Adaptively Calibrated Critic Estimates for Deep Reinforcement Learning

Nicolai Dorka[1] and Tim Welschehold[1] and Joschka Bödecker[1] and Wolfram Burgard[2]

Abstract— **Accurate value estimates are important for off-policy reinforcement learning. Algorithms based on temporal difference learning typically are prone to an over- or underestimation bias building up over time. In this paper, we propose a general method called Adaptively Calibrated Critics (ACC) that uses the most recent high variance but unbiased on-policy rollouts to alleviate the bias of the low variance temporal difference targets. We apply ACC to Truncated Quantile Critics [1], which is an algorithm for continuous control that allows regulation of the bias with a hyperparameter tuned per environment. The resulting algorithm adaptively adjusts the parameter during training rendering hyperparameter search unnecessary and sets a new state of the art on the OpenAI gym continuous control benchmark among all algorithms that do not tune hyperparameters for each environment. ACC further achieves improved results on different tasks from the Meta-World robot benchmark. Additionally, we demonstrate the generality of ACC by applying it to TD3 [2] and showing an improved performance also in this setting.**

I. Introduction

Off-policy reinforcement learning is an important research direction as the reuse of old experience promises to make these methods more sample efficient than their on-policy counterparts. This is an important property for many applications such as robotics where interactions with the environment are very time- and cost-intensive. Many successful off-policy methods make use of a learned Q-value function [2], [3], [4], [5]. If the action space is discrete the Q-function can be directly used to generate actions while for continuous action spaces it is usually used in an actor-critic setting where the policy is trained to choose actions that maximize the Q-function. In both cases accurate estimates of the Q-values are of crucial importance.

Unfortunately, learning the Q-function off-policy can lead to an overestimation bias [6]. Especially when a nonlinear function approximator is used to model the Q-function, there are many potential sources of bias. Different heuristics were proposed for their mitigation, such as the double estimator in the case of discrete action spaces [7] or taking the minimum of two estimates in the case of continuous actions [2]. While these methods successfully prevent extreme overestimation, due to their coarse nature, they can still induce under- or overestimation bias to a varying degree depending on the environment [8].

To overcome these problems we propose a principled and general method to alleviate the bias called Adaptively Calibrated Critics (ACC). Our algorithm uses the most recent on-policy rollouts to determine the current bias of the Q-estimates and adjusts a bias controlling parameter accordingly. This parameter adapts the size of the temporal difference (TD) targets such that the bias can be corrected in the subsequent updates. As the parameter changes slower than the rollout returns, our method still benefits from stable and low-variance temporal difference targets, while it incorporates the information from unbiased but high variance samples from the recent policy to reduce the bias.

We apply ACC to Truncated Quantile Critics (TQC) [1], which is a recent off-policy actor-critic algorithm for continuous control showing strong performance on various tasks. In TQC the bias can be controlled in a finegrained way with the help of a hyperparameter that has to be tuned for every environment. ACC allows to automatically adjusts this parameter online during the training in the environment. As a result, it eliminates the need to tune this hyperparameter in a new environment, which is very expensive or even infeasible for many applications.

We evaluate our algorithm on a range of continuous control tasks from OpenAI gym [9] and robotic tasks from the meta world benchmark [10] and exceed the current state-of-the-art results among all algorithms that do not need tuning of environment-specific hyperparameters. For each environment, ACC matches the performance of TQC with the optimal hyperparameter for that environment. Further, we show that the automatic bias correction allows to increase the number of value function updates performed per environment step, which results in even larger performance gains in the sample-efficient regime. We additionally apply ACC to the TD3 algorithm [2] where it also leads to notably improved performance, underscoring the generality of our proposed method. To summarize, the main contributions of this work are:

1) We propose Adaptively Calibrated Critics, a new general algorithm that reduces the bias of value estimates in a principled fashion with the help of the most recent unbiased on-policy rollouts.
2) As a practical implementation we describe how ACC can be applied to learn a bias-controlling hyperparameter of the TQC algorithm and show that the resulting algorithm sets a new state of the art on the OpenAI continuous control benchmark suite.
3) ACC achieves strong performance on robotics tasks.
4) We demonstrate that ACC is a general algorithm with respect to the adjusted parameter by additionally applying it successfully to TD3.

To allow for reproducibility of our results we describe our

algorithm in detail, report all hyperparameters, use a large number of random seeds for evaluation, and made the source code publicly available[1].

II. BACKGROUND

We consider model-free reinforcement learning for episodic tasks with continuous state and action spaces $\mathcal{S}$ and $\mathcal{A}$. An agent interacts with its environment by selecting an action $a_t \in \mathcal{A}$ in state $s_t \in \mathcal{S}$ for every discrete time step t. The agent receives a scalar reward r_t and transitions to a new state s_{t+1}. To model this in a mathematical framework we use a Markov decision process, defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$. Given an action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$ the unknown state transition density $\mathcal{P}$ defines a distribution over the next state. Rewards come from the reward function $\mathcal{R}$ and future rewards are discounted via the discount factor $\gamma \in [0, 1]$.

The goal is to learn a policy π mapping a state s to a distribution over actions such that the sum of future discounted rewards $R_t = \sum_{i=t}^{T} \gamma^{i-t} r_i$ is maximized. We use the term π_ϕ for the policy with parameters ϕ trained to maximize the expected return $J(\phi) = \mathbb{E}_{s_i \sim \mathcal{P}, a_i \sim \pi}[R_0]$. The value function for a given state-action pair (s, a) is defined as $Q^\pi(s, a) = \mathbb{E}_{s_i \sim \mathcal{P}, a_i \sim \pi}[R_t | s, a]$, which is the expected return when executing action a in state s and following π afterwards.

A. Soft Actor Critic

TQC extends Soft Actor-Critic (SAC) [3], which is a strong off-policy algorithm for continuous control using entropy regularization. While in the end we are interested in maximizing the performance with respect to the total amount of reward collected in the environment, SAC maximizes for an auxiliary objective that augments the original reward with the entropy of the policy $J(\phi) = \mathbb{E}_{s_t \sim \mathcal{P}, a_t \sim \pi}[\sum_t \gamma^t(r_t + \alpha \mathcal{H}(\pi(\cdot|s_t)))]$, where $\mathcal{H}$ denotes the entropy.

A critic is learned that evaluates the policy π in terms of its Q-value of the entropy augmented reward. The policy—called actor—is trained to choose actions such that the Q-function is maximized with an additional entropy regularization

$$J_\pi(\phi) = \mathbb{E}_{s_t \sim \mathcal{D}, a_t \sim \pi_\phi}[Q_\theta(s_t, a_t) - \alpha \log \pi_\phi(a_t|s_t)]. \quad (1)$$

The weighting parameter α of the entropy term can be automatically adjusted during the training [11]. Both the training of actor and critic happen off-policy with transitions sampled from a replay buffer.

B. Truncated Quantile Critics

The TQC algorithm uses distributional reinforcement learning [12] to learn a distribution over the future augmented reward instead of a Q-function which is a point estimate for the expectation of this quantity. To do so TQC utilizes quantile regression [13] to approximate the distribution with Dirac delta functions $Z_\theta(s_t, a_t) = \frac{1}{M} \sum_{m=1}^{M} \delta(\theta^m(s_t, a_t))$. The Diracs are located at the quantile locations for fractions $\tau_m = \frac{2m-1}{m}, m \in \{1, \ldots, M\}$. The network is trained to learn the quantile locations $\theta^m(s, a)$ by regressing the predictions $\theta^m(s_t, a_t)$ onto the Bellman targets $y_m(s_t, a_t) =$

$r_t + \gamma(\theta^m(s_{t+1}, a_{t+1}) - \alpha \log \pi_\phi(a_{t+1}|s_{t+1}))$ via the Huber quantile loss.

