

Synthesizing Priority Planning Formulae for Multi-Agent Pathfinding

Shuwei Wang¹, Vadim Bulitko¹, Taoan Huang², Sven Koenig², Roni Stern³

¹Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada

²University of Southern California, Los Angeles, USA

³Ben Gurion University of the Negev

{shuwei4, bulitko}@ualberta.ca, {taoanhua, skoenig}@usc.edu, roni.stern@gmail.com

Abstract

Prioritized planning is a popular approach to multi-agent pathfinding. It prioritizes the agents and then repeatedly invokes a single-agent pathfinding algorithm for each agent such that it avoids the paths of higher-priority agents. Performance of prioritized planning depends critically on cleverly ordering the agents. Such an ordering is provided by a priority function. Recent work successfully used machine learning to automatically produce such a priority function given good orderings as the training data. In this paper we explore a different technique for synthesizing priority functions, namely program synthesis in the space of arithmetic formulae. We synthesize priority functions expressed as arithmetic formulae over a set of meaningful problem features via a genetic search in the space induced by a context-free grammar. Furthermore we regularize the fitness function by formula length to synthesize short, human-readable formulae. Such readability is an advantage over previous numeric machine-learning methods and may help explain the importance of features and how to combine them into a good priority function for a given domain. Moreover, our experimental results show that our formula-based priority functions outperform existing machine-learning methods on the standard benchmarks in terms of success rate, run time and solution quality without using more training data.

Introduction

Multi-agent pathfinding (MAPF) is the problem of moving a group of agents to a set of goal locations while avoiding collisions. The two common objectives of MAPF are to minimize the makespan (the largest arrival time of any agent at its goal location) and to minimize the sum of costs (the sum of the arrival times of all agents at their goal locations). In this paper, we use the sum of costs as our optimization objective. MAPF is NP-hard to solve optimally (Yu and LaValle 2013; Surynek 2010; Ma et al. 2016) but has numerous real-world applications, including moving game characters in formation in video games (Ma et al. 2017), transporting goods in automated warehouses (Wurman, D’Andrea, and Mountz 2008), routing pipes in gas plants (Belov et al. 2020), co-ordinating self-driving cars in intersections (Li et al. 2023) and embedding virtual network requests in computer networks (Zheng et al. 2023).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Prioritized planning (PP) (Latombe 1991; Bennewitz, Burgard, and Thrun 2002; Silver 2005; Van Den Berg and Overmars 2005) is a common approach for solving MAPF suboptimally. It prioritizes the agents according to a given priority function and then plans shortest paths for them from their start locations to their goal locations in the order of their priorities such that the path of each agent collides neither with the already planned paths of all agents with higher priorities (moving obstacles) nor with the blocked cells in the environment (static obstacles). An effective priority function yields in a small sum of costs while an ineffective one could even prevent one from finding a set of collision-free paths.

Some researchers have manually designed priority functions (Van Den Berg and Overmars 2005; Wang et al. 2019; Wu, Bhattacharya, and Prorok 2020; Buckley 1988; Li et al. 2019). Such functions are understandable by humans but since they are designed for specific problem settings, they cannot be applied across a wide spectrum of MAPF instances. Thus research from the last year automatically learned priority functions with support vector machines from a set of features of MAPF instances and training data (Zhang et al. 2022). Such machine learned priority functions may be able to achieve better performance than manually designed ones but are typically less readable to humans and thus less explainable.

In this paper we attempt to combine human-readability of manually designed priority functions and high performance of machine-learned ones. We adopt the approach by Bulitko et al. (2022) and learn priority functions automatically. They are expressed as arithmetic formulae that map MAPF features to priorities. They are synthesized via a genetic search in the space of formulae specified by a context-free grammar. We regularize the fitness function by formula length to synthesize compact readable formulae. We then compare the performance of our formula synthesis method with the support vector machine learning method by Zhang et al. (2022).

Overall, we make the following **contributions**. First, we describe how we adapt an existing automated program synthesis method to the important problem of creating a priority function for solving MAPF with prioritized planning. Second, we show empirically that our formula-based priority functions often outperform the state-of-the-art support vector machine on the standard benchmarks in terms of success

rate, run time, and solution quality without requiring more training data. Third, we show how the readability of formulae may help explain the importance of specific MAPF features.

Problem Formulation

We define the MAPF problem and the PP approach for solving it. Then we frame finding a high-performance priority function as an optimization problem which is the focus of this work.

The **multi-agent pathfinding (MAPF) problem** (Stern 2019) p is defined by a tuple (\mathcal{A}, G) where $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$ is a set of agents and $G = (V, E)$ is a non-weighted undirected graph with a set of vertices V and a set of edges E . Each agent a_i has a start location $s_i \in V$ and a goal location $g_i \in V$. Time advances in discrete steps. At each step an agent can either move to an adjacent location or wait at its current location. We consider two types of conflicts: vertex conflicts and edge conflicts. A vertex conflict occurs when two agents attempt to be at the same vertex $v_i \in V$ at the same time step. An edge conflict occurs when two agents attempt to traverse the same edge $(v_i, v_j) \in E$ in the opposite direction at the same time step. A solution to a MAPF problem p is a set of conflict-free paths $\{\rho_1, \rho_2, \dots, \rho_n\}$ to move each agent from its start location to its goal location. The cost of the agent a_i 's path in the solution is the number of time steps needed for a_i to complete its path from s_i to g_i and remain motionless at g_i . The performance measure is the *sum of the costs* of all agents in a solution to a MAPF problem p . Solutions with lower sums of costs are preferred.

In PP a *priority function* assigns priority to each agent which is used to order the agents. Then a single-agent pathfinding algorithm, such as A* (Hart, Nilsson, and Raphael 1968), is repeatedly invoked in the order of agent priorities to compute a plan for each agent. When searching for a plan for an agent a_i the algorithm considers and avoids conflicts with the plans of higher priority agents $\{a_1, a_2, \dots, a_{i-1}\}$ and ignores all the lower priority agents $\{a_{i+1}, a_{i+2}, \dots, a_n\}$. To do so the single-agent search algorithm is modified to take into account not only the static structure of the search graph G but also already planned paths of the higher priority agents.

The order of agents in PP can have a substantial effect on the sum of costs of the computed paths as well as whether a valid path is found for each agent at all. An optimal ordering allows PP to succeed in computing solutions of lowest sum of costs. Such an ordering normally depends on the problem instance at hand. Thus the problem we tackle in this paper is to find a priority function that minimizes sums of costs on a set of MAPF problem instances. Such a priority function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ maps a vector of agent features and returns the agent priority which is used to order the agents. Formally we attempt to solve the following optimization problem:

$$f_{\min} = \underset{f \in F}{\operatorname{argmin}} \ell(f, P) \quad (1)$$

where F is a space of priority formulae, P is a set of MAPF

problems and ℓ is a *loss function*:

$$\ell(f, P) = \frac{1}{|P|} \sum_{i=1}^{|P|} \ln(\zeta(f, p_i)). \quad (2)$$

Here $\zeta(f, p_i)$ computes the sum of costs of the MAPF problem p_i when PP uses the priority function f . We scale the sum of costs logarithmically to emphasize the impact of the regularizer (Section).

