

GHOST: Solving the Traveling Salesman Problem on Graphs of Convex Sets

Jingtao Tang, Hang Ma

Simon Fraser University
{jingtao.tang, hangma}@sfu.ca

Abstract

We study GCS-TSP, a variant of the Traveling Salesman Problem (TSP) defined over a Graph of Convex Sets (GCS)—a powerful representation for trajectory planning that decomposes the configuration space into convex regions connected by a sparse graph. In GCS-TSP, edge costs are not fixed but depend on the specific trajectory passing through each convex region, making classical TSP methods inapplicable. We introduce GHOST, a hierarchical framework that optimally solves GCS-TSP by combining combinatorial tour search with convex trajectory optimization. GHOST systematically explores tours on a complete graph induced by the GCS, using a novel abstract-path-unfolding algorithm to compute admissible lower bounds that guide best-first search at both the high level (over tours) and the low level (over feasible GCS paths realizing the tour). These bounds provide strong pruning power, reducing unnecessary optimization calls. We prove that GHOST guarantees optimality and present a bounded-suboptimal variant for time-critical settings. Experiments show that GHOST is orders-of-magnitude faster than unified mixed-integer convex programming baseline while uniquely handling complex problems involving high-order continuity constraints and incomplete GCSs.

Homepage — <https://sites.google.com/view/ghost-gcs-tsp>

Code — <https://github.com/reso1/ghost>

Introduction

Robotic motion planning often relies on trajectory optimization, particularly in high-dimensional spaces subject to kinodynamic constraints. However, even a single obstacle in the configuration space can render the problem nonconvex and computationally challenging (Ratliff et al. 2009; Schulman et al. 2014). A common approach is to reformulate such problems as convex programs for scalability. To this end, Marcucci et al. (2023) proposed the *Graph of Convex Sets* (GCS) framework, which decomposes the configuration space into convex regions connected via a graph structure. This representation enables transforming nonconvex trajectory planning into convex optimization subproblems while preserving feasibility and continuity. Building on this foundation, prior work has extensively studied shortest-path problems on GCSs (Marcucci et al. 2024b), enabling

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: GCS-TSP solution by GHOST. Gray polygons are convex sets. Trajectory color follows the bar (top-bottom).

globally optimal navigation under convex constraints. Yet, many real-world robotics tasks—such as area coverage, inspection, and structured exploration—require visiting multiple regions in a cost-effective order. This naturally leads to the *Traveling Salesman Problem on a GCS* (GCS-TSP).

The input to GCS-TSP is a GCS whose vertices represent convex subsets of continuous configuration space and whose edges represent feasible transitions. The goal is to compute a tour that visits each convex set at least once and a continuous trajectory through representative points within these sets, satisfying trajectory constraints. This problem tightly couples combinatorial tour optimization with continuous kinodynamic feasibility. Fig. 1 depicts an example solution generated by our method in a 2D environment. Applications include coverage and inspection planning, constrained multi-target visitation tasks, and task and motion planning.

Related Work

Graphical TSP: The TSP seeks the shortest Hamiltonian cycle on a weighted graph, given pairwise distances between vertices. When the input graph is incomplete and may not admit a Hamiltonian cycle, the problem is typically solved on its *metric closure* (Cornuéjols, Fonlupt, and Naddef 1985)—a complete graph with edges weighted by shortest-path distances in the original graph. However, edge weights in a GCS are not fixed: Each transition cost depends on the specific trajectory endpoints within convex sets and may be subject to trajectory constraints. Thus, the GCS setting does not admit a static metric closure, and the TSP must be solved over dynamically evaluated edge costs.

Combinatorial Optimization in GCSs: Advances in constructing collision-free convex sets (Werner et al. 2024; Dai et al. 2024) have enabled efficient shortest-path solv-

ing in GCSs (Marcucci et al. 2024b), supporting applications such as non-Euclidean motion planning (Cohn et al. 2025), quadrotor navigation (Marcucci et al. 2024a), and multi-robot planning (Tang et al. 2025). Marcucci (2024) formulated a Mixed Integer Convex Program (MICP) for GCS-TSP, later extended to moving-target settings (Philip et al. 2024), but restricted to Hamiltonian cycles on complete GCSs with poor scalability. Bhat et al. (2025) proposed a two-level approach for 3D moving-target TSP, where the high level incrementally expands a partial GCS path, relying on problem-specific time-window assumptions, small branching factors, and piecewise-linear motion models.

Coverage and Inspection Planning: TSP is central to coverage planning (Galceran and Carreras 2013), typically applied to grid graphs (Gabriely and Rimon 2001; Tang, Mao, and Ma 2025) or other workspace decompositions (Choset and Pignon 1998; Acar et al. 2002), followed by tour computation over the resulting regions. These methods often constrain robots to disk-shaped sweeps, limiting flexibility in complex environments. Recent global approaches (Wu et al. 2019; Tang and Ma 2024) use smooth space-filling curves to cover workspaces without decomposition, but rely on predefined trajectory templates and do not consider kinodynamic constraints. Inspection planning, or the Watchman Route Problem (Ntafos 1992), is conceptually similar to GCS-TSP, allowing coverage by visibility in polygonal environments. However, existing methods typically rely on dense waypoint sampling on fixed graphs (Fu et al. 2019) or are limited to grid-based decompositions (Skylar et al. 2022).

Our Contributions

To address the limitations of prior work, we formulate GCS-TSP and present **GCS-HIERARCHICAL OPTIMAL SEARCH FOR TOURS (GHOST)**, a novel algorithmic framework that solves GCS-TSP optimally. GHOST employs a two-level hierarchical search: At the high level, it systematically explores TSP tours on the complete graph induced by the input GCS, ordered by lower bounds on their trajectory costs; At the low level, it searches for feasible GCS paths that realize each tour and computes the optimal trajectory for each path via convex optimization.

A central component of GHOST is a general *abstract-path-unfolding* algorithm that computes admissible lower bounds on the cost of realizing a given abstract tour—or more generally, any sequence of convex sets (e.g., abstract paths) in which consecutive sets may not be directly connected in the GCS. These lower bounds can be precomputed and used to guide the best-first exploration of feasible GCS paths, significantly improving efficiency. They also serve to prioritize high-level search for pruning unpromising tour candidates early, enabling provably optimal search while avoiding unnecessary convex optimization calls.

We prove GHOST is optimal under this search strategy. We also present a bounded-suboptimal variant that trades optimality for improved scalability. We evaluate and demonstrate GHOST across various planning tasks in robotics. Compared to the unified MICP baseline, GHOST is orders of magnitude faster in simple cases and uniquely capable of solving general instances previously out of reach.

Problem Formulation

In this section, we formally define GCS and GCS-TSP.

