

Trajectory Planning Using Safe Ellipsoidal Corridors as Projections of Orthogonal Trust Regions

Akshay Jaitly¹, Jon Arrizabalaga², and Guanrui Li¹

Abstract—Planning collision free trajectories in complex environments remains a core challenge in robotics. Existing corridor based planners which rely on decomposition of the free space into collision free subsets scale poorly with environmental complexity and require explicit allocations of time windows to trajectory segments. We introduce a new trajectory parameterization that represents trajectories in a nonconvex collision free corridor as being in a convex cartesian product of balls. This parameterization allows us to decouple problem size from geometric complexity of the solution and naturally avoids explicit time allocation by allowing trajectories to evolve continuously inside ellipsoidal corridors. Building on this representation, we formulate the Orthogonal Trust Region Problem (Orth-TRP), a specialized convex program with separable block constraints, and develop a solver that exploits this parallel structure and the unique structure of each parallel subproblem for efficient optimization. Experiments on a quadrotor trajectory planning benchmark show that our approach produces smoother trajectories and lower runtimes than state-of-the-art corridor based planners, especially in highly complicated environments.

I. INTRODUCTION

Collision free trajectory optimization is a core challenge in robotics, useful in mobile robot navigation and mobile manipulation [1]–[3]. By mathematically defining sets of feasible trajectories that a system can undergo, and assigning a cost to each element of the set, trajectory optimization can be posed as a mathematical programming problem. Several methods exist to characterize these sets to facilitate efficient searches for optimal feasible trajectories.

Conventional sampling based planners [4]–[6] build tree or graph like structures in the configuration space, with nodes corresponding to collision free states and edges representing locally feasible motions. Paths through these structures define potential trajectories, which can then be searched efficiently using graph search or shortest path algorithms. On the other hand, some optimization based approaches [7]–[12] impose constraints on robot’s states (enforcing that a barrier between collision objects is respected) to define sets of feasible solutions, then minimize a cost function within this nonconvex feasible set. These enable powerful nonconvex optimization solvers for trajectory optimization. In more complicated scenarios, with larger dimensional spaces or

¹A. Jaitly and G. Li are with the Robotics Engineering Department, Worcester Polytechnic Institute, Worcester, MA 01609, USA {ajaitly, gli7}@wpi.edu

²J. Arrizabalaga is with the Robotics Institute at Carnegie Mellon University jarizab@andrew.cmu.edu

We would like to thank Ethan Chandler regarding discussions about the acceleration of Trust Region methods using Eigendecomposition and thank Arvind Raghunathan for discussions about the dual of QCQPs.

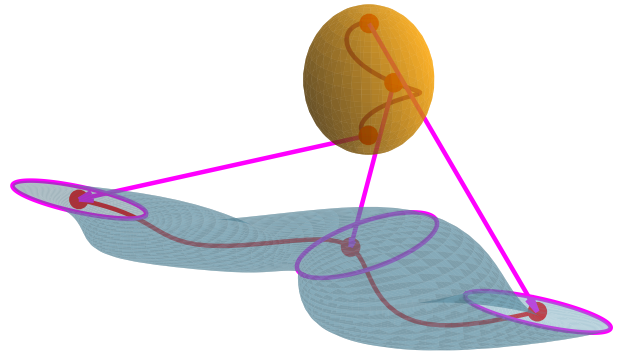


Fig. 1: A path (red) within an ellipsoidal corridor (blue) as a time varying projection of a **trajectory** within a lifted hypersphere (orange). Note that projected points are restricted to be within the relevant ellipsoidal cross section.

with numerous obstacles, the above methods inevitably face the curse of dimensionality.

Recent work has introduced convex approximations of free space through “safe corridors”, using either polytopic or ellipsoidal regions [13]–[16]. In these approaches, a trajectory is divided into segments, and each segment is constrained to lie within a corresponding convex subset of the obstacle free space at a specific point in time. By enforcing convex constraints on the robot’s configuration, these methods convert a difficult global planning problem into convex optimization, which can be solved using efficient convex programming techniques [1], [15], [16].

However, a key challenge in these methods is allocating time across the convex regions. Each region requires a specified time interval, and if this interval is too short, the system must accelerate sharply to meet continuity constraints, while overly long intervals produce inefficient motions. As the number of convex subsets increases, these timing decisions become increasingly difficult to manage, often introducing sensitivity and limiting the scalability of convex corridor methods.

[17] introduced Differentiable Collision Free Parametric Corridors which model free space as a smooth nonconvex corridor, made of a continuously deforming convex set rather than discrete segments, offering a more unified description of collision free regions. Building on this idea, we view these nonconvex corridors as time-varying projections of orthogonal trust regions in a higher dimensional parameter space. By lifting trajectories into a space where feasible solutions form a cartesian product of high dimensional balls,

each point in this lifted space naturally corresponds to a collision free trajectory within the corridor. This decouples problem size from environmental complexity while allowing representations of paths in the safe corridor as points in the lifted space.

This representation results in a favorable convex feasible set for trajectory optimization. We pose the resulting problem as the Orthogonal Trust Region Problem (Orth-TRP), which can be expressed as a collection of interconnected trust region subproblems (TRPs). Because each block of variables has its own separable constraint, the Orth-TRP naturally supports a parallelizable algorithm, where each block can be updated efficiently using trust region steps that resemble simple one dimensional line searches.

In summary, our contributions are:

- We develop a convex representation of sets of trajectories within a nonconvex set of configurations using a product of multiple high dimensional balls, decoupling the solution complexity from the size of the optimization problem.
- We create a solver that exploits both, the parallelizable structure and the unique Trust-Region-like structure of our resulting problem to solve quadratically constrained quadratic optimization problems where constraints are enforced on separable variables.

