

Diff-DAC: Distributed Actor-Critic for Average Multitask Deep Reinforcement Learning

Sergio Valcarcel Macua* and Aleksi Tukiainen*

Daniel García-Ocaña Hernández† and David Baldazo†

Enrique Munoz de Cote* and Santiago Zazo†

*PROWLER.io

†Information Processing and Telecommunications Center - Universidad Politécnica de Madrid

Abstract

We propose a fully distributed actor-critic algorithm approximated by deep neural networks, named *Diff-DAC*, with application to single-task and to average multitask reinforcement learning (MRL). Tasks share state-action sets and reward function, but have different state transition distribution. Each agent has access to data from its local task only, but it aims to learn a policy that performs well on average for the whole set of tasks. During the learning process, agents communicate their value-policy parameters to their neighbors, diffusing the information across the network, so that they converge to a common policy, with no need for a central node. The method is scalable, since the computational and communication costs per agent grow with its number of neighbors. We derive Diff-DAC's 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 that approximates the solution of a linear program. Experiments suggest that Diff-DAC can outperform the single previous distributed MRL approach (i.e., Dist-MTLPS) and even the centralized architecture.¹

1 Introduction

Within a decade, billions of interconnected devices will be processing and exchanging data throughout the global economy (van der Meulen 2015). Centralised reinforcement learning (RL) architectures where all devices interact with a central station might be unfeasible. Fully distributed RL algorithms, where agents communicate only with neighbors and without central control, offer a solution to this problem, since the communication cost per agent scales linearly with its number of neighbors. In this distributed approach, each agent learns by interacting with its own environment, but is able to cooperate and benefit from the learning process of the whole network. When all agents' environments are equal, they learn to perform a single task; when environments are different but related, they learn to generalize across all tasks (Taylor and Stone 2009). The latter is known as the multitask reinforcement learning (MRL) problem. We propose an algorithm named *Diffusion-based Distributed Actor-Critic* (Diff-DAC) for both single-task and MRL problems.

Most previous MRL approaches assumed access to data from all tasks (Bou-Ammar et al. 2014; Parisotto, Ba, and

Salakhutdinov 2016; Teh et al. 2017). But if the number of tasks is large and their data are geographically distributed, the communication cost of transmitting the data to a central station might be prohibitive.

The idea of making scalable MRL with distributed optimization was first proposed by El Bsat, Bou-Ammar, and Taylor (2017) with the Dist-MTLPS method, which extended a distributed implementation of ADMM due to Wei and Ozdaglar (2012). Our work improves over Dist-MTLPS in a number of ways: *i*) Dist-MTLPS 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, uses deep learning architectures to avoid costly feature engineering, and is able to learn more expressive policies. *ii*) The distributed ADMM updates of the agents are done in sequential order, requiring finding a cyclic path that visits all agents, which is generally an NP-hard problem (Karp 1972). Diff-DAC uses a diffusion strategy (Sayed 2014), where each agent interacts with its neighbors, with no ordering, and possibly asynchronously (Zhao and Sayed 2015). *iii*) Sequential strategies are sensitive to agent or link failures, since they stop the information flow; while diffusion strategies are robust since the agents can still operate even if parts of the network become isolated.

As far as we know, all other previous works on distributed RL only considered tabular or linear function approximations (e.g., (Kar, Moura, and Poor 2013; Valcarcel Macua et al. 2015; Tutunov, Bou-Ammar, and Jadbabaie 2016)), and do not apply immediately to expressive nonlinear approximations. In particular, reference (Kar, Moura, and Poor 2013) added a consensus rule to tabular Q-learning; a principled nonlinear extension raises questions like whether we should we add consensus to the target network updates, and would be an alternative contribution to our actor-critic approach. The Dist-GTD method due to (Valcarcel Macua et al. 2015) is for policy evaluation with linear approximation, and its extension to control and nonlinear approximations isn't trivial even for the single-agent GTD. Finally, reference (Tutunov, Bou-Ammar, and Jadbabaie 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 related works suffer from similar drawbacks.

¹Presented at Adaptive Learning Agents workshop (ALA2018), July 14th, 2018, Stockholm, Sweden.

Contributions. (1) We propose a fully distributed actor-critic deep reinforcement learning algorithm named Diff-DAC for the single and average multitask problem that scales gracefully to large number of tasks. (2) We re-derive the actor-critic framework from duality theory and show that it is an instance of *dual-ascent* to approximate the saddle-point of the Lagrangian of a linear program (LP). This derivation formalizes previous intuitions (Pfau and Vinyals 2016) and provides novel insights, like a policy gradient that includes the advantage function explicitly, rather than as a variance reduction technique. (3) Experimental results suggest that Diff-DAC outperforms Dist-MTLPS, and that it is more stable and achieves better local optima than the centralized approach, without replay memory or target networks.

