PolyFly's trajectories are shown in Fig. 4 and Fig. 5.

We include an additional baseline, PolySingle, which mirrors our method but adopts the modeling of [7], where a single polytope represents the entire robot. AutoTrans fails to produce a collision-free trajectory in 5/10 environments, as indicated by a $\times$. We also observed that in some cases, AutoTrans generated shorter paths in terms of distance, yet required more time to traverse them. This indicates that our method is able to better optimize for total flight time, even if that entails following a slightly longer path at higher speeds. Compared to PolySingle, our method successfully navigates the narrow gaps in Env. 7, 9 and 10, whereas PolySingle fails due to its modeling over-approximation. PolyFly's boost in performance compared to the AutoTrans may stem from the geometric modeling of the workspace and the use of the dual formulation, which produces smooth, non-conservative collision constraints [19]. These properties enable it to plan paths around obstacles while maintaining speed.

Q2 Independent Polytopic Models for Robot Components

Modeling each robot component as a separate polytope allows the trajectory generator to accurately represent the system's geometry. This fidelity is important when producing




Fig. 4: Visualization of trajectories generated using PolyFly in 8 tightly cluttered environments. The quadrotor is in yellow, payload in green, cable in gray and obstacles in light blue. Envs. 1-4 are shown in the top row and 5-8 in the bottom row.

TABLE II: PolyFly vs. PolySingle vs AutoTrans. Statistics are calculated after running 10 trials in each environment. T is the mean trajectory duration in seconds, Path is mean path length in meters, Comp. is the mean optimization wall time in seconds, and SR is the success ratio. % $T \uparrow$ w/ A denotes the % by which PolyFly's trajectory duration T is shorter than method A 's. $\times$ marks a failure.

Env.	PolyFly					PolySingle (PS)					AutoTrans (AT) [6]				% $T \uparrow$ w/ AT	% $T \uparrow$ w/ PS
	T	Path	Comp.	SR	α_{to}	T	Path	Comp.	SR	α_{to}	T	Path	Comp.	SR		
1	3.90	4.51	4.83	1.0	5000	3.96	4.50	5.26	1.0	5000	4.10	4.45	0.06	1.0	4.9	1.5
2	4.60	6.24	3.97	1.0	6000	4.68	6.34	2.99	1.0	6000	$\times$	$\times$	$\times$	0.0	∞	1.7
3	4.92	6.24	7.11	1.0	5000	5.13	6.35	7.51	1.0	5000	5.68	6.47	0.11	1.0	13.4	4.1
4	7.09	8.82	6.80	1.0	6000	7.22	9.16	5.16	1.0	6000	7.47	8.74	0.89	0.9	5.1	1.8
5	4.05	5.82	11.14	1.0	5000	4.12	5.82	11.20	1.0	5000	4.30	4.92	1.63	0.6	5.8	1.7
6	5.31	7.95	4.90	1.0	6000	5.48	8.08	7.10	1.0	6000	6.43	8.02	49.44	0.0	17.4	3.1
7	2.86	3.07	3.12	1.0	800	$\times$	$\times$	$\times$	0.0	800	$\times$	$\times$	$\times$	0.0	∞	∞
8	4.92	9.40	3.31	1.0	10000	4.96	9.51	3.21	1.0	10000	$\times$	$\times$	$\times$	0.0	∞	0.8
9	3.72	5.89	3.34	1.0	2000	$\times$	$\times$	$\times$	0.0	2000	$\times$	$\times$	$\times$	0.0	∞	∞
10	4.46	6.89	9.92	1.0	2000	$\times$	$\times$	$\times$	0.0	2000	$\times$	$\times$	$\times$	0.0	∞	∞




Fig. 5: PolyFly's trajectories for Env. 7 (top), 9 (bottom-left) and 10 (bottom-right) showcasing narrow-gap navigation.

agile trajectories to navigate through narrow gaps and in tightly constrained environments where the robot must cut corners aggressively while still avoiding obstacles.

