

Physics-Informed Multi-Agent Reinforcement Learning for Distributed Multi-Robot Problems

Eduardo Sebastián, Thai Duong, Nikolay Atanasov, Eduardo Montijano and Carlos Sagüés

Abstract—The networked nature of multi-robot systems presents challenges in the context of multi-agent reinforcement learning. Centralized control policies do not scale with increasing numbers of robots, whereas independent control policies do not exploit the information provided by other robots, exhibiting poor performance in cooperative-competitive tasks. In this work we propose a physics-informed reinforcement learning approach able to learn distributed multi-robot control policies that are both scalable and make use of all the available information to each robot. Our approach has three key characteristics. First, it imposes a port-Hamiltonian structure on the policy representation, respecting energy conservation properties of physical robot systems and the networked nature of robot team interactions. Second, it uses self-attention to ensure a sparse policy representation able to handle time-varying information at each robot from the interaction graph. Third, we present a soft actor-critic reinforcement learning algorithm parameterized by our self-attention port-Hamiltonian control policy, which accounts for the correlation among robots during training while overcoming the need of value function factorization. Extensive simulations in different multi-robot scenarios demonstrate the success of the proposed approach, surpassing previous multi-robot reinforcement learning solutions in scalability, while achieving similar or superior performance (with averaged cumulative reward up to $\times 2$ greater than the state-of-the-art with robot teams $\times 6$ larger than the number of robots at training time). We also validate our approach on multiple real robots in the Georgia Tech Robotarium under imperfect communication, demonstrating zero-shot sim-to-real transfer and scalability across number of robots.

Index Terms—Cooperative control, distributed systems, multi-robot systems, physics-informed neural networks, reinforcement learning.

I. INTRODUCTION

MULTI-ROBOT systems promise improved efficiency and reliability compared to single robots in many applications, including exploration and mapping [3], [4], agriculture and herding [5]–[8], and search and rescue [9]. However, the complexity of describing mathematically the objective and constraints in many of these problems makes the design of analytical controllers a challenging task. Multi-agent reinforcement learning [10]–[14] addresses this issue by only requiring

E. Sebastián, E. Montijano and C. Sagüés are with the DIIS - I3A, Universidad de Zaragoza, Spain (e-mails: {esebastian, emonti, csagues}@unizar.es).

T. Duong and N. Atanasov are with the Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093 USA (e-mails: {tduong, natanasov}@ucsd.edu).

This work has been supported by ONR N00014-23-1-2353, N62909-24-1-2081 and NSF CCF-2402689 (ExpandAI), by Spanish projects PID2021-125514NB-I00, PID2021-124137OB-I00 and TED2021-130224B-I00 funded by MCIN/AEI/10.13039/501100011033, by ERDF A way of making Europe and by the European Union NextGenerationEU/PRTR, DGA T45-23R, a Spanish grant FPU19-05700 and a US-Spain Fulbright grant.

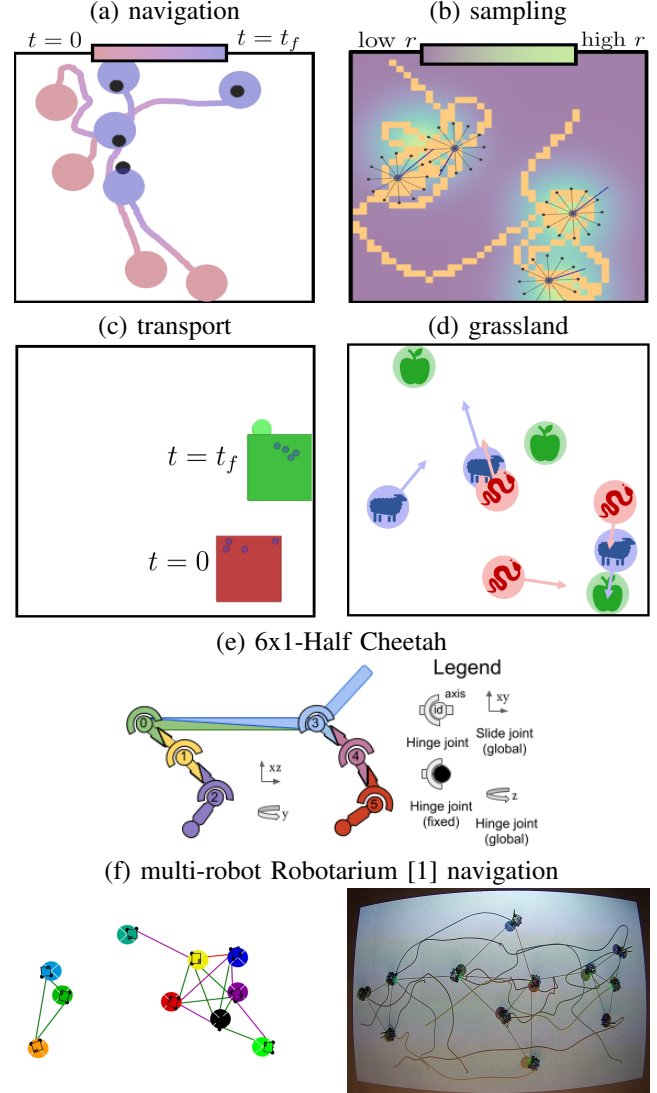


Fig. 1. Examples of scenarios addressed by our physics-informed multi-agent reinforcement learning approach. The scenarios cover a wide variety of cooperative/competitive behaviors and levels of coordination complexity: (a) robots learn to navigate to landmarks (black dots) while avoiding collisions with other robots (colored trajectories from pink at $t = 0$ to blue at $t = t_f$); (b) robots cooperate to do active LiDAR (star pattern) sampling (orange areas refer to regions already sampled) in an unknown environment with regions of low reward (magenta areas) and high reward (green areas); (c) robots (blue dots) cooperate to transport a square box (red) with unknown mass and that follows Newtonian dynamics towards a desired region (green dot, box turns green when it reaches the goal); (d) the robots (blue) collect food (green) while avoiding collisions with attackers (red), since collision with attackers lead to deactivation of the robots; (e) joints learn how to coordinate to make the Half-Cheetah robot walk, where each joint is an independent agent that can only interact with the adjacent agents (image from [2]); (f) a team of robots learns to navigate and avoid collisions in a real-world setting characterized by a tight space and high density of robots.

a high-level mathematical specification of the task (the reward function), which is commonly available.

A fundamental limitation of existing multi-agent reinforcement learning approaches is the poor scalability they offer against increasing and time-varying numbers of robots. Centralized control policies do not scale whereas independent control policies neglect the information that other robots can offer. How to design and train control policies based on neural networks that are distributed and use all the available information is still an open problem [13], [15], [16]. The approach explored in this paper leverages physical knowledge about the robot system [17]–[22] to make this possible. Physics-informed neural networks are becoming popular in different fields, e.g., climate science, quantum mechanics or fluid dynamics [23], [24]. They need less data for training and allow encoding general constraints found in physical systems. Nonetheless, they have not been used to learn distributed multi-robot control policies in multi-agent reinforcement learning.

The main contribution of our work is a novel physics-informed multi-agent reinforcement learning approach suitable for general multi-robot problems (Sec. III). The techniques to make this possible (Sec. IV) are summarized as follows.

- We develop a novel distributed and scalable by design neural network architecture to describe multi-robot control policies. This is achieved by combining a physics-informed port-Hamiltonian description of the multi-robot system (Sec. IV-A) with self-attention (Sec. IV-B). The former naturally encodes the distributed nature of the policy and respects the energy conservation laws of the individual dynamics of the robots. The latter handles the information coming from communication or perception in time-varying neighborhoods.
- We integrate the port-Hamiltonian self-attention policy in a soft actor-critic reinforcement learning algorithm that exploits the physics-informed description of the robot team to impose the sparsity pattern from the interaction graph in the policy function (Sec. V). To handle the networked nature of the multi-robot policy and avoid non-stationarity issues during training, we modify the acquisition of experience to keep track of the correlations among robots.

Extensive simulations and real experiments in multi-robot scenarios (Sec. VI) demonstrate that the combination of multi-agent reinforcement learning techniques and a physics-informed description of the system achieves scalable control policies with similar or superior performance as the state-of-the-art (Sec. VII).

This paper is an extension of our prior conference paper [19]. The main difference of this work compared to [19] is that it addresses a multi-agent reinforcement learning problem (where only reward samples are provided), whereas [19] solves a learning-from-demonstrations problem (where data from an expert control policy is available). This leads to a new multi-agent reinforcement learning problem formulation and a substantially different approach beyond the policy parameterization, since our new approach involves modifications on the parameterization of the actor and the replay buffer of the actor-critic, besides some particularities in the reward function

design and the critic. The neural networks that model the control policies are different as well, since the multi-agent reinforcement learning scenarios are continuous-time while those in [19] are discrete-time. Finally, all experiments in this paper are new, including an extensive scalability analysis in a real multi-robot platform involving up to 16 robots in a tight space.

II. RELATED WORK

A. Learning multi-robot control policies from data

The design of multi-robot control policies deals with two challenges.

The first challenge is the mathematical formulation of the problem, where we must consider the task and the constraints inherited from the cooperative-competitive nature of the multi-robot team. To address this, recent works exploit machine learning and focus on learning control policies for optimal control or reinforcement learning problems [25]–[27]. When demonstrated data from an expert are available, inverse reinforcement learning [28] can be used to learn centralized [29], [30] or distributed [31], [32] policies from task demonstrations. However, finding experts for multi-robot applications is difficult. It is also possible to apply supervised learning approaches to learn multi-robot control policies [33], [34], but, again, collecting labeled trajectories is hard. Therefore, in this work we use reinforcement learning, because a reward function is a high-level description of the task that is easy to build and is typically available.

The second challenge is that the learning and execution of control policies for multi-robot systems should scale favorably with an increasing numbers of robots. Learning a joint policy function is challenging due to the exponential growth of the state and action space [35]. Attention mechanisms are widely used in multi-agent reinforcement learning problems [36]–[39] to enhance the performance of centralized training settings where the agents are isolated from each other and do not consider communication during deployment. Graph neural networks [13] have been utilized as a stable [40], scalable and communication-aware policy representation in path planning, coverage, exploration, and flocking problems [41]–[46]. These approaches assume discrete robot dynamics, fixed or known communication topology, or prior knowledge on the formulation of the control policy. In contrast, by using a port-Hamiltonian formulation and self-attention mechanisms, our approach directly learns control policies that handle time-varying neighbors, do not constrain the size of the neighborhoods, and learn constraints such as collision avoidance without specific mathematical formulation.

B. Multi-agent reinforcement learning for robotic problems

Multi-agent reinforcement learning extends reinforcement learning approaches to problems where multiple agents interact in the environment [47]–[52]. The first multi-agent reinforcement learning approaches considered centralized policies where the states, actions and observations of every agent are globally known by a central unit during deployment [10]. The complexity of these approaches exponentially scales with the

number of agents, so they are not feasible in practice. In the absence of a centralized policy, the Markov game modeling the problem becomes non-stationary [53]–[55]. The task is no longer stationary because the policies of all the other agents are changing, and so the environment itself. As a consequence, it is not possible to reach an equilibrium in training.

To solve the non-stationarity issue, a common approach is to factorize the value and/or policy functions. The most common factorization, the so-called centralized-training decentralized-execution, departs from centralized-training reinforcement learning approaches [56]–[58] and considers that the agents act independently. To provide some global information to the agents, agents learn an approximation of the other agents' control policies [59]–[61]. However, this proxy is only used during training and, since the resulting control policies only use the individual observations, the performance deteriorates the number of robots change. This aspect is partially overcome by considering heterogeneous agents [62] but scalability is deteriorated again in large teams because the agents do not learn how to exploit neighboring information effectively. Another option is to provide each agent with a global estimator that predicts the trajectories of the whole multi-robot team [63] or to implicitly learn to coordinate with other agents through the value function [55]. In mean field approaches [64], the multi-agent game is reduced to an interaction between an agent and the average of the other agents, which is not practical. To overcome these issues, we propose a new formulation of the policy function that is distributed by design and captures the underlying interaction graph that describes the multi-robot system.

Other works focus on more complex factorizations that account for local interactions among neighboring agents in the value function. For instance, the value function can be approximated as dependent on the neighbors only [65]. Either using analytical [16], [35], [66] or learned [67]–[69] factorizations, the control policies are still restricted to the same team size used during training because the policies do not explicitly consider communication. Closely related to ours, the authors of [70] propose an end-to-end neural network based on attention to achieve scalability of the control policy in quadrotor swarm navigation problems. Similarly, by using a port-Hamiltonian description of the system, our approach avoids factorizing the reward or the value functions, thus fully exploiting the available information in the experience while learning scalable distributed control policies.

C. Physics-informed neural networks for robotic problems

While black-box neural networks are widely used for learning control policies, they do not encode energy conservation and kinematic constraints satisfied by physical robot systems. Failing to infer them from data may result in unstable behaviors. Moreover, other constraints are present that come from the perception and communication modules of the robots and which restrict the information available to the robot. Similar issues are found in other physical applications [24], leading to physics-informed neural network [23], neural networks that use the differential equations that model physical systems as

building blocks. The use of physics-informed machine learning for robotics and control is very recent [20]–[22], [71], [72] and focuses on centralized controllers. Nevertheless, physics-informed neural network can also be used to address the learning of cooperative distributed control policies.

Many dynamical systems—from robots interacting with their surroundings to large-scale multi-physics systems—involve a number of interacting subsystems. This compositional property can be exploited [18] to train neural network sub-models from data generated by simple sub-systems, and the dynamics of more complex composite systems are then predicted without requiring additional data. The systems are represented as a port-Hamiltonian neural network [17], a class of neural ordinary differential equations that uses a port-Hamiltonian dynamics formulation as inductive bias [73]. A key contribution of our work is to represent the robot team as a port-Hamiltonian system and learn distributed control policies by modeling robot interactions as energy exchanges. The use of Hamiltonian mechanics has been explored for centralized control policies or distributed but fixed-time known topologies [74]–[76]. Meanwhile, our work achieves scalability with a time-varying topology by modeling robot interactions using self-attention [77].

