

# QD-MAPPER: A Quality Diversity Framework to Automatically Evaluate Multi-Agent Path Finding Algorithms in Diverse Maps\*

Cheng Qian<sup>†</sup>  
Carnegie Mellon University  
Pittsburgh, United States  
chengqia@andrew.cmu.edu

Yulun Zhang<sup>†</sup>  
Carnegie Mellon University  
Pittsburgh, United States  
yulunzhang@cmu.edu

Varun Bhatt  
University of Southern California  
Los Angeles, United States  
vsbhatt@usc.edu

Matthew C. Fontaine  
University of Southern California  
Los Angeles, United States  
mfontain@usc.edu

Stefanos Nikolaidis  
University of Southern California  
Los Angeles, United States  
nikolaid@usc.edu

Jiaoyang Li  
Carnegie Mellon University  
Pittsburgh, United States  
jiaoyangli@cmu.edu

## ABSTRACT

We use the Quality Diversity (QD) algorithm with Neural Cellular Automata (NCA) to automatically evaluate Multi-Agent Path Finding (MAPF) algorithms by generating diverse maps. Previously, researchers typically evaluate MAPF algorithms on a set of specific, human-designed maps at their initial stage of algorithm design. However, such fixed maps may not cover all scenarios, and algorithms may overfit to the small set of maps. To seek further improvements, systematic evaluations on a diverse suite of maps are needed. In this work, we propose **Quality-Diversity Multi-Agent Path Finding Performance EvaluatoR** (QD-MAPPER), a general framework that takes advantage of the QD algorithm to comprehensively understand the performance of MAPF algorithms by generating maps with patterns, be able to make fair comparisons between two MAPF algorithms, providing further information on the selection between two algorithms and on the design of the algorithms. Empirically, we employ this technique to evaluate and compare the behavior of different types of MAPF algorithms, including search-based, priority-based, rule-based, and learning-based algorithms. Through both single-algorithm experiments and comparisons between algorithms, researchers can identify patterns that each MAPF algorithm excels and detect disparities in runtime or success rates between different algorithms. Our generated maps are available at <https://airtclick.github.io/qdmapper>. Our code is available at <https://github.com/AirTClick/QD-MAPPER>.

## KEYWORDS

Multi-Agent Path Finding; Procedural Content Generation; Quality Diversity Optimization

### ACM Reference Format:

Cheng Qian, Yulun Zhang, Varun Bhatt, Matthew C. Fontaine, Stefanos Nikolaidis, and Jiaoyang Li. 2026. QD-MAPPER: A Quality Diversity Framework to Automatically Evaluate Multi-Agent Path Finding Algorithms in

\*The full arXiv version: <https://arxiv.org/abs/2409.06888>.

<sup>†</sup>Equal contribution.

This work is licensed under a Creative Commons Attribution International 4.0 License.

Diverse Maps. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/UZVN6052>

## 1 INTRODUCTION

We study the problem of evaluating the performance of Multi-Agent Path Finding (MAPF) algorithms by generating diverse maps. Given a map and a group of agents, MAPF is the problem of finding collision-free paths from their start to goal locations. MAPF has wide applications in coordinating hundreds of robots in automated warehouses [50, 51], moving characters in video games [25], managing drone traffic [8], and controlling multi-robotic arms [33].

Given the wide applicability of MAPF, many algorithms have been proposed to solve MAPF, and evaluating these algorithms becomes an important task. At the beginning of the algorithm design, researchers prefer to use benchmark maps to have an overview on the performance of their algorithms. Stern et al. [40] have proposed a set of 33 MAPF benchmark maps, covering a diverse spectrum of map sizes, layouts, and difficulties. These human-designed maps are widely used to test newly proposed MAPF algorithms [23, 29, 38]. However, to further discover the pros and cons and improve algorithms, such fixed maps are not sufficient.

Meanwhile, Quality Diversity (QD) algorithms have been used to generate a diverse set of high-quality solutions by optimizing a given objective function and a set of diversity measure functions. A recent work [51] has used QD algorithms to optimize layouts for automated warehouses. To ensure that the optimized layouts possess human-explainable and regularized patterns, a follow-up work [50] then leverages Neural Cellular Automata (NCA), a CNN-based neural generator, to generate the layouts and uses QD algorithms to optimize the parameters of the NCA. Cellular Automata [18] iteratively generates complex cell-based structures from a simple one through local interaction between cells. Each cell decides its next state based on its neighbors via a fixed rule. NCA then represents the rules using a convolutional neural network that can be optimized.

In this paper, we adapt the layout optimization approach from the previous work [50], proposing **Quality-Diversity Multi-Agent Path Finding Performance EvaluatoR** (QD-MAPPER), a general framework with the goal of evaluating MAPF algorithms by generating diverse maps. To demonstrate that our approach can evaluate a broad spectrum of MAPF algorithms, we present evaluation results

on 6 representative MAPF algorithms, namely CBS [34], EECBS [23], PBS [24], LaCAM3 [28], PIBT [29], and Learn-to-Follow (LTF) [38]. We present a number of interesting results undiscovered in prior works. For example, we generate new maps with patterns that LaCAM3 runs out of time, while it was reported to solve 99% of the MAPF benchmarks [40]. We also expose the incompleteness of PBS by presenting maps in which PBS fails to find any solution instead of running out of time on benchmark maps.

We make the following contributions: (1) we adapt the layout optimization approach to propose QD-MAPPER to evaluate MAPF algorithms by generating diverse maps, (2) we apply QD-MAPPER on 6 representative MAPF algorithms of different categories, finding new challenging maps for MAPF algorithms, and providing new insights about their performance on maps of different patterns, and (3) we propose two concrete realizations of QD-MAPPER that are useful for evaluating MAPF algorithms: *one-algorithm* realization and *two-algorithm* realization. As a generic framework, researchers can extend other realizations, such as generating easy maps or comparing more than two algorithms.

## 2 PRELIMINARIES

### 2.1 Multi-Agent Path Finding (MAPF)

**DEFINITION 1 (MAP).** *A map is a four-neighbored 2D grid, where each tile can either be an empty space or an obstacle. A map is valid if all empty spaces are connected.*

**DEFINITION 2 (MAPF).** *Given a valid map and a set of agents with their start and goal locations, MAPF aims to find collision-free paths from their start to goal locations. Agents can either move to their adjacent locations or stay at their current locations at each timestep. Two agents collide if they are at the same location or swap locations at the same timestep. The objective of MAPF is minimizing the sum-of-cost, defined as the sum of travel time of all agents.*

We present a short summary of existing MAPF algorithms and the chosen algorithms for experiments in this paper.

**Search-Based.** This category includes exponential-time algorithms that exhaustively explore the solution space of MAPF. They usually have theoretical guarantees such as optimality and bounded suboptimality, but suffer from long computational times. Examples include  $M^*$  [44], BCP [22], ICTS [35], CBS [34], ECBS [3], and EECBS [23]. We choose CBS [34] and EECBS [23] as the representatives of this category. CBS is an optimal algorithm. It starts by planning a shortest path for each agent, ignoring the collisions, and then resolving the collisions with a two-level search. The high-level search iteratively selects unresolved collisions and adds constraints to tackle them. The low-level search runs single-agent planning to compute new paths that satisfy the constraints. The search continues until all collisions are resolved. EECBS is a bounded-suboptimal variant of CBS that uses Explicit Estimation Search [42] on the high level and focal search [30] on the low level.

