

Fully distributed actor-critic architecture for multitask deep reinforcement learning

SERGIO VALCARCEL MACUA , IAN DAVIES, ALEKSI TUKIAINEN and ENRIQUE MUNOZ DE COTE

Secondmind, Cambridge, CB2 1LA, UK

e-mails: sergiovalmac@gmail.com, davies.ian.r@gmail.com, aleksi.tukiainen@gmail.com, enrique@people-ai.com

Abstract

We propose a fully distributed actor-critic architecture, named diffusion-distributed-actor-critic *Diff-DAC*, with application to multitask reinforcement learning (MRL). During the learning process, agents communicate their value and policy parameters to their neighbours, diffusing the information across a network of agents with no need for a central station. Each agent can only access data from its local task, but aims to learn a common policy that performs well for the whole set of tasks. The architecture is scalable, since the computational and communication cost per agent depends on the number of neighbours rather than the overall number of agents. We derive Diff-DAC from duality theory and provide novel insights into the actor-critic framework, showing that it is actually an instance of the dual-ascent method. We prove almost sure convergence of Diff-DAC to a common policy under general assumptions that hold even for deep neural network approximations. For more restrictive assumptions, we also prove that this common policy is a stationary point of an approximation of the original problem. Numerical results on multitask extensions of common continuous control benchmarks demonstrate that Diff-DAC stabilises learning and has a regularising effect that induces higher performance and better generalisation properties than previous architectures.

1 Introduction

Within a decade, billions of interconnected devices will be sensing, processing, and exchanging data throughout the global economy (van der Meulen, 2015). Centralised reinforcement learning (RL) architectures, where multiple devices collect data that is sent to a central agent, may be infeasible at such a scale due to communication costs, excessive delays in communication, lack of resilience to link failures (e.g., consider a sensor network where the gateway that connects to the Internet fails), or even privacy concerns when sensitive data have to hop through multiple devices before reaching the central point of processing.

Diffusion-based fully distributed algorithms, offer a solution to these problems. Under such systems, agents learn from their local data and in the absence of a central coordinator. Agents can communicate only with their neighbours. Therefore, the communication cost per agent scales linearly with its number of neighbours; all agents learn in parallel and are able to adapt to any number of neighbours, so the process is naturally more robust against device or link failure; and as the agents do not exchange data samples, privacy is naturally preserved.

In addition, *diffusion* schemes enjoy theoretical convergence guarantees; and have shown the ability to outperform standard centralised schemes in the context of non-convex optimisation (Valcarcel Macua, 2017, Chapter 4).

* Work is done while affiliated with Secondmind.

The main motivation of this work is to show that the *diffusion* mechanism can be used to develop fully distributed *actor-critic* methods that enjoy the same favourable properties observed in distributed optimisation problems. The proposed architecture, named *diffusion-distributed-actor-critic* (Diff-DAC), is very flexible and in principle can be applied to any gradient-based single-agent method, as explained in Section 3.

1.1 Related work

The problem of learning policies that generalise across tasks is, in general, known as the multitask reinforcement learning (MRL) problem (Taylor & Stone, 2009). There are multiple approaches to MRL. In the context of model-free deep RL, one approach is to learn a single policy network that is able to control all tasks by learning shared low-level features (Parisotto *et al.*, 2016). An alternative approach is to learn policies that are related but distinct for each task. This can be achieved architecturally, by learning modular policy networks (Andreas *et al.*, 2017); by learning latent features that correlate the policies (Bou-Ammar *et al.*, 2014; El Bsat *et al.*, 2017); or by ensuring that the individual policies remain close to one another, for example, their KL divergence or the Euclidean distance of their parameters are small (Teh *et al.*, 2017). Model-based MRL approaches include learning a common model of the dynamics and then using this prior knowledge to solve new tasks (Fu *et al.*, 2016).

Most previous MRL approaches are centralised, in the sense that they assume an agent with access to data from all observed tasks (Bou-Ammar *et al.*, 2014; Parisotto *et al.*, 2016; Fu *et al.*, 2016; Teh *et al.*, 2017; Andreas *et al.*, 2017). But if the number of tasks is large or their data are *geographically distributed*, the communication cost of transmitting the data to a central station, together with the sensitivity to link or node failure and other issues, may be prohibitive, as discussed above.

Previous works have scaled single-task RL by distributing the sampling of training data. The Asynchronous Advantage Actor-Critic (A3C) algorithm proposed by Mnih *et al.* (2016) enables RL across multiple machines through the decentralised collection of data samples. Under A3C, decentralised agents pass locally calculated gradients to a central node that performs the global gradient update. Then, the decentralised agents periodically update their local policies by copying the parameters of the global policy maintained at the central node. Espeholt *et al.* (2018) take an alternative approach by passing trajectories of experience in the place of calculated gradients. Both approaches have been shown to make training more stable and enable RL agents to train faster and to higher final performance than traditional non-decentralised sampling schemes.

Our approach goes further and decentralises learning as well as sampling so that parameter sharing between agents is not centrally mediated. Furthermore, we develop our approach for the MRL problem and investigate the characteristics of policies learned in a decentralised manner.

The idea of scaling MRL with distributed optimisation was first proposed by El Bsat *et al.* (2017) with the Distributed Multitask Policy Search (Dist-MTPS) method. Dist-MTPS assumes a factored policy model composed of two terms: a linear transformation that represents latent knowledge shared across all tasks and a task-specific coefficient vector. The dot product of these terms forms the parameter vector. This parameter vector linearly combines task-dependent features representing the state to give the mean of a task-specific Gaussian distribution from which actions are drawn. The MRL problem is formulated as a non-convex optimisation problem over the linear transformation and the task-specific coefficients. A local optimum is then obtained with an iterative bi-level optimisation approach: at each iteration, the task-specific parameters are fixed and the agents find a consensus on the shared linear transformation, then the linear transformation is fixed and each agent finds its own task-specific parameter from local data. In order to find the shared linear transformation in a scalable manner, Dist-MTPS extends a distributed implementation of the alternating direction method of multipliers (ADMM) due to Wei and Ozdaglar (2012).

Our work improves over Dist-MTPS in a number of ways: (i) Dist-MTPS relies on linear function approximation, which requires finding salient features, and it only considers policies in the natural exponential family of distributions. Diff-DAC, on the other hand, allows nonlinear parametrisations, like deep learning architectures, which avoid costly feature engineering, and are able to learn more expressive

policies. (ii) Distributed ADMM updates the agents in sequential order which requires finding a cyclic path that visits all agents. Finding such a path is generally an NP-hard problem (Karp, 1972). Diff-DAC, however, uses a diffusion strategy (Sayed, 2014), where each agent interacts with its neighbours, with no ordering, and possibly asynchronously (Zhao & Sayed, 2015). (iii) Sequential updates, as in Dist-MTPS, are sensitive to agent or link failures, since they stop the information flow. Diffusion strategies, however, are robust to such failures as the agents can still operate even if parts of the network become isolated.

To the best of our knowledge, all other previous works on fully distributed RL consider only tabular or linear function approximations (Kar *et al.*, 2013; Valcarcel Macua *et al.*, 2015; Tutunov *et al.*, 2016), and do not apply immediately to expressive nonlinear approximations. For example, Kar *et al.* (2013) added a consensus rule to tabular Q-learning. Its nonlinear extension, therefore, raises questions such as whether we should add a diffusion step to the target network updates. A principled response to this question would be an alternative contribution to our actor-critic approach. The Dist-GTD method presented by Valcarcel Macua *et al.* (2015) is for policy evaluation with linear approximation, and its extension to control and nonlinear approximations is far from trivial even for the single-agent GTD. Finally, Tutunov *et al.* (2016) proposed a second-order method, implying the inversion of the Hessian at every agent, which might be problematic for neural networks with hundreds of neurons. Other earlier works suffer from similar drawbacks.

More recently¹, independent variants of the Diff-DAC architecture have been proposed for complementary scenarios. Zhang *et al.* (2018) studied a similar algorithm adapted to the case where the agents interact with each other in the same environment. Assran *et al.* (2019) studied the computational efficiency, in terms of GPU utilisation, of a similar algorithm when all agents aim to solve the same task.

1.2 Contributions

1. We propose a fully distributed architecture, named Diff-DAC, that allows us to transform single-agent actor-critic algorithms into fully distributed implementations, which can be applied to the MRL problem, and scales gracefully to large number of tasks even with geographically distributed data.
2. We derive Diff-DAC from duality theory and provide novel insights into the standard actor-critic framework, showing that it is actually an instance of the dual-ascent method.
3. We study the asymptotic convergence of the Diff-DAC architecture and prove that all agents converge to a common policy even for nonlinear policy and value function representations. Under more restrictive assumptions, we also show that this common policy approximates a stationary point of the multitask RL objective.
4. We apply the Diff-DAC architecture to two algorithms: a simple Actor-Critic (SiAC) and the Advantage Actor-Critic (A2C)—a synchronous version of A3C from Mnih *et al.* (2016)—to study the stability and regularisation capabilities of the diffusion mechanism.
5. We perform multiple experiments over robotic control tasks, illustrating that even the Diff-DAC implementation of SiAC (Diff-SiAC) outperforms Dist-MTPS for fully distributed multitask RL; that the Diff-DAC architecture is more stable and usually achieves better local optima than the centralised approach; and that Diff-DAC exhibits interesting generalisation properties.

2 Problem formulation

In this section, we formalise tasks as Markov decision processes (MDPs), define a family of tasks, and introduce the multitask optimisation problem.

Consider a parametric family of MDPs defined over state-action sets, \mathbb{S} and \mathbb{A} . Each MDP of the family has a different state transition distribution, $\mathcal{P}_\theta : \mathbb{S} \times \mathbb{S} \times \mathbb{A} \mapsto \mathbb{R}^+$, that depends on some parameter $\theta \in \Theta$, where Θ is a measurable set, and \mathbb{R}^+ denotes the non-negative real values. We assume the reward function, $r : \mathbb{S} \times \mathbb{A} \times \mathbb{S} \mapsto \mathbb{R}$, and the distribution over the initial state, $\mu : \mathbb{S} \mapsto \mathbb{R}^+$, are the same for all tasks. A family of tasks is defined by a probability distribution over the parameter set, f , so that the

¹ After the appearance of a preliminary version of this draft: Valcarcel Macua *et al.* (2017).

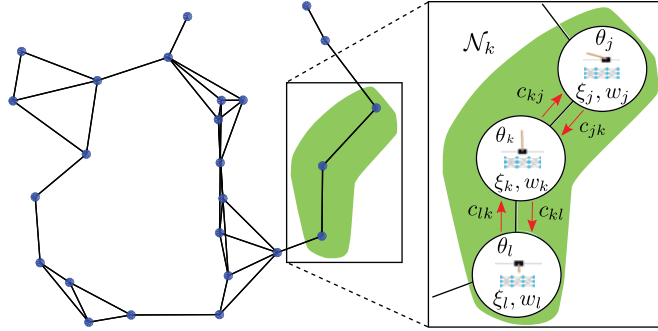


Figure 1 Example of network and detailed neighbourhood. Blue nodes represent agents, and edges represent their connectivity. On the right, the figure zooms in on neighbourhood \mathcal{N}_k (the green area), where each agent k runs its own instance of the environment (here illustrated as swing-up cart-poles, each with different pole length and mass). As explained in Section 3, agent k transmits its critic and actor parameters, $\xi_{k,i}$ and $w_{k,i}$, to its neighbours j and l ; and it receives their parameters $\xi_{j,i}, w_{j,i}$ and $\xi_{l,i}, w_{l,i}$, and combines them with weights c_{jk} and c_{lk} , respectively.

parameter is a random variable²: $\theta = \theta \sim f$. Let $\pi : \mathbb{S} \times \mathbb{A} \mapsto \mathbb{R}^+$ be a stationary policy, such that $\pi(a|s)$ denotes the conditional probability of taking action a at state s .

For a given θ , the single-task infinite-horizon discounted return objective is defined as

$$\begin{aligned} J_\theta(\pi) &\triangleq \mathbb{E}_{\mathcal{P}_\theta, \mu, \pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, \mathbf{a}_t, \mathbf{s}_{t+1}) \right] \\ &= \int_{\mathbb{S}} \rho_\gamma^\pi(s) \int_{\mathbb{A}} \pi(a|s) \int_{\mathbb{S}} \mathcal{P}_\theta(s'|s, a) r(s, a, s') ds' da ds, \end{aligned} \quad (1)$$

where $0 < \gamma < 1$ is the discount factor, and ρ_γ^π is the discounted state measure³ under policy π .

We consider the multitask RL problem of finding a single policy that performs well on average for the whole distribution of tasks. More formally, our goal is to solve the following problem:

$$\underset{\pi}{\text{maximise}} \quad J(\pi) \triangleq \int_{\Theta} f(\theta) J_\theta(\pi) d\theta. \quad (2)$$

In practice, the agents will have to learn by only observing a subset of N tasks that correspond to some parameters $\{\theta_k\}_{k=1}^N$. Moreover, for large (or continuous) state-action sets, it is convenient to approximate the policy with a parametric function: $\pi_w(a|s) \approx \pi(a|s), \forall (s, a) \in \mathbb{S} \times \mathbb{A}$, where $w \in \mathbb{W} \subseteq \mathbb{R}^{M_\pi}$ is the parameter vector, the length of which is denoted by M_π . Therefore, the actual problem the agents have to solve is to maximise the parametric empirical risk:

$$\underset{w}{\text{maximise}} \quad \bar{J}(w) \triangleq \frac{1}{N} \sum_{k=1}^N J_{\theta_k}(\pi_w). \quad (3)$$

Existence of a solution to (3) is guaranteed under standard assumptions on the policy set and state transition kernel, such that the induced state Markov chain has positive steady-state probabilities (Bertsekas, 2012), as well as on the reward function such that the objective satisfies the Weirstrass conditions (Bertsekas, 2009, Proposition 3.2.1).

When all task parameters $\{\theta_k\}_{k=1}^N$ are equal, (3) is the single-task RL problem; when they differ, (3) becomes an MRL problem, where we aim to learn a single policy that performs well for the whole set of tasks. Although a single policy might perform well for some tasks but poorly for others, experiments show that our solution to (3) can outperform Dist-MTPS, even when the latter learns related but distinct task-specialised policies.

