
Supplementary Material for Cooperative Multi-Agent Reinforcement Learning with Sequential Credit Assignment

Anonymous Author(s)

Affiliation

Address

email

1 A Proof of Claim 1

2 **Claim 1.** *The proposed sequential credit assignment achieves additive advantage-decomposition.*

3 *Proof.* We discuss a multi-agent system with n agents identified by $a_i (i \in \{1, \dots, n\})$ under one
4 specific sequence $\langle a_1, a_2, \dots, a_n \rangle$. Claim 1 can also be proved from the rest $(n! - 1)$ orders in the
5 same way. Here we denote $\mathbf{u}_{a_1}^{a_i} = [u^{a_1}, u^{a_2}, \dots, u^{a_i}] (i = 1, 2, 3, \dots, n)$.

6 For better understanding, we analyze in the reverse order (from agent a_n to agent a_1). As for the last
7 agent a_n in the sequence, we evaluate its action based on all the preceding agents' fixed behaviors, so
8 there is no need to calculate the expectations on the others' actions, and the advantage function for
9 agent a_n is the same as in COMA:

$$A^{a_n}(s, \mathbf{u}) = Q(s, \mathbf{u}) - \sum_{u^{a_n}} \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (\mathbf{u}^{-a_n}, u^{a_n})). \quad (\text{A1})$$

10 When evaluating the second-to-last agent a_{n-1} , the actions of agent a_1 to a_{n-2} are fixed. We only
11 consider the expectation on agent a_n 's action for the first term in Equ.(6) and the expectation on the
12 actions of agent a_{n-1} and a_n for the second term in Equ.(6) in our main paper:

$$\begin{aligned} A^{a_{n-1}}(s, \mathbf{u}) &= \sum_{u^{a_n}} \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (\mathbf{u}^{-a_n}, u^{a_n})) \\ &\quad - \sum_{u^{a_{n-1}}} \sum_{u^{a_n}} \pi^{a_{n-1}}(u^{a_{n-1}} | \tau^{a_{n-1}}) \cdot \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (\mathbf{u}_{a_1}^{a_{n-2}}, u^{a_{n-1}}, u^{a_n})). \end{aligned} \quad (\text{A2})$$

13 We can conclude the advantage function of each agent $a_i (i \in \{2, \dots, n\})$ as:

$$\begin{aligned} A^{a_i}(s, \mathbf{u}) &= \sum_{u^{a_{i+1}}} \dots \sum_{u^{a_n}} \pi^{a_{i+1}}(u^{a_{i+1}} | \tau^{a_{i+1}}) \dots \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (\mathbf{u}_{a_1}^{a_i}, u^{a_{i+1}}, \dots, u^{a_n})) \\ &\quad - \sum_{u^{a_i}} \dots \sum_{u^{a_n}} \pi^{a_i}(u^{a_i} | \tau^{a_i}) \dots \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (\mathbf{u}_{a_1}^{a_{i-1}}, u^{a_i}, \dots, u^{a_n})), \end{aligned} \quad (\text{A3})$$

14 and the advantage of the first agent in the sequence a_1 is:

$$\begin{aligned} A^{a_1}(s, \mathbf{u}) &= \sum_{u^{a_2}} \dots \sum_{u^{a_n}} \pi^{a_2}(u^{a_2} | \tau^{a_2}) \dots \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (u^{a_1}, u^{a_2}, \dots, u^{a_n})) \\ &\quad - \sum_{u^{a_1}} \dots \sum_{u^{a_n}} \pi^{a_1}(u^{a_1} | \tau^{a_1}) \dots \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (u^{a_1}, u^{a_2}, \dots, u^{a_n})) \end{aligned} \quad (\text{A4})$$

15 The second term of the above equation, i.e., the expected Q value of all possible joint actions, is
 16 recognized as the value of the state $V(s)$. We then rewrite Equ.(A4) as:

$$A^{a_1}(s, \mathbf{u}) = \sum_{u^{a_2}} \cdots \sum_{u^{a_n}} \pi^{a_2}(u^{a_2} | \tau^{a_2}) \cdots \pi^{a_n}(u^{a_n} | \tau^{a_n}) \cdot Q(s, (u^{a_1}, u^{a_2}, \dots, u^{a_n})) - V(s) \quad (\text{A5})$$