TQC uses an ensemble of N networks $(\theta_1, \cdots, \theta_N)$ where each network θ_n predicts the distribution $Z_{\theta_n}(s_t, a_t) = \frac{1}{M} \sum_{m=1}^{M} \delta(\theta_n^m(s_t, a_t))$. A single Bellman target distribution is computed for all networks. This happens by first computing all targets for all networks, pooling all targets together in one set and sorting them in ascending order. Let $k \in \{1, \ldots, M\}$, then the kN smallest of these targets y_i are used to define the target distribution $Y(s_t, a_t) = \frac{1}{kN} \sum_{i=1}^{kN} \delta(y_i(s_t, a_t))$. The networks are trained by minimizing the quantile Huber loss which in this case is given by

$$L(s_t, a_t; \theta_n) = \frac{1}{kNM} \sum_{m,i=1}^{M,kN} \rho_{\tau_m}^H(y_i(s_t, a_t) - \theta_n^m(s_t, a_t)) \quad (2)$$

where $\rho_\tau^H(u) = |\tau - \mathbf{1}(u < 0)|\mathcal{L}_H^1(u)$ and $\mathcal{L}_H^1(u)$ is the Huber loss with parameter 1.

The rationale behind truncating some quantiles from the target distribution is to prevent overestimation bias. In TQC the number of dropped targets per network $d = M - k$ is a hyperparameter that has to be tuned per environment but allows for a finegrained control of the bias.

The policy is trained as in SAC by maximizing the entropy penalized estimate of the Q-value which is the expectation over the distribution obtained from the critic

$$J(\phi) = \mathbb{E}_{\substack{s \sim \mathcal{D} \\ a \sim \pi}}\left[\frac{1}{NM} \sum_{m,n=1}^{M,N} \theta_n^m(s, a) - \alpha \log \pi_\phi(a|s)\right]. \quad (3)$$

III. ADAPTIVELY CALIBRATED CRITICS

In this section, we will introduce the problem of estimation bias in TD learning, present our method ACC and demonstrate how it can be applied to TQC.

A. Over- and Underestimation Bias

The problem of overestimation bias in temporal difference learning with function approximation has been known for a long time [6]. In Q-learning [14] the predicted Q-value $Q(s_t, a_t)$ is regressed onto the target given by $y = r_t + \gamma \max_a Q(s_{t+1}, a)$. In the tabular case and under mild assumptions the Q-values converge to that of the optimal policy [14] with this update rule. However, using a function approximator to generate the Q-value introduces an approximation error. Even under the assumption of zero mean noise corruption of the Q-value $\mathbb{E}[\epsilon_a] = 0$, an overestimation bias occurs in the computation of the target value because of Jensen's inequality

$$\max_a Q(s_{t+1}, a) = \max_a \mathbb{E}[Q(s_{t+1}, a) + \epsilon_a]$$
$$\leq \mathbb{E}\left[\max_a \{Q(s_{t+1}, a) + \epsilon_a\}\right]. \quad (4)$$

In continuous action spaces it is impossible to take the maximum over all actions. The most successful algorithms rely on an actor-critic structure where the actor is trained to choose actions that maximize the Q-value [2], [3], [15]. So the actor can be interpreted an approximation to the argmax of the Q-value.

With deep neural networks as function approximators other problems such as over-generalization [5], [16] can occur

where the updates to $Q(s_t, a_t)$ also increases the target through $Q(s_{t+1}, a)$ for all a which could lead to divergence. There are many other potential sources for overestimation bias such as stochasticity of the environment [17] or computing the Q-target from actions that lie outside of the current training data distribution [18].

While for discrete action spaces the overestimation can be controlled with the double estimator [7], [17], it was shown that this estimator does not prevent overestimation when the action space is continuous [2]. As a solution the TD3 algorithm [2] uses the minimum of two separate estimators to compute the critic target. This approach was shown to prevent overestimation but can introduce an underestimation bias. In TQC [1] the problem is handled by dropping some targets from the pooled set of all targets of an ensemble of distributional critics. This allows for more finegrained control of over- or underestimation by choosing how many targets are dropped. TQC is able to achieve an impressive performance but the parameter d determining the number of dropped targets has to be set for each environment individually. This is highly undesirable for many applications since the hyperparameter sweep to determine a good choice of the parameter increases the actual number of environment interactions proportional to the number of hyperparameters tested. For many applications like robotics this makes the training prohibitively expensive.

B. Dynamically Adjusting the Bias

In the following we present a new general approach to adaptively control bias emerging in TD targets regardless of the source of the bias. Let $R^\pi(s, a)$ be the random variable denoting the sum of future discounted rewards when the agent starts in state s, executes action a and follows policy π afterwards. This means that the Q-value is defined as its expectation $Q^\pi(s, a) = \mathbb{E}[R^\pi(s, a)]$. For notational convenience we will drop the dependency on the policy π in the following. We start with the tabular case. Suppose for each state-action pair (s, a) we have a family $\{\hat{Q}_\beta(s, a)\}_{\beta \in [\beta_{min}, \beta_{max}] \subset \mathbb{R}}$ of estimators for $Q(s, a)$ with the property that $\hat{Q}_{\beta_{min}}(s, a) \leq Q(s, a) \leq \hat{Q}_{\beta_{max}}(s, a)$, where $Q(s, a)$ is the true Q-value of the policy π and Q_β a continuous monotone increasing function in β.

If we have samples $R_i(s, a)$ of the discounted returns $R(s, a)$, an unbiased estimator for $Q(s, a)$ is given by the average of the R_i through Monte Carlo estimation [19]. We further define the estimator $\hat{Q}_{\beta^*}(s, a)$, where β^* is given by

$$\beta^*(s, a) = \arg\min_{\beta \in [\beta_{min}, \beta_{max}]} \left| \hat{Q}_\beta(s, a) - \frac{1}{N} \sum_{i=1}^N R_i(s, a) \right|. \quad (5)$$

The following Theorem, which we prove in the appendix, shows that the estimator is unbiased under some assumptions.

Theorem 1: Let $Q_\beta(s, a)$ be a continuous monotone increasing function in β and assume that for all (s, a) it holds $\hat{Q}_{\beta_{min}}(s, a) \leq Q(s, a) \leq \hat{Q}_{\beta_{max}}(s, a)$, the returns $R(s, a)$ follow a symmetric probability distribution and that $\hat{Q}_{\beta_{min}}(s, a)$ and $\hat{Q}_{\beta_{max}}(s, a)$ have the same distance to $Q(s, a)$. Then Q_{β^*} from Equation 5 is an unbiased estimator for the true value Q for all (s, a).

Algorithm 1 ACC - General

Initialize: bias controlling parameter β, steps between β updates T_β, $t_\beta = 0$
for $t = 1$ **to** total number of environment steps **do**
    Interact with environment according to π, store transitions in replay buffer $\mathcal{B}$ and store observed returns $R(s, a)$, increment $t_\beta \mathrel{+}= 1$
    if episode ended **and** $t_\beta >= T_\beta$ **then**
        Update β with Eq. 6 using the most recent experience and set $t_\beta = 0$
    end if
    Sample mini-batch b from $\mathcal{B}$
    Update Q with target computed from Q_β and b
end for

The symmetry and same distance assumption can be replaced by assuming that $\hat{Q}_{\beta_{min}}(s, a) \leq R_i \leq \hat{Q}_{\beta_{max}}(s, a)$ with probability one. In this case the proof is straightforward since Q_β can take any value for which R_i has positive mass.

