
Sample-Efficient Reinforcement Learning with Stochastic Ensemble Value Expansion

Jacob Buckman*
buckman@google.com

Danijar Hafner
danijar@google.com

George Tucker
gjt@google.com

Eugene Brevdo
ebrevdo@google.com

Honglak Lee
honglak@google.com

Google Brain, Mountain View, CA, USA

Abstract

Integrating model-free and model-based approaches in reinforcement learning has the potential to achieve the high performance of model-free algorithms with low sample complexity. However, this is difficult because an imperfect dynamics model can degrade the performance of the learning algorithm, and in sufficiently complex environments, the dynamics model will almost always be imperfect. As a result, a key challenge is to combine model-based approaches with model-free learning in such a way that errors in the model do not degrade performance. We propose stochastic ensemble value expansion (STEVE), a novel model-based technique that addresses this issue. By dynamically interpolating between model rollouts of various horizon lengths for each individual example, STEVE ensures that the model is only utilized when doing so does not introduce significant errors. Our approach outperforms model-free baselines on challenging continuous control benchmarks with an order-of-magnitude increase in sample efficiency, and in contrast to previous model-based approaches, performance does not degrade in complex environments.

1 Introduction

Deep model-free reinforcement learning has had great success in recent years, notably in playing video games Mnih et al. [2013] and strategic board games Silver et al. [2016]. However, training agents using these algorithms requires tens to hundreds of millions of samples, which makes

many practical applications infeasible, particularly in real-world control problems (e.g., robotics) where data collection is expensive.

Model-based approaches aim to reduce the environment samples required to learn a policy by modeling the dynamics. A dynamics model can be used to increase sample efficiency in various ways, for example training the policy on rollouts from the dynamics model Sutton [1990], using rollouts to improve targets for TD learning Feinberg et al. [2018], and using rollout statistics as inputs to the policy Weber et al. [2017]. Model-based algorithms such as PILCO Deisenroth and Rasmussen [2011] have shown that it is possible to learn from orders-of-magnitude fewer samples.

These successes have mostly been limited to environments where the dynamics are simple to model. In noisy, complex environments, it is difficult to learn an accurate model of the environment. When the model makes mistakes, it can cause the wrong policy to be learned, hindering performance. Recent work has begun to address this issue Sutton and Barto [1998], Kurutach et al. [2018], Kalweit and Boedecker [2017], Weber et al. [2017], Gal et al., Depeweg et al. [2016], Gu et al. [2016a]; see Appendix G for a more in-depth discussion of these approaches.

We propose *stochastic ensemble value expansion* (STEVE), an extension to *model-based value expansion* (MVE), proposed by Feinberg et al. [2018]. Both techniques use a dynamics model to compute “rollouts” that are used to improve the targets for temporal difference learning. MVE rolls out a fixed length into the future, potentially accumulating model errors or increasing value estimation error along the way. In contrast, STEVE interpolates between many different horizon lengths, favoring those whose estimates have lower uncertainty, and thus lower error. To compute the interpolated target, we replace both the model and Q-function with ensembles, approximating the uncertainty of an estimate by computing its variance under samples from the ensemble. Through

*Work completed as part of the Google AI Residency Program.

these uncertainty estimates, STEVE dynamically utilizes the model rollouts only when they do not introduce significant errors. We systematically evaluate STEVE on several challenging continuous control benchmarks and demonstrate that STEVE significantly outperforms model-free baselines with an order-of-magnitude increase in sample efficiency, and that in contrast to previous model-based approaches, the performance of STEVE does not degrade as the environment gets more complex.

2 Background

Reinforcement learning aims to learn an agent policy that maximizes the expected (discounted) sum of rewards [Sutton and Barto, 1998]. We focus on the deterministic case for exposition; however, our method is applicable to the stochastic case as well. The agent starts at an initial state $s_0 \sim p(s_0)$. Then, the agent chooses an action a_t according to its policy $\pi_\phi(s_t)$ with parameters ϕ , receives a reward $r_t = r(s_t, a_t)$, and transitions to a subsequent state s_{t+1} according to the Markovian dynamics $T(a_t, s_t)$ of the environment. This generates a trajectory of states, actions, and rewards $\tau = (s_0, a_0, r_0, s_1, a_1, \dots)$. We abbreviate the trajectory by τ . The goal is to maximize the expected discounted sum of rewards along sampled trajectories $J(\theta) = \mathbb{E}_\tau [\sum_{t=0}^{\infty} \gamma^t r_t]$ where $\gamma \in [0, 1)$ is a discount parameter.

