

Fluid-Agent Reinforcement Learning

Shishir Sharma

Mila- Quebec Artificial Intelligence
Institute
McGill University
Montreal, Canada
shishir.sharma@mail.mcgill.ca

Doina Precup

Mila- Quebec Artificial Intelligence
Institute
McGill University
Montreal, Canada
doina.precup@mcgill.ca

Theodore J. Perkins

Ottawa Hospital Research Institute
University of Ottawa
Ottawa, Canada
tperkins@ohri.ca

ABSTRACT

The primary focus of multi-agent reinforcement learning (MARL) has been to study interactions among a fixed number of agents embedded in an environment. However, in the real world, the number of agents is neither fixed nor known a priori. Moreover, an agent can decide to create other agents (for example, a cell may divide, or a company may spin off a division). In this paper, we propose a framework that allows agents to create other agents; we call this a fluid-agent environment. We present game-theoretic solution concepts for fluid-agent games and empirically evaluate the performance of several MARL algorithms within this framework. Our experiments include fluid variants of established benchmarks such as Predator–Prey and Level-Based Foraging, where agents can dynamically spawn, as well as a new environment we introduce that highlights how fluidity can unlock novel solution strategies beyond those observed in fixed-population settings. We demonstrate that this framework yields agent teams that adjust their size dynamically to match environmental demands.

KEYWORDS

Multi-Agent Coordination; Reinforcement Learning; Nash Equilibrium

ACM Reference Format:

Shishir Sharma, Doina Precup, and Theodore J. Perkins. 2026. Fluid-Agent Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/TAXB8518>

1 INTRODUCTION

Historically, reinforcement learning (RL) [31], has primarily focused on the viewpoint of a single agent embedded in some environment, acting to maximize its long-term return. Much has been achieved using this viewpoint, giving rise to algorithms that have varied applications in robotics [33], healthcare [38], and economics [27], among other fields. Multi-agent reinforcement learning (MARL) scenarios have also been explored, especially in the context of game-playing [6]. World-class or even world-best players have been developed for two- or multi-player games such as Chess [26], GO [25] and Dota [4]. The vast majority of MARL research has focused on scenarios where there is a fixed number of agents, cooperating and/or competing with each other.

However, in real-world circumstances, the number of agents with which one interacts can vary widely over time. Thus, while assuming that the number of agents is fixed over time leads to mathematical simplicity and allows a clean study of equilibria, sample complexity and other problems, it is also limiting. Moreover, in many real-world situations, agents take actions that *actively modify* the number of agents in the environment, for example by creating other agents or destroying them. For instance, living beings procreate, thus creating new agents. A business might acquire another business, or spin-off a subsidiary, or even decide to close down. In short, many real-world systems include a fluid collection of decision-making agents that may be in some mixture of cooperation and competition with each other, and that collection is constantly changing—in part due to the choices of the agents themselves, which may create or eliminate other agents.

In this paper, we formalize this setting by defining the notion of a *fluid-agent* environment. We study a simplified version of this setting, in which agents are only able to create (spawn) other agents. We propose the Partially Observable Fluid Stochastic Games (POFSG) framework for modeling such fluid-agent games and prove the existence of Nash equilibria (NE) as well as subgame-perfect Nash equilibria (SPNE) in such games. We adapt existing popular environments, Predator-Prey [11] and Level-Based Foraging [22], to allow agent spawning and also introduce *PuddleBridge*, a goal-based environment centered around dynamically leveraging fluidity to achieve optimal performance. Finally, we conduct an experimental analysis to examine several facets of spawn-only fluidity, including: (a) the interplay between algorithmic inductive biases—joint versus per-agent optimization—and reward structure under fluidity; (b) the ability of agents to adapt team size in response to task variability; (c) the capacity to optimize the characteristics of spawned agents; (d) the access to a richer policy space due to fluidity, allowing agents to switch between fluid and non-fluid strategies.

2 RELATED WORK

To the best of our knowledge, no prior work has provided a formal treatment of fluid-agent environments as defined in this study. The closest practical analogue is AlphaStar [35], which permits unit creation and destruction, though without an explicit formalization. However, the AlphaStar game playing agent employs a single centralized policy to control all units. Thus, creating new units does not introduce additional decision-making entities. In contrast, our approach features decentralized, autonomous agents: multiple teammates act concurrently based on local observations (under CTDE training), and spawning introduces new decision-makers rather than merely expanding the action load of a single controller.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/TAXB8518>

Our framework is also conceptually related to *ad-hoc teamwork* (AHT) [3, 17, 28], which tackles the problem of agents coordinating with previously unseen teammates. However, there are several key differences between AHT and the setting we consider. Mirsky et al. [17] characterize AHT by three assumptions: (i) no prior coordination such as through a joint training phase, (ii) no control of agents over teammates, and (iii) a collaborative objective. Our setting is not consistent with (i) and (ii). First, we permit centralized training with decentralized execution (CTDE), which means that teammates can coordinate during training. Second, agents directly determine team composition: they decide whether and when to spawn additional teammates and optimize for their capabilities. There is somewhat greater conceptual overlap between our work and open ad-hoc teamwork (OAHT) setting [36], which permits team composition to evolve over time; however, newly introduced teammates are still externally provided rather than generated through strategic decisions within the system itself.

Our notion of population fluidity also relates to ideas from evolutionary game theory (EGT), where the size and composition of a population change over time as successful strategies replicate and weaker ones vanish. Classical models such as replicator dynamics [34, 37] describe this process at the population level, but agents in those frameworks do not make explicit decisions to reproduce or die—the dynamics arise automatically from payoff differences. In contrast, our formulation endogenizes this process: spawning is itself a strategic action taken by autonomous learners within the game. Thus, while EGT explains how advantageous behaviors can spread over time, our work examines how agents can intentionally control population change to improve collective performance.

