

Guidance Graph Optimization for Lifelong Multi-Agent Path Finding

Yulun Zhang¹, He Jiang¹, Varun Bhatt², Stefanos Nikolaidis² and Jiaoyang Li¹

¹Robotics Institute, Carnegie Mellon University

²Thomas Lord Department of Computer Science, University of Southern California

{yulunzhang,hejiangrivers}@cmu.edu, {vsbhatt,nikolaid}@usc.edu, jiaoyangli@cmu.edu

Abstract

We study how to use guidance to improve the throughput of lifelong Multi-Agent Path Finding (MAPF). Previous studies have demonstrated that, while incorporating guidance, such as highways, can accelerate MAPF algorithms, this often results in a trade-off with solution quality. In addition, how to generate good guidance automatically remains largely unexplored, with current methods falling short of surpassing manually designed ones. In this work, we introduce the guidance graph as a versatile representation of guidance for lifelong MAPF, framing Guidance Graph Optimization as the task of optimizing its edge weights. We present two GGO algorithms to automatically generate guidance for arbitrary lifelong MAPF algorithms and maps. The first method directly optimizes edge weights, while the second method optimizes an update model capable of generating edge weights. Empirically, we show that (1) our guidance graphs improve the throughput of three representative lifelong MAPF algorithms in eight benchmark maps, and (2) our update model can generate guidance graphs for as large as 93×91 maps and as many as 3,000 agents. We include the source code at: https://github.com/lunjohnzhang/ggo_public. All optimized guidance graphs are available online at: <https://yulunzhang.net/publication/zhang2024ggo>.

1 Introduction

We study the problem of leveraging a guidance graph with optimized edge weights to guide agent movement, thereby improving the throughput of lifelong Multi-Agent Path Finding (MAPF). MAPF [Stern *et al.*, 2019] aims to plan collision-free paths for a set of agents from their start to goal locations on a given map, depicted as a graph G . Lifelong MAPF [Li *et al.*, 2021] extends MAPF by assigning new goals to agents as soon as they reach their current ones. Example applications include character control in video games [Ma *et al.*, 2017b; Jansen and Sturtevant, 2008] and automated warehouses in which hundreds of robots are continually assigned new tasks to transport inventory pods [Varambally *et al.*, 2022]. Driven




Figure 1: Comparison of no guidance, human-designed crisscross guidance, and our guidance with a simulation of 240 agents in a 33×36 warehouse map, shown in Figure 4b. The heatmaps show the tile-usage (the frequency that each tile is occupied). Our guidance results in the most balanced traffic with the highest throughput.

by these real-world demands, numerous studies have focused on improving the throughput, namely the average number of reached goals per timestep, by developing better lifelong MAPF algorithms [Ma *et al.*, 2017a; Li *et al.*, 2020; Kou *et al.*, 2020; Damani *et al.*, 2021] or optimizing map layouts [Zhang *et al.*, 2023a; Zhang *et al.*, 2023b].

Given that lifelong MAPF requires online computation of new paths as agents are continuously assigned to new goal locations, lifelong MAPF algorithms always decompose the problem into a series of (one-shot) MAPF instances and solve them sequentially. However, such methods are myopic because each MAPF instance involves only the current goal locations. Achieving (near-)optimal solutions for individual MAPF instances does not necessarily result in the best throughput. In this work, we propose to foster implicit cooperation among agents over the long term by introducing global guidance for agent movement. Our guidance takes the form of a directed weighted graph that alters the costs of agents moving along each edge and waiting at each vertex of graph G . Intuitively, such a guidance graph serves two purposes. First, by amplifying the cost difference of traveling through an edge in opposite directions, we encourage agents to move in the same direction, reducing the number of *head-on* collisions, which happen when two agents try to traverse through the same edge in opposite directions. Second, by increasing the cost of moving in areas prone to congestion, we motivate agents to navigate through less congested areas, ultimately reducing traffic congestion. Figure 1 illustrates the traffic resulting from different guidance strategies.

One closely related work to our guidance graph is highways [Cohen *et al.*, 2015; Li and Sun, 2023], which are a subset of edges selected from graph G with assigned directions and a lower traversal cost. This strategy incentivizes agents to move along the highways, reducing the number of collisions to be resolved by MAPF algorithms. However, the question of how to select these edges and determine their directions and costs remains largely unexplored. The *crisscross* approach [Cohen, 2020], where edge directions alternate in even and odd rows and columns, is common but not optimal. Li and Sun [2023] show that, while crisscross highways speed up lifelong MAPF algorithms, they do not always improve throughput. Additionally, Cohen *et al.* [2016] proposed two methods to select edges and directions for highways, but neither of them outperforms crisscross highways.

Therefore, we introduce the Guidance Graph Optimization (GGO) problem to improve throughput by optimizing the edge weights of a guidance graph. We present two automatic GGO methods. The first applies Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) [Hansen, 2016], a state-of-the-art black-box optimizer, to solve GGO, but its solution is map-specific. The second method, Parameterized Iterative Update (PIU), uses CMA-ES to optimize an update model. This update model, represented by a neural network, starts with an unweighted guidance graph and iteratively updates it with traffic information obtained from a lifelong MAPF simulator. It is capable of optimizing guidance graphs for different maps with similar layouts.

We make the following contributions: (1) introducing the guidance graph, a versatile representation of guidance for lifelong MAPF, and guidance graph optimization (GGO) to improve its throughput, (2) conducting an in-depth study of various existing guidance works in MAPF, and (3) proposing two automatic GGO methods, CMA-ES and PIU, showcasing their superior performance over unweighted graphs and previous guidance methods, along with the transferability of PIU to larger maps with similar layouts.

2 Problem Definition and Preliminaries

2.1 Lifelong Multi-Agent Path Finding

Definition 1 (MAPF). *The (one-shot) MAPF problem takes as inputs a graph $G(V, E)$ and k agents with their start and goal locations. At each timestep, an agent can move to an adjacent vertex or stay at its current vertex. Two agents collide when they arrive at the same vertex or swap locations at the same timestep. The MAPF problem searches for collision-free paths that move all agents from their start to goal locations with minimum sum-of-cost, defined as the total number of move and wait actions that the agents need to take.*

Definition 2 (Lifelong MAPF). *Lifelong MAPF extends one-shot MAPF by constantly assigning new goals to agents when they reach their current ones. Lifelong MAPF searches for collision-free paths that maximize throughput, namely the average number of reached goals per timestep.*

Lifelong MAPF Algorithms

Solving MAPF optimally is known to be NP-hard [Yu and LaValle, 2013]. Lifelong MAPF poses an even greater challenge as agents consistently receive new goals, requiring the

continuous computation of new paths. Consequently, state-of-the-art algorithms approach lifelong MAPF by decomposing it into a series of (modified) one-shot MAPF instances, usually one at each timestep, assuming that minimizing their sum-of-costs enhances lifelong MAPF throughput. They can be divided into three categories. To show the generality of our GGO algorithms, we select a leading algorithm from each category to conduct our experiments.