Taken together, these components create a scalable and geometrically intuitive framework for collision-free trajectory optimization in complex environments. Our experiments show that this approach consistently outperforms state-of-the-art implementations on challenging benchmarks. In direct comparisons with the method of [15], even when using a powerful solver like OSQP [18], our method solves long-horizon planning problems faster while producing comparable or smoother results. To the best of our knowledge, this is the first continuous parameterization of trajectories that enables optimization within a collision-free corridor where problem size and runtime are agnostic to horizon length and solution complexity. We achieve this by combining a new convex programming formulation with a solver that fully exploits the problem’s separable trust-region structure, eliminating the scaling and time-allocation issues inherent to existing corridor-based approaches and enabling both speed and robustness at scale.

II. BACKGROUND

In this work, we develop a parameterization of paths within Parametric Collision-Free Corridors. These are time-varying convex (ellipsoidal) sets embedded in the obstacle-free space. By construction, any path inside these corridors is guaranteed to be collision free. We formulate trajectory optimization within this corridor as minimizing a quadratic objective over a ball-constrained set, matching the structure of the classical Trust Region Problem. Accordingly, this section reviews key concepts from Parametric Collision Free Corridors and the classical Trust Region Problem, which form the basis of our approach. The notation used throughout is summarized in Table I.

TABLE I: Notation

Symbol	Meaning
$q \in \mathbb{R}^{N_q}$	Configuration in the configuration space.
$\epsilon \in [0, 1]$	Central “time” variable along the trajectory; $\epsilon = 1$ denotes the end of the trajectory.
$y \in \mathbb{R}^{N_y}$	Vector in the lifted parameter space, with $N_y \geq N_q$. A configuration space trajectory is given by $q(t) = L(t)y + \bar{q}(t)$.
\mathcal{S}^{N_a}	Unit ball in \mathbb{R}^{N_a} , defined as $\mathcal{S}^{N_a} = \{a \in \mathbb{R}^{N_a} : \ a\ \leq 1\}$ (in contrast to using $\ a\ = 1$ to describe only the surface).
$\mathcal{A} = BC$	For set valued \mathcal{A}, \mathcal{C} and matrix valued B , represents $\mathcal{A} = \{a = Bc \mid c \in \mathcal{C}\}$.
$F(t)$ and F_t	Used interchangeably to refer to some function evaluated at t .




Fig. 2: The optimal curve in the Trust Region Problem. Each point on the curve respects the “stationarity conditions” of optimality.

A. Convex Ellipsoidal Corridors

[17] proposed a framework for constructing parametric collision free corridors in which each cross section is an ellipsoid whose center, orientation, and shape vary smoothly with a path parameter ϵ . By expressing these parameters as low degree polynomials of ϵ , they obtain a differentiable corridor that follows the reference path and maximizes free space volume while remaining obstacle free. At each ϵ the cross section is restricted to the plane orthogonal to the reference trajectory and can be written as

$$\mathcal{C}_\epsilon = \left\{ q = r_\epsilon + Tq_\perp \mid q_\perp^\top E_\epsilon q_\perp + d_\epsilon^\top q_\perp \leq 1, q_\perp \in \mathbb{R}^{N_q-1} \right\},$$

where T defines the axes that span the plane perpendicular to \dot{r}_ϵ . $E_\epsilon \succ 0$, d_ϵ define the ellipsoid. The corridor parameters are obtained by solving a linear program that maximizes corridor size subject to collision avoidance constraints.

B. Solving the Trust Region Problem

The classical Trust Region Problem (TRP) minimizes a quadratic objective subject to a norm bound

$$\begin{aligned} \min_x \quad & \frac{1}{2}x^\top Qx + g^\top x \\ \text{s.t.} \quad & \|x\| \leq \Delta. \end{aligned} \quad (1)$$

The TRP is crucial to many problems in robotics, and has enabled better methods for nonconvex optimization, optimization on Riemannian Manifolds [19], optimization with local approximations of dynamics [20], and even in




Fig. 3: A decomposition of the space with polytopic corridors (in green). The polytopes are inflated about linear segments (red) to exclude collision points (grey).

highly efficient quadratic program solvers like HPIPM [21] because of its exploitation of problem structure for enhanced solving time. From the KKT conditions, if $-Q^{-1}g$ is not in the interior of the feasible set, the trust region solution satisfies

$$(Q + \lambda I)x^* = -g, \quad \|x^*\| = \Delta$$

for some scalar $\lambda \geq 0$. The resulting curve of points respecting ‘stationarity conditions’,

$$x(\lambda) = -(Q + \lambda I)^{-1}g \quad (2)$$

is the *optimal curve*, spanning from the unconstrained minima $x(0) = -Q^{-1}g$ to the origin as $\lambda \rightarrow \infty$ as shown in Fig. 2. If the unconstrained optimum lies outside the trust region, finding the constrained optimum reduces to a one dimensional search over λ to meet the norm bound. TRP solvers were explored in [22], [23].

This search can be accelerated with an eigendecomposition $Q = V\Lambda V^T$, which yields

$$(Q + \lambda I)^{-1} = V(\Lambda + \lambda I)^{-1}V^T,$$

allowing efficient evaluation of $\|(Q + \lambda I)^{-1}g\|^2 = V(\Lambda + \lambda I)^{-2}V^T$ at each iteration. $(\Lambda + \lambda I)$ remains a diagonal matrix, making this inversion an elementwise operation.

