
UneVEN: Universal Value Exploration for Multi-Agent Reinforcement Learning

Tarun Gupta¹ Anuj Mahajan¹ Bei Peng¹ Wendelin Böhmer² Shimon Whiteson¹

Abstract

VDN and QMIX are two popular value-based algorithms for cooperative MARL that learn a centralized action value function as a monotonic mixing of per-agent utilities. While this enables easy decentralization of the learned policy, the restricted joint action value function can prevent them from solving tasks that require significant coordination between agents at a given timestep. We show that this problem can be overcome by improving the *joint exploration* of all agents during training. Specifically, we propose a novel MARL approach called Universal Value Exploration (UneVEN) that learns a set of *related* tasks simultaneously with a linear decomposition of universal successor features. With the policies of already solved related tasks, the joint exploration process of all agents can be improved to help them achieve better coordination. Empirical results on a set of exploration games, challenging cooperative predator-prey tasks requiring significant coordination among agents, and StarCraft II micromanagement benchmarks show that UneVEN can solve tasks where other state-of-the-art MARL methods fail.

1. Introduction

Learning control policies for cooperative multi-agent reinforcement learning (MARL) remains challenging as agents must search the joint action space, which grows exponentially with the number of agents. Current state-of-the-art value-based methods such as VDN (Sunehag et al., 2017) and QMIX (Rashid et al., 2020b) learn a *centralized* joint action value function as a *monotonic* factorization of *decentralized* agent utility functions. Due to this monotonic

factorization, the joint action value function can be decentralizedly maximized as each agent can simply select the action that maximizes its corresponding utility function, known as the Individual Global Maximum principle (IGM, Son et al., 2019).

This monotonic restriction cannot represent the value of all joint actions and an agent’s utility depends on the policies of the other agents (nonmonotonicity, Mahajan et al., 2019). Even in collaborative tasks this can exhibit *relative overgeneralization* (RO, Panait et al., 2006), when the optimal action’s utility falls below that of a suboptimal action (Wei et al., 2018; 2019). While this pathology depends in practice on the agents’ random experiences, we show in Section 2 that in expectation RO prevents VDN from learning a large set of predator-prey games during the critical phase of uniform exploration.

QTRAN (Son et al., 2019) and WQMIX (Rashid et al., 2020a) show that this problem can be avoided by weighting the *joint actions of the optimal policy* higher. They propose to deduce this weighting from an unrestricted joint value function that is learned simultaneously. However, this unrestricted value is only a critic of the factored model, which itself is prone to RO, and often fails in practice due to insufficient ϵ -greedy exploration. MAVEN (Mahajan et al., 2019) improves exploration by learning an ensemble of monotonic joint action value functions through committed exploration and maximizing diversity in the joint team behavior. However, it does not specifically target optimal actions and may not work in tasks with strong RO.

The core idea of this paper is that even when a *target task* exhibits RO under value factorization, there may be *similar tasks* that do not. If their optimal actions overlap in some states with the target task, executing these simpler tasks can implicitly weight exploration to overcome RO. We call this novel paradigm Universal Value Exploration (UneVEN). To learn different MARL tasks simultaneously, we extend Universal Successor Features (USFs, Borsa et al., 2018) to Multi-Agent USFs (MAUSFs), using a VDN decomposition. During execution, UneVEN samples task descriptions from a Gaussian distribution centered around the target and executes the task with the highest value. This biases exploration towards the optimal actions of tasks that are similar but have

¹Department of Computer Science, University of Oxford, Oxford, United Kingdom ²Department of Software Technology, Delft University of Technology, Delft, Netherlands. Correspondence to: Tarun Gupta <tarun.gupta@cs.ox.ac.uk>.




Figure 1. Simplified predator-prey environment (left), where two agents are only rewarded when they both stand next to prey (right).

already been solved. Sampling these reduces RO for every task that shares the same optimal action and thus increases the number of solvable tasks, eventually overcoming RO for the target task, too. This is different from exploration methods like MAVEN, which keep the task same but reweigh explorative behaviours using the task returns. We show in Section 2 on a classic RO example that UneVEN can gradually increase the set of *solvable tasks* until it includes the target task.

We evaluate our novel approach against several ablations in predator-prey tasks that require significant coordination amongst agents and highlight the RO pathology, as well as other commonly used benchmarks like StarCraft II (Samvelyan et al., 2019). We also show empirically that UneVEN significantly outperforms current state-of-the-art value-based methods on target tasks that exhibit strong RO and in zero-shot generalization (Borsa et al., 2018) to other reward functions.

2. Illustrative Example

Figure 1 sketches a simplified predator-prey task where two agents (blue and magenta) can both move left or right (L/R) and execute a special ‘catch’ (C) action when they both stand next to the stationary prey (red). The agents are collaboratively rewarded (shown on the right of Figure 1) $+r$ when they catch the prey together and punished $-p$ if they attempt it alone, both ending the episode. For large p , both VDN and QMIX can lead to relative overgeneralization (RO, Panait et al., 2006) in the rewarded state s when agent 1’s utility $Q^1(s, u^1)$ of the catch action $u^1 = C$ drops below that of the movement actions $u^1 \in \{L, R\}$. At the beginning of training, when the value estimates are near zero and both agents explore random actions, we have:

$$Q^1(s, C) < Q^1(s, L/R) \Rightarrow r < p \left(\frac{1}{\pi^2(C|s)} - 2 \right) + c,$$

where c is a constant that depends mainly on future values. See Appendix A for a formal derivation. The threshold at which p yields RO depends strongly on agent 2’s probability of choosing $u^2 = C$, i.e., $\pi^2(C|s)$. For uniform exploration, this criterion is fulfilled if $p > r$, but reinforcing other actions than $u^2 = C$ can lower this threshold significantly. However, if agent 2 chooses $u^2 = C$ for more than half of its actions, no amount of punishment can prevent learning




Figure 2. Task space of Figure 1. Tasks solvable under uniform exploration ($\pi^2(C|s) = \frac{1}{3}$) are green, shades of red represent tasks that on average exhibit RO for different $\pi^2(C|s)$.

of the correct greedy policy.