² We use boldface font to denote random variables and regular font to denote instances or deterministic variables.

³ See (19) below.

3 Networked multi-agent setting and diffusion mechanism

In this section, we introduce the networked multi-agent setting wherein policies are learned in a fully distributed manner.

We have a network of agents, each one learning from data coming from its own task⁴. Let $\mathcal{N} \triangleq \{1, \dots, N\}$ denote the set of agents.

The network is expressed as a graph, Γ , where nodes represent agents and edges represent communication links. The graph can be represented by a non-negative matrix of size $N \times N$, denoted $C \triangleq (c_{lk})_{l,k=1}^N$, such that the element $c_{lk} \geq 0$ represents the weight given by agent k to information coming from l . Each agent $k \in \mathcal{N}$ has data coming from its own task only, with parameter $\theta_k \sim f$; and it is only allowed to communicate within its own neighbourhood, \mathcal{N}_k , which is defined as the set of agents to which it is directly connected, including k itself:

$$\mathcal{N}_k \triangleq \{l \in \{1, \dots, N\} : c_{lk} > 0\}. \quad (4)$$

We rely on the *diffusion* mechanism for fully distributed optimisation (Chen & Sayed, 2013; Sayed, 2014), which typically consists of two steps: local *adaptation* and in-neighbourhood *combination*. During the adaptation step, each agent updates its parameters in the direction of its stochastic gradient calculated from local data. In the combination step, each agent averages its local approximation with those coming from its neighbours. For problem (3), these two steps are described by the following updates, which run in parallel for all agents $k = 1, \dots, N$:

$$\widehat{w}_{k,i+1} = w_{k,i} + \beta_{i+1} \widehat{\nabla}_w J_{\theta_k}(\pi_{w_{k,i}}) \quad (\text{adaptation}) \quad (5a)$$

$$w_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \widehat{w}_{l,i+1}, \quad (\text{combination}) \quad (5b)$$

where i is the iteration index; β_i is the step size; $w_{k,i}$ is the approximate solution to the global problem (3) available at agent k at iteration i ; $\widehat{\nabla}_w J_{\theta_k}(\pi_w)$ is the stochastic gradient of its local objective function evaluated at its current parameter $w_{k,i}$; and $\widehat{w}_{k,i+1}$ is an intermediate parameter resulting from the local adaptation step.

In order to ensure that the information flows through the network, we assume that the graph, Γ , is *strongly connected* (i.e., there is at least one path between every pair of agents), and require the following standard conditions on the connectivity matrix C , which together with strong connectivity make it doubly stochastic and primitive (Sayed, 2014; Valcarcel Macua *et al.*, 2015):

$$C^\top \mathbb{1}_N = \mathbb{1}_N, \quad C \mathbb{1}_N = \mathbb{1}_N, \quad \text{and} \quad c_{lk} \geq 0, \quad \forall l, k \in \{1, \dots, N\}, \quad (6)$$

$$\text{trace}[C] > 0, \quad (7)$$

where $\mathbb{1}_N$ is a vector of ones of length N . There are procedures for every agent k to find the weights $\{c_{lk}\}_{l \in \mathcal{N}_k}$ in a fully distributed manner, such that the resulting C satisfies (6)–(7). One such procedure is the Hastings rule Zhao & Sayed (2012); Sayed (2014), 492.

4 Architecture for fully distributed actor-critic algorithms

Equations (5a)–(5b) constitute the basis of a diffusion-based distributed policy gradient algorithm where we only optimise over the policy space, typically using an unbiased but high-variance Monte Carlo estimate of the policy gradient, like REINFORCE (Williams, 1992).

Actor-critic methods were introduced to reduce the variance of the policy gradient estimate (see, e.g., Grondman *et al.*, 2012 and references therein). In particular, it has been proven that an optimal control

⁴ For ease of exposition, we assume that each agent is allocated with one task, similar to El Bsati *et al.* (2017). The extension to multiple tasks per agent is trivial.

variate is given by the value function, which is defined for some given policy π and some given task parameter θ as follows:

$$v_\theta^\pi(s) \triangleq \mathbb{E}_{\mathcal{P}_\theta, \pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, \mathbf{a}_t, s_{t+1}) \mid s_0 = s \right]. \quad (8)$$

In this context, the policy (used by the agent to operate in the environment) is known as the *actor*, and the value function (which evaluates the goodness of a policy) is known as the *critic*.

In this section, we propose a diffusion-based fully distributed actor-critic architecture, named Diff-DAC, where the agents cooperate to both estimate the multitask critic and optimise the actor. We derive Diff-DAC from first principles, as a primal-dual scheme to find the saddle point of an approximate multitask Lagrangian. Our primal-dual derivation makes clear that the agents can benefit from applying diffusion to solve both the primal and dual problems in a cooperative manner. This is not obvious when the critic is motivated only from a variance reduction point of view. Let's proceed to explain the details.

4.1 Primal-dual derivation of actor-critic methods

Throughout this subsection, we assume finite state-action sets for the sake of simplicity. Hence, the single-task objective (1) can be expressed as a dot product:

$$J_\theta(\pi) = \mu^\top v_\theta^\pi, \quad (9)$$

where $v_\theta^\pi \triangleq (v_\theta^\pi(s))_{s \in \mathbb{S}} \in \mathbb{R}^{|\mathbb{S}|}$, is a vector with the value function for every state.

We can find the maximum of (9) as the solution to the following linear program (LP) (Puterman, 2005, Section 9.1):

$$\begin{aligned} & \underset{v}{\text{minimise}} \quad \mu^\top v \\ & \text{s.t.} \quad v(s) \geq \sum_{s' \in \mathbb{S}} \mathcal{P}_\theta(s'|s, a) (r(s, a, s') + \gamma v(s')), \quad \forall (s, a) \in \mathbb{S} \times \mathbb{A}. \end{aligned} \quad (10)$$

The Lagrangian of (10) is given by

$$L_\theta(v, d) = \mu^\top v + \sum_{(s, a) \in \mathbb{S} \times \mathbb{A}} d(s, a) \left(\sum_{s' \in \mathbb{S}} \mathcal{P}_\theta(s'|s, a) (r(s, a, s') + \gamma v(s')) - v(s) \right), \quad (11)$$

where the dual variable, $d \triangleq (d(s, a))_{(s, a) \in \mathbb{S} \times \mathbb{A}} \geq 0$, is a non-negative vector of length $|\mathbb{S}| |\mathbb{A}|$. Let us introduce the multitask Lagrangian that integrates over the distribution of tasks:

$$\begin{aligned} L(v, d) & \triangleq \int_{\Theta} L_\theta(v, d) f(\theta) d\theta \\ & = \mu^\top v + \sum_{(s, a) \in \mathbb{S} \times \mathbb{A}} d(s, a) \left(\sum_{s' \in \mathbb{S}} \mathcal{P}(s'|s, a) (r(s, a, s') + \gamma v(s')) - v(s) \right), \end{aligned} \quad (12)$$

where we have introduced a shortcut for the expected state transition distribution:

$$\mathcal{P}(s'|s, a) \triangleq \int_{\Theta} \mathcal{P}_\theta(s'|s, a) f(\theta) d\theta. \quad (13)$$

Note that (12) can be thought of as the Lagrangian of another LP, similar to (10) but where \mathcal{P}_θ has been replaced with \mathcal{P} :

$$\begin{aligned} & \underset{v}{\text{minimise}} \quad \mu^\top v \\ & \text{s.t.} \quad v(s) \geq \sum_{s' \in \mathbb{S}} \mathcal{P}(s'|s, a) (r(s, a, s') + \gamma v(s')), \quad \forall (s, a) \in \mathbb{S} \times \mathbb{A}. \end{aligned} \quad (14)$$

Indeed, problem (14) corresponds to a single MDP with transitions given by \mathcal{P} , which represents a valid distribution, as shown by the following proposition.

PROPOSITION 1. \mathcal{P} is a row-stochastic matrix.

Proof. Stochastic matrices lie in a compact convex set (Horn & Johnson, 1990, Theorem 8.7). Hence, their convex combination lies in the same set (Boyd & Vandenberghe, 2004, 24). \square

Since (14) satisfies Slater's condition, strong duality holds (Boyd & Vandenberghe, 2004, Section 5.2.3), and optimal primal and dual variables are attained and form a saddle point of (12):

$$\min_v \max_d L(v, d) = L(v^*, d^*) = \max_d \min_v L(v, d), \quad (15)$$

where v^* and d^* denote optimal primal and dual variables⁵ of (14), respectively.

There are multiple approaches to find a saddle point that satisfies optimality condition (15). We focus on the *dual-ascent* scheme (Arrow *et al.*, 1958), which consists in alternating between:

1. finding a primal solution, given the dual variable and
2. ascending in the direction of the gradient of the Lagrangian w.r.t. the dual variable, given the primal variable.

First, we show how to update the *primal* variable given d , that is

$$v \leftarrow \arg \min_v L(v, d). \quad (16)$$

The Karush–Kuhn–Tucker (KKT) conditions are sufficient for optimality in convex problems that satisfy Slater's condition and have differentiable objective and constraints (Boyd & Vandenberghe, 2004, Section 5.5.3). These KKT conditions include the feasibility constraints, non-negativity of the dual variable for inequality constraints, *complementary slackness*, and null gradient of the Lagrangian w.r.t. the primal variable. Since problem (14) is linear, first-order conditions do not depend on v . Thus, the only condition that depends on v is *complementary slackness*:

$$d(s, a) \left(\sum_{s' \in \mathbb{S}} \mathcal{P}(s'|s, a) (r(s, a, s') + \gamma v(s')) - v(s) \right) = 0, \quad \forall (s, a) \in \mathbb{S} \times \mathbb{A}. \quad (17)$$

Similar to the standard single-task problem (recall Proposition 1), it can be shown (Puterman, 2005, Section 6.9) that our dual variable is the discounted state-action visitation measure:

$$d(s, a) = \sum_{j \in \mathbb{S}} \mu(j) \sum_{t=0}^{\infty} \gamma^t p(s_t = s, a_t = a | s_0 = j), \quad (18)$$

such that:

$$\rho_\gamma^\pi(s) = \sum_{a \in \mathbb{A}} d(s, a). \quad (19)$$

Therefore, finding d allows us to obtain the corresponding policy:

$$\pi(a|s) = \frac{d(a, s)}{\sum_{a \in \mathbb{A}} d(s, a)}. \quad (20)$$

⁵ Note v^* is unique, while there might be multiple optimal dual variables.

Since $\pi(a|s) \geq 0$ and $\sum_{s' \in \mathbb{S}} \mathcal{P}(s'|s, a) (r(s, a, s') + \gamma v(s')) - v(s) \leq 0$ for all $(s, a) \in \mathbb{S} \times \mathbb{A}$, we conclude that the Bellman equation, typically used to derive the *critic* in actor-critic methods, and given by

$$v(s) = \sum_{a \in \mathbb{A}} \pi(a|s) \left(\sum_{s' \in \mathbb{S}} \mathcal{P}(s'|s, a) (r(s, a, s') + \gamma v(s')) \right), \quad \forall s \in \mathbb{S}, \quad (21)$$

is sufficient for (17), and thus to find the primal variable that optimises the Lagrangian.

Second, for the *dual* variable, we simply perform gradient ascent in the Lagrangian, yielding a recursion of the form:

$$d \leftarrow [d + \beta \nabla_d L(v, d)]^+, \quad (22)$$

where β is the step size, $[\cdot]^+$ denotes projection on the non-negative orthant, and ∇_d denotes the gradient w.r.t. the dual variable d :

$$\nabla_d L(v, d) = \left(\frac{\partial L(v, d)}{\partial d(s, a)} \right)_{(s, a) \in \mathbb{S} \times \mathbb{A}}. \quad (23)$$

Interestingly, note that the partial derivatives of the Lagrangian in (23) are equal to the so-named *advantage function*, $A : \mathbb{S} \times \mathbb{A} \rightarrow \mathbb{R}$, which was originally motivated for learning to control continuous time systems approximated with small discrete time steps (Baird III, 1993) and, more recently, as a general variance reduction technique (Bhatnagar *et al.*, 2009):

$$\frac{\partial L(v, d)}{\partial d(s, a)} = A(s, a) \triangleq \sum_{s' \in \mathbb{S}} \mathcal{P}(s'|s, a) (r(s, a, s') + \gamma v(s')) - v(s). \quad (24)$$

Since by learning d^* , we can obtain π^* from (20), the recursion in (22) can be seen as an *actor* update.

In summary, (21) and (22) define an *actor-critic* method for the tabular setting, with no need for parametric policies. This novel derivation shows that by finding a saddle point of the multitask Lagrangian (12), we obtain a solution to the original problem (2), as formally stated by the following proposition.

PROPOSITION 2. *Let v^* and d^* satisfy the saddle point optimality condition (15). Then, they constitute a solution to the multitask problem (2):*

- i. $\mu^\top v^* = \max_{\pi} \int_{\Theta} f(\theta) J_{\theta}(\pi) d\theta$.
- ii. Let $\pi^*(a|s) \triangleq \frac{d^*(a, s)}{\sum_{a \in \mathbb{A}} d^*(a, s)}$, $\forall (s, a) \in \mathbb{S} \times \mathbb{A}$, then $\pi^* \in \arg \max_{\pi} \int_{\Theta} f(\theta) J_{\theta}(\pi) d\theta$.