2.1 Value Estimation with TD-learning

The action-value function $Q^\pi(s_0, a_0) = \sum_{t=0}^{\infty} \gamma^t r_t$ is a critical quantity to estimate for many learning algorithms. Using the fact that $Q^\pi(s, a)$ satisfies a recursion relation

$$Q^\pi(s, a) = r(s, a) + \gamma Q^\pi(s', \pi(s')),$$

where $s' = T(s, a)$, we can estimate $Q^\pi(s, a)$ off-policy with collected transitions of the form (s, a, r, s') sampled uniformly from a replay buffer [Sutton and Barto, 1998]. We approximate $Q^\pi(s, a)$ with a deep neural network, $\hat{Q}_\theta^\pi(s, a)$. We learn parameters θ to minimize the mean squared error (MSE) between Q-value estimates of states and their corresponding temporal difference targets Sutton and Barto [1998]:

$$\begin{aligned} \mathcal{T}^{TD}(r, s') &= r + \gamma \hat{Q}_\theta^\pi(s', \pi(s')) \\ \mathcal{L}_\theta &= \mathbb{E}_{(s,a,r,s')} \left[(\hat{Q}_\theta^\pi(s, a) - \mathcal{T}^{TD}(r, s'))^2 \right]. \end{aligned} \quad (1)$$

Note that we use an older copy of the parameters, θ^- , when computing targets [Mnih et al., 2013].

To approximate a policy which maximizes our Q-function, we use a neural network Lillicrap et al. [2015], learning

parameters ϕ to minimize the negative Q-value:

$$\mathcal{L}_\phi = -\hat{Q}_\theta^\pi(s, \pi_\phi(s)). \quad (2)$$

2.1.1 Model-Based Value Expansion

Recently, Feinberg et al. [2018] showed that a learned dynamics model can be used to improve value estimation. Their method, model-based value expansion, combines a short term value estimate formed by unrolling the model dynamics and a long term value estimate using the learned \hat{Q}_θ^π function. When the model is accurate, this reduces the bias of the targets, leading to improved performance.

The learned dynamics model consists of three learned functions: the transition function $\hat{T}_\xi(s, a)$, which returns a successor state s' ; a termination function $\hat{d}_\xi(s)$, which returns the probability that s is a terminal state; and the reward function $\hat{r}_\psi(s, a, s')$, which returns a scalar reward. This model is trained to minimize

$$\begin{aligned} \mathcal{L}_{\xi, \psi} &= \mathbb{E}_{(s,a,r,s')} [\|\hat{T}_\xi(s, a) - s'\|^2 \\ &\quad + \mathbb{H} \left(d(t | s'), \hat{d}_\xi(t | \hat{T}_\xi(s, a)) \right) + (\hat{r}_\psi(s, a, s') - r)^2], \end{aligned} \quad (3)$$

where the expectation is over collected transitions (s, a, r, s') , $d(t | s')$ is an indicator function which returns 1 when s' is a terminal state and 0 otherwise, and \mathbb{H} is the cross-entropy. In this work, we consider continuous environments; for discrete environments, the first term can be replaced by a cross-entropy loss term.

To incorporate the model into value estimation, we replace our standard Q-learning target with an improved target, $\mathcal{T}_H^{\text{MVE}}$, computed by rolling the learned model out for H steps.

$$\begin{aligned} s'_0 &= s', \quad a'_i = \pi_\phi(s'_i), \quad s'_i = \hat{T}_\xi(s'_{i-1}, a'_{i-1}), \\ D^i &= \prod_{j=0}^i (1 - d(s'_j)) \\ \mathcal{T}_H^{\text{MVE}}(r, s') &= r + \left(\sum_{i=1}^H D^i \gamma^i \hat{r}_\psi(s'_{i-1}, a'_{i-1}, s'_i) \right) + \\ &\quad D^{H+1} \gamma^{H+1} \hat{Q}_\theta^\pi(s'_H, a'_H). \end{aligned} \quad (4)$$

To use this target¹, we substitute $\mathcal{T}_H^{\text{MVE}}$ in place of \mathcal{T}^{TD} when training θ using Equation 1. Note that when $H = 0$, MVE reduces to TD-learning (i.e., $\mathcal{T}^{\text{TD}} = \mathcal{T}_0^{\text{MVE}}$). See Appendix B for further discussion of MVE.

¹This formulation is a minor generalization of the original MVE objective in that we additionally model the reward function and termination function; Feinberg et al. [2018] consider fully observable environments in which the reward function and termination condition were known, deterministic functions of the observations.