17 It can be seen from Equ.(A3) that the second term of $A^{a_i}(s, \mathbf{u})$ is the same as $A^{a_{i-1}}(s, \mathbf{u})$'s first term
 18 ($i \in \{2, 3, \dots, n\}$). We can eliminate most of the terms by summing from $A^{a_1}(s, \mathbf{u})$ to $A^{a_n}(s, \mathbf{u})$,
 19 acquiring the following equation that only retains the first term of A^{a_n} and the second term of A^{a_1} .
 20 Therefore, we derive the following equation from Equ.(A1) and Equ.(A5):

$$\sum_{i=1}^n A^{a_i}(s, \mathbf{u}) = Q(s, \mathbf{u}) - V(s) = A(s, \mathbf{u}). \quad (\text{A6})$$

21 From Equ.(A6), we can see that our explicit credit assignment through the proposed sequential
 22 advantage function decomposes the total advantage function $A(s, \mathbf{u})$ in an additive form. \square

23 B Pseudocode

24 We provide SeCA's pseudocode below for a better understanding of its optimization procedure.

Algorithm 1 Sequential Credit Assignment Optimization Procedure

```

1: Randomly initialize the policy network parameter  $\theta$  and the critic network parameter  $\phi$ .
2: Initialize target critic network parameter  $\phi^- \leftarrow \phi$ .
3: while not terminated do
4:   Sample  $b$  episodes  $\{\tau^i\}$  ( $i \in \{1, 2, \dots, b\}$ ) where  $\tau^i = \{s_0^i, z_0^i, \mathbf{u}_0^i, r_0^i, \dots, s_T^i, z_T^i, \mathbf{u}_T^i, r_T^i\}$ .
5:   for each episode  $i = 1$  to  $b$  do
6:     for timestep  $t = T$  to  $0$  do
7:       Compute target  $y_t^i = r_t^i + \gamma(\lambda y_{t+1}^i + (1 - \lambda)f_{\phi^-}(s_{t+1}^i, \mathbf{u}_{t+1}^i))$ .
8:       Compute critic loss  $\mathcal{L}_t(\phi)^i = (y_t^i - f_{\phi}(s_t^i, \mathbf{u}_t^i))^2$ .
9:     end for
10:    for all agents  $a \in \{1, 2, \dots, n\}$  do
11:      Compute  $a$ 's contribution to  $\tau^i$ :  $c_a^i = \sum_{x_j \in \pi^a} \text{PathIG}_j^{\tau^i}(\pi^a)$  according to Equ.(1).
12:    end for
13:    Decide the sequence based on  $c_a^i$ . Arrange agents with higher contributions in the front.
14:    for all agents in the sequence do
15:      Compute each agent  $a$ 's sequential advantage  $A_a^i(s, \mathbf{u})$  according to Equ.(7) and (8).
16:    end for
17:  end for
18:  // Critic Learning:
19:  Update the critic parameter  $\phi$  by descending the gradient  $\nabla_{\phi} \frac{1}{bT} \sum_i \sum_t \mathcal{L}_t(\phi)^i$ .
20:  // Policy Learning:
21:  Update the policy parameter  $\theta$  by by maximizing the following objective:
      
$$\frac{1}{n} \sum_a \left[ \log \pi_a(u_a | \tau_a) \left( \frac{1}{b} \sum_i \left( \frac{1}{T} \sum_t A_a^i(s_t, \mathbf{u}_t) \right) \right) + \mathcal{H}(\pi^a(\cdot | \tau^a)) \right]$$

22:  if at target update interval then
23:    Update the target critic parameter  $\phi^- \leftarrow \phi$ .
24:  end if
25: end while

```

25 C Experiment Details on Multi-Agent Particle Environments

26 We evaluate our proposed sequential advantage and COMA's advantage in two multi-agent particle
 27 environments [3] in Section 3.3. The code for these experiment environments is available in the
 28 Codes folder in our Supplementary Material. Here we introduce these two environments and our
 29 settings in detail.

30 C.1 Environments Introduction

31 **Predator-Prey.** Three slower predators cooperate to chase a faster prey that acts randomly in an
32 area containing two obstacles impeding the way. The action space of the predators is [move_up,
33 move_down, move_left, move_right, stay]. Each predator can observe its current position
34 and velocity and the displacement to the prey, other predators, and the obstacles. The predators’
35 goal is to capture the prey in as few steps as possible, and the shared reward is the negative minimal
36 distance between any predator and the prey. When any predator captures the prey, an additional
37 positive reward is given to the team, and the game terminates.