Proof. Consider \mathcal{P} as the state transition of some MDP, then Theorem 6.9.4 from Puterman (2005) states the following:

$$\mu^\top v^* = \sum_{s, a \in \mathbb{S} \times \mathbb{A}} d^*(s, a) \sum_{s' \in \mathbb{S}} \mathcal{P}(s'|s, a) r(s, a, s') = \max_{\pi} \mu^\top v^{\pi}, \quad (25)$$

where $v^{\pi} \triangleq \mathbb{E}_{\mathcal{P}, \pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, \mathbf{a}_t, s_{t+1}) \right]$. From (25), applying (17) in (12), and using a relationship similar to (9) but for the multitask transition kernel, \mathcal{P} , we have:

$$\begin{aligned} L(v^*, d^*) &= \max_{\pi} \mu^\top v^{\pi} \\ &= \max_{\pi} \int_{\mathbb{S}} \rho_{\gamma}^{\pi}(s) \int_{\mathbb{A}} \pi(a|s) \int_{\mathbb{S}} \mathcal{P}(s'|s, a) r(s, a, s') ds' da ds \\ &= \max_{\pi} \int_{\Theta} f(\theta) J_{\theta}(\pi) d\theta. \end{aligned} \quad \square$$

When we consider the empirical risk (3), the expected state transition distribution, \mathcal{P} , can be approximated with this unbiased estimate:

$$\bar{\mathcal{P}}(s'|s, a) \triangleq \frac{1}{N} \sum_{k=1}^N \mathcal{P}_{\theta_k}(s'|s, a), \quad \forall (s, a, s') \in (\mathbb{S} \times \mathbb{A} \times \mathbb{S}). \quad (26)$$

From Proposition 1, we conclude that this is also a valid state transition distribution for some MDP. Hence, the same arguments hold for this average distribution, and we can derive a dual-ascent method to obtain the saddle point of an empirical approximation of the multitask Lagrangian (12) that averages (11) for the set of tasks $\{\theta_k\}_{k=1}^N$:

$$\bar{L}(v, d) = \frac{1}{N} \sum_{k=1}^N L_{\theta_k}(v, d). \quad (27)$$

It is trivial to see that the only difference between (12) and (27) is that the former uses the expected distribution \mathcal{P} , while the latter uses its estimate $\bar{\mathcal{P}}$. This yields the following corollary.

COROLLARY 1. *Let v^* and d^* be a saddle point of (27). Then, they constitute a solution to the empirical risk problem:*

- i. $\mu^\top v^* = \max_{\pi} \frac{1}{N} \sum_{k=1}^N J_{\theta_k}(\pi)$.
- ii. Let $\pi^*(a|s) \triangleq \frac{d^*(a,s)}{\sum_{a \in \mathbb{A}} d^*(s,a)}$, $\forall (s, a) \in \mathbb{S} \times \mathbb{A}$, then $\pi^* \in \arg \max_{\pi} \frac{1}{N} \sum_{k=1}^N J_{\theta_k}(\pi)$.

In the following subsections, we propose a diffusion-based fully distributed actor-critic architecture for large (possibly infinite) state-action sets, with parametric value function and policy that approximate the saddle point of (27).

4.2 Distributed critic

In the previous subsection, we have seen that the primal update of the dual-ascent method is equivalent to finding a value function that satisfies the multitask Bellman Equation (21).

When computing this value function for large or continuous state sets, it is common⁶ to rely on some parametric function approximation: $v_{\xi}(s) \approx v(s)$, where $\xi \in \mathbb{R}^{M_v}$ is the parameter vector of length M_v . Thus, for some given policy, π , we can learn value function parameters ξ by transforming (21) into a regression problem:

$$\underset{\xi}{\text{minimise}} \quad \bar{J}(\xi) \triangleq \mathbb{E}_{\rho_{\bar{\mathcal{P}}}} \left[(v_{\xi}(s_t) - \bar{y}_t)^2 \right], \quad (28)$$

where the target values are given by:

$$\bar{y}_t \triangleq \int_{\mathbb{A}} \pi(a|s) \int_{\mathbb{S}} \bar{\mathcal{P}}(s'|s_t, a_t) (r(s_t, a_t, s') + \gamma v_{\xi}(s')) ds' da, \quad (29)$$

This is equivalent to finding ξ , such that the advantage function is as close to zero as possible.

In order to derive a diffusion-based distributed critic, we have to reformulate the problem as minimising the average cost over all tasks. The cost for each individual task takes the form:

$$J_k(\xi) \triangleq \mathbb{E}_{\rho_{\bar{\mathcal{P}}}} \left[(v_{\xi}(s_t) - y_{k,t})^2 \right], \quad k = 1, \dots, N, \quad (30)$$

⁶ See, for example, Ng *et al.* (1999), Konda Tsitsiklis (2003), Melo and Lopes (2008), Bhatnagar *et al.* (2009), Powell and Ma (2011), Van Hasselt (2012), Weinstein and Littman (2012), Wierstra *et al.* (2014), Lillicrap *et al.* (2015), Heess *et al.* (2015), Schulman *et al.* (2015).

where $\mathbf{y}_{k,t}$ is the target for the k th task at time t :

$$\mathbf{y}_{k,t} \triangleq \int_{\mathbb{A}} \pi(a|s) \int_{\mathbb{S}} \mathcal{P}_{\theta_k}(s'|s_t, a_t) (r(s_t, a_t, s') + \gamma v_{\xi}(s')) ds' da, \quad (31)$$

such that: $\bar{\mathbf{y}}_t = 1/N \sum_{k=1}^N \mathbf{y}_{k,t}$.

We can use Jensen's inequality to upper bound $J(\xi)$ by another function, $\tilde{J}(\xi)$, and use this upper bound as surrogate cost:

$$\tilde{J}(\xi) \triangleq \frac{1}{N} \sum_{k=1}^N J_k(\xi) = \frac{1}{N} \sum_{k=1}^N \mathbb{E} \left[(v_{\xi}(s_t) - \mathbf{y}_{k,t})^2 \right] \geq \mathbb{E} \left[\left(\frac{1}{N} \sum_{k=1}^N (v_{\xi}(s_t) - \mathbf{y}_{k,t}) \right)^2 \right] = \bar{J}(\xi). \quad (32)$$

Now, we can apply *diffusion* to minimise $\tilde{J}(\xi)$ in a distributed fashion, with every agent $k = 1, \dots, N$, applying *adaptation* and *combination* steps in parallel:

$$\hat{\xi}_{k,i+1} = \xi_{k,i} - \alpha_{i+1} \hat{\nabla}_{\xi} J_k(\xi_{k,i}) \quad (33a)$$

$$\xi_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \hat{\xi}_{l,i+1}, \quad (33b)$$

where i is the iteration index; α_i is the step size; $\hat{\nabla}_{\xi} J_k(\xi_{k,i})$ is the stochastic gradient evaluated at $\xi_{k,i}$ and estimated from a batch of $T_{k,i}$ local samples, $\{(s_{k,t}, a_{k,t}, r_{k,t+1}, s_{k,t+1})\}_{t=0}^{T_{k,i}}$, drawn from the stationary distribution induced by policy π .

Equations (33a)–(33b) represent a template for a fully distributed critic update. Depending on how the targets are estimated, the sampling process and other factors, we can build different distributed critic methods. In addition, instead of using a simple stochastic gradient update in the adaptation step (33a), we can use more sophisticated methods, like momentum, adaptive learning rate, and other variants. In this paper, we will introduce two specific algorithms that use this template, Diff-SiAC and Diff-A2C. But before discussing their details, we present a template for the fully distributed actor update.

4.3 Distributed actor

Recall from Section 4.1 that the dual step of the dual-ascent method involves ascending in the direction of the gradient of the Lagrangian w.r.t. the dual variable.

For large or continuous state-action sets, it is convenient to approximate the policy with a parametric function. Let $\pi_w \approx \pi$ denote the parametric approximation of the actual policy, where $w \in \mathbb{R}^{M_{\pi}}$ is the parameter vector of length M_{π} .

Replacing π with π_w in (27), we obtain an approximate Lagrangian, $\tilde{L}(v, w) \approx \bar{L}(v, d)$:

$$\tilde{L}(v, w) = \mu^{\top} v + \int_{\mathbb{S}} \int_{\mathbb{A}} d_w(s, a) \bar{A}(s, a) ds da, \quad (34)$$

where $d_w(s, a)$ is the resulting approximation of the dual variable:

$$d_w(s, a) \triangleq \pi_w(a|s) \rho_{\gamma}^{\pi_w}(s), \quad (35)$$

and $\bar{A}(\cdot) \triangleq 1/N \sum_{k=1}^N A_k(\cdot)$ refers to the average advantage function, with:

$$A_k(s, a) \triangleq \int_{\mathbb{S}} \mathcal{P}_{\theta_k}(s'|s, a) (r(s, a, s') + \gamma v(s')) ds' - v(s). \quad (36)$$

Thus, in order to approximate the optimal dual variable, we can move in the ascent direction of the gradient of (34) w.r.t. the policy parameter. The following theorem provides the required gradient.

THEOREM 1. *The gradient of the Lagrangian w.r.t. the policy parameter is given by*

$$\nabla_w \tilde{\mathcal{L}}(v, w) = \int_{\mathbb{S}} \rho_{\gamma}^{\pi_w}(s) \int_{\mathbb{A}} \pi_w(a|s) \nabla_w \log \pi_w(a|s) \bar{A}(s, a) ds da. \quad (37)$$

Proof. See Appendix C.1. \square

Interestingly, (37) is similar to previous *policy gradient* theorems (Sutton *et al.*, 1999), with the important difference that it yields the advantage function explicitly; while previous works motivated the advantage—as opposed to the state-action value function—as a variance reduction technique (Williams, 1992; Bhatnagar *et al.*, 2009).

In order to derive a fully *distributed* actor, we introduce the approximate Lagrangian for each individual task:

$$\tilde{\mathcal{L}}_k(v, w) \triangleq \mu^\top v + \int_{\mathbb{S}} \int_{\mathbb{A}} \rho_{\gamma}^{\pi_w}(s) \pi_w(a|s) A_k(s, a) ds da, \quad (38)$$

such that $\tilde{\mathcal{L}}(\cdot) = 1/N \sum_{k=1}^N \tilde{\mathcal{L}}_k(\cdot)$.

Similar to the critic, we can apply diffusion to perform the actor update, with step size β_{i+1} :

$$\hat{w}_{k,i+1} = w_{k,i} + \beta_{i+1} \hat{\nabla}_w \tilde{\mathcal{L}}_k(v_{\xi_{k,i}}, w_{k,i}), \quad (39a)$$

$$w_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \hat{w}_{l,i+1}, \quad (39b)$$

where $\hat{\nabla}_w \tilde{\mathcal{L}}_k(v_{\xi_{k,i}}, w_{k,i})$ is an estimate of each agent’s local policy gradient—which is similar to (37) but replacing the average advantage function, \bar{A} , with each agent’s estimate of its local advantage function, A_k —evaluated at its local critic estimate, $v_{\xi_{k,i}}$, obtained from (33a)–(33b).

Similar to the critic update, there are many ways of estimating the stochastic gradient in (39a). Depending on the way of estimating the advantage function (see Schulman *et al.*, 2015 for multiple estimators); the sampling process; and how the value parameters, $\xi_{k,i}$, relate to the policy parameters, we can obtain different fully distributed actor updates. These options lead to different actor-critic methods.

In the following subsections, we apply the Diff-DAC architecture to propose two actor-critic algorithms, which will allow us to demonstrate the benefits of Diff-DAC in terms of stabilisation, performance, generalisation and robustness.

4.4 Simple Actor-Critic (SiAC)

The first algorithms we introduce are centralised and Diff-DAC implementations of a simple actor-critic algorithm (SiAC) algorithm, characterised by having a very simple and highly biased estimator of the advantage function that allows us to evaluate the stability provided by the diffusion mechanism. In particular, the parameters are only updated once the agent has completed an episode on each task. In order to be sample efficient, we perform as many updates to both critic and actor as available transitions in the episode. That is, if at the i -th iteration, the k -th agent samples an episode with $T_{k,i}$ transitions, then we obtain a set of $T_{k,i}$ advantage estimates given by

$$\hat{A}_{k,t} \triangleq \sum_{j=t}^{T_{k,i}} \gamma^{j-t} r_{k,j+1} - v_{\xi_{k,i}}(s_{k,t}), \quad 0 \leq t < T_{k,i} - 1. \quad (40)$$

Note that the advantage estimates $\{\hat{A}_{k,t}\}_{t=1}^{T_{k,i}}$ will be highly correlated, and can therefore cause learning to become unstable. This is done by design in order to conveniently study the stabilisation properties of the Diff-DAC architecture.

In order to compare the centralised and Diff-DAC architectures, we build two algorithms *Centralised SiAC* and *Diff-SiAC* where we introduce the following two changes: (i) Centralised SiAC has a single set of value and policy parameters, whereas Diff-SiAC has N sets of value and policy parameters (one per

agent) and (ii) Centralised SiAC updates its set of value and policy parameters with advantage estimates calculated across all tasks whereas Diff-SiAC updates each agent’s parameters with its own local set of samples and then combines parameters with diffusion.

For Centralised SiAC, the value and policy parameters, ξ_i and w_i , are updated with stochastic gradients that average all samples and all advantage estimates from all tasks:

$$\widehat{\nabla}_{\xi} \widetilde{J}(\xi_i, \pi_{w_i}) = \frac{1}{T_i} \sum_{k=1}^N \sum_{t=0}^{T_{k,i}-1} \nabla_{\xi} v_{\xi_i}(s_{k,t}) \widehat{A}_{k,t}, \quad (41)$$

$$\widehat{\nabla}_w \widetilde{L}(v_{\xi_i}, w_i) = \frac{1}{T_i} \sum_{k=1}^N \sum_{t=0}^{T_{k,i}-1} \nabla_w \log \pi_{w_i}(a_{k,t}|s_{k,t}) \widehat{A}_{k,t}, \quad (42)$$

where $T_i \triangleq \sum_{k=1}^N T_{k,i}$, denotes the total number of samples from the i th episode of all tasks.

For Diff-SiAC, each agent updates its parameters, $\xi_{k,i}$ and $w_{k,i}$, with stochastic gradients that average its local samples and advantage estimates:

$$\widehat{\nabla}_{\xi} J_k(\xi_{k,i}, \pi_{w_{k,i}}) = \frac{1}{T_{k,i}} \sum_{t=0}^{T_{k,i}-1} \nabla_{\xi} v_{\xi_{k,i}}(s_{k,t}) \widehat{A}_{k,t}, \quad (43)$$

$$\widehat{\nabla}_w \widetilde{L}_k(v_{\xi_{k,i}}, w_{k,i}) = \frac{1}{T_{k,i}} \sum_{t=0}^{T_{k,i}-1} \nabla_w \log \pi_{w_{k,i}}(a_{k,t}|s_{k,t}) \widehat{A}_{k,t}. \quad (44)$$

While we have considered a standard stochastic gradient for the derivations, in practice we are free to use any gradient-based optimiser to perform parameter updates. In the experiments with Centralised SiAC and Diff-SiAC, we use the Adam optimiser (Kingma & Ba, 2015).