3 Stochastic Ensemble Value Expansion

From a single rollout of H timesteps, we can compute $H + 1$ distinct *candidate targets* by considering rollouts to various horizon lengths: $\mathcal{T}_0^{\text{MVE}}, \mathcal{T}_1^{\text{MVE}}, \mathcal{T}_2^{\text{MVE}}, \dots, \mathcal{T}_H^{\text{MVE}}$. Standard TD learning uses $\mathcal{T}_0^{\text{MVE}}$ as the target, while MVE uses $\mathcal{T}_H^{\text{MVE}}$ as the target. We propose interpolating all of the candidate targets to produce a target which is better than any individual. Naïvely, one could average the candidate targets, or weight the candidate targets in an exponentially-decaying fashion, similar to TD- λ Sutton and Barto [1998]. However, we show that we can do still better by weighting the candidate targets in a way that balances errors in the learned Q -function and errors from longer model rollouts. STEVE provides a computationally-tractable and theoretically-motivated algorithm for choosing these weights. We describe the algorithm for STEVE in Section 3.1, and justify it in Section 3.2.

3.1 Algorithm

To estimate uncertainty in our learned estimators, we maintain an ensemble of parameters for our Q -function, reward function, and model: $\theta = \{\theta_1, \dots, \theta_L\}$, $\psi = \{\psi_1, \dots, \psi_N\}$, and $\xi = \{\xi_1, \dots, \xi_M\}$, respectively. Each parameterization is initialized independently and trained on different subsets of the data in each minibatch. See Appendix D for further discussion of our choice to use ensembles.

We rollout an H step trajectory with each of the M models, $\tau^{\xi_1}, \dots, \tau^{\xi_M}$. Each trajectory consists of $H + 1$ states, $\tau_0^{\xi_m}, \dots, \tau_H^{\xi_m}$, which correspond to s'_0, \dots, s'_H in Equation 4 with the transition function parameterized by ξ_m . Similarly, we use the N reward functions and L Q -functions to evaluate Equation 5 for each τ^{ξ_m} at every rollout-length $0 \leq i \leq H$. This gives us $M \cdot N \cdot L$ different values of $\mathcal{T}_i^{\text{MVE}}$ for each rollout-length i .

Using these values, we can compute the empirical mean \mathcal{T}_i^μ and variance $\mathcal{T}_i^{\sigma^2}$ for each partial rollout of length i . In order to form a single target, we use an inverse variance weighting of the means:

$$\mathcal{T}_H^{\text{STEVE}}(r, s') = \sum_{i=0}^H \frac{\tilde{w}_i}{\sum_j \tilde{w}_j} \mathcal{T}_i^\mu, \quad \tilde{w}_i^{-1} = \mathcal{T}_i^{\sigma^2} \quad (6)$$

To learn a value function with STEVE, we substitute in $\mathcal{T}_H^{\text{STEVE}}$ in place of \mathcal{T}^{TD} when training θ using Equation 1.

3.2 Justification

We wish to find weights w_i , where $\sum_i w_i = 1$ that minimize the mean-squared error between the weighted-

average of candidate targets $\mathcal{T}_0^{\text{MVE}}, \mathcal{T}_1^{\text{MVE}}, \mathcal{T}_2^{\text{MVE}}, \dots, \mathcal{T}_H^{\text{MVE}}$ and the true Q -value.

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{i=0}^H w_i \mathcal{T}_i^{\text{MVE}} - Q^\pi(s, a) \right)^2 \right] \\ &= \text{Bias} \left(\sum_i w_i \mathcal{T}_i^{\text{MVE}} \right)^2 + \text{Var} \left(\sum_i w_i \mathcal{T}_i^{\text{MVE}} \right) \\ &\approx \text{Bias} \left(\sum_i w_i \mathcal{T}_i^{\text{MVE}} \right)^2 + \sum_i w_i^2 \text{Var}(\mathcal{T}_i^{\text{MVE}}), \end{aligned}$$

where the expectation considers the candidate targets as random variables conditioned on the collected data and minibatch sampling noise, and the approximation is due to assuming the candidate targets are independent. Experiments suggested that including the covariance terms was not effective in practice, so we dropped the extra terms for simplicity.

Our goal is to minimize this with respect to w_i . We can estimate the variance terms using empirical variance estimates from the ensemble. Unfortunately, we could not devise a reliable estimator for the bias terms, and this is a limitation of our approach and an area for future work. In this work, we ignore the bias terms and minimize the weighted sum of variances

$$\sum_i w_i^2 \text{Var}(\mathcal{T}_i^{\text{MVE}}).$$

With this approximation, which is equivalent to inverse-variance weighting Fleiss [1993], we achieve state-of-the-art results. Setting each w_i equal to $\frac{1}{\sqrt{\text{Var}(\mathcal{T}_i^{\text{MVE}})}}$ and normalizing yields the formula for $\mathcal{T}_H^{\text{STEVE}}$ given in Equation 6.

