SAMPLE EFFICIENT DEEP REINFORCEMENT LEARNING VIA UNCERTAINTY ESTIMATION

Anonymous authors
Paper under double-blind review

ABSTRACT

In model-free deep reinforcement learning (RL) algorithms, using noisy value estimates to supervise policy evaluation and optimization is detrimental to the sample efficiency. As this noise is heteroscedastic, its effects can be mitigated using uncertainty-based weights in the optimization process. Previous methods rely on sampled ensembles, which do not capture all aspects of uncertainty. We provide a systematic analysis of the sources of uncertainty in the noisy supervision that occurs in RL, and introduce inverse-variance RL, a Bayesian framework which combines probabilistic ensembles and Batch Inverse Variance weighting. We propose a method whereby two complementary uncertainty estimation methods account for both the Q-value and the environment stochasticity to better mitigate the negative impacts of noisy supervision. Our results show significant improvement in terms of sample efficiency on discrete and continuous control tasks.

1 INTRODUCTION

Deep reinforcement learning (DRL) methods have proven to be powerful at solving sequential decision-making tasks across domains (Silver et al., 2016; OpenAI et al., 2019). Combining the flexibility of the reinforcement learning framework with the representational power of deep neural networks enables policy optimization in complex and high-dimensional environments with unknown dynamics models to maximize the expected cumulative reward (Sutton & Barto, 2018).

An important limitation of DRL methods is their sample inefficiency: an enormous amount of data is necessary and makes training expensive.

Among the current state-of-the-art approaches to improve learning efficiency, a promising direction is to exploit the prevalence of uncertainty in the underlying DRL algorithm. By adopting a Bayesian framework, we can consider the sampled quantities in DRL as random variables and leverage information about their distributions to improve the learning process (Osband et al., 2018). In this paper, we consider the particular problem of unreliable supervision in the temporal difference update and the policy optimization process. In DRL, value predictions are used to supervise the training: in temporal difference-based algorithms, they are included in bootstrapped target values which are used as labels; in actor-critic frameworks, the policy is trained to optimize them. That these value predictions are noisy slows the learning and brings instability (Kumar et al., 2019; 2020). The amount of noise in the supervision depends on the uncertainty of the value prediction, which evolves during the training process and depends on the state (and action) evaluated. It is therefore *heteroscedastic*.

While there is an extensive body of literature focused on using the uncertainty of the value prediction to guide the exploration/exploitation trade-off (Dearden et al., 1998; Strens, 2001; Osband et al., 2016; Pathak et al., 2017; Chen et al., 2017; Osband et al., 2018; Fortunato et al., 2019; Osband et al., 2019; Flennerhag et al., 2020; Clements et al., 2020; Jain et al., 2021; Aravindan & Lee, 2021), there are very few works focused in leveraging it to mitigate the impact of unreliable supervision.

Distributional RL (Bellemare et al., 2017) considers the value function as a distribution to be learned as such. It is orthogonal to our proposition: we consider the uncertainty of the labels used to learn a scalar value function. In the offline RL setting, where the dataset is limited, uncertainty-weighted actor-critic (UWAC) (Wu et al., 2021) uses inverse-variance weighting to discard out-of-distribution state-action pairs using Monte Carlo dropout (Gal & Ghahramani, 2016) for uncertainty estimation.

Closer to our work, Lee et al. (2020) propose SUNRISE, in which each sample of the Bellman backup in the TD update step is weighted to lower the importance of the targets which have a high standard deviation. The weights $w(s', a')$ are computed based on a sigmoïd of the negative standard deviation $\hat{Q}_{\text{std}}(s', a')$ scaled by a temperature hyperparameter T, and then offset such that they are between 0.5 and 1: $w(s, a) = \sigma(-\hat{Q}_{\text{std}}(s', a') * T) + 0.5$. The uncertainty of the target is estimated using sampled ensembles. While SUNRISE proposes other contributions such as an exploration bonus, the heuristic weighting scheme and the limitations of sampled ensembles in capturing the predictive uncertainty leave space for improvement in the mitigation of the effects of unreliable supervision.

We propose inverse-variance reinforcement learning (IV-RL). IV-RL also uses weights to reduce the importance of uncertain targets in training. It does so by addressing the problem from two viewpoints. First, we use probabilistic networks (Kendall & Gal, 2017), whose loss function for regression is the negative log-likelihood instead of the L2 distance. For a given state-action pair (s, a), the network learns the target's noise, due for example to the stochasticity of the environment or the update of the policy. It then naturally down-weights the highly noisy samples in the training process. Second, we use probabilistic ensembles (Lakshminarayanan et al., 2017) to estimate the uncertainty of the target due to the prediction of $Q(s', a')$ during the temporal-difference update. We merge the predicted variances of several probabilistic networks through a mixture of Gaussians, which has been shown to be a very reliable method to capture predictive uncertainty (Ovadia et al., 2019). We then use Batch Inverse-Variance (BIV) (Mai et al., 2021), which has been shown to significantly improve the performance of supervised learning with neural networks in the case of heteroscedastic regression. BIV is normalized, which makes it ideal to cope with different and time-varying scales of variance. We show analytically that these two different variance predictions for the target are complementary and their combination leads to consistent and significant improvements in the sample efficiency and overall performance of the learning process.

