

ACE: An Actor Ensemble Algorithm for Continuous Control with Tree Search

Shangdong Zhang,^{*1} Hengshuai Yao²

¹ Department of Computing Science, University of Alberta

² Reinforcement Learning for Autonomous Driving Lab, Noah's Ark Lab, Huawei
shangdong.zhang@ualberta.ca, hengshuai.yao@huawei.com

Abstract

In this paper, we propose an actor ensemble algorithm, named ACE, for continuous control with a deterministic policy in reinforcement learning. In ACE, we use actor ensemble (i.e., multiple actors) to search the global maxima of the critic. Besides the ensemble perspective, we also formulate ACE in the option framework by extending the option-critic architecture with deterministic intra-option policies, revealing a relationship between ensemble and options. Furthermore, we perform a look-ahead tree search with those actors and a learned value prediction model, resulting in a refined value estimation. We demonstrate a significant performance boost of ACE over DDPG and its variants in challenging physical robot simulators.

Introduction

In this paper, we propose an actor ensemble algorithm, named ACE, for continuous control in reinforcement learning (RL). In continuous control, a deterministic policy (Silver et al. 2014) is a recent approach, which is a mapping from state to action. In contrast, a stochastic policy is a mapping from state to a probability distribution over the actions.

Recently, neural networks has achieved great success as function approximators in various challenging domains (Tesauro 1995; Mnih et al. 2015; Silver et al. 2016). A deterministic policy parameterized by a neural network is usually trained via gradient ascent to maximize the critic, which is a state-action value function parameterized by a neural network (Silver et al. 2014; Lillicrap et al. 2015; Barth-Maron et al. 2018). However, gradient ascent can be easily trapped by local maxima during the search for the global maxima. We utilize the ensemble technique to mitigate this issue. We train multiple actors (i.e., deterministic policies) in parallel, and each actor has a different initialization. In this way, each actor is in charge of maximizing the state-action value function in a local area. Different actors may be trapped in different local maxima. By considering the maximum state-action value of the actions proposed by all the actors, we are more likely to find the global maxima of the state-action value function than a single actor.

ACE fits in with the option framework (Sutton, Precup, and Singh 1999). First, each option has its intra-option policy, which maximizes the return in a certain area of the state space. Similarly, an actor in ACE maximizes the critic in a certain area of the domain of the critic. It may be difficult for a single actor to maximize the critic in the whole domain due to the complexity of the manifold of the critic. However, in contrast, the job for action search is easier if we ask an actor to find the best action in a local neighborhood of the action dimension. Second, we chain the outputs of all the actors to the critic, enabling a selection over the locally optimal action values. In this way, the critic works similar to the policy over options in the option framework. We quantify this similarity between ensemble and options by extending the option-critic architecture (OCA, Bacon, Harb, and Precup 2017) with deterministic intra-option policies. Particularly, we provide the Deterministic Intra-option Policy Gradient theorem, based on which we show the actor ensemble in ACE is a special case of the general option-critic framework.

To make the state-action value function more accurate, which is essential in the actor selection, we perform a look-ahead tree search with the multiple actors. The look-ahead tree search has achieved great success in various discrete action problems (Knuth and Moore 1975; Browne et al. 2012; Silver et al. 2016; Oh, Singh, and Lee 2017; Farquhar et al. 2018). Recently, look-ahead tree search was extended to continuous-action problems. For example, Mansley, Weinstein, and Littman (2011) combined planning with adaptive discretization of a continuous action space, resulting in a performance boost in continuous bandit problems. Nitti, Belle, and De Raedt (2015) utilized probability programming in planning in continuous action space. Yee, Lisỳ, and Bowling (2016) used kernel regression to generalize the values between explored actions and unexplored actions, resulting in a new Monte Carlo tree search algorithm. However, to our best knowledge, a general tree search algorithm for continuous-action problems is a gap. In ACE, we use the multiple actors as meta-actions to perform a tree search with the help of a learned value prediction model (Oh, Singh, and Lee 2017; Farquhar et al. 2018).

We demonstrate the superiority of ACE over DDPG and its variants empirically in Roboschool¹, a challenging phys-

^{*}Work done during an internship at Huawei
Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://github.com/openai/roboschool>

ical robot environment.

In the rest of this paper, we first present some preliminaries of RL. Then we detail ACE and show some empirical results, after which we discuss some related work, followed by closing remarks.

Preliminaries

We consider a Markov Decision Process (MDP) which consists of a state space \mathcal{S} , an action space \mathcal{A} , a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, a transition kernel $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, and a discount factor $\gamma \in [0, 1]$. We use $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ to denote a stochastic policy. At each time step t , an agent is at state s_t and selects an action $a_t \sim \pi(\cdot|s_t)$. Then the environment gives the reward r_{t+1} and leads the agent to the next state s_{t+1} according $p(\cdot|s_t, a_t)$. We use q^π to denote the state action value of a policy π , i.e., $q^\pi(s, a) \doteq \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r_{t+1} | s_0 = s, a_0 = a]$. In an RL problem, we are usually interested in finding an optimal policy π^* s.t. $q^{\pi^*}(s, a) \geq q^\pi(s, a)$ for $\forall(\pi, s, a)$. All the optimal policies share the same optimal state action value function q^* , which is the unique fixed point of the Bellman optimality operator \mathcal{T} ,

$$\mathcal{T}Q(s, a) \doteq r(s, a) + \gamma \mathbb{E}_{s' \sim p(\cdot|s, a)} [\max_{a'} Q(s', a')] \quad (1)$$

where we use Q to indicate an estimation of q^* . With tabular representation, Q-learning (Watkins and Dayan 1992) is a commonly used method to find this fixed point. The per step update is

$$\Delta Q(s_t, a_t) \propto r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \quad (2)$$

