

Multi-Goal Motion Memory

Yuanjie Lu¹, Dibyendu Das¹, Erion Plaku², and Xuesu Xiao¹

Abstract—Autonomous mobile robots (e.g., warehouse logistics robots) often need to traverse complex, obstacle-rich, and changing environments to reach multiple fixed goals (e.g., warehouse shelves). Traditional motion planners need to calculate the entire multi-goal path from scratch in response to changes in the environment, which result in a large consumption of computing resources. This process is not only time-consuming but also may not meet real-time requirements in application scenarios that require rapid response to environmental changes. In this paper, we provide a novel Multi-Goal Motion Memory technique that allows robots to use previous planning experiences to accelerate future multi-goal planning in changing environments. Specifically, our technique predicts collision-free and dynamically-feasible trajectories and distances between goal pairs to guide the sampling process to build a roadmap, to inform a Traveling Salesman Problem (TSP) solver to compute a tour, and to efficiently produce motion plans. Experiments conducted with a vehicle and a snake-like robot in obstacle-rich environments show that the proposed Motion Memory technique can substantially accelerate planning speed by up to 90%. Furthermore, the solution quality is comparable to state-of-the-art algorithms and even better in some environments.

I. INTRODUCTION

Robots used in applications such as logistics, inspection, surveillance, and transportation often have to reach multiple fixed locations to carry out assigned tasks [1], [2]. Their workspaces, such as warehouses, factories, and power plants, are usually unstructured and contain numerous obstacles, requiring the robots to navigate around them and pass through several narrow passages. Therefore, working in these settings give rise to challenging multi-goal motion-planning problems, where robots need to quickly plan their motions to reach the goal locations while avoiding obstacles.

In a multi-goal motion-planning setting, a robot has to reason at a high level to determine the order in which to reach the goals, bearing similarities with Traveling Salesman Problems (TSPs). However, in this case, the costs and trajectories of moving from one location to another are not known beforehand, as using straight-line paths and Euclidean distances as a proxy does not account for the obstacles or robot dynamics (i.e., limiting the direction of motion, velocity, turn radius, acceleration, etc.). Therefore, multi-goal motion planning requires solving numerous single-goal motion-planning problems. Another challenge is that high-level reasoning and motion planning cannot be decoupled.

¹Department of Computer Science, George Mason University, Fairfax, VA 22030, USA.

²National Science Foundation, Alexandria, VA 22314, USA

The work by E. Plaku is supported by (while serving at) the National Science Foundation. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

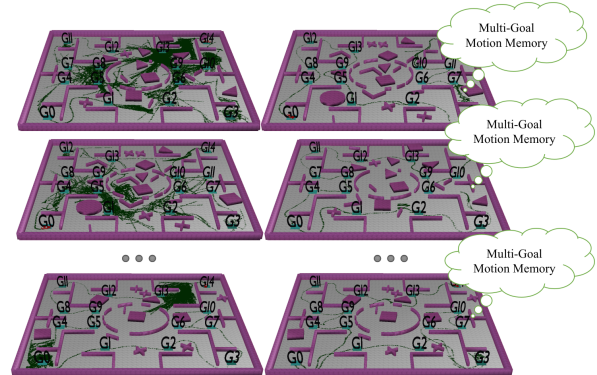


Fig. 1. In multi-goal motion planning, the robot must visit each goal ($G_i, \forall i$) while avoiding collisions. Classical multi-goal motion planners require extensive computation when facing a new planning problem (shown in green), which can be significantly reduced by Multi-Goal Motion Memory.

Computing a goal ordering and then running a motion planner to reach the goals in succession will likely lead to suboptimal solutions, since the goal ordering has not necessarily accounted for the obstacles and robot dynamics. Thus, the vast spaces of feasible motion plans and goal orderings must be simultaneously explored.

To make matters worse, while classical motion planners can solve such multi-goal motion planning problems in conjunction with TSP solvers [3] and satisfy the physical constraints imposed by the robot dynamics, they need to start planning from scratch each time the environment changes, regardless of how small the change is (e.g., some items are moved on a factory floor). Such repetitive planning not only leads to a waste of computation but also limits the robot’s ability to respond quickly when facing environmental changes, which are not uncommon in real-world applications.

To address these challenges, we present a new technique called Multi-Goal Motion Memory, which incorporates past planning experiences into a planner and uses this prior knowledge to accelerate new planning tasks (Fig. 1). Specifically, leveraging historical planning experiences, Multi-Goal Motion Memory quickly generates the best trajectory with dynamics and the predicted distance between each goal pair to guide the sampling process for building a new motion map for the new environment. Then, utilizing the generated goal-to-goal distances, a TSP solver efficiently computes a low-cost tour that determines the optimal order to visit each goal. After determining the order of accessing the goals, the predicted best trajectory is used to guide the motion-tree expansion in a motion planner, which allows the motion tree to more effectively explore low-cost areas, considerably

reducing the motion-planning runtime. Experiments are conducted in simulation using a vehicle and snake-like robot operating in unstructured, obstacle-rich environments with an increasing number of goals. The results show that Multi-Goal Motion Memory can significantly accelerate planning speed by up to 90%. Moreover, the plan quality in terms of solution distance remains comparable to state-of-the-art algorithms and is even better in many environments.

II. RELATED WORK

We review related work in multi-goal motion planning and machine learning for motion planning.

