

# Making Friends in the Dark: Ad Hoc Teamwork Under Partial Observability

João G. Ribeiro<sup>a,\*</sup>, Cassandro Martinho<sup>a</sup>, Alberto Sardinha<sup>a, b</sup> and Francisco S. Melo<sup>a</sup>

<sup>a</sup>INESC-ID & Instituto Superior Técnico, University of Lisbon

<sup>b</sup>Pontifícia Universidade Católica do Rio de Janeiro

**Abstract.** This paper introduces a formal definition of the setting of ad hoc teamwork under partial observability and proposes a first-principled model-based approach which relies only on prior knowledge and partial observations of the environment in order to perform ad hoc teamwork. We make three distinct assumptions that set it apart from previous works, namely: i) the state of the environment is always partially observable, ii) the actions of the teammates are always unavailable to the ad hoc agent and iii) the ad hoc agent has no access to a reward signal which could be used to learn the task from scratch. Our results in 70 POMDPs from 11 domains show that our approach is not only effective in assisting unknown teammates in solving unknown tasks but is also robust in scaling to more challenging problems. Supplementary material is available at <https://github.com/jmribeiro/adhoc-teamwork-under-partial-observability>

## 1 Introduction

One of many aspects of human intelligence is the ability to cooperate with strangers when unexpected events, such as emergencies, occur. Introduced by Stone et al. [26], the setting of ad hoc teamwork, studies how autonomous agents can efficiently handle this same scenario. Instead of being evaluated according to their ability to learn and/or perform a given task, agents are evaluated according to how efficient they are able to assist unknown teammates in performing unknown tasks (without being able to pre-coordinate or communicate beforehand).

Even though the setting of ad hoc teamwork has recently been under great attention by the multi-agent systems community [20], current approaches [3] rely on assumptions unfeasible to attain in real-world scenarios, such as full observability of the state of the environment [3, 23] and being able to observe the actions of the teammates [3]. In a more recent line of work, Gu et al. [14] showcase that by learning via trial-and-error how to perform tasks on a fully observable environment (where both the states and actions of the teammates are fully observable), an ad hoc agent is then able to transfer this knowledge to a partial observable environment where only a single observation is available.

In this paper, we follow both Ribeiro et al. [23] and Gu et al. [14], and formalize the problem of ad hoc teamwork under partial observability, relying on the novel assumptions that an agent never has access to the full state of the environment nor the actions of its teammates. We propose a novel, first-principled Bayesian online inference algorithm for the setting of ad hoc teamwork, named ATPO

(Ad hoc Teamwork for Partially Observable environments), which relies on a library of possible models (in our case, Partially Observable Markov Decision Processes (POMDPs)). To the best of our knowledge, ATPO represents the first model-based approach for ad hoc teamwork capable of assisting unknown teams in completing unknown tasks, relying only on partial observations of the environment. These assumptions also make ATPO suited to address tasks involving *ad hoc robotic agents*, accommodating the natural perceptual limitations of robotic platforms [12] and allowing for the interaction, not only with other robots but also humans as teammates.

We evaluate ATPO by conducting an empirical evaluation, resorting to a total of 70 POMDPs from 11 benchmark multiagent problems. Our results demonstrate the effectiveness of ATPO in performing ad hoc teamwork under partial observability by not only efficiently assisting unknown teammates, but also identifying, from partial observations, which teammates it is interacting with and which tasks they are currently performing.

To summarize, our contributions are threefold:

- We contribute to the formalization of ad hoc teamwork under partial observability;
- We propose ATPO, a novel approach for ad hoc teamwork under partial observability that is capable of inferring the underlying target task and teammate's behaviour from the agent's history of observations;
- We illustrate the applicability of our approach in a set of eleven distinct domains.

## 2 Related Work

The setting of ad hoc teamwork was first introduced by Stone et al. [26]. In ad hoc teamwork, an agent (named the *ad hoc agent*) has the goal of joining an already formed team and assisting it on-the-fly. The ad hoc agent is not assumed to be able to pre-coordinate or pre-communicate with the team.

Earlier work in the setting of ad hoc teamwork assumed that the ad hoc agent knew, beforehand, the teams' behavior and the task they were performing. Given this information, the only remaining goal for the ad hoc agent is to understand its role and act accordingly. Stone and Kraus [27] showcase an example of this scenario, introducing a finite-horizon decision problem where an ad hoc agent uses a planner in order to compute its own actions, assisting a single teammate in performing a single task. The authors are then followed by Barrett and Stone [4], who extend the approach to discounted infinite horizon problems. Later on, Agmon and Stone [1] approach the same

\* Corresponding Author. Email: joao.g.ribeiro@tecnico.ulisboa.pt

scenario, but now considering  $N$  teammates instead of one. All these approaches, however, by assuming the ad hoc agent knows the teammates' behavior and the task they are performing beforehand, reduce ad hoc teamwork to a planning problem (where the only challenge is to figure out the optimal role for the ad hoc agent).

It is later on that Chakraborty and Stone [8] drop the assumption that the teammates' behavior is known a priori. In their work, the authors assume that the teammates follow an unknown Markov policy that the ad hoc agent must identify to plan its actions. The authors, however, still assume that the task being performed is known to the ad hoc agent. From this perspective, ad hoc teamwork is closely related to *learning in games*, a problem for which an extensive literature already exists [11].

More recently, Melo and Sardinha [18] formalize the setting of ad hoc teamwork by breaking it down into three sub-problems: (i) task identification (the ad hoc agent has to identify the task the teammates are performing), (ii) teammate identification (the ad hoc agent has to identify the behavior of its teammates) and (iii) planning/execution (after identifying the task and teammate behavior, the ad hoc agent has to plan its best response accordingly). The authors are then followed by Barrett et al. [3], who follow this formalization and introduce the PLASTIC framework [3, 5]. The PLASTIC framework is an approach which, relying on a set of previous interactions, is able to identify the current team (and task being performed) by performing Bayesian updates. This is done by first assigning each possible model a belief prior, and, using as evidence the observed actions of the teammates for the current state, updating the beliefs in each timestep until. Using the most likely model, a pre-trained policy for the ad hoc agent can then be used to compute the best response.