Replan All. We replan all agents at every timestep (or every few timesteps) [Wan *et al.*, 2018; Li *et al.*, 2021]. In each replanning cycle, we solve a MAPF instance with the start locations being the current locations of all agents and the goals being their current goals. We select RHCR [Li *et al.*, 2021] as a representative algorithm from this category.

Replan New. This category is similar to the previous one except that, at every timestep, we replan only agents that have just reached their current goals and have been assigned new goals [Cap *et al.*, 2015; Ma *et al.*, 2017a; Grenouilleau *et al.*, 2019; Liu *et al.*, 2019]. Since agents being replanned must avoid collisions with agents not being replanned, methods in this category need to impose constraints on the map structure and goal locations, often denoted as well-formed maps, to ensure the existence of collision-free paths. We select Dummy Path Planning (DPP) [Liu *et al.*, 2019; Li *et al.*, 2021] as a representative algorithm from this category.

Reactive. In contrast to the previous two categories, reactive methods plan paths for each agent without considering collisions with other agents (resulting in paths with no wait actions) and then resolve collisions reactively through predefined rules [Wang and Botea, 2008; Okumura *et al.*, 2019; Yu and Wolf, 2023], such as inserting wait actions or taking short detours. We select PIBT [Okumura *et al.*, 2019] as a representative algorithm from this category. It is complete on biconnected graphs and runs significantly faster than algorithms in other categories.

2.2 Guidance Graph Optimization

To maintain generality, we consider graph $G(V, E)$ to be either directed, undirected, or mixed. We use $\{u, v\}$ and (u, v) to denote an undirected edge and a directed edge, respectively. We use E_{und} and E_{dir} to denote the subsets of edges in E that are undirected and directed, respectively.

Definition 3 (Guidance Graph). *Given a graph $G(V, E)$ for lifelong MAPF, we define a guidance graph as a directed weighted graph $G_g(V_g, E_g, \omega)$ with the same vertex set $V_g = V$. Each edge in E_g corresponds to an action that an agent can take at each vertex, with the edge weight indicating the action cost. Formally, we define $E_g = E_{wait} \cup E_{move}$ with*

$$E_{wait} = \bigcup_{v \in V} \{(v, v)\} \quad (1)$$

$$E_{move} = \left\{ \bigcup_{\{u, v\} \in E_{und}} \{(u, v), (v, u)\} \right\} \cup E_{dir}. \quad (2)$$

All edges weights are collectively represented as a vector $\omega \in \mathbb{R}_{>0}^{|E_g|}$.

Planning with guidance graphs. To utilize guidance graphs in lifelong MAPF, we redefine the sum-of-costs of the underlying (one-shot) MAPF instances as the sum of the action

		Representation			Generation			Usage Method
		Edge direction	Move cost	Wait cost	Design	MAPF	Online Update	
1	Jansen and Sturtevant [2008]	soft	\mathbb{R}^+	N/A	handcrafted procedure	lifelong	Yes	reactive
2	Wang and Botea [2008]	strict	1	N/A	crisscross	one-shot	No	reactive
3	Cohen <i>et al.</i> [2015]	soft	1 or c	1	crisscross	one-shot	No	ECBS
4	Cohen <i>et al.</i> [2016]	soft	1 or c	1	handcrafted procedure	one-shot	No	ECBS
5	Han and Yu [2022]	soft	\mathbb{R}^+	N/A	handcrafted procedure	both	Yes	reactive & RHCR
6	Li and Sun [2023]	soft	1 or c	1	crisscross	lifelong	No	RHCR
7	Yu and Wolf [2023]	soft	\mathbb{R}^+	N/A	handcrafted procedure	lifelong	Yes	reactive
8	Chen <i>et al.</i> [2024]	soft	\mathbb{R}^+	N/A	handcrafted procedure	both	Yes	PIBT
9	GGO (ours)	soft	\mathbb{R}^+	\mathbb{R}^+	automatic	lifelong	No	many

Table 1: Overview of previous works on representation, generation, and usage of guidance in MAPF. For edge direction, strict means movement is unidirectional along each edge, and soft means bidirectional. Move cost refers to the cost of moving to an adjacent vertex and wait cost refers to the cost of waiting at the current vertex. For move cost, “1 or c ” means that the value considers 1 and a scalar $c > 1$ only. For design, handcrafted procedure refers to using a manually designed procedure to generate guidance, while crisscross refers to the popular human-designed guidance [Cohen, 2020].

costs across all paths for all agents (instead of the total number of actions). This leads to a minor modification to existing lifelong MAPF algorithms. Specifically, when planning paths for each agent, instead of seeking the shortest path on G , we aim to find a cost-minimal path on G_g . This modification alters the MAPF objective without compromising feasibility.

Definition 4 (Guidance Graph Optimization (GGO)). *Given a graph $G(V, E)$, an objective function $f : \mathbb{R}^{|E_g|} \rightarrow \mathbb{R}$, as well as predefined lower and upper bounds ω_{lb} and ω_{ub} ($0 < \omega_{lb} \leq \omega_{ub}$) for edge weights, the GGO problem searches for the optimal guidance graph $G_g^*(V_g, E_g, \omega^*)$ with*

$$\omega^* = \arg \max_{\omega_{lb} \leq \omega \leq \omega_{ub}} f(\omega). \quad (3)$$

In this paper, our objective function f is a simulator that runs a given lifelong MAPF algorithm on a given guidance graph and returns the throughput.

3 Guidance in MAPF

While the term “guidance” has not been explicitly proposed in the MAPF literature, the concept of enforcing global guidance and rules to enhance MAPF has been employed by numerous works in various ways. In this section, we present a summary of these works and provide a comprehensive review of how they represent, generate, and utilize guidance. Table 1 shows a comparison between them and our GGO. We refer to them by their indices in the table for the rest of this section.

Representing Guidance. Previous works primarily represent guidance through modified edge directions or movement costs. They are all particular cases within the definition of our guidance graph, and none of them consider varying wait costs. We roughly divide them into 4 categories. (1) Inspired by potential-field and flow-field methods used in swarm robotics, Work 1 represents guidance through a direction map. This map assigns a direction vector to every vertex of G and sets the movement cost along an edge to be inversely related to the dot product of its direction vector and the vector of the edge, encouraging agents to move along the direction vectors. Consequently, a direction map can be transformed into a guidance graph with edge weights defined as dot products mentioned above. (2) Work 2 turns

some undirected edges into unidirectional, strictly prohibiting agents from moving against the assigned edge directions. This can be seen as a special guidance graph with infinite weights for constrained edges and 1 for others. (3) Works 3, 4, and 6 use the highway idea that converts undirected edges into directed edges in both directions and then selects a subset of directed edges to be highways. They assign a weight of 1 for all highway edges and a weight of a predefined constant $c > 1$ for other edges, encouraging agents to move along the highway edges. Highways are special guidance graphs with restrictions on the values of movement and wait costs. (4) Work 5 uses a temporal heuristic function to estimate the movement cost between adjacent vertices at each timestep. This temporal heuristic function can be viewed as a time-extended guidance graph where edge weights can be different at different timesteps. However, the time-extended guidance graph is only applicable while an online update mechanism is incorporated (see below for details). (5) Works 7 and 8 represent guidance similarly to our guidance graphs, with the distinction that they do not allow self-edges and thus cannot represent wait costs.