2 Problem Formulation

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

Consider a parametric family of MDPs defined over finite² state-action sets, \mathbb{S} and \mathbb{A} . Each MDP of the family has different state transition distribution, $\mathcal{P}_\theta(s'|s, a)$, that depend on some parameter $\theta \in \Theta$, where Θ is a measurable compact set. The task family is given as a probability distribution over the parameter set, f , so that the parameter is a random variable³: $\theta = \theta \sim f$. All MDPs in the family share the reward function, $r(s, a)$, and the distribution over the initial state, $\mu(s)$, $\forall s, s' \in \mathbb{S}, a \in \mathbb{A}$,

Let $\pi : \mathbb{S} \times \mathbb{A} \mapsto [0, 1]$ be a stationary policy, such that $\pi(a|s)$ denotes the probability of taking action a at state s . Let $v : \mathbb{S} \mapsto \mathbb{R}$ denote the value function, such that $v_\theta^\pi(s)$ is the value at state s when following policy π :

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

where $\mathbb{E}_{\pi, \mathcal{P}_\theta} [\cdot]$ is the expected value when $a_t \sim \pi(\cdot|s_t)$ and $s_{t+1} \sim \mathcal{P}_\theta(\cdot|s_t, a_t)$; and $0 < \gamma < 1$ is the discount factor. Introduce the vector of values: $v_\theta^\pi \triangleq (v_\theta^\pi(s))_{s \in \mathbb{S}} \in \mathbb{R}^{|\mathbb{S}|}$.

Suppose we have observed N tasks that correspond to parameters $\{\theta_k\}_{k=1}^N$. Our goal is to learn a stationary policy that maximizes the *global* average value:

$$\underset{\pi}{\text{maximize}} \quad \mu^\top \left(\sum_{k=1}^N v_{\theta_k}^\pi \right), \quad (2)$$

We assume bounded rewards:

$$|r(s, a)| \leq R_{\max} < \infty, \quad \forall (s, a) \in \mathbb{S} \times \mathbb{A}, \quad (3)$$

for some scalar R_{\max} . Under this assumption we can easily ensure existence of solution to (2).

When all task parameters $\{\theta_k\}_{k=1}^N$ are equal, (2) is the single-task RL problem; when they differ, (2) becomes an MRL problem where we aim to learn a single policy that performs optimally in average for the whole set of tasks.

²The proposed Diff-DAC algorithm uses function approximation so that is able to work in continuous state-action sets as well.

³We use boldface font to denote random variables and regular font to denote instances or deterministic variables.




Figure 1: Example of network and detailed neighborhood. Blue nodes represent agents, and edges represent their connectivity. This network consists of $N = 25$ agents, with average neighborhood size: $|\bar{\mathcal{N}}| = \frac{1}{N} \sum_{k=1}^N |\mathcal{N}_k| = 4.2$. On the right, the figure zooms over neighborhood \mathcal{N}_k (green area), where each agent k runs its own task instance of the swing-up cart-pole task (with different pole length and mass, defined by parameter θ_k). As explained later, agent k transmits its neural network weights, $\xi_{k,i}$ and $w_{k,i}$, to its neighbors j and l ; and it receives their weights $\xi_{j,i}$ and $\xi_{l,i}$, $w_{l,i}$, and combines them with coefficients c_{jk} and c_{lk} , respectively.

3 Networked Multiagent Setting

In this section we introduce the networked multiagent setting that learns in a fully distributed manner.

We have a network of N agents, which is expressed as a graph, \mathcal{N} . Each node, denoted $k = 1, \dots, N$, corresponds to an agent that learns from data coming from its own task⁴, with parameter $\theta_k \sim f$. The edges in the graph represent communication links. We assume that the graph is *connected* (i.e., there is at least one path between every pair of agents).

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 is only allowed to communicate within its own neighborhood, \mathcal{N}_k , which is defined as the set of agents to which it is directly connected, including k itself: $\mathcal{N}_k \triangleq \{l \in \{1, \dots, N\} : c_{lk} > 0\}$. In order to ensure that the information flows through the network, we require the following standard conditions on the connectivity matrix C , which make it doubly-stochastic and primitive (Sayed 2014; Valcarcel Macua et al. 2015):

$$C^\top \mathbf{1} = \mathbf{1}, \quad C \mathbf{1} = \mathbf{1}, \quad \text{and} \quad c_{lk} \geq 0, \quad k, l = 1, \dots, N, \quad (4)$$

$$\text{trace}[C] > 0, \quad (5)$$

where $\mathbf{1}$ is a vector of ones. Although conditions (4)–(5) seem restrictive, it turns out that 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 C satisfies the required conditions. One of such procedures is the Hastings rule (Zhao and Sayed 2012), (Sayed 2014, p.492).

⁴For simplicity, we assume that each agent is allocated with one task, similar to (El Bsat, Bou-Ammar, and Taylor 2017). The extension to multiple tasks per agent is trivial.