We are interested in the case where $\hat{Q}$ is given by a function approximator such that there is generalization between state-action pairs and that it is possible to generate estimates for pairs for which there are no samples of the return available. Consider off-policy TD learning where the samples for updates of the Q-function are sampled from a replay buffer of past experience. While the above assumptions might not hold anymore in this case, we have an estimator for all state-action pairs and not just the ones for which we have samples of the return. Also in practice rolling out the policy several times from each state action pair is undesirable and so we set $N = 1$ which allows the use of the actual exploration rollouts. Our proposed algorithm starts by initializing the bias-controlling parameter β to some value. After a number of environment steps and when the next episode is finished, the Q-value estimates and actual observed returns are compared. Depending on the difference β is adjusted according to

$$\beta_{new} = \beta_{old} + \alpha \sum_{t=1}^{T_\beta} \left[R(s_t, a_t) - \hat{Q}(s_t, a_t) \right], \quad (6)$$

where α is a step size parameter and $(s_t, a_t)_{t=1}^{T_\beta}$ are the $T_\beta \in \mathbb{N}$ most recent state-action pairs. As a result β is decreased in the case of overestimation, where the Q-estimates are larger than the actual observed returns, and increased in the case of underestimation. We assumed that Q_β is continuous and monotonically increasing in β. Hence, increasing β increases Q_β and vice versa. For updating the Q-function the target will be computed from Q_β.

Only performing one update step and not the complete minimization from Equation 5 has the advantage that β is changing relatively slow which means the targets are more stable. Through this mechanism our method can incorporate the high variance on-policy samples to correct for under- or overestimation bias. At the same time our method can benefit from the low variance TD targets. ACC in its general form is summarized in Algorithm 1.

Other algorithms that attempt to control the bias arising in

TD learning with non-linear function approximators usually use some kind of heuristic that includes more than one estimator. Some approaches use them to decouple the choice of the maximizing action and the evaluation of the maximum in the computation of the TD targets [7]. Alternative approaches take the minimum, maximum or a combination of both over the different estimators [2], [8], [20], [21]. All of these have in common that the same level of bias correction is done for every environment and for all time steps during training. In the deep case there are many different sources that can influence the tendency of TD learning building up bias in non-trivial ways. ACC is more principled in the regard that it allows to dynamically adjust the magnitude and direction of bias correction during training. Regardless of the source and amount of bias ACC provides a way to alleviate it. This makes ACC promising to work robustly on a wide range of different environments.

One assumption of ACC is that there is a way to adjust the estimated Q-value with a parameter β such that $\hat{Q}_\beta$ is continuous and monotonically increasing in β. There are many different functions that are in accordance with this assumption. We give one general example of how such a $\hat{Q}_\beta$ can be easily constructed for any algorithm that learns a Q-value. Let $\hat{Q}$ be the current estimate. Then define $\hat{Q}_\beta = \beta |\hat{Q}| / K + \hat{Q}$, where K is a constant (e.g. 100) and $[\beta_{min}, \beta_{max}]$ is some interval around 0. In the following section we will present an application of ACC in a more sophisticated way.

C. Applying ACC to TQC

As a practical instantiation of the general ACC algorithm we apply it to adjust the number of targets dropped from the set of all targets in TQC. Denote with $d_{max} \in \{0, \ldots, M\}$ some upper limit of targets to drop per network. Define $\beta_{min} = 0$, $\beta_{max} = d_{max}$ and let $d = d_{max} - \beta$ be the current number of targets dropped for each network. Further, we write Q_β for the TQC estimate with dN targets dropped from the pooled set of all targets. If d_{max} is set high enough the TQC estimate without dropped targets $Q_{\beta_{max}}$ induces overestimation while the TQC estimate with d_{max} dropped targets per net $Q_{\beta_{min}}$ induces underestimation.

In general, $\beta \in [0, d_{max}]$ is continuous and hence also d is a continuous value. As the number of dropped targets from the pooled set of all targets has to be a discrete number in $\{0, \ldots, NM\}$ we round the total number of dropped targets dN to the nearest integer in the computation of the TD target. When updating β with Equation 6, we divide the expectation by the moving average of the absolute value of the difference between returns and estimated Q-values for normalization.

IV. EXPERIMENTS

We evaluate our algorithm on a range of continuous control tasks from OpenAI Gym [9] and the meta world benchmark [10] that both use the physics engine MuJoCo [22] (version 1.5). First, we benchmark ACC against strong methods that do not use environment specific hyerparameters. Then we compare the performance of TQC with a fixed number of dropped targets per network with that of ACC. Next, we

evaluate the effect of more critic updates for ACC and show results in the sample efficient regime. Further, we study the effect of ACC on the accuracy of the value estimate, and investigate the generality of ACC by applying it to TD3.

We implemented ACC on top of the PyTorch code published by the authors[2] to ensure a fair comparison. While in general a safe strategy is to use a very high value for d_{max} as it gives ACC more flexibility in choosing the right amount of bias correction we set it to $d_{max} = 5$, which is the maximum value used by TQC for the number of dropped targets in the original publication. At the beginning of the training we initialize $\beta = 2.5$ and set the step size parameter to $\alpha = 0.1$. After $T_\beta = 1000$ steps since the last update and when the next episode finishes, β is updated with a batch that stores the most recent state-action pairs encountered in the environment and their corresponding observed discounted returns. After every update of β the oldest episodes in this stored batch are removed until there are no more than 5000 state-action pairs left. This means that on average β is updated with a batch whose size is a bit over 5000. The updates of β are started after 25000 environment steps and the moving average parameter in the normalization of the $\beta-$update is set to 0.05. The first 5000 environment interactions are generated with a random policy after which learning starts. We did not tune most of these additional hyperparameters and some choices are directly motivated by the environment (e.g. setting T_β to the maximum episode length). Only for α we tested a few different choices but found that for reasonable values it does not have a noticeable influence on performance. All hyperparameters of the underlying TQC algorithm with $N = 5$ critic networks were left unchanged.

Compared to TQC the additional computational overhead caused by ACC is minimal because there is only one update to β that is very cheap compared to one training step of the actor-critic and there are at least $T_\beta = 1000$ training steps in between one update to β.

During training, the policy is evaluated every 1,000 environment steps by averaging the episode returns of 10 rollouts with the current policy. For each task and algorithm we run 10 trials each with a different random seed.

A. Comparative Evaluation

We compare ACC to the state of the art continuous control methods SAC [3] (with learned temperature parameter [11]) and TD3 [2] on six OpenAI Gym continuous control environments. To make the different environments comparable we normalize the scores by dividing the achieved return by the best achieved return among all evaluations points of all algorithms for that environment.

Figure 1a) shows the aggregated data efficiency curve over all 6 tasks computed with the method of [23], where the interquantile mean (IQM) ignores the bottom and top 25% of the runs across all games and computes the mean over the remaining. The absolute performance of ACC for each single task can be seen in Figure 2. Overall, ACC reaches a much higer performance than SAC and TD3.

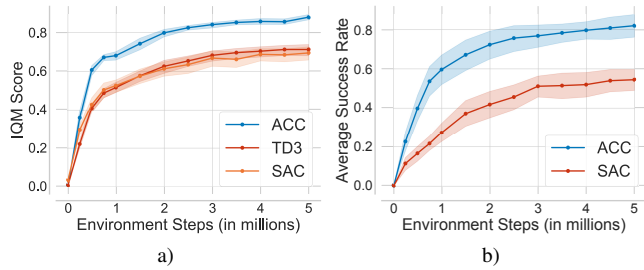[2]https://github.com/bayesgroup/tqc_pytorch

Fig. 1. Sample efficiency curves aggregated from the results over several environments. The normalized IQM score and the mean of the success rate respectively is plotted against the number of environment steps. Shaded regions denote pointwise 95% stratified bootstrap confidence intervals according to the method of [23]. **(a)** Aggregated results over the 6 gym continuous control tasks. **(b)** Aggregated results over the 12 metaworld tasks.