III. NONCONVEX OBSTACLE FREE CORRIDOR PARAMETERIZATIONS

In this section we present a method for parameterizing trajectories in collision-free space using the ellipsoidal corridors described above. We represent the corridor as a set-valued trajectory $\mathcal{C}(\epsilon)$, shown as the blue region in Fig. 4, where $0 \leq \epsilon \leq 1$ is the normalized ‘time’ along the path ($\epsilon = 0$ at the start and $\epsilon = 1$ at the end).

In this work, we take an approach that defines trajectories in the configuration space by their offset from a reference trajectory. This reference trajectory serves as the centerline of our ellipsoidal corridor, and is denoted by \tilde{q}_ϵ . Although \tilde{q}_ϵ can naturally be a feasible candidate of the trajectory of the system, it doesn’t fully exploit the full obstacle-free configuration space, leading to conservative and often suboptimal trajectories. Hence, we would like to leverage a proxy of the obstacle-free space to optimize for the trajectory, which is treated as a variation q_ϵ of the centerline of the

ellipsoidal corridor \tilde{q}_ϵ . We can set our trajectory in the configuration space to be

$$q_\epsilon = L_\epsilon y + \tilde{q}_\epsilon, \quad (3)$$

where $y \in \mathbb{R}^{N_y}$ is a vector of trajectory parameters our proposed method optimizes for and L_ϵ is a possibly non-polynomial function of ϵ that we will construct in the following Section. III-A, III-B, III-C.

To ensure that resulting trajectories are collision free, we will restrict $L_\epsilon y$ to be an offset from \tilde{q}_ϵ such that the trajectory point is still within our collision-free ellipsoidal cross section. With this parameterization, optimizing a quadratic objective on the configuration or its derivatives is

$$\min_y \sum_{\epsilon_i \in [0,1]} \|Q_i (L_{\epsilon_i}^{(n)} y + \tilde{q}_{\epsilon_i}^{(n)}) + g_i\| \quad (4)$$

$$\text{s.t. } q_{\epsilon_i} = L_{\epsilon_i} y + \tilde{q}_{\epsilon_i} \quad (5)$$

$$\tilde{q}_{\epsilon_i} \in \mathcal{C}_{\epsilon_i} \quad (6)$$

where $q_\epsilon^{(n)} = L_\epsilon^{(n)} y + \tilde{q}_\epsilon^{(n)}$ is the n th derivative of the trajectory q_ϵ . These constraints are complicated, and scale with the number of samples (ϵ_i) used. We will instead define a set of valid parameter values \mathcal{A} and parameterization L_ϵ , such that

$$\mathcal{C}_\epsilon = L_\epsilon \mathcal{A} + \tilde{q}_\epsilon.$$

From this parameterization, any $y \in \mathcal{A}$ results in a configuration within collision free space, reflected in

$$q_\epsilon = L_\epsilon y + \tilde{q}_\epsilon \in \mathcal{C}_\epsilon \text{ if } y \in \mathcal{A},$$

reducing the constraints in the optimization problem ((5)-(6)) to simply $y \in \mathcal{A}$. Note that \tilde{q}_ϵ serves as the center of \mathcal{C}_ϵ .

We now describe how we parameterize collision-free trajectories within the time-varying ellipsoidal corridors. We build up the parameterization in three steps. First, we show how a simple ball can be mapped to an ellipsoidal corridor cross-section. Next, we extend this to a lifted higher-dimensional space to introduce more trajectory diversity. Finally, we generalize to a Cartesian product of balls to capture boundary-hugging trajectories. This section thus moves from the simplest to the richest representations, culminating in our full parameterization.

Unlike [17], which restricts ellipsoidal cross-sections to degenerate planar slices perpendicular to a reference trajectory, our approach allows full-volume ellipsoids in the original configuration space. We define each cross-section by

$$\mathcal{C}_\epsilon = \{q \mid \|C_\epsilon(q - \tilde{q}_\epsilon)\|^2 \leq 1\}, \quad (7)$$

and use this definition to construct the mappings from \mathcal{A} to \mathcal{C}_ϵ . Here, C_ϵ represents the axes of the ellipsoid. In the following, we are going to present how we construct the set \mathcal{A} and the corresponding mappings from \mathcal{A} to \mathcal{C}_ϵ .




Fig. 4: A corridor (blue) as a time-varying projection of a lifted hypersphere (orange). Example shown for $N_q = 2$ and $N_y = 3$.

A. Corridors as transformations of balls ($\mathcal{A} = \mathcal{S}^{N_q}$)

We begin with the simplest case, mapping a unit ball \mathcal{S}^{N_q} in configuration space into the time-varying ellipsoidal corridor cross sections.

Our cross section is defined by (7) which is equivalent to

$$\mathcal{C}_\epsilon = C_\epsilon^{-1} \mathcal{S}^{N_q} + \tilde{q}_\epsilon = \left\{ q = C_\epsilon^{-1} y + \tilde{q}_\epsilon \mid \|y\|^2 \leq 1 \right\}. \quad (8)$$

$\mathcal{S}^{N_q} = \{y \in \mathbb{R}^{N_q} \mid \|y\|^2 \leq 1\}$ is the unit ball in \mathbb{R}^{N_q} . Equation (8) defines \mathcal{C}_ϵ as a varying transformation of a ball, \mathcal{S}^{N_q} , and defines the associated map as $L_\epsilon = C_\epsilon^{-1}$.