Figure 2 plots the entire task space w.r.t. p and r , marking the set of tasks solvable under uniform exploration (green area) and tasks exhibiting RO on average for varying $\pi^2(C|s)$. MAUSFs uses a VDN decomposition of successor features to learn all tasks in the black circle simultaneously, which initially can only solve monotonic tasks in the green area. UneVEN changes this by evaluating random tasks (blue dots), and exploring those already solved (magenta dots). This increases the fraction of observed $u^2 = C$, and thus the set of solvable tasks, which eventually reaches the *target task* (cross).

3. Background

A fully cooperative multi-agent task can be formalized as a *decentralized partially observable Markov decision process* (Dec-POMDP, Oliehoek et al., 2016) consisting of a tuple $G = \langle \mathcal{S}, \mathcal{U}, P, R, \Omega, O, n, \gamma \rangle$. $s \in \mathcal{S}$ describes the true state of the environment. At each time step, each agent $a \in \mathcal{A} \equiv \{1, \dots, n\}$ chooses an action $u^a \in \mathcal{U}$, forming a joint action $\mathbf{u} \in \mathcal{U} \equiv \mathcal{U}^n$. This causes a transition in the environment according to the state transition kernel $P(s'|s, \mathbf{u}) : \mathcal{S} \times \mathcal{U} \times \mathcal{S} \rightarrow [0, 1]$. All agents are collaborative and share therefore the same reward function $R(s, \mathbf{u}) : \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}$. $\gamma \in [0, 1)$ is a discount factor.

Due to *partial observability*, each agent a cannot observe the true state s , but receives an observation $o^a \in \Omega$ drawn from observation kernel $o^a \sim O(s, a)$. At time t , each agent a has access to its action-observation history $\tau_t^a \in \mathcal{T}_t \equiv (\Omega \times \mathcal{U})^t \times \Omega$, on which it conditions a stochastic policy $\pi^a(u_t^a | \tau_t^a)$. $\tau_t \in \mathcal{T}_t^n$ denotes the histories of all agents. The joint stochastic policy $\pi(\mathbf{u}_t | s_t, \tau_t) \equiv \prod_{a=1}^n \pi^a(u_t^a | \tau_t^a)$ induces a joint action value function: $Q^\pi(s_t, \tau_t, \mathbf{u}_t) = \mathbb{E}[G_t | s_t, \tau_t, \mathbf{u}_t]$, where $G_t = \sum_{i=0}^{\infty} \gamma^i r_{t+i}$ is the *discounted return*.

CTDE: We adopt the framework of *centralized training and decentralized execution* (CTDE, Kraemer & Banerjee, 2016), which assumes access to all action-observation histories τ_t and global state s_t during training, but each

agent’s decentralized policy π^a can only condition on its own action-observation history τ^a . This approach can exploit information that is not available during execution and also freely share parameters and gradients, which improves the sample efficiency (see e.g., Foerster et al., 2018; Rashid et al., 2020b; Böhmer et al., 2020).

Value Decomposition Networks: A naive way to learn in MARL is *independent Q-learning* (IQL, Tan, 1993), which learns an independent action-value function $Q^a(\tau_t^a, u_t^a; \theta^a)$ for each agent a that conditions only on its local action-observation history τ_t^a . To make better use of other agents’ information in CTDE, *value decomposition networks* (VDN, Sunehag et al., 2017) represent the joint action value function Q_{tot} as a sum of per-agent *utility functions* Q^a : $Q_{tot}(\tau, \mathbf{u}; \theta) \equiv \sum_{a=1}^n Q^a(\tau^a, \mathbf{u}^a; \theta)$. Each Q^a still conditions only on individual action-observation histories and can be represented by an agent network that shares parameters across all agents. The joint action-value function Q_{tot} can be trained using Deep Q-Networks (DQN, Mnih et al., 2015). Unlike VDN, QMIX (Rashid et al., 2020b) represents the joint action-value function Q_{tot} with a nonlinear *monotonic* combination of individual utility functions. The greedy joint action in both VDN and QMIX can be computed in a decentralized fashion by individually maximizing each agent’s utility. See OroojlooyJadid & Hajinezhad (2019) for a more in-depth overview of cooperative deep MARL.

Task based Universal Value Functions: In this paper, we consider tasks that differ only in their reward functions $R_w(s, \mathbf{u}) \equiv \mathbf{w}^\top \phi(s, \mathbf{u})$, which are linear combinations of a set of basis functions $\phi: \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}^d$. Intuitively, the basis functions ϕ encode potentially rewarded events, such as opening a door or picking up an object. We use the weight vector \mathbf{w} to denote the task with reward function R_w . Universal Value Functions (UVFs, Schaul et al., 2015) extend DQN to learn a *generalizable* value function conditioned on tasks. UVFs are typically of the form $Q^\pi(s_t, \mathbf{u}_t, \mathbf{w})$ to indicate the action-value function of task \mathbf{w} under policy π at time t as:

$$\begin{aligned} Q^\pi(s_t, \mathbf{u}_t, \mathbf{w}) &= \mathbb{E}^\pi \left[\sum_{i=0}^{\infty} \gamma^i R_w(s_{t+i}, \mathbf{u}_{t+i}) \mid s_t, \mathbf{u}_t \right] \\ &= \mathbb{E}^\pi \left[\sum_{i=0}^{\infty} \gamma^i \phi(s_{t+i}, \mathbf{u}_{t+i})^\top \mathbf{w} \mid s_t, \mathbf{u}_t \right]. \end{aligned} \quad (1)$$

Successor Features: The Successor Representation (Dayan, 1993) has been widely used in single-agent settings (Barreto et al., 2017; 2018; Borsa et al., 2018) to generalize across tasks with given reward specifications. By simply rewriting the definition of the action value function $Q^\pi(s_t, \mathbf{u}_t, \mathbf{w})$ of task \mathbf{w} from Equation 1 we have:

$$\begin{aligned} Q^\pi(s_t, \mathbf{u}_t, \mathbf{w}) &= \mathbb{E}^\pi \left[\sum_{i=0}^{\infty} \gamma^i \phi(s_{t+i}, \mathbf{u}_{t+i}) \mid s_t, \mathbf{u}_t \right]^\top \mathbf{w} \\ &\equiv \boldsymbol{\psi}^\pi(s_t, \mathbf{u}_t)^\top \mathbf{w}, \end{aligned} \quad (2)$$

where $\boldsymbol{\psi}^\pi(s, \mathbf{u})$ are the Successor Features (SFs) under policy π . For the optimal policy π_z^* of task z , the SFs $\boldsymbol{\psi}^{\pi_z^*}$ summarize the dynamics under this policy, which can then be weighted with any reward vector $\mathbf{w} \in \mathbb{R}^d$ to instantly evaluate policy π_z^* on it: $Q^{\pi_z^*}(s, \mathbf{u}, \mathbf{w}) = \boldsymbol{\psi}^{\pi_z^*}(s, \mathbf{u})^\top \mathbf{w}$.

