

# GUIDE ACTOR-CRITIC FOR CONTINUOUS CONTROL

**Voot Tangkaratt**

RIKEN AIP, Tokyo, Japan  
voot.tangkaratt@riken.jp

**Abbas Abdolmaleki**

The University of Aveiro, Aveiro, Portugal  
abbas.a@ua.pt

**Masashi Sugiyama**

RIKEN AIP, Tokyo, Japan  
The University of Tokyo, Tokyo, Japan  
masashi.sugiyama@riken.jp

## ABSTRACT

Actor-critic methods solve reinforcement learning problems by updating a parameterized policy known as an actor in a direction that increases an estimate of the expected return known as a critic. However, existing actor-critic methods only use values or gradients of the critic to update the policy parameter. In this paper, we propose a novel actor-critic method called the guide actor-critic (GAC). GAC firstly learns a guide actor that locally maximizes the critic and then it updates the policy parameter based on the guide actor by supervised learning. Our main theoretical contributions are two folds. First, we show that GAC updates the guide actor by performing second-order optimization in the action space where the curvature matrix is based on the Hessians of the critic. Second, we show that the deterministic policy gradient method is a special case of GAC when the Hessians are ignored. Through experiments, we show that our method is a promising reinforcement learning method for continuous controls.

## 1 INTRODUCTION

The goal of reinforcement learning (RL) is to learn an optimal policy that lets an agent achieve the maximum cumulative rewards known as the return (Sutton & Barto, 1998). Reinforcement learning has been shown to be effective in solving challenging artificial intelligence tasks such as playing games (Mnih et al., 2015; Silver et al., 2016) and controlling robots (Deisenroth et al., 2013; Levine et al., 2016).

Reinforcement learning methods can be classified into three categories: *value-based*, *policy-based*, and *actor-critic* methods. Value-based methods learn an optimal policy by firstly learning a *value function* that estimates the expected return. Then, they infer an optimal policy by choosing an action that maximizes the learned value function. Choosing an action in this way requires solving a maximization problem which is not trivial for continuous controls. While extensions to continuous controls were considered recently, they are restrictive since specific structures of the value function are assumed (Gu et al., 2016; Amos et al., 2017).

On the other hand, policy-based methods, also called policy search methods (Deisenroth et al., 2013), learn a parameterized policy maximizing a sample approximation of the expected return without learning the value function. For instance, policy gradient methods such as REINFORCE (Williams, 1992) use gradient ascent to update the policy parameter so that the probability of observing high sample returns increases. Compared with value-based methods, policy search methods are simpler and naturally applicable to continuous problems. Moreover, the sample return is an unbiased estimator of the expected return and methods such as policy gradients are guaranteed to converge to a locally optimal policy under standard regularity conditions (Sutton et al., 1999). However, sample returns usually have high variance and this makes such policy search methods converge too slowly.

Actor-critic methods combine the advantages of value-based and policy search methods. In these methods, the parameterized policy is called an actor and the learned value-function is called a critic.

The goal of these methods is to learn an actor that maximizes the critic. Since the critic is a low variance estimator of the expected return, these methods often converge much faster than policy search methods. Prominent examples of these methods are *actor-critic* (Sutton et al., 1999; Konda & Tsitsiklis, 2003), *natural actor-critic* (Peters & Schaal, 2008), *trust-region policy optimization* (Schulman et al., 2015a), and *asynchronous advantage actor-critic* (Mnih et al., 2016). While their approaches to learn the actor are different, they share a common property that they only use the *value* of the critic, i.e., the zero-th order information, and ignore higher-order ones such as gradients and Hessians w.r.t. actions of the critic<sup>1</sup>. To the best of our knowledge, the only actor-critic methods that use gradients of the critic to update the actor are *deterministic policy gradients* (DPG) (Silver et al., 2014) and *stochastic value gradients* (Heess et al., 2015). However, these two methods do not utilize the second-order information of the critic.

In this paper, we argue that the second-order information of the critic is useful and should not be ignored. A motivating example can be seen by comparing gradient ascent to the Newton method: the Newton method which also uses the Hessian converges to a local optimum in a fewer iterations when compared to gradient ascent which only uses the gradient (Nocedal & Wright, 2006). This suggests that the Hessian of the critic can accelerate actor learning which leads to higher data efficiency. However, the computational complexity of second-order methods is at least quadratic in terms of the number of optimization variables. For this reason, applying second-order methods to optimize the parameterized actor directly is prohibitively expensive and impractical for deep reinforcement learning which represents the actor by deep neural networks.

Our contribution in this paper is a novel actor-critic method for continuous controls which we call *guide actor-critic* (GAC). Unlike existing methods, the actor update of GAC utilizes the second-order information of the critic in a computationally efficient manner. This is achieved by separating actor learning into two steps. In the first step, we learn a *non-parameterized Gaussian actor* that locally maximizes the critic under a Kullback-Leibler (KL) divergence constraint. Then, the Gaussian actor is used as a guide for learning a parameterized actor by supervised learning. Our analysis shows that learning the mean of the Gaussian actor is equivalent to performing a second-order update in the action space where the curvature matrix is given by Hessians of the critic and the step-size is controlled by the KL constraint. Furthermore, we establish a connection between GAC and DPG where we show that DPG is a special case of GAC when the Hessians and KL constraint are ignored.

## 2 BACKGROUND

In this section, we firstly give a background of reinforcement learning. Then, we discuss existing second-order methods for policy learning and their issue in deep reinforcement learning.

### 2.1 REINFORCEMENT LEARNING

We consider discrete-time Markov decision processes (MDPs) with continuous state space  $\mathcal{S} \subseteq \mathbb{R}^{d_s}$  and continuous action space  $\mathcal{A} \subseteq \mathbb{R}^{d_a}$ . We denote the state and action at time step  $t \in \mathbb{N}$  by  $\mathbf{s}_t$  and  $\mathbf{a}_t$ , respectively. The initial state  $\mathbf{s}_1$  is determined by the initial state density  $\mathbf{s}_1 \sim p(\mathbf{s})$ . At time step  $t$ , the agent in state  $\mathbf{s}_t$  takes an action  $\mathbf{a}_t$  according to a policy  $\mathbf{a}_t \sim \pi(\mathbf{a}|\mathbf{s}_t)$  and obtains a reward  $r_t = r(\mathbf{s}_t, \mathbf{a}_t)$ . Then, the next state  $\mathbf{s}_{t+1}$  is determined by the transition function  $\mathbf{s}_{t+1} \sim p(\mathbf{s}'|\mathbf{s}_t, \mathbf{a}_t)$ . A trajectory  $\tau = (\mathbf{s}_1, \mathbf{a}_1, r_1, \mathbf{s}_2, \dots)$  gives us the cumulative rewards or return defined as  $\sum_{t=1}^{\infty} \gamma^{t-1} r(\mathbf{s}_t, \mathbf{a}_t)$ , where the discount factor  $0 < \gamma < 1$  assigns different weights to rewards given at different time steps. The expected return of  $\pi$  after executing an action  $\mathbf{a}$  in a state  $\mathbf{s}$  can be expressed through the *action-value function* which is defined as

$$Q^\pi(\mathbf{s}, \mathbf{a}) = \mathbb{E}_{\pi(\mathbf{a}_t|\mathbf{s}_t)_{t \geq 2}, p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)_{t \geq 1}} \left[ \sum_{t=1}^{\infty} \gamma^{t-1} r(\mathbf{s}_t, \mathbf{a}_t) | \mathbf{s}_1 = \mathbf{s}, \mathbf{a}_1 = \mathbf{a} \right], \quad (1)$$

where  $\mathbb{E}_p[\cdot]$  denotes the expectation over the density  $p$  and the subscript  $t \geq 1$  indicates that the expectation is taken over the densities at time steps  $t \geq 1$ . We can define the expected return as

$$\mathcal{J}(\pi) = \mathbb{E}_{p(\mathbf{s}_1), \pi(\mathbf{a}_t|\mathbf{s}_t)_{t \geq 1}, p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)_{t \geq 1}} \left[ \sum_{t=1}^{\infty} \gamma^{t-1} r(\mathbf{s}_t, \mathbf{a}_t) \right] = \mathbb{E}_{p(\mathbf{s}), \pi(\mathbf{a}|\mathbf{s})} [Q^\pi(\mathbf{s}, \mathbf{a})]. \quad (2)$$

<sup>1</sup>This is different from using the gradient of the critic w.r.t. critic parameters to update the critic itself.

The goal of reinforcement learning is to find an optimal policy that maximizes the expected return.

The policy search approach (Deisenroth et al., 2013) parameterizes  $\pi$  by a parameter  $\theta \in \mathbb{R}^{d_\theta}$  and finds  $\theta^*$  which maximizes the expected return:

$$\theta^* = \operatorname{argmax}_{\theta} \mathbb{E}_{p(s), \pi_\theta(a|s)} [Q^{\pi_\theta}(s, a)]. \quad (3)$$

*Policy gradient* methods such as REINFORCE (Williams, 1992) solve this optimization problem by gradient ascent:

$$\theta \leftarrow \theta + \alpha \mathbb{E}_{p(s), \pi_\theta(a|s)} [\nabla_{\theta} \log \pi_{\theta}(a|s) Q^{\pi_{\theta}}(s, a)], \quad (4)$$

where  $\alpha > 0$  is a step-size. In policy search, the action-value function is commonly estimated by the *sample return*:  $Q^{\pi_\theta}(s, a) \approx \frac{1}{N} \sum_{n,t} \gamma^{t-1} r(s_{t,n}, a_{t,n})$  obtained by collecting  $N$  trajectories using  $\pi_\theta$ . The sample return is an unbiased estimator of the action-value function. However, it often has high variance which leads to slow convergence.

An alternative approach is to estimate the action-value function by a *critic* denoted by  $\hat{Q}(s, a)$  whose parameter is learned such that  $\hat{Q}(s, a) \approx Q^{\pi_\theta}(s, a)$ . By replacing the action-value function in Eq.(3) with the critic, we obtain the following optimization problem:

$$\theta^* = \operatorname{argmax}_{\theta} \mathbb{E}_{p(s), \pi_\theta(a|s)} [\hat{Q}(s, a)]. \quad (5)$$

The *actor-critic* method (Sutton et al., 1999) solves this optimization problem by gradient ascent:

$$\theta \leftarrow \theta + \alpha \mathbb{E}_{p(s), \pi_\theta(a|s)} [\nabla_{\theta} \log \pi_{\theta}(a|s) \hat{Q}(s, a)]. \quad (6)$$

The gradient in Eq.(6) often has less variance than that in Eq.(4), which leads to faster convergence<sup>2</sup>. A large class of actor-critic methods is based on this method (Peters & Schaal, 2008; Mnih et al., 2016). As shown in these papers, these methods only use the value of the critic to learn the actor.

The *deterministic policy gradients* (DPG) method (Silver et al., 2014) is an actor-critic method that uses the first-order information of the critic. DPG updates a deterministic actor  $\pi_\theta(s)$  by

$$\theta \leftarrow \theta + \alpha \mathbb{E}_{p(s)} [\nabla_{\theta} \pi_{\theta}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\pi_{\theta}(s)}]. \quad (7)$$

A method related to DPG is the *stochastic value gradients* (SVG) method (Heess et al., 2015) that is able to learn a stochastic policy but it requires learning a model of the transition function.

The actor-critic methods described above only use up to the first-order information of the critic when learning the actor and ignore higher-order ones. Below, we discuss existing approaches that utilize the second-order information by applying second-order optimization methods to solve Eq.(5).

## 2.2 SECOND-ORDER METHODS FOR POLICY LEARNING

The actor-critic methods described above are *first-order methods* which update the optimization variables based on the gradient of the objective function. First-order methods are popular in deep learning thanks to their computational efficiency. However, it is known in machine learning that *second-order methods* often lead to faster learning because they use the curvature information to compute a better update direction, i.e., the steepest ascent direction along the curvature<sup>3</sup>.

The main idea of second-order methods is to rotate the gradient by the inverse of a *curvature matrix*. For instance, second-order updates for the actor-critic method in Eq.(6) are in the following form:

$$\theta \leftarrow \theta + \alpha \mathbf{G}^{-1} \left\{ \mathbb{E}_{p(s), \pi_\theta(a|s)} [\nabla_{\theta} \log \pi_{\theta}(a|s) \hat{Q}(s, a)] \right\}, \quad (8)$$

<sup>2</sup>This gradient is a biased estimator of the policy gradient in Eq.(4). However, it is unbiased under some regularity conditions such as compatible function conditions (Sutton et al., 1999).

<sup>3</sup>We use the term “second-order methods” in a broad sense here, including quasi-Newton and natural gradient methods which approximate the curvature matrix by the first-order information.

where  $\mathbf{G} \in \mathbb{R}^{d_\theta \times d_\theta}$  is a curvature matrix. The behavior of second-order methods depend on the definition of a curvature matrix. The most well-known second-order method is the *Newton method* where its curvature matrix is the Hessian of the objective function w.r.t. the optimization variables:

$$\mathbf{G}_{\text{Hessian}} = \mathbb{E}_{p(\mathbf{s}), \pi_\theta(\mathbf{a}|\mathbf{s})} \left[ (\nabla_{\boldsymbol{\theta}} \log \pi_\theta(\mathbf{a}|\mathbf{s}) \nabla_{\boldsymbol{\theta}} \log \pi_\theta(\mathbf{a}|\mathbf{s})^\top + \nabla_{\boldsymbol{\theta}}^2 \log \pi_\theta(\mathbf{a}|\mathbf{s})) \hat{Q}(\mathbf{s}, \mathbf{a}) \right]. \quad (9)$$

The *natural gradient method* is another well-known second-order method which uses the *Fisher information matrix* (FIM) as the curvature matrix (Amari, 1998):

$$\mathbf{G}_{\text{FIM}} = \mathbb{E}_{p(\mathbf{s}), \pi_\theta(\mathbf{a}|\mathbf{s})} \left[ \nabla_{\boldsymbol{\theta}} \log \pi_\theta(\mathbf{a}|\mathbf{s}) \nabla_{\boldsymbol{\theta}} \log \pi_\theta(\mathbf{a}|\mathbf{s})^\top \right]. \quad (10)$$

Unlike the Hessian matrix, FIM provides information about changes of the policy measured by an approximated KL divergence:  $\mathbb{E}_{p(\mathbf{s})} [\text{KL}(\pi_\theta(\mathbf{a}|\mathbf{s}) || \pi_{\theta'}(\mathbf{a}|\mathbf{s}))] \approx (\boldsymbol{\theta} - \boldsymbol{\theta}')^\top \mathbf{G}_{\text{FIM}} (\boldsymbol{\theta} - \boldsymbol{\theta}')$  (Kakade, 2001). We can see that  $\mathbf{G}_{\text{Hessian}}$  and  $\mathbf{G}_{\text{FIM}}$  are very similar but the former also contains the critic and the Hessian of the actor while the latter does not. This suggests that the Hessian provides more information than that in FIM. However, FIM is always positive semi-definite while the Hessian may be indefinite. Please see Furnston et al. (2016) for detailed comparisons between the two curvature matrices in policy search<sup>4</sup>. Nonetheless, actor-critic methods based on natural gradient were shown to be very efficient (Peters & Schaal, 2008; Schulman et al., 2015a).