We prefer solutions to the optimization problem to be (i) generated automatically, (ii) be human-readable and (iii) be portable (i.e., yield low sum of costs solutions on novel MAPF problem instances not seen during its generation).

We adopt the evaluation settings of Zhang et al. (2022): deterministic ranking and stochastic ranking with random restarts. We reproduce them below for the reader's convenience.

In the case of **deterministic ranking** we rank the agents by their predicted priority scores: $a_i < a_j$ if and only if $\hat{y}(\mathbf{x}_i) > \hat{y}(\mathbf{x}_j)$ where $\hat{y} : \mathbb{R}^n \rightarrow \mathbb{R}$ is the priority function that takes as input the feature vector $\mathbf{x} \in \mathbb{R}^n$ of an agent and outputs a priority score z_i and $\mathbf{x}_i, \mathbf{x}_j$ are feature vectors of agent a_i, a_j respectively.

In the case of **stochastic ranking** with random restarts we use the computed priority scores to form a probability distribution and generate a total priority ordering from agents with high priority to agents with low priority. Specifically, we normalize the predicted scores \hat{y}_I using the softmax function. Agents with higher normalized scores are more likely to be selected earlier and thus can plan earlier.

To test the performance of the PP methods, given a map $M \in \mathcal{M}$ and a number of agents n , we generate a set of test MAPF instances $\mathcal{I}_{Test}^{(M)}$, one from each scenario (Stern et al. 2019), by using the first n pairs of start and goal indices. We use the following measures for additional evaluation: (i) *success rate* (the percentage of successfully solved problems), (ii) *run time to first solution* (only applicable to stochastic ranking with random restarts, the time it takes to find a solution to a problem), (iii) *normalized sum of costs** and (iv) *average solution rank* (we rank the PP methods in ascending order by their sum of costs starting from index 0, and the index of a PP method on a problem is the method's solution rank on that problem).

Related Work

While PP is not guaranteed to be complete (i.e., to always find a solution when one exists) or optimal (i.e., to minimize the sum of costs), it is widely adopted and can produce solutions with near-optimal sum of costs (Morag et al. 2022). Therefore improvements to PP have been proposed over the years. Windowed Hierarchical Cooperative A* (Silver 2005) and Rolling-Horizon Collision Resolution (Li et al. 2021)

*The average normalized sum-of-costs measure only considers the problem instances that are solved and ignores the ones that are not solved. For instance, if out of 25 problems a PP method solves only one, its average normalized sum of costs is the sum of costs of that one problem.

use PP iteratively in an online manner, considering in every planning period only conflicts in the next few steps.[†]

Bennewitz, Burgard, and Thrun (2002) run PP multiple times, each with a different, randomly generated priority over the agents. Priority-Based Search (PBS) applies a conflict-directed approach to search in the space of possible agent priorities, defining a partial order over the agents to resolve conflicts between the agents (Ma et al. 2019).

Van Den Berg and Overmars (2005) proposed to use the distances between an agent’s start and goal locations as a priority function. Bnaya and Felner (2014) proposed a greedy conflict-oriented prioritization mechanism, where an agent is given a priority based on how much reserving its paths will increase the sum of costs of all other agents.

Learning in the context of MAPF has been explored in different ways including learning which MAPF algorithm to use for a given problem (Sigurdson et al. 2019; Ren et al. 2021; Alkazzi et al. 2022; Kaduri, Boyarski, and Stern 2020), learning which sets of agents should re-plan (Huang et al. 2022), learning which conflicts should be resolved first (Huang, Koenig, and Dilkina 2021) or which search tree node to expand first (Huang, Dilkina, and Koenig 2021) in tree search algorithms for MAPF (Sharon et al. 2015; Barer et al. 2014). The closest learning approach to the work in this paper was to learn a support vector machine as the priority function (Zhang et al. 2022). In their work, they tested the performance of models trained on MAPF PP instances of 500 agents on problems with number of agents larger than 500 and demonstrated that the models can be applied to problem never seen during synthesis. We show that our trained formulae also have generalizability by testing the performance of a synthesized formula trained on a particular map with a specific number of agents on other maps with other numbers of agents thus demonstrate that the formulae are generalizable because it performs well in problems not seen during synthesis.

Program synthesis in heuristic search has been explored by Bulitko (2020); Bulitko, Hernandez, and Lelis (2021); Hernandez and Bulitko (2021); Bulitko et al. (2022) who ran various algorithms to search a space of arithmetic formulae to represent a heuristic function in single-agent pathfinding. A similar approach by Bulitko and Bulitko (2020) was used to synthesize Artificial Life (A-life) agents represented as arithmetic formulae. Both are cases of program synthesis, a growing research area attempting to synthesize/learn programs (or their parts) from data.

Our Approach

We adopt the approach of Bulitko et al. (2022) to synthesize priority functions for PP expressed as arithmetic formulae. We define the space of the formulae using a context-free grammar and then search this space via a genetic algorithm with regularized loss ℓ as the fitness function. The resulting synthesized formula computes a priority value for each agent, taking as input a set of the agent’s features (defined

[†]RHCR is a general framework that has been used with PP as well as with other MAPF solvers.

Algorithm 1: Parallel multi-trial synthesis; adopted from the work of Bulitko et al. (2022).

input : training problem set P , formula space F , regularized loss function ℓ^λ , number of trials T , per-trial time limit κ , limit of number of consecutive stagnate generations μ , population size n
output: synthesized formula f
1 **for** $t = 1, \dots, T$ **in parallel do**
2 $f_t \leftarrow \text{trial}(F, \ell, \lambda, P, \kappa, \mu, n)$
3 $\ell_t^\lambda \leftarrow \ell^\lambda(f_t, P)$
4 **return** $f = \arg\min_f \ell_t^\lambda$

by Zhang et al. (2022) and explained in Table 1). This priority value is then used by PP to order the agents and solve a given MAPF instance.

Multiple Synthesis Trials

In line with Bulitko et al. (2022) we take advantage of parallel computing hardware. Each candidate formula from the space F is evaluated with the PP algorithm on multiple MAPF problem instances in parallel by utilizing multiple cores on a single cluster node. Each of the genetic algorithm (or trial) run in parallel on different cluster nodes. Each trial synthesizes a single priority function. The best of them is selected from the results of multiple evolution trials.