III. PROBLEM FORMULATION

Consider a team of robots, indexed by $\mathcal{V} = \{1, \dots, n\}$. The robot team motion is governed by *known* continuous-time control-affine stochastic dynamics:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) + \mathbf{L}\omega(t), \quad (1)$$

where $\mathbf{x}(t) = [(\mathbf{x}^1(t))^\top, \dots, (\mathbf{x}^n(t))^\top]^\top \in \mathcal{X} \subseteq \mathbb{R}^{n \times n_x}$ is the joint state of the robot team at time $t \geq 0$, with $\mathbf{x}^i(t)$ the state of robot i at time t . On the other hand, $\mathbf{u} = [(\mathbf{u}^1(t))^\top, \dots, (\mathbf{u}^n(t))^\top]^\top \in \mathcal{U} \subseteq \mathbb{R}^{n \times n_u}$ is the joint input, with $\mathbf{u}^i(t)$ the input of robot i at time t . The term $\omega(t)$ is white noise modeling the uncertainty in the robot sensors and actuators and $\Xi = \mathbf{L}\mathbf{L}^\top$ is the noise diffusion matrix. Let $\{t_\tau\}_{\tau=0}^\infty$ be a sequence of discrete time instants such that $t_{\tau+1} - t_\tau = T_\tau > 0$ and assume zero-order hold inputs such that $\mathbf{u}(t) = \mathbf{u}(t_\tau), \forall t \in [t_\tau, t_{\tau+1})$. Then, from Eq. (1), we can obtain an Euler discretization of the dynamics with discrete-time state $\mathbf{s}_\tau = \mathbf{x}(t_\tau)$, action $\mathbf{a}_\tau = \mathbf{u}(t_\tau)$, and dynamics:

$$\mathbf{s}_{\tau+1} = \mathbf{s}_\tau + T_\tau \mathbf{f}(\mathbf{s}_\tau, \mathbf{a}_\tau) + \mathbf{n}_\tau, \quad (2)$$

where \mathbf{n}_τ is zero-mean Gaussian noise with covariance $T_\tau \Xi$. We will formulate the dynamics \mathbf{f} using port-Hamiltonian mechanics as its modularity in terms of energy effectively describes the networked interactions in a robot team.

The multi-robot task and the interactions among the robots are modeled as a Markov Decision Process (MDP), defined as a tuple $(\mathcal{X}, \mathcal{U}, p, r, \gamma)$. In the tuple, $p : \mathcal{X} \times \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ is the probability density of the next joint state $\mathbf{s}_{\tau+1}$ conditioned on the current joint state \mathbf{s}_τ and joint action \mathbf{a}_τ , $r : \mathcal{X} \times \mathcal{U} \rightarrow [r_{\min}, r_{\max}]$ is a reward function encoding the objectives of the multi-robot task, and $\gamma \in (0, 1)$ is the discount factor. Whereas the robot dynamics are formulated in continuous time, the MDP is formulated in discrete time

based on the zero-order hold discretization in Eq. (2). Accordingly, $p(\mathbf{s}_{\tau+1}|\mathbf{s}_\tau, \mathbf{a}_\tau)$ is a Gaussian density with mean $\mathbf{s}_\tau + T\mathbf{f}(\mathbf{s}_\tau, \mathbf{a}_\tau)$ and covariance $T\Xi$.

The robots interact in a distributed manner, described by a time-varying undirected graph $\mathcal{G}_\tau = (\mathcal{V}, \mathcal{E}_\tau)$, where $\mathcal{E}_\tau \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. An edge $(i, j) \in \mathcal{E}_\tau$ exists when robots i and j interact at time τ . Let $\mathbf{A}_\tau \in \{0, 1\}^{n \times n}$ be the adjacency matrix associated to \mathcal{G}_τ , such that $[\mathbf{A}_\tau]_{ij} = 1$ if and only if $(i, j) \in \mathcal{E}_\tau$, and 0 otherwise. The set of k -hop neighbors of robot i at time τ is $\mathcal{N}_\tau^{i,k} = \{j \in \mathcal{V} \mid [\mathbf{A}_\tau^k]_{ij} \neq 0\}$, where \mathbf{A}_τ^k is the k -th power of matrix \mathbf{A}_τ . We remark that $\mathcal{N}_\tau^{i,k}$ includes robot i .

The goal of this work is to learn distributed control policies that solve a given multi-robot task, such that they respect the networked structure of the multi-robot team and the MDP model. We represent the policies as stochastic Markov control policies that depend on the k -hop neighbors of robot i :

$$\mathbf{a}_\tau^i \sim \pi_\theta \left(\mathbf{a}_\tau^i | \mathbf{s}_{\mathcal{N}_\tau^{i,k}}^i \right). \quad (3)$$

Here, $\mathbf{s}_{\mathcal{N}_\tau^{i,k}}^i = \{\mathbf{s}_\tau^j | j \in \mathcal{N}_\tau^{i,k}\}$ denotes the states of robot i and its k -hop neighbors, and θ denotes the control policy parameters. The use of a stochastic Markov policy is not only motivated by the MDP model but also by the fact that distributed control policies are prone to uncertainty from the communication-perception modules, control goals, and interaction with the environment. We assume that the policy in Eq. (3) is the same for all the robots. The joint control policy of all the robots is denoted by $\Pi_\theta = [\pi_\theta^\top(\mathbf{a}_\tau^1 | \mathbf{s}_{\mathcal{N}_\tau^{1,k}}^1), \dots, \pi_\theta^\top(\mathbf{a}_\tau^n | \mathbf{s}_{\mathcal{N}_\tau^{n,k}}^n)]^\top$.

Learning a distributed control policy π_θ that solves a certain multi-robot task is equivalently posed as learning the parameters θ such that π_θ maximizes the expected sum of rewards over time, i.e.,

$$\begin{aligned} \max_{\theta} Q_{\Pi_\theta}(\mathbf{s}, \mathbf{a}) = \\ \max_{\theta} \mathbb{E}_{\mathbf{s}_\tau \sim p} \mathbb{E}_{\mathbf{a}_\tau \sim \Pi_\theta(\cdot | \mathbf{s}_\tau)} \left[\sum_{\tau=0}^{\infty} \gamma^\tau r(\mathbf{s}_\tau, \mathbf{a}_\tau) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a} \right]. \end{aligned} \quad (4)$$

In Eq. (4), the function $Q_{\Pi_\theta}(\mathbf{s}, \mathbf{a})$ is known as the action-value (Q) function associated with the policy Π_θ [48].

We do not make any assumptions on the reward function or action-value function, such as a specific factorization [16], [35]. The purpose of this work is to design a control policy that maximizes $Q_{\Pi_\theta}(\mathbf{s}, \mathbf{a})$ and enforces the distributed factorization expressed in Eq. (3).

IV. PHYSICS-INFORMED MULTI-AGENT REINFORCEMENT LEARNING

In this section, we present a novel physics-informed multi-agent reinforcement learning approach to find distributed control policies that solve multi-robot tasks as defined in Eq. (4), under the restrictions on the robot dynamics and available information defined in Eq. (3). Our formulation is done in continuous time, following the continuous-time definition of the robot dynamics in Eq. (1). Later, zero-order hold control is used to utilize a discrete-time version of the control policies in the reinforcement learning algorithm.

First, we present a port-Hamiltonian formulation of the multi-robot dynamics and an energy-based distributed control design that can shape the interactions and the Hamiltonian of the closed-loop system (Sec. IV-A). Given experience from trial and error simulations, we employ a self-attention mechanism to learn the parameters of the control policy that maximize the cumulative reward (Sec. IV-B). To simplify the notation, we omit the time dependence of the states \mathbf{x} and controls \mathbf{u} in the reminder of the paper. To facilitate the exposition, in Sec. IV-A and IV-B we assume that the robot dynamics and control policies are deterministic. In Sec. V, we explain how to return to the stochastic setting.

A. Port-Hamiltonian dynamics for multi-robot energy conservation

Port-Hamiltonian mechanics are a general yet interpretable modeling approach for learning and control. On the one hand, many physical networked systems can be described as a port-Hamiltonian system [74] using the same formulation and with a modular and distributed interpretation. Meanwhile, the port-Hamiltonian description allows to derive general energy-based controllers with closed-loop stability guarantees. Since robots are physical systems that satisfy Hamiltonian mechanics, we model each robot as a port-Hamiltonian system [73]:

$$\dot{\mathbf{x}}^i = (\mathbf{J}^i(\mathbf{x}^i) - \mathbf{R}^i(\mathbf{x}^i)) \frac{\partial H^i(\mathbf{x}^i)}{\partial \mathbf{x}^i} + \mathbf{F}^i(\mathbf{x}^i) \mathbf{u}^i, \quad (5)$$

where the skew-symmetric interconnection matrix $\mathbf{J}^i(\mathbf{x}^i)$ represents energy exchange within a robot, the positive-semidefinite dissipation matrix $\mathbf{R}^i(\mathbf{x}^i)$ represents energy dissipation, the Hamiltonian $H^i(\mathbf{x}^i)$ represents the total energy, and the matrix $\mathbf{F}^i(\mathbf{x}_i)$ is the input gain.

The interconnection of port-Hamiltonian systems leads to another port-Hamiltonian system [78]. Therefore, if the control and state of each robot are considered as input and output energy ports, then, due to the modularity of port-Hamiltonian dynamics, the multi-robot system with joint state \mathbf{x} also follows port-Hamiltonian dynamics:

$$\dot{\mathbf{x}} = (\mathbf{J}(\mathbf{x}) - \mathbf{R}(\mathbf{x})) \frac{\partial H(\mathbf{x})}{\partial \mathbf{x}} + \mathbf{F}(\mathbf{x}) \mathbf{u}, \quad (6)$$

where $\mathbf{J}(\mathbf{x})$, $\mathbf{R}(\mathbf{x})$, and $\mathbf{F}(\mathbf{x})$ are block-diagonal:

$$\begin{aligned} \mathbf{J}(\mathbf{x}) &= \text{diag}(\mathbf{J}^1(\mathbf{x}^1), \dots, \mathbf{J}^n(\mathbf{x}^n)), \\ \mathbf{R}(\mathbf{x}) &= \text{diag}(\mathbf{R}^1(\mathbf{x}^1), \dots, \mathbf{R}^n(\mathbf{x}^n)), \\ \mathbf{F}(\mathbf{x}) &= \text{diag}(\mathbf{F}^1(\mathbf{x}^1), \dots, \mathbf{F}^n(\mathbf{x}^n)), \end{aligned} \quad (7)$$

and $H(\mathbf{x}) = \sum_{i=1}^n H^i(\mathbf{x}^i)$. It is noteworthy that the expression in Eq. (6) is control-affine and follows the definition of Eq. (1), with $\mathbf{f}(\mathbf{x}, \mathbf{u}) = \mathbf{h}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}$, $\mathbf{h}(\mathbf{x}) = (\mathbf{J}(\mathbf{x}) - \mathbf{R}(\mathbf{x})) \frac{\partial H(\mathbf{x})}{\partial \mathbf{x}}$, $\mathbf{g}(\mathbf{x}) = \mathbf{F}(\mathbf{x})$, and $\mathbf{L} = \mathbf{0}$ since we are considering a deterministic setting for now.

Without control, the trajectories of the open-loop system in (6) would follow the dynamics of the robots in the absence of interactions with the environment or other robots. The dynamics need to be controlled by a policy in order to accomplish the desired task. We propose to design a control policy $\mu_\theta(\mathbf{x})$ and, then, obtain the desired policy in Eq. (3). Policy

$\mu_\theta(\mathbf{x})$ is designed using an interconnection and damping assignment passivity-based control (IDA-PBC) approach [73], which injects additional energy in the system through the control input \mathbf{u} to obtain closed-loop dynamics that achieve the desired task:

$$\dot{\mathbf{x}} = (\mathbf{J}_\theta(\mathbf{x}) - \mathbf{R}_\theta(\mathbf{x})) \frac{\partial H_\theta(\mathbf{x})}{\partial \mathbf{x}}, \quad (8)$$

with Hamiltonian $H_\theta(\mathbf{x})$, skew-symmetric interconnection matrix $\mathbf{J}_\theta(\mathbf{x})$, and positive semidefinite dissipation matrix $\mathbf{R}_\theta(\mathbf{x})$, which depend on the control policy $\mu_\theta(\mathbf{x})$. By matching the terms in (6) and (8), one obtains the joint policy:

$$\mathbf{u} = \Pi_\theta(\mathbf{x}) = \mathbf{F}^\dagger(\mathbf{x}) \left((\mathbf{J}_\theta(\mathbf{x}) - \mathbf{R}_\theta(\mathbf{x})) \frac{\partial H_\theta(\mathbf{x})}{\partial \mathbf{x}} - (\mathbf{J}(\mathbf{x}) - \mathbf{R}(\mathbf{x})) \frac{\partial H(\mathbf{x})}{\partial \mathbf{x}} \right), \quad (9)$$

where $\mathbf{F}^\dagger(\mathbf{x}) = (\mathbf{F}^\top(\mathbf{x})\mathbf{F}(\mathbf{x}))^{-1}\mathbf{F}^\top(\mathbf{x})$ is the pseudo-inverse of $\mathbf{F}(\mathbf{x})$.

If the robots are fully-actuated, i.e., $\mathbf{F}(\mathbf{x})$ is full-rank, then the input \mathbf{u} in (9) exactly transforms the open loop system in (6) to the closed-loop system in (8). For underactuated systems, the transformation may not be exact [79]. Being able to maximize Q_{Π_θ} is, hence, related to whether the robot configurations that solve the task are realizable by the class of control policies in (9). Even if goals of the task are not realizable, the policy parameters θ may still be optimized to achieve a behavior as good as possible to solve the task.

Let $[\mathbf{J}_\theta(\mathbf{x})]_{ij}$ and $[\mathbf{R}_\theta(\mathbf{x})]_{ij}$ denote the $n_x \times n_x$ blocks with index (i, j) , representing the energy exchange between robot i and j and the energy dissipation of robot i caused by robot j , respectively. Since the input gain $\mathbf{F}(\mathbf{x})$ in (7) is block-diagonal, the individual control policy of robot i is:

$$\mu_\theta(\mathbf{x}) = (\mathbf{F}^i)^{\dagger}(\mathbf{x}^i) \left(\sum_{j \in \mathcal{N}^i} ([\mathbf{J}_\theta(\mathbf{x})]_{ij} - [\mathbf{R}_\theta(\mathbf{x})]_{ij}) \frac{\partial H_\theta(\mathbf{x})}{\partial \mathbf{x}_j} - (\mathbf{J}^i(\mathbf{x}^i) - \mathbf{R}^i(\mathbf{x}^i)) \frac{\partial H^i(\mathbf{x})}{\partial \mathbf{x}^i} \right). \quad (10)$$

Note that the individual control policy $\mu_\theta(\mathbf{x})$ in (10) does not necessarily respect the hops in the communication network as desired in (3) because it depends on the structure of $\mathbf{J}_\theta(\mathbf{x})$, $\mathbf{R}_\theta(\mathbf{x})$, and $H_\theta(\mathbf{x})$. In Sec. IV-B, we impose conditions on these terms to ensure that they respect the communication topology and are skew-symmetric, positive semidefinite and positive respectively, as required for a valid port-Hamiltonian system and to find a policy μ_θ that only depends on the k -hop neighbors.