We are not aware of existing work that considers second-order updates for DPG or SVG. However, their second-order updates can be trivially derived. For example, a Newton update for DPG is

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \alpha \mathbf{D}^{-1} \left\{ \mathbb{E}_{p(\mathbf{s})} \left[ \nabla_{\boldsymbol{\theta}} \pi_\theta(\mathbf{s}) \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a}) \Big|_{\mathbf{a}=\pi_\theta(\mathbf{s})} \right] \right\}, \quad (11)$$

where the  $(i, j)$ -th entry of the Hessian matrix  $\mathbf{D} \in \mathbb{R}^{d_\theta \times d_\theta}$  is

$$D_{ij} = \mathbb{E}_{p(\mathbf{s})} \left[ \frac{\partial \pi_\theta(\mathbf{s})}{\partial \theta_i} \nabla_{\mathbf{a}}^2 \hat{Q}(\mathbf{s}, \mathbf{a}) \Big|_{\mathbf{a}=\pi_\theta(\mathbf{s})} \frac{\partial \pi_\theta(\mathbf{s})}{\partial \theta_j} + \frac{\partial^2 \pi_\theta(\mathbf{s})}{\partial \theta_i \partial \theta_j} \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a}) \Big|_{\mathbf{a}=\pi_\theta(\mathbf{s})} \right]. \quad (12)$$

Note that  $\partial \pi_\theta(\mathbf{s}) / \partial \theta$  and  $\partial^2 \pi_\theta(\mathbf{s}) / \partial \theta \partial \theta'$  are vectors since  $\pi_\theta(\mathbf{s})$  is a vector-valued function. Interestingly, the Hessian of DPG contains the Hessians of the actor and the critic. In contrast, the Hessian of the actor-critic method contains the Hessian of the actor and the value of the critic.

Second-order methods are appealing in reinforcement learning because they have high data efficiency. However, inverting the curvature matrix (or solving a linear system) requires cubic computational complexity in terms of the number of optimization variables. For this reason, the second-order updates in Eq.(8) and Eq.(11) are impractical in deep reinforcement learning due to a large number of weight parameters in deep neural networks. In such a scenario, an approximation of the curvature matrix is required to reduce the computational burden. For instance, Furnston et al. (2016) proposed to use only diagonal entries of an approximated Hessian matrix. However, this approximation clearly leads to a loss of useful curvature information since the gradient is scaled but not rotated. More recently, Wu et al. (2017) proposed a natural actor-critic method that approximates block-diagonal entries of FIM. However, this approximation corresponds to ignoring useful correlations between weight parameters in different layers of neural networks.

### 3 GUIDE ACTOR-CRITIC

In this section, we propose the guide actor-critic (GAC) method that performs second-order updates without the previously discussed computational issue. Unlike existing methods that directly learn the parameterized actor from the critic, GAC separates the problem of learning the parameterized actor into problems of 1) learning a *guide actor* that locally maximizes the critic, and 2) learning a parameterized actor based on the guide actor. This separation allows us to perform a second-order update for the guide actor where the dimensionality of the curvature matrix is independent of the parameterization of the actor.

We formulate an optimization problem for learning the guide actor in Section 3.1 and present its solution in Section 3.2. Then in Section 3.3 and Section 3.4, we show that the solution corresponds to performing second-order updates. Finally, Section 3.5 presents the learning step for the parameterized actor using supervised learning. The pseudo-code of our method is provided in Appendix B and the source code is available at <https://github.com/voot-t/guide-actor-critic>.

<sup>4</sup> Furnston et al. (2016) proposed an approximate Newton method for policy search. Their policy search method was shown to perform better than methods based on gradient ascent and natural gradient ascent.

### 3.1 OPTIMIZATION PROBLEM FOR GUIDE ACTOR

Our first goal is to learn a guide actor that maximizes the critic. However, greedy maximization should be avoided since the critic is a noisy estimate of the expected return and a greedy actor may change too abruptly across learning iterations. Such a behavior is undesirable in real-world problems, especially in robotics (Deisenroth et al., 2013). Instead, we maximize the critic with additional constraints:

$$\begin{aligned} \max_{\tilde{\pi}} \quad & \mathbb{E}_{p^\beta(\mathbf{s}), \tilde{\pi}(\mathbf{a}|\mathbf{s})} [\hat{Q}(\mathbf{s}, \mathbf{a})], \\ \text{subject to} \quad & \mathbb{E}_{p^\beta(\mathbf{s})} [\text{KL}(\tilde{\pi}(\mathbf{a}|\mathbf{s})||\pi_\theta(\mathbf{a}|\mathbf{s}))] \leq \epsilon, \\ & \mathbb{E}_{p^\beta(\mathbf{s})} [\text{H}(\tilde{\pi}(\mathbf{a}|\mathbf{s}))] \geq \kappa, \\ & \mathbb{E}_{p^\beta(\mathbf{s})} \tilde{\pi}(\mathbf{a}|\mathbf{s}) = 1, \end{aligned} \tag{13}$$

where  $\tilde{\pi}(\mathbf{a}|\mathbf{s})$  is the guide actor to be learned,  $\pi_\theta(\mathbf{a}|\mathbf{s})$  is the current parameterized actor that we want to improve upon, and  $p^\beta(\mathbf{s})$  is the state distribution induced by past trajectories. The objective function differs from the one in Eq.(5) in two important aspects. First, we maximize for a policy function  $\tilde{\pi}$  and not for the policy parameter. This is more advantageous than optimizing for a policy parameter since the policy function can be obtained in a closed form, as will be shown in the next subsection. Second, the expectation is defined over a state distribution from past trajectories and this gives us off-policy methods with higher data efficiency. The first constraint is the Kullback-Leibler (KL) divergence constraint where  $\text{KL}(p(\mathbf{x})||q(\mathbf{x})) = \mathbb{E}_{p(\mathbf{x})} [\log p(\mathbf{x}) - \log q(\mathbf{x})]$ . The second constraint is the Shannon entropy constraint where  $\text{H}(p(\mathbf{x})) = -\mathbb{E}_{p(\mathbf{x})} [\log p(\mathbf{x})]$ . The KL constraint is commonly used in reinforcement learning to prevent unstable behavior due to excessively greedy update (Peters & Schaal, 2008; Peters et al., 2010; Levine & Koltun, 2013; Schulman et al., 2015a). The entropy constraint is crucial for maintaining stochastic behavior and preventing premature convergence (Ziebart et al., 2010; Abdolmaleki et al., 2015; Mnih et al., 2016; Haarnoja et al., 2017). The final constraint ensures that the guide actor is a proper probability density. The KL bound  $\epsilon > 0$  and the entropy bound  $-\infty < \kappa < \infty$  are hyper-parameters which control the exploration-exploitation trade-off of the method. In practice, we fix the value of  $\epsilon$  and adaptively reduce the value of  $\kappa$  based on the current actor’s entropy, as suggested by Abdolmaleki et al. (2015). More details of these tuning parameters are given in Appendix C.

This optimization problem can be solved by the method of Lagrange multipliers. The solution is

$$\tilde{\pi}(\mathbf{a}|\mathbf{s}) \propto \pi_\theta(\mathbf{a}|\mathbf{s}) \frac{\eta^*}{\eta^* + \omega^*} \exp\left(\frac{\hat{Q}(\mathbf{s}, \mathbf{a})}{\eta^* + \omega^*}\right), \tag{14}$$

where  $\eta^* > 0$  and  $\omega^* > 0$  are dual variables corresponding to the KL and entropy constraints, respectively. The dual variable corresponding to the probability density constraint is contained in the normalization term and is determined by  $\eta^*$  and  $\omega^*$ . These dual variables are obtained by minimizing the dual function:

$$g(\eta, \omega) = \eta\epsilon - \omega\kappa + (\eta + \omega) \mathbb{E}_{p^\beta(\mathbf{s})} \left[ \log \int \pi_\theta(\mathbf{a}|\mathbf{s}) \frac{\eta}{\eta + \omega} \exp\left(\frac{\hat{Q}(\mathbf{s}, \mathbf{a})}{\eta + \omega}\right) d\mathbf{a} \right]. \tag{15}$$

All derivations and proofs are given in Appendix A. The solution in Eq.(14) tells us that the guide actor is obtained by weighting the current actor with  $\hat{Q}(\mathbf{s}, \mathbf{a})$ . If we set  $\epsilon \rightarrow 0$  then we have  $\tilde{\pi} \approx \pi_\theta$  and the actor is not updated. On the other hand, if we set  $\epsilon \rightarrow \infty$  then we have  $\tilde{\pi}(\mathbf{a}|\mathbf{s}) \propto \exp(\hat{Q}(\mathbf{s}, \mathbf{a})/\omega^*)$ , which is a softmax policy where  $\omega^*$  is the temperature parameter.

### 3.2 LEARNING GUIDE ACTOR

Computing  $\tilde{\pi}(\mathbf{a}|\mathbf{s})$  and evaluating  $g(\eta, \omega)$  are intractable for an arbitrary  $\pi_\theta(\mathbf{a}|\mathbf{s})$ . We overcome this issue by imposing two assumptions. First, we assume that the actor is the Gaussian distribution:

$$\pi_\theta(\mathbf{a}|\mathbf{s}) = \mathcal{N}(\mathbf{a}|\phi_\theta(\mathbf{s}), \Sigma_\theta(\mathbf{s})), \tag{16}$$

where the mean  $\phi_\theta(\mathbf{s})$  and covariance  $\Sigma_\theta(\mathbf{s})$  are functions parameterized by a policy parameter  $\theta$ . Second, we assume that Taylor’s approximation of  $\hat{Q}(\mathbf{s}, \mathbf{a})$  is locally accurate up to the second-order. More concretely, the second-order Taylor’s approximation using an arbitrary action  $\mathbf{a}_0$  is

given by

$$\widehat{Q}(\mathbf{s}, \mathbf{a}) \approx \widehat{Q}(\mathbf{s}, \mathbf{a}_0) + (\mathbf{a} - \mathbf{a}_0)^\top \mathbf{g}_0(\mathbf{s}) + \frac{1}{2}(\mathbf{a} - \mathbf{a}_0)^\top \mathbf{H}_0(\mathbf{s})(\mathbf{a} - \mathbf{a}_0) + \mathcal{O}(\|\mathbf{a}\|^3), \quad (17)$$

where  $\mathbf{g}_0(\mathbf{s}) = \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0}$  and  $\mathbf{H}_0(\mathbf{s}) = \nabla_{\mathbf{a}}^2 \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0}$  are the gradient and Hessian of the critic w.r.t.  $\mathbf{a}$  evaluated at  $\mathbf{a}_0$ , respectively. By assuming that the higher order term  $\mathcal{O}(\|\mathbf{a}\|^3)$  is sufficiently small, we can rewrite Taylor’s approximation at  $\mathbf{a}_0$  as

$$\widehat{Q}_0(\mathbf{s}, \mathbf{a}) = \frac{1}{2} \mathbf{a}^\top \mathbf{H}_0(\mathbf{s}) \mathbf{a} + \mathbf{a}^\top \boldsymbol{\psi}_0(\mathbf{s}) + \xi_0(\mathbf{s}), \quad (18)$$

where  $\boldsymbol{\psi}_0(\mathbf{s}) = \mathbf{g}_0(\mathbf{s}) - \mathbf{H}_0(\mathbf{s})\mathbf{a}_0$  and  $\xi_0(\mathbf{s}) = \frac{1}{2} \mathbf{a}_0^\top \mathbf{H}_0(\mathbf{s}) \mathbf{a}_0 - \mathbf{a}_0^\top \mathbf{g}_0(\mathbf{s}) + \widehat{Q}(\mathbf{s}, \mathbf{a}_0)$ . Note that  $\mathbf{H}_0(\mathbf{s})$ ,  $\boldsymbol{\psi}_0(\mathbf{s})$ , and  $\xi_0(\mathbf{s})$  depend on the value of  $\mathbf{a}_0$  and do not depend on the value of  $\mathbf{a}$ . This dependency is explicitly denoted by the subscript. The choice of  $\mathbf{a}_0$  will be discussed in Section 3.3.

Substituting the Gaussian distribution and Taylor’s approximation into Eq.(14) yields another Gaussian distribution  $\tilde{\pi}(\mathbf{a}|\mathbf{s}) = \mathcal{N}(\mathbf{a}|\boldsymbol{\phi}_+(\mathbf{s}), \boldsymbol{\Sigma}_+(\mathbf{s}))$ , where the mean and covariance are given by

$$\boldsymbol{\phi}_+(\mathbf{s}) = \mathbf{F}^{-1}(\mathbf{s})\mathbf{L}(\mathbf{s}), \quad \boldsymbol{\Sigma}_+(\mathbf{s}) = (\eta^* + \omega^*)\mathbf{F}^{-1}(\mathbf{s}). \quad (19)$$

The matrix  $\mathbf{F}(\mathbf{s}) \in \mathbb{R}^{d_a \times d_a}$  and vector  $\mathbf{L}(\mathbf{s}) \in \mathbb{R}^{d_a}$  are defined as

$$\mathbf{F}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbf{H}_0(\mathbf{s}), \quad \mathbf{L}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) \boldsymbol{\phi}_\theta(\mathbf{s}) + \boldsymbol{\psi}_0(\mathbf{s}). \quad (20)$$

The dual variables  $\eta^*$  and  $\omega^*$  are obtained by minimizing the following dual function:

$$\begin{aligned} \widehat{g}(\eta, \omega) = & \eta\epsilon - \omega\kappa + (\eta + \omega) \mathbb{E}_{p^\beta(\mathbf{s})} \left[ \log \sqrt{\frac{|2\pi(\eta + \omega)\mathbf{F}_\eta^{-1}(\mathbf{s})|}{|2\pi\boldsymbol{\Sigma}_\theta(\mathbf{s})|^{\frac{\eta}{\eta + \omega}}}} \right] \\ & + \frac{1}{2} \mathbb{E}_{p^\beta(\mathbf{s})} \left[ \mathbf{L}_\eta(\mathbf{s})^\top \mathbf{F}_\eta^{-1}(\mathbf{s}) \mathbf{L}_\eta(\mathbf{s}) - \eta \boldsymbol{\phi}_\theta(\mathbf{s})^\top \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) \boldsymbol{\phi}_\theta(\mathbf{s}) \right] + \text{const}, \end{aligned} \quad (21)$$

where  $\mathbf{F}_\eta(\mathbf{s})$  and  $\mathbf{L}_\eta(\mathbf{s})$  are defined similarly to  $\mathbf{F}(\mathbf{s})$  and  $\mathbf{L}(\mathbf{s})$  but with  $\eta$  instead of  $\eta^*$ .

