
Randomized Entity-wise Factorization for Multi-Agent Reinforcement Learning

Shariq Iqbal¹ Christian A. Schroeder de Witt² Bei Peng² Wendelin Böhmer³ Shimon Whiteson² Fei Sha^{1,4}

Abstract

Multi-agent settings in the real world often involve tasks with varying types and quantities of agents and non-agent entities; however, common patterns of behavior often emerge among these agents/entities. Our method aims to leverage these commonalities by asking the question: “What is the expected utility of each agent when only considering a randomly selected subgroup of its observed entities?” By posing this counterfactual question, we can recognize state-action trajectories within sub-groups of entities that we may have encountered in another task and use what we learned in that task to inform our prediction in the current one. We then reconstruct a prediction of the full returns as a combination of factors considering these disjoint groups of entities and train this “randomly factorized” value function as an auxiliary objective for value-based multi-agent reinforcement learning. By doing so, our model can recognize and leverage similarities across tasks to improve learning efficiency in a multi-task setting. Our approach, **Randomized Entity-wise Factorization for Imagined Learning (REFIL)**, outperforms all strong baselines by a significant margin in challenging multi-task StarCraft micromanagement settings.

1. Introduction

Multi-agent reinforcement learning techniques often focus on learning in settings with fixed groups of agents and entities; however, many real-world multi-agent settings contain tasks across which an agent must deal with varying quan-

¹Department of Computer Science, University of Southern California ²Department of Computer Science, University of Oxford ³Department of Software Technology, Delft University of Technology ⁴Google Research. Correspondence to: Shariq Iqbal <shariqiq@usc.edu>.




Figure 1: Breakaway sub-scenario in soccer. Agents in the yellow square can generalize this experience to similar subsequent experiences, regardless of the state of agents outside the square.

ties and types of cooperative agents, antagonists, or other entities. This variability in type and quantity of entities results in a combinatorial growth in the number of possible configurations, aggravating the challenge of learning control policies that generalize. For example, the sport of soccer exists in many forms, from casual 5 vs. 5 to full scale 11 vs. 11 matches, with varying formations within each consisting of different quantities of player types (defenders, midfielders, forwards, etc.). Within these varied tasks, however, exist common patterns. For instance, a “breakaway” occurs in soccer when an attacker with the ball passes the defense and only needs to beat the goalkeeper in order to score (Figure 1). The goalkeeper and attacker can apply what they have learned in a breakaway to the next one, regardless of the task (e.g., 5 vs. 5). If players can disentangle their understanding of common patterns from their surroundings, they should be able to learn more efficiently as well as share their experiences across *all* forms of soccer. These repeated patterns within sub-groups of entities can, in fact, be found in a wide variety of multi-agent tasks (e.g., heterogeneous swarm control (Prorok et al., 2017) and StarCraft unit micromanagement (Samvelyan et al., 2019)). Our work¹ aims to develop a methodology for artificial agents to incorporate knowledge of these shared patterns to accelerate learning in a multi-task setting.

One way to leverage structural independence among agents, as in our soccer example, is to represent value functions as a combination of factors that depend on disjunct subsets of the state and action spaces (Koller & Parr, 1999). These subsets are typically fixed in advance using domain knowledge and thus do not scale to complex domains where dependencies are unknown and may shift over time. Recent ap-

¹Code available at: <https://github.com/shariqiqbal2810/REFIL>

proaches (e.g. VDN (Sunehag et al., 2018), QMIX (Rashid et al., 2018)) in cooperative deep multi-agent reinforcement learning (MARL) factor value functions into separate components for each agent’s action and observation space in order to enable decentralized execution. These approaches learn a utility function for each agent that depends on the agent’s own action and observations, resulting in a unique observation space for each task and exacerbating the challenge of learning in a multi-task setting.

How can we teach agents to be “situationally aware” of common patterns that are not pre-specified, such that they can share knowledge across tasks? Our main idea is as follows: Given observed trajectories in a real task, we randomly partition entities into sub-groups to “imagine” that agents only observe a (random) subset of the entities they actually observe. Then, in addition to estimating utility of their actions given the full observations, we use the same model to predict utilities in the imagined scenario, providing an opportunity to discover sub-group patterns that appear across tasks. For example, we might sample a breakaway (or a superset of the breakaway entities) in both 5v5 soccer and 11v11, allowing our model to share value function factors across tasks. We can then use these factors to construct a prediction of the full returns.

Of course, the possibility of sampling sub-groups that do not contain independent behavior exists. Imagine a sub-group that looks like a breakaway, but in reality a defender is closing in on the attacker’s left. In such cases, we must include factors that account for the effect interactions outside of sampled sub-groups on each agent’s utility. Crucially however, the estimated utility derived from imagining a breakaway often provides at least *some* information as to the agent’s utility given the full observation (i.e., the agent knows there is value in dribbling toward the goal). Imagined sub-group factors are combined with interaction factors to produce an estimate of the value function that we train as an auxiliary objective on top of a standard value function loss. As such, our approach allows models to exploit shared inter-task patterns via factorization without losing any expressivity. We emphasize that this approach does not rely on sampling an “optimal” sub-group. In other words, there is no requirement to sample sub-groups that are independent from one another (c.f. Section 4.1). In fact, it is useful to learn a utility function for *any* sub-group state that may appear in another task.

Our approach: **Randomized Entity-wise Factorization for Imagined Learning (REFIL)** can be implemented easily in practice by using masks in attention-based models. We evaluate our approach on complex StarCraft Multi-Agent Challenge (SMAC) (Samvelyan et al., 2019) multi-task settings with varying agent teams, finding **REFIL** attains improved performance over state-of-the-art methods.

2. Background and Preliminaries

In this work, we consider the *decentralized partially observable Markov decision process* (Dec-POMDP) (Oliehoek et al., 2016) with entities (Schroeder de Witt et al., 2019), which describes fully cooperative multi-agent tasks.