B. Self-attention parameterization to enforce communication patterns

We seek to learn distributed control policies that follow the structure of Eqs. (9)-(10) and (i) scale with the number of robots, (ii) handle time-varying communication and (iii) guarantee the port-Hamiltonian constraints. To do so, we first derive conditions on the port-Hamiltonian terms of the controller, $\mathbf{J}_\theta(\mathbf{x})$, $\mathbf{R}_\theta(\mathbf{x})$ and $H_\theta(\mathbf{x})$, which are the terms to be learned from the reinforcement learning experience. Then, we develop a novel architecture based on self-attention to ensure

that the learned control policies guarantee the desired requirements. We summarize the overall neural network architecture in Fig. 2.

To respect the robot team topology defined by the graph \mathcal{G} , we first impose $\mathbf{J}_\theta(\mathbf{x})$ and $\mathbf{R}_\theta(\mathbf{x})$ to be block-sparse,

$$[\mathbf{J}_\theta(\mathbf{x})]_{ij} = [\mathbf{R}_\theta(\mathbf{x})]_{ij} = \mathbf{0}, \quad \forall j \notin \mathcal{N}^{i,k}. \quad (11)$$

From the perspective of robot i , this means that the controller only considers information from its k -hop neighbors. Moreover, we require that the desired Hamiltonian factorizes over the k -hop neighborhoods:

$$H_\theta(\mathbf{x}) = \sum_{i=0}^n H_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i), \quad (12)$$

with $\mathbf{x}_{\mathcal{N}^{i,k}}^i = \{\mathbf{x}^j | j \in \mathcal{N}^{i,k}\}$. The factorization in (12) ensures that each robot i can calculate

$$\frac{\partial H_\theta(\mathbf{x})}{\partial \mathbf{x}^i} = \sum_{j \in \mathcal{N}^{i,k}} \frac{\partial H_\theta^j(\mathbf{x}_{\mathcal{N}^{j,k}}^j)}{\partial \mathbf{x}^i} \quad (13)$$

by gathering $\partial H_\theta^j(\mathbf{x}_{\mathcal{N}^{j,k}}^j)/\partial \mathbf{x}^i$ from its k -hop neighbors. Then, the control policy μ_θ of robot i becomes:

$$\mu_\theta(\mathbf{x}) = (\mathbf{F}^i(\mathbf{x}^i))^{\dagger} \left(\sum_{j \in \mathcal{N}^{i,k}} ([\mathbf{J}_\theta(\mathbf{x})]_{ij} - [\mathbf{R}_\theta(\mathbf{x})]_{ij}) \frac{\partial H_\theta(\mathbf{x})}{\partial \mathbf{x}^j} - (\mathbf{J}^i(\mathbf{x}^i) - \mathbf{R}^i(\mathbf{x}^i)) \frac{\partial H^i(\mathbf{x})}{\partial \mathbf{x}^i} \right). \quad (14)$$

Imposing the requirements in (11)-(12) is a first step towards making the control policy in (14) distributed. Note that the terms $[\mathbf{J}_\theta(\mathbf{x})]_{ij}$ and $[\mathbf{R}_\theta(\mathbf{x})]_{ij}$ might still depend on the joint state vector \mathbf{x} even though the sum runs over the k -hop neighbors in $\mathcal{N}^{i,k}$. Next, we discuss how to remove this dependence and achieve a similar factorization as (12).

First, we model $[\mathbf{J}_\theta(\mathbf{x})]_{ij}$, $[\mathbf{R}_\theta(\mathbf{x})]_{ij}$, and $H_\theta^i(\mathbf{x})$ in Eq. (14) with the parameters θ shared across the robots, so that the team can handle time-varying communication graphs. Specifically, we propose a novel architecture based on self-attention [77]. Self-attention layers extract the relationships among the inputs of a sequence by calculating the importance associated to each input using an attention map. The length of the sequences can vary as the number of parameters of the self-attention is constant with the number of inputs. Our key idea is to consider the self and neighboring states as the sequence, where each neighbor's state is an input. We now detail how to model each of the port-Hamiltonian terms.

To learn $[\mathbf{R}_\theta(\mathbf{x})]_{ij}$, robot i will use, at instant t , the state \mathbf{x}^j from all k -hop neighbors $j \in \mathcal{N}^{i,k}$. The sequence of states is given by $\mathbf{x}_{\mathcal{N}^{i,k}}^i$. The proposed architecture is composed by a sequence of layers, indexed using the subscript $w = 1, \dots, W$.

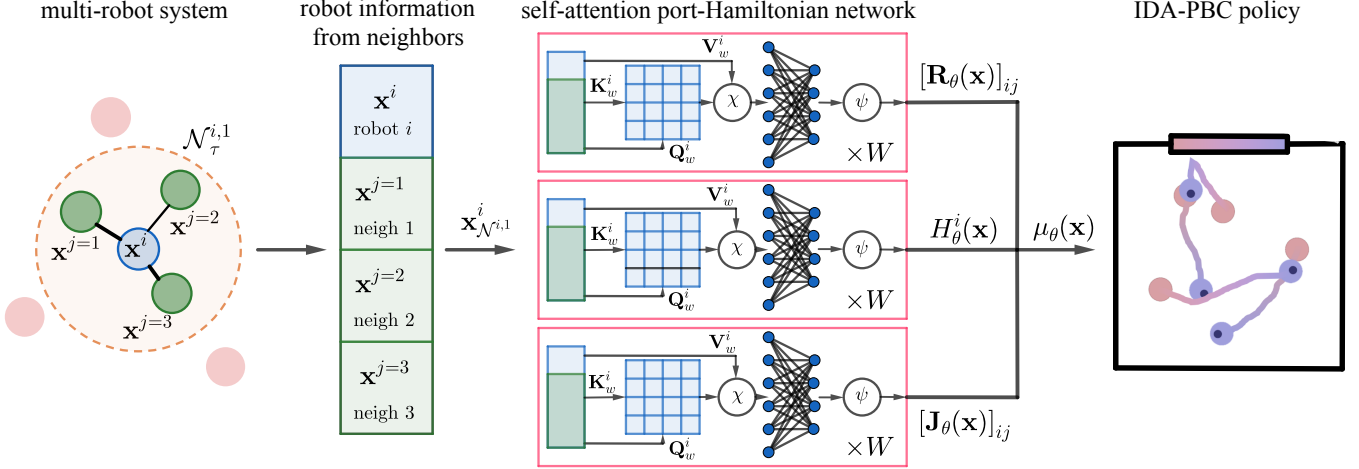


Fig. 2. Physics-informed policy parameterization. At each instant, robot i receives state information from its neighbors, typically associated to a perception or communication range. Then, three self-attention-based modules, each one associated to a component of the desired closed-loop port-Hamiltonian dynamics ($[\mathbf{R}_\theta(\mathbf{x})]_{ij}$, $[\mathbf{J}_\theta(\mathbf{x})]_{ij}$, $H_\theta^i(\mathbf{x})$), use the information from the neighbors to compute the parameters of an interconnection and damping assignment passivity-based control (IDA-PBC) policy, which is then used to execute the desired control action.

Following a self-attention mechanism, each layer computes the operations presented below:

$$\begin{aligned} \mathbf{Q}_1^i &= \mathbf{A}_1^{\mathbf{R}} \mathbf{x}_{\mathcal{N}^{i,k}}^i, \quad \mathbf{K}_1^i = \mathbf{B}_1^{\mathbf{R}} \mathbf{x}_{\mathcal{N}^{i,k}}^i, \quad \mathbf{V}_1^i = \mathbf{C}_1^{\mathbf{R}} \mathbf{x}_{\mathcal{N}^{i,k}}^i, \\ \mathbf{Y}_1^i &= \chi \left(\text{softmax} \left(\frac{\beta(\mathbf{Q}_1^i) \beta((\mathbf{K}_1^i)^\top)}{\sqrt{|\mathcal{N}^{i,k}|}} \right) \beta(\mathbf{V}_1^i) \right), \\ \mathbf{X}_1^i &= \psi(\mathbf{D}_1^{\mathbf{R}} \mathbf{Y}_1^i), \\ &\vdots \\ \mathbf{Q}_w^i &= \mathbf{A}_w^{\mathbf{R}} \mathbf{x}_{w-1}^i, \quad \mathbf{K}_w^i = \mathbf{B}_w^{\mathbf{R}} \mathbf{x}_{w-1}^i, \quad \mathbf{V}_w^i = \mathbf{C}_w^{\mathbf{R}} \mathbf{x}_{w-1}^i, \\ \mathbf{Y}_w^i &= \chi \left(\text{softmax} \left(\frac{\beta(\mathbf{Q}_w^i) \beta((\mathbf{K}_w^i)^\top)}{\sqrt{|\mathcal{N}^{i,k}|}} \right) \beta(\mathbf{V}_w^i) \right), \\ \mathbf{X}_w^i &= \psi(\mathbf{D}_w^{\mathbf{R}} \mathbf{Y}_w^i), \end{aligned} \quad (15)$$

where $\beta(\cdot)$, $\chi(\cdot)$, and $\psi(\cdot)$ are nonlinear activation functions. In the aforementioned operations, $\mathbf{A}_w^{\mathbf{R}}, \mathbf{B}_w^{\mathbf{R}}, \mathbf{C}_w^{\mathbf{R}} \in \mathbb{R}^{r_w \times h_w}$ and $\mathbf{D}_w^{\mathbf{R}} \in \mathbb{R}^{n_x \times r_w}$ for $w = 1, \dots, W$ are matrices to be learned and shared across robots; and $h_w, r_w, d_w > 0$, with $d_W = n_x$ and $h_1 = n_x$ for valid matrix multiplications. Matrices $\mathbf{A}_w^{\mathbf{R}}, \mathbf{B}_w^{\mathbf{R}}, \mathbf{C}_w^{\mathbf{R}}$ are of fixed size, so they are independent of the number of robots and neighbors. Thus, robot i can deal with time-varying neighbors. In particular, $\mathbf{A}_1^{\mathbf{R}}, \mathbf{B}_1^{\mathbf{R}}, \mathbf{C}_1^{\mathbf{R}}$ transforms the states to features encoded in the query \mathbf{Q}_1^i , key \mathbf{K}_1^i and value \mathbf{V}_1^i matrices. Then, $[\mathbf{R}_\theta(\mathbf{x})]_{ij}$ is constructed as a weighted matrix that models the interactions of robot i with its k -hop neighbors, and a diagonal positive semidefinite matrix that accounts for the self-interactions:

$$\begin{aligned} \mathbf{Z}_{ij}^{\mathbf{R}} &= \text{diag}(\mathbf{x}_W^{i,j}), \\ [\mathbf{R}_\theta(\mathbf{x})]_{ij} &= -(\mathbf{Z}_{ij}^{\mathbf{R}} + \mathbf{Z}_{ji}^{\mathbf{R}}), \quad \forall j \in \mathcal{N}^{i,k}, \\ [\mathbf{R}_\theta(\mathbf{x})]_{ii} &= \mathbf{Z}_{ii}^{\mathbf{R}} + \sum_{j \in \mathcal{N}^{i,k}} (\mathbf{Z}_{ij}^{\mathbf{R}} + \mathbf{Z}_{ji}^{\mathbf{R}}), \end{aligned} \quad (16)$$

where $\mathbf{x}_W^{i,j}$ is the column that corresponds to neighbor j in \mathbf{X}_W^i , and $\text{diag}(\cdot)$ is the operator that reshapes the $n_x \times 1$

vector to a $n_x \times n_x$ diagonal matrix. If the nonlinear activation function $\psi(\cdot)$ is designed such that the elements of the output \mathbf{X}_W^i are positive, then $\mathbf{Z}_{ij}^{\mathbf{R}}$ is a diagonal positive semidefinite matrix and $\mathbf{R}_\theta(\mathbf{x})$ is a diagonally dominant matrix. This way, $\mathbf{R}_\theta(\mathbf{x})$ is positive semidefinite by design. In fact, note the similarities between (16) and a weighted Laplacian matrix. More importantly, each element $[\mathbf{R}_\theta(\mathbf{x})]_{ij}$ depends only on the information from robot i and its neighbor j , so the computation is distributed.

To learn $[\mathbf{J}_\theta(\mathbf{x})]_{ij}$, we follow the same steps in (15), with parameters $\mathbf{A}_w^{\mathbf{J}}, \mathbf{B}_w^{\mathbf{J}}, \mathbf{C}_w^{\mathbf{J}}, \mathbf{D}_w^{\mathbf{J}}$ instead of $\mathbf{A}_w^{\mathbf{R}}, \mathbf{B}_w^{\mathbf{R}}, \mathbf{C}_w^{\mathbf{R}}, \mathbf{D}_w^{\mathbf{R}}$, to obtain encodings $\mathbf{Z}_{ij}^{\mathbf{J}}$ instead of $\mathbf{Z}_{ij}^{\mathbf{R}}$. Due to the reciprocal communication between robots i and j , we enforce the skew-symmetry of $\mathbf{J}_\theta(\mathbf{x})$ by:

$$[\mathbf{J}_\theta(\mathbf{x})]_{ij} = \mathbf{Z}_{ij}^{\mathbf{J}} - \mathbf{Z}_{ji}^{\mathbf{J}} \quad \forall j \in \mathcal{N}^{i,k}. \quad (17)$$

Since $[\mathbf{J}_\theta(\mathbf{x})]_{ii} = \mathbf{0}$, the interconnection matrix is such that $\mathbf{J}_\theta(\mathbf{x}) + \mathbf{J}_\theta^\top(\mathbf{x}) = \mathbf{0}$ and, thus, is skew-symmetric by design. Again, each element $[\mathbf{J}_\theta(\mathbf{x})]_{ij}$ depends only on the information from robot i and its neighbor j , so the computation is distributed.

Finally, to learn $H_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i)$, we represent it as follows:

$$H_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i) = \text{vec}(\mathbf{x}_{\mathcal{N}^{i,k}}^i)^\top \mathbf{M}_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i) \text{vec}(\mathbf{x}_{\mathcal{N}^{i,k}}^i) + U_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i). \quad (18)$$

The first term in the right-hand side of Eq. (18) is a kinetic-like energy function and the second terms is a potential energy function with

$$\mathbf{M}_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i) = \text{diag}(\mathbf{1}^\top \mathbf{Z}_i^{\mathbf{M}}) \quad \text{and} \quad U_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i) = \mathbf{1}^\top \mathbf{Z}_i^{\mathbf{U}} \mathbf{1}.$$

The encodings $\mathbf{Z}_i^{\mathbf{M}}$ and $\mathbf{Z}_i^{\mathbf{U}}$ are calculated using the same steps in (15), with parameters $\mathbf{A}_w^{\mathbf{M}}, \mathbf{B}_w^{\mathbf{M}}, \mathbf{C}_w^{\mathbf{M}}, \mathbf{D}_w^{\mathbf{M}}$ and $\mathbf{A}_w^{\mathbf{U}}, \mathbf{B}_w^{\mathbf{U}}, \mathbf{C}_w^{\mathbf{U}}, \mathbf{D}_w^{\mathbf{U}}$, respectively. Once we have $H_\theta^i(\mathbf{x}_{\mathcal{N}^{i,k}}^i)$, we obtain $\partial H_\theta^j(\mathbf{x}_{\mathcal{N}^{j,k}}^j) / \partial \mathbf{x}^i$ from all the neighboring robots and compute Eq. (13). Therefore, since all the operations are only dependent on the available information in robot i and

its neighbor j , the computations involving the Hamiltonian function are distributed by design.