The practical advantage of using the Gaussian distribution and Taylor’s approximation is that the guide actor can be obtained in a closed form and the dual function can be evaluated through matrix-vector products. The expectation over  $p^\beta(\mathbf{s})$  can be approximated by e.g., samples drawn from a replay buffer (Mnih et al., 2015). We require inverting  $\mathbf{F}_\eta(\mathbf{s})$  to evaluate the dual function. However, these matrices are computationally cheap to invert when the dimension of actions is not large.

As shown in Eq.(19), the mean and covariance of the guide actor is computed using both the gradient and Hessian of the critic. Yet, these computations do not resemble second-order updates discussed previously in Section 2.2. Below, we show that for a particular choice of  $\mathbf{a}_0$ , the mean computation corresponds to a second-order update that rotates gradients by a curvature matrix.

### 3.3 GUIDE ACTOR LEARNING AS SECOND-ORDER OPTIMIZATION

For now we assume that the critic is an accurate estimator of the true action-value function. In this case, the quality of the guide actor depends on the accuracy of sample approximation in  $\widehat{g}(\eta, \omega)$  and the accuracy of Taylor’s approximation. To obtain an accurate Taylor’s approximation of  $\widehat{Q}(\mathbf{s}, \mathbf{a})$  using an action  $\mathbf{a}_0$ , the action  $\mathbf{a}_0$  should be in the vicinity of  $\mathbf{a}$ . However, we did not directly use any individual  $\mathbf{a}$  to compute the guide actor, but we weight  $\pi_\theta(\mathbf{a}|\mathbf{s})$  by  $\exp(\widehat{Q}(\mathbf{s}, \mathbf{a}))$  (see Eq.(14)). Thus, to obtain an accurate Taylor’s approximation of the critic, the action  $\mathbf{a}_0$  needs to be similar to actions sampled from  $\pi_\theta(\mathbf{a}|\mathbf{s})$ . Based on this observation, we propose two approaches to perform Taylor’s approximation.

**Taylor’s approximation around the mean.** In this approach, we perform Taylor’s approximation using the mean of  $\pi_\theta(\mathbf{a}|\mathbf{s})$ . More specifically, we use  $\mathbf{a}_0 = \mathbb{E}_{\pi_\theta(\mathbf{a}|\mathbf{s})}[\mathbf{a}] = \boldsymbol{\phi}_\theta(\mathbf{s})$  for Eq.(18). In this case, we can show that the mean update in Eq.(19) corresponds to performing a second-order update in the action space to maximize  $\widehat{Q}(\mathbf{s}, \mathbf{a})$ :

$$\boldsymbol{\phi}_+(\mathbf{s}) = \boldsymbol{\phi}_\theta(\mathbf{s}) + \mathbf{F}_{\boldsymbol{\phi}_\theta}^{-1}(\mathbf{s}) \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\boldsymbol{\phi}_\theta(\mathbf{s})}, \quad (22)$$

where  $\mathbf{F}_{\phi_\theta}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbf{H}_{\phi_\theta}(\mathbf{s})$  and  $\mathbf{H}_{\phi_\theta}(\mathbf{s}) = \nabla_{\mathbf{a}}^2 \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_\theta(\mathbf{s})}$ . This equivalence can be shown by substitution and the proof is given in Appendix A.2. This update equation reveals that the guide actor maximizes the critic by taking a step in the action space similarly to the Newton method. However, the main difference lies in the curvature matrix where the Newton method uses Hessians  $\mathbf{H}_{\phi_\theta}(\mathbf{s})$  but we use a *damped Hessian*  $\mathbf{F}_{\phi_\theta}(\mathbf{s})$ . The damping term  $\eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s})$  corresponds to the effect of the KL constraint and can be viewed as a trust-region that controls the step-size. This damping term is particularly important since Taylor’s approximation is accurate only locally and we should not take a large step in each update (Nocedal & Wright, 2006).

**Expectation of Taylor’s approximations.** Instead of using Taylor’s approximation around the mean, we may use an expectation of Taylor’s approximation over the distribution. More concretely, we define  $\widetilde{Q}(\mathbf{s}, \mathbf{a})$  to be an expectation of  $\widehat{Q}_0(\mathbf{s}, \mathbf{a})$  over  $\pi_\theta(\mathbf{a}_0|\mathbf{s})$ :

$$\widetilde{Q}(\mathbf{s}, \mathbf{a}) = \frac{1}{2} \mathbf{a}^\top \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s})] \mathbf{a} + \mathbf{a}^\top \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\boldsymbol{\psi}_0(\mathbf{s})] + \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\xi_0(\mathbf{s})]. \quad (23)$$

Note that  $\mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s})] = \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\nabla_{\mathbf{a}}^2 \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0}]$  and the expectation is computed w.r.t. the distribution  $\pi_\theta$  of  $\mathbf{a}_0$ . We use this notation to avoid confusion even though  $\pi_\theta(\mathbf{a}_0|\mathbf{s})$  and  $\pi_\theta(\mathbf{a}|\mathbf{s})$  are the same distribution. When Eq.(23) is used, the mean update does not directly correspond to any second-order optimization step. However, under an (unrealistic) independence assumption  $\mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s}) \mathbf{a}_0] = \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s})] \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{a}_0]$ , we can show that the mean update corresponds to the following second-order optimization step:

$$\phi_+(\mathbf{s}) = \phi_\theta(\mathbf{s}) + \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{F}_0(\mathbf{s})]^{-1} \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0} \right], \quad (24)$$

where  $\mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{F}_0(\mathbf{s})] = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s})]$ . Interestingly, the mean is updated by rotating an expected gradient using an expected Hessians. In practice, the expectations can be approximated using sampled actions  $\{\mathbf{a}_{0,i}\}_{i=1}^S \sim \pi_\theta(\mathbf{a}|\mathbf{s})$ . We believe that this sampling can be advantageous for avoiding local optima. Note that when the expectation is approximated by a single sample  $\mathbf{a}_0 \sim \pi_\theta(\mathbf{a}|\mathbf{s})$ , we obtain the update in Eq.(24) regardless of the independence assumption.

In the remainder, we use  $\mathbf{F}(\mathbf{s})$  to denote both of  $\mathbf{F}_{\phi_\theta}(\mathbf{s})$  and  $\mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{F}_0(\mathbf{s})]$ , and use  $\mathbf{H}(\mathbf{s})$  to denote both of  $\mathbf{H}_{\phi_\theta}(\mathbf{s})$  and  $\mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s})]$ . In the experiments, we use GAC-0 to refer to GAC with Taylor’s approximation around the mean, and we use GAC-1 to refer to GAC with Taylor’s approximation by a single sample  $\mathbf{a}_0 \sim \pi_\theta(\mathbf{a}|\mathbf{s})$ .

### 3.4 GAUSS-NEWTON APPROXIMATION OF HESSIAN

The covariance update in Eq.(19) indicates that  $\mathbf{F}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbf{H}(\mathbf{s})$  needs to be positive definite. The matrix  $\mathbf{F}(\mathbf{s})$  is guaranteed to be positive definite if the Hessian matrix  $\mathbf{H}(\mathbf{s})$  is negative semi-definite. However, this is not guaranteed in practice unless  $\widehat{Q}(\mathbf{s}, \mathbf{a})$  is a concave function in terms of  $\mathbf{a}$ . To overcome this issue, we firstly consider the following identity:

$$\mathbf{H}(\mathbf{s}) = \nabla_{\mathbf{a}}^2 \widehat{Q}(\mathbf{s}, \mathbf{a}) = -\nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a}) \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})^\top + \nabla_{\mathbf{a}}^2 \exp(\widehat{Q}(\mathbf{s}, \mathbf{a})) \exp(-\widehat{Q}(\mathbf{s}, \mathbf{a})). \quad (25)$$

The proof is given in Appendix A.3. The first term is always negative semi-definite while the second term is indefinite. Therefore, a negative semi-definite approximation of the Hessian can be obtained as

$$\mathbf{H}_0(\mathbf{s}) \approx - \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a}) \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})^\top \right]_{\mathbf{a}=\mathbf{a}_0}. \quad (26)$$

The second term in Eq.(25) is proportional to  $\exp(-\widehat{Q}(\mathbf{s}, \mathbf{a}))$  and it will be small for high values of  $\widehat{Q}(\mathbf{s}, \mathbf{a})$ . This implies that the approximation should get more accurate as the policy approach a local maxima of  $\widehat{Q}(\mathbf{s}, \mathbf{a})$ . We call this approximation *Gauss-Newton approximation* since it is similar to the Gauss-Newton approximation for the Newton method (Nocedal & Wright, 2006).

### 3.5 LEARNING PARAMETERIZED ACTOR

The second step of GAC is to learn a parameterized actor that well represents the guide actor. Below, we discuss two supervised learning approaches for learning a parameterized actor.

### 3.5.1 FULLY-PARAMETERIZED GAUSSIAN POLICY

Since the guide actor is a Gaussian distribution with a state-dependent mean and covariance, a natural choice for the parameterized actor is again a parameterized Gaussian distribution with a state-dependent mean and covariance:  $\pi_{\theta}(\mathbf{a}|\mathbf{s}) = \mathcal{N}(\mathbf{a}|\phi_{\theta}(\mathbf{s}), \Sigma_{\theta}(\mathbf{s}))$ . The parameter  $\theta$  can be learned by minimizing the expected KL divergence to the guide actor:

$$\begin{aligned} \mathcal{L}^{\text{KL}}(\theta) &= \mathbb{E}_{p^{\beta}(\mathbf{s})} [\text{KL}(\pi_{\theta}(\mathbf{a}|\mathbf{s})||\tilde{\pi}(\mathbf{a}|\mathbf{s}))] \\ &= \mathbb{E}_{p^{\beta}(\mathbf{s})} \left[ \frac{\text{Tr}(\mathbf{F}(\mathbf{s})\Sigma_{\theta}(\mathbf{s}))}{\eta^* + \omega^*} - \log |\Sigma_{\theta}(\mathbf{s})| \right] + \frac{\mathcal{L}^{\text{W}}(\theta)}{\eta^* + \omega^*} + \text{const}, \end{aligned} \quad (27)$$

where  $\mathcal{L}^{\text{W}}(\theta) = \mathbb{E}_{p^{\beta}(\mathbf{s})} [\|\phi_{\theta}(\mathbf{s}) - \phi_{+}(\mathbf{s})\|_{\mathbf{F}(\mathbf{s})}^2]$  is the weighted-mean-squared-error (WMSE) which only depends on  $\theta$  of the mean function. The const term does not depend on  $\theta$ .

Minimizing the KL divergence reveals connections between GAC and deterministic policy gradients (DPG) (Silver et al., 2014). By computing the gradient of the WMSE, it can be shown that

$$\begin{aligned} \frac{\nabla_{\theta} \mathcal{L}^{\text{W}}(\theta)}{2} &= -\mathbb{E}_{p^{\beta}(\mathbf{s})} \left[ \nabla_{\theta} \phi_{\theta}(\mathbf{s}) \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(\mathbf{s})} \right] + \mathbb{E}_{p^{\beta}(\mathbf{s})} \left[ \nabla_{\theta} \phi_{\theta}(\mathbf{s}) \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{+}(\mathbf{s})} \right] \\ &\quad + \eta^* \mathbb{E}_{p^{\beta}(\mathbf{s})} \left[ \nabla_{\theta} \phi_{\theta}(\mathbf{s}) \Sigma^{-1}(\mathbf{s}) \mathbf{H}_{\phi_{\theta}^{-1}}^{-1}(\mathbf{s}) \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(\mathbf{s})} \right] \\ &\quad - \eta^* \mathbb{E}_{p^{\beta}(\mathbf{s})} \left[ \nabla_{\theta} \phi_{\theta}(\mathbf{s}) \Sigma^{-1}(\mathbf{s}) \mathbf{H}_{\phi_{\theta}^{-1}}^{-1}(\mathbf{s}) \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{+}(\mathbf{s})} \right]. \end{aligned} \quad (28)$$

The proof is given in Appendix A.4. The negative of the first term is precisely equivalent to DPG. Thus, updating the mean parameter by minimizing the KL loss with gradient descent can be regarded as updating the mean parameter with *biased* DPG where the bias terms depend on  $\eta^*$ . We can verify that  $\nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{+}(\mathbf{s})} = 0$  when  $\eta^* = 0$  and this is the case of  $\epsilon \rightarrow \infty$ . Thus, all bias terms vanish when the KL constraint is ignored and the mean update of GAC coincides with DPG. However, unlike DPG which learns a deterministic policy, we can learn both the mean and covariance in GAC.

### 3.5.2 GAUSSIAN POLICY WITH PARAMETERIZED MEAN

While a state-dependent parameterized covariance function is flexible, we observe that learning performance is sensitive to the initial parameter of the covariance function. For practical purposes, we propose using a parametrized Gaussian distribution with state-independent covariance:  $\pi_{\theta}(\mathbf{a}|\mathbf{s}) = \mathcal{N}(\mathbf{a}|\phi_{\theta}(\mathbf{s}), \Sigma)$ . This class of policies subsumes deterministic policies with additive independent Gaussian noise for exploration. To learn  $\theta$ , we minimize the mean-squared-error (MSE):

$$\mathcal{L}^{\text{M}}(\theta) = \frac{1}{2} \mathbb{E}_{p^{\beta}(\mathbf{s})} [\|\phi_{\theta}(\mathbf{s}) - \phi_{+}(\mathbf{s})\|_2^2]. \quad (29)$$

For the covariance, we use the average of the guide covariances:  $\Sigma = (\eta^* + \omega^*) \mathbb{E}_{p^{\beta}(\mathbf{s})} [\mathbf{F}^{-1}(\mathbf{s})]$ . For computational efficiency, we execute a single gradient update in each learning iteration instead of optimizing this loss function until convergence.