4 Multitask Actor-Critic from Duality Theory

In this section, we reformulate MRL as a linear program (LP), and show that by applying dual ascent to the Lagrangian, we get a tabular model-based actor-critic method that solves (2).

Introduce the average dynamics:

$$\bar{\mathcal{P}}(s'|s, a) \triangleq \frac{1}{N} \sum_{k=1}^N \mathcal{P}_{\theta_k}(s'|s, a). \quad (6)$$

The following lemma is key in our derivations, since it allow us to consider the multitask problem as a single MDP, with state-action transitions given $\bar{\mathcal{P}}$.

Lemma 1. $\bar{\mathcal{P}}$ is a row-stochastic matrix.

Proof. Stochastic matrices lie in a compact convex set (Horn and Johnson 1990, p.527), so that their convex combination lies in the same set (Boyd and Vandenberghe 2004, p.24). \square

Thus, we can use standard (i.e., single-task) optimal control results (Puterman 2005, pp. 143–151) for our setting. Let \mathbb{V} denote the set of bounded real functions on \mathbb{S} , with componentwise partial order and norm $\|v\| \triangleq \sup_{s \in \mathbb{S}} |v(s)|$.

Corollary 1. For any stationary policy, π , the Bellman equation for the new MDP defined by $\bar{\mathcal{P}}$ is given by:

$$v(s) = \sum_{a \in \mathbb{A}} \pi(a|s) \left(r(s, a) + \gamma \sum_{s' \in \mathbb{S}} \bar{\mathcal{P}}(s'|s, a) v(s') \right), \quad (7)$$

Introduce the multitask Bellman operator $\bar{T} : \mathbb{V} \mapsto \mathbb{V}$:

$$(\bar{T}v)(s) \triangleq \max_{a \in \mathbb{A}} \left[r(s, a) + \gamma \sum_{s' \in \mathbb{S}} \bar{\mathcal{P}}(s'|s, a) v(s') \right]. \quad (8)$$

Similar to single task optimal control theory (Puterman 2005, Sec. 9.1), we can reformulate (2) as an LP.

$$\begin{aligned} & \text{minimize } \mu^\top v \\ & v \in \mathbb{R}^{|\mathbb{S}|} \\ \text{s.t. } & v(s) \geq r(s, a) + \gamma \sum_{s' \in \mathbb{S}} \bar{\mathcal{P}}(s'|s, a) v(s') \\ & \forall (s, a) \in \mathbb{S} \times \mathbb{A}, \end{aligned} \quad (9)$$

Since problem (9) satisfies Slater condition, strong-duality holds (Boyd and Vandenberghe 2004, Sec. 5.2.3) and the primal and dual optimal values are attained and equal. The Lagrangian of (9) is given by:

$$\begin{aligned} L(v, d) = & \mu^\top v + \sum_{(s, a) \in \mathbb{S} \times \mathbb{A}} d(s, a) \left(r(s, a) \right. \\ & \left. + \gamma \sum_{s' \in \mathbb{S}} \bar{\mathcal{P}}(s'|s, a) v(s') - v(s) \right), \end{aligned} \quad (10)$$

where the dual variable $d \triangleq (d(s, a))_{(s, a) \in \mathbb{S} \times \mathbb{A}} \in \mathbb{R}^{|\mathbb{S}| |\mathbb{A}|}$ is a nonnegative vector of length $|\mathbb{S}| |\mathbb{A}|$. Let v^* denote the optimal primal variable. Let d^* denote an optimal dual variable,

which might not be unique. The saddle point condition of (10) is given by:

$$\min_v \max_d L(v, d) = L(v^*, d^*) = \max_d \min_v L(v, d). \quad (11)$$

There are multiple approaches to find a saddle point. We focus on the *dual-ascent* scheme (Arrow, Hurwicz, and Uzawa 1958), which consists in alternating between: *i*) Finding a primal solution, given the dual variable; and *ii*) ascending in the direction of $\nabla_d L(v, d)$, given the primal variable.

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

$$v \leftarrow \arg \min_{v \in \mathbb{R}^{|\mathbb{S}|}} L(v, d). \quad (12)$$

Since problem (9) is linear, if we derive the KKT conditions, we can see that the first-order condition holds $\forall v \in \mathbb{R}^{|\mathbb{S}|}$. Thus, the only condition that depends on v is *complementary slackness*:

$$\begin{aligned} & \sum_{(s, a) \in \mathbb{S} \times \mathbb{A}} d(s, a) \left(r(s, a) + \gamma \sum_{s' \in \mathbb{S}} \bar{\mathcal{P}}(s'|s, a) v(s') - v(s) \right) \\ & = 0. \end{aligned} \quad (13)$$

Similar to the standard single-task problem (Puterman 2005, Sec. 6.9), it can be shown that our dual variable is the discounted steady-state state-action visitation measure, so that we can obtain the policy from d :

$$\pi(a|s) = \frac{d(a, s)}{\rho(s)}, \quad \text{where } \rho(s) \triangleq \sum_{a' \in \mathbb{A}} d(a', s). \quad (14)$$

Hence, the Bellman equation (7), typically used to derive the *critic* in actor-critic methods, is sufficient to guarantee (13).