The advantage is evident in Env. 7, 9 and 10, where the robot must pass through a narrow opening wider than the cable but thinner than the payload and quadrotor. Both the baselines fail, but our method successfully produces trajectories that thread the cable through the gaps, allowing the robot to safely navigate the environment. The real world experiment snapshots are shown in Fig. 1.

PolySingle fails in these narrow-gap environments due to its single-polytope representation that over-approximates the robot's geometry. On the other hand, AutoTrans models the cable as a discrete set of points and performs collision checks on each point, which should geometrically allow the cable to

pass through the narrow gaps. Even though a solution exists, AutoTrans fails to produce a trajectory, even after extensive tuning of its optimization weights. Therefore, rather than the underlying geometry, the cause of the failures in AutoTrans may stem from other elements of the pipeline. First, the MINCO polynomial state parameterization, cannot, in general, represent optimal solutions [17] which limits its performance in tightly cluttered scenarios. And second, distance-field-based collision checking methods are not well-suited for applications that require accurate modeling of geometry [23]. These methods compute the gradients to prevent collisions per point of the robot/obstacle, rather than the exact gradient needed to move the polytope out of collision. These two gradients do not always agree, and can cause issues when the robot travels in close proximity to obstacles.

Q3 Incorporating the Quadrotor's Orientation Into Polytopic Representation

Accurately capturing the quadrotor's orientation further increases the fidelity of the polytopic model and therefore navigation reliability. This detail becomes critical when the vehicle flies close to a ceiling or floor, particularly for a platform with long rotor arms that form a thin yet wide rectangular footprint. Since high-acceleration maneuvers induce large orientation changes, aggressive trajectories close to surfaces can cause rotor collisions with the environment. As shown in Fig. 2 and 7, most planners over-approximate the vehicle with a sphere or




Fig. 6: Left: An orientation-agnostic planner produces a trajectory that collides with the ceiling. Right: By accounting for the robot’s orientation, the optimized path descends to prevent collisions while accelerating.



Fig. 7: Optimized trajectories by PolyFly for the ceiling-floor (left) and ceiling-wall (right) environments. The dashed lines are the polytopic approximation by PolySingle and the semi-transparent ellipse is the ellipsoidal approximation.

ellipsoid. By employing a tight polytope that accounts for the robot’s orientation, we eliminate this conservatism and ensure collision free navigation in near-ceiling flight.

We demonstrate these benefits in Fig. 6 and 7. Fig. 6 illustrates how incorporating the orientation into the planner prevents collisions. In Fig. 6 (A), the quadrotor strikes the ceiling during a high-acceleration maneuver because the planner ignores the orientation change that accompanies acceleration. In contrast, Fig. 6 (B) shows how the orientation-aware planner descends during the acceleration phase to avoid a collision. The ceiling-floor scenario in Fig. 7 shows the PolyFly trajectory and overlays the PolySingle polytope to illustrate the over-approximation introduced by [7]. PolySingle fails in these scenarios because [7] uses a common orientation for all components due to its single polytope representation of the entire system, which prevents navigation in such tightly constrained environments.

Q4 Regularizing Proximity to Initial Guess

We perform an ablation study to determine the importance of regularizing the proximity to the initial guess. We run the optimizer with 3 different α_g values. The results are shown in Table III. PolyFly fails in 5/10 environments if $\alpha_g = 0$. Trajectories generated with $\alpha_g = 1$ are on average slower than those generated with $\alpha_g = 5$. This highlights the value of this regularization term. We believe this aids convergence by incentivizing the optimizer to remain in the vicinity of the initial guess, preventing it from entering regions of local minima that are difficult to escape from.

Q5 Payload Velocity Initialization

To evaluate the effectiveness of our solver initialization strategy, we conduct an ablation study that compares the payload-velocity initialization of eq. (20) against a variant where velocities are initialized to zero. Table IV summarizes the results. The solver fails in 2/10 environments without the velocity initialization strategy. For all other scenarios, our strategy reduces the solver iterations by an average of 59%.

Q6 Real World Experiments

We run real world experiments with Env. 1, 3 and 7 to validate that our method produces trajectories that can be

TABLE III: T for different values of α_g . % $T \uparrow$ (5 w/ 1) indicates improvement of $\alpha_g = 5$ over $\alpha_g = 1$.