B. Robotics Benchmark

To investigate, if ACCs strong performance also translates into robotics environments, we evaluate ACC and SAC on 12 of the more challenging tasks in the Meta-World benchmark [10], which consists of several manipulation tasks with a Sawyer arm. We use version V2 and use the following 12 tasks: sweep, stick-pull, dial-turn, door-open, peg-insert-side, push, pick-out-of-hole, push-wall, faucet-open, hammer, stick-push, soccer. We evaluate the single tasks in the in the MT1 version of the benchmark, where the goal and object positions change across episodes. Different to the gym environments, β is updated every 500 environment steps as this is the episode length for these tasks. Figure 1b) shows the aggregated data efficiency curve in terms of success rate over all 12 tasks computed with the method of [23].

The curves demonstrate that ACC achieves drastically stronger results than SAC both in terms of data efficiency and asymptotic performance. After 2 million steps ACC already achieves a close to optimal task success rate which is even considerably higher than what SAC achieves at the end of the training. This shows, that ACC is a promising approach for real world robotics applications.

C. Fixing the Number of Dropped Targets

In this experiment we evaluate how well ACC performs when compared to TQC where the number of dropped targets per network d is fixed to some value. Since in the original publication for each environment the optimal value was one of the three values 0, 2, and 5, we evaluated TQC with d fixed to one of these values for each environment. To ensure comparability we used the same codebase as for ACC. The results in Figure 2 show that it is not possible to find one value for d that performs well on all environments. With $d = 0$, TQC is substantially worse on three environments and unstable on the *Ant* environment. Setting $d = 2$ is overall the best choice but still performs clearly worse for two environments and is also slightly worse for *Humanoid*. Dropping $d = 5$ targets per network leads to an algorithm that can compete with ACC only on two of the six environments. Furthermore, even if there would be one tuned parameter that performs equally well as ACC on a given set of environments we hypothesize there are likely very different environments for which the specific parameter choice will not perform well.
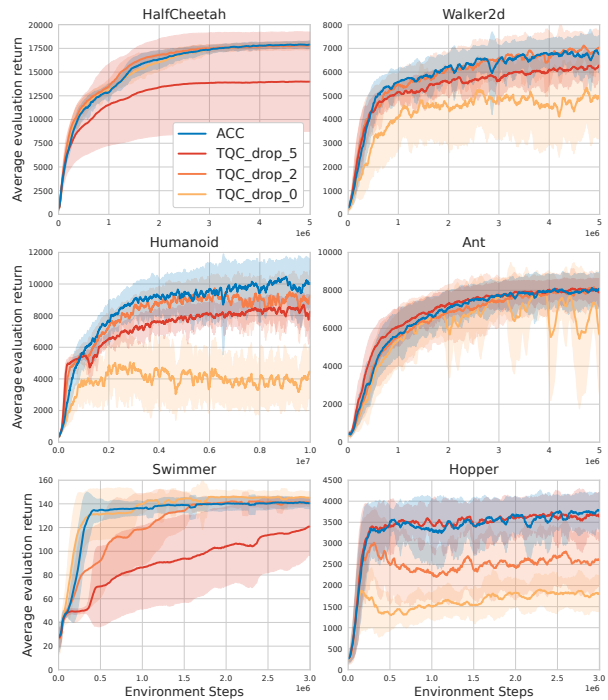

Fig. 2. Learning curves of ACC applied to TQC and TQC with different fixed choices for the number of dropped atoms d on six OpenAi gym environments. We used version *v3*. The shaded area represents mean $\pm$ standard deviation over the 10 trials. For readability the curves showing the mean are filtered with a uniform filter of size 15.

The principled nature of ACC on the other hand provides reason to believe that it can perform robustly on a wide range of different environments. This is supported by the robust performance on all considered environments.

D. Evaluation of Sample Efficient Variant

In principle more critic updates per environment step should make learning faster. However, because of the bootstrapping in the target computation this can easily become unstable. The problem is that as targets are changing faster, bias can build up easier and divergence becomes more likely. ACC provides a way to detect upbuilding bias in the TD targets and to correct the bias accordingly. This motivates to increase the number of gradient updates of the critic. In TD3, SAC and TQC one critic update is performed per environment step. We conducted an experiment to study the effect of increasing this rate up to 4. ACC using 4, 2 and 1 updates are denoted with ACC_4q, ACC_2q and ACC_1q respectively. ACC_1q is equal to ACC from the previous experiments. We use the same notation also for TD3 and SAC.

Scaling the number of critic updates by a factor of 4 increases the computation time by a factor of 4. But this can be worthwhile in the sample efficient regime, where a huge number of environment interactions is not possible or the interaction cost dominate the computational costs as it is the case when training robots in the real world. The results in Figure 3a) show that in the sample efficient regime ACC4q further increases over plain ACC. ACC4q reaches the final performance of TD3 and SAC in less than a third of the number of steps for five environments and for
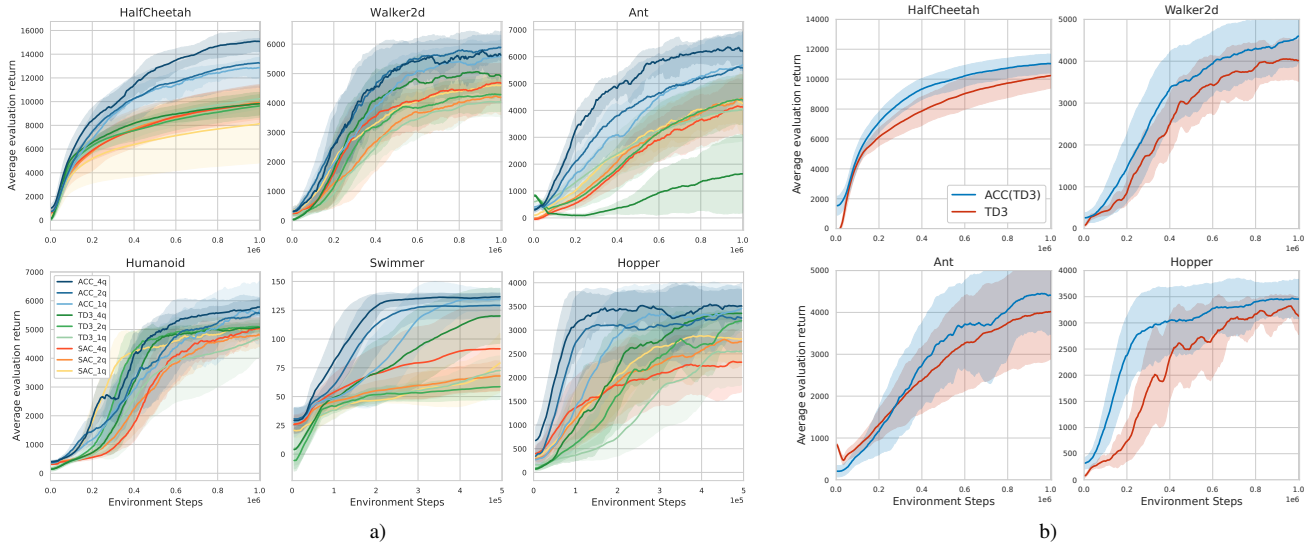
Fig. 3. The mean ± standard deviation over 10 trials. **(a)** Results in the sample efficient regime where tuning of hyperparameters in an inner loop is too costly with different choices for the number of value function updates per environment step. **(b)** Results for ACC applied to TD3 compared to pure TD3.