4.5 Diffusion A2C (Diff-A2C)

In this subsection, we introduce a Diff-DAC implementation of A2C that will allow us to evaluate a new regularisation effect that emerges from fully distributed sampling and learning.

A2C is a variant of the A3C algorithm described in Section 1 that uses distributed sampling but synchronous learning updates. In addition, A2C typically uses a more advanced multistep advantage estimator that reduces variance and can lead to faster learning:

$$\widehat{A}_{k,t} \triangleq \sum_{j=t}^{T_{k,i}} \gamma^{j-t} r_{k,j+1} + \gamma^{T_{k,i}-t} v_{\xi_{k,i}}(s_{k,T_{k,i}}) - v_{\xi_{k,i}}(s_{k,t}), \quad 0 \leq t < T_k - 1, \quad (45)$$

where $T_{k,i}$ is a parameter and denotes the number of transitions per parameter update. Note that the multistep estimator (45) allows A2C agents to set $T_{k,i}$ to a fixed small value such that the parameters are updated at a fixed interval in terms of environment steps, instead of at the end of each episode. The distributed sampling approach, the multistep advantage estimates, and the quick updates allows A2C to offer stable and sample efficient learning. In typical A2C approaches, the training of a single set of parameters is centralised.

Since the advantage estimator in (45) relies on local samples and parameters, we are able to apply it to agents in the Diff-DAC architecture, with stochastic gradients similar to (41)–(44). This yields an algorithm we refer to as Diff-A2C wherein both experience collection and training are decentralised.

As with SiAC and Diff-SiAC we are free to choose the gradient-based optimiser. For the experiments with A2C and Diff-A2C we use RMSProp (Tieleman & Hinton, 2012) to demonstrate Diff-DAC working with varied optimisers⁷. See Appendix D for pseudocode and a discussion contrasting Diff-SiAC and Diff-A2C.

⁷ During experimentation we observed similar results across runs with RMSProp and Adam optimisers.

5 Convergence analysis

In this section, we illustrate some asymptotic properties of the distributed actor-critic iterations, (33a)–(33b) and (39a)–(39b). We provide two asymptotic properties of the Diff-DAC architecture. First, we show that all agents converge to an average recursion. Second, we show that this recursion converges to a stationary point of the parametric approximation of the MRL problem.

Let us start the analysis by rewriting the updates as follows:

$$\widehat{\boldsymbol{\phi}}_{k,i+1} = \boldsymbol{\phi}_{k,i} - \alpha_{i+1} (g_k(\boldsymbol{\phi}_{k,i}) + \mathbf{F}_{k,i+1}) \quad (46a)$$

$$\boldsymbol{\phi}_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \widehat{\boldsymbol{\phi}}_{l,i+1}, \quad (46b)$$

where we aggregated the critic and actor parameters into a single vector, and we expressed the stochastic gradient as the sum of the exact expected gradient, g_k , plus a random vector, $\mathbf{F}_{k,i+1}$:

$$C\boldsymbol{\phi}_{k,i} \triangleq \begin{pmatrix} \boldsymbol{\xi}_{k,i} \\ \mathbf{w}_{k,i} \end{pmatrix}, \quad g_k(\boldsymbol{\phi}_{k,i}) \triangleq \begin{pmatrix} \nabla_{\boldsymbol{\xi}} J_k(\boldsymbol{\xi}_{k,i}, \boldsymbol{\pi}_{\mathbf{w}_{k,i}}) \\ \frac{\beta_{i+1}}{\alpha_{i+1}} \nabla_{\mathbf{w}} \widetilde{L}_k(v_{\boldsymbol{\xi}_{k,i}}, \mathbf{w}_{k,i}) \end{pmatrix}, \quad (47)$$

$$\mathbf{F}_{k,i+1} \triangleq \begin{pmatrix} \widehat{\nabla}_{\boldsymbol{\xi}} J_k(\boldsymbol{\xi}_{k,i}, \boldsymbol{\pi}_{\mathbf{w}_{k,i}}) \\ \frac{\beta_{i+1}}{\alpha_{i+1}} \widehat{\nabla}_{\mathbf{w}} \widetilde{L}_k(v_{\boldsymbol{\xi}_{k,i}}, \mathbf{w}_{k,i}) \end{pmatrix} - g_k(\boldsymbol{\phi}_{k,i}). \quad (48)$$

We require the following standard assumptions:

ASSUMPTION 1. *The graph Γ is strongly connected and C satisfies conditions (6)–(7).*

ASSUMPTION 2. *Gradients g_k are Lipschitz for all $k \in \mathcal{N}$:*

$$\|g_k(\boldsymbol{\phi}) - g_k(\boldsymbol{\phi}')\| \leq L_1 \|\boldsymbol{\phi} - \boldsymbol{\phi}'\|, \quad 0 < L_1 < \infty, \quad \forall \boldsymbol{\phi}, \boldsymbol{\phi}' \in \mathbb{R}^{M_v + M_\pi}. \quad (49)$$

ASSUMPTION 3. *Step-size sequences $\{\alpha_i\}_{i=0}^\infty$ and $\{\beta_i\}_{i=0}^\infty$ consist of positive scalars satisfying:*

$$\sum_{i=0}^\infty \alpha_i = \sum_{i=0}^\infty \beta_i, \quad \sum_{i=0}^\infty (\alpha_i^2 + \beta_i^2) < \infty, \quad \lim_{i \rightarrow \infty} \frac{\beta_i}{\alpha_i} = 0. \quad (50)$$

ASSUMPTION 4. *$\{\mathbf{F}_{k,i+1}\}_{i=0}^\infty$ are square-integrable w.r.t. the increasing families of σ -fields $\mathcal{F}_{k,i} \triangleq \sigma(\boldsymbol{\phi}_{k,j}, \mathbf{F}_{k,j}, j \leq i) = \sigma(\boldsymbol{\phi}_{k,0}, \mathbf{F}_{k,1}, \dots, \mathbf{F}_{k,i})$.*⁸

$$\mathbb{E}[\|\mathbf{F}_{k,i+1}\|^2 | \mathcal{F}_{k,i}] \leq L_2 (1 + \|\boldsymbol{\phi}_{k,i}\|^2) \quad \text{a.s.} \quad i \geq 0, \quad 0 < L_2 < \infty, \quad \forall k \in \mathcal{N}. \quad (51)$$

ASSUMPTION 5. *The iterates of (46a)–(46b) remain bounded:*

$$\sup_i \|\boldsymbol{\phi}_{k,i}\| < \infty \quad \text{a.s.} \quad \forall k \in \mathcal{N}. \quad (52)$$

Assumption 1 can be satisfied by letting the agents find their combination weights in a fully distributed manner (as explained in Section 3). Moreover, it can be easily relaxed such that the connectivity matrix is left stochastic; or drawn from a probability distribution, with its elements being i.i.d. random variables so that the assumption holds in expectation, that is, $C = \mathbb{E}[C]$.

Assumption 2 holds for very expressive function approximations, like deep neural networks with common activation functions.

Assumption 3 is easily imposed in practice. The rightmost term of (50) ensures that the step size for the actor goes to zero faster than the critic. This implies that the actor updates on a slower timescale than the critic.

⁸ a.s. stands for almost surely.

Assumption 4 is standard (and holds easily, e.g., when rewards are bounded) and allows the noise to be averaged by the diminishing step sizes (Borkar, 2008, 12).

Assumption 5 typically requires a stability analysis (Borkar & Meyn, 1999; Borkar, 2008; Lakshminarayanan & Bhatnagar, 2017). One simple way to ensure this assumption would be that each agent projects its parameters to some compact set after performing the local adaptation step (e.g., by clipping the parameters).

Let us introduce the average parameter and gradient vectors:

$$\bar{\phi}_i \triangleq \frac{1}{N} \sum_{k=1}^N \phi_{k,i}, \quad \bar{g}(\cdot) \triangleq \frac{1}{N} \sum_{k=1}^N g_k(\cdot) = \begin{pmatrix} \nabla_{\xi} J(\cdot) \\ \nabla_w \tilde{L}(\cdot) \end{pmatrix}. \quad (53)$$

The following theorem shows that all agents converge to a single average parameter vector.

THEOREM 2. *Under Assumptions 1–5, each agent’s parameters obtained from distributed recursion (46a)–(46b) converge almost surely to their average recursion:*

$$\lim_{i \rightarrow \infty} \phi_{k,i} = \lim_{i \rightarrow \infty} \bar{\phi}_i \quad a.s. \quad \forall k \in \mathcal{N}. \quad (54)$$

Proof. Appendix C.2 includes the proof, which is composed of the following three steps: First, we introduce agreement and disagreement vectors, the former being the average of all agents’ parameters, and the latter being the difference between each agent’s parameter and the agreement vector. Second, we build a recursion on the the disagreement vector and show that it converges almost surely to some point. Finally, we show that the convergent point is the null vector, implying that the agents do not deviate from their average. \square

Our next result shows that $\bar{\phi}_i$ converges to a stationary point of \bar{g} . To show this, we use the ordinary differential equation (ODE) method (Borkar, 2008), which requires three additional assumptions. Consider the following ODE:

$$\dot{\xi}(t) = \nabla_{\xi} J(\xi(t), \pi_w(t)), \quad (55)$$

$$\dot{w}(t) = \nabla_w \tilde{L}(v_{\xi}(t), w(t)). \quad (56)$$

ASSUMPTION 6. *The sequence $\{\mathbf{F}_{k,i+1}\}_{i=0}^{\infty}$ is a martingale difference sequence w.r.t. $\mathcal{F}_{k,i}$:*

$$\mathbb{E}[\mathbf{F}_{k,i+1} | \mathcal{F}_{k,i}] = 0 \quad a.s. \quad i \geq 0, \quad \forall k \in \mathcal{N}. \quad (57)$$

ASSUMPTION 7. *The ODE in (55) has a globally asymptotically stable equilibrium $\lambda(w)$ (uniformly in w), where $\lambda : \mathbb{R}^{M_{\pi}} \rightarrow \mathbb{R}^{M_v}$ is a Lipschitz map.*

ASSUMPTION 8. *The ODE (56) has a globally asymptotically stable attractor \mathbb{W}^* .*

Although Assumptions 6–7 do not hold for the deep neural network representations of the value function used for the experiments, they are convenient to simplify the analysis and rely on the two timescale stochastic approximation analysis due to Borkar (1997). This allows us to illustrate the behaviour of the Diff-DAC architecture with little effort. More general conditions can be found in, for example, Tadic (2004), Ramaswamy and Bhatnagar (2017), Yaji and Bhatnagar (2020).

Moreover, condition (57) in Assumption 6 would hold for a variation of Diff-A2C that used a linear critic with the so-named *compatible features* condition, which is given by

$$\nabla_{\xi} v_{\xi,k,i}(s) = \int_{\mathbb{A}} \nabla_w \log \pi_{w,k,i}(s, a) da, \quad (58)$$

and ensures unbiased policy gradients even with a parametric value function approximation (Sutton *et al.*, 1999; Konda & Tsitsiklis, 2003; Tomczak *et al.*, 2019). Assumption 7 would also hold for multiple linear critics (Sutton *et al.*, 2009; Bhatnagar *et al.*, 2009; Scherrer, 2010).

Assumption 8 is more general than Assumption 7; it holds for parametric policies for which $\nabla_w \tilde{\mathcal{L}}$ has multiple stationary points.

We are ready to state the main result of this section.

THEOREM 3. *Under Assumptions 2–8, the average parameter converges to a stationary point of the average gradient:*

$$\lim_{i \rightarrow \infty} \bar{\phi}_i \in \mathbb{W}^* \quad a.s. \quad \text{such that} \quad \bar{g}(\phi^*) = 0, \quad \forall \phi^* \in \mathbb{W}^*.$$

Proof. See Appendix C.3. □

6 Numerical experiments

In this section, we illustrate the benefits of the diffusion mechanism when applied to actor-critic algorithms. First, we illustrate its stabilising effect on the sampling distribution. Second, we show that Diff-SiAC is able to outperform Dist-MTPS. Third, we show a regularisation effect that can lead to higher performance and better generalisation properties than the centralised architecture. Finally, we explore the influence of the network topology and noisy links on the performance of the algorithm.

6.1 Stability

We evaluate the stabilising properties of the Diff-DAC architecture by comparing Centralised SiAC versus Diff-SiAC on multitask variants of two classic control problems: inverted pendulum and cart-pole swing-up.

The inverted pendulum consists of a rigid pole and an actuated joint with limited torque. The goal is, starting from a random position, to take the pendulum to the upright position and balance. Cart-pole swing-up is an extension of cart-pole balance where the pole starts from the bottom and the goal is to swing the pole to the upright position and balance. The MRL problem consists of 25 tasks where we vary the pole mass and length for the inverted pendulum; and the pole mass, pole length and cart mass for cart-pole swing-up (see task details in Appendix A.1).

Our goal with these experiments is not to compare with state of the art centralised algorithms—which use several advancements to stabilise or improve performance—but to evaluate whether diffusion is an effective mechanism to stabilise learning. Thus, we compare Centralised SiAC versus Diff-SiAC, which are two versions of the same algorithm, described in Section 4.4, with the same advantage estimates and the same neural network architectures (agent details in Appendix A.1).

For Diff-SiAC, the network consists of $N = 25$ agents (each agent is allocated one task), in a random strongly connected graph with average degree $|\bar{\mathcal{N}}| \triangleq \sum_{k=1}^N |\mathcal{N}_k| \approx 4.2$. The connectivity matrix C was obtained with the Hastings rule (Sayed, 2014, 492) so that (6)–(7) hold. The random network topology aims to reflect the sparse connectivity that appears naturally when agents and data are geographically distributed—it is not related to any form of task similarity.