A. Multi-Goal Motion Planning

Multi-goal motion planning is the problem of finding a collision-free and dynamically-feasible trajectory to visit a set of goals. Researchers have addressed different aspects of this problem. At a high level, which does not account for the obstacles or the vehicle dynamics, the focus has been on computing optimal tours based on goal-pair distances via TSP solvers [4]–[8]. When considering obstacles but not robot dynamics, general motion planners often construct a roadmap based on PRM [9]. Researchers have used TSP solvers on a cost matrix based on the shortest path distance to find paths to solve multi-goal motion planning [10]–[12]. TSP solvers have also been combined with controllers for Dubins vehicles [13], [14]. Sampling-based motion planning [15]–[17] has often been used for solving multi-goal problems with robot dynamics.

Despite the potential similarity among planning problems in the workspace, these classical multi-goal motion planners re-plan from scratch in response to any environmental changes. Multi-Goal Motion Memory is designed to address such a limitation by leveraging previous experiences to more efficiently guide the current motion-tree expansion.

B. Machine Learning for Motion Planning

In recent years, researchers have started to use machine learning to solve single-goal motion planning [18]–[26]. Reinforcement and imitation learning [19] has been used for motion planning in complex environments, involving highly-cluttered obstacles [24]–[33], dynamic decision-making agents [34]–[38], and off-road conditions [39]–[42]. The work in [43] presents a neural-network-based planning algorithm to encode the given workspaces from a point cloud measurement and generate collision-free paths from the start to the goal. The work in [44] combines machine learning and geometric approximation to build a new collision-free configuration space, using parallel k -nearest neighbors and parallel collision detection algorithms to accelerate motion planning. In terms of sampling-based motion planners, the work in [45] presents a policy-search-based method to learn implicit sampling distributions for different environments, making it possible to generate better distributions and reduce cost. Machine learning has also been used to guide the motion-tree expansion by first predicting if a given sample is in a collision and then estimating if it is a promising sample

integrated into RRT [46]. The work in [47] proposed the idea of training a model to predict the runtime on single-goal problems, which helps to facilitate the construction of cost matrices for multi-goal problems.

The most relevant work is our previous approach [48] to improve the planning efficiency of DROMOS [49]. Given the same environment, the work in [48] learns to predict goal-to-goal distances for any arbitrary goal pair using a large number of collision-free configurations and motion plans by DROMOS. Our Multi-Goal Motion Memory is based on the observation that in many real-world robotics applications, the goal locations remain constant while the environment keeps changing, necessitating a way to quickly plan in response to different obstacle configurations.

III. PRELIMINARIES

This section defines the problem of multi-goal motion planning with robot dynamics. It also presents a general pipeline for many classical sampling-based multi-goal motion planners and sets the stage for potential data-driven improvement. The section concludes with a description of the single-start-goal motion memory technique, which our Multi-Goal Motion Memory generalizes.

A. Problem Definition

A robot operates in a world \mathcal{W} containing obstacles $\mathcal{O} = \{\mathcal{O}_1, \dots, \mathcal{O}_m\}$ and goal regions $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_n\}$, as shown in Fig. 1. The robot is modeled as $\mathcal{R} = \langle \mathcal{P}, \mathcal{S}, \mathcal{U}, f \rangle$ in terms of its geometric shape \mathcal{P} , state space \mathcal{S} , action space \mathcal{U} , and dynamics f expressed as a set of differential equations:

$$\dot{s} = f(s, a), s \in \mathcal{S}, a \in \mathcal{U}.$$

A vehicle and snake-like robot are shown in Fig. 2 as examples. The vehicle robot is defined as $s = (x, y, \theta, \psi, v)$ in terms of the position (x, y) , orientation θ , steering angle ψ ($|\psi| \leq 1.5\text{rad}$), and speed v ($|v| \leq 2.25\text{m/s}$). The vehicle is controlled by setting the acceleration a_{acc} ($|a_{\text{acc}}| \leq 1\text{m/s}^2$) and the steering turning rate a_{ω} ($|\omega| \leq 2.7\text{rad/s}$). The motion equations are defined as

$$\begin{aligned} \dot{x} &= v \cos(\theta) \cos(\psi), \dot{y} = v \sin(\theta) \cos(\psi), \\ \dot{\theta} &= v \sin(\psi)/L, \dot{v} = a_{\text{acc}}, \dot{\psi} = a_{\omega}, \end{aligned}$$

where L is the distance from the back to the front wheels. The snake-like robot, modeled as a car pulling trailers, is defined as $s = (x, y, \theta, \psi, v, \theta_1, \dots, \theta_N)$, and f includes

$$\dot{\theta}_i = \frac{v}{H} (\sin \theta_{i-1} - \sin \theta_0) \prod_{j=1}^{i-1} \cos(\theta_{j-1} - \theta_j),$$

where $\theta_0 = \theta$ and the hitch distance H is set to a small value so that the robot resembles a snake.

A state $s \in \mathcal{S}$ reaches a goal \mathcal{G}_i if the position defined by s is inside \mathcal{G}_i . The state s is in collision if the robot overlaps with an obstacle when placed according to the position and orientation defined by s .