38 **Cooperative Navigation.** This environment initializes three agents and three landmarks with random
39 locations in an area. The agents aim to acquire a bigger shared reward and cooperate to cover all the
40 landmarks with action space [move_up, move_down, move_left, move_right, stay]. Each
41 agent observes its own position, velocity, and the displacement to the other agents and the targets.
42 The global reward is the negative sum of the distance between each target and the nearest agent to it.
43 If the agents collide, the team will receive a penalty, so agents must avoid collisions.

44 C.2 Training Settings

45 We train agents for 5000 episodes for these two environments. Each episode has a maximum of 200
46 steps. We adapt the original open-source environment implementation¹ and the training framework²
47 provided by [7]. We follow the pre-set environment hyperparameters, including the shared reward for
48 capturing the prey in Predator-Prey, the collision penalty in Cooperative Navigation, and the spawn
49 regions size of the agents and obstacles. The prey acts randomly in Predator-Prey is a random agent
50 with a uniform action probability distribution.

51 We utilize the default setting of the training framework. The policy networks and critic network
52 implement one hidden layer MLPs, with 128 units and 32 units for Predator-Prey and Cooperative
53 Navigation respectively. Agents share parameters in these environments. The number of transitions
54 for each update is 128 and 32, and the discount factor is 0.99 and 0.9. We leverage target networks
55 that update every 200 training iteration. Our models are trained by Adam Optimizer with a learning
56 rate of 0.0002 for critic learning in Predator-Prey and 0.0001 in Cooperative Navigation, and 0.0001
57 for policy network in both environments.

58 Note, we only compare our sequential advantage with COMA’s counterfactual advantage in Figure 2
59 in the main paper, so we do not implement SeCA’s whole architecture. The three agents’ advantages
60 for policy learning here are directly calculated according to Equ.(5) or Equ.(7) in our main paper.
61 The credit assignment sequence is fixed and initialized as <Agent 1, Agent 2, Agent 3>.

62 D Details of StarCraft Multi-Agent Challenge Experiments

63 We mainly evaluate methods on StarCraft II micromanagement in Section 4 and follow the default
64 setup of the StarCraft Multi-Agent Challenge (SMAC) [5].³ We utilize the open-source implementa-
65 tions of the baseline algorithms, including COMA, QMIX, QTRAN⁴, and LICA⁵. All these methods
66 are all based on the PyMARL framework [5]. The code for SeCA in SMAC is available in the Codes
67 folder in Supplementary Material.

68 D.1 Detailed Information about SMAC and Scenarios

69 A group of units controlled by decentralized agents cooperates to defeat the enemy agent system
70 controlled by handcrafted heuristics in each SMAC micromanagement problem. Each agent’s
71 partial observation comprises of the attributes (such as health, location, unit_type) of all units
72 shown up in its view range. The global state information includes all agents’ positions and health,
73 and allied units’ last actions and cooldown, which is only available to agents during centralized

¹<https://github.com/openai/multiagent-particle-envs>

²<https://github.com/hsvgbkghbv/SQDDPG>

³<https://github.com/oxwhirl/smac>

⁴<https://github.com/oxwhirl/pymarl>

⁵<https://github.com/mzho7212/LICA>

74 training. The agents’ discrete action space consists of `attack[enemy_id]`, `move[direction]`,
75 `stop`, and `no-op` for the dead agents only. Particular unit Medivac has no action `attack[enemy_id]`
76 but has `heal[enemy_id]`. Agents can only attack enemies within their shooting range. Proper
77 micromanagement requires agents to maximize the damage to the enemies and take as little damage
78 as possible in combat, so they need to cooperate with each other or even sacrifice themselves. We
79 follow the default setup of SMAC in our experiments, and more settings, including rewards and
80 detailed observation/state information, can be acquired from the original paper or implementation.