Similarly to the above analysis, the gradient of the MSE w.r.t.  $\theta$  can be expanded and rewritten into

$$\nabla_{\theta} \mathcal{L}^{\text{M}}(\theta) = \mathbb{E}_{p^{\beta}(\mathbf{s})} \left[ \nabla_{\theta} \phi_{\theta}(\mathbf{s}) \mathbf{H}^{-1}(\mathbf{s}) \left( \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(\mathbf{s})} - \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{+}(\mathbf{s})} \right) \right]. \quad (30)$$

Again, the mean update of GAC coincides with DPG when we minimize the MSE and set  $\eta^* = 0$  and  $\mathbf{H}(\mathbf{s}) = -\mathbf{I}$  where  $\mathbf{I}$  is the identity matrix. We can also substitute these values back into Eq.(22). By doing so, we can interpret DPG as a method that performs first-order optimization in the action space:

$$\phi_{+}(\mathbf{s}) = \phi_{\theta}(\mathbf{s}) + \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(\mathbf{s})}, \quad (31)$$

and then uses the gradient in Eq.(30) to update the policy parameter. This interpretation shows that DPG is a first-order method that only uses the first-order information of the critic for actor learning. Therefore in principle, GAC, which uses the second-order information of the critic, should learn faster than DPG.

### 3.6 POLICY EVALUATION FOR CRITIC

Beside actor learning, the performance of actor-critic methods also depends on the accuracy of the critic. We assume that the critic  $\widehat{Q}_\nu(\mathbf{s}, \mathbf{a})$  is represented by neural networks with a parameter  $\nu$ . We adopt the approach proposed by Lillicrap et al. (2015) with some adjustment to learn  $\nu$ . More concretely, we use gradient descent to minimize the squared Bellman error with a slowly moving target critic:

$$\nu \leftarrow \nu - \alpha \nabla_\nu \mathbb{E}_{p^\beta(\mathbf{s}), \beta(\mathbf{a}|\mathbf{s}), p(\mathbf{s}'|\mathbf{s}, \mathbf{a})} \left[ \left( \widehat{Q}_\nu(\mathbf{s}, \mathbf{a}) - y \right)^2 \right], \quad (32)$$

where  $\alpha > 0$  is the step-size. The target value  $y = r(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_{\pi(\mathbf{a}'|\mathbf{s}')} [\widehat{Q}_{\bar{\nu}}(\mathbf{s}', \mathbf{a}')] is computed by the target critic  $\widehat{Q}_{\bar{\nu}}(\mathbf{s}', \mathbf{a}')$  whose parameter  $\bar{\nu}$  is updated by  $\bar{\nu} \leftarrow \tau \nu + (1 - \tau) \bar{\nu}$  for  $0 < \tau < 1$ . As suggested by Lillicrap et al. (2015), the target critic improves the learning stability and we set  $\tau = 0.001$  in experiments. The expectation for the squared error is approximated using mini-batch samples  $\{(\mathbf{s}_n, \mathbf{a}_n, r_n, \mathbf{s}'_n)\}_{n=1}^N$  drawn from a replay buffer. The expectation over the current actor  $\pi(\mathbf{a}'|\mathbf{s}')$  is approximated using samples  $\{\mathbf{a}'_{n,m}\}_{m=1}^M \sim \pi_\theta(\mathbf{a}'|\mathbf{s}'_n)$  for each  $\mathbf{s}'_n$ . We do not use a target actor to compute  $y$  since the KL upper-bound already constrains the actor update and a target actor will further slow it down. Note that we are not restricted to this evaluation method and more efficient methods such as *Retrace* (Munos et al., 2016) can also be used.$

Our method requires computing  $\nabla_{\mathbf{a}} \widehat{Q}_\nu(\mathbf{s}, \mathbf{a})$  and its outer product for the Gauss-Newton approximation. The computational complexity of the outer product operation is  $\mathcal{O}(d_{\mathbf{a}}^2)$  and is inexpensive when compared to the dimension of  $\nu$ . For a linear-in-parameter model  $\widehat{Q}_\nu(\mathbf{s}, \mathbf{a}) = \nu^\top \boldsymbol{\mu}(\mathbf{s}, \mathbf{a})$ , the gradient can be efficiently computed for common choices of the basis function  $\boldsymbol{\mu}$  such as the Gaussian function. For deep neural network models, the gradient can be computed by the automatic-differentiation (Goodfellow et al., 2016) where its cost depends on the network architecture.

## 4 RELATED WORK

Besides the connections to DPG, our method is also related to existing methods as follows.

A similar optimization problem to Eq.(13) was considered by the *model-free trajectory optimization* (MOTO) method (Akrouer et al., 2016). Our method can be viewed as a non-trivial extension of MOTO with two significant novelties. First, MOTO learns a sequence of time-dependent log-linear Gaussian policies  $\pi_t(\mathbf{a}|\mathbf{s}) = \mathcal{N}(\mathbf{a}|\mathbf{B}_t \mathbf{s} + \mathbf{b}_t, \boldsymbol{\Sigma}_t)$ , while our method learns a log-nonlinear Gaussian policy. Second, MOTO learns a time-dependent critic given by  $\widehat{Q}_t(\mathbf{s}, \mathbf{a}) = \frac{1}{2} \mathbf{a}^\top \mathbf{C}_t \mathbf{a} + \mathbf{a}^\top \mathbf{D}_t \mathbf{s} + \mathbf{a}^\top \mathbf{c}_t + \xi_t(\mathbf{s})$  and performs policy update with these functions. In contrast, our method learns a more complex critic and performs Taylor’s approximation in each training step.

Besides MOTO, the optimization problem also resembles that of *trust region policy optimization* (TRPO) (Schulman et al., 2015a). TRPO solves the following optimization problem:

$$\max_{\theta'} \mathbb{E}_{p^{\pi_{\theta'}(\mathbf{s}), \pi_{\theta'}(\mathbf{a}|\mathbf{s})}} [\widehat{Q}(\mathbf{s}, \mathbf{a})] \quad \text{subject to} \quad \mathbb{E}_{p^{\pi_{\theta'}(\mathbf{s})}} [\text{KL}(\pi_{\theta'}(\mathbf{a}|\mathbf{s}) || \pi_{\theta'}(\mathbf{a}|\mathbf{s}))] \leq \epsilon, \quad (33)$$

where  $\widehat{Q}(\mathbf{s}, \mathbf{a})$  may be replaced by an estimate of the advantage function (Schulman et al., 2015b). There are two major differences between the two problems. First, TRPO optimizes the policy parameter while we optimize the guide actor. Second, TRPO solves the optimization problem by conjugate gradient where the KL divergence is approximated by the Fisher information matrix, while we solve the optimization problem in a closed form with a quadratic approximation of the critic.

Our method is also related to *maximum-entropy RL* (Ziebart et al., 2010; Azar et al., 2012; Haarnoja et al., 2017; Nachum et al., 2017), which maximizes the expected cumulative reward with an additional entropy bonus:  $\sum_{t=1}^{\infty} \mathbb{E}_{p^{\pi}(\mathbf{s})} [r(\mathbf{s}_t, \mathbf{a}_t) + \alpha \text{H}(\pi(\mathbf{a}_t|\mathbf{s}_t))]$ , where  $\alpha > 0$  is a trade-off parameter. The optimal policy in maximum-entropy RL is the softmax policy given by

$$\pi_{\text{MaxEnt}}^*(\mathbf{a}|\mathbf{s}) = \exp \left( \frac{Q_{\text{soft}}^*(\mathbf{s}, \mathbf{a}) - V_{\text{soft}}^*(\mathbf{s})}{\alpha} \right) \propto \exp \left( \frac{Q_{\text{soft}}^*(\mathbf{s}, \mathbf{a})}{\alpha} \right), \quad (34)$$

where  $Q_{\text{soft}}^*(\mathbf{s}, \mathbf{a})$  and  $V_{\text{soft}}^*(\mathbf{s})$  are the optimal soft action-value and state-value functions, respectively (Haarnoja et al., 2017; Nachum et al., 2017). For a policy  $\pi$ , these are defined as

$$Q_{\text{soft}}^\pi(\mathbf{s}, \mathbf{a}) = r(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_{p(\mathbf{s}'|\mathbf{s}, \mathbf{a})} [V_{\text{soft}}^\pi(\mathbf{s}')], \quad (35)$$

$$V_{\text{soft}}^\pi(\mathbf{s}) = \alpha \log \int \exp\left(\frac{Q_{\text{soft}}^\pi(\mathbf{s}, \mathbf{a})}{\alpha}\right) d\mathbf{a}. \quad (36)$$

The softmax policy and the soft state-value function in maximum-entropy RL closely resemble the guide actor in Eq.(14) when  $\eta^* = 0$  and the log-integral term in Eq.(15) when  $\eta = 0$ , respectively, except for the definition of action-value functions. To learn the optimal policy of maximum-entropy RL, Haarnoja et al. (2017) proposed *soft Q-learning* which uses importance sampling to compute the soft value functions and approximates the intractable policy using a separate policy function. Our method largely differs from soft Q-learning since we use Taylor’s approximation to convert the intractable integral into more convenient matrix-vector products.

The idea of firstly learning a non-parameterized policy and then later learning a parameterized policy by supervised learning was considered previously in *guided policy search* (GPS) (Levine & Koltun, 2013). However, GPS learns the guide policy by trajectory optimization methods such as an iterative linear-quadratic Gaussian regulator (Li & Todorov, 2004), which requires a model of the transition function. In contrast, we learn the guide policy via the critic without learning the transition function.

## 5 EXPERIMENTAL RESULTS

We evaluate GAC on the OpenAI gym platform (Brockman et al., 2016) with the Mujoco Physics simulator (Todorov et al., 2012). The actor and critic are neural networks with two hidden layers of 400 and 300 units, as described in Appendix C. We compare GAC-0 and GAC-1 against deep DPG (DDPG) (Lillicrap et al., 2015), Q-learning with a normalized advantage function (Q-NAF) (Gu et al., 2016), and TRPO (Schulman et al., 2015a;b). Figure 1 shows the learning performance on 9 continuous control tasks. Overall, both GAC-0 and GAC-1 perform comparably with existing methods and they clearly outperform the other methods in Half-Cheetah.

The performance of GAC-0 and GAC-1 is comparable on these tasks, except on Humanoid where GAC-1 learns faster. We expect GAC-0 to be more stable and reliable but easier to get stuck at local optima. On the other hand, the randomness introduced by GAC-1 leads to high variance approximation but this could help escape poor local optima. We conjecture GAC-S that uses  $S > 1$  samples for the averaged Taylor’s approximation should outperform both GAC-0 and GAC-1. While this is computationally expensive, we can use parallel computation to reduce the computation time.

The expected returns of both GAC-0 and GAC-1 have high fluctuations on the Hopper and Walker2D tasks when compared to TRPO as can be seen in Figure 1g and Figure 1h. We observe that they can learn good policies for these tasks in the middle of learning. However, the policies quickly diverge to poor ones and then they are quickly improved to be good policies again. We believe that this happens because the step-size  $\mathbf{F}^{-1}(\mathbf{s}) = (\eta^* \boldsymbol{\Sigma}^{-1} - \mathbf{H}(\mathbf{s}))^{-1}$  of the guide actor in Eq. (22) can be very large near local optima for Gauss-Newton approximation. That is, the gradients near local optima have small magnitude and this makes the approximation  $\mathbf{H}(\mathbf{s}) = \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a}) \nabla_{\mathbf{a}} \hat{Q}(\mathbf{s}, \mathbf{a})^\top$  small as well. If  $\eta^* \boldsymbol{\Sigma}^{-1}$  is also relatively small then the matrix  $\mathbf{F}^{-1}(\mathbf{s})$  can be very large. Thus, under these conditions, GAC may use too large step sizes to compute the guide actor and this results in high fluctuations in performance. We expect that this scenario can be avoided by reducing the KL bound  $\epsilon$  or adding a regularization constant to the Gauss-Newton approximation.

Table 1 in Appendix C shows the wall-clock computation time. DDPG is computationally the most efficient method on all tasks. GAC has low computation costs on tasks with low dimensional actions and its cost increases as the dimensionality of action increases. This high computation cost is due to the dual optimization for finding the step-size parameters  $\eta$  and  $\omega$ . We believe that the computation cost of GAC can be significantly reduced by letting  $\eta$  and  $\omega$  be external tuning parameters.

## 6 CONCLUSION AND FUTURE WORK

Actor-critic methods are appealing for real-world problems due to their good data efficiency and learning speed. However, existing actor-critic methods do not use second-order information of the

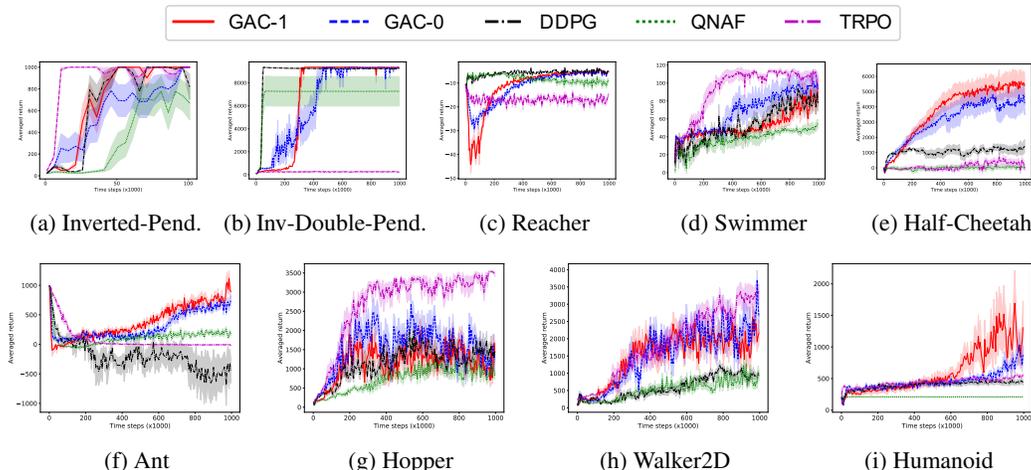


Figure 1: Expected returns averaged over 10 trials. The x-axis indicates training time steps. The y-axis indicates averaged return and higher is better. More clear figures are provided in Appendix C.2.

critic. In this paper, we established a novel framework that distinguishes itself from existing work by utilizing Hessians of the critic for actor learning. Within this framework, we proposed a practical method that uses Gauss-Newton approximation instead of the Hessians. We showed through experiments that our method is promising and thus the framework should be further investigated.

