

CTRM: Learning to Construct Cooperative Timed Roadmaps for Multi-agent Path Planning in Continuous Spaces

Keisuke Okumura*
Tokyo Institute of Technology
Tokyo, Japan
okumura.k@coord.c.titech.ac.jp

Ryo Yonetani, Mai Nishimura
OMRON SINIC X
Tokyo, Japan
{ryo.yonetani, mai.nishimura}@sinicx.com

Asako Kanezaki
Tokyo Institute of Technology
Tokyo, Japan
kanezaki@c.titech.ac.jp

ABSTRACT

Multi-agent path planning (MAPP) in continuous spaces is a challenging problem with significant practical importance. One promising approach is to first construct graphs approximating the spaces, called *roadmaps*, and then apply multi-agent pathfinding (MAPF) algorithms to derive a set of conflict-free paths. While conventional studies have utilized roadmap construction methods developed for single-agent planning, it remains largely unexplored how we can construct roadmaps that work effectively for multiple agents. To this end, we propose a novel concept of roadmaps called *cooperative timed roadmaps (CTRM)*. CTRMs enable each agent to focus on its important locations around potential solution paths in a way that considers the behavior of other agents to avoid inter-agent collisions (*i.e.*, “cooperative”), while being augmented in the time direction to make it easy to derive a “timed” solution path. To construct CTRMs, we developed a machine-learning approach that learns a generative model from a collection of relevant problem instances and plausible solutions and then uses the learned model to sample the vertices of CTRMs for new, previously unseen problem instances. Our empirical evaluation revealed that the use of CTRMs significantly reduced the planning effort with acceptable overheads while maintaining a success rate and solution quality comparable to conventional roadmap construction approaches.[†]

KEYWORDS

Biased Sampling; Multi-agent Pathfinding; Data-driven Planning

ACM Reference Format:

Keisuke Okumura*, Ryo Yonetani, Mai Nishimura, and Asako Kanezaki. 2022. CTRM: Learning to Construct Cooperative Timed Roadmaps for Multi-agent Path Planning in Continuous Spaces. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022, IFAAMAS*, 10 pages.

1 INTRODUCTION

Multi-agent path planning (MAPP) is a fundamental problem in multi-agent systems with attractive applications such as warehouse automation [54] and coordination of self-driving cars [9, 24]. Over the last decade, significant progress has been made on MAPP in *discretized* environments, which is also known as *multi-agent path finding (MAPF)*, aimed at finding a set of conflict-free paths on a graph given *a priori*, typically as a grid. Despite its optimization

intractability in various criteria [2, 37, 56], recent work has shown MAPF can obtain near-optimal solutions in a short amount of time (*e.g.*, less than a minute), even for hundreds of agents [29, 32, 40].

In this work, we tackle another largely unexplored research area: MAPP in *continuous* spaces, which has been recognized as tremendously challenging since the 1980s [17, 50]. On the one hand, allowing agents to move in a continuous space leads to better solutions compared to when their movements are limited in a non-deliberative discretized space such as lattice grids. On the other hand, it is not obvious how to search for these better solutions in the continuous space.

One possible approach to MAPP in continuous spaces consists of two steps: (1) approximating the continuous spaces by constructing graphs called *roadmaps*, and then (2) using powerful MAPF algorithms on those roadmaps to derive a solution. While such approaches have been widely used for single-agent planning [30], doing the same for multiple agents is non-trivial. This is primarily due to the necessity of constructing sparse roadmaps, as doing otherwise would make it dramatically more difficult to find a combination of plausible paths due to having to manage a higher number of inter-agent collisions. Nevertheless, there is generally a trade-off between roadmap density and solution quality; roadmaps should be sufficiently dense to ensure a high planning success rate and better solutions. This begs our key question: “*What are the characteristics to consider in order to effectively construct roadmaps for MAPP in continuous spaces?*”

To address this question, we propose a novel concept of graph representations of the space called *cooperative timed roadmaps (CTRM)*. CTRMs consist of directed acyclic graphs constructed in the following fashion. (1) *Agent-specific*: Each CTRM is specialized for individual agents to focus on their important locations. (2) *Cooperative*: Each CTRM is aware of the behaviors of other agents so as to make it easier for the subsequent MAPF algorithms to find collision-free paths. (3) *Timed*: Each vertex in CTRMs is augmented in the time direction to represent not only “where,” as commonly done in conventional roadmaps, but also “when,” because solutions of MAPP are a set of “timed” paths. By considering these properties collectively, CTRMs aim to provide a small search space that still contains plausible solutions for MAPF algorithms.

To construct CTRMs, we develop a machine learning (ML)-based approach (Fig. 1). Suppose we are given a collection of MAPP demonstrations (a pair of problem instances and plausible solution paths) obtained by intensive offline computation with conventional roadmaps, *e.g.*, those constructed with uniform random sampling. The proposed ML model learns from how agents behave cooperatively at each unit of time (*i.e.*, timestep), which is implicitly included in the demonstrations, to predict how they will move in

*Work done during an internship at OMRON SINIC X.

[†]Supplementary material is available at <https://omron-sinicx.github.io/ctrm/>.




Figure 1: Learning to construct CTRMs. (a) From MAPF demonstrations, we learn a conditional variational autoencoder (CVAE) that predicts how agents move to their goals while avoiding collisions with others. (b) For a new problem instance, we use the learned model as a vertex sampler F_{CTRM} to construct CTRMs, on which an MAPF algorithm is invoked to derive solution paths efficiently.

the next timesteps (Fig. 1a). For a new, previously unseen problem instance, the learned model can then be used to sample a small set of agent-specific vertices (*i.e.*, space-time pairs) for generating multiple solution path candidates, and eventually, for constructing CTRMs by compositing those candidates. As a result, the CTRMs can serve a smaller but promising search space and enable the MAPF algorithms to derive a solution more efficiently than when using conventional roadmaps (Fig. 1b). Technically, we extend a class of generative models called the conditional variational autoencoder (CVAE) [47], a popular choice for ML sampling-based motion planning [19, 27], to learn a conditional probability distribution of the vertices of CTRMs for each agent, given the observations of multiple agents. Our model can work with an arbitrary number of agents and even within a heterogeneous setting where agents are diverse in their spatial size and kinematic constraints.

We extensively evaluate our approach on a variety of MAPP problems with several different setups in terms of the number of agents (21–40), the presence of obstacles, and the heterogeneity in agent sizes and motion speeds. Our results consistently demonstrate that, compared to standard roadmap construction strategies, planning by learning to construct CTRMs is several orders of magnitude more efficient in the planning effort (*e.g.*, assessed by search node expansions or runtime), while maintaining a comparable planning success rate and solution quality, with acceptable overheads. Our main contributions are as follows. (1) We present cooperative timed roadmaps (CTRMs), a novel concept of graph representations tailored to multi-agent path planning in continuous spaces. (2) We develop an ML-based approach to construct CTRMs and demonstrate its effectiveness in solving a variety of MAPP problems.

2 RELATED WORK

Our study is broadly categorized into multi-agent path planning (MAPP), which has evolved in multiple directions and also shares some techniques with single-agent planning, as described below.

Multi-agent path finding (MAPF) is a problem of MAPP in discrete spaces, with the objective of finding a set of timed paths on graphs while preventing collisions with other agents [51]. MAPF has been studied extensively over the last decade and has produced numerous powerful algorithms (*e.g.*, [3, 29, 39, 44]). More recently, some works have attempted to leverage machine-learning techniques for solving MAPF. These techniques learn from planning demonstrations collected offline to directly predict the next actions of agents given the current observations by means of reinforcement learning [7, 38, 43] or using graph neural networks [33, 34].