4 Experiments

4.1 Implementation

We use DDPG Lillicrap et al. [2015] as our baseline model-free algorithm. We train two deep feedforward neural networks, a Q -function network $\hat{Q}_\theta^\pi(s, a)$ and a policy network $\pi_\phi(s)$, by minimizing the loss functions given in Equations 1 and 2. We also train another three deep feedforward networks to represent our worldmodel, corresponding to function approximators for the transition $\hat{T}_\xi(s, a)$, termination $\hat{d}_\xi(t | s)$, and reward $\hat{r}_\psi(s, a, s')$, and minimize the loss function given in Equation 3.

We use a distributed implementation to parallelize computation. In the style of ApeX Horgan et al. [2018], IMPALA Espeholt et al. [2018], and D4PG Barth-Maron et al. [2018], we use a centralized learner with several

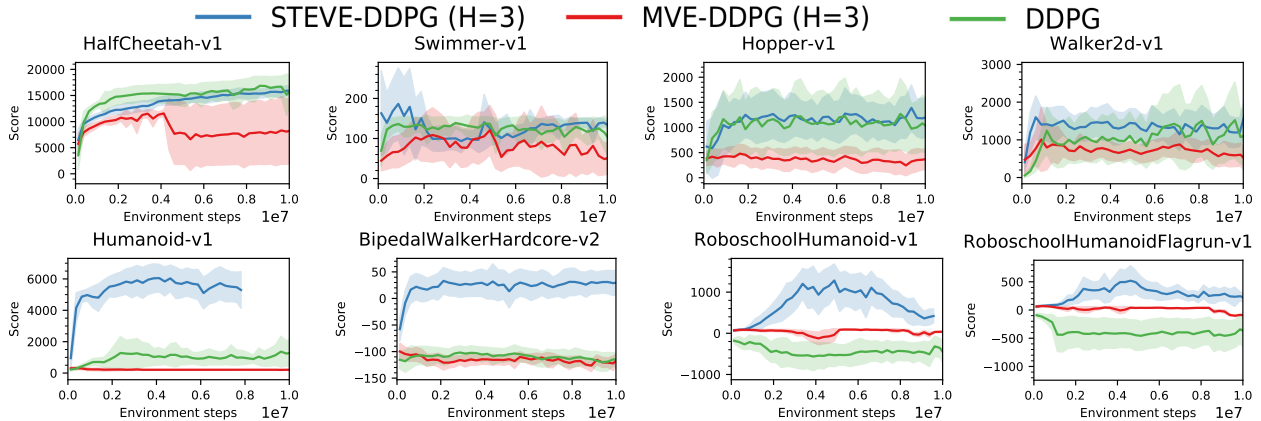


Figure 1: Learning curves comparing our method to both model-free and model-based baselines. Each experiment was run four times.

agents operating in parallel. Each agent periodically loads the most recent policy, interacts with the environment, and sends its observations to the central learner. The learner stores received frames in a replay buffer, and continuously loads batches of frames from this buffer to use as training data for a model update. In the algorithms with a model-based component, there are two learners, a policy-learner and a model-learner. In these cases, the policy-learner periodically reloads the latest copy of the model.

All baselines reported in this section were re-implementations of existing methods. This allowed us to ensure that the various methods compared were consistent with one another, and that the differences reported are fully attributable to the independent variables in question. Our baselines are competitive with state-of-the-art implementations of these algorithms Haarnoja et al. [2018], Feinberg et al. [2018]. All MVE experiments utilize the TD-k trick. For hyperparameters and additional implementation details, please see Appendix C.²

4.2 Comparison of Performance

We tested on a variety of continuous control tasks [Brockman et al., 2016, Klimov and Schulman]; the results can be seen in Figure 2. We found that STEVE yields significant improvements in both performance and sample efficiency across a wide range of environments. Importantly, the gains are more substantial in the more complex environments. On all of the most challenging environments, we see that STEVE is the only algorithm to begin to learn a better-than-random policy within 10M frames. STEVE also outperforms other recently-published results on these tasks in terms of sample efficiency Gu et al.