Generating Guidance. An important distinction between our work and previous works is that we are the first to propose an automated method for generating guidance. All previous works either use handcrafted guidance, such as crisscross highways (Works 2, 3, and 6), or use handcrafted procedures to generate guidance from a heatmap or a similar data structure that predicts traffic flows (Works 1, 4, 5, 7, and 8). More specifically, Work 1 computes direction vectors (of the direction map) from past traffic flows and then uses a handcrafted equation to convert them into movement costs. Work 4 introduces two methods, GM and HM, for generating highways. GM uses a graphic model [Koller and Friedman, 2009] with a number of handcrafted features obtained from the estimated traffic flow. HM converts the estimated traffic flow into a score for each edge using a handcrafted score function and selects edges based on a predefined score threshold. Works 5 and 8 collect the planned paths of all agents and convert them into movement costs through handcrafted equations. Last, Work 7 uses a data-driven model to predict traffic flow, or more specifically, the delays that the agents will encounter (due to collision avoidance etc.) and directly uses the




Figure 2: CMA-ES for GGO. The edge weights are iteratively sampled from a Gaussian distribution and then evaluated by a lifelong MAPF simulator. The simulated results are used to update the Gaussian distribution towards high-throughput regions.

predicted delays as movement costs.

Therefore, we select 4 baseline methods to generate guidance graphs in our experiments: (1) *Unweighted*, where no guidance is used, (2) *Crisscross*, (3) *HM Cost*, adapted from HM in Work 4, and (4) *Traffic Flow*, adapted from Work 8. We did not compare Work 1, as HM from Work 4 is inspired by it. We did not compare GM from Work 4, as its performance is similar to HM. We did not compare Work 5 because the non-temporal version of their proposed guidance is similar to Works 4 and 8. We did not compare Work 7, because, while it obtains predicted traffic flow differently from Work 4, the procedure of converting predicted traffic flow to guidance is similar.

Please note that the HM Cost and Traffic Flow used in our experiments do not use the original traffic flow models in their papers. This is because Work 4 tackles one-shot MAPF and predicts traffic flows by planning shortest paths between the start and goal locations of the agents, which is not realistic in lifelong MAPF as goal locations are unknown in advance. Work 8 tackles lifelong MAPF but assumes the guidance graph can be updated on the fly using real-time traffic information, while we assume that our guidance graph is optimized offline, and thus we do not have access to real-time traffic information. We focus on optimizing the guidance graphs offline because adding online adaptation not only requires additional computation to update the guidance graph but can also dramatically slow down path planning. This is because we need to either update the heuristic for the single-agent A* search when the guidance graph is updated or use a less informed heuristic without update. For example, Work 8 incorporates an online adaptation mechanism in PIBT, but it slows down the algorithm by 2-10 times. Thus, in this paper, we use the same traffic flow model, namely the tile-usage map obtained from simulation, for both HM Cost and Traffic Flow methods. More details can be found in Appendix A.

Using Guidance. All previous works study their guidance methods with a specific MAPF algorithm, such as ECBS [Barer *et al.*, 2014] and RHCR [Li *et al.*, 2021], so it remains unclear whether and how well their methods can generalize to other MAPF algorithms. For example, an evident limitation of methods designed for reactive (lifelong) MAPF algorithms (namely Works 1, 2, 5, 7, and 8) is that, since paths planned by reactive methods do not include wait actions, these guidance methods, by design, do not define wait costs, making it non-trivial to extend them to (lifelong) MAPF algorithms in other categories. In contrast, we assess our GGO methods with three leading lifelong MAPF algo-

gorithms from different categories, thereby demonstrating their generality.

4 Approach

We first introduce CMA-ES to solve GGO directly. Then we introduce Parameterized Iterative Update (PIU), which uses CMA-ES to optimize an update model that iteratively generates a guidance graph based on simulated traffic information.

4.1 CMA-ES

CMA-ES [Hansen, 2016] is a derivative-free, black-box optimization algorithm based on covariance matrix adaptation. Figure 2 gives an overview of using CMA-ES to solve GGO. Specifically, we model the edge weight vector ω as a multivariate Gaussian distribution. We then iteratively sample from the distribution for a new batch of b edge weight vectors, forming b guidance graphs. We normalize each ω to meet the bound constraint $\omega_{lb} \leq \omega \leq \omega_{ub}$. We then evaluate each guidance graph by running $N_{e.cma}$ simulations in a given lifelong MAPF simulator and computing the average throughput. The evaluated guidance graphs are ranked based on their throughput, and the top N_{best} of them are used to update the mean and covariance of the Gaussian distribution. We run CMA-ES for I iterations and return the guidance graph with the highest throughput as the solution.

Handling Bounds through Normalization. We use min-max normalization to enforce the bound constraint because it does not affect path-planning solutions. To prove it, consider two guidance graphs with edge weights ω_1 and $\omega_2 = C \cdot \omega_1$, where $C \in \mathbb{R}^+$. Since the weight of every edge is scaled by the same scalar C , the paths returned by the lifelong MAPF algorithms with low-level single agent solvers minimizing the sum of edge weights do not change. We show additional experiments in Appendix B.1 that min-max normalization yields better solutions than representative bounds handling methods introduced in a prior study Biedrzycki [2020].

4.2 Parameterized Iterative Update

CMA-ES is known to scale poorly to high dimensional search spaces, making it challenging to optimize guidance graph for large maps. Therefore, we propose Parameterized Iterative Update (PIU). Figure 3 gives an overview of PIU, and Algorithm 1 provides the pseudocode. On a high level, PIU leverages a parameterized *update model* to iteratively update the edge weights of the guidance graph using traffic information obtained from lifelong MAPF simulations. PIU can work with a wide variety of optimization methods. In this work, we choose to use CMA-ES to optimize the update model.