Despite such progress, it remains challenging to determine how these techniques should be applied to MAPP in continuous spaces due to the inherent limitation that assumes the search space to be given *a priori* (typically as a grid map).

Path planning in continuous spaces has long been studied for the single-agent case, typically in the form of sampling-based motion planning (SBMP) [11] that iteratively samples state points from the space to seek a solution. Popular approaches include probabilistic roadmaps (PRM) [23], rapidly-exploring random trees (RRT) [31], or their asymptotically optimal versions (PRM*, RRT*) [22]. Most of the works on MAPP in continuous spaces have been built on these techniques. Roadmaps (or simply pre-defined grid maps) are either constructed or given *a priori* to be shared across agents [3, 16, 53, 55] or individually for each agent [14, 28, 48]. A pre-defined set of motion primitives has also been used to discretize possible agent actions [6]. Other works extend RRT* to deal with joint state-spaces [4] or the tensor-product of roadmaps constructed for each agent [45, 49]. In contrast to these previous attempts, we explore a different research direction that delves into what roadmap representations are effective for solving MAPP in continuous spaces and how they should be constructed.

Our approach learns how to construct effective roadmaps from MAPP demonstrations. Such a learning-based approach is gaining attention as a promising method for biased sampling in SBMP [5, 19, 20, 27, 41] to sample state points from important regions of given problem instances to derive a solution efficiently. In contrast, only a few studies have examined how learning can improve roadmaps for multi-agent cases. The most relevant work is [1], which proposed learning to bias sampling of PRMs so as to effectively avoid collisions with dynamic obstacles. However, this approach can only be applied to MAPP with homogeneous agents having the same size and kinematic constraints, since it generates a single roadmap shared across all agents. Moreover, it biases samples based on learned obstacle layouts (*e.g.*, narrow passages), making it hard to run in obstacle-free or random-obstacle environments. Another related work [15] optimizes already constructed roadmaps by means of online learning, which is orthogonal to the above studies that focus on the roadmap construction.

3 PRELIMINARIES

Notations: We consider a problem of path planning for a team of N agents $\mathcal{A} = \{a_1, \dots, a_N\}$ in the 2D continuous space $\mathcal{W} \subset \mathbb{R}^2$; hereafter, we simply refer to it as MAPP. Each agent a_i has a body modeled by a convex region $\mathcal{R}_i(q) \in \mathcal{W}$ around its position $q \in \mathcal{W}$,

which remains fixed regardless of position but may vary from agent to agent. The space \mathcal{W} may contain a set of convex obstacles $\{O_1, \dots, O_M\} \subset \mathcal{W}$, where $\mathcal{O} = \bigcup_{j \in \{1, \dots, M\}} O_j$. Then, the obstacle space for an agent a_i is represented by $C_i^{\text{obs}} = \mathcal{O} \ominus \mathcal{R}_i(\mathbf{0})$, where \ominus is the Minkowski difference and $\mathbf{0}$ is the origin of \mathbb{R}^2 . The free space for a_i is $C_i^{\text{free}} = \mathcal{W} \setminus C_i^{\text{obs}}$. A trajectory for agent a_i is defined by a continuous mapping $\sigma_i : [0, T] \mapsto \mathcal{W}$, where $T \in \mathbb{R}_{>0}$. Each agent a_i has its kinematic constraints \mathcal{K}_i , e.g., each agent has a maximum velocity and moves based on a constant acceleration model.

Problem Description: A problem instance of MAPP is defined by a tuple $I = (\mathcal{A}, \mathcal{W}, \mathcal{O}, \mathcal{R}, \mathcal{K}, \mathcal{S}, \mathcal{G})$, where $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_N\}$ and $\mathcal{K} = \{\mathcal{K}_1, \dots, \mathcal{K}_N\}$. $\mathcal{S} = \{s_1, \dots, s_N \mid s_i \in \mathcal{W}\}$ and $\mathcal{G} = \{g_1, \dots, g_N \mid g_i \in \mathcal{W}\}$ are a set of initial positions and that of goal positions, respectively. A solution for the problem instance is a set of N trajectories $\sigma = \{\sigma_1, \dots, \sigma_N\}$ that satisfy the following four conditions:

- (Endpoint) $\sigma_i(0) = s_i \wedge \sigma_i(T) = g_i$
- (Obstacle) $\sigma_i(\tau) \in C_i^{\text{free}}, 0 \leq \tau \leq T$
- (Inter-agent) $\mathcal{R}_i(\sigma_i(\tau)) \cap \mathcal{R}_j(\sigma_j(\tau)) = \emptyset, i \neq j, 0 \leq \tau \leq T$
- (Kinematic-aware) σ_i satisfies constraints of \mathcal{K}_i

Under these conditions, the quality of solution σ is measured by *sum-of-costs*: $\sum_i T_i$ where $\forall i \in \{1, \dots, N\}, T_i \leq \tau \leq T : \sigma_i(\tau) \in g_i$.

Local Planner: For simplicity, we assume that each agent a_i has a local planner $\mu_i : C_i^{\text{free}} \times C_i^{\text{free}} \times [0, 1] \mapsto C_i^{\text{free}} \cup \{\perp\}$. This continuous mapping interpolates a_i 's locations between distinct space-time pairs under the constraints of \mathcal{K}_i , e.g., $\mu_i(p, q, \epsilon) := (1 - \epsilon)p + \epsilon q$ or Dubins paths [10], while ensuring $\mu_i(p, q, 0) = p$ and $\mu_i(p, q, 1) = q$. When such mapping does not exist, the local planner returns \perp as “not found.” This is a natural assumption used in planning in continuous space either explicitly or implicitly [30].

Timed Roadmaps: To solve MAPP problems, our approach generates a *timed roadmap* D_i for each agent a_i , which approximates the original space C_i^{free} with a finite set of vertices augmented in the time direction, similar to [12, 13]. Each timed roadmap D_i is defined as a directed acyclic graph $D_i = (V_i, E_i)$, where each vertex $v = (p, t) \in V_i$ is a tuple of space $p \in \mathcal{W}$ and discrete time $t \in \mathbb{N}$, representing that an agent a_i is at location p at timestep t . An edge $(u, v) \in E_i$ (to be more precise, arc) exists only when $u = (p, t)$ and $v = (p', t + 1)$. The roadmap D_i is regarded as *consistent* with a_i when (1) $p \in C_i^{\text{free}}$ for all $(p, t) \in V_i$ and (2) the local planner returns a continuous mapping for all $((p, t), (p', t + 1)) \in E_i$.

Multi-Agent Path Planner: On the set of timed roadmaps that are consistent with respective agents, $\{D_1, \dots, D_N\}$, a *multi-agent path planner* is invoked to find a set of agent paths defined on discrete and synchronized time: $\pi = (\pi_1, \dots, \pi_N)$, where $\pi_i = (\pi_i[0], \dots, \pi_i[\lceil T \rceil]) \in (C_i^{\text{free}})^{\lceil T \rceil}$. We can then obtain a trajectory $\sigma_i(\tau)$ that meets the solution conditions by applying the local planner that interpolates between consecutive points in the path π_i . The local planner is also used to check the inter-agent condition. Any typical MAPF algorithm, such as conflict-based search [44] or prioritized planning [3, 12, 46, 53], is potentially applicable for a multi-agent path planner.