To deploy the control policy (14), we design a message $\mathbf{m}_{\tau,j}^{i,j}$, encoding information that robot i needs from robot j at time τ to calculate $[\mathbf{J}_{\theta}(\mathbf{x})]_{ij}$, $[\mathbf{R}_{\theta}(\mathbf{x})]_{ij}$, and $H_{\theta}^i(\mathbf{x})$. When there is no communication among robots, i.e., $k = 0$, no message is needed. For $k \geq 1$, robot i uses the following communication protocol:

- 1) Robot i receives messages $\mathbf{m}_{\tau,1}^{i,j} = \mathbf{x}_j$ from its k -hop neighbors in $\mathcal{N}^{i,k}$. Then, robot i computes $\mathbf{Z}_{ij}^{\mathbf{R}}$, $\mathbf{Z}_{ij}^{\mathbf{J}}$, H_{θ}^i , and $\partial H_{\theta}^i / \partial \mathbf{x}_j$.
- 2) Robot i receives messages $\mathbf{m}_{\tau,2}^{i,j} = \{\partial H_{\theta}^j / \partial \mathbf{x}_i, \mathbf{Z}_{ji}^{\mathbf{R}}, \mathbf{Z}_{ji}^{\mathbf{J}}\}$ from its k -hop neighbors in $\mathcal{N}^{i,k}$ and computes $\partial H_{\theta} / \partial \mathbf{x}_i$, $[\mathbf{J}_{\theta}]_{ij}$, $[\mathbf{R}_{\theta}]_{ij}$.
- 3) Robot i receives messages $\mathbf{m}_{\tau,3}^{i,j} = \partial H_{\theta} / \partial \mathbf{x}_j$ from its k -hop neighbors in $\mathcal{N}^{i,k}$ and computes the control input \mathbf{a}_i .

In summary, each robot i receives a message $\mathbf{m}_{\tau,j}^{i,j} = [\mathbf{m}_{\tau,1}^{i,j}, \mathbf{m}_{\tau,2}^{i,j}, \mathbf{m}_{\tau,3}^{i,j}]$ in 3 communication rounds from its neighboring robot j . We assume negligible delays between communication rounds. If the delay is large, Wang et al. [65] suggest to learn a function that predicts quantities such as $\partial H_{\theta}(\mathbf{x}) / \partial \mathbf{x}_j$, $\mathbf{Z}_{ji}^{\mathbf{J}}$, $\mathbf{Z}_{ji}^{\mathbf{R}}$, leading to one communication round. We leave this for future work.

V. PHYSICS-INFORMED MULTI-ROBOT SOFT ACTOR-CRITIC

The previous section describes how to parameterize the distributed control policies to be learned. In this section, we present a soft actor-critic algorithm [80] to train the self-attention port-Hamiltonian neural network, detailing the main features that allow the integration of our physics-informed policy representation. The key components are presented in Fig. 3, namely: (a) actor, (b) environment, (c) reward, and (d) critic.

A. The actor

First, we model the actor with the port-Hamiltonian system detailed in Sec. IV-A and we use the IDA-PBC policy in Eq. (14) parameterized by the self-attention architecture proposed in Sec. IV-B. The known robot dynamics are typically provided by the simulation environments or the hardware specifications of the robots [81]. Nevertheless, two aspects must be adapted to match the soft actor-critic formulation: (i) the MDP formulation and the stochastic Markov policy are in discrete time while the robot dynamics and the IDA-PBC are in continuous-time, and (ii) the actions in the soft actor-critic method are stochastic while the port-Hamiltonian formulation is deterministic.

To address the continuous versus discrete time mismatch, we use the zero-order hold control approach described in Sec. III. More specifically, $\mathbf{u}(t) = \mathbf{a}_{\tau} = \mu_{\theta}(\mathbf{x}(t_{\tau}))$ for $t \in [t_{\tau}, t_{\tau+1})$. This means that we can write the control policy as dependent on the discrete-time state \mathbf{s}_{τ} rather than the continuous-time state $\mathbf{x}(t_{\tau})$, leading to $\mu_{\theta}(\mathbf{s}_{\tau})$ with the same expression in Eq. (14).

To address stochasticity and maintain the desired distributed structure already derived in Sec. IV-A and IV-B, we model

the distributed control policies as (squashed) Gaussian distributions, whose mean is given by the IDA-PBC controller in (14), therefore respecting the distributed policy factorization. Meanwhile, the variance of the policy distribution is provided by a neural network that is learned during training. Overall, we obtain a distributed and stochastic Markov control policy of the form:

$$\mathbf{a}_{\tau}^i \sim \pi_{\theta}(\mathbf{a}_{\tau}^i | \mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i) = \tanh(\mu_{\theta}(\mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i) + \sigma_{\theta}(\mathbf{a}_{\tau}^i, \mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i)\xi), \quad (19)$$

where $\mu_{\theta}(\mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i)$ is the IDA-PBC policy parameterized by self-attention neural networks under the restrictions imposed by the networked structure of the robot team.

On the other hand, $\sigma_{\theta}(\mathbf{a}_{\tau}^i, \mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i)$ is a vector of standard deviations given by a neural network that approximates the variance of the squashed Gaussian distribution with parameters θ . Besides, $\xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, so the control policy is a Gaussian distribution with mean $\mu_{\theta}(\mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i)$ and diagonal covariance matrix whose diagonal is equal to $(\sigma_{\theta}(\mathbf{a}_{\tau}^i, \mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i))^2$. The control is constrained to $\mathbf{a}_{\tau}^i \in [-1, 1]^{n_u}$ by means of a tanh function [56], [80], [82], leading to a squashed Gaussian policy. In practice, the control input can be constrained to $\mathbf{a}_{\tau}^i \in [a_{\min}, a_{\max}]^{n_u}$ with $-\infty < a_{\min} \leq a_{\max} < \infty$ by scaling the output of the tanh function. It is also important to remark that $\sigma_{\theta}(\mathbf{a}_{\tau}^i, \mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i)$ only depends on the available information at each robot, such that the desired distributed factorization of the control policy is preserved.

During training, the control inputs are sampled from the Gaussian distribution, as stated in Eq. (3). After training, the mean control input is chosen, i.e., $\mathbf{a}_{\tau}^i = \tanh(\mu_{\theta}(\mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i))$. In the case that an additional layer of robust, adaptive or active control is desired, then $\sigma_{\theta}(\mathbf{a}_{\tau}^i, \mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i)$ can be used as a proxy of the uncertainty in the control policy, since it only depends on the neighboring information at robot i . The design of $\sigma_{\theta}(\mathbf{a}_{\tau}^i, \mathbf{s}_{\mathcal{N}_{\tau}^{i,k}}^i)$ is free to choose, but in this work we opt for the same architecture in Eq. (15).

B. The environment

The second modification over the soft actor-critic algorithm is in the collection of experiences from the environment. The environment is determined by the desired task to be solved, with examples found in Fig. 1. The main difference with other multi-agent reinforcement learning works is that, to build the replay buffer that stores the trial and error experience used for training the control policies, we take into account all the interactions among the robots. Differently from centralized-training decentralized-deployment approaches, where the experience of each agent is recorded independently of the other agents, we record for each experience all the robot states, actions, and interaction graph together. Thanks to that, we keep track of the correlation among robots during training. This also allows us to condition the trained control policies on the available information at each robot, therefore only providing the information robots will have access during deployment.

C. The reward

The reward function can be shared across robots and include global terms because the actor is the whole multi-robot

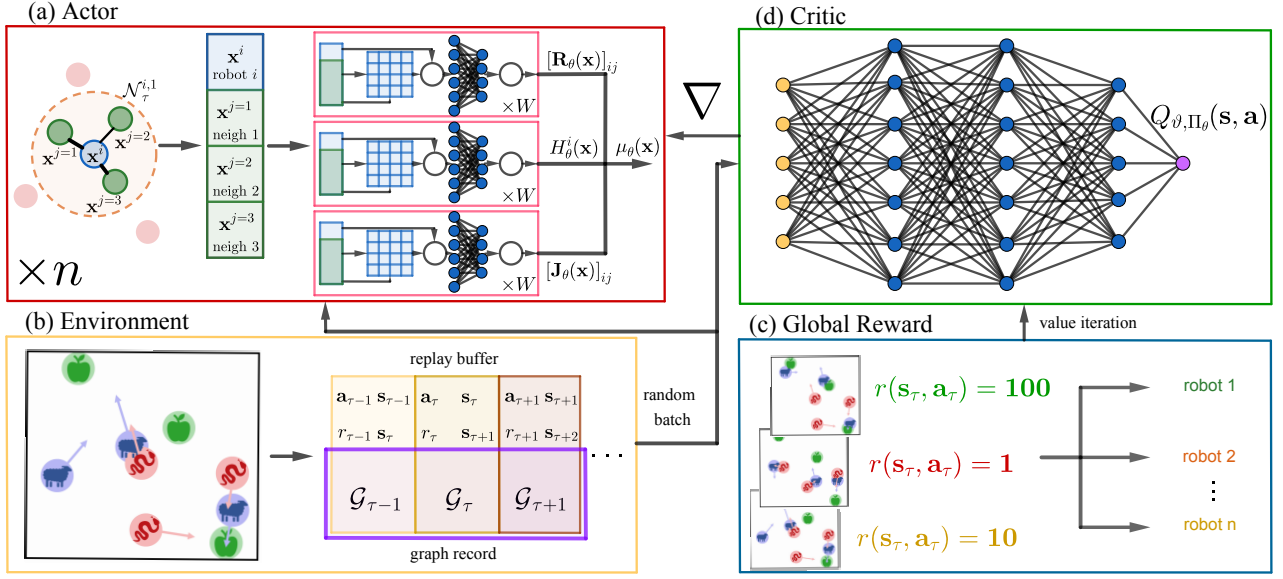


Fig. 3. Overview of our physics-informed soft actor-critic multi-agent reinforcement learning approach. The main differences with respect of other actor-critic methods are the following: (a) the actor is the multi-robot network modeled as a port-Hamiltonian system with a self-attention-based IDA-PBC policy; (b) the replay buffer not only stores actions, states and rewards, but also the graph structure of the multi-robot network to enforce the desired distributed structure; (c) the reward is global because the actor is the whole multi-robot team and the physics-informed parameterization conditions the policy on the graph structure of the team; (d) the output of the critic is shared across robots and the policy parameters are the same for all robots, so for the same critic gradient step, n policy gradient steps are taken.

network. This is important because, from the perspective of the soft actor-critic algorithm during training, the whole robot network is a single centralized agent. Other state-of-the-art multi-agent reinforcement learning approaches (see Sec. II) require a factorization of the reward function because each agent is considered as an isolated learning unit which does not exploit neighboring information during the execution of the control policy. In contrast, in our approach the shared distributed control policy is simultaneously learned at all the robots used in training, and since the experience records the underlying communication graph, we can seamlessly associate global rewards with distributed cooperative control policies through communication. It is worth to note that, as in any classical soft actor-critic approach, the maximization objective is changed to include an entropy term $\mathcal{H}(\Pi_\theta(\mathbf{a}|\mathbf{s}))$ that measures the entropy of the control policy. This entropy term, weighted by a temperature parameter $\alpha > 0$, trades off exploration (high value) and exploitation (low value). The value of α changes over time according to a gradient descend law which manages automatically the exploration/exploitation dichotomy. The details of this automatic temperature adjustment rule can be found in [80].

D. The critic

Finally, regarding the critic, its main purpose is to learn the action-value function of the environment, encoded in $Q_{\Pi_\theta}(\mathbf{s}, \mathbf{a})$. This approximation steers the training of the control policy towards the maximization objective. Nevertheless and importantly, the critic is only used during training. Thus, it is not necessary to design it to be distributed and an existing centralized neural network architecture can be used. In this

work, we opt for a multi-layer perceptron with parameters ϑ to learn $Q_{\vartheta, \Pi_\theta}(\mathbf{s}, \mathbf{a})$. As with the reward, since, from the perspective of the soft actor-critic algorithm, the single agent is the whole robot network described by the port-Hamiltonian dynamics, a single centralized critic can learn the appropriate action-value function to condition the training of the distributed control policy.

A shared centralized action value function implies that each optimization step considers, simultaneously, all the policies gathered in the joint policy Π_θ . Since the policies are homogeneous, each optimization step is simultaneously updating the policy π_θ n times. This is achieved without any particular specification of the gradients nor factorization of the action value function. This is one of the key properties and advantages of our proposed soft actor-critic algorithm compared to other state-of-the-art algorithms. Typical multi-agent reinforcement learning approaches, by factorizing the policy and/or the value functions, arrive to a different set of policy parameters, one per agent used in training. In cooperative tasks where the robots share the same goal, there should be a single policy that resolves the task independently on the configuration of the robot. In our case, by appropriately integrating the modular port-Hamiltonian description of the multi-robot system with the soft actor-critic algorithm, we train a single distributed control policy which takes into account the current available information at each robot. Using the soft actor-critic update rules [56], [80], [82], the gradients with respect to θ are, simultaneously for all the robots, already conditioned on the multi-robot communication topology.

In conclusion, in contrast with other solutions, ours explicitly considers the exchanges of information among robots by

modeling the multi-robot system as a graph. This allows to use a reinforcement learning algorithm for single-robot problems, where the single agent is the multi-robot network. From the perspective of the reinforcement learning algorithm, the actor is centralized. However, by means of a physics-informed self-attention parameterization of the dynamics and control of the robots, the learned policies are distributed by design, modeling the robot team as a modular port-Hamiltonian system.

