

Revisiting the Gumbel-Softmax in MADDPG

Callum Rhys Tilbury
University of Edinburgh
Edinburgh, United Kingdom
callum@tilbury.co.za

Filippos Christianos
University of Edinburgh
Edinburgh, United Kingdom
f.christianos@ed.ac.uk

Stefano V. Albrecht
University of Edinburgh
Edinburgh, United Kingdom
s.albrecht@ed.ac.uk

ABSTRACT

MADDPG is an algorithm in multi-agent reinforcement learning (MARL) that extends the popular single-agent method, DDPG, to multi-agent scenarios. Importantly, DDPG is an algorithm designed for continuous action spaces, where the gradient of the state-action value function exists. For this algorithm to work in discrete action spaces, discrete gradient estimation must be performed. For MADDPG, the Gumbel-Softmax (GS) estimator is used—a reparameterisation which relaxes a discrete distribution into a similar continuous one. This method, however, is statistically biased, and a recent MARL benchmarking paper suggests that this bias makes MADDPG perform poorly in grid-world situations, where the action space is discrete. Fortunately, many alternatives to the GS exist, boasting a wide range of properties. This paper explores several of these alternatives and integrates them into MADDPG for discrete grid-world scenarios. The corresponding impact on various performance metrics is then measured and analysed. It is found that one of the proposed estimators performs significantly better than the original GS in several tasks, achieving up to 55% higher returns, along with faster convergence.

KEYWORDS

Discrete Gradient Estimation, Gumbel-Softmax, Multi-Agent Reinforcement Learning, MADDPG

1 INTRODUCTION

In recent years, interest in the field of reinforcement learning (RL) has grown markedly. Though in existence for over three decades, the discipline’s recent integration with deep learning, often called deep RL, has catalysed a renewed hope for its capabilities. Such excitement is certainly warranted: deep RL algorithms have been excelling consistently on a wide range of challenges, many of which seemed unthinkable in the past. Commonly cited feats include conquering popular games, both modern and ancient [5, 37, 44, 47].

An important type of problem in RL is where not only a single agent acts, but multiple agents. These agents act together, either adversarially, co-operatively, or some combination thereof. Broadly, this paradigm is termed *multi-agent* RL (MARL). Algorithms developed for single-agent contexts can be applied for multiple agents, where each agent simply learns independently, *e.g.* independent Q-learning [41]. Though suitable for some tasks, this approach struggles to learn desired behaviours in certain complex environments, such as those with partial observability [32]. As an alternative, researchers have developed MARL-specific algorithms—either by extending extant single-agent approaches to multi-agent scenarios, or by developing new algorithms altogether.

One of the earliest algorithms proposed for deep MARL (that is, MARL with the integration of deep learning) was MADDPG, by Lowe et al. [27]. In this work, the authors extended the single-agent DDPG [24] method, which is itself an extension of the DPG [38] method, to multi-agent scenarios. Crucially, DPG and its descendants are designed to work only with *continuous* action spaces, where each action comes from an uncountable, continuous domain; *e.g.* the torque applied to motor. The alternative is a *discrete* action space, which has countable set of possibilities; *e.g.* choosing to go up or down. The restriction to continuous domains is because the gradient of the state-action value function, taken with respect to the action, must exist. In a discrete action context, this gradient does not exist.

Despite this restriction, it seems that the authors of MADDPG desired a unified algorithm, which could be applied to both continuous and discrete problems, while still building on the foundations of DPG. To enable MADDPG to work in discrete situations, then, a mathematical trick was applied: the *Gumbel-Softmax* (GS) reparameterisation [18, 28]. Essentially, this trick ‘relaxes’ the discrete, non-differentiable action space into a somewhat equivalent, continuous space—thus allowing an approximation of the gradient to exist. Relaxing the space in this way, however, introduces statistical bias into the gradient computation.

Recently, a benchmarking paper by Papoudakis et al. [32] found that MADDPG achieved decent performance in certain MARL environments, but performed markedly worse in grid-world situations, where the action space is discrete. The authors suggested that this degradation of performance may be due to the bias from the GS.

Interestingly, this field of *discrete gradient estimation* appears in a host of contexts outside of MARL [21, 35, 48]. As a result, a wealth of alternatives has been proposed for the GS, many of which focus on lowering the bias it introduces [10, 23, 33]. As of yet, though, it seems that not many of these techniques have been integrated into MARL, and certainly not into MADDPG. Accordingly, we ask: *Can alternative discrete gradient estimation methods improve the performance of MADDPG in grid-world environments, when compared to the original Gumbel-Softmax reparameterisation?*

In this paper, we study four alternative estimators: two with simple changes to the existing GS (decreasing or annealing the relaxation temperature used), and two novel methods from the literature [10, 33]. We test these estimators with MADDPG in nine grid-world tasks, as a subset of those in the benchmarking paper by Papoudakis et al. [32]. We find that the Gapped Straight Through (GST) [10] estimator yields the most significant improvements, with faster convergence and up to 55% higher returns. Finally, we look at the variance of the gradients from the GST compared to those from the GS for one of the tasks, to help us understand the performance gains observed. We release our code online¹.

Proc. of the Adaptive and Learning Agents Workshop (ALA 2023), Cruz, Hayes, Wang, Yates (eds.), May 29-30, 2023, London, UK, <https://alaworkshop2023.github.io/>. 2023.

¹<https://github.com/uoe-agents/revisiting-madddpg>

2 BACKGROUND

We consider a multi-agent learning problem, which we model as a partially-observable stochastic game (POSG) [16, 36], operating in discrete time-steps, with a set of N agents, $\mathcal{N} = \{1, \dots, N\}$. Denote the state-space as \mathcal{S} , the joint-action space as $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_N$, and the joint-observation space as $\mathcal{O} = \mathcal{O}_1 \times \dots \times \mathcal{O}_N$. At each time-step, each agent $i \in \mathcal{N}$ takes an action $a_i \in \mathcal{A}_i$, and perceives a local observation, $o_i \in \mathcal{O}_i$, which depends on the current state and the joint-action taken. We define a transition function, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$, which describes the probability of transitioning from one state to another, given a joint action. We further define a reward function for each agent, $\mathcal{R}_i : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$. Let the reward given to agent i at time-step t be denoted as $r_i^{(t)}$. We define the *return* for an agent i as its discounted cumulative reward, $G_i = \sum_{t=0}^T \gamma^t r_i^{(t)}$, where T is the number of time-steps in an episode, and $\gamma \in (0, 1]$ is a discounting factor—controlling how much we care about future rewards relative to current rewards. The game begins in an initial state, which depends on the distribution $\rho = \mathcal{S} \mapsto [0, 1]$.