Definition 5 (Update model). *Given a guidance graph $G_g(V_g, E_g, \omega)$, an update model is a function $\pi_\theta : \mathbb{R}^{|E_g|} \times \mathbb{R}^{|E_g|} \rightarrow \mathbb{R}^{|E_g|}$ that computes the updated edge weights $\omega' \in \mathbb{R}_{>0}^{|E_g|}$ given the current edge weights $\omega \in \mathbb{R}_{>0}^{|E_g|}$ and edge usage $U_{E_g} \in \mathbb{R}_{>0}^{|E_g|}$. The edge usage is the frequency with which each edge is used by the agents in the lifelong MAPF simulation. The model π_θ is parameterized by a vector $\theta \in \Theta$, where Θ is the space of all parameters.*

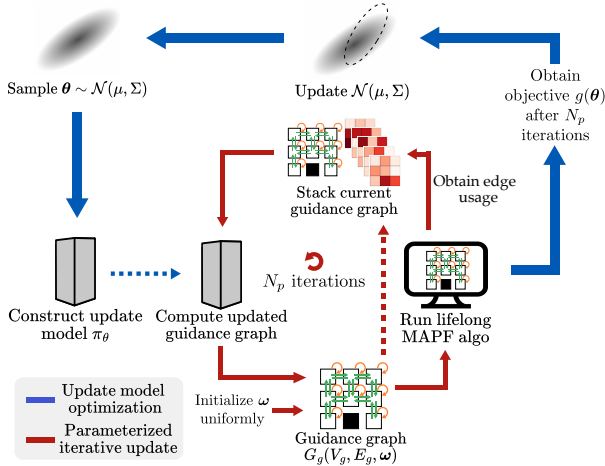


Figure 3: PIU for GGO. Starting with a guidance graph with uniform edge weights, we run MAPF simulations to get the edge usage. We then use an update model π_θ to update the edge weights. We run this process iteratively for N_p iterations. The update model π_θ is optimized using CMA-ES.

PIU. The red loop in Figure 2 gives an overview of PIU, and Lines 9 to 19 of Algorithm 1 provides the pseudocode. We first form an update model π_θ parameterized by a given parameter vector θ (Line 10). We then start an iterative update procedure (Lines 11 to 18). In each iteration, we either initialize the edge weights to 1 (Line 12) in the first iteration or use the update model π_θ to update the edge weights (Line 13) in the following iterations. We then construct the current guidance graph using ω (Line 14). We then run N_{e_piu} lifelong MAPF simulations (Lines 15 and 16), computing the average throughput f and edge usage U_{E_g} (Lines 17 and 18). We run PIU for N_p iterations. Finally, we return the throughput f of the last iteration (Line 19).

Update Model Optimization. To train the update model, we run the PIU algorithm for N_p iterations with update model π_θ given parameters θ . We search for optimal parameters $\theta^* = \arg \max_{\theta \in \Theta} \text{PIU}(\theta, N_p)$ using CMA-ES. Lines 1 to 8 of Algorithm 1 show the pseudocode. Starting with a given initial multivariate Gaussian distribution (Line 1), the algorithm samples b parameter vectors (Line 3) and runs PIU with them (Lines 4 and 5). Based on the returned throughput values, it keeps track of the best update model (Line 6) and updates the Gaussian distribution (Line 7), starting a new iteration. The optimization ends after running the above process for I iterations (Line 2). Finally, the algorithm returns the best update model and the corresponding throughput (Line 8).

On the Advantage of PIU. Compared to directly using CMA-ES, the advantage of optimizing the update model and using PIU to generate guidance graph is two-folds. First, optimizing the update model reduces the dimension of search space. Although solving GGO directly with CMA-ES is versatile and applicable to various lifelong MAPF algorithms and maps, its effectiveness diminishes in high-dimensional search spaces, making it challenging to use CMA-ES to search for edge weights directly for large maps. Specifically,

Algorithm 1: Update model optimization

Input: μ_0, Σ_0 : initial mean and covariance matrix of the multivariate Gaussian distribution.
 N_p : number of iterations to run in PIU.
 $update_gaussian$: function to update μ and Σ according to evaluated parameter vectors.
 $simulate$: function to run lifelong MAPF simulation and return edge usage U_{E_g} and throughput f .

```

1 Initialize  $\mu \leftarrow \mu_0, \Sigma \leftarrow \Sigma_0, \theta^* \leftarrow \text{NULL}, g^* \leftarrow -\infty$ 
2 for  $i \leftarrow 1$  to  $I$  do
3   Sample  $b$  parameter vectors  $\theta_1, \dots, \theta_b \sim \mathcal{N}(\mu, \Sigma)$ 
4   for  $k \leftarrow 1$  to  $b$  do
5      $g_k \leftarrow \text{PIU}(\theta_k, N_p)$ 
6     if  $g_k > g^*$  then  $g^* \leftarrow g_k, \theta^* \leftarrow \theta_k$ 
7    $\mu, \Sigma \leftarrow update\_gaussian(g_{1 \sim b}, \theta_{1 \sim b})$ 

```

```
8 return  $g^*, \theta^*$ 
```

```
9 Function  $\text{PIU}(\theta, N_p)$ :
```

```

10 Construct update model  $\pi_\theta$ 
11 for  $j \leftarrow 1$  to  $N_p$  do
12   if  $j = 1$  then  $\omega \leftarrow 1$ 
13   else  $\omega \leftarrow \pi_\theta(\omega, U_{E_g})$ 
14   Construct guidance graph  $G_g(V_g, E_g, \omega)$ 
15   for  $q \leftarrow 1$  to  $N_{e\_piu}$  do
16      $U_{E_g}^{(q)}, f^{(q)} \leftarrow simulate(G_g)$ 
17    $U_{E_g} \leftarrow \frac{1}{N_{e\_piu}} \sum_{q=1}^{N_{e\_piu}} U_{E_g}^{(q)}$ 
18    $f \leftarrow \frac{1}{N_{e\_piu}} \sum_{q=1}^{N_{e\_piu}} f^{(q)}$ 
19 return  $f$ 

```

CMA-ES employs a full-rank $n \times n$ covariance matrix to model its Gaussian distribution in an n -dimensional space, leading to quadratic increases in both time and space complexity [Varelas *et al.*, 2018]. In the case of GGO, the number of edge weights of a guidance graph increases at least linearly with the number of vertices, while our update model maintains a consistent number of parameters regardless of the size of the guidance graph, offering a more scalable solution than directly applying CMA-ES. Second, the optimized update model is not specific to the map it is optimized on. Different maps with similar layouts could potentially have similar high-throughput guidance graphs that can be generated by the same update model. The guidance graph optimized by CMA-ES, on the other hand, consists of edge weights for a specific map.

5 Experimental Evaluation

In this section, we compare guidance graphs optimized by CMA-ES and PIU with various baselines and assess the capability of PIU to generate high-throughput guidance graphs for maps of larger sizes with similar layouts.