Graph of Convex Sets (GCS): A GCS $G = (V, E)$ is a directed graph. Each vertex $v \in V$ is associated with a non-empty, compact convex set \mathcal{X}_v . Each edge $(u, v) \in E$ exists iff transitions are allowed from vertices u to v . A path $\pi = (v_1, \dots, v_{|\pi|})$ is a sequence of vertices where $(v_i, v_{i+1}) \in E$ for all $i = 1, \dots, |\pi| - 1$. A trajectory $\tau = (\mathbf{x}_{v_1}, \dots, \mathbf{x}_{v_{|\pi|}})$ conditioned on π is a sequence of points where each $\mathbf{x}_{v_i} \in \mathcal{X}_{v_i}$. Additional convex constraints may be imposed on the trajectory, e.g., kinodynamic feasibility. Typically, for each edge $(v_i, v_{i+1}) \in E$, constraints of the form $(\mathbf{x}_{v_i}, \mathbf{x}_{v_{i+1}}) \in \mathcal{X}_{(v_i, v_{i+1})}$ may apply, where $\mathcal{X}_{(v_i, v_{i+1})}$ is a closed convex set. Unlike in standard graphs with fixed edge weights, the cost of traversing π depends on the optimal trajectory conditioned on it. Specifically, we define the trajectory cost as $c(\tau) = \sum_{i=1}^{|\pi|-1} c(\mathbf{x}_{v_i}, \mathbf{x}_{v_{i+1}})$, where $c(\mathbf{x}_{v_i}, \mathbf{x}_{v_{i+1}})$ is a proper, closed, convex, positive, and bounded-away-from-zero cost function.

Graph optimization problems on GCSs involve both discrete decisions (selecting sets and their visitation order) and continuous decisions (choosing trajectory points within those sets). When formulated as network flows (Marcucci 2024), they yield mixed-integer nonconvex programs, which are generally intractable. Marcucci et al. (2023) introduced the “GCS convex restriction” trick that solves the shortest-path problem on a GCS from a given starting point $\mathbf{x}_s \in \mathcal{X}_s$ to a target point $\mathbf{x}_t \in \mathcal{X}_t$ by fixing a path π —typically simple (i.e., without vertex repetition)—and solving a convex program for the optimal trajectory conditioned on π . Paths are explored iteratively, often guided by a convex relaxation over discrete set-selection variables. Recent GCS planners (Natarajan et al. 2024; Chia et al. 2024; Tang et al. 2025) extend this idea with best-first search over feasible paths, leveraging the fact that trajectory optimization along a fixed path is convex and efficient.

Traveling Salesman Problem on a GCS (GCS-TSP): GCS-TSP generalizes the classical TSP to the GCS setting. Given a GCS $G = (V, E)$, the goal is to find a tour—a closed path $\pi = (v_1, v_2, \dots, v_{|\pi|} = v_1)$ that visits each vertex in V at least once—along with its associated trajectory $\tau = (\mathbf{x}_{v_1}, \mathbf{x}_{v_2}, \dots, \mathbf{x}_{v_{|\pi|}} = \mathbf{x}_{v_1})$ through the convex sets that starts and ends at the same point, such that the trajectory cost is minimized. Formally, we define GCS-TSP as:

$$\begin{aligned} \min_{\pi, \mathbf{x}} \quad & c(\tau) := \sum_{i=1}^{|\pi|-1} c(\mathbf{x}_{v_i}, \mathbf{x}_{v_{i+1}}) & (1) \\ \text{s.t.} \quad & \mathbf{x}_{v_1} = \mathbf{x}_{v_{|\pi|}}, & (2) \\ & \mathbf{x}_{v_i} \in \mathcal{X}_{v_i}, & \forall i = 1, \dots, |\pi|, & (3) \\ & (\mathbf{x}_{v_i}, \mathbf{x}_{v_{i+1}}) \in \mathcal{X}_{(v_i, v_{i+1})}, & \forall i = 1, \dots, |\pi| - 1. & (4) \end{aligned}$$

In this paper, we assume the starting point and set are not fixed, following the classical TSP setting. However, our formulation and results naturally extend to cases with a fixed starting set \mathcal{X}_s and point \mathbf{x}_s .

Lower-Bound Graph and Path Unfolding

In this section, we establish a static lower-bound metric for a given GCS $G = (V, E)$ to guide both levels of search in




Figure 2: LBG and path unfolding. (a) The optimal trajectory (red dashed line) yielding lower-bound cost lb_p for any trajectories (black dashed curves) conditioned on each possible p along $\hat{\pi}$. (b)(c) Trajectories τ and τ' (blue solid curves) conditioned on π and π' unfolded from $\hat{\pi}$, respectively.

our GHOST framework. This metric is also required by the TSP subroutine (see next section), which generates tour candidates in best-first order. The key idea is to define a lower-bound distance metric that underestimates the true trajectory cost between any vertex pair. Summing these lower bounds over an abstract path yields a provable lower bound cost of the trajectories conditioned on it. This enables GHOST to explore abstract tours in order of increasing lower-bound estimates, enabling early pruning while preserving optimality. **Notations:** Our tour search operates over an abstract complete graph $\hat{G} = (V, \hat{E})$ induced by G , where $E = \{(u, v) | u, v \in V, u \neq v\}$. Each tour candidate is an *abstract path* (or tour), i.e., a vertex sequence where consecutive vertices may not be adjacent in G . Any abstract path $\hat{\pi} = (v_1, \dots, v_k)$ can be *unfolded* into (i.e., realized by) a valid GCS path $\pi = \pi_{v_1, v_2} \circ \dots \circ \pi_{v_{k-1}, v_k}$, where each subpath $\pi_{v_i, v_{i+1}} = (v_i = w_1^{(i)}, w_2^{(i)}, \dots, w_{\mathcal{K}_i}^{(i)} = v_{i+1})$, and \circ denotes concatenation with repeated endpoints omitted.

Lower-Bound Graph (LBG) Construction

We construct a lower-bound graph (LBG) to efficiently estimate path costs. Unlike standard graphs where edge costs provide meaningful path estimates, the cost of traversing a GCS edge $(u, v) \in E$ depends on trajectory endpoints $\mathbf{x}_u \in \mathcal{X}_u$ and $\mathbf{x}_v \in \mathcal{X}_v$, which are unconstrained unless the full path is fixed. As a result, edge costs in G yield loose lower bounds. Instead, we define the LBG on each triplet (u, v, w) , representing a vertex $v \in V$ along with its predecessor $u \in V$ and successor $w \in V$, allowing for a tighter lower bound that captures the “volume” of set \mathcal{X}_v .

Given a GCS $G = (V, E)$, the LBG $\mathcal{H} = (\mathcal{P}, \mathcal{F})$ is a directed hypergraph. A triplet $(u, v, w) \in \mathcal{P}$ is a passage formed from every pair of adjacent edges $(u, v), (v, w) \in E$, and a hyperedge $(p, p') \in \mathcal{F}$ connects each $p = (\cdot, u, v)$ to $p' = (u, v, \cdot)$. Each triplet $p = (u, v, w)$ is labeled with a lower-bound cost lb_p , computed as the optimal trajectory cost conditioned on p via a convex program. Since this program relaxes any constraints on $\mathbf{x}_u \in \mathcal{X}_u$ and $\mathbf{x}_w \in \mathcal{X}_w$, lb_p never overestimates the cost portion $c(\mathbf{x}_u, \mathbf{x}_v) + c(\mathbf{x}_v, \mathbf{x}_w)$ for any trajectory conditioned on a path using triplet p . Fig. 2-(a) demonstrates an example of lb_p . The LBG has also been used in Natarajan et al. (2024) to obtain an admissible heuristic for accelerating A*-like search in GCSs.