Note that, in the above definition, columns of C_ϵ^{-1} define the ellipsoid axes. In the special case of a ball, $C_\epsilon^{-1} y = r I^{N_q} y$, effectively scales y by r .

let us find the trajectory that a single point in this ϵ -varying ellipsoid traces. Given some point $y \in \mathcal{S}^{N_q}$, its corresponding trajectory is

$$q_\epsilon = C_\epsilon^{-1} y + \tilde{q}_\epsilon.$$

Because y is in $\mathcal{A} = \mathcal{S}^{N_q}$, the trajectory value q_ϵ remains in \mathcal{C}_ϵ for all ϵ . Conversely, only a single unique trajectory (y) achieves configuration q at ϵ . The parameter for this trajectory is given by the inverse mapping,

$$y = C_\epsilon(q - \tilde{q}_\epsilon).$$

B. Corridors as Projections of lifted balls ($\mathcal{A} = \mathcal{S}^{N_y}$)

However, restricting ourselves to $y \in \mathcal{S}^{N_q}$ limits the diversity of trajectories. With ϵ fixed, the above mapping between q and y is bijective, meaning each q corresponds to exactly one trajectory (y). To overcome this, we lift to a higher-dimensional ball \mathcal{S}^{N_y} where $N_y > N_q$, as in Fig. 4. We define our new map L_ϵ such that

$$\mathcal{C}_\epsilon = L_\epsilon \mathcal{S}^{N_y} + \tilde{q}_\epsilon, \quad N_y > N_q,$$

where $L_\epsilon \in \mathbb{R}^{N_q \times N_y}$ now defines an injective map from y to q . This creates an affine subspace

$$Y(\epsilon, q) = \{y \mid q = L_\epsilon y + \tilde{q}_\epsilon\} \quad (9)$$

of parameter values that achieve q at ϵ . The intersection of this space and the feasible set ($Y(\epsilon, q) \cap \mathcal{S}^{N_y}$) reflects




Fig. 5: Projection of a lifted hypersphere to an ellipsoidal cross section: rotation in \mathbb{R}^{N_y} , projection to \mathbb{R}^{N_q} , and stretching to form \mathcal{C}_ϵ .

the volume of feasible trajectories in $\mathcal{A} = \mathcal{S}^{N_y}$ that pass through q at ϵ . By using this lifted space, we introduce a multitude of feasible trajectories that achieve a possible desired configuration at ϵ . By constraining $y \in \mathcal{S}^{N_y}$ we guarantee the resulting trajectory stays inside the collision free corridor, while also obtaining a single norm constraint suitable for efficient Trust Region Subproblem solvers.

Building L_ϵ . We will build L_ϵ by composing two transformation operations. The first is a projection from \mathcal{S}^{N_y} into \mathcal{S}^{N_q} . To project \mathcal{S}^{N_y} into \mathcal{S}^{N_q} , we require a projection matrix L_ϵ^{proj} . Combining L_ϵ^{proj} with C_ϵ^{-1} yields the ellipsoidal cross section of the corridor at ϵ ,

$$\mathcal{C}_\epsilon = L_\epsilon \mathcal{S}^{N_y} + \tilde{q}_\epsilon, \quad L_\epsilon = C_\epsilon^{-1} L_\epsilon^{\text{proj}}, \quad (10)$$

where

$$q_\epsilon = L_\epsilon y + \tilde{q}_\epsilon, \quad q(\epsilon, y) \in \mathcal{C}_\epsilon \text{ if } \|y\|^2 \leq 1.$$

We first show the conditions for the projection to \mathcal{S}^{N_y} to be valid, then we attempt to find a valid L_ϵ^{proj} that respects the presented conditions for all ϵ .

Validity of Projection Matrices. We show here that any projection matrix $L_\epsilon^{\text{proj}} \in \mathbb{R}^{N_q \times N_y}$ with **orthonormal rows** defines a valid mapping from the unit ball in \mathbb{R}^{N_y} to the unit ball in \mathbb{R}^{N_q} . Geometrically, applying this matrix to \mathcal{S}^{N_y} is equivalent to rotating the ball in \mathbb{R}^{N_y} and then taking its N_q -dimensional cross section, as illustrated in Fig. 5.

These operations can be written as

$$L_\epsilon^{\text{proj}} = [I_{N_q} \ 0] R, \quad \text{with } R \in O(N_y),$$

which corresponds to selecting the first N_q rows of an orthogonal matrix R , resulting in an $N_q \times N_y$ matrix with orthonormal rows. Conversely, given the top N_q rows, we can always reconstruct an orthogonal $N_y \times N_y$ matrix R (for example, via Gram–Schmidt completion).

Constructing L_ϵ^{proj} . We parameterize L_ϵ^{proj} as

$$L_\epsilon^{\text{proj}} = I_{N_q} \otimes \bar{B}_\epsilon^\top, \quad \bar{B}_\epsilon \in \mathbb{R}^{\frac{N_y}{N_q}}, \quad \|\bar{B}_\epsilon\| = 1,$$

where \otimes is the Kronecker product. This guarantees each row has unit norm and rows are mutually orthogonal. We generate \bar{B}_ϵ as the normalized basis of a Bézier curve to obtain a smooth time-varying projection.

C. Trajectories in the Lifted Space ($\mathcal{A} = (\mathcal{S}^{N_y})^{N_J}$)

In spite of the ability of the above parameterization ($\mathcal{A} = \mathcal{S}^{N_y}, L_\epsilon = C_\epsilon^{-1} L_\epsilon^{\text{proj}}$) to describe multiple trajectories that pass through a given configuration, it does not have the




Fig. 6: Feasible lifted parameters y (black) projecting to two different q values at ϵ . Points closer to the corridor center admit a larger set of y , while boundary points admit fewer.

ability to describe qualitatively ideal trajectories, motivating a new representation of \mathcal{A} as a product of balls, $\mathcal{A} = \mathcal{S}^{N_y} \times \mathcal{S}^{N_y} \times \dots \times \mathcal{S}^{N_y}$.