Our analysis showed that the proposed method is closely related to deterministic policy gradients (DPG). However, DPG was also shown to be a limiting case of the stochastic policy gradients when the policy variance approaches zero (Silver et al., 2014). It is currently unknown whether our framework has a connection to the stochastic policy gradients as well, and finding such a connection is our future work.

Our main goal in this paper was to provide a new actor-critic framework and we do not claim that our method achieves the state-of-the-art performance. However, its performance can still be improved in many directions. For instance, we may impose a KL constraint for a parameterized actor to improve its stability, similarly to TRPO (Schulman et al., 2015a). We can also apply more efficient policy evaluation methods such as Retrace (Munos et al., 2016) to achieve better critic learning.

#### ACKNOWLEDGMENTS

MS was partially supported by KAKENHI 17H00757.

#### REFERENCES

- Abbas Abdolmaleki, Rudolf Lioutikov, Jan Peters, Nuno Lau, Luís Paulo Reis, and Gerhard Neumann. Model-Based Relative Entropy Stochastic Search. In *Advances in Neural Information Processing Systems 28, December 7-12, 2015, Montreal, Quebec, Canada*, 2015.
- Riad Akrou, Gerhard Neumann, Hany Abdulsamad, and Abbas Abdolmaleki. Model-Free Trajectory Optimization for Reinforcement Learning. In *Proceedings of the 33rd International Conference on Machine Learning, June 19-24, 2016, New York City, NY, USA*, 2016.
- Shun-ichi Amari. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10(2): 251–276, 1998.
- Brandon Amos, Lei Xu, and J. Zico Kolter. Input Convex Neural Networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, August 6-11, 2017, Sydney, Australia*, 2017.
- Mohammad Gheshlaghi Azar, Vicenç Gómez, and Hilbert J. Kappen. Dynamic Policy Programming. *Journal of Machine Learning Research*, 13:3207–3245, 2012.

- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *CoRR*, abs/1606.01540, 2016.
- Marc Peter Deisenroth, Gerhard Neumann, and Jan Peters. A Survey on Policy Search for Robotics. *Foundations and Trends in Robotics*, 2(1-2):1–142, 2013.
- Thomas Furst, Guy Lever, and David Barber. Approximate Newton Methods for Policy Search in Markov Decision Processes. *Journal of Machine Learning Research*, 17(227):1–51, 2016.
- Xavier Glorot and Yoshua Bengio. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In Yee Whye Teh and Mike Titterton (eds.), *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics, May 13-15, 2010, Sardinia, Italy*, 2010.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- Shixiang Gu, Timothy P. Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous Deep Q-Learning with Model-based Acceleration. In *Proceedings of the 33rd International Conference on Machine Learning, June 19-24, 2016, New York City, NY, USA*, 2016.
- Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement Learning with Deep Energy-Based Policies. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, August 6-11, 2017, Sydney, Australia*, 2017.
- Nicolas Heess, Gregory Wayne, David Silver, Timothy P. Lillicrap, Tom Erez, and Yuval Tassa. Learning Continuous Control Policies by Stochastic Value Gradients. In *Advances in Neural Information Processing Systems 28, December 7-12, 2015, Montreal, Quebec, Canada*, 2015.
- Sham Kakade. A Natural Policy Gradient. In *Advances in Neural Information Processing Systems 14, December 3-8, 2001, Vancouver, British Columbia, Canada*, 2001.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *CoRR*, abs/1412.6980, 2014.
- Vijay R. Konda and John N. Tsitsiklis. On Actor-Critic Algorithms. *SIAM Journal on Control and Optimization*, 42(4):1143–1166, April 2003.
- Sergey Levine and Vladlen Koltun. Guided Policy Search. In *Proceedings of the 30th International Conference on Machine Learning, June 16-21, 2013, Atlanta, GA, USA*, 2013.
- Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end Training of Deep Visuomotor Policies. *Journal of Machine Learning Research*, 17(1):1334–1373, January 2016. ISSN 1532-4435.
- Weiwei Li and Emanuel Todorov. Iterative Linear Quadratic Regulator Design for Nonlinear Biological Movement Systems. In *Proceedings of the 1st International Conference on Informatics in Control, Automation and Robotics, August 25-28, 2004, Setúbal, Portugal*, pp. 222–229, 2004.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous Control with Deep Reinforcement Learning. *CoRR*, abs/1509.02971, 2015.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellefleur, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-Level Control Through Deep Reinforcement Learning. *Nature*, 518(7540):529–533, February 2015. ISSN 00280836.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous Methods for Deep Reinforcement Learning. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of the 33rd International Conference on Machine Learning, June 19-24, 2016, New York City, NY, USA*, 2016.

- Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc G. Bellemare. Safe and Efficient Off-Policy Reinforcement Learning. In *Advances in Neural Information Processing Systems 29, December 5-10, 2016, Barcelona, Spain*, 2016.
- Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the Gap Between Value and Policy Based Reinforcement Learning. In *Advances in Neural Information Processing Systems 30, 4-9 December 2017, Long Beach, CA, USA*, 2017.
- Jorge Nocedal and Stephen J. Wright. *Numerical Optimization, Second Edition*. World Scientific, 2006.
- Jan Peters and Stefan Schaal. Natural Actor-Critic. *Neurocomputing*, 71(7-9):1180–1190, 2008.
- Jan Peters, Katharina Mülling, and Yasemin Altun. Relative Entropy Policy Search. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence, July 11-15, 2010, Atlanta, Georgia, USA*, 2010.
- John Schulman, Sergey Levine, Philipp Moritz, Michael Jordan, and Pieter Abbeel. Trust Region Policy Optimization. In *Proceedings of the 32nd International Conference on Machine Learning, July 6-11, 2015, Lille, France*, pp. 1889–1897, 2015a.
- John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-Dimensional Continuous Control Using Generalized Advantage Estimation. *CoRR*, abs/1506.02438, 2015b.
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin A. Riedmiller. Deterministic Policy Gradient Algorithms. In *Proceedings of the 31st International Conference on Machine Learning, June 21-26, 2014, Beijing, China*, 2014.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneshelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature*, 529(7587):484–489, 2016.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning - an Introduction*. Adaptive computation and machine learning. MIT Press, 1998.
- Richard S. Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems 12, November 29 - December 4, 1999, Colorado, USA*, 1999.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A Physics Engine for Model-Based Control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, October 7-12, 2012, Vilamoura, Algarve, Portugal*, 2012.
- Ronald J. Williams. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine Learning*, 8(3):229–256, 1992. doi: 10.1007/BF00992696.
- Yuhuai Wu, Elman Mansimov, Roger B Grosse, Shun Liao, and Jimmy Ba. Scalable Trust-Region Method for Deep Reinforcement Learning using Kronecker-factored Approximation. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30, 4-9 December 2017, Long Beach, CA, USA*, 2017.
- Brian D. Ziebart, J. Andrew Bagnell, and Anind K. Dey. Modeling Interaction via the Principle of Maximum Causal Entropy. In *Proceedings of the 27th International Conference on Machine Learning, June 21-24, 2010, Haifa, Israel*, 2010.

## A DERIVATIONS AND PROOFS

### A.1 DERIVATION OF SOLUTION AND DUAL FUNCTION OF GUIDE ACTOR

The solution of the optimization problem:

$$\begin{aligned} & \max_{\tilde{\pi}} \mathbb{E}_{p^\beta(\mathbf{s}), \tilde{\pi}(\mathbf{a}|\mathbf{s})} [\hat{Q}(\mathbf{s}, \mathbf{a})], \\ & \text{subject to } \mathbb{E}_{p^\beta(\mathbf{s})} [\text{KL}(\tilde{\pi}(\mathbf{a}|\mathbf{s}) || \pi_\theta(\mathbf{a}|\mathbf{s}))] \leq \epsilon, \\ & \mathbb{E}_{p^\beta(\mathbf{s})} [\text{H}(\tilde{\pi}(\mathbf{a}|\mathbf{s}))] \geq \kappa, \\ & \mathbb{E}_{p^\beta(\mathbf{s})} \tilde{\pi}(\mathbf{a}|\mathbf{s}) = 1, \end{aligned} \quad (37)$$

can be obtained by the method of Lagrange multipliers. The derivation here follows the derivation of similar optimization problems by Peters et al. (2010) and Abdolmaleki et al. (2015). The Lagrangian of this optimization problem is

$$\begin{aligned} \mathcal{L}(\tilde{\pi}, \eta, \omega, \nu) = & \mathbb{E}_{p^\beta(\mathbf{s}), \tilde{\pi}(\mathbf{a}|\mathbf{s})} [\hat{Q}(\mathbf{s}, \mathbf{a})] + \eta(\epsilon - \mathbb{E}_{p^\beta(\mathbf{s})} [\text{KL}(\tilde{\pi}(\mathbf{a}|\mathbf{s}) || \pi_\theta(\mathbf{a}|\mathbf{s}))]) \\ & + \omega(\mathbb{E}_{p^\beta(\mathbf{s})} [\text{H}(\tilde{\pi}(\mathbf{a}|\mathbf{s}))] - \kappa) + \nu(\mathbb{E}_{p^\beta(\mathbf{s})} \tilde{\pi}(\mathbf{a}|\mathbf{s}) - 1), \end{aligned} \quad (38)$$

where  $\eta$ ,  $\omega$ , and  $\nu$  are the dual variables. Then, by taking derivative of  $\mathcal{L}$  w.r.t.  $\tilde{\pi}$  we obtain

$$\partial_{\tilde{\pi}} \mathcal{L} = \mathbb{E}_{p^\beta(\mathbf{s})} \left[ \int \left( \hat{Q}(\mathbf{s}, \mathbf{a}) - (\eta + \omega) \log \tilde{\pi}(\mathbf{a}|\mathbf{s}) + \eta \log \pi_\theta(\mathbf{a}|\mathbf{s}) \right) d\mathbf{a} \right] - (\eta + \omega - \nu). \quad (39)$$

We set this derivation to zero in order to obtain

$$\begin{aligned} 0 = & \mathbb{E}_{p^\beta(\mathbf{s})} \left[ \int \left( \hat{Q}(\mathbf{s}, \mathbf{a}) - (\eta + \omega) \log \tilde{\pi}(\mathbf{a}|\mathbf{s}) + \eta \log \pi_\theta(\mathbf{a}|\mathbf{s}) \right) d\mathbf{a} \right] - (\eta + \omega - \nu) \\ = & \hat{Q}(\mathbf{s}, \mathbf{a}) - (\eta + \omega) \log \tilde{\pi}(\mathbf{a}|\mathbf{s}) + \eta \log \pi_\theta(\mathbf{a}|\mathbf{s}) - (\eta + \omega - \nu). \end{aligned} \quad (40)$$

Then the solution is given by

$$\tilde{\pi}(\mathbf{a}|\mathbf{s}) = \pi_\theta(\mathbf{a}|\mathbf{s})^{\frac{\eta}{\eta+\omega}} \exp\left(\frac{\hat{Q}(\mathbf{s}, \mathbf{a})}{\eta+\omega}\right) \exp\left(-\frac{\eta+\omega-\nu}{\eta+\omega}\right) \quad (41)$$

$$\propto \pi_\theta(\mathbf{a}|\mathbf{s})^{\frac{\eta}{\eta+\omega}} \exp\left(\frac{\hat{Q}(\mathbf{s}, \mathbf{a})}{\eta+\omega}\right). \quad (42)$$

To obtain the dual function  $g(\eta, \omega)$ , we substitute the solution to the constraint terms of the Lagrangian and this gives us

$$\begin{aligned} \mathcal{L}(\eta, \omega, \nu) = & \mathbb{E}_{p^\beta(\mathbf{s}), \tilde{\pi}(\mathbf{a}|\mathbf{s})} [\hat{Q}(\mathbf{s}, \mathbf{a})] \\ & - (\eta + \omega) \mathbb{E}_{p^\beta(\mathbf{s}), \tilde{\pi}(\mathbf{a}|\mathbf{s})} \left[ \frac{\hat{Q}(\mathbf{s}, \mathbf{a})}{\eta + \omega} + \frac{\eta}{\eta + \omega} \log \pi_\theta(\mathbf{a}|\mathbf{s}) - \frac{\eta + \omega - \nu}{\eta + \omega} \right] \\ & + \eta \mathbb{E}_{p^\beta(\mathbf{s}), \tilde{\pi}(\mathbf{a}|\mathbf{s})} [\log \pi_\theta(\mathbf{a}|\mathbf{s})] + \nu (\mathbb{E}_{p^\beta(\mathbf{s}), \tilde{\pi}(\mathbf{a}|\mathbf{s})} - 1) + \eta\epsilon - \omega\kappa. \end{aligned} \quad (43)$$

After some calculation, we obtain

$$\begin{aligned} \mathcal{L}(\eta, \omega, \nu) = & \eta\epsilon - \omega\kappa + \mathbb{E}_{p^\beta(\mathbf{s})} [\eta + \omega - \nu] \\ = & \eta\epsilon - \omega\kappa + (\eta + \omega) \mathbb{E}_{p^\beta(\mathbf{s})} \left[ \int \pi_\theta(\mathbf{a}|\mathbf{s})^{\frac{\eta}{\eta+\omega}} \exp\left(\frac{\hat{Q}(\mathbf{s}, \mathbf{a})}{\eta+\omega}\right) d\mathbf{a} \right] \\ = & g(\eta, \omega), \end{aligned} \quad (44)$$

where in the second line we use the fact that  $\exp(-\frac{\eta+\omega-\nu}{\eta+\omega})$  is the normalization term of  $\tilde{\pi}(\mathbf{a}|\mathbf{s})$ .