The results are reported as the return for all tasks, which is averaged over 10 test episodes per epoch and scaled by the maximum value achieved by any algorithm. Each epoch consists of five episodes per agent in Diff-SiAC, and $5N$ episodes in total for Centralised SiAC, so that both algorithms simulate the same number of episodes. Each experiment was repeated with six seeds. Figure 2 shows the median and first and third quartiles of the distribution of the return.

We observe that Centralised SiAC starts learning faster than Diff-SiAC in multitask inverted pendulum (left), and achieves near maximum performance for both problems. Faster learning under the centralised architecture is expected since it can compute the gradients with data from all tasks at every iteration, while Diff-SiAC has to wait until the parameters are diffused across the network. However, Centralised SiAC is unstable in both problems, which is also expected since the samples used to estimate the gradients are highly correlated and Centralised SiAC has no decorrelation mechanism.

Diff-SiAC converges more slowly but is much more stable in both problems. Recall that Centralised SiAC and Diff-SiAC have the same network architecture, sample from multiple environments in parallel

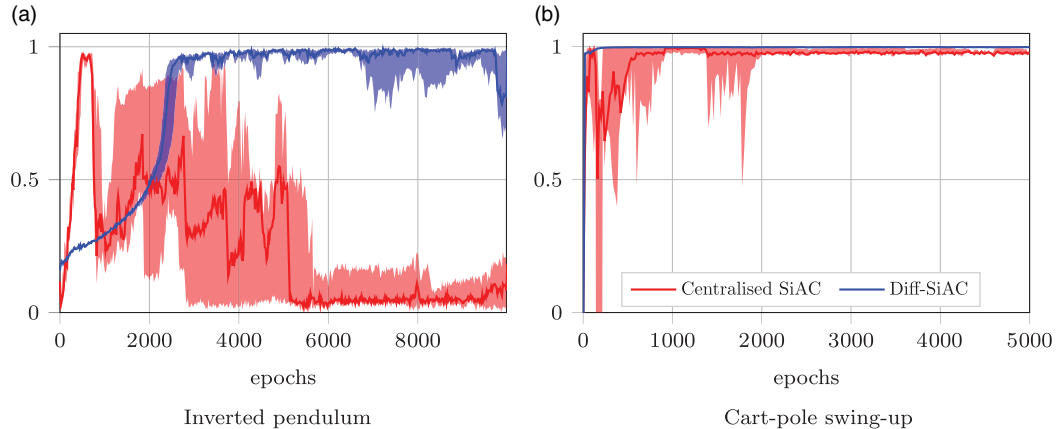


Figure 2 Stability. Learning curves from training with 25 randomly sampled tasks of inverted pendulum (left) and cart-pole swing-up (right). Plots show median and first and third quartiles of the return averaged over all tasks, and six seeds. Centralised SiAC (red) is unstable, while Diff-SiAC (blue) learns robustly.

and are trained with the same amount of data in total. We therefore believe that this enhanced stability is due to a bias removal effect on the sampling distribution introduced by the diffusion mechanism, similar to the distributed exploration capabilities exhibited when applying diffusion to off-policy evaluation (Valcarcel Macua *et al.*, 2015).

We performed similar experiments for the single-task problem—where Centralised SiAC learns from multiple copies of the same task, and all Diff-SiAC agents train on copies of the same task—and observed the same stability improvement (see results in Appendix B).

The instability of Centralised SiAC might be alleviated in multiple ways, like adding a *replay buffer* (Mnih *et al.*, 2013; Lillicrap *et al.*, 2015). Our goal with these experiments is precisely to show that diffusion can be used as an alternative stabilisation technique.

6.2 Comparison with Dist-MTPS

We compare Diff-SiAC with Dist-MTPS on an MRL extension of the standard cart-pole balance problem, which is the only environment considered in this paper that is controllable with linear policies learnt from raw data. The agents have to learn $N = 25$ tasks, characterised by different pole masses and lengths (see Appendix A.2 for details).

In particular, we compare against two variants of Dist-MTPS, which vary by their use of differing policy search methods: REINFORCE (Williams, 1992) and PoWER (Kober & Peters, 2009).

Figure 3 shows that Diff-SiAC learns faster and reaches higher return than Dist-MTPS REINFORCE. Dist-MTPS PoWER converges faster than Diff-SiAC, however, the return of the latter is significantly higher. This is remarkable since Dist-MTPS learns a different policy for each task, while Diff-DAC agents learn a single policy for all tasks.

We also compare with Centralised SiAC in this example. Consistent with the stability experiments, Figure 3 shows that Diff-SiAC learns more slowly but with lower variance than Centralised SiAC. In addition, we observe a new effect: Diff-SiAC clearly achieves higher performance than the centralised architecture. This is a very interesting feature of the diffusion mechanism that has been already reported for non-convex optimisation problems (Valcarcel Macua, 2017, Chapter 4). We believe this is due to a new form of regularisation, which we explore further in the next subsection.

6.3 Regularisation

We extend our analysis to multitask variants of two more challenging and higher dimensional environments: Acrobot and MuJoCo Hopper.




Figure 3 Multitask cart-pole balance with continuous action space. Centralised SiAC learns faster than the distributed approaches and achieves similar performance to Dist-MTPS REINFORCE, but Diff-SiAC achieves the best asymptotic performance and with lower variance. Each epoch consists of five episodes per agent in Diff-SiAC, and $5N$ episodes for Dist-MTPS and Centralised SiAC, so that all algorithms run the same number of episodes. Results are averaged over tasks and five seeds.

In Acrobot the agent must learn a policy to swing a two-link robot from its initial position, approximately hanging down, to a goal position of balancing straight upwards. The randomisation of the environment dynamics generates tasks by sampling the length, mass and inertia of the robot uniformly from a distribution centred on the values of the standard Acrobot environment. This formulation was originally introduced by Packer *et al.* (2018). We design a MRL problem by sampling 25 tasks (see Appendix A.3 for details).

The objective in MuJoCo Hopper is to learn a policy to keep the agent from falling over and to enable the agent to move as quickly as possible in any direction. We implement 2 MRL problems, consisting of 25 tasks each. In one set of experiments, we change the mass of the agent for each task. The second set of experiments introduces wind. We change the direction and intensity of the wind for each task (see Appendix A.3 for details).

The randomisation of the environment parameters—pole mass, pole length, and inertia for Acrobot and agent mass and wind speed and direction for Hopper—is calibrated to generate an MRL problem more challenging than the original single-task environments.

For these MRL problems, we compare 3 agent implementations: (i) Diff-A2C on a network of $N = 25$ agents, each one with a single process sampling from one of the tasks; (ii) Centralised A2C where 1 agent has 25 processes, 1 per task, so it simultaneously learns from experience collected on all tasks; and (iii) Specialised A2C agents which train on each of the tasks individually but each one with the equivalent resources to the Centralised A2C agent (i.e., each of the 25 specialised agents has access to 25 copies of one of the sampled environments).

The comparison of Diff-A2C and Centralised A2C enables us to compare two learning architectures using approximately the same amount of computation. Centralised A2C centralises learning while maintaining decentralised sampling from environments run in parallel. Diff-A2C decentralises both sampling and learning. The specialised agents provide an upper bound to performance during training thereby offering an opportunity to evaluate the performance gap incurred by the multitask agents.

For this set of experiments, all agents are trained using the RMSProp optimiser (Tieleman & Hinton, 2012). This provides an opportunity to show that the Diff-DAC architecture is optimiser agnostic in practice—although it is derived and analysed assuming vanilla stochastic gradient descent. Indeed, we attain similar results when running the experiments presented here with the Adam optimiser. Full details of agent hyperparameters are provided in Appendix A.

The ring topology is used for Diff-A2C in this section. This topology can be challenging due to the sparsity of the connectivity, which implies a slower diffusion rate (it takes up to 12 iterations to diffuse each agent’s parameters to every other agent) and less robust communication (there are only 2 possible paths between every pair of agents).

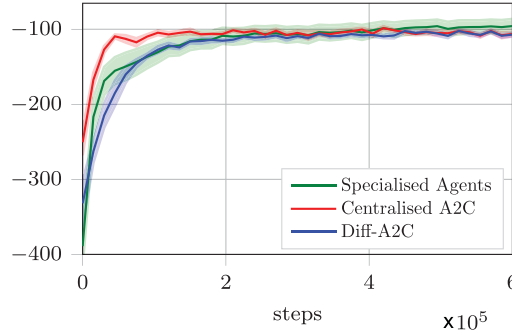


Figure 4 Regularisation Acrobot. Learning curves from training with 25 randomly sampled tasks of Random Extreme Acrobot. For each experiment, we average performance over all 25 tasks. The plot shows the average ± 1 standard error over 8 seeds. Centralised A2C learns faster. Diff-A2C achieves similar performance to Centralised A2C. Specialised agents achieve the highest performance, but with small gap.




Figure 5 Regularisation MuJoCo Hopper. Learning curves from training with 25 randomly sampled tasks of Random Mass (left) and Random Wind (right) MuJoCo Hopper. For each experiment, we average performance over all 25 tasks. The plots show the average ± 1 standard error (of averages) over 8 seeds. Specialised agents achieve the highest performance, with a significant gap w.r.t. to the other agents. Diff-A2C agents clearly outperform Centralised A2C.

Figure 4 displays the results for experiments on the Acrobot environments. Diff-A2C and Centralised A2C achieve the same final performance in training while the specialised agents achieve slightly greater ultimate performance. The optimality gap is small, meaning that both Diff-A2C and Centralised A2C learn a single policy that is expressive enough to control the 25 tasks near optimally. Consistent with previous experiments, we see that Centralised A2C learns faster initially. In this case, however, Centralised A2C does not exhibit stability issues, as A2C’s decentralised sampling stabilises the sampling distribution, similar to the stabilising effect of the Diff-DAC architecture.

Figure 5 shows results for the two MRL problems with MuJoCo Hopper: random mass and random wind. These are significantly more challenging tasks than Acrobot. We observe the pattern of initially fast learning for the centralised architecture. But now we can clearly see that the optimality gap is smaller for Diff-A2C in both set of experiments. This supports the hypothesis that diffusion induces some form of regularisation effect that allows the agents to find a better policy that can control the set of tasks more effectively.

However, since each Diff-A2C agent maintains its own policy, it is not clear whether the performance improvement over the centralised architecture is due to finding a better policy or to being able to exploit policies specialised to each task. In order to investigate, we look at how similar the policies learnt by the Diff-A2C agents are in terms of parameters and behaviour on the tasks used in training. If the policies of all Diff-A2C agents are equal, then we can be sure that the policy found via Diff-A2C is better than that found by Centralised A2C.

Table 1 Parameter similarity. Relative deviation of final parameters of trained agents. This is determined by calculating the mean and the standard deviation of each parameter of the trained agent over 25 agents (1 per task). The standard deviations are then divided by the mean. The reported values are the average over the parameters and over eight seeds.

	Diff-A2C (%)	Specialised agents (%)
Randomised Acrobot	3.56	59.27
Random Mass Hopper	3.10	114.08
Random Wind Hopper	3.14	116.23

Table 2 Zero-shot learning. Evaluation performance for trained agents on unseen Acrobot tasks. The values are averages over 8 seeds and 10 episodes each. For the specialised and Diff-A2C agents, a single agent from each experiment was chosen at random and evaluated on the environment they were trained on as well as the same previously unseen easy and hard environments. Values in brackets denote a 95% confidence interval.

	Training Environment	Out-of-sample Easy environment	Out-of-sample Hard environment
Specialised A2C	-82.1 (-98.3, -66.0)	-80.0(-79.1, -62.8)	-163.8(-182.3, -145.3)
Diff-A2C	-109.4(-128.0, -90.9)	-48.8 (-57.4, -40.2)	-122.6 (-142.9, -102.3)
Centralised A2C	-115.6(-138.6, -92.6)	-53.4(-63.6, -43.3)	-141.8(-166.2, -117.3)

Table 1 shows the relative deviation of the parameters as a measure of the similarity of the agents at the end of training. For Diff-A2C and specialised agents, the value is calculated by dividing the standard deviation by the mean of each model parameter across all agents, then averaging the results over all parameters, and finally averaging over eight runs with differing random seeds. The results show that all Diff-A2C agents converge to very similar policy networks. On the other hand, the specialised agents converge to policy networks that are significantly different from each other. This is indeed consistent with Theorem 2, which shows almost certain asymptotic agreement amongst the agents, and supports the hypothesis that diffusion allows us to find a single policy that performs better than the one found by the centralised architecture.

In order to further validate the convergence of Diff-A2C agents to a common solution, we ran each of the 25 trained agents on the task they were trained on as well as the tasks of all other agents in the network. On average, agents achieved 0.07% (± 3.4) higher return, over 10 episodes, in the task they were trained on compared to the tasks of their peers (result averaged over all agents and 8 random seeds). Since the performance difference is not statistically significant, we conclude the small deviations in the policy parameters are not significant.

The above analysis suggests that the policy learnt by diffusion should be better able to generalise across tasks than the one found by the centralised architecture. To consider the question of generalisation, we performed a zero-shot learning experiment in the Acrobot environment, evaluating our trained agents on newly sampled environments (details of these environments are provided in Appendix A.4). Results are presented in Table 2, and show that Diff-A2C yields agents better able to generalise to new environments. The Centralised A2C agents show some ability to generalise to new environments but achieve lower returns in the evaluation when compared to Diff-A2C despite attaining near identical performance at the end of training. Unsurprisingly, the specialised agents do not generalise well. Specialised agents attain returns approximately 33% higher than Diff-A2C agents on the environment in which they specialise but are outperformed by 31% and 25% when transferred to held out easy and hard environments, respectively⁹.

⁹ Table A.1 in Appendix C.4 presents the results in terms of performance relative to specialised agents in full. Relative performance values are calculated from Table 2.




Figure 6 Network topology. Influence of network topology in single-task cart-pole balance with continuous action space. Diffusion is able to combine the experience of all agents in a way that is relatively insensitive to network sparsity. Moreover, by increasing the number of agents, we obtain higher performance, suggesting Diff-SiAC is able to find a better local optimum. Results are averaged over all agents and five seeds.