²Our code is available open-source at: <https://github.com/tensorflow/models/tree/master/research/steve>

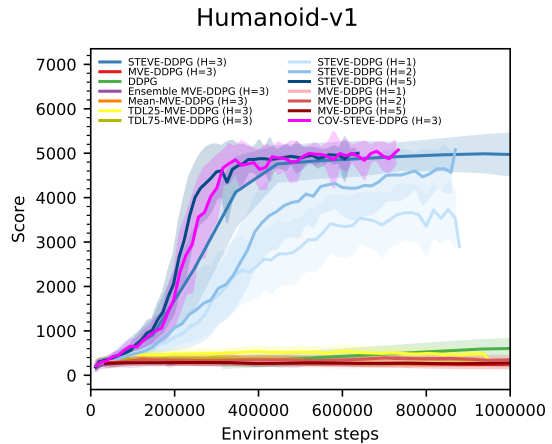


Figure 2: Ablation experiments. Each experiment was run twice.

[2016b], Haarnoja et al. [2018], Schulman et al. [2017].

4.3 Ablation Study

In order to verify that STEVE’s gains in sample efficiency are due to the reweighting, and not simply due to the additional parameters of the ensembles of its components, we examine several ablations. Ensemble MVE is the regular MVE algorithm, but the model and Q-functions are replaced with with ensembles. Mean-MVE uses the exact same architecture as STEVE, but uses a simple uniform weighting instead of the uncertainty-aware reweighting scheme. Similarly, TDL25 and TDL75 correspond to TD- λ reweighting schemes with $\lambda = 0.25$ and $\lambda = 0.75$, respectively. We also investigate the effect of the horizon parameter on the performance of both STEVE and MVE. These results, which can be seen in Figure 2, support the hypothesis that the increased performance is due

to the uncertainty-dependent reweighting of targets, and demonstrate that the performance of STEVE consistently increases with longer horizon lengths.

4.4 Additional Analysis and Discussion

See Appendix E for a comparison of wall-clock time and Appendix F for an analysis of model usage.

5 Conclusion

In this work, we introduced STEVE, an uncertainty-aware approach for merging model-free and model-based reinforcement learning, which outperforms model-free approaches while reducing sample complexity on several challenging tasks. We believe that this is a strong step towards enabling RL for practical, real-world applications.

Acknowledgements

The authors would like to thank the following individuals for their valuable insights and discussion: David Ha, Prajit Ramachandran, Tuomas Haarnoja, Dustin Tran, Matt Johnson, Matt Hoffman, Ishaan Gulrajani, and Sergey Levine. Also, we would like to thank Jascha Sohl-Dickstein, Joseph Antognini, Shane Gu, and Samy Bengio for their feedback during the writing process. Finally, we would like to thank Erwin Coumans for his help on PyBullet environments.

References

- Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.
- Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I Jordan, Joseph E Gonzalez, and Sergey Levine. Model-based value estimation for efficient model-free reinforcement learning. *arXiv preprint arXiv:1803.00101*, 2018.
- JL Fleiss. Review papers: The statistical basis of meta-analysis. *Statistical methods in medical research*, 2(2):121–145, 1993.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- Yarin Gal, Rowan McAllister, and Carl Edward Rasmussen. Improving pilco with bayesian neural network dynamics models.
- Shixiang Gu, Sergey Levine, Ilya Sutskever, and Andriy Mnih. Muprop: Unbiased backpropagation for stochastic neural networks. *International Conference on Learning Representations (ICLR)*, 2016a.
- Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E Turner, and Sergey Levine. Q-prop: Sample-efficient policy gradient with an off-policy critic. *arXiv preprint arXiv:1611.02247*, 2016b.
- Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning*, pages 2829–2838, 2016c.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, 2018. URL <https://openreview.net/forum?id=HJjvxl-Cb>.
- Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado van Hasselt, and David Silver. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.
- Gabriel Kalweit and Joschka Boedecker. Uncertainty-driven imagination for continuous deep reinforcement learning. In *Conference on Robot Learning*, pages 195–206, 2017.
- Oleg Klimov and John Schulman. Roboschool. <https://github.com/openai/roboschool>.
- Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. Model-ensemble trust-region policy optimization. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SJJinbWRZ>.
- 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. *arXiv preprint arXiv:1509.02971*, 2015.
- David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.
- Gabriel Barth-Maron, Matthew W. Hoffman, David Budden, Will Dabney, Dan Horgan, Dhruva TB, Alistair Muldal, Nicolas Heess, and Timothy Lillicrap. Distributional policy gradients. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SyZipzbCb>.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472, 2011.
- Stefan Depeweg, José Miguel Hernández-Lobato, Finale Doshi-Velez, and Steffen Udluft. Learning and policy search in stochastic dynamical systems with bayesian neural networks. *arXiv preprint arXiv:1605.07127*, 2016.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim

- Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. pages 4026–4034, 2016.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Richard S Sutton. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Machine Learning Proceedings 1990*, pages 216–224. Elsevier, 1990.
- Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*, volume 1. MIT Press Cambridge, 1998.
- Théophane Weber, Sébastien Racanière, David P Reichert, Lars Buesing, Arthur Guez, Danilo Jimenez Rezende, Adria Puigdomènech Badia, Oriol Vinyals, Nicolas Heess, Yujia Li, et al. Imagination-augmented agents for deep reinforcement learning. *arXiv preprint arXiv:1707.06203*, 2017.

