

Manifold-constrained Hamilton-Jacobi Reachability Learning for Decentralized Multi-Agent Motion Planning

Qingyi Chen¹, Ruiqi Ni¹, Junyoung Kim¹ and Ahmed H. Qureshi¹

Abstract—Safe multi-agent motion planning (MAMP) under task-induced constraints is a critical challenge in robotics. Many real-world scenarios require robots to navigate dynamic environments while adhering to manifold constraints imposed by tasks. For example, service robots must carry cups upright while avoiding collisions with humans or other robots. Despite recent advances in decentralized MAMP for high-dimensional systems, incorporating manifold constraints remains difficult. To address this, we propose a manifold-constrained Hamilton-Jacobi reachability (HJR) learning framework for decentralized MAMP. Our method solves HJR problems under manifold constraints to capture task-aware safety conditions, which are then integrated into a decentralized trajectory optimization planner. This enables robots to generate motion plans that are both safe and task-feasible without requiring assumptions about other agents’ policies. Our approach generalizes across diverse manifold-constrained tasks and scales effectively to high-dimensional multi-agent manipulation problems. Experiments show that our method outperforms existing constrained motion planners and operates at speeds suitable for real-world applications. Video demonstrations and source code are available at <https://youtu.be/RyCEHMnPTH8> and <https://github.com/qingyichen/hammar>.

I. INTRODUCTION

Real-world operations often require multiple robots to function in a shared workspace while dynamically avoiding collisions and adhering to task-induced manifold constraints. For instance, teams of robots carrying large objects on a factory floor must maintain their closed-loop kinematic constraints while also avoiding collisions with other robots, whether those are moving independently or under similar constraints. Likewise, in restaurants, multiple service robots must maintain the orientation of trays or objects to prevent spillage while navigating to different tables without interfering with each other. The ability to plan and react in decentralized, multi-agent environments under manifold constraints is essential for realizing practical, large-scale applications of multi-robot systems.

Recent work on decentralized multi-agent motion planning (MAMP) has made promising progress in scaling to high-dimensional systems [1]–[7]. However, these methods generally do not account for manifold constraints and assume the robots operate in the full configuration space, limiting their

*This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-24-1-0233. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

¹Qingyi Chen, Ruiqi Ni, Junyoung Kim and Ahmed H. Qureshi are with the Department of Computer Science at Purdue University, West Lafayette, IN 47907, USA. {chen5221, ni1117, kim3722, ahqureshi}@purdue.edu




Fig. 1: Our method, HaMMAR, successfully solving a real-world cup-holding task with two UR5 manipulators. Transparent models indicate the goal configurations. The sequence illustrates how the manipulators move from their initial states and avoid potential collisions to reach their goals.

applications in constrained settings. A complementary line of research focuses on developing centralized approaches to solving MAMP problems with manifold constraints. These methods typically extend sampling-based constrained motion planners [8]–[12] to multi-agent settings [13], [14], combining constrained sampling with centralized coordination or scheduling to generate feasible motions under task constraints. While effective in enforcing manifold constraints, their reliance on centralization limits their applicability in dynamic environments where agents move independently.

In contrast to the above approaches, Hamilton-Jacobi reachability (HJR) offers a principled way to analyze safety and achieve dynamic optimal control of multi-agent systems. Recent advances have leveraged the strong approximation ability of neural networks to overcome the exponential scaling limitation of traditional HJR solvers and paved the way for various applications in safe MAMP [15]–[21]. Despite these advances, existing approaches do not often consider task-specific manifold constraints, thus motivating our integration of constrained motion planning with HJR.

This paper proposes HaMMAR, **H**amilton–**J**acobi with **M**anifold constraints for **M**ulti–**A**gent **R**eachability, a framework that learns manifold-constrained HJR for decentralized multi-agent motion planning. Our approach extends learning-based HJR solvers (i.e., DeepReach [15]) to learn reachability value functions on constraint manifolds, ensuring safety analysis under task-specific requirements. The value function is then incorporated into a receding-horizon trajectory planning framework to generate motion plans that both account for collision avoidance and respect task constraints. HaMMAR works as a multi-agent motion planner in a

decentralized manner without assuming the control policies of other agents. Our contributions are summarized as follows:

- 1) A theoretical extension to DeepReach [15] that solves HJR under manifold constraints.
- 2) A decentralized multi-agent trajectory planner that generate motion plans that avoid collisions while respecting task-induced manifold constraints.
- 3) A demonstration of our method on various constrained MAMP tasks, such as object-carrying, cup-holding, and doorway-crossing, to illustrate the applicability of our approach to realistic daily-life scenarios.

Our results demonstrate that HaMMAR generalizes across a range of constrained MAMP tasks, generating decentralized motion plans that are both safety-aware and task-compliant. Figure 1 shows snippets of HAMMAR performing a real-world MAMP task of cup-holding under orientation constraints. Corresponding demonstration videos are provided in the supplementary materials.

II. RELATED WORK

Decentralized multi-agent motion planning studies the problem of enabling multiple agents to independently generate motion plans that guide them from start to goal configurations without colliding with each other. Recent advances have seen promising results in scaling decentralized methods to high-dimensional systems such as robot manipulators. Model Predictive Control (MPC) and optimization-based methods [4], [5] achieve decentralized collision avoidance by embedding safety constraints into the optimization, while reinforcement learning-based approaches [1], [2] extract coordinated motion policies from experience. However, these methods typically assume robots operate in the full configuration space and do not explicitly account for task-induced manifold constraints. Reactive methods based on potential functions [6], [7] enable local collision avoidance and can in principle be adapted to constrained settings, but their decisions are often myopic. In general, ensuring safety in dynamic environments requires systematically reasoning about the influence of manifold constraints on the robot’s ability to avoid collisions and carefully generating corresponding motion plans.