Dec-POMDPs with Entities are described as tuples: $(\mathbf{S}, \mathbf{U}, \mathbf{O}, P, r, \mathcal{E}, \mathcal{A}, \Phi, \mu)$. \mathcal{E} is the set of entities in the environment. Each entity e has a state representation s^e , and the global state is the set $\mathbf{s} = \{s^e | e \in \mathcal{E}\} \in \mathbf{S}$. Some entities can be agents $a \in \mathcal{A} \subseteq \mathcal{E}$. Non-agent entities are parts of the environment that are not controlled by learning policies (e.g., landmarks, obstacles, agents with fixed behavior). The state features of each entity consist of two parts: $s^e = [f^e, \phi^e]$ where f^e represents the description of an entity’s current state (e.g., position, orientation, velocity, etc.) while $\phi^e \in \Phi$ represents the entity’s type (e.g., outfield player, goalkeeper, etc.), of which there is a discrete set. An entity’s type affects the state dynamics as well as the reward function and, importantly, it remains fixed for the duration of the entity’s existence. Not all entities may be visible to each agent, so we define a binary observability mask: $\mu(s^a, s^e) \in \{1, 0\}$, where agents can always observe themselves $\mu(s^a, s^a) = 1, \forall a \in \mathcal{A}$. Thus, an agent’s observation is defined as $o^a = \{s^e | \mu(s^a, s^e) = 1, e \in \mathcal{E}\} \in \mathbf{O}$. Each agent a can execute actions u^a , and the joint action of all agents is denoted $\mathbf{u} = \{u^a | a \in \mathcal{A}\} \in \mathbf{U}$. P is the state transition function which defines the probability $P(s' | \mathbf{s}, \mathbf{u})$ and $r(\mathbf{s}, \mathbf{u})$ is the reward function that maps the global state and joint actions to a single scalar reward.

Entities cannot be added during an episode, but they may become inactive (e.g., a unit dying in StarCraft) and no longer affect transitions and rewards. Since \mathbf{s} and \mathbf{u} are sets, their ordering does not matter, and our modeling construct should account for this (e.g., by modeling with permutation invariant/equivariant attention models (Lee et al., 2019)). In many domains, the set of entity types present $\{\phi^e | e \in \mathcal{E}\}$ is fixed. We are particularly interested in a multi-task setting where the quantity and types of entities are varied between tasks, as identifying patterns within sub-groups of entities is crucial to generalizing experience effectively in these cases.

Learning We aim to learn a set of policies that maximize expected discounted reward (returns). Q -learning is specifically concerned with learning an accurate action-value function Q^{tot} (defined below), and using this function to select the actions that maximize expected returns. The optimal Q -function for our setting is defined as:

$$\begin{aligned}
 Q^{\text{tot}}(\mathbf{s}, \mathbf{u}) &:= \\
 &\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(\mathbf{s}_t, \mathbf{u}_t) \mid \begin{array}{l} \mathbf{s}_0 = \mathbf{s}, \mathbf{u}_0 = \mathbf{u}, \mathbf{s}_{t+1} \sim P(\cdot | \mathbf{s}_t, \mathbf{u}_t) \\ \mathbf{u}_{t+1} = \arg \max Q^{\text{tot}}(\mathbf{s}_{t+1}, \cdot) \end{array} \right] \quad (1) \\
 &= r(\mathbf{s}, \mathbf{u}) + \gamma \mathbb{E} \left[\max Q^{\text{tot}}(\mathbf{s}', \cdot) \mid \mathbf{s}' \sim P(\cdot | \mathbf{s}, \mathbf{u}) \right]
 \end{aligned}$$

Partial observability is typically handled by using the history of actions and observations as a proxy for state, often processed by a recurrent neural network (Hausknecht & Stone, 2015): $Q_\theta^{\text{tot}}(\tau_t, \mathbf{u}_t) \approx Q^{\text{tot}}(s_t, \mathbf{u}_t)$ where the trajectory is $\tau_t^a := (o_0^a, u_0^a, \dots, o_t^a)$ and $\tau_t := \{\tau_t^a\}_{a \in \mathcal{A}}$. To learn the Q -function, deep reinforcement learning uses neural networks as function approximators trained to minimize the loss function:

$$\mathcal{L}_Q(\theta) := \mathbb{E} \left[\left(y_t^{\text{tot}} - Q_\theta^{\text{tot}}(\tau_t, \mathbf{u}_t) \right)^2 \mid (\tau_t, \mathbf{u}_t, r_t, \tau_{t+1}) \sim \mathcal{D} \right]$$

$$y_t^{\text{tot}} := r_t + \gamma Q_{\bar{\theta}}^{\text{tot}}(\tau_{t+1}, \arg \max Q_{\bar{\theta}}^{\text{tot}}(\tau_{t+1}, \cdot)) \quad (2)$$

where $\bar{\theta}$ are the parameters of a target network that is copied from θ periodically to improve stability (Mnih et al., 2015) and \mathcal{D} is a replay buffer (Lin, 1992) that stores transitions collected by an exploratory policy (typically ϵ -greedy). Double deep Q -learning (van Hasselt et al., 2016) mitigates overestimation of the learned values by using actions that maximize Q_θ^{tot} as inputs for the target network $Q_{\bar{\theta}}^{\text{tot}}$.

Value Function Factorization Centralized training for decentralized execution (CTDE) has been a major focus in recent efforts in deep multi-agent RL (Lowe et al., 2017; Foster et al., 2018; Sunehag et al., 2018; Rashid et al., 2018; Iqbal & Sha, 2019). Some methods achieve CTDE through factoring Q -functions into monotonic combinations of per-agent utilities, with each depending only on a single agent’s history of actions and observations $Q^a(\tau^a, u^a)$. This factorization allows agents to independently maximize their local utility functions in a decentralized manner with their selected actions combining to form the optimal joint action. While factored value functions can only represent a limited subset of all possible value functions (Böhmer et al., 2020), they tend to perform better empirically than those that learn unfactored joint action value functions (Oliehoek et al., 2008).