Properties of LBG: Any path π on G can be equivalently written as (p_1, \dots, p_k) , where each $p_i = (u_i, v_i, w_i)$, and

$\pi = (u_1, v_1 = u_2, w_1 = v_2 = u_3, \dots, w_{k-2} = v_{k-1} = u_k, w_{k-1} = v_k, w_k)$. We define $\mathcal{L}(\pi)$ as the cost accumulating the lower-bound costs of triplets along π :

$$\mathcal{L}(\pi) = \sum_{p \text{ along } \pi} lb_p + \begin{cases} 0, & \text{if } \pi \text{ is open (i.e., } u_1 = w_k) \\ lb_{(v_k, u_1, v_1)}, & \text{otherwise } (\pi \text{ is closed)} \end{cases}$$

Theorem 3 in Natarajan et al. (2024) show that for any trajectory τ from \mathcal{X}_u to $\mathcal{X}_{w'}$, conditioned on an open path π , we have $\mathcal{L}(\pi_{u, w'}^*) \leq \mathcal{L}(\pi) \leq c(\tau)$, where $\pi_{u, w'}^*$ is the *lower-bound-cost-minimal path* from \mathcal{X}_u to $\mathcal{X}_{w'}$. Theorem 1 generalizes this result to any path unfolded from an abstract path.

Theorem 1 (Lower-Bound Path Cost). *For any abstract path $\hat{\pi}$ on \hat{G} and the lower-bound-cost-minimal π^* unfolded from it, $\mathcal{L}(\pi^*) \leq \mathcal{L}(\pi) \leq c(\tau)$ holds for any trajectory τ conditioned on an arbitrary π unfolded from $\hat{\pi}$.*

Proof. The first inequality follows by definition. The second follows from Theorem 3 in Natarajan et al. (2024) for open paths and extends to closed paths by noting that $\mathcal{L}(\pi)$ is the cost of a relaxed trajectory optimization conditioned on $\pi = (p_1, \dots, p_k)$ by ignoring “connection constraints” between consecutive triplets $p_i \rightarrow p_{i+1}$. The same argument applies to a closed path π , where the additional triplet cost term similarly ignores these constraints for the wraparound connections $p_{k-1} \rightarrow p_k = p_1$ and $p_k = p_1 \rightarrow p_2$. \square

Path Unfolding

We present our (abstract-)path-unfolding algorithm, which uses a multi-label A*-like best-first search (Grenouilleau, Van Hove, and Hooker 2019) to enumerate all paths unfolded from a given abstract path $\hat{\pi} = (v_1, v_2, \dots, v_k)$ on \hat{G} in non-decreasing order of their lower-bound costs. Each search state consists of a triplet p and a label ℓ , indicating progress toward the next endpoint $v_{\ell+1}$. To allow for cyclic and next-best paths to a state, each search node n consists of its path $n.\pi$ (instead of a triplet) and label $n.\ell$. The node also stores the current cost $n.g$ and an estimated final cost $n.f = n.g + h_{\hat{\pi}}(p, n.\ell)$, where $h_{\hat{\pi}}(p, n.\ell)$ is an admissible heuristic underestimating the cost to sequentially visit all remaining endpoints $v_{\ell+1}, \dots, v_k$. Fig. 2-(bc) shows an example in which the optimal trajectory conditioned on π can have a smaller cost than π' , even when $\mathcal{L}(\pi) < \mathcal{L}(\pi')$. This motivates exhaustively exploring all unfolded paths for $\hat{\pi}$ to guarantee finding the one that yields the optimal trajectory.

Pseudocode (Alg. 1) The search maintains an OPEN list of nodes n , prioritized by $n.f$ (Line 1). Nodes are initialized for every triplet p that matches the prefix of $\hat{\pi}$, with their labels set according to how many initial vertices of $\hat{\pi}$ are matched (Lines 2-8). At each step, n with the lowest $n.f$ is popped (Line 10). For an open $\hat{\pi}$, if $n.\pi$ ends at v_k and $n.\ell = k$, the search outputs $n.\pi$ as the next-best unfolded path (Line 12). For a closed $\hat{\pi}$, if $n.\pi$ starts and ends at the same triplet (\cdot, v_k, \cdot) and $n.\ell = k$, the search outputs $n.\pi$ as the next-best unfolded path (Line 14), with $n.g$ correctly capturing the wraparound triplet cost. The search then branches on all triplets p consistent with the suffix of $n.\pi$. Fig. 3 shows an example. Each child node n_c increments the label if appending p reaches endpoint $v_{n.\ell+1}$ (Line 18), updates $n_c.g$ by adding lb_p , and sets $n_c.f$ as the sum of g and an admissible heuristic value $h_{\hat{\pi}}$ (Line 21). Optionally, nodes with




Figure 3: Illustration of Alg. 1 for $\hat{\pi} = (v_1, \dots)$ on the GCS from Fig. 2: OPEN evolution (nodes in gray blocks sorted top-bottom in non-decreasing f), node expansions (dashed arrows), and appended triplets (blue).

$n_c.f > \bar{c}$ are pruned (Line 22), supporting efficient hierarchical search in GHOST. We employ a multi-label heuristic inspired by Li et al. (2021); Zhong et al. (2022) as defined in Lemma 2. Fig. 3 illustrates Alg. 1 with a concrete example.

Lemma 2. *The multi-label heuristic value $h_{\hat{\pi}}(p, n.l) = \mathcal{L}(\pi_{w, v_{\ell+1}}^*) + \sum_{i=\ell+1}^{k-1} \mathcal{L}(\pi_{v_i, v_{i+1}}^*)$ never overestimates the actual lower-bound cost-to-go for any path π unfolded from $\hat{\pi} = (v_1, v_2, \dots, v_k)$ via $p = (u, v, w)$ with label $n.l$.*

Proof. Consider any open path $\pi = \pi_{v_1, v_2} \circ \dots \circ \pi_{v_{k-1}, v_k}$ unfolded from $\hat{\pi}$, where each subpath $\pi_{v_i, v_{i+1}} = (v_i = w_1^{(i)}, w_2^{(i)}, \dots, w_{\mathcal{K}_i}^{(i)} = v_{i+1})$. Now, assume the triplet $p = (u, v, w)$ at label $n.l$ is followed by some $p' = (v, w, \cdot)$. The cost-to-go from p can be decomposed as $lb_{p'} + \mathcal{L}(\pi_{w, v_{\ell+1}}) + \sum_{i=\ell+1}^{k-1} \left(lb_{(w_{\mathcal{K}_{i-1}-1, v_i, w_2^{(i+1)}})} + \mathcal{L}(\pi_{v_i, v_{i+1}}) \right)$, which is at least $\mathcal{L}(\pi_{w, v_{\ell+1}}^*) + \sum_{i=\ell+1}^{k-1} \mathcal{L}(\pi_{v_i, v_{i+1}}^*) = h_{\hat{\pi}}(p, n.l)$. For closed paths, the inequality still holds with an additional non-negative term $lb_{(w_{\mathcal{K}_{k-1}-1, v_k, w_2^{(1)}})}$ on the left. In summary, $h_{\hat{\pi}}$ aggregates the minimum possible subpath costs but does not account for the triplet costs for p' and trajectory endpoints in the actual cost-to-go. \square

Restricted-TSP over Abstract Triplets

In this section, we describe the key subroutine used to compute the next best abstract tour on the induced complete graph \hat{G} . The procedure solves the *Restricted-TSP* (RTSP) (Hamacher and Queyranne 1985) on $\hat{G} = (V, \hat{E})$. A feasible RTSP solution is a simple abstract tour $\hat{\pi}$ that visit all vertices. Unlike standard TSP, RTSP specifies two disjoint abstract edge sets: $\hat{E}^+ \subseteq \hat{E}$ (edges that must be included) and $\hat{E}^- \subseteq \hat{E}$ (edges that must be excluded). Standard TSP is a special case with $\hat{E}^+ = \hat{E}^- = \emptyset$.