Another line of research addresses the problem of constrained motion planning, which aims at finding collision-free paths under kinematic constraints. A representative example is AtlasRRT [10], which incrementally explores the constraint manifold using local charts. Extensions of sampling-based methods [13], [14], [22] have also demonstrated promising results in multi-agent cooperative tasks. However, these approaches generally emphasize offline planning and lack the ability to adapt online in dynamic or uncertain environments. Moreover, task constraints and inter-agent safety are often handled in centralized frameworks, limiting their applicability in decentralized scenarios.

Hamilton-Jacobi reachability provides a complementary perspective to solve MAMP by characterizing safe sets and optimal control policies in multi-agent scenarios [23]–[25]. Recently, neural HJR solvers have demonstrated promising

results in scaling HJR analysis to higher-dimensional systems [15], [16], [26] and have enabled applications in safety-critical planning, including decentralized settings [19], [21]. Inspired by this line of research, our method extends HJR to explicitly incorporate manifold constraints, enabling decentralized planning that is both safety-aware and task-feasible.

III. PRELIMINARIES AND BACKGROUND

A. Problem Setup

We investigate the problem of safe decentralized MAMP under manifold constraints, where agents are tasked to independently navigate toward their goals while avoiding collisions with other agents and adhering to their manifold task constraints. We assume that each agent can perfectly perceive the states of other agents and know their dynamics. However, we assume no knowledge of the control policies of other agents.

B. Notation

We denote the state (configuration) of a robot by $x \in \mathcal{X}$ where $\mathcal{X} \subseteq \mathbb{R}^{n_d}$ is its state (configuration) space with dimensionality n_d . We use $\mathcal{M} \subset \mathcal{X}$ to denote a manifold embedded in the ambient space \mathcal{X} and use $T_x\mathcal{M}$ to denote the tangent space at $x \in \mathcal{M}$,

When multiple agents are involved, we use superscripts to distinguish them by writing $x^i \in \mathcal{X}$ as the state of the i -th agent. We use $\langle \cdot, \cdot \rangle$ to denote a vector inner product. Other notations will be introduced and explained within the context of the discussion.

C. Hamilton-Jacobi Reachability

Hamilton-Jacobi Reachability is a foundational method in optimal control theory, offering mathematical guarantees in computing the Backward Reachable Set (BRS), the set of initial states that can reach a target set given a specified time instance and system dynamics.

Consider the optimal control problem of a system with dynamics $\dot{x} = f(x, u)$ and a cost function $\text{cost}(x(t), u(t))$, where $u \in \mathcal{U}$ is the control input with \mathcal{U} being compact. The corresponding cost of a trajectory is then defined by

$$J(t, x(t), u(t)) = l(x(T)) + \int_{t_0}^T \text{cost}(x(t), u(t)) dt \quad (1)$$

where T is the terminal time and $l(x(T))$ is a bounded and Lipschitz-continuous function that computes the terminal cost. The value function $V : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}$

$$V(t, x(t)) = \min_{u(\cdot)} J(t, x(t), u(t)). \quad (2)$$

then defines the optimal cost that can be achieved from state x at time t . Applying the principle of dynamic programming yields the Hamilton-Jacobi Bellman (HJB) partial differential equation (PDE) to solve for $V(t, x)$

$$\frac{\partial V}{\partial t} + \min_u \{ \langle \nabla_x V(t, x), f(x, u) \rangle + \text{cost}(x, u) \} = 0 \quad (3)$$

$$V(T, x(T)) = l(x(T)). \quad (4)$$

In motion planning, the target set can be described by a level set $\mathcal{L} = \{x \mid l(x) \leq 0\}$. The solution $V(t, x) \leq 0$ if and only if the system can reach \mathcal{L} by time T . The set $\mathcal{V}(t) = \{x \mid V(t, x) \leq 0\}$ thus defines the backward reachable set. The same principle extends to target-avoiding or reach-avoid tasks, as well as multi-agent settings. We refer interested readers to [27]–[29] for further background on HJR.

IV. PROPOSED METHOD

In this section, we introduce HaMMAR, which formulates a constrained HJR problem, solves it in a data-driven framework, and constructs a trajectory optimization planner that generates safe, task-feasible motions.

A. Manifold-Constrained Hamilton-Jacobi Reachability

Consider a dynamical system $\dot{x} = f(x, u)$, $x \in \mathcal{X} \subseteq \mathbb{R}^{n_d}$, $u \in \mathcal{U}$ constrained by n_c smooth equality constraints $C(x) = [c_1(x) \ c_2(x) \ \cdots \ c_{n_c}(x)]^T = 0$, with $0 < n_c < n_d$. These constraints $C: \mathcal{X} \rightarrow \mathbb{R}^{n_c}$ implicitly define a constrained configuration space (manifold) $\mathcal{M} = \{x \in \mathcal{X} \mid C(x) = 0\}$ embedded in the ambient configuration space \mathcal{X} (e.g., a vehicle constrained to a road curve, or a manipulator constrained to maintain an orientation).