Recently, Mnih et al. (2015) used a deep convolutional neural network θ to parameterize an estimation Q , resulting in the Deep-Q-Network (DQN). At each time step t , DQN samples the transition $(s_t, a_t, r_{t+1}, s_{t+1})$ from a replay buffer (Lin 1992) and performs stochastic gradient descent to minimize the loss

$$\frac{1}{2} (r_{t+1} + \gamma \max_a Q_\theta(s_{t+1}, a) - Q_\theta(s_t, a_t))^2$$

where θ^- is a target network (Mnih et al. 2015), which is a copy of θ and is synchronized with θ periodically.

Continuous Action Control

In continuous control problems, it is not straightforward to apply Q-learning. The basic idea of the deterministic policy algorithms (Silver et al. 2014) is to use a deterministic policy $\mu : \mathcal{S} \rightarrow \mathcal{A}$ to approximate the greedy action $\arg \max_a Q(s, a)$. The deterministic policy is trained via gradient ascent according to chain rule. The gradient per step is

$$\nabla_{\theta^\mu} Q(s_t, \mu(s_t)) = \nabla_a Q(s, a)|_{s=s_t, a=\mu(s_t)} \nabla_{\theta^\mu} \mu(s)|_{s=s_t} \quad (3)$$

where we assume μ is parameterized by θ^μ . With this actor μ , we are able to update the state action value function Q as usual. To be more specific, assuming Q is parameterized by

θ^Q , we perform gradient decent to minimize the following loss

$$\frac{1}{2} (r_{t+1} + \gamma Q(s_{t+1}, \mu(s_{t+1})) - Q(s_t, a_t))^2$$

Recently, Lillicrap et al. (2015) used neural networks to parameterize μ and Q , resulting in the Deep Deterministic Policy Gradient (DDPG) algorithm. DDPG is an off-policy control algorithm with experience replay and a target network.

In the on-policy setting, the gradient in Equation (3) guarantees policy improvement thanks to the Deterministic Policy Gradient theorem (Silver et al. 2014). In off-policy setting, the policy improvement of this gradient is based on Off-policy Policy Gradient theorem (OPG, Degris, White, and Sutton 2012). However, Errata of Degris, White, and Sutton (2012) shows that OPG only holds for tabular function representation and certain linear function approximation. So far no policy improvement is guaranteed for DDPG. However, DDPG and its variants has gained great success in various challenging domains (Lillicrap et al. 2015; Barth-Maron et al. 2018; Tassa et al. 2018). This success may be attributed to the gradient ascent via chain rule in Equation (3).

Option

An option ω is a triple, $(\mathcal{I}_\omega, \pi_\omega, \beta_\omega)$, and we use Ω to indicate the option set. We use $\mathcal{I}_\omega \subseteq \mathcal{S}$ to denote the initiation set of ω , indicating where the option ω can be initiated. In this paper, we consider $\mathcal{I}_\omega \equiv \mathcal{S}$ for $\forall \omega$, meaning that each option can be initiated at all states. We use $\pi_\omega : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ to denote the intra-option policy of ω . Once the option ω is committed, the action selection is based on π_ω . We use $\beta_\omega : \mathcal{S} \rightarrow [0, 1]$ to denote the termination function of ω . At each time step t , the agent terminates the previous option ω_{t-1} with probability $\beta_{\omega_{t-1}}(s_t)$. In this paper, we consider the *call-and-return* option execution model (Sutton, Precup, and Singh 1999), where an agent executes an option until the option terminates.

An MDP augmented with options forms a Semi-MDP (Sutton, Precup, and Singh 1999). We use $\pi_\Omega : \mathcal{S} \times \Omega \rightarrow [0, 1]$ to denote the policy over options. We use $Q_\Omega : \mathcal{S} \times \Omega \rightarrow \mathbb{R}$ to denote the option-value function and $V_\Omega : \mathcal{S} \rightarrow \mathbb{R}$ to denote the value function of π_Ω . Furthermore, we use $U : \Omega \times \mathcal{S} \rightarrow \mathbb{R}$ to denote the option value upon arrival at the state option pair (ω, s') and $Q_U : \mathcal{S} \times \Omega \times \mathcal{A} \rightarrow \mathbb{R}$ to denote the value of executing an action a in the context of a state-option pair (s, ω) . They are related as

$$Q_\Omega(s, \omega) = \sum_a \pi_\omega(a|s) Q_U(s, \omega, a)$$

$$Q_U(s, \omega, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim p(\cdot|s, a)} U(\omega, s') \quad (4)$$

$$U(\omega, s') = (1 - \beta_\omega(s')) Q_\Omega(s', \omega) + \beta_\omega(s') V_\Omega(s') \quad (5)$$

$$V_\Omega(s) = \sum_\omega \pi_\Omega(\omega|s) Q_\Omega(s, \omega) \quad (6)$$

Bacon, Harb, and Precup (2017) proposed a policy gradient method, the Option-Critic Architecture (OCA), to learn stochastic intra-option policies $\{\pi_\omega\}$ (parameterized by θ^π)

and termination functions $\{\beta_\omega\}$ (parameterized by θ^β). The objective is to maximize the expected discounted return per episode, i.e., $Q_\Omega(s_0, \omega_0)$. Based on their Intro-option Policy Gradient Theorem and Termination Gradient Theorem, the per step updates for θ^π and θ^β are

$$\begin{aligned}\Delta\theta^\pi &\propto \nabla_{\theta^\pi} \log \pi_{\omega_t}(a_t|s_t) Q_U(s_t, \omega_t, a_t) \\ \Delta\theta^\beta &\propto -\nabla_{\theta^\beta} \beta_{\omega_{t-1}}(s_t) (Q_\Omega(s_t, \omega_{t-1}) - V_\Omega(s_t))\end{aligned}$$