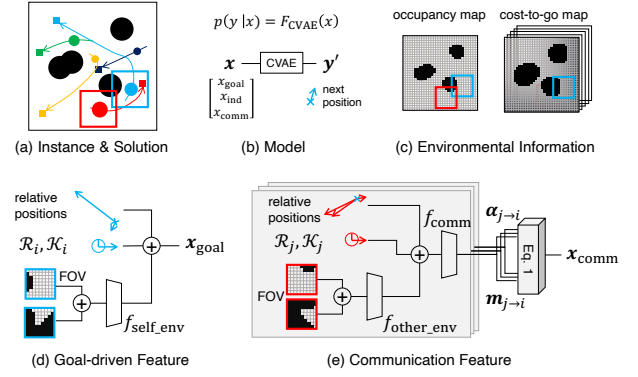


Figure 2: The model architecture and its components. “ \oplus ” represents the concatenation of multiple vectors.

4 LEARNING GENERATIVE MODEL FOR COOPERATIVE TIMED ROADMAPS

Our main technical challenge lies in the construction of timed roadmaps for each agent in a way that effectively reduces the computational effort of multi-agent path planners. We want the roadmaps to provide the planners with a small search space that contains a solution path with plausible quality. To this end, we argue that roadmaps should be “cooperative” in that the planners can easily find a conflict-free path by taking into account the presence of other agents during the roadmap construction. This leads to our proposal of CTRMs, which are agent-specific timed roadmaps aware of agent interactions.

4.1 Model

In this work, we cast roadmap construction as a machine learning problem. Suppose we are given a collection of MAPP demonstrations consisting of problem instances and their solutions by offline intensive computation of MAPF algorithms on sufficiently dense roadmaps constructed with uniform random or grid sampling (Fig. 2a). Using MAPP demonstrations as training data, we learn a generative model that approximates a probability distribution for vertices constituting a solution path for each agent. The learned model can then be used as a “vertex sampler” to construct CTRMs that are more effective for solving a new, previously unseen problem instance efficiently.

As a generative model, we extend the conditional variational auto-encoder (CVAE) [47], which is known to be effective for sequence modeling and prediction [21, 42]. Our CVAE model takes as input (*i.e.*, condition) a feature vector \mathbf{x} extracted from the observations of agent a_i and its neighbors at current timestep t in a solution π to provide a conditional probability distribution $p(\mathbf{y} \mid \mathbf{x})$ of the vector \mathbf{y} that informs how that agent should move in the next timestep $t + 1$ (Fig. 2b). Using a pair of positions $\pi_i[t]$ and $\pi_i[t + 1]$ extracted from solution path π_i , we aim to learn the distribution of $\mathbf{y} = \xi(\pi_i[t + 1] - \pi_i[t])$, where $\xi(v) = [||v||, (v/||v||)^T]^T$ is the magnitude and relative direction of $\pi_i[t + 1]$ with respect to $\pi_i[t]$.

Our model consists of two neural networks: an *encoder* and a *decoder*. The encoder $Enc(\mathbf{x}; \theta)$, parameterized by θ , takes \mathbf{x} as input to produce a conditional probability distribution $p_\theta(z \mid \mathbf{x})$.

The variable z is referred to as a *latent variable* that represents x in a lower-dimensional space. Then, the decoder $Dec(x, z; \phi)$, parameterized by ϕ , accepts x and z drawn from $p_\theta(z | x)$ to form another conditional distribution $p_\phi(y | x, z)$. The output distribution $p(y | x)$ can be obtained by $p(y | x) = \sum_z p_\theta(z | x) p_\phi(y | x, z)$. Doing so is particularly effective when x is high-dimensional, as in our case.

Given a collection of input features x_k and ground-truth outputs y_k , i.e., $\{(x_k, y_k)\}_{k=1}^K$, extracted across agents and timesteps from problem instances in the training data, we can train the encoder and decoder jointly in a supervised learning fashion (see Sec. A in the supplementary material for the details). Once learned, the CVAE can be used as a vertex sampler taking x as input to yield a particular sample y' in accordance with the learned probability distribution $p(y | x)$. In the follows sections, we refer to this sampling process with the learned CVAE as a sampler function $y' \sim F_{CVAE}(x)$.

4.2 Features

Extracting informative features to form x is essential for constructing effective CTRMs. An inherent challenge in MAPP is that each agent needs to move toward its goal while avoiding collisions with other agents that are also in motion. In this section, we first present *goal-driven* features that inform agent a_i of the way to its goal (Sec. 4.2.1). Then, we introduce *communication* features extracted using an attention network that encodes the information of other agents (Sec. 4.2.2). Finally, we propose an *indicator feature* that takes into account high-level choices for which direction to move in when obstacles and other agents are present (Sec. 4.2.3).

4.2.1 Goal-driven Features. The most basic feature that helps agent a_i move toward its goal g_i is the magnitude and relative direction of g_i with respect to the current position $\pi_i[t]$, i.e., $\xi(g_i - \pi_i[t])$. We also include a single-step motion history $\xi(\pi_i[t-1] - \pi_i[t])$, parameters of body region \mathcal{R}_i , and kinematic constraints \mathcal{K}_i , as they affect how agent a_i moves.

Furthermore, we encode environmental information into the field of view (FOV) of agent a_i around its position $\pi_i[t]$ to take into account nearby obstacles. We first discretize the original world $\mathcal{W} \subset \mathbb{R}^2$ by the L^2 grid, where $L \in \mathbb{N}_{>0}$. We then create the FOV with size l^2 , where $l \in \mathbb{N}_{>0}$, centered around $\pi_i[t]$. Within the FOV, we extract two binary maps (illustrated in Fig. 2c: *local occupancy map*) indicating if each grid cell is occupied by obstacles and a *cost-to-go map* that shows if each cell is closer to the goal compared to the current position, pre-computed by means of a breadth-first search on the grid. We feed these two maps to a neural network f_{self_env} to transform them into a compact vector. As shown in Fig. 2d, we then concatenate these features to form a goal-driven feature vector x_{goal} .

4.2.2 Communication Features. To learn the cooperative behaviors between agents observed in ground-truth paths, it is critical to extract features about other agents a_j , such as where they are in the current and previous timesteps, i.e., $\xi(\pi_j[t] - \pi_i[t])$, $\xi(\pi_j[t-1] - \pi_i[t])$, where they will be moving, i.e., their goal position $\xi(g_j - \pi_i[t])$, and how their trajectory could be affected by agent body \mathcal{R}_j and kinematic constraints \mathcal{K}_j . We also take into account the occupancy and cost-to-go maps around the position of a_j by

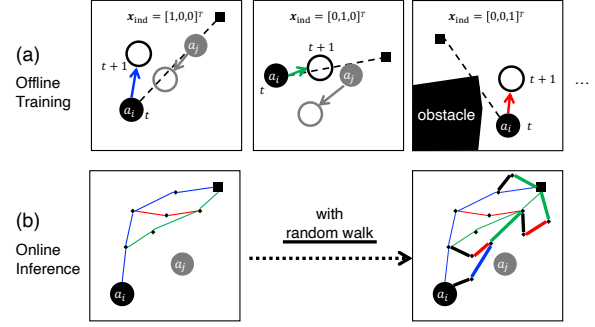


Figure 3: Illustration of the indicator function. Agents are represented by solid black/gray circles.

feeding them into a neural network f_{other_env} to be represented by a compact vector. Let $x_{j \rightarrow i}$ be a feature vector concatenating all this information about a_j with respect to a_i , as shown in Fig. 2e.