Universal Successor Features and Generalized Policy Improvement: Borsa et al. (2018) introduce universal successor features (USFs) that learn SFs conditioned on tasks using the *generalization* power of UVFs. Specifically, they define UVFs of the form $Q(s, \mathbf{u}, z, \mathbf{w})$ which represents the value function of policy π_z evaluated on task $\mathbf{w} \in \mathbb{R}^d$. These UVFs can be factored using the SFs property (Equation 2) as: $Q(s, \mathbf{u}, z, \mathbf{w}) = \boldsymbol{\psi}(s, \mathbf{u}, z)^\top \mathbf{w}$, where $\boldsymbol{\psi}(s, \mathbf{u}, z)$ are the USFs that generate the SFs induced by task-specific policy π_z . One major advantage of using SFs is the ability to *efficiently* do generalized policy improvement (GPI, Barreto et al., 2017), which allows a new policy to be computed for *any unseen* task based on instant policy evaluation of a *set* of policies on that unseen task with a simple dot-product. Formally, given a set $\mathcal{C} \subseteq \mathbb{R}^d$ of tasks and their corresponding SFs $\{\boldsymbol{\psi}(s, \mathbf{u}, z)\}_{z \in \mathcal{C}}$ induced by corresponding policies $\{\pi_z\}_{z \in \mathcal{C}}$, a new policy π'_w for any unseen task $\mathbf{w} \in \mathbb{R}^d$ can be derived using:

$$\begin{aligned} \pi'_w(s) &\in \arg \max_{\mathbf{u} \in \mathcal{U}} \max_{z \in \mathcal{C}} Q(s, \mathbf{u}, z, \mathbf{w}) \\ &\in \arg \max_{\mathbf{u} \in \mathcal{U}} \max_{z \in \mathcal{C}} \boldsymbol{\psi}(s, \mathbf{u}, z)^\top \mathbf{w}. \end{aligned} \quad (3)$$

Setting $\mathcal{C} = \{\mathbf{w}\}$ allows us to revert back to UVFs, as we evaluate SFs induced by policy π_w on task \mathbf{w} itself. However, we can use any set of tasks that are similar to \mathbf{w} based on some similarity distribution $\mathcal{D}(\cdot | \mathbf{w})$. The computed policy π'_w is guaranteed to perform no worse on task \mathbf{w} than *each* of the policies $\{\pi_z\}_{z \in \mathcal{C}}$ (Barreto et al., 2017), but often performs much better. SFs thus enable efficient use of GPI, which allows *reuse* of learned knowledge for zero-shot generalization.

4. Multi-Agent Universal Successor Features

In this section, we introduce Multi-Agent Universal Successor Features (MAUSFs), extending single-agent USFs (Borsa et al., 2018) to multi-agent settings and show how we can learn generalized *decentralized* greedy policies for agents. The USFs based centralized joint action value function $Q_{tot}(\tau, \mathbf{u}, z, \mathbf{w})$ allows evaluation of joint policy $\pi_z = \langle \pi_z^1, \dots, \pi_z^n \rangle$ comprised of local agent policies π_z^a of the *same* task z on task \mathbf{w} . However, each agent a may execute a different policy $\pi_{z^a}^a$ of different task $z^a \in \mathcal{C}$, resulting in a combinatorial set of joint policies. Maximizing over all combinations $\bar{z} \equiv \langle z^1, \dots, z^n \rangle \in \mathcal{C}^n$ should therefore enormously improve GPI. To enable this flexibility, we define the joint action-value function (Q_{tot}) of joint policy $\pi_{\bar{z}} = \{\pi_{z^a}^a\}_{z^a \in \mathcal{C}}$ evaluated on any task $\mathbf{w} \in \mathbb{R}^d$ as:




Figure 3. Schematic illustration of the MAUSFs training and UneVEN exploration with GPI policy.

$Q_{tot}(\tau, \mathbf{u}, \bar{z}, \mathbf{w}) = \psi_{tot}(\tau, \mathbf{u}, \bar{z})^\top \mathbf{w}$, where $\psi_{tot}(\tau, \mathbf{u}, \bar{z})$ are the MAUSFs of (τ, \mathbf{u}) summarizing the joint dynamics of the environment under joint policy $\pi_{\bar{z}}$. However, training centralized MAUSFs and using centralized GPI to maximize over a combinatorial space of \bar{z} becomes impractical when there are more than a handful of agents, since the joint action space (\mathcal{U}) and joint task space (\mathcal{C}^n) of the agents grows exponentially with the number of agents. To leverage CTDE and enable decentralized execution by agents, we therefore propose novel *agent-specific SFs* for each agent a following local policy $\pi_{z^a}^a$, which condition only on its own local action-observation history and task z^a .

Decentralized Execution: We define local *utility* functions for each agent a as $Q^a(\tau^a, u^a, z^a, \mathbf{w}) = \psi^a(\tau^a, u^a, z^a; \theta)^\top \mathbf{w}$, where $\psi^a(\tau^a, u^a, z^a; \theta)$ are the local agent-specific SFs induced by local policy $\pi_{z^a}^a(u^a|\tau^a)$ of agent a sharing parameters θ . Intuitively, $Q^a(\tau^a, u^a, z^a, \mathbf{w})$ is the utility function for agent a when local policy $\pi_{z^a}^a(u^a|\tau^a)$ of task z^a is executed on task \mathbf{w} . We use VDN decomposition to represent MAUSFs ψ_{tot} as a sum of local agent-specific SFs for each agent a :

$$\begin{aligned} Q_{tot}(\tau, \mathbf{u}, \bar{z}, \mathbf{w}) &= \sum_{a=1}^n Q^a(\tau^a, u^a, z^a, \mathbf{w}) \\ &= \underbrace{\sum_{a=1}^n \psi^a(\tau^a, u^a, z^a; \theta)^\top \mathbf{w}}_{\psi_{tot}(\tau, \mathbf{u}, \bar{z}; \theta)} \end{aligned} \quad (4)$$

We can now learn local agent-specific SFs ψ^a for each agent a that can be instantly weighted with any task vector $\mathbf{w} \in \mathbb{R}^d$ to generate local utility functions Q^a , thereby allowing agents to use the GPI policy in a decentralized manner.