And at each time step t , OCA takes one gradient descent step minimizing

$$\frac{1}{2} (g_t - Q_U(s_t, \omega_t, a_t))^2 \quad (7)$$

to update the critic Q_U , where

$$\begin{aligned}g_t &\doteq r_{t+1} + \gamma(1 - \beta_{\omega_t}(s_{t+1})) Q_\Omega(s_{t+1}, \omega_t) \\ &\quad + \gamma \beta_{\omega_t}(s_{t+1}) \max_{\omega'} Q_\Omega(s_{t+1}, \omega')\end{aligned}$$

The update target g_t is also used in Intro-option Q-learning (Sutton, Precup, and Singh 1999).

Model-based RL

In RL, a transition model usually takes as inputs a state-action pair and generates the immediate reward and the next state. A transition model reflects the dynamics of an environment and can either be given or learned. A transition model can be used to generate imaginary transitions for training the agent to increase data efficiency (Sutton 1990; Yao et al. 2009; Gu et al. 2016). A transition model can also be used to reason about the future. When making a decision, an agent can perform a look-ahead tree search (Knuth and Moore 1975; Browne et al. 2012) with a transition model to maximize the possible rewards. A look-ahead tree search can be performed with either a perfect model (Coulom 2006; Sturtevant 2008; Silver et al. 2016) or a learned model (Weber et al. 2017).

A latent state is an abstraction of the original state. A latent state is also referred to as an abstract state (Oh, Singh, and Lee 2017) or an encoded state (Farquhar et al. 2018). A latent state can be used as the input of a transition model. Correspondingly, the transition model then predicts the next latent state, instead of the original next state. A latent state is particularly useful for high dimensional state space (e.g., images), where a latent state is usually a low dimensional vector.

Recently, some works demonstrated that learning a value prediction model instead of a transition model is effective for a look-ahead tree search. For example, VPN (Oh, Singh, and Lee 2017) predicted the value of the next latent state, instead of the next latent state itself. Although this value prediction was explicitly done in two phases (predicting the next latent state and then predicting its value), the loss of predicting the next latent state was not used in training. Instead, only the loss of the value prediction for the next latent state was used. Oh, Singh, and Lee (2017) showed this value prediction model is particularly helpful for non-deterministic environments. TreeQN (Farquhar et al. 2018) adopted a similar idea, where only the outcome value of a look-ahead

tree search is grounded in the loss. Farquhar et al. (2018) showed that grounding the predicted next latent state did not bring in a performance boost. Although a value prediction model predicts much fewer information than a transition model, VPN and TreeQN demonstrated improved performance over baselines in challenging domains. This value prediction model is particularly helpful for a look-ahead tree search in non-deterministic environments. In ACE, we followed these works and built a value prediction model similar to TreeQN.

The Actor Ensemble Algorithm

As discussed earlier, it is important for the actor to maximize the critic in DDPG. Lillicrap et al. (2015) trained the actor via gradient ascent. However, gradient ascent can easily be trapped by local maxima or saddle points.

To mitigate this issue, we propose an ensemble approach. In our work, we have N actors $\{\mu_1, \dots, \mu_N\}$. At each time step t , ACE selects the action over the proposals by all the actors,

$$a_t \doteq \operatorname{argmax}_{a \in \{\mu_i(s_t)\}_{i=1, \dots, N}} Q(s_t, a) \quad (8)$$

Our actors are trained in parallel. Assuming those actors are parameterized by θ , each actor μ_i is trained such that given a state s the action $a = \mu_i(s)$ maximizes $Q(s, a)$. We adopted similar gradient ascent as DDPG. The gradient at time step t is

$$\nabla_{\theta} Q(s_t, \mu_i(s_t)) = \nabla_a Q(s, a)|_{s=s_t, a=\mu_i(s_t)} \nabla_{\theta} \mu_i(s)|_{s=s_t} \quad (9)$$

for all $i \in \{1, \dots, N\}$. ACE initializes each actor independently. So that the actors are likely to cover different local maxima of $Q(s, a)$. By considering the maximum action value of the proposed actions, the action in Equation (8) is more likely to find the global maxima of the critic Q . To train the critic Q , ACE takes one gradient descent step at each time t minimizing

$$\frac{1}{2} (r_{t+1} + \gamma \max_{i \in \{1, \dots, N\}} Q(s_{t+1}, \mu_i(s_{t+1})) - Q(s_t, a_t))^2 \quad (10)$$

Note our actors are not independent. They reinforce each other by influencing the critic update, which in turn gives them better policy gradient.