a) Density of Describable Trajectories.: To see how the previous parameterization limits feasible trajectories, consider $Y(\epsilon, q)$, the set of trajectories achieving configuration q at ϵ . From (9), a basic solution for $y \in Y(\epsilon, q)$ is $y^* = (L_\epsilon)^\dagger (q - \tilde{q}_\epsilon)$. All other solutions in this set are offset from y^* by vectors in $\ker(L_\epsilon)$ so that

$$Y(\epsilon, q) = y^* + z \quad \text{where } y^* \perp z, \quad z \in \ker(L_\epsilon).$$

The feasible set of trajectories in $Y(\epsilon, q)$, given by $Y(\epsilon, q) \cap \mathcal{S}^{N_y}$ yields a ball in $\ker(L_\epsilon)$ of radius $r(\epsilon, q) = \sqrt{1 - \|y^*\|^2} = \sqrt{1 - \|(L_\epsilon)^\dagger (q - \tilde{q}_\epsilon)\|^2}$. The radius decreases with the Mahalanobis distance from the center, $\|(L_\epsilon)^\dagger (q - \tilde{q}_\epsilon)\|$, as visualized in Fig. 6. This shows that the previous parameterization naturally concentrates trajectories near the center of \mathcal{C}_ϵ while limiting exploration near its edges, which is undesirable when boundary-hugging paths are required.

b) Time-Varying Parameters: To counter this concentration effect, we expand the parameter space to $\mathcal{A} = (\mathcal{S}^{N_y})^{N_J}$, the cartesian product of N_J unit balls in \mathbb{R}^{N_y} .

$$\mathbf{y} = [y_1, y_2, \dots, y_{N_J}], \quad y_j \in \mathcal{S}^{N_y} \quad \forall j.$$

We define an ϵ -varying transformation from $(\mathcal{S}^{N_y})^{N_J}$ to \mathbb{R}^{N_y} that traces a trajectory contained within the ball \mathcal{S}^{N_y} . This specifically results in $y_\epsilon = B_\epsilon \mathbf{y}$, where $y_\epsilon \in \mathcal{S}^{N_y}$, as visualized in Fig. 1. We can generate this mapping by taking an ϵ -varying convex combination of y_1, \dots, y_{N_J} (e.g., a Bézier curve), which automatically stays inside the unit ball. This richer parameterization increases the amount of feasible trajectories that pass near the boundary of the corridor, capturing curved or boundary-hugging paths while maintaining the simple ball constraint structure that enables efficient Trust Region methods for optimization.

With $y(\epsilon) = B_\epsilon \mathbf{y}$ defining this trajectory, the projection from $(\mathcal{S}^{N_y})^{N_J}$ is given by

$$C_\epsilon = L_\epsilon^Y (\mathcal{S}^{N_y})^{N_J} + \tilde{q}_\epsilon, \quad L_\epsilon^Y = C_\epsilon^{-1} L_\epsilon^{\text{proj}} B_\epsilon,$$

resulting in a trajectory in the configuration space given by

$$q(\epsilon, \mathbf{y}) = L_\epsilon^Y \mathbf{y} + \tilde{q}_\epsilon.$$

IV. SOLVING SPECIALIZED QCQPs

We have now introduced three main parameterizations of trajectories. In the first two, parameters are restricted to a single ball ($y \in \mathcal{S}^{N_y}$ or $y \in \mathcal{S}^{N_y}$). In the third, parameters restricted to a product of balls ($\mathbf{y} \in (\mathcal{S}^{N_y})^{N_J}$).

In the first two cases, our optimization problem is a Trust Region Problem. Based on this observation, let us develop specialized optimization methods that exploit the structure arising in the third case.

The Orthogonal Trust Region Problem. With the parameterization using a product of balls, the trajectory optimization takes the form

$$\min_{\mathbf{y}=[y_1, \dots, y_{N_J}]^\top} \frac{1}{2} \mathbf{y}^\top Q \mathbf{y} + g^\top \mathbf{y} \quad (11)$$

$$\text{s.t. } \|y_j\|^2 \leq 1 \quad \forall 1 \leq j \leq N_J \quad (12)$$

The feasible set of this problem is made of ball constraints enforced on orthogonal sets of axes. We dub this problem the ‘Orthogonal Trust Regions Problem’.

Separability. An important observation is that the norm constraints (12) are separable, as each block variable y_j has its own independent constraint $y_j \in \mathcal{S}^{N_y}$.

Let us consider the optimization problem along only a cross section of the variables (a single block), with other variables fixed. The resulting subproblem is

$$\min_{y_i} \frac{1}{2} y_i^\top Q_{i,i} y_i + \left(g_i + \sum_{j \neq i} Q_{i,j} y_j \right)^\top y_i \quad (13)$$

$$\text{s.t. } \|y_i\|^2 \leq 1$$

which is a traditional trust region problem. Note here that $Q_{i,j}$ is a submatrix of Q whose columns correspond to the block y_i and rows to y_j . Since Q is symmetric, $Q_{i,j}^\top = Q_{j,i}$.

Due to the separability of the problem, the feasible set for y_i is constant, regardless of values for $y_{j \neq i}$. This is a desirable structure, as the iteration on the solution for each block can be done by solving (or approximately solving) a TRP.

This motivates a block coordinate descent (BCD) strategy for solving (11). In each iteration we cycle through the blocks, updating y_i by solving (13). Rather than solving each subproblem to exact optimality, we may perform only a single dual update step for the corresponding multiplier λ_i , followed by the primal update y_i . This yields an inexact BCD strategy that exploits the Trust Region like structure to make each step computationally inexpensive.

Geometric intuition. Let us look at the subproblem in each separable direction. From (2) and (13), the optimal curve in direction i is given by

$$y_i(\lambda_i) = -(Q_{i,i} + \lambda_i I)^{-1} \left(g_i + \sum_{1 \leq j \leq N_J} Q_{i,j} y_j \right)$$

while the feasible set is constant.