While the ILP formulation for TSP (Miller, Tucker, and Zemlin 1960) can be adapted to RTSP with a static pairwise distance metric over vertices, this is insufficient for GCS, where edge costs do not provide meaningful estimates or capture the “volume” of tour endpoint sets. Instead, we leverage the triplet-based distance metric established above and reformulate RTSP over all abstract triplets $\hat{p} \in \hat{\mathcal{P}}$. Each $\hat{p} = (u, v, w)$ is formed from every pair of abstract edges $(u, v), (v, w) \in \hat{E}$, with a lower-bound cost defined

Algorithm 1: (Next-)Best-First Path Unfolding

Input: abstract path $\hat{\pi} = (v_1, \dots, v_k)$, pruning cost \bar{c}

- 1 OPEN $\leftarrow \emptyset$ \triangleright prioritizing nodes by their $n.f$
- 2 **foreach** $p = (v_1, u, w) \in \mathcal{P}$ **do**
- 3 create node n with $n.l \leftarrow 1, n.\pi \leftarrow p$, and $n.g \leftarrow lb_p$
- 4 **if** $u = v_2$ **then**
- 5 $n.l \leftarrow 2$
- 6 **if** $w = v_3$ **then**
- 7 $n.l \leftarrow 3$
- 8 $n.f \leftarrow n.g + h_{\hat{\pi}}(p, n.l)$ and OPEN \leftarrow OPEN $\cup \{n\}$
- 9 **while** OPEN $\neq \emptyset$ **do**
- 10 $n \leftarrow$ OPEN.pop()
- 11 **if** $v_1 \neq v_k$ **and** $n.l = k$ **and** $n.\pi$ ends at (\cdot, \cdot, v_k) **then**
- 12 **yield** $n.\pi$ as the next-best unfolded (open) path
- 13 **if** $v_1 = v_k$ **and** $n.l = k$ **and** $n.\pi$ starts and ends at the same (\cdot, v_k, \cdot) **then**
- 14 **yield** $n.\pi$ as the next-best unfolded (closed) path
- 15 **foreach** $p = (u, v, w) \in \mathcal{P}$ where $n.\pi$ ends at (\cdot, u, v) **do**
- 16 $n_c \leftarrow$ new child node of n with $n_c.\pi$ appending p to $n.\pi$
- 17 **if** $w = v_{n.l+1}$ **then**
- 18 $n_c.l \leftarrow n.l + 1$
- 19 **else**
- 20 $n_c.l \leftarrow n.l$
- 21 $n_c.g \leftarrow n.g + lb_p$, and $n_c.f \leftarrow n_c.g + h_{\hat{\pi}}(p, n.l)$
- 22 **if** $n_c.f \leq \bar{c}$ **then**
- 23 $\text{OPEN} \leftarrow \text{OPEN} \cup \{n_c\}$

as $b(\hat{p}) = \frac{1}{2} \mathcal{L}(\pi_{u,v}^*) + b_{\text{mid}}(\hat{p}) + \frac{1}{2} \mathcal{L}(\pi_{v,w}^*)$, where

$$b_{\text{mid}}(\hat{p}) = \begin{cases} lb_{(u,v,w)}, & \text{if } (u,v), (v,w) \in E \\ \min_{p=(u,v,\cdot) \in \mathcal{P}} lb_p, & \text{if } (u,v) \in E, (v,w) \notin E \\ \min_{p=(v,w,\cdot) \in \mathcal{P}} lb_p, & \text{if } (u,v) \notin E, (v,w) \in E \\ \min_{p=(\cdot,v,\cdot) \in \mathcal{P}} lb_p, & \text{if } (u,v), (v,w) \notin E \end{cases}$$

The lower-bound path cost $\mathcal{L}(\pi_{u,v}^*)$ is typically zero when $(u, v) \in E$ (and thus $\mathcal{X}_u, \mathcal{X}_v$ are adjacent sets), unless with complex vertex or edge constraints. For notational simplicity, we use the same symbol \mathcal{L} to denote the cumulative cost of the abstract triplets along an abstract path $\hat{\pi}$, that is, we define $\mathcal{L}(\hat{\pi}) = \sum_{\hat{p} \text{ along } \hat{\pi}} b(\hat{p})$ for any abstract path $\hat{\pi}$. Theorem 3 establishes $\mathcal{L}(\hat{\pi})$ as a valid lower-bound tour cost.

Theorem 3 (Lower-Bound Tour Cost). *For an abstract tour $\hat{\pi}$ on \hat{G} , $\mathcal{L}(\hat{\pi}) \leq \mathcal{L}(\pi) \leq c(\tau)$ holds for any closed path π unfolded from $\hat{\pi}$ and any trajectory τ conditioned on π .*

Proof. The second inequality follows Theorem 1. Consider $\hat{\pi} = (v_1, \dots, v_k)$ and an unfolded $\pi = \pi_{v_1, v_2} \circ \dots \circ \pi_{v_{k-1}, v_k}$, where each subpath $\pi_{v_i, v_{i+1}} = (v_i = w_1^{(i)}, w_2^{(i)}, \dots, w_{\mathcal{K}_i}^{(i)} = v_{i+1})$. We decompose $\mathcal{L}(\pi) = \sum_{i=1}^k \mathcal{L}(\pi_{v_i, v_{i+1}}) + \sum_{i=1}^k lb_{(w_{\mathcal{K}_{i-1}-1, v_i, w_2^{(i)}})}$, where $v_{k+1} = v_1$ (wraparound) and $w_{\mathcal{K}_0}^{(0)} = w_{\mathcal{K}_{k-1}-1}^{(k-1)}$ for notation convenience. As each $\mathcal{L}(\pi_{v_i, v_{i+1}}) \geq \mathcal{L}(\pi_{v_i, v_{i+1}}^*)$ by definition and $lb_{(w_{\mathcal{K}_{i-1}-1, v_i, w_2^{(i)}})} \geq b_{\text{mid}}(v_{i-1}, v_i, v_{i+1})$ by construction, summing over all $i = 1, \dots, k$ concludes the proof. \square