5.1 Experiment Setup

General Setups. Table 2 outlines our experimental setup. Column 2 shows the lifelong MAPF algorithms. Follow-

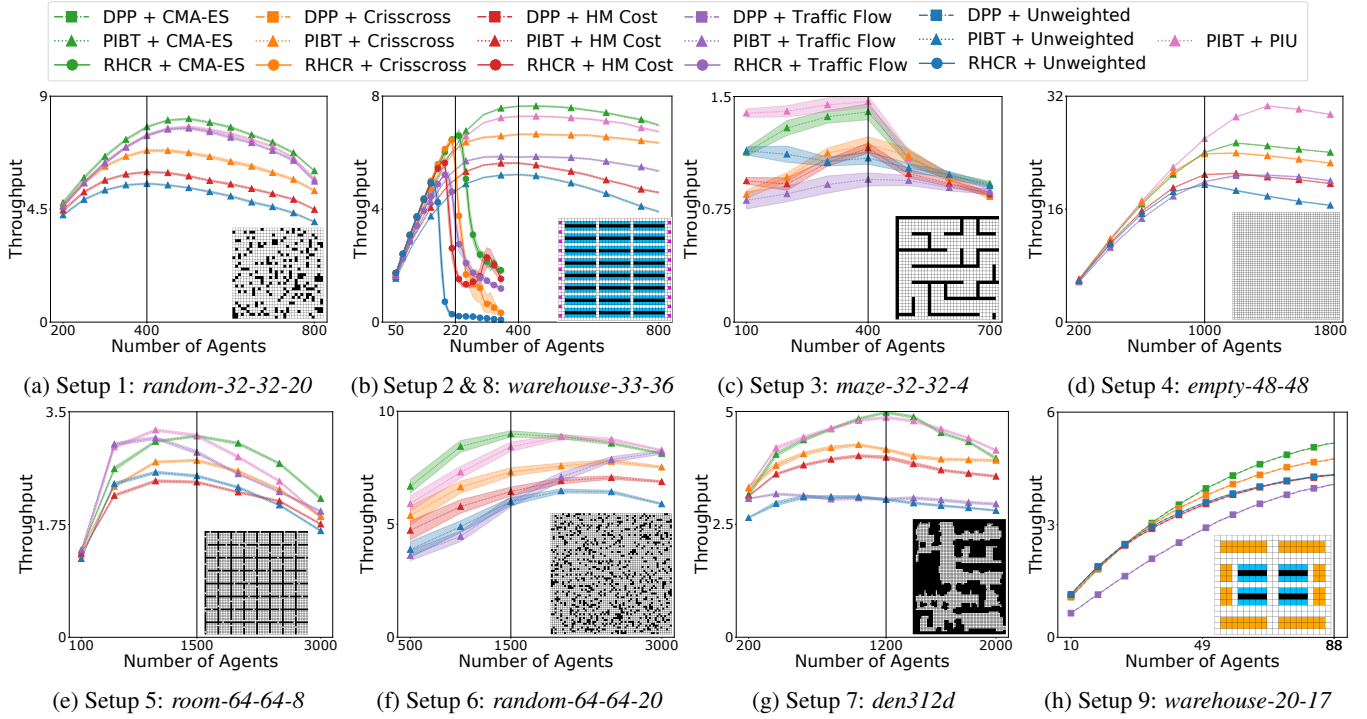


Figure 4: Throughput with different numbers of agents. The guidance graphs are optimized with N_a agents, which is indicated by the black vertical lines. In (b), the black vertical lines at 220 and 400 agents indicate N_a for setups 8 and 2, respectively.

Setup	MAPF	Map	$ E_g $	$ E_{wait} $	$ E_{move} $	N_a	GGO
1		<i>random-32-32-20</i>	3,359	819	2,540		
2		<i>warehouse-33-36</i>	4,074	948	3,126	400	
3		<i>maze-32-32-4</i>	3,484	790	2,694		
4	PIBT	<i>empty-48-48</i>	11,328	2,304	9,024	1,000	CMA-ES & PIU
5		<i>room-64-64-8</i>	14,340	3,232	11,108		
6		<i>random-64-64-20</i>	13,568	3,270	10,298	1,500	
7		<i>den312d</i>	11,227	2,445	8,782	1,200	
8	RHCR	<i>warehouse-33-36</i>	4,074	948	3,126	220	CMA-ES
9	DPP	<i>warehouse-20-17</i>	1,478	320	1,158	88	CMA-ES
10	PIBT	<i>warehouse-33-36</i>	4,074	948	3,126	150	CMA-ES & PIU

Table 2: Summary of the experiment setup. N_a is the number of agents. $|E_g|$, $|E_{wait}|$, and $|E_{move}|$ are the number of wait edges, movement edges, and all edges in the guidance graph, respectively. Setups 1 to 9 compare our optimized guidance graphs with the baselines. Setup 10 compares PIBT with GGO against RHCR without GGO when there are fewer agents.

ing the recommendations of previous works [Li *et al.*, 2021; Zhang *et al.*, 2023b], we use PBS [Ma *et al.*, 2019] and SIPP [Phillips and Likhachev, 2011] as the MAPF solver and the single-agent solver, respectively, in both RHCR and DPP and use $h = 5$ and $w = 10$ in RHCR.

Column 3 outlines the maps, all being 4-neighbor grids, including two warehouse maps (*warehouse-33-36* and *warehouse-20-17*) from previous works [Li *et al.*, 2021; Zhang *et al.*, 2023b] and six maps (*random-32-32-20*, *maze-32-32-4*, *empty-48-48*, *room-64-64-8*, *random-64-64-20*, and *den312d*) from the MAPF benchmark [Stern *et al.*, 2019]. We choose multiple maps for PIBT to demonstrate that both CMA-ES and PIU work for different maps. We show the

maps at the corners of Figure 4. In all maps, black tiles are obstacles, and non-black tiles are traversable. In warehouse maps, orange tiles are home locations, blue tiles are endpoints, and purple tiles are workstations. In *warehouse-20-17*, agents start from orange tiles and move constantly between blue tiles. In *warehouse-33-36*, agents start from non-black tiles and move between blue and purple tiles alternatively. In all other maps, agents start from and move between white tiles. The agents' goals are assigned uniformly randomly.

Columns 4 to 6 show the number of edges in the guidance graphs of the corresponding maps. In setups 1 to 7, we optimize all $|E_g| = |E_{wait}| + |E_{move}|$ edges. In setups 8 and 9, however, SIPP cannot handle different wait costs at different vertices. Therefore, we optimize the wait costs of all vertices as one variable, resulting in $|E_{move}| + 1 = 3, 127$ and $1, 159$ variables to be optimized in the guidance graphs in setups 8 and 9, respectively. Column 7 is the number of agents used in lifelong MAPF simulations, with a larger N_a for PIBT compared to RHCR and DPP to demonstrate that both CMA-ES and PIU work for congested scenarios.

Column 8 shows the GGO algorithms we run for each setup. We apply CMA-ES across all setups to demonstrate its versatility. However, due to computational constraints, we focus on using PIU primarily with PIBT. While both PIU and CMA-ES conduct the same number of simulations, there is a notable difference in their execution. In CMA-ES, all $N_{e.cma}$ simulations in each guidance graph evaluation can be parallelized. In contrast, PIU runs the simulations sequentially, resulting in slower runtime.