Decentralized Local GPI: Our novel agent-specific SFs allow each agent a to locally perform decentralized GPI by instant policy evaluation of a set \mathcal{C} of local task policies

$\{\pi_{z^a}^a\}_{z^a \in \mathcal{C}}$ on any unseen task \mathbf{w} to compute a local GPI policy. Due to the linearity of the VDN decomposition, this is equivalent to maximization over all combinations of $\bar{z} \equiv \langle z^1, \dots, z^n \rangle \in \mathcal{C} \times \dots \times \mathcal{C} \equiv \mathcal{C}^n$ as:

$$\begin{aligned} \pi_{\mathbf{w}}^a(\tau) &\in \arg \max_{\mathbf{u} \in \mathcal{U}} \max_{\bar{z} \in \mathcal{C}^n} Q_{tot}(\tau, \mathbf{u}, \bar{z}, \mathbf{w}) \\ &\in \left\{ \arg \max_{u^a \in \mathcal{U}} \max_{z^a \in \mathcal{C}} \psi^a(\tau^a, u^a, z^a; \theta)^\top \mathbf{w} \right\}_{a=1}^n. \end{aligned} \quad (5)$$

As all of the above relies on the linearity of the VDN decomposition, it cannot be directly applied to nonlinear mixing techniques like QMIX (Rashid et al., 2020b).

Training: MAUSFs for task combination \bar{z} are trained end-to-end by gradient descent on the loss:

$$\begin{aligned} \mathcal{L}(\theta, \bar{z}) &= \mathbb{E}_{\mathcal{B}} \left[\left\| \phi(s_t, \mathbf{u}_t) + \gamma \psi_{tot}(\tau_{t+1}, \mathbf{u}'_{\bar{z}}, \bar{z}; \theta^-) \right. \right. \\ &\quad \left. \left. - \psi_{tot}(\tau_t, \mathbf{u}_t, \bar{z}; \theta) \right\|_2^2 \right], \end{aligned} \quad (6)$$

where the expectation is over a minibatch of samples $\{(s_t, \mathbf{u}_t, \tau_t)\}$ from the replay buffer \mathcal{B} (Lin, 1992), θ^- denotes the parameters of a target network (Mnih et al., 2015) and joint actions $\mathbf{u}'_{\bar{z}} = \{u'_{z^a}\}_{a=1}^n$ are selected individually by each agent network using the current parameters θ (called Double Q-learning, van Hasselt et al., 2016): $u'_{z^a} = \arg \max_{u \in \mathcal{U}} \psi^a(\tau_{t+1}^a, u, z^a; \theta)^\top z^a$. Each agent learns therefore local agent-specific SFs $\psi^a(\tau^a, u^a, z; \theta)$ by gradient descent on $\mathcal{L}(\theta, \bar{z})$ for all $z \in \mathcal{C} \equiv \nu \cup \{w\}$, where $\nu \sim \mathcal{D}(\cdot|w)$ is drawn from a distance measure around target task w . The green region of Figure 3 shows a CTDE based architecture to train MAUSFs for a given target task w . A detailed algorithm is presented in Appendix B.

5. UneVEN

In this section, we present UneVEN (red region of Figure 3), which leverages MAUSFs and decentralized GPI to help

overcome relative overgeneralization on the target task w . At the beginning of every exploration episode, we sample a set of related tasks $\nu = \{z \sim \mathcal{D}(\cdot|w)\}$, containing potentially simpler reward functions, from a distribution \mathcal{D} around the target task. The basic idea is that *some* of these related tasks can be efficiently learned with a factored value function. These tasks are therefore solved early and exploration concentrates on the state-action pairs that are useful to them. If other tasks close to those already solved have the same optimal actions, this implicit weighting allows to solve them too (shown by Rashid et al., 2020a). Furthermore, tasks closer to w are sampled more frequently, which biases the process to eventually overcome relative overgeneralization on the target task itself.

Our method assumes that the basis functions ϕ and the reward-weights w of the target task are known, but Barreto et al. (2020) show that both can be learned using multi-task regression. Many choices for \mathcal{D} are possible, but in the following we sample related tasks using a normal distribution centered around the target task $w \in \mathbb{R}^d$ with a small variance σ as $\mathcal{D} = \mathcal{N}(w, \sigma \mathbf{I}_d)$. This samples similar tasks closer to w more frequently. As long as σ is large enough to cover tasks that do not induce RO (see Figure 2), our method works well and therefore, does not rely on any domain knowledge. The resulting task vectors weight the basis functions ϕ differently and represent different reward functions. In particular the varied reward functions can make these tasks much easier, but also harder, to solve with monotonic value functions. The consequences of sampling harder tasks on learning are discussed with the below action-selection schemes.

Action-Selection Schemes: UneVEN uses two novel schemes to enable action selection based on related tasks. To emphasize the importance of the target task, we define a probability α of selecting actions based on the target task. Therefore, with probability $1 - \alpha$, the action is selected based on the related task. Similar to other exploration schemes, α is annealed from 0.3 to 1.0 in our experiments over a fixed number of steps at the beginning of training. Once this exploration stage is finished (i.e., $\alpha = 1$), actions are always taken based on the target task’s joint action value function. Each action-selection scheme employs a local decentralized GPI policy, which maximizes over a set of policies π_z based on $z \in \mathcal{C}_1$ (also referred to as the *evaluation* set) to estimate the Q -values of another set of tasks $k \in \mathcal{C}_2$ (also referred to as the *target* set) using:

$$u_t = \left\{ u_t^a = \arg \max_{u \in \mathcal{U}} \max_{k \in \mathcal{C}_2} \max_{z \in \mathcal{C}_1} \overbrace{\psi^a(\tau_t^a, u, z; \theta)^\top k}^{Q^a(\tau_t^a, u, z, k)} \right\}_{a \in \mathcal{A}}. \quad (7)$$

Here $\mathcal{C}_1 = \nu \cup \{w\}$ is the set of target and related tasks that induce the policies that are evaluated (dot-product) on the

set of tasks \mathcal{C}_2 , which varies with different action-selection schemes. The red box in Figure 3 illustrates UneVEN exploration. For example, Q -learning always picks actions based on the target task, i.e., the target set $\mathcal{C}_2 = \{w\}$. However, this scheme does not favour important joint actions. We call this default action-selection scheme **target GPI** and execute it with probability α . We now propose two novel action-selection schemes based on related tasks with probability $1 - \alpha$, and thereby implicitly weighting joint actions during learning.