Now suppose that we are given a collection of features $\{x_{j \rightarrow i} | j \in Neigh(i)\}$ for all the agents nearby target a_i , where $Neigh(i)$ is a set of indices for the predefined number of neighboring agents of a_i . The question then is how to aggregate them as a part of feature vector x . Obviously, just concatenating them all is not scalable and would not be able to deal with problem instances affecting variable numbers of agents. Instead, we leverage the recent progress in multi-agent interaction modeling that deals with agent communications using an attention network [18]. Specifically, we feed $x_{j \rightarrow i}$ to a neural network f_{comm} that outputs two variables: attention vector $\alpha_{j \rightarrow i}$ and message vector $m_{j \rightarrow i}$. Message vectors are then aggregated while weighted using attention vectors to provide a communication feature vector x_{comm} , as follows:

$$x_{comm} = \sum_{j \in Neigh(i)} m_{j \rightarrow i} \cdot w_{j \rightarrow i}, \quad (1)$$

$$w_{j \rightarrow i} = \frac{\exp[-\|\alpha_{j \rightarrow i} - \alpha_{i \rightarrow i}\|^2]}{\sum_{k \in Neigh(i)} \exp[-\|\alpha_{k \rightarrow i} - \alpha_{i \rightarrow i}\|^2]}$$

where $w_{j \rightarrow i}$ is a scalar weight for the message vector $m_{j \rightarrow i}$, which is defined by the L2 distance between two attention vectors $\alpha_{j \rightarrow i}$ and $\alpha_{i \rightarrow i}$ normalized across $j \in Neigh(i)$ using the softmax function. With Eq. (1), agent a_i can consider message vectors only from selected agents who are close to a_i in terms of attention vectors.

4.2.3 Indicator feature. Figure 3a depicts typical situations where we want to sample the next motion of target agent a_i in the presence of another agent a_j or obstacles in front of a_i . As shown in these examples, agents are often presented with multiple and discrete choices about their moving direction (such as moving left, straight, or right) depending on the layout of the surrounding agents and obstacles. As such, we expect the sampling of next motions to be affected by such high-level choices, which is indeed crucial for constructing effective CTRMs, as we will empirically demonstrate in our experiments.

To this end, we propose augmenting an input feature x with a discrete feature called an *indicator feature* x_{ind} that is learned to indicate promising choices of the next moving directions based on

Algorithm 1 construct_CTRMs

Input: a problem instance \mathcal{I}
Output: a set of timed roadmaps $\mathcal{D} = \{D_1, \dots, D_N\}$
Hyperparameters: $N_{\text{traj}} \in \mathbb{N}_{>0}, T_{\text{max}} \in \mathbb{N}_{>0}$

- 1: $D_i \leftarrow (\{(s_i, 0)\}, \emptyset)$: for each $i = 1, \dots, N$
- 2: $T_{\text{makespan}} \leftarrow 0$
- 3: **for** $j = 1, \dots, N_{\text{traj}}$ **do**
- 4: initialize table $loc \triangleright loc[i, t]$ is location of a_i at timestep t
- 5: $loc[i, 0] \leftarrow s_i$: for each $i = 1, \dots, N$
- 6: **for** $t = 1, \dots, T_{\text{max}} - 1$ **do** \triangleright create one set of trajectories
- 7: **for** $i = 1, \dots, N$ **do**
- 8: $p \leftarrow \text{sample_next_vertex}(\mathcal{I}, t, i, loc) \triangleright p \in C_i^{\text{free}}$
- 9: $q \leftarrow \text{find_compatible_vertex}(\mathcal{I}, t, i, p, D_i)$
- 10: **if** q is not found **then**
- 11: $loc[i, t] \leftarrow p$; insert(p, t, D_i)
- 12: **else**
- 13: $loc[i, t] \leftarrow q$
- 14: **end if**
- 15: **end for**
- 16: **if** check_reachability_to_goals(\mathcal{I}, t, loc) **then**
- 17: $T_{\text{makespan}} \leftarrow \max\{T + 1, T_{\text{makespan}}\}$; **break**
- 18: **end if**
- 19: **end for**
- 20: **end for**
- 21: insert(g_i, t, D_i): for each $t = 1, \dots, T_{\text{makespan}}, i = 1, \dots, N$
- 22: **return** $\mathcal{D} = \{D_1, \dots, D_N\}$

the current observations. Specifically, we define \mathbf{x}_{ind} by the relative direction of $\pi_i[t + 1]$ from $\pi_i[t]$ with respect to the goal direction $g_i - \pi_i[t]$, in the one-hot form such as left, straight, or right, respectively indicated by $[1, 0, 0]^\top$, $[0, 1, 0]^\top$, and $[0, 0, 1]^\top$ (concrete implementations are presented in Sec. 6.2.) While calculating this feature from ground-truth path π_i in the training phase, we also learn a neural network f_{ind} taking the concatenation of \mathbf{x}_{goal} and \mathbf{x}_{comm} as input to predict it. The learned network is then used in the inference phase to provide \mathbf{x}_{ind} , where $\pi_i[t + 1]$ is unknown.

4.2.4 Putting Everything Together. Finally, we concatenate all features to form \mathbf{x} , i.e., $\mathbf{x} = [\mathbf{x}_{\text{goal}}^\top, \mathbf{x}_{\text{comm}}^\top, \mathbf{x}_{\text{ind}}^\top]^\top$. Note that the neural networks used for the feature extraction, i.e., $f_{\text{self_env}}$, $f_{\text{other_env}}$, f_{comm} , and f_{ind} , can be trained end-to-end with the CVAE.

5 CONSTRUCTING COOPERATIVE TIMED ROADMAPS WITH LEARNED MODEL

5.1 Algorithm Overview

In this section, we explain how to construct CTRMs using the learned model as a vertex sampler. Algorithm 1 is our proposed method which takes a problem instance \mathcal{I} as input to build a set of CTRMs consistent with respective agents, i.e., $\mathcal{D} = \{D_1, \dots, D_N\}$.

We generate a sequence of vertices (i.e., path) with a maximum length $T_{\text{max}} \in \mathbb{N}_{>0}$ for each agent using `sample_next_vertex` presented in Sec. B.1 of the supplementary material, which is repeated $N_{\text{traj}} \in \mathbb{N}_{>0}$ times to construct CTRMs. Each path being generated is temporally stored in the table loc , where $loc[i, t]$ refers to the location of agent a_i sampled for timestep t . This table is used by

the learned model in `sample_next_vertex` to sample the next vertices for each agent while being aware of the locations of other agents; we use $loc[i, t]$ as a proxy of $\pi_i[t]$ for the feature extraction described in Sec. 4.2.

During each path generation, we invoke the `insert` function to add a sampled vertex (p, t) to the current CTRM $D_i = (V_i, E_i)$ and update E_i properly (Line 11). This procedure is done only when we confirm that the sampled vertex is not compatible with those we already have in V_i using `find_compatible_vertex`, as explained in Sec. 5.3. Further, we keep updating T_{makespan} to be the smallest number of timesteps to move all the agents to their goals by checking whether they can reach their goals from their last locations with the `check_reachability_to_goals` function (Line 16). After all the iterations, the algorithm inserts each agent’s goal g_i into D_i for time $t = 1, \dots, T_{\text{makespan}}$ (Line 21) and returns a set of CTRMs \mathcal{D} .

5.2 Combining Sampling from Learned Model with Random Walk

In the sub-routine `sample_next_vertex`, we leverage a learned model F_{CTRM} to sample vertices such that CTRMs can efficiently cover a variety of possible solution paths. The complete algorithm description is presented in Sec. B.1 of the supplementary material.

The crucial point here is that we replace the learned model with a random walk centered at the current location ($loc[i, t - 1]$) at probability $1 - P_{\text{model}}$. This contributes to improving the expressiveness of CTRMs beyond what has been learned from the training data, as illustrated in Fig. 3b. Similar techniques are commonly introduced in learning-based sampling for SBMP [5, 19, 20]. We set P_{model} to be low in the initial timestep and gradually increase it from there.