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

$$d \leftarrow [d + \alpha \nabla_d L(v, d)]^+, \quad (15)$$

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

$$\nabla_d L(v, d) = \left(\frac{\partial L(v, d)}{\partial d(s, a)} \right)_{(s, a) \in \mathbb{S} \times \mathbb{A}}. \quad (16)$$

Interestingly, note that the partial derivatives of the Lagrangian in (16) are indeed the so named *advantage function* extended to our multitask problem:

$$\begin{aligned} A(s, a) & \triangleq r(s, a) + \gamma \sum_{s' \in \mathbb{S}} \bar{\mathcal{P}}(s'|s, a) v(s') - v(s) \\ & = \frac{\partial L(v, d)}{\partial d(s, a)}. \end{aligned} \quad (17)$$

If we learn d^* , we can use (14) to obtain π^* , so that recursion (15) can be seen as an *actor* update. Thus, (7) and (15) define a novel *tabular model-based actor-critic* method. In the following, we extend this approach to a model-free distributed actor-critic method with neural network approximations.

5 Distributed Deep Actor-Critic

In order to use *diffusion* strategies (Sayed 2014) to derive a distributed optimization method, we have to express the global objective as a convex combination of each agent’s local objective. Thus every agent can optimize its objective from local data; and by communicating with their neighbors, all agents converge to a common solution that optimizes the global objective. We do this for both critic and actor.

5.1 Distributed policy evaluation: Critic

When computing the critic for large (or continuous) state-action sets, it is common to approximate the value function with some parametric function $v_\xi(s) \approx v(s)$, where $\xi \in \mathbb{R}^{M_v}$ denotes the parameter vector of length M_v . We choose neural networks with multiple hidden layers (i.e., deep learning) as parametric approximators. Hence, we can learn the network weights, ξ , by transforming (7) into a non-linear regression problem:

$$\underset{\xi \in \mathbb{R}^{M_v}}{\text{minimize}} \quad J(\xi) \triangleq \mathbb{E} \left[(v_\xi(\mathbf{s}_t) - \bar{\mathbf{y}}_t)^2 \right], \quad (18)$$

where the target values are given by:

$$\bar{\mathbf{y}}_t \triangleq r(\mathbf{s}_t, \mathbf{a}_t) + \gamma \sum_{s' \in \mathcal{S}} \bar{\mathcal{P}}(s' | \mathbf{s}_t, \mathbf{a}_t) v_\xi(s'). \quad (19)$$

In order to derive a diffusion-based distributed critic, we have to reformulate the problem as minimizing the convex combination of costs that depend only on a single task each. The cost for each individual task takes the form:

$$\tilde{J}_k(\xi) \triangleq \mathbb{E} \left[(v_\xi(\mathbf{s}_t) - \mathbf{y}_{k,t})^2 \right], \quad k = 1, \dots, N, \quad (20)$$

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

$$\mathbf{y}_{k,t} = r_{\theta_k}(\mathbf{s}_t, \mathbf{a}_t) + \gamma \sum_{s' \in \mathcal{S}} \mathcal{P}_{\theta_k}(s' | \mathbf{s}_t, \mathbf{a}_t) v_\xi(s'),$$

such that $\bar{\mathbf{y}}_t = 1/N \sum_{k=1}^N \mathbf{y}_{k,t}$. Now, in order to obtain a cost that is a combination of the individual costs, 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:

$$\begin{aligned} \tilde{J}(\xi) &\triangleq \frac{1}{N} \sum_{k=1}^N \tilde{J}_k(\xi) = \frac{1}{N} \sum_{k=1}^N \mathbb{E} \left[(v_\xi(\mathbf{s}_t) - \mathbf{y}_{k,t})^2 \right] \\ &\geq \mathbb{E} \left[\left(\frac{1}{N} \sum_{k=1}^N (v_\xi(\mathbf{s}_t) - \mathbf{y}_{k,t}) \right)^2 \right] = J(\xi). \end{aligned} \quad (21)$$