Since the trajectory of the system must remain on \mathcal{M} rather than the full configuration space \mathcal{X} , reachability must be analyzed intrinsically on this manifold. In particular, the principle of dynamic programming for the value function over a short trajectory $V: [0, T] \times \mathcal{M} \rightarrow \mathbb{R}$ satisfies

$$\begin{aligned} & V(t, x(t)) \\ &= \min_u \{V(t + \delta, x(t + \delta)) + \int_t^{t+\delta} \text{cost}(x(\tau), u(\tau)) d\tau\} \\ & \quad \text{s.t. } x(\tau) \in \mathcal{M} \ \forall \tau \in [t, t + \delta] \quad (5) \\ &= \min_u \{V(t, x(t)) + \frac{\partial V}{\partial t}(t, x(t)) \cdot \delta + \langle \nabla V(t, x(t)), \frac{dx}{dt} \rangle \cdot \delta \\ & \quad + \text{cost}(x(t), u(t)) \delta\} \\ & \quad \text{s.t. } x(\tau) \in \mathcal{M} \ \forall \tau \in [t, t + \delta] \quad (6) \end{aligned}$$

with a first-order Taylor Expansion. Taking the limit of δ to zero yields a constrained form of the HJB PDE

$$\begin{aligned} & \frac{\partial V(t, x)}{\partial t} + \min_u \{\langle \nabla V(t, x), f(x, u) \rangle + \text{cost}(x, u)\} = 0 \\ & \quad \text{s.t. } \dot{C}(x) = \begin{bmatrix} \langle \nabla c_1(x), f(x, u) \rangle \\ \langle \nabla c_2(x), f(x, u) \rangle \\ \vdots \\ \langle \nabla c_{n_c}(x), f(x, u) \rangle \end{bmatrix} = 0 \quad (7) \end{aligned}$$

with the terminal condition $V(T, x) = l(x)$, $x \in \mathcal{M}$. The corresponding Hamiltonian is then

$$\begin{aligned} H_{\mathcal{M}}(t, x, p) &= \min_u \{ \langle p, f(x, u) \rangle + \text{cost}(x, u) \} \\ & \quad \text{s.t. } \dot{C}(x) = 0 \quad (8) \end{aligned}$$

where the costate $p = \nabla V(t, x)$. Importantly, since V is only defined on the manifold \mathcal{M} , the Euclidean gradient $\nabla V(t, x)$ is not globally meaningful. Therefore, when we compute the Hamiltonian and write ∇V , we are referring to a directional

derivative of the form $DV(t, x)[v] = \langle \nabla_{\mathcal{M}} V, v \rangle$, $\forall v \in T_x \mathcal{M}$ where $\nabla_{\mathcal{M}} V$ denotes a manifold gradient. We also assume satisfactions of the standard conditions for the HJR framework [27] and that $\nabla c_i(x)$ is bounded $\forall i = 1, \dots, n_c \ \forall x \in \mathcal{M}$. Essentially, the equality constraints restrict admissible dynamics to the tangent bundle

$$T\mathcal{M} = \{(x, v) \in (\mathcal{M}, \mathbb{R}^{n_d}) \mid J_C(x)v = 0\} \quad (9)$$

where $J_C(x)$ is the Jacobian of the constraints, so that feasible trajectories evolve intrinsically on \mathcal{M} . As in the unconstrained setting, the resulting PDE must be understood in the viscosity sense. We note that a majority of the derivation follows the one in the unconstrained case and refer interested readers to [27], [30] for more details.

B. Solving Manifold-Constrained HJR with DeepReach

This section discusses how we solve manifold-constrained HJR problems (i.e., Equation 7) as a safety constraint for multi-agent motion planning. In safe robot motion planning, HJR is often concerned with the Backward Reachable Tube (BRT) problem formulated by

$$\begin{aligned} & \frac{\partial V(t, x)}{\partial t} + \min\{H(t, x, \nabla V(t, x)), 0\} = 0 \\ & \quad V(T, x) = l(x). \quad (10) \end{aligned}$$

where H is the Hamiltonian of the system. For systems with disturbances,

$$H(t, x, \nabla V(t, x)) = \max_u \min_d \langle \nabla V(t, x), f(x, u, d) \rangle \quad (11)$$

with $u \in \mathcal{U}$ being the control to the system and $d \in \mathcal{D}$ being the disturbance to the system. The inclusion of a non-anticipative disturbance component reflects a zero-sum differential game [28] setting that well models the interaction between an agent and the other adversary agent(s), thus producing a safety value function for multi-agent motion planning. The solution ensures that by acting optimally, the agent can remain safe for the horizon $[t, T]$ given that $V(t, x) > 0$. In these cases, the boundary condition $l(x)$ is often chosen as a signed distance function that distinguish safe and unsafe states by its sign.

DeepReach [15] is a leading approach that solves Equation 10 using neural networks, overcoming the exponential scaling limitation of traditional methods. For a given training point (t, x) , it defines the loss function $L(t, x; \theta)$ as

$$L(t, x; \theta) = L_1(t, x; \theta) + \lambda L_2(t, x; \theta) \quad (12)$$

where

$$\begin{aligned} L_1(t, x; \theta) &= |V_{\theta}(t, x) - l(x)| \mathbf{1}(t = T) \quad (13) \\ L_2(t, x; \theta) &= \left| \frac{\partial V_{\theta}(t, x)}{\partial t} + \min\{0, H(t, x, \nabla V(t, x))\} \right|. \quad (14) \end{aligned}$$

The network weights θ are optimized to minimize the loss function. For further details on DeepReach, we refer interested readers to [15].

We now discuss how to extend this framework to manifold-constrained settings. Recall from Equations 7 and

8 that the key difference is the restriction of system dynamics to the tangent bundle of the constraint manifold. This therefore requires replacing the unconstrained Hamiltonian in Equation 14 with its constrained counterpart. While this generally entails solving a constrained optimization problem at each step and may lead to inefficiency in a training, there are important cases where a closed-form constrained Hamiltonian can be derived.