Once a vertex (p, t) is sampled either from the model or randomly, we validate whether this is reachable from the current location $loc[i, t - 1]$ using the local planner. If p is invalid, we repeat the random-walk sampling up to the maximum times specified for resampling, and return $p = loc[i, t - 1]$ (which is trivially valid) if p is still invalid.

5.3 Finding Compatible Vertices

At Line 9 of Alg. 1, we invoke `find_compatible_vertex` to search the current CTRMs for a vertex that is compatible with the sampled vertex (p, t) . By replacing the original vertex (p, t) with a compatible one $(q, t) \in V_i$, we can reduce the time required for connectivity checking in the `insert` function and, more importantly, reduce the search space for multi-agent path planners. The specific algorithm is presented in Sec. B.2 of the supplementary material. Intuitively, two vertices p, q are regarded as compatible when they (1) are spatially close enough and (2) share the same connectivity to other vertices.

6 EVALUATION

In this section, we evaluate CTRMs in various MAPP problems and clarify the merits of our method.

6.1 Experimental Setups

6.1.1 MAPP Problems. Recall that an MAPP problem instance is a tuple $\mathcal{I} = (\mathcal{A}, \mathcal{W}, \mathcal{O}, \mathcal{R}, \mathcal{K}, \mathcal{S}, \mathcal{G})$ (see Sec. 3). For all the MAPP problem instances, we model agent bodies \mathcal{R} and obstacles \mathcal{O} as




Figure 4: Constructed roadmaps for one agent in the Basic scenario. The blue line is a path in a solution. Parameters are: *CTRM*s: $N_{\text{traj}} = 25$; *random*: 3000 samples; *grid*: 32×32 ; *square*: the middle density.

a circle in the 2D closed continuous space ($\mathcal{W} = [0, 1]^2$) to detect collisions easily; hence, an agent body $\mathcal{R}_i \in \mathcal{R}$ is characterized by a scalar value r_i representing the radius. As kinematic constraints $\mathcal{K}_i \in \mathcal{K}$, we assume that each agent has the maximum velocity k_i and moves linearly in two vertices (space-time pairs) at constant velocity. With these settings, we generated 100 different instances, hereafter referred to as *test* instances, for each of the following five scenarios of MAPP (see also the top of Fig. 5).

► **Basic scenario:** A baseline scenario referred to as **(1) Basic** that corresponds to a discrete setting of a 32×32 grid and contains 21–30 homogeneous agents as \mathcal{A} and ten non-uniform obstacles as \mathcal{O} . For each timestep, an agent can move a maximum distance of $k_i = 1/32$. The radius of agents is set to half of their maximum speed, i.e., $r_i = 1/64$. For each problem instance, the number of agents N was randomly determined from the range $N \in \{21, 22, \dots, 30\}$. The initial positions \mathcal{S} , goal positions \mathcal{G} , and the positions of the obstacles were set randomly for each instance

► **Variants of the basic scenario:** Three scenarios that each change one of the settings from those of **Basic**: **(2) More Agents** with $N \in \{31, 32, \dots, 40\}$, **(3) No Obstacles** with no obstacles, and **(4) More Obstacles** with 20 non-uniform obstacles as \mathcal{O} , to see how robust the proposed approach against these parameters.

► **Heterogeneous agent scenario:** An advanced scenario called **(5) Hetero Agents** that contains $N \in \{21, 22, \dots, 30\}$ agents whose size r_i and speed k_i are each multiplied randomly by $\times 1$, $\times 1.25$, or $\times 1.5$ from the original ones defined in **Basic**.

6.1.2 Baselines for Roadmap Construction Methods. To evaluate the effectiveness of CTRMs, we implemented four other roadmap construction methods (see also Fig. 4) that are non-timed and non-cooperative, which have commonly been used for single-agent SBMP and previous work on MAPP in continuous spaces.

► **Random sampling (random)** that samples agent locations uniformly at random from the space. This is equivalent to a simplified version of PRM [22] and has been used as part of the procedure to solve MAPP [48, 53]. The numbers of samples were set to $\{3000, 5000, 7000\}$.

► **Grid sampling (grid)** like the ones the conventional MAPP assumes as a discretized environment [51]. The grid sizes were set to $\{32 \times 32, 64 \times 64, 84 \times 84\}$.

► **SPARS (SPArse Roadmap Spanner algorithm)** [8], an algorithm for roadmap construction that attempts to reduce both vertices and edges. SPARS has been developed for single-agent planning and has also been used in MAPP [16]. We used the implementation in the Open Motion Planning Library [52].

► **Square sampling (square)** as a variant of the random sampling focusing on the square region with its diagonal line given by the start and goal for a single agent (with a margin). We introduce this sampling as a simpler approach to provide a set of agent-specific roadmaps. The number of samples was determined by the length of the diagonal line times a parameter given adaptively to generate low-, middle-, and high-density roadmaps.

More details for SPARS and square sampling are presented in Sec. C.1 of the supplementary material. Given a set of samples as vertices, we built roadmaps by creating edges at two vertices where agents can travel within a single timestep given their size r_i and speed k_i parameters. For the former three methods, a single roadmap was shared across all the agents in the scenarios (1–4) since their size and speed were set identically. For scenario (5) with heterogeneous agents, we constructed different roadmaps for individual agents, as done for CTRMs and the square sampling.

6.1.3 Evaluation Metrics. Since the objective of MAPP is to plan conflict-free trajectories, the effectiveness of roadmap construction methods should be evaluated in terms of the subsequent planning results. We used standard prioritized planning (PP) [46, 53] as the planner and computed the following metrics.

► **Success rate of the planning:** Percentage of successful planning among 100 instances. We regard two cases as a failure: when the PP yielded failure or when the planning time reached a 10-minute timeout.

► **Sum-of-costs:** Solution quality defined in Sec. 3 averaged across all problem instances that resulted in planning success.

► **The number of expanded search nodes in the planning:** Metric of the planning effort (the smaller, the better), which was also averaged across all problem instances that resulted in planning success. We use this metric as a proxy of runtime, since actual runtimes rely heavily on implementations. Specifically, we simply counted the number of search nodes expanded in PP.

► **Runtime:** Reference record of the execution time for the roadmap construction and the subsequent planning.

Note that we implemented all the methods in Python with partial use of C++. The results were obtained on a desktop PC with Intel Core i7-8700 CPU and 32GB RAM.