Expanding upon Barrett et al. [3], Ribeiro et al. [23] later on tackle these limitations by dropping the assumption that the actions of the teammates are visible to the ad hoc agent. In their work, the authors also rely on a set of prior interactions, however this time modeled as Markov Decision Processes. Gu et al. [14] afterwards introduce an approach relying on reinforcement learning techniques to allow an agent to first learn how to perform tasks from scratch in a fully observable environment (in which the agent is also able to observe the actions of its teammates) and is then deployed into a partial observable scenario, where it has to assist unknown teammates. In this paper, we expand upon Ribeiro et al. [23] and Gu et al. [14] by relying on the assumptions that the state of the environment is never observable and the actions of the teammates never visible to the ad hoc agent. Furthermore, we consider a setting of closed ad hoc teamwork, where teams are fixed and teammates do not dynamically join and leave. We formalize ad hoc teamwork under partial observability and propose a novel algorithm (ATPO) which assumes the environment is partially observable and the ad hoc agent has no access to the actions of its teammates. Similar to Ribeiro et al. [23], ATPO relies on a library of prior interactions, this time modeled as Partially Observable Markov Decision Processes (POMDPs). The beliefs over each possible model are then updated using Bayesian updates, using as evidence the partial observations. ATPO then uses the beliefs to compute the best response for the ad hoc agent, by relying on the solutions for the POMDPs, obtained using an arbitrary POMDP solver (in our case, Perseus [25]).

## 2.1 Similar Lines of Research

Four similar lines of work, parallel to ad hoc teamwork, are the assistance framework [9], the problem of zero-shot coordination [15], the IPOMDP framework [13] and the problem of few-shot coordina-

tion [10]. The *assistance* framework [9] models a scenario where an agent must help a teammate in solving a given sequential task under uncertainty. The goal, however, is for an agent to assist unknown teammates in performing known tasks. Even though the authors do not refer to the problem as ad hoc teamwork, one could argue that the identification of unknown teammates with known tasks falls under the scope of ad hoc teamwork. In their work, however, Fern et al. [9] also consider that the teammate's actions to be accessible to the assistant. The problem of *zero-shot coordination* [15], unlike ad hoc teamwork, studies how independently trained agents may interact with one another on first-attempt [6, 17, 28]. This can be seen as a flavor of ad hoc teamwork where there is only one teammate which the ad hoc agent knows to have computed an optimal solution to the task being performed. It also assumes the ad hoc agent knows the task being performed, reducing it only to a problem of *teammate identification*. The IPOMDP framework [13], considers how an agent can augment the state space of a POMDP taking into account all possible, unknown, teammates. Not only does this approach grow exponentially with the number of possible teammates, but, similar to the framework of assistance, it also assumes that all teammates perform the same task (modeled by a single reward function). In their work, however, Gmytrasiewicz and Doshi [13] do not assume that the teammates' actions are visible to the agent, therefore sharing a common assumption with our work. Unlike Gmytrasiewicz and Doshi [13], our approach grows linearly with the number of possible teammates and tasks (instead of exponentially, like with IPOMDPs). Finally, the problem of few-shot teamwork (FST) considers a scenario where teams of agents are tasked with adapting to one another, similar to ad hoc teamwork. FST, however, unlike ad hoc teamwork, allows for an adaptation phase where teams can interact with one another before the final evaluation takes place, while ad hoc teamwork, requires agents to adapt on-the-fly without any pre-coordination or pre-communication.

## 3 Background

A *Markov decision problem* [22], or MDP, is denoted as a tuple  $(\mathcal{X}, \mathcal{A}, \{\mathbf{P}_a, a \in \mathcal{A}\}, r, \gamma)$ , where  $\mathcal{X}$  is the state space,  $\mathcal{A}$  is the action space,  $\mathbf{P}_a$  is a transition probability matrix, where  $\mathbf{P}_a(x' | x)$  is the probability of moving from state  $x$  to  $x'$  given action  $a \in \mathcal{A}$ ,  $r$  is the expected reward function, and  $\gamma \in [0, 1]$  is a discount factor.

A *policy*  $\pi$  maps states to distributions over actions. We write  $\pi(a | x)$  to denote the probability of selecting action  $a$  in state  $x$  according to policy  $\pi$ . Solving an MDP consists of determining a policy  $\pi$  to maximize the value

$$v^\pi(x) \triangleq \mathbb{E}_{A_t \sim \pi(X_t)} \left[ \sum_{t=0}^{\infty} \gamma^t R_t \mid X_0 = x \right], \quad (1)$$

for any initial state  $x \in \mathcal{X}$ . In the above expression,  $X_t$ ,  $A_t$  and  $R_t$  denote the (random) state, action and reward at time step  $t$ . The function  $v^\pi : \mathcal{X} \rightarrow \mathbb{R}$  is called a *value function*, and a policy  $\pi^*$  is *optimal* if, given any policy  $\pi$ ,  $v^{\pi^*}(x) \geq v^\pi(x)$ , for all  $x \in \mathcal{X}$ . The value function associated with an optimal policy is denoted as  $v^*$  and can be computed using, for example, dynamic programming. An optimal policy,  $\pi^*$ , is such that  $\pi^*(a | x) > 0$  only if  $a \in \operatorname{argmax} q^*(x, \cdot)$ , where

$$q^*(x, a) = r(x, a) + \gamma \sum_{x' \in \mathcal{X}} \mathbf{P}_a(x' | x) v^*(x').$$

A multiagent MDP (MMDP) is an extension of MDPs to multiagent settings and can be described as a tuple

$$\mathcal{M} = (N, \mathcal{X}, \{\mathcal{A}^n, n = 1, \dots, N\}, \{\mathbf{P}_a, a \in \mathcal{A}\}, r, \gamma),$$

where  $N$  is the number of agents,  $\mathcal{A}^n$  is the individual action space for agent  $n$ , and  $\mathbf{P}_a$  is the transition probability matrix associated with joint action  $a$ . We write  $\mathcal{A}$  to denote the set of all joint actions, corresponding to the Cartesian product of all individual action spaces  $\mathcal{A}^n$ . We also denote an element of  $\mathcal{A}^n$  as  $a^n$  and an element of  $\mathcal{A}$  as a tuple  $a = (a^1, \dots, a^N)$ , with  $a^n \in \mathcal{A}^n$ . We write  $a^{-n}$  to denote a reduced joint action, i.e., a tuple  $a^{-n} = (a^1, \dots, a^{n-1}, a^{n+1}, \dots, a^N)$ , and thus  $\mathcal{A}^{-n}$  is the set of all reduced joint actions. We adopt, for policies, a similar notation. Specifically, we write  $\pi^n$  to denote an individual policy for agent  $n$ ,  $\pi = (\pi^1, \dots, \pi^N)$  to denote a joint policy, and  $\pi^{-n}$  to denote a reduced joint policy.