Uniform GPI: At the beginning of each episode, this action-selection scheme uniformly picks *one* related task, i.e., the target set $\mathcal{C}_2 = \{k \sim \text{Uniform}(\nu)\}$, and selects actions based on that task using the GPI policy throughout the episode. This uniform task selection explores the learned policies of all related tasks in \mathcal{D} . This works well in practice as there are often enough simpler tasks to induce the required bias over important joint actions. However, if the sampled related task is harder than the target task, the action-selection based on these harder tasks might hurt learning on the target task and lead to higher variance during training.

Greedy GPI: At every time-step t , this action-selection scheme picks the task $k \in \nu \cup \{w\}$ that gives the highest Q -value amongst the related and target tasks, i.e., the target set becomes $\mathcal{C}_2 = \nu \cup \{w\}$. Due to the greedy nature of this action-selection scheme, exploration is biased towards solved tasks, as those have larger values. We are thus exploring the solutions of tasks that are both solvable and similar to the target task w , which should at least in some states lead to the same *optimal joint actions* as w .

NO-GPI: To demonstrate the influence of GPI on the above schemes, we also investigate ablations, where we define the evaluation set $\mathcal{C}_1 = \{k\}$ to only contain the currently estimated task k , i.e., using $u_t = \{u_t^a = \arg \max_{u \in \mathcal{U}} \max_{k \in \mathcal{C}_2} \psi^a(\tau_t^a, u, k; \theta)^\top k\}_{a \in \mathcal{A}}$ for action selection.

6. Experiments

In this section, we evaluate UneVEN on a variety of complex domains. For evaluation, all experiments are carried out with six random seeds and results are shown with \pm standard error across seeds. We compare our method against a number of SOTA value-based MARL approaches: IQL (Tan, 1993), VDN (Sunehag et al., 2017), QMIX (Rashid et al., 2020b), MAVEN (Mahajan et al., 2019), WQMIX (Rashid et al., 2020a), QTRAN (Son et al., 2019), and QPLEX (Wang et al., 2020a).

Domain 1 : m -Step Matrix Game

We first evaluate UneVEN on the m -step matrix game proposed by Mahajan et al. (2019). This task is difficult to

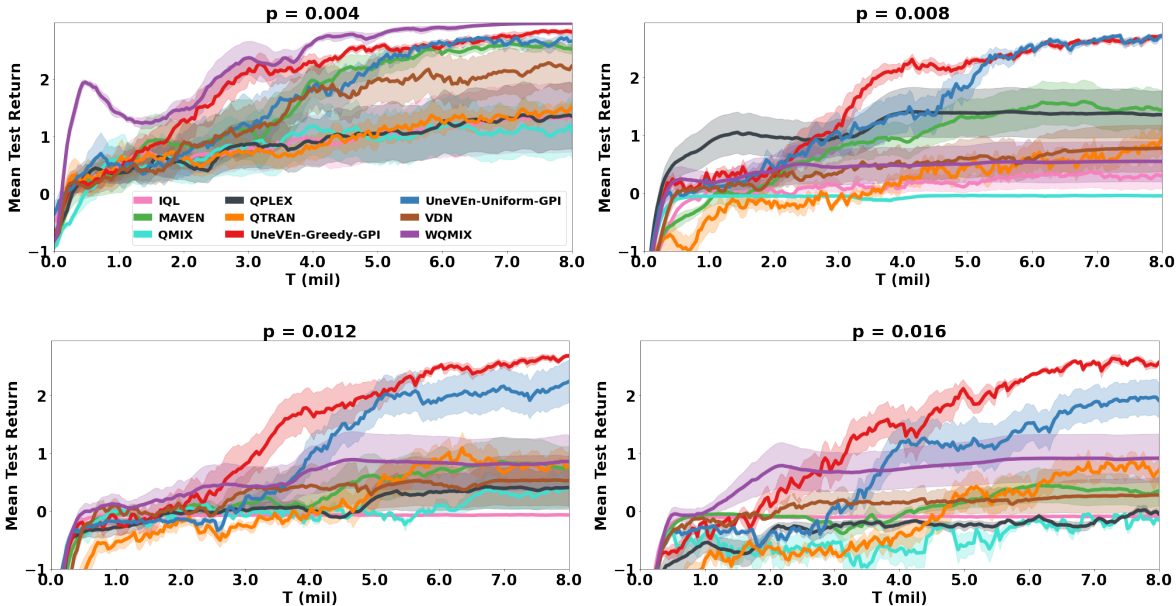


Figure 5. Comparison between UneVEN and SOTA MARL baselines with $p \in \{0.004, 0.008, 0.012, 0.016\}$.

monotonic action-value functions, but does not concentrate on joint actions that would overcome RO if explored more.

Ablations: Figure 7 shows ablation results for higher penalty tasks, i.e., $p = \{0.012, 0.016\}$. To contrast the effect of UneVEN on exploration, we compare our two novel action-selection schemes to UneVEN-Target-GPI, which only selects the greedy actions of the target task. The results clearly show that UneVEN-Target-GPI fails to solve the higher penalty RO tasks as the employed monotonic joint value function of the target task fails to accurately represent the values of different joint actions. This demonstrates the critical role of UneVEN and its action-selection schemes.

Next we evaluate the effect of GPI by comparing against UneVEN with MAUSFs without using the GPI policy, i.e., setting the evaluation set $\mathcal{C}_1 = \{k\}$ in Equation 7. First, UneVEN using a NOGPI policy with both uniform (Uniform-NOGPI) and greedy (Greedy-NOGPI) action selection out-

performs Target-NOGPI, further strengthening the claim that UneVEN with its novel action-selection scheme enables efficient exploration and bias towards optimal joint actions. In addition, the left and middle plots of Figure 7 shows that for each corresponding action-selection scheme (uniform, greedy, and target), using a GPI policy (*-GPI) is consistently favourable as it performs either similarly to the NOGPI policy (*-NOGPI) or much better. GPI appears to improve zero-shot generalization of MAUSFs across tasks, which in turn enables good action selection for related tasks during UneVEN exploration.