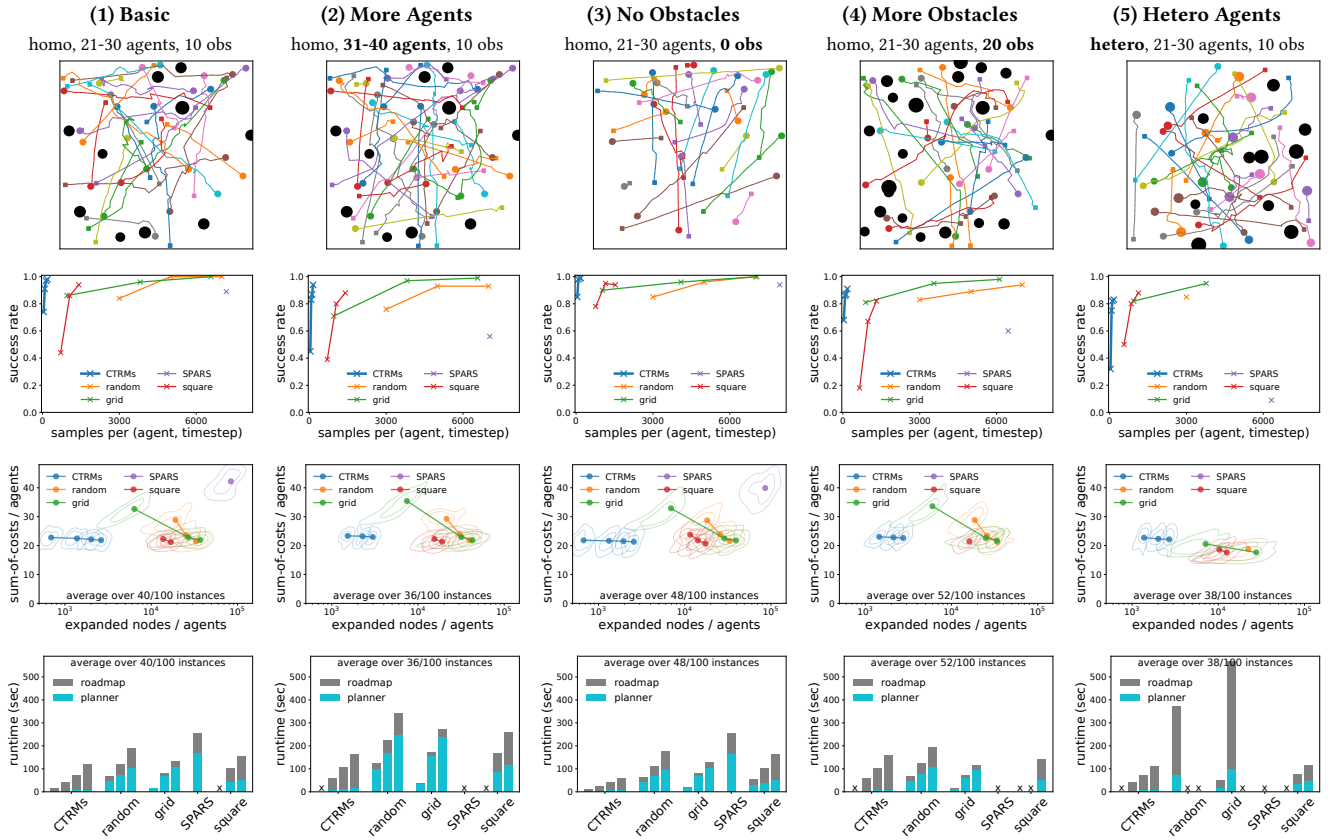


Figure 5: Summary of results. The columns correspond to respective scenarios. For each method, plots to the right correspond to the denser roadmaps (e.g., CTRMs: $N_{\text{traj}} = 25 \rightarrow N_{\text{traj}} = 100$, grid: $32 \times 32 \rightarrow 84 \times 84$). *First row:* Solution examples planned by PP on CTRMs ($N_{\text{traj}} = 25$). *Second row:* Success rate vs. the number of samples per agent and per timestep in the constructed roadmaps. *Third row:* Sum-of-costs vs. the number of expanded search nodes in the planning phase. Each value is normalized by dividing by the number of agents. *Fourth row:* Overall runtime, consisting of the planning and the roadmap construction phases. For the results in the third and fourth rows, in order to obtain meaningful insights, we excluded methods with success rates below 70% and calculated the average of the metrics for problem instances where all the remaining methods were successful in planning. We also omitted some baselines with specific parameter settings resulting in extremely long computation times (over 15 minutes per problem instance) from the evaluation on the Hetero Agent scenario. All these omitted methods and settings are marked with “x” in the fourth row.

6.2 Implementation Details for CTRMs

Model Architectures: All neural networks used in our method are standard multilayer perceptrons (see Sec. C.2 in the supplementary material for details). The FOV was set to $l = 19$ and $L = 160$. The number of neighboring agents of $Neigh(i)$ was set to 15 based on the distance between the current locations of agents, $\pi_i[t]$. The indicator feature \mathbf{x}_{ind} was defined by a three-dimensional one-hot vector indicating left, straight, and right, which were determined based on whether the sine of the angle between two vectors $g_i - \pi_i[t]$ and $\pi_i[t+1] - \pi_i[t]$ was in $[-1, -\frac{1}{3}]$, $(-\frac{1}{3}, \frac{1}{3}]$, or $(\frac{1}{3}, 1]$.

Training Setups: We created 1,000 problem instances for training data and 100 instances for validation, where the trained model with the minimum loss on the validation data was stored and used to construct CTRMs for test instances. Regardless of test scenarios, we used the same training and validation data so that the parameters followed the **Hetero-Agent** scenario. This is because the

Hetero-Agent scenario naturally contains diverse agent behaviors that could also be observed in the other scenarios, which helps reduce the time needed for training with little empirical performance degradation. Note that no identical instances were shared among training, validation, and test instances. Solutions to those instances were obtained by means of roadmap construction based on random sampling with 3,000 samples and path planning by PP. Model parameters were updated using the Adam optimizer [25] with the mini-batch size, number of epochs, and learning rate set to (50, 1000, 0.001). We conducted the data creation and model training on a single workstation equipped with an Intel Xeon CPU E5-2698 v4 and NVIDIA Tesla V100 GPU to complete the procedures in a reasonable amount of time (data creation: 1 day with $\times 40$ multiprocessing; training: 2 hours).

Weighted Loss: We observed that the training and validation data created as above contained many agents that trivially moved

along the shortest path from their start to goal positions, while we expect to observe actions of agents dodging each other to learn their interactions. As such data imbalance typically makes training difficult [35, 36], we introduced a weighted loss technique that gives different weights to each training sample based on the angle Δ between $g_i - \pi_i[t]$ and $\pi_i[t+1] - \pi_i[t]$ with the following criterion: $1 - \exp(-\gamma\Delta^2)$, where $\gamma = 50$. Intuitively, we gave smaller weights to samples that just move agents forward.

Roadmap Construction: We set the parameters for constructing CTRMs in Alg. 1 as follows. N_{traj} was set to $\{25, 50, 75, 100\}$, and T_{max} was 64. In `sample_next_vertex`, we set P_{model} adaptively to timestep t with $1 - \exp(-\gamma t / \min(T_{\text{max}}, T_{\text{makespan}}))$, where $\gamma = 5$, until each agent arrives at its goal. P_{model} was then fixed to 0.1 after agents reach their goal so that they can move around to avoid other agents still moving toward their goal.

6.3 Results

Figure 5 shows our main results, and examples of constructed roadmaps are provided in Fig. 4. Overall, we confirmed that our approach substantially reduced the planning effort to obtain plausible solutions while maintaining a high success rate and sum-of-costs. Specific findings are summarized below.

► *CTRMs contain solutions in a small search space.* The size of the search space can be assessed by the number of vertices in each roadmap per agent and per timestep, as compared to the success rate in the second row of Fig. 5. CTRMs have significantly fewer vertices compared to the others, but still allow the subsequent planner to successfully find a solution at a high rate. The roadmaps created by the random, grid, and SPARS sampling provided much larger search spaces since they are shared among agents and should thus cover the entire space. The square sampling was a bit better than those baselines, but was consistently outperformed by CTRMs.

► *While keeping solution quality, CTRMs contribute to reducing the planning effort by several orders of magnitude* compared to the baselines, as indicated in the third and fourth rows, simply thanks to the small search space. The sum-of-cost metric normalized by the number of agents was comparable to the baseline methods and even better than when the number of samples for the baseline methods is limited. This is primarily because our approach learns where to sample vertices from the solution paths of MAPP demonstrations.