Humanoid in roughly half the number of steps. Increasing the number of critic updates for TD3 and SAC shows mixed results, increasing performance for some environments while decreasing it for others. Only ACC benefits from more updates on all environments, which supports the hypothesis that ACC is successful at calibrating the critic estimate.

E. Analysis of ACC

To evaluate the effect of ACC on the bias of the value estimate, we analyze the difference between the value estimate and the corresponding observed return when ACC is applied to TQC. For each state-action pair encountered during exploration, we compute its value estimate at that time and at the end of the episode compare it with the actual discounted return from that state onwards. Hence, the state-action pair was not used to update the value function at the point when the value estimate has been computed. If an episode ends because the maximum number of episode time-steps has been reached, which is 1,000 for the considered environments, we ignore the last 100 state-action pairs. The reason is that in TQC the value estimator is trained to ignore the episode timeout and uses a bootstrapped target also at the end of the episode. We normalize for different value scales by computing the absolute error between the value estimate and the observed discounted return and divide that by the absolute value of the discounted return. Every 1,000 steps, the average over the errors of the last 1,000 state-action pairs is computed. The aggregated results in Figure 4b) show that averaged over all environments ACC indeed achieves a lower value error than TQC with the a fixed number of dropped atoms d. This supports our hypothesis that the strong performance of ACC applied to TQC indeed stems from better values estimates.

To better understand the hidden training dynamics of ACC we show in Figure 4a) how the number of dropped targets per network $d = d_{max} - \beta$ evolves during training. Interestingly, the relatively low standard deviation indicates a similar behaviour across runs for a specific environment. However, there are large differences between the environments

which indicates that it might not be possible to find a single hyperparameter that works well on a wide variety of different environments. Further, the experiments shows that the optimal amount of overestimation correction might change over time during the training even on a single environment.

F. Beyond TQC: Improving TD3 with ACC

To demonstrate the generality of ACC, we additionally applied it to the actor-critic style TD3 algorithm [2], which uses two critics. These are initialized differently but trained with the same target value, which is the minimum over the two targets computed from the two critics. While this successfully prevents the overestimation bias, using the minimum of the two target estimates is very coarse and can instead lead to an underestimation bias. We applied ACC to TD3 by defining the target for each critic network to be a convex combination between its own target and the minimum over both targets. Let $Q_i = Q_{\bar{\theta}_i}(s_{t+1}, \pi_{\bar{\phi}}(s_{t+1}))$, we define the k-th critic target

$$y_k = r + \gamma \Big(\beta \, Q_k + (1 - \beta) \min_{i=1,2} Q_i \Big),$$

where $\beta \in [0, 1]$ is the ACC parameter that is adjusted to balance between under- and overestimation. The results are displayed in Figure 3b) and show that ACC also improves the performance of TD3.

V. RELATED WORK

A. Overestimation in Reinforcement Learning

The problem of overestimation in Q-learning with function approximation was introduced by [6]. For discrete actions the double estimator has been proposed [17] where two Q-functions are learned and one is used to determine the maximizing action, while the other evaluates the Q-function for that action. The Double DQN algorithm extended this to neural networks [7]. However, Zhang *et al.* [24] observed that the double estimator sometimes underestimates the Q-value and propose to use a weighted average of the single and the double estimator as target. This work is similar to

Fig. 4. **(a)** Mean (thick line) and standard deviation (shaded area) over 10 trials of the number of dropped targets per network $d = d_{max} - \beta$ in ACC over time for different environments with a uniform filter of size 15. **(b)** The normalized absolute error of the value estimate aggregated over the 6 environments. Shown are the mean with stratified bootstrapped confidence intervals computed from the results of 5 trials per environment. We used a uniform filter of size 401 for readability.

ours in the regard that depending on the parameter over- or underestimation could be corrected. A major difference to our algorithm is that the weighting parameter is computed from the maximum and minimum of the estimated Q-value and does not use unbiased rollouts. Similarly, the weighted estimator [25], [26] estimates the maximum over actions in the TD target as the sum of values weighted by their probability of being the maximum. In continuous action spaces this can be done through Gaussian process regression [26] and for discrete actions via dropout variational inference [25]. Different to ACC the weighting is computed from the same off-policy data used to compute the single quantities while ACC adjusts the weighting parameter β in a separate process using the latest on-policy rollouts. Lv *et al.* [27] use a similar weighting but suggest a stochastic selection of either the single or double estimator. The probability of choosing one or the other follows a predefined schedule. Other approaches compute the weighted average of the minimum and maximum over different Q-value estimates [21], [18]. However, the weighting parameter is a fixed hyperparameter. The TD3 algorithm [2] uses the minimum over two Q-value estimates as TD target. Maxmin Q-learning is another approach for discrete action spaces using an ensemble of Q-functions. For the TD target, first the minimum of over all Q-functions is computed followed by maximization with respect to the action [8]. Decreasing the ensemble size increases the estimated targets while increasing the size decreases the targets. Similarly to TQC this provides a way to control the bias in a more fine-grained way; the respective hyperparameter has to be set before the start of the training for each environment, however. Cetin *et al.* [28] propose to learn a pessimistic penalty to overcome the overestimation bias.

What sets ACC apart from the previously mentioned works is that unbiased on-policy rollouts are used to adjust a term that controls the bias correction instead of using some predefined heuristic.

B. Combining On- and Off-Policy Learning

There are many approaches that combine on- and off-policy learning by combining policy gradients with off-policy samples [29], [30], [31]. In [32] an actor-critic is used where the critic is updated off-policy and the actor is updated with a mixture of policy gradient and Q-gradient. This differs from our work in that we are interested only in better critic estimates through the information of on-policy samples. To learn better value estimates by combining on- and off-policy

data prior works proposed the use of some form of importance sampling [33], [34]. In [35] the TD target is computed by mixing Monte Carlo samples with the bootstrap estimator. These methods provide a tradeoff between variance and bias. They differ from our work in using the actual returns directly in the TD targets while we incorporate the returns indirectly via another parameter. Bhatt *et al.* [36] propose the use of a mixture of on- and off-policy transitions to generate a feature normalization that can be used in off-policy TD learning. Applied to TD3, learning becomes more stable eliminating the need to use a delayed target network.

C. Hyperparameter Tuning for Reinforcement Learning

Most algorithms that tune hyperparameters of RL algorithms use many different instances of the environment to find a good setting [37], [38], [39]. There is, however, also work that adjusts a hyperparameter online during training [40]. In this work the meta-gradient (i.e., the gradient of the update rule) is used to adjust the discount factor and the length of bootstrapping intervals. However, it would not be straightforward to apply this method to control the bias of the value estimate. Their method also differs from ours in that they do not use a combination of on- and off-policy data.

VI. CONCLUSION

We present Adaptively Calibrated Critics (ACC), a general off-policy algorithm that learns a Q-value function with bias calibrated TD targets. The bias correction in the targets is determined via a parameter that is adjusted by comparing the current value estimates with the most recently observed on-policy returns. Our method incorporates information from the unbiased sample returns into the TD targets while keeping the high variance of the samples out. We apply ACC to TQC, a recent off-policy continuous control algorithm that allows fine-grained control of the TD target scale through a hyperparameter tuned per environment. With ACC, this parameter can automatically be adjusted during training, obviating the need for extensive tuning. The strong experimental results suggest that our method provides an efficient and general way to control the bias occurring in TD learning.

Interesting directions for future research are to evaluate the effectiveness of ACC applied to algorithms that work with discrete action spaces and when learning on a real robot where tuning of hyperparameters is very costly.