We choose the hyperparameters of CMA-ES and PIU such

Setup	MAPF + GGO	SR	Throughput	CPU Runtime (s)
1	PIBT + CMA-ES	100%	7.78 ± 0.02	1.31 ± 0.01
	PIBT + PIU	100%	7.46 ± 0.02	1.29 ± 0.02
	PIBT + Crisscross	100%	6.84 ± 0.02	1.24 ± 0.02
	PIBT + HM Cost	100%	5.98 ± 0.02	1.17 ± 0.02
	PIBT + Traffic Flow	100%	7.43 ± 0.02	1.19 ± 0.02
	PIBT + Unweighted	100%	5.52 ± 0.01	1.20 ± 0.02
2	PIBT + CMA-ES	100%	7.64 ± 0.01	1.27 ± 0.01
	PIBT + PIU	100%	7.28 ± 0.01	1.22 ± 0.01
	PIBT + Crisscross	100%	6.65 ± 0.01	1.21 ± 0.01
	PIBT + HM Cost	100%	5.63 ± 0.01	1.24 ± 0.01
	PIBT + Traffic Flow	100%	5.84 ± 0.01	1.23 ± 0.01
	PIBT + Unweighted	100%	5.22 ± 0.01	1.25 ± 0.01
3	PIBT + CMA-ES	100%	1.40 ± 0.03	0.60 ± 0.01
	PIBT + PIU	100%	1.47 ± 0.02	0.60 ± 0.01
	PIBT + Crisscross	100%	1.18 ± 0.03	0.59 ± 0.01
	PIBT + HM Cost	100%	1.16 ± 0.02	0.65 ± 0.01
	PIBT + Traffic Flow	100%	0.95 ± 0.02	0.63 ± 0.01
	PIBT + Unweighted	100%	1.09 ± 0.02	0.58 ± 0.01
4	PIBT + CMA-ES	100%	24.04 ± 0.01	2.86 ± 0.05
	PIBT + PIU	100%	25.98 ± 0.01	2.28 ± 0.03
	PIBT + Crisscross	100%	23.84 ± 0.02	2.82 ± 0.04
	PIBT + HM Cost	100%	20.90 ± 0.02	2.79 ± 0.05
	PIBT + Traffic Flow	100%	19.90 ± 0.02	2.73 ± 0.04
	PIBT + Unweighted	100%	19.48 ± 0.03	2.78 ± 0.04
5	PIBT + CMA-ES	100%	3.12 ± 0.01	4.91 ± 0.06
	PIBT + PIU	100%	3.13 ± 0.01	4.59 ± 0.05
	PIBT + Crisscross	100%	2.75 ± 0.01	4.67 ± 0.07
	PIBT + HM Cost	100%	2.41 ± 0.01	4.71 ± 0.05
	PIBT + Traffic Flow	100%	2.87 ± 0.01	4.55 ± 0.05
	PIBT + Unweighted	100%	2.51 ± 0.01	4.95 ± 0.06
6	PIBT + CMA-ES	100%	9.00 ± 0.07	3.52 ± 0.06
	PIBT + PIU	100%	8.43 ± 0.11	4.25 ± 0.07
	PIBT + Crisscross	100%	7.31 ± 0.09	3.55 ± 0.05
	PIBT + HM Cost	100%	6.45 ± 0.10	3.75 ± 0.06
	PIBT + Traffic Flow	100%	6.00 ± 0.12	3.39 ± 0.05
	PIBT + Unweighted	100%	6.01 ± 0.09	3.45 ± 0.05
7	PIBT + CMA-ES	100%	4.98 ± 0.01	2.71 ± 0.04
	PIBT + PIU	100%	4.87 ± 0.01	2.61 ± 0.03
	PIBT + Crisscross	100%	4.16 ± 0.01	3.27 ± 0.05
	PIBT + HM Cost	100%	3.99 ± 0.02	2.90 ± 0.06
	PIBT + Traffic Flow	100%	3.06 ± 0.01	2.73 ± 0.06
	PIBT + Unweighted	100%	3.05 ± 0.01	3.17 ± 0.05
8	RHCR + CMA-ES	100%	6.58 ± 0.04	91.24 ± 15.66
	RHCR + Crisscross	100%	5.59 ± 0.20	3771.71 ± 659.20
	RHCR + HM Cost	0%	N/A	N/A
	RHCR + Traffic Flow	32%	3.57 ± 0.20	324.12 ± 63.91
	RHCR + Unweighted	0%	N/A	N/A
9	DPP + CMA-ES	100%	5.17 ± 0.00	28.75 ± 0.53
	DPP + Crisscross	100%	4.76 ± 0.01	17.36 ± 0.31
	DPP + HM Cost	100%	4.32 ± 0.01	29.94 ± 1.39
	DPP + Traffic Flow	100%	4.07 ± 0.00	578.43 ± 53.61
	DPP + Unweighted	100%	4.34 ± 0.01	20.30 ± 0.49
10	PIBT + CMA-ES	100%	4.74 ± 0.01	0.58 ± 0.00
	PIBT + PIU	100%	4.77 ± 0.01	0.58 ± 0.01
	PIBT + Unweighted	100%	3.75 ± 0.00	0.50 ± 0.00
	RHCR + Unweighted	100%	4.95 ± 0.00	87.37 ± 0.72

Table 3: Success rates (SR), throughput, and CPU runtimes of the simulations on different guidance graphs. For RHCR and DPP, the success rate is the percentage of simulations that end without congestion. For PIBT, it is the percentage of simulations that end without timeout. We measure the throughput and CPU runtime over only successful simulations.

that they run the same number of simulations to ensure a fair comparison. In particular, we set batch size $b = 100$ and the number of iterations $I = 100$ for both CMA-ES and PIU, resulting in a total of $b \times I = 10k$ objective function evaluations for both algorithms. In each iteration, we select the top $N_{best} = 50$ solutions to update the Gaussian distribution. For CMA-ES, each evaluation runs $N_{e_cma} = 5$ simulations, re-

sulting in 50k simulations. For PIU, each evaluation runs the PIU algorithm for $N_p = 5$ iterations and each iteration runs $N_{e_piu} = 1$ simulation, resulting in 50k simulations, identical to CMA-ES. We run each simulation for 1,000 timesteps. For RHCR and DPP, we stop the simulation early in case of congestion, which happens if more than half of the agents wait at their current location.

Update Model. Given our use of grid maps in the experiments, we use a Convolutional Neural Network (CNN) as our update model, which can generate guidance graphs for maps of arbitrary sizes. The CNN has 3 convolutional layers of kernel sizes 3, 1, 1, respectively. Each layer is followed by a ReLU activation and a batch normalization layer. The update model has 4,231 parameters. For a map of dimension $h \times w$, we represent the edge weights ω of the guidance graph as a tensor of size $h \times w \times 5$, where the first four channels are the movement costs and the last channel is the wait costs.

Baselines. As mentioned in Section 3, we have 4 baseline guidance graphs, namely (1) *Unweighted*, (2) *Crisscross* [Cohen, 2020], (3) *HM Cost* [Cohen *et al.*, 2016], and (4) *Traffic Flow* [Chen *et al.*, 2024]. We discuss the methods of generating the baseline guidance graphs in Appendix A.