Algorithm 1 General Multi-Goal Motion Planner

Input: world \mathcal{W} ; obstacles \mathcal{O} ; goals \mathcal{G} ; robot model \mathcal{R} ; initial state s_{init} ; TSPSOLVER

Output: collision-free and dynamically-feasible trajectory that starts at s_{init} and reaches each goal in \mathcal{G}

- 1: $\mathcal{RM} \leftarrow \text{CREATEROADMAP}(\mathcal{W}, \mathcal{O}, \mathcal{G}, s_{\text{init}})$
 - 2: $\text{tour} \leftarrow \text{TSPSOLVER}(\mathcal{RM}, \mathcal{G}, \mathcal{R}, s_{\text{init}})$
 - 3: $p \leftarrow \text{PLANWITHMOTIONTREE}(\text{tour}, \mathcal{R})$
-

Applying a control action $a \in \mathcal{U}$ to a state $s \in \mathcal{S}$ yields a new state $s_{\text{new}} \in \mathcal{S}$, which is obtained by numerically integrating the motion equations f for one time step dt , i.e.,

$$s_{\text{new}} \leftarrow \text{SIMULATE}(s, a, f, dt).$$

A dynamically-feasible trajectory $\zeta : \{0, \dots, \ell\} \rightarrow \mathcal{S}$ is obtained by applying a sequence of actions $\langle a_0, \dots, a_{\ell-1} \rangle$ in succession, where $\zeta_0 \leftarrow s$ and $\forall j \in \{1, \dots, \ell\}$:

$$\zeta_j \leftarrow \text{SIMULATE}(\zeta_{j-1}, a_{j-1}, f, dt). \quad (1)$$

The multi-goal motion-planning problem is defined as follows: Given a world \mathcal{W} containing obstacles $\mathcal{O} = \{\mathcal{O}_1, \dots, \mathcal{O}_m\}$ and goals $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_n\}$, a robot model $\mathcal{R} = \langle \mathcal{P}, \mathcal{S}, \mathcal{U}, f \rangle$, and an initial state $s_{\text{init}} \in \mathcal{S}$, compute a collision-free and dynamically-feasible trajectory $\zeta : \{0, \dots, \ell\} \rightarrow \mathcal{S}$ that starts at s_{init} and reaches each goal. The objective of this work is to reduce the planning time.

B. Classical Sampling-Based Multi-Goal Motion Planning

Classical sampling-based multi-goal motion planning is primarily divided into three parts (Algorithm 1). In line 1, a roadmap is constructed and represented as an undirected weighted graph by sampling collision-free states. The roadmap retains information about the coordinates of each state and the cost of each edge. Subsequently, through Dijkstra’s algorithm, we can obtain the path from each collision-free state to any goal and its cost. In line 2, the cost from the robot position to each goal and from each goal to every other goal is sent to a TSP solver, which then computes a tour visiting all the goals. In line 3, a motion planner will plan the path to reach all the goals based on the order specified by the tour. Our Multi-Goal Motion Memory aims at improving all three parts in a data-driven manner.

C. Single-Start-Goal Motion Memory

Motion Memory is a recently proposed approach [50] with an experience augmentation technique and a representation learning method that enable robots to reflect on prior planning experiences for efficient future planning for a single start-goal pair [50]. Given a set of existing motion planning experience for a fixed start and goal, i.e., a mapping from an obstacle configuration to a motion plan, motion memory first produces a variety of similar obstacle configurations where the same motion plan would still be optimal. Subsequently, this augmented dataset is utilized to learn an efficient latent representation space by minimizing a triplet loss function. As

a result, the representation space is organized into several latent clusters, with each cluster representing a distinct motion plan from past experiences. When a new planning problem arises, it is projected into this latent space. The motion plan associated with the nearest cluster centroid is then considered the most appropriate approximation to the actual motion planning solution for the new problem.

Our Multi-Goal Motion Memory is a generalization of the specific single-goal case. By generalizing motion memory to consider different start-goal pairs, it can be efficiently used to solve multi-goal motion planning problems by accelerating the solution time of multiple components (Algorithm 1).

IV. METHODOLOGY

Motivated by many real-world multi-goal motion planning problems with a fixed goal set (e.g., stations or targets the robot needs to visit) but changing obstacles, Multi-Goal Motion Memory aims to leverage past planning experiences to accelerate future planning given new environments. In principle, Multi-Goal Motion Memory can be used in conjunction with any multi-goal motion planner. We first introduce Multi-Goal Motion Memory in a planner-agnostic manner, and then we integrate Multi-Goal Motion Memory into sampling-based multi-goal motion planners based on TSP solvers following Algorithm 1 at different stages.

A. Multi-Goal Motion Memory

A particular environment’s State Space (S-space, S) can be decomposed into reachable (S_{free}) and unreachable (S_{obst} determined by \mathcal{O}) states, considering various constraints like obstacles, nonholonomic constraints, and velocity limits. Given a start state ($S_s \in S$) and a goal state ($S_g \in S$), a motion plan, denoted as $P = \{a_i | 1 \leq i \leq T\}$, where $a_i \in \mathcal{U}$, drives the robot through S_{free} by optimizing a cost function and observing robot dynamics.