We run T independent synthesis trials in parallel in line 1 of Algorithm 1. We regularize the loss function defined in Section in the same way as Bulitko et al. (2022):

$$\ell^\lambda(f, P) = \ell(f, P) + \lambda|f| \quad (3)$$

to bias the genetic search towards shorter formulae which may be easier to read and less likely to overfit. Here $|f|$ is the number of vertices in an abstract syntax tree that represents the formula and λ is the regularizer constant.

A Single Synthesis Trial

Each synthesis trial (i.e., a run of the genetic algorithm) is carried out by Algorithm 2. Our adaptations to the original algorithm by Bulitko et al. (2022) lie with the addition of κ , ν and using a single problem set P .

Henceforth we use $x_{1, \dots, n}$ as a shorthand for (x_1, \dots, x_n) . We start with a population of n formulae randomly drawn from the space F in line 1 of Algorithm 2. In line 2 we initialize a historically best/lowest loss to ∞ so that it gets updated immediately during the synthesis. The main loop ending in line 14 terminates if either the time allowance κ expires or the number of consecutive stagnate generations ν (i.e., generations which did not produce a better solution than the current historic best formula) exceeds μ .

In each generation all formulae in the population are evaluated by computing $\ell^\lambda(f_i, P)$ and the formula f' with the lowest loss is selected in line 5 as the candidate. The historic best formula is updated and the number of stagnate generations ν is reset.

Feature	Description
x_1, x_2, x_3	width of each level (excluding the first and the last levels) of MDD_i : their mean, max and min
x_4, x_5, x_6	graph distance between s_i and the start locations of the other agents: their mean, max and min
x_7, x_8, x_9	graph distance between g_i and the goal locations of the other agents: their mean, max and min
$x_{10}, x_{11}, x_{12}, x_{13}$	graph and Manhattan distances between s_i and g_i : graph distance, Manhattan distance, the ratio of the graph distance over the Manhattan distance and the absolute difference between the graph distance and the Manhattan distance
x_{14}	the sum of the widths of all levels of MDD_i
x_{15}	the number of locations in MDD_i that are also in the MDD of at least one other agent
x_{16}	the number of unit-width levels of MDD_i
x_{17}, x_{18}	the number of vertex conflicts between any shortest path of a_i and any shortest path of one of the other agents: counted by agent pair and counted by raw conflict count
x_{19}	the number of the other agents whose goal locations are in MDD_i
x_{20}	the number of the other agents whose start locations are in MDD_i
x_{21}	the number of the other agents whose MDDs contain g_i
x_{22}	the number of the other agents whose MDDs contain s_i
x_{23}, x_{24}	the number of edge conflicts between any shortest path of a_i and any shortest path of one of the other agents: counted by agent pair and counted by raw conflict count
x_{25}, x_{26}	the number of cardinal conflicts between any shortest path of a_i and any shortest path of one of the other agents: counted by agent pair and counted by raw conflict count

Table 1: The 26 features $\Phi = \{x_1, \dots, x_{26}\}$ for the agent a_i as defined by Zhang et al. (2022).

Algorithm 2: Single synthesis trial; adopted from the work of Bulitko et al. (2022).

```

input :  $F, \ell, \lambda, P, \kappa, \mu, n$ 
output: synthesized formula  $f$ 
1  $f_{1,\dots,n} \sim F$ 
2  $l_{\text{historic best}} \leftarrow \infty$ 
3  $\nu \leftarrow 0$ 
4 repeat
5    $f', l' \leftarrow \text{best-of}_{\ell^\lambda(f_i, P)}(f_{1,\dots,n})$ 
6   if  $l' < l_{\text{historic best}}$  then
7      $l_{\text{historic best}} \leftarrow l'$ 
8      $f_{\text{historic best}} \leftarrow f'$ 
9      $\nu \leftarrow 0$ 
10  else
11     $\nu \leftarrow \nu + 1$ 
12   $f_1 \leftarrow f'$ 
13   $f_{2,\dots,n} \leftarrow \text{offspring}(f_1)$ 
14 until  $\kappa$  is reached or  $\nu > \mu$ 
15 return  $f_{\text{historic best}}$ 

```

The champion f' is put in the next generation (line 12) while the rest of the population is formed from the offspring of f' (line 13). Each offspring is f' mutated as follows. A node in f' is selected by random and either deleted or modified while keeping the new formula valid. For instance, if the selected node is a unary operator (i.e., $| \cdot |$), we may either replace it with another unary operator, replace it with

a random binary operator with randomly initialized children nodes or bypass it by deleting the current node and connect its child node with its parent node (only delete the node if it is the root node).

Empirical Evaluation

We compare performance of our synthesized formulae to the five methods of ordering agents in PP evaluated by Zhang et al. (2022) on the same 6 maps with the same numbers of agents. The maps have different sizes and structures and they are from the MAPF benchmark suite (Stern et al. 2019). The five priority functions are: (1) **LH**: a query-distance heuristic where agents with longer start-goal graph distances have higher priority (Van Den Berg and Overmars 2005); (2) **SH**: a query-distance heuristic where agents with shorter start-goal graph distance have higher priority (Ma et al. 2019); (3) **RND**: a heuristic that generates a random total priority ordering (Bennewitz, Burgard, and Thrun 2002); (4) **ML-T**: a machine-learned total priority ordering (Zhang et al. 2022); (5) **ML-P**: a machine-learned partial priority ordering (Zhang et al. 2022).

We generated training data for ML-T and ML-P in the same way as Zhang et al. (2022). That is we ran PP 100 times, once with LH, once with SH, and 98 times with RND to solve each MAPF instance $I \in \mathcal{I}_{\text{Train}}^{(M)}$. We picked the PP run with the least sum of costs for ML-T and the top 5 PP runs with the least sum of costs for ML-P. We used LH and SH to generate agent rankings as the training data for the support vector machines.

We extended the C++ implementation used by Zhang et al. (2022) with a support for arithmetic formulae used as priority functions within PP. We then used the training data for ML-T and ML-P used by Zhang et al. (2022) and additionally generated training data for higher number of agents for two large maps. Specifically we generated training data for lak303d and ost003d with number of agents $n \in \{600, 700, 800, 900\}$ while the previous work only went up to 500 agents.

Unlike the supervised machine learning methods used before, our synthesis algorithm does not require *a priori* solved MAPF problems. This is due to the fact that it solves MAPF problems using PP *during* the synthesis to compute the loss of each candidate formula (line 5 in Algorithm 2). While this removes the requirement to build supervised-style training data the actual synthesis itself is made slower.

In line with Zhang et al. (2022) we experimented in both the deterministic ranking and stochastic ranking with random restarts settings for the baseline PP algorithms LH, SH, RND, ML-T, ML-P and our formula synthesis method on MAPF test instances on all six maps.

The control variables for the synthesis algorithm are set as follows: the population size $n = 20$, the regularizer constant $\lambda = 0.05$, the per-trial time limit $\kappa = 1$ hour, the limit of consecutive stagnate generations $\mu = 15$.