Denote each agent’s policy (that is, what action it should take in a given state) as π_i , with the set of all policies being $\pi = \{\pi_1, \dots, \pi_N\}$. The objective in MARL, then, is to find policies such that the return of each agent i , following π_i , is maximised with respect to the other agents’ policies, $\pi_{-i} := \{\pi \setminus \pi_i\}$. That is, we aim to find an optimal set of policies, π , such that

$$\forall i : \pi_i \in \arg \max_{\hat{\pi}_i} \mathbb{E}[G_i \mid \hat{\pi}_i, \pi_{-i}] \quad (1)$$

Our focus here is on *policy gradient* methods, where each agent’s policy is explicitly encoded as a parametric distribution over actions given the state: $\pi_i(a \mid s; \theta)$. Assuming the policy is differentiable with respect to its parameters (*i.e.* $\frac{\partial \pi(a \mid s; \theta)}{\partial \theta}$ exists), an optimal policy can be found through gradient ascent of the expected return. To optimise this expectation, one can apply the *Stochastic Policy Gradient Theorem* [40]:

$$\nabla_{\theta} \mathbb{E}_{\pi} [r(s, a)] = \mathbb{E}_{\pi} [\nabla_{\theta} \log \pi(a \mid s; \theta) Q^{\pi}(s, a)] \quad (2)$$

where $Q^{\pi}(s_i, a_i) := \mathbb{E}_{\pi} [G_i \mid s_i, a_i]$, the state-action value function.

Notice that the gradient, ∇_{θ} , is taken with respect to the policy parameters, which exists if our policy is designed to be sufficiently smooth. This approach is simple and has been popular; *e.g.* estimating $Q^{\pi}(s, a)$ with the sampled return yields the REINFORCE algorithm [46]. However, it has been shown to suffer from high variance [30].

Taking a different angle, Silver et al. [38] introduced the Deterministic Policy Gradient (DPG) method. Here, instead of trying to learn a stochastic policy, the policy is deterministic: $a = \mu(s; \theta)$. The result is the *Deterministic Policy Gradient Theorem*:

$$\nabla_{\theta} \mathbb{E}_{\mu} [r(s, a = \mu(s; \theta))] = \mathbb{E}_{\mu} [\nabla_{\theta} \mu(s; \theta) \nabla_a Q^{\mu}(s, a)] \quad (3)$$

Notice the key difference now: a gradient is also taken with respect to the actions, ∇_a . Importantly, this difference means one cannot use DPG in discrete action problems, for this gradient does not exist.

DDPG [24] extends the DPG method by integrating it with deep neural networks, and incorporating techniques introduced in the DQN paper [29]: separate ‘behaviour’ and ‘target’ networks with Polyak averaging, and the use of a replay buffer. MADDPG [27] then applies DDPG to the multi-agent setting. At the heart of MADDPG is

the *Centralised Training, Decentralised Execution* (CTDE) paradigm. Here, each agent’s state-action value function (*i.e.* critic) is learned in a centralised manner, endowed with the joint observations and joint actions taken: $Q_i(\mathbf{o}, \mathbf{a})$. These critic networks, parameterised by ϕ_i , are updated by minimising the loss, \mathcal{L}_c :

$$\nabla_{\phi_i} \mathcal{L}_c(\phi_i) = \nabla_{\phi_i} \left(r_i + \gamma \bar{Q}_i(\mathbf{o}', \bar{\mu}_1(o'_1), \dots, \bar{\mu}_N(o'_N)) - Q_i(\mathbf{o}, \mathbf{a}) \right)^2$$

where $\bar{\mu}_i$ denotes the target policy networks, and \bar{Q}_i the target critic networks.

Whereas the critic networks in MADDPG are centralised, notice that the policy networks are decentralised—with each agent considering only their local observation, $\mu_i(o_i)$. These networks are updated using the sampled policy gradient:

$$-\nabla_{\theta_i} \mathcal{L}_a(\theta_i) = \nabla_{\theta_i} \mu_i(o_i) \nabla_{a_i} Q_i(\mathbf{o}, a_1, \dots, a_i, \dots, a_N) \Big|_{a_i = \mu_i(o_i)}$$

Notice again, as in (3), that a gradient is taken with respect to the action, $\nabla_{a_i} Q_i$, which does not exist with discrete actions. The authors of both DPG and DDPG paid no attention to this restriction, for their methods were presented explicitly for continuous action problems. In contrast, the MADDPG algorithm was presented for both continuous and discrete cases. To enable the gradient to exist in the latter, the authors used the *Gumbel-Softmax* (GS) trick [18, 28] to the discrete actions taken, thus enabling an approximation of $\nabla_{a_i} Q_i$ to exist.

3 DISCRETE GRADIENT ESTIMATION

In this work, we consider four alternatives to the original GS trick: two of which apply simple changes to the extant method, and two of which are novel methods drawn from the literature. We now provide brief explanations of these techniques, starting with the original GS method.

We consider a situation of a parametric discrete distribution, $p(a; \zeta)$, specified by an unconstrained vector of parameters, $\zeta \in \mathbb{R}^N$. In this context, these parameters represent the outputs of a policy network, where actions are sampled as $a \sim p(a; \zeta)$. We seek an estimation of $\nabla_a p(a; \zeta)$, but since a is discrete, we must ‘relax’ the distribution for this gradient to exist.

3.1 Baseline (STGS-1)

The GS method was introduced concurrently by Jang et al. [18] and Maddison et al. [28], as a differentiable approximation of the arg max function. As its name suggests, a tempered *softmax* is used, with a temperature parameter, $\tau > 0$: $\text{softmax}_{\tau}(x) := \text{softmax}(\frac{x}{\tau})$. In the limit of $\tau \rightarrow 0$, this operation is equivalent to the arg max, and thus the GS approaches the original distribution. Conversely, as $\tau \rightarrow \infty$, the GS approaches a uniform distribution, where each category is equally-likely. The temperature thus controls the degree of relaxation.

Using $\xi_{GS}(\cdot)$ to notate the GS, the relaxed distribution is:

$$\xi_{GS}(p(a_i; \zeta)) = \text{softmax}_{\tau}(\zeta_i + g_i) \quad , \quad g_i \sim G(0, 1) \quad (4)$$

where g_i is noise sampled from the Gumbel distribution [31].

By relaxing the distribution in this way, it becomes differentiable, meaning we can incorporate it into a gradient-based optimisation procedure. There is a downside, however: in relaxing, we introduce statistical bias [33]. To understand this bias intuitively, consider

again the limit $\tau \rightarrow \infty$, where the distribution becomes uniform. In such a case, we have removed all parametric information, ζ , about our problem—each category simply has a probability of $1/N$. Hence, as we relax, we also steer further away from the original distribution. Herein lies a trade-off: turning the temperature too low means having extreme gradients (or non-existent gradients when $\tau = 0$), but turning the temperature too high means introducing a large bias.

Though such a bias is inevitable when relaxing the distribution, there is an easy improvement to the vanilla GS. By naively applying the relaxation, we introduce bias in both the forward pass (when we sample from the distribution) and in the backward pass (when we calculate the gradients, *e.g.* for updating our neural network). However, it is only the latter that requires differentiability. Hence, building on the so-called *Straight-Through* estimator proposed by Bengio et al. [4], Jang et al. [18] also introduce the STGS estimator—where, in the backward pass, the GS relaxation is applied, but in the forward pass, the original arg max operation is used. MADDPG uses this variant, and all further discussions focus on it.

In both the original MADDPG paper [27] and the benchmarking paper by Papoudakis et al. [32], it seems that the authors simply use a temperature of 1.0 for the STGS relaxation². Thus, we use this configuration as our baseline, denoted as **STGS-1**.

3.2 Lower Temperature Gumbel Softmax (STGS-T)

Recall that Papoudakis et al. [32] suggest it is the bias that is problematic in the STGS, which is positively correlated to the temperature of relaxation—lowering the temperature lowers the bias. Accordingly, our first alternative estimator is the simplest: the STGS estimator with a temperature of $\tau < 1.0$, where τ is a tunable hyperparameter, denoted as **STGS-T**.