An Option Perspective

Intuitively, each actor in ACE is similar to an option. To quantify the relationship between ACE and the option framework, we first extend OCA with deterministic intra-option policies, referred to as OCAD. For each option ω , we use $\mu_\omega : \mathcal{S} \rightarrow \mathcal{A}$ to denote its intra-option policy, which is a deterministic policy. The intro-option policies $\{\mu_\omega\}$ are parameterized by θ . The termination functions $\{\beta_\omega\}$ are parameterized by ν . We have

Theorem 1 (Deterministic Intra-option Policy Gradient)
Given a set of Markov options with deterministic intra-option policies, the gradient of the expected discounted

return objective w.r.t. θ is:

$$\sum_{s,\omega} \rho_{\Omega}(s,\omega|s_0,\omega_0) \nabla_a Q_U(s,\omega,a)|_{a=\mu_{\omega}(s)} \nabla_{\theta} \mu_{\omega}(s)$$

where $\rho_{\Omega}(s,\omega|s_0,\omega_0) = \sum_{k=0}^{\infty} P_{\gamma}^{(k)}(s,\omega|s_0,\omega_0)$ is the limiting state-option pair distribution. Here $P_{\gamma}^{(k)}$ represents the γ -discounted probability of transitioning to (s,ω) from (s_0,ω_0) in k steps.

Theorem 2 (Termination Policy Gradient) *Given a set of Markov options with deterministic intra-option policies, the gradient of the expected discounted return objective w.r.t. to ν is:*

$$\sum_{\omega,s'} \rho_{\Omega}(s',\omega|s_1,\omega_0) \nabla_{\nu} \beta_{\omega}(s') (V_{\Omega}(s') - Q_{\Omega}(s',\omega))$$

The proof of Theorem 1 follows a similar scheme of Sutton et al. (2000), Silver et al. (2014), and Bacon, Harb, and Precup (2017). The proof of Theorem 2 follows a similar scheme of Bacon, Harb, and Precup (2017). The conditions and proofs of both theorems are detailed in Supplementary Material. The critic update of OCAD remains the same as OCA (Equation 7).

We now show that the actor ensemble in ACE is a special setting of OCAD. The gradient update of the actors in ACE (Equation 9) can be justified via Theorem 1. The critic update in ACE (Equation 10) is equivalent to the critic update in OCAD (Equation 13). We first consider a special setting, where $\beta_{\omega}(s) \equiv 1$ for $\forall(\omega,s)$, which means each option terminates at every time step. In this setting, the value of $Q_U(s,\omega,a)$ does not depend on ω (Equations 4 and 5). Based on this observation, we rewrite $Q_U(s,\omega,a)$ as $Q(s,a)$. We have

$$Q_{\Omega}(s,\omega) = Q_U(s,\omega,\mu_{\omega}(s)) = Q(s,\mu_{\omega}(s)) \quad (11)$$

With the three Q s being the same, we rewrite the intra-policy gradient update in OCAD according to Theorem 1 as

$$\nabla_a Q(s,a)|_{s_t=s,a_t=\mu_{\omega_t}(s_t)} \nabla_{\theta} \mu_{\omega_t}(s)|_{s_t=s} \quad (12)$$

And we rewrite the critic update in OCAD (Equation 7) as

$$\frac{1}{2} (r_{t+1} + \gamma \max_{\omega'} Q(s_{t+1}, \mu_{\omega'}(s_{t+1})) - Q(s_t, a_t))^2 \quad (13)$$

Now the actor-critic updates in OCAD (Equations 12 and 13) recover the actor-critic updates in ACE (Equations 9 and 10), revealing the relationship between the ensemble in ACE and OCAD.

Note that in the intra-option policy update of OCAD (Equation 12) only one intra-option policy is updated at each time step, while in the actor ensemble update (Equation 9) all actors are updated. Based on the intra-option policy update of OCAD, we propose a variant of ACE, named Alternative ACE (ACE-Alt), where only the selected actor is updated at each time step. In practice, we add exploration noise for each action a_t and use experience replay to stabilize the training of the neural network function approximator like DDPG, resulting in off-policy learning.

Model-based Enhancements

To refine the state-action value function estimation, which is essential for actor selection, we utilize a look-ahead tree search method with a learned value prediction model similar to TreeQN. TreeQN was developed for discrete action space. We extend TreeQN to continuous control problems via the actor ensemble by searching over the actions proposed by the actors.

Formally speaking, we first define the following learnable functions:

- $f_{enc} : \mathcal{S} \rightarrow \mathbb{R}^n$, an encoding function that transforms a state into an n -dimensional latent state, parameterized by θ^{enc}
- $f_{rew} : \mathbb{R}^n \times \mathcal{A} \rightarrow \mathbb{R}$, a reward prediction function that predicts the immediate reward given a latent state and an action, parameterized by θ^{rew}
- $f_{trans} : \mathbb{R}^n \times \mathcal{A} \rightarrow \mathbb{R}^n$, a transition function that predicts the next latent state given a latent state and an action, parameterized by θ^{trans}
- $f_q : \mathbb{R}^n \times \mathcal{A} \rightarrow \mathbb{R}$, a value prediction function that computes the value for a pair of a latent state and an action, parameterized by θ^q
- $f_{\mu_i} : \mathbb{R}^n \rightarrow \mathcal{A}$: an actor that computes an action given a latent state, parameterized by θ^{μ_i} , for each $i \in \{1, \dots, N\}$

We use θ^Q to denote $\{\theta^{enc}, \theta^{rew}, \theta^{trans}, \theta^q\}$ and θ^{μ} to denote $\{\theta^{enc}, \theta^{\mu_1}, \dots, \theta^{\mu_N}\}$. Note the encoding function is shared in our implementation.