A Toy Problem: A Tabular FSM with Model Noise

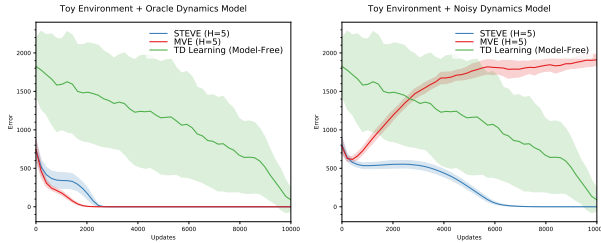


Figure 3: Value error per update on a value-estimation task in a toy environment. When given access to a perfect dynamics model, model-based approaches solve this task with $5\times$ fewer samples than model-free TD learning. However, when only given access to a noisy dynamics model, MVE diverges due to model errors. In contrast, STEVE converges to the correct solution, and does so with a $2\times$ speedup over TD learning. This is because STEVE dynamically adapts its rollout horizon to accommodate model error. See Appendix A for more details.

To demonstrate the benefits of Bayesian model-based value expansion, we evaluated it on a toy problem. We used a finite state environment with states $\{s_0, \dots, s_{100}\}$, and a single action A available at every state which always moves from state s_t to s_{t+1} , starting at s_0 and terminating at s_{100} . The reward for every action is -1 , except when moving from s_{99} to s_{100} , which is $+100$. Since this environment is so simple, there is only one possible policy π , which is deterministic. It is possible to compute the true action-value function in closed form, which is $Q^\pi(s_i, A) = i$.

We estimate the value of each state using tabular TD-learning. We maintain a tabular function $\hat{Q}^\pi(s_i, A)$, which is just a lookup table matching each state to its estimated value. We initialize all values to random integers between 0 and 99, except for the terminal state s_{100} , which we initialize to 0 (and keep fixed at 0 at all times). We update using the standard undiscounted one-step TD update, $\hat{Q}^\pi(s_i, A) = r + \hat{Q}^\pi(s_{i+1}, A)$. For each update, we sampled a nonterminal state and its corresponding transition (s_i, r, s_{i+1}) at random. For experiments with an ensemble of Q-functions, we repeat this update once for each member of the ensemble at each timestep.

The transition and reward function for the oracle dynamics model behaved exactly the same as the true environment. In the “noisy” dynamics model, noise was added in the following way: 10% of the time, rather than correctly moving from s_t to s_{t+1} , the model transitions to a random state. (Other techniques for adding noise gave qualitatively similar results.)

On the y-axis of Figure 3, we plot the mean squared error between the predicted values and the true values of each state: $\frac{1}{100} \sum_{i=0}^{99} (\hat{Q}^\pi(s_i, A) - Q^\pi(s_i, A))^2$.

For both the STEVE and MVE experiments, we use an ensemble of size 8 for both the model and the Q-function. To compute the MVE target, we average across all ensemble rollouts and predictions.

B More on Model-Based Value Expansion

While the results of Feinberg et al. [2018] are promising, they rely on task-specific tuning of the rollout horizon H . This sensitivity arises from the difficulty of modeling transitions and the Q-function, which are task-specific and may change throughout training as the policy explores different parts of the state space. Complex environments require much smaller rollout horizon H , which limits the effectiveness of the approach (e.g., Feinberg et al. [2018] used $H = 10$ for HalfCheetah-v1, but had to reduce to $H = 3$ on Walker2d-v1).

When the model is perfect and the learned Q-function has similar bias on all states and actions, Feinberg et al. [2018] show that the MVE target with rollout horizon H will decrease the target error by a factor of γ^{2H} Feinberg et al. [2018]. Errors in the learned model can lead to worse targets, so in practice, we must tune H to balance between the errors in the model and the Q-function estimates. An additional challenge is that the bias in the learned Q-function is not uniform across states and actions [Feinberg et al., 2018]. In particular, they find that the bias in the Q-function on states sampled from the replay buffer is lower than when the Q-function is evaluated on states generated from model rollouts. They term this the *distribution mismatch problem*, and propose the *TD-k trick* as a solution.