Space of Priority Functions

We define the formula space as $F = \{f(\mathbf{x}) = S\}$ where $\mathbf{x} = (x_1, \dots, x_{26})$ is the feature vector of an agent. The formula body S is generated by the following context-free grammar, adopted and modified from previous work (Bulitko et al. 2022). This context-free grammar was designed in attempt to balance formula expressiveness and formula space size.

$$\begin{aligned} S &\rightarrow T \mid U \mid B \\ T &\rightarrow x_1 \mid x_2 \mid \dots \mid x_{26} \mid C \\ U &\rightarrow \sqrt{S} \mid |S| \mid -S \mid S^2 \\ B &\rightarrow S + S \mid S - S \mid S \times S \mid \frac{S}{S} \mid \max\{S, S\} \mid \min\{S, S\} \end{aligned}$$

Here x_1 to x_{26} are the agent features (Table 1), $C \in \{1.0, 1.1, 1.2, \dots, 10.0\}$. Note that any support vector machine learned by ML-T and ML-P approaches (Zhang et al. 2022) can be represented in our space of formulae albeit with numeric weights rounded to the values in C . Also note that our formula space includes the two types of query-distance heuristics LH and SH as x_{10} and $-x_{10}$.

Synthesis of Priority Formulae

In order to generate training MAPF instances that follow a similar distribution as the test MAPF instances, given a scenario with a map $M \in \mathcal{M}$ and a number of agents n , we generated a training instance $I \in \mathcal{I}^{(M)}_{\text{Train}}$ by randomly selecting n start locations from all start locations in the scenario, randomly selecting n goal locations from all goal locations in the scenario, and then randomly combining them into n pairs of start and goal locations. Unlike the way Zhang et al.

(2022) prepared the training data by solving each instance before training, we did not solve each generated training instance because we only needed the features of the n agents of each problem instance. We then normalized the feature values of the agents to make them to be between 0 and 1 across all agents.

For the three small maps random-32-32-20, room-32-32-4 and maze-32-32-2 (referred to as random, room and maze in the tables) we generated 10 training MAPF instances from each of the 25 scenarios, thus $|\mathcal{I}^{(M)}_{\text{Train}}| = 250$. For the medium sized map warehouse-10-20-10-2-1 (referred to as warehouse in the tables) and the two large maps lak303d and ost003d, we generated 2 training MAPF instances from each of the 25 scenarios, thus $|\mathcal{I}^{(M)}_{\text{Train}}| = 50$. The number of problems decreases with the map size as problem instances become computationally more expensive to solve.

We ran 32 trials for each number of agents of each map and selected the trial with the lowest/best loss as the best trial. The formula produced by each best trial was used to compute agent priority scores for testing. Each trial ran on a separate 16-core cluster node with 64Gb of RAM.

Results: Deterministic Ranking

In terms of both success rate and average solution rank, the formula synthesis method outperformed the other methods for 3 out of 6 maps. Our formula synthesis method achieved average normalized sum of costs values similar to others.

We omit results for random-32-32-20 with number of agents 250 because none of the PP methods solved any test problem instances in the deterministic setting.

Results: Stochastic Ranking with Random Restarts

We report the success rate, solution ranks and run time to find the first solution in Table 3. We adopted the value for the hyper parameter from Zhang et al. (2022) ($\beta = 0.5$) when testing. We omit results of small numbers of agents for the six maps because the success rates and average solution ranks are similar across all PP algorithms.

The random restart technique is used to boost the success rate of the two ML-guided methods while preserving their average solution ranks (Zhang et al. 2022). While it is clear that the two ML-guided methods have the best performance in all three measures among the PP algorithms, the synthesized formulae are not benefiting as much from the random restart scheme as the two ML-guided methods are and sometimes the formulae perform even worse with random restarts compared to without.

In the stochastic setting with random restarts, ML-P remained the best-performing algorithm in success rate, average run time to the first solution and average solution rank. Although not shown in Table 3, ML-P also performed the best among all methods in terms of the average normalized sum of costs.

Explainability and Portability of Synthesized Formulae

Synthesized formulae for all maps and all numbers of agents in Table 4. Consider that a trained support vector machine