Figure 7 (Noisy links) Learning trajectories of Diff-A2C agents with differing probabilities of links being dropped at an update step. The probability, denoted by p , parameterises independent Bernoulli distributions, one per connection. Experiments conducted with 14 agents connected in a ring topology. The plots show performance averaged over all agents in the graph and averaged over 8 seeds.

6.4 Network dependence

In this section, we perform experiments to study the influence of the communications network on the performance of the proposed Diff-DAC architecture. First, we evaluate the impact of the network topology by varying the size of the graph and its neighbourhoods. We then explore the impact of noisy or intermittent links on agent performance.

We evaluate the impact of the network topology with Diff-SiAC agents learning the *single-task* cart-pole balance problem, where all agents operate in different copies of exactly the same environment (see details in Appendix A.4). Figure 6 shows that for the same network size, $N = 25$, a relatively sparse network, $|\bar{N}| \approx N/6$, learns slightly more slowly but achieves performance similar to a more dense network, $|\bar{N}| \approx N/3$. This is consistent with the theoretical analysis that shows that as long as the connectivity matrix, C , satisfies conditions (6)–(7), asymptotic performance is guaranteed independent on the network sparsity. From a practical point of view, this is a promising result that suggests that communication and computational costs per agent can be reduced without penalising performance, since they scale with the number of neighbours. In addition, we see that larger number of agents $N = 100$ improves the final performance, suggesting that the regularisation effect discussed in Section 6.3 is stronger with a larger number of agents.

Finally, we consider the impact of failing links on training in the original (single-task) Acrobot and MuJoCo Hopper environments. These experiments emulate issues of noisy or interrupted message passing, typically seen in real deployments. As shown in Figure 7, as the probability of any link in the network failing at a given update step increases, the Diff-A2C agents’ learning is slowed. However, the agents

converge towards similar final performance. This shows that diffusion is robust to failed links even when the failure rate is as high as 80%. This was predicted in Section 5, where we discuss that conditions on the connectivity matrix need only hold in expectation. The Diff-A2C experiments were conducted with a ring topology meaning that dropped links easily cause the graph to become disconnected. We would therefore expect this topology to be particularly vulnerable to dropped links. In the case of a centralised agent, however, should a single environment instance fail, the agent’s learning is liable to fail completely unless explicit countermeasures are introduced.

7 Conclusions

We considered the MRL problem where tasks are MDPs with parametrised dynamics where parameters are drawn from some distribution.

We used standard optimal control and convex optimisation theory to derive a tabular actor-critic algorithm as an instance of dual ascent for finding the saddle point of a Lagrangian. We also re-derived the policy gradient theorem from this same formulation. These derivations are interesting in themselves and provide novel insights in the actor-critic framework. By approximating the exact method with parametric approximations, we obtained the Diff-DAC architecture, which can be applied to transform many actor-critic algorithms into fully distributed implementations that scale to a large number of tasks. We also proved convergence of Diff-DAC to a common policy under mild assumptions, and to stationary point of the joint value and policy gradients under more restrictive assumptions.

Results from our experiments showed that the Diff-DAC architecture can achieve higher performance than Dist-MTPS. This is a remarkable result since the Diff-DAC agents converge to a single policy that behaves better than the task-dependent policies obtained by Dist-MTPS. Moreover, Diff-DAC can solve complex problems that are uncontrollable from raw data by linear policies, while Dist-MTPS requires feature engineering.

We also demonstrate that Diff-DAC is very stable and achieves similar or higher performance and shows better generalisation than a comparable centralised approach. This suggests that the sparse connectivity amongst agents induces a regularisation effect that helps them to achieve better local optima.

Finally, the experiments with intermittent connections between agents provide empirical evidence of the robustness to communication disruption Diff-DAC achieves in distributed learning systems.

Acknowledgements

We thank David Baldazo, Daniel García-Ocaña Hernández, and Santiago Zazo for insightful preliminary discussions; Felix Leibfried for his support with the experiments; Peter Vrancx, Haitham Bou-Ammar, and Rasul Tutunov for helpful comments; and the anonymous reviewers and the Deputy Editor Patrick Mannion for their comments and suggestions that have helped to improve the presentation of the paper.

Competing interests

The authors declare none.

References

- Andreas, J., Klein, D. & Levine, S. 2017. Modular multitask reinforcement learning with policy sketches. In *Proceedings of the International Conference on Machine Learning (ICML)*, 166–175.
- Arrow, K. J., Hurwicz, L. & Uzawa, H. 1958. *Studies in Linear and Non-linear Programming*. Stanford University Press.
- Assran, M., Romoff, J., Ballas, N., Pineau, J. & Rabbat, M. 2019. Gossip-based actor-learner architectures for deep reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, 13320–13330.
- Baird III, L. C. 1993. *Advantage Updating*, Technical report, Wright Lab Wright-Patterson AFB OH.
- Bertsekas, D. P. 2009. *Convex Optimization Theory*. Athena Scientific.

- Bertsekas, D. P. 2012. *Dynamic Programming and Optimal Control*, 4th edition, 2. Athena Scientific.
- Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M. & Lee, M. 2009. Natural actor-critic algorithms. *Automatica* **45**(11), 2471–2482.
- Bianchi, P. & Jakubowicz, J. 2013. Convergence of a multi-agent projected stochastic gradient algorithm for non-convex optimization. *IEEE Transactions on Automatic Control* **58**(2), 391–405.
- Borkar, V. 2008. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press.
- Borkar, V. S. 1997. Stochastic approximation with two time scales. *Systems and Control Letters* **29**(5), 291–294.
- Borkar, V. S. & Meyn, S. 1999. The O.D.E. method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization* **38**, 447–469.
- Bou-Ammar, H., Eaton, E., Ruvolo, P. & Taylor, M. 2014. Online multi-task learning for policy gradient methods. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1206–1214.
- Boyd, S. & Vandenberghe, L. 2004. *Convex Optimization*. Cambridge University Press.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J. & Zaremba, W. 2016. OpenAI Gym.
- Chen, J. & Sayed, A. H. 2013. Distributed Pareto optimization via diffusion strategies. *IEEE Journal of Selected Topics in Signal Processing* **7**(2), 205–220.
- El Bsati, S., Bou-Ammar, H. & Taylor, M. E. 2017. Scalable multitask policy gradient reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 1847–1853.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I. et al. 2018. Impala: scalable distributed deep-rl with importance weighted actor-learner architectures. In *ICML*.
- Fu, J., Levine, S. & Abbeel, P. 2016. One-shot learning of manipulation skills with online dynamics adaptation and neural network priors. In *IEEE RSJ International Conference on Intelligent Robots and Systems (IROS)*, 4019–4026.
- Golub, G. & Van Loan, C. 1996. *Matrix Computations*. Johns Hopkins University Press.
- Grondman, I., Busoniu, L., Lopes, G. A. D. & Babuska, R. 2012. A survey of actor-critic reinforcement learning: standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* **42**(6), 1291–1307.
- Heess, N., Wayne, G., Silver, D., Lillicrap, T., Erez, T. & Tassa, Y. 2015. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems (NIPS)*, 2926–2934.
- Horn, R. & Johnson, C. 1990. *Matrix Analysis*. Cambridge University Press.
- Kar, S., Moura, J. M. F. & Poor, H. V. 2013. QD-learning: a collaborative distributed strategy for multi-agent reinforcement learning through consensus + innovations. *IEEE Transactions on Signal Processing* **61**(7), 1848–1862.
- Karp, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*. Springer, 85–103.
- Kingma, D. & Ba, J. L. 2015. Adam: a method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Kober, J. & Peters, J. R. 2009. Policy search for motor primitives in robotics. In *Advances in Neural Information Processing Systems (NIPS)*, 849–856.
- Konda, V. R. & Tsitsiklis, J. N. 2003. On actor-critic algorithms. *SIAM Journal on Control and Optimization* **42**(4), 1143–1166.
- Lakshminarayanan, C. & Bhatnagar, S. 2017. A stability criterion for two timescale stochastic approximation schemes. *Automatica* **79**, 108–114.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D. & Wierstra, D. 2015. Continuous control with deep reinforcement learning.
- Melo, F. S. & Lopes, M. 2008. Fitted natural actor-critic: a new algorithm for continuous state-action MDPs. In *Machine Learning and Knowledge Discovery in Databases*, **5212**. Springer, 66–81.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D. & Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 1928–1937.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D. & Riedmiller, M. 2013. Playing atari with deep reinforcement learning. arXiv preprint [arXiv:1312.5602](https://arxiv.org/abs/1312.5602).
- Ng, A. Y., Parr, R. & Koller, D. 1999. Policy search via density estimation. In *Advances in Neural Information Processing Systems (NIPS)*, Denver, CO, USA, 1022–1028.
- Packer, C., Gao, K., Kos, J., Krähenbühl, P., Koltun, V. & Song, D. 2018. Assessing generalization in deep reinforcement learning.
- Parisotto, E., Ba, J. L. & Salakhutdinov, R. 2016. Actor-mimic: deep multitask and transfer reinforcement learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

- Powell, W. B. & Ma, J. 2011. A review of stochastic algorithms with continuous value function approximation and some new approximate policy iteration algorithms for multidimensional continuous applications. *Journal of Control Theory and Applications* **9**(3), 336–352.
- Puterman, M. L. 2005. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 2nd edition. John Wiley & Sons.
- Ramaswamy, A. & Bhatnagar, S. 2017. A generalization of the Borkar-Meyn theorem for stochastic recursive inclusions. *Mathematics of Operations Research* **42**(3), 648–661.
- Sayed, A. H. 2014. Adaptation, learning, and optimization over networks. *Foundations and Trends in Machine Learning* **7**(4–5), 311–801.
- Scherrer, B. 2010. Should one compute the temporal difference fix point or minimize the Bellman residual? The unified oblique projection view. In *Proceedings of the International Conference on Machine Learning (ICML)*, 959–966.
- Schulman, J., Moritz, P., Levine, S., Jordan, M. & Abbeel, P. 2015. High-dimensional continuous control using generalized advantage estimation. arXiv preprint [arXiv:1506.02438](https://arxiv.org/abs/1506.02438).
- Sutton, R. S., Maei, H. R., Precup, D., Bhatnagar, S., Silver, D., Szepesvari, C. & Wiewiora, E. 2009. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proceedings of the International Conference on Machine Learning (ICML)*, 993–1000.
- Sutton, R. S., Mcallester, D., Singh, S. & Mansour, Y. 1999. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems (NIPS)*, 1057–1063.
- Tadic, V. B. 2004. Almost sure convergence of two time-scale stochastic approximation algorithms. In *IEEE American Control Conference*, **4**, 3802–3807.
- Taylor, M. E. & Stone, P. 2009. Transfer learning for reinforcement learning domains: a survey. *Journal of Machine Learning Research* **10**, 1633–1685.
- Teh, Y. W., Bapst, V., Czarnecki, W. M., Quan, J., Kirkpatrick, J., Hadsell, R., Heess, N. & Pascanu, R. 2017. Distal: robust multitask reinforcement learning. arXiv preprint [arXiv:1707.04175](https://arxiv.org/abs/1707.04175).
- Tieleman, T. & Hinton, G. 2012. Lecture 6.5- RMSProp: divide the gradient by a running average of its recent magnitude.
- Tomczak, M. B., Valcarcel Macua, S., de Cote, E. M. & Vrancx, P. 2019. Compatible features for monotonic policy improvement.
- Tutunov, R., Bou-Ammar, H. & Jadbabaie, A. 2016. An exact distributed newton method for reinforcement learning. In *IEEE Conference on Decision and Control (CDC)*, 1003–1008.
- Valcarcel Macua, S. 2017. *Distributed Optimization, Control and Learning in Multiagent Networks*. PhD thesis, Universidad Politécnica de Madrid.
- Valcarcel Macua, S., Chen, J., Zazo, S. & Sayed, A. H. 2015. Distributed policy evaluation under multiple behavior strategies. *IEEE Transactions on Automatic Control* **60**(5), 1260–1274.
- Valcarcel Macua, S., Tukiainen, A., Hernández, D. G.-O., Baldazo, D., de Cote, E. M. & Zazo, S. 2017. Diff-DAC: distributed actor-critic for average multitask deep reinforcement learning. arXiv preprint [arXiv:1710.10363](https://arxiv.org/abs/1710.10363).
- van der Meulen, R. 2015. Gartner says 6.4 billion connected ‘things’ will be in use in 2016, up 30 percent from 2015. <http://www.gartner.com/newsroom/id/3165317>.
- Van Hasselt, H. 2012. Reinforcement learning in continuous state and action spaces. In *Reinforcement Learning*. Springer, 207–251.
- Wei, E. & Ozdaglar, A. 2012. Distributed alternating direction method of multipliers. In *IEEE Annual Conf. Decision and Control (CDC)*, 5445–5450.
- Weinstein, A. & Littman, M. L. 2012. Bandit-based planning and learning in continuous-action markov decision processes. In *International Conference on Automated Planning and Scheduling (ICAPS)*.
- Wierstra, D., Schaul, T., Glasmachers, T., Sun, Y., Peters, J. & Schmidhuber, J. 2014. Natural evolution strategies. *Journal of Machine Learning Research* **15**(1), 949–980.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* **8**(3–4), 229–256.
- Yaji, V. G. & Bhatnagar, S. 2020. Stochastic recursive inclusions in two timescales with nonadditive iterate-dependent Markov noise. *Mathematics of Operations Research* **45**(4), 1405–1444.
- Zhang, K., Yang, Z., Liu, H., Zhang, T. & Basar, T. 2018. Fully decentralized multi-agent reinforcement learning with networked agents. In *Proceedings International Conference on Machine Learning (ICML)*, 5872–5881.
- Zhao, X. & Sayed, A. H. 2012. Performance limits for distributed estimation over LMS adaptive networks. *IEEE Transactions on Signal Processing* **60**(10), 5107–5124.
- Zhao, X. & Sayed, A. H. 2015. Asynchronous adaptation and learning over networks—part i: modeling and stability analysis. *IEEE Transactions on Signal Processing* **63**(4), 811–826.