The previous single-start-goal motion memory assumes all planning problems have the same start and goal pair. Therefore, solving a motion planning problem is to produce a motion planning solution based on the S-space alone. A dataset of motion planning problems and solutions for a fixed start-goal pair becomes $\mathcal{D}^S = \{d_i\}_{i=1}^N = \{S_i, P_i\}_{i=1}^N$ and the motion planner becomes a function $f(\cdot)$ that maps from S-spaces to motion plans, i.e., $P = f(S)$. Single-start-goal motion memory also utilizes an augmentation technique by hallucinating other M_i planning problems S_i^j for which the existing motion plan P_i would also be optimal: $\mathcal{D}^{S^*} = \{d_i^*\}_{i=1}^N = \{\{S_i^j\}_{j=1}^{M_i}, P_i\}_{i=1}^N$.

However, for multi-goal motion planning, such an assumed mapping from S-spaces alone to motion plans does not hold. A motion planner needs to map from an S-space, start s_s , and goal s_g to a motion plan P , i.e., $P = f(S, s_s, s_g)$. Thus, we generalize the previous augmented single-start-goal motion memory dataset \mathcal{D}^{S^*} into a multi-start-goal dataset:

$$\begin{aligned} \mathcal{D}^{M^*} &= \{k \mathcal{D}^{S^*}\}_{k=1}^K = \{\{k d_i^*\}_{i=1}^{kN}\}_{k=1}^K \\ &= \{\{\{k S_i^j\}_{j=1}^{kM_i}, k P_i\}_{i=1}^{kN}\}_{k=1}^K, \end{aligned}$$

where K is the number of start-goal pairs, $K = |\mathcal{G}|P_2$ (permutation of selecting and ordering 2 goals from $|\mathcal{G}|$ goals). Left subscript k denotes the k th start-goal pair. Each start and goal are denoted as ${}_k s_s \in \mathcal{G}$ and ${}_k s_g \in \mathcal{G}$.

The augmented dataset (\mathcal{D}^{M*}) is used to learn the motion memory so that given a specific start-goal pair, s_s and s_g , and a S-space S it will predict a most likely collision-free and near-optimal path P . Despite many different data-driven approaches to approximate this function, inspired by the success of single-start-goal motion memory, we utilize a similar representation learning technique for each start-goal pair k : An encoder, $e_{\theta_k}(\cdot)$, parameterized by θ_k , maps S-spaces to latent representations by minimizing a triplet loss:

$$L_k(\langle {}_k S^a, {}_k S^s, {}_k S^d \rangle) = \max(|e_{\theta_k}({}_k S^a) - e_{\theta_k}({}_k S^s)| - |e_{\theta_k}({}_k S^a) - e_{\theta_k}({}_k S^d)| + \delta, 0),$$

where ${}_k S^a$ and ${}_k S^s$ are sampled from the same set of ${}_k M_i$ planning problems for start-goal pair k where ${}_k P_i$ is optimal, and sample ${}_k S^d$ sampled from other problem sets where P_i is not optimal. δ is a margin that ensures a minimum distance between similar and dissimilar planning problem pairs. This process creates K different representation spaces, each with ${}_k N$ different latent clusters, which correspond to the ${}_k N$ existing motion plans $\{{}_k P_i\}_{i=1}^{{}_k N}$ for each of the K start-goal pairs. The cluster centroids are computed as

$${}_k c_i = \frac{1}{{}_k M_i} \sum_{j=1}^{{}_k M_i} e_{\theta_k}({}_k S_i^j).$$

Upon encountering a new planning problem ${}_k S_{N+1}$ for the k th start-goal pair, the latent representation $l_{N+1} = e_{\theta_k}({}_k S_{N+1})$ is calculated and compared against the centroids of existing clusters (${}_k c_i$) in the latent space. The closest cluster indicates the most suitable pre-existing motion plan (${}_k P_{i^*}$):

$$i^* = \underset{i}{\operatorname{argmin}} \|e_{\theta_k}({}_k S_{N+1}) - {}_k c_i\|.$$

where i^* identifies the selected motion plan that closely approximates the optimal solution for ${}_k S_{N+1}$ from start ${}_k s_s \in \mathcal{G}$ to goal ${}_k s_g \in \mathcal{G}$.

B. Multi-Goal Motion Planner Integration

Using the predicted motion plans ${}_k P_{i^*}, \forall k$, dynamically-feasible trajectories ζ_k can be produced by Eqn. (1) for each start-goal pair k . Then a motion map Φ is constructed to facilitate the discovery of paths connecting all goals by sampling along the coordinates of all ζ_k within the free space S_{free} , rather than the entire space, and linking these sampled regions using Dijkstra's algorithm to establish goal-to-goal paths (Algorithm 1 line 1).

To compute a TSP tour τ originating from s_{init} and targeting the remaining goals, we first evaluate the distance set $\text{DIS} = \{\text{dis}_k\}_{k=1}^K$ for each predicted trajectory ζ_k . Beginning from s_{init} , the tour τ is planned by iteratively moving towards the next goal with the lowest dis_k from the current goal (Algorithm 1 line 2).

With the tour τ determined efficiently, a motion tree, \mathcal{T} , with its root at s_{init} , is incrementally developed within S_{free} ,




Fig. 2. Different Planning Problem Classes: random, curve, maze, and storage (Left to Right) and Car and Snake Robot.

taking into account the dynamics of the robot and potential obstacles. In this process, targets around ζ_k are sampled to extend \mathcal{T} from any given node, $\eta \in \mathcal{T}$, towards any goal $g \in \mathcal{G}$. The motion tree (\mathcal{T}) is expanded to sequentially reach all goals according to τ , guided by the Multi-Goal Motion Memory. This involves prioritizing states near the predicted optimal trajectories ζ_k , thereby focusing the exploration on promising areas of S_{free} and enhancing planning efficiency (Algorithm 1 line 3).