**Priority-Based.** Priority-based algorithms plan paths for each agent following a priority order, forcing agents with lower priorities to avoid colliding with those with higher ones. The priority planning (PP) [11, 39] algorithm plans paths for agents in a pre-defined priority order. Monte-Carlo PP [17] pre-defines a set of randomized

priority orders and selects the one with the best paths. Representative algorithms include priority planning (PP) [11], Monte-Carlo PP [17], and PBS [24]. We choose PBS as the representative of this category. PBS combines CBS with PP to explicitly search for a good priority order with a two-level search. The high-level search of PBS is similar to CBS except that PBS constrains one agent to have a higher priority order than the other. The low-level search then plans paths with the order from the high-level search.

**Rule-Based.** Rule-based algorithms [45] leverage pre-defined rules to move agents to their goals. They usually run much faster, but produce worse solutions than search-based and priority-based methods. We choose PIBT [29] and LaCAM3 [28] as representatives. PIBT uses an iterative one-timestep rule to move agents. At each timestep, each agent plans a single-step action towards its goal. In case of collisions, a pre-defined rule is applied to resolve the collision. LaCAM3 is an anytime algorithm that combines PIBT and search-based methods. As a hybrid algorithm that combines PIBT with search-based methods, It performs the search by sequentially expanding search nodes, and when it encounters the goal configuration, it derives the solution by backtracking parent pointers.

**Learning-Based.** Learning-based algorithms either formulate MAPF as a Multi-Agent Reinforcement Learning (MARL) [32] problem or leverage Imitation Learning (IL) to imitate the behavior of an oracle MAPF algorithm. MARL-based methods [32, 38, 48] typically train a shared policy for each agent, taking its local field of view as input and deciding its next action. IL-based methods, on the other hand, attempt to imitate an oracle algorithm. For example, MAPF-GPT [2] tokenizes MAPF instances and trains a transformer to imitate LaCAM3 [28]. SILLM [20] trains a learnable PIBT by imitating the behavior of WPPL [21]. We choose Learn to Follow (LTF) [38] as the representative of the learning-based algorithms. LTF starts by searching for a guide path for each agent without considering collisions. Then it uses a shared learned policy to move the agents to their goals along the guide paths while avoiding collisions. LTF is developed for lifelong MAPF, a variant of MAPF that constantly assigns new goals to agents. We modify it for MAPF by asking agents to stop when they reach their first goals.

### 2.2 MAPF Algorithm Evaluation

In early works, researchers typically used randomly generated benchmark maps [44] or maps from the single-agent path planning benchmark set [41] to evaluate new MAPF algorithms. Stern et al. [40] established the major set of MAPF benchmark maps with 33 maps, which are used extensively in MAPF research [7, 23, 53]. These maps can provide an overview of the performance of MAPF algorithms at the initial stage. However, to further discover the pros and cons of the algorithm, the benchmark set is insufficient with its fixed patterns and limited size. For example, LaCAM3, despite prior claims of solving 99% of standard benchmarks, also struggles under our generated maps with long corridors and narrow entries, which are patterns underrepresented in existing benchmarks.

To the best of our knowledge, very few works use automatic generation methods to evaluate the performance of MAPF algorithms. The POGEMA platform [37] leverages procedural content generation [9] to train and evaluate MARL agents, but it does not target specific MAPF algorithms. One recent work [31] also uses

QD to generate maps for MAPF. However, it has two major differences with our work. First, they focus on generating maps of fixed difficulties, quantified by an approximate metric based on map connectivity, while we are interested in evaluating the performance of specific MAPF algorithms by generating diverse maps with different difficulties. Second, they directly optimize tile types (i.e., whether it is an empty space or an obstacle), while we optimize a map generator based on NCA, which has been shown to be more effective for generating maps with diverse regularized patterns [50].

### 2.3 Quality Diversity (QD) Algorithms and Automatic Scenario Generation

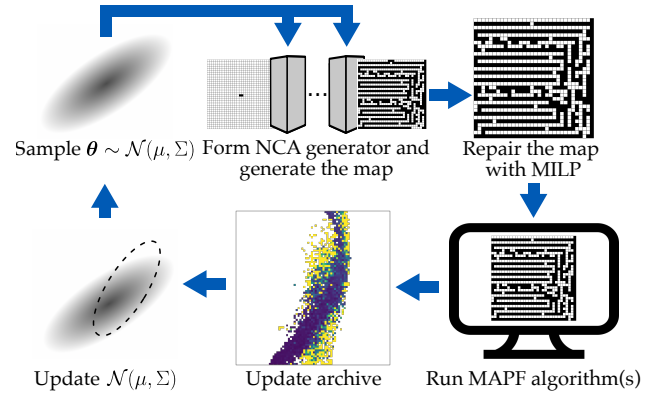
QD algorithms [26] are inspired by evolutionary algorithms with diversity search to generate a diverse collection of high-quality solutions by optimizing an objective function and diversifying a set of diversity measure functions. QD algorithms maintain an *archive*, which is a tessellated measure space defined by the measure functions. The archive stores the best solution in each tessellated cell. QD algorithms then optimize the sum of objective values of all solutions in the archive, defined as QD-score. We choose Covariance Matrix Adaptation MAP-Annealing (CMA-MAE) [13] as the method to generate maps because it is the state-of-the-art QD algorithm specialized for continuous search domains. CMA-MAE is an extension of MAP-Elites [26] that incorporates the covariance matrix adaptation mechanism of CMA-ES [19], which is a derivative-free single-objective optimizer. CMA-ES maintains a multi-variate Gaussian distribution and iteratively samples from it for new solutions. It evaluates the solutions and updates the Gaussian toward high-objective regions. CMA-MAE adapts this mechanism to optimize the QD-score. QD algorithms are widely applied in automatic scenario generation scenarios. In multi-robot systems, Zhang et al. [50] optimize NCA via QD algorithms to generate arbitrarily large warehouse layouts. In autonomous driving, researchers generate scenarios to evaluate developed autonomous driving systems [1, 27]. In human-robot interaction, QD algorithms are used to generate shared autonomy scenarios [16] and diverse kitchen layouts [14] to study the coordination behavior between humans and robots. A follow-up work [4] leverages model-based QD methods [5, 52] to improve the sample efficiency of QD algorithms. In reinforcement learning, prior works have generated maps to benchmark [9] or continuously improve [46] trained agents.

## 3 MAP GENERATION APPROACH

### 3.1 Overview

Figure 1 shows the pipeline of QD-MAPPER. We adapt previous works to use CMA-MAE [13] to search for and update a diverse collection of NCA [50] generators to generate diverse maps with the objective and measures computed by running MAPF algorithms. We start by sampling a batch of  $b$  parameter vectors  $\theta$  from a multi-variate Gaussian distribution, forming  $b$  NCA generators. Starting from a fixed seed map, each NCA generator iteratively updates the input map to generate a map with complex local patterns, leading to  $b$  maps. Appendix A exemplifies the process of generating a map.