Consider a velocity-controlled system $\dot{x} = f(x, u) = u$, $\|u\|_2 \leq \bar{u}$. This for example corresponds to a particle system controlled by a velocity input, up to a maximum speed \bar{u} . The manifold-constrained Hamiltonian of this system in a HJB formulation is then

$$H_{\mathcal{M}}(t, x, \nabla V(t, x)) \quad (15)$$

$$= \min_u \langle \nabla V(t, x), u \rangle \quad \text{s.t. } \dot{C}(x) = 0, \|u\|_2 \leq \bar{u} \quad (16)$$

$$= \min_u \langle \nabla V(t, x), u \rangle \quad \text{s.t. } J_C(x)u = 0, \|u\|_2 \leq \bar{u} \quad (17)$$

where $J_C(x) \in \mathbb{R}^{n_c \times n_d}$ is the Jacobian of the constraint function $C(x)$. This can be simplified as

$$H_{\mathcal{M}}(t, x, \nabla V(t, x)) \quad (18)$$

$$= \min_u \langle \nabla V(t, x), P(x)u \rangle \quad \text{s.t. } \|u\|_2 \leq \bar{u} \quad (19)$$

$$= -\bar{u} \|P(x)\nabla V(t, x)\|_2 \quad (20)$$

with the minimizer

$$u^* = -\bar{u} \frac{P(x)\nabla V(t, x)}{\|P(x)\nabla V(t, x)\|_2} \quad (21)$$

and the projection matrix

$$P(x) = \mathbb{I} - J_C(x)^T (J_C(x)J_C(x)^T)^{-1} J_C(x). \quad (22)$$

Here $P(x)$ projects vectors onto the nullspace of $J_C(x)$, ensuring that control inputs lie in the tangent space of the manifold. We notes that $J_C(x)$ is required to be full row rank so that $J_C J_C^T$ is invertible. Although derived for a velocity-controlled system, such closed-form solution to constrained Hamiltonian can be similarly obtained for linear control-affine systems.

C. Trajectory Optimization with Neural HJR

In this section, we extend NeHMO [21] to manifold-constrained scenarios to generate motion plans that guide an agent toward its goal, satisfy the specified manifold constraints, and avoid collisions in the presence of other agents. At each planning step, the computed trajectory must remain on the manifold \mathcal{M} , making sure the value function V_θ is well-defined. Moreover, the value function $V_\theta(t, x)$ is required to stay above a specified safety threshold ϵ to achieve collision avoidance. These requirements lead to a receding-horizon trajectory optimization problem (OPT) formulated

as follows:

$$\min_{u \in \mathcal{U}} \int_0^{t_{\text{plan}}} \text{cost}((x^1(t), u, x_{\text{goal}}^1)) dt \quad (23)$$

$$\text{s.t. } x^1(t + \Delta t) = f_d(x^1(t), u)$$

$$x^i(t + \Delta t) = f_d(x^i(t), d^i)$$

$$C(x(t_{\text{plan}})) = 0$$

$$\tau = T - t_{\text{plan}}$$

$$d^i = \arg \min_{d \in \mathcal{D}} \max_{u \in \mathcal{U}} \langle \nabla V_\theta(\tau, x^1(t), x^i(t)), f(x(t), u, d) \rangle$$

$$V_\theta(T - t_{\text{safe}}, x^1(t_{\text{plan}}), x^i(t_{\text{plan}})) > \epsilon$$

$$\text{for } i = 2, \dots, m$$

Here, m is the number of agents in the environment, ϵ is a safety buffer, and $f_d(x^i(t), u)$ is the discrete-time model propagation function that computes the next state of an agent given its current state and a control input. The manifold constraints ensure that the agent's next state lies on \mathcal{M} , while the HJR value function is computed pair-wise with every other agent to enforce the planned trajectory to be collision-free under worst-case actions from other agents.

The solution of the optimization provides the control input for the next planning horizon. To ensure safety, we adopt a fail-safe strategy: if the optimization fails to solve, the agent falls back to a conservative control given by the most collision-evading input computed similar to Equation 21. During execution, the agent follows its current plan while simultaneously computing the next one, enabling real-time adaptation to environmental changes and the states of other agents.

V. EXPERIMENTS AND VALIDATION

This section presents two sets of analyses to evaluate our proposed HaMMAR framework. First, we compare the HJR value functions computed with and without the manifold-constrained formulation on a 2-dimensional particle system. Second, we evaluate HaMMAR as a motion planner in three representative MAMP tasks that must operate under manifold constraints. In the object-carrying task (Fig. 3), sub-teams of robots operate in a decentralized manner to transport a large object while respecting closed-loop manifold constraints. In the cup-holding task (Fig. 4), multiple robots independently transport their cups while keeping it upright and avoiding collisions with one another. Finally, in the doorway-crossing task (Fig. 6), we extend the classical doorway problem to high-DOF manipulators, requiring them to pass through a doorway while maintaining several end-effector constraints. All experiments are conducted using UR5 robots. For each problem setup, we generate 100 test cases with distinct start and goal configurations for the multi-agent system.