V. IMPLEMENTATION

We implement Multi-Goal Motion Memory with a multi-goal motion planner, which follows the same structure defined in Sec. III, to showcase the improved planning time without sacrificing plan quality.

A. Guided Motion Map Construction

To build a motion map guided by Multi-Goal Motion Memory, the approach first generates all trajectories for all start-goal and goal-goal pairs, as shown in Alg. 2a. This process is efficient as it relies on retrieving the most suitable trajectories from previous experiences. Since the current environment is possibly different from environments used during training, there is some possibility that parts of the retrieved trajectories could collide with the obstacles. For this reason, a trajectory repair process is used which samples collision-free states around all the retrieved trajectories. These collision-free states then serve as nodes for our motion map. The next step is to connect neighboring states, discarding those edges that are in collision, giving rise to a motion map that captures the connectivity of the environment. Dijkstra's algorithm is then used to find shortest paths for each goal pair. If the motion map does not contain the path between a certain goal pair, we call the motion memory again to use the predicted trajectory as an alternative path. This process is repeated until all goal pairs are connected by paths. Pseudocode can be found in Alg. 2b.

B. Multi-Goal Motion Planning Augmented with Motion Memory

The overall approach starts by rooting the motion tree \mathcal{T} at the initial state. The first stage (Alg. 2(a)) leverages the motion-memory framework to predict paths for each start-goal and goal-goal pair. The predicted paths are stored in a set Λ and their distances are stored in a cost matrix Δ . The second stage (Alg. 2(b)), as described in Section V-A, builds the motion map by utilizing the predicted paths Λ , repairing via sampling when necessary, in order to connect each start-goal and goal-goal pair with a collision-free path.

Algorithm 2 Multi-Goal Motion Memory

Input: goals \mathcal{G} ; a set of coordinates for the predicted trajectory Λ ; length of predicted trajectories Δ

a) CALLMEMORY($s_{\text{init}}, \mathcal{G}$)

- 1: $\Lambda \leftarrow \text{null}; \Delta \leftarrow \text{null}$
- 2: **for** $i \in \mathcal{G}$ **do** $\Lambda_{s_{\text{init}}, i}, \Delta_{s_{\text{init}}, i} \leftarrow \text{MEMORY}(s_{\text{init}}, i)$
- 3: **for** $(i, j) \in \mathcal{G} \times \mathcal{G}$ **do** $\Lambda_{i, j}, \Delta_{i, j} \leftarrow \text{MEMORY}(i, j)$

b) GENERATEMOTIONMAP($s_{\text{init}}, \mathcal{G}, \Lambda$)

- 1: **for** $i \in \mathcal{G}$ **do**
- 2: **for** $p \in \Lambda_{s_{\text{init}}, i}$ **do**
- 3: $p_{\text{near}} \leftarrow \text{SAMPLEINVICINITY}(p)$
- 4: **if** CHECKCOLLISION(p_{near}) **then**
- 5: ADDMOTIONMAP(p_{near})
- 6: **for** $(i, j) \in \mathcal{G} \times \mathcal{G}$ **do**
- 7: **for** $p \in \Lambda_{i, j}$ **do**
- 8: $p_{\text{near}} \leftarrow \text{SAMPLEINVICINITY}(p)$
- 9: **if** CHECKCOLLISION(p_{near}) **then**
- 10: ADDMOTIONMAP(p_{near})
- 11: GENERATEPATHS(\mathcal{G}, \mathcal{G})
- 12: **for** $(i, j) \in \mathcal{G} \times \mathcal{G}$ **do**
- 13: **if** $\neg \text{FINDPATH}(i, j)$ **then** $\text{path}_{i, j} \leftarrow \Lambda_{i, j}$

c) SELECTTARGET(η, Λ)

- 1: $p_{\text{target}} \leftarrow \text{null}; g \leftarrow \text{FIRSTGOALFROMTOUR}(\eta, \text{tour})$
- 2: **for** several times **do**
- 3: **if** $\Lambda_{\eta, s, g} = \text{null}$ **then**
- 4: **return** $p_{\text{target}} \leftarrow \text{RANDOMSAMPLE}()$
- 5: $p \leftarrow \text{SELECTFIRSTCOORDINATE}(\Lambda_{\eta, s, g})$
- 6: $p_{\text{target}} \leftarrow \text{GENERATESAMPLEINVICINITY}(p)$
- 7: **if** COLLISIONCHECK($\eta, s, p_{\text{target}}$) **then**
- 8: REMOVEFIRSTCOORDINATE($\Lambda_{\eta, s, g}$)
- 9: **return** p_{target}
- 10: **return** $p_{\text{target}} \leftarrow \text{RANDOMSAMPLE}()$

The third stage utilizes the motion map, the cost matrix Δ , and a TSP solver to effectively guide the motion-tree expansion. The motion-tree is partitioned into groups, denoted by Γ , based on the goals and nodes in the motion map that have been reached. New groups are created when new goals or new nodes in the motion map are reached. More details about the group partition can be found in [48], [49].