Maps generated by NCA might not be valid. We then adapt a Mixed Integer Linear Programming (MILP) solver [49, 51] to repair the map to enforce connectivity and domain-specific constraints



**Figure 1: Overview of QD-MAPPER: (1) sample parameter vectors  $\theta$  from multi-variate Gaussian of CMA-MAE, (2) form NCA generators, (3) use a Mixed Integer Linear Programming (MILP) solver to ensure map connectivity and the correct number of obstacles, (4) run MAPF algorithms, (5) update the archive, and (6) update multi-variate Gaussian.**

while making minimal modifications to the generated map. We then evaluate repaired maps with MAPF algorithms to calculate the average objectives and measure values, and add the evaluated maps to an archive. For each evaluation, we run the given MAPF algorithm in  $N_e$  instances. Finally, we update the parameters of the multivariate Gaussian distribution, sampling a new batch of  $b$  parameter vectors and starting a new iteration. We repeat this process until we evaluate  $N_{eval}$  maps.

**NCA generator.** We define an NCA generator as a function  $g(s; \theta; C) : M \rightarrow M$ , where  $M$  is the space of possible maps,  $s \in M$  represents a fixed seed map, and  $C \in \mathbb{Z}^+$  is a hyperparameter controlling the number of iterations to run the NCA generator. The generator is parameterized by  $\theta \in \Theta$ , where each parameter vector  $\theta$  corresponds to a distinct NCA generator that then defines a distinct map  $m \in M$ . We then use CMA-MAE to optimize  $\theta$ , attempting to find the best NCA generator in each cell of the archive.

Our NCA generator is a convolutional neural network (CNN) following the architecture in prior work [50]. It consists of three convolutional layers with kernel size  $3 \times 3$ , each followed by either ReLU or sigmoid activations. We represent the 2D tile-based maps as 3D tensors, with each tile encoded as a one-hot vector, representing its tile type (i.e. obstacle or empty space). Since the output of the NCA has the same size as its input, we can repeatedly pass the output back into the network.

**MILP Solver.** Maps generated by NCA might not be valid (e.g., a wall partitioning two open sections). To ensure that the maps are valid, we formulate and solve a MILP [49, 51] to obtain a repaired map. The MILP constraints are set to be: (1) all empty spaces are connected, and (2) the number of obstacles falls within a pre-defined range  $[O_{lb}, O_{ub}]$ . We set the objective to minimize the Hamming distance between unrepaired and repaired maps to ensure that the patterns of the generated map are maintained as much as possible. **Generation of MAPF Instances From Maps.** To generate a MAPF instance based on a map, we follow the bucket method [40] to generate evenly distributed start and goal locations of the agents with a

distance constraint, which was used for generating the *even* scenarios in the MAPF benchmark. For each map, we generate  $N_e$  MAPF instances with different starts and goals, run MAPF algorithms on each instance, and compute the average objective and measures of all instances as the objective and measures of the map. Notably, our MAPF instance generation strategy serves only as an example. In applications where a different set of instances is needed, users can customize the instance generation method inside QD-MAPPER.

**Framework Realization.** Our proposed framework is flexible at a high level. In this paper, we propose two concrete realizations of QD-MAPPER that are useful for evaluating MAPF algorithms with carefully designed objectives and measures. We first provide a *one-algorithm* realization, intending to evaluate five representative MAPF algorithms by generating diverse and difficult maps for each of them. We then provide a *two-algorithm* realization, aiming to compare the performance of two algorithms by generating maps that are easy for one algorithm and hard for the other. We highlight that other realizations, such as generating easy maps or comparing more than two algorithms, are possible.

### 3.2 Objective

In all experiments, our objective is a function  $f : X \rightarrow \mathbb{R}$ , where  $X$  is the space of all maps. The function  $f$  runs one or two MAPF algorithms on  $N_e$  different MAPF instances with  $N_a$  agents. For CBS, EECBS, LaCAM3, and PBS, we set a time limit  $T$  for each run. For PIBT and LTF, we set a maximum makespan of  $M$  timesteps.

**One-Algorithm Experiments.** We intend to generate maps that are challenging for a given MAPF algorithm  $\phi \in A$ , where  $A = \{\text{CBS, EECBS, LaCAM3, PBS, PIBT, LTF}\}$  in our experiments. The objective quantifies the empirical hardness of a map  $x \in X$  to  $\phi$ .

For CBS, EECBS, and PBS, the objective is to maximize their *CPU runtime* because they are exponential-time algorithms, the major limitation being frequently running out of time. We do not consider the solution quality as the objective because CBS and EECBS are proven to be optimal and bounded-suboptimal algorithms. PBS, while being unbounded suboptimal, empirically finds near-optimal solutions. For example, Ma et al. [24] shows that PBS never finds solutions over 4% worse than optimal on game maps from the MAPF benchmark. We include LaCAM3 because it can solve almost all instances in the MAPF benchmark. Our objective is to identify more challenging maps for it. Therefore, we maximize *CPU runtime* to find failure cases for CBS, EECBS, PBS, and LaCAM3. Concretely, suppose that the function  $t_\phi : X \rightarrow \mathbb{R}_{>0}$  computes the average CPU runtime by running  $\phi \in A$  on  $N_e$  instances, the objective  $f$  for the algorithm  $\phi$  is  $f(x) = t_\phi(x)$ ,  $\phi \in \{\text{CBS, EECBS, PBS, LaCAM3}\}$ . If  $\phi$  cannot find solutions within the time limit  $T$ , we set the CPU runtime to  $T$ .

PIBT and LTF, on the other hand, suffer mainly from deadlocks, meaning that some agents might never reach their goals. Therefore, to find failure cases, we minimize the *Regularized Success Rate (RSR)*, defined as follows:

$$RSR_\phi(x) = \begin{cases} SR_\phi(x) & \text{if } SR < 1 \\ SR_\phi(x) \cdot C - SoC_\phi(x) & \text{if } SR = 1 \end{cases} \quad (1)$$

$SR_\phi : X \rightarrow \mathbb{R}_{>0}$  computes the success rate of running algorithm  $\phi$  on one MAPF instance, which is computed as the percentage of

agents that successfully reach their goals within the given makespan  $M$ .  $SoC_\phi : X \rightarrow \mathbb{R}_{>0}$  computes the sum-of-cost of the solution if  $SR = 1$ .  $C$  is a large constant making sure that the RSR score of maps with  $SR = 1$  always dominates that of maps with  $SR < 1$ . Intuitively, if not all agents reach their goals ( $SR < 1$ ), we maximize the success rate. If all agents reach their goals ( $SR = 1$ ), we regularize the success rate with the sum-of-cost to quantify the quality of the solution. Let the function  $r_\phi : X \rightarrow \mathbb{R}_{>0}$  return the average regularized success rate by running algorithm  $\phi$  in  $N_e$  MAPF instances. Concretely,  $r_\phi(x) = \sum_{i=1}^{N_e} RSR_\phi^{(i)}(x)$ , where  $i$  denotes the  $i$ -th run of algorithm  $\phi$ . Then the objective of PIBT and LTF are  $f(x) = r_{\text{PIBT}}(x)$  and  $f(x) = r_{\text{LTF}}(x)$ , respectively. We do not consider runtime as the objective of PIBT and LTF because they are both one-step planners and thus their runtime is strongly correlated with success rates and solution quality.