Now, we can apply *diffusion* stochastic-gradient-descent (SGD) strategies (Sayed 2014), which consist of two steps: *adaptation* and *combination*. During the *adaptation* step, each agent performs SGD on its individual cost, \tilde{J}_k , to obtain some intermediate-parameter update. Then, each agent *combines* the intermediate-parameters from its neighbors. These two steps are described by the following updates, which run in parallel for all agents $k = 1, \dots, N$:

$$\hat{\xi}_{k,i+1} = \xi_{k,i} - \alpha_{i+1} \widehat{\nabla}_\xi \tilde{J}_k(\xi_{k,i}), \quad (22a)$$

$$\xi_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \hat{\xi}_{l,i+1}, \quad (22b)$$

where i is the iteration index; α_i is the step-size; and $\widehat{\nabla}_\xi \tilde{J}_k(\xi_{k,i})$ is the stochastic gradient evaluated at $\xi_{k,i}$, estimated from samples $\{(s_{k,t}, a_{k,t}, r_{k,t+1}, s_{k,t+1})\}_{t=0}^{T_{k,i}}$ of the i -th episode, of length $T_{k,i}$, gathered by the k -th agent. We use Monte Carlo estimates for the target $y_{k,t} = \sum_{j=t}^{T_{k,i}} \gamma^{j-t} r_{k,j+1}$ (for simplicity), where $r_{k,j+1} \triangleq r_{\theta_k}(s_{k,j}, a_{k,j})$ is a shorthand. Then, the stochastic gradient is given by:

$$\widehat{\nabla}_\xi \tilde{J}_k(\xi_{k,i}) = \frac{1}{T_{k,i}} \sum_{t=0}^{T_{k,i}} \nabla_\xi v_{\xi_{k,i}}(s_{k,t}) (v_{\xi_{k,i}}(s_{k,t}) - y_{k,t}). \quad (23)$$

We remark that each agent learns from its current episode, without replay buffer, similar to A3C (Mnih et al. 2016), but in a fully distributed fashion, as opposed as having multiple threads updating the same neural network at a single location.

5.2 Distributed policy gradient: Actor

For large state-action sets, it is convenient to approximate the policy with a parametric function. Again, we consider expressive deep neural networks for the policy. From (14), we can rewrite the Lagrangian as:

$$L(v, \pi, \rho) = \mu^\top v + \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \pi(a|s) \rho(s) A(s, a). \quad (24)$$

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 (24), we obtain an approximate Lagrangian, $\tilde{L}(v, w, \rho) \approx L(v, \pi, \rho)$, of the form:

$$\tilde{L}(v, w, \rho) = \mu^\top v + \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \pi_w(a|s) \rho(s) A(s, a). \quad (25)$$

Thus, in order to approximate the saddle point condition (11), we can move in the ascent direction of the gradient of (25) w.r.t. the policy parameter, which is given by:

$$\begin{aligned} \nabla_w \tilde{L}(v, w, \rho) &= \nabla_{\pi_w} \tilde{L}(v, w, \rho) \nabla_w \pi_w \\ &= (\nabla_w \pi_w(a_1 | s_1), \dots, \nabla_w \pi_w(a_{|\mathcal{A}|} | s_{|\mathcal{S}|}))^\top \\ &\quad \left(\frac{\partial \tilde{L}(v, w, \rho)}{\partial \pi_w(a|s)} \right)_{(s,a) \in \mathcal{S} \times \mathcal{A}} \\ &= \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \nabla_w \pi(a|s) \frac{\partial L(v, w, \rho)}{\partial \pi_w(a|s)} \\ &= \sum_{s \in \mathcal{S}} \rho(s) \sum_{a \in \mathcal{A}} \nabla_w \pi_w(a|s) A(s, a) \\ &= \sum_{s \in \mathcal{S}} \rho(s) \sum_{a \in \mathcal{A}} \pi_w(a|s) \nabla_w \log \pi_w(a|s) A(s, a), \end{aligned} \quad (26)$$

where we used: $\nabla_w \pi_w(a|s) = \pi_w(a|s) \nabla_w \log \pi_w(a|s)$.

Interestingly, (26) 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 *baseline* mainly as a variance reduction technique (Williams 1992; Bhatnagar et al. 2009).

In order to derive a fully *distributed* actor, let us write the multitask advantage function (17) as the convex combination of advantage functions for the individual tasks:

$$A(s, a) = \frac{1}{N} \sum_{k=1}^N A_k(s, a), \quad (27)$$

$$A_k(s, a) \triangleq r_{\theta_k}(s, a) + \gamma \sum_{s' \in \mathcal{S}} \mathcal{P}_{\theta_k}(s'|s, a)v(s') - v(s). \quad (28)$$

Hence, we write the approximate Lagrangian for each task:

$$\tilde{L}_k(v, w, \rho) \triangleq \mu_{\theta_k}^\top v + \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \pi_w(a|s)\rho(s)A_k(s, a), \quad (29)$$

such that $\tilde{L}(v, w, \rho) = \frac{1}{N} \sum_{k=1}^N \tilde{L}_k(v, w, \rho)$.