When a new group is created, the TSP solver is invoked to compute an optimal tour of how to visit the remaining goals associated with that group. The TSP solver relies on the cost matrix Δ , which stores the set of distances produced by the motion memory predictions. The cost of the tour is used to define the group's priority, giving higher priority to groups associated with low-cost tours.

Starting from the initial state, the motion-tree \mathcal{T} is expanded incrementally. During each iteration, the group with the highest priority is selected. A node η then is selected at random from the group and attempts are then made to extend the motion tree from η along the path to the first goal in the group's tour. Specifically, the objective is to reach the points along the path in succession. For this reason, as described




Fig. 3. Example Goal-to-Goal Path Predictions by Multi-Goal Motion Memory in random, curve, maze, and storage (Left to Right).

in Alg. 2c, the first point from the path is retrieved and the target is generated near it (by sampling at random inside a small circle). If the target is in collision, the sampling is repeated again. If no collision-free target is generated after several attempts, then the target is generated by sampling a collision-free point at random from the entire space.

Once the target is generated, a PID controller is used to extend the motion tree from the selected η toward the target. If the target is reached, the first point is removed from the path (so that the expansion can continue toward the next point in the path). Otherwise, this process is repeated several times. If after several times, the target is not reached, the group's priority is lowered and the expansion from this group is abandoned. In this way, the approach proceeds with a new iteration, possibly selecting a different group for expansion. As demonstrated by the experiments in the next section, this results in an effective way to expand the motion-tree along tours that leverage previous planning experiences, as generated by the motion memory framework.

VI. EXPERIMENTS AND RESULTS

Experiments are conducted using a vehicle and snake-like robot operating in unstructured, obstacle-rich environments. The scenes are shown in Figs. 1 and 2, while the robot model is described in Section III-A. Fig. 1 left depicts the area a classical motion planners have to explore when facing a new environment, while Fig. 1 right shows the region explored using Multi-Goal Motion Memory, revealing a significant reduction in the resources required for planning.

A. Experimental Setup

1) *Methods for Comparisons:* Multi-Goal Motion Memory (Alg. 2) is compared to DROMOS [49], a state-of-the-art motion planner specifically designed for multi-goal problems. DROMOS was shown to be significantly faster than other planners. DROMOS uses a roadmap, shortest roadmap paths, and a TSP solver to guide the motion-tree expansion. As a baseline, we also used a sequential version of RRT [51], denoted by SequentialRRT, which aims to reach the goals one at a time, by going first to the nearest goal.

2) *Multi-Goal Problem Instances and Measuring Performance:* Our experiments include four scenes: random, curve, maze, and storage. For each scene, the experiments are conducted with 4 (2×2), 9 (3×3), and 16 (4×4) goals. These goal locations are fixed and the robot starts from one randomly-chosen goal. For a given scene and number of goals, we generate 400 problem instances. Each method is run on each of the instances. An upper bound of




Fig. 4. Runtime Results with Different Numbers of Goals: (a) Multi-Goal Motion Memory [car]; (b) DROMOS [car]; (c) SequentialRRT [car]; (d) Multi-Goal Motion Memory [snake]; (e) DROMOS [snake]; and (f) SequentialRRT [snake]. Entries marked with X indicate failure.




Fig. 5. Solution Distance with Different number of Goals: (a) Multi-Goal Motion Memory (MM); (b) DROMOS.

10s is set for each run. When calculating the performance statistics, we remove the lower and upper 25% instances in terms of runtime to avoid the influence of outliers. The runtime includes everything from reading the input files to reporting that a solution is found. The solution length is measured as the distance traveled by the robot along the solution trajectory.

3) Training Datasets For Multi-Goal Motion Memory:

For each different scene and each pair of goals, we generate 100 original planning problems and subsequently 5000 augmented environments for these different problems by slightly rearranging obstacles close to the motion plan and randomly shuffling obstacles in other places. These images are represented as 2D grids and input into convolutional neural networks for training. Each model predicts path for different goal pairs in a new set of 100 test problems. Four example predictions are shown in Fig. 3. Despite some collisions, the prediction can guide motion planners to efficiently find a path. In our experiments, the motion memory model trained for the car is directly used for the snake-like robot.

4) Computing Resources:

The experiments are run on an AMD Ryzen 9 5900X (3.7 GHz CPU) using Ubuntu 20.04. The motion planning code is written in C++ and compiled using g++-9.4.0, while the code for training the models is written in Python 3.8.

B. Results

1) Runtime Results with Different Number of Goals:

Fig. 4 shows the runtime results for 4, 9 and 16 goals. The results indicate that Multi-Goal Motion Memory is significantly faster than DROMOS and SequentialRRT.

The predicted trajectories can effectively guide motion-map construction, estimate goal-to-goal distance in the TSP solver, and bias the expansion of the motion tree along low-cost tours. SequentialRRT is the least effective among the three approaches. It lacks guidance and often wastes many resources during the motion-tree expansion process due to its randomized characteristics. DROMOS can solve multi-goal motion planning problems by relying on roadmaps to guide the motion-tree expansion, but it has difficulty in challenging environments since the roadmap paths do not necessarily account for the robot dynamics.

Parts d, e, and f in Fig. 4 show the snake-like results produced by these three approaches. Experimental results also show that our motion memory method has a significant acceleration effect even on a different robot.