The TD-k trick, proposed by Feinberg et al. [2018], involves training the Q-function using every intermediate state of the rollout:

$$s'_{-1} = s$$

$$\mathbf{L}_\theta = \mathbb{E}_{(s,a,r,s')} \left[\frac{1}{H} \sum_{i=-1}^{H-1} (\hat{Q}_\theta^\pi(s'_i, a_i) - \mathcal{T}_H^{\text{MVE}}(r, s'))^2 \right]$$

To summarize Feinberg et al. [2018], the TD-k trick is helpful because the off-policy states collected by the replay buffer may have little overlap with the states encountered during on-policy model rollouts. Without the TD-k trick, the Q-function approximator is trained to minimize error only on states collected from the replay buffer, so it is likely to have high error on states not present in the

replay buffer. This would imply that the Q-function has high error on states produced by model rollouts, and that this error may in fact continue to increase the more steps of on-policy rollout we take. By invoking the TD-k trick, and training the Q-function on intermediate steps of the rollout, Feinberg et al. [2018] show that we can decrease the Q-function bias on frames encountered during model-based rollouts, leading to better targets and improved performance.

The TD-k trick is orthogonal to STEVE. STEVE tends to ignore estimates produced by states with poorly-learned Q-values, so it is not hurt nearly as much as MVE by the distribution mismatch problem. However, better Q-values will certainly provide more information with which to compute STEVE’s target, so in that regard the TD-k trick seems beneficial. An obvious question is whether these two approaches are complimentary. STEVE+TD-k is beyond the scope of this work, and we did not give it a rigorous treatment; however, initial experiments were not promising. In future work, we hope to explore the connection between these two approaches more deeply.

C Implementation Details

All algorithms were implemented in Tensorflow [Abadi et al., 2016] and run on Google Cloud Engine.

All models were feedforward neural networks with ReLU nonlinearities. The policy network, reward model, and termination model each had 4 layers of size 128, while the transition model had 8 layers of size 512. All environments were reset after 1000 timesteps. Parameters were trained with the Adam optimizer with a learning rate of $3e-4$.

Policies were trained using minibatches of size 512 sampled uniformly at random from a replay buffer of size $1e6$. The first $1e5$ frames were sampled via random interaction with the environment; after that, 4 policy updates were performed for every frame sampled from the environment. (In wall-clock experiments, the policy updates and frames were instead de-synced.) Policy checkpoints were saved every 500 updates; these checkpoints were also frozen and used as θ^- . For model-based algorithms, the most recent checkpoint of the model was loaded every 500 updates as well.

Each policy training had 8 agents interacting with the environment to send frames back to the replay buffer. These agents typically took the greedy action predicted by the policy, but with probability $\epsilon = 0.5$, instead took an action sampled from a normal distribution surrounding the pre-tanh logit predicted by the policy. In addition, each policy had 2 purely-greedy agents interacting with the environment for evaluation.

Dynamics models were trained using minibatches of size 1024 sampled uniformly at random from a replay buffer of size $1e6$. The first $1e5$ frames were sampled via random interaction with the environment; the dynamics model was then pre-trained for $1e5$ updates. After that, 4 model updates were performed for every frame sampled from the environment. (In wall-clock experiments, the model updates and frames were instead de-synced.) Model checkpoints were saved every 500 updates.

All ensembles were of size 4. During training, each ensemble member was trained on a disjoint subset of each minibatch, i.e. each sample in the minibatch was “assigned” to one member of the ensemble uniformly at random. Additionally, $M, N, L = 4$ for all experiments.

D A Note On Ensembles

This technique for calculating uncertainty estimates will work equally well for any family of models from which we can sample. For example, we could train a Bayesian neural network for each model MacKay [1992], or use dropout as a Bayesian approximation by resampling the dropout masks each time we wish to sample a new model Gal and Ghahramani [2016]. These options could potentially give better diversity of various samples from the family, and thus better uncertainty estimates; exploring them further is a promising direction for future work. However, we found that these methods degraded the accuracy of the base models. An ensemble is far easier to train, and so we focus on that in this work. This is a common choice, as the use of ensembles in the context of uncertainty estimations for deep reinforcement learning has seen wide adoption in the literature. It was first proposed by Osband et al. [2016] as a technique to improve exploration, and subsequent work showed that this approach gives a good estimate of the uncertainty of both value functions Kalweit and Boedecker [2017] and models Kurutach et al. [2018].

E Wall-Clock Comparison




Figure 4: Wall-clock time. The x-axis is hours elapsed since start of training; the y-axis is mean reward, shaded area is one standard deviation. Each experiment was run three times.