**Two-Algorithm Experiments.** To compare two given MAPF algorithms, we aim to generate maps that maximize the performance gap between the them. We consider two pairs of comparisons: (1) EECBS vs. PBS and (2) PIBT vs. LTF. We compare EECBS and PBS because (1) both suboptimal algorithms use a two-level search based on collision avoidance, (2) both find near-optimal solutions empirically, and (3) few prior works have compared them thoroughly. We compare PIBT and LTF because (1) both use a pre-defined policy, either rule-based or learned, to move the agents step by step to their goals while avoiding collisions, (2) both run fast but suffer from deadlocks, and (3) while the LTF work [38] has compared them on warehouse maps, showing that LTF outperforms PIBT, we want to find cases where PIBT outperforms LTF. To compare EECBS and PBS, we set the objective as the absolute difference in the average CPU runtime. Concretely, the objective of a generated map  $x \in X$  is computed as  $f(x) = |t_{\text{EECBS}}(x) - t_{\text{PBS}}(x)|$ . Similarly, to compare PIBT and LTF, we set the objective as the absolute difference in the regularized success rate, i.e.,  $f(x) = |r_{\text{PIBT}}(x) - r_{\text{LTF}}(x)|$ .

### 3.3 Diversity Measures

We select diversity measures such that the generated maps are diverse in (1) general hardness towards all algorithms to analyze whether the general hardness of a map aligns with the hardness to a specific algorithm, and (2) spatial arrangement of the obstacles to generate maps of different patterns.

**General Hardness.** To measure the general hardness of a map, the most commonly used metric is the number of obstacles. The number of obstacles in a map affects the space available to resolve collisions as well as the search space of the algorithms. Therefore, diversifying it would give a range of difficulties. We also consider two metrics that are proposed specifically to reflect the hardness of a map: the standard deviation of Betweenness Centrality (Std of BC) [12] and  $\lambda_2$  [31].

The Betweenness Centrality of an empty space is the fraction of all possible shortest paths in the map that pass through it. We use the Std of BC here to measure the congestion in the map. The second smallest eigenvalue (known as  $\lambda_2$ ) of the normalized Laplacian matrix of different maps is also claimed to be correlated to the difficulty of the maps [31].

**Spatial Arrangement.** In addition to the hardness of the maps, we want to diversify the spatial arrangement of the maps, making




Figure 2: Examples of one-entry space and one-tile entry.

the generated maps stylistically diverse. We first consider the KL divergence of the tile pattern distribution [15] between the tile pattern distribution of the generated map and a set of maps selected from the MAPF benchmark [40].

A *tile pattern* is one possible arrangement of obstacles and empty spaces in a  $3 \times 3$  grid. To compute the tile pattern distribution of a map, we count the number of possible tile patterns in the map. We then compute the KL divergence between the tile pattern distribution of the map and a fixed distribution calculated from all maze maps<sup>1</sup> of the single-agent path finding benchmark [41]. The smaller the KL divergence, the more similar in local patterns our generated map is to the selected set of benchmark maps. We choose maze maps to compare the KL divergence with because they empirically yield the most complex patterns.

In addition, we consider the entropy of the tile pattern distribution [50]. The entropy of the tile pattern distribution is a diversity measure to indicate the level of pattern regularization. The lower the entropy, the more regularized patterns the map possesses.

Finally, we consider the KL divergence of the Weisfeiler-Lehman (WL) graph feature [36] between the generated map and the maze maps. The WL graph feature indicates the similarity between graphs. Similar to the tile pattern distribution, a smaller KL divergence of the WL graph feature implies more similarities between the local patterns of the generated map and the maze maps.

In our experiments, we choose **the number of obstacles and KL divergence of tile pattern distribution** as the diversity measures. We empirically observe that, among the measures considered, the number of obstacles is the most correlated with the general hardness of the maps, and the KL divergence of tile pattern distribution yields the most diverse map patterns. We justify our choice with experimental results in Appendix B.

## 4 ONE ALGORITHM ANALYSIS

For one-algorithm experiments, we intend to generate diverse difficult maps for CBS, EECBS, PBS, LaCAM3, PIBT, and LTF.

### 4.1 Experiment Setup

**Hyperparameters.** Defining  $w$  as the suboptimality bound away from the optimal solution, we use  $w = 1.5$  for EECBS. For CBS, we use the implementation of EECBS with  $w = 1$ . We set  $T = 20$  seconds for CBS, EECBS, LaCAM3, and PBS,  $M = 1000$  for PIBT, and  $M = 512$  for LTF. For all algorithms, we fix the map size as  $32 \times 32$ . The range for the number of obstacles is  $[O_{lb} = 307, O_{ub} = 717]$ , which corresponds to 30% to 70% of the map size. For each evaluation, we run the given MAPF algorithm in  $N_e = 5$  different instances, each with  $N_a$  agents. For CBS, EECBS, and PBS,  $N_a =$

<sup>1</sup><https://movingai.com/benchmarks/grids.html>

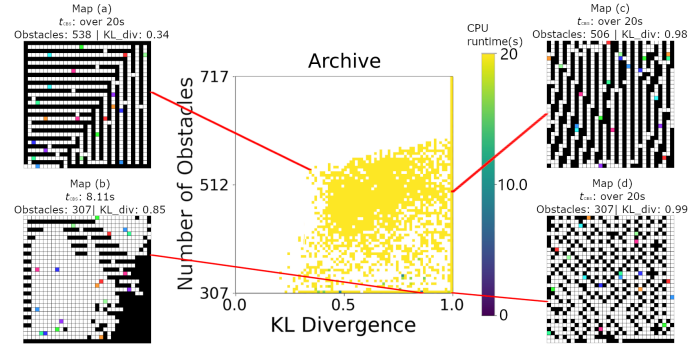


Figure 3: Archive for CBS with sample maps.

50, for LaCAM3,  $N_a = 100$ , and for PIBT and LTF,  $N_a = 150$ . We use more agents for LaCAM3, PIBT, and LTF because they empirically have higher success rates in congested MAPF instances. For all algorithms, we evaluated  $N_{eval} = 10,000$  maps. We present computational costs and resources in Appendix C.

**Implementation.** We implement CMA-MAE with Pyribs [43]. For MAPF algorithms, we use the open source implementations of CBS and EECBS<sup>2</sup>, PBS<sup>3</sup>, LaCAM3<sup>4</sup>, PIBT<sup>5</sup>, and LTF<sup>6</sup>.