2) Solution Distance Results:

Fig. 5 shows the solution distances, i.e., plan quality. SequentialRRT is not included due to its frequent failure. The results show that Multi-Goal Motion Memory produces similar solution distances as DROMOS, and even outperforms DROMOS in many environments. Such results highlight that Multi-Goal Motion Memory significantly improves multi-goal motion planning

efficiency, without sacrificing plan quality.

VII. CONCLUSIONS AND FUTURE WORK

Leveraging past planning experiences, Multi-Goal Motion Memory is an efficient approach to augment multi-goal motion planners to solve challenging problems in unstructured, obstacle-rich environments while considering robot dynamics. The approach's scalability is demonstrated by successfully handling a variety of numbers of goals with a large number of changing environments.

This work suggests several avenues for future research. Enhancing the machine-learning models may further improve the overall performance and reduce extra computation overhead. We demonstrate that machine-learning models trained on one robot model has the potential to transfer to other robot models. Therefore, future research can focus on generalizing training across similar and different robot models. Additionally, another interesting direction is to explore extending Multi-Goal Motion Memory to multiple robots.

REFERENCES

- [1] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, 2005.
- [2] S. M. LaValle, *Planning Algorithms*. Cambridge, MA: Cambridge University Press, 2006.
- [3] J. McMahon and E. Plaku, "Dynamic multi-goal motion planning with range constraints for autonomous underwater vehicles following surface vehicles," in *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2021, pp. 704–711.
- [4] G. Gutin and A. P. Punnen, *The traveling salesman problem and its variations*. Springer Science & Business Media, 2006, vol. 12.
- [5] D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook, "The traveling salesman problem," in *The Traveling Salesman Problem*. Princeton University Press, 2011.
- [6] K. Helsing, "An extension of the Lin-Kernighan-Helsing TSP solver for constrained traveling salesman and vehicle routing problems," *Roskilde: Roskilde University*, pp. 24–50, 2017.
- [7] J. Faigl, P. Váňa, and J. Deckerová, "Fast heuristics for the 3-d multi-goal path planning based on the generalized traveling salesman problem with neighborhoods," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2439–2446, 2019.
- [8] A. Bogrybayeva, T. Yoon, H. Ko, S. Lim, H. Yun, and C. Kwon, "A deep reinforcement learning approach for solving the traveling salesman problem with drone," *Transportation Research Part C: Emerging Technologies*, p. 103981, 2023.
- [9] L. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [10] T. Danner and L. E. Kavraki, "Randomized planning for short inspection paths," in *IEEE International Conference on Robotics and Automation*, vol. 2, 2000, pp. 971–976.
- [11] M. Saha, T. Roughgarden, J.-C. Latombe, and G. Sánchez-Ante, "Planning tours of robotic arms among partitioned goals," *International Journal of Robotics Research*, vol. 25, no. 3, pp. 207–223, 2006.
- [12] B. J. Englot and F. S. Hover, "Sampling-based coverage path planning for inspection of complex structures," in *International Conference on Automated Planning and Scheduling*, 2012, pp. 29–37.
- [13] D.-S. Jang, H.-J. Chae, and H.-L. Choi, "Optimal control-based uav path planning with dynamically-constrained tsp with neighborhoods," in *2017 17th International Conference on Control, Automation and Systems (ICCAS)*. IEEE, 2017, pp. 373–378.
- [14] J. Janoš, R. Pěnička, and V. Vonásek, "Randomized multi-goal path planning for dubins vehicles," in *International Conference on Emerging Technologies and Factory Automation*, 2022, pp. 1–4.
- [15] S. Edelkamp and E. Plaku, "Multi-goal motion planning with physics-based game engines," in *IEEE Conference on Computational Intelligence and Games*, 2014, pp. 1–8.
- [16] S. Rashidian, E. Plaku, and S. Edelkamp, "Motion planning with rigid-body dynamics for generalized traveling salesman tours," in *International Conference on Motion in Games*, 2014, pp. 87–96.
- [17] S. Edelkamp, M. Lahijanian, D. Magazzeni, and E. Plaku, "Integrating temporal reasoning and sampling-based motion planning for multi-goal problems with dynamics and time windows," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3473–3480, 2018.
- [18] T. McMahon, A. Sivaramakrishnan, E. Granados, K. E. Bekris *et al.*, "A survey on the integration of machine learning with sampling-based motion planning," *Foundations and Trends® in Robotics*, vol. 9, no. 4, pp. 266–327, 2022.
- [19] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Motion planning and control for mobile robot navigation using machine learning: a survey," *Autonomous Robots*, vol. 46, no. 5, pp. 569–597, 2022.
- [20] J. Wang, W. Chi, C. Li, C. Wang, and M. Q.-H. Meng, "Neural RRT*: Learning-based optimal path planning," *IEEE Transactions on Automation Science and Engineering*, vol. 17, no. 4, pp. 1748–1758, 2020.
- [21] Y. Li, R. Cui, Z. Li, and D. Xu, "Neural network approximation based near-optimal motion planning with kinodynamic constraints using RRT," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 11, pp. 8718–8729, 2018.
- [22] B. Ihter, E. Schmerling, T.-W. E. Lee, and A. Faust, "Learned critical probabilistic roadmaps for robotic motion planning," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 9535–9541.
- [23] J. Rajamäki and P. Hämäläinen, "Augmenting sampling based controllers with machine learning," in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2017, pp. 1–9.
- [24] Z. Wang, X. Xiao, A. J. Nettekoven, K. Umasankar, A. Singh, S. Bommakanti, U. Topcu, and P. Stone, "From agile ground to aerial navigation: Learning from learned hallucination," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 148–153.
- [25] X. Xiao, B. Liu, and P. Stone, "Agile robot navigation through hallucinated learning and sober deployment," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 7316–7322.
- [26] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Toward agile maneuvers in highly constrained spaces: Learning from hallucination," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1503–1510, 2021.
- [27] X. Xiao, Z. Xu, Z. Wang, Y. Song, G. Warnell, P. Stone, T. Zhang, S. Ravi, G. Wang, H. Karnan *et al.*, "Autonomous ground navigation in highly constrained spaces: Lessons learned from the benchmark autonomous robot navigation challenge at icra 2022 [competitions]," *IEEE Robotics & Automation Magazine*, vol. 29, no. 4, pp. 148–156, 2022.
- [28] X. Xiao, Z. Wang, Z. Xu, B. Liu, G. Warnell, G. Dhamankar, A. Nair, and P. Stone, "Appl: Adaptive planner parameter learning," *Robotics and Autonomous Systems*, vol. 159, p. 104132, 2022.
- [29] X. Xiao, B. Liu, G. Warnell, J. Fink, and P. Stone, "Appld: Adaptive planner parameter learning from demonstration," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4541–4547, 2020.
- [30] Z. Wang, X. Xiao, B. Liu, G. Warnell, and P. Stone, "Appli: Adaptive planner parameter learning from interventions," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 6079–6085.
- [31] Z. Wang, X. Xiao, G. Warnell, and P. Stone, "Apple: Adaptive planner parameter learning from evaluative feedback," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7744–7749, 2021.
- [32] Z. Xu, B. Liu, X. Xiao, A. Nair, and P. Stone, "Benchmarking reinforcement learning techniques for autonomous navigation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9224–9230.
- [33] Z. Xu, X. Xiao, G. Warnell, A. Nair, and P. Stone, "Machine learning methods for local motion planning: A study of end-to-end vs. parameter learning," in *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2021, pp. 217–222.
- [34] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning,"