We use $f_{\mu_i}(z_{t|0})$ to represent $\mu_i(s_t)$ and $f_q(z_{t|0}, a_{t|0})$ to represent $Q(s_t, a_t)$, where $z_{t|0} = f_{enc}(s_t)$ is a latent state and $a_{t|0} = a_t$. Furthermore, $f_q(z_{t|0}, a_{t|0})$ can also be decomposed into the sum of the predicted immediate reward and the value of the predicted next latent state, i.e.,

$$f_q(z_{t|0}, a_{t|0}) \leftarrow f_{rew}(z_{t|0}, a_{t|0}) + \gamma f_q(z_{t|1}, a_{t|1}) \quad (14)$$

where

$$z_{t|1} = f_{trans}(z_{t|0}, a_{t|0}) \quad (15)$$

$$a_{t|1} = \operatorname{argmax}_{a \in \{f_{\mu_i}(z_{t|1})\}_{i=1, \dots, N}} f_q(z_{t|1}, a) \quad (16)$$

We apply Equation (14) recursively d times with state prediction and action selection in Equations (15), (16), resulting in a new estimator $f_q^d(z_{t|0}, a_{t|0})$ for $Q(s_t, a_t)$, which is defined as

$$f_q^d(z_{t|l}, a_{t|l}) = \begin{cases} f_q(z_{t|l}, a_{t|l}) & (d=0) \\ f_{rew}(z_{t|l}, a_{t|l}) + \gamma f_q^{d-1}(z_{t|l+1}, a_{t|l+1}) & (l < d) \end{cases} \quad (17)$$

where

$$z_{t|0} \doteq f_{enc}(s_t), \quad a_{t|0} \doteq a_t$$

$$z_{t|l} \doteq f_{trans}(z_{t|l-1}, a_{t|l-1}) \quad (l \geq 1)$$

$$a_{t|l} \doteq \operatorname{argmax}_{a \in \{f_{\mu_i}(z_{t|l})\}_{i=1, \dots, N}} f_q^{d-l}(z_{t|l}, a) \quad (l \geq 1)$$

The look-ahead tree search and backup process (Equation 17) are illustrated in Figure 1. The value of $f_q^d(z_{t|l}, a_{t|l})$

stands for the state-action value estimation for the predicted latent state and action after l steps from t , with Equation (14) applied d times.

As f_q^d is fully differentiable w.r.t. θ^Q , we plug in f_q^d whenever we need Q . We also ground the predicted reward in the first recursive expansion as suggested by Farquhar et al. (2018). To summarize, given a transition $(s_t, a_t, r_{t+1}, s_{t+1})$, the gradients for θ^Q and θ^μ are

$$\nabla_{\theta^Q} \left(\frac{1}{2} (f_q^d(z_{t|0}, a_{t|0}) - \max_i f_q^d(z_{t+1|0}, f_{\mu_i}(z_{t+1|0})))^2 + \frac{1}{2} (f_{rew}(z_{t|0}, a_{t|0}) - r_{t+1})^2 \right)$$

and

$$\sum_{i=1}^N \nabla_a f_q^d(z_{t|0}, a)|_{a=f_{\mu_i}(z_{t|0})} \nabla_{\theta^\mu} f_{\mu_i}(z_{t|0})$$

respectively. ACE also utilizes experience replay and a target network similar to DDPG. The pseudo-code of ACE is provided in Supplementary Material.

Experiments

We designed experiments to answer the following questions:

- Does ACE outperform baseline algorithms?
- If so, how do the components of ACE contribute to the performance boost?

We used twelve continuous control tasks from Roboschool, a free port of Mujoco² by OpenAI. A state in Roboschool contains joint information of a robot (e.g., speed or angle) and is presented as a vector of real numbers. An action in Roboschool is a vector, with each dimension being a real number in $[-1, 1]$. All the implementations are made publicly available.³

ACE Architecture

In this section we describe the parameterization of f_{enc} , f_{rew} , f_{trans} , f_q and $\{f_{\mu_i}\}_{i=1, \dots, N}$ for Roboschool tasks. For each state s , we first transformed it into a latent state $z \in \mathbb{R}^{400}$ by f_{enc} , which was parameterized as a single neural network layer. This latent state z was used as the input for all other functions (i.e., f_{rew} , f_q , f_{trans} , $f_{\mu_1}, \dots, f_{\mu_N}$).

The networks for f_{rew} , f_q , f_{trans} are single hidden layer networks with $400 + m$ input units and 300 hidden units, taking as inputs the concatenation of a latent state and an m -dimensional action. Particularly, the network for f_{trans} used two residual connections similar to Farquhar et al. (2018). The networks for $\{f_{\mu_i}\}_{i=1, \dots, N}$ are single hidden layer networks with 400 input units and 300 hidden units, and all the N networks of $\{f_{\mu_i}\}$ shared a common first layer.

We used tanh as the activation functions for the hidden units. (This selection will be further discussed in the next section.) We set the number of actors to 5 (i.e., $N = 5$) and the planning depth to 1 (i.e. $d = 1$).