### 4.2 Results

**Generated Maps.** We first define *one-entry space* and *one-tile entry* (shown in Figure 2) as features of the maps. A one-entry space represents the empty space surrounded by three obstacles. A one-tile entry represents a single empty space separating two corridors. In Figures 3 to 5, the yellow and dark blue cells in the archives indicate maps with high and low CPU runtime, respectively. In Figures 7 and 8, the yellow and dark blue cells indicate maps with low and high success rates, respectively. The colored tiles in the maps display 10 out of 150 pairs of start and goal locations. While we are primarily interested in maps that are challenging for the algorithms, we also present easy maps as the by-products of the experiments. We summarize the typical patterns of the maps and show more maps with similar patterns in Appendix D.

**CBS.** Figure 3 shows the archive of maps for CBS. It is easy to generate challenging maps for CBS. Map (a), Map (c), and Map (d) show the most typical hard maps for CBS. Map (a) and Map (c) contain many long corridors, and Map (d) contains many one-entry spaces with long corridors.

Some maps that CBS can solve are maps with fewer obstacles and a large chunk of empty spaces. Map (b) is one of these typical maps. Map (b) contains a large chunk of empty spaces and short corridor components in between, providing more space for agents to avoid collisions.

**EECBS.** Figure 4 shows the archive of maps for EECBS. Compared to the archive of CBS, it is harder to generate challenging maps for EECBS. Map (b) and Map (c) are hard maps with typical patterns

<sup>2</sup><https://github.com/Jiaoyang-Li/EECBS>

<sup>3</sup><https://github.com/Jiaoyang-Li/PBS>

<sup>4</sup><https://github.com/Kei18/lacam3>

<sup>5</sup><https://github.com/Kei18/piibt2>

<sup>6</sup><https://github.com/Cognitive-AI-Systems/learn-to-follow>




Figure 4: Archive for EECBS with sample maps.




Figure 6: Archive for LaCAM3 with sample maps.

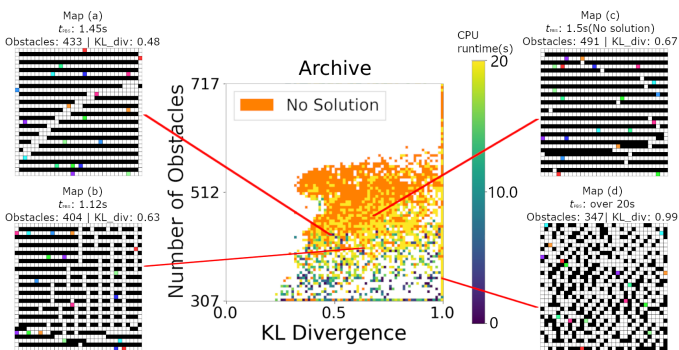


Figure 5: Archive for PBS with sample maps.

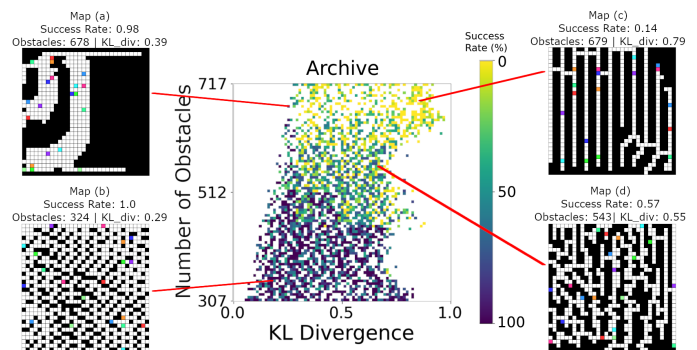


Figure 7: Archive for PIBT with sample maps.

for EECBS. Map (b) contains a large chunk of empty space, but with many long corridors and one-tile entries. Map (c) contains many long corridors and one-tile entries. EECBS runs out of time on both maps.

On the other hand, we observe that EECBS performs well in maps with more empty spaces between each long obstacle component and maps with short obstacle components, where obstacle components refer to clusters of two or more obstacles. Map (a) and Map (d) are two maps with typical patterns. Map (a) contains many long obstacle components with one-tile entries; however, the map has two or more columns of empty space between each long obstacle component. Map (d) contains many short obstacle components with many entries in between, providing more space for agents to avoid collision. Other maps where EECBS performs well are of similar patterns or have a large chunk of empty spaces. We show more maps with similar patterns in Appendix D.

**PBS.** Figure 5 shows the archive of maps for PBS. We empirically discover that PBS returns no solutions in many maps. Therefore, we show the maps in which PBS returns no solution in at least one out of five MAPF instances during the evaluation in the archive as orange cells. Map (c) and Map (d) are hard maps with typical patterns for PBS. Map (c) contains many long corridors with one-tile entries in between, and Map (d) contains many one-entry spaces. PBS always reports no solution on Map (c) and runs out of time on Map (d). On the other hand, Map (a) and (b) show two typical patterns where PBS performs well. Map (a) contains long corridors

but with more entry spaces in between. Map (b) contains many short obstacle components, providing more space for agents to avoid collisions. To validate our observation, we run PBS in 200 MAPF instances on Map (a). The results in Table 2 of Appendix E indicate that PBS can effectively solve maps with long corridors with increased entry spaces in between.

**LaCAM3.** Figure 6 shows the archive of maps for LaCAM3. Notably, the previous work [28] claims that LaCAM3 solves 99% of the maps in the MAPF benchmark [40], while we find a number of challenging maps for the algorithm. Map (a), Map (b), and Map (c) are hard maps with typical patterns for LaCAM3. Map (a) and Map (b) contain many long corridors with one entry. Map (c) contains a large chunk of empty space, but with many long corridors and one-tile entries. LaCAM3 runs out of time on all three maps. On the other hand, we observe that LaCAM3 performs well in maps with more empty spaces between each long obstacle component, such as Map (d).

**PIBT.** Figure 7 shows the archive of maps of PIBT. The success rate of PIBT decreases with an increased number of obstacles and KL divergence. Similar to EECBS and CBS, PIBT performs poorly in maps with long corridors, such as Map (c). With a similar number of obstacles and smaller KL divergence, we can find easy maps for PIBT. They have large chunks of empty spaces with few corridors. Map (a) is an example. In the middle of the archive, we observe maps with success rates of around 50%. They usually have long corridors and one-entry spaces. Map (d) is an example. In the bottom left corner, we find Map (b), which has many one-entry spaces but




Figure 8: Archive for LTF with sample maps.




Figure 10: Archive for EECBS vs PBS with sample maps.

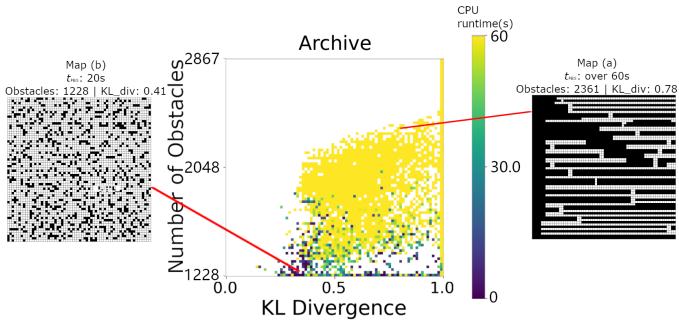


Figure 9: Archive for PBS with sample maps of size  $64 \times 64$ .