- in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3052–3059.
- [35] X. Xiao, T. Zhang, K. M. Choromanski, T.-W. E. Lee, A. Francis, J. Varley, S. Tu, S. Singh, P. Xu, F. Xia, S. M. Persson, L. Takayama, R. Frostig, J. Tan, C. Parada, and V. Sindhwani, “Learning model predictive controllers with real-time attention for real-world navigation,” in *Conference on robot learning*. PMLR, 2022.
- [36] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone, “Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 807–11 814, 2022.
- [37] D. M. Nguyen, M. Nazeri, A. Payandeh, A. Datar, and X. Xiao, “Toward human-like social robot navigation: A large-scale, multi-modal, social human navigation dataset,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023.
- [38] R. Mirsky, X. Xiao, J. Hart, and P. Stone, “Conflict avoidance in social navigation—a survey,” *ACM Transactions on Human-Robot Interaction*, 2024.
- [39] X. Xiao, J. Biswas, and P. Stone, “Learning inverse kinodynamics for accurate high-speed off-road navigation on unstructured terrain,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 6054–6060, 2021.
- [40] H. Karnan, K. S. Sikand, P. Atreya, S. Rabiee, X. Xiao, G. Warnell, P. Stone, and J. Biswas, “Vi-ikd: High-speed accurate off-road navigation using learned visual-inertial inverse kinodynamics,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3294–3301.
- [41] K. S. Sikand, S. Rabiee, A. Uccello, X. Xiao, G. Warnell, and J. Biswas, “Visual representation learning for preference-aware path planning,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 11 303–11 309.
- [42] A. Datar, C. Pan, M. Nazeri, and X. Xiao, “Toward wheeled mobility on vertically challenging terrain: Platforms, datasets, and algorithms,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024.
- [43] A. H. Qureshi, A. Simeonov, M. J. Bency, and M. C. Yip, “Motion planning networks,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2118–2124.
- [44] J. Pan and D. Manocha, “Efficient configuration space construction and optimization for motion planning,” *Engineering*, vol. 1, no. 1, pp. 046–057, 2015.
- [45] C. Zhang, J. Huh, and D. D. Lee, “Learning implicit sampling distributions for motion planning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 3654–3661.
- [46] O. Arslan and P. Tsiotras, “Machine learning guided exploration for sampling-based motion planning algorithms,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015, pp. 2646–2652.
- [47] H.-D. Bui, Y. Lu, and E. Plaku, “Improving the efficiency of sampling-based motion planners via runtime predictions for motion-planning problems with dynamics,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 4486–4491.
- [48] Y. Lu and E. Plaku, “Leveraging single-goal predictions to improve the efficiency of multi-goal motion planning with dynamics,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 850–857.
- [49] E. Plaku, S. Rashidian, and S. Edelkamp, “Multi-group motion planning in virtual environments,” *Computer Animation and Virtual Worlds*, vol. 29, no. 6, p. e1688, 2018.
- [50] D. Das, Y. Lu, E. Plaku, and X. Xiao, “Motion memory: Leveraging past experiences to accelerate future motion planning,” 2024.
- [51] S. M. LaValle and J. J. Kuffner Jr, “Randomized kinodynamic planning,” *The International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.