²<http://www.mujoco.org/>

³<https://github.com/ShangtongZhang/DeepRL>

Baseline Algorithms

DDPG In DDPG, Lillicrap et al. (2015) used two separate networks to parameterize the actor and the critic. Each network had 2 hidden layers with 400 and 300 hidden units respectively. Lillicrap et al. (2015) used ReLU activation function (Nair and Hinton 2010) and applied a L_2 regularization to the critic. However, our analysis experiments found that tanh activation function outperformed ReLU with L_2 regularization. So throughout all our experiments, we always used tanh activation function (without L_2 regularization) for all algorithms. All other hyper-parameter values were the same as Lillicrap et al. (2015). All the other compared algorithms inherited the hyper-parameter values from DDPG without tuning. We used the same Ornstein-Uhlenbeck process (Uhlenbeck and Ornstein 1930) as Lillicrap et al. (2015) for exploration in all the compared algorithms.

Wide-DDPG ACE had more parameters than DDPG. To investigate the influence of the number of parameters, we implemented Wide-DDPG, where the hidden units were doubled (i.e., the two hidden layers had 800 and 600 units respectively). Wide-DDPG had a comparable number of parameters as ACE and remained the same depth as ACE.

Shared-DDPG DDPG used separate networks for actor and critic, while the actor and critic in ACE shared a common representation layer. To investigate the influence of a shared representation, we implemented Shared-DDPG, where the actor and the critic shared a common bottom layer in the same manner as ACE.

Ensemble-DDPG To investigate the influence of the tree search in ACE, we removed the tree search in ACE by setting $d = 0$, giving an ensemble version of DDPG, called Ensemble-DDPG. We still used 5 actors in Ensemble-DDPG.

TM-ACE To investigate the usefulness of the value prediction model, we implemented Transition-Model-ACE (TM-ACE), where we learn a transition model instead of a value prediction model. To be more specific, f_{trans} and f_{rew} were trained to fit sampled transitions from the replay buffer to minimize the squared loss of the predicted reward and the predicted next latent state. This model was then used for a look-ahead tree search. The pseudo-code of TM-ACE is detailed in Supplementary Material.

The architectures of all the above algorithms are illustrated in Supplementary Material.

Results

For each task, we trained each algorithm for 1 million steps. At every 10 thousand steps, we performed 20 deterministic evaluation episodes without exploration noise and computed the mean episode return. We report the best evaluation performance during training in Table 1, which is averaged over 5 independent runs. The full evaluation curves are reported in Supplementary Material.

In a summary, either ACE or ACE-Alt was placed among the best algorithms in 11 out of the 12 games. ACE itself was placed among the best algorithms in 8 games taking

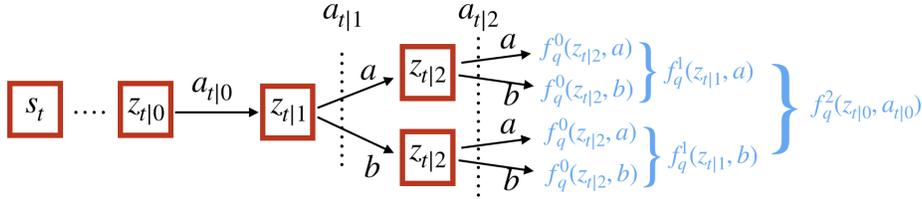


Figure 1: An example of the tree search (Equation 17) with $d = 2$ and $N = 2$. We use $f_q^2(z_{t|0}, a_{t|0})$ as the estimation $Q(s_t, a_t)$. Here, $z_{t|l}$ ($l = 1, 2$) represents the predicted latent state after l steps from t , a and b are actions proposed by the two actors and depend on latent state. Arrows represent the transition kernel f_{trans} , and brackets represent maximization (backup) operations. Reward prediction is omitted for simplicity.

ACE-Alt into comparison. Without considering ACE-Alt, ACE was placed among the best algorithms in 10 games. ACE-Alt itself was placed among the best algorithms in 7 games taking ACE into comparison. Without considering ACE, ACE-Alt was placed among the best algorithms in 10 games. Overall, ACE was slightly better than ACE-Alt. However, ACE-Alt enjoyed lower variance than ACE. We conjecture this is because ACE had more off-policy learning than ACE-Alt. Off-policy learning improved data efficiency but increased variances.

Wide-DDPG outperformed DDPG in only 1 game, indicating that naively increasing the parameters does not guarantee performance improvement. Shared-DDPG outperformed DDPG in only 2 games (lost in 2 games and tied in 8 games), showing shared representation contributes little to the overall performance in ACE. Ensemble-DDPG outperformed DDPG in 6 games (lost in 3 games and tied in 3 games), indicating the DDPG agent benefits from an actor ensemble. This may be attributed to that multiple actors are more likely to find the global maxima of the critic. ACE further outperformed Ensemble-DDPG in 9 games, indicating the agent benefits from the look-ahead tree search with a value prediction model. In contrast, TM-ACE outperformed Ensemble-DDPG only in 2 games (lost in 3 games and tied in 7 games), indicating that a value prediction model is better than a transition model for a look-ahead tree search. This was also consistent with the results observed earlier in VPN and TreeQN.

In conclusion, the actor ensemble and the look-ahead tree search with a learned value prediction model are key to the performance boost.

ACE and ACE-Alt increase performance in terms of environment steps while require more computation than vanilla DDPG. We benchmarked the wall time for the algorithms. The results are reported in Supplementary Material. We also verified the diversity of the actors in ACE and ACE-Alt in Supplementary Material.