We benchmark **HaMMAR** against the following baselines: (i) **(Opt) w/o HJR**, an ablation without safety constraints, (ii) **AtlasRRT (decentralized)**, and (iii) **AtlasRRT (centralized)**. **AtlasRRT (decentralized)** runs AtlasRRT for each robot independently to generate manifold-constrained waypoints that connect them respectively from start to goal configurations. This baseline serves as a reference where

planned paths are not coordinated by a trajectory planner, to highlight the necessity of a trajectory planner that reacts to the environment in real-time. **AtlasRRT (centralized)** plans the motion of all manipulators in a centralized manner. This baseline has the privilege to globally control the motion of all robots to avoid collisions, however at the cost of higher dimensionality. This baseline serves as a reference of the task difficulties. Evaluation of the MAMP experiments is based on the following metrics:

- **Success Rate (SR%)**: The ratio of trials where all agents reach their goals without any collision.
- **Collision Rate (CR%)**: The ratio of trials where collisions happen at any time.
- **Planning Time (Time)**: The mean time taken for each agent to generate a plan.
- **Path Length (PL)**: The average path length across successful trials.

All experiments are conducted on a desktop with Intel(R) Xeon(R) w5-2455X CPUs and an NVIDIA GeForce RTX 4070 GPU. The optimization problem (Opt) is solved with IPOPT [31]. The multi-arm environment is adapted from [32] using `trimesh` [33] for collision detection. We assign a 5-second time limit to **AtlasRRT (decentralized)** and a 20-second time limit to **AtlasRRT (centralized)**.

A. 2D Circle-Constrained Particle

This section analyzes the solution quality of the learned value function on a 2-dimensional circle-constrained particle system. As an ablation, we compare the BRS obtained from the value functions learned with and without the manifold-constrained formulation.

We consider a 2D velocity-controlled particle system with state $x = (x_1, x_2) \in \mathbb{R}^2$ and dynamics $\dot{x} = f(x, u) = u \in \mathbb{R}^2$, $\|u\|_2 \leq 1$. The particle is constrained to a circle of radius 0.5 centered at the origin, thus defining the manifold constraint $c(x) = \|x\|_2 - 0.5 = 0$. We now study a goal-reaching reachability problem with the goal located at $x_{\text{goal}} = (0.5, 0.0)$ and terminal condition $V(T, x) = l(x) = \|x - x_{\text{goal}}\|_2$. This corresponds to a geodesic distance problem and gives rise to the ground truth BRS $\mathcal{V}_{\text{gt}}(t) = \{(x_1, x_2) \in \mathcal{M} \mid \frac{1}{2} \arccos(2x_1) - (T - t) \leq 0\}$.

Figure 2 illustrates the BRS obtained from the value functions trained with and without the constrained formulation. As the time horizon increases from $\pi/8$ to $3\pi/8$, the BRS computed without the constrained formulation progressively over-approximates the BRS. For a quantitative analysis, we treat BRS prediction as a binary classification problem and report the two value functions’ accuracy, recall, precision, and F-1 scores in Table I. The unconstrained version over-approximates the BRS, achieving high recall by capturing all true positives. However, this over-approximation also introduces false positives, lowering its precision. In contrast, the constrained version closely matches the ground truth and achieves high performance across all metrics.




Fig. 2: A comparison of the BRS obtained from the ground truth (left), the manifold-constrained HJR value function (middle), and the unconstrained HJR value function (right), each shown at three time slices. Backward reachable regions appear in white, backward unreachable regions in blue, and off-manifold regions in orange. The BRS obtained with the constrained formulation closely aligns the ground truth, while the unconstrained one over-approximates the BRS.

Time	Manifold	Acc. \uparrow	Rec. \uparrow	Prec. \uparrow	F-1 \uparrow
$T - \frac{\pi}{8}$	w/	99.4	99.6	98.2	98.9
	w/o	97.4	100.0	90.7	95.1
$T - \frac{\pi}{4}$	w/	99.5	99.3	99.7	99.5
	w/o	89.7	100.0	83.0	90.7
$T - \frac{3\pi}{8}$	w/	99.5	99.4	99.9	99.7
	w/o	75.0	100.0	75.0	85.7

TABLE I: The accuracy, recall, precision, and F-1 score of the BRS prediction obtained from the learned value functions with and without the manifold constraint.

B. Object-Carrying UR5 Manipulators

This section evaluates HaMMAR as a multi-agent motion planner in a factory-inspired dual-arm manipulation scenario. The setup consists of two UR5 leader manipulators placed face-to-face, each with an additional follower manipulator positioned at a distance. The followers match the motion of their corresponding leaders to carry an object collaboratively under closed-loop manifold constraints.

We simulate a task where the leader robots aim to transport their respective objects to a shared workspace while avoiding collisions with each other. Both the initial positions of the objects and their goal positions in the shared workspace are randomized for each trial to vary task configurations. All manipulators are constrained to maintain a fixed end-

effector orientation, introducing a manifold constraint on the system $C(x) = \text{orientation}(x) = R_{\text{desired}}$. This constraint simulates industrial requirements where the object must remain upright or aligned during handling. Figure 3 illustrates an example of the setup.

We solve for the leaders’ motion to ensure collision avoidance and adherence to the manifold, while the followers match the leaders’ trajectories using inverse kinematics to maintain object stability. Table II summarizes the experiment result of the object-carrying task. **HaMMAR** outperforms the baselines both in achieving success and in avoiding collisions. **AtlasRRT (decentralized)** performs similarly with **(Opt) w/o HJR**, highlighting the need of a trajectory planner in dynamic environments. When the two leader manipulators are jointly considered in a centralized manner, **AtlasRRT (centralized)** achieves more successes with zero collisions. However, the high-dimensionality of the problem prevents it from either frequently finding a valid solution or solving the problem in real-time.