The common goal of the agents in an MMDP is to select a joint policy,  $\pi^*$ , such that  $v^{\pi^*}(x) \geq v^\pi(x)$ , where, as before,

$$v^\pi(x) = \mathbb{E}_{\mathcal{A}_t \sim \pi(X_t)} \left[ \sum_{t=0}^{\infty} \gamma^t R_t \mid X_0 = x \right]. \quad (2)$$

In other words, an MMDP is just an MDP in which the action selection process is distributed across  $N$  agents, and can be solved by computing  $\pi^*$  from  $v^*$  as standard MDPs.

Finally, a *partially observable MDP*, or POMDP, is an extension of MDPs to partially observable settings. A POMDP can be described as  $(\mathcal{X}, \mathcal{A}, \mathcal{Z}, \{\mathbf{P}_a, a \in \mathcal{A}\}, \{\mathbf{O}_a, a \in \mathcal{A}\}, r, \gamma)$ , where  $\mathcal{X}$ ,  $\mathcal{A}$ ,  $\{\mathbf{P}_a, a \in \mathcal{A}\}$ ,  $r$ , and  $\gamma$ , are the same as in MDPs,  $\mathcal{Z}$  is the *observation space*, and  $\mathbf{O}_a$  is the observation probability matrix, where  $\mathbf{O}_a(z \mid x) = \mathbb{P}[Z_{t+1} = z \mid X_{t+1} = x, A_t = a]$ . The *belief* at time step  $t$  is a distribution  $b_t$  such that

$$b_t(x) \triangleq \mathbb{P}[X_t = x \mid X_0 \sim b_0, A_0 = a_0, Z_1 = z_1, \dots, Z_t = z_t],$$

where  $b_0$  is the initial state distribution. Given the action  $a_t$  and the observation  $z_{t+1}$ , we can update the belief  $b_t$  to incorporate the new information yielding

$$b_{t+1}(x') = \text{Bel}(b_t, a_t, z_{t+1}) \triangleq \frac{1}{\rho_{t+1}} \sum_{x \in \mathcal{X}} b_t(x) \mathbf{P}(x' \mid x, a_t) \mathbf{O}(z_{t+1} \mid x', a_t), \quad (3)$$

where  $\rho_{t+1}$  is a normalization factor. Every finite POMDP admits an equivalent *belief-MDP* with  $b_t$  being the state of this new MDP at time step  $t$ . A policy in a POMDP can thus be seen as mapping  $\pi$  from beliefs to distributions over actions, and we define

$$v^\pi(b) \triangleq \mathbb{E}_{\mathcal{A}_t \sim \pi(b_t)} \left[ \sum_{t=0}^{\infty} \gamma^t R_t \mid b_0 = b \right]. \quad (4)$$

As in MDPs, the value function associated with an optimal policy is denoted as  $v^*$  and can be computed using, for example, point-based approaches [21]. From  $v^*$ , the optimal  $Q$ -function can now be computed as

$$q^*(b, a) = \sum_{x \in \mathcal{X}} b(x) \left[ r(x, a) + \gamma \sum_{z \in \mathcal{Z}} \sum_{y \in \mathcal{X}} \mathbf{P}(y \mid x, a) \mathbf{O}(z \mid y, a) v^*(\text{Bel}(b, a, z)) \right],$$

yielding as optimal any policy  $\pi^*$  such that  $\pi^*(a \mid b) > 0$  only if  $a \in \text{argmax}_{a \in \mathcal{A}} q^*(b, a)$ .

## 4 Ad Hoc Teamwork under Partial Observability

We consider a team of  $N$  agents engaged in a cooperative task (henceforth referred as “target task”), described as an MMDP  $m = (N, \mathcal{X}, \{\mathcal{A}_n\}, \{\mathbf{P}_a\}, r, \gamma)$ . One of the agents does not know the task beforehand but must, nevertheless, engage in ad hoc teamwork with the remaining agents to complete the unknown task. We refer to such agent as the “ad hoc agent” and denote it as  $\alpha$ , and refer to the remaining  $N - 1$  agents collectively as the “teammates”. Formally, we treat the teammates as a “meta-agent” and denote it as  $-\alpha$ .

We assume that the teammates all know the target task. The ad hoc agent, however, may not know the target task nor the teammates’ policies. Instead, it knows only that the combination between target task and teammate policies is one among  $K$ , modeled as an MMDP

$$m_k = (2, \mathcal{X}_k, \{\mathcal{A}^\alpha, \mathcal{A}_k^{-\alpha}\}, \{\mathbf{P}_{k,a}, a \in \mathcal{A}\}, r_k, \gamma_k). \quad (5)$$

Note that we require the action space of the ad hoc agent,  $\mathcal{A}^\alpha$ , to be the same in all tasks. Other than that, we impose no restrictions on the state space, dynamics, or reward describing these tasks (in particular, they may all be different).

Let  $\pi_k^{-\alpha}$  denote a teammates policy for  $m_k$ ,  $k = 1, \dots, K$ . We have

$$\mathbf{P}_k(y \mid x, a^\alpha) \triangleq \mathbb{P}[X_{t+1} = y \mid X_t = x, A^\alpha = a^\alpha, A^{-\alpha} \sim \pi_k^{-\alpha}(x), M = m_k], \quad (6)$$

for  $x, y \in \mathcal{X}_k$ , where we write  $M = m_k$  to denote the fact that the transitions in (6) concerns task  $k$ .

Let us now suppose that, at each moment, the ad hoc agent cannot observe the underlying state of the environment. Instead, at each time step  $t$ , the agent can only access an observation  $Z_t$ . We assume that the observations  $Z_t$  take values in a (task-independent) set  $\mathcal{Z}$  and depend both on the underlying state of the environment and the previous action of the agent (not the teammates). Specifically, for each task  $k = 1, \dots, K$ , we assume that there is a family of task-dependent observation probability matrices,  $\mathbf{O}_{k,a^\alpha}$ ,  $k = 1, \dots, K$ ,  $a^\alpha \in \mathcal{A}^\alpha$ , with

$$\mathbf{O}_k(z \mid x, a^\alpha) = \mathbb{P}[Z_t = z \mid X_t = x, A_{t-1}^\alpha = a^\alpha, M = m_k]. \quad (7)$$

The elements  $[\mathbf{O}_{k,a^\alpha}]_{xz}$  are only defined for  $x \in \mathcal{X}_k$ . Thus, from the ad hoc agent’s perspective, each task  $k$  defines a POMDP  $\hat{m}_k$  corresponding to the tuple

$$(\mathcal{X}_k, \mathcal{A}^\alpha, \mathcal{Z}, \{\mathbf{P}_{k,a^\alpha}, a^\alpha \in \mathcal{A}^\alpha\}, \{\mathbf{O}_{k,a^\alpha}, a^\alpha \in \mathcal{A}^\alpha\}, r_k, \gamma_k).$$

We denote the solution to  $\hat{m}_k$  as  $\hat{\pi}_k$ .