Ensemble Size and Planning Depth

In this section, we investigate how the ensemble size N and the planning depth d in ACE influence the performance. We performed experiments in HalfCheetah with various N and d and used the same evaluation protocol as before. As a large N and d induced a significant computation increase, we only

used N up to 10 and d up to 2. The results are reported in Figure 2.

To summarize, $N = 5$ and $d = 1$ achieved the best performance. We hypothesize there is a trade-off in the selection of both the ensemble size and the planning depth. On the one hand, a single actor can easily be trapped into local maxima during training. The more actors we have, the more likely we find the global maxima. On the other hand, all the actors share the same encoding function with the critic. A large number of actors may dominate the training of the encoding function to damage the critic learning. So a medium ensemble size is likely to achieve the best performance. A possible solution is to normalize the gradient according to the ensemble size, and we leave this for future work. With a perfect model, the more planning steps we have, the more accurate Q estimation we can get. However, with a learned value prediction model, there is a compound error in unrolling. So a medium planning depth is likely to achieve the best performance. Similar phenomena were also observed in the multi-step Dyna planning (Yao et al. 2009).

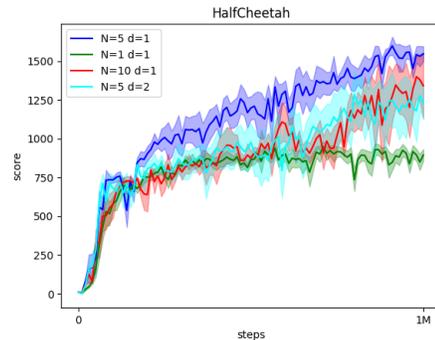


Figure 2: Evaluation curves of ACE in HalfCheetah with different N and d . Each curve is averaged over 5 independent runs, and standard errors are plotted in shadow.

Related Work

Continuous-action RL Klissarov et al. (2017) extended OCA into continuous action problems with the Proximal Policy Option Critic (PPOC) algorithm. However, PPOC considered stochastic intra-option policies, and each intra-

	ACE	ACE-Alt	TM-ACE	Ensemble-DDPG	Shared-DDPG	Wide-DDPG	DDPG
Ant	1041 (70.8)	983 (36.8)	1031 (55.6)	1026 (87.2)	796(16.8)	871(19.9)	875(14.2)
HalfCheetah	1667 (40.4)	1023(60.4)	800(28.8)	812(49.4)	771(79.6)	733(52.5)	703(37.3)
Hopper	2136 (86.4)	1923(88.3)	1586(85.0)	1972(63.4)	2047 (76.7)	2090 (118.6)	2133 (99.0)
Humanoid	380 (56.1)	441 (90.1)	61(6.8)	76(11.4)	53(2.2)	54(1.0)	54(1.7)
HF	311 (30.3)	289 (20.8)	126(39.6)	85(5.6)	53(1.0)	55(1.6)	53(1.4)
HFH	22 (2.4)	20 (2.2)	-4(2.1)	2(7.5)	6(6.4)	15(3.2)	15(2.5)
IDP	7555(1610.9)	9356 (1.1)	7549(1613.7)	4102(1923.2)	7549(1613.6)	7548(1618.7)	5662(1945.9)
IP	417(212.8)	1000 (0.0)	415(213.4)	1000 (0.0)	1000 (0.0)	1000 (0.0)	1000 (0.0)
IPS	892 (0.1)	892(0.2)	891(0.2)	891(0.2)	891(0.4)	546(308.7)	891(0.4)
Pong	12 (0.3)	11(0.1)	6(0.7)	8(0.9)	4(0.2)	4(0.1)	5(0.8)
Reacher	16(0.7)	17(0.2)	17(0.5)	17(0.3)	20 (0.7)	15(2.2)	18(0.9)
Walker2d	1659(65.9)	1864 (21.4)	1086(97.0)	1142(99.5)	1142(146.3)	1185(121.2)	815(11.4)

Table 1: Best evaluation performance during training. Mean and standard error are reported. Bold numbers indicate the best performance. Scores are averaged over 5 independent runs. Numbers are rounded for the ease of display. HF, HFH, IDP, IP, and IPS stand for HumanoidFlagrun, HumanoidFlagrunHarder, InvertedDoublePendulum, InvertedPendulum, and InvertedPendulumSwingup respectively.

option policy was trained via a policy search method. In ACE, we consider deterministic intra-option policies, and the intra-option policies are optimized under the same objective as OCA.

Gu et al. (2016) parameterized the Q function in a quadric form to deal with continuous control problems. In this way, the global maxima can be determined analytically. However, in general, the optimal Q value does not necessarily fall into this quadric form. In ACE, we use an actor ensemble to search the global maxima of the Q function. Gu et al. (2016) utilized a transition model to generate imaginary data, which is orthogonal to ACE.

Ensemble in RL Wiering and Van Hasselt (2008) designed four ensemble methods combining five RL algorithms with a voting scheme based on value functions of different RL algorithms. Hans and Udluft (2010) used a network ensemble to improve the performance of Fitted Q-Iteration. Osband et al. (2016) used a Q ensemble to approximate Thomas’ sampling, resulting in improved exploration and performance boost in challenging video games. Huang et al. (2017) used both an actor ensemble and a critic ensemble in continuous control problems. However, to our best knowledge, the present work is the first to relate ensemble with options and to use an ensemble for a look-ahead tree search in continuous control problems.