Evaluation. To evaluate PIU, we use the optimized update model to run PIU with $N_{e_piu} = 1$ to generate the guidance graph. In Appendix B.2, we show that the choice of N_{e_piu} does not have significant impact on the throughput of the generated guidance graphs. When we evaluate a guidance graph from CMA-ES, PIU, or baselines with a given number of agents, we run 50 simulations, each for 1,000 timesteps, and report the results with both means and standard errors.

Implementation. We implement the update model in PyTorch [Paszke *et al.*, 2019], CMA-ES in Pyribs [Tjanaka *et al.*, 2023], and Traffic Flow and HM Cost guidance graph generation in Python. We implement the lifelong MAPF algorithms in C++ based on openly available implementation from previous works [Li *et al.*, 2021; Okumura *et al.*, 2019].

Compute Resource. We run our experiments on two machines: (1) a local machine with a 64-core AMD Ryzen Threadripper 3990X CPU, 192 GB of RAM, and an Nvidia RTX 3090Ti GPU, and (2) a high-performing cluster with numerous 64-core AMD EPYC 7742 CPUs, each with 256 GB of RAM. We measure all CPU runtime on machine (1).

5.2 Results

GGO vs Baselines. We first compare our optimized guidance graphs with the baseline guidance graphs. For each guidance graph, we run 50 simulations and report the numerical results in Table 3 in the format of $x \pm y$, where x is the average and y is the standard error. Both CMA-ES and PIU outperform all baseline guidance graphs in all setups in terms of throughput. Specifically, CMA-ES outperforms all baseline guidance graphs in all setups, showing the versatility of the algorithm. For the baseline methods, the human-designed crisscross guidance performs quite well in setups 2, 3, 4, 6, 7, 8, and 9, outperforming all other baselines. Traffic Flow is more competitive in setups 1 and 5. When comparing the throughput of CMA-ES and PIU, no clear winner emerges: CMA-ES wins in setups 1, 2, 6, and 7, PIU wins in 3 and 4,




Figure 5: Setup 10: An optimized guidance graph enables PIBT to have competitive throughput with RHCR despite the advantage of RHCR with fewer agents.

and they perform similarly in setup 5. Appendix C visualizes all the optimized guidance graphs.

We also report CPU runtimes in Table 3, although runtimes are not the optimization objectives of our GGO algorithms. PIBT runs very fast and finishes all 1,000-timestep simulations in 5 seconds, implying each timestep takes less than 0.005 seconds. Therefore, the runtime difference between different methods is negligible in practice. RHCR and DPP are significantly slower than PIBT. For RHCR, CMA-ES leads to the best runtime. For DPP, it runs PBS to solve a one-shot MAPF instance for all agents in the first timestep and then, in each timestep, replans only for agents that have just reached their goals. We conjecture that the slower runtime of DPP with CMA-ES than baselines comes from the slower runtime in the first timestep. This is because the optimized guidance graph could encourage the agents to take longer paths in order to avoid congestion.

To further understand the performance of the optimized guidance graphs, we vary the number of agents and plot the throughput in Figure 4. The trends are similar in all maps, with CMA-ES and PIU generally outperforming all baselines, except that Traffic Flow matches PIU with fewer agents in *random-32-32-20* and *room-64-64-8*. However, Traffic Flow is less competitive in all other maps, indicating that the performance of Traffic Flow depends on the map structures.

RHCR vs PIBT+GGO. Given the considerable runtime advantage of PIBT over RHCR (and DPP), we conduct an additional experiment, as detailed in setup 10 of Table 2, to explore whether PIBT with an optimized guidance graph can achieve higher throughput than RHCR without guidance graphs. We choose $N_a = 150$ because it is the largest number with which RHCR without guidance graphs can maintain a 100% success rate. Setup 10 in Table 3 shows the numerical results. While RHCR still has the highest throughput, both GGO methods significantly reduce the throughput gap between PIBT and RHCR, from 24.2% to less than 4.2%. Therefore, with the help of our optimized guidance graph, we can enable a greedy, distributed, yet extremely fast rule-based MAPF algorithm (PIBT) to achieve performance comparable to a centralized, computationally heavy, search-based MAPF algorithm (RHCR). Figure 5 further compares PIBT and RHCR with various numbers of agents. The throughput of RHCR quickly drops after 150 agents, while that of PIBT maintains an increasing trend with more agents.

Update Model Transferability. We attempt to transfer the update model optimized with setup 2 to larger warehouse




Figure 6: Max throughput and the number of agents at this maximum. PIU Transfer refers to using the update model optimized in setup 2 to generate guidance graphs.

maps with similar layouts. We mimic the layout pattern of *warehouse-33-36* and design larger maps with sizes up to 93×91 by repeating blocks of 10 shelves and endpoints and placing workstations on the left and right borders. Figure 9 in Appendix B.3 plots the resulting maps. We use the optimized update model from setup 2 to generate guidance graphs for these maps with an increasing number of agents, ranging from around 10% to 90% of the non-black tiles in the maps. We then run 50 simulations in each of the generated guidance graphs with PIBT. We plot the maximum throughput achieved in each map and the corresponding number of agents in Figure 6, comparing PIU Transfer with baseline guidance graphs. We observe that PIU Transfer dominates all baselines with all sizes in terms of throughput. Notably, while crisscross has the second highest throughput across different map sizes, PIU Transfer can achieve higher throughput with an equal or smaller number of agents. Appendix B.3 shows the throughput with different numbers of agents in larger warehouse maps.

6 Conclusion

We define guidance graphs and GGO to maximize the throughput of lifelong MAPF, reviewing previous works on guidance in MAPF and highlighting the generality of our guidance graph. We propose CMA-ES and PIU that optimize guidance graphs across different algorithms and maps. We also show that the update model can generate guidance graphs for larger maps with similar patterns.

Our work is limited in many ways, yielding numerous future directions. First, both CMA-ES and PIU are computationally expensive, requiring a large number of simulations, taking 1.2 hours (setup 1) to 55 hours (setup 8) to run on machine (2) in Section 5.1. Future works can focus on reducing the computational requirements of both methods. Second, our optimized guidance graphs improve throughput but such improvement lacks explainability. Future works can focus on either generating more explainable guidance graphs or analyzing the explainability of our optimized guidance graphs. Third, although our guidance graph can be used with on-line update mechanisms introduced in previous works [Chen *et al.*, 2024; Yu and Wolf, 2023], we limit our experiment settings without such mechanisms. Integrating these mechanisms with our GGO algorithms could further enhance MAPF guidance utility.