VI. RESULTS

To assess our physics-informed multi-agent reinforcement learning approach, we present seven multi-robot scenarios. The first three are extracted from the VMAS simulator [81]. The next three are adaptations from the MPE simulator [59], [83] that can be found in [65]. The last scenario comes from the Multi-Agent MuJoCo benchmark [2]. The overview of the scenarios are detailed in the following:

- a) **Reverse transport:** the robots are randomly spawned inside a box that they must push towards a desired landmark in the arena. The initial position of the box and the landmark is random. Compared to [81], we decrease the mass of the box to 1 kg to ensure that the box can be moved even with a small number of robots. Also, when the box is in the landmark, the reward is set to 1.
- b) **Sampling:** the robot team is randomly spawned in an arena with an underlying Gaussian density function composed of 3 modes. The field is discretized to a grid. Robots must collect samples of the field such that once a robot visits a cell its sample is collected without replacement and given as reward to the team. Robots use a LiDAR to sense each other, and they observe the samples in the 3×3 grid around it.
- c) **Navigation:** each robot has a landmark to reach. The initial position of the robots and landmarks are randomly spawned in a 2×2 m square arena. Robots must navigate to reach their corresponding landmarks while avoiding collisions with other robots. Compared to [81], we encourage collision avoidance by changing the collision penalty from -1 to -5 . Besides, each robot only observes its desired landmark instead of all the landmarks.
- d) **Food collection:** it is a version of the simple spread scenario [59] where a team of robots and food landmarks is randomly spawned in an arena. There are as many landmarks as robots. Robots must cooperate to cover as many food landmarks as possible. Each time a robot covers a new landmark, the whole team is rewarded.
- e) **Grassland:** A team of robots must collect food resources that are randomly spawned in an arena, while evading multiple predators. There are as many robots as predators, and the former move twice faster than the latter. The predator team is positively rewarded when some member captures a robot, whereas the robot team is negatively rewarded and the robot is deactivated. On the other hand, the robots receive a positive reward if they reach a food landmark, which is then spawned again in a new position.
- f) **Adversarial:** two robotic teams compete for the same food landmarks. Both teams have the same number of

robots. When a member of a team reaches a food landmark, the landmark is randomly spawned and the team is positively rewarded. When two members of a team collide with one member of the other team, then the first team is positively rewarded, while the second team is penalized and the robot deactivated.

- g) **6x1-Half Cheetah:** a robotic cheetah with two legs and 6 joints, where each joint is a different agent. The multi-agent team is distributed in the sense that each agent can communicate only with the adjacent joints, enforced by a ring graph topology. The task is to make the robot run as fast as possible. The reward is global and is composed by a first term that favours forward movements and a second term that penalizes too large actions.

The first three scenarios are used for ablation studies, where the main goal is to evaluate how our physics-informed policy parameterization improves upon existing standard policy parameterizations. The second three are used to compare our method with other state-of-the-art multi-agent reinforcement learning approaches that are not physics-informed. The last scenario is used to validate our approach in a realistic robotic platform.

All robots have a communication radius $r_{comm} > 0$ which allows to exchange information with 1-hop neighbors. The communication radius and other hyperparameters of the soft actor-critic algorithm are specified in Appendix B. Each robot observes its position, velocity, position and velocity of the landmarks or objects of interest (e.g., the box in reverse transport). The particularities of the observation space of each robot can be found in [59], [65], [81]. In some scenarios, the robot observation vector changes its dimension depending on the number of robots. Specifically, in the food collection, grassland and adversarial scenarios the number of landmarks changes with the number of robots. Therefore, to accommodate the proposed self-attention port-Hamiltonian neural network with an observation vector which may change its size, we use an additional neural network to pre-process the observation vector. In particular, we concatenate a self-attention layer with a dense layer which receives, as input, the food landmarks' positions and outputs a feature vector of constant dimension that is used as part of the state vector to build \mathbf{S}_t^i . The details can be found in Appendix A. Supplementary material can be found in our repository¹ and the supplementary video².

A. Ablation results

We conduct ablation studies using the reverse transport, sampling and navigation scenarios. The goal is to evaluate whether our physics-informed policy parameterization improves upon existing standard policy parameterizations. We compare the proposed self-attention port-Hamiltonian neural network (pH-MARL) with a classical multi-layer perceptron (MLP), a modular self-attention-based neural network (MSA) and a graph- and attention-based neural network (GSA). The implementation details can be found in Appendix A.

¹<https://github.com/EduardoSebastianRodriguez/phMARL>

²<https://youtu.be/pSzP3LBVyZg>




Fig. 4. Comparison of the performance of the ablated control policies when we scale the number of robots in deployment. In all the scenarios, our proposed combination of a port-Hamiltonian modeling and self-attention-based neural networks achieves the best cumulative reward without further training the control policy. Each bar displays the mean and standard deviation of \bar{R} over 10 evaluation episodes.

The MLP network is an unstructured neural network that receives, as input, the current states and actions of the robots and outputs the action vector. The MSA network replaces the dense layers of the MLP network with an architecture based on self-attention, but has the same input and output. The GSA network is similar to the MSA network but has the adjacency matrix as additional input. The three neural networks constitute a sequential improvement from a standard neural network to our port-Hamiltonian formulation. It is important to note that, with the current formulation, it is not possible to try a port-Hamiltonian neural network which is not based on self-attention because each element of the processed input sequence is employed to compute the port-Hamiltonian terms, that are also a sequence spanned from $i \in \mathcal{N}_t^{i,k}$. This is key for the scalability of the policy and is a feature that is not provided by other architectures based on, e.g., multi-layer perceptrons or convolutions.

Finally, we employ the same architecture used in pH-MARL to learn the standard deviation for all the options, along with the same critic and soft actor-critic algorithm (Appendix A).

Table I shows the averaged cumulative reward $\bar{R} = \frac{1}{n} R = \frac{1}{n} \sum_{\tau=0}^{\tau_{\max}-1} r(\mathbf{s}_\tau, \mathbf{a}_\tau)$ after convergence of the training, where τ_{\max} is the maximum number of steps per episode. For the three scenarios and the four neural network architectures, we use $n = 4$ robots. The table reports the mean of \bar{R} over 10 evaluation episodes. In all the cases, pH-MARL surpasses the other three architectures, with a particularly significant difference in the reverse transport and navigation scenarios. Given the same number of parameters in the four architectures, the use of a physics-informed formulation of the neural network leads to a structured learning with efficient sampling in training. The difference in the reverse transport scenario is due to the fact that the other networks sometimes fail in reaching the landmark, despite arriving to a very close position. In the sampling scenario, pH-MARL is the fastest in inspecting the environment, therefore covering more informative cells. It is followed, in decreasing order, by GSA, MSA and MLP, which is reasonable since the three neural networks are in decreasing order of architecture complexity in terms of neural network modules and input information. The navigation scenario is the simplest among the three. Nevertheless, pH-MARL is still the best because it better learns the collision avoidance constraint.

To assess the scalability of the different architectures, in

TABLE I
AVERAGED TRAINING CONVERGENCE CUMULATIVE REWARD FOR THE FOUR ABLATED METHODS. IN ALL THE CASES, $n = 4$ ROBOTS ARE USED.

Method	Mean and std of \bar{R} over 10 evaluation episodes		
	Reverse transport	Sampling	Navigation
pH-MARL	213 \pm 21	161 \pm 41	-53 \pm 101
MLP	64 \pm 38	73 \pm 25	-280 \pm 98
MSA	57 \pm 43	82 \pm 32	-353 \pm 99
GSA	90 \pm 49	89 \pm 38	-204 \pm 87

Fig. 4 each trained neural network is evaluated with a different number of robots, showing the mean and standard deviation of \bar{R} over 10 episodes. The first conclusion is that pH-MARL and GSA scale well with the number of robots, achieving the same cumulative reward per robot for all team sizes. In contrast, MSA and MLP do not scale because their architectures considers directly the global observation and action vectors, so their modules are not ready to process an input vector of different dimension. On the other hand, pH-MARL achieves much better performance than the other networks in all the scenarios and team sizes, thus confirming the importance of using physics-informed priors to ease the learning. The single exception is in the navigation example, where the performance deteriorates when the number of robots increases.

Fig. 5 provides some qualitative examples for the three scenarios and helps to understand some limitations of our approach. Beginning with the navigation scenario, we can see that the size of the arena does not scale with the number of robots, it is always 2×2 m. Hence, taking into account that the robots have a radius of 15cm, the space is very small to navigate towards the landmarks without colliding. According to the reward, the robots search for safe motions, so they prefer to avoid collisions rather than moving to their landmarks (Fig. 5, navigation with 8 robots). There are other features that do not scale in the reverse transport and sampling scenarios: size of the box, weight of the box, size of the arena and number of Gaussian modes. Nonetheless, pH-MARL is capable of generalizing, scaling and circumventing these issues and achieve the desired tasks. In the reverse transport scenario, even with 16 robots that are highly packed inside the box, the robots manage to push the box and avoid erratic movements caused by the low weight of the box. In the sampling scenario, robots learn to spread and coordinate to cover more cells when

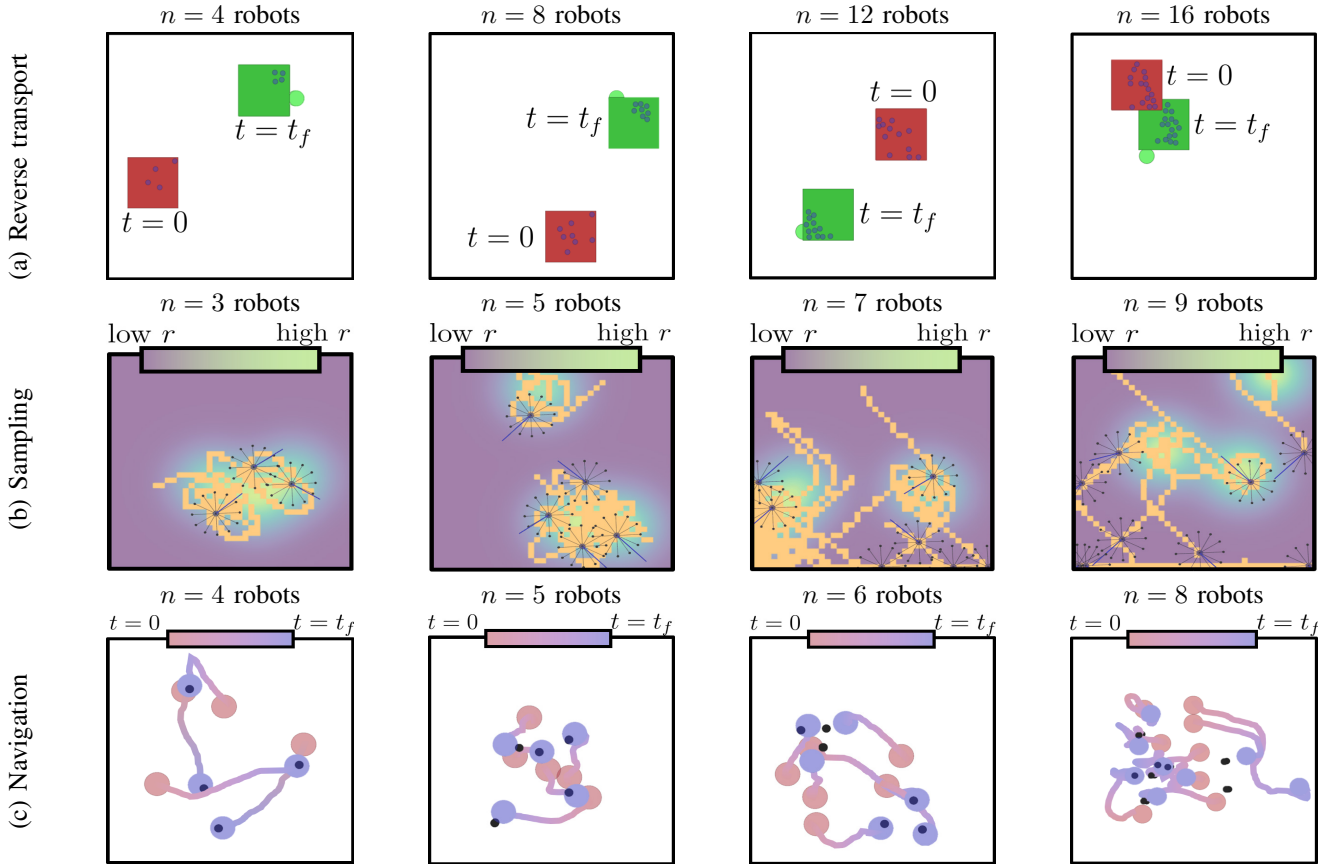


Fig. 5. Examples of multi-robot scenarios for different initial conditions and numbers of robots. In the reverse transport scenario, the red box turns green when it reaches the desired goal (green dot), and the robots are depicted as small blue dots. In the sampling scenario, LiDAR measurements are depicted as points forming a star pattern centered at the position of the robot, and regions already explored are shown in orange. In the navigation scenario, evolution over time of the trajectories is colored from pink ($t = 0$) to blue ($t = t_f$), whereas the goals are depicted as black dots. It is interesting to see that some aspects of the environment do not scale with the team size, e.g., the size and weight of the box in the reverse transport scenario, the number of hot spots in the sampling scenario or the size of the arena in the navigation scenario.

there are more robots. Interestingly, when there are a lot of robots ($n = 9$), some robots move to the corners because they annoy the others by increasing the potential number of collisions.

Summing up, the combination of a port-Hamiltonian formulation and self-attention mechanisms leads to superior performance and the desired scalability with the number of robots, learning control policies that are distributed by design and fully exploit the graph structure of the multi-robot system.

B. Comparative results

We compare the performance of our proposed physics-informed multi-agent reinforcement learning approach with other state-of-the-art multi-agent reinforcement learning approaches that are not physics-informed. In particular, we compare with: Multi-Agent Deep Deterministic Policy Gradient (MADDPG) [59], Mean Field Actor Critic (MFAC) [64], Evolutionary Population Curriculum (EPC) [84], Distributed multi-Agent Reinforcement Learning with One-hop Neighbors (DARL1N) [65] and Multi-Agent Proximal Policy Optimization (MAPPO) [60]. The comparison results of the first four methods reported in this paper are extracted from [65].

We use the food collection, grassland and adversarial scenarios for the comparison. Following the same procedure in

[65], for the scenarios with an adversarial team (grassland and adversarial), we first use MADDPG to learn the control policies of both teams. Then, the adversarial team control policy is frozen and the policy of the other team is learned using any of the aforementioned approaches.