81 Based on baseline algorithms’ performances, the scenarios in SMAC are broadly grouped into three
82 categories: *Easy*, *Hard*, and *Super Hard*. The key point to win some *Hard* or *Super Hard* battles is
83 mastering specific micro techniques, such as *focus fire*, *avoid overkill*, *kiting*, et cetera. The battles
84 can be both symmetric or asymmetric, and the group of agents can be homogeneous or heteroge-
85 neous. We consider six scenarios with different difficulties and characteristics: *2s3z*, *1c3s5z* (*Easy*,
86 heterogeneous, symmetric); *2c_vs_64zg*, *3s_vs_5z* (*Hard*, homogeneous, asymmetric); and *MMM2*,
87 *3s5z_vs_3s6z* (*Super Hard*, heterogeneous, asymmetric). Here we provide some characteristics of
88 each scenario to help gain insights into the good or poor performance of the methods:

- 89 • Both *2s3z* and *1c3s5z* are symmetric combats where two heterogeneous teams battle
90 against each other. *s* represents Stalkers that can attack enemies at a distance. *z* represents
91 Zealots, melee units with a short attack range. *c* for Colossus. It is a ranged and durable
92 unit that can harm an area instead of a single agent. These two *Easy* scenarios that do not
93 need to cooperate too much can be solved easily by most of the recognized methods.
- 94 • *2c_vs_64zg* is a *Hard* asymmetric scenario where two Colossi battle against 64 Zerglings,
95 which are melee units with low health and low attack damage. The number of the enemy
96 units in this map is the largest in the SMAC benchmark, making the agents’ action space
97 much larger than other maps. It can be utilized as an example to test methods’ performance
98 on large action spaces.
- 99 • *3s_vs_5z* is also a *Hard* asymmetric battle between two different homogeneous teams. The
100 allied Stalkers have to master the *kiting* technique and disperse in the area to kill the Zealots
101 that chase them one after another. This map faces the delayed reward problem; however, it
102 is not very strict about micro-cooperation between agents because of their scattering.
- 103 • *MMM2* is a representative *Super Hard* asymmetric battle between two heterogeneous teams
104 with three kinds of units. One Medivac, two Marauders, and seven Marines have to battle
105 against a team with one more Marine. Marauder has greater attack damage and health
106 than Marine but with a longer attack `cooldown`. Medivac has no damage but can heal any
107 other agent in the team. This map with three kinds of units and many agents requires more
108 cooperation between agents, so we picked this map for our ablation studies.
- 109 • *3s5z_vs_3s6z* is another *Super Hard* map that requires breaking the bottleneck of explo-
110 ration, where three Stalkers and five Zealots battle against three Stalkers and six Zealots.

111 D.2 Training Settings

112 We follow most of the training hyperparameters in the original PyMARL implementation. Figure
113 3 in the main paper illustrates SeCA’s network structure. The policy network consists of two FC
114 layers and a GRU layer between them. The critic network maps the state into two weight matrices
115 and biases, which, in turn, maps the concatenated action-policy vector into the Q estimate. We follow
116 existing methods [1, 9] where all methods align on the batch size, the number of batch updates, and
117 the total number of environment steps. We train all the methods for 32 million steps in *Easy* scenarios
118 and 64 million steps in *Hard* and *Super Hard* scenarios. We use 32 actors to generate the trajectories
119 in parallel and use one NVIDIA Titan V GPU for training. More training details of our method
120 utilized in SMAC are shown in Table A1 and can be found in our code.

121 D.3 Additional Results on SMAC

122 We do not compare QPD’s learning curves with other methods in Figure 5 of our main paper, as
123 QPD modifies the original SMAC’s implementation, and it is unfair to compare QPD’s learning
124 speed with other methods that trained in the original environment. Here we follow the original paper
125 of QPD, providing a test win percentage table of median and mean performance at the end of the

Table A1: Hyperparameters of SeCA in SMAC.