Figure 11: Archive for PIBT vs LTF with sample maps.

fewer obstacles, achieving a success rate of 100%. Overall, PIBT can efficiently solve maps that include long corridors and one-entry spaces with sufficient space.

**LTF.** Figure 8 shows the archive of maps for LTF. Similar to PIBT, LTF performs worse with more obstacles and higher KL divergence, as shown in Map (a). Similar to challenging maps for CBS, EECBS, and PIBT, Map (a) contains long corridors. With fewer obstacles, LTF performs poorly on maps with many one-entry spaces, such as Map (b) and Map (d). These two maps contain different numbers of obstacles, but they have similar patterns of one-entry spaces. On the other hand, LTF performs better with fewer obstacles with large chunks of empty spaces, such as Map (c).

**Validation Results.** We only evaluate each map by running  $N_e = 5$  MAPF instances during the QD search. To further validate our observations on the hardness of the maps, we additionally present *validation results* in Appendix E by running 200 MAPF instances on the selected maps.

**On the Generation of Larger Maps.** To demonstrate that QD-MAPPER can evaluate MAPF algorithms by generating maps larger than  $32 \times 32$ , we further extend our evaluation by generating maps with a size of  $64 \times 64$  using PBS and PIBT. Figure 9 shows the results of PBS. The results of PIBT are shown in Appendix F. Both archives have a similar distribution compared to experiments of generating benchmark maps with a size of  $32 \times 32$ .

## 5 COMPARING TWO ALGORITHMS

Given the different categories and properties of the MAPF algorithms, hard maps for one algorithm might not be hard for the other. Therefore, we aim to generate maps that are easy for one algorithm and hard for the other by maximizing the performance gap between them.

### 5.1 Experiment Setup

We run EECBS with  $w = 1.5$ . For EECBS and PBS, we use  $T = 20$  seconds. For PIBT and LTF, we use  $M = 512$ , which is used to train the LTF policy [38]. We use the same  $O_{lb}$ ,  $O_{ub}$ ,  $N_e$ ,  $N_{eval}$ ,  $N_a$ , and map sizes as the one-algorithm experiments. The compute resources are specified in Appendix C.

### 5.2 Results

**Generated Maps.** We show the archives and representative maps in Figures 10 and 11. To better visualize the comparisons of the algorithms, we plot the differences, instead of the absolute differences, between the CPU runtime or success rate of the algorithms. In Figure 10, the yellow and dark blue cells represent EECBS having lower and higher CPU runtime than PBS, respectively. In Figure 11, the yellow and blue cells represent PIBT having higher and lower success rates than LTF, respectively.

**EECBS vs. PBS.** Figure 10 shows the archive comparing EECBS and PBS. In most cases, EECBS and PBS have similar performances.

**Table 1: Validation of observation by running maps with 200 different instances. Left: average CPU runtime  $T$  and success rate  $SR$  for EECBS and PBS. Right: success rate  $SR$  for PIBT and LTF.**

EECBS vs. PBS				PIBT vs. LTF			
Map	$T_{EECBS}$	$SR_{EECBS}$	$T_{PBS}$	$SR_{PBS}$	Map	$SR_{PIBT}$	$SR_{LTF}$
Map (a)	4.3s	95.5%	12.5s	69.5%	Map (a)	70%	56%
Map (b)	9.8s	56.5%	2.0s	96.5%	Map (b)	80%	96%
Map (c)	18.3s	14.5%	13.7s	43.0%	Map (c)	52%	80%
Map (d)	1.1s	100.0%	7.4s	77.5%	Map (d)	96%	96%

Maps with more than 512 obstacles are usually too hard for both algorithms. With fewer than 400 obstacles, however, EECBS generally performs better than PBS in two types of maps, exemplified in Map (a) and (d). Map (a) is similar to Map (a) of Figure 4, with long but wide corridors and one-entry spaces between adjacent corridors. Map (d) shows a random pattern with many one-entry spaces. Meanwhile, Map (b) and (c) are examples where PBS outperforms EECBS. Notably, while Map (c) contains narrow long corridors, both ends of the corridors are open, potentially making PBS resolve conflicts easier. To validate our observation, we run 200 instances on maps shown in Figure 10 and show the result in Table 1. Our results serve as a more comprehensive comparison between EECBS and PBS. Notably, only one work [7] has systematically compared EECBS and PBS. They conclude that EECBS outperforms PBS in five out of six maps selected from the MAPF benchmark [40]. In contrast, our results indicate that EECBS and PBS exhibit similar behavior on most maps, with each having distinct advantages on different map patterns.

**PIBT vs. LTF.** Figure 11 shows the archive that compares PIBT and LTF. In most cases with less than 50% of obstacles, PIBT and LTF have a similar success rate. In maps with large chunks of empty space such as Map (d), both PIBT and LTF perform well. With more than 50% of obstacles, PIBT generally has a 5% to 10% higher success rate on average than LTF. In most cases, these maps contain many one-entry spaces, such as Map (a). Upon running 200 instances on Map (a) with PIBT and LTF, PIBT outperforms LTF by an average success rate of 14%, as shown in Table 1. The result aligns with our observations in Section 4. Meanwhile, the maps in which LTF outperforms PIBT are concentrated in the upper right and bottom left, such as Map (b) and (c). Upon running 200 instances on both maps, LTF has an average of 20% higher success rate on these maps compared to PIBT. Our results provide additional comparisons between PIBT and LTF to the experiments conducted by Skrynnik et al. [38], which conclude that LTF outperforms PIBT by comparing them on a set of human-designed warehouse maps.

## 6 DISCUSSION AND CONCLUSION

We propose QD-MAPPER based on QD algorithms to systematically evaluate the performance of MAPF algorithms by generating diverse maps. We provide two concrete realizations of QD-MAPPER, presenting experimental results with diverse maps of different levels of hardness and patterns. For the one-algorithm experiments, we explore the limitations of five representative MAPF algorithms by generating hard maps by either maximizing the CPU runtime

or minimizing the regularized success rate. For the two-algorithm experiments, we compare two algorithms by generating diverse maps that are hard for one algorithm and easy for the other.

**Main Observations.** In one-algorithm experiments, PBS frequently fails to return any solution, not just timing out, on maps with long corridors and sparse or one-entry points, exposing a form of incompleteness that standard benchmark maps have failed to reveal. LaCAM3, despite prior claims of solving 99% of standard benchmarks [28], also struggles under our generated maps with long corridors and narrow entries, which are patterns underrepresented in existing benchmarks. From PIBT-LTF comparison, as obstacle density rises above 50%, PIBT typically outperforms LTF by 5–10%, particularly on maps rich in one-entry spaces. Meanwhile, LTF holds a performance advantage on maps with over 60% obstacles and high complexity. These findings extend previous evaluations [38] with LTF and PIBT on a set of human-designed warehouse maps.