Unlike classical TSP ILP, which uses variables for edges, our RTSP ILP introduces binary variables $y_{\hat{p}}$ for triplets $\hat{p} \in \hat{\mathcal{P}}$, minimizing the lower-bound tour cost:

$$\min_{y_{\hat{p}}} \sum_{\hat{p} \in \hat{\mathcal{P}}} b(\hat{p}) y_{\hat{p}} \quad (5)$$

$$\begin{aligned} \text{s.t.} \quad & \sum_{\hat{p}=(\cdot, v, \cdot) \in \hat{\mathcal{P}}} y_{\hat{p}} = 1, & \forall v \in V & \quad (6) \\ & \sum_{z \neq v} y_{zuv} = \sum_{w \neq u} y_{uvw}, & \forall (u, v) \in \hat{E} & \quad (7) \\ & \sum_{\hat{p} \in \hat{\mathcal{P}}^{\text{ind}}} y_{\hat{p}} = 2, & \forall e = (u, v) \in \hat{E}^+ & \quad (8) \\ & y_{\hat{p}} = 0, & \forall \hat{p} \in \hat{\mathcal{P}}_e^{\text{ind}}, e \in \hat{E}^- & \quad (9) \end{aligned}$$

where $\hat{\mathcal{P}}_{e=(u,v)}^{\text{ind}} = \{\hat{p} \in \hat{\mathcal{P}} \mid \hat{p} = (\cdot, u, v) \text{ or } \hat{p} = (u, v, \cdot)\}$ denotes all triplets induced by (u, v) . Eqn. (6) ensures that each $v \in V$ is visited exactly once. Eqn. (7) enforces flow conservation. Eqn. (8-9) enforce inclusion and exclusion constraints. Whenever a new subtour C is found in the solution, we add a subtour elimination constraint similar to the DFJ constraint (Dantzig, Fulkerson, and Johnson 1954):

$$\sum_{\hat{p} \in C} y_{\hat{p}} \leq |C| - 1 \quad (10)$$

and re-solve until only a single tour remains.

The GHOST Framework

In this section, we present the GHOST framework for solving GCS-TSP. GHOST leverages the LBG \mathcal{H} and path unfolding to compute lower-bound costs that guide and prune the search at both levels. At the high level, GHOST explores abstract tours $\hat{\pi}$ on the induced complete graph \hat{G} in non-decreasing order of their lower-bound costs. At the low level, for each abstract tour, it explores unfolded paths π from $\hat{\pi}$, also in non-decreasing order of lower-bound cost, and computes the optimal trajectory for each π via convex optimization. The resulting trajectory cost is used to prune any abstract tours or paths whose lower-bound costs are no smaller, guaranteeing optimality as formally shown later.

High-Level Tour Search

At the high level, GHOST generalizes the *Lawler-Murty* procedure (Murty 1968; Lawler 1972), originally developed to find the K -best solutions to TSP and related combinatorial optimization problems (Hamacher and Queyranne 1985; Ren, Rathinam, and Choset 2023). The procedure can be conceptualized as a tree search where each node corresponds to an RTSP instance. Fig. 4 illustrates the high-level search. **Pseudocode** (Alg. 2): Given a GCS G , each search node n corresponds to solving an RTSP instance on the induced complete graph \hat{G} with specified inclusion \hat{E}^+ and exclusion \hat{E}^- edge sets, and stores the resulting optimal abstract tour $n.\hat{\pi}$. Function `EvalNode`(n) unfolds $n.\hat{\pi}$ and evaluates the trajectory cost. GHOST performs a best-first search with an OPEN list prioritized by the lower-bound tour cost $n.\underline{c}$ (Line 4). The root node r is initialized with the unrestricted TSP, n^* tracks the best node found so far (Lines 1-3). In each iteration, the node n with the smallest $n.\underline{c}$ is popped (Line 6). If $n.\underline{c}$ is no smaller than the cost of the current best trajectory $c(n^*.\tau)$ (Line 7), GHOST returns $n^*.\tau$ as optimal (Line 18). Otherwise, GHOST calls `EvalNode` to unfold $n.\hat{\pi}$ and compute $c(n.\tau)$ (Line 9). If $c(n.\tau)$ improves upon $c(n^*.\tau)$, n^* is updated (Line 10). For each edge e_i in $n.\hat{\pi}$, GHOST creates a child node n_c , (Lines 11-12) by adding




Figure 4: High level of GHOST. Each T_i denotes a low-level search tree that explores unfolded paths for abstract tour $\hat{\pi}_i$. The search forest is pruned at T_j where $\mathcal{L}(\hat{\pi}_j)$ exceeds the current best trajectory cost.

Algorithm 2: GHOST for GCS-TSP

Input: GCS $G = (V, E)$

- 1 $r \leftarrow$ a root node with $r.\hat{E}^+ = r.\hat{E}^- = \emptyset$ and empty $r.\tau$
- 2 $r.\hat{\pi} \leftarrow$ optimal TSP tour and $r.\underline{c} \leftarrow \mathcal{L}(r.\hat{\pi})$
- 3 initialize the best n^* as a dummy node with $n^*.\tau = \infty$
- 4 OPEN $\leftarrow \{r\}$ ▷ prioritizing nodes by their $n.\underline{c}$
- 5 **while** OPEN $\neq \emptyset$ **do**
- 6 $n \leftarrow$ OPEN.pop()
- 7 **if** $n.\underline{c} \geq c(n^*.\tau)$ **then**
- 8 | **break**
- 9 EvalNode($n, c(n^*.\tau)$)
- 10 update n^* to n if $c(n.\tau) < c(n^*.\tau)$
- 11 **foreach** edge e_i in $n.\hat{\pi} = (1, \dots, e_{|n.\hat{\pi}|})$ **do**
- 12 $n_c \leftarrow$ new child node of n
- 13 $n_c.\hat{E}^- \leftarrow \{e_i\} \cup n.\hat{E}^-$
- 14 $n_c.\hat{E}^+ \leftarrow \{e_j\}_{j=1}^{i-1} \cup n.\hat{E}^+$
- 15 **if** a feasible RTSP tour exists for $(n_c.\hat{E}^+, n_c.\hat{E}^-)$ **then**
- 16 | $n_c.\hat{\pi} \leftarrow$ optimal RTSP tour for $(n_c.\hat{E}^+, n_c.\hat{E}^-)$
- 17 | $n_c.\underline{c} \leftarrow \mathcal{L}(n_c.\hat{\pi})$ and OPEN \leftarrow OPEN $\cup \{n_c\}$
- 18 **return** optimal GCS-TSP trajectory $n^*.\tau$
- 19 **Function** EvalNode(n, \bar{c}):
- 20 $c(n.\tau) \leftarrow \infty$
- 21 **foreach** $\pi \leftarrow$ next-best path unfolded from $\hat{\pi}$ (by calling
- 22 | Alg. 1 with pruning cost \bar{c}) **do**
- 23 | $\tau \leftarrow$ optimal trajectory conditioned on π
- 24 | **if** $\mathcal{L}(\pi) \geq c(n.\tau)$ **then**
- 25 | **break**
- 25 | update $n.\tau$ to τ if $c(\tau) < c(n.\tau)$

new constraints to its parent: It excludes e_i (Line 13) and includes all edges with indices less than i (Line 14). If the resulting RTSP instance is feasible, n_c is inserted into OPEN with its tour and lower bound (Lines 15-17).