Appendix A Numerical experiment details

A.1 Stability experiments (Section 6.1)

Inverted pendulum. The pendulum consists of a rigid pole and an actuated joint, with maximum torque clipped to the interval $[-2, 2]$. The pendulum starts at a random angle sampled from a uniform distribution on $[-\pi, \pi]$. The goal is to take the pendulum to the upright position and balance there. In the original single-task setting, pole mass and length are $(1.0, 1.0)$. Our MRL problem consists of 25 tasks with mass in $\{0.8, 0.9, 1.0, 1.1, 1.2\}$, and length in $\{0.8, 0.9, 1.0, 1.1, 1.2\}$.

Cart-pole swing-up. We extend cart-pole balance to the case where the pole starts from the bottom and the task is to swing the pole to the upright position and balance. The reward function is $r = \frac{2}{1+e^d} + \cos(\psi)$, where d is the Euclidean distance of the pole from the upright position at the centre of the track, and ψ is the pole angle. The original single-task uses parameters $(0.5, 0.25, 0.5)$ for the pole mass, pole half-length, and cart mass, respectively. Our cart-pole swing-up MRL problem consists of 25 tasks, where pole mass is in $\{0.1, 0.2, 0.3, 0.4, 0.5\}$, pole half-length is in $\{0.2, 0.4, 0.6, 0.8, 1.0\}$, and cart mass is 0.5.

We use $\gamma = 0.99$ in all cases.

Agent. The actor and critic neural networks for both algorithms consist of 2 hidden layers of 400 neurons each with ReLU activation functions. The output layer for the critic network is linear. The output of the actor network includes a *tanh* activation function that determines the mean of a normal distribution, and a *Softplus* activation function that determines the variance for the normal distribution. We have also included a penalty in the loss function equal to the entropy of the policy, with penalty coefficient 0.0005.

We use the ADAM optimiser (Kingma & Ba, 2015), with learning rate 0.01 for the critic and 0.001 for the actor. Diff-SIAC performs a learning step ($i \leftarrow i + 1$) every fifth episode.

A.2 Comparison with Dist-MTLS (Section 6.2)

Cart-pole balance. We use the OpenAI Gym (Brockman *et al.*, 2016) implementation, but with continuous force. The action follows a Gaussian distribution with mean in the interval $[-10, 10]$. The episode finishes when either the pole is beyond 12 degrees from vertical, the cart moves more than 2.4 units from the centre, or the simulation reaches 200 time steps. The original single-task setting uses parameters $(0.1, 0.5, 1.0)$ for the pole mass, pole half-length, and cart mass, respectively. Our MRL problem consists of 25 tasks with the following parameters: pole mass in $\{0.1, 0.325, 0.55, 0.775, 1\}$, pole length in $\{0.05, 0.1625, 0.275, 0.3875, 0.5\}$, and cart mass 1 in all cases.

We also use $\gamma = 0.99$ in this problem.

Agent. We use the same agent architecture and hyperparameters as described in Appendix A.1. above for the stability experiments from Section 6.1.

A.3 Regularisation (Section 6.3)

Acrobot. We use the Extreme Acrobot multitask environment introduced by Packer *et al.* (2018). Based on the original OpenAI Gym implementation of the Acrobot environment (Brockman *et al.*, 2016), this setting is as the original Acrobot environment but the robot’s mass, length, and inertia are all sampled from independent uniform distributions on $[0.5, 0.75] \cup [1.25, 1.5]$. These environment parameters are sampled at the start of each experiment and fixed throughout the duration of training. In the original single-task setting, the robot’s mass, length and inertia parameters are all set to 1. The agent is penalised by receiving a reward of -1 for each time step until it reaches the goal state. For more details, see Packer *et al.* (2018).

MuJoCo Hopper. Inspired by Packer *et al.* (2018), we adapt the continuous control MuJoCo Hopper environment to develop a continuum of environments. We implement two separate extensions, one where

Table A.1 Zero-shot learning on Acrobot. Here we show performance relative to the specialised agents. Negative values show the performance gap, which is the case on the training environments. Positive values denote outperforming the specialised agents, which is the case for out-of-sample environments. For the specialised and Diff-A2C agents, a single agent from each experiment was chosen at random and evaluated on the environment they were trained on as well as the same held out easy and hard environments. Values in brackets denote a 95% confidence interval. The values are averages over 8 seeds, 10 episodes each.

	Training Environment (%)	Out-of-sample Easy environment (%)	Out-of-sample Hard environment (%)
Diff-A2C	-33.2 (-55.9, -10.6)	31.3 (19.2, 43.3)	25.1 (12.8, 37.5)
Centralised A2C	-40.8(-68.7, -12.8)	24.7 (10.4, 39.0)	13.4 (-1.5, 28.4)

the agent is in a setting with a randomly chosen constant wind and another where the agent’s mass is scaled by a random value. In the setting with wind, we first adapt the medium through which the agent moves to have the same density and viscosity as air. Then, wind is implemented in the xy -plane. Wind is, therefore, a random 2-vector in which each element is sampled from an independent beta distribution on the interval $[-15, 15]$. Once sampled, the direction and strength of the wind are held constant. The beta distributions are parameterised with $\alpha = \beta = 0.4$ in order to encourage more extreme conditions. We found that this was necessary to make the environment meaningfully different from the standard hopper environment. In the environment where the agent’s mass is randomly scaled, a scalar is sampled at the environment’s inception and is used to scale the agent’s mass. This scalar is fixed thereafter. The scalar is sampled from a Beta distribution transformed to be over the interval $[0.05, 0.95]$. We use $\alpha = \beta = 0.1$ to again encourage more extreme settings and ensure a meaningfully different setting to that of the unmodified hopper environment.

We again use a discount factor, $\gamma = 0.99$ in these settings.

Agent. The actor and critic are both implemented as neural networks with 2 hidden layers of 64 units with \tanh activation function. The output layer is linear for both actor (producing logits parameterising the action distribution) and critic.

To allow for fair comparison across the different agents, we tuned the hyperparameters of training for both Diff-A2C and Centralised A2C separately. The specialised agents use the hyperparameters for Centralised A2C as they are the same agents but without environment randomisation. Under Diff-A2C each agent only has one environment, as opposed to Centralised A2C where the agent samples experience from 25 environments in parallel. We found that Diff-A2C was most effective with a lower learning rate and a greater number of samples between updates when compared to Centralised A2C. This is consistent with findings of recent works on fully distributed RL (Assran *et al.*, 2019). In all the experiments of Section 6.3, the agents are trained with the RMSProp optimiser (Tieleman & Hinton, 2012) with $\alpha = 0.99$ and $\epsilon = 0.00001$. There is an entropy term included in the loss function with coefficient 0.01.

Through tuning the learning rate for Centralised A2C and Diff-A2C independently we arrived at the following hyperparameters used in all experiments relating to Acrobot or MuJoCo Hopper. Centralised A2C: Learning Rate 0.002, 5 steps per update. Diff-A2C: Learning Rate 0.0007, 60 steps per update.

For all the experiments in this section, the connectivity matrix C used in Diff-A2C represents a ring, with every agent connected with two neighbours, such that equal weight is given to parameters from each agent involved in the combination step.

Generalisation. The experiments used to consider generalisation in Acrobot found that the different sampled environments for Acrobot vary in difficulty as noted by the average performance of all agents during evaluation. In light of this, we chose an easy and a hard Acrobot environment to evaluate on. The easy environment has a robot arm length of 0.7046, an agent mass of 0.5259 and an inertia parameter of 0.6346. Contrastingly, the hard environment has a robot arm length of 1.3963, an agent mass of 1.3929, and an inertia parameter of 0.6256. We refer the interested reader to the work of Packer *et al.* (2018) for a detailed discussion regarding the influence of these environment parameters. Table A.1 presents the

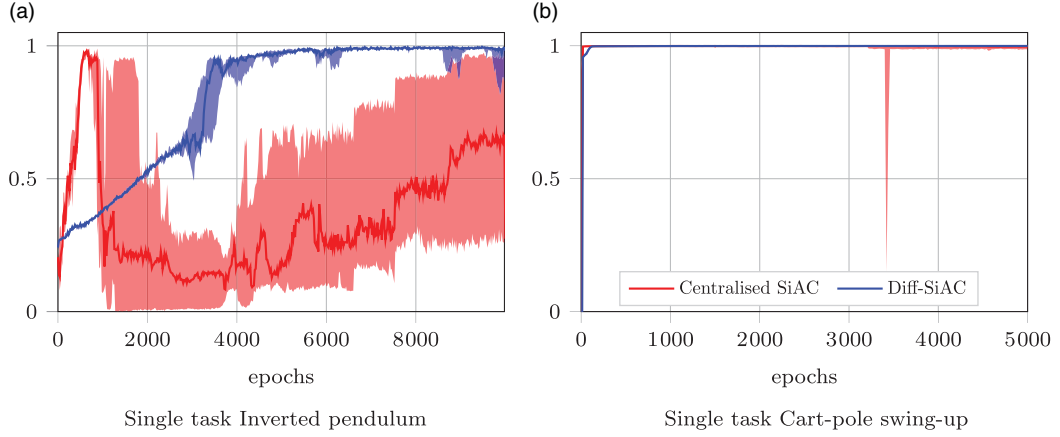


Figure B.1 Stability in single task. Learning curves from training with 25 copies of the standard version of the inverted pendulum (left) and cart-pole swing-up (right) environments. Plots show median and first and third quartiles of the return averaged over all tasks, and 6 seeds.

results from Table 2 of Section 6.3 but explicitly converted to performance relative to the specialised A2C agents (which represent 0% in all cases).

A.4 Network dependence (Section 6.4)

Network topology. We use the same implementation of the cart-pole balance described in Appendix A.2, but with the original single-task parameters (0.1, 0.5, 1.0) for the pole mass, pole half-length, and cart mass, respectively. We use the same agent described in Appendix A.1.

Noisy links. We use the original (single-task) Gym implementation of Acrobot (Brockman *et al.*, 2016). We use the same agent described in Appendix A.3 again with agents arranged in a ring topology.

Appendix B Stability experiments in the single-task problems

In this appendix, we compare Centralised SiAC and Diff-SiAC in the standard single-task problem for the inverted pendulum and cart-pole swing-up environments. Figure B.1 shows that Centralised SiAC follows the same trend as for the multitask setting, it learns faster but is more unstable in both problems, especially in inverted pendulum. On the other hand, Diff-SiAC learns robustly in both environments despite the highly correlated advantage estimates at each agent. This result further supports our hypothesis that diffusion introduces a stabilising effect in the parameter update rule.

Appendix C Proofs

C.1 Proof of Theorem 1

We expand the chain rule for vector functions, where $\langle \cdot, \cdot \rangle$ denotes inner product:

$$\begin{aligned}
 \nabla_w \tilde{L}(v, w) &= \left\langle \nabla_w d_w, \frac{\partial L(v, w)}{\partial d_w} \right\rangle \\
 &= \int_{\mathbb{S}} \int_{\mathbb{A}} \nabla_w d(s, a) \frac{\partial L(v, w)}{\partial d_w(s, a)} ds da \\
 &\stackrel{(a)}{=} \int_{\mathbb{S}} \int_{\mathbb{A}} (\rho_\gamma^{\pi_w}(s) \nabla_w \pi_w(a|s) + \pi_w(a|s) \nabla_w \rho_\gamma^{\pi_w}(s)) \bar{A}(s, a) ds da \\
 &\stackrel{(b)}{=} \int_{\mathbb{S}} \rho_\gamma^{\pi_w}(s) \int_{\mathbb{A}} \nabla_w \pi_w(a|s) \bar{A}(s, a) ds da \\
 &\stackrel{(c)}{=} \int_{\mathbb{S}} \rho_\gamma^{\pi_w}(s) \int_{\mathbb{A}} \pi_w(a|s) \nabla_w \log \pi_w(a|s) \bar{A}(s, a) ds da. \tag{C1}
 \end{aligned}$$

In step (a), we expanded the gradient of the product. In (b), we used Bellman equation that says that the expected value of the advantage function over the policy distribution is zero, so that:

$$\nabla_w \rho_{\gamma}^{\pi_w}(s) \int_{\mathbb{A}} \pi_w(a|s) \bar{A}(s, a) da = 0. \quad (\text{C2})$$

In (c), we used a standard reparameterisation trick: $\nabla_w \pi_w(a|s) = \pi_w(a|s) \nabla_w \log \pi_w(a|s)$. \square

C.2 Proof of Theorem 2

Let $M \triangleq M_v + M_\pi$ be the length of each agent's parameter vector, $\phi_{k,i}$, and let I_M denote the identity matrix of size M . We introduce a network-wide combination matrix: $\mathcal{C} \triangleq C \otimes I_M$, where \otimes denotes the Kronecker product.

In order to prove Theorem 2, we first present a useful lemma. For this, we introduce the following *agreement* and *disagreement* projection matrices, which will be used to compute the average and the difference from the average of a set of vectors (stacked on top of each other in a long vector), respectively:

$$\mathbf{P}^{\text{agr}} \triangleq \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N} \otimes I_M, \quad \mathbf{P}^{\text{dis}} \triangleq I_{MN} - \mathbf{P}^{\text{agr}} = I_{MN} - \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N} \otimes I_M, \quad (\text{C3})$$

LEMMA 1. Let C be primitive and doubly stochastic. Then: $\|\mathbf{P}^{\text{dis}} \mathcal{C}\| < 1$.