3 BACKGROUND

A MARL problem has two parts: (a) a game model defining how the agents interact and (b) a solution concept specifying what is desired or what can be achieved by the joint behavior of the agents [1]. *Partially Observable Stochastic Games* (POSGs) provide a natural framework for modeling multi-agent interactions with a fixed number of agents. Formally, a POSG is defined as a tuple $G = \langle \mathcal{I}, \mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{Z}, \mathcal{R}, \gamma \rangle$, where $\mathcal{I} = \{1 \dots N\}$ is a finite set of agents; \mathcal{S} is the set of states; \mathcal{A}_i is the set of actions for agent i ; \mathcal{O}_i is the set of observations of agent i ; \mathcal{T} is the transition function, mapping the current state and joint action to a probability distribution over next states; \mathcal{Z}_i is the observation function from which agent i draws its observations; \mathcal{R}_i is the reward function for agent i ; $\gamma \in (0, 1)$ is the discount factor.

There exists a rich literature of solution concepts proposed for various kinds of games. The most widely studied solution concept is perhaps the Nash equilibrium (NE) [20].

Given a joint policy $\pi = (\pi^1, \pi^2, \dots, \pi^N)$, the state-value function for agent i is defined as the expected discounted return obtained when all agents follow π :

$$V^i(\pi; s) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r^i(s_t, \pi(s_t)) \mid s_0 = s \right] \quad (1)$$

A joint policy π is a Nash equilibrium if, for every state $s \in \mathcal{S}$ and every agent $i \in \mathcal{I}$,

$$V^i(\pi; s) \geq V^i(\pi^i, \pi^{-i}; s), \quad \forall \pi^i, \quad (2)$$

where π^i denotes any alternative policy available to agent i .

While NE as a solution concept has broad applicability, in sequential and dynamic settings, it suffers from limitations such as failing to ensure rationality at every decision point as well as allowing non-credible threats, i.e., a strategic action that a player claims they will take but which they would have no incentive to follow through if the relevant decision point were actually reached. To account for these limitations, solution concepts such as subgame-perfect Nash equilibria (SPNE) [24] have been proposed in the game theory literature. SPNE refines NE by requiring that the players' strategies form a NE in every subgame of the original game, where a subgame is any subtree of the original game that starts at a single decision node and contains all of its successors.

The problem of coordination among groups of agents has received considerable attention [5, 7, 14, 16]. Research has explored learning communication protocols to improve coordination [9, 29], as well as the emergence of natural language in such communication [15, 18]. In cases of no communication, agents have to be deployed independently and achieve coordination implicitly. MARL algorithms differ through the amount and type of communication that the agents are allowed among themselves. Algorithms belonging to the *decentralized training decentralized execution* (DTDE) paradigm operate by having each agent consider all other agents as part of the environment; in this case, each agent trains independently. Despite their simplicity, DTDE algorithms such as Independent-Q learning (IQL) [32] and Proximal Policy Optimization (PPO) [39] have been shown to perform competitively on several MARL tasks [22]. Another popular approach is to allow agents to access joint information in the training phase, but be completely independent during execution. These approaches fall under the category of *centralized training decentralized execution* (CTDE) [12, 21]. Value-based CTDE methods factor the joint value function into decentralized components [23, 30]. A specific implementation of this idea is Value Decomposition Networks (VDN) [30], in which the value function is decomposed into additive components, each of which is used by an individual agent. There also exists a class of algorithms that employ a centralized critic while maintaining decentralized actors. These methods, such as Multi-Agent PPO (MAPPO), leverage centralized information during training to improve coordination, yet execute policies independently at test time—embodying the CTDE framework.

4 FLUID-AGENT RL

We begin by detailing a game model that allows agents to create other agents. Subsequently, we analyze existing solution concepts and their application to the proposed fluid-agent model.

4.1 Partially Observable Fluid Stochastic Game

To model interactions in a fluid-agent environment, we modify the POSG framework and define *Partially Observable Fluid Stochastic Games* (POFSG) $G = \langle \mathcal{I}, \mathcal{S}, \mathcal{L}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{Z}, \mathcal{R}, \gamma \rangle$, where

- $\mathcal{I} = \{1 \dots N\}$ is the set of all agents that could be present in the environment at any time
- \mathcal{S} is the set of states
- $\mathcal{L} : \mathcal{S} \rightarrow 2^{\mathcal{I}}$ is the *alive* function, which returns the set of all agents present in the environment at a given state

- $\mathcal{A} = \prod_{i \in \mathcal{L}(s)} \mathcal{A}_i$ is the set of joint actions for all agents which are alive at state s ; note that each agent’s action space has a special “spawning” action, which allows it to add to the number of alive agents
- \mathcal{O}_i is the set of observations of agent i
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \text{Dist}(\mathcal{S})$ is the transition function, mapping the current state and joint action to a probability distribution over next states
- $\mathcal{Z}_i : \mathcal{S} \rightarrow \text{Dist}(\mathcal{O}_i)$ is the observation distribution of agent i
- $\mathcal{R}_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function for agent i
- $\gamma \in (0, 1)$ is the discount factor

We denote the population size at state s by $|\mathcal{L}(s)|$. From the definition of \mathcal{I} , it follows that the maximum number of agents that could be present in a fluid environment is finite. This population ceiling is enforced when an alive agent chooses to spawn in state s_t , resulting in two possible scenarios. If $|\mathcal{L}(s_t)| = N$, then the action of spawning does not change the state. If $|\mathcal{L}(s_t)| < N$, another agent is added to the population.¹

We note that the POFSG definition is related to POSG, but with two important differences: (a) the set of alive agents—and thus the joint action space—varies over time; (b) spawning directly modifies the action set and indirectly alters the set of reachable states. Consequently, spawning can affect optimal and equilibrium solutions.

Despite these differences, transitions, observations, rewards, and spawning remain Markovian: the population at the next state depends only on the current state and joint action. Thus, it is natural to consider applying existing POSG solution concepts and algorithms to this setting.