REFERENCES

[1] A. Kuznetsov, P. Shvechikov, A. Grishin, and D. Vetrov, "Controlling overestimation bias with truncated mixture of continuous distributional quantile critics," in *International Conference on Machine Learning*. PMLR, 2020, pp. 5556–5566.

[2] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *International Conference on Machine Learning*, 2018, pp. 1582–1591.

[3] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 1861–1870.

[4] M. Hessel, J. Modayil, H. Van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. Azar, and D. Silver, "Rainbow: Combining improvements in deep reinforcement learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.

[5] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, p. 529, 2015.

[6] S. Thrun and A. Schwartz, "Issues in using function approximation for reinforcement learning," in *Proceedings of the 1993 Connectionist Models Summer School*, 1993, pp. 255–263.

[7] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

[8] Q. Lan, Y. Pan, A. Fyshe, and M. White, "Maxmin q-learning: Controlling the estimation bias of q-learning," in *International Conference on Learning Representations*, 2020. [Online]. Available: https://openreview.net/forum?id=Bkg0u3Etwr

[9] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.

[10] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, "Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning," in *Conference on Robot Learning*. PMLR, 2020, pp. 1094–1100.

[11] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine, "Soft actor-critic algorithms and applications," *CoRR*, vol. abs/1812.05905, 2018. [Online]. Available: http://arxiv.org/abs/1812.05905

[12] M. G. Bellemare, W. Dabney, and R. Munos, "A distributional perspective on reinforcement learning," in *International Conference on Machine Learning*, 2017, pp. 449–458.

[13] W. Dabney, M. Rowland, M. Bellemare, and R. Munos, "Distributional reinforcement learning with quantile regression," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.

[14] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.

[15] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.

[16] I. Durugkar and P. Stone, "Td learning with constrained gradients," in *Proceedings of the Deep Reinforcement Learning Symposium, NIPS 2017*, Long Beach, CA, USA, December 2017. [Online]. Available: http://www.cs.utexas.edu/users/ai-lab?NIPS17-ishand

[17] H. V. Hasselt, "Double q-learning," in *Advances in Neural Information Processing Systems*, 2010, pp. 2613–2621.

[18] A. Kumar, J. Fu, M. Soh, G. Tucker, and S. Levine, "Stabilizing off-policy q-learning via bootstrapping error reduction," in *Advances in Neural Information Processing Systems*, vol. 32, 2019, pp. 11 784–11 794.

[19] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT Press, 2018.

[20] R. Agarwal, D. Schuurmans, and M. Norouzi, "An optimistic perspective on offline reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2020, pp. 104–114.

[21] S. Fujimoto, D. Meger, and D. Precup, "Off-policy deep reinforcement learning without exploration," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2052–2062.

[22] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control." in *IROS*. IEEE, 2012, pp. 5026–5033.

[23] R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. Bellemare, "Deep reinforcement learning at the edge of the statistical precipice," *Advances in Neural Information Processing Systems*, vol. 34, 2021.

[24] Z. Zhang, Z. Pan, and M. J. Kochenderfer, "Weighted double q-learning," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, ser. IJCAI'17. AAAI Press, 2017, pp. 3455–3461. [Online]. Available: http://dl.acm.org/citation.cfm?id=3172077.3172372

[25] A. Cini, C. D'Eramo, J. Peters, and C. Alippi, "Deep reinforcement learning with weighted q-learning," *arXiv preprint arXiv:2003.09280*, 2020.

[26] C. D'Eramo, A. Nuara, M. Pirotta, and M. Restelli, "Estimating the maximum expected value in continuous reinforcement learning problems," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.

[27] P. Lv, X. Wang, Y. Cheng, and Z. Duan, "Stochastic double deep q-network," *IEEE Access*, vol. 7, pp. 79 446–79 454, 2019.

[28] E. Cetin and O. Celiktutan, "Learning pessimism for robust and efficient off-policy reinforcement learning," *arXiv preprint arXiv:2110.03375*, 2021.

[29] T. Degris, M. White, and R. Sutton, "Off-policy actor-critic," in *International Conference on Machine Learning*, 2012.

[30] T. Jie and P. Abbeel, "On a connection between importance sampling and the likelihood ratio policy gradient," in *Advances in Neural Information Processing Systems*, J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, Eds., vol. 23, 2010, pp. 1000–1008.

[31] B. O'Donoghue, R. Munos, K. Kavukcuoglu, and V. Mnih, "Combining policy gradient and q-learning," in *ICLR*, 2016.

[32] S. S. Gu, T. Lillicrap, R. E. Turner, Z. Ghahramani, B. Schölkopf, and S. Levine, "Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 3846–3855.

[33] A. R. Mahmood, H. P. van Hasselt, and R. S. Sutton, "Weighted importance sampling for off-policy learning with linear function approximation," in *Advances in Neural Information Processing Systems*, vol. 27, 2014, pp. 3014–3022.

[34] D. Precup, "Eligibility traces for off-policy policy evaluation," *Computer Science Department Faculty Publication Series*, p. 80, 2000.

[35] M. Hausknecht and P. Stone, "On-policy vs. off-policy updates for deep reinforcement learning," in *Deep Reinforcement Learning: Frontiers and Challenges, IJCAI 2016 Workshop*, 2016.

[36] A. Bhatt, M. Argus, A. Amiranashvili, and T. Brox, "Crossnorm: Normalization for off-policy td reinforcement learning," *arXiv preprint arXiv:1902.05605*, 2019.

[37] H.-T. L. Chiang, A. Faust, M. Fiser, and A. Francis, "Learning navigation behaviors end-to-end with autorl," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2007–2014, 2019.

[38] S. Falkner, A. Klein, and F. Hutter, "BOHB: Robust and efficient hyperparameter optimization at scale," in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 1437–1446.

[39] M. Jaderberg, V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, K. Simonyan, *et al.*, "Population based training of neural networks," *arXiv preprint arXiv:1711.09846*, 2017.

[40] Z. Xu, H. P. van Hasselt, and D. Silver, "Meta-gradient reinforcement learning," in *NeurIPS*, 2018.

Adaptively Calibrated Critic Estimates for Deep Reinforcement Learning

- Appendix -

Nicolai Dorka   Tim Welschehold   Joschka Bödecker   Wolfram Burgard

A.1. PROOF OF THEOREM 1

The estimator $\hat{Q}^{\pi}_{\beta^*}(s,a)$ was defined via

$$\beta^*(s,a) = \arg\min_{\beta \in [\beta_{min}, \beta_{min}]} \left| \hat{Q}_\beta(s,a) - \frac{1}{N}\sum_{i=1}^{N} R_i(s,a) \right|. \quad (7)$$