Zero-Shot Generalization: Lastly, we evaluate this zero-shot generalization for all methods to check if the learnt policies are useful for unseen high penalty test tasks. We train all methods for 8 million environmental steps on a task with $p = 0.004$, and test 60 rollouts of the resulting policies of all methods that are able to solve the training task, i.e., UneVEN-Greedy-GPI, UneVEN-Uniform-GPI, VDN, MAVEN, and WQMIX, on tasks with $p \in \{0.2, 0.5\}$. For policies trained with UneVEN-Greedy-GPI and UneVEN-Uniform-GPI, we use the NOGPI policy for the zero-shot testing, i.e., $\mathcal{C}_1 = \mathcal{C}_2 = \{w\}$. The right plot of Figure 7 shows that UneVEN with both uniform and greedy schemes exhibits great zero-shot generalization and solves both test tasks even with very high penalties. As MAUSFs learn the reward’s basis functions, rather than the reward itself, zero-shot generalization to larger penalties follows naturally. Furthermore, using UneVEN exploration allows the agents to collect enough diverse behaviour to come up with a near optimal policy for the test tasks. On the other hand, the




Figure 6. Baseline predator-prey results without RO ($p = 0$).




Figure 7. Ablation results. (left, middle): Comparison between different action selection schemes of UneVEN for $p \in \{0.012, 0.016\}$. (right): Zero-shot generalization comparison; training on $p = 0.004$, testing on $p \in \{0.2, 0.5\}$.




Figure 8. Comparison between UneVEN and SOTA MARL baselines on SMAC maps with $p = 1.0$.

learnt policies for all other methods that solve the target task with $p = 0.004$ are ineffective in these higher penalty RO tasks, as they do not learn to avoid unsuccessful capture attempts. See Appendix D for details about the implementations and Appendix E for additional ablation experiments.

Domain 3 : Starcraft Multi-Agent Challenge (SMAC)

We now evaluate UneVEN on challenging cooperative StarCraft II (SC2) maps from the popular SMAC benchmark (Samvelyan et al., 2019). Our method aims to overcome relative overgeneralization (RO) which happens very often in practice with cooperative games (Wei & Luke, 2016). However, the default reward function in SMAC does not suffer from RO as it has been designed with QMIX in mind, deliberately dropping any punishment for loosing ally units and thereby making it solvable for all considered baselines. We therefore consider SMAC maps where each ally agent unit is additionally penalized (p) for being killed or suffering damage from the enemy, in addition to receiving positive reward for killing/damage on enemy units. A detailed discussion about different reward functions in SMAC along with their implications are discussed in Appendix E. The additional penalty to the reward function (similar to QTRAN++, Son et al. (2020)) for losing our own ally units induces RO in the same way as in our predator prey example. This makes maps that were classified as “easy” by default benchmark (e.g., “8m” in Fig. 8) very hard to solve for methods like VDN and QMIX for $p = 1.0$ (equally weighting the lives of ally and enemy units).

Prior work on SC2 has established that VDN/QMIX can

solve nearly all SMAC maps in the absence of any punishment (i.e., negative reward scaling $p = 0$) and we confirm this also holds for low punishments (see results in Figure 13 for $p = 0.5$ in Appendix E). However, this puts little weight on the lives of your own allies and cautious generals may want to learn a more conservative policy with higher punishment for losing their ally units. Similarly to the predator-prey example, Figure 8 shows that increasing the punishment to $p = 1$ (equally weighting the lives of allies and enemies) leads VDN/QMIX and other SOTA MARL methods like QTRAN, WQMIX and QPLEX to fail in many maps, whereas our novel method UneVEN, in particular with greedy action selection, reliably outperforms baselines on all tested SMAC maps.

7. Related Work

Improving monotonic value function factorization in CTDE: MAVEN (Mahajan et al., 2019) shows that the monotonic joint action value function of QMIX and VDN suffers from suboptimal approximations on nonmonotonic tasks. It addresses this problem by learning a diverse ensemble of monotonic joint action-value functions conditioned on a latent variable by optimizing the mutual information between the joint trajectory and the latent variable. Deep Coordination Graphs (DCG) (Böhmer et al., 2020) uses a predefined coordination graph (Guestrin et al., 2002) to represent the joint action value function. However, DCG is not a fully decentralized approach and specifying the coordination graph can require significant domain knowledge. Son et al. (2019) propose QTRAN, which addresses

the monotonic restriction of QMIX by learning a (decentralizable) VDN-factored joint action value function along with an unrestricted centralized critic. The corresponding utility functions are distilled from the critic by solving a linear optimization problem involving all joint actions, but its exact implementation is computationally intractable and the corresponding approximate versions have unstable performance. TESSERACT (Mahajan et al., 2021) proposes a tensor decomposition based method for learning arbitrary action-value functions, they also provide sample complexity analysis for their method. QPLEX (Wang et al., 2020a) uses a duplex dueling (Wang et al., 2016) network architecture to factorize the joint action value function with linear decomposition structure. WQMIX (Rashid et al., 2020a) learns a QMIX-factored joint action value function along with an unrestricted centralized critic and proposes explicit weighting mechanisms to bias the monotonic approximation of the optimal joint action value function towards important joint actions, which is similar to our work. However, in our work, the weightings are implicitly done through action-selection based on related tasks, which are easier to solve.