### 4.1 Algorithm

We adopt a Bayesian framework and treat the target task/teammate model as a random variable,  $M$ , taking values in the set of possible model descriptions,  $\mathcal{M} = \{m_1, \dots, m_K\}$ . For  $m_k \in \mathcal{M}$ , let  $p_0(m_k)$  denote the ad hoc agent’s prior over  $\mathcal{M}$ . Additionally, let  $H_t$  denote the random variable corresponding to the history of the agent up to time step  $t$ , defined as

$$H_t = \{a_0^\alpha, z_1, a_1^\alpha, z_2, \dots, a_{t-1}^\alpha, z_t\}. \quad (8)$$

Then, given a history  $h_t$ , we define

$$p_t(m_k) \triangleq \mathbb{P}[M = m_k \mid H_t = h_t], \quad m_k \in \mathcal{M}. \quad (9)$$

The distribution  $p_t$  corresponds to the agent’s belief about the target model at time step  $t$ . The action for the ad hoc agent at time step  $t$  can be computed within our Bayesian setting as

$$\pi_t(a^\alpha | h_t) = \sum_{k=1}^K \hat{\pi}_k(a^\alpha | b_{k,t}) p_t(m_k),$$

where

$$b_{k,t}(x) \triangleq \mathbb{P}[X_t = x | H_t = h_t, M = m_k], \quad (10)$$

for  $x \in \mathcal{X}_k$ . Upon selecting an action  $a_t^\alpha$  and making a new observation  $z_{t+1}$ , we can update  $p_t$  by noting that

$$p_{t+1}(m_k) = \frac{1}{\rho} p_t(m_k) \cdot \mathbb{P}[A_t^\alpha = a_t^\alpha, Z_{t+1} = z_{t+1} | M = m_k, H_t = h_t],$$

where  $\rho$  is some normalization constant. Moreover,

$$\begin{aligned} \mathbb{P}[A_t^\alpha = a_t^\alpha, Z_{t+1} = z_{t+1} | M = m_k, H_t = h_t] \\ = \sum_{y \in \mathcal{X}_k} \mathbf{O}_k(z_{t+1} | y, a_t^\alpha) \pi_t(a^\alpha | h_t) \\ \cdot \mathbb{P}[X_{t+1} = y | A_t^\alpha = a_t^\alpha, H_t = h_t, M = m_k], \end{aligned}$$

where the last equality follows from the fact that the agent’s action selection given the history does not depend on the model  $M$ . Therefore,

$$\begin{aligned} \mathbb{P}[A_t^\alpha = a_t^\alpha, Z_{t+1} = z_{t+1} | M = m_k, H_t = h_t] \\ = \sum_{x, y \in \mathcal{X}_k} \mathbf{O}_k(z_{t+1} | y, a_t^\alpha) \mathbf{P}_k(y | x, a_t^\alpha) b_{k,t}(x) \pi_t(a^\alpha | h_t). \end{aligned}$$

Putting everything together, we get

$$p_{t+1}(m_k) = \frac{1}{\rho} \sum_{x, y \in \mathcal{X}_k} \mathbf{O}_k(z_{t+1} | y, a_t^\alpha) \mathbf{P}_k(y | x, a_t^\alpha) \cdot b_{k,t}(x) \pi_t(a^\alpha | h_t) p_t(m_k), \quad (11)$$

with

$$\rho = \sum_{k=1}^K \sum_{x, y \in \mathcal{X}_k} \mathbf{O}_k(z_{t+1} | y, a_t^\alpha) \mathbf{P}_k(y | x, a_t^\alpha) \cdot b_{k,t}(x) \pi_t(a^\alpha | h_t) p_t(m_k).$$

Note that the update (11) requires the ad hoc agent to keep track of the beliefs  $b_{k,t}$  for the different POMDPs  $\hat{m}_k$ . In other words, at each time step  $t$ , upon executing its individual action  $a_t^\alpha$  and observing  $z_{t+1}$ , the agent updates each belief  $b_{k,t}$  using (3), yielding

$$b_{k,t+1}(y) = \frac{1}{\rho} \sum_{x \in \mathcal{X}_k} b_{t,k}(x) \mathbf{P}_k(y | x, a_t^\alpha) \mathbf{O}_k(z_{t+1} | y, a_t^\alpha), \quad (12)$$

where  $\rho$  is the corresponding normalization constant. Since some of the computations in the update (11) are common to the update (12), some computational savings can be achieved by caching the intermediate values.

Finally, at each time step  $t$ , we define the *loss* for selecting an action  $a^\alpha \in \mathcal{A}^\alpha$  when the target model is  $m_k$  to be

$$\ell_t(a^\alpha | m_k) = v^{\hat{\pi}_k}(b_{k,t}) - q^{\hat{\pi}_k}(b_{k,t}, a^\alpha), \quad (13)$$

where  $\hat{\pi}_k$  is the solution to the POMDP  $\hat{m}_k$ .

It is important to note that both  $v^{\hat{\pi}_k}(b_{k,t})$  and  $q^{\hat{\pi}_k}(b_{k,t}, a^\alpha)$  can be computed from  $\hat{v}^{\hat{\pi}_k}$  at runtime, while the latter can be computed offline when solving the POMDP  $\hat{m}_k$ . A bound on  $\ell_t(a^\alpha | m_k)$  can be found in Section A of the Appendix. Note also that  $\ell_t(a^\alpha | m_k) \geq 0$  for all  $a^\alpha$ , and  $\ell_t(a^\alpha | m_k) = 0$  only if  $\hat{\pi}_k(a^\alpha | b_{k,t}) > 0$ .