Methods	SR% \uparrow	CR% \downarrow	Time [s] \downarrow	PL [rad] \downarrow
HaMMAR (ours)	85	0	0.16 \pm 0.15	2.7 \pm 1.3
(Opt) w/o HJR	68	32	0.063 \pm 0.023	2.1 \pm 0.6
AtlasRRT (decentralized)	68	32	0.22 \pm 0.14	2.4 \pm 0.6
AtlasRRT (centralized)	75	0	7.6 \pm 8.6	2.5 \pm 0.6

TABLE II: The success rate, collision rate, planning time, and mean path length of each method on the **object-carrying** task. Best results for each metric are shown in bold.

C. Cup-Holding UR5 Manipulators

We further test our framework on a goal-reaching task in which multiple manipulators aim to reach their respective goal poses without colliding. Each manipulator holds a cup that must stay upright throughout the motion, introducing a manifold constraint on the end effector’s orientation in z-direction $C(x) = \text{orientation}_z(x) = [0 \ 0 \ 1]^T$. Unlike the factory manipulation scenario, this task emphasizes pose-to-pose motion planning and involves longer-range motion, making collision-free path-finding more challenging. Figure 4 illustrates an example of the setup.

Table III reports the results of the cup-holding experiment with two manipulators. **HaMMAR** outperforms the decentralized baselines both in achieving success and in avoiding collisions. As the required range of motion increases compared to the object-carrying task, **AtlasRRT (centralized)** starts to struggle to find a path due to the necessity to adequately sample a high-dimensional manifold. For **real-world**




Fig. 5: An example of the **five-UR5 cup-holding** task. Start configurations are shown in solid gray with goals shown in transparent green.

validation, we evaluated our method on five challenging cup-carrying tasks with varying start and goal configurations. In all cases, **HaMMAR** executed successfully, achieving an average planning time of 0.10 ± 0.051 seconds per step and a mean path length of 11 ± 3.3 . A visualization of the hardware experiment is shown in Figure 1, with videos provided in the supplementary materials.

To further assess the scalability of **HaMMAR**, we extend the cup-holding task to involve more agents. In these experiments, we surround a manipulator with two and four additional UR5 manipulators, thereby ensuring that each added agent meaningfully increases task difficulty. Planning time is therefore only measured on the surrounded manipulator, which faces the most challenging coordination requirement. An example configuration of the five-UR5 cup-holding is illustrated in Figure 5. With three UR5 manipulators, **HaMMAR** achieves 78/100 successful trials with 2 collisions, taking on average 0.12 ± 0.13 seconds per planning step. When scaled up to five manipulators, it records 57/100 successes and 4 collisions while maintaining a reasonable planning time of 0.17 ± 0.18 seconds. In comparison, **AtlasRRT (centralized)** achieves fewer than 40 successes in the three-manipulator setting and fails to return any valid trajectory within the 20-second limit for five manipulators. The lower success rates in this setting reflect the higher risk of entanglement and tighter collision-avoidance constraints due to the presence of more agents. Still, **HaMMAR** is able to maintain safety to a large extent with efficiently scaled planning time. A video example of the five-manipulator setup is included in the supplementary materials.

Methods	SR% \uparrow	CR% \downarrow	Time [s] \downarrow	PL [rad] \downarrow
HaMMAR (ours)	82	2	0.10 \pm 0.087	7.7 \pm 3.1
(Opt) w/o HJR	67	33	0.056 \pm 0.010	6.8 \pm 2.4
AtlasRRT (decentralized)	67	33	1.9 \pm 1.7	7.0 \pm 2.4
AtlasRRT (centralized)	60	1	15 \pm 7.4	6.8 \pm 2.2

TABLE III: The success rate, collision rate, planning time, and mean path length of each method on the **cup-holding** task. Best results for each metric are shown in bold.

D. Doorway-Crossing for UR5 Manipulators

Finally, we evaluate our method on a doorway problem designed to involve various manifold constraints. In this task, two UR5 manipulators are equipped with sticks and must coordinate to cross a doorway defined by grooves on a block. This geometry introduces narrow passages and tight clearances, making collision avoidance particularly challenging.

The task enforces several manifold constraints on the system. First, the stick must be held upright such that $c_1(x) = \text{orientation}_x(x) = [0 \ 0 \ -1]^T$. Then two alignment constraints $c_2(x) = \text{z_position}(x) = z_{\text{desired}}$ and $c_3(x) = \text{y_position}(x) \pm \text{x_position}(x) + \text{offset} = 0$ enforce the end of the sticks to properly align with the grooves on the block. Each trial is randomized by sampling the initial and goal locations of the sticks within a neighborhood around the grooves’ endpoints to encourage variation in the setup




Fig. 3: An example of HaMMAR successfully solving an **object-carrying** task. The robots’ current configurations are shown in solid gray, while each object is shown in red or blue, with transparent counterparts indicating their respective goal positions. The sub-figures illustrate a single trial at different time steps, where the manipulators start from their initial configurations and get close to each other (a-b), adjust their configurations to avoid a potential collision (c-d), and then safely move towards their goals (e). The corresponding video demonstration is available in the supplementary materials.




Fig. 4: An example of HaMMAR successfully solving a **cup-holding** task. The robots’ current configurations are shown in solid gray with their goals shown in transparent green. The sub-figures illustrate a single trial at different time steps, where the manipulators start from their initial configurations and get close to each other (a-b), adjust their configurations to avoid a potential collision (c-d), and then safely move towards their goals (e). The corresponding video demonstration is available in the supplementary materials.