### When and How to Use QD-MAPPER to Evaluate Performance.

While developing new MAPF algorithms, researchers can use our one-algorithm pipeline to evaluate their performance by finding maps their algorithm has difficulty solving. In this case, it helps reveal the weakness of new algorithms. Researchers can also leverage our two-algorithm pipeline to compare their proposed algorithm with an existing algorithm by identifying map types where their algorithm excels or struggles. We argue that this is a more systematic way of comparing two algorithms than comparing on a fixed set of benchmark maps because researchers might intentionally or unintentionally cherry-pick maps that favor their algorithm or overfit the design of their algorithms to the selected set of maps.

QD-MAPPER is a modular framework that separates instance generation, algorithm evaluation, and quality-diversity optimization. It allows users to customize in multiple parts. Researchers can tailor QD-MAPPER for alternative purposes by designing different objectives and diversity measures. For example, researchers can *minimize* rather than maximize CPU runtime of CBS, EECBS, and PBS with a larger number of agents to generate maps that are *easy*, instead of *hard*, for these algorithms. If researchers are interested in both hard and easy maps, using the CPU runtime or regularized success rate as diversity measures is a better choice. Researchers can also generalize the two-algorithm pipeline to an arbitrary number of MAPF algorithms to rank the performance of more than two algorithms. In addition, researchers can leverage learning-based diversity measures [10, 47] to evaluate MAPF algorithms on maps that meet criteria defined by humans. Future research can also extend QD-MAPPER by generating *instances* with varied numbers of agents, map sizes, as well as start and goal locations to provide a comprehensive evaluation of algorithm behavior under diverse configuration settings.

## ACKNOWLEDGMENTS

This work used Bridge-2 at Pittsburgh Supercomputing Center (PSC) [6] 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. This work was partially supported by the National Science Foundation under grant #2441629.