Nevertheless, fluidity fundamentally changes the strategic landscape. In fully cooperative POSGs, where agents share a common reward, spawning alters both the attainable performance and the cooperative structure itself. In particular, the maximum achievable return is no longer fixed by the initial population size: agents can increase the effective capacity of the team by spawning additional agents when tasks require greater collective effort. However, excessive spawning can deplete shared resources. Moreover, when total reward is divided among alive agents, increasing population size can reduce per-agent return, effectively breaking full cooperativity. Consequently, optimal solutions and equilibrium behavior in fluid-agent systems can differ substantially from those in fixed-population POSGs.

4.2 Nash Equilibrium

Given the parallels between the proposed POFSG framework and existing POSG framework, we now discuss whether existing solution concepts for POSG can be leveraged to analyze fluid-agent games. First, we prove the existence of a Nash equilibrium in a POFSG, by exploiting the flexibility of Fink’s proof [8] w.r.t allowing action spaces of agents to not only be variable, but to arbitrarily depend on the environment state. This allows us to transform a POFSG into a fixed-player stochastic game, with the unborn agents simply having a “Do Nothing (N)” action. Thus, spawning can be viewed as an action that transitions the environment into a state with a modified action space for spawned agents.

¹In our experiments, this is done by choosing the agent with the smallest ID that is not alive at the moment and adding it to the set of alive agents.

THEOREM 4.1 (EXISTENCE OF STATIONARY NASH EQUILIBRIUM). *Every POFSG possesses a stationary mixed–strategy Nash equilibrium.*

PROOF. We embed the POFSG into an equivalent N –player discounted stochastic game by augmenting each agent’s action space with a null action when the agent is not alive, and by defining transitions and rewards to depend only on the actions of agents alive in the current state. This construction yields a standard discounted stochastic game with fixed player set and finite state and action spaces.²

By the fixed–point argument of Fink [8], the augmented game admits a stationary Nash equilibrium. Restricting the equilibrium mixed strategies to the agents alive in each state induces a Nash equilibrium of the original fluid game:

$$\pi_{\text{fluid}}^*(s) = \{ \pi_i^*(s) \mid i \in \mathcal{L}(s) \},$$

where $\pi_i^*(s)$ denotes the equilibrium strategy of agent i in the augmented game. \square

4.3 Subgame–Perfect Nash Equilibrium

THEOREM 4.2 (EXISTENCE OF SUBGAME–PERFECT NASH EQUILIBRIUM). *Every finite-horizon POFSG with publicly observed joint actions and perfect recall possesses a subgame–perfect Nash equilibrium.*

PROOF. We prove the theorem in three steps.

Step 1: Sequentialization: Consider a stage t with state s_t and alive agents $\mathcal{L}(s_t)$. Transform the simultaneous-move stage into a sequential-move extensive form by fixing a common-knowledge ordering of agents (e.g., by increasing ID). Let G^{seq} denote the resulting extensive-form game, whose decision nodes are tuples (t, s_t, i) indicating stage t , state s_t , and the turn of agent i . While the assumption of observed joint actions allows the information set at any decision node to contain information of all agents’ moves in the previous stages, agents remain unaware of any moves made by other agents in the current stage, hence preserving the information structure of the simultaneous-move game.

Step 2: Perfect recall and subgame definition: We assume agents have *perfect recall*: each agent remembers its own past actions and private observations. Together with public monitoring of joint actions, this implies that in G^{seq} every agent knows both the complete public history and its own private history at each of its decision points. Consequently, G^{seq} is a finite-horizon extensive-form game with perfect recall.

A *stage-start subgame* is rooted at a node (t, s_t, i_1) , where i_1 is the first mover of stage t . Since the public history is common knowledge and private histories are known to the respective agents, this node belongs to a singleton information set. Thus every stage-start node defines a proper subgame of G^{seq} .

Step 3: Existence and lifting to the original POFSG: The game G^{seq} is finite (finite horizon, bounded population, finite action sets). Because G^{seq} is a finite extensive-form game, every proper subgame is itself a finite extensive-form game. By Nash’s theorem [19], every finite game admits a mixed-strategy Nash equilibrium. Applying this result to each subgame and proceeding by backward induction over the finite horizon yields a mixed-strategy subgame–perfect

²A complete construction of the augmented game is provided in the extended version of this paper, available at <https://arxiv.org/abs/2602.14559>




Figure 1: Fluid-agent Environments

Nash equilibrium of G^{seq} . As G^{seq} has perfect recall, Kuhn’s theorem [13] guarantees the existence of an *outcome-equivalent* behavioral-strategy SPNE.

Finally, we map this equilibrium back to the original POFSG. In G^{seq} , an agent’s strategy depends only on information that is available in the original POFSG by construction. Hence the behavioral strategy of G^{seq} induces a well-defined policy for the simultaneous-move POFSG. Because the equilibrium is subgame-perfect in G^{seq} , it is sequentially rational at every information set of the game. The induced behavioral strategy profile therefore constitutes a subgame-perfect Nash equilibrium of the original POFSG. □

5 FLUID-AGENT ENVIRONMENTS

5.1 Fluid Predator-Prey

We work with a predator-prey class of environments [11], an example of which is shown in Fig. 1a.

5.1.1 Dynamics: The environment consists of a square grid, with predators depicted as blue circles and preys as red. Blue squares in the cardinal direction of each predator denote their range for capturing prey. However, a prey must be in range of two different predators in order to be captured; thus, the predators must cooperate. At any time, a square can only be occupied by a single entity. Movement by either a predator or a prey is therefore only possible if the target square is empty. The preys follow a fixed policy of either staying at the same position, with probability 0.3, or randomly moving in any direction with probability 0.175, unless the randomly selected choice would move it into range of a predator. In this case, the prey simply stays still. Once the team of agents takes an action, the target position for each agent, starting with the agent with ID 1, is checked and if empty, the agent is moved. After all agents are at their new positions, any preys now within range of at least 2 agents are captured and removed from the grid. Finally, the preys that are still alive are moved to their new positions. An episode ends when either all the preys are captured or when 100 time steps have elapsed.