Fig. 6: An example of HaMMAR successfully solving a **doorway-crossing** task. The robots’ current configurations are shown in solid gray with their goals shown in transparent green. The sub-figures illustrate a single trial at different time steps, where the robots start from their initial configurations and get close to each other (a-b), wait until clearance (c), and then safely move towards their goals (d). The corresponding video demonstration is available in the supplementary materials.

while keeping the problem challenging. Figure 6 illustrates an example of the setup.

Table IV records the results of the doorway-crossing experiment. Without an effective safety constraint or a trajectory planner, **(Opt) w/o HJR** and **AtlasRRT (decentralized)** frequently collide. The presence of several manifold constraints also causes **AtlasRRT (centralized)** to constantly fail by hitting the 20-second time limit and only return a partial trajectory that does not lead to the goal. In contrast, **HaMMAR** outperforms the baselines in achieving success by a large margin for its ability to identify unsafe situations under complex manifold constraints.

VI. CONCLUSION AND DISCUSSION

This paper introduces HaMMAR, a framework that solves constrained HJR to enable decentralized manifold-constrained multi-agent motion planning. The resulting value function captures safety conditions under manifold constraints and enables a receding-horizon trajectory optimization scheme that avoids collisions while respecting task-

Methods	SR% \uparrow	CR% \downarrow	Time [s] \downarrow	PL [rad] \downarrow
HaMMAR (ours)	71	1	0.14 \pm 0.14	4.7 \pm 1.3
(Opt) w/o HJR	9	91	0.072 \pm 0.011	4.5 \pm 0.49
AtlasRRT (decentralized)	8	92	1.7 \pm 1.9	4.7 \pm 0.46
AtlasRRT (centralized)	0	0	22 \pm 0.72	-

TABLE IV: The success rate, collision rate, planning time, and mean path length of each method on the **doorway-crossing** task. Best results for each metric are shown in bold.

induced equality constraints. We demonstrate the effectiveness of our method on challenging multi-agent motion planning problems involving diverse task constraints.

Several limitations remain and point to future directions. First, discretized trajectories may drift off the constraint manifold due to solver approximations or inaccurate dynamics, and trajectory segments between planning updates are not explicitly constrained to the manifold. Addressing this issue may require integrating a high-frequency low-level controller with the planner. Second, while the adversarial

formulation of HJR removes assumptions on other agents' control policies for MAMP, it also introduces conservatism, as reflected in the longer paths observed in our experiments. Under the velocity-control setting considered here, applying the fail-safe control inferred from the HJR formulation can also lead to jittery motions. Finally, several sources of uncertainty may affect real-world deployment. For example, the use of neural networks sacrifices formal safety guarantees, and solution quality cannot be directly verified due to the lack of ground-truth solutions in high-dimensional HJR problems. In addition, perception noise, delayed or inaccurate state estimates, and dynamics mismatch may introduce further challenges. Our empirical validation is also limited to velocity-controlled agents and relatively structured scenarios; evaluating the approach on more complex robotic systems remains an important direction. Future work aims to address these challenges, for example by incorporating calibration and uncertainty quantification methods [34], [35].