## REFERENCES

- [1] Yasasa Abeyirigoonawardena, Florian Shkurti, and Gregory Dudek. 2019. Generating Adversarial Driving Scenarios in High-Fidelity Simulators. In *ICRA*. 8271–8277.
- [2] Anton Andreychuk, Konstantin Yakovlev, Aleksandr Panov, and Alexey Skrynnik. 2025. MAPF-GPT: Imitation Learning for Multi-Agent Pathfinding at Scale. In *AAAI*. 23126–23134.
- [3] Max Barer, Guni Sharon, Roni Stern, and Ariel Felner. 2014. Suboptimal Variants of the Conflict-Based Search Algorithm for the Multi-Agent Pathfinding Problem. In *SoCS*. 19–27.
- [4] Varun Bhatt, Heramb Nemlekar, Matthew Christopher Fontaine, Bryon Tjanaka, Hejia Zhang, Ya-Chuan Hsu, and Stefanos Nikolaidis. 2023. Surrogate Assisted Generation of Human-Robot Interaction Scenarios. In *CoRL*. 513–539.
- [5] Varun Bhatt, Bryon Tjanaka, Matthew Fontaine, and Stefanos Nikolaidis. 2022. Deep Surrogate Assisted Generation of Environments. In *NeurIPS*. 37762–37777.
- [6] Shawn T. Brown, Paola Buitrago, Edward Hanna, Sergiu Sanielevici, Robin Scibek, and Nicholas A. Nystrom. 2021. Bridges-2: A Platform for Rapidly-Evolving and Data Intensive ResearchBridges-2: A Platform for Rapidly-Evolving and Data Intensive Research. In *PEARC*. 1–4.
- [7] Shao-Hung Chan, Roni Stern, Ariel Felner, and Sven Koenig. 2023. Greedy Priority-Based Search for Suboptimal Multi-Agent Path Finding. In *SoCS*. 11–19.
- [8] Shushman Choudhury, Kiril Solovey, Mykel Kochenderfer, and Marco Pavone. 2022. Coordinated Multi-Agent Pathfinding for Drones and Trucks over Road Networks. In *AAMAS*. 272–280.
- [9] Karl Cobbe, Christopher Hesse, Jacob Hilton, and John Schulman. 2020. Leveraging procedural generation to benchmark reinforcement learning. In *ICML*. 2048–2056.
- [10] Li Ding, Jenny Zhang, Jeff Clune, Lee Spector, and Joel Lehman. 2024. Quality Diversity Through Human Feedback: Towards Open-Ended Diversity-Driven Optimization. In *ICML*. 11072–11090.
- [11] Michael A. Erdmann and Tomás Lozano-Pérez. 1987. On Multiple Moving Objects. *Algorithmica* 2 (1987), 477–521.
- [12] Eric Ewing, Jingyao Ren, Dhvani Kansara, Vikraman Sathiyarayanan, and Nora Ayanian. 2022. Betweenness Centrality in Multi-Agent Path Finding. In *AAMAS*. 400–408.
- [13] Matthew Fontaine and Stefanos Nikolaidis. 2023. Covariance Matrix Adaptation MAP-Annealing. In *GECCO*. 456–465.
- [14] Matthew C. Fontaine, Ya-Chuan Hsu, Yulun Zhang, Bryon Tjanaka, and Stefanos Nikolaidis. 2021. On the Importance of Environments in Human-Robot Coordination. In *RSS*.
- [15] Matthew C. Fontaine, Ruilin Liu, Ahmed Khalifa, Jignesh Modi, Julian Togelius, Amy K. Hoover, and Stefanos Nikolaidis. 2021. Illuminating Mario Scenes in the Latent Space of a Generative Adversarial Network. In *AAAI*. 5922–5930.
- [16] Matthew C. Fontaine and Stefanos Nikolaidis. 2021. A Quality Diversity Approach to Automatically Generating Human-Robot Interaction Scenarios in Shared Autonomy. In *RSS*.
- [17] Paul Friedrich, Yulun Zhang, Michael Curry, Ludwig Dierks, Stephen McAleer, Jiaoyang Li, Tuomas Sandholm, and Sven Seuken. 2024. Scalable Mechanism Design for Multi-Agent Path Finding. In *IJCAI*. 58–66.
- [18] Martin Gardner. 1970. Mathematical Games – The Fantastic Combinations of John Conway’s New Solitaire Game “Life”. *Scientific American* 223, 4 (1970), 120–123.
- [19] Nikolaus Hansen. 2016. The CMA Evolution Strategy: A Tutorial. *ArXiv abs/1604.00772* (2016).
- [20] He Jiang, Yutong Wang, Rishi Veerapaneni, Tanishq Harish Duhan, Guillaume Adrien Sartoretti, and Jiaoyang Li. 2025. Deploying Ten Thousand Robots: Scalable Imitation Learning for Lifelong Multi-Agent Path Finding. In *ICRA*. 1–7.
- [21] He Jiang, Yulun Zhang, Rishi Veerapaneni, and Jiaoyang Li. 2024. Scaling Lifelong Multi-Agent Path Finding to More Realistic Settings: Research Challenges and Opportunities. In *SoCS*. 234–242.
- [22] Edward Lam, Pierre Le Bodic, Daniel Harabor, and Peter J. Stuckey. 2019. Branch-and-Cut-and-Price for Multi-Agent Pathfinding. In *IJCAI*. 1289–1296.
- [23] Jiaoyang Li, Wheeler Ruml, and Sven Koenig. 2021. EECBS: A Bounded-Suboptimal Search for Multi-Agent Path Finding. In *AAAI*. 12353–12362.
- [24] Hang Ma, Daniel Harabor, Peter J. Stuckey, Jiaoyang Li, and Sven Koenig. 2019. Searching with Consistent Prioritization for Multi-Agent Path Finding. In *AAAI*. 7643–7650.
- [25] Hang Ma, Jingxing Yang, Liron Cohen, T. K. Satish Kumar, and Sven Koenig. 2017. Feasibility Study: Moving Non-Homogeneous Teams in Congested Video Game Environments. In *AIIDE*. 270–272.
- [26] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating Search Spaces by Mapping Elites. *ArXiv abs/1504.04909* (2015).
- [27] Galen E. Mullins, Paul G. Stankiewicz, R. Chad Hawthorne, and Satyandra K. Gupta. 2018. Adaptive Generation of Challenging Scenarios for Testing and Evaluation of Autonomous Vehicles. *Journal of Systems and Software* 137 (2018), 197–215.
- [28] Keisuke Okumura. 2024. Engineering LaCAM\*: Towards Real-Time, Large-Scale, and Near-Optimal Multi-Agent Pathfinding. In *AAMAS*. 1501–1509.
- [29] Keisuke Okumura, Manao Machida, Xavier Défago, and Yasumasa Tamura. 2019. Priority Inheritance with Backtracking for Iterative Multi-agent Path Finding. In *IJCAI*. 535–542.
- [30] Judea Pearl and Jin H Kim. 1982. Studies in Semi-Admissible Heuristics. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 4, 4 (1982), 392–399.
- [31] Jingyao Ren, Eric Ewing, T. K. Satish Kumar, Sven Koenig, and Nora Ayanian. 2024. Map Connectivity and Empirical Hardness of Grid-based Multi-Agent Pathfinding Problem. In *ICAPS*. 484–488.
- [32] Guillaume Sartoretti, Justin Kerr, Yunfei Shi, Glenn Wagner, T. K. Satish Kumar, Sven Koenig, and Howie Choset. 2019. PRIMAL: Pathfinding via Reinforcement and Imitation Multi-Agent Learning. *IEEE Robotics and Automation Letters* 4, 3 (2019), 2378–2385.
- [33] Yorai Shaoul, Itamar Mishani, Maxim Likhachev, and Jiaoyang Li. 2024. Accelerating Search-Based Planning for Multi-Robot Manipulation by Leveraging Online-Generated Experiences. In *ICAPS*. 523–531.
- [34] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R. Sturtevant. 2015. Conflict-Based Search for Optimal Multi-Agent Pathfinding. *Artificial Intelligence* 219 (2015), 40–66.
- [35] Guni Sharon, Roni Stern, Meir Goldenberg, and Ariel Felner. 2013. The Increasing Cost Tree Search for Optimal Multi-Agent Pathfinding. *Artificial Intelligence* 195 (2013), 470–495.
- [36] Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M. Borgwardt. 2011. Weisfeiler-Lehman Graph Kernels. *Journal of Machine Learning Research* 12, 77 (2011), 2539–2561.
- [37] Alexey Skrynnik, Anton Andreychuk, Anatolii Borzilov, Alexander Chernyavskiy, Konstantin Yakovlev, and Aleksandr Panov. 2025. POGEMA: A Benchmark Platform for Cooperative Multi-Agent Pathfinding. In *ICLR*.
- [38] Alexey Skrynnik, Anton Andreychuk, Maria Nesterova, Konstantin Yakovlev, and Aleksandr Panov. 2024. Learn to Follow: Decentralized Lifelong Multi-Agent Pathfinding via Planning and Learning. In *AAAI*. 17541–17549.
- [39] Trevor Scott Standley. 2010. Finding Optimal Solutions to Cooperative Pathfinding Problems. In *AAAI*. 173–178.
- [40] 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. 2019. Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks. In *SoCS*. 151–159.
- [41] N. Sturtevant. 2012. Benchmarks for Grid-Based Pathfinding. *Transactions on Computational Intelligence and AI in Games* 4, 2 (2012), 144–148.
- [42] Jordan T. Thayer and Wheeler Ruml. 2011. Bounded Suboptimal Search: A Direct Approach Using Inadmissible Estimates. In *IJCAI*. 674–679.
- [43] Bryon Tjanaka, Matthew C. Fontaine, David H. Lee, Yulun Zhang, Nivedit Reddy Balam, Nathaniel Dennler, Sujay S. Garlanka, Nikitas Dimitri Klapsis, and Stefanos Nikolaidis. 2023. pyribs: A Bare-Bones Python Library for Quality Diversity Optimization. In *GECCO*. 220–229.
- [44] G. Wagner and H. Choset. 2015. Subdimensional Expansion for Multirobot Path Planning. *Artificial Intelligence* 219 (2015), 1–24.
- [45] K. Wang and A. Botea. 2011. MAPP: A Scalable Multi-Agent Path Planning Algorithm with Tractability and Completeness Guarantees. *Journal of Artificial Intelligence Research* 42 (2011), 55–90.
- [46] Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O. Stanley. 2019. Paired Open-Ended Trailblazer (POET): Endlessly Generating Increasingly Complex and Diverse Learning Environments and Their Solutions. *ArXiv abs/1901.01753* (2019).
- [47] Ren-Jian Wang, Ke Xue, Yu-Tong Wang, Peng Yang, Hao-Bo Fu, Qiang Fu, and Chao Qian. 2025. Diversity from Human Feedback. *Frontiers of Computer Science* 20, 2 (2025), 2002320.
- [48] Yutong Wang, Bairan Xiang, Shinan Huang, and Guillaume Sartoretti. 2023. SCRIMP: Scalable Communication for Reinforcement- and Imitation-Learning-Based Multi-Agent Pathfinding. In *IROS*. 9301–9308.
- [49] Hejia Zhang, Matthew C. Fontaine, Amy K. Hoover, Julian Togelius, Bistra Dilkina, and Stefanos Nikolaidis. 2020. Video Game Level Repair via Mixed Integer Linear Programming. In *AIIDE*. 151–158.
- [50] Yulun Zhang, Matthew C. Fontaine, Varun Bhatt, Stefanos Nikolaidis, and Jiaoyang Li. 2023. Arbitrarily Scalable Environment Generators via Neural Cellular Automata. In *NeurIPS*. 57212–57225.
- [51] Yulun Zhang, Matthew C. Fontaine, Varun Bhatt, Stefanos Nikolaidis, and Jiaoyang Li. 2023. Multi-Robot Coordination and Layout Design for Automated Warehousing. In *IJCAI*. 5503–5511.
- [52] Yulun Zhang, Matthew C. Fontaine, Amy K. Hoover, and Stefanos Nikolaidis. 2022. Deep Surrogate Assisted MAP-Elites for Automated Hearthstone Deckbuilding. In *GECCO*. 158–167.
- [53] Yulun Zhang, He Jiang, Varun Bhatt, Stefanos Nikolaidis, and Jiaoyang Li. 2024. Guidance Graph Optimization for Lifelong Multi-Agent Path Finding. In *IJCAI*. 311–320.