5.1.2 Observation Space: We restrict the visual field of each agent to an 11×11 window centered around it, as in previous works [32]. The agents observe their own coordinates as well as their ID, the latter being necessary for performing parameter sharing [10]. While these three components might suffice for solving the Predator-Prey

environment, adding spawning dynamics to the environment makes it much harder to solve. Consequently, in order for the agents to learn to spawn, we augment the observation space of each agent with the team’s joint action taken in the previous step, a *parent indicator* i.e. the number of children the agent itself has spawned, and the number of preys and agents present in the environment. The motivation behind conditioning the agents on the previous joint action and the number of agents spawned is to keep the credit assignment problem tractable. The information about the number of preys and agents present is added as it is crucial for fluid-agent groups to accurately judge whether spawning is useful.

5.1.3 Action Space: Each agent’s action space consists of actions move[each direction], do nothing, and spawn. After a spawning action, new agents are spawned randomly at an empty position. There are no restrictions on how many agents can be spawned in one step, but the total number of agents is bounded.

5.1.4 Reward Function: We retain a common-reward structure i.e. all alive predators receive the same per-agent reward. The design choice is whether to normalize the payoff for capturing prey by the current population size. Let P be the prey_capture_reward, c_{spawn} the spawn penalty, c_{step} the per-step cost per alive agent, $N_{\text{cap}}(s, a) \in \mathbb{N}_0$ the number of captures on (s, a) , and $N_{\text{sp}}(s, a) \in \mathbb{N}_0$ the number of successful spawns. Each alive agent $i \in \mathcal{L}(s)$ receives

$$r_i(s, a) = \underbrace{\text{Payoff}(s, a; \mathcal{L}(s))}_{\text{predation reward}} - \underbrace{\left(\frac{c_{\text{spawn}} N_{\text{sp}}(s, a)}{|\mathcal{L}(s)|} + c_{\text{step}} \right)}_{\text{costs (spawn split per capita; step per agent)}},$$

where $\mathcal{L}(s)$ is measured pre-transition.

Size-constant payoff (SCP) does not change with population size:

$$\text{Payoff}(s, a; \mathcal{L}(s)) = P N_{\text{cap}}(s, a)$$

Here, the per-agent capture reward is invariant to $|\mathcal{L}(s)|$, while the joint capture reward scales linearly with the number of agents.

Size-inverse payoff (SIP) performs a normalization per capita:

$$\text{Payoff}(s, a; \mathcal{L}(s)) = \frac{P N_{\text{cap}}(s, a)}{|\mathcal{L}(s)|}$$

Here, the *per-agent* capture payoff decreases with $|\mathcal{L}(s)|$, while the *joint* capture payoff remains invariant (summing identical per-agent rewards over $|\mathcal{L}(s)|$ recovers $P N_{\text{cap}}$). Spawn cost is per-capita in both variants; c_{step} is a per-alive linear pressure in both.

5.2 Fluid Level-Based Foraging

We adapt the Level-Based Foraging (LBF) environment [2], illustrated in Fig.1b, to a fluid paradigm.

5.2.1 Dynamics: The world is a square grid populated by food items (red apples labeled with their level) and agents (black silhouettes labeled with their levels). Agents can only attempt to collect food in the four neighboring cells. Any collection attempt of a level ℓ food is only successful if the sum of levels of surrounding agents is $\geq \ell$.

5.2.2 Observation Space: Unlike Predator-Prey, here we consider the fully observed setting. Each agent receives an observation, encoding the positions and levels of all food items (x, y, ℓ) , and for each agent, the positions, levels, alive flags, parent indicators, and previous actions $(x, y, \ell, \text{alive}, \text{parent}, a_{\text{prev}})$. Two additional scalars record the environment’s population capacity and the agent’s own ID. This yields an observation of dimension, $\text{obs_dim} = 3N_{\text{food}} + 6N_{\text{agents}} + 2$, allowing agents to reason jointly over food distribution, team composition, and population constraints when deciding whether to spawn or act cooperatively.

5.2.3 Action Space: Originally, the action space of each agent consisted of actions `do nothing`, `move[each direction]` and `load`. In the fluid version, we also include a `spawn` action that lets any alive agent spawn another agent at a random empty location in the grid. The spawned agent inherits the level of its parent agent. Every square can hold at most one entity, and movement (N, S, E, W) succeeds only when the destination cell is empty. As with our Predator-Prey variant, entity updates are executed sequentially in order of increasing agent ID, such that collision resolution is deterministic.

5.2.4 Reward Function: Loading a food of level ℓ results in each agent receiving a share of ℓ weighted by their level. As in Predator-Prey, the spawn cost c_s is shared among all alive agents. However, agents that do not participate in collecting a food do not receive any reward. Episodes terminate when all foods are consumed or after 100 steps. A step cost c_{step} applies to all alive agents.

5.3 PuddleBridge

We introduce the *PuddleBridge* environment (Fig. 1c), an 8×8 grid world featuring land (checkered blocks), walls (grey blocks with diagonal lines), puddles (blue blocks with circles), a designated spawn cell (represented by \star), and a goal cell (represented by \blacktriangleright). Agents (represented by white circles) begin in the spawn cell and must coordinate to reach the goal on the opposite side of a central wall barrier. The land, spawn cell and goal are single-occupancy; Walls are impassable.

5.3.1 Puddle Dynamics:

- **Capacity.** Entering a puddle is permitted iff its current occupancy is < 2 .
- **Stacking.** If an agent moves into a puddle cell already holding one agent, given that the already present agent stays in the same position, the recently moved agent becomes the *top*; the cell is then a 2-stack.

- **Bottom lock.** While a cell is a 2-stack, the bottom agent is immobilized for that entire step and can only spawn; the top agent may move and spawn.³
- **Puddle \rightarrow puddle mobility.** A puddle \rightarrow puddle move is permitted *only* for the *top* agent of a 2-stack. In that case, the move is treated like a land \rightarrow puddle move: any destination puddle with free capacity is admissible (it need not already be occupied). A lone occupant of a puddle cell cannot perform a puddle \rightarrow puddle move.
- **Unstacking.** When a 2-stack loses its top (the only one allowed to leave), the cell reverts to single occupancy.