Closing Remarks

In this paper, we propose the ACE algorithm for continuous control problems. From an ensemble perspective, ACE utilizes an actor ensemble to search the global maxima of a critic function. From an option perspective, ACE is a special option-critic algorithm with deterministic intra-option policies. Thanks to the actor ensemble, ACE is able to perform a look-ahead tree search with a learned value prediction model in continuous control problems, resulting in a significant performance boost in challenging robot manipulation tasks.

Acknowledgement

Due to a technical mistake in submission, we failed to input the whole author list in this camera-ready version, which should be Shangtong Zhang, Hao Chen, Hengshuai Yao. The authors thank Bo Liu for feedbacks of the first draft of this paper. We also appreciate the insightful reviews of the anonymous reviewers.

References

- Bacon, P.-L.; Harb, J.; and Precup, D. 2017. The option-critic architecture. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*.
- Barth-Maron, G.; Hoffman, M. W.; Budden, D.; Dabney, W.; Horgan, D.; Muldal, A.; Heess, N.; and Lillicrap, T. 2018. Distributed distributional deterministic policy gradients. *arXiv preprint arXiv:1804.08617*.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*.
- Coulom, R. 2006. Efficient selectivity and backup operators in monte-carlo tree search. In *Proceedings of the International Conference on Computers and Games*.
- Degrís, T.; White, M.; and Sutton, R. S. 2012. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*.
- Farquhar, G.; Rocktäschel, T.; Igl, M.; and Whiteson, S. 2018. Treeqn and atreec: Differentiable tree-structured models for deep reinforcement learning. *arXiv preprint arXiv:1710.11417*.
- Gu, S.; Lillicrap, T.; Sutskever, I.; and Levine, S. 2016. Continuous deep q-learning with model-based acceleration. In *Proceedings of the 33rd International Conference on Machine Learning*.
- Hans, A., and Udluft, S. 2010. Ensembles of neural networks for robust reinforcement learning. In *Proceedings of the 9th International Conference on Machine Learning and Applications*.

- Huang, Z.; Zhou, S.; Zhuang, B.; and Zhou, X. 2017. Learning to run with actor-critic ensemble. *arXiv preprint arXiv:1712.08987*.
- Klissarov, M.; Bacon, P.-L.; Harb, J.; and Precup, D. 2017. Learnings options end-to-end for continuous action tasks. *arXiv preprint arXiv:1712.00004*.
- Knuth, D. E., and Moore, R. W. 1975. An analysis of alpha-beta pruning. *Artificial Intelligence*.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Lin, L.-J. 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*.
- Mansley, C. R.; Weinstein, A.; and Littman, M. L. 2011. Sample-based planning for continuous action markov decision processes. In *Proceedings of the 21st International Conference on Automated Planning and Scheduling*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature*.
- Nair, V., and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning*.
- Nitti, D.; Belle, V.; and De Raedt, L. 2015. Planning in discrete and continuous markov decision processes by probabilistic programming. In *Proceedings of the 17th Joint European Conference on Machine Learning and Knowledge Discovery in Databases*.
- Oh, J.; Singh, S.; and Lee, H. 2017. Value prediction network. In *Advances in Neural Information Processing Systems*.
- Osband, I.; Blundell, C.; Pritzel, A.; and Van Roy, B. 2016. Deep exploration via bootstrapped dqn. In *Advances in Neural Information Processing Systems*.
- Silver, D.; Lever, G.; Heess, N.; Degris, T.; Wierstra, D.; and Riedmiller, M. 2014. Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *Nature*.
- Sturtevant, N. 2008. An analysis of uct in multi-player games. In *Proceedings of the International Conference on Computers and Games*.
- Sutton, R. S.; McAllester, D. A.; Singh, S. P.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*.
- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*.
- Sutton, R. S. 1990. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proceedings of the 7th International Conference on Machine Learning*.
- Tassa, Y.; Doron, Y.; Muldal, A.; Erez, T.; Li, Y.; Casas, D. d. L.; Budden, D.; Abdolmaleki, A.; Merel, J.; Lefrancq, A.; et al. 2018. Deepmind control suite. *arXiv preprint arXiv:1801.00690*.
- Tesauro, G. 1995. Temporal difference learning and td-gammon. *Communications of the ACM*.
- Uhlenbeck, G. E., and Ornstein, L. S. 1930. On the theory of the brownian motion. *Physical review*.
- Watkins, C. J., and Dayan, P. 1992. Q-learning. *Machine Learning*.
- Weber, T.; Racanière, S.; Reichert, D. P.; Buesing, L.; Guez, A.; Rezende, D. J.; Badia, A. P.; Vinyals, O.; Heess, N.; Li, Y.; et al. 2017. Imagination-augmented agents for deep reinforcement learning. *arXiv preprint arXiv:1707.06203*.
- Wiering, M. A., and Van Hasselt, H. 2008. Ensemble algorithms in reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*.
- Yao, H.; Bhatnagar, S.; Diao, D.; Sutton, R. S.; and Szepesvári, C. 2009. Multi-step dyna planning for policy evaluation and control. In *Advances in Neural Information Processing Systems*.
- Yee, T.; Lisý, V.; and Bowling, M. H. 2016. Monte carlo tree search in continuous action spaces with execution uncertainty. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*.