## A.2 PROOF OF SECOND-ORDER OPTIMIZATION IN ACTION SPACE

Firstly, we show that GAC performs second-order optimization in the action space when Taylor's approximation is performed with  $\mathbf{a}_0 = \mathbb{E}_{\pi(\mathbf{a}|\mathbf{s})}[\mathbf{a}] = \phi_\theta(\mathbf{s})$ . Recall that Taylor's approximation with  $\phi_\theta$  is given by

$$\widehat{Q}(\mathbf{s}, \mathbf{a}) = \frac{1}{2} \mathbf{a}^\top \mathbf{H}_{\phi_\theta}(\mathbf{s}) \mathbf{a} + \mathbf{a}^\top \boldsymbol{\psi}_{\phi_\theta}(\mathbf{s}) + \xi_{\phi_\theta}(\mathbf{s}), \quad (45)$$

where  $\boldsymbol{\psi}_{\phi_\theta}(\mathbf{s}) = \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_\theta(\mathbf{s})} - \mathbf{H}_{\phi_\theta}(\mathbf{s}) \phi_\theta(\mathbf{s})$ . By substituting  $\boldsymbol{\psi}_{\phi_\theta}(\mathbf{s})$  into  $\mathbf{L}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) \phi_\theta(\mathbf{s}) - \boldsymbol{\psi}_{\phi_\theta}(\mathbf{s})$ , we obtain

$$\begin{aligned} \mathbf{L}(\mathbf{s}) &= \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) \phi_\theta(\mathbf{s}) + \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_\theta(\mathbf{s})} - \mathbf{H}_{\phi_\theta}(\mathbf{s}) \phi_\theta(\mathbf{s}) \\ &= (\eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbf{H}_{\phi_\theta}(\mathbf{s})) \phi_\theta(\mathbf{s}) + \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_\theta(\mathbf{s})} \\ &= \mathbf{F}(\mathbf{s}) \phi_\theta(\mathbf{s}) + \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_\theta(\mathbf{s})}. \end{aligned} \quad (46)$$

Therefore, the mean update is equivalent to

$$\begin{aligned} \phi_+(\mathbf{s}) &= \mathbf{F}^{-1}(\mathbf{s}) \mathbf{L}(\mathbf{s}) \\ &= \phi_\theta(\mathbf{s}) + \mathbf{F}^{-1}(\mathbf{s}) \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\phi_\theta(\mathbf{s})}, \end{aligned} \quad (47)$$

which is a second-order optimization step with a curvature matrix  $\mathbf{F}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbf{H}_{\phi_\theta}(\mathbf{s})$ .

Similarly, for the case where a set of samples  $\{\mathbf{a}_0\} \sim \pi_\theta(\mathbf{a}_0|\mathbf{s}) = \mathcal{N}(\mathbf{a}_0|\phi_\theta(\mathbf{s}), \boldsymbol{\Sigma}(\mathbf{s}))$  is used to compute the averaged Taylor's approximation, we obtain

$$\mathbf{L}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) \phi_\theta(\mathbf{s}) + \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0} \right] - \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s}) \mathbf{a}_0(\mathbf{s})]. \quad (48)$$

Then, by assuming that  $\mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s}) \mathbf{a}_0(\mathbf{s})] = \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0] \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{a}_0]$ , we obtain

$$\begin{aligned} \mathbf{L}(\mathbf{s}) &= \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) \phi_\theta(\mathbf{s}) + \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0} \right] - \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0] \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{a}_0] \\ &= \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) \phi_\theta(\mathbf{s}) + \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0} \right] - \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0] \phi_\theta(\mathbf{s}) \\ &= (\eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s})]) \phi_\theta(\mathbf{s}) + \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0} \right] \\ &= \mathbf{F}(\mathbf{s}) \phi_\theta(\mathbf{s}) + \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0} \right]. \end{aligned} \quad (49)$$

Therefore, we have a second-order optimization step

$$\phi_+(\mathbf{s}) = \phi_\theta(\mathbf{s}) + \mathbf{F}^{-1}(\mathbf{s}) \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} \left[ \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_0} \right], \quad (50)$$

where  $\mathbf{F}^{-1}(\mathbf{s}) = \eta^* \boldsymbol{\Sigma}_\theta^{-1}(\mathbf{s}) - \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s})]$  is a curvature matrix. As described in the main paper, this interpretation is only valid when the equality  $\mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0(\mathbf{s}) \mathbf{a}_0(\mathbf{s})] = \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{H}_0] \mathbb{E}_{\pi_\theta(\mathbf{a}_0|\mathbf{s})} [\mathbf{a}_0]$  holds. While this equality does not hold in general, it holds when only one sample  $\mathbf{a}_0 \sim \pi_\theta(\mathbf{a}_0|\mathbf{s})$  is used. Nonetheless, we can still use the expectation of Taylor's approximation to perform policy update regardless of this assumption.

## A.3 PROOF OF GAUSS-NEWTON APPROXIMATION

Let  $f(\mathbf{s}, \mathbf{a}) = \exp(\widehat{Q}(\mathbf{s}, \mathbf{a}))$ , then the Hessian  $\mathbf{H}(\mathbf{s}) = \nabla_{\mathbf{a}}^2 \widehat{Q}(\mathbf{s}, \mathbf{a})$  can be expressed as

$$\begin{aligned} \mathbf{H}(\mathbf{s}) &= \nabla_{\mathbf{a}} [\nabla_{\mathbf{a}} \log f(\mathbf{s}, \mathbf{a})] \\ &= \nabla_{\mathbf{a}} [\nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1}] \\ &= \nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}) (\nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a})^{-1})^\top + \nabla_{\mathbf{a}}^2 f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1} \\ &= \nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}) \nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a})^{-1} (\nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}))^\top + \nabla_{\mathbf{a}}^2 f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1} \\ &= -\nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-2} (\nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}))^\top + \nabla_{\mathbf{a}}^2 f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1} \\ &= -(\nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1}) (\nabla_{\mathbf{a}} f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1})^\top + \nabla_{\mathbf{a}}^2 f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1} \\ &= -\nabla_{\mathbf{a}} \log f(\mathbf{s}, \mathbf{a}) \nabla_{\mathbf{a}} \log f(\mathbf{s}, \mathbf{a})^\top + \nabla_{\mathbf{a}}^2 f(\mathbf{s}, \mathbf{a}) f(\mathbf{s}, \mathbf{a})^{-1} \\ &= -\nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a}) \nabla_{\mathbf{a}} \widehat{Q}(\mathbf{s}, \mathbf{a})^\top + \nabla_{\mathbf{a}}^2 \exp(\widehat{Q}(\mathbf{s}, \mathbf{a})) \exp(-\widehat{Q}(\mathbf{s}, \mathbf{a})), \end{aligned} \quad (51)$$

which concludes the proof.

Beside Gauss-Newton approximation, an alternative approach is to impose a special structure on  $\hat{Q}(s, \mathbf{a})$  so that Hessians are always negative semi-definite. In literature, there exists two special structures that satisfies this requirement.

**Normalized advantage function (NAF)** (Gu et al., 2016): NAF represents the critic by a quadratic function with a negative curvature:

$$\hat{Q}^{\text{NAF}}(s, \mathbf{a}) = \frac{1}{2}(\mathbf{a} - \mathbf{b}(s))^{\top} \mathbf{W}(s)(\mathbf{a} - \mathbf{b}(s)) + V(s), \quad (52)$$

where a negative-definite matrix-valued function  $\mathbf{W}(s)$ , a vector-valued function  $\mathbf{b}(s)$  and a scalar-valued function  $V(s)$  are parameterized functions whose their parameters are learned by policy evaluation methods such as Q-learning (Sutton & Barto, 1998). With NAF, negative definite Hessians can be simply obtained as  $\mathbf{H}(s) = \mathbf{W}(s)$ . However, a significant disadvantage of NAF is that it assumes the action-value function is quadratic regardless of states and this is generally not true for most reward functions. Moreover, the Hessians become action-independent even though the critic is a function of actions.

**Input convex neural networks (ICNNs)** (Amos et al., 2017): ICNNs are neural networks with special structures which make them convex w.r.t. their inputs. Since Hessians of concave functions are always negative semi-definite, we may use ICNNs to represent a negative critic and directly use its Hessians. However, similarly to NAF, ICNNs implicitly assume that the action-value function is concave w.r.t. actions regardless of states and this is generally not true for most reward functions.

#### A.4 GRADIENT OF THE LOSS FUNCTIONS FOR LEARNING PARAMETERIZED ACTOR

We first consider the weight mean-squared-error loss function where the guide actor is  $\mathcal{N}(\mathbf{a}|\phi_+(s), \Sigma_+(s))$  and the current actor is  $\mathcal{N}(\mathbf{a}|\phi_{\theta}(s), \Sigma_{\theta}(s))$ . Taylor’s approximation of  $\hat{Q}(s, \mathbf{a})$  at  $\mathbf{a}_0 = \phi_{\theta}(s)$  is

$$\hat{Q}(s, \mathbf{a}) = \frac{1}{2} \mathbf{a}^{\top} \mathbf{H}_{\phi_{\theta}}(s) \mathbf{a} + \mathbf{a}^{\top} \psi_{\phi_{\theta}}(s) + \xi_{\phi_{\theta}}(s). \quad (53)$$

By assuming that  $\mathbf{H}_{\phi_{\theta}}(s)$  is strictly negative definite<sup>5</sup>, we can take a derivative of this approximation w.r.t.  $\mathbf{a}$  and set it to zero to obtain  $\mathbf{a} = \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a}) - \mathbf{H}_{\phi_{\theta}}^{-1}(s) \psi_{\phi_{\theta}}(s)$ . Replacing  $\mathbf{a}$  by  $\phi_{\theta}(s)$  and  $\phi_+(s)$  yields

$$\phi_{\theta}(s) = \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(s)} - \mathbf{H}_{\phi_{\theta}}^{-1}(s) \psi_{\phi_{\theta}}(s), \quad (54)$$

$$\phi_+(s) = \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_+(s)} - \mathbf{H}_{\phi_{\theta}}^{-1}(s) \psi_{\phi_{\theta}}(s). \quad (55)$$

Recall that the weight mean-squared-error is defined as

$$\mathcal{L}^{\text{W}}(\theta) = \mathbb{E}_{p^{\beta}(s)} \left[ \|\phi_{\theta}(s) - \phi_+(s)\|_{\mathbf{F}(s)}^2 \right]. \quad (56)$$

<sup>5</sup>This can be done by subtracting a small positive value to the diagonal entries of Gauss-Newton approximation.

Then, we consider its gradient w.r.t.  $\theta$  as follows:

$$\begin{aligned}
\nabla_{\theta} \mathcal{L}^W(\theta) &= 2\mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \mathbf{F}(s) (\phi_{\theta}(s) - \phi_{+}(s)) \right] \\
&= 2\mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) (\eta^* \Sigma_{\theta}^{-1}(s) - \mathbf{H}_{\phi_{\theta}}(s)) (\phi_{\theta}(s) - \phi_{+}(s)) \right] \\
&= 2\mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) (\eta^* \Sigma_{\theta}^{-1}(s) - \mathbf{H}_{\phi_{\theta}}(s)) \right. \\
&\quad \left. \times \left( \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(s)} - \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} \right) \right] \\
&= 2\eta^* \mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \Sigma_{\theta}^{-1}(s) \mathbf{H}_{\phi_{\theta}}^{-1}(s) \left( \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(s)} - \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} \right) \right] \\
&\quad + 2\mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \left( \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} - \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(s)} \right) \right] \\
&= -2\mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(s)} \right] + 2\mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} \right] \\
&\quad + 2\eta^* \mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \Sigma^{-1}(s) \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(s)} \right] \\
&\quad - 2\eta^* \mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \Sigma^{-1}(s) \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} \right]. \tag{57}
\end{aligned}$$

This concludes the proof for the gradient in Eq.(28). Note that we should not directly replace the mean functions in the weight mean-square-error by Eq.(54) and Eq.(55) before expanding the gradient. This is because the analysis would require computing the gradient w.r.t.  $\theta$  inside the Hessians and this is not trivial. Moreover, when we perform gradient descent in practice, the mean  $\phi_{+}(s)$  is considered as a constant w.r.t.  $\theta$  similarly to an output function in supervised learning.

The gradient for the mean-squared-error can be derived similarly. Let the mean-squared-error be defined as

$$\mathcal{L}^M(\theta) = \frac{1}{2} \mathbb{E}_{p^{\beta}(s)} \left[ \|\phi_{\theta}(s) - \phi_{+}(s)\|_2^2 \right]. \tag{58}$$

Its gradient w.r.t.  $\theta$  is given by

$$\begin{aligned}
\nabla_{\theta} \mathcal{L}^M(\theta) &= \mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) (\phi_{\theta}(s) - \phi_{+}(s)) \right] \\
&= \mathbb{E}_{p^{\beta}(s)} \left[ \nabla_{\theta} \phi_{\theta}(s) \mathbf{H}^{-1}(s) \left( \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{\theta}(s)} - \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} \right) \right], \tag{59}
\end{aligned}$$

which concludes the proof for the gradient in Eq.(30).

To show that  $\nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} = 0$  when  $\eta^* = 0$ , we directly substitute  $\eta^* = 0$  into  $\phi_{+}(s) = (\eta^* \Sigma^{-1}(s) - \mathbf{H}_{\phi_{\theta}}(s))^{-1} \left( \eta^* \Sigma^{-1}(s) \phi_{\theta}(s) + \psi_{\phi_{\theta}}(s) \right)$  and this yield

$$\phi_{+}(s) = -\mathbf{H}_{\phi_{\theta}}^{-1}(s) \psi_{\phi_{\theta}}(s). \tag{60}$$

Since  $\phi_{+}(s) = \mathbf{H}_{\phi_{\theta}}^{-1}(s) \nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} - \mathbf{H}_{\phi_{\theta}}^{-1}(s) \psi_{\phi_{\theta}}(s)$  from Eq.(55) and the Hessians are non-zero, it has to be that  $\nabla_{\mathbf{a}} \hat{Q}(s, \mathbf{a})|_{\mathbf{a}=\phi_{+}(s)} = 0$ . This is intuitive since without the KL constraint, the mean of the guide actor always be at the optima of the second-order Taylor's approximation and the gradients are zero at the optima.