Env.	$\alpha_g = 0$	$\alpha_g = 1$	$\alpha_g = 5$	% $T \uparrow$ (5 w/ 1)
1	×	3.92	3.92	0.0%
2	×	5.02	4.61	8.2%
3	4.94	4.96	4.97	-0.2%
4	×	7.13	7.15	0.3%
5	4.20	4.03	4.05	-0.5%
6	×	5.48	5.37	2.0%
7	3.15	4.59	3.16	31.2%
8	×	4.97	4.99	-0.4%
9	3.93	3.93	3.94	-0.3%
10	4.67	4.72	4.69	0.6%

TABLE IV: Performance with and without $\dot{\mathbf{x}}_L$ initialization.

Env.	With Velocity Initialization		W/o Velocity Initialization	
	T (s)	Iterations	T (s)	Iterations
1	3.90	201 ↓	3.98	745
2	4.60	223 ↓	4.61	640
3	4.92	133 ↓	4.92	677
4	7.09	185 ↓	×	×
5	4.05	353 ↓	×	×
6	5.31	206 ↓	5.31	559
7	2.86	234 ↓	2.87	309
8	4.92	161 ↓	4.90	617
9	3.72	112 ↓	3.72	206
10	4.46	306 ↓	4.43	652




Fig. 8: System architecture for real-world deployment.

TABLE V: Tracking error statistics (meters)

Axis	Env. 1		Env. 3		Env. 7	
	Mean	Std	Mean	Std	Mean	Std
x	0.027	0.017	0.029	0.016	0.021	0.094
y	0.058	0.047	0.059	0.048	0.052	0.036
z	0.024	0.012	0.039	0.029	0.029	0.087

deployed on real robots. Our system architecture is shown in Fig. 8. For real-world experiments only, we use the Non-Linear Model Predictive Control (NMPC) setup from [24] as our controller. The NMPC uses the dynamics eqs. (1)–(5) with inputs f_{1-4} . After generating a trajectory with PolyFly, we send the desired τ to the NMPC at 200 Hz. As in [24], the predicted collective thrust and bodyrates from the MPC’s second stage are sent to the low-level controller. The robot’s states are measured using a motion capture system and are upsampled using Kalman Filters. Snapshots of our experiments are shown in Fig. 1. The position errors for the quadrotor and payload are shown in Fig. 9 and Table V. Our results indicate that PolyFly produces dynamically feasible trajectories that can be tracked by an onboard controller with low errors. The supplementary video contains real-world footage that highlights the agile trajectories generated by our method.




Fig. 9: Desired vs. Measured positions of the quadrotor and payload during real-world experiments for Env. 1, 3, and 7.

VI. CONCLUSION

This paper presented PolyFly, a global planner for a single robot aerial transportation system that models the robot and environment using a set of polytopes and enables the generation of collision-free trajectories in cluttered scenarios. Our proposed method yielded consistently faster trajectories than a popular state-of-the-art baseline in ten different maze-like environments. Experimental results confirm the trajectories' dynamic feasibility and the applicability of the proposed solution in real-world settings. One limitation is that PolyFly does not yet run at real-time frequencies or account for onboard obstacle sensing. Future work will address this by efficiently updating trajectories using onboard local obstacle information and leveraging advances in polytopic collision-avoidance checks [25] to reduce computation.