REFERENCES

- [1] H. Ha, J. Xu, and S. Song, "Learning a decentralized multi-arm motion planner," 2020. [Online]. Available: <https://arxiv.org/abs/2011.02608>
- [2] M. Lai, K. Go, Z. Li, T. Kröger, S. Schaal, K. Allen, and J. Scholz, "Roboballet: Planning for multirobot reaching with graph neural networks and reinforcement learning," *Science Robotics*, vol. 10, no. 106, Sept. 2025. [Online]. Available: <http://dx.doi.org/10.1126/scirobotics.ads1204>
- [3] N. Gafur, G. Kanagalingam, A. Wagner, and M. Ruskowski, "Dynamic collision and deadlock avoidance for multiple robotic manipulators," *IEEE Access*, vol. 10, p. 55766–55781, 2022. [Online]. Available: <http://dx.doi.org/10.1109/ACCESS.2022.3176626>
- [4] A. Tika and N. Bajcinca, "Synchronous minimum-time cooperative manipulation using distributed model predictive control," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2020, pp. 7675–7681.
- [5] V. N. Hartmann, A. Orthey, D. Driess, O. S. Oguz, and M. Toussaint, "Long-horizon multi-robot rearrangement planning for construction assembly," *IEEE Trans. on Robotics*, vol. 39, no. 1, p. 239–252, Feb. 2023. [Online]. Available: <http://dx.doi.org/10.1109/TRO.2022.3198020>
- [6] S. Bakker, L. Knoedler, M. Spahn, W. Böhmer, and J. Alonso-Mora, "Multi-robot local motion planning using dynamic optimization fabrics," 2023. [Online]. Available: <https://arxiv.org/abs/2310.12816>
- [7] E. J. Rodríguez-Seda and M. D. Kutzler, "Guaranteed real-time cooperative collision avoidance for n-dof manipulators," *Robotica*, vol. 42, no. 9, pp. 3149–3173, 2024.
- [8] D. Berenson, S. Srinivasa, D. Ferguson, and J. Kuffner, "Manipulation planning on constraint manifolds," pp. 625 – 632, May 2009.
- [9] Z. Kingston, M. Moll, and L. E. Kavraki, "Sampling-based methods for motion planning with constraints," *Annual review of control, robotics, and autonomous systems*, vol. 1, no. 1, pp. 159–185, 2018.
- [10] L. Jaillet and J. M. Porta, "Asymptotically-optimal path planning on manifolds," in *Robotics: Science and Systems*, 2012, pp. 145–152.
- [11] D. Berenson, S. Srinivasa, and J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *The Int. Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.
- [12] B. Kim, T. T. Um, C. Suh, and F. C. Park, "Tangent bundle rrt: A randomized algorithm for constrained motion planning," *Robotica*, vol. 34, no. 1, pp. 202–225, 2016.
- [13] M. Gharbi, J. Cortes, and T. Simeon, "A sampling-based path planner for dual-arm manipulation," in *2008 IEEE/ASME Int. Conf. on Advanced Intelligent Mechatronics*, 2008, pp. 383–388.
- [14] W. Guo, Z. Kingston, K. Hang, and L. E. Kavraki, "Efficient multi-robot motion planning for manifold-constrained manipulators by randomized scheduling and informed path generation," 2024. [Online]. Available: <https://arxiv.org/abs/2412.00366>
- [15] S. Bansal and C. Tomlin, "DeepReach: A deep learning approach to high-dimensional reachability," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2021.
- [16] J. F. Fisac, N. F. Lugovoy, V. Rubies-Royo, S. Ghosh, and C. J. Tomlin, "Bridging hamilton-jacobi safety analysis and reinforcement learning," in *2019 Int. Conf. on Robotics and Automation (ICRA)*, 2019, pp. 8550–8556.
- [17] K.-C. Hsu, D. P. Nguyen, and J. F. Fisac, "Isaacs: Iterative soft adversarial actor-critic for safety," 2024. [Online]. Available: <https://arxiv.org/abs/2212.03228>
- [18] K.-C. Hsu*, V. Rubies-Royo*, C. Tomlin, and J. Fisac, "Safety and liveness guarantees through reach-avoid reinforcement learning," in *Robotics: Science and Systems XVII*, ser. RSS2021. Robotics: Science and Systems Foundation, July 2021. [Online]. Available: <http://dx.doi.org/10.15607/RSS.2021.XVII.077>
- [19] H. J. Jeong, Z. Gong, S. Bansal, and S. Herbert, "Parameterized fast and safe tracking (FaSTrack) using DeepReach," in *Proceedings of the 6th Annual Learning for Dynamics and Control Conf.*, ser. Proceedings of Machine Learning Research, A. Abate, M. Cannon, K. Margellos, and A. Papachristodoulou, Eds., vol. 242. PMLR, 15–17 Jul 2024, pp. 1006–1017. [Online]. Available: <https://proceedings.mlr.press/v242/jeong24a.html>
- [20] S. Bansal, M. Chen, S. L. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," *CoRR*, vol. abs/1709.07523, 2017. [Online]. Available: <http://arxiv.org/abs/1709.07523>
- [21] Q. Chen and A. H. Qureshi, "Nehmo: Neural hamilton-jacobi reachability learning for decentralized safe multi-agent motion planning," 2025. [Online]. Available: <https://arxiv.org/abs/2507.13940>
- [22] A. Agrawal, P. Mayer, Z. Kingston, and G. A. Hollinger, "Constrained nonlinear kaczmarz projection on intersections of manifolds for coordinated multi-robot mobile manipulation," in *2025 IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2025, pp. 7726–7732.
- [23] K. Leung, E. Schmerling, M. Zhang, M. Chen, J. Talbot, J. C. Gerdes, and M. Pavone, "On infusing reachability-based safety assurance within planning frameworks for human-robot vehicle interactions," 2020. [Online]. Available: <https://arxiv.org/abs/2012.03390>
- [24] M. Chen, J. C. Shih, and C. J. Tomlin, "Multi-vehicle collision avoidance via hamilton-jacobi reachability and mixed integer programming," in *2016 IEEE 55th Conf. on Decision and Control (CDC)*. IEEE, 2016, pp. 1695–1700.
- [25] S. Bansal, M. Chen, K. Tanabe, and C. J. Tomlin, "Provably safe and scalable multivehicle trajectory planning," *IEEE Trans. on Control Systems Technology*, vol. 29, no. 6, pp. 2473–2489, 2021.
- [26] Z. Feng, L. Qiu, and S. Bansal, "Bridging model predictive control and deep learning for scalable reachability analysis," *arXiv preprint arXiv:2505.03830*, 2025.
- [27] I. M. Mitchell, *Application of level set methods to control and reachability problems in continuous and hybrid systems*. stanford university, 2002.
- [28] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry, "Reach-avoid problems with time-varying dynamics, targets and constraints," 2014. [Online]. Available: <https://arxiv.org/abs/1410.6445>
- [29] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," 2017. [Online]. Available: <https://arxiv.org/abs/1709.07523>
- [30] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [31] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, pp. 25–57, 2006.
- [32] J. Michaux, A. Li, Q. Chen, C. Chen, and R. Vasudevan, "Safe Planning for Articulated Robots Using Reachability-based Obstacle Avoidance With Spheres," in *Proceedings of Robotics: Science and Systems*, 2024.
- [33] Dawson-Haggerty et al., "trimesh." [Online]. Available: <https://trimesh.org/>
- [34] A. Lin and S. Bansal, "Verification of neural reachable tubes via scenario optimization and conformal prediction," in *Proceedings of the 6th Annual Learning for Dynamics and Control Conf.*, ser. Proceedings of Machine Learning Research, A. Abate, M. Cannon, K. Margellos, and A. Papachristodoulou, Eds., vol. 242. PMLR, 15–17 Jul 2024, pp. 719–731. [Online]. Available: <https://proceedings.mlr.press/v242/lin24a.html>
- [35] A. Lin, A. Pinto, and S. Bansal, "Robust verification of controllers under state uncertainty via hamilton-jacobi reachability analysis," *arXiv preprint arXiv:2511.14755*, 2025.