Map	n	Success rate (%)						Average solution rank						Avg. normalized sum of costs					
		LH	SH	RND	ML-T	ML-P	FML	LH	SH	RND	ML-T	ML-P	FML	LH	SH	RND	ML-T	ML-P	FML
random	50	96	12	76	92	52	96	2.56	3.76	1.48	1.60	1.92	2.00	1.48	1.10	1.29	1.46	1.09	1.47
	100	96	24	48	40	16	92	1.76	2.12	1.96	1.76	2.64	1.32	1.59	1.16	1.43	1.15	1.15	1.47
	150	52	4	20	64	4	60	1.12	1.88	1.76	1.28	1.92	0.72	1.60	1.19	1.34	1.67	1.24	1.41
	175	36	0	8	44	40	48	1.20	1.76	1.48	1.12	0.96	0.76	1.71	∞	1.40	1.64	1.66	1.45
	200	12	0	0	32	28	40	0.96	1.12	1.12	0.68	0.84	0.40	1.70	∞	∞	1.71	1.68	1.64
	225	0	0	0	4	4	8	0.16	0.16	0.16	0.12	0.12	0.08	∞	∞	∞	1.82	1.39	1.44
room	50	88	16	40	60	16	88	1.92	2.28	1.96	1.16	2.24	1.32	1.66	1.38	1.44	1.45	1.35	1.57
	75	88	0	4	24	32	80	1.16	2.28	2.24	1.60	1.48	0.60	1.69	∞	2.09	1.58	1.60	1.63
	100	52	0	4	56	64	68	1.72	2.44	2.36	1.40	0.88	0.60	1.78	∞	1.87	1.80	1.73	1.75
	125	20	0	0	24	20	28	0.48	0.92	0.92	0.60	0.64	0.48	1.88	∞	∞	1.91	1.88	1.86
	150	0	0	0	0	4	8	0.12	0.12	0.12	0.12	0.08	0.04	∞	∞	∞	∞	2.13	1.81
maze	50	84	0	16	76	76	84	1.60	3.36	2.68	1.36	1.16	1.56	3.88	∞	3.55	4.00	3.99	4.00
	70	84	0	0	80	80	84	1.40	3.28	3.28	1.56	1.16	1.40	4.11	∞	∞	4.03	3.94	3.95
	90	72	0	0	56	64	68	1.20	2.60	2.60	1.48	1.20	1.08	4.25	∞	∞	4.49	4.25	4.30
	110	48	0	0	24	48	24	0.64	1.44	1.44	0.92	0.76	1.00	4.68	∞	∞	5.12	4.56	4.52
	130	12	0	0	12	16	8	0.36	0.48	0.48	0.32	0.24	0.40	5.51	∞	∞	4.86	4.63	5.17
warehouse	100	92	24	80	88	80	20	2.84	2.84	1.56	2.00	0.52	3.08	1.76	1.01	1.39	1.62	1.02	1.27
	200	92	36	40	52	56	32	2.16	1.36	2.20	1.40	0.96	2.40	1.89	1.03	1.53	1.09	1.06	1.40
	300	64	20	24	16	36	28	1.24	1.16	1.48	1.56	1.00	1.36	1.88	1.05	1.49	1.07	1.08	1.49
	350	36	8	16	20	12	28	0.84	1.00	0.96	0.84	0.92	0.80	1.96	1.07	1.77	1.29	1.22	1.43
	400	12	4	4	32	8	32	0.80	0.80	0.88	0.48	0.72	0.36	2.00	1.10	1.53	1.94	1.18	1.45
	450	16	0	4	12	8	24	0.48	0.64	0.56	0.48	0.56	0.32	2.04	∞	1.62	1.97	1.46	1.47
	500	0	0	4	0	0	4	0.08	0.08	0.04	0.08	0.08	0.04	∞	∞	1.87	∞	∞	1.56
	550	0	0	0	0	4	4	0.08	0.08	0.08	0.08	0.04	0.04	∞	∞	∞	∞	2.06	1.53
lak303d	300	100	24	100	96	80	96	3.68	3.64	1.16	2.76	1.20	2.32	2.73	1.71	2.15	2.60	1.76	2.58
	400	100	40	96	88	76	100	3.56	2.72	2.32	1.76	1.44	2.72	2.65	1.71	2.36	2.07	1.75	2.64
	500	96	24	80	84	72	92	3.16	3.12	2.24	1.44	1.32	2.44	2.71	1.74	2.37	2.03	1.79	2.59
	600	92	48	84	92	64	88	3.32	1.88	1.92	2.56	1.52	2.36	2.71	1.78	2.54	2.73	1.82	2.62
	700	92	36	76	52	48	88	2.68	2.08	2.04	1.44	2.08	2.28	2.86	1.82	2.66	1.82	1.87	2.72
	800	80	16	56	76	40	88	2.04	2.76	1.80	2.16	1.88	1.88	2.99	1.81	2.68	2.97	1.94	2.84
	900	60	8	28	56	16	64	1.28	1.92	1.44	1.48	1.80	1.00	3.09	1.83	2.73	3.18	1.97	3.05
	300	100	32	96	100	80	84	3.68	3.16	1.80	2.52	1.56	1.92	2.70	1.71	2.17	2.56	1.78	1.98
ost003d	400	100	36	96	84	88	92	3.76	3.00	2.56	1.80	1.20	2.24	2.74	1.75	2.41	1.84	1.76	2.08
	500	100	36	80	80	72	100	3.20	2.56	1.88	2.04	1.56	2.52	2.79	1.75	2.18	2.25	1.81	2.60
	600	96	40	80	76	68	92	3.08	2.28	1.88	1.88	1.56	2.88	2.67	1.78	2.32	2.34	1.79	2.83
	700	100	36	72	88	60	100	2.76	2.44	2.00	2.52	1.80	2.36	2.79	1.82	2.37	2.90	1.82	2.88
	800	88	24	52	88	40	96	2.52	2.56	2.20	2.28	2.24	1.24	2.91	1.87	2.30	2.92	1.84	2.82
	900	88	12	44	20	12	88	1.28	2.12	1.44	2.08	2.20	1.12	2.95	1.87	2.33	1.86	1.85	2.93

Table 2: Success rate, average solution rank and average normalized sum of costs for deterministic ranking. The best results achieved among all algorithms are shown in bold. The results are obtained by training and testing on the same map with the same number of agents n . ∞ indicates that no problem was solved.

$f(\Phi) = \mathbf{w}^\top \Phi$ (Zhang et al. 2022) when converted into a formula in space F , contains 26 features, 26 weights multiplied to each feature, 26 multiplication operators and 25 addition operators, thus in total $26 + 26 + 26 + 25 = 103$ nodes (assuming no weights are 0). In contrast, the average number of nodes in the synthesized formulae in Table 4 is 6.95. Therefore the synthesized formulae are on average shorter and thus may be more readable than the support vector machines.

Synthesized formulae that outperform other algorithms in terms of the success rate and average solution rank frequently include x_7 or $-x_7$. The feature x_7 is the mean graph distance between the agent i 's goal g_i and the goal locations of other agents. This suggests that proximity of the agent's

goal to other agents' goals is an important factor in determining the agent's priority. Whether x_7 or $-x_7$ is more effective depends on the problem configuration. Indeed, $-x_7$ appears most often in formulae synthesized for random-32-32-20 while x_7 appears most often in formulae synthesized for warehouse-10-20-10-2-1.

Consider the following formula from Table 4:

$$f_6 = - \left(\frac{x_7}{10 - x_1 + x_{18}^2} \right)^2 = \frac{x_7^2}{-(10 - x_1 + x_{18}^2)^2}$$

which outperforms other PP methods in success rate and average solution rank on random-32-32-20 with number of agents 225 (Table 2).