Tour Unfolding and Trajectory Optimization

The low-level `EvalNode` function calls our path-unfolding algorithm (Alg. 1) to generate path π unfolded from $\hat{\pi}$, considering each in non-decreasing order of $\mathcal{L}(\pi)$ (Line 21) and pruning once $\mathcal{L}(\pi)$ exceeds the best trajectory cost $c(n.\tau)$ among all previously considered paths for $\hat{\pi}$ (Lines 23-24). For each π , the function computes the optimal trajectory τ conditioned on π via convex optimization (Line 22) and updates $n.\tau$ if a lower cost is found (Line 25).

The convex optimization follows “GCS convex restriction” of Marcucci et al. (2023), but with a key extension: Our formulation allows π to revisit vertices and edges multiple times (which can occur when G is incomplete). We handle this by introducing separate variables for each occurrence of

a vertex and defining edge traversal constraints according to the specific sequence in π . To improve efficiency, we cache the optimal trajectory for each π since the same π may be unfolded from different abstract tours.

Bounded-Suboptimal ϵ -GHOST

We introduce a bounded-suboptimal variant, ϵ -GHOST, which trades optimality for computational efficiency. Given a parameter $\epsilon \in [0, 1)$, we modify GHOST in two ways: On Line 7, the termination condition is relaxed from $c(n^*. \tau) - n.\underline{c} < 0$ to $c(n^*. \tau) - n.\underline{c} \leq \epsilon \cdot c(n^*. \tau)$. In `EvalNode`, Alg.1 is called with a pruning cost $(1 - \epsilon) \cdot c(n^*. \tau)$ instead of $c(n^*. \tau)$ (Line 21), and the pruning condition is updated to $\mathcal{L}(\pi) \geq (1 - \epsilon) \cdot c(n^*. \tau)$ instead of $\mathcal{L}(\pi) \geq c(n^*. \tau)$ (Line 23). Setting $\epsilon = 0$ recovers the original, optimal GHOST.

Properties of (ϵ -)GHOST

As in Van der Poort et al. (1999); Ren, Rathinam, and Choset (2023), the Lawler-Murty procedure at the high level of GHOST guarantees to explore TSP tours in non-decreasing order of their lower-bound costs. Theorem 4 and Corollary 5 establish the (bounded-sub)optimality of (ϵ -)GHOST.

Theorem 4. *ϵ -GHOST produces a GCS-TSP solution with cost at most $\frac{1}{1-\epsilon}$ times the optimal cost.*

Proof. Consider any node n_i popped at iteration i , where `EvalNode` is called with pruning cost $(1 - \epsilon) \cdot c(n_i^*. \tau)$, and n_i^* stores the best node over all previous iterations. Let $\bar{\tau}_i$ denote any pruned trajectory conditioned on some path π unfolded from $n_i.\hat{\pi}$. Such a trajectory is pruned either (1) in Alg. 1 if $\mathcal{L}(\pi) \geq (1 - \epsilon) \cdot c(n_i^*. \tau)$, or (2) in `EvalNode` if $\mathcal{L}(\pi) \geq (1 - \epsilon) \cdot c(n_i^*. \tau)$. Since unfolded paths are considered in non-decreasing order of $\mathcal{L}(\pi)$, in either case we have

$c(\bar{\tau}_i) \stackrel{\text{Theorem 1}}{\geq} \mathcal{L}(\pi) \geq (1 - \epsilon) \cdot c(n^*. \tau)$, where $n^*. \tau$ is the best trajectory across all iterations.

When ϵ -GHOST terminates and returns $n^*. \tau$, the termination condition ensures $n.\underline{c} \geq (1 - \epsilon) \cdot c(n^*. \tau)$. Since the Lawler-Murty procedure explores tours in non-decreasing order of their lower-bound costs, any pruned trajectory $\hat{\tau}$ conditioned on some path unfolded from an unexplored tour

$\hat{\pi}$ satisfies $c(\hat{\tau}) \stackrel{\text{Theorem 3}}{\geq} \mathcal{L}(\hat{\pi}) \geq n.\underline{c} \geq (1 - \epsilon) \cdot c(n^*. \tau)$.

Therefore, $c(n^*. \tau) \leq \frac{1}{1-\epsilon} \cdot c(\hat{\tau})$ for any pruned trajectory $\hat{\tau}$, which is at most $\frac{1}{1-\epsilon}$ times the optimal cost. \square

Corollary 5. *GHOST returns an optimal GCS-TSP solution.*

Proof. This follows from Theorem 4 with $\epsilon = 0.0$. \square

Implementation Details and GHOST Variants

GHOST inherently considers all paths, including cyclic ones, as required by rare, complex constraints (e.g., Nataraajan et al. (2024)). Since such cases are uncommon, many GCS shortest-path planners treat cyclic paths as optional and typically restrict to simple paths. Our implementation similarly restricts path unfolding to simple subpaths between tour endpoints, which limits (bounded-sub)optimality to trajectories in this class (each vertex visited at most $|V|$ times), but yields improved efficiency.

GHOST sequentially explores tours, fully evaluating all the unfolded paths for one tour before considering the next.

While it is possible to compare paths from different tours in parallel, we found it offers little practical benefit, as pruning often eliminates most paths from each tour—especially when ϵ is used.

Finally, extending the current triplet-based LBG and the RTSP ILP to higher-order tuples (e.g., 4-tuples) could yield tighter bounds. However, we observed that this incurs substantial computational overhead with little improvement.

Empirical Evaluation

This section presents our experimental results on an *Apple*[®] M4 CPU machine with 16GB RAM. We use *Drake* (Tedrake and the Drake Development Team 2019) with *Gurobi* solver for solving the GCS convex restriction and the RTSP ILP.

Experiment Setup

To evaluate our **GHOST** framework with LBG (Alg. 2) for GCS-TSP, we compare against the following baselines:

ECG (non-optimal) replaces the LBG in GHOST with a complete edge-cost graph (ECG), assigning each abstract edge (u, v) the Euclidean distance between the Chebyshev centers (Boyd and Vandenberghe 2004) of \mathcal{X}_u and \mathcal{X}_v . ECG solves the original RTSP ILP over abstract edges (Van der Poort et al. 1999) and unfolds tours using edge costs. It is non-optimal as these costs do not yield valid lower bounds.

Greedy (non-optimal) uses greedy search on G with its LBG. Given an initial vertex, a depth-first search chooses the next unvisited vertex v from the current vertex u such that $p = (\cdot, u, v) \in \mathcal{P}$ has the smallest lb_p . After visiting all vertices, an abstract tour is unfolded via Alg. 1, and the search backtracks to explore other tours. The best trajectory found is recorded throughout the greedy search.

MICP solves GCS-TSP via an unified optimization (Marcucci 2024), assuming $G = (V, E)$ is complete. For incomplete G , it is replaced with $|V|$ copies, denoted as G^i , $i = 1, 2, \dots, |V|$. Each G^i represents GCS G at a different step i and is connected to G^{i+1} as described in Marcucci (2024).