Similar to the critic, once we have expressed the multitask approximate Lagrangian as the convex combination of the approximate Lagrangian of each individual task, we can apply diffusion SGD to perform the actor update, with smaller step-size, $\beta_{i+1} \leq \alpha_{i+1}$, to approximate convergence of the critic at every actor update:

$$\hat{w}_{k,i+1} = w_{k,i} + \beta_{i+1} \hat{\nabla}_w \tilde{L}_k(v_{\xi_{k,i}}, w_{k,i}, \rho), \quad (30a)$$

$$w_{k,i+1} = \sum_{l \in \mathcal{N}_k} c_{lk} \hat{w}_{l,i+1}, \quad (30b)$$

where each agent estimates its stochastic gradient as:

$$\hat{\nabla}_w \tilde{L}_k(v_{\xi_{k,i}}, w, \rho) = \frac{1}{T_{k,i}} \sum_{t=0}^{T_{k,i}} \nabla_w \log \pi_w(a_{k,t}|s_{k,t}) \hat{A}_{k,t}, \quad (31)$$

and $\hat{A}_{k,t}$ can be any approximation of the advantage function (Schulman et al. 2017). We use the simple estimate:

$$\hat{A}_{k,t} = \sum_{j=t}^{T_{k,i}} \gamma^{j-t} r_{k,j+1} - v_{\xi_{k,i}}(s_{k,t}). \quad (32)$$

Two remarks: *i*) In order to simplify the implementation, we set the target $y_{k,t}$ to be the empirical return, so that the stochastic gradient of the critic in (23) is the negative advantage estimate: $\hat{\nabla}_\xi \tilde{J}_k(\xi_{k,i}) = -\hat{A}_{k,t}$. *ii*) Note that replacing $\xi_{k,i}$ with $\xi_{k,i+1}$ in (30a)–(30b) and (31)–(32), in a Gauss-Seidel fashion, might lead to faster convergence.

A detailed description of Diff-DAC is given in Algorithm 1.

6 Numerical Experiments

We evaluate the performance of Diff-DAC on three MRL problems of varying levels of difficulty. We use $\gamma = 0.99$ for all tasks:

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 the pole is beyond 12 degrees from vertical, cart moves more than 2.4 units from the center, or run for 200 time-steps. The single task uses parameters (0.1, 0.5, 1.0) for the pole mass, pole half-length and cart mass, respectively. The MRL problem consists of 25 tasks: pole mass in $\{0.1, 0.325, 0.55, 0.775, 1\}$,

Algorithm 1: Diff-DAC. This algorithm runs in parallel at every agent $k = 1, \dots, N$.

Input: Maximum number of episodes E , maximum number of steps per episode T , learning rate sequences (α_i, β_i) .

- 1: Initialize critic, $v_{\xi_{k,0}}$, and actor, $\pi_{w_{k,0}}$, networks, $\forall k \in \mathcal{N}$.
- 2: Initialize episode counter, $i = 0$.
- 3: **while** $i < E$:
- 4: Initialize empty trajectory, $\mathbb{M}_k = \{\}$.
- 5: Initialize step counter: $t = 0$.
- 6: Observe $s_{k,0}$.
- 7: **while** $t < T$ and not terminal state:
- 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:
 $\hat{A}_{k,t} = \sum_{j=t}^{|\mathbb{M}_k|} \gamma^{j-t} r_{k,j+1} - v_{\xi_{k,i}}(s_{k,t})$
- 15: **end for**
- 16: Compute distributed critic gradient:

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

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

- 17: Compute distributed actor update:

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

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

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

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

pole length in $\{0.05, 0.1625, 0.275, 0.3875, 0.5\}$, and cart mass 1.

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

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 up 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 track center and upright position, and ψ is the pole angle. This is a much more difficult task than standard cart-pole and more difficult than the inverted pendulum due to more complex dynamics. The 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. The single task uses parameters $(0.5, 0.25, 0.5)$ for the pole mass, pole half-length and cart mass respectively.

We compare Diff-DAC with Dist-MTLPS for the MRL problem in the cart-pole balance environment. In particular, we compare against two variants of Dist-MTLPS, which consist in using two standard policy search methods, namely Reinforce (Williams 1992) and PoWER (Kober and Peters 2009), for solving the individual tasks. We only compare Diff-DAC with Dist-MTLPS in the cart-pole balance task, since the other two environments are uncontrollable with linear policies from raw data. Our goal is to compare the performance of the single but expressive neural network policy provided by Diff-DAC with the task-specialised but less expressive linear policies provided by Dist-MTLPS.

We also compare Diff-DAC with Cent-AC, which has only one agent (central coordinator) that gathers and process samples from all the tasks synchronously and has the same neural network architectures and hyperparameters as Diff-DAC. We remark that we use two versions of exactly the same vanilla actor-critic algorithm, where their only difference is whether there is a single agent with access to all the data (Cent-AC) or multiple networked agents with access to local datasets (Diff-DAC).