In the previous experiments, we synchronized data collection, policy updates, and model updates. However, when we run these steps asynchronously, we can reduce the wall-clock time at the risk of instability. To evaluate this configuration, we compare DDPG, MVE-DDPG, and STEVE-DPPG on Humanoid-v1 and RoboschoolHumanoidFlagrun-v1. Both were trained on a P100 GPU and had 8 CPUs collecting data; STEVE-DPPG additionally used a second P100 to learn a model in parallel. We plot reward as a function of wall-clock time for these tasks in Figure 4. These results demonstrate that in spite of the additional computation per epoch, the gains in sample efficiency are enough that STEVE is competitive with model-free algorithms in terms of wall-clock time. The speed gains associated with improved sample efficiency will only be exacerbated as samples become more expensive to collect, making STEVE a promising choice for applications involving real-world interaction. Moreover, in future work, STEVE could be accelerated by parallelizing training of each component of the ensemble.

F Model Usage




Figure 5: Average model usage for STEVE on each environment. The x-axis represents the number of frames sampled from the environment. The y-axis represents the amount of probability mass assigned to weights that were not w_0 , i.e. the probability mass assigned to candidate targets that include at least one step of model rollout.

Given that the improvements stem from the dynamic reweighting between horizon lengths, it may be interesting to examine the choices that the model makes about which candidate targets to favor most heavily. In Figure 5, we plot the average model usage over the course of training. Intriguingly, most of the lines seem to remain stable at around 50% usage, with two notable exceptions: Humanoid-v1, the most complex environment tested (with an observation-space of size 376); and Swimmer-v1, the least complex environment tested (with an observation-space of size 8). This supports the hypothesis that STEVE is trading off between Q-function bias and model bias; it chooses to ignore the model almost immediately when the environment is too complex to learn, and gradually begins ignore the model as the Q-function improves if an optimal environment model is learned quickly.

G Related Work

Sutton and Barto [1998] describe TD- λ , a family of Q-learning variants in which targets from multiple timesteps are merged via exponentially decay. STEVE is similar in that it is also computing a weighted average between targets. However, our approach is significantly more powerful because it adapts the weights to the specific characteristics of each individual rollout, rather than being constant between examples and throughout training. Our approach can be thought of as a generalization of TD- λ , in that the two approaches are equivalent in the specific case where the overall uncertainty grows exponentially at rate λ at every timestep.

Kurutach et al. [2018] propose *model-ensemble trust-region policy optimization* (ME-TRPO), which is motivated similarly to this work in that they also propose an algorithm which uses an ensemble of models to mitigate the deleterious effects of model bias. However, the algorithm proposed is quite different. ME-TRPO is a purely model-based policy-gradient approach, and uses the ensemble to avoid overfitting to any one model. In contrast, STEVE interpolates between model-free and model-based estimates, uses a value-estimation approach, and uses the ensemble to explicitly estimate uncertainty.

Kalweit and Boedecker [2017] train on a mix of real and imagined rollouts, and adjust the ratio over the course of training by tying it to the variance of the Q-function. Similarly to our work, this variance is computed via an ensemble. However, they do not adapt to the uncertainty of individual estimates, only the overall ratio of real to imagined data. Additionally, they do not take into account model bias, or uncertainty in model predictions.

Weber et al. [2017] use rollouts generated by the dynamics model as inputs to the policy function, by “summarizing” the outputs of the rollouts with a deep neural network. This second network allows the algorithm to implicitly calculate uncertainty over various parts of the rollout and use that information when making its decision. However, I2A has only been evaluated on discrete domains. Additionally, the lack of explicit model use likely tempers the sample-efficiency benefits gained relative to more traditional model-based learning.

Gal et al. use a deep neural network in combination with the PILCO algorithm Deisenroth and Rasmussen [2011] to do sample-efficient reinforcement learning. They demonstrate good performance on the continuous-control task of cartpole swing-up. They model uncertainty in the learned neural dynamics function using dropout as a Bayesian approximation, and provide evidence that maintaining these uncertainty estimates is very important for model-based reinforcement learning.

Depeweg et al. [2016] use a Bayesian neural network as the environment model in a policy search setting, learning a policy purely from imagined rollouts. This work also demonstrates that modeling uncertainty is important for model-based reinforcement learning with neural network models, and that uncertainty-aware models can escape many common pitfalls.

Gu et al. [2016c] propose a continuous variant of Q-learning known as *normalized advantage functions* (NAF), and show that learning using NAF can be accelerated by using a model-based component. They use a variant of Dyna-Q Sutton [1990], augmenting the experience available to the model-free learner with imaginary on-policy data generated via environment rollouts. They use an iLQG controller and a learned locally-linear model to plan over small, easily-modelled regions of the environment, but show that using more complex neural network models of the environment can yield errors that negate the benefits of model-based rollouts.