To declutter the notation we drop the dependencies on the state-action pairs (s,a) and the policy π. Further we write $\bar{R} = \frac{1}{N}\sum_{i=1}^{N} R_i$. First note that the average of symmetrically distributed random variables is still a symmetric distributed random variable and hence $\bar{R}$ is symmetrically distributed. By assumption $\hat{Q}_{\beta_{min}}$ and $\hat{Q}_{\beta_{max}}$ have the same distance to the true Q-value which is the mean $Q = \mathbb{E}[\bar{R}]$, i.e. there is a distance real valued value d such that $Q = \hat{Q}_{\beta_{min}} + d = \hat{Q}_{\beta_{max}} - d$ Denote the tail probabilty $P(\bar{R} < \hat{Q}_{\beta_{min}}) = p_t$. Because of the symmetry and the same distance to the mean we also have that $P(\bar{R} > \hat{Q}_{\beta_{max}}) = p_t$. In the computation of $\mathbb{E}[\hat{Q}_{\beta^*}]$ we can differentiate three events. If $\hat{Q}_{\beta_{min}} \leq \bar{R} \leq \hat{Q}_{\beta_{max}}$ then $\hat{Q}_{\beta^*} = \bar{R}$, if $\hat{Q}_{\beta_{min}} \geq \bar{R}$ then $\hat{Q}_{\beta^*} = \hat{Q}_{\beta_{min}}$ and if $\hat{Q}_{\beta_{min}} \geq \bar{R}$ then $\hat{Q}_{\beta^*} = \hat{Q}_{\beta_{max}}$. We denote the indicator function with $1[A]$, which is equal to 1 if the event A is true and 0 otherwise. Then we get

$$\mathbb{E}\left[\hat{Q}_{\beta^*}\right] = \mathbb{E}\left[1\left[\hat{Q}_{\beta_{min}} \leq \bar{R} \leq \hat{Q}_{\beta_{max}}\right]\bar{R}\right]$$

$$+ \mathbb{E}\left[1\left[\hat{Q}_{\beta_{min}} \geq \bar{R}\right]\hat{Q}_{\beta_{min}}\right]$$

$$+ \mathbb{E}\left[1\left[\hat{Q}_{\beta_{max}} \leq \bar{R}\right]\hat{Q}_{\beta_{max}}\right]$$

$$= (1 - 2p_t) \cdot \mathbb{E}[\bar{R}] + p_t \mathbb{E}\left[\hat{Q}_{\beta_{min}}\right] + p_t \mathbb{E}\left[\hat{Q}_{\beta_{max}}\right]$$

$$= (1 - 2p_t)Q + p_t\hat{Q}_{\beta_{min}} + p_t\hat{Q}_{\beta_{max}}$$

$$= (1 - 2p_t)Q + p_t(Q - d) + p_t(Q + d)$$

$$= (1 - 2p_t)Q + 2p_tQ + p_t(d - d)$$

$$= Q.$$

A.2. PSEUDOCODE

The pseudocode for ACC applied to TQC is in Algorithm 2. As the number of dropped targets per network is given by $d = d_{\max} - \beta$, we state the pseudocode in terms of the parameter d instead of β.

Algorithm 2 ACC - Applied to TQC

Initialize: d the bias controlling parameter, α the learning rate for d, T_d the minimum number of steps between updates to d, T_d^{init} the initial steps before d is updated, S_R the size from which on episodes are removed from the batch storing the most recent returns, moving average parameter τ_d, $t_d = 0$

for $t = 1$ **to** total number of environment steps **do**

   Interact with environment according to π, store transitions in replay buffer $\mathcal{B}$ and, increment $t_d += 1$

   if episode ended **then**

     Store observed returns $R(s,a)$ and corresponding state-action pairs (s,a) in $\mathcal{B}_R$

     if $t_d >= T_d$ **and** $t > T_d^{init}$ **then**

       $C = \sum_{(s,a,R) \in \mathcal{B}_R}\left[Q(s,a) - R(s,a)\right]$, $ma = (1 - \tau_d)ma + \tau_d C$

       $d = d + \alpha\frac{C}{ma}$, clip d in interval $[0, d_{max}]$, set $t_d = 0$

       Remove the oldest episodes from $\mathcal{B}_R$ until there are at most S_R left

     end if

   end if

   Sample mini-batch from $\mathcal{B}$

   Update critic Q as in TQC, where dN (rounded to the next integer) number of targets are dropped from the set of pooled targets

   Update policy π as in TQC

end for

A.3. USING FEWER CRITIC NETWORKS FOR FASTER RUNTIME

Using 5 critic networks - the default in TQC - to approximate the value function leads to a high runtime of the algorithm. It is possible to trade off performance against runtime by changing the number of critic networks. We evaluated ACC applied to TQC with 2 networks and compare it to the standard setting with 5 networks in Figure 5. The results show that reducing the number of critic networks to 2 leads only to a small drop in performance while the runtime is more than 2 times faster.

A.4. HYPERPARAMETERS

At the beginning of the training we initialize $\beta = 2.5$ and set the step size parameter to $\alpha = 0.1$. After $T_\beta = 1000$ steps since the last update and when the next episode finishes, β is
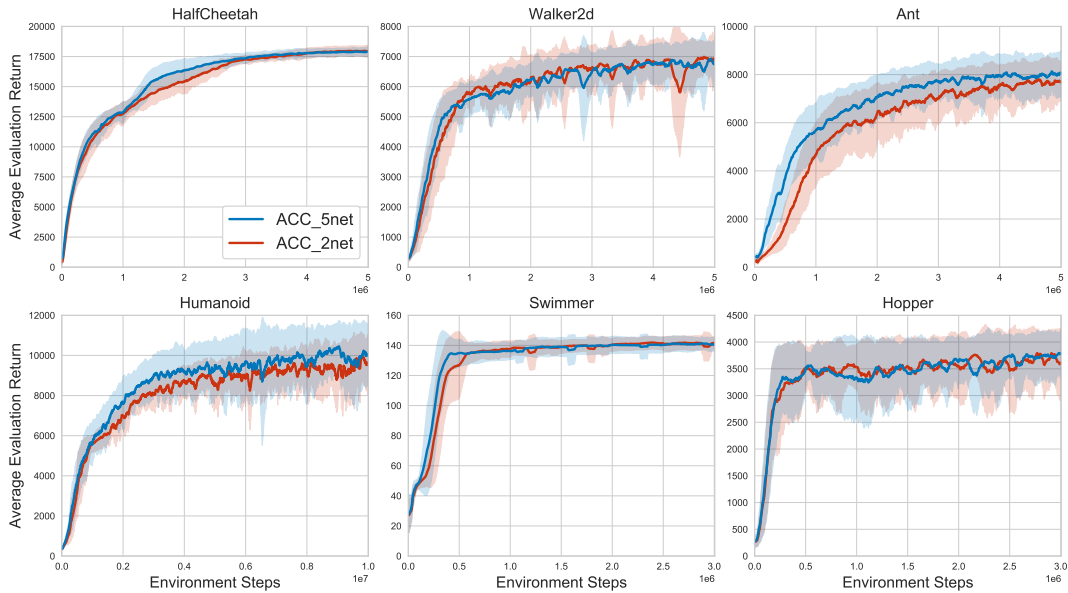
Fig. 5. The mean ± standard deviation over 10 trials. Results with different choices for the number of critic networks for each algorithm.

updated with a batch that stores the most recent state-action pairs encountered in the environment and their corresponding observed discounted returns. The choice of T_β was motivated by the fact that the maximum duration of an episode is 1000 steps for the considered environments. After every update of β the oldest episodes in this stored batch are removed until there are no more than 5000 state-action pairs left. This means that on average β is updated with a batch whose size is a bit over 5000. The updates of β are started as soon as 25000 environment steps as completed and the moving average parameter in the normalization of the $\beta-$update is set to 0.05. The first 5000 environment interactions are generated with a random policy after which learning starts. Apart from that all hyperparameters are the same as in TQC with $N = 5$ critic networks. In Table I we list all hyperparameters of ACC applied to TQC.