Although experimenting with more environments and benchmarking with more algorithms would be helpful, testing in these simple environments is already useful, since they illustrate the behavior of the algorithms, without having to handle complex neural network architectures.

The network consists of $N = 25$ agents, randomly deployed in a 2D world, with average degree $|\bar{\mathcal{N}}| \triangleq \sum_{k=1}^N |\mathcal{N}_k| \approx 4.2$ (connectivity is determined by the distance between agents). For Figure 5, we also include two additional networks of $N = 25$ and $|\bar{\mathcal{N}}| = 7.4$, and $N = 100$ and $|\bar{\mathcal{N}}| = 20$. Matrix C was obtained with the Hastings-rule (Sayed 2014, p.492) in all cases, so that (4)–(5) hold. We remark that the network topologies used in the experiments are not related to any form of task similarity, but they just reflect the sparse connectivity that appears naturally when agents and data are geographically distributed.

The critic and actor neural networks 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 also included an extra penalty in the loss function equal to the entropy of the policy, with penalty coefficient 0.0005. Thus, both the mean and the variance are learned for the policy. We use ADAM optimizer (Kingma and Ba 2015), with learning rate 0.01 for critic and 0.001

for actor. Diff-DAC performs a learning step ($i \leftarrow i + 1$) every fifth episode.

The return of the tasks is reported as the (undiscounted) total rewards every 20 episodes and is averaged over 10 test episodes at each point. Figures show the median and first and third quartiles of the distribution of the average return of all the tasks. Each epoch consists of 5 episodes per epoch and per agent in Diff-DAC, and $5N$ episodes in total per epoch for Dist-MTLPS and Cent-AC, so that the three algorithms simulate the same number of episodes. Every experiment was repeated at least 6 times.

In Figure 2 (bottom), we observe that Diff-DAC learns faster than Dist-MTLPS Reinforce and reaches better asymptotic performance. Dist-MTLPS PoWER converges faster than Diff-DAC, however the asymptotic performance of the latter is much better. This is remarkable since Dist-MTLPS learns one different policy per task, while Diff-DAC learns a single policy common to all tasks.

We also observe that Diff-DAC converges slower than the Cent-AC, which was expected since the latter can compute the gradients with data from all tasks at every iteration, while the former has to wait until the parameters are diffused across the network. However, Diff-DAC usually achieves higher asymptotic performance and less variance, in both single task and multitask problems, which was also expected due to the already reported enhanced robustness against local optima of diffusion strategies for nonconvex optimization problems (Valcarcel Macua 2017, Ch. 4). This is shown in Figure 2 (top), where Cent-AC tends to reach the optimal faster, but is unstable. We guess that this may be alleviated by adding a *replay buffer* (Mnih et al. 2013; Lillicrap et al. 2015) or randomizing agents’ samples to reduce their correlation, simulating asynchronicity (similar to asynchronous methods like A3C (Mnih et al. 2016)). However, our goal with these experiments is not to compare with SOTA centralized algorithms—which use several advancements to stabilise or improve performance—but to evaluate whether diffusion is not only a feasible distributed strategy but also a valid alternative to stabilize learning. Thus, we decided to use vanilla central actor-critic vs. vanilla distributed actor-critic, where their only difference is having a single agent with all the data vs. having multiple agents with local datasets.

Finally, Figure 5 shows a simple experiment that studies the influence of the network topology. We evaluate Diff-DAC for the single-task cart-pole balance problem and see that for the same network size, $N = 25$, a relatively sparse network, $|\bar{\mathcal{N}}| \approx N/6$, achieves performance similar to a more dense network, $|\bar{\mathcal{N}}| \approx N/3$. In addition, we see that larger number of agents $N = 100$ improves the asymptotic performance.

7 Conclusions

We considered MRL where tasks are parametrized MDPs with parameters drawn from some distribution, and we derived an algorithm that learns a policy that performs well on average for the observed set of tasks. We defined average global variables that allowed us to use standard optimal




Figure 2: Cart-pole balance with continuous action for single-task (top) and multitask (bottom). Cent-AC is faster than distributed approaches, but Diff-DAC achieves the best asymptotic performance.




Figure 4: Cart-pole swing-up for single-task (top) and multitask (bottom). Diff-DAC learns to swing-up and balance the pole consistently, while Cent-AC achieves much inferior performance.




Figure 3: Inverted pendulum for single-task (top) and multitask (bottom). Diff-DAC learns in both the single-task and multitask robustly. The central method learns the task quickly but is unstable.

control theory and reformulate our MRL problem as an LP. From this LP, we derived an exact, model-based actor-critic algorithm as an instance of dual ascent for finding the sad-




Figure 5: Influence of network topology in single-task cart-pole balance with continuous action. Diff-DAC combines the experience of all agents, relatively insensitive to network sparsity.

dle point of the Lagrangian. This saddle-point derivation is interesting in itself and provides novel insights in the actor-critic framework. By approximating the exact method with deep neural networks, we obtained the Diff-DAC algorithm, which can scale to large number of tasks.