► *CTRMs achieve efficient MAPP solving from the end-to-end perspective.* The total runtime with the overhead from roadmap constructions is shown in the fourth row. Although the runtime relies heavily on the implementation, our method can produce CTRMs in a realistic computation time. Once CTRMs are constructed, the planning can be finished immediately, unlike the baselines, despite the use of the same planning algorithm and implementation. Even so, there is room for further optimization in the roadmap construction, including improvement to the implementation of connectivity checks between vertices, as it is often a bottleneck of SBMP [11].

► *Our approach is generalizable for various scenarios.* Even though our model was trained on just a single scenario with heterogeneous agents, the model once trained can be used for other scenarios with different numbers of agents and obstacles.

Table 1: Results of ablation study performed on the Basic scenario. Sum-of-costs and expanded nodes are normalized by the number of agents and averaged over 11 instances that succeeded in all methods except “no random walk” (resulting in an extremely low success rate), accompanied by 95% confidence intervals.

	Success rate	Sum-of-costs	Expanded nodes
CTRMs ($N_{\text{traj}} = 25$)	0.80	21.2 (20.3, 22.0)	612.4 (547.4, 674.2)
no x_{comm}	0.23	28.7 (27.3, 30.2)	996.3 (923.1, 1068.5)
no x_{ind}	0.33	31.3 (30.5, 32.2)	1058.6 (993.8, 1117.3)
no random walk	0.03	N/A	N/A

Ablation Study: We further assessed the effect of each technical component used in our approach. Specifically, we evaluated degraded versions of our approach that omit 1) communication features (no x_{comm}) or 2) indicator features (no x_{ind}) when learning the model, and 3) no random walk (*i.e.*, $P_{\text{model}} = 1$) when constructing CTRMs. All the variants were evaluated on the Basic scenario in terms of success rate, sum-of-costs, and the number of expanded nodes with $N_{\text{traj}} = 25$. As shown in Table 1, involving the random walk is particularly critical to ensure a high success rate. This is consistent with the common practice of previous studies [5, 19, 20] that use learned sampling and random sampling together. On top of that, the lack of communication feature x_{comm} greatly limits the success rate, sum-of-costs, and the number of expanded nodes, as agents are then required to avoid others solely by random walk. Furthermore, we confirmed that the indicator feature x_{ind} is equally important to improve the overall performance.

Limitations: Although CTRMs generally worked well in various scenarios, we informally confirmed that the method sometimes failed in “bug trap” situations where SBMP generally struggles [30]. In general, our approach is applicable only if each agent can find an obstacle-free path to reach the goal when constructing CTRMs. Failing to do so will make it impossible to perform MAPF in the later stage. One possible resolution would be to introduce bi-directional path generation like a variant of RRT [26]. Improving the overall performances by replacing the multi-agent path planner from PP is a promising direction for future work.

7 CONCLUSION

This paper presented cooperative timed roadmaps (CTRMs), a novel graph representation tailored to multi-agent path planning in continuous spaces. The CTRMs are constructed to provide the subsequent planner with a small search space for each agent while simultaneously being aware of other agents to avoid potential inter-agent collisions and contain plausible solution paths. This can be done by learning a generative model from relevant demonstrations and using it as an effective vertex sampler. The experimental results demonstrate the effectiveness of CTRMs for a variety of MAPP problems. Future work will seek to extend the proposed method to anytime planning, planning in continuous time, or multi-agent motion planning in higher-dimensional spaces.

ACKNOWLEDGMENTS

We are grateful to Kazumi Kasaura for his insightful comments throughout this study.