QMIX (Rashid et al., 2018) improves over value decomposition networks (VDN) (Sunehag et al., 2018) by using a more expressive factorization than a summation of factors:

$$Q^{\text{tot}} = g(Q^1(\tau^1, u^1; \theta_Q), \dots, Q^{|\mathcal{A}|}(\tau^{|\mathcal{A}|}, u^{|\mathcal{A}|}; \theta_Q); \theta_g)$$

The parameters of the monotonic mixing function θ_g are generated by a hyper-network (Ha et al., 2017) conditioning on the global state \mathbf{s} : $\theta_g = h(\mathbf{s}; \theta_h)$. Every state can therefore have a different mixing function; however, the mixing functions’s monotonicity maintains decentralizability, as agents can greedily maximize Q^{tot} without communication. All parameters $\theta = \{\theta_Q, \theta_h\}$ are trained with the DQN loss of Equation 2.

Attention Mechanisms for MARL Attention models have recently generated intense interest due to their ability to incorporate information across large contexts, including in MARL (Jiang & Lu, 2018; Iqbal & Sha, 2019; Long et al.,

2020). Importantly for our purposes, they can process variable sized sets of fixed length vectors (in our case entities). At the core of these models is a parameterized transformation known as multi-head attention (Vaswani et al., 2017) that allows entities to selectively extract information from other entities based on their local context.

We define \mathbf{X} as a matrix where each row corresponds to the state representation (or its transformation) of an entity. The global state \mathbf{s} is represented in matrix form as $\mathbf{X}^\mathcal{E}$ where $\mathbf{X}_{e,*} = s^e$. Our models consist of entity-wise feedforward layers $\text{eFF}(\mathbf{X})$, which apply an identical linear transformation to all input entities and multi-head attention layers $\text{MHA}(\mathcal{A}, \mathbf{X}, \mathbf{M})$, which integrate information across entities. The latter take three arguments: the set of agents \mathcal{A} for which to compute an output vector, the matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{E}| \times d}$ where d is the dimensionality of the input representations, and a mask $\mathbf{M} \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{E}|}$. The layer outputs a matrix $\mathbf{H} \in \mathbb{R}^{|\mathcal{A}| \times h}$ where h is the hidden dimension of the layer. The row $\mathbf{H}_{a,*}$ corresponds to a weighted sum of linearly transformed representations from all entities selected by agent a . Importantly, if the entry of the mask $\mathbf{M}_{a,e} = 0$, then entity e ’s representation is not included in $\mathbf{H}_{a,*}$. Masking enables 1) decentralized execution by providing the mask $\mathbf{M}_{a,e}^\mu = \mu(s^a, s^e)$, such that agents can only see entities observable by them in the environment, and 2) “imagination” of the returns among sub-groups of entities. We integrate entity-wise feedforward layers and multi-head attention into QMIX in order to share a model across tasks where the number of agents and entities is variable and build our approach from there. The exact process of computing attention layers, as well as the specifics of our attention-augmented version of QMIX are described in detail in the Supplement.

3. REFIL

We now propose **Randomized Entity-wise Factorization for Imagined Learning (REFIL)**. We observe that common patterns often emerge in sub-groups of entities within complex multi-agent tasks (cf. soccer breakaway example in §1) and hypothesize that learning to predict agents’ utilities within sub-groups of entities is a strong inductive bias that allows models to share information more freely across tasks. We instantiate our approach by constructing an estimate of the value function from factors based on randomized sub-groups, sharing parameters with the full value function, and training this factorized version of the value function as an auxiliary objective.

3.1. Main Idea

Random Partitioning Given an episode sampled from a replay buffer, we first randomly partition all entities in \mathcal{E} into two disjoint groups, held fixed for the episode. We denote

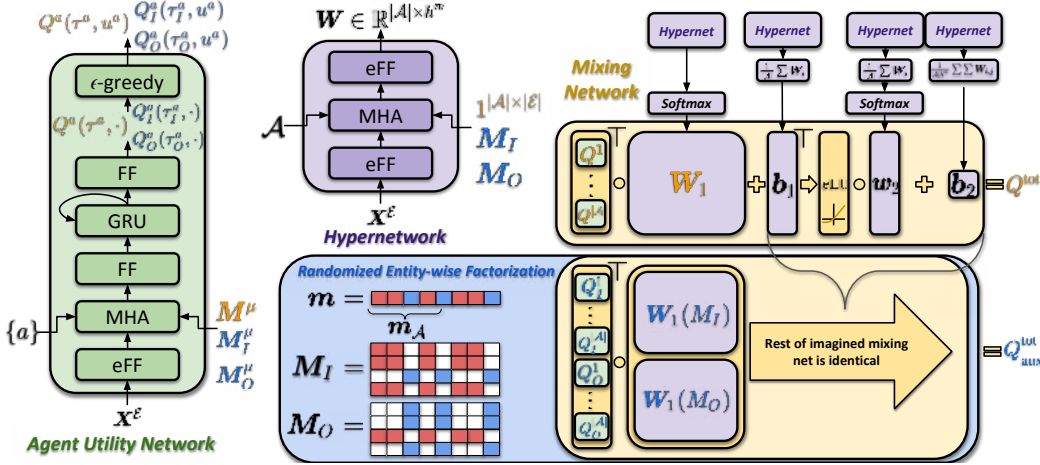


Figure 2: Schematic for **REFIL**. Values colored orange or blue are used for computing Q^{tot} and $Q_{\text{aux}}^{\text{tot}}$ respectively. **(left)** Agent-specific utility networks. These are decentralizable due to the use of an observability mask (M^μ). We include Gated Recurrent Units (Chung et al., 2014) to retain information across timesteps in order to handle partial observability. **(top center)** Hypernetworks used to generate weights for the mixing network. We use a softmax function on the weights across the hidden dimension to enforce non-negativity, which we find empirically to be more stable than the standard absolute value function. Hypernetworks are not restricted by partial observability since they are only required during training and not execution. **(top right)** The mixing network used to calculate Q^{tot} . **(bottom right)** Procedure for performing randomized entity-wise factorization. For masks M_I and M_O , colored spaces indicate a value of 1 (i.e., the agent designated by the row will be able to see the entity designated by the column), while white spaces indicate a value of 0. The color indicates which group the entity belongs to, so agents in the red group see red entities in M_I and blue entities in M_O . Agents are split into sub-groups and their utilities are calculated for both interactions within their group, as well as to account for the interactions outside of their group, then monotonically mixed to predict $Q_{\text{aux}}^{\text{tot}}$.

the partition by a random binary² vector $\mathbf{m} \in \{0, 1\}^{|\mathcal{E}|}$. m_e indicates whether entity e is in the first group. The negation $-m_e$ represents whether e is in the second group. The subset of all agents is denoted $\mathbf{m}_A := [m_a]_{a \in \mathcal{A}}$. With these vectors, we construct binary attention masks M_I and M_O :

$$M_I := \mathbf{m}_A \mathbf{m}^\top \vee \neg \mathbf{m}_A \neg \mathbf{m}^\top, M_O := \neg M_I. \quad (3)$$

where $M_I^\mu[a, e]$ indicates whether agent a and entity e are in the same group, and $M_O^\mu[a, e]$ indicates the opposite. They are further combined with a partial observability mask M^μ , which is provided by the environment, to generate the final attention masks

$$M_I^\mu := M^\mu \wedge M_I, M_O^\mu := M^\mu \wedge M_O \quad (4)$$

These matrices are of size $|\mathcal{A}| \times |\mathcal{E}|$ and will be used by the multi-head attention layers to constrain which entities can be observed by agents.

Counterfactual Reasoning Given an *imagined* partition \mathbf{m} , an agent a can examine its history of observations and actions and reason *counterfactually* what its utility would be had it *solely* observed the entities in its group. We call this quantity *in-group utility* and denote it by $Q_I^a(\tau_I^a, u^a; \theta_Q)$. In order to account for the potential interactions with entities outside of the agents group, we calculate an *out-group*

²We first draw $p \in (0, 1)$ uniformly, followed by $|\mathcal{E}|$ independent draws from a Bernoulli(p) distribution. Partitioning into two groups induces a uniform distribution over all possible sub-groups.

utility: $Q_O^a(\tau_O^a, u^a; \theta_Q)$. Note that the real and imagined utilities share the same parameters θ_Q , allowing us to leverage imagined experience to improve utility prediction in real scenarios and vice versa. Breaking the fully observed utilities Q^a into these randomized sub-group factors is akin to breaking an image into cut-outs of the comprising entities. While the “images” (i.e. states) from each task are a unique set, it is likely that the pieces comprising them share similarities.

Since we do not know the returns within the imagined sub-groups, we must ground our predictions in the observed returns. Just as QMIX learns a value function with n factors (Q^a for each agent), we learn an imagined value function with $2n$ factors (Q_I^a and Q_O^a for each agent) that estimates the same value:

$$\begin{aligned} Q^{\text{tot}} &= g(Q^1, \dots, Q^{|\mathcal{A}|}; h(\mathbf{s}; \theta_h, \mathbf{M})) \\ &\approx Q_{\text{aux}}^{\text{tot}} = g(Q_I^1, \dots, Q_I^{|\mathcal{A}|}, Q_O^1, \dots, Q_O^{|\mathcal{A}|}; \\ &\quad h(\mathbf{s}; \theta_h, \mathbf{M}_I), h(\mathbf{s}; \theta_h, \mathbf{M}_O)) \end{aligned} \quad (5)$$

Where $g(\cdot)$ are mixing networks whose parameters are generated by hypernetworks $h(\mathbf{s}; \theta_h, \mathbf{M})$. This network’s first layer typically takes n inputs, one for each agent. Since we have $2n$ factors, we simply concatenate two generated versions of the input layer (using M_I and M_O). We then apply the network to the concatenated utilities $Q_I^a(\tau_I^a, u^a)$ and $Q_O^a(\tau_O^a, u^a)$ of all agents a , to compute the predicted value $Q_{\text{aux}}^{\text{tot}}$. This procedure is visualized in Figure 2 and described in more detail in the Supplement.




Figure 3: Group Matching Game. We use the values $n_a = 8$, $n_c = 6$, and $n_g = 2$ in our experiments. Shaded region is a 95% confidence interval across 24 runs.

Importantly, since the mixing network is generated by the *full state context*, our model can weight factors contextually. For example, if the agent a 's sampled sub-group contains all relevant information to compute its utility such that $Q_I^a \approx Q^a$, then the mixing networks can weight Q_I^a more heavily than Q_O^a . Otherwise, the networks learn to balance Q_I^a and Q_O^a for each agent, in order to estimate Q^{tot} . In this way, we can share knowledge in similar sub-group states across tasks while accounting for the differences in utility that result from the out-of-group context.

Learning We now describe the overall learning objective of **REFIL**. To enforce Equation 5, we replace Q^{tot} in Equation 2 with $Q_{\text{aux}}^{\text{tot}}$, resulting a new loss \mathcal{L}_{aux} . We combine the standard QMIX loss (Eq. 2) \mathcal{L}_Q with this auxiliary loss to form:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_Q + \lambda\mathbb{E}_m\mathcal{L}_{\text{aux}} \quad (6)$$

where λ controls the tradeoff between the two losses. Note that we randomly partition in each episode, hence the expectation with respect to the partition vector m . We emphasize that the sub-groups are *imagined*. While we compute $Q_{\text{aux}}^{\text{tot}}$ and its related quantities, we do not use them to select actions in Equation 2. Action selection is performed by each agent maximizing Q^a given their local observations. This greedy local action selection is guaranteed to maximize Q^{tot} due to the monotonic structure of the mixing network (Rashid et al., 2018). Moreover, our auxiliary objective is only used in training, and execution in the environment does not use random factorization. Treating random factorization as an auxiliary task, rather than as a representational constraint, allows us to retain the expressivity of QMIX value functions (without sub-group factorization) while exploiting the existence of shared sub-group states across tasks.

3.2. Implementation Details

The model architecture is shown in Figure 2, with more details described in the supplement. ‘‘Imagination’’ can be implemented efficiently using attention masks. Specifically two additional passes through the network are needed, with

M_O^μ and M_I^μ as masks instead of M^μ , per training step. These additional passes can be parallelized by computing all necessary quantities in one batch on GPU. It is feasible to split entities into an arbitrary number i of random sub-groups without using more computation by sampling several disjunct vectors m^i and combining them in the same way as we combine m and $-m$ in Equation 3 to form M_I and M_O . Doing so could potentially bias agents towards considering patterns within smaller subsets of entities.

4. Experimental Results

In our experiments, we aim to answer the following questions: 1) Are randomized counterfactuals an efficient means for leveraging common patterns? 2) Does our approach improve generalization in a multi-task setting? 3) Is training as an auxiliary objective justified? We begin with experiments in a simple domain we construct such that agents’ decisions rely only on a *known subset* of all entities, so we can compare our approach to those that use this domain knowledge. Then, we move on to testing on complex StarCraft micro-management tasks to demonstrate our method’s ability to scale to complex domains.

4.1. Group Matching Game

In order to answer our first question, we construct a group matching game, pictured in Figure 3a, where each agent only needs to consider a subset of other agents to act effectively and we know that subset as ground truth (unlike in more complex domains such as StarCraft). Agents (of which there are n_a) are randomly placed in one of n_c cells and assigned to one of n_g groups (represented by the different colors) at the start of each episode. Each unique group assignment corresponds to a task. Agents can choose from three actions: move clockwise, stay, and move counter-clockwise. Their ultimate goal is to be located in the same cell as the rest of their group members, at which point an episode ends. There is no restriction on which cell agents form a group in (e.g., both groups can form in the same cell). All agents share a reward of 2.5 when any group is

completed (and an equivalent penalty for a formed group breaking) as well as a penalty of -0.1 for each time step in order to encourage agents to solve the task as quickly as possible. Agents’ entity-state descriptions s^e include the cell that the agent is currently occupying as well as the group it belongs to (both one-hot encoded), and the task is fully-observable. Notably, agents can act optimally while only considering a subset of observed entities.

Ground-truth knowledge of relevant entities enables us to disentangle two aspects of our approach: the use of entity-wise factorization in general and specifically using randomly selected factors. We would like to answer the question: does our method rely on sampling the “right” groups of entities (i.e., those with no interactions between them), or is the randomness of our method a feature that promotes generalization? We construct two approaches that use this knowledge to build factoring masks M_I and M_O that are used in place of randomly sampled groups (otherwise the methods are identical to **REFIL**). **REFIL (Fixed Oracle)** directly uses the ground truth group assignments (different for each task) to build masks. **REFIL (Randomized Oracle)** randomly samples sub-groups from the ground truth groups only, rather than from all possible entities. We additionally train **REFIL** and *QMIX (Attention)* (i.e., **REFIL** with no auxiliary loss).

Figure 3b shows that using domain knowledge does not significantly improve performance in this domain (*QMIX (Attention)* vs. **REFIL (Fixed Oracle)**). In fact our *randomized* factorization approach outperforms the use of domain knowledge. The randomization in **REFIL** therefore appears to be crucial. Our hypothesis is that randomization of sub-group factors enables better knowledge sharing across tasks. For example, the situation where two agents from the same group are located in adjacent cells occurs within *all* possible group assignments. When sampling randomly, our approach occasionally samples these two agents alone in their own group. Even if the rest of the context in a given episode has never been seen before, as long as this sub-scenario has been seen, the model has some indication of the value associated with each action. Even when restricting the set of entities to form sub-groups with those that we know can be relevant to each agent (**REFIL (Randomized Oracle)**) we find that performance does not significantly improve. These results suggest that randomized sub-group formation for **REFIL** is a viable strategy, and the main benefit of our approach is to promote generalization across tasks by breaking value function predictions into reusable components, even when the sampled sub-groups are not completely independent.

4.2. STARCRAFT

We next test on the StarCraft multi-agent challenge (SMAC) (Samvelyan et al., 2019). The tasks in SMAC involve mi-

Table 1: Comparison of tested methods.

Name	Imagined Learning	Model	Base Algorithm
REFIL	✓	MHA ¹	QMIX ²
QMIX (Attention)		MHA	QMIX
REFIL (VDN)	✓	MHA	VDN ³
VDN (Attention)		MHA	VDN
QMIX (Max Pooling)		Max-Pool	QMIX
QMIX (EMP)		EMP ⁴	QMIX
ROMA (Attention)		MHA	ROMA ⁵
Qatten (Attention)		MHA	Qatten ⁶
QTRAN (Attention)		MHA	QTRAN ⁷
REFIL (UPDeT)	✓	UPDeT ⁸	QMIX
QMIX (UPDeT)		UPDeT	QMIX

¹: Vaswani et al. (2017) ²: Rashid et al. (2018) ³: Sunehag et al. (2018)
⁴: Agarwal et al. (2019) ⁵: Wang et al. (2020a) ⁶: Yang et al. (2020)
⁷: Son et al. (2019) ⁸: Hu et al. (2021)

cromanagement of units in order to defeat a set of enemy units in battle. Specifically, we consider a multi-task setting where we train our models simultaneously on tasks with variable types and quantities of agents. We hypothesize that our approach is especially beneficial in this setting, as it should encourage models to learn utilities for common patterns and generalize to more diverse settings as a result. The dynamic setting involves minor modifications to SMAC but we change the environment as little as possible to maintain the challenging nature of the tasks. In the standard version of SMAC, both state and action spaces depend on a fixed number of agents and enemies, so our modifications, discussed in detail in the supplement, alleviate these problems.

In our tests we evaluate on three settings we call 3-8sz, 3-8csz, and 3-8MMM. 3-8sz pits symmetrical teams of between 3 and 8 agents against each other where the agents are a combination of Zealots and Stalkers (inspired by the 2s3z and 3s5z tasks in the original SMAC), resulting in 39 unique tasks. 3-8csz pits symmetrical teams of between 0 and 2 Colossi and 3 to 6 Stalkers/Zealots against each other (inspired by 1c3s5z), resulting in 66 tasks. 3-8MMM pits symmetrical teams of between 0 and 2 Medics and 3 to 6 Marines/Marauders against each other (inspired by MMM and MMM2, again resulting in 66 tasks).

Ablations and Baselines We introduce several ablations of our method, as well as adaptations of existing methods to handle variable sized inputs. These comparisons are summarized in Table 1. *QMIX (Attention)* is our method without the auxiliary loss. **REFIL (VDN)** is our approach using summation to combine all factors as in VDN rather than a non-linear monotonic mixing network. *VDN (Attention)* does not include the auxiliary loss and uses summation for factor mixing. *QMIX (Mean Pooling)* is *QMIX (Attention)* with attention layers replaced by mean pooling. We also test max pooling but find the performance to be marginally worse than mean pooling. Importantly, for pooling layers we add entity-wise linear transformations prior to the pooling operations such that the total number of parameters is comparable to attention layers.




Figure 4: Test win rate over time on multi-task STARCRAFT environments. Tasks are sampled uniformly at each episode. Shaded region is a 95% confidence interval across 5 runs. (top row) Ablations of our method. (bottom row) Baseline methods.

For baselines we consider some follow-up works to QMIX that improve the mixing network’s expressivity: QTRAN (Son et al., 2019) and Qatten (Yang et al., 2020). We also compare to a method that builds on QMIX by attempting to learn dynamic roles that depend on the context each agent observes: ROMA (Wang et al., 2020a). We additionally consider an alternative mechanism for aggregating information across variable sets of entities, known as Entity Message Passing (EMP) (Agarwal et al., 2019). We specifically use the restricted communication setting where agents can only communicate with agents they observe, and we set the number of message passing steps to three. Finally, we consider the UPDeT architecture (Hu et al., 2021), a recent work that also targets the multi-task MARL setting. UPDeT utilizes domain knowledge of the environment to map entities to the specific actions that they correspond to. We train UPDeT with QMIX as well as REFIL. For all approaches designed for the standard single-task SMAC setting, we extend them with the same multi-head attention architecture that our approach uses.

Ablation Results Our results on challenges in multi-task STARCRAFT settings can be found in Figure 4. Tasks are sampled uniformly at each episode, so the curves represent average win rate across all tasks. We find that REFIL outperforms all ablations consistently in these settings. REFIL (VDN) performs much worse than our approach and VDN (Attention), highlighting the importance of the mixing network handling contextual dependencies between entity partitions. Since the trajectory of a subset of entities can play out differently based on the surrounding context, it is

important for our factorization approach to recognize and adjust for these situations. The use of mean-pooling in place of attention also performs poorly, indicating that attention is valuable for aggregating information from variable length sets of entities.

Baseline Results We find that algorithms designed to improve on QMIX for the single task MARL setting (ROMA, Qatten, QTRAN), naively applied to the multi-task setting, do not see the same improvements. REFIL, on the other hand, consistently outperforms other methods, highlighting the unique challenge of learning in multi-task settings. In Fig. 5 (top right) we investigate the performance of REFIL compared to the two next best methods in the 3-8sz setting on a task-by-task basis. We evaluate each method on each task individually and rank the tasks by performance, plotting from left to right. We find that the performance gain of REFIL comes from generalizing performance across a wider range of tasks, hence the reduced rate of decay in task performance from best to worst.

The entity aggregation method of EMP underperforms relative to the MHA module that we use. UPDeT is a related work that focuses on designing an architecture compatible with multiple tasks and variable entities and action spaces by utilizing domain knowledge to map entities to their corresponding actions. Despite adding this domain knowledge, QMIX (UPDeT) surprisingly underperforms in 2 of 3 settings, while performing similarly to REFIL on 3-8MMM; however, since UPDeT is an attention-based architecture, it is amenable to our proposed auxiliary training scheme. We find applying random factorization to QMIX (UPDeT) im-

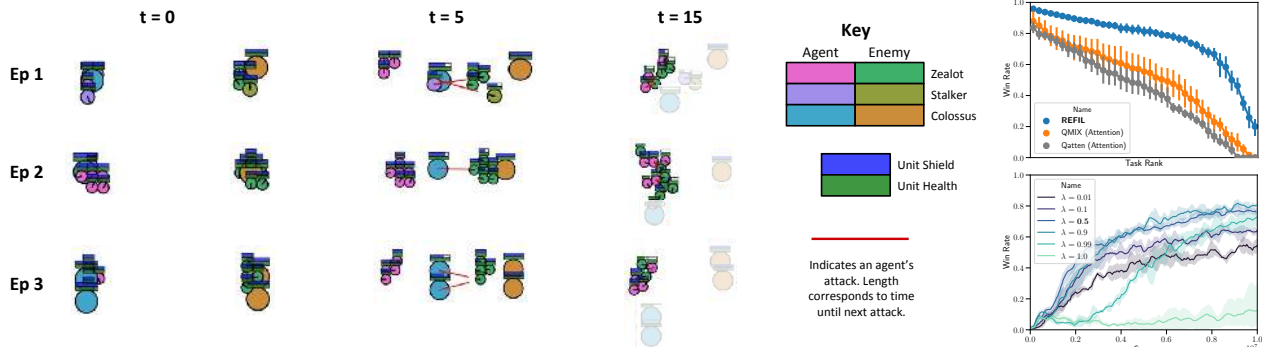


Figure 5: (left) Simplified rendering of a common pattern that emerges across tasks in the 3-8csz SMAC setting, highlighted at $t = 15$. **REFIL** enables learning from each task to inform behavior in the others. (top right) Task-by-task performance on 3-8sz. **REFIL** generalizes better across a wider range of tasks. (bottom right) Varying λ for **REFIL** in 3-8sz.

proves its performance further in 3-8MMM as well as in 3-8sz. In the case of 3-8MMM, where the asymptotic win rate of **REFIL** (UPDeT) and QMIX (UPDeT) are similar, we find that **REFIL** (UPDeT) wins on average in 22% fewer time steps by targeting enemy Medivacs, a unit capable of healing its teammates. Targeting of Medivacs is an example of a common pattern that emerges across tasks which **REFIL** is able to leverage.

Role of Auxiliary Objective In order to understand the role of training as an auxiliary objective (rather than entirely replacing the objective) we vary the value of λ to interpolate between two modes: $\lambda = 0$ is simply *QMIX (Attention)*, while $\lambda = 1$ trains exclusively with random factorization. Our results on 3-8sz (Figure 5 (bottom right)) show that, similar to regularization methods such as Dropout (Srivastava et al., 2014), there is a sweet spot where performance is maximized before collapsing catastrophically. Training exclusively with random factorization does not learn anything significant. This failure is likely due to the fact that we use the full context in our targets for learning with imagined scenarios as well as when executing our policies, so we still need to learn with it in training.

Qualitative Example of Common Pattern Finally, we visualize an example of the sort of common patterns that **REFIL** is able to leverage (Fig. 5 (left)). Zealots (the only melee unit present) are weak to Colossi, so they learn to hang back and let other units engage first. Then, they jump in and intercept the enemy Zealots while all other enemy units are preoccupied, leading to a common pattern of a Zealot vs. Zealot skirmish (highlighted at $t=15$). **REFIL** enables behaviors learned in these types of sub-groups to be applied more effectively across all tasks. By sampling groups from all entities randomly, we will occasionally end up with sub-groups that include only Zealots, and the value function predictions learned in these sub-groups can be applied not only to the task at hand, but to any task where a similar pattern emerges.

5. Related Work

Multi-agent reinforcement learning (MARL) is a broad field encompassing cooperative (Foerster et al., 2018; Rashid et al., 2018; Sunehag et al., 2018), competitive (Bansal et al., 2018; Lanctot et al., 2017), and mixed (Lowe et al., 2017; Iqbal & Sha, 2019) settings. This paper focuses on cooperative MARL with centralized training and decentralized execution (Oliehoek et al., 2016, CTDE). Our approach utilizes value function factorization, an approach aiming to simultaneously overcome limitations of both joint (Hausknecht, 2016) and independent learning (Claus & Boutilier, 1998) paradigms. Early attempts at value function factorisation require a priori knowledge of suitable per-agent team reward decompositions or interaction dependencies. These include optimising over local compositions of individual Q -value functions learnt from individual reward functions (Schneider et al., 1999), as well as summing individual Q -functions with individual rewards before greedy joint action selection (Russell & Zimdars, 2003). Recent approaches from cooperative deep multi-agent RL learn value factorisations from a single team reward function by treating all agents as independent factors, requiring no domain knowledge and enabling decentralized execution. Value-Decomposition Networks (VDN) (Sunehag et al., 2018) decompose the joint Q -value function into a sum of local utility functions used for greedy action selection. QMIX (Rashid et al., 2018; 2020) extends such additive decompositions to general monotonic functions. Some works extend QMIX to improve the expressivity of mixing functions (Son et al., 2019; Yang et al., 2020), learn latent embeddings to help exploration (Mahajan et al., 2019) or learn dynamic roles (Wang et al., 2020a), and encode knowledge of action semantics into network architectures (Wang et al., 2020b).

Several recent works have addressed the topic of generalization and transfer across related tasks with varying agent quantities, though the learning paradigms considered and assumptions made differ from our approach. Carion et al. (2019) devise an approach for assigning agents to tasks, assuming the existence of low-level controllers to carry out the

tasks, and show that it can scale to much larger tasks than those seen in training. Burden (2020) propose a transfer learning approach using convolutional neural networks and grid-based state representations to scale to tasks of arbitrary size. Wang et al. (2021) introduce a method to decompose action spaces into roles, which they show can transfer to tasks with larger numbers of agents by grouping new actions into existing clusters. They do not propose a model to handle the larger observation sizes, instead using a euclidean distance heuristic to observe a fixed number of agents. Several approaches devise attention or graph-neural-network based models for handling variable sized inputs and focus on learning curricula to progress on increasingly large/challenging settings (Long et al., 2020; Baker et al., 2019; Wang et al., 2020c; Agarwal et al., 2019). Most recently, Hu et al. (2021) introduce a method for handling variable-size inputs and action spaces and evaluate their model on single-task to single-task transfer. In contrast to these curriculum and transfer learning approaches, we focus on training simultaneously on multiple tasks and specifically develop a training paradigm for improving knowledge sharing across tasks.

6. Conclusion

In this paper we considered a multi-task MARL setting where we aim to learn control policies for variable-sized teams of agents. We proposed **REFIL**, an approach that regularizes value functions to share factors comprised of sub-groups of entities, in turn promoting generalization and knowledge transfer within and across complex cooperative multi-agent tasks. Our results showed that our contributions yield significant average performance improvements across these tasks when training on them concurrently, specifically through improving generalization across a wider variety of tasks.

Acknowledgements

This work is partially supported by NSF Awards IIS-1513966/ 1632803/1833137, CCF-1139148, DARPA Award#: FA8750-18-2-0117, FA8750-19-1-0504, DARPA-D3M - Award UCB-00009528, Google Research Awards, gifts from Facebook and Netflix, and ARO# W911NF-12-1-0241 and W911NF-15-1-0484. This project has also received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation programme (grant agreement number 637713). The experiments were made possible by a generous equipment grant from NVIDIA.

We thank Greg Farquhar for developing the initial version of the StarCraft environment with dynamic team compositions. We also thank Luisa Zintgraf, Natasha Jaques, and Robby Costales for their helpful feedback on the draft.

References

- Agarwal, A., Kumar, S., and Sycara, K. Learning transferable cooperative behavior in multi-agent teams. *arXiv preprint arXiv:1906.01202*, 2019.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autotutorials. In *International Conference on Learning Representations*, 2019.
- Bansal, T., Pachocki, J., Sidor, S., Sutskever, I., and Mordatch, I. Emergent complexity via multi-agent competition. In *International Conference on Learning Representations*, 2018.
- Böhmer, W., Kurin, V., and Whiteson, S. Deep coordination graphs. In *Proceedings of Machine Learning and Systems (ICML)*, pp. 2611–2622, 2020.
- Burden, N. Deep Multi-Agent Reinforcement Learning in Starcraft II. Master’s thesis, University of Oxford, 2020.
- Carion, N., Usunier, N., Synnaeve, G., and Lazaric, A. A structured prediction approach for generalization in cooperative multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 8128–8138, 2019.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning, December 2014*, 2014.
- Claus, C. and Boutilier, C. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI*, 1998(746-752):2, 1998.
- Foerster, J., Farquhar, G., Afouras, T., Nardelli, N., and Whiteson, S. Counterfactual multi-agent policy gradients. In *AAAI Conference on Artificial Intelligence*, 2018.
- Ha, D., Dai, A. M., and Le, Q. V. Hypernetworks. In *International Conference on Learning Representations*, 2017.
- Hausknecht, M. J. *Cooperation and Communication in Multiagent Deep Reinforcement Learning*. PhD thesis, The University of Texas at Austin, 2016.
- Hausknecht, M. J. and Stone, P. Deep recurrent q-learning for partially observable mdps. In *2015 AAAI Fall Symposium*, pp. 29–37, 2015.
- Hu, S., Zhu, F., Chang, X., and Liang, X. {UPD}et: Universal multi-agent {rl} via policy decoupling with transformers. In *International Conference on Learning Representations*, 2021.

- Iqbal, S. and Sha, F. Actor-attention-critic for multi-agent reinforcement learning. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2961–2970, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- Jiang, J. and Lu, Z. Learning attentional communication for multi-agent cooperation. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31, pp. 7254–7264. Curran Associates, Inc., 2018.
- Koller, D. and Parr, R. Computing factored value functions for policies in structured mdps. In *IJCAI*, volume 99, pp. 1332–1339, 1999.
- Lanctot, M., Zambaldi, V., Gruslys, A., Lazaridou, A., Tuyls, K., Perolat, J., Silver, D., and Graepel, T. A unified Game-Theoretic approach to multiagent reinforcement learning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 4193–4206. Curran Associates, Inc., 2017.
- Lee, J., Lee, Y., Kim, J., Kosiorok, A., Choi, S., and Teh, Y. W. Set transformer: A framework for attention-based permutation-invariant neural networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3744–3753, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- Lin, L.-J. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*, 8(3):293–321, 1992.
- Long, Q., Zhou, Z., Gupta, A., Fang, F., Wu, Y., and Wang, X. Evolutionary population curriculum for scaling multi-agent reinforcement learning. In *International Conference on Learning Representations*, 2020.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, O. P., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pp. 6382–6393, 2017.
- 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.
- 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, 2015.
- Oliehoek, F. A., Spaan, M. T., Vlassis, N., and Whiteson, S. Exploiting locality of interaction in factored dec-pomdps. In *Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, pp. 517–524, 2008.
- Oliehoek, F. A., Amato, C., et al. *A concise introduction to decentralized POMDPs*, volume 1. Springer, 2016.
- Prorok, A., Hsieh, M. A., and Kumar, V. The impact of diversity on optimal control policies for heterogeneous robot swarms. *IEEE Transactions on Robotics*, 33(2): 346–358, 2017.
- Rashid, T., Samvelyan, M., Schroeder, C., Farquhar, G., Foerster, J., and Whiteson, S. QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4295–4304, Stockholmssan, Stockholm Sweden, 10–15 Jul 2018.
- 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. *Journal of Machine Learning Research*, 21(178):1–51, 2020.
- Russell, S. and Zimdars, A. L. Q-decomposition for reinforcement learning agents. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning, ICML’03*, pp. 656–663, Washington, DC, USA, August 2003. AAAI Press. ISBN 978-1-57735-189-4.
- Samvelyan, M., Rashid, T., Schroeder de Witt, C., Farquhar, G., Nardelli, N., Rudner, T. G., Hung, C.-M., Torr, P. H., Foerster, J., and Whiteson, S. The starcraft multi-agent challenge. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 2186–2188. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- Schneider, J., Wong, W.-K., Moore, A., and Riedmiller, M. Distributed Value Functions. In *In Proceedings of the Sixteenth International Conference on Machine Learning*, pp. 371–378. Morgan Kaufmann, 1999.
- Schroeder de Witt, C., Foerster, J., Farquhar, G., Torr, P., Boehmer, W., and Whiteson, S. Multi-agent common knowledge reinforcement learning. In *Advances in Neural Information Processing Systems 32*, pp. 9927–9939. Curran Associates, Inc., 2019.

- 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. In *International Conference on Machine Learning*, pp. 5887–5896, 2019.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., and Graepel, T. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18*, pp. 2085–2087, Richland, SC, 2018. International Foundation for Autonomous Agents and Multiagent Systems.
- van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *Proceedings of the 13th AAI Conference on Artificial Intelligence*, pp. 2094–2100, 2016.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 6000–6010, 2017.
- Wang, T., Dong, H., Lesser, V., and Zhang, C. Roma: Multi-agent reinforcement learning with emergent roles. In *Proceedings of the 37th International Conference on Machine Learning*, 2020a.
- Wang, T., Gupta, T., Mahajan, A., Peng, B., Whiteson, S., and Zhang, C. {RODE}: Learning roles to decompose multi-agent tasks. In *International Conference on Learning Representations*, 2021.
- Wang, W., Yang, T., Liu, Y., Hao, J., Hao, X., Hu, Y., Chen, Y., Fan, C., and Gao, Y. Action semantics network: Considering the effects of actions in multiagent systems. In *International Conference on Learning Representations*, 2020b.
- Wang, W., Yang, T., Liu, Y., Hao, J., Hao, X., Hu, Y., Chen, Y., Fan, C., and Gao, Y. From few to more: Large-scale dynamic multiagent curriculum learning. In *AAAI Conference on Artificial Intelligence*, 2020c.
- Yang, Y., Hao, J., Liao, B., Shao, K., Chen, G., Liu, W., and Tang, H. Qatten: A general framework for cooperative multiagent reinforcement learning. *arXiv preprint arXiv:2002.03939*, 2020.