Exploration: There are many techniques for exploration in model-free single-agent RL, based on intrinsic novelty reward (Bellemare et al., 2016; Tang et al., 2017), predictability (Pathak et al., 2017), pure curiosity (Burda et al., 2019) or Bayesian posteriors (Osband et al., 2016; Gal et al., 2017; Fortunato et al., 2018; O’Donoghue et al., 2018). In the context of multi-agent RL, Böhmer et al. (2019) discuss the influence of unreliable intrinsic reward and Wang et al. (2020c) quantify the influence that agents have on each other’s return. Similarly, Wang et al. (2020b) propose a learnable action effect based role decomposition which eases exploration in the joint action space. Zheng & Yue (2018) propose to coordinate exploration between agents by shared latent variables, whereas Jaques et al. (2018) investigate the social motivations of competitive agents. However, these techniques aim to visit as much of the state-action space as possible, which exacerbates the relative overgeneralization pathology. Approaches that use state-action space abstraction can speed up exploration, these include those which can automatically learn the abstractions (e.g., Mahajan & Tulabandhula, 2017a;b) and those which require prior knowledge (e.g., Roderick et al., 2018), however they are difficult to scale for multi-agent scenarios. In contrast, UneVEN explores similar *tasks*. This guides exploration to states and actions that prove *useful*, which restricts the explored space and overcomes relative overgeneralization. To the best of our knowledge, the only other work that explores the task space is Leibo et al. (2019): they use the evolution of competing agents as an auto-curriculum of harder and harder tasks. Collaborative agents cannot compete against each other, though, and their approach therefore is not suitable to learn cooperation.

Successor Features: Most of the work on SFs have been focused on single-agent settings (Dayan, 1993; Kulkarni et al., 2016; Lehnert et al., 2017; Barreto et al., 2017; 2018; Borsa et al., 2018; Lee et al., 2019; Hansen et al., 2019) for transfer learning and zero-shot generalization across tasks with different reward functions. Gupta et al. (2019) use single-agent SFs in a multi-agent setting to estimate the probability of events, but they only consider transition independent MARL (Becker et al., 2004; Gupta et al., 2018).

8. Conclusion

This paper presents a novel approach decomposing multi-agent universal successor features (MAUSFs) into local agent-specific SFs, which enable a decentralized version of the GPI to maximize over a combinatorial space of agent policies. We propose UneVEN, which leverages the generalization power of MAUSFs to perform action selection based on simpler related tasks to address the suboptimality of the target task’s monotonic joint action value function in current SOTA methods. Our experiments show that UneVEN significantly outperforms VDN, QMIX and other state-of-the-art value-based MARL methods on challenging tasks exhibiting severe RO.

9. Future Work

UneVEN relies on the assumption that both the basis functions and the reward-weights of the target task are known which restricts the applicability of the method. Moreover, our SFs based method requires learning in d dimensions instead of learning scalar values for each action in the case of Q -learning, which decreases the scalability of our method for domains where d is very large. More efficient neural network architectures with hyper networks (Ha et al., 2016) can be leveraged to handle larger dimensional features. Finally, the paradigm of UneVEN with related task based action selection can be directly applied with universal value functions (UVFs) (Schaul et al., 2015) to enable other state-of-the-art nonlinear mixing techniques like QMIX (Rashid et al., 2020b) and QPLEX (Wang et al., 2020a) to overcome RO, at the cost of losing ability to perform GPI.

Acknowledgements

We thank Christian Schroeder de Witt, Tabish Rashid and other WhiRL members for their feedback. Tarun Gupta is supported by the Oxford University Clarendon Scholarship. Anuj Mahajan is funded by the J.P. Morgan A.I. fellowship. This project has received funding from the European Research Council under the European Unions Horizon 2020 research and innovation programme (grant agreement number 637713). The experiments were made possible by a generous equipment grant from NVIDIA.

References

- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H. P., and Silver, D. Successor features for transfer in reinforcement learning. In *Advances in neural information processing systems*, pp. 4055–4065, 2017.
- Barreto, A., Borsa, D., Quan, J., Schaul, T., Silver, D., Hessel, M., Mankowitz, D., Zidek, A., and Munos, R. Transfer in deep reinforcement learning using successor features and generalised policy improvement. In *International Conference on Machine Learning*, pp. 501–510. PMLR, 2018.
- Barreto, A., Hou, S., Borsa, D., Silver, D., and Precup, D. Fast reinforcement learning with generalized policy updates. *Proceedings of the National Academy of Sciences*, 2020.
- Becker, R., Zilberstein, S., Lesser, V., and Goldman, C. V. Solving transition independent decentralized markov decision processes. *Journal of Artificial Intelligence Research*, 22:423–455, 2004.
- Bellemare, M. G., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems (NIPS) 29*, pp. 1471–1479, 2016.
- Böhmer, W., Rashid, T., and Whiteson, S. Exploration with unreliable intrinsic reward in multi-agent reinforcement learning. *CoRR*, abs/1906.02138, 2019. URL <http://arxiv.org/abs/1906.02138>. Presented at the ICML Exploration in Reinforcement Learning workshop.
- Böhmer, W., Kurin, V., and Whiteson, S. Deep coordination graphs. In *Proceedings of Machine Learning and Systems (ICML)*, pp. 2611–2622, 2020. URL <https://arxiv.org/abs/1910.00091>.
- Borsa, D., Barreto, A., Quan, J., Mankowitz, D., Munos, R., van Hasselt, H., Silver, D., and Schaul, T. Universal successor features approximators. *arXiv preprint arXiv:1812.07626*, 2018.
- Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., and Efros, A. A. Large-scale study of curiosity-driven learning. In *International Conference on Learning Representations (ICLR)*, 2019.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Dayan, P. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5(4):613–624, 1993.
- Foerster, J. N., Farquhar, G., Afouras, T., Nardelli, N., and Whiteson, S. Counterfactual multi-agent policy gradients. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- Fortunato, M., Azar, M. G., Piot, B., Menick, J., Hessel, M., Osband, I., Graves, A., Mnih, V., Munos, R., Hassabis, D., Pietquin, O., Blundell, C., and Legg, S. Noisy networks for exploration. In *International Conference on Learning Representations (ICLR)*, 2018. URL <https://openreview.net/forum?id=rywHCPkAW>.
- Gal, Y., Hron, J., and Kendall, A. Concrete dropout. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 3584–3593, 2017.
- Guestrin, C., Lagoudakis, M., and Parr, R. Coordinated reinforcement learning. In *ICML*, volume 2, pp. 227–234. Citeseer, 2002.
- Gupta, T., Kumar, A., and Paruchuri, P. Planning and learning for decentralized mdps with event driven rewards. In *AAAI*, 2018.
- Gupta, T., Kumar, A., and Paruchuri, P. Successor features based multi-agent rl for event-based decentralized mdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 6054–6061, 2019.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- Hansen, S., Dabney, W., Barreto, A., Van de Wiele, T., Warde-Farley, D., and Mnih, V. Fast task inference with variational intrinsic successor features. *arXiv preprint arXiv:1906.05030*, 2019.
- Jaques, N., Lazaridou, A., Hughes, E., Gülçehre, Ç., Ortega, P. A., Strouse, D., Leibo, J. Z., and de Freitas, N. Intrinsic social motivation via causal influence in multi-agent RL. *CoRR*, abs/1810.08647, 2018. URL <https://arxiv.org/abs/1810.08647>.
- Kraemer, L. and Banerjee, B. Multi-agent reinforcement learning as a rehearsal for decentralized planning. *Neuro-computing*, 190:82–94, 2016.
- Kulkarni, T. D., Saedi, A., Gautam, S., and Gershman, S. J. Deep successor reinforcement learning. *arXiv preprint arXiv:1606.02396*, 2016.
- Lee, D., Srinivasan, S., and Doshi-Velez, F. Truly batch apprenticeship learning with deep successor features. *arXiv preprint arXiv:1903.10077*, 2019.
- Lehnert, L., Tellex, S., and Littman, M. L. Advantages and limitations of using successor features for transfer in reinforcement learning. *arXiv preprint arXiv:1708.00102*, 2017.

- Leibo, J. Z., Hughes, E., Lanctot, M., and Graepel, T. Autocurricula and the emergence of innovation from social interaction: A manifesto for multi-agent intelligence research. *CoRR*, abs/1903.00742, 2019. URL <http://arxiv.org/abs/1903.00742>.
- Lin, L.-J. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*, 8(3):293–321, 1992.
- Mahajan, A. and Tulabandhula, T. Symmetry detection and exploitation for function approximation in deep rl. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pp. 1619–1621, 2017a.
- Mahajan, A. and Tulabandhula, T. Symmetry learning for function approximation in reinforcement learning. *arXiv preprint arXiv:1706.02999*, 2017b.
- Mahajan, A., Rashid, T., Samvelyan, M., and Whiteson, S. Maven: Multi-agent variational exploration. In *Advances in Neural Information Processing Systems*, pp. 7613–7624, 2019.
- Mahajan, A., Samvelyan, M., Mao, L., Makoviychuk, V., Garg, A., Kossaiji, J., Whiteson, S., Zhu, Y., and Anandkumar, A. Tesseract: Tensorised actors for multi-agent reinforcement learning. *arXiv preprint arXiv:2106.00136*, 2021.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- O’Donoghue, B., Osband, I., Munos, R., and Mnih, V. The uncertainty Bellman equation and exploration. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pp. 3836–3845, 2018. URL <http://proceedings.mlr.press/v80/o-donoghue18a.html>.
- Oliehoek, F. A., Amato, C., et al. *A concise introduction to decentralized POMDPs*, volume 1. Springer, 2016.
- OroojlooyJadid, A. and Hajinezhad, D. A review of cooperative multi-agent deep reinforcement learning. *arXiv preprint arXiv:1908.03963*, 2019.
- Osband, I., Van Roy, B., and Wen, Z. Generalization and exploration via randomized value functions. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning (ICML)*, pp. 2377–2386, 2016.
- Panait, L., Luke, S., and Wiegand, R. P. Biasing coevolutionary search for optimal multiagent behaviors. *IEEE Transactions on Evolutionary Computation*, 10(6):629–645, 2006.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017.
- Rashid, T., Farquhar, G., Peng, B., and Whiteson, S. Weighted qmix: Expanding monotonic value function factorisation. *arXiv preprint arXiv:2006.10800*, 2020a.
- Rashid, T., Samvelyan, M., De Witt, C. S., Farquhar, G., Foerster, J., and Whiteson, S. Monotonic value function factorisation for deep multi-agent reinforcement learning. *arXiv preprint arXiv:2003.08839*, 2020b.
- Roderick, M., Grimm, C., and Tellex, S. Deep abstract Q-networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pp. 131–138, 2018.
- Samvelyan, M., Rashid, T., de Witt, C. S., Farquhar, G., Nardelli, N., Rudner, T. G., Hung, C.-M., Torr, P. H., Foerster, J., and Whiteson, S. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*, 2019.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal value function approximators. In *International conference on machine learning*, pp. 1312–1320, 2015.
- Son, K., Kim, D., Kang, W. J., Hostallero, D. E., and Yi, Y. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:1905.05408*, 2019.
- Son, K., Ahn, S., Reyes, R. D., Shin, J., and Yi, Y. Qtran++: Improved value transformation for cooperative multi-agent reinforcement learning, 2020.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., et al. Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296*, 2017.
- Tan, M. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pp. 330–337, 1993.
- Tang, H., Houthoofd, R., Foote, D., Stooke, A., Xi Chen, O., Duan, Y., Schulman, J., DeTurck, F., and Abbeel, P. #Exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS) 30*, pp. 2753–2762. 2017.

- van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *Proceedings of the 13th AAAI Conference on Artificial Intelligence*, pp. 2094–2100, 2016. URL <https://arxiv.org/pdf/1509.06461.pdf>.
- Wang, J., Ren, Z., Liu, T., Yu, Y., and Zhang, C. Qplex: Duplex dueling multi-agent q-learning. *arXiv preprint arXiv:2008.01062*, 2020a.
- Wang, T., Gupta, T., Mahajan, A., Peng, B., Whiteson, S., and Zhang, C. Rode: Learning roles to decompose multi-agent tasks. *arXiv preprint arXiv:2010.01523*, 2020b.
- Wang, T., Wang, J., Wu, Y., and Zhang, C. Influence-based multi-agent exploration. In *International Conference on Learning Representations*, 2020c. URL <https://openreview.net/forum?id=BJgy96EYvr>.
- Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., and Freitas, N. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, pp. 1995–2003, 2016.
- Wei, E. and Luke, S. Lenient learning in independent-learner stochastic cooperative games. *The Journal of Machine Learning Research*, 17(1):2914–2955, 2016.
- Wei, E., Wicke, D., Freelan, D., and Luke, S. Multiagent soft q-learning. *arXiv preprint arXiv:1804.09817*, 2018.
- Wei, E., Wicke, D., and Luke, S. Multiagent adversarial inverse reinforcement learning. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, pp. 2265–2266, 2019.
- Zheng, S. and Yue, Y. Structured exploration via hierarchical variational policy networks, 2018. URL <https://openreview.net/forum?id=HyunpgrR->.