Simulation results showed that Diff-DAC can be faster and achieve higher asymptotic performance than the state of the art distributed algorithm for solving the MRL problem (i.e., Dist-MTLPS). This is a remarkable result since the Diff-DAC agents converge to a single common policy that behaves better than the task-dependent linear policies obtained by Dist-MTLPS. Moreover, Diff-DAC can solve complex problems that are uncontrollable from raw

data by linear policies, while Dist-MTLPS requires (usually costly) feature engineering. Diff-DAC is also very stable and achieves similar or usually higher asymptotic performance than the centralized approach in both single and multitask problems. This suggests that the sparse connectivity among agents induces a regularization effect that helps them to achieve better local optimum. We consider this form of regularization an interesting line of research.

Diff-DAC can be also extended to *zero-shot learning* by taking the task parameter θ_k as an additional input to each agent's value/policy networks, so that when a new task appears, it can input its parameter in the neural networks and adapt its behavior to this task without further training.

Finally, it would be interesting to apply the proposed framework to derive distributed variants of other algorithms like PPO (Schulman et al. 2017).

8 Acknowledgements

We thank Haitham Bou-Ammar and Peter Vrancx for insightful discussions.

References

- Arrow, K. J.; Hurwicz, L.; and Uzawa, H. 1958. *Studies in Linear and Non-linear Programming*. Stanford University Press.
- Bhatnagar, S.; Sutton, R. S.; Ghavamzadeh, M.; and Lee, M. 2009. Natural actor-critic algorithms. *Automatica* 45(11):2471–2482.
- Bou-Ammar, H.; Eaton, E.; Ruvolo, P.; and Taylor, M. 2014. Online multi-task learning for policy gradient methods. In *Proc. Int. Conf. on Machine Learning (ICML)*, 1206–1214.
- Boyd, S., and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge University Press.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym.
- El Bsati, S.; Bou-Ammar, H.; and Taylor, M. E. 2017. Scalable multitask policy gradient reinforcement learning. In *AAAI Conf. on Artificial Intelligence (AAAI)*, 1847–1853.
- Horn, R., and Johnson, C. 1990. *Matrix Analysis*. Cambridge University Press.
- Kar, S.; Moura, J. M. F.; and 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., and Ba, J. L. 2015. Adam: A method for stochastic optimization. In *Proc. Int. Conf. on Learning Representations (ICLR)*.
- Kober, J., and Peters, J. R. 2009. Policy search for motor primitives in robotics. In *Advances in Neural Information Processing Systems (NIPS)*, 849–856.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 1928–1937.
- Parisotto, E.; Ba, J. L.; and Salakhutdinov, R. 2016. Actor-mimic: Deep multitask and transfer reinforcement learning. In *Proc. Int. Conf. on Learning Representations (ICLR)*.
- Pfau, D., and Vinyals, O. 2016. Connecting generative adversarial networks and actor-critic methods. *arXiv preprint arXiv:1610.01945*.
- Puterman, M. L. 2005. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2nd edition.
- Sayed, A. H. 2014. Adaptation, learning, and optimization over networks. *Foundations and Trends in Machine Learning* 7(4-5):311–801.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 1999. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems (NIPS)*, 1057–1063.
- Taylor, M. E., and Stone, P. 2009. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research* 10(Jul):1633–1685.
- Teh, Y. W.; Bapst, V.; Czarnecki, W. M.; Quan, J.; Kirkpatrick, J.; Hadsell, R.; Heess, N.; and Pascanu, R. 2017. Distal: Robust multitask reinforcement learning. *arXiv preprint arXiv:1707.04175*.
- Tutunov, R.; Bou-Ammar, H.; and Jadbabaie, A. 2016. An exact distributed newton method for reinforcement learning. In *IEEE Conf. on Decision and Control (CDC)*, 1003–1008.
- Valcarcel Macua, S.; Chen, J.; Zazo, S.; and Sayed, A. H. 2015. Distributed policy evaluation under multiple behavior strategies. *IEEE Transactions on Automatic Control* 60(5):1260–1274.
- Valcarcel Macua, S. 2017. *Distributed optimization, control and learning in multiagent networks*. Ph.D. Dissertation, Universidad Politécnica de Madrid.
- 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>.
- Wei, E., and Ozdaglar, A. 2012. Distributed alternating direction method of multipliers. In *IEEE Annual Conf. Decision and Control (CDC)*, 5445–5450.

Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4):229–256.

Zhao, X., and Sayed, A. H. 2012. Performance limits for distributed estimation over LMS adaptive networks. *IEEE Transactions on Signal Processing* 60(10):5107–5124.

Zhao, X., and 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.