REFERENCES

- [1] M. A. Trujillo, J. R. Martinez-de Dios, C. Martin, A. Viguria, and A. Ollero, "Novel aerial manipulator for accurate and robust industrial ndt contact inspection: A new tool for the oil and gas inspection industry," *Sensors*, vol. 19, no. 6, 2019.
- [2] G. Loianno and V. Kumar, "Cooperative transportation using small quadrotors using monocular vision and inertial sensing," *IEEE Robotics and Automation Letters*, pp. 680–687, April 2018.
- [3] K. Sreenath, T. Lee, and V. Kumar, "Geometric control and differential flatness of a quadrotor uav with a cable-suspended load," in *52nd IEEE Conference on Decision and Control (CDC)*, 2013, pp. 2269–2274.
- [4] E. N. Barmounakis, E. I. Vlahogianni, and J. C. Golias, "Unmanned aerial aircraft systems for transportation engineering: Current practice and future challenges," *International Journal of Transportation Science and Technology*, vol. 5, no. 3, pp. 111–122, 2016.
- [5] T. A. Rodrigues, J. Patrikar, N. L. Oliveira, H. S. Matthews, S. Scherer, and C. Samaras, "Drone flight data reveal energy and greenhouse gas emissions savings for very small package delivery," *Patterns*, vol. 3, no. 8, 2022.
- [6] H. Li, H. Wang, C. Feng, F. Gao, B. Zhou, and S. Shen, "Autotrans: A complete planning and control framework for autonomous uav payload transportation," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6859–6866, 2023.
- [7] J. Zeng, P. Kotaru, M. W. Mueller, and K. Sreenath, "Differential flatness based path planning with direct collocation on hybrid modes for a quadrotor with a cable-suspended payload," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3074–3081, 2020.
- [8] S. Tang and V. Kumar, "Mixed integer quadratic program trajectory generation for a quadrotor with a cable-suspended payload," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [9] H. Wang, H. Li, B. Zhou, F. Gao, and S. Shen, "Impact-aware planning and control for aerial robots with suspended payloads," *IEEE Transactions on Robotics*, vol. 40, pp. 2478–2497, 2024.
- [10] C. Y. Son, H. Seo, D. Jang, and H. J. Kim, "Real-time optimal trajectory generation and control of a multi-rotor with a suspended load for obstacle avoidance," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1915–1922, 2020.
- [11] P. Foehn, D. Falanga, C. Luethi, M. Achtelik, R. Siegwart, and D. Scaramuzza, "Fast trajectory optimization for agile quadrotor maneuvers with a cable-suspended payload," in *Robotics: Science and Systems XIII. RSS Foundation*, July 2017.
- [12] C. Y. Son, D. Jang, H. Seo, T. Kim, H. Lee, and H. J. Kim, "Real-time optimal planning and model predictive control of a multi-rotor with a suspended load," in *International Conference on Robotics and Automation (ICRA)*, 2019, pp. 5665–5671.
- [13] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [14] G. Yu, D. Cabecinhas, R. Cunha, and C. Silvestre, "Aggressive maneuvers for a quadrotor-slung-load system through fast trajectory generation and tracking," *Autonomous Robots*, Apr. 2022.
- [15] A. Pallar, G. Li, M. Sarvaiya, and G. Loianno, "Optimal trajectory planning for cooperative manipulation with multiple quadrotors using control barrier functions," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 2808–2814.
- [16] A. Thirugnanam, J. Zeng, and K. Sreenath, "Safety-critical control and planning for obstacle avoidance between polytopes with control barrier functions," in *International Conference on Robotics and Automation (ICRA)*, 2022, pp. 286–292.
- [17] Z. Wang, X. Zhou, C. Xu, and F. Gao, "Geometrically constrained trajectory optimization for multicopters," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 3259–3278, 2022.
- [18] T. Lee, M. Leok, and N. H. McClamroch, "Geometric tracking control of a quadrotor uav on $se(3)$," in *49th IEEE Conference on Decision and Control (CDC)*, 2010, pp. 5420–5425.
- [19] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-based collision avoidance," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 3, pp. 972–983, 2020.
- [20] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [21] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi - A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, 2019.
- [22] A. Wachter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, Mar. 2006.
- [23] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization," in *Robotics: science and systems*, vol. 9, no. 1. Berlin, Germany, 2013, pp. 1–10.
- [24] M. Sarvaiya, G. Li, and G. Loianno, "Hpa-mpc: Hybrid perception-aware nonlinear model predictive control for quadrotors with suspended loads," *IEEE Robotics and Automation Letters*, 2024.
- [25] Z. Wu, Z. Wang, and H. Zhang, "Gpu-accelerated optimization-based collision avoidance," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 7561–7567.