The lightly shaded blocks in Fig. 1c represent walls that may or may not be present in each episode, simulating a gating mechanism. When this group of walls is absent, i.e. the gate is open, the agent can reach the goal state by itself, by going across the central wall barrier. If the walls are present, i.e. the gate is closed, this route is blocked. There is an alternative route through the puddles, but it requires coordination with another agent to move from puddle \rightarrow puddle. The agent stacking mechanic enables a physical “bridge” over the puddle, where the top agent can move onward while the bottom agent remains submerged, effectively turning local cooperation into a spatial resource.

5.3.2 Observation Space: For every grid cell, each agent’s observation includes a one-hot encoding of the tile type and a scalar ID feature representing the base occupant (bottom agent if stacked). Each puddle cell contributes an additional virtual slot recording the top agent’s ID. Agent-specific information is appended at the end: the agent’s normalized position (r, c) , its own ID, and a one-hot vector of previous actions for every agent in the population.

5.3.3 Action Space: Each agent selects one of six discrete actions: none, north, south, west, east, or spawn. The spawn action attempts to create a new agent at the spawn cell if it is unoccupied and the episode’s population cap is not yet reached. Agents act in ascending order of IDs to ensure deterministic collision resolution.

5.3.4 Reward Function: Reaching the goal yields a team reward of 10, shared equally among all currently alive agents. Each alive agent pays a per-step cost c_{step} , and successful spawns incur a team spawn cost c_{spawn} , equally divided among agents alive at the start of the step. Episodes terminate when any agent reaches the goal or after 100 steps.

5.4 Exploration Paradigm

Introducing fluidity exacerbates the challenge of exploration. Beyond exploring the environment itself, agents must now also explore how learned behaviors interact with different population sizes, as these alter the environment’s dynamics. Naïve exploration strategies often become trapped in suboptimal extremes of population size. To address this, during training, we randomly sample both the initial population and a population ceiling (less than or equal to the true ceiling) at the start of each episode. This forces agents to experience and adapt to a wider range of population sizes that would

³This is evaluated from the pre-move state each step; once a cell is a 2-stack at step start, the bottom remains locked for that step even if the top leaves earlier in the move ordering.




Figure 2: Algorithmic comparison on Predator-Prey environment under different payoff normalizations.

otherwise be underexplored. For all fluid-agent environments discussed thus far, we additionally include this sampled ceiling in the agents’ observation space. During evaluation, the sampled ceiling is replaced with the true population limit. While randomizing the initial population is straightforward in Predator-Prey, owing to the complete homogeneity among agents, the process is more involved in LBF and PuddleBridge. In LBF, we handle this by sampling both population size and agent levels randomly at initialization, constraining the sampled levels to the maximum level specified in the experimental configuration. In PuddleBridge, where spatial relationships are critical, any agents beyond the first are placed only in cells adjacent to the spawn cell, ensuring consistent starting configurations.

Moreover, for algorithms employing simple exploration schemes such as ϵ -greedy (e.g., IQL and VDN), we introduce a separate exploration parameter ϵ_{spawn} for spawning, which increases linearly throughout training. When randomly choosing an action, we choose spawn with probability ϵ_{spawn} and other actions with probability $(1 - \epsilon_{spawn}) / (|\mathcal{A}| - 1)$. This approach, inspired by curriculum learning, allows agents to first learn effective behaviors at smaller population sizes before progressively adapting their policies to settings with larger and more variable populations.

6 EMPIRICAL ANALYSIS

6.1 Experimental Design

This section outlines the design and rationale behind our experimental setup.⁴ We design environments and evaluation protocols that allow us to probe specific questions about how fluidity interacts with learning dynamics, coordination structure, and policy optimality.

6.1.1 Algorithmic Inductive bias vs. Fluid reward: Rather than adopting the conventional classification of algorithms into CTDE and DTDE, we instead categorize them by their optimization target i.e. whether they optimize the per-agent or joint return. This distinction is particularly salient under fluidity, since unlike in static settings where the two objectives coincide for cooperative agents, fluid-agent environments introduce an additional means of improving the joint return: adapting the team size itself. This distinction is illustrated by the case of a size-constant payoff: agents may achieve

an optimal per-agent return yet fail to maximize the joint return if they neglect to spawn additional agents whose inclusion would yield a net positive contribution to the joint return. To examine how an algorithm’s optimization target interacts with the reward structure under fluidity, we evaluate performance under two distinct payoff formulations. With a size-constant payoff, algorithms optimizing the joint return are incentivized to spawn additional agents as a means of increasing total reward, whereas per-agent optimizers lack such incentive unless the spawned agents contribute a net positive return to their parent. The dynamics become more balanced under a size-inverse payoff, where spawning provides no direct advantage in joint return. In this case, the only incentive to spawn is to increase the average per-agent return experienced by the alive agents. We conduct experiments using IQL, VDN, PPO, and MAPPO, evaluating the latter under two critic configurations: one with concatenated observations as input and another using the global state (the full predator-prey grid), which we denote as MAPPO_state.

6.1.2 Task Variability: A key advantage expected of fluid agents over a fixed group is their ability to dynamically adapt to task variability. To test this hypothesis, we conduct an ablation study in the Predator-Prey environment, where the number of preys is randomly sampled from a multinomial distribution at the start of each episode. With a substantial spawn cost, a fluid group should learn to limit spawning when preys are scarce and expand when they are abundant. Since fixed groups incur no spawn cost, we subtract these costs post hoc for fairness. Although this adjustment introduces different learning dynamics, the fluid group nonetheless faces a significantly harder problem—it must learn not only to capture preys efficiently but also to optimize its own population size.

6.1.3 Adaptive Team Composition: We next examine whether agents can learn to optimize team composition, deciding not only whether to spawn but which type of agent to add. Because all spawned agents in Predator-Prey are indistinguishable, we turn to LBF, where we introduce an inherited parent level rule—each child inherits the level of its parent. We then construct a scenario with two initial agents of levels $\{1, 2\}$, food items of levels $\{2, 3, 4, 5\}$, and a maximum of four agents. With a fixed spawning cost per agent, the optimal policy is to spawn exactly one additional agent, specifically of level 2, since this configuration alone enables the team to collect

⁴Full hyperparameter specifications and implementation details are available in the extended version of the paper (<https://arxiv.org/abs/2602.14559>).