Map	n	Success rate (%)							Average solution rank							Avg. run time to the first solution						
		LH	SH	RND	ML-T	ML-P	FML	LH	SH	RND	ML-T	ML-P	FML	LH	SH	RND	ML-T	ML-P	FML			
random	175	100	100	100	100	100	100	3.16	1.88	2.84	2.40	1.96	2.76	0.61	3.64	3.82	1.67	1.41	1.38			
	200	88	84	92	100	100	100	2.92	2.36	2.60	2.76	2.24	2.00	14.57	21.30	15.79	3.58	2.91	3.70			
	225	32	16	20	84	96	16	2.00	2.20	1.84	1.28	0.48	2.08	50.36	52.78	51.32	24.49	13.39	52.72			
	250	0	4	0	12	80	0	0.96	0.88	0.96	0.72	0.16	0.96	60.00	58.31	60.00	56.65	20.77	60.00			
room	75	100	100	100	100	100	100	3.40	2.44	2.36	2.08	2.48	2.16	0.21	0.29	0.59	0.55	0.38	0.33			
	100	80	80	80	96	100	92	2.80	2.40	2.92	2.00	1.80	2.24	12.56	19.48	20.66	5.00	1.21	10.08			
	125	24	0	12	76	92	12	1.44	2.16	1.92	0.92	0.56	1.88	47.74	60.00	55.45	27.70	8.60	55.10			
	150	0	0	0	8	76	76	1.60	1.60	1.60	1.48	0.64	0.32	60.00	60.00	60.00	57.58	27.41	16.76			
maze	50	100	100	100	100	100	100	3.52	2.52	1.52	2.80	2.20	2.36	0.40	1.63	4.30	1.36	0.26	0.74			
	70	88	76	72	96	100	100	3.00	2.40	2.36	1.96	2.68	2.04	7.37	20.49	21.75	5.78	0.45	7.03			
	90	64	20	20	84	100	40	1.64	2.64	2.76	1.40	0.80	2.24	24.92	49.35	52.59	17.64	0.57	39.37			
	110	40	0	0	40	100	28	1.36	2.08	2.08	1.40	0.40	1.36	36.04	60.00	60.00	43.63	1.86	50.97			
	130	8	0	0	16	84	0	1.00	1.08	1.08	0.84	0.08	1.08	55.22	60.00	60.00	53.89	15.81	60.00			
warehouse	350	96	96	92	96	96	96	3.28	2.12	2.36	2.24	2.20	2.20	9.39	14.76	11.23	8.69	11.25	14.56			
	400	80	88	84	80	84	84	3.20	2.04	2.48	2.28	2.04	1.84	19.38	16.41	23.30	23.04	19.92	22.93			
	450	64	64	68	60	84	68	2.40	2.08	2.12	2.44	1.56	1.96	31.90	35.31	32.71	37.21	30.30	31.70			
	500	36	20	32	56	20	28	1.36	1.60	0.88	0.92	1.24	1.08	44.60	55.25	46.89	42.88	53.22	49.29			
	550	16	12	8	24	16	12	0.52	0.60	0.76	0.48	0.60	0.64	55.96	57.82	58.56	52.44	56.65	56.98			
lak303d	500	100	100	100	100	96	100	4.48	2.08	2.48	1.72	2.24	2.00	5.25	11.56	22.66	9.41	34.79	8.91			
	600	100	100	100	100	100	100	4.64	2.56	1.88	2.40	1.12	2.40	7.49	40.40	36.11	32.90	99.87	21.10			
	700	96	100	100	96	100	100	4.32	1.48	1.76	2.04	2.40	3.00	41.98	38.08	73.10	101.03	72.38	68.01			
	800	96	96	92	96	100	92	4.12	2.68	2.08	2.24	1.56	2.04	60.65	148.05	158.82	117.32	126.01	124.54			
	900	92	92	84	84	80	76	2.88	1.80	2.00	2.44	1.88	2.64	130.59	232.18	249.02	260.38	276.37	255.47			
ost003d	500	100	96	100	96	92	100	4.44	2.52	1.84	2.40	1.16	2.48	2.40	39.21	20.80	31.62	60.03	10.79			
	600	100	100	96	100	96	100	4.08	2.44	2.28	1.80	1.36	3.00	4.46	16.58	37.82	18.79	62.19	13.55			
	700	100	96	96	100	96	96	4.32	1.96	2.40	2.00	1.40	2.64	13.79	47.78	46.36	44.20	83.66	47.17			
	800	100	96	96	96	96	96	3.56	1.92	2.60	2.72	1.44	2.36	32.53	95.01	82.01	80.12	110.79	72.39			
	900	100	92	96	96	92	92	3.68	2.32	2.16	2.08	2.04	2.32	72.49	164.68	153.91	160.27	172.58	169.36			

Table 3: Success rate, average solution rank and average run time to find the first solution for stochastic ranking with random restarts. The best results achieved among all algorithms are shown in bold. The results are obtained by training and testing on the same map with the same number of agents n . ∞ indicates that no problem was solved.

The numerator x_7^2 gives low priority to the agents that have goals close to other agents' goals. Consider the denominator $-(10 - x_1 + x_{18}^2)^2$. Suppose $x_1 < 10$ then $10 - x_1$ becomes smaller as x_1 grows and because of the negation sign, it makes the denominator larger which leads to smaller priority scores. So as x_1 grows (given $x_1 < 10$) the agent's priority score becomes larger. Suppose $x_1 > 10$ then $10 - x_1 < 0$ but the square makes it positive. Therefore by the same reasoning, as x_1 grows, the agent's priority score becomes larger. When $x_1 = 10$ the denominator becomes $-x_{18}^4$. Changes in x_1 is less significant to the value of the priority score than that of x_{18} because x_{18} appears with a power of 4 in the formula while x_1 does not so we focus more on x_{18} .

As x_{18}^2 grows the denominator becomes smaller because of the negation sign and the priority score becomes larger. Thus the larger x_{18} the larger the priority score and agents that are more likely to have conflicts with other agents will be prioritized and will plan earlier.

We will now evaluate performance of f_6 on the three small maps and the one medium sized map with deterministic ranking to demonstrate an example of the formula's portability. Instead of computing the agents' priority scores using the formulae synthesized specifically for that num-

ber of agents and that map, we use f_6 on all the maps and with all numbers of agents. As per Figure 1 the formula f_6 had the highest success rate and the best average solution rank for most numbers of agents on random-32-32-20 and room-32-32-4. It performed worse on maze-32-32-2 and warehouse-10-20-10-2-1. This shows that the synthesized formulae can outperform existing PP algorithms even on maps not seen during synthesis.

Future Work

While performing well, our priority functions were synthesized for a given map and a given number of agents. Future work will synthesize priority functions with training data from different maps and/or numbers of agents together. Effectiveness of our approach may also be increased by considering additional building blocks for the arithmetic formulae. In particular, these building blocks can include previously synthesized formulae (Bulitko et al. 2022), leading to iterative expansion of the space of formulae. Finally, future work may aim to scale our approach to larger training sets by using other synthesis methods (Shah et al. 2020; Mouret and Clune 2015; Gallotta, Arulkumaran, and Soros 2022;