## 5 Evaluation

We setup the ad hoc evaluation method from Stone et al. [26]. A single trial therefore consists on running an agent  $\alpha_0$  (representing the ad hoc agent) with an agent  $\alpha_1$  (representing the teammate) in an environment which is modeled as a POMDP. We then register the total accumulated reward in order to assess the overall team’s performance. Since the assumption that there is no available reward signal nor visible teammate’s action has never been explored, we compare our approach, ATPO, against five other agents which take on the role of  $\alpha_0$ :

- *Value Iteration (non-ad hoc; full observability)*: an agent which is able to perfectly observe the state of the environment and knows the underlying MMDP’s optimal policy. Performance reference not as baseline, but as an ‘oracle’ approach for the best theoretical performance in each domain;
- *Perseus (non-ad hoc; partial observability)*: an agent, which like ATPO, is only able to observe a partial observation of the environment, but unlike ATPO, knows the task it is performing and the teammate’s behavior;
- *Random-Picker (ad hoc, partial observability)*: an agent which, like ATPO, also has the library of models and Perseus solutions but selects one randomly in each step (allowing an evaluation of the similarity between possible models);
- *BOPA (ad hoc; full observability)*: the current state-of-the-art in ad hoc teamwork which relies on the state of the environment;
- *Random (non-ad hoc)*: an agent which selects its action randomly in every step. Reference as the worst possible behavior in each domain;

We then ask three main research questions:

1. Is ATPO effective in assisting the teammates in completing the tasks?
2. Is ATPO robust in scaling to POMDPs with large state spaces?
3. Is ATPO robust in scaling to large sets of POMDPs?

### 5.1 Domains

In order to answer all three research questions, we setup a total of 70 POMDPs from 11 different domains. Further details regarding each domain can be found in Appendix B. Table 1 summarizes the characteristics of all POMDPs from each of the eleven domains. Some domains allow for different teammates (varying in POMDP’s transition probabilities  $\mathbf{P}$ , which incorporate the policy of the teammate  $\pi_1$ ), others for different task goals (by varying the reward function  $r$ ), and others for both. One domain (abandoned power plant) even allows for different state spaces between different tasks (varying both  $r$  and  $\mathcal{X}$ ).

#### 5.1.1 Gridworld

The gridworld domain represents a standard multi-agent navigation problem. In the gridworld domain,  $N$  agents must reach  $N$  static destination (goal cells). Goal cells are fixed and not encoded in the state, only in the reward. Different possible configurations for the goals thus correspond to the different tasks in  $\mathcal{M}$ .

### 5.1.2 Pursuit

The pursuit domain used in our work is a modification of the original Pursuit domain where we add partial state observability. In it,  $N$  agents ( $N = 2$  in our experiments) must capture a single moving prey by surrounding it in a coordinated way.

### 5.1.3 Abandoned Power Plant

The abandoned power plant domain models a scenario where two agents—a robot and a human—must cooperate in order to secure an abandoned power plant with six rooms, where three are yet to be explored and two are filled with toxic waste. The ad hoc agent takes in the role of the robot and the teammate the role of the human. In this domain, there are two possible tasks which the human may be performing—charting the entire plant (exploring the three unexplored rooms) and securing the entire plant (by cleaning the two dirty rooms of all available toxic material). Each task has different modelling requirements, and as such, even though both are modeled as POMDPs, they have different state spaces.

### 5.1.4 NTU, ISR, MIT, Pentagon and CIT

The ntu, isr, mit, pentagon and cit domains [16, 19] all model navigation tasks where two agents— $\alpha_0$  and  $\alpha_1$  are spawned in random location of a map and must reach two destinations. Unlike the grid-world domain, which models an open gridworld, the ntu, isr, mit, pentagon and cit domains all model close quarters gridworlds, with walls and collision between the agents. The domains abide by the same base rules yet differ in terms of the map layout (with some being smaller and easier to navigate and others being larger and harder to navigate). They therefore have different state spaces, however, with states modeled the same way.

### 5.1.5 Overcooked

The Overcooked domain [7] models a scenario where two agents, a helper and a cook, are required to cooperate in order to cook soups. In this case, tasks represent different teammates which the ad hoc agent must (i) identify; and (ii) adapt to, in order to cook soups. Our ad hoc agent plays the role of the *helper*, which has the goal of providing the cook with onions and plates, placing them on a kitchen counter and the teammate plays the role of the *cook*, which has the goal of cooking a soup using three onions and serving it in a plate (which is then dispatched through a window).

Given that we are evaluating a total of six agents in the role of  $\alpha_0$  and three of them (ATPO, Perseus, and Random-Picker) require the approximate solutions to the POMDPs, we solve all 70 POMDPs from each domain using the Perseus algorithm [25]. Hyperparameters for the Perseus algorithm can be found in Table 5 of the Appendix and the times it takes to both setup and solve each POMDP can be found in Table 2.

## 5.2 Results

We run, for all 11 environments, 32 independent trials for each of the six agents in the role of  $\alpha_0$ . In each trial, we randomly select one of the possible POMDPs  $m \in \mathcal{M}$  and run the interaction for a horizon of  $h$  timesteps (the same horizon used for the Perseus algorithm, which can be found in Table 5 of the Appendix). We then compare the average accumulated reward over the  $h$  timesteps and 32 trials for all agents. Table 3 displays the obtained results.

At first glance, we may from these results conclude which domains are easier and which domains are harder. This can be seen by looking at the performance obtained by the Random Agent. Domains such as ntu and abandoned power plant, in which the Random Agent was able to achieve its highest possible results, are also the domains with the smallest state space. On the other hand, domains like cit and mit, where the Random Agent obtained its worst results, are also the domains with the highest number of possible states, showcasing that the problem difficulty is associated with the number of states in the POMDP.

From these results we are able to observe that all domains were effectively solved by the Perseus algorithm, providing a solid foundation for ATPO’s library. In one domain (abandoned power plant, one of the easiest), the agent which knows the task and teammate behavior, but is unable to observe the state of the environment and therefore resorts the Perseus policy (Perseus agent), was even able to, on average, slightly outperform the Value Iteration agent (although an insignificant difference).

Given that ATPO is required to identify, within its library  $\mathcal{M}$ , the POMDP which best models the environment (unlike the Perseus agent which already knows which one it is), we expected it to underperform the Perseus agent. In all domains this phenomenon was observed, further showcasing that our testbed is, in fact, a fair testbed for ATPO’s evaluation.

We can also compare ATPO against the Random-Picker agent in order to identify which environments contain the most similar tasks and teammates. Given that the Random-Picker can be seen as an instance of ATPO with a constant uniform belief over the possible models, it is expected in most cases to underperform ATPO, since Random-Picker will most likely use the wrong policy. In domains where this doesn’t happen, and Random-Picker is able to obtain similar results to ATPO, it means the tasks and teammates are very similar. As expected, ATPO outperforms Random-Picker in the different environments, showcasing the impact of the task identification ability of ATPO. In some scenarios, the difference is not very large (e.g., MIT, NTU), indicating that the task identification component is not critical to attain a good performance. However, in several other tasks, the difference is quite significant (e.g., Gridworld, Pentagon, Pursuit), and clearly shows the advantages of ATPO. As for BOPA, since the agent has full state observability, it is unsurprising that it attains better performance than ATPO - the difference attesting, to some extent, the impact that partial state observability can have in the ability of the agent to act.