In the following we also desribe the process of hyperparameter selection. The range of values d is allowed to take is set to the interval $[0, 5]$ as it includes the optimal hyperparameters for TQC from all environments, which are in the set $\{0, 2, 5\}$. We did not try higher values than 5. The initial value for number of dropped targets per network was set to 2.5 as this value is in the middle of the allowed range and did not evaluated other choices. The learning rate α of d was set to 0.1 based on visual inspection of how fast d changes. We evaluated $\alpha = 0.05$ for a small subset of tasks and seeds, but $\alpha = 0.1$ gave slightly better results. T_d was set to 1000 as the episode length is 1000 and we did not evaluate other choices. For T_d^{init} we evaluated the choices 10000 and 25000 on a small subset of environments and seeds and did not found a big impact on performance. As d changes very quickly in the beginning we chose $T_d^{init} = 25000$. For S_R we evaluated the choices 1000 and 5000 also on a small subset of environments and seeds and found 5000 to perform

slightly better. We did not tune the moving average parameter and set it to $\tau_d = 0.05$. For all hyperparameters for which we evaluated more than one choice we do not have definite results as the number of seeds and environments were limited. The hyperparameters shared with TQC were not changed. For TD3 and SAC we used the hyperparameters from the respective papers.

A.5. POTENTIAL LIMITATIONS

One limitation of our work is that ACC can not be applied in the offline RL setting, as ACC also uses on-policy data. Furthermore, in the stated form ACC relies on the episodic RL setting. However, we believe that ACC could potentially be adapted to that setting. It is also not entirely clear how the algorithm would perform in the terminal reward setting, where a reward of for example 1 is given upon successful completion of a specific task. While we do not have experiments for such environments we imagine that the positive effect of ACC could diminish as the true Q-values of states closer to the start of the episode are almost zero because of the discounting.

A.6. ANALYSIS OF THE ACC PARAMETER

To better understand the hidden training dynamics of ACC we show in Figure 6 how the number of dropped targets per network $d = d_{max} - \beta$ evolves during training. To do so we plotted d after every 5000 steps during the training of ACC. From the top row the first observation is that per environment the results are similar over the 10 seeds as can be seen from the relatively low standard deviation. We show the single runs for all seeds in the appendix to further support this observation. However, there are large differences between the environments which supports the argument that it might not be possible to find a single hyperparameter that works well on a wide variety of different environments. Another

HYPERPARAMETER	ACC		
OPTIMIZER	ADAM		
LEARNING RATE	3×10^{-4}		
DISCOUNT γ	0.99		
REPLAY BUFFER SIZE	1×10^6		
NUMBER OF CRITICS N	5		
NUMBER OF ATOMS M	25		
HUBER LOSS PARAMETER	1		
NUMBER OF HIDDEN LAYERS IN CRITIC NETWORKS	3		
SIZE OF HIDDEN LAYERS IN CRITIC NETWORKS	512		
NUMBER OF HIDDEN LAYERS IN POLICY NETWORK	2		
SIZE OF HIDDEN LAYERS IN POLICY NETWORK	256		
MINIBATCH SIZE	256		
ENTROPY TARGET	$-\dim \mathcal{A}$		
NONLINEARITY	ReLU		
TARGET SMOOTHING COEFFICIENT	0.005		
TARGET UPDATES PER CRITIC GRADIENT STEP	1		
CRITIC GRADIENT STEPS PER ITERATION	1		
ACTOR GRADIENT STEPS PER ITERATION	1		
ENVIRONMENT STEPS PER ITERATION	1		
INITIAL VALUE FOR NUMBER OF DROPPED TARGETS PER NETWORK	2.5		
MAXIMUM VALUE FOR d DENOTED $d_{\max}$	5		
MINIMUM VALUE FOR d DENOTED $d_{\min}$	0		
LEARNING RATE FOR d DENOTED α	0.1		
MINIMUM NUMBER OF STEPS BETWEEN UPDATES TO d DENOTED T_d	1000		
INITIAL NUMBER OF STEPS BEFORE d IS UPDATED DENOTED T_d^{init}	25000		
LIMITING SIZE FOR BATCH USED TO UPDATE d DENOTED S_R	5000		
MOVING AVERAGE PARAMETER τ_d	0.05		
HYPERPARAMETER IN SAMPLE EFFICIENT EXPERIMENT	ACC_1Q	ACC_2Q	ACC_4Q
CRITIC GRADIENT STEPS PER ITERATION	1	2	4
ACTOR GRADIENT STEPS PER ITERATION	1	1	1
TARGET UPDATES PER CRITIC GRADIENT STEP	1	1	1

point that becomes clear from the plots is that the optimal amount of overestimation correction might change over time during the training even on a single environment.

In the bottom row of Figure 6 we plotted the evolution of d for one of the 10 trials in order to shed light on the actual training mechanics of a single run without lost information due to averaging. For each environment there is a trend but d is also fluctuating to a certain degree. While this shows that the initial value of d is not very important as the value quickly changes, this also highlights another interesting aspect of ACC. The rollouts give highly fluctuating returns. The parameter $d = d_{max} - \beta$ is changing more slowly and picks up the trend. So a lot of the variance of the returns is filtered out in ACC by incorporating on-policy samples via the detour over β. This leads to relatively stable TD targets computed from Q_β while an upbuilding under- or overestimation is prevented as β picks up the trend. On the other hand, if β would change too slowly the upbuilding of the bias might not be stopped.
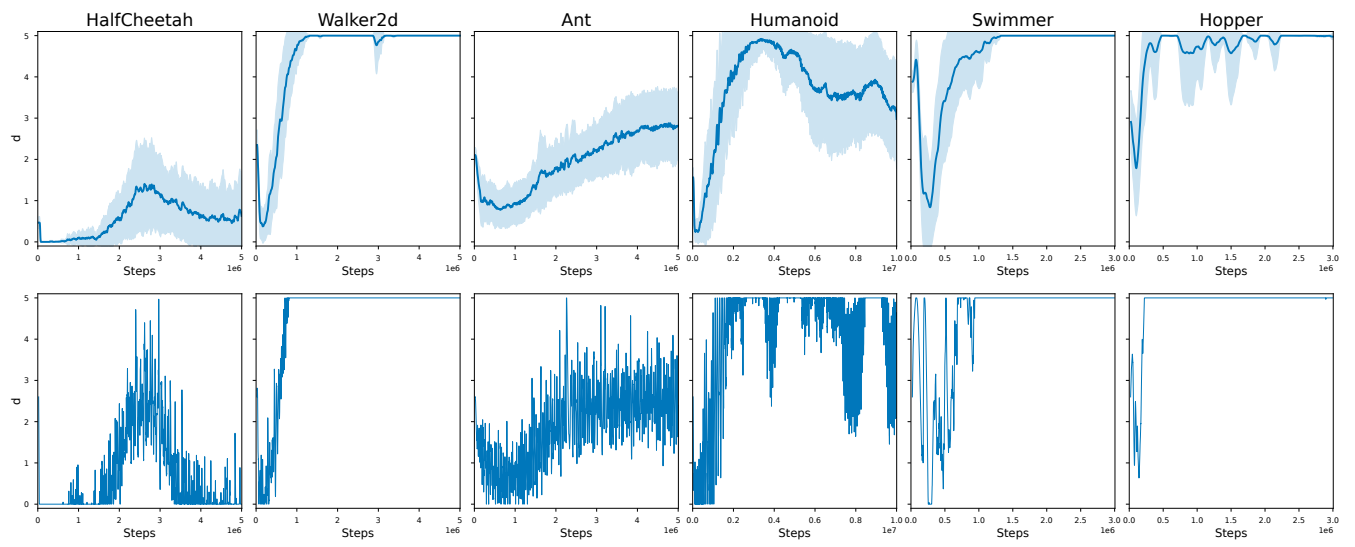
Fig. 6. Development of the number of dropped targets per network $d = d_{max} - \beta$ in ACC over time for different environments. The top row shows the mean (thick line) and standard deviation (shaded area) over the 10 trials where for readability a uniform filter of size 15 is used. The bottom row shows the unfiltered development for one of the seeds.