Hyperparameters	#	Description
hidden units	64	Hidden units number for policy and critic network
batch size b	32	The number of transitions for each update
parallel runners	32	Number of environments to run in parallel
discount factor γ	0.99	The importance of future rewards
λ in TD(λ)	0.8	TD(λ) parameter for critic training
entropy regularization ξ	0.005	Weight or regularization for exploration
policy network’s initial lr	0.0025	Initial learning rate for policy network
critic network’s initial lr	0.0005	Initial learning rate for critic network
each IG step number	5	Summation of # intervals to approximate integrated gradients
target critic update frequency	200	Target network updates every # gradient steps
test interval	320000	Test the model every # steps
test episode number	32	Number of episodes to test for

126 training period. We use results from the SMAC paper [5] and the original QPD paper [8] because
 127 these reports show higher performance than the original works [6, 2, 4] and our implementation on
 128 QPD. Table A2 shows the evaluation results, where \tilde{m} is the median win percentage and \bar{m} is the
 129 mean win percentage. In general, the median performance is more persuasive as it avoids the effect
 130 of any outliers [5]. We could see from Table A2 that our method SeCA achieves state-of-the-art
 131 performances on all these scenarios that mentioned in the original QPD paper, demonstrating our
 132 improvement on QPD that also utilizes integrated gradients for the credit assignment problem.

Table A2: Median and mean performance of the test win percentage.

Map	IQL		COMA		QMIX		QTRAN		QPD		SeCA (Ours)	
	\tilde{m}	\bar{m}	\tilde{m}	\bar{m}	\tilde{m}	\bar{m}	\tilde{m}	\bar{m}	\tilde{m}	\bar{m}	\tilde{m}	\bar{m}
3m	100	97	91	92	100	99	100	100	95	92	100	100
8m	91	90	95	94	100	96	100	97	94	93	100	100
2s3z	39	42	66	64	100	97	77	80	95	94	100	99
3s5z	0	3	0	0	16	25	0	4	85	81	98	97
1c3s5z	7	8	30	30	89	89	31	33	92	92	100	100

133 E Visualization

134 We visualize two 3s_vs_5z battles and show the learned sequences to provide insights into our
 135 sequence adjustment result. We focus on this map because it contains only three agents, and its
 136 sequence analysis would be clear and illuminating. The allied Stalkers must master *kiting* technique
 137 and disperse in the map to kill the Zealots that chase them one after another. We number these three
 138 agents as Agent 1, Agent 2, and Agent 3 and print their positions at every step to follow their moves.
 139 The replays are available in the Visualization folder in our Supplementary Material.

140 Figure A1 shows the mini-maps of representative moments in one testing episode. Agent 3 only
 141 attracts one Zealot in this battle and kills it quickly, while Agent 1 and Agent 2 kite two enemies.
 142 After defeating the following Zealot, Agent 3 goes supporting alliances and observes Agent 1 in
 143 danger. Agent 3 blocks the enemies’ way to shelter Agent 1, and they kill the following two Zealots
 144 together. Agent 2, who fights alone, defeats both of the following enemies and does not show in its
 145 alliances’ field of view. The learned sequence of this episode is <Agent 3, Agent 2, Agent 1>.

146 We visualize another testing episode of 3s_vs_5z in Figure A2. Agent 3 also kites only one Zealot,
 147 and both Agent 1 and Agent 2 attract two enemies. When going to rescue alliances, Agent 3 misses
 148 Agent 1 this time and observes Agent 2 afterward. Then Agent 3 support Agent 2, who already killed
 149 a Zealot but has low health. However, Agent 1, who fights alone, only defeats one enemy and is
 150 killed by another Zealot. This survived Zealot is defeated by Agent 3, who protects Agent 2 at death’s
 151 door at last. The credit assignment sequence learned for this episode is <Agent 3, Agent 1, Agent 2>.



Figure A1: Some critical moments in a 3s_vs_5z battle. The big red squares represent Stalkers controlled by the agents, and each small blue square is an enemy Zealot. Agents win this battle by scattering and kiting the melee enemy units. Agent 3 kites only one Zealot and kills it quickly, and then it goes to the top of the map, protecting Agent 1 nearby. Agent 3 supports Agent 1 in defeating two following Zealots, while Agent 2 kites and kills two enemies by itself.

152 Although winning these two battles similarly, SeCA learns two different sequences in this map,
 153 illustrating that our dynamic adjustment algorithm (IG-episode) decides the sequence based on the
 154 battle's real-time circumstances. We find some similarities in these two learned sequences: Agent
 155 3 that supports and protects alliances in both battles, is arranged at the front, while the agent who
 156 receives backup and flees to keep alive is always the last in the sequence. These similarities are
 157 consistent with our common understanding of team contribution. Thus, these two visualizations
 158 demonstrate the rationality of our methods in practice.