Regarding the impact of the ATPO’s Bayesian inference, our experiments already hint to the impact of model identification in the performance of ATPO. However, we should note that during our experiments we also keep track of the beliefs over tasks, as ATPO interacts with the environment. By analyzing how the beliefs evolve over time, we can highlight a couple of more salient aspects. Specifically, if we look at the average probability associated to the correct model against that associated with the other tasks, we note that i) the correct task is identified, in average, after 5 timesteps, ii) in average, after 10 timesteps, the correct task has an associated probability of around 70% and iii) in average, after 20 timesteps, the correct task has an associated probability of around 90%. These values are not uniform across domains, but roughly indicate that the proposed approach is indeed able to effectively “zero-in” on the correct task. Analyzing this same results in more detail, now considering how the beliefs evolve for each scenario individually, we can report that, in the most difficult domains, it takes around 10-15 timesteps for the correct task to reach the highest probability (becoming the most likely one), and

**Table 1:** Multi-agent domains used for the ad hoc evaluation.  $\pi_1$  represents whether the domain allows for varying the teammate’s policy,  $r$  represents whether the domain allows for varying the target task,  $|\mathcal{X}|$ ,  $|\mathcal{A}|$  and  $|\mathcal{Z}|$  represent the total number of states, actions and observations per POMDP from each domain (respectively),  $\epsilon$  represents the amount of noise used in computing the transition and observation probabilities for each POMDP,  $h$  represents the number of steps in the trajectory horizon during which each POMDP is ran as an environment and, finally,  $|\mathcal{M}|$  represents the total number of different setup POMDPs for each domain (which also corresponds to the size of our approach’s model library).

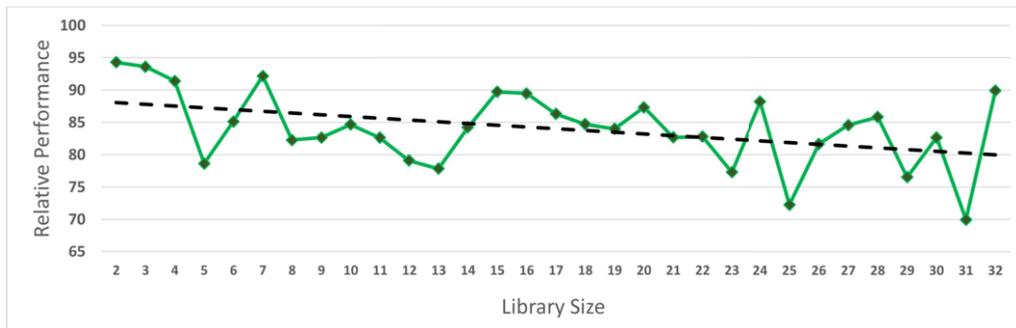
Environment	$\pi_1$	$r$	$ \mathcal{X} $	$ \mathcal{A} $	$ \mathcal{Z} $	$\epsilon$	$h$	$ \mathcal{M} $
gridworld	non-varying	varying	626	5	81	0.20	50	2 $\rightarrow$ 32
pursuit-task	non-varying	varying	626	5	81	0.20	75	4
pursuit-teammate	varying	non-varying	626	5	81	0.15	85	2
pursuit-both	varying	varying	626	5	81	0.15	85	8
abandoned power plant	non-varying	varying	varying (97, 105)	6	6	0.20	50	2
ntu	non-varying	varying	241	5	81	0.20	75	4
overcooked	varying	non-varying	1730	4	1730	0.00	50	4
isr	non-varying	varying	1807	5	81	0.20	75	3
mit	non-varying	varying	2163	5	81	0.20	75	3
pentagon	non-varying	varying	2653	5	81	0.20	75	3
cit	non-varying	varying	4831	5	81	0.10	85	3

**Table 2:** Times to setup and solve POMDPs from each domain. Setup times represent the total time it takes to create the data structures of the POMDP and solve times represent the total time it takes the Perseus algorithm [25] to obtain a policy.

Environment	Avg. Time to Setup	(Std. Dev.)	Avg. Time to Solve	(Std. Dev.)
gridworld	27s 411ms	( $\pm 668$ ms)	41m 02s 934ms	( $\pm 41$ m 10s 857ms)
pursuit-task	03m 38s 023ms	( $\pm 02$ s 316ms)	2h 30m 54s 820ms	( $\pm 1$ h 03m 45s 600ms)
pursuit-teammate	03m 21s 236ms	( $\pm 112$ ms)	2h 31m 23s 526ms	( $\pm 14$ m 45s 856ms)
pursuit-both	03m 31s 369ms	( $\pm 03$ s 085ms)	3h 03m 02s 597ms	( $\pm 1$ h 16m 14s 350ms)
abandoned power plant	01s 406ms	( $\pm 136$ ms)	43s 322ms	( $\pm 19$ s 601ms)
ntu	14s 004ms	( $\pm 502$ ms)	10m 05s 246ms	( $\pm 01$ m 24s 147ms)
overcooked	04m 00s 960ms	( $\pm 19$ s 409ms)	05m 22s 543ms	( $\pm 01$ m 16s 879ms)
isr	08m 22s 774ms	( $\pm 19$ s 985ms)	6h 02m 58s 980ms	( $\pm 1$ h 15m 47s 677ms)
mit	11m 41s 286ms	( $\pm 52$ s 980ms)	5h 43m 18s 890ms	( $\pm 4$ h 20m 11s 880ms)
pentagon	16m 14s 605ms	( $\pm 13$ s 107ms)	4h 55m 49s 658ms	( $\pm 1$ h 31m 04s 993ms)
cit	51m 45s 357ms	( $\pm 58$ s 266ms)	11h 53m 30s 001ms	( $\pm 5$ h 06m 50s 992ms)

**Table 3:** Average accumulated reward for the six agents in the eleven domains over 32 trials. BOPA requires a common state-space, so it cannot be run in the abandoned power plant scenario.