Next, we assess the scalability of our proposed approach in this second series of scenarios. We train the other state-of-the-art approaches for each specific number of robots, ranging from 3 to 48 robots. In contrast, we train pH-MARL $n = 4$ robots and evaluate the trained control policies with the other number of robots, without further training. Fig. 6 demonstrates that pH-MARL achieves better or similar results than the other approaches without further training the control policies. When the team size is close to the one used during training ($n = 3, 6, 12$) pH-MARL surpasses the other methods in the three scenarios, proving the accuracy of combining a physics-informed description of the multi-robot team with self-attention mechanisms. In this sense, MAPPO in the grassland scenario with $n = 6$ robots is the only case that obtains better results than pH-MARL. When the team size is much greater than the one used during training (24, 48), pH-MARL outperforms (MADDPG, MFAC) or achieves similar results (EPC, DARL1N, MAPPO) as the state-of-the-art. Notably, pH-MARL outperforms all the state-of-the-art methods in the




Fig. 6. Comparison of our proposed physics-informed multi-agent reinforcement learning approach with other state-of-the-art approaches. To measure performance, we use the averaged cumulative reward $\bar{R} = \frac{1}{n} R = \frac{1}{n} \sum_{\tau=0}^{T_{\max}-1} r(s_{\tau}, a_{\tau})$. pH-MARL is only trained with $n = 4$ robots and deployed with different team sizes, while the state-of-the-art control policies are trained for each specific number of robots. For team sizes similar to the one used in training, our policy outperforms the other approaches. For team sizes much larger than the training team size, our policy still outperforms or achieves similar results as the state-of-the-art approaches.

adversarial scenario even when $n = 48$, which is significant considering that they are all trained with $n = 48$ robots except pH-MARL, which is trained with $n = 4$ robots. Regarding MAPPO, despite its strong performance in some cases (e.g., food collection with $n = 3$ robots, grassland with $n = 6$ robots or grassland with $n = 48$ robots), it generally exhibits poor results when the number of robots is larger than 10 robots. Besides, MAPPO is unable to learn successful policies that avoid collisions and find the food resources for any of the teams in the adversarial scenario.

The reason why pH-MARL exhibits a slight drop in performance when $n = 24, 48$ robots is the same discussed in the ablation experiments. There are some features of the scenarios, like the size of the arena, that do not scale with the number of robots. Robots are not trained to move in configurations that far from those experienced during training. How to design the control policy to be invariant to environmental changes is part of the future work. Qualitative examples can be found in the repository associated to the paper.

Regarding the computational cost, it is worth noting that, given the same scenario, optimizer and replay buffer hyperparameters, the main differences among methods come from the parameterization of the actor. In this sense, our approach involves greater training times *per sample and robot* compared to other non-physics-informed methods. Specifically, the increase in computational burden comes from the differentiation of the Hamiltonian $\partial H_{\theta}(\mathbf{x})/\partial \mathbf{x}_j$ at each time step. However, by restricting the space of admissible policies, we have better sample efficiency during training compared to non-physics-informed parameterizations. Therefore, the increase of computational burden per sample and robot is compensated by the sample efficiency of the method, leading to similar overall computational times in robot tasks during training. Besides, as it has been observed in the scalability results, our approach needs a very small number of robots to reproduce or even outperform the other methods when they are trained with larger numbers of robots, saving additional computation. Importantly, the computational cost during deployment scales with the number of neighbors rather than the network size, therefore preserving the desirable properties coming from

a distributed multi-robot setting. Alternatively, to improve computational efficiency, one could consider a neural network estimator that learns to infer $\partial H_{\theta}(\mathbf{x})/\partial \mathbf{x}_j$ from $\mathbf{x}_{\mathcal{N}^i, k}^i$, which bypasses automatic differentiation of $H_{\theta}(\mathbf{x})$ and the three message-passing communication protocol. Avoiding automatic differentiation to reduce inference time is an active area of research [85], and we leave its exploration to optimize the training time of our method for future work.

C. Sensitivity analysis

Regarding the evaluation of the training scalability of our approach, we have conducted a series of experiments using the same scenarios of Section VI-A and Section VI-B. Given those scenarios, we have trained our pH-MARL with an increasing number of robots, namely, $n = \{2, 4, 8, 12\}$ (see Fig. 7). The first interesting conclusion is that the increase in number of robots does not necessarily lead to better performance because the interactions among robots always happen at the robot neighborhood level. In the reverse transport scenario, more robots inside the cage implies a greater number of robot contacts with the cage and, therefore, a drop in averaged cumulative reward. In the sampling scenario, a greater number of robots leads to an improvement in performance because, for the same time horizon, robots can explore the arena faster and localize the informative regions faster. In the case of the navigation scenario, due to space constraints, an increase in the number of robots during training leads to a significant drop in averaged cumulative reward because the space available to each robot is smaller, so more collisions happen (specially with $n = 12$). In all the scenarios, $n = 2$ is a too small multi-robot team size to capture the richness of distributed multi-robot interactions, so the training process is unable to learn a successful policy in terms of averaged cumulative reward.

Therefore, in general, the optimal number of training robots depends on constraints of the scenario and the task. Nonetheless, as a rule of thumb, increasing the number of training robots towards a large multi-robot team only leads to greater sample complexity, whereas decreasing the number of training robots does not allow to learn sufficiently varied behaviors to

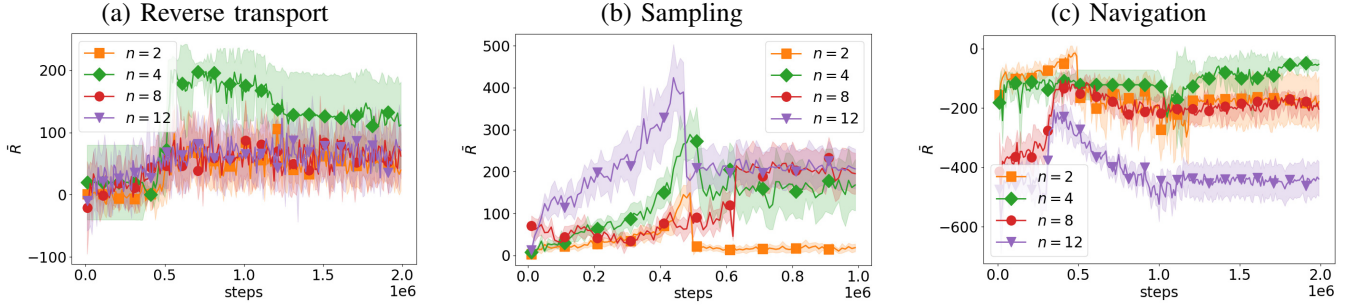


Fig. 7. Sensitivity analysis of pH-MARL with respect to the number of robots at training. To measure performance, we use the averaged cumulative reward $\bar{R} = \frac{1}{n} R = \frac{1}{n} \sum_{\tau=0}^{\tau_{\max}-1} r(\mathbf{s}_\tau, \mathbf{a}_\tau)$. Compared to the nominal value ($n = 4$), increasing the number of training robots only leads to greater sample complexity, whereas decreasing the number of training robots does not allow to learn sufficiently varied behaviors to scale with the number of robots.

TABLE II
PHYSICAL PROPERTIES OF THE JOINTS OF THE MUJoCo HALF CHEETAH.

Joint	Type	Range [rad]	Damping [Ns/m]	Stiffness [N/m]
Back thigh	hinge	$[-0.52, 1.05]$	6	240
Back shin	hinge	$[-0.785, 0.785]$	4.5	180
Back foot	hinge	$[-0.4, 0.785]$	3	120
Front thigh	hinge	$[-1, 0.7]$	4.5	180
Front shin	hinge	$[-1.2, 0.87]$	3	120
Front foot	hinge	$[-0.5, 0.5]$	1.5	60

scale with the number of robots. Henceforth, we conclude that our approach is best suited for training with a small number of training robots (relative to the task/scenario) but still sufficiently large to capture all possible neighborhood interactions.

From a behavioral analysis perspective, the results in Fig. 7 suggest that the learned distributed multi-robot policies benefit from small neighborhoods, irrespective of the network size. In scenarios with a fixed arena size, an increase in the number of robots indirectly leads to an increase in neighborhood sizes, with a subsequent drop in averaged cumulative reward both in training and evaluation (see Fig. 6). The appropriate number of training robots depends on the scenario, e.g., in the reverse transport scenario the best performance is achieved with $n = 4$ while in the sampling scenario this is obtained with $n = 8$. In any case, as concluded above, our approach requires a sufficiently large team size to capture a variety of local interactions. On the other hand, different from other centralized-training decentralized-execution approaches, our approach considers a single set of policy parameters that is invariant to the number of robots in the team due to its distributed structure. Therefore, behavioral analysis to understand individual influence in the overall team reduces to understanding how different neighborhood sizes affect overall performance, which is already given by the sensitivity analysis with respect to the number of training robots for a fixed environment size, and the scalability tests shown in Fig. 4 and Fig. 6.

D. Validation on a realistic robot platform

We further validate pH-MARL in a realistic robot setting using MuJoCo [86]. MuJoCo is a general purpose physics




Fig. 8. Average episode reward obtained by pH-MARL and other MARL approaches in the 6x1-Half Cheetah environment.

engine that replicates multi-joint dynamics with contact, allowing for fast yet realistic physical behaviors. Specifically, we use the Multi-agent MuJoCo benchmark [2], where the task is to learn coordination policies for multi-joint robots that make them walk. Our approach was evaluated using the 6x1-Half Cheetah environment (see Fig. 1e). Since our fundamental contribution is the physics-informed distributed policy parameterization, we integrated our policy design in a Trust-Region Policy Optimization (TRPO) reinforcement learning algorithm [87]. This demonstrates the flexibility of the physics-informed policy parameterization in the sense that it can be integrated in different reinforcement learning algorithms. The benchmark provides results from state-of-the-art multi-agent reinforcement learning approaches for comparison, including Heterogeneous Agent Trust Region Policy Optimization (HATRPO) and Heterogeneous-Agent Proximal Policy Optimization (HAPPO) [62], MAPPO, IPPO [66], and MADDPG. Importantly, in all these state-of-the-art methods, the policy of each agent/joint depends on the global observation vector, whereas our policy only depends on the neighboring information, enforcing a ring topology.

The 6x1-Half Cheetah model is composed of $n = 6$ joints, corresponding to: {back thigh, back shin, back foot, front thigh, front shin, front foot}. The input to each joint is a torque with [Nm] units. The local observation of each agent is the angle, angular velocity, linear velocity, angular acceleration and linear acceleration of the joint; and the (x,z)-position, angular velocity and (x,z)-linear velocity of front tip. The physical properties of the joints follow those of the original MuJoCo Half Cheetah [86], detailed in Table II.




Fig. 9. Quantitative results of the pH-MARL policy for multi-robot navigation in the Robotarium simulator. From left to right, the following metrics are reported for different numbers of robots and communication quality (baseline means perfect communications): travel distance, travel time and success rate. All the metrics are the average over 10 random runs.




Fig. 10. Qualitative results of the pH-MARL policy for multi-robot navigation in the Robotarium simulator. In (a), communication is perfect, whereas (b), (c) and (d) operate under imperfect communication. In the imperfect communication cases, magenta links denote delays, red links denote packet loss, yellow links denote disturbances, and green links denote unperturbed links.

Importantly, physical interactions like friction of the joints with the armature of the main body are computed numerically from the stiffness and damping of the joint. Stiffness is modeled as a spring with equilibrium at the nominal position of the joint, whereas the damping is modeled as a force linear in velocity.

As it can be observed in Fig. 8, our method, although less informed than the others due to the networked constraints, improves upon IPPO, MADDPG and MAPPO in average episode reward, and obtains a similar final average episode reward to HAPPO and HATRPO. Besides, by fully exploiting the physical dynamics of the robots, our method presents better sample efficiency than MADDPG, IPPO and MAPPO, and similar sample efficiency than HAPPO and HATRPO. In this sense, our method leverages the fact that the acceleration of the joints is available, so in the open-loop feed-forward term $\dot{\mathbf{x}}$ is known without explicit knowledge on the values of $\mathbf{J}(\mathbf{x})$, $\mathbf{R}(\mathbf{x})$ and $\partial H(\mathbf{x})/\partial \mathbf{x}$. This experiment shows that our approach is a promising method for multi-joint dynamical robots. The sample efficiency achieved by distributing the policy according to the imposed ring topology is reduced when the connectivity of the network increases. The ring topology restricts the neighborhood of joint i to $\{i-1 \bmod 6, i, i+1 \bmod 6\}$, which implies a very sparse graph. This serves as a prior for how the other joints affect joint i that restricts the class of policies that can be learned, leading to improved sample efficiency. In the extreme case where the topology is a fully connected graph, our approach would recover the centralized topology of the benchmarked methods. The performance in terms of average episode reward would, therefore, be at least the same, since more information is provided to the joints; however, the sample efficiency would be hampered by computing 6 centralized policies, one per joint.

E. Real multi-robot experiments

Finally, we validate the sim-to-real transfer capabilities of pH-MARL using the Georgia Tech Robotarium [1]. The goal of this validation is twofold. First, in all previous sections we assumed perfect communication channels, whereas physical communication in real settings is subject to packet losses, disturbances in the messages and delays, leading to potentially adverse effects on the control policies. Second, we aim to evaluate how pH-MARL policies trained in an ideal simulated environment handle the transfer gap to real multi-robot settings subject to imperfect actuation in the differential-drive robot dynamics of the Robotarium robots. We pose a navigation problem with collision avoidance akin to the “Navigation” scenario used for ablations in Section VI-A. The only difference is that, to better represent the Robotarium arena, we train over a 3.2×2 m space and increase the physical radius of the agents to 11 cm to fit the size of the Robotarium robots. We train a pH-MARL navigation policy with $n = 4$ as in Section VI-A in VMAS.

After that, we integrate the trained policy in the Robotarium simulator, designed to be as similar as possible to the Robotarium arena to guarantee that policies submitted to the real Robotarium platform comply with all requirements. Two settings are evaluated. The first one evaluates the pH-MARL policy with perfect communications with $r_{comm} = 0.75$ m. The second evaluates the pH-MARL policy with realistic communication, as follows.

- **Packet losses:** An existing communication link has probability 0.1 of being removed from the graph.
- **Communication disturbances:** Messages sent through a communication channel have probability 0.1 of being distorted by an additive zero-mean Gaussian noise with variance 0.05 I.




Fig. 11. Quantitative results of the pH-MARL policy for multi-robot navigation in the Robotarium arena under imperfect communication. From left to right: travel distance, travel time and success rate. All the metrics are the average over 3 random runs.

- **Delays:** Messages sent through a communication channel have probability 0.1 of suffer a delay uniformly distributed in the range $[1, 10]$ sampling times (sampling time is 0.033s).

Fig. 9 compares the travel distance, travel time and success rate for the pH-MARL policy under perfect and imperfect communication for different number of robots, averaged over 10 random runs each. The pH-MARL policy is able to scale to teams up to $\times 4$ larger than the one used for training, achieving constant travel distances and almost constant travel times with respect to n . The success rates are always above 95%, proving that the pH-MARL policy consistently transfers to different robot dynamics (training is done with a damped second-order integrator, whereas the Robotarium robots follow first-order differential-drive dynamics) and environment. Importantly, the pH-MARL policy is robust against imperfect communication, achieving the same results obtained with perfect communication. Qualitatively speaking, Fig. 10 shows that the robots achieve smooth trajectories and they reach their goals avoiding collisions.