We can view the problem as N_J interconnected trust region problems. As we update the guess for the optimal value in some direction, y_j , the objective of the other $N_J - 1$




Fig. 7: A visualization of the Orth-TRP problem as $N_J = 3$ interconnected Trust Region Problems. As the guess for x_1 is updated, the optimal curves for x_2 and x_3 are influenced, changing the optimal solution for those parallel problems.

Algorithm 1 Orth-TRP solver with Moré–Sorensen block updates

Require: $Q_{ij} \in \mathbb{R}^{N_y \times N_y}$ with $Q_{ii} \succ 0$, vectors $g_i \in \mathbb{R}^{N_y}$, $y_i \in \mathcal{S}^{N_y}$ for $i = 1, \dots, N_J$.

- 1: Precompute eigendecomposition of each Q_{ii} .
- 2: Initialize $y_i^{(0)} \leftarrow 0$; $\lambda_i^{(0)} \leftarrow \infty$ set $k \leftarrow 0$.
- 3: **repeat**
- 4: **for** $i = 1$ to N_J **do** \triangleright Gauss–Seidel sweep
- 5: $c_i^{(k)} \leftarrow g_i + \sum_{j \neq i} Q_{ij} y_j^{(k + \mathbb{1}_{[j < i]})}$
- 6: $y_i^{(k+1)}, \lambda_i^{(k+1)} \leftarrow \text{TRP_Step}_{\text{Moré–Sorensen}}(Q_{ii}, c_i^{(k)}, \lambda_i^{(k)})$
- 7: **end for**
- 8: $k \leftarrow k + 1$
- 9: **until** converged
- 10: **return** $[y_1^{(k)}; \dots; y_{N_J}^{(k)}]$

subproblems (and thus the relevant optimal curve) is changed slightly, as shown in Fig. 7.

Algorithm. We can implement the proposed solver in Algorithm 1, using a Moré–Sorensen style trust region step for each block. Only $c_i = g_i + \sum_{j \neq i} Q_{ij} y_j$ changes between iterations, while Q_{ii} remains constant; its eigendecomposition can be precomputed and reused. Each TRP_STEP updates the dual variable, λ_i for the direction i .

V. RESULTS

While our approach is broadly applicable to collision free planning, we evaluate our method on a quadrotor trajectory optimization problem designed to highlight how solver performance scales with trajectory complexity.

Experimental Setup. We evaluate our approach on a quadrotor trajectory optimization task to show how solver performance scales with trajectory complexity. The environment is a 3D grid of adjacent rectangular “rooms,” with collision points sampled along shared edges (Fig. 8a). Each experiment attempts to generate a trajectory moving from one room to the next, minimizing acceleration. We measure two main outcomes. The first is solver time, contrasting a large polytope based QP (solved with OSQP a state of the art QP solver) with our Orth-TRP solver. The second is trajectory quality, quantified by the relative improvement in

average acceleration over a naively generated trajectory. By varying the number of waypoints and homotopy classes, we analyze the performance.

Initial Trajectories and Homotopy Classes. We generate a linearly interpolated initial guess for each trial, where a trajectory visiting rooms “A–B–C” interpolates through the shared walls between A and B, then between B and C. Each initial guess defines a distinct homotopy class. The number of waypoints, varied from 11 to 30, controls trajectory curvature and acts as a proxy for topological complexity.

Smooth Reference Trajectory and Corridors. We fit a 15 control point Bézier curve to the linearly interpolated trajectory to generate a smooth reference trajectory with zero initial and final velocity. Around each sampled point we solve a linear program to fit a large, axis aligned ellipsoid. Interpolating between the parameters of these ellipsoids helps generate the time-varying corridor \mathcal{C} .

Decision Variables. Our planner uses $N_y = 6$ and $N_J = 14$ orthogonal blocks, fixing the first and last two blocks to enforce zero velocity at the boundaries. The remaining variables capture free directions within the ellipsoidal corridor. For the polytope baseline, we enforce continuity of position, velocity, and acceleration, with 5 control points (resulting in $5N_q = 15$ variables) per polytope; with N waypoints, this yields $15(N - 1)$ variables in total.

Baseline Methods. *Polytopic Corridor Method [15]:* We implemented the planner in [15], which builds convex polytopes along the initial trajectory, and uses those to generate a polytopic corridor. We allocate time across polytopes by computing the ‘start time’ for each segment in the polytope with

$$s_n^* = \frac{\text{arc length for segments before } n}{\text{total arc length}}$$


$$\text{start time}^n = \text{smoothstep}^{-1}(s_n^*) \times \text{horizon}$$

which proportionally enlarges time for longer segments. Inverting the ‘smoothstep’ function allows this approach to allocate larger amounts of time for the start and end of the trajectory to account for acceleration from (or deceleration to) rest. Although this is a heuristic, it produces smoother timing than uniform allocation.

Bubble Planning: We also reproduced the “Bubble Planning” method, which uses overlapping spheres along the path and solves a QCQP with NLOPT. Waypoints are chosen from the intersections of these spheres, and the solution is formed by interpolating between them. As the waypoint count increased (necessary for the complex trajectories considered here), the interpolated paths exhibited unpredictable, oscillatory behavior, and the resulting trajectories became infeasible in all trials with more than 13 waypoints.

A. Solving Time vs. Trajectory Complexity

In contrast to existing corridor-based planners, the runtime and problem size of our method does not scale with regards to the solution complexity or the geometric complexity of the environment.