REFERENCES

- [1] Felipe Felix Arias, Brian Andrew Ichter, Aleksandra Faust, and Nancy M. Amato. 2021. Avoidance Critical Probabilistic Roadmaps for Motion Planning in Dynamic Environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- [2] Jacopo Banfi, Nicola Basilico, and Francesco Amigoni. 2017. Intractability of Time-Optimal Multirobot Path Planning on 2D Grid Graphs with Holes. *IEEE Robotics and Automation Letters (RA-L)* 2, 4 (2017), 1941–1947.
- [3] Michal Čáp, Peter Novák, Alexander Kleiner, and Martin Selecký. 2015. Prioritized Planning Algorithms for Trajectory Coordination of Multiple Mobile Robots. *IEEE Transactions on Automation Science and Engineering (T-ASE)* 12, 3 (2015), 835–849.
- [4] Michal Čáp, P. Novák, J. Vokřínek, and M. Pechoucek. 2013. Multi-Agent RRT: Sampling-based Cooperative Pathfinding. In *Proceedings of the International Joint Conference on Autonomous Agents & Multiagent Systems (AAMAS)*.
- [5] Binghong Chen, Bo Dai, Qinjie Lin, Guo Ye, Han Liu, and Le Song. 2020. Learning to Plan in High Dimensions via Neural Exploration-Exploitation Trees. In *Proceedings of the International Conference on Learning and Representation (ICLR)*.
- [6] Liron Cohen, Tansel Uras, TK Satish Kumar, and Sven Koenig. 2019. Optimal and Bounded-Suboptimal Multi-Agent Motion Planning. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*.
- [7] Mehul Damani, Zhiyao Luo, Emerson Wenzel, and Guillaume Sartoretti. 2021. PRIMAL_2: Pathfinding Via Reinforcement and Imitation Multi-Agent Learning-Lifelong. *IEEE Robotics and Automation Letters (RA-L)* 6, 2 (2021), 2666–2673.
- [8] Andrew Dobson and Kostas E Bekris. 2014. Sparse Roadmap Spanners for Asymptotically Near-Optimal Motion Planning. *International Journal of Robotics Research (IJRR)* 33, 1 (2014), 18–47.
- [9] Kurt Dresner and Peter Stone. 2008. A Multiagent Approach to Autonomous Intersection Management. *Journal of Artificial Intelligence Research (JAIR)* 31 (2008), 591–656.
- [10] Lester E Dubins. 1957. On Curves of Minimal Length with a Constraint on Average Curvature, and with Prescribed Initial and Terminal Positions and Tangents. *American Journal of Mathematics* 79, 3 (1957), 497–516.
- [11] Mohamed Elbanhawi and Milan Simic. 2014. Sampling-based Robot Motion Planning: A Review. *IEEE Access* 2 (2014), 56–77.
- [12] Michael Erdmann and Tomas Lozano-Perez. 1987. On Multiple Moving Objects. *Algorithmica* 2, 1 (1987), 477–521.
- [13] Thierry Fraichard. 1998. Trajectory Planning in a Dynamic Workspace: A 'State-Time Space' Approach. *Advanced Robotics* 13, 1 (1998), 75–94.
- [14] Mokhtar Gharbi, Juan Cortés, and Thierry Siméon. 2009. Roadmap Composition for Multi-Arm Systems Path Planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2471–2476.
- [15] Christian Henkel and Marc Toussaint. 2020. Optimized Directed Roadmap Graph for Multi-Agent Path Finding using Stochastic Gradient Descent. In *Proceedings of the Annual ACM Symposium on Applied Computing (SAC)*. 776–783.
- [16] Wolfgang Hönig, James A Preiss, TK Satish Kumar, Gaurav S Sukhatme, and Nora Ayanian. 2018. Trajectory Planning for Quadrotor Swarms. *IEEE Transactions on Robotics (T-RO)* 34, 4 (2018), 856–869.
- [17] John E Hopcroft, Jacob Theodore Schwartz, and Micha Sharir. 1984. On the Complexity of Motion Planning for Multiple Independent Objects; PSPACE-Hardness of the Warehouseman's Problem. *International Journal of Robotics Research (IJRR)* 3, 4 (1984), 76–88.
- [18] Yedid Hoshen. 2017. VAIN: Attentional Multi-Agent Predictive Modeling. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*. 2698–2708.
- [19] Brian Ichter, James Harrison, and Marco Pavone. 2018. Learning Sampling Distributions for Robot Motion Planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 7087–7094.
- [20] Brian Ichter, Edward Schmerling, Tsang-Wei Edward Lee, and Aleksandra Faust. 2020. Learned Critical Probabilistic Roadmaps for Robotic Motion Planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 9535–9541.
- [21] Boris Ivanovic, Karen Leung, Edward Schmerling, and Marco Pavone. 2021. Multimodal Deep Generative Models for Trajectory Prediction: A Conditional Variational Autoencoder Approach. *IEEE Robotics and Automation Letters (RA-L)* 6, 2 (2021), 295–302.
- [22] Sertac Karaman and Emilio Frazzoli. 2011. Sampling-based Algorithms for Optimal Motion Planning. *International Journal of Robotics Research (IJRR)* 30, 7 (2011), 846–894.
- [23] Lydia E Kavvaki, Petr Svestka, J-C Latombe, and Mark H Overmars. 1996. Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. *IEEE Transactions on Robotics and Automation* 12, 4 (1996), 566–580.
- [24] Mohammad Khayatian, Mohammadreza Mehrabian, Edward Andert, Rachel Dedinsky, Sarthake Choudhary, Yingyan Lou, and Aviral Shirvastava. 2020. A Survey on Intersection Management of Connected Autonomous Vehicles. *ACM Transactions on Cyber-Physical Systems* 4, 4 (2020), 1–27.
- [25] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the International Conference on Learning and Representation (ICLR)*.
- [26] James J Kuffner and Steven M LaValle. 2000. RRT-Connect: An Efficient Approach to Single-Query Path Planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 2. 995–1001.
- [27] Rahul Kumar, Aditya Mandalika, Sanjiban Choudhury, and Siddhartha Srinivasa. 2019. LEGO: Leveraging Experience in Roadmap Generation for Sampling-based Planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1488–1495.
- [28] Sandip Kumar and Suman Chakravorty. 2012. Multi-Agent Generalized Probabilistic Roadmaps: MAGPRM. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3747–3753.
- [29] Edward Lam, Pierre Le Bodic, Daniel Damir Harabor, and Peter J Stuckey. 2019. Branch-and-Cut-and-Price for Multi-Agent Pathfinding. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 1289–1296.
- [30] Steven M LaValle. 2006. *Planning algorithms*. Cambridge University Press.
- [31] Steven M LaValle et al. 1998. *Rapidly-Exploring Random Trees: A New Tool for Path Planning*. Technical Report. Computer Science Department, Iowa State University (TR 98–11).
- [32] Jiaoyang Li, Wheeler Ruml, and Sven Koenig. 2021. EECBS: A Bounded-Suboptimal Search for Multi-Agent Path Finding. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- [33] Qingbiao Li, Fernando Gama, Alejandro Ribeiro, and Amanda Prorok. 2020. Graph Neural Networks for Decentralized Multi-Robot Path Planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 11785–11792.
- [34] Qingbiao Li, Weizhe Lin, Zhe Liu, and Amanda Prorok. 2021. Message-aware Graph Attention Networks for Large-scale Multi-Robot Path Planning. *IEEE Robotics and Automation Letters (RA-L)* 6, 3 (2021), 5533–5540.
- [35] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal Loss for Dense Object Detection. In *Proceeding of the International Conference on Computer Vision (ICCV)*. 2980–2988.
- [36] Xiankai Lu, Chao Ma, Bingbing Ni, Xiaokang Yang, Ian Reid, and Ming-Hsuan Yang. 2018. Deep Regression Tracking with Shrinkage Loss. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 353–369.
- [37] Hang Ma, Craig Tovey, Guni Sharon, TK Satish Kumar, and Sven Koenig. 2016. Multi-Agent Path Finding with Payload Transfers and the Package-Exchange Robot-Routing Problem. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- [38] Ziyuan Ma, Yudong Luo, and Hang Ma. 2021. Distributed Heuristic Multi-Agent Path Finding with Communication. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- [39] Keisuke Okumura, Manao Machida, Xavier Défago, and Yasumasa Tamura. 2019. Priority Inheritance with Backtracking for Iterative Multi-agent Path Finding. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 535–542.
- [40] Keisuke Okumura, Yasumasa Tamura, and Xavier Défago. 2021. Iterative Refinement for Real-Time Multi-Robot Path Planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [41] Ahmed H Qureshi, Yinglong Miao, Anthony Simeonov, and Michael C Yip. 2020. Motion Planning Networks: Bridging the Gap Between Learning-based and Classical Motion Planners. *IEEE Transactions on Robotics (T-RO)* (2020), 1–9.
- [42] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. 2020. Trajectron++: Dynamically-Feasible Trajectory Forecasting with Heterogeneous Data. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 683–700.
- [43] Guillaume Sartoretti, Justin Kerr, Yunfei Shi, Glenn Wagner, TK Satish Kumar, Sven Koenig, and Howie Choset. 2019. PRIMAL: Pathfinding via Reinforcement and Imitation Multi-Agent Learning. *IEEE Robotics and Automation Letters (RA-L)* 4, 3 (2019), 2378–2385.
- [44] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. 2015. Conflict-based Search for Optimal Multi-Agent Pathfinding. *Artificial Intelligence (AIJ)* 219 (2015), 40–66.
- [45] Rahul Shome, Kiril Solovey, Andrew Dobson, Dan Halperin, and Kostas E Bekris. 2020. dRRT*: Scalable and Informed Asymptotically-Optimal Multi-Robot Motion Planning. *Autonomous Robots* 44, 3 (2020), 443–467.
- [46] David Silver. 2005. Cooperative Pathfinding. *Proceedings of the Artificial Intelligence for Interactive Digital Entertainment Conference (AIIDE)* (2005), 117–122.
- [47] Kihyuk Sohn, Honglak Lee, and Kinchen Yan. 2015. Learning Structured Output Representation using Deep Conditional Generative Models. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, Vol. 28. 3483–3491.
- [48] Irving Solis, James Motes, Read Sandström, and Nancy M Amato. 2021. Representation-Optimal Multi-Robot Motion Planning using Conflict-based Search. *IEEE Robotics and Automation Letters (RA-L)* 6, 3 (2021), 4608–4615.
- [49] Kiril Solovey, Oren Salzman, and Dan Halperin. 2016. Finding a Needle in an Exponential Haystack: Discrete RRT for Exploration of Implicit Roadmaps in Multi-Robot Motion Planning. *International Journal of Robotics Research (IJRR)* 35, 5 (2016), 501–513.

- [50] P. Spirakis and C. Yap. 1984. Strong NP-Hardness of Moving Many Discs. *Inform. Process. Lett.* 19 (1984), 55–59.
- [51] Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Kumar, et al. 2019. Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*. 151–159.
- [52] Ioan A. Şucan, Mark Moll, and Lydia E. Kavraki. 2012. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine (RAM)* 19, 4 (2012), 72–82.
- [53] Jur P Van Den Berg and Mark H Overmars. 2005. Prioritized Motion Planning for Multiple Robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 430–435.
- [54] Peter R Wurman, Raffaello D’Andrea, and Mick Mountz. 2008. Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses. *AI Magazine* 29, 1 (2008), 9.
- [55] Konstantin Yakovlev and Anton Andreychuk. 2017. Any-Angle Pathfinding for Multiple Agents based on SIPP Algorithm. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*.
- [56] Jingjin Yu and Steven M LaValle. 2013. Structure and Intractability of Optimal Multi-Robot Path Planning on Graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.