Once we verified that the pH-MARL policy is scalable and robust against communication disturbances in the Robotarium simulator, we evaluated it in the real Robotarium arena under imperfect communication. All the parameters are the same, and every configuration of number of robots is averaged over 3 random runs. As show in Fig. 11, the travel distance and travel time are consistent across number of robots, even though the space in the arena becomes increasingly tight with fleets of 12 and 16 robots. The case of $n = 4$ is much easier than the others, and consequently the travel distance and time are smaller than in the other configurations. The success rate is always 100%, allowing us to conclude that our pH-MARL policy supports zero-shot sim-to-real transfer. This is reinforced by the qualitative results in Fig. 12, where robots follow smooth paths and avoid collisions akin to what the pH-MARL policy achieves in the Robotarium simulator.

VII. CONCLUSIONS

We proposed a novel reinforcement learning formulation where the single agent is the multi-robot graph. On the one hand, this allowed to explicitly consider the potential networked interactions among agents, going beyond the classical assumptions of independent execution found in other multi-agent reinforcement learning approaches. On the other hand, this allowed to avoid non-stationarity issues found during training in other multi-robot learning setting. In particular,

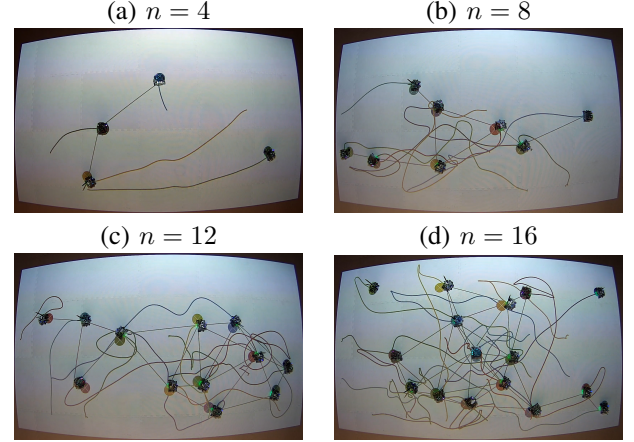


Fig. 12. Qualitative results of the pH-MARL policy for multi-robot navigation in the Robotarium arena under imperfect communication. After training a navigation policy with $n = 4$ robots in the VMAS simulator, we directly deployed it in teams of $n = \{4, 8, 12, 16\}$ robots. The figures display coloured trajectories and circles for each robot and associated goal. The robots achieve safe navigation with collision avoidance even in tight spaces with teams $4\times$ larger than those seen during training despite the imperfect communications.

we designed a soft actor-critic algorithm to manage the networked and stochastic nature of distributed multi-robot policies, learning simultaneously the homogeneous distributed control policy for all the robots while respecting the policy k -hop factorization and the correlation among robots. Besides, the method exploits the collective knowledge of a centralized critic during training. To achieve the learning of scalable and distributed control policies by design, we proposed an inter-connection and damping passivity-based control policy based on a port-Hamiltonian description of the multi-robot dynamics that preserves energy conservation laws and individual robot dynamics. To parameterized the controller, we proposed a set of self-attention-based neural networks that respects the desired distributed structure of the control policy and handles the time-varying available information at each robot.

We conducted ablation studies and simulations in comparison to state-of-the-art multi-agent reinforcement learning approaches in seven scenarios, covering a wide variety of cooperative and competitive behaviors such as collision avoidance, navigation, transport, evasion and monitoring, and including a validation experiment in a realistic robotic platform. In all the cases, our proposed approach exhibits superior performance in terms of cumulative reward per robot and scalability. Our approach, without further training, scales and achieves the same performance of other methods that are trained, ad hoc,

with a specific number of robots. Extensive validation in a real multi-robot navigation setting demonstrated that policies trained in an ideal simulator exhibit great zero-shot sim-to-real transfer capabilities, even under imperfect communication. Nevertheless, there is still room for improvement in terms of generalization with respect to environmental conditions such as stage size. These results lead to the following conclusion: the combination of physics-informed neural networks and reinforcement learning techniques is a promising research line to address multi-robot problems.

A important future direction is to utilize the physics-informed port-Hamiltonian parameterization of the control policies to provide performance or safety guarantees. For instance, passivity theory or control barrier functions can be employed to ensure convergence properties of the control policy or design safety modules on top of the control policy that provide collision avoidance guarantees. It also remains to be explored how to bypass the three-message passing scheme to compute $\partial H_{\theta}(\mathbf{x})/\partial \mathbf{x}_j$ and rely only on single-message broadcasting protocols or perception cues since it is currently a computation and implementation bottleneck.

APPENDIX A NETWORK PARAMETERS

The self-attention-based control policy in ph-MARL is parameterized as follows:

- $[\mathbf{R}_{\theta}]_{ij}$: $W = 3$, $h_w = [n_x, 16, 8]$, $r_w = [8, 16, 8]$, $d_w = [16, 8, 16]$; functions $\beta = \text{sigmoid}$, $\chi = \psi = \text{swish}$ [88].
- $[\mathbf{J}_{\theta}]_{ij}$: $W = 3$, $h_w = [n_x, 16, 8]$, $r_w = [8, 16, 8]$, $d_w = [16, 8, 16]$; functions $\beta = \text{sigmoid}$, $\chi = \psi = \text{swish}$ [88].
- H_{θ}^i : $W = 3$ layers, $h_w = [n_x, 16, 8]$, $r_w = [16, 8, 8]$, $d_w = [16, 8, 25]$; functions $\beta = \text{sigmoid}$, $\chi = \psi = \text{swish}$ [88].

For the variance network we use the same architecture of $[\mathbf{R}_{\theta}]_{ij}$ but with $h_1 = n_x + n_u$ and $d_W = 2$. The other networks used in the ablation studies are as follows:

- MLP: the policy is a single multi-layer perceptron of size $[n \times n_x, n \times n_u, n \times n_u]$ with swish hidden activation function and linear output. For the variance network we use the same architecture but with size $[n(n_x + n_u), n \times n_u, n \times n_u]$.
- MSA: the policy is a single multi-layer perceptron of size $[n \times n_x, n \times n_x]$ with linear output followed by a self-attention layer where query, key and values are directly the feature vector from the multi-layer perceptron, and an additional multi-layer perceptron of size $[n \times n_x, n \times n_u]$ with linear output. For the variance network we use the same architecture but with size $[n(n_x + n_u), n \times n_u]$ for the first multi-layer perceptron.
- GSA: the policy is the same one used to predict $[\mathbf{R}_{\theta}]_{ij}$ but with $d_W = 2$. For the variance network we use the same architecture but $h_1 = n_x + n_u$.

The $Q_{\Pi_{\theta}}(\mathbf{s}_t, \mathbf{a}_t)$ function is always parameterized as a multi-layer perceptron with layers of size $[n(n_x + n_u), 2n(n_x + n_u), n(n_x + n_u), (n_x + n_u), 1]$ with swish [88] activation functions except the last layer, that is linear.

The food collection, grassland, and adversarial scenarios use a neural network to pre-process the observation vector to move from a time-varying observation size to a fixed state size compatible with the control policies. To do so, the state is a concatenation of the position (2-dimensional vector), velocity (2-dimensional vector), aliveness (boolean quantity), closest goal relative distance (2-dimensional vector) and a 2-dimensional feature vector provided by a neural network. The neural network is composed by a multi-layer perceptron with no hidden layers of size $[n_x - 2, 2]$ and swish activation function, a self-attention layer of size $h_w = r_w = d_w = 2$, and another multi-layer perceptron with no hidden layers of size $[2, 2]$ and linear output.

APPENDIX B SOFT ACTOR-CRITIC HYPERPARAMETERS

The following table details the parameterization of the soft actor-critic algorithm for the different scenarios.

TABLE III: Soft actor-critic hyperparameters.

Parameter	Scenario	Value
optimizer	all	Adam [89]
r_{comm}	reverse transport, navigation	0.45m
	sampling	0.75m
	food collection, grassland, adversarial	0.15m
n training	reverse transport, sampling, navigation	4
	food collection, grassland, adversarial	8
# parallel environments	all	96
shared r	all	True
maximum steps per episode	reverse transport, navigation, food collection, grassland, adversarial	400
	sampling	1000
replay buffer size	all	2×10^6
initial random steps	all	10^3
γ	all	0.99
α_0	all	5
α_{min}	all	0.1
α_{max}	all	10
ρ	all	0.005
learning rate α	all	10^{-5}
learning rate	all	10^{-4}

Continued on next page

TABLE III: Soft actor-critic hyperparameters. (Continued)

batch size	all	1024
# training steps	navigation, reverse transport, sampling, food collection, grassland	2×10^6
	adversarial	6×10^5
clip gradients	all	False
reward scaling	all	False
σ_{\min}	all	e^{-5}
σ_{\max}	all	e^2
landmark mass	sampling, navigation, food collection, grassland, adversarial	default
	reverse transport	1
evaluation interval	all	10^4
# evaluation episodes per interval	all	10

REFERENCES

- [1] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "The Robotarium: A remotely accessible swarm robotics research testbed," in *IEEE International Conference on Robotics and Automation*, 2017, pp. 1699–1706.
- [2] B. Peng, T. Rashid, C. Schroeder de Witt, P.-A. Kamienny, P. Torr, W. Böhmer, and S. Whiteson, "FACMAC: Factored multi-agent centralised policy gradients," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12208–12221, 2021.
- [3] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, "Decentralized active information acquisition: Theory and application to multi-robot SLAM," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 4775–4782.
- [4] Y. Tian, Y. Chang, F. H. Arias, C. Nieto-Granda, J. P. How, and L. Carlone, "Kimera-multi: Robust, distributed, dense metric-semantic SLAM for multi-robot systems," *IEEE Transactions on Robotics*, 2022.
- [5] X. Kan, T. C. Thayer, S. Carpin, and K. Karydis, "Task planning on stochastic aisle graphs for precision agriculture," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3287–3294, 2021.
- [6] A. Pierson and M. Schwager, "Bio-inspired non-cooperative multi-robot herding," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 1843–1849.
- [7] E. Sebastián and E. Montijano, "Multi-robot implicit control of herds," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 1601–1607.
- [8] E. Sebastián, E. Montijano, and C. Sagüés, "Adaptive multirobot implicit control of heterogeneous herds," *IEEE Transactions on Robotics*, 2022.
- [9] L. Heintzman, A. Hashimoto, N. Abaid, and R. K. Williams, "Anticipatory planning and dynamic lost person models for human-robot search and rescue," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 8252–8258.
- [10] M. J. Matarić, "Reinforcement learning in the multi-robot domain," *Autonomous Robots*, vol. 1, no. 4, pp. 73–83, 1997.
- [11] L. Matignon, G. J. Laurent, and N. Le Fort-Piat, "Hysteretic q-learning: an algorithm for decentralized reinforcement learning in cooperative multi-agent teams," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 64–69.
- [12] L. Matignon, L. Jeanpierre, and A.-I. Mouaddib, "Coordinated multi-robot exploration under communication constraints using decentralized markov decision processes," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 26, no. 1, 2012, pp. 2017–2023.
- [13] S. Munikoti, D. Agarwal, L. Das, M. Halappanavar, and B. Natarajan, "Challenges and opportunities in deep reinforcement learning with graph neural networks: A comprehensive review of algorithms and applications," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [14] Á. Serra-Gómez, H. Zhu, B. Brito, W. Böhmer, and J. Alonso-Mora, "Learning scalable and efficient communication policies for multi-robot collision avoidance," *Autonomous Robots*, pp. 1–23, 2023.
- [15] Y. L. Lo, C. S. de Witt, S. Sokota, J. N. Foerster, and S. Whiteson, "Cheap talk discovery and utilization in multi-agent reinforcement learning," in *International Conference on Learning Representations*, 2023.
- [16] G. Qu, A. Wierman, and N. Li, "Scalable reinforcement learning for multiagent networked systems," *Operations Research*, vol. 70, no. 6, pp. 3601–3628, 2022.
- [17] T. Beckers, T. Z. Jiahao, and G. J. Pappas, "Learning switching port-hamiltonian systems with uncertainty quantification," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 525–532, 2023.
- [18] C. Neary and U. Topcu, "Compositional learning of dynamical system models using port-hamiltonian neural networks," in *Learning for Dynamics and Control Conference*. PMLR, 2023, pp. 679–691.
- [19] E. Sebastián, T. Duong, N. Atanasov, E. Montijano, and C. Sagüés, "LEMURS: Learning Distributed Multi-Robot Interactions," in *IEEE International Conference on Robotics and Automation*, 2023, pp. 7713–7719.
- [20] T. X. Nghiem, J. Dragoña, C. Jones, Z. Nagy, R. Schwan, B. Dey, A. Chakrabarty, S. Di Cairano, J. A. Paulson, A. Carron, et al., "Physics-informed machine learning for modeling and control of dynamical systems," in *IEEE American Control Conference*, 2023, pp. 3735–3750.
- [21] S. Sanyal and K. Roy, "Ramp-net: A robust adaptive mpc for quadrotors via physics-informed neural network," in *IEEE International Conference on Robotics and Automation*, 2023, pp. 1019–1025.
- [22] C. Rodwell and P. Tallapragada, "Physics-informed reinforcement learning for motion control of a fish-like swimming robot," *Scientific Reports*, vol. 13, no. 1, p. 10754, 2023.
- [23] S. Cuomo, V. S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli, "Scientific machine learning through physics-informed neural networks: Where we are and what's next," *Journal of Scientific Computing*, vol. 92, no. 3, p. 88, 2022.
- [24] Y. Xu, S. Kohtz, J. Boakye, P. Gardoni, and P. Wang, "Physics-informed machine learning for reliability and systems safety applications: State of the art and challenges," *Reliability Engineering & System Safety*, p. 108900, 2022.
- [25] D. Bloembergen, K. Tuyls, D. Hennes, and M. Kaisers, "Evolutionary dynamics of multi-agent learning: A survey," *Journal of Artificial Intelligence Research*, vol. 53, pp. 659–697, 2015.
- [26] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 6252–6259.
- [27] S. H. Semnani, H. Liu, M. Everett, A. De Ruiter, and J. P. How, "Multi-agent motion planning for dense and dynamic environments via deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3221–3226, 2020.
- [28] A. Y. Ng, S. Russell, et al., "Algorithms for inverse reinforcement learning," in *International Conference on Machine Learning*, vol. 1, 2000, p. 2.
- [29] S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, K. Schmeckpeper, S. Singh, S. Levine, and C. Finn, "Robonet: Large-scale multi-robot learning," in *Conference on Robot Learning*, 2020, pp. 885–897.
- [30] K. Bogert and P. Doshi, "Multi-robot inverse reinforcement learning under occlusion with estimation of state transitions," *Artificial Intelligence*, vol. 263, pp. 46–73, 2018.
- [31] R. Han, S. Chen, and Q. Hao, "Cooperative multi-robot navigation in dynamic environment with deep reinforcement learning," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 448–454.
- [32] I. Gharbi, J. Kuckling, D. G. Ramos, and M. Birattari, "Show me what you want: Inverse reinforcement learning to automatically design robot swarms by demonstration," in *IEEE International Conference on Robotics and Automation*, 2023, pp. 5063–5070.
- [33] H. Zhu, F. M. Claramunt, B. Brito, and J. Alonso-Mora, "Learning interaction-aware trajectory predictions for decentralized multi-robot motion planning in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2256–2263, 2021.
- [34] S. Zhou, M. J. Phielipp, J. A. Sefair, S. I. Walker, and H. B. Amor, "Clone swarms: Learning to predict and control multi-robot systems by