Environment	Value Iteration		Perseus		ATPO (ours)		Random-Picker		BOPA		Random	
gridworld	95.62	( $\pm 1.54$ )	93.44	( $\pm 3.13$ )	86.53	( $\pm 9.12$ )	39.91	( $\pm 61.34$ )	93.38	( $\pm 3.14$ )	5.88	( $\pm 64.13$ )
pursuit-both	90.19	( $\pm 8.32$ )	84.44	( $\pm 11.76$ )	82.72	( $\pm 12.81$ )	35.34	( $\pm 63.18$ )	90.06	( $\pm 7.47$ )	20.72	( $\pm 69.47$ )
pursuit-task	92.72	( $\pm 4.87$ )	88.03	( $\pm 7.51$ )	82.59	( $\pm 11.18$ )	47.59	( $\pm 55.46$ )	92.81	( $\pm 3.48$ )	14.59	( $\pm 71.34$ )
pursuit-teammate	89.88	( $\pm 8.38$ )	86.97	( $\pm 12.04$ )	86.41	( $\pm 14.52$ )	81.34	( $\pm 31.56$ )	91.38	( $\pm 6.13$ )	-7.66	( $\pm 78.7$ )
abandoned power plant	94.91	( $\pm 1.99$ )	95.12	( $\pm 2.1$ )	94.75	( $\pm 2.35$ )	93.81	( $\pm 4.97$ )	N.A.	N.A.	47.44	( $\pm 62.31$ )
ntu	98.00	( $\pm 0.0$ )	97.25	( $\pm 0.43$ )	97.19	( $\pm 0.53$ )	97.28	( $\pm 0.51$ )	98.0	( $\pm 0.0$ )	88.00	( $\pm 7.42$ )
overcooked	39.12	( $\pm 60.73$ )	39.00	( $\pm 60.63$ )	25.31	( $\pm 66.43$ )	27.94	( $\pm 64.58$ )	24.81	( $\pm 66.19$ )	-46.75	( $\pm 18.1$ )
isr	92.41	( $\pm 2.68$ )	91.16	( $\pm 4.68$ )	82.62	( $\pm 12.19$ )	60.59	( $\pm 59.15$ )	91.66	( $\pm 4.22$ )	-23.31	( $\pm 68.05$ )
mit	84.50	( $\pm 0.87$ )	84.09	( $\pm 0.58$ )	83.72	( $\pm 0.57$ )	83.75	( $\pm 0.75$ )	84.31	( $\pm 0.88$ )	-71.41	( $\pm 20.01$ )
pentagon	91.00	( $\pm 0.0$ )	90.62	( $\pm 0.6$ )	81.19	( $\pm 11.04$ )	53.28	( $\pm 56.53$ )	85.5	( $\pm 6.71$ )	-45.94	( $\pm 55.14$ )
cit	86.22	( $\pm 2.27$ )	85.09	( $\pm 2.53$ )	75.66	( $\pm 9.52$ )	69.91	( $\pm 14.46$ )	84.88	( $\pm 2.94$ )	-81.25	( $\pm 20.88$ )



**Figure 1:** Varying the total number of possible POMDPs from the gridworld domain in ATPO’s library. Results correspond to the relative performance, normalized between the best possible performance (100%, obtained by the Value Iteration agent) and worst possible performance (0%, obtained by the Random Agent).

by the end of the interaction, ATPO had near 100% certainty it was in fact the correct one. In the smaller and easier domains, it took only around 3-4 timesteps for the correct task to be the most likely one and, by the end of the interaction, ATPO also reached around 100% certainty that it was the correct one. A visualization of the average entropy of the beliefs for ATPO and BOPA can be found in Figure 6 in the Appendix.

Moving on, we can treat the results obtained by the Random Agent as a lower bound to an agent’s performance and the performance of the Value Iteration agent (which knows the task being performed, the behavior of the teammate, and the current state of the environment) as an upper bound. If we normalize our results into a relative scale (where 100% corresponds to the performance of the Value Iteration agent and 0% corresponds to the performance of the Random agent) we are now able to more effectively evaluate our approach. Table 4 displays the same results from Table 3, but normalized.

**Table 4:** ATPO’s results from Table 3 normalized between the performance obtained by the Value Iteration agent (100%) and performance obtained by the Random Agent (0%).

Environment	ATPO
gridworld	89.87%
pursuit-both	89.25%
pursuit-task	87.03%
pursuit-teammate	96.44%
abandoned power plant	99.66%
ntu	91.90%
overcooked	83.92%
isr	91.54%
mit	99.50%
pentagon	92.84%
cit	93.69%

Having now established a solid evaluation testbed, we can start looking at our approach’s results (ATPO). Overall, ATPO was able to obtain above 83.92% performance on all domains (showcasing its effectiveness). Its worst results were on the overcooked domain (with 83.92% relative performance) and its best results were on the abandoned power plant and mit domains (with 99.66% and 99.50%, respectively). These results therefore showcase that our approach is effective in assisting unknown teammates in solving unknown tasks resorting only to a partial observation of the environment, and robust in adapting to domains with large state spaces.

### 5.3 Increasing ATPO’s Library

We now evaluate ATPO’s robustness in scaling to larger libraries of possible POMDPs (our third and final research question). We resort to a single domain—the gridworld domain—and conduct an experiment where we increase the number of POMDPs in ATPO’s library  $|\mathcal{M}|$  from 2 to 32. We then run, per each of the 31 library sizes, 32 independent trials on the environment (adding up to a total of 992 additional trials). Figure 1 reports our obtained relative results (normalized between the best possible performance and worst possible performance, taken from the results obtained in Table 3 by the Value Iteration and Random Agent, respectively).

As we can see from these results, as we increase the number of possible POMDPs in ATPO’s library, its performance follows a small downward trend. This is expected, given that the more possible POMDPs there might be to explain the environment, the harder it will be to identify the correct one. Nevertheless, ATPO’s performance with up to 32 possible models was always above 70%, strongly showcasing our approach’s robustness in scaling to a higher number of possible models.

## 6 Conclusion

This paper makes three main contributions to the ad hoc teamwork literature: (i) formalizes the setting of ad hoc teamwork under partial observability using a Markovian framework, a setting where an ad hoc agent is unable observe the full state of the environment and the actions of its teammates, (ii) proposes a first-principled, novel model-based approach for this setting and (iii) showcases its effectiveness by conducting an empirical evaluation in a total of 70 POMDPs from 11 different domains. Our results show that our approach (ATPO), relying only on its model library of possible POMDPs describing the multi-agent systems and their respective solutions and partial observations, was able to i) effectively assist unknown teammates in performing unknown tasks, ii) robustly scale to environments with large state spaces (up to 4831 states), iii) robustly scale to a high number of possible models (up to 32 possible POMDPs) and iv) identify which model better describes the environment it is interacting with alongside its respective teammates.