We generate GCS-TSP instances on a 5×5 grid with 12 random seeds (see Fig. 5-(a)), where the convex sets are centered at some intersections of the grid. The following three specific types of GCS-TSP problems are used for evaluation.

Point-GCS: We generate complete GCSs of $N \in [5, 25]$ vertices (convex sets). Each convex set $\mathcal{X}_v \subseteq \mathbb{R}^2$ is a 2d point and each edge $e = (u, v)$ is unconstrained (i.e., $\mathcal{X}_e = \mathcal{X}_u \times \mathcal{X}_v$). The trajectory cost $c(\tau)$ is defined as the sum of the Euclidean distances between consecutive points along τ .

Linear-GCS: We generate incomplete GCS of $M \in [10, 50]$ edges. Each convex set $\mathbf{x}_v \in \mathcal{X}_v \subseteq \mathbb{R}^{2+2}$ comprises of two points $\mathbf{a}_v, \mathbf{b}_v \in \mathbb{R}^2$. Each edge (u, v) indicates $\mathcal{X}_u \cap \mathcal{X}_v \neq \emptyset$ with edge constraint $\mathcal{X}_e := \{(\mathbf{x}_u, \mathbf{x}_v) | \mathbf{b}_v = \mathbf{a}_u\} \subseteq \mathcal{X}_u \times \mathcal{X}_v$. The trajectory cost $c(\tau)$ is defined as the sum of the Euclidean distance between \mathbf{a}_v and \mathbf{b}_v for each v along τ .

Bézier-GCS: We generate incomplete GCS of $M \in [10, 50]$ edges. Each $\mathbf{x}_v \in \mathcal{X}_v \subseteq \mathbb{R}^{3 \times 5}$ consists of 5 control points for two 4-order Bézier curves: one in 2D space and the other in 1D time that reconstruct a continuous trajectory piece bounded by \mathcal{X}_v . An edge (u, v) connects \mathcal{X}_u and \mathcal{X}_v if




Figure 5: A collection of GCS-TSP problem solutions produced by GHOST. (a) GCS-TSP with 23 convex sets (colored polygons) and 50 edges on Point-GCS, Linear-GCS, and Bézier-GCS. (b) Ground vehicle inspection planning for target location sets (squares) in a 2D maze. (c) Quadrotor coverage planning in a 3D obstacle-rich (gray cuboids) bounded region. (d) Task and motion planning for a 7-DoF manipulator, where black and red spheres indicate the trajectory and task poses, respectively.

$\mathcal{X}_u \cap \mathcal{X}_v \neq \emptyset$, with \mathcal{X}_e defined as trajectory continuity constraints (see Marcucci et al. (2023)). The cost $c(\tau)$ is defined as the sum of the durations of each 1D time Bézier curve.

Comparison, Ablation Study, and Demonstration

Fig. 6 summarizes the performance comparisons.

Metrics: We compare all methods in terms of solution trajectory cost $c(\tau)$ and runtime. Each method is given a 100-second limit per random GCS-TSP instance, and metrics for each instance are averaged over 12 random seeds. We also report the optimality gap $\rho = (c(\tau^*) - lb^*)/c(\tau^*)$ for GHOST and MICP, where lb^* is the greatest provable lower bound at termination. For GHOST, $lb^* = \mathcal{L}(\hat{\pi}_i)$ is the cost of the next-best abstract tour when $c(\tau^*)$ is returned. Note that $\mathcal{L}(\hat{\pi}_i) > c(\tau^*)$ is possible for an optimal $c(\tau^*)$, so negative ρ values are capped at zero. For MICP, $lb^* \leq c(\tau^*)$ is the current-best fractional solution cost in branch-and-bound.

Result Analysis: GHOST consistently outperforms ECG and Greedy in both trajectory cost and runtime. Even without LBG, ECG is a strong baseline, as it uses the edge cost estimates to guide the GHOST search, highlighting the flexibility and generality of our GHOST framework. The comparison with Greedy, which replaces systematic search with a heuristic approach, further demonstrates the advantage of GHOST: Systematic exploration yields better solutions and more reliable performance. A key strength of GHOST is the integration of LBG, which provides a provable lower-bound cost for the solution trajectory. This feature is particularly valuable in practice, as it allows the optimality gap to be monitored throughout the search, enabling users to gauge solution quality even if the search is terminated before optimality is achieved. In contrast, ECG and Greedy lack such guarantees (unless all tours are explicitly enumerated), leading to longer runtimes and no certificate of optimality—especially noticeable in the Point-GCS cases. GHOST also shows clear scalability benefits over MICP, which fails to produce feasible solutions for Point-GCS with $N > 13$, Linear-GCS with $M > 19$, or any Bézier-GCS instances. The bounded-suboptimal 0.5-GHOST variant further highlights the flexibility of our approach, often obtaining solutions of the same quality as those of GHOST, but with improved efficiency—particularly when runtime is limited.

Applications: We demonstrate the versatility of GHOST on Bézier-GCS-TSP instances for coverage, inspection, and

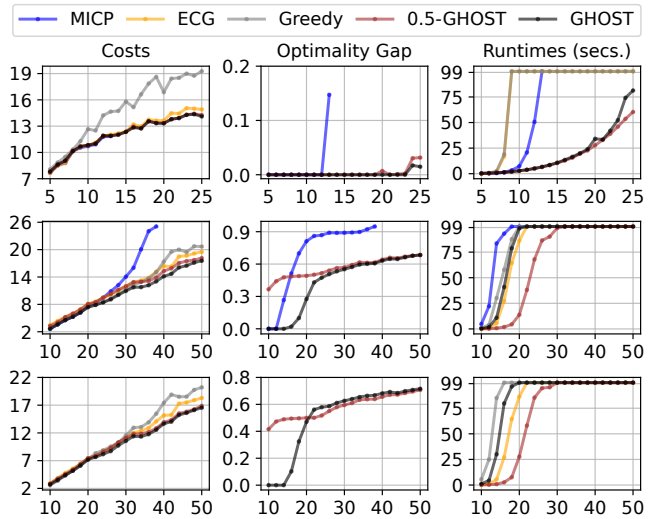


Figure 6: Comparison and ablation study on Point-GCS, Linear-GCS, and Bézier-GCS (from top to bottom row), where their x-axes indicate the number of vertices, number of edges, and number of edges, respectively.

task and motion planning. Fig. 1 and Fig. 5-(c) show continuous trajectories covering precomputed collision-free convex regions. Fig. 5-(b) shows inspection of a subset of the convex regions, achieved by modifying our RTSP ILP to require only subset coverage. Finally, Fig. 5-(d) demonstrates a 7-DoF manipulator (KUKA LBR iiwa) reaching eight convex sets, each corresponding to an independent task pose.

Conclusions and Future Work

We presented GHOST, an optimal hierarchical framework for GCS-TSP that combines combinatorial tour search with convex trajectory optimization. GHOST systematically explores tours on GCS while computing admissible lower bounds that guide best-first search at both the tour level and the path unfolding level for feasible trajectories realizing each tour, enabling effective pruning while maintaining optimality guarantees. Future work includes extending GHOST to general graph optimization problems in GCS, improving its scalability, and adapting it for multi-agent systems.