Map	n	Formula
random	50	$f_1 = \sqrt{x_7} - x_{15} + \sqrt{x_{21}}$
	100	$f_2 = -x_7 - x_{12} \cdot x_{15}$
	150	$f_3 = -21.2 \cdot x_7 - x_{15}$
	175*	$f_4 = -x_7 - x_{12}^2 \cdot x_{15}$
	200*	$f_5 = -8.3 \cdot x_7 + \sqrt{5.5 + x_{15}}$
	225*	$f_6 = -\left(\frac{x_7}{10-x_1+x_{18}}\right)^2$
room	50	$f_7 = x_4 + 28.1 \cdot x_7 - x_{14}$
	75	$f_8 = (x_7 - 4.7)^8 + x_{15}$
	100*	$f_9 = -x_7 - (0.2 \cdot x_{15} \cdot x_{16})^2$
	125*	$f_{10} = \left(\frac{x_4 \cdot x_{15}}{7.9 + \max\{x_{14}, \sqrt{x_{16}}\}}\right)^2 - \sqrt{x_7}$
	150*	$f_{11} = -(789.1 \cdot x_7^2 + \sqrt{x_{11}} + x_{17} + x_{18})$
maze	50	$f_{12} = -x_{10} - x_{15} + x_{21}$
	70	$f_{13} = -x_7^2 - x_{10} + \sqrt{x_{21}}$
	90	$f_{14} = -\sqrt{x_7 + x_{10} \cdot \sqrt{x_{16}}} + x_{21}$
	110	$f_{15} = -x_{15} + 4.3 \cdot \sqrt{\sqrt{x_{21}}}$
	130	$f_{16} = -x_7 - x_{15}$
warehouse	100	$f_{17} = x_7 + x_{12} \cdot x_{14}$
	200	$f_{18} = 61.5 \cdot x_8^2 - x_{15}$
	300	$f_{19} = 181 \cdot x_7 + x_{15}$
	350	$f_{20} = x_7 - x_{12} \cdot x_{14}$
	400*	$f_{21} = x_3 + x_8 + x_{12} \cdot x_{15}$
	450*	$f_{22} = x_8 + (x_{20} - x_4)^4 \cdot x_{14}^2$
	500*	$f_{23} = x_7 + 1.8 \cdot x_{11}^2$
lak3d	550*	$f_{24} = x_7 + x_8^2$
	300	$f_{25} = \sqrt{x_4} - x_{15}$
	400	$f_{26} = x_1 - x_{14} + x_{21}$
	500	$f_{27} = -x_{10}(x_{10} + 1) + x_{15}$
	600	$f_{28} = -7 \cdot x_7^2 - \sqrt{x_{15}}$
	700	$f_{29} = x_1^2 - x_{15} + \sqrt{x_{21}}$
	800	$f_{30} = -x_8 - x_{15}^2 - x_{16}^2$
	900*	$f_{31} = \min\{x_{23}, -x_7 - (x_{10} + (x_{12} - x_3) \cdot x_{22})^2\}$
ost003d	300	$f_{32} = x_7 - x_{18}^2$
	400	$f_{33} = x_7 - x_{11}$
	500	$f_{34} = -x_7 - x_{15}$
	600	$f_{35} = x_1 - 1.7 \cdot x_{15}$
	700	$f_{36} = -x_7 - \sqrt{x_{15}}$
	800*	$f_{37} = -x_7^2 - x_{12}^2$
	900*	$f_{38} = -\sqrt{3.2 + x_{10} + x_{12} + x_{15}}$

Table 4: Formulae synthesized for each number of agents of each map. All formulae are manually simplified and the numeric constants are rounded. When PP with a formula outperforms all other PP methods in both success rate and average solution rank for a given map and a number of agents we mark the line with an asterisk.

Kimbrough et al. 2008).

Conclusions

We adopted the approach by Bulitko et al. (2022) to learn priority functions to solve multi-agent pathfinding problems with prioritized planning. The priority functions are expressed as arithmetic formulae and synthesized via a genetic search. They are short and human-readable and often outperform the state-of-the-art machine-learning approach in terms

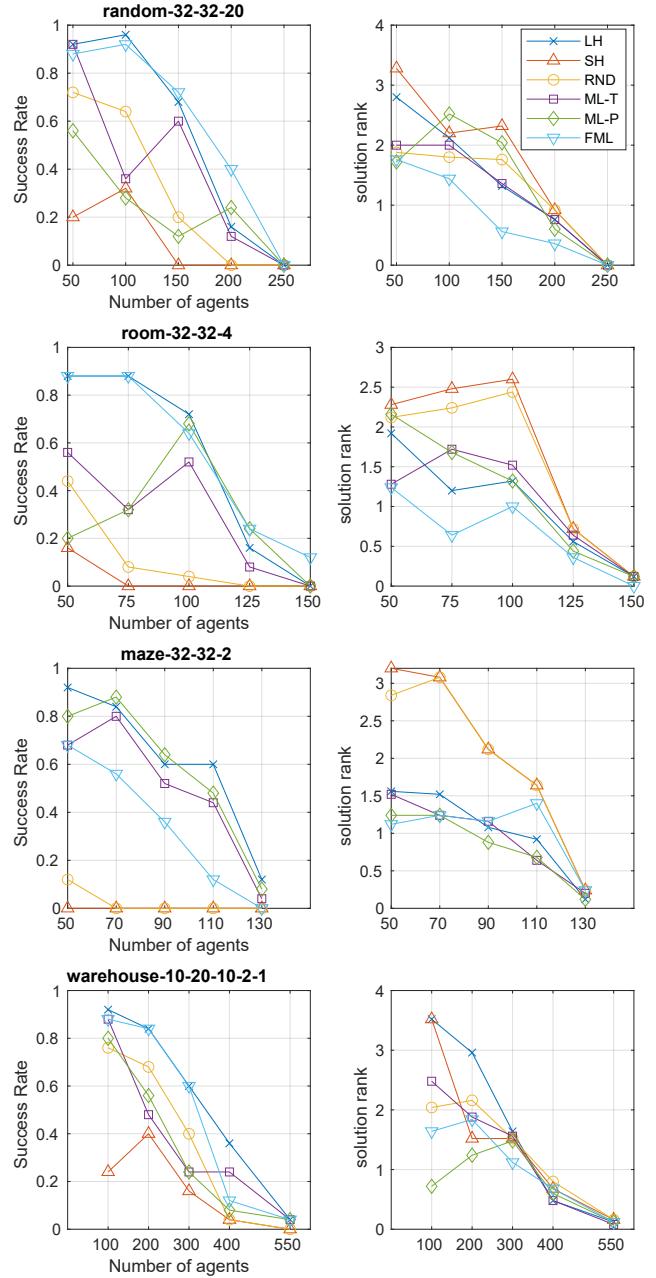


Figure 1: Success rate and average solution rank of f_6 on the three small maps and the medium sized map compared to existing PP algorithms.

of success rate, run time and solution quality without requiring more training data.

We also showed that our synthesized formula can outperform existing PP algorithms in both success rate and average solution rank even on maps and problems not seen during synthesis and can provide insight in their operation.

Acknowledgments

We appreciate support from Compute Canada and NSERC.

References

Alkazzi, J.-M.; Rizk, A.; Salomon, M.; and Makhoul, A. 2022. MAPFASTER: A Faster and Simpler take on Multi-Agent Path Finding Algorithm Selection. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 10088–10093.

Barer, M.; Sharon, G.; Stern, R.; and Felner, A. 2014. Sub-optimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, volume 5, 19–27.

Belov, G.; Du, W.; de la Banda, M. G.; Harabor, D.; Koenig, S.; and Wei, X. 2020. From Multi-Agent Pathfinding to 3D Pipe Routing. In *Proceedings of the Symposium on Combinatorial Search (SoCS)*, 11–19.

Bennewitz, M.; Burgard, W.; and Thrun, S. 2002. Finding and optimizing solvable priority schemes for decoupled path planning techniques for teams of mobile robots. *Robotics and autonomous systems*, 41(2-3): 89–99.

Bnaya, Z.; and Felner, A. 2014. Conflict-oriented windowed hierarchical cooperative A. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 3743–3748.

Buckley, S. J. 1988. Fast motion planning for multiple moving robots. Technical Report RC 14099 (#63205), IBM Thomas J. Watson Research Division.