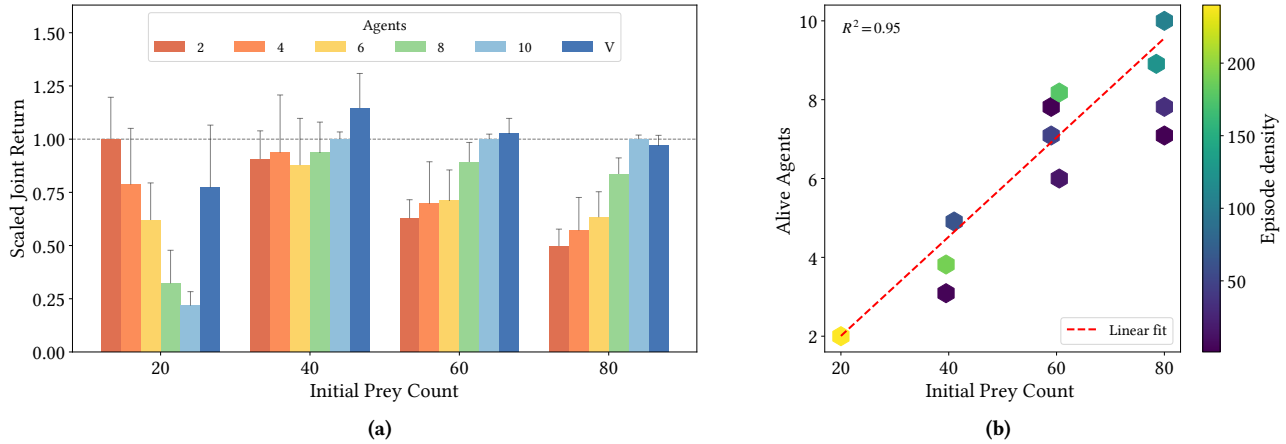


Figure 3: (a) Fluid and fixed groups on tasks with preys sampled from a distribution. Error bars represent one standard deviation. (b) Density plot showing the relationship between episodic prey availability and converged fluid-agent populations.

all four foods while paying the smallest total spawn cost. We set the spawn and step costs to $c_{\text{spawn}} = 1.0$ and $c_{\text{step}} = 0.025$.

6.1.4 Emergence of Richer Policies through Fluidity: Up to this point, our analysis has assumed that spawned agents are policy clones of their parent, reflecting parameter sharing among all agents. We now relax this assumption and consider the case in which agents possess distinct policies. This raises several natural questions. Can an agent learn not only when to spawn, but also how to coordinate effectively with a heterogeneous teammate? Can it dynamically alternate between fluid and fixed strategies depending on context? More broadly, does fluidity lead to qualitatively different behaviors compared to fixed-agent settings? To investigate these questions, we conduct experiments in the PuddleBridge environment, simulating agent independence by disabling parameter sharing and analyzing the resulting behavior of VDN algorithm. The spawn and step costs are fixed at $c_{\text{spawn}} = 1.0$ and $c_{\text{step}} = 0.1$.

6.2 Evaluation

For all algorithmic comparisons, we report the normalized, undiscounted joint return unless specified otherwise. This is done by computing the joint return over 1000 test episodes evaluated at 100 uniformly spaced checkpoints during training and normalizing using the following formula: $R_{\text{norm}} = (R - R_{\text{min}}) / (R_{\text{max}} - R_{\text{min}})$ where R denotes the observed joint return, and R_{min} and R_{max} are the minimum and maximum returns across all algorithms and timesteps, respectively. All results are averaged over 5 seeds, with shaded area denoting one standard deviation.

6.3 Results

6.3.1 Algorithmic Comparison: To compare the nature of different MARL algorithms in the fluid-agent setting, we run experiments with a 21×21 grid with 42 preys, $\text{prey_capture_reward } P = 5$ and $c_{\text{step}} = 0.01$. For all experiments on Predator-Prey, the fluid group of agents begins with 2 agents and is restricted to a maximum population of 10 agents.

- **SIP:** We first consider the size-inverse payoff, with $c_{\text{spawn}} = 10$. Fig. 2(a) displays the normalized joint return and the agents alive at the end of each test episode. As expected, under a size-inverse payoff, per-agent optimization algorithms are discouraged from spawning, since each additional agent sharply reduces the per-agent return—even when doing so could improve the joint return. In contrast, value decomposition methods such as VDN, which optimize directly for the joint return, perform well in this setting. Notably, VDN does not spawn excessively; instead, it effectively balances the cost–benefit trade-off of adding agents and converges to spawning approximately six agents.
- **SCP:** For the size-constant payoff, we counterbalance the clear incentive to spawn—since any agent capturing a prey yields a shared reward P for all—by setting a relatively high spawn cost of $c_{\text{spawn}} = 50$. Fig. 2(b) displays the normalized joint return and the agents alive at the end of each test episode. Despite this penalty, the optimal per-agent return still corresponds to spawning the maximum number of agents. Empirically, we observe that IQL and PPO fail to reach this optimum, whereas CTDE-based methods like VDN and MAPPO achieve it with ease.

6.3.2 Ablation: To perform this fluidity ablation, we use a 25×25 grid and evaluate fixed groups of 2, 4, 6, 8, and 10 agents alongside a fluid group, all trained using VDN. We consider 2 agents as the baseline population, such that spawn cost is incurred only when the population exceeds this baseline. Thus, for fixed groups of n agents, the spawn cost $c_{\text{spawn}} \times (n - 2)$, is applied *post hoc*. We maintain the $\text{prey_capture_reward } P = 5$, $c_{\text{step}} = 0.01$, $c_{\text{spawn}} = 10$ and employ the size-inverse payoff for fluid agents. Preys are sampled from a multinomial distribution of $\{20, 40, 60, 80\}$ each episode. For each prey population, we measure the joint return over 1000 post-convergence episodes for all agent groups.