Acknowledgments

This work used Bridge-2 at Pittsburgh Supercomputing Center (PSC) through allocation CIS220115 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

References

- [Barer *et al.*, 2014] Max Barer, Guni Sharon, Roni Stern, and Ariel Felner. Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *Proceedings of the Annual Symposium on Combinatorial Search (SoCS)*, pages 19–27, 2014.
- [Biedrzycki, 2020] Rafał Biedrzycki. Handling bound constraints in CMA-ES: An experimental study. *Swarm and Evolutionary Computation*, 52:100627, 2020.
- [Cáp *et al.*, 2015] Michal Cáp, Jirí Vokřínek, and Alexander Kleiner. Complete decentralized method for on-line multi-robot trajectory planning in well-formed infrastructures. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, pages 324–332, 2015.
- [Chen *et al.*, 2024] Zhe Chen, Daniel Harabor, Jiaoyang Li, and Peter Stuckey. Traffic flow optimisation for lifelong multi-agent path finding. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 20674–20682, 2024.
- [Cohen *et al.*, 2015] Liron Cohen, Tansel Uras, and Sven Koenig. Feasibility study: Using highways for bounded-suboptimal multi-agent path finding. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, pages 2–8, 2015.
- [Cohen *et al.*, 2016] Liron Cohen, Tansel Uras, T. K. Satish Kumar, Hong Xu, Nora Ayanian, and Sven Koenig. Improved solvers for bounded-suboptimal multi-agent path finding. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3067–3074, 2016.
- [Cohen, 2020] Liron Cohen. *Efficient Bounded-Suboptimal Multi-Agent Path Finding and Motion Planning via Improvements to Focal Search*. PhD thesis, University of Southern California, 2020.
- [Damani *et al.*, 2021] Mehul Damani, Zhiyao Luo, Emerson Wenzel, and Guillaume Sartoretti. PRIMAL₂: Pathfinding via reinforcement and imitation multi-agent learning - lifelong. *IEEE Robotics and Automation Letters*, 6(2):2666–2673, 2021.
- [Grenouilleau *et al.*, 2019] Florian Grenouilleau, Willem-Jan van Hoeve, and John N. Hooker. A multi-label A* algorithm for multi-agent pathfinding. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, pages 181–185, 2019.
- [Han and Yu, 2022] Shuai D. Han and Jingjin Yu. Optimizing space utilization for more effective multi-robot path planning. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 10709–10715, 2022.
- [Hansen, 2016] Nikolaus Hansen. The CMA evolution strategy: A tutorial. *ArXiv*, abs/1604.00772, 2016.
- [Jansen and Sturtevant, 2008] M. Renee Jansen and Nathan R Sturtevant. Direction maps for cooperative pathfinding. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, pages 185–190, 2008.
- [Koller and Friedman, 2009] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009.
- [Kou *et al.*, 2020] Ngai Meng Kou, Cheng Peng, Hang Ma, T. K. Satish Kumar, and Sven Koenig. Idle time optimization for target assignment and path finding in sortation centers. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 9925–9932, 2020.
- [Li and Sun, 2023] Ming-Feng Li and Min Sun. The study of highway for lifelong multi-agent path finding. *ArXiv*, 2304.04217, 2023.
- [Li *et al.*, 2020] Jiaoyang Li, Kexuan Sun, Hang Ma, Ariel Felner, T. K. Satish Kumar, and Sven Koenig. Moving agents in formation in congested environments. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 726–734, 2020.
- [Li *et al.*, 2021] Jiaoyang Li, Andrew Tinka, Scott Kiesel, Joseph W. Durham, T. K. Satish Kumar, and Sven Koenig. Lifelong multi-agent path finding in large-scale warehouses. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 11272–11281, 2021.
- [Liu *et al.*, 2019] Minghua Liu, Hang Ma, Jiaoyang Li, and Sven Koenig. Task and path planning for multi-agent pickup and delivery. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1152–1160, 2019.
- [Ma *et al.*, 2017a] Hang Ma, Jiaoyang Li, T. K. Satish Kumar, and Sven Koenig. Lifelong multi-agent path finding for online pickup and delivery tasks. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 837–845, 2017.
- [Ma *et al.*, 2017b] Hang Ma, Jingxing Yang, Liron Cohen, T. K. Satish Kumar, and Sven Koenig. Feasibility study: Moving non-homogeneous teams in congested video game environments. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, pages 270–272, 2017.
- [Ma *et al.*, 2019] Hang Ma, Daniel Harabor, Peter J. Stuckey, Jiaoyang Li, and Sven Koenig. Searching with consistent prioritization for multi-agent path finding. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 7643–7650, 2019.

- [Okumura *et al.*, 2019] Keisuke Okumura, Manao Machida, Xavier Défago, and Yasumasa Tamura. Priority inheritance with backtracking for iterative multi-agent path finding. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 535–542, 2019.
- [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 8024–8035, 2019.
- [Phillips and Likhachev, 2011] Mike Phillips and Maxim Likhachev. SIPP: Safe interval path planning for dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 5628–5635, 2011.
- [Stern *et al.*, 2019] Roni Stern, Nathan R. Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne T. Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, T. K. Satish Kumar, Roman Barták, and Eli Boyarski. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, pages 151–159, 2019.
- [Tjanaka *et al.*, 2023] Bryon Tjanaka, Matthew C. Fontaine, David H. Lee, Yulun Zhang, Nivedit Reddy Balam, Nathaniel Dennler, Sujay S. Garlanka, Nikitas Dimitri Klapsis, and Stefanos Nikolaidis. pyribs: A bare-bones python library for quality diversity optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pages 220–229, 2023.
- [Varambally *et al.*, 2022] Sumanth Varambally, Jiaoyang Li, and Sven Koenig. Which MAPF model works best for automated warehousing? In *Proceedings of the Symposium on Combinatorial Search (SoCS)*, pages 190–198, 2022.
- [Varelas *et al.*, 2018] Konstantinos Varelas, Anne Auger, Dimo Brockhoff, Nikolaus Hansen, Ouassim Ait ElHara, Yann Semet, Rami Kassab, and Frédéric Barbaresco. A comparative study of large-scale variants of CMA-ES. In *Proceedings of the International Conference on Parallel Problem Solving from Nature (PPSN)*, pages 3–15, 2018.
- [Wan *et al.*, 2018] Qian Wan, Chonglin Gu, Sankui Sun, Mengxia Chen, Hejiao Huang, and Xiaohua Jia. Lifelong multi-agent path finding in a dynamic environment. In *Proceedings of the International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pages 875–882, 2018.
- [Wang and Botea, 2008] Ko-Hsin Cindy Wang and Adi Botea. Fast and memory-efficient multi-agent pathfinding. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, pages 380–387, 2008.
- [Yu and LaValle, 2013] Jingjin Yu and Steven M. LaValle. Structure and intractability of optimal multi-robot path planning on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1444–1449, 2013.
- [Yu and Wolf, 2023] Ge Yu and Michael Wolf. Congestion prediction for large fleets of mobile robots. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 7642–7649, 2023.
- [Zhang *et al.*, 2023a] Yulun Zhang, Matthew C. Fontaine, Varun Bhatt, Stefanos Nikolaidis, and Jiaoyang Li. Arbitrarily scalable environment generators via neural cellular automata. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 57212–57225, 2023.
- [Zhang *et al.*, 2023b] Yulun Zhang, Matthew C. Fontaine, Varun Bhatt, Stefanos Nikolaidis, and Jiaoyang Li. Multi-robot coordination and layout design for automated warehousing. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 5503–5511, 2023.