
Supplementary Material for Loaded DiCE: Trading off Bias and Variance in Any-Order Score Function Estimators for Reinforcement Learning

Anonymous Author(s)

Affiliation

Address

email

1 Derivation of Value Function formulation

2 We start out with the J_\diamond objective:

$$J_\diamond = \sum_{t=0}^T \gamma^t \left(\mathbb{Q}(a_{\leq t}) - \mathbb{Q}(a_{< t}) \right) R_{t+1}. \quad (1)$$

3 We evaluate this objective by taking an expectation over the trajectories τ as induced by the policy π .
 4 Here τ is a complete sequence of states, actions and rewards, $\tau = \{s_0, a_0, r_1, \dots, s_T, a_T\}$. Note that
 5 for convenience in the following derivation we have defined the reward, r_{t+1} , to arrive at the next
 6 time step, after action a_t was taken.

$$\mathbb{E}_\pi[J_\diamond] = \sum_{\tau} P(\tau) J_\diamond(\tau) \quad (2)$$

$$= \sum_{\tau} P(\tau) \left(\sum_{t=0}^T \gamma^t (\mathbb{Q}(a_{\leq t}) - \mathbb{Q}(a_{< t})) R_{t+1} \right) \quad (3)$$

$$= \sum_{t=0}^T \gamma^t \left(\sum_{\tau} P(\tau) (\mathbb{Q}(a_{\leq t}) - \mathbb{Q}(a_{< t})) R_{t+1} \right) \quad (4)$$

$$= \sum_{t=0}^T \gamma^t J_t \quad (5)$$

(6)

7 We note that for each time step the term, J_t is of the form:

$$J_t = \sum_{\tau} P(\tau) f(\tau_{\leq t}) g(\tau_{> t}), \quad (7)$$

8 where $f(a_{\leq t}, s_{\leq t}) = (\mathbb{Q}(a_{\leq t}) - \mathbb{Q}(a_{< t}))$ and $g(\tau_{> t}) = R_{t+1}$.

9 Next we use:

$$P(\tau) = P(\tau_{\leq t}) P(\tau_{> t} | \tau_{\leq t}) \quad (8)$$

$$= P(\tau_{\leq t}) P(\tau_{> t} | s_t, a_t), \quad (9)$$

10 where in the last step we have used the Markov property. Substituting we obtain:

$$J_t = \sum_{\tau} P(\tau \leq t) P(\tau > t | s_t, a_t) f(a_{\leq t}, s_{\leq t}) g(\tau > t) \quad (10)$$

$$= \sum_{\tau \leq t} P(\tau \leq t) f(a_{\leq t}, s_{\leq t}) \sum_{\tau > t} P(\tau > t | s_t, a_t) g(\tau > t) \quad (11)$$

$$(12)$$

11 If we substitute back for g and f we obtain:

$$J_t = \sum_{\tau \leq t} P(\tau \leq t) (\mathbb{E}[a_{\leq t}] - \mathbb{E}[a_{< t}]) \sum_{\tau > t} P(\tau > t | s_t, a_t) R_{t+1} \quad (13)$$

$$= \sum_{\tau \leq t} P(\tau \leq t) (\mathbb{E}[a_{\leq t}] - \mathbb{E}[a_{< t}]) \mathbb{E}[R_{t+1} | s_t, a_t] \quad (14)$$

$$= \sum_{\tau \leq t} P(\tau \leq t) (\mathbb{E}[a_{\leq t}] - \mathbb{E}[a_{< t}]) Q(s_t, a_t) \quad (15)$$

$$(16)$$

12 Putting all together we obtain the final form:

$$\mathbb{E}_{\pi}[J_{\diamond}] = \mathbb{E}_{\pi} \left[\sum_{t=0}^T \gamma^t (\mathbb{E}[a_{\leq t}] - \mathbb{E}[a_{< t}]) Q(s_t, a_t) \right] \quad (17)$$

13 2 Experimental Details

14 2.1 Random MDPs

15 We use the `mdptoolbox.example.rand()` function from PyMDPToolbox to generate random MDP
16 transition functions with five states and four actions per state.

17 The reward is a function only of state, and is sampled from $\mathcal{N}(5, 10)$. We use $\gamma = 0.95$. When
18 sampling for the stochastic estimators, we use batches of 512 rollouts of length 50 steps unless the
19 batch size is otherwise specified.

20 We only compute higher order derivatives of the derivative of the first parameter at each order, to
21 save computation.

22 For the sweeps over λ and τ we use 200 batches for each value of λ or τ . To simulate function
23 approximation error in our analysis of the impact of τ , we add a gaussian noise with standard
24 deviation 10 to the true value function.

25 2.2 MAML experiments

26 We use the following hyperparameters for our MAML experiments:

Parameter	Value
γ	0.97
hidden layer size	100
number of layers	2
task batch size	20 trajectories
meta batch size	40 tasks
inner loop learning rate	0.1
outer loop optimiser	Adam
outer loop learning rate	0.0005
outer loop τ	1.0
reward noise	Uniform(-0.01, 0.01) at each timestep

27 We also normalise all advantages in each batch (per task).