Fig. 3(a) compares the joint returns achieved by fixed-size predator groups (2–10 agents) and a fluid population (“V”) across varying levels of prey availability. Returns are scaled by the maximum




Figure 4: Algorithmic comparison on the LBF environment.

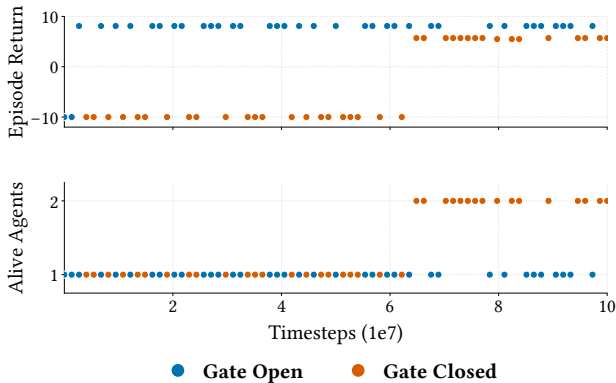


Figure 5: Episode return (top) and alive agents (bottom) during training in PuddleBridge. Points are colored by gate condition: open (blue) or closed (orange).

fixed-group performance to highlight relative efficiency. At low prey counts, small fixed groups (2–4 agents) achieve near-optimal returns due to minimal spawn and step costs. However, as prey abundance increases, their performance saturates, failing to capitalize on the additional resources. Larger fixed groups show the opposite trend: they underperform in low-resource settings due to higher per-step and spawn costs, but improve with prey richness. In contrast, the fluid group consistently adapts its population size to the resource level and performs competitively with the best fixed group across prey densities. Fig. 3(b) depicts the relationship between initial prey availability and the population size adopted by the fluid agents. Each hexagon represents the density of episodes, while the red dashed line traces the average equilibrium population attained across varying prey counts. This demonstrates that agents collectively learn to regulate their population size in response to environmental abundance—expanding when resources are plentiful and refraining from spawning when scarce.

6.3.3 *Level-Based Foraging*: Figure 4 shows the normalized joint return and the number of level-1 and level-2 agents spawned at the end of each episode for different MARL algorithms. In the optimal configuration, agents should spawn exactly one additional agent, with the level-2 agent initiating the spawn so that the new

teammate also has level 2. This setup enables capturing level-5 food with only three agents while minimizing spawn cost. From the plot of level-2 agents spawned, it is evident that PPO, MAPPO and VDN all learn the optimal strategy—spawning exactly one additional level-2 agent—while IQL approaches similar behavior with slightly higher variance. This confirms that fluid agents can optimize not only population size, but also the type of agent spawned, enabling precise and role-aware spawning strategies.

6.3.4 *PuddleBridge*: Figure 5 presents 75 equidistantly sampled training checkpoints, reporting the corresponding unnormalized joint episode return (top) and the number of agents alive at the end of each episode (bottom), with points colored by gate condition (open: blue; closed: orange). The open-gate task is simple and quickly solved, as indicated by positive single-agent episodes (blue). In contrast, the closed-gate scenario requires both spawning and coordination, making it substantially harder to master. Moreover, while two agents are required to achieve a positive return when the gate is closed, the accumulated spawn and step costs render this policy suboptimal when the gate is open. The agents not only learn the fluid coordination policy required for the closed-gate case, but also correctly distinguish between the two regimes—evidenced by the absence of blue points near two agents—switching between fluid and non-fluid strategies to maximize return.

7 CONCLUSION AND FUTURE WORK

We propose Fluid-Agent Reinforcement Learning as a framework for modeling real-world interactions where the assumption of a fixed number of agents is overly restrictive. In this setting, a team of agents can dynamically expand its population through a spawning mechanism integrated into the environment. We evaluate the performance of different MARL algorithms under various reward structures and show that the fluid-agent setting enables teams to dynamically adjust their population size in response to task variability, optimize the characteristics of spawned agents and also discover richer policies that are inaccessible in fixed-agent settings.

This work primarily introduces the fluid-agent framework and focuses on the aspect of agent creation. Future directions include studying the more general case where environments support agent creation and death, as well as investigating the theoretical properties of fluid systems—such as conditions for equilibrium convergence and sample complexity under fluid dynamics.