#### A.5 RELATION TO Q-LEARNING WITH NORMALIZED ADVANTAGE FUNCTION

The *normalized advantage function* (NAF) (Gu et al., 2016) is defined as

$$\hat{Q}^{\text{NAF}}(s, \mathbf{a}) = \frac{1}{2} (\mathbf{a} - \mathbf{b}(s))^{\top} \mathbf{W}(s) (\mathbf{a} - \mathbf{b}(s)) + V(s), \tag{61}$$

where  $\mathbf{W}(s)$  is a negative definite matrix. Gu et al. (2016) proposed to perform Q-learning with NAF by using the fact that  $\arg\max_{\mathbf{a}} \hat{Q}^{\text{NAF}}(s, \mathbf{a}) = \mathbf{b}(s)$  and  $\max_{\mathbf{a}} \hat{Q}^{\text{NAF}}(s, \mathbf{a}) = V(s)$ .

Here, we show that GAC includes the Q-learning method by Gu et al. (2016) as its special case. This can be shown by using NAF as a critic instead of performing Taylor's approximation of the critic.

Firstly, we expand the quadratic term of NAF as follows:

$$\begin{aligned}\widehat{Q}^{\text{NAF}}(\mathbf{s}, \mathbf{a}) &= \frac{1}{2}(\mathbf{a} - \mathbf{b}(\mathbf{s}))^\top \mathbf{W}(\mathbf{s})(\mathbf{a} - \mathbf{b}(\mathbf{s})) + V(\mathbf{s}) \\ &= \frac{1}{2}\mathbf{a}^\top \mathbf{W}(\mathbf{s})\mathbf{a} - \mathbf{a}^\top \mathbf{W}(\mathbf{s})\mathbf{b}(\mathbf{s}) + \frac{1}{2}\mathbf{b}(\mathbf{s})^\top \mathbf{W}(\mathbf{s})\mathbf{b}(\mathbf{s}) + V(\mathbf{s}) \\ &= \frac{1}{2}\mathbf{a}^\top \mathbf{W}(\mathbf{s})\mathbf{a} + \mathbf{a}^\top \boldsymbol{\psi}(\mathbf{s}) + \xi(\mathbf{s}),\end{aligned}\tag{62}$$

where  $\boldsymbol{\psi}(\mathbf{s}) = -\mathbf{W}(\mathbf{s})\mathbf{b}(\mathbf{s})$  and  $\xi(\mathbf{s}) = \frac{1}{2}\mathbf{b}(\mathbf{s})^\top \mathbf{W}(\mathbf{s})\mathbf{b}(\mathbf{s}) + V(\mathbf{s})$ . By substituting the quadratic model obtained by NAF into the GAC framework, the guide actor is now given by  $\tilde{\pi}(\mathbf{a}|\mathbf{s}) = \mathcal{N}(\mathbf{a}|\phi_+(\mathbf{s}), \Sigma_+(\mathbf{s}))$  with

$$\phi_+(\mathbf{s}) = (\eta^* \Sigma^{-1}(\mathbf{s}) - \mathbf{W}(\mathbf{s}))^{-1}(\eta^* \Sigma^{-1}(\mathbf{s})\phi_\theta(\mathbf{s}) - \mathbf{W}(\mathbf{s})\mathbf{b}(\mathbf{s}))\tag{63}$$

$$\Sigma_+(\mathbf{s}) = (\eta^* + \omega^*)(\eta^* \Sigma^{-1}(\mathbf{s}) - \mathbf{W}(\mathbf{s})).\tag{64}$$

To obtain Q-learning with NAF, we set  $\eta^* = 0$ , i.e., we perform a greedy maximization where the KL upper-bound approaches infinity, and this yields

$$\begin{aligned}\phi_+(\mathbf{s}) &= -\mathbf{W}(\mathbf{s})^{-1}(-\mathbf{W}(\mathbf{s})\mathbf{b}(\mathbf{s})) \\ &= \mathbf{b}(\mathbf{s}),\end{aligned}\tag{65}$$

which is the policy obtained by performing Q-learning with NAF. Thus, NAF with Q-learning is a special case of GAC if Q-learning is also used in GAC to learn the critic.

## B PSEUDO-CODE OF GAC

The pseudo-code of GAC is given in Algorithm 1. The source code is available at <https://github.com/voot-t/guide-actor-critic>.

## C EXPERIMENT DETAILS

### C.1 IMPLEMENTATION

We try to follow the network architecture proposed by the authors of each baseline method as close as possible. For GAC and DDPG, we use neural networks with two hidden layers for the actor network and the critic network. For both networks the first layer has 400 hidden units and the second layer has 300 units. For NAF, we use neural networks with two hidden layers to represent each of the functions  $\mathbf{b}(\mathbf{s})$ ,  $\mathbf{W}(\mathbf{s})$  and  $V(\mathbf{s})$  where each layer has 200 hidden units. All hidden units use the relu activation function except for the output of the actor network where we use the tanh activation function to bound actions. We use the Adam optimizer (Kingma & Ba, 2014) with learning rate 0.001 and 0.0001 for the critic network and the actor network, respectively. The moving average step for target networks is set to  $\tau = 0.001$ . The maximum size of the replay buffer is set to 1000000. The mini-batches size is set to  $N = 256$ . The weights of the actor and critic networks are initialized as described by Glorot & Bengio (2010), except for the output layers where the initial weights are drawn uniformly from  $U(-0.003, 0.003)$ , as described by Lillicrap et al. (2015). The initial covariance  $\Sigma$  in GAC is set to be an identity matrix. DDPG and QNAF use the OU-process with noise parameters  $\theta = 0.15$  and  $\sigma = 0.2$  for exploration.

For TRPO, we use the implementation publicly available at <https://github.com/openai/baselines>. We also use the provided network architecture and hyper-parameters except the batch size where we use 1000 instead of 1024 since this is more suitable in our test setup.

For GAC, the KL upper-bound is fixed to  $\epsilon = 0.0001$ . The entropy lower-bound  $\kappa$  is adjusted heuristically by

$$\kappa = \max(0.99(E - E_0) + E_0, E_0),\tag{71}$$

where  $E \approx \mathbb{E}_{p_\beta(\mathbf{s})} [\mathbb{H}(\pi_\theta(\mathbf{a}|\mathbf{s}))]$  denotes the expected entropy of the current policy and  $E_0$  denotes the entropy of a base policy  $\mathcal{N}(\mathbf{a}|\mathbf{0}, 0.01\mathbf{I})$ . This heuristic ensures that the lower-bound gradually decreases but the lower-bound cannot be too small. We apply this heuristic update once every 5000 training steps. The dual function is minimized by the sequential least-squares quadratic programming (SLSQP) method with an initial values  $\eta = 0.05$  and  $\omega = 0.05$ . The number of samples for computing the target critic value is  $M = 10$ .

**Algorithm 1** Guide actor critic

---

```

1: Input: Initial actor  $\pi_{\theta}(\mathbf{a}|s) = \mathcal{N}(\mathbf{a}|\phi_{\theta}(s), \Sigma)$ , critic  $\widehat{Q}_{\nu}(s, \mathbf{a})$ , target critic network
 $\widehat{Q}_{\bar{\nu}}(s, \mathbf{a})$ , KL bound  $\epsilon$ , entropy bound  $\kappa$ , learning rates  $\alpha_1, \alpha_2$ , and data buffer  $\mathcal{D} = \emptyset$ .
2: for  $t = 1, \dots, T_{\max}$  do
3:   procedure COLLECT TRANSITION SAMPLE
4:     Observe state  $s_t$  and sample action  $\mathbf{a}_t \sim \mathcal{N}(\mathbf{a}|\phi_{\theta}(s_t), \Sigma)$ .
5:     Execute  $\mathbf{a}_t$ , receive reward  $r_t$  and next state  $s'_t$ .
6:     Add transition  $\{(s_t, \mathbf{a}_t, r_t, s'_t)\}$  to  $\mathcal{D}$ .
7:   end procedure
8:   procedure LEARN
9:     Sample  $N$  mini-batch samples  $\{(s_n, \mathbf{a}_n, r_n, s'_n)\}_{n=1}^N$  uniformly from  $\mathcal{D}$ .
10:    procedure UPDATE CRITIC
11:      Sample actions  $\{\mathbf{a}'_{n,m}\}_{m=1}^M \sim \mathcal{N}(\mathbf{a}|\phi_{\theta}(s_n), \Sigma)$  for each  $s_n$ .
12:      Compute  $y_n$ , update  $\nu$  by, e.g., Adam, and update  $\bar{\nu}$  by moving average:

$$y_n = r_n + \gamma \frac{1}{M} \sum_{m=1}^M \widehat{Q}_{\bar{\nu}}(s'_n, \mathbf{a}'_{n,m}), \quad (66)$$


$$\nu \leftarrow \nu - \alpha_1 \frac{1}{N} \sum_{n=1}^N \nabla_{\nu} \left( \widehat{Q}_{\nu}(s_n, \mathbf{a}_n) - y_n \right)^2, \quad (67)$$


$$\bar{\nu} \leftarrow \tau \nu + (1 - \tau) \bar{\nu}. \quad (68)$$

13:    end procedure
14:    procedure LEARN GUIDE ACTOR
15:      Compute  $\mathbf{a}_{n,0}$  for each  $s_n$  by  $\mathbf{a}_{n,0} = \phi_{\theta}(s_n)$  or  $\mathbf{a}_{n,0} \sim \mathcal{N}(\mathbf{a}|\phi_{\theta}(s_n), \Sigma)$ .
16:      Compute  $\mathbf{g}_0(s) = \nabla_{\mathbf{a}} \widehat{Q}(s_n, \mathbf{a})|_{\mathbf{a}=\mathbf{a}_{n,0}}$  and  $\mathbf{H}_0(s_s) = -\mathbf{g}_0(s_n) \mathbf{g}_0(s_n)^{\top}$ .
17:      Solve for  $(\eta^*, \omega^*) = \operatorname{argmin}_{\eta>0, \omega>0} \widehat{g}(\eta, \omega)$  by a non-linear optimization method.
18:      Compute the guide actor  $\tilde{\pi}(\mathbf{a}|s_n) = \mathcal{N}(\mathbf{a}|\phi_{+}(s_n), \Sigma_{+}(s_n))$  for each  $s_n$ .
19:    end procedure
20:    procedure UPDATE PARAMETERIZED ACTOR
21:      Update policy parameter by, e.g., Adam, to minimize the MSE:

$$\theta \leftarrow \theta - \alpha_2 \frac{1}{N} \sum_{n=1}^N \nabla_{\theta} \|\phi_{\theta}(s_n) - \phi_{+}(s_n)\|_2^2. \quad (69)$$

22:      Update policy covariance by averaging the guide covariances:

$$\Sigma \leftarrow \frac{1}{N} \Sigma_{+}(s_n). \quad (70)$$

23:    end procedure
24:  end procedure
25: end for
26: Output: Learned actor  $\pi_{\theta}(\mathbf{a}|s)$ .

```

---

## C.2 ENVIRONMENTS AND RESULTS

We perform experiments on the OpenAI gym platform (Brockman et al., 2016) with Mujoco Physics simulator (Todorov et al., 2012) where all environments are v1. We use the state space, action space and the reward function as provided and did not perform any normalization or gradient clipping. The maximum time horizon in each episode is set to 1000. The discount factor  $\gamma = 0.99$  is only used for learning and the test returns are computed without it.

Experiments are repeated for 10 times with different random seeds. The total computation time are reported in Table 1. The figures below show the results averaged over 10 trials. The y-axis indicates the averaged test returns where the test returns in each trial are computed once every 5000 training time steps by executing 10 test episodes without exploration. The error bar indicates standard error.

Table 1: The total computation time for training the policy for 1 million steps (0.1 million steps for the Invert-Pendulum task). The mean and standard error are computed over 10 trials with the unit in hours. TRPO is not included since it performs a lesser amount of update using batch data samples.

Task	GAC-1	GAC-0	DDPG	QNAF
Inv-Pend.	1.13(0.09)	0.80(0.04)	0.45(0.02)	0.40(0.03)
Inv-Double-Pend.	15.56(0.93)	15.67(0.77)	9.04(0.29)	7.47(0.22)
Reacher	17.23(0.87)	12.64(0.38)	10.67(0.37)	30.91(2.09)
Swimmer	12.43(0.74)	11.94(0.74)	9.61(0.52)	32.44(2.21)
Half-Cheetah	29.82(1.76)	32.64(1.41)	10.13(0.41)	27.94(2.37)
Ant	37.84(1.80)	40.57(3.06)	9.75(0.37)	27.09(0.94)
Hopper	18.99(1.06)	14.22(0.56)	8.74(0.45)	26.48(1.42)
Walker2D	33.42(0.96)	31.71(1.84)	7.92(0.11)	26.94(1.99)
Humanoid	111.07(2.99)	106.80(6.80)	13.90(1.60)	30.43(2.21)

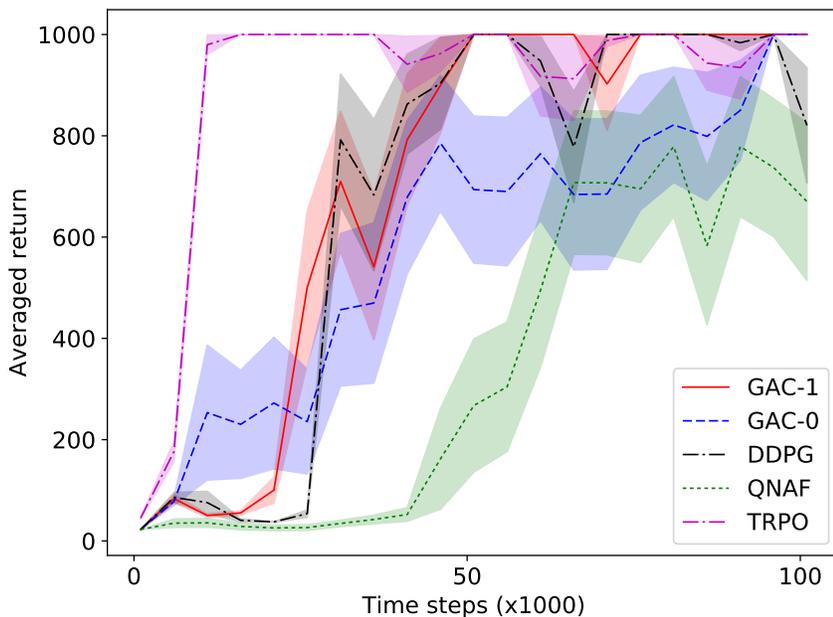


Figure 2: Performance averaged over 10 trials on the Inverted Pendulum task.