- imitation,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019, pp. 4092–4099.
- [35] G. Qu, A. Wierman, and N. Li, “Scalable reinforcement learning of localized policies for multi-agent networked systems,” in *Learning for Dynamics and Control*. PMLR, 2020, pp. 256–266.
- [36] V. Zambaldi, D. Raposo, A. Santoro, V. Bapst, Y. Li, I. Babuschkin, K. Tuyls, D. Reichert, T. Lillicrap, E. Lockhart, et al., “Deep reinforcement learning with relational inductive biases,” in *International Conference on Learning Representations*, 2018.
- [37] S. Iqbal and F. Sha, “Actor-attention-critic for multi-agent reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 2961–2970.
- [38] G. Li, B. Jiang, H. Zhu, Z. Che, and Y. Liu, “Generative attention networks for multi-agent behavioral modeling,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 7195–7202.
- [39] P. Parnika, R. B. Diddigi, S. K. R. Danda, and S. Bhatnagar, “Attention actor-critic algorithm for multi-agent constrained co-operative reinforcement learning,” in *International Conference on Autonomous Agents and Multiagent Systems*, 2021.
- [40] A. Marino, C. Pacchierotti, and P. R. Giordano, “On the stability of gated graph neural networks,” *arXiv preprint arXiv:2305.19235*, 2023.
- [41] Q. Li, W. Lin, Z. Liu, and A. Prorok, “Message-aware graph attention networks for large-scale multi-robot path planning,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5533–5540, 2021.
- [42] A. Khan, E. Tolstaya, A. Ribeiro, and V. Kumar, “Graph policy gradients for large scale robot control,” in *Conference on Robot Learning*, 2020, pp. 823–834.
- [43] E. Tolstaya, F. Gama, J. Paulos, G. Pappas, V. Kumar, and A. Ribeiro, “Learning decentralized controllers for robot swarms with graph neural networks,” in *Conference on Robot Learning*, 2020, pp. 671–682.
- [44] E. Tolstaya, J. Paulos, V. Kumar, and A. Ribeiro, “Multi-robot coverage and exploration using spatial graph neural networks,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 8944–8950.
- [45] F. Yang and N. Matni, “Communication topology co-design in graph recurrent neural network based distributed control,” in *IEEE Conference on Decision and Control*, 2021, pp. 3619–3626.
- [46] F. Gama, Q. Li, E. Tolstaya, A. Prorok, and A. Ribeiro, “Synthesizing decentralized controllers with graph neural networks and imitation learning,” *IEEE Transactions on Signal Processing*, vol. 70, pp. 1932–1946, 2022.
- [47] L. Kuyer, S. Whiteson, B. Bakker, and N. Vlassis, “Multiagent reinforcement learning for urban traffic control using coordination graphs,” in *Machine Learning and Knowledge Discovery in Databases: European Conference*. Springer, 2008, pp. 656–671.
- [48] L. Buşoni, R. Babuška, and B. De Schutter, “Multi-agent reinforcement learning: An overview,” *Innovations in multi-agent systems and applications-I*, pp. 183–221, 2010.
- [49] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser, et al., “Starcraft ii: A new challenge for reinforcement learning,” *arXiv preprint arXiv:1708.04782*, 2017.
- [50] B. Ellis, J. Cook, S. Moalla, M. Samvelyan, M. Sun, A. Mahajan, J. Foerster, and S. Whiteson, “Smacv2: An improved benchmark for cooperative multi-agent reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [51] S. Gronauer and K. Diepold, “Multi-agent deep reinforcement learning: a survey,” *Artificial Intelligence Review*, pp. 1–49, 2022.
- [52] A. Oroojlooy and D. Hajinezhad, “A review of cooperative multi-agent deep reinforcement learning,” *Applied Intelligence*, vol. 53, no. 11, pp. 13 677–13 722, 2023.
- [53] L. Matignon, G. J. Laurent, and N. Le Fort-Piat, “Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems,” *The Knowledge Engineering Review*, vol. 27, no. 1, pp. 1–31, 2012.
- [54] G. Papoudakis, F. Christianos, A. Rahman, and S. V. Albrecht, “Dealing with non-stationarity in multi-agent deep reinforcement learning,” *arXiv preprint arXiv:1906.04737*, 2019.
- [55] W. Böhmer, V. Kurin, and S. Whiteson, “Deep coordination graphs,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 980–991.
- [56] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al., “Soft actor-critic algorithms and applications,” *arXiv preprint arXiv:1812.05905*, 2018.
- [57] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *arXiv preprint arXiv:1509.02971*, 2015.
- [58] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [59] R. Lowe, Y. I. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, “Multi-agent actor-critic for mixed cooperative-competitive environments,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [60] C. Yu, A. Velu, E. Vinitsky, J. Gao, Y. Wang, A. Bayen, and Y. Wu, “The surprising effectiveness of ppo in cooperative multi-agent games,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 611–24 624, 2022.
- [61] M. Bettini, A. Prorok, and V. Moens, “BenchMARL: Benchmarking multi-agent reinforcement learning,” *Journal of Machine Learning Research*, vol. 25, no. 217, pp. 1–10, 2024.
- [62] J. G. Kuba, R. Chen, M. Wen, Y. Wen, F. Sun, J. Wang, and Y. Yang, “Trust region policy optimisation in multi-agent reinforcement learning,” in *International Conference on Learning Representations*, 2021.
- [63] J. Bloom, P. Paliwal, A. Mukherjee, and C. Pinciroli, “Decentralized multi-agent reinforcement learning with global state prediction,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023, pp. 8854–8861.
- [64] Y. Yang, R. Luo, M. Li, M. Zhou, W. Zhang, and J. Wang, “Mean field multi-agent reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 5571–5580.
- [65] B. Wang, J. Xie, and N. Atanasov, “DARLIN: Distributed multi-agent reinforcement learning with one-hop neighbors,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022.
- [66] C. S. De Witt, T. Gupta, D. Makoviichuk, V. Makoviychuk, P. H. Torr, M. Sun, and S. Whiteson, “Is independent learning all you need in the starcraft multi-agent challenge?” *arXiv preprint arXiv:2011.09533*, 2020.
- [67] Y. Motokawa and T. Sugawara, “Interpretability for conditional coordinated behavior in multi-agent reinforcement learning,” in *IEEE International Joint Conference on Neural Networks*, 2023, pp. 1–8.
- [68] R. Kortvelesy and A. Prorok, “QGNN: Value Function Factorisation with Graph Neural Networks,” *arXiv preprint arXiv:2205.13005*, 2022.
- [69] Y. Hu, J. Fu, and G. Wen, “Graph soft actor-critic reinforcement learning for large-scale distributed multirobot coordination,” *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [70] Z. Huang, Z. Yang, R. Krupani, B. Şenbaşlar, S. Batra, and G. S. Sukhatme, “Collision avoidance and navigation for a quadrotor swarm using end-to-end deep reinforcement learning,” in *IEEE International Conference on Robotics and Automation*, 2024, pp. 300–306.
- [71] P. Zhao and Y. Liu, “Physics informed deep reinforcement learning for aircraft conflict resolution,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 8288–8301, 2021.
- [72] G. Sartoretti, Y. Wu, W. Paivine, T. S. Kumar, S. Koenig, and H. Choset, “Distributed reinforcement learning for multi-robot decentralized collective construction,” in *Distributed Autonomous Robotic Systems: The 14th International Symposium*. Springer, 2019, pp. 35–49.
- [73] A. Van Der Schaft and D. Jeltsema, “Port-Hamiltonian systems theory: An introductory overview,” *Foundations and Trends in Systems and Control*, vol. 1, no. 2-3, pp. 173–378, 2014.
- [74] L. Furieri, C. L. Galimberti, M. Zakwan, and G. Ferrari-Trecate, “Distributed neural network control with dependability guarantees: a compositional port-hamiltonian approach,” in *Learning for Dynamics and Control Conference*, 2022, pp. 571–583.
- [75] C. L. Galimberti, L. Furieri, L. Xu, and G. Ferrari-Trecate, “Hamiltonian deep neural networks guaranteeing nonvanishing gradients by design,” *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 3155–3162, 2023.
- [76] G. Shi, W. Hönig, Y. Yue, and S.-J. Chung, “Neural-swarm: Decentralized close-proximity multirobot control using learned interactions,” in *IEEE International Conference on Robotics and Automation*, 2020, pp. 3241–3247.
- [77] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [78] A. J. van der Schaft, “Port-Hamiltonian systems: network modeling and control of nonlinear physical systems,” in *Advanced dynamics and control of structures and machines*. Springer, 2004, pp. 127–167.
- [79] G. Blankenstein, R. Ortega, and A. J. Van Der Schaft, “The matching conditions of controlled lagrangians and ida-passivity based control,” *International Journal of Control*, vol. 75, no. 9, pp. 645–665, 2002.

- [80] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1861–1870.
- [81] M. Bettini, R. Kortvelesy, J. Blumenkamp, and A. Prorok, "VMAS: a vectorized multi-agent simulator for collective robot learning," in *International Symposium on Distributed Autonomous Robotic Systems*. Springer, 2022, pp. 42–56.
- [82] T. Haarnoja, S. Ha, A. Zhou, J. Tan, G. Tucker, and S. Levine, "Learning to walk via deep reinforcement learning," *arXiv preprint arXiv:1812.11103*, 2018.
- [83] I. Mordatch and P. Abbeel, "Emergence of grounded compositional language in multi-agent populations," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
- [84] Q. Long, Z. Zhou, A. Gupta, F. Fang, Y. Wu, and X. Wang, "Evolutionary population curriculum for scaling multi-agent reinforcement learning," in *International Conference on Learning Representations*, 2019.
- [85] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic differentiation in machine learning: a survey," *Journal of Machine Learning Research*, vol. 18, no. 153, pp. 1–43, 2018.
- [86] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 5026–5033.
- [87] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International conference on machine learning*. PMLR, 2015, pp. 1889–1897.

- [88] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," *arXiv preprint arXiv:1710.05941*, 2017.
- [89] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference for Learning Representations*, 2015.



Nikolay Atanasov (S'07-M'16-SM'23) is an Associate Professor of Electrical and Computer Engineering at the University of California San Diego, La Jolla, CA, USA. He obtained a B.S. degree in Electrical Engineering from Trinity College, Hartford, CT, USA in 2008 and M.S. and Ph.D. degrees in Electrical and Systems Engineering from the University of Pennsylvania, Philadelphia, PA, USA in 2012 and 2015, respectively. Dr. Atanasov's research focuses on robotics, control theory, and machine learning with emphasis on active perception problems for autonomous mobile robots. He works on probabilistic models for simultaneous localization and mapping (SLAM) and on optimal control and reinforcement learning algorithms for minimizing probabilistic model uncertainty. Dr. Atanasov's work has been recognized by the Joseph and Rosaline Wolf award for the best Ph.D. dissertation in Electrical and Systems Engineering at the University of Pennsylvania in 2015, the Best Conference Paper Award at the IEEE International Conference on Robotics and Automation (ICRA) in 2017, the NSF CAREER Award in 2021, and the IEEE RAS Early Academic Career Award in Robotics and Automation in 2023.




Eduardo Sebastián (Graduate Student Member, IEEE) received the B.Eng. (Hons.) degree in electronic and automatic engineering, the M.Eng. (Hons.) degree in electronics, and the Ph.D. (cum laude) degree in computer science and systems engineering in 2019, 2020 and 2025, respectively, from the Universidad de Zaragoza, Zaragoza, Spain. He currently holds a postdoctoral position at the University of Cambridge, UK. He has been a visiting scholar at the University of California San Diego, USA, in 2022 and 2024. His research interests

include nonlinear control and distributed multirobot systems. Eduardo is a Fulbright Scholar and a DAAD Alnet Fellow.



Eduardo Montijano (Member, IEEE) received the M.Sc. degree in computer science and the Ph.D. degree in robotics and systems engineering from the Universidad de Zaragoza, Zaragoza, Spain, in 2008 and 2012, respectively. From 2012 to 2016, he was a faculty member with Centro Universitario de la Defensa, Zaragoza.

He is currently an Associate Professor with the Departamento de Informática e Ingeniería de Sistemas, Universidad de Zaragoza. His current research focuses on the design of scene understanding algorithms for planning and control of multiple robots.



Thai Duong (S'18-M'24) is a postdoctoral researcher in Computer Science at Rice University, Houston, TX, USA. He obtained a Ph.D. degree in Electrical and Computer Engineering at the University of California San Diego, La Jolla, CA, USA in 2024, an M.S. degree in Electrical and Computer Engineering from Oregon State University, Corvallis, OR, USA, in 2013, and a B.S. degree in Electronics and Telecommunications from Hanoi University of Science and Technology, Hanoi, Vietnam in 2011. His research interests include robotics,

machine learning, control theory and optimization. His work focuses on efficient robot learning for task planning, motion planning and control of different robot platforms such as mobile and legged robots, and manipulators.



Carlos Sagiés (Senior Member, IEEE) received the M.Sc. degree in computer science and systems engineering and the Ph.D. degree in industrial engineering from the University of Zaragoza, Zaragoza, Spain, in 1989 and 1992, respectively.

In 1994, he joined as an Associate Professor with the Departamento de Informática e Ingeniería de Sistemas, University of Zaragoza, where he became a Full Professor in 2009 and also the Head Teacher. He was engaged in research on force and infrared sensors for robots. His current research interests include control systems and industry applications, computer vision, visual control, and multi-vehicle cooperative control.