Figure A2: Another episode in 3s_vs_5z. Similar to Figure A1, Agent 3 in this map still kites only one Zealot and kills it quickly. However, Agent 3 misses Agent 1 this time and supports Agent 2 in defeating the following Zealot. Agent 1 kills one enemy and does not survive from the other one, which is killed by Agent 3 later. Agent 2, with low health, choose to run away to keep alive.

159 F Additional Statement

160 We are very sorry for a typo in our main paper that may affect reviewers' understanding of our work
 161 and correct it here. We intended to show the objective to maximize when updating the policy network
 162 parameter θ in Equ.(11) in line 209 but gave a gradient form. Here we provide the correct version:

163 **Policy Learning.** We optimize each agent a 's policy parameter θ_a by maximizing the following
 164 objective, which contains our proposed advantage function and an entropy regularization term \mathcal{H} :

$$J^a(\theta) = \mathbb{E}_{\tau \sim \pi} [\log \pi^a(u^a | \tau^a) A^a(s, \mathbf{u}) + \mathcal{H}(\pi^a(\cdot | \tau^a))], \quad (\text{A7})$$

165 where the derivative of the adaptive entropy regularization term $\mathcal{H}(\pi^a(\cdot | \tau^a))$ [9] with respect to the
 166 i -th action probability p_i^a is given by:

$$d\mathcal{H}_i := -\xi \cdot (\log p_i^a + 1) / H(\pi^a(\cdot | \tau^a)), \text{ and } H(\pi^a(\cdot | \tau^a)) = \mathbb{E}_{u^a \sim \pi^a} [-\log \pi^a(u^a | \tau^a)]. \quad (\text{A8})$$

167 We share parameters among agents, and the gradient we use to train the actor shared by all agents is:

$$g = \mathbb{E}_{\tau \sim \pi} [\mathbb{E}_a [\nabla_{\theta_a} (\log \pi^a(u^a | \tau^a) A^a(s, \mathbf{u}) + \mathcal{H}(\pi^a(\cdot | \tau^a)))]]. \quad (\text{A9})$$

168 We sincerely apologize to the reviewers. Sorry for the inconvenience caused by this careless slip-up.

169 References

- 170 [1] Yali Du, Lei Han, Meng Fang, Ji Liu, Tianhong Dai, and Dacheng Tao. Liir: Learning individual
171 intrinsic reward in multi-agent reinforcement learning. In *Advances in Neural Information*
172 *Processing Systems*, pages 4403–4414, 2019.
- 173 [2] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson.
174 Counterfactual multi-agent policy gradients. In *AAAI Conference on Artificial Intelligence*, pages
175 2974–2982, 2018.
- 176 [3] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent
177 actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information*
178 *Processing Systems*, pages 6382–6393, 2017.
- 179 [4] Tabish Rashid, Gregory Farquhar, Bei Peng, and Shimon Whiteson. Weighted qmix: Expanding
180 monotonic value function factorisation. In *Advances in Neural Information Processing Systems*,
181 pages 10199–10210, 2020.
- 182 [5] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas
183 Nardelli, Tim G. J. Rudner, Chia-Man Hung, Philip H. S. Torr, Jakob Foerster, and Shimon
184 Whiteson. The starcraft multi-agent challenge. *CoRR*, abs/1902.04043, 2019.
- 185 [6] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Internat-*
186 *ional Conference on Machine Learning*, pages 330–337, 1993.
- 187 [7] Jianhong Wang, Yuan Zhang, Tae-Kyun Kim, and Yunjie Gu. Shapley q-value: A local reward
188 approach to solve global reward games. In *AAAI Conference on Artificial Intelligence*, pages
189 7285–7292, 2020.
- 190 [8] Yaodong Yang, Jianye Hao, Guangyong Chen, Hongyao Tang, Yingfeng Chen, Yujing Hu,
191 Changjie Fan, and Zhongyu Wei. Q-value path decomposition for deep multiagent reinforcement
192 learning. In *International Conference on Machine Learning*, pages 10706–10715, 2020.
- 193 [9] Meng Zhou, Ziyu Liu, Pengwei Sui, Yixuan Li, and Yuk Ying Chung. Learning implicit
194 credit assignment for cooperative multi-agent reinforcement learning. In *Advances in Neural*
195 *Information Processing Systems*, 2020.