(a) An example solution trajectory (in green) using ellipsoidal corridors in our grid of rooms. This is contrasted with the smooth reference trajectory (red). Note that the solution is able to ‘hug’ the boundary of the corridor, and pass near obstacles. (b) An example ellipsoidal corridor in the collision free space of our example environment. Trajectories are restricted to be within this corridor.

Fig. 9 directly demonstrates this. Because the lifted, ellipsoidal parameterization and the Orth-TRP solver decouple problem size from geometric complexity, solver runtime stays nearly constant even as the number of waypoints increases dramatically. At low complexity ($N = 11$ waypoints), our method is already faster than the polytope-based QP. As the path becomes more complex, this gap widens. By contrast, the solving time of the polytope-based solver scales with the number of waypoints, reflecting the burden of utilizing increasingly numerous convex regions to describe environments with complicated geometries. This result confirms the core contribution of our approach; by lifting trajectories into a product-of-balls representation and solving an Orth-TRP, we achieve substantial computational gains over traditional planners.

The Orth-TRP solver itself consists of two stages. The first is a parallelizable eigendecomposition of $N_J N_y \times N_y$ matrices (which dominates computation time). The second is an inexpensive per-block trust-region updates with matrix multiplication and elementwise inversion operations. In our experiments, eigendecomposition was about $8\times$ more expensive than the iterations, highlighting how the solver’s structure concentrates computational effort in a single, parallelizable step. This makes it feasible to plan smooth, dynamically feasible trajectories in environments far more complex than those tractable with polytopic corridors.

B. Comparison of Solution Optimality

Our method avoids the sharp acceleration peaks characteristic of other corridor-based approaches that assign fixed time intervals to each segment, as illustrated in Fig. 10, which highlights how our approach resolves this issue. Across all tested complexities, Orth-TRP yields consistently smoother, more dynamically feasible trajectories compared to both the naive smooth reference and the polytope-based baseline. Representative acceleration profiles in Fig. 11 show that our method avoids the pronounced peaks that occur at discrete switching points (marked by red dots) in polytopic corridors.

This improvement stems directly from our key contribution, replacing discrete, time-allocated polytopes with a

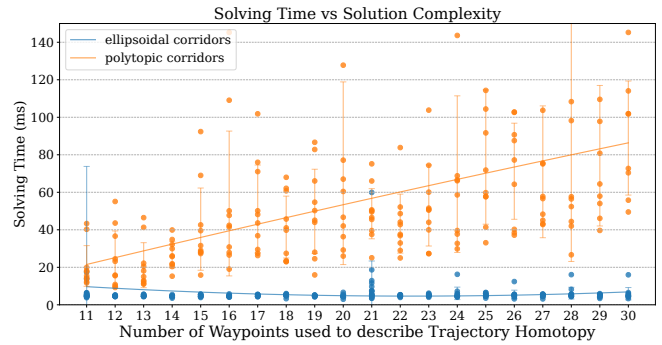


Fig. 9: Solver runtimes versus trajectory complexity. Polytopic corridors (orange) become slower as waypoint count increases, while ellipsoidal corridors with Orth-TRP (blue) maintain nearly constant runtime across all tested complexities.

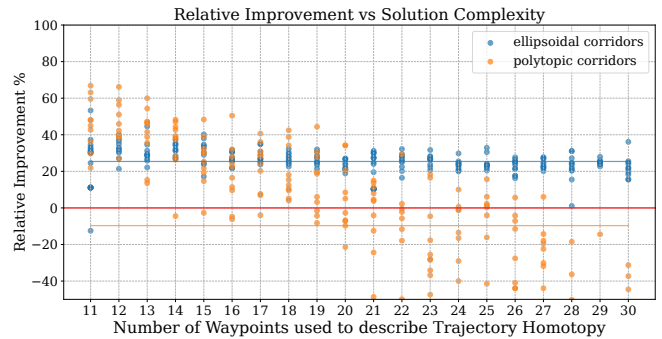


Fig. 10: Average acceleration of optimized trajectories versus a naive reference. Orth-TRP trajectories (blue) remain smoother across all complexities than polytopic corridors (orange), with the advantage increasing as the path becomes more complex. The average cost improvement across all trials (horizontal lines) was negative for the polytopic corridors.

continuously deforming ellipsoidal corridor parameterization. By allowing the trajectory to evolve inside smoothly varying ellipsoidal cross sections and lifting to a Cartesian product of balls, we create a feasible set that naturally accommodates boundary-hugging and curved paths without forcing discontinuous accelerations. This explains why our method achieves lower acceleration costs and better scaling with increasing geometric complexity.

Overall, these results demonstrate that Orth-TRP not only solves the scalability problem but also mitigates the time-allocation and discontinuity issues that have historically limited corridor-based planners. By combining an efficient, block-separable QCQP solver with a rich parameterization of feasible paths, our approach produces trajectories that are both faster to compute and higher in quality in highly complicated environments.