Acknowledgements

This work was supported by the NSERC under grant number RGPIN2020-06540 and a CFI JELF award.

References

- Acar, E. U.; Choset, H.; Rizzi, A. A.; Atkar, P. N.; and Hull, D. 2002. Morse decompositions for coverage tasks. *The International Journal of Robotics Research*, 21(4): 331–344.
- Bhat, A.; Gutow, G.; Vundurthy, B.; Ren, Z.; Rathinam, S.; and Choset, H. 2025. A Complete and Bounded-Suboptimal Algorithm for a Moving Target Traveling Salesman Problem with Obstacles in 3D. In *ICRA*, 6132–6138.
- Boyd, S. P.; and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Chia, S. Y. C.; Jiang, R. H.; Graesdal, B. P.; Kaelbling, L. P.; and Tedrake, R. 2024. GCS*: Forward heuristic search on implicit graphs of convex sets. *arXiv preprint arXiv:2407.08848*.
- Choset, H.; and Pignon, P. 1998. Coverage path planning: The boustrophedon cellular decomposition. In *Field and Service Robotics*, 203–209.
- Cohn, T.; Petersen, M.; Simchowitz, M.; and Tedrake, R. 2025. Non-Euclidean motion planning with graphs of geodesically convex sets. *The International Journal of Robotics Research*, 44(10-11): 1840–1862.
- Cornuéjols, G.; Fonlupt, J.; and Naddef, D. 1985. The traveling salesman problem on a graph and some related integer polyhedra. *Mathematical Programming*, 33(1): 1–27.
- Dai, H.; Amice, A.; Werner, P.; Zhang, A.; and Tedrake, R. 2024. Certified polyhedral decompositions of collision-free configuration space. *The International Journal of Robotics Research*, 43(9): 1322–1341.
- Dantzig, G.; Fulkerson, R.; and Johnson, S. 1954. Solution of a large-scale traveling-salesman problem. *Journal of the Operations Research Society of America*, 2(4): 393–410.
- Fu, M.; Kuntz, A.; Salzman, O.; and Alterovitz, R. 2019. Toward Asymptotically-Optimal Inspection Planning via Efficient Near-Optimal Graph Search. *RSS*, 2019: 10–15607.
- Gabriely, Y.; and Rimon, E. 2001. Spanning-tree based coverage of continuous areas by a mobile robot. *Annals of Mathematics and Artificial Intelligence*, 31: 77–98.
- Galceran, E.; and Carreras, M. 2013. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12): 1258–1276.
- Grenouilleau, F.; Van Hoes, W.-J.; and Hooker, J. N. 2019. A multi-label A* algorithm for multi-agent pathfinding. In *ICAPS*, 181–185.
- Hamacher, H. W.; and Queyranne, M. 1985. K best solutions to combinatorial optimization problems. *Annals of Operations Research*, 4(1): 123–143.
- Lawler, E. L. 1972. A procedure for computing the k best solutions to discrete optimization problems and its application to the shortest path problem. *Management Science*, 18(7): 401–405.
- Li, J.; Tinka, A.; Kiesel, S.; Durham, J. W.; Kumar, T. S.; and Koenig, S. 2021. Lifelong multi-agent path finding in large-scale warehouses. In *AAAI*, 11272–11281.
- Marcucci, T. 2024. *Graphs of convex sets with applications to optimal control and motion planning*. Ph.D. thesis, Massachusetts Institute of Technology.
- Marcucci, T.; Nobel, P.; Tedrake, R.; and Boyd, S. 2024a. Fast path planning through large collections of safe boxes. *IEEE Transactions on Robotics*.
- Marcucci, T.; Petersen, M.; von Wrangel, D.; and Tedrake, R. 2023. Motion planning around obstacles with convex optimization. *Science Robotics*, 8(84): eadf7843.
- Marcucci, T.; Umenberger, J.; Parrilo, P.; and Tedrake, R. 2024b. Shortest paths in graphs of convex sets. *SIAM Journal on Optimization*, 34(1): 507–532.
- Miller, C. E.; Tucker, A. W.; and Zemlin, R. A. 1960. Integer programming formulation of traveling salesman problems. *Journal of the ACM*, 7(4): 326–329.
- Murty, K. G. 1968. An algorithm for ranking all the assignments in order of increasing cost. *Operations Research*, 16(3): 682–687.
- Natarajan, R.; Liu, C.; Choset, H.; and Likhachev, M. 2024. Implicit Graph Search for Planning on Graphs of Convex Sets. *RSS*.
- Ntafos, S. 1992. Watchman routes under limited visibility. *Computational Geometry*, 1(3): 149–170.
- Philip, A. G.; Ren, Z.; Rathinam, S.; and Choset, H. 2024. A mixed-integer conic program for the moving-target traveling salesman problem based on a graph of convex sets. In *IROS*, 8847–8853.
- Ratliff, N.; Zucker, M.; Bagnell, J. A.; and Srinivasa, S. 2009. CHOMP: Gradient optimization techniques for efficient motion planning. In *ICRA*, 489–494.
- Ren, Z.; Rathinam, S.; and Choset, H. 2023. CBSS: A new approach for multiagent combinatorial path finding. *IEEE Transactions on Robotics*, 39(4): 2669–2683.
- Schulman, J.; Duan, Y.; Ho, J.; Lee, A.; Awwal, I.; Bradlow, H.; Pan, J.; Patil, S.; Goldberg, K.; and Abbeel, P. 2014. Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 33(9): 1251–1270.
- Skyler, S.; Atzmon, D.; Yaffe, T.; and Felner, A. 2022. Solving the Watchman Route Problem with Heuristic Search. *Journal of Artificial Intelligence Research*, 75: 747–791.
- Tang, J.; and Ma, H. 2024. Multi-robot connected fermat spiral coverage. In *ICAPS*, 579–587.
- Tang, J.; Mao, Z.; and Ma, H. 2025. Large-Scale Multi-Robot Coverage Path Planning on Grids with Path Deconfliction. *IEEE Transactions on Robotics*.
- Tang, J.; Mao, Z.; Yang, L.; and Ma, H. 2025. Space-Time Graphs of Convex Sets for Multi-Robot Motion Planning. In *IROS*.
- Tedrake, R.; and the Drake Development Team. 2019. Drake: Model-based design and verification for robotics.
- Van der Poort, E. S.; Libura, M.; Sierksma, G.; and van der Veen, J. A. 1999. Solving the k-best traveling salesman problem. *Computers & Operations Research*, 26(4): 409–425.
- Werner, P.; Amice, A.; Marcucci, T.; Rus, D.; and Tedrake, R. 2024. Approximating robot configuration spaces with few convex sets using clique covers of visibility graphs. In *ICRA*, 10359–10365.
- Wu, C.; Dai, C.; Gong, X.; Liu, Y.-J.; Wang, J.; Gu, X. D.; and Wang, C. C. 2019. Energy-efficient coverage path planning for general terrain surfaces. *IEEE Robotics and Automation Letters*, 4(3): 2584–2591.
- Zhong, X.; Li, J.; Koenig, S.; and Ma, H. 2022. Optimal and bounded-suboptimal multi-goal task assignment and path finding. In *ICRA*, 10731–10737.