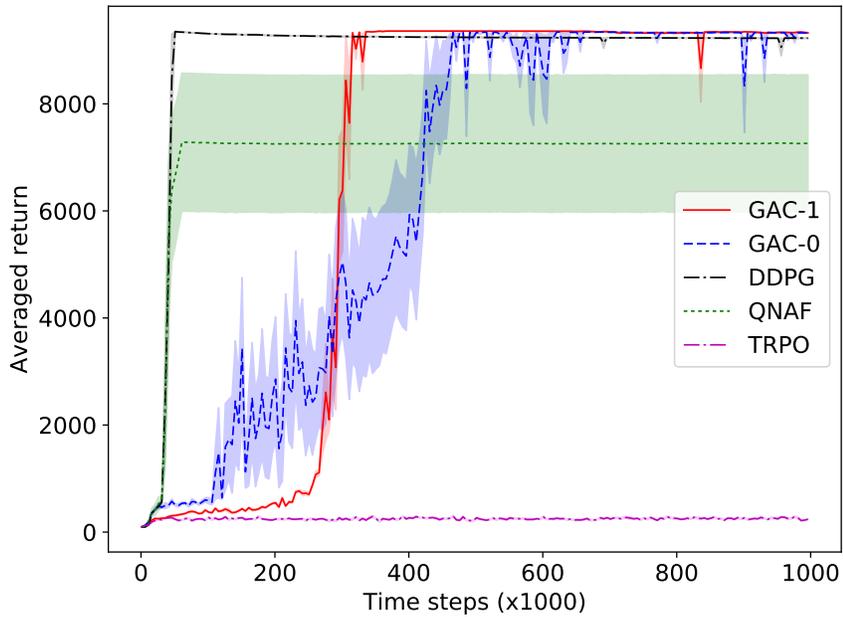


Figure 3: Performance averaged over 10 trials on the Inverted Double Pendulum task.

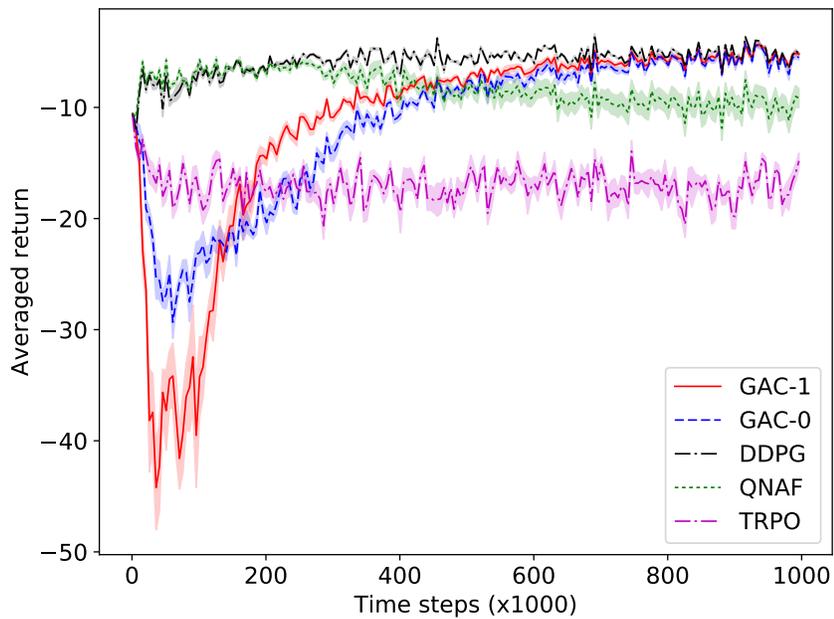


Figure 4: Performance averaged over 10 trials on the Reacher task.

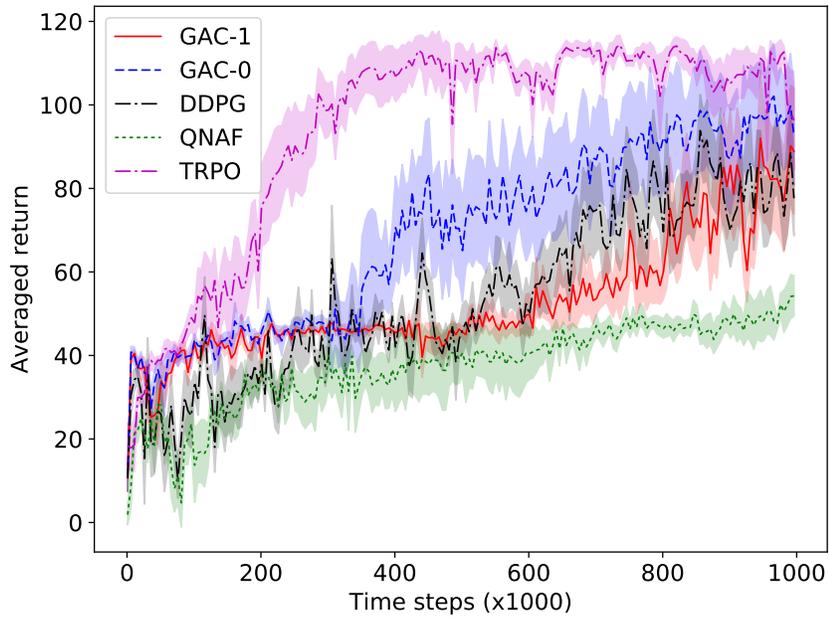


Figure 5: Performance averaged over 10 trials on the Swimmer task.

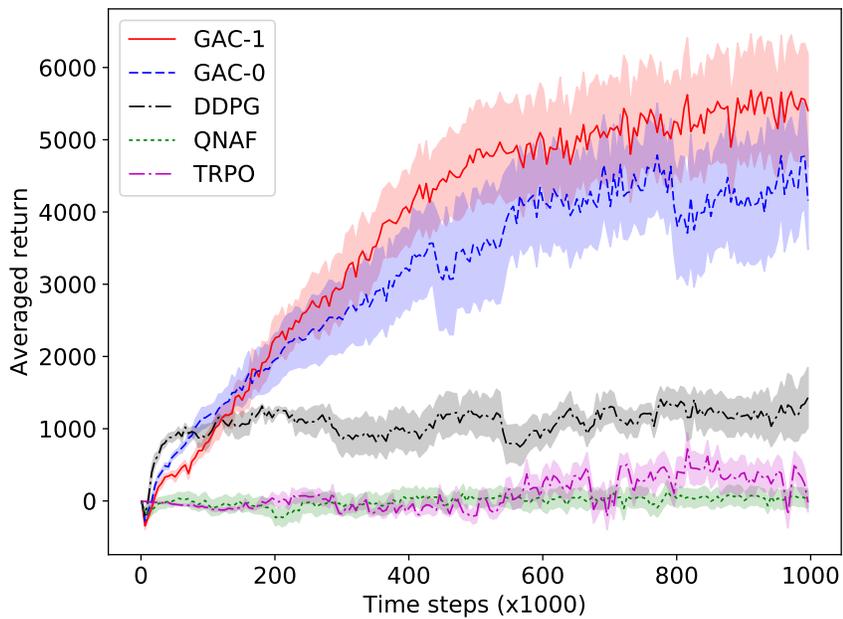


Figure 6: Performance averaged over 10 trials on the Half-Cheetah task.

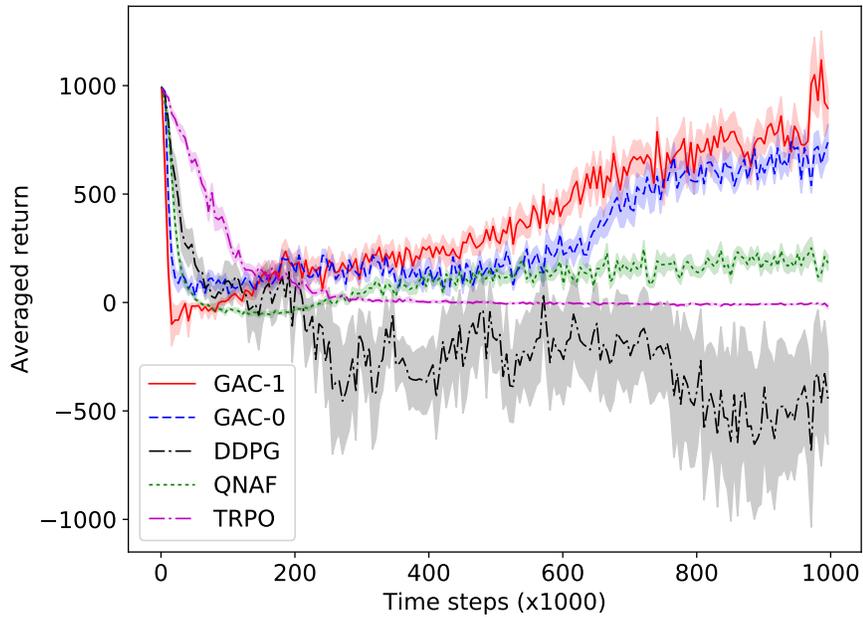


Figure 7: Performance averaged over 10 trials on the Ant task.

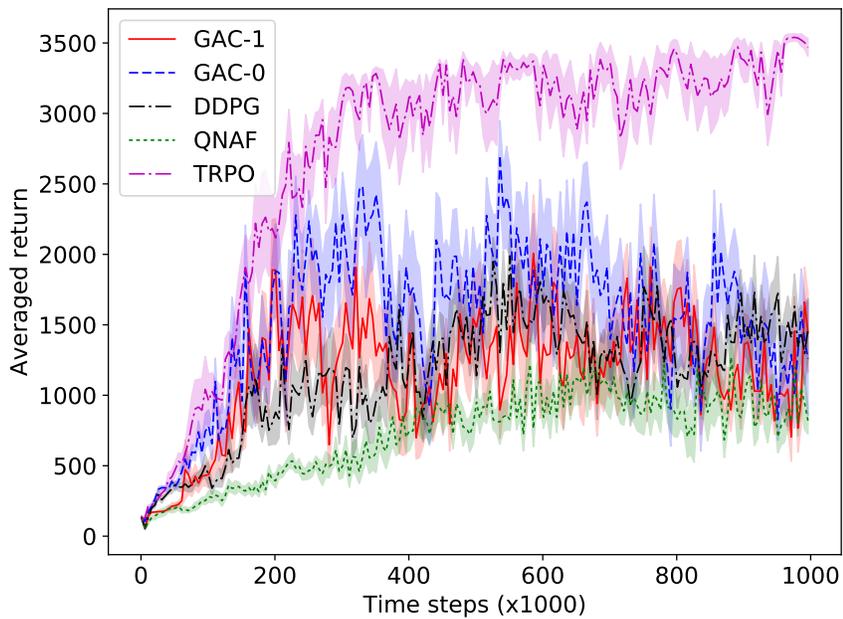


Figure 8: Performance averaged over 10 trials on the Hopper task.

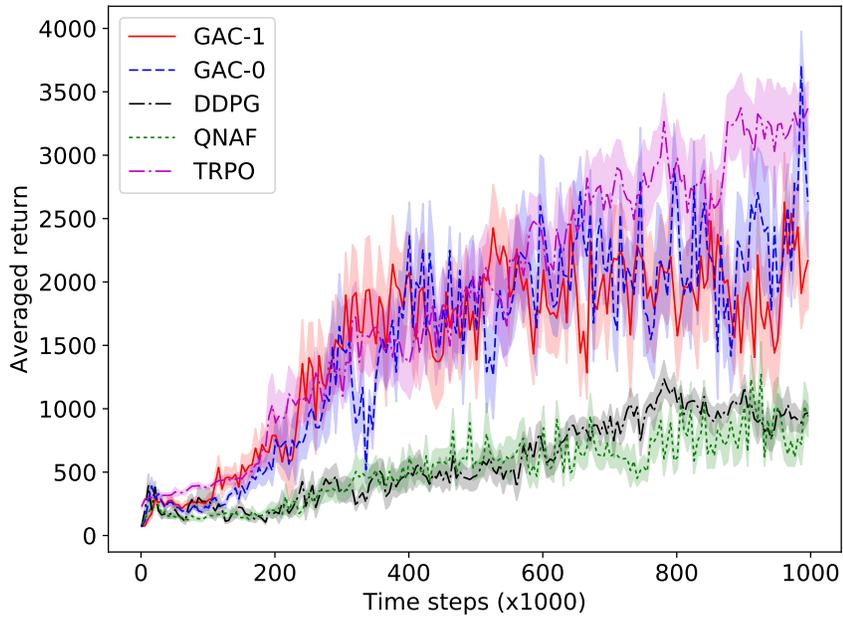


Figure 9: Performance averaged over 10 trials on the Walker2D task.

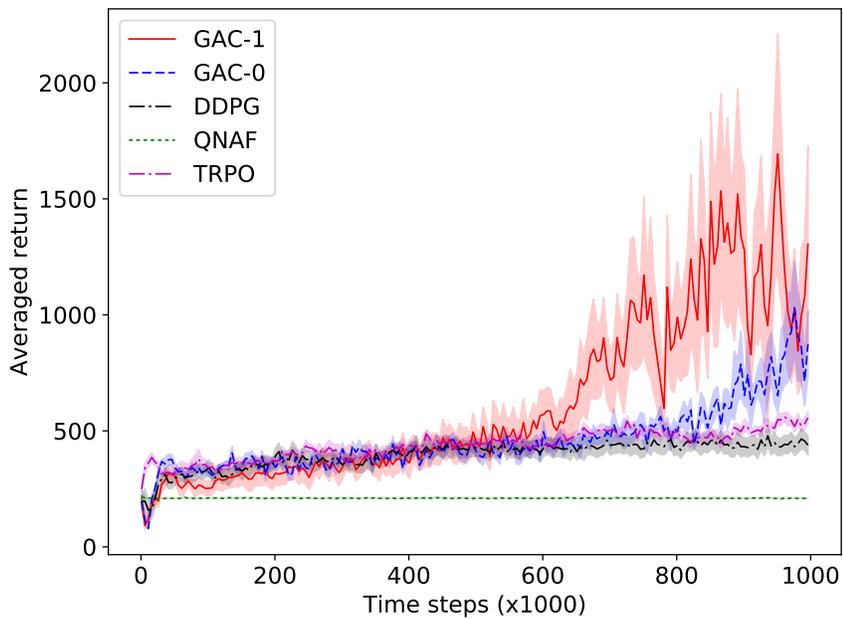


Figure 10: Performance averaged over 10 trials on the Humanoid task.