A logical future line of work will be to expand our approach to domains with continuous or arbitrarily large state spaces, requiring ATPO to be adapted using function approximation solutions such as feedforward recurrent neural networks, combined with model-based planning approaches such as MuZero [24].

## Acknowledgments

This work was partially supported by national funds through FCT, Fundação para a Ciência e a Tecnologia, under project UIDB/50021/2020 (INESC-ID multi-annual funding), the HOTSPOT project, with reference PTDC/CCI-COM/7203/2020 and the RELEVANT project, with reference PTDC/CCI-COM/5060/2021. In addition, this material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-22-1-0475, and by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215. The first author acknowledges the PhD grant 2020.05151.BD from FCT. The authors thankfully acknowledge the comments from the anonymous reviewers.

## References

- [1] N. Agmon and P. Stone, ‘Leading ad hoc agents in joint action settings with multiple teammates’, in *Proc. 11th Int. Conf. Autonomous Agents and Multiagent Systems*, pp. 341–348, (2012).
- [2] A. Banerjee, ‘On Bayesian bounds’, in *Proc. 23rd Int. Conf. Machine Learning*, pp. 81–88, (2006).
- [3] S. Barrett, A. Rosenfeld, S. Kraus, and P. Stone, ‘Making friends on the fly: Cooperating with new teammates’, *Artificial Intelligence*, **242**, 132–171, (2017).
- [4] S. Barrett and P. Stone, ‘Ad hoc teamwork modeled with multi-armed bandits: An extension to discounted infinite rewards’, in *Proc. AAMAS Workshop on Adaptive Learning Agents*, (2011).
- [5] S. Barrett and P. Stone, ‘Cooperating with unknown teammates in complex domains: A robot soccer case study of ad hoc teamwork’, *Proc. 29th AAAI Conf. Artificial Intelligence*, 2010–2016, (2015).
- [6] Kalesha Bullard, Douwe Kiela, Joelle Pineau, and Jakob Foerster, ‘Quasi-equivalence discovery for zero-shot emergent communication’, *arXiv preprint arXiv:2103.08067*, (2021).
- [7] M. Carroll, R. Shah, M. Ho, T. Griffiths, S. Seshia, P. Abbeel, and A. Dragan, ‘On the utility of learning about humans for human-AI coordination’, in *Adv. Neural Information Processing Systems 32*, (2019).
- [8] D. Chakraborty and P. Stone, ‘Cooperating with a Markovian ad hoc teammate’, in *Proc. 12th Int. Conf. Autonomous Agents and Multiagent Systems*, (2013).
- [9] Alan Fern, Sriraam Natarajan, Kshitij Judah, and Prasad Tadepalli, ‘A decision-theoretic model of assistance.’, in *IJCAI*, pp. 1879–1884, (2007).
- [10] Elliot Fosong, Arrasy Rahman, Ignacio Carlucho, and Stefano V Albrecht, ‘Few shot teamwork’, *arXiv preprint arXiv:2207.09300*, (2022).
- [11] D. Fudenberg and D. Levine, *The Theory of Learning in Games*, MIT Press, 2000.
- [12] K. Genter, T. Laue, and P. Stone, ‘Three years of the RoboCup standard platform league drop-in player competition’, *J. Autonomous Agents and Multi-Agent Systems*, **31**(4), 790–820, (2017).
- [13] Piotr J Gmytrasiewicz and Prashant Doshi, ‘A framework for sequential planning in multi-agent settings’, *Journal of Artificial Intelligence Research*, **24**, 49–79, (2005).
- [14] Pengjie Gu, Mengchen Zhao, Jianye Hao, and Bo An, ‘Online ad hoc teamwork under partial observability’, in *International Conference on Learning Representations*, (2021).
- [15] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster, ‘“other-play” for zero-shot coordination’, in *International Conference on Machine Learning*, pp. 4399–4410. PMLR, (2020).
- [16] Yujing Hu, Yang Gao, and Bo An, ‘Learning in multi-agent systems with sparse interactions by knowledge transfer and game abstraction.’, in *AAMAS*, pp. 753–761, (2015).
- [17] Andrei Lupu, Brandon Cui, Hengyuan Hu, and Jakob Foerster, ‘Trajectory diversity for zero-shot coordination’, in *International Conference on Machine Learning*, pp. 7204–7213. PMLR, (2021).
- [18] F. Melo and A. Sardinha, ‘Ad hoc teamwork by learning teammates’ task’, *Autonomous Agents and Multi-Agent Systems*, **30**(2), 175–219, (2016).
- [19] Francisco S Melo and Manuela Veloso, ‘Learning of coordination: Exploiting sparse interactions in multiagent systems’, in *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pp. 773–780. Citeseer, (2009).
- [20] Reuth Mirsky, Ignacio Carlucho, Arrasy Rahman, Elliot Fosong, William Macke, Mohan Sridharan, Peter Stone, and Stefano V Albrecht, ‘A survey of ad hoc teamwork: Definitions, methods, and open problems’, *arXiv preprint arXiv:2202.10450*, (2022).
- [21] J. Pineau, G. Gordon, and S. Thrun, ‘Anytime point-based approximations for large POMDPs’, *J. Artificial Intelligence Res.*, **27**, 335–380, (2006).
- [22] M. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley-Interscience, 2005.
- [23] João G Ribeiro, Miguel Faria, Alberto Sardinha, and Francisco S Melo, ‘Helping people on the fly: Ad hoc teamwork for human-robot teams’, in *EPIA Conference on Artificial Intelligence*, pp. 635–647. Springer, (2021).
- [24] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al., ‘Mastering atari, go, chess and shogi by planning with a learned model’, *Nature*, **588**(7839), 604–609, (2020).
- [25] Matthijs TJ Spaan and Nikos Vlassis, ‘Perseus: Randomized point-based value iteration for pomdps’, *Journal of artificial intelligence research*, **24**, 195–220, (2005).
- [26] P. Stone, G. Kaminka, S. Kraus, and J. Rosenschein, ‘Ad hoc autonomous agent teams: Collaboration without pre-coordination’, in *Proc. 24th AAAI Conf. Artificial Intelligence*, pp. 1504–1509, (2010).
- [27] P. Stone and S. Kraus, ‘To teach or not to teach?: Decision making under uncertainty in ad hoc teams’, in *Proc. 9th Int. Conf. Autonomous Agents and Multiagent Systems*, pp. 117–124, (2010).
- [28] Johannes Treutlein, Michael Dennis, Caspar Oesterheld, and Jakob Foerster, ‘A new formalism, method and open issues for zero-shot coordination’, *arXiv preprint arXiv:2106.06613*, (2021).