3.3 Temperature-Annealed Gumbel Softmax (TAGS)

The *exploitation-exploration dilemma* [39]—which describes the trade-off between taking actions that yield known, good rewards (exploiting), and taking actions which may or may not yield better rewards (exploring)—is often discussed in RL. In the continuous-action formulation of MADDPG [27, Appx: Alg. 1], exploration is achieved via the addition of noise to the policy output: $a_i = \mu_i(o_i) + \eta_i$, where η is drawn from some random process (originally discussed in DDPG [24]). However, in discrete cases, the STGS itself provides some degree of exploration, since relaxing the distribution places some probability mass onto other actions. As a result, the amount of exploration is controlled by the temperature parameter: more relaxation implies more exploration. Notice, then, the coupling between the exploration achieved and the bias introduced.

Since exploration is usually desirable in the beginning of a training procedure, we propose setting the temperature to be high early-on, and then annealing it to be lower over time. This strategy allows agents to explore, while still reducing the bias in later stages of

²Nothing is explicitly stated about the temperature used in these papers; we are making such conclusions by looking at their code implementations: Link to snippet from Lowe et al. [27]; Link to snippet from Papoudakis et al. [32].

training. Huijben et al. [17] highlight temperature-annealing as a strategy incorporated by several authors in various experiments with the STGS. Specifically, they mention using an exponentially-decaying annealing scheme, which we adopt here. We define this estimator as the Temperature-Annealed Gumbel Softmax (**TAGS**).

3.4 Gumbel-Rao Monte Carlo (GRMCK)

The next estimator is drawn from the literature, entitled the Gumbel-Rao Monte Carlo (GRMC), by Paulus et al. [33]. Here, the authors seek a way to lower the STGS’s variance. Notating the gradient of the original STGS estimator as $\nabla_{\text{STGS}} = \nabla_a \zeta_{\text{STGS}}$, we have:

$$\nabla_{\text{STGS}} := \frac{d\text{softmax}_\tau(\zeta + g)}{da} \quad (5)$$

With this notation, the authors propose the Gumbel-Rao estimator:

$$\nabla_{\text{GR}} := \mathbb{E} \left[\frac{d\text{softmax}_\tau(\zeta + g)}{da} \mid a \right] \quad (6)$$

That is, $\nabla_{\text{GR}} = \mathbb{E}[\nabla_{\text{STGS}} \mid a]$. This estimator is a Rao-Blackwell [6] version of the original STGS estimator. It can be shown that it thus enjoys the same mean as the STGS, but with lower (or at most, the same) variance:

$$\mathbb{E}[\|\nabla_{\text{GR}} - \nabla_\zeta\|^2] \leq \mathbb{E}[\|\nabla_{\text{STGS}} - \nabla_\zeta\|^2] \quad (7)$$

where ∇_ζ is the true gradient. For rigorous mathematical details about the estimator’s impact on variance, the reader is encouraged to see the full paper [33]. Recall that Papoudakis et al. [32] consider the bias of the estimator to be the problem for MADDPG, not the variance. Though guarantees are only made about the latter, the authors of the GRMC argue that with a lower variance, one can safely train the estimator at lower temperatures—*i.e.* with a lower bias. Empirically, they show this idea to be true.

Though theoretically appealing, there is still the challenge of actually computing $\mathbb{E}[d\text{softmax}_\tau(\zeta + g)/da \mid a]$ —indeed, a closed-form expression is shown to be difficult. Therefore, the authors provide a Monte Carlo estimate, with K samples, which they term the **GRMCK** estimator. They first show a distributional equivalence:

$$(\zeta_j + g_j \mid a) \stackrel{d}{=} \begin{cases} -\log(E_j) + \log Z(\zeta) & \text{if } j = i \\ -\log\left(\frac{E_j}{\exp(\zeta_j)} + \frac{E_i}{Z(\zeta)}\right) & \text{otherwise} \end{cases} \quad (8)$$

where a is a one-hot sample with a 1 at index i , E_j are i.i.d. samples from the exponential distribution, and $Z(\zeta) = \sum_j \exp(\zeta_j)$.

They then define the GRMCK estimator as:

$$\nabla_{\text{GRMCK}} := \frac{1}{K} \sum_k \frac{d\text{softmax}_\tau(\zeta + g_k)}{da}, \quad g_k \sim (\zeta + g \mid a) \quad (9)$$

In other words, we first sample $a \sim p(a; \zeta)$, and then average over K Gumbel noise samples conditioned on a .

3.5 Gapped Straight-Through (GST)

The final estimator considered is the most recently introduced: the Gapped Straight Through (GST), by Fan et al. [10]. Here, the authors find that the Gumbel randomness used in the STGS and GRMCK can be replaced with two deterministic perturbations, resulting in an estimator with lower variance.

As in the GRMC estimator, we first draw $a \sim p(a; \zeta)$ —a one-hot representation of the selected action—for the straight-through sample. In GRMC, we would then perturb each of the logits, ζ , with Gumbel noise conditioned on a ; now, we perturb with two functions, $m_1(\zeta, a)$ and $m_2(\zeta, a)$. Detailed justifications of these functions are given in the paper itself [10].

Firstly, we desire *consistency* in the estimator: we want the sample conditioned on a to have the same largest logit as the input distribution; *i.e.* $\max_j \zeta_j$. To this end, the first perturbation, m_1 , pushes the sample to the correct realisation, if necessary:

$$m_1(\zeta, a) = (\max_j \zeta_j - \langle \zeta, a \rangle) \cdot a \quad (10)$$

where $\langle \cdot, \cdot \rangle$ indicates the inner product. Consider how this works: if a has already selected the largest logit, then $\langle \zeta, a \rangle = \max_j \zeta_j$, and $m_1 = 0$. If not, then $m_1 \neq 0$, and the sample is moved in the direction of the largest logit.

If non-zero, the first perturbation makes the largest logit the same as the a -selected logit. However, we also want a *strict gap* between these values—that is, we want the unselected logits to be smaller. Accordingly, we define m_2 to create a gap of κ between them:

$$m_2(\zeta, a) = -(\kappa + \zeta - \max_j \zeta_j)_+ \odot (1 - a) \quad (11)$$

where $(x)_+ := \max(0, x)$, \odot indicates the Hadamard product. The value of κ can usually be set to 1.0 [10]. Here, the term $(1 - a)$ takes all the unselected logits in the one-hot representation, and moves their parameter values away from the selected logit, with a gap of at least κ .

With these perturbation functions defined, the GST estimator is then:

$$\check{\zeta}_{\text{GST}}(p(a; \zeta)) = \text{softmax}_\tau(\zeta + m_1(\zeta, a) + m_2(\zeta, a)) \quad (12)$$

4 EXPERIMENTAL METHODS

4.1 Environments

To test the performance of the proposed gradient estimators compared to the original STGS estimator, we train on two grid-world environments, across a total of nine tasks (*i.e.* configurations). For simplicity, we choose to focus solely on co-operative contexts, where agents are working together to maximise their cumulative reward—we feel this approach is sufficient to answer our research question. We use a sensible subset of the choices made by Papoudakis et al. [32] in their benchmarking paper: seven tasks in Level-Based Foraging (LBF)³ [1, 8], and two in Multi-Robot Warehouse (RWARE)⁴ [2, 8].

4.2 Evaluation Metrics

To understand the success (or failure) of the new gradient estimation techniques compared to the original STGS approach, metrics for evaluation are now defined.

Returns: Maximum & Average. Recall that we define our MARL goal, in (1), as trying to find an optimal set of policies, such that each agent maximises their expected return with respect to the other agents’ policies. Since we focus on co-operative situations

for this paper, we simply consider the sum of the achieved returns from all agents. Importantly, we are not concerned with the returns achieved here relative to those achieved in, *e.g.*, the MARL benchmarking paper by Papoudakis et al. [32]. Instead, for cogent and consistent analysis, we focus solely on the relative performance of the various estimators against each other.

For this evaluation metric, we run the MADDPG algorithm in each task, with each of the proposed gradient estimators. We train the algorithm for a fixed number of time-steps, updating the networks with a defined period. Throughout training, we evaluate the achieved returns 100 times every 50 000 time-steps. Each training iteration is done over five random seeds and a 95% confidence interval is calculated over the results.

Under this heading, we consider two distinct aspects of the achieved returns, following the lead of Papoudakis et al. [32]. Firstly, we consider the *maximum* return: the evaluation time-step at which the return, averaged over the five seeds, is highest—indicating the peak performance of the algorithm when using a given estimator. Secondly, we consider the *average* return: the mean of the evaluation returns over all time-steps and seeds, for a given estimator in a given task—a proxy for understanding not just the magnitude of the returns, but how quickly the training converges.

Compute Time. Though this paper revolves around—and is motivated by—the MADDPG algorithm, notice that the gradient estimators can also be compared in isolation. That is, when comparing the computational burden of the various estimation procedures, we need not integrate them into the broader MADDPG problem. Instead, we can take a closer look solely at each estimator’s performance, unhindered by potential bottlenecks elsewhere.

Accordingly, we define here a simple, toy problem for the estimators. We define a set of input logits, ζ , of various dimensionalities, and measure the time it takes for each estimator to calculate the corresponding relaxations. Because STGS-1, STGS-T, and TAGS all have the same underlying mechanics, we consider these under the single umbrella of the STGS. For the GRMCK, we consider three values of K : 1, 10, and 50. For each dimensionality, the estimation procedure is repeated 10 000 times, over five different logit instances. These results are reported over a 95% confidence interval.

Gradient Variance. Suppose one of the alternative gradient estimation techniques performs significantly better or worse than the original STGS, based on the returns achieved. The natural follow-up question is: why? As an initial step to answering this question, we choose one of the tasks where there is a notable difference in performance between two estimators: between the baseline STGS-1 method and one which performs much better (or worse). We then retrain the MADDPG algorithm in this task, using each of the two estimators, now logging the variance of the computed gradients across each training mini-batch, over the course of training.

We hypothesise that *uninformative* gradients, *i.e.* those due to a poor discrete gradient estimator, will yield a mini-batch with low variance, since there are no elements in particular which ‘stand out’. In contrast, we believe that *informative* gradients will have higher variance across the mini-batch, for the opposite reason. We hope for this metric to stimulate future discussion into why we might observe a difference in estimator performance.

³Code: <https://github.com/semitable/lb-foraging>

⁴Code: <https://github.com/semitable/robotic-warehouse>

4.3 Training Details

Hyperparameter tuning is an important, though time-consuming, component of training RL algorithms. For simplicity, then, the optimal hyperparameters for the core MADDPG algorithm suggested by Papoudakis et al. [32] are adopted here, mostly without any changes. Our hyperparameter search is thus limited to be over the novel gradient estimation techniques and their associated parameters. Bayesian optimisation [11] is performed for this search, and we use search-range suggestions from the literature, when available [17, 33]. Each parameter is optimised for one task in a particular environment, and then used for all other tasks in that environment. The resulting hyperparameters can be found online⁵, along with the code used for all experiments.

5 EXPERIMENTAL RESULTS & DISCUSSION

5.1 Returns: Maximum & Average

Table 1 shows the maximum and average returns achieved using each gradient estimation technique, across each of the nine tasks. Discussion of these results follows, per environment, with plots shown when relevant.

Table 1: Maximum returns (Average returns) shown across all tasks and all algorithms, presented with a 95% confidence interval over 5 seeds. Bold indicates the best performing metric for a situation. An asterisk (*) indicates that a given metric is not significantly different from the best performing metric in that situation, based on a heteroscedastic, two-sided t-test with 5% significance. Under each task name is the number of time-steps used for training.

	Tasks	STGS-1	STGS-T	TAGS	GRMCK	GST
LBF	8x8-2p-2f-c [5M]	1.00 ± 0.00 (0.88 ± 0.05)	1.00 ± 0.01 (0.91 ± 0.04)	1.00 ± 0.01 (0.87 ± 0.05)	1.00 ± 0.00 (0.88 ± 0.05)	1.00 ± 0.00 (0.89 ± 0.04)
	8x8-2p-2f-2s-c [5M]	0.79 ± 0.07* (0.65 ± 0.04)	0.83 ± 0.03 (0.66 ± 0.04)	0.78 ± 0.03* (0.62 ± 0.04)	0.81 ± 0.05* (0.67 ± 0.04)	0.81 ± 0.02* (0.68 ± 0.03)
	10x10-3p-3f [6M]	0.75 ± 0.03 (0.58 ± 0.04)	0.75 ± 0.03 (0.59 ± 0.03)	0.74 ± 0.06 (0.51 ± 0.04)	0.71 ± 0.07 (0.57 ± 0.03)	0.79 ± 0.04 (0.66 ± 0.03)
	10x10-3p-3f-2s [6M]	0.55 ± 0.05* (0.48 ± 0.01)	0.58 ± 0.06 (0.48 ± 0.01)	0.54 ± 0.03 (0.46 ± 0.01)	0.56 ± 0.03* (0.49 ± 0.01)	0.56 ± 0.05* (0.50 ± 0.01)
	15x15-3p-5f [7.5M]	0.24 ± 0.02 (0.12 ± 0.01)	0.28 ± 0.06 (0.15 ± 0.01)	0.20 ± 0.03 (0.08 ± 0.01)	0.26 ± 0.05 (0.16 ± 0.01)	0.31 ± 0.04 (0.20 ± 0.01)
	15x15-4p-3f [7.5M]	0.79 ± 0.03 (0.54 ± 0.04)	0.79 ± 0.06 (0.58 ± 0.04)	0.77 ± 0.06 (0.45 ± 0.04)	0.79 ± 0.04 (0.58 ± 0.04)	0.83 ± 0.04 (0.67 ± 0.03)
	15x15-4p-5f [7.5M]	0.33 ± 0.06 (0.13 ± 0.02)	0.46 ± 0.12 (0.22 ± 0.02)	0.24 ± 0.05 (0.10 ± 0.01)	0.43 ± 0.06 (0.21 ± 0.02)	0.48 ± 0.06 (0.30 ± 0.02)
RWARE	tiny 2ag [7.5M]	1.37 ± 0.22* (0.55 ± 0.07)	1.37 ± 0.49* (0.64 ± 0.08)	1.50 ± 0.46 (0.60 ± 0.08)	1.37 ± 0.40* (0.65 ± 0.07)	1.40 ± 0.58* (0.65 ± 0.07)
	tiny 4ag [7.5M]	2.68 ± 0.49 (0.84 ± 0.12)	3.18 ± 0.60 (1.09 ± 0.15)	2.23 ± 0.50 (0.78 ± 0.11)	3.17 ± 0.85 (1.15 ± 0.15)	4.16 ± 0.97 (1.82 ± 0.21)

Level-Based Foraging. We consider now the seven LBF tasks. Firstly, we look at two tasks with two agents over an 8 × 8 grid: one with full-observability (8x8-2p-2f), and one with partial-observability (8x8-2p-2f-2s). Notice in these results, in Table 1, that performance differences across the estimation techniques is statistically insignificant. In 8x8-2p-2f, we see that STGS-T trains marginally faster than

the other approaches, and in 8x8-2p-2f-2s, we see that TAGS trains marginally slower than the other approaches—both based on the average returns observed. Nonetheless, each algorithm arrives at a similar maximum return. Due to the insignificance of this result, the training curves are uninteresting, and are not plotted here.

We next look at two tasks with three agents over an 10 × 10 grid, with similar situations as before: one with full-observability (10x10-3p-3f), and one with partial-observability (10x10-3p-3f-2s). Plots of the evaluation returns over the duration of training are given in Figure 1.

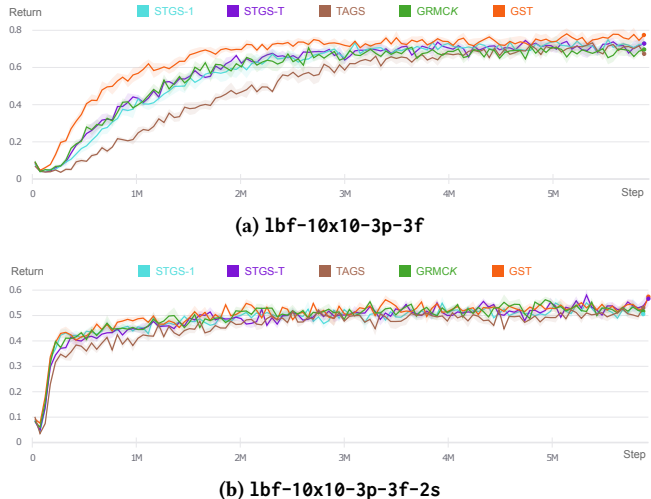


Figure 1: Evaluation returns for two LBF tasks (10 × 10) over the training period, where the shaded region indicates the standard error as calculated over 5 seeds.

In the first task, seen in Figure 1a, we see an improvement with a novel gradient estimation technique: the GST achieves the highest maximum and average returns for the task, beating the baseline STGS-1 method with statistical significance. It is clear in the figure how training with the GST converges faster than with the other methods. STGS-T and GRMCK perform similarly to the baseline. TAGS, however, performs much worse in average returns—*i.e.* it converges slower for the task—though it eventually achieves similar maximum returns.

In the task with partial observability, seen in Figure 1b, we are less successful—the alternative techniques achieve statistically similar returns to the baseline, across both maximum and average metrics. The exception is TAGS, which again performs *worse* than the baseline, though not markedly so.

We consider now the remaining three tasks in LBF, with a fully-observable, 15 × 15 grid, with three or four agents: 15x15-3p-5f, 15x15-4p-3f, and 15x15-4p-5f. We note that MADDPG performed particularly poorly in these larger, more-complex LBF situations, according to the benchmarking paper by Papoudakis et al. [32]. The training curves for each of these tasks is given in Figure 2.

We notice here significant improvements over the baseline. Yet again, TAGS markedly underperforms, both in maximum and average returns; but the other estimators perform well. STGS-T and GRMCK beat the baseline in average returns for 15x15-3p-5f and

⁵<https://github.com/uoe-agents/revisiting-madpg>

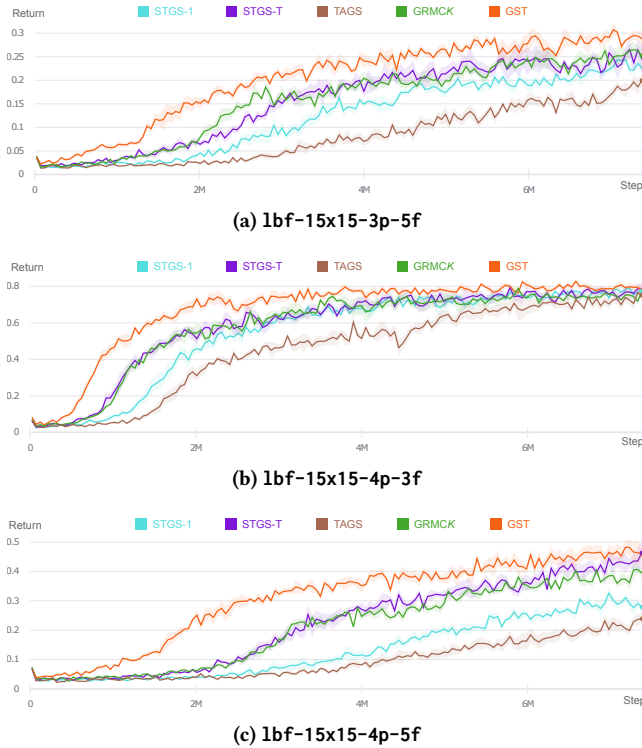


Figure 2: Evaluation returns for three LBF tasks (15×15) over the training period, where the shaded region indicates the standard error as calculated over 5 seeds.

15x15-4p-3f, and in both average and maximum returns for 15x15-4p-5f. GST is superior throughout: across all three tasks, it yields significantly higher returns and converges faster than the baseline (and the other techniques). Indeed, these improvements are clearly noticeable in the plots provided.

Multi-Robot Warehouse. Next, we consider the RWARE environment for two tasks over a 10×11 grid: one with two agents (tiny-2ag) and one with four agents (tiny-4ag). Figure 3 shows the returns for these two environments, over the training period.

In tiny-2ag, we see insignificant differences across the estimation techniques, with each achieving similar maximum returns. The alternative techniques do converge slightly faster, particularly GRMCK and GST, with marginally higher average returns, but not by much.

In tiny-4ag, we see the most significant improvements yet. Barring TAGS, which somewhat underperforms, we notice substantial improvements from the other proposed estimators, for both average and maximum returns. GST triumphs once more, achieving 55% higher maximum returns over the baseline, and over double the average returns. This result is again clear in the plot, in Figure 3b.

Discussion. This section presented the results from training MADDPG with each of the proposed gradient estimation techniques, across nine tasks from two grid-world environments. In simpler tasks, the alternative techniques do not make a significant difference to the returns achieved. We suspect this outcome is because informative gradients are not as crucial in simple tasks. That

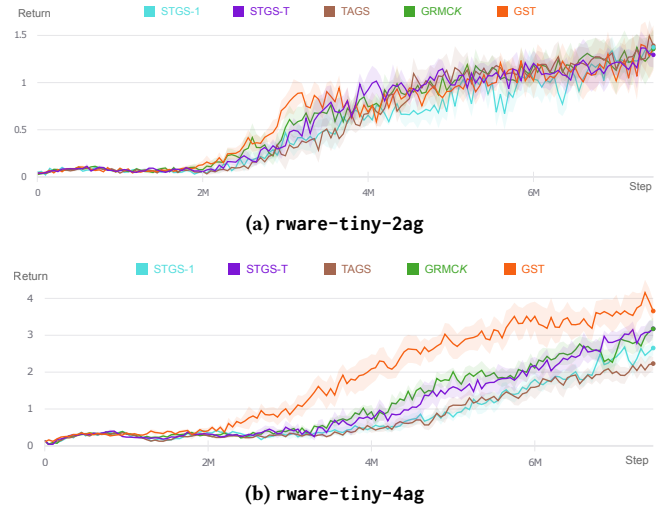


Figure 3: Evaluation returns for two RWARE tasks (tiny grid) over the training period, where the shaded region indicates the standard error as calculated over 5 seeds.

is, the gradient estimation is not a problematic aspect of the training, and limitations arise elsewhere in the mechanics of MADDPG.

Interestingly, in some of the more challenging tasks, particularly in LBF with a grid-size of 15×15 , and tiny-4ag in RWARE, we see significant improvements. We note that simply lowering the temperature (and hence, the gradient estimator’s bias), as in STGS-T, can improve the results somewhat—supporting the hypothesis that the bias introduced by the STGS is a problem for MADDPG. The Rao-Blackwellisation procedure of GRMCK also sees better returns and faster convergence. Much better than these alternatives, though, is the GST. With this estimator, we consistently see marked improvements across the two return metrics, and these are statistically significant.

Though a lower temperature seems to yield better returns, our results suggest that annealing the temperature, as in TAGS, performs poorly. This result may be due to the coupling of exploration and exploitation, as highlighted earlier, but more investigation is required. Alternatively, it may simply be the hyperparameters chosen—the annealing start and end points were taken from the advice of Huijben et al. [17]. It is conceivable that using lower values here may yield better returns, especially considering the improvements seen with STGS-T. Future work could also explore using alternative annealing schemes, or annealing with a different underlying estimator—*e.g.* one could try a temperature-annealed GST.

From these results, considering both the maximum returns and the time to convergence, we note that alternative gradient estimation techniques can indeed yield better returns when incorporated into MADDPG, particularly the recently proposed GST, from Fan et al. [10].

5.2 Compute Time

We now consider the computational requirements for each of the algorithms, using the toy problem outlined earlier. Recall that we

perform these tests for three classes of estimator: STGS (which accounts for STGS-1, STGS-T, and TAGS); GRMC (with three different K values); and GST. Table 2 shows the outcome of these tests.

Table 2: Time-per-relaxation, in μ s, for the three classes of gradient estimators, when using logits of various dimensionality as input. Results are given over a 95% confidence interval from 5 different logit instances, where each procedure is repeated 10 000 times. Underneath each metric, using round brackets, (\cdot) , we indicate how much slower the alternative techniques are, when compared to the baseline STGS.

↓ Dimensionality	STGS	GRMC-1	GRMC-10	GRMC-50	GST
3	135.28 ± 0.19 (1.0)	445.91 ± 11.38 (3.3)	446.65 ± 0.86 (3.3)	486.36 ± 1.7 (3.6)	357.56 ± 1.04 (2.64)
5	135.65 ± 0.51 (1.0)	438.52 ± 0.63 (3.23)	446.66 ± 0.7 (3.29)	501.9 ± 1.16 (3.7)	356.95 ± 0.6 (2.63)
10	135.83 ± 0.48 (1.0)	438.95 ± 1.01 (3.23)	451.65 ± 0.82 (3.33)	531.97 ± 0.45 (3.92)	356.52 ± 0.55 (2.62)
50	134.05 ± 0.89 (1.0)	440.46 ± 1.01 (3.29)	484.77 ± 1.29 (3.62)	765.56 ± 2.09 (5.71)	356.77 ± 1.77 (2.66)
100	139.23 ± 0.25 (1.0)	455.54 ± 2.97 (3.27)	520.06 ± 0.78 (3.74)	1055.18 ± 4.16 (7.58)	359.26 ± 1.26 (2.58)
1000	154.17 ± 0.45 (1.0)	533.12 ± 1.01 (3.46)	1060.74 ± 2.92 (6.88)	6171.59 ± 8.67 (40.03)	386.57 ± 1.36 (2.51)

We notice firstly that STGS scales well with dimensionality—the computational overhead when increasing the dimension does not change significantly. Even in the high-dimensional case of 1 000, the technique is only marginally slower. These benefits are common to the baseline STGS-1 approach, as well as the proposed techniques of STGS-T and TAGS.

Next, we see that GRMCK is at least three times slower than the baseline approach. Moreover, using a larger K value does, understandably, increase the computational burden of the relaxation. Though this effect is not substantial for low-dimensional inputs, for higher-dimensional problems, K has a marked impact—e.g. with $K = 50$, computation slows down considerably, becoming 40 times slower than the baseline for an input dimension of 1 000.

The computational burden of GST sits somewhat in-between the baseline, STGS, and the GRMCK approach. Importantly, though, this method also scales well with dimensionality, staying at just over 2.5 times slower than the baseline, irrespective of the input size—an attractive property.

From these results, and the insights drawn from the previous section, we can draft general guidelines for choosing an alternative estimator: if minimising the computational burden is paramount for a given problem, it may be worth using the STGS-T, for it has the same overhead as the STGS-1, and it does yield improvements in both the achieved returns and convergence time. However, if one can afford a more expensive relaxation procedure, the GST is a good fit—it is somewhat slower, but the benefits are significant. Since GRMCK is more expensive than the GST, yet usually yields lower returns, it does not seem like a sensible option as an estimator in either case. Granted, it could be that better performance comes from increasing K further (e.g. Paulus et al. [33] use $K = 1000$ in some of

their experiments), but the computational burden will only worsen in such a case, which is undesirable. TAGS, as it was presented here, should not be used.

5.3 Gradient Variance

We have previously seen marked improvements in some tasks when using the proposed gradient estimators, particularly the GST. We now stimulate further discussion by presenting a cursory look into *why*. We reconsider the LBF task of 15x15-4p-5f (see Figure 2c), and retrain with two algorithms: the baseline STGS-1, and the best performing alternative, GST. Figure 4 shows the variance of the gradients across mini-batches, for each of the layers in the policy networks, over the course of the training.




Figure 4: Plots showing the gradient variance (left: gradients of *weight* parameters, right: gradients of *bias* parameters), for each layer in the policy networks, for the 15x15-4p-5f task in LBF. The results are aggregated across the 4 agents in this task—the shaded region indicates the maximum and minimum values across agents, the solid line indicates the mean.

Immediately, we notice a trend in these graphs: the variance of the gradients, taken across a mini-batch, increases more rapidly for the GST algorithm than those for the baseline. Though not definitive, such results indicate that more informative gradients are being propagated through the policy networks. Informative gradients, in turn, allow the algorithm to achieve higher returns and converge faster—as evidenced in the results seen previously.

6 RELATED WORK

A core insight of this work is that discrete gradient estimation does not exist solely in the domain of RL. In fact, the original GS papers [18, 28] demonstrated the technique on problems such as structured output prediction and density estimation. Indeed, discrete gradients arise in a wide variety of contexts—including discrete variational auto-encoders [19, 35], hard attention [15, 48],

generative adversarial networks for text [22, 49], and convolutional networks [43]. As a result, the biased reparameterisation of the GS is problematic in a wide variety of domains, and accordingly, a significant research effort has focused on improving the method. In this paper, we drew two methods from the literature: the GRMCK [33] and the GST [10]. Various other approaches could, conceivably, be integrated into the MADDPG algorithm to handle discrete-action environments. We outline them briefly here.

The STGS, GRMCK, and GST are all instances of the *pathwise-derivative* approach to gradient estimation, which is underpinned by a more-general form of the Deterministic Policy Gradient Theorem, seen in (3). The Invertible Gaussian Reparameterisation [34] also belongs to this class, where Gaussian noise is used instead of Gumbel noise. Andriyash et al. [3], too, move away from using Gumbel noise in their method, and propose a simple piecewise-linear relaxation instead. Other methods steer further away from the STGS. For example, Lee et al. [23] generalise the reparameterisation trick through manifold sampling, and are able to create an unbiased and reduced-variance estimator. Lorberbom et al. [25] avoid the need to relax the categorical distribution altogether by applying the technique of *direct optimisation*.

The other approach to gradient estimation is *score-function* methods, which is underpinned by a more-general form of the Stochastic Policy Gradient Theorem, seen in (2). Alternatives here include: subtracting a baseline [13, 45]; using a Taylor expansion of a mean-field network, as in MuProp [14]; and using copula-based sampling, as in CARMS [9]. However, it is unclear how using a score-function method would change the underlying DPG mechanics when applied in MADDPG.

Authors have also combined score function and pathwise derivative methods, leveraging desirable qualities from both approaches. For example: using both REINFORCE and the GS in conjunction, as in REBAR [42]; training a surrogate neural network as a control variate, as in RELAX [12]; and using sampling without replacement [20].

7 CONCLUSION

This paper explored the impact of the Gumbel-Softmax (GS) reparameterisation [18, 28] on MADDPG [27] when applied to grid-world environments. Firstly, some necessary theoretical foundations were presented and the problem was framed in the context of the broader literature. Thereafter, we looked closely at the straight-through GS (STGS) and discrete gradient estimation more generally, highlighting the key concepts therein. After presenting a handful of candidate STGS alternatives—two with simple tweaks to the STGS, and two from the literature [10, 33]—these estimators were implemented into the MADDPG algorithm. A suite of nine MARL tasks across two environments was used for testing, and various metrics were analysed.

On some of the tasks—particularly the simpler ones, where MADDPG already performed well—no significant changes were observed, in terms of returns achieved and the speed to convergence. On other tasks though, particularly in the more challenging ones, substantial improvements occurred. It was found that even an easy change to the original STGS estimator, simply lowering the temperature parameter, yielded good results. The proposed temperature-annealing

scheme in TAGS, however, was shown to be a bad choice for the estimator—though we acknowledge a different set of hyperparameters may have helped here. The GRMCK estimator [33] showed promising results, but was hindered by a below-par computational burden. Finally, far superior to the other methods was the GST estimator [10]. This method achieved the best results across a range of tasks, with up to 55% higher returns, as well as faster convergence, when compared to the original STGS. Though it did introduce additional computational burden, at around 2.5 times slower than the STGS, the method nonetheless scaled well with dimensionality, and is certainly a viable technique for many use-cases.

We are now in a good position to support the suggestions made by Papoudakis et al. [32] in their benchmarking paper. Based on the empirical data observed, we agree that the bias of the STGS method is indeed problematic for MADDPG. As a result, by improving the estimator used—*i.e.* by lowering its bias—we can improve the returns achieved by MADDPG. To answer our research question from Section 1, then: yes, alternative discrete gradient-estimation techniques *can* improve the performance of MADDPG in discrete grid-worlds.

Notice the benefit of our findings. We can take the extant MADDPG algorithm, replace *only* the gradient estimation technique—that is, swap out, *e.g.*, the STGS for the GST, and leave everything else the same—and the resulting performance may likely improve. Though our algorithm becomes slightly more expensive computationally, we witness faster convergence and higher returns, with minimal development overhead.

8 FUTURE WORK

Many avenues of future work extend from this paper—we highlight a handful here.

Firstly, it would be useful to evaluate the performance of the proposed algorithms on a wider variety of tasks—in particular, the results in RWARE were promising but distinctly limited, with only two tasks tested. Moreover, only co-operative tasks were tested here, and adversarial configurations should be explored too.

Secondly, much more investigation ought to be done into why the GST is boasting better performance. Though an interesting foray, the analysis into the gradient variance was just a first step. Future research should continue to focus on the mechanics of the algorithms, and probe at various points.

Thirdly, the core MADDPG algorithm designed for this paper did not incorporate various extensions suggested in the literature, *e.g.* parameter sharing [7, 8]. Combining the benefits observed here with other strong extensions elsewhere would be an interesting exercise. Furthermore, a wider hyperparameter search, now with the alternative estimators involved too, may be helpful.

Finally, we note that only two alternative methods from the literature were presented here—the GRMCK [33] and the GST [10]. Though sufficient for our analysis, it would be useful to explore the other options synthesised from the literature (*e.g.* [3, 23, 26]). Some of these, though more complex, boast many attractive properties, and may prove to be even more fruitful.

REFERENCES

- [1] Stefano V. Albrecht and Subramanian Ramamoorthy. 2013. A Game-Theoretic Model and Best-Response Learning Method for Ad Hoc Coordination in Multi-agent Systems. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems* (St. Paul, MN, USA) (AAMAS '13). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1155–1156.
- [2] Stefano V. Albrecht and Subramanian Ramamoorthy. 2016. Exploiting causality for selective belief filtering in dynamic Bayesian networks. *Journal of Artificial Intelligence Research* 55 (2016), 1135–1178.
- [3] Evgeny Andriyash, Arash Vahdat, and Bill Macready. 2019. Improved Gradient-Based Optimization Over Discrete Distributions. *arXiv:1810.00116 [cs, stat]* (June 2019). arXiv:1810.00116 [cs, stat]
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. 2013. Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation. *CoRR abs/1308.3432* (2013). arXiv:1308.3432 <http://arxiv.org/abs/1308.3432>
- [5] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique Pondé de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. 2019. Dota 2 with Large Scale Deep Reinforcement Learning. *CoRR abs/1912.06680* (2019). <http://arxiv.org/abs/1912.06680>
- [6] David Blackwell. 1947. Conditional expectation and unbiased sequential estimation. *The Annals of Mathematical Statistics* (1947), 105–110.
- [7] Filippos Christianos, Georgios Papoudakis, Muhammad A Rahman, and Stefano V. Albrecht. 2021. Scaling multi-agent reinforcement learning with selective parameter sharing. In *International Conference on Machine Learning*. PMLR, 1989–1998.
- [8] Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. 2020. Shared Experience Actor-Critic for Multi-Agent Reinforcement Learning. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 10707–10717.
- [9] Alek Dimitriev and Mingyuan Zhou. 2021. CARMS: Categorical-Antithetic-REINFORCE Multi-Sample Gradient Estimator. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 13217–13229.
- [10] Ting-Han Fan, Ta-Chung Chi, Alexander I. Rudnicky, and Peter J Ramadge. 2022. Training Discrete Deep Generative Models via Gapped Straight-Through Estimator. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 6059–6073. <https://proceedings.mlr.press/v162/fan22a.html>
- [11] Roman Garnett. 2022. *Bayesian Optimization*. Cambridge University Press. in preparation.
- [12] Will Grathwohl, Dami Choi, Yuhuai Wu, Geoffrey Roeder, and David Duvenaud. 2018. Backpropagation through the Void: Optimizing Control Variates for Black-Box Gradient Estimation. *ICLR 2018* (2018). arXiv:1711.00123
- [13] Evan Greensmith, Peter Bartlett, and Jonathan Baxter. 2001. Variance Reduction Techniques for Gradient Estimates in Reinforcement Learning. In *Advances in Neural Information Processing Systems*, T. Dietterich, S. Becker, and Z. Ghahramani (Eds.), Vol. 14. MIT Press. <https://proceedings.neurips.cc/paper/2001/file/584b98aac2d59ee2cf19ca4ccb75e-Paper.pdf>
- [14] Shixiang Gu, Sergey Levine, Ilya Sutskever, and Andriy Mnih. 2016. MuProP: Unbiased Backpropagation for Stochastic Neural Networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1511.05176>
- [15] Caglar Gulcehre, Sarath Chandar, and Yoshua Bengio. 2017. Memory augmented neural networks with wormhole connections. *arXiv preprint arXiv:1701.08718* (2017).
- [16] Eric A. Hansen, Daniel S. Bernstein, and Shlomo Zilberstein. 2004. Dynamic Programming for Partially Observable Stochastic Games. In *Proceedings of the 19th National Conference on Artificial Intelligence* (San Jose, California) (AAAI'04). AAAI Press, 709–715.
- [17] Iris A.M. Huijben, Wouter Kool, Max Benedikt Paulus, and Ruud JG Van Sloun. 2022. A Review of the Gumbel-max Trick and Its Extensions for Discrete Stochasticity in Machine Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022), 1–1. <https://doi.org/10.1109/TPAMI.2022.3157042>
- [18] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparameterization with Gumbel-Softmax. *ICLR 2017* (2017).
- [19] Diederik P Kingma, Max Welling, et al. 2019. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning* 12, 4 (2019), 307–392.
- [20] Wouter Kool, Herke van Hoof, and Max Welling. 2020. Estimating Gradients for Discrete Random Variables by Sampling without Replacement. *ICLR 2020* (2020). arXiv:2002.06043
- [21] Matt J Kusner and José Miguel Hernández-Lobato. 2016. Gans for sequences of discrete elements with the gumbel-softmax distribution. *arXiv preprint arXiv:1611.04051* (2016).
- [22] Matt J. Kusner and José Miguel Hernández-Lobato. 2016. GANS for Sequences of Discrete Elements with the Gumbel-softmax Distribution. <https://doi.org/10.48550/ARXIV.1611.04051>
- [23] Wonyeol Lee, Hangyeol Yu, and Hongseok Yang. 2018. Reparameterization Gradient for Non-differentiable Models. In *Advances in Neural Information Processing Systems*, Vol. 31. Curran Associates, Inc.
- [24] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous Control with Deep Reinforcement Learning. In *ICLR 2016*.
- [25] Guy Lorberbom, Andreea Gane, Tommi Jaakkola, and Tamir Hazan. 2019. Direct Optimization through Argmax for Discrete Variational Auto-Encoder. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc.
- [26] Guy Lorberbom, Chris J Maddison, Nicolas Heess, Tamir Hazan, and Daniel Tarlow. 2020. Direct Policy Gradients: Direct Optimization of Policies in Discrete Action Spaces. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 18076–18086.
- [27] Ryan Lowe, Wu Yi, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *Advances in Neural Information Processing Systems*, Vol. 30.
- [28] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. *ICLR 2017* (2017). arXiv:1611.00712
- [29] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing Atari with Deep Reinforcement Learning. *NIPS Deep Learning Workshop* (Dec. 2013). arXiv:1312.5602 [cs]
- [30] Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. 2020. Monte Carlo Gradient Estimation in Machine Learning. *Journal of Machine Learning Research* 21, 132 (2020), 1–62.
- [31] Kevin P. Murphy. 2023. *Probabilistic Machine Learning: Advanced Topics*. MIT Press. probml.ai
- [32] Georgios Papoudakis, Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. 2021. Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1* (Dec. 2021).
- [33] Max B. Paulus, Chris J. Maddison, and Andreas Krause. 2021. Rao-Blackwellizing the Straight-Through Gumbel-Softmax Gradient Estimator. *ICLR 2021* (2021). arXiv:2010.04838
- [34] Andres Potapczynski, Gabriel Loaiza-Ganem, and John P Cunningham. 2020. Invertible Gaussian Reparameterization: Revisiting the Gumbel-Softmax. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 12311–12321.
- [35] Jason Tyler Rolfe. 2017. Discrete Variational Autoencoders. *ICLR 2017* (2017). arXiv:1609.02200
- [36] L. S. Shapley. 1953. Stochastic Games. *Proceedings of the National Academy of Sciences* 39, 10 (Oct. 1953), 1095–1100. <https://doi.org/10.1073/pnas.39.10.1095>
- [37] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* 529, 7587 (Jan. 2016), 484–489. <https://doi.org/10.1038/nature16961>
- [38] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. 2014. Deterministic Policy Gradient Algorithms. In *Proceedings of the 31st International Conference on Machine Learning*. PMLR, 387–395.
- [39] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction* (second edition ed.). The MIT Press, Cambridge, Massachusetts.
- [40] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. 1999. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems*, Vol. 12. MIT Press.
- [41] Ming Tan. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the Tenth International Conference on Machine Learning*. 330–337.
- [42] George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. 2017. REBAR: Low-variance, Unbiased Gradient Estimates for Discrete Latent Variable Models. In *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc.
- [43] Andreas Veit and Serge J. Belongie. 2019. Convolutional Networks with Adaptive Inference Graphs. *International Journal of Computer Vision* 128 (2019), 730–741.
- [44] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sutsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney,

- Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. 2019. Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning. *Nature* 575, 7782 (Nov. 2019), 350–354. <https://doi.org/10.1038/s41586-019-1724-z>
- [45] Lex Weaver and Nigel Tao. 2001. The Optimal Reward Baseline for Gradient-Based Reinforcement Learning. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence* (Seattle, Washington) (*UAI'01*). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 538–545.
- [46] Ronald J. Williams. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine Learning* 8, 3 (May 1992), 229–256. <https://doi.org/10.1007/BF00992696>
- [47] Peter R. Wurman, Samuel Barrett, Kenta Kawamoto, James MacGlashan, Kaushik Subramanian, Thomas J. Walsh, Roberto Capobianco, Alisa Devlic, Franziska Eckert, Florian Fuchs, Leilani Gilpin, Piyush Khandelwal, Varun Kompella, HaoChih Lin, Patrick MacAlpine, Declan Oller, Takuma Seno, Craig Sherstan, Michael D. Thomure, Houmeh Aghabozorgi, Leon Barrett, Rory Douglas, Dion Whitehead, Peter Dürr, Peter Stone, Michael Spranger, and Hiroaki Kitano. 2022. Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning. *Nature* 602, 7896 (Feb. 2022), 223–228. <https://doi.org/10.1038/s41586-021-04357-7>
- [48] Shiyang Yan, Jeremy S Smith, Wenjin Lu, and Bailing Zhang. 2018. Hierarchical multi-scale attention networks for action recognition. *Signal Processing: Image Communication* 61 (2018), 73–84.
- [49] Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ricardo Henao, Dinghan Shen, and Lawrence Carin. 2017. Adversarial Feature Matching for Text Generation. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 4006–4015. <https://proceedings.mlr.press/v70/zhang17b.html>