VI. CONCLUSIONS

We have presented a new framework for collision free trajectory optimization that treats safe corridors as projections of orthogonal trust regions. By lifting trajectories into a

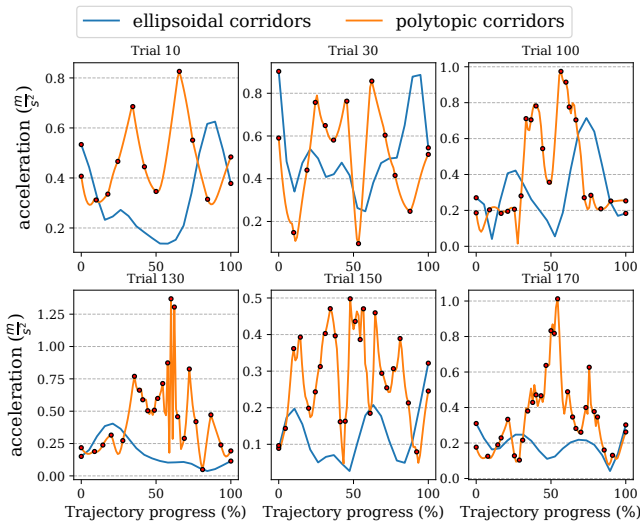


Fig. 11: Representative acceleration profiles showing how Orth-TRP trajectories (blue) avoid the sharp peaks visible at discrete switching points (red dots) in polytopic corridors (orange).

higher dimensional parameter space and representing feasible sets as Cartesian products of balls, we project this structure into time-varying ellipsoidal corridors. This decouples geometric complexity from problem size while maintaining a convex and scalable formulation. Unlike traditional corridor methods, which require difficult time allocation across convex regions, our approach allows trajectories to evolve smoothly inside ellipsoidal corridors without assigning fixed time intervals. Building on this representation, we formulated the Orthogonal Trust Regions Problem (Orth-TRP), a specialized QCQP whose separable nature results in optimization over interconnected Trust Region subproblems. Our Orth-TRP solver exploits this structure by reusing per-block eigendecompositions and solving each trust region dual via a simple scalar search, producing a parallelizable solver.

Future work will extend this parameterization to higher-DOF systems, cooperative manipulation, and receding-horizon planning with inexact Newton updates over the eigendecomposition, enabling a sequential updating method for similar problems. In summary, this work demonstrates a principled, convex, and scalable approach to planning highly curved trajectories in complex environments, opening new possibilities for advanced robotic motion planning.

REFERENCES

- [1] T. Marcucci, M. Petersen, D. von Wrangel, and R. Tedrake, "Motion planning around obstacles with convex optimization," *Science Robotics*, vol. 8, no. 84, p. eadf7843, 2023.
- [2] A. Pallar, G. Li, and G. Loianno, "Optimal trajectory planning for cooperative manipulation with multiple quadrotors using control barrier functions," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [3] J. Mao, G. Li, S. Nogar, C. Kroninger, and G. Loianno, "Aggressive visual perching with quadrotors on inclined surfaces," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2021.

- [4] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," Department of Computer Science, Iowa State University, Ames, IA, USA, Tech. Rep. TR 98-11, oct 1998.
- [5] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, Jun. 2011.
- [6] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [7] K. Tracy, T. A. Howell, and Z. Manchester, "Differentiable collision detection for a set of convex primitives," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3663–3670.
- [8] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Science Robotics*, vol. 6, no. 56, p. eabh1221, Jul. 2021.
- [9] G. Li and G. Loianno, "Nonlinear model predictive control for cooperative transportation and manipulation of cable suspended payloads with multiple quadrotors," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [10] A. Thirugnanam, J. Zeng, and K. Sreenath, "Safety-Critical Control and Planning for Obstacle Avoidance between Polytopes with Control Barrier Functions," in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 286–292.
- [11] M. Goarin, G. Li, A. Saviolo, and G. Loianno, "Decentralized nonlinear model predictive control for safe collision avoidance in quadrotor teams with limited detection range," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [12] A. Jaitly, D. Jha, K. Ota, and Y. Shirai, "Analytic conditions for differentiable collision detection in trajectory optimization."
- [13] R. Deits and R. Tedrake, "Computing Large Convex Regions of Obstacle-Free Space Through Semidefinite Programming," in *Algorithmic Foundations of Robotics XI*, H. L. Akin, N. M. Amato, V. Isler, and A. F. van der Stappen, Eds. Cham: Springer International Publishing, 2015, vol. 107, pp. 109–124.
- [14] M. Petersen and R. Tedrake, "Growing convex collision-free regions in configuration space using nonlinear programming," *arXiv preprint arXiv:2303.14737*, 2023, preprint.
- [15] S. Liu, M. Watterson, K. Mohta, K. Sun, S. Bhattacharya, C. J. Taylor, and V. Kumar, "Planning Dynamically Feasible Trajectories for Quadrotors Using Safe Flight Corridors in 3-D Complex Environments," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1688–1695, Jul. 2017.
- [16] Y. Ren, F. Zhu, W. Liu, Z. Wang, Y. Lin, F. Gao, and F. Zhang, "Bubble Planner: Planning High-speed Smooth Quadrotor Trajectories using Receding Corridors," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2022, pp. 6332–6339.
- [17] J. Arrizabalaga, Z. Manchester, and M. Ryll, "Differentiable Collision-Free Parametric Corridors," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2024, pp. 1839–1846.
- [18] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: an operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.
- [19] P.-A. Absil, C. G. Baker, and K. A. Gallivan, "Trust-region methods on riemannian manifolds," *Foundations of Computational Mathematics*, vol. 7, no. 3, pp. 303–330, 2007.
- [20] H. J. T. Suh, T. Pang, T. Zhao, and R. Tedrake, "Dexterous contact-rich manipulation via the contact trust region," 2025. [Online]. Available: <https://arxiv.org/abs/2505.02291>
- [21] G. Frison and M. Diehl, "HPIPM: A high-performance quadratic programming framework for model predictive control," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 6563–6569, 2020.
- [22] D. M. Gay, "Computing optimal steps for trust region steps," Bell Laboratories, Murray Hill, NJ, USA, Tech. Rep. no. 6, 1981, this is a technical memorandum that was influential in the development of trust-region methods.
- [23] J. J. Moré and D. C. Sorensen, "Computing a trust region step," *SIAM Journal on Scientific and Statistical Computing*, vol. 4, no. 3, pp. 553–572, 1983.