Proof. We can easily obtain the following equivalence:

$$\mathbf{P}^{\text{dis}} (C \otimes I_M) = \left(C - \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N} \right) \otimes I_M. \quad (\text{C4})$$

Let λ_n , u_n , and v_n denote the n th eigenvalue, and n th right and left eigenvectors of C , respectively. From Assumption 1 we have: $C = \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N} + \sum_{n=2}^N \lambda_n u_n v_n^\top$. Hence, we bound the 1-norm:

$$\left\| C - \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N} \right\|_1 = \left\| \sum_{n=2}^N \lambda_n u_n v_n^\top \right\|_1 < \left\| \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N} + \sum_{n=2}^N \lambda_n u_n v_n^\top \right\|_1 = \|C\|_1 = 1.$$

We obtain the same result for infinity norm. We conclude by applying a well known inequality for the 2-norm (Golub & Van Loan, 1996, Corollary 2.3.2): $\|\mathbf{P}^{\text{dis}} \mathcal{C}\| \leq \sqrt{\|\mathbf{P}^{\text{dis}} \mathcal{C}\|_1 \|\mathbf{P}^{\text{dis}} \mathcal{C}\|_\infty} < 1$. \square

We adapt previous works (Bianchi & Jakubowicz, 2013; Valcarcel Macua, 2017, Chapter 4.6.2.) to our case. Let's start introducing the following network-wide variables:

$$\phi_i \triangleq \begin{pmatrix} \phi_{1,i} \\ \vdots \\ \phi_{N,i} \end{pmatrix}, \quad \mathcal{G}(\phi_i) \triangleq \begin{pmatrix} g_1(\phi_{1,i}) \\ \vdots \\ g_N(\phi_{N,i}) \end{pmatrix}, \quad \mathbf{F}_{i+1} \triangleq \begin{pmatrix} \mathbf{F}_{1+1,i} \\ \vdots \\ \mathbf{F}_{N+1,i} \end{pmatrix}. \quad (\text{C5})$$

Now, we can stack all agents' recursions as follows:

$$\phi_{i+1} = \mathcal{C} (\phi_i - \alpha_{i+1} (\mathcal{G}(\phi_i) + \mathbf{F}_{i+1})). \quad (\text{C6})$$

To aid our explanation, we introduce network-wide *agreement* and *disagreement* vectors:

$$\phi_i^{\text{agr}} \triangleq \mathbf{P}^{\text{agr}} \phi_i = (\mathbb{1}_N \otimes I_M) \bar{\phi}_i, \quad \phi_i^{\text{dis}} \triangleq \mathbf{P}^{\text{dis}} \phi_i = \phi_i - (\mathbb{1}_N \otimes I_M) \bar{\phi}_i, \quad (\text{C7})$$

Note that ϕ_i , ϕ_i^{agr} and ϕ_i^{dis} have length MN , while $\phi_{k,i}$ and $\bar{\phi}_i$ have length M .

From (65), we have: $\phi_i = \phi_i^{\text{agr}} + \phi_i^{\text{dis}}$. Hence, our goal is to show: $\|\phi_i^{\text{dis}}\| \rightarrow 0$ a.s., so that $\phi_i \rightarrow \phi_i^{\text{agr}}$. We build a recursion on ϕ_i^{dis} by multiplying both sides of (64) with \mathbf{P}^{dis} :

$$\phi_{i+1}^{\text{dis}} = \mathbf{P}^{\text{dis}} \mathcal{C} (\phi_i^{\text{dis}} - \alpha_{i+1} \mathbf{P}^{\text{dis}} (\mathcal{G}(\phi_i) + \mathbf{F}_{i+1})), \quad (\text{C8})$$

where we used: $\mathbf{P}^{\text{dis}}\mathcal{C} = \mathbf{P}^{\text{dis}}\mathcal{C}\mathbf{P}^{\text{dis}}$. Now, we bound $\|\phi_{i+1}^{\text{dis}}\|$:

$$\begin{aligned} \|\phi_{i+1}^{\text{dis}}\| &= \|\mathbf{P}^{\text{dis}}\mathcal{C}(\phi_i^{\text{dis}} - \alpha_{i+1}\mathbf{P}^{\text{dis}}(\mathcal{G}(\phi_i) + \mathbf{F}_{i+1}))\| \\ &\leq \|\mathbf{P}^{\text{dis}}\mathcal{C}\| \|\phi_i^{\text{dis}} - \alpha_{i+1}\mathbf{P}^{\text{dis}}(\mathcal{G}(\phi_i) + \mathbf{F}_{i+1})\| \\ &< \|\phi_i^{\text{dis}} - \alpha_{i+1}\mathbf{P}^{\text{dis}}(\mathcal{G}(\phi_i) + \mathbf{F}_{i+1})\| \\ &\leq \|\phi_i^{\text{dis}}\| + \alpha_{i+1} \|\mathbf{P}^{\text{dis}}(\mathcal{G}(\phi_i) + \mathbf{F}_{i+1})\|, \end{aligned} \quad (\text{C9})$$

where the first inequality is due to the submultiplicative property of the norm; the second inequality is due to Lemma 1; and the third one is the Minkowski's inequality. Rearranging terms and squaring and taking the limit on both sides we have:

$$\begin{aligned} \lim_{i \rightarrow \infty} (\|\phi_{i+1}^{\text{dis}}\| - \|\phi_i^{\text{dis}}\|)^2 &\leq \lim_{i \rightarrow \infty} \alpha_{i+1}^2 \|\mathbf{P}^{\text{dis}}(\mathcal{G}(\phi_i) + \mathbf{F}_{i+1})\|^2 \\ &\leq \lim_{i \rightarrow \infty} \alpha_{i+1}^2 \|\mathbf{P}^{\text{dis}}\|^2 (\|\mathcal{G}(\phi_i)\|^2 + \|\mathbf{F}_{i+1}\|^2) \\ &= 0 \quad a.s., \end{aligned} \quad (\text{C10})$$

where we used the submultiplicative property of the norm and Minkowski's inequality for the second inequality; the bounded norm of the projector; Assumptions 2, 4, and 5 to bound the norms of the gradient and the noise; and Assumption 3 to converge to zero and average the noise.

We conclude that the sequence $\{\|\phi_i^{\text{dis}}\|\}_{i=0}^{\infty}$ converges *a.s.* Let us introduce a shorthand for this limit: $\|\phi_{\infty}^{\text{dis}}\| \triangleq \lim_{i \rightarrow \infty} \|\phi_i^{\text{dis}}\|$. Squaring the second line of (C9), taking the limit, applying Minkowski's inequality, and rearranging terms yields:

$$\|\phi_{\infty}^{\text{dis}}\|^2 (1 - \|\mathbf{P}^{\text{dis}}\mathcal{C}\|^2) \leq \lim_{i \rightarrow \infty} \alpha_{i+1} \|\mathbf{P}^{\text{dis}}\mathcal{C}\|^2 \|\mathbf{P}^{\text{dis}}(\mathcal{G}(\phi_i) + \mathbf{F}_{i+1})\|^2 \quad a.s. \quad (\text{C11})$$

From Lemma 1 and using the same arguments as in (68), we conclude: $\|\phi_{\infty}^{\text{dis}}\| = 0$ *a.s.* \square

C.3 Proof of Theorem 3

Let's build a recursion on ϕ_i^{agr} by multiplying (64) by \mathbf{P}^{agr} :

$$\phi_{i+1}^{\text{agr}} = \phi_i^{\text{agr}} - \alpha_{i+1} (\mathbf{P}^{\text{agr}}\mathcal{G}(\phi_i) + \mathbf{P}^{\text{agr}}\mathbf{F}_{i+1}), \quad (\text{C12})$$

where we used $\mathbf{P}^{\text{agr}}\mathcal{C} = \mathbf{P}^{\text{agr}}$. We will use the following notation for the average difference between the average of the gradient of each agent's parameter and the gradient of the average parameter:

$$\mathbf{r}_i \triangleq \mathbf{P}^{\text{agr}} (\mathcal{G}(\phi_i) - \mathcal{G}(\phi_i^{\text{agr}})). \quad (\text{C13})$$

Further, we define the average of the blocks inside \mathbf{r}_i , which are all equal:

$$\bar{\mathbf{r}}_i \triangleq \frac{1}{N} \mathbb{1}_N^{\top} \otimes I_M \mathbf{r}_i. \quad (\text{C14})$$

Note that \mathbf{r}_i and $\bar{\mathbf{r}}_i$ are random vectors of length MN and M , respectively. Finally, we introduce the average Martingale noise:

$$\bar{\mathbf{F}}_{i+1} \triangleq \frac{1}{N} \sum_{k=1}^N \mathbf{F}_{k,i+1}. \quad (\text{C15})$$

Rearranging terms in (C12), and from (C7), we see that convergence of (C12) is equivalent to convergence of:

$$\bar{\phi}_{i+1} = \bar{\phi}_i - \alpha_{i+1} (\bar{\mathcal{G}}(\bar{\phi}_i) + \bar{\mathbf{r}}_i + \bar{\mathbf{F}}_{i+1}), \quad (\text{C16})$$

where \bar{g} is given by (53). From Assumptions 2, 3, and 5, we conclude that: $\lim_{i \rightarrow \infty} \alpha_{i+1} \bar{F}_i = 0$ a.s. Thus, the convergence of (C6) is equivalent to the convergence of this average recursion:

$$\bar{\phi}_{i+1} = \bar{\phi}_i - \alpha_{i+1} \left(\bar{g}(\bar{\phi}_i) + \bar{F}_{i+1} \right). \quad (\text{C17})$$

At this point, we can rely on Assumptions 2–8 to use the ODE method for the analysis of two timescales stochastic approximations, in particular Theorem 1.1—Remark 2 from Borkar (1997) concludes the proof (see also Borkar, 2008, Chapter 6). \square

Appendix D Diff-DAC pseudocode

Algorithm 1 Diff-DAC. This algorithm generalises the Diff-DAC architecture and therefore has Diff-SiAC and Diff-A2C as variants. The algorithm runs in parallel at every agent $k = 1, \dots, N$.

Input: Training termination condition (`terminate_training`), stop sampling condition (`stop_sampling`), maximum number of steps per update T , learning rate sequences $\{\alpha_i, \beta_i\}_{i=1}^{\infty}$.

- 1: Initialise critic, $v_{\xi_{k,0}}$, and actor, $\pi_{w_{k,0}}$, networks, $\forall k \in \mathcal{N}$.
- 2: Initialise iteration counter, $i = 0$.
- 3: **while** not `terminate_training`:
- 4: Initialise empty trajectory, $\mathbb{M}_k = \{\}$.
- 5: Initialise step counter: $t = 0$.
- 6: **while** $t < T$ and not `stop_sampling`:
- 7: Observe environment state, $s_{k,t}$.
- 8: Select action $a_{k,t} \sim \pi_{w_{k,t}}(\cdot | s_{k,t})$.
- 9: Execute $a_{k,t}$ and observe $r_{k,t+1}$ and $s_{k,t+1}$.
- 10: Store tuple $(s_{k,t}, a_{k,t}, r_{k,t+1}, s_{k,t+1})$ in \mathbb{M}_k .
- 11: Update step counter: $t \leftarrow t + 1$.
- 12: **end while**
- 13: **for** each sample $t \in \mathbb{M}_k$:
- 14: Compute advantage function (see Equations (40) and (45))
- 15: **end for**
- 16: Compute distributed critic update:

$$\widehat{\xi}_{k,i+1} = \text{SGD_optimizer} \left(\alpha_{i+1}, \xi_{k,i}, \frac{1}{|\mathbb{M}_k| - 1} \sum_{t=0}^{|\mathbb{M}_k|} \nabla_{\xi} v_{\xi_{k,i}}(s_{k,t}) \widehat{A}_{k,t} \right)$$

$$\xi_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \widehat{\xi}_{l,i+1}$$

- 17: Compute distributed actor update:

$$\widehat{w}_{k,i+1} = \text{SGD_optimizer} \left(\beta_{i+1}, w_{k,i}, \frac{1}{|\mathbb{M}_k|} \sum_{t=0}^{|\mathbb{M}_k|-1} \nabla_w \log \pi_{w_{k,i}}(a_{k,t} | s_{k,t}) \widehat{A}_{k,t} \right)$$

$$w_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \widehat{w}_{l,i+1}$$

- 18: Update iteration counter: $i \leftarrow i + 1$.
- 19: **end while**

Return: Critic and actor weights: $\xi_{k,\text{end}}, w_{k,\text{end}}$.

Algorithm 1 shows the pseudocode for the Diff-DAC architecture, which can be adapted to multiple actor-critic algorithms, including Diff-SiAC and Diff-A2C amongst others. The choice of advantage

function, the termination conditions (for training and sampling) and the optimiser differentiates our implementations of Diff-SiAC and Diff-A2C.

Under Diff-SiAC the advantage function in (40) is used, while Diff-A2C uses the more advanced advantage estimator in (45).

Diff-SiAC also terminates training after a fixed number of episodes and updates the parameters at the end of each episode. In this case, the termination condition is based on an episode count. Our implementation of Diff-A2C, however, terminates training and sampling based on the number of steps taken in the environment.

At each iteration, Diff-SiAC resets the environment, takes samples until the end of the episode or when reaching a maximum number of environment steps, and then updates parameters. On the other hand, Diff-A2C updates at a fixed interval in terms of environment steps and may, therefore, encounter terminal states during data collection. In such cases, data collection continues while marking the terminal states in the trajectory \mathbb{M}_k to aid in calculating the advantage estimates. We note that, in the case of an environment with fixed episode length, it is trivial to fix the termination and update interval conditions to be equivalent for Diff-SiAC and Diff-A2C.

`SGD_optimizer` is a placeholder for different stochastic gradient descent variants, that take the learning rate, current parameter value, and gradient estimate and output the updated parameter. Diff-SiAC uses Adam (Kingma & Ba, 2015), while Diff-A2C uses RMSProp (Tieleman & Hinton, 2012). We note that the results do not change significantly when using Adam for Diff-A2C.

The implementation of Diff-SiAC uses a single process, so all agents update synchronously. The implementation of Diff-A2C however exploits multiprocessing. We have performed all simulations on a single machine with a single GPU. Due to the multiprocessing, each Diff-A2C agent performs its local update asynchronously. Hence, some of its neighbours might not be ready when an agent begins the diffusion update. In this scenario, we allow each agent to perform its diffusion update with the most recently received values for the neighbours' parameters, provided that no value for a neighbour's parameters is more than five iterations old. Where this staleness limit is reached, an agent waits for a set of up-to-date values from its lagging neighbour(s). In practice, this limit of five update iterations of lag between neighbours is very rarely hit and our results are therefore not sensitive to this value. For the experiments exploring the sensitivity of Diff-A2C to noisy links in Section 6.4, we set this limit to zero so agents do not wait for their neighbours. Our implementation of Diff-A2C adapts the code put forward by Assran *et al.* (2019).