In summary, our contribution is threefold:

1. We present a systematic analysis of the sources of uncertainty in the supervision of model-free DRL algorithms. We show that the variance of the supervision noise can be estimated with two complementary methods: negative log-likelihood and probabilistic ensembles.
2. We introduce IV-RL, a framework that accounts for the uncertainty of the supervisory signal by weighting the samples in a mini-batch during the agent's training. IV-RL uses BIV, a weighting scheme which is robust to poorly calibrated variance estimation.[1]
3. Our experiments show that IV-RL can lead to significant improvements in sample efficiency when applied to Deep Q-Networks (DQN) (Mnih et al., 2013) and Soft-Actor Critic (SAC) (Haarnoja et al., 2018).

In section 2, we introduce BIV as a weighting scheme for heteroscedastic regression, and probabilistic ensembles as an uncertainty estimation method. We analyse the sources of uncertainty in the target in section 3, where we also introduce our IV-RL framework. We finally present our experimental results in section 4.

2 Background and preliminaries

2.1 Batch Inverse-Variance Weighting

In supervised learning with deep neural networks, it is assumed that the training dataset consists of inputs $\mathbf{x}_k$ and labels y_k. However, depending on the label generation process, the label may be noisy. In regression, we can model the noise as a normal distribution around the true label: $\tilde{y}_k = y_k + \delta_k$ with $\delta_k \sim \mathcal{N}(0, \sigma_k^2)$. In some cases, the label generation process leads to different variances for the label noises. When these variances can be estimated, each sample is a triplet $\left(\mathbf{x}_k, \tilde{y}_k, \sigma_k^2\right)$.

Batch Inverse-Variance (BIV) weighting (Mai et al., 2021) leverages the additional information σ_k^2, which is assumed to be provided, to learn faster and obtain better performance in the case of heteroscedastic noise on the labels. It optimizes the neural network parameters θ using the following

[1] Our code is available as supplementary material. We plan to open-source our codebase in case of acceptance/publication of the paper.

loss function for a mini-batch D_i of size K [2]:

$$\mathcal{L}_{\text{BIV}}(D_i, \theta) = \left(\sum_{k=0}^{K} \frac{1}{\sigma_k^2 + \xi} \right)^{-1} \sum_{k=0}^{K} \frac{\mathcal{L}\left(f(\mathbf{x}_k, \theta), \tilde{y}_k\right)}{\sigma_k^2 + \xi} \tag{1}$$

This is a normalized weighted sum with weights $w_k = 1/(\sigma_k^2 + \xi)$. Normalizing in the mini-batch enables control of the effective learning rate, especially in cases where the training data changes over time, such as in DRL. By focusing on the relative scale of the variances instead of their absolute value, it also provides robustness to poor scale-calibration of the variance estimates.

As explained in Mai et al. (2021), ξ is a hyperparameter which is important for the stability of the optimization process. A higher ξ limits the highest weights, thus preventing very small variance samples from dominating the loss function for a mini-batch. However, by controlling the discrimination between the samples, ξ is also key when the variance estimation is not completely trusted. It provides control of the effective mini-batch size EBS, according to:

$$EBS = \frac{\left(\sum_k^K w_k\right)^2}{\sum_k^K w_k^2} = \frac{\left(\sum_k^K \frac{1}{(\sigma_k^2 + \xi)}\right)^2}{\sum_k^K \frac{1}{(\sigma_k^2 + \xi)^2}} \tag{2}$$

For example, imagine a mini-batch where most samples have very high variances, and only one has a very low variance. If $\xi = 0$, this one low-variance sample is effectively the only one to count in the mini-batch, and EBS tends towards 1. Increasing ξ would give more relative importance to the other samples, thus increasing EBS. With a very high ξ compared to the variances, all weights are equal, and EBS tends towards K; in this case, the BIV loss tends towards $L2$.

Tuning the ξ parameter The simplest way to set ξ is to choose a constant value as an additional hyperparameter. However, the best value is difficult to evaluate a priori and can change when the profile of variances changes during a task, as is the case in DRL.

It is instead possible to numerically compute the value of ξ which ensures a minimal EBS for each mini-batch. This method allows ξ to automatically adapt to the different scales of variance, while ensuring a minimal amount of information from the dataset to be accounted for by the algorithm. The minimal EBS is also a hyper-parameter, but it is easier to set and to transfer among environments, as it is simply a fraction of the original batch size. As such, it can be set as a batch size ratio.

2.2 Estimating the uncertainty of a neural network prediction

The predictive uncertainty of a neural network can be considered as the combination of aleatoric and epistemic uncertainties Kendall & Gal (2017). Aleatoric uncertainty is irreