
Off-Beat Multi-Agent Reinforcement Learning

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We investigate model-free multi-agent reinforcement learning (MARL) in environ-
2 ments where *off-beat* actions are prevalent, *i.e.*, all actions have pre-set execution
3 durations. During execution durations, the environment changes are influenced
4 by, but not synchronised with, action execution. Such a setting is ubiquitous in
5 many real-world problems. However, most MARL methods assume actions are
6 executed immediately after inference, which is often unrealistic and can lead to
7 catastrophic failure for multi-agent coordination with off-beat actions. In order to
8 fill this gap, we develop an algorithmic framework for MARL with off-beat actions.
9 We then propose a novel episodic memory, LeGEM, for model-free MARL
10 algorithms. LeGEM builds agents' episodic memories by utilizing agents' indi-
11 vidual experiences. It boosts multi-agent learning by addressing the challenging
12 temporal credit assignment problem raised by the off-beat actions via our novel
13 reward redistribution scheme, alleviating the issue of non-Markovian reward. We
14 evaluate LeGEM on various multi-agent scenarios with off-beat actions, including
15 Stag-Hunter Game, Quarry Game, Afforestation Game, and StarCraft II microman-
16 agement tasks. Empirical results show that LeGEM significantly boosts multi-agent
17 coordination and achieves leading performance and improved sample efficiency.

18 1 Introduction

19 In Multi-Agent Reinforcement Learning (MARL), multiple agents act interactively and complete tasks
20 in a sequential decision-making manner with Reinforcement Learning (RL). It has made remarkable
21 advances in many domains, including autonomous systems [8, 19, 72] and real-time strategy (RTS)
22 video games [58]. By the virtue of the *centralised training with decentralised execution* (CTDE) [33]
23 paradigm, which aims to tackle the scalability and partial observability challenges in MARL, many
24 CTDE-based MARL methods are proposed [13, 49, 41, 62, 47, 63, 23, 35]. With these methods, an
25 agent executes actions only via feeding its individual observations independently and optimizes its
26 own policy with access to global trajectories centrally.

27 Despite the recent successes of MARL, learning effective multi-agent coordination policies for
28 complex multi-agent systems remains challenging. One key challenge is the *off-beat* actions, *i.e.*, all
29 actions have pre-set execution durations¹ and during the execution durations, the environment changes
30 are influenced by, but not synchronised with, action execution (an illustrative scenario is shown in
31 Fig. 1). However, Dec-POMDP [32], which underpins many CTDE-based MARL methods, hinges
32 on the assumption that actions are executed momentarily after inference, leading to catastrophic
33 failure for *centralized training* on various off-beat multi-agent scenarios (OBMAS). To fill this gap,
34 we study MARL in settings where off-beat actions are prevalent. Such setting is very common in
35 many real-world problems. For example, in the traffic light control problem, traffic lights in the
36 conjunctions of the road network have pre-set execution time which is set asynchronously.

¹In the RL literature [39, 6], action execution durations are called *delays of actions*. In this paper, we use the term *execution durations*, which is self-consistent with off-beat actions defined in Sec. 3.




Figure 1: **An illustrative scenario:** two-agent stag-hunter game, where two agents (hunters) have only partial observations, different durations of the shoot action, and cannot communicate. The goal is to catch the stag and they are rewarded when their shot hits – as in, completion of the action is synchronised, the stag at the same time. Both agents can see the stag. As the shoot action durations of the two agents are different, to catch the stag, the two agents should shoot the arrow at different timesteps given the distances. Though the scenario is easy for human beings, it is hard for MARL agents due to the action duration. **Experiment results:** in this scenario, the optimal policy for agent 0 is to shoot the arrow at timestep 0 while the optimal policy for agent 1 is to shoot the arrow at timestep 5. Such asynchronous property of OBMA motivates agents to learn tacit policies. The curves show that VDN and IQL fail to learn coordination policies even in this simple scenario. With LeGEM, MARL methods gain enhanced performances as well as improved sample efficiency.

37 The problem of off-beat actions in MARL has yet to be investigated and tackled. Training MARL
 38 policies in OBMA is challenging: (i) Each agent’s actions can have a variety of execution durations,
 39 which augments the order of complexity of OBMA during decentralized execution, resulting
 40 in failure of the coordination; (ii) The action durations are unknown to agents during individual
 41 executions, and communication is constrained and not always feasible, making it non-trivial to model
 42 the environment; (iii) During training, both the temporal credit assignment with TD-learning [51]
 43 and the *inter-agent* credit assignment with value decomposition methods [41] cannot perform well
 44 due to the displaced rewards in multi-agent replay. With off-beat actions, the nonstationarity issue,
 45 which mainly stems from rewards’ time dependency on the agents’ past actions, is exacerbated.

46 While actions durations are ubiquitous, existing works only focus on single-agent settings, *i.e.*, delay,
 47 in RL. Many approaches [59, 39, 66] augment the state space with the queuing actions to be executed
 48 into the environment. However, such state-augmentation trick leads to exponentially increasing
 49 training samples with the growing action duration, making training intractable [11]. Chen et al. [10]
 50 extend the delayed MDP [39] and propose Delayed Markov Game for MARL. However, on one
 51 hand, such state-augmentation treatment is confined to short delays, *e.g.*, one timestep delay; on
 52 the other hand, the delayed timestep of the actions is privileged information, which is not available
 53 in many scenarios. Recent works on macro-actions [67, 68] introduce asynchronous actions by
 54 designing macro-actions with prior environment knowledge. Macro-actions are different from options
 55 in hierarchical RL (HRL) [52, 3] in that the later is not manually designed but learned. The key
 56 difference between macro-actions and off-beat actions is that macro-actions are high-level actions
 57 while off-beat actions are primitive actions. Unfortunately, the *inter-agent* credit assignment is still
 58 a challenge of HRL in OBMA and the asynchronous² nature of off-beat actions undermines the
 59 temporal credit assignment of *centralized training*, causing poor sample efficiency and unsatisfactory
 60 performance (more discussions can be found in the related works section in Sec. 7).

61 We aim to address the aforementioned issues. We first propose off-beat Dec-POMDP. We then instan-
 62 tiate a new class of episodic memory, LeGEM, for model-free MARL algorithms. LeGEM boosts
 63 multi-agent learning by addressing the challenging temporal credit assignment problem raised by the
 64 off-beat actions via our novel levelled graph-based temporal recency reward redistribution scheme.
 65 Specifically, each agent maintains LeGEM and during centralized training, each agent searches the
 66 pivot timestep given observations from its graph. The pivot timestep is the timestep wherein the off-
 67 beat reward relates to the given node. The pivot timesteps of each agent are ranked, in which the final
 68 pivot timestep will be chosen by recency and later used for reward redistribution and target estimation
 69 in TD-learning. We evaluate our method on Stag-Hunter Game, Quarry Game, Afforestation Game,
 70 and StarCraft II micromanagement tasks. Empirical results show that our method significantly boosts
 71 multi-agent coordination and achieves leading performance as well as improved sample efficiency.

²We clarify the term *asynchronous*: actions that simultaneously committed into the environment by all agents in MARL will not complete their respective action durations at the same time in future timesteps.

72 2 Preliminaries

73 **Dec-POMDP.** A cooperative MARL problem can be modeled as a *decentralised partially observable*
 74 *Markov decision process* (Dec-POMDP) which can be formulated as a tuple $\langle \mathcal{S}, \mathcal{U}, \mathcal{P}, R, O, \mathcal{N}, \gamma \rangle$,
 75 where $\mathbf{s} \in \mathcal{S}$ denotes the state of the environment. Each agent $i \in \mathcal{N} := \{1, \dots, N\}$ chooses an
 76 action $u_i \in \mathcal{U}$ at each timestep, forming a joint action vector, $\mathbf{u} := [u_i]_{i=1}^N \in \mathcal{U}^N$. The Markovian
 77 transition function can be defined as $\mathcal{P}(\mathbf{s}' | \mathbf{s}, \mathbf{u}) : \mathcal{S} \times \mathcal{U}^N \times \mathcal{S} \mapsto [0, 1]$, transiting one state of current
 78 timestep to the state of next timestep conditioned on current state and joint action. Every agent shares
 79 the reward and the reward function is $R(\mathbf{s}, \mathbf{u}) : \mathcal{S} \times \mathcal{U}^N \mapsto \mathcal{R}$. $\gamma \in [0, 1)$ is the discount factor.
 80 Due to *partial observability*, each agent has individual partial observation $o \in \mathcal{O}$, according to the
 81 observation function $O(\mathbf{s}, i) : \mathcal{S} \times \mathcal{N} \mapsto \mathcal{O}$. The goal of each agent is to optimize its own policy
 82 $\pi_i(u_i | \tau_i) : \mathcal{T} \times \mathcal{U} \mapsto [0, 1]$ given its action-observation-reward history $\tau_i \in \mathcal{T} := (\mathcal{O} \times \mathcal{U})$.

83 **Multi-Agent Reinforcement Learning.** MARL aims to learn optimal policies for all the agents
 84 in the team. With TD-learning and a global Q value proxy Q^{tot} for the optimal Q^* , $\{Q_i\}_{i=1}^N$ are
 85 optimized via minimizing the loss [65, 31]: $\theta^* = \arg \min_{\theta} \mathcal{L}(\theta) := \mathbb{E}_{D' \sim \mathcal{D}} [(y_t^{\text{tot}} - Q_{\theta}^{\text{tot}}(\mathbf{s}_t, \mathbf{u}_t))^2]$,
 86 where $y_t^{\text{tot}} = r_t + \gamma \max_{\mathbf{u}'} Q_{\theta}^{\text{tot}}(\mathbf{s}_{t+1}, \mathbf{u}')$ and θ is the parameters of the agents. $\tilde{\theta}$ is the parameter
 87 of the target Q^{tot} and is periodically copied from θ . D' is a sample from the replay buffer \mathcal{D} .

88 3 Off-Beat Dec-POMDP

89 We introduce our formulation for OBMAS. We first define the off-beat actions³ for multi-agent
 90 scenarios; then we propose the Off-Beat Dec-POMDP. All the proofs can be found in Appx. A.

Definition 1 (Off-Beat Actions). Off-beat action $\tilde{u} \in \mathcal{U}$ characterizes OBMAS where the action
 \tilde{u}_i taken by agent i has execution duration $m_{\tilde{u}_i} \sim A(m | \tilde{u}_i, i)$, $A \in \mathcal{A}$, $m \in \{0, 1, 2, \dots, M\}$
 and $M \leq T$, where T is the maximum duration and A is the action duration distribution. It is a
 distribution and takes \tilde{u}_i and the index of the agent as parameters. A can be either stochastic or
 deterministic. The joint off-beat action is $\tilde{\mathbf{u}} = [\tilde{u}_i]_{i=1}^N$. The execution duration is decided at the time
 the action was committed to the environment. Thus, the execution duration of an action $\tilde{\mathbf{u}}_t$ initiated
 at timestep t is $\mathbf{m}_t = \{m_{\tilde{u}_i}^t\}_{i=1}^N$.

91

92 Note that for each agent, $m_{\tilde{u}_i}^t$ ⁴ can be different. At timestep t , there are at least 1 action⁵ and at most
 93 N actions being initiated (committed to the environment for execution), leading to asynchronicity of
 94 the joint actions. Next, we propose the Off-Beat Dec-POMDP for OBMAS and discuss its properties.

Definition 2 (Off-Beat Dec-POMDP). Off-Beat Dec-POMDP extends Dec-POMDP, such that
 (1) state space is \mathcal{S} ; (2) joint action space is \mathcal{U}^N ; (3) action duration space is \mathcal{A}^N ;
 (4) transition function is $\mathcal{P}(\mathbf{s}' | \mathbf{s}, \tilde{\mathbf{u}}, \mathbf{m}) : \mathcal{S} \times \mathcal{U}^N \times \mathcal{S} \times \mathcal{A}^N \mapsto [0, 1]$, and \mathbf{m} is the action durations
 of the joint action;
 (5) the reward function is $R(\mathbf{s}, \tilde{\mathbf{u}}, \mathbf{m}) : \mathcal{S} \times \mathcal{U}^N \times \mathcal{A}^N \mapsto \mathcal{R}$;
 (6) we call a reward r as off-beat reward when any its $m_{\tilde{u}_i} \geq 1$, $m_{\tilde{u}_i} \in \mathbf{m}$, and $r \neq 0$.

95

96 In OBMAS, at each timestep t , the environment receives actions that agent initiates for execution
 97 in the environment. The initiated actions $\tilde{\mathbf{u}}_t$ are instantaneous actions inferred by agents' policies
 98 given individuals' observations. The joint reward is the consequence of the committed joint actions
 99 of current timestep and previous timesteps, depending on the actions' duration. The asynchronicity
 100 is an inherent feature of the environment, which is different from asynchronicity incurred by com-
 101 munication delays in many video games (asynchronous gameplay⁶). We discuss some properties of
 102 Off-Beat Dec-POMDP below.

³Asynchronicity is prevalent in real-world multi-agent scenarios, including asynchronicity in observations,
 actions and communication, etc. In this paper, we focus on the asynchronicity of actions in multi-agent scenarios.
 For brevity, we name the asynchronicity of actions in MARL as *off-beat*.

⁴We will omit t in the rest of the paper for brevity.

⁵We note that agents have a special NO-OP action available.

⁶<https://www.whatgamesare.com/2011/08/synchronous-or-asynchronous-definitions.html>

103 **Remark 1.** When the durations for all actions are identical, off-beat Dec-POMDP reduces to
 104 Delayed Dec-POMDP and there is no off-beat actions in it.

105 **Remark 2.** There exists \tilde{u} that is synchronous since duration of agents' actions can be $m = 0$.
 106 When m of all actions is zero, off-beat Dec-POMDP reduces to Dec-POMDP.

107 In Delayed Dec-POMDP, actions have the same delayed timesteps, which is different from off-
 108 beat actions where actions have different action durations or delays. In order to investigate the
 109 problem, we consider the deterministic setting of the transition function and the reward function.

110 **Remark 3** (Non-episodic Reward). In our formulation, the reward is not episodic reward [16].

111 **Remark 4** (Non-Markovian Reward). With off-beat actions, the Markovian property of the reward
 112 function $R(s, \tilde{u}, m)$ does not hold.

113 With off-beat actions, the shared rewards can be readily displaced, causing non-Markovian rewards.
 114 Solving Off-Beat Dec-POMDP is challenging as discussed in Sec. 1. We propose our methods to
 115 tackle aforementioned challenges.

116 4 The Journey is the Reward: A Collective Mental Time Travel Method

117 We propose two methodological elements for Off-Beat MARL. The first, LeGEM, presented in
 118 this section, is a form of episodic memory that facilitates discovery of a pivotal timestep for off-
 119 beat rewards; and the second, presented in Sec. 5, is redistribution of the off-beat reward to the pivot
 120 timestep when the relevant off-beat actions were initialised.

121 4.1 LeGEM: A Levelled Graph Episodic Memory for Off-Beat MARL

122 Human learning relies on retrospecting our detailed memory of the past [55, 48]. For example,
 123 while exploring a new scenic area, we do not just remember a multitude of specific spots there,
 124 but can recall the paths that connect them with junctions and turns. However, there is no MARL
 125 method that can explicitly recall the past and identify key states that lead to future rewards. Such
 126 “mental time travel” [24] ability is vital for tackling the challenges in OBMAS. Inspired by the recent
 127 progress in RL with episodic memory [18, 5, 17] that is based on the memory prosthesis proposed
 128 by neuroscientists [55, 48], we propose our method of episodic memory representation for MARL.
 129 Unlike previous episodic memory methods that train a parameterized memory by either augmenting
 130 the policy inputs for execution [18] or regularizing the TD learning [17] for RL, our method utilizes
 131 the levelled graph data structure [4], a well established structure for data storage and retrieval, to
 132 represent an agent’s individual episodic memory.

133 We propose our novel episodic memory, Levelled Graph Episodic
 134 Memory method (LeGEM), via the levelled graph data structure.
 135 LeGEM memorizes each agent’s past trajectories which are partial
 136 observations and the unilateral action of the agent. During training,
 137 each agent i collects its individual trajectories τ_i . We then define τ_i
 138 of agent i as $\tau_i = [(o_i^0, \tilde{u}_i^0, r^0), \dots, (o_i^{T-1}, \tilde{u}_i^{T-1}, r^{T-1})]$, where T is
 139 the length of the trajectory and the triplet $(o_i^t, \tilde{u}_i^t, r^t)$ represents the
 140 observation, action and reward of timestep t . Note that r^t is globally
 141 shared between agents. We define agent i ’s LeGEM as a directed graph
 142 $\phi_i^t \in \Phi_i$ where Φ_i is the set of graphs of agent i and ϕ_i^t is the t -th graph
 143 of Φ_i , $t \in \{0, \dots, T-1\}$. Each ϕ_i^t consists of a tuple of (Ψ, Ξ) where
 144 Ψ is the set of nodes and Ξ represents the set of edges that connect
 145 nodes in the graph. To model an agent’s behaviour explicitly and make
 146 the trajectories of agents easy to represent, we create T graphs for each
 147 agent and let $\Phi_i = \{\phi_i^t\}_{t=0}^{T-1}$ where T is the maximum level of all graphs and the maximum length
 148 of the episode as well. The maximum level of ϕ_i^t is $t + 1$. The node contains key, visit count and
 149 pointers connecting the precursors (node at the previous level) and the successors (node at the next
 150 level). Unlike many parameterized episodic memory using state/observation as the key [18, 24], we
 151 resort to *afterstate* [36]. That is, we use agent i ’s observation o_i^t and action at timestep t , \tilde{u}_i^t , to define
 152 the key (o_i^t, \tilde{u}_i^t) . We provide an example to showcase the relationship between sub-graph and the
 153 graph in Fig. 2. For complex and continuous state scenarios, for example StarCraft II scenarios, we




Figure 2: The maximum level of the graph is 7. Circles indicate the nodes and numbers indicate the visit count.

Algorithm 1: SearchPivotTimesteps (ρ)

```
1 Input:  $\tau, \Phi, \Upsilon$  and Search (scheme I or II);
2 Initialize:  $\kappa$ : an empty list to store pivot timesteps;
   // Length of  $\tau$  and  $\tau_i$  are equal.
3  $l \leftarrow \text{length}(\tau_i) - 1$ ;
4 for  $t \leftarrow 0$  to  $\text{length}(\tau) - 1$  do
5   if  $r^t \neq 0$  ( $r^t \in \tau$ ) then
6     // Off-beat reward
7     for  $i \leftarrow 1$  to  $N$  do
8       Get  $\tau_i$  from  $\tau$ ;
9        $\phi_i^l \leftarrow \Phi_i[l]$ ;
10       $\psi \leftarrow \phi_i^l \cdot \text{getNode}(o_i^t, \tilde{u}_i^t)$ ;
11      Find all the paths  $\Lambda_i^{t,l}$  from node  $\psi$  to
        the node at level 0;
12      Get the discretized episode return  $r^{l,i}$ ;
13      Get the index  $\omega$  from  $\Upsilon$  with  $r^{l,i}$ ;
14       $e_i^t \leftarrow \text{Search}(\omega, \Lambda_i^t, \tau_i, r^{l,i}, \Upsilon, \Phi_i)$ ;
15      Get  $e_t$  (Eqn. 1) and append  $e_t$  to  $\kappa$ ;
16 Return:  $\kappa$ .
```

Algorithm 2: Search Scheme I

```
1 Input:  $\omega, \Lambda_i^{t,l}, \tau_i, r^{l,i}, \Upsilon$  and  $\Phi_i$ ;
2 Initialize:  $e_t^i$ : a list whose values are all  $t$ 
   and its size is the number of paths in  $\Lambda_i^{t,l}$ ;
3  $\phi_i^{l,\omega} \leftarrow \Phi_i^{l,\Omega}[\omega]$ ;
4  $\text{vc} \leftarrow \text{VisitCount}(\Lambda_i^{t,l})$  (Alg. 4);
5 foreach path  $\Lambda_i^{t,l}[j] \in \Lambda_i^{t,l}$  do
6    $e_t^{i,j,\downarrow} \leftarrow \text{UL}(\Lambda_i^{t,l}[j], \text{vc}, \tau_i)$  (Alg. 5);
7    $e_t^{i,j,\uparrow} \leftarrow \text{LU}(\Lambda_i^{t,l}[j], \text{vc}, \tau_i)$  (Alg. 6);
8   if  $e_t^{i,j,\downarrow} \neq -1$  then
9      $e_t^i[j] \leftarrow e_t^{i,j,\downarrow}$ ;
10  else if  $e_t^{i,j,\uparrow} \neq -1$  then
11     $e_t^i[j] \leftarrow e_t^{i,j,\uparrow}$ ;
12  else
13     $e_t^i[j] \leftarrow t$ ;
14  $e_t^i \leftarrow \text{Summarize}(e_t^i)$  (Alg. 7);
15 Return:  $e_t^i$ .
```

154 use SimHash [9] to discretize the key (o_i^t, \tilde{u}_i^t) . This technique has been widely used in commercial
155 search engines and RL [54]. Visit count indicates the total visits made by agent i to the node. Its initial
156 value is 1. Note that nodes are bidirectional since it is helpful for searching (see Sec. 4.2).

157 Given a τ_i with the length of T , if the node is already in the graph at level t , we then increase the
158 visit count by 1. Otherwise, we create a new node for level t of the graph and update its pointers.
159 Meanwhile, sub-graphs will be also created and updated. The process of updating LeGEM is in Alg.
160 3. We provide an example of Alg. 3 in Fig. 9, Appx. B.1. It is worth noting that τ_i is generated via
161 the interaction of the agent with the environment, and there is no extra interaction needed to collect τ_i .
162 The generated trajectories are saved in the experience replay and later sampled for MARL training.

163 4.2 Multi-Agent Collective Mental Time Travel with LeGEM

164 With structured agent's past experiences, it can be used to search the pivot timestep when actions
165 that triggered the rewarded state were executed. For example, with LeGEM, we can find the pivot
166 timestep, $e_t = 5$, when agent 1 shoots the arrow in Fig. 1.

167 **Fact 1.** (Action-Reward Association) *When an off-beat reward r_t exists in the trajectory τ_i ($i \in$
168 $\{1, \dots, N\}$), $r_t \in \tau_i$, off-beat action $u_{t'}$ exists ($t' < t$) in the trajectory set $\{\tau_j\}_{j=1}^N$, where $\{\tau_j\}_{j=1}^N$
169 constitutes the global trajectory of all agents.*

170 As the reward function and transition function are deterministic in our setting, Fact 1 holds. Intuitively,
171 once we find an off-beat reward in a trajectory, we are sure that the action which triggered the reward
172 can be found in the trajectory. With more experiences collected by the agents, such pattern is obvious
173 and significant. It motivates us to propose a method to leverage the association property of the
174 off-beat action-reward data and search the pivot timestep for timesteps when off-beat rewards occur,
175 which can further help to redistribute the reward backward to mitigate the temporal credit assignment
176 issue (c.f. Sec. 5). Therefore, we first propose a search method to search the pivot timestep and then
177 propose a proximal ranking method to estimate the pivot timestep that invokes the future reward.

178 **Collective Mental Time Travel.** The displaced rewards in the replay buffer hinder multi-agent
179 learning. It is essential for each agent to search the pivot timestep when the potential off-beat action
180 that triggered the rewarded state was committed to the environment. Therefore, we propose two
181 search schemes to find the pivot timestep for all agents given an off-beat reward.

182 *Scheme I:* For agent i , given $r_t \in \tau_i$, episode return $r^{l,i}$ of τ_i , $\phi_i^l = \Phi_i[l]$ and $\phi_i^{l,\omega} = \Phi_i^{l,\Omega}[\omega]$,
183 agent i searches from the node (the key is (o_i^t, \tilde{u}_i^t) and $o_i^t \in \tau_i, u_i^t \in \tau_i$) at level t in sub-graph $\phi_i^{l,\omega}$ to
184 find the pivot timestep e_t for r_t . Concretely, we propose our bi-directional search method. The first
185 one is called Low-Up (LU) search, which traverses from the given node at level t upwards to the node
186 at level 0. The second one is named Up-Low (UL) search which traverses from the node at level 0
187 downwards to the given node at level t . LU traversing ends when the pattern of increasing visit count

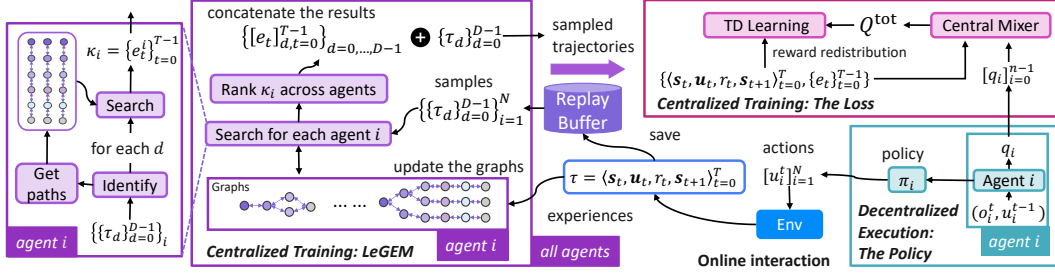


Figure 3: Our framework: LeGEM, the loss and the agent’s policy.

188 ends and the corresponding level is the candidate pivot timestep. On the contrary, UL traversing
 189 ends when the pattern of decreasing visit count ends and the corresponding level is the candidate
 190 pivot timestep. In Alg. 2, we first get visit count (Line 4) and then apply UL traversing (Line 6)
 191 and LU traversing (Line 7). We summarize the results (Line 14) by select the pivot timestep that
 192 has the maximum count. UL traversing has a higher searching priority than its counterpart. The
 193 reason is that there exists pattern that the visit count is decreasing from the node at level 0 and
 194 such pattern ends at the pivot timestep. In practise, it works well in scenarios whose trajectories are
 195 single-off-beat-reward trajectories (there is only one off-beat reward) and the accuracy of Scheme I is
 196 over 90% in grid world scenarios. For scenarios, especially complex scenarios, whose trajectories are
 197 multiple-off-beat-reward trajectories, we apply Scheme II. We put Alg. 4, Alg. 5, Alg. 6 and Alg. 7 in
 198 Appx. B.1 as these algorithms are intuitive and easy to understand literally. The time complexity is
 199 $\mathcal{O}(n \cdot m)$ (a slight notation abuse) where n is the size of each $\Lambda_i^{t,l}$ and m ($1 \leq m \leq n$) is the average
 200 distance between the level of the given node to the level of the node at the pivot timestep.

201 **Scheme II:** Scheme II is a simplified version of scheme I for scenarios that have multiple-off-beat-
 202 reward trajectories, which searches the pivot timestep by finding the nearest timestep in the most
 203 visited path. The node of the nearest timestep has the maximum visitcount in that path. Despite the
 204 simplicity, it works effective and the time complexity is $\mathcal{O}(n)$ where n is the number of paths in $\Lambda_i^{t,l}$.
 205 The pseudo code is shown in Alg. 8 in Appx. B.1.

206 Given a node at level t , agents collectively search from the node to find the pivot time step (Line 13
 207 in Alg. 1). The visit count is vital for search methods. In MARL, we use ϵ -greedy [31] for agents to
 208 explore the environment and collect individual trajectories. The collected trajectories will be used to
 209 build the memory and train the policy. We apply annealing to ϵ (in Appx. E).

210 **Ranking the Pivot Timesteps.** With our two search schemes, we can search the pivot timesteps
 211 for each global trajectory $\tau = \{(s^t, \tilde{u}^t, r^t, s^{t+1})\}_{t=0}^{T-1}$. We define the pivot timesteps κ of each
 212 global trajectory τ as $\kappa = \{e_t\}_{t=0}^{T-1}$, $0 \leq e_t \leq t$, where e_t indicates the pivot timestep of t when
 213 r_t is the consequence of actions committed before timestep t . We first get e_t by aggregating all the
 214 searching outcomes (Line 13 in Alg. 1). Then, each agent gets $\kappa_i = \{e_t^i\}_{t=0}^{T-1}$. In order to subserve
 215 the inter-agent credit assignment [13, 41], κ can be collectively calculated via proximity:

$$e_t = \min_{e_t^i} \left[t - e_t^1, \dots, t - e_t^N \right], i \in \{1, \dots, N\} \quad (1)$$

216 The pseudo code is shown in Alg. 1. For each sampled global trajectory τ , we extract τ_i for each
 217 agent in Line 7; then we get e_t for each agent and aggregate κ in line 14 and line 15, respectively.

218 5 Reward Redistribution for Off-Beat Multi-Agent Reinforcement Learning

219 Searching in LeGEM leverages the collective intelligence [25, 15] in OBMA. We utilize TD learning
 220 to train MARL policies. The TD error is the difference between the TD target and the prediction.
 221 TD targets can be estimated with n -step target, TD(λ) and other techniques [12, 56]. Unfortunately,
 222 current n -step target and TD(λ) methods are far from accurate estimating TD targets. They even
 223 incur underestimation with off-beat trajectories. In essence, to train MARL policies in OBMA, one
 224 should accurately estimate the TD target where the reward plays the key role [46, 70]. We resolve the
 225 aforementioned conundrum by redistributing rewards to their pivot timesteps. The key idea is that we
 226 can pull the outcome of one joint off-beat action back to the timestep when it was committed to the
 227 environment, which can dramatically enhance learning despite the long-term reward delays incurred
 228 by off-beat actions. We utilize e_t to update the reward of the transit $(s^{e_t}, \tilde{u}^{e_t}, r^{e_t}, s^{e_t+1})$:

$$\hat{r}^{e_t} = \mathbb{I}(e_t \geq t) \cdot r^{e_t} + \mathbb{I}(e_t < t) \cdot r_t, \quad (2)$$

229 where $\mathbb{1}(\cdot)$ is the indicator function. Such update rule is conducted iteratively from $t = 0$ to $t = T - 1$.
 230 β is a very small positive hyperparameter. To stabilize learning and circumvent the overestimation
 231 of the TD target, r_t is also updated after Eqn. 2 via $r_t = (1 - \mathbb{1}(e_t < t) \cdot (1 - \beta)) \cdot r_t$. It also
 232 avoids aggregated biased/wrong estimation of TD target being back propagated in Bellman Equation.
 233 Formally, we define the reward redistribution operator as Π_Φ , i.e., $e_t = \Pi_\Phi \rho(r^t, \mathbf{s}, \tilde{\mathbf{u}})$, and then
 234 define the Off-Beat Bellman operator Γ :

$$(\Gamma Q^{\text{tot}})(\mathbf{s}, \tilde{\mathbf{u}}) := \mathbb{E}[\Pi_\Phi R(\mathbf{s}, \tilde{\mathbf{u}}, \mathbf{m}) + \gamma \max_{\tilde{\mathbf{u}}'} Q^{\text{tot}}(\mathbf{s}', \tilde{\mathbf{u}}')] \quad (3)$$

235 With the Off-Beat Bellman operator Γ , we propose its contraction property.

236 **Proposition 1.** $\Gamma : \mathcal{Q} \mapsto \mathcal{Q}$ is a γ -contraction.

237 Therefore, we can utilize \hat{r}_{e_t} for *centralized training* in TD-learning:

$$\mathcal{L}^{\text{TD}}(\theta) := \mathbb{E}_{\mathcal{D}' \sim \mathcal{D}}[(\hat{y}_{e_t}^{\text{tot}} - Q_\theta^{\text{tot}}(\mathbf{s}^{e_t}, \tilde{\mathbf{u}}^{e_t}))^2], \text{ where } \hat{y}_{e_t}^{\text{tot}} = \hat{r}_{e_t} + \gamma \max_{\tilde{\mathbf{u}}'} Q_\theta^{\text{tot}}(\mathbf{s}^{e_{t+1}}, \tilde{\mathbf{u}}'). \quad (4)$$

238 Our method can be easily incorporated into any model-free MARL method for OBMAS. We present
 239 the pseudo code of incorporating our method into model-free MARL methods in Alg. 9, Appx. E.
 240 We also provide a pictorial view of our framework in Fig. 3 to show the whole pipeline.

241 6 Experiments

242 We perform experiments on various multi-agent scenarios with off-beat actions. We introduce
 243 off-beat actions in Stag-Hunter Game, Quarry Game, Afforestation Game and StarCraft II microman-
 244 agement tasks [44] and use them as testbeds in our experiments. We aim to answer the following
 245 questions: **Q1:** *Can our LeGEM improve the multi-agent coordination of many MARL methods*
 246 *in OBMAS?* **Q2:** *Can our LeGEM outperform previous parameterized episodic memory (EM) for*
 247 *MARL?* **Q3:** *Can bootstrapping method of RL help?* **Q4:** *Can our LeGEM outperform the multi-agent*
 248 *exploration and multi-agent risk-sensitive (Ex-Risk) methods?*




Figure 4: The test catch rate of the stag on the Stag-Hunter Game with off-beat actions.

249 6.1 Experiment Setup




Figure 5: Stag-Hunter Game, Quarry Game and Afforestation Game. More information can be found in Appx. C.

Categories	Methods
MARL (Q1)	QMIX [41], VDN [49] IQL [53], QTRAN [47] QPLEX [60]
EM (Q2)	EMC [71]
Bootstrap (Q3)	N-step & λ -Return [51]
Ex-Risk (Q4)	MAVEN [28], EMC [71] RMIX [38]

Table 1: Baseline algorithms.

251 **Baselines and scenarios.** We list all baselines in table 1, including the corresponding research
 252 questions to be answered. We implement our method on PyMARL [44] and use 10 random seeds
 253 to train each method on all environments. We do not use macro-action methods [67, 68] as the
 254 baseline because it is hard to make a fair comparison between macro-actions methods and our method.
 255 As discussed in Sec. 1, macro-actions rely on manually designed macro-actions, i.e., designing
 256 the macro-actions by utilizing the simulator settings and domain knowledge, which is different
 257 from learning options [52, 3]. Designing macro-actions is not feasible in scenarios where domain
 258 knowledge and simulator settings are unknown, such as the OBMAS scenarios. In OBMAS, the agent
 259 has no idea of the durations of other agents' actions, which is challenging for designing macro-actions.
 260 We conduct experiments on Stag-Hunter Game, Quarry Game, Afforestation Game (Fig. 5) and
 261 StarCraft II micromanagement tasks [44] where off-beat action are introduced.

262 **Training settings.** We use opensourced code of baselines publicly by the corresponding authors on
 263 Github in all experiments. We resort to mean-std values as our performance evaluation measurement
 264 in all figures where the bold line and the shaded area indicate the mean value and one standard
 265 deviation of the episode return, respectively. Readers can refer to Appx. C, D, E and F for more
 266 information on our environment, baselines, training method, training platform and empirical results.




Figure 6: The test task completion rate of the Quarry Game with off-beat actions.




Figure 7: Performance of MARL methods

267 6.2 Experiment Results

268 **The Effectiveness of LeGEM.** We answer **Q1**. With LeGEM, MARL methods get enhanced
 269 performance as shown in Fig. 4. Without LeGEM, all methods perform poorly in Stag-Hunter Game;
 270 IQL and VDN’s final final results are even 0. By incorporating LeGEM, all of them can get converged
 271 performance and improved sample efficiency. We are also interested in finding if LeGEM could
 272 reinforce the performance of simple methods. As depicted in Fig. 7, with LeGEM, both VDN and
 273 QMIX outperforms QPLEX, which is a state-of-the-art MARL method armed with various advanced
 274 techniques, including attention network [57], dueling network [64] and advantage function.

275 **Performance of Episodic Memory method.** We answer **Q2** by presenting the performance curves
 276 of EMC in Fig. 7. EMC is an episodic memory MARL method with curiosity-driven exploration. It
 277 utilizes the episodic memory from RL [74, 17]. With LeGEM, QMIX outperforms EMC. EMC even
 278 fails to converge in Stag-Hunter Game.

Table 2: Results (mean and std) of n -step return (left) and $TD(\lambda)$ (right) on Stag-Hunter Game.

n	1	5	10	15	λ	0.8	0.9	0.99	1
QMIX	60.0 ± 40%	0 ± 0	0 ± 0	0 ± 0	QMIX	100 ± 0%	100 ± 0%	89 ± 10%	61 ± 37%
VDN	0 ± 0	0 ± 0	0 ± 0	0 ± 0	VDN	0 ± 0	0 ± 0	0 ± 0	0 ± 0

279 **Performance of n -step return and $TD(\lambda)$ methods.** To answer **Q3**, we use n -step return and
 280 $TD(\lambda)$ to estimate the TD-target. As shown in Table. 2, with n -step return, both QMIX and VDN
 281 fail to learn good policies even with $n = 15$. Surprisingly, with $TD(\lambda)$, QMIX can achieve good
 282 performance with $\lambda \in \{0.8, 0.9, 0.99, 1\}$. However, we cannot find such outcome on VDN and there
 283 is no guarantee of good results on using $TD(\lambda)$.

284 **Performance of Multi-Agent Exploration and Risk-Sensitive MARL methods.** We also provide results of
 285 exploration methods for MARL and risk-sensitive MARL method to answer **Q4**. MAVEN utilizes mutual information
 286 to learn latent space for exploration and RMIX aims to learning risk-sensitive policies for MARL. In Fig. 7, RMIX
 287 even fails to learn. Mainly because the potential loss of reward is displaced by off-beat actions. Overall, MAVEN
 288 is stabler than EMC and RMIX. QMIX-LeGEM is stable in all scenarios and outperforms MAVEN. With LeGEM,
 289 even simple method such VDN can perform well and outperforms many MARL methods with complex and advanced components. Indeed, exploration
 290 in OBMA is beneficial for multi-agent learning. However, the key challenge of temporal credit
 291 assignment can not be easily addressed merely with exploration.

292 **SMAC.** We also conduct experiments on SMAC [44]. We train MARL methods and our method
 293 on 2m_vs_1z where are two agents combating with one opponent. To overcome the issue of
 294




Figure 8: The performance of MARL methods on 2m_vs_1z.

300 high dimension continuous state space, We utilize simhash [9] to calculate the hash value of the
301 key. We only select the attack action and set the action duration with 9. As illustrated in Fig. 8,
302 incorporated with our novel episodic memory, QMIX, IQL and VDN illustrate enhanced performance,
303 demonstrating the superiority of our method on complex multi-agent scenarios.

304 7 Related Works

305 **Action Delay in RL.** Conventionally, the execution of actions in RL is instantaneous and the execution
306 duration is neglected. Katsikopoulos et al. [20] propose the Delayed MDP where actions have delays
307 and Walsh et al [59] propose a model-based method for the Delayed MDP. To optimize the delayed
308 MDP, many RL approaches [59, 39, 66, 69] augment the state space with the queuing actions to
309 be executed into the environment. However, this state-augmentation trick is intractable [11]. Chen
310 et al. [10] extend the delayed MDP [39] and propose a Delayed Markov Game. However, the
311 state-augmentation treatment is confined to short delays and neglects the off-beat actions in multi-
312 agent scenarios. Recently, Bouteiller et al. [6] apply replay buffer correction method. However, the
313 delayed timestep is privileged information. It is not available for agents in many scenarios. Simply
314 applying this single-agent trajectory correction in MARL cannot attain satisfactory performance due
315 to off-beat actions; devising inter-agent trajectory correction methods for OBMAS is non-trivial.

316 **Credit Assignment in RL.** Credit assignment [50, 52] tackles long-horizon sequential decision-
317 making problem by distributing the contribution of each single step over the temporal interval. TD
318 learning [51] is the most established credit assignment method, which is the basis of many RL methods.
319 RUDDER [2] redistributes the episodic return to key timesteps in the episode [14, 42, 40]. Klissarov
320 et al. [22] propose a reward propagation method via graph convolutional neural network [21]. Another
321 line of works utilize episodic memory (EM) [37, 5, 73, 27, 74] to recall key events and aggregate
322 information of the past for decision-making or learning. However, simply applying EM of RL to
323 MARL cannot perform well in OBMAS due to the non-stationarity and the displaced rewards.

324 **Multi-Agent RL.** Many MARL methods focus on factorizing the global Q value to train agents’
325 policies via CTDE [13, 49, 41, 47, 60, 63, 35]. However, these existing works assume actions are
326 executed synchronously. Messias et al. [30] propose an event-driven, asynchronous formulation of
327 the multi-agent POMDP. However, the assumption of free communication [61] is limited and the
328 asynchronous execution [34] in the paper is confined to the design of events and did not propose
329 methods on solving challenging credit assignment issue in OBMAS. Recently, Amato et al. [1]
330 and Xiao et al. [67, 68] propose macro-action methods, which are similar to hierarchical methods.
331 Macro-actions are manually designed via abstracting primitive actions. However, macro-action
332 methods mainly focus on macro-action selection during multi-timestep decision-making and assume
333 the environment can use manually pre-defined methods for state transition. Unfortunately, the above
334 works either focus on synchronous actions or defining specific asynchronous execution components
335 with human knowledge. Learning coordination in OBMAS remains a challenge.

336 8 Conclusion

337 In this paper, we investigate model-free MARL with off-beat actions. To address challenges in
338 OBMAS, we first propose Off-Beat Dec-POMDP. Then, we propose a new class of episodic memory,
339 LeGEM, for model-free MARL algorithms. LeGEM addresses the challenging temporal credit
340 assignment problem raised by off-beat actions in TD-learning via the novel reward redistribution
341 scheme. We evaluate our method on various OBMAS scenarios. Empirical results show that our
342 method significantly boosts the multi-agent coordination and achieves leading performance as well as
343 improved sample efficiency.

344 **Limitations and Future Work.** Searching from a graph-structured episodic memory takes much
345 overhead in LeGEM. Scaling up LeGEM to complex OBMAS is our future direction. Recently, there
346 is a growing interest in model-based planing [45]. Leveraging LeGEM for model-based planning
347 is also our future work. Our paper focuses on Dec-POMDP-based MARL methods. We leave it to
348 future work for investigating off-beat actions in frameworks like Markov Game [26] and MMDP [7].
349 We are also interested in finding the merit of our method in real-world problem in our future work,
350 such as scheduling [29] with off-beat settings.

References

- 351
- 352 [1] C. Amato, G. Konidaris, L. P. Kaelbling, and J. P. How. Modeling and planning with macro-
353 actions in decentralized pomdps. *Journal of Artificial Intelligence Research*, 64:817–859,
354 2019.
- 355 [2] J. A. Arjona-Medina, M. Gillhofer, M. Widrich, T. Unterthiner, J. Brandstetter, and S. Hochreiter.
356 RUDDER: Return decomposition for delayed rewards. In *Advances in Neural Information*
357 *Processing Systems*, volume 32, 2019.
- 358 [3] P.-L. Bacon, J. Harb, and D. Precup. The option-critic architecture. In *Proceedings of the AAAI*
359 *Conference on Artificial Intelligence*, volume 31, 2017.
- 360 [4] N. Biggs, E. K. Lloyd, and R. J. Wilson. *Graph Theory, 1736-1936*. Oxford University Press,
361 1986.
- 362 [5] M. Botvinick, S. Ritter, J. X. Wang, Z. Kurth-Nelson, C. Blundell, and D. Hassabis. Reinforce-
363 ment learning, fast and slow. *Trends in Cognitive Sciences*, 23(5):408–422, 2019.
- 364 [6] Y. Bouteiller, S. Ramstedt, G. Beltrame, C. Pal, and J. Binas. Reinforcement learning with
365 random delays. In *International Conference on Learning Representations*, 2020.
- 366 [7] C. Boutilier. Planning learning and coordination in multiagent decision processes. In *Proceed-*
367 *ings of the 6th Conference on the Theoretical Aspects of Rationality and Knowledge*, pages
368 195–210, 1996.
- 369 [8] Y. Cao, W. Yu, W. Ren, and G. Chen. An overview of recent progress in the study of distributed
370 multi-agent coordination. *IEEE Transactions on Industrial Informatics*, 9(1):427–438, 2012.
- 371 [9] M. S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of*
372 *the thirty-fourth Annual ACM Symposium on Theory of Computing*, pages 380–388, 2002.
- 373 [10] B. Chen, M. Xu, Z. Liu, L. Li, and D. Zhao. Delay-aware multi-agent reinforcement learning
374 for cooperative and competitive environments. *arXiv e-prints*, pages arXiv–2005, 2020.
- 375 [11] E. Derman, G. Dalal, and S. Mannor. Acting in delayed environments with non-stationary
376 Markov policies. In *International Conference on Learning Representations*, 2020.
- 377 [12] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley,
378 I. Dunning, et al. Impala: Scalable distributed deep-RL with importance weighted actor-learner
379 architectures. In *International Conference on Machine Learning*, pages 1407–1416, 2018.
- 380 [13] J. Foerster, G. Farquhar, T. Afouras, N. Nardelli, and S. Whiteson. Counterfactual multi-agent
381 policy gradients. *arXiv preprint arXiv:1705.08926*, 2017.
- 382 [14] T. Gangwani, Y. Zhou, and J. Peng. Learning guidance rewards with trajectory-space smoothing.
383 *Advances in Neural Information Processing Systems*, 33:822–832, 2020.
- 384 [15] D. Ha and Y. Tang. Collective intelligence for deep learning: A survey of recent developments.
385 *arXiv preprint arXiv:2111.14377*, 2021.
- 386 [16] B. Han, Z. Ren, Z. Wu, Y. Zhou, and J. Peng. Off-policy reinforcement learning with delayed
387 rewards. *arXiv preprint arXiv:2106.11854*, 2021.
- 388 [17] H. Hu, J. Ye, G. Zhu, Z. Ren, and C. Zhang. Generalizable episodic memory for deep
389 reinforcement learning. In *Proceedings of the 38th International Conference on Machine*
390 *Learning*, pages 4380–4390, 18–24 Jul 2021.
- 391 [18] C.-C. Hung, T. Lillicrap, J. Abramson, Y. Wu, M. Mirza, F. Carnevale, A. Ahuja, and G. Wayne.
392 Optimizing agent behavior over long time scales by transporting value. *Nature Communications*,
393 10(1):1–12, 2019.
- 394 [19] M. Hüttenrauch, A. Šošić, and G. Neumann. Guided deep reinforcement learning for swarm
395 systems. In *AAMAS 2017 Autonomous Robots and Multirobot Systems (ARMS) Workshop*,
396 2017.

- 397 [20] K. V. Katsikopoulos and S. E. Engelbrecht. Markov decision processes with delays and
398 asynchronous cost collection. *IEEE Transactions on Automatic Control*, 48(4):568–574, 2003.
- 399 [21] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks.
400 *arXiv preprint arXiv:1609.02907*, 2016.
- 401 [22] M. Klissarov and D. Precup. Reward propagation using graph convolutional networks. *Advances*
402 *in Neural Information Processing Systems*, 33, 2020.
- 403 [23] J. G. Kuba, R. Chen, M. Wen, Y. Wen, F. Sun, J. Wang, and Y. Yang. Trust region policy
404 optimisation in multi-agent reinforcement learning. *arXiv preprint arXiv:2109.11251*, 2021.
- 405 [24] A. K. Lampinen, S. C. Chan, A. Banino, and F. Hill. Towards mental time travel: a hierarchical
406 memory for reinforcement learning agents. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W.
407 Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- 408 [25] J. Z. Leibo, E. D. nez Guzmán, A. S. Vezhnevets, J. P. Agapiou, P. Sunehag, R. Koster, J. Matyas,
409 C. Beattie, I. Mordatch, and T. Graepel. Scalable evaluation of multi-agent reinforcement
410 learning with melting pot. PMLR, 2021.
- 411 [26] M. L. Littman. Markov games as a framework for multi-agent reinforcement learning. In
412 *Machine Learning Proceedings 1994*, pages 157–163. Elsevier, 1994.
- 413 [27] X. Ma and W.-J. Li. State-based episodic memory for multi-agent reinforcement learning. *arXiv*
414 *preprint arXiv:2110.09817*, 2021.
- 415 [28] A. Mahajan, T. Rashid, M. Samvelyan, and S. Whiteson. MAVEN: Multi-agent variational
416 exploration. In *Advances in Neural Information Processing Systems*, pages 7613–7624, 2019.
- 417 [29] H. Mao, M. Schwarzkopf, S. B. Venkatakrishnan, Z. Meng, and M. Alizadeh. Learning
418 scheduling algorithms for data processing clusters. In *Proceedings of the ACM special interest*
419 *group on data communication*, pages 270–288. 2019.
- 420 [30] J. V. Messias, M. T. Spaan, and P. U. Lima. Multiagent pomdps with asynchronous execution.
421 In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent*
422 *Systems*, pages 1273–1274, 2013.
- 423 [31] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves,
424 M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep rein-
425 forcement learning. *Nature*, 518(7540):529–533, 2015.
- 426 [32] F. A. Oliehoek, C. Amato, et al. *A Concise Introduction to Decentralized POMDPs*, volume 1.
427 Springer, 2016.
- 428 [33] F. A. Oliehoek, M. T. Spaan, and N. Vlassis. Optimal and approximate q-value functions for
429 decentralized POMDPs. *Journal of Artificial Intelligence Research*, 32:289–353, 2008.
- 430 [34] S. Omidshafiei, A.-A. Agha-Mohammadi, C. Amato, and J. P. How. Decentralized control of
431 partially observable markov decision processes using belief space macro-actions. In *2015 IEEE*
432 *International Conference on Robotics and Automation (ICRA)*, pages 5962–5969, 2015.
- 433 [35] L. Pan, T. Rashid, B. Peng, L. Huang, and S. Whiteson. Regularized softmax deep multi-agent
434 q-learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- 435 [36] W. B. Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*,
436 volume 703. John Wiley & Sons, 2007.
- 437 [37] A. Pritzel, B. Uria, S. Srinivasan, A. P. Badia, O. Vinyals, D. Hassabis, D. Wierstra, and
438 C. Blundell. Neural episodic control. In *International Conference on Machine Learning*, pages
439 2827–2836, 2017.
- 440 [38] W. Qiu, X. Wang, R. Yu, R. Wang, X. He, B. An, S. Obraztsova, and Z. Rabinovich. RMIX:
441 Learning risk-sensitive policies for cooperative reinforcement learning agents. In *Advances in*
442 *Neural Information Processing Systems*, 2021.

- 443 [39] S. Ramstedt and C. Pal. Real-time reinforcement learning. *Advances in Neural Information*
444 *Processing Systems*, 32:3073–3082, 2019.
- 445 [40] D. Raposo, S. Ritter, A. Santoro, G. Wayne, T. Weber, M. Botvinick, H. van Hasselt, and
446 F. Song. Synthetic returns for long-term credit assignment. *arXiv preprint arXiv:2102.12425*,
447 2021.
- 448 [41] T. Rashid, M. Samvelyan, C. Schroeder, G. Farquhar, J. Foerster, and S. Whiteson. QMIX:
449 Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International*
450 *Conference on Machine Learning*, pages 4295–4304, 2018.
- 451 [42] Z. Ren, R. Guo, Y. Zhou, and J. Peng. Learning long-term reward redistribution via randomized
452 return decomposition. *arXiv e-prints*, pages arXiv–2111, 2021.
- 453 [43] R. T. Rockafellar, S. Uryasev, et al. Optimization of conditional value-at-risk. *Journal of Risk*,
454 2:21–42, 2000.
- 455 [44] M. Samvelyan, T. Rashid, C. S. de Witt, G. Farquhar, N. Nardelli, T. G. J. Rudner, C.-M. Hung,
456 P. H. S. Torr, J. Foerster, and S. Whiteson. The StarCraft Multi-Agent Challenge. *CoRR*,
457 abs/1902.04043, 2019.
- 458 [45] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lock-
459 hart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a
460 learned model. *Nature*, 588(7839):604–609, 2020.
- 461 [46] D. Silver, S. Singh, D. Precup, and R. S. Sutton. Reward is enough. *Artificial Intelligence*,
462 299:103535, 2021.
- 463 [47] K. Son, D. Kim, W. J. Kang, D. E. Hostallero, and Y. Yi. QTRAN: Learning to factorize with
464 transformation for cooperative multi-agent reinforcement learning. In *International Conference*
465 *on Machine Learning*, pages 5887–5896, 2019.
- 466 [48] T. Suddendorf, D. R. Addis, and M. C. Corballis. Mental time travel and the shaping of
467 the human mind. *Philosophical Transactions of the Royal Society B: Biological Sciences*,
468 364(1521):1317–1324, 2009.
- 469 [49] P. Sunehag, G. Lever, A. Grusl, W. M. Czarnecki, V. Zambaldi, M. Jaderberg, M. Lanctot,
470 N. Sonnerat, J. Z. Leibo, K. Tuyls, et al. Value-decomposition networks for cooperative
471 multi-agent learning. *arXiv preprint arXiv:1706.05296*, 2017.
- 472 [50] R. S. Sutton. *Temporal credit assignment in reinforcement learning*. PhD thesis, University of
473 Massachusetts Amherst, 1984.
- 474 [51] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.
- 475 [52] R. S. Sutton, D. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal
476 abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- 477 [53] A. Tampuu, T. Matiisen, D. Kodelja, I. Kuzovkin, K. Korjus, J. Aru, J. Aru, and R. Vicente.
478 Multiagent cooperation and competition with deep reinforcement learning. *PLoS ONE*, 12(4),
479 2017.
- 480 [54] H. Tang, R. Houthoofd, D. Foote, A. Stooke, O. Xi Chen, Y. Duan, J. Schulman, F. DeTurck, and
481 P. Abbeel. # exploration: A study of count-based exploration for deep reinforcement learning.
482 *Advances in neural information processing systems*, 30, 2017.
- 483 [55] E. Tulving. Memory and consciousness. *Canadian Psychology/Psychologie Canadienne*,
484 26(1):1, 1985.
- 485 [56] H. van Hasselt, S. Madjiheurem, M. Hessel, D. Silver, A. Barreto, and D. Borsa. Expected
486 eligibility traces. In *Proceedings of the AAAI Conference on Artificial Intelligence*, number 11,
487 pages 9997–10005, 2021.

- 488 [57] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and
489 I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*,
490 pages 5998–6008, 2017.
- 491 [58] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi,
492 R. Powell, T. Ewalds, P. Georgiev, et al. Grandmaster level in StarCraft II using multi-agent
493 reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- 494 [59] T. J. Walsh, A. Nouri, L. Li, and M. L. Littman. Learning and planning in environments with
495 delayed feedback. *Autonomous Agents and Multi-Agent Systems*, 18(1):83–105, 2009.
- 496 [60] J. Wang, Z. Ren, T. Liu, Y. Yu, and C. Zhang. QPLEX: Duplex dueling multi-agent q-learning.
497 *arXiv preprint arXiv:2008.01062*, 2020.
- 498 [61] R. Wang, X. He, R. Yu, W. Qiu, B. An, and Z. Rabinovich. Learning efficient multi-agent
499 communication: An information bottleneck approach. In *International Conference on Machine
500 Learning*, pages 9908–9918. PMLR, 2020.
- 501 [62] W. Wang, T. Yang, Y. Liu, J. Hao, X. Hao, Y. Hu, Y. Chen, C. Fan, and Y. Gao. Action
502 semantics network: Considering the effects of actions in multiagent systems. *arXiv preprint
503 arXiv:1907.11461*, 2019.
- 504 [63] Y. Wang, B. Han, T. Wang, H. Dong, and C. Zhang. DOP: Off-policy multi-agent decomposed
505 policy gradients. In *International Conference on Learning Representations*, 2021.
- 506 [64] Z. Wang, T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas. Dueling network
507 architectures for deep reinforcement learning. In *International Conference on Machine Learning*,
508 pages 1995–2003, 2016.
- 509 [65] C. J. Watkins and P. Dayan. Q-Learning. *Machine Learning*, 8(3-4):279–292, 1992.
- 510 [66] T. Xiao, E. Jang, D. Kalashnikov, S. Levine, J. Ibarz, K. Hausman, and A. Herzog. Think-
511 ing while moving: Deep reinforcement learning with concurrent control. In *International
512 Conference on Learning Representations*, 2019.
- 513 [67] Y. Xiao, J. Hoffman, and C. Amato. Macro-action-based deep multi-agent reinforcement
514 learning. In *Conference on Robot Learning*, pages 1146–1161. PMLR, 2020.
- 515 [68] Y. Xiao, J. Hoffman, T. Xia, and C. Amato. Multi-agent/robot deep reinforcement learning
516 with macro-actions (student abstract). In *Proceedings of the AAAI Conference on Artificial
517 Intelligence*, volume 34, pages 13965–13966, 2020.
- 518 [69] Y. Yuan and R. Mahmood. Asynchronous reinforcement learning for real-time control of
519 physical robots. *arXiv preprint arXiv:2203.12759*, 2022.
- 520 [70] T. Zahavy, B. O’Donoghue, G. Desjardins, and S. Singh. Reward is enough for convex mdps.
521 *Advances in Neural Information Processing Systems*, 34, 2021.
- 522 [71] L. Zheng, J. Chen, J. Wang, J. He, Y. Hu, Y. Chen, C. Fan, Y. Gao, and C. Zhang. Episodic
523 multi-agent reinforcement learning with curiosity-driven exploration. *Advances in Neural
524 Information Processing Systems*, 34, 2021.
- 525 [72] M. Zhou, J. Luo, J. Vilella, Y. Yang, D. Rusu, J. Miao, W. Zhang, M. Alban, I. Fadarar, Z. Chen,
526 et al. Smarts: Scalable multi-agent reinforcement learning training school for autonomous
527 driving. *arXiv preprint arXiv:2010.09776*, 2020.
- 528 [73] Y. Zhou, D. E. Asher, N. R. Waytowich, and J. A. Shah. On memory mechanism in multi-agent
529 reinforcement learning. *arXiv e-prints*, pages arXiv–1909, 2019.
- 530 [74] G. Zhu, Z. Lin, G. Yang, and C. Zhang. Episodic reinforcement learning with associative
531 memory. In *International Conference on Learning Representations*, 2020.

532 **Checklist**

- 533 1. For all authors...
- 534 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
535 contributions and scope? [Yes] See Sec. 4.
- 536 (b) Did you describe the limitations of your work? [Yes] See Sec. 8.
- 537 (c) Did you discuss any potential negative societal impacts of your work? [No] Our method
538 does not have negative social impacts.
- 539 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
540 them? [Yes]
- 541 2. If you are including theoretical results...
- 542 (a) Did you state the full set of assumptions of all theoretical results? [Yes] See the
543 statements.
- 544 (b) Did you include complete proofs of all theoretical results? [Yes] See Appx. A.
- 545 3. If you ran experiments...
- 546 (a) Did you include the code, data, and instructions needed to reproduce the main ex-
547 perimental results (either in the supplemental material or as a URL)? [Yes] See the
548 supplementary files.
- 549 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
550 were chosen)? [Yes] See Appx. E.
- 551 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
552 ments multiple times)? [Yes] See Sec. 6 and Appx. E.
- 553 (d) Did you include the total amount of compute and the type of resources used (e.g., type
554 of GPUs, internal cluster, or cloud provider)? [Yes] See Appx. E.
- 555 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 556 (a) If your work uses existing assets, did you cite the creators? [Yes] We cited, see Sec. 6.
- 557 (b) Did you mention the license of the assets? [Yes] See Appx. E.
- 558 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
559 See the supplementary files.
- 560 (d) Did you discuss whether and how consent was obtained from people whose data you're
561 using/curating? [N/A]
- 562 (e) Did you discuss whether the data you are using/curating contains personally identifiable
563 information or offensive content? [N/A]
- 564 5. If you used crowdsourcing or conducted research with human subjects...
- 565 (a) Did you include the full text of instructions given to participants and screenshots, if
566 applicable? [N/A]
- 567 (b) Did you describe any potential participant risks, with links to Institutional Review
568 Board (IRB) approvals, if applicable? [N/A]
- 569 (c) Did you include the estimated hourly wage paid to participants and the total amount
570 spent on participant compensation? [N/A]