REFERENCES

- [1] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press. <https://www.marl-book.com>
- [2] Stefano V. Albrecht and Subramanian Ramamoorthy. 2015. A Game-Theoretic Model and Best-Response Learning Method for Ad Hoc Coordination in Multi-agent Systems. arXiv:1506.01170 [cs.GT] <https://arxiv.org/abs/1506.01170>
- [3] Samuel Barrett, Peter Stone, and Sarit Kraus. 2011. Empirical evaluation of ad hoc teamwork in the pursuit domain. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2 (Taipei, Taiwan) (AAMAS '11)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 567–574.
- [4] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakob W. Pachocki, Michael Petrov, Henrique Pond'e de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. 2019. Dota 2 with Large Scale Deep Reinforcement Learning. ArXiv abs/1912.06680 (2019).
- [5] Craig Boutilier. 1996. Planning, learning and coordination in multiagent decision processes. In *Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge (The Netherlands) (TARK '96)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 195–210.
- [6] Lucian Busoni, Robert Babuska, and Bart De Schutter. 2008. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38, 2 (2008), 156–172.
- [7] Georgios Chalkiadakis and Craig Boutilier. 2003. Coordination in multiagent reinforcement learning: a Bayesian approach. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (Melbourne, Australia) (AAMAS '03)*. Association for Computing Machinery, New York, NY, USA, 709–716. <https://doi.org/10.1145/860575.860689>
- [8] A. M. Fink. 1964. Equilibrium in a stochastic n -person game. *Journal of Science of the Hiroshima University, Series A-1 (Mathematics)* 28, 1 (1964), 89 – 93. <https://doi.org/10.32917/hmj/1206139508>
- [9] Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. *Advances in neural information processing systems* 29 (2016).
- [10] Jayesh K. Gupta, Maxim Egorov, and Mykel J. Kochenderfer. 2017. Cooperative Multi-agent Control Using Deep Reinforcement Learning. In *AAMAS Workshops*. <https://api.semanticscholar.org/CorpusID:9421360>
- [11] Anurag Koul. 2019. ma-gym: Collection of multi-agent environments based on OpenAI gym. <https://github.com/koulanurag/ma-gym>.
- [12] Landon Kraemer and Bikramjit Banerjee. 2016. Multi-agent reinforcement learning as a rehearsal for decentralized planning. *Neurocomputing* 190 (2016), 82–94. <https://doi.org/10.1016/j.neucom.2016.01.031>
- [13] H. W. Kuhn. 1953. *11. Extensive Games and the Problem of Information*. Princeton University Press, Princeton, 193–216. <https://doi.org/doi:10.1515/9781400881970-012>
- [14] Martin Lauer and Martin Riedmiller. 2000. An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In *In Proceedings of the Seventeenth International Conference on Machine Learning*. Citeseer.
- [15] Angeliki Lazaridou, Alexander Pearsakhovich, and Marco Baroni. 2017. Multi-Agent Cooperation and the Emergence of (Natural) Language. arXiv:1612.07182 [cs.CL] <https://arxiv.org/abs/1612.07182>
- [16] Laëtitia Matignon, Guillaume J Laurent, and Nadine Le Fort-Piat. 2007. Hysteretic q-learning: an algorithm for decentralized reinforcement learning in cooperative multi-agent teams. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 64–69.
- [17] Reuth Mirsky, Ignacio Carlucho, Arrasy Rahman, Elliot Fosong, William Macke, Mohan Sridharan, Peter Stone, and Stefano V. Albrecht. 2022. A Survey of Ad Hoc Teamwork Research. In *Multi-Agent Systems: 19th European Conference, EUMAS 2022, Düsseldorf, Germany, September 14–16, 2022, Proceedings* (Düsseldorf, Germany). Springer-Verlag, Berlin, Heidelberg, 275–293. https://doi.org/10.1007/978-3-031-20614-6_16
- [18] Igor Mordatch and Pieter Abbeel. 2018. Emergence of Grounded Compositional Language in Multi-Agent Populations. arXiv:1703.04908 [cs.AI] <https://arxiv.org/abs/1703.04908>
- [19] John Nash. 1951. Non-Cooperative Games. *Annals of Mathematics* 54, 2 (1951), 286–295. <http://www.jstor.org/stable/1969529>
- [20] John F. Nash. 1950. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences* 36, 1 (1950), 48–49. <https://doi.org/10.1073/pnas.36.1.48> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.36.1.48>
- [21] F. A. Oliehoek, M. T. J. Spaan, and N. Vlassis. 2008. Optimal and Approximate Q-value Functions for Decentralized POMDPs. *Journal of Artificial Intelligence Research* 32 (May 2008), 289–353. <https://doi.org/10.1613/jair.2447>
- [22] Georgios Papoudakis, Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. 2021. Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS)*
- [23] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of Machine Learning Research* 21, 178 (2020), 1–51.
- [24] R. Selten. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory* 4, 1 (March 1975), 25–55. <https://doi.org/10.1007/BF01766400>
- [25] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* 529, 7587 (Jan. 2016), 484–489. <https://doi.org/10.1038/nature16961>
- [26] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharrshan Kumaran, Thore Graepel, Timothy P. Lillicrap, Karen Simonyan, and Demis Hassabis. 2017. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. CoRR abs/1712.01815 (2017). arXiv:1712.01815 <http://arxiv.org/abs/1712.01815>
- [27] Farzan Soleymani and Eric Paquet. 2021. Deep graph convolutional reinforcement learning for financial portfolio management—DeepPocket. *Expert Systems with Applications* 182 (2021), 115127.
- [28] Peter Stone, Gal A. Kaminka, Sarit Kraus, and Jeffrey S. Rosenschein. 2010. Ad hoc autonomous agent teams: collaboration without pre-coordination. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (Atlanta, Georgia) (AAAI'10)*. AAAI Press, 1504–1509.
- [29] Sainbayar Sukhbaatar, Rob Fergus, et al. 2016. Learning multiagent communication with backpropagation. *Advances in neural information processing systems* 29 (2016).
- [30] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z. Leibo, Karl Tuyls, and Thore Graepel. 2017. Value-Decomposition Networks For Cooperative Multi-Agent Learning. arXiv:1706.05296 [cs.AI] <https://arxiv.org/abs/1706.05296>
- [31] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA.
- [32] Ming Tan. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*. 330–337.
- [33] Chen Tang, Ben Abbatematteo, Jiaheng Hu, Rohan Chandra, Roberto Martin-Martin, and Peter Stone. 2024. Deep Reinforcement Learning for Robotics: A Survey of Real-World Successes. arXiv:2408.03539 [cs.RO] <https://arxiv.org/abs/2408.03539>
- [34] Peter D. Taylor and Leo B. Jonker. 1978. Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences* 40, 1–2 (1978), 145–156.
- [35] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *nature* 575, 7782 (2019), 350–354.
- [36] Jianhong Wang, Yang Li, Yuan Zhang, Wei Pan, and Samuel Kaski. 2024. Open Ad Hoc Teamwork with Cooperative Game Theory. arXiv:2402.15259 [cs.MA] <https://arxiv.org/abs/2402.15259>
- [37] Jörgen W. Weibull. 1995. *Evolutionary Game Theory*. MIT Press.
- [38] Chao Yu, Jiming Liu, and Shamim Nemati. 2020. Reinforcement Learning in Healthcare: A Survey. arXiv:1908.08796 [cs.LG] <https://arxiv.org/abs/1908.08796>
- [39] Chao Yu, Akash Velu, Eugene Vinyals, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in neural information processing systems* 35 (2022), 24611–24624.