Bulitko, V. 2020. Evolving Initial Heuristic Functions for Agent-Centered Heuristic Search. In *Proceedings of the IEEE Conference on Games (CoG)*, 534–541.

Bulitko, V.; and Bulitko, V. K. 2020. Complexity of Agents in Non-stationary Environments: A Partial Computational Model. In *Proceedings of the Annual Conference on Advances in Cognitive Systems (ACS), Poster Collection*.

Bulitko, V.; Hernandez, S. P.; and Lelis, L. H. 2021. Fast Synthesis of Algebraic Heuristic Functions for Video-game Pathfinding. In *Proceedings of the IEEE Conference on Games (CoG)*, 01–05. IEEE.

Bulitko, V.; Wang, S.; Stevens, J.; and Lelis, L. H. 2022. Portability and Explainability of Synthesized Formula-based Heuristics. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, volume 15, 29–37.

Gallotta, R.; Arulkumaran, K.; and Soros, L. B. 2022. Surrogate Infeasible Fitness Acquisition FI-2Pop for Procedural Content Generation. In *Proceedings of IEEE Conference on Games (CoG)*, 500–503. IEEE.

Hart, P.; Nilsson, N.; and Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.

Hernandez, S. P.; and Bulitko, V. 2021. Speeding Up Heuristic Function Generation via Automatically Extending the Formula Grammar. In *Proceedings of the Symposium on Combinatorial Search (SoCS)*, 233–235.

Huang, T.; Dilkina, B.; and Koenig, S. 2021. Learning node-selection strategies in bounded suboptimal conflict-based search for multi-agent path finding. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 611–619.

Huang, T.; Koenig, S.; and Dilkina, B. 2021. Learning to resolve conflicts for multi-agent path finding with conflict-based search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11246–11253.

Huang, T.; Li, J.; Koenig, S.; and Dilkina, B. 2022. Anytime Multi-Agent Path Finding via Machine Learning-Guided Large Neighborhood Search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9368–9376.

Kaduri, O.; Boyarski, E.; and Stern, R. 2020. Algorithm selection for optimal multi-agent pathfinding. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, volume 30, 161–165.

Kimbrough, S. O.; Koehler, G. J.; Lu, M.; and Wood, D. H. 2008. On a Feasible-Infeasible Two-Population (FI-2Pop) genetic algorithm for constrained optimization: Distance tracing and no free lunch. *European Journal of Operational Research*, 190(2): 310–327.

Latombe, J.-C. 1991. Multiple Moving Objects. In *Robot motion planning*, 1–57. Springer.

Li, H.; Long, T.; Xu, G.; and Wang, Y. 2019. Coupling-degree-based heuristic prioritized planning method for UAV swarm path generation. In *Proceedings of Chinese Automation Congress (CAC)*, 3636–3641.

Li, J.; Hoang, T.; Lin, E.; Vu, H.; and Koenig, S. 2023. Intersection Coordination with Priority-Based Search for Autonomous Vehicles. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 11578–11585.

Li, J.; Tinka, A.; Kiesel, S.; Durham, J. W.; Kumar, T. S.; and Koenig, S. 2021. Lifelong multi-agent path finding in large-scale warehouses. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11272–11281.

Ma, H.; Harabor, D.; Stuckey, P. J.; Li, J.; and Koenig, S. 2019. Searching with consistent prioritization for multi-agent path finding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7643–7650.

Ma, H.; Tovey, C.; Sharon, G.; Kumar, T.; and Koenig, S. 2016. Multi-agent path finding with payload transfers and the package-exchange robot-routing problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 3166–3173.

Ma, H.; Yang, J.; Cohen, L.; Kumar, T.; and Koenig, S. 2017. 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)*, volume 13, 270–272.

Morag, J.; Felner, A.; Stern, R.; Atzmon, D.; and Boyarski, E. 2022. Online Multi-Agent Path Finding: New Results. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, volume 15, 229–233.

Mouret, J.-B.; and Clune, J. 2015. Illuminating search spaces by mapping elites. *arXiv:1504.04909*.

Ren, J.; Sathiyanarayanan, V.; Ewing, E.; Senbaslar, B.; and Ayanian, N. 2021. MAPFAST: A Deep Algorithm Selector for Multi Agent Path Finding using Shortest Path Embeddings. In *Proceedings of the International Conference*

on Autonomous Agents and Multiagent Systems (AAMAS), 1055–1063.

Shah, A.; Zhan, E.; Sun, J.; Verma, A.; Yue, Y.; and Chaudhuri, S. 2020. Learning Differentiable Programs with Admissible Neural Heuristics. In *Proceedings of the conference on Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 4940–4952.

Sharon, G.; Stern, R.; Felner, A.; and Sturtevant, N. R. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219: 40–66.

Sigurdson, D.; Boutilko, V.; Koenig, S.; Hernandez, C.; and Yeoh, W. 2019. Automatic algorithm selection in multi-agent pathfinding. *arXiv:1906.03992*.

Silver, D. 2005. Cooperative pathfinding. In *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment (AIIDE)*, volume 1, 117–122.

Stern, R. 2019. Multi-agent path finding—an overview. *Artificial Intelligence: 5th RAAI Summer School, Dolgoprudny, Russia, July 4–7, 2019, Tutorial Lectures*, 96–115.

Stern, R.; Sturtevant, N.; Felner, A.; Koenig, S.; Ma, H.; Walker, T.; Li, J.; Atzmon, D.; Cohen, L.; Kumar, T.; et al. 2019. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, volume 10, 151–158.

Surynek, P. 2010. An optimization variant of multi-robot path planning is intractable. In *Proceedings of the AAAI conference on artificial intelligence*, volume 24, 1261–1263.

Van Den Berg, J. P.; and Overmars, M. H. 2005. Prioritized motion planning for multiple robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 430–435.

Wang, J.; Li, J.; Ma, H.; Koenig, S.; and Kumar, S. 2019. A new constraint satisfaction perspective on multi-agent path finding: Preliminary results. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2253–2255.

Wu, W.; Bhattacharya, S.; and Prorok, A. 2020. Multi-robot path deconfliction through prioritization by path prospects. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 9809–9815.

Wurman, P. R.; D’Andrea, R.; and Mountz, M. 2008. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine*, 29: 9–9.

Yu, J.; and LaValle, S. 2013. Structure and intractability of optimal multi-robot path planning on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, 1443–1449.

Zhang, S.; Li, J.; Huang, T.; Koenig, S.; and Dilkina, B. 2022. Learning a Priority Ordering for Prioritized Planning in Multi-Agent Path Finding. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, volume 15, 208–216.

Zheng, Y.; Ma, H.; Koenig, S.; Kline, E.; and Kumar, S. 2023. Priority-Based Search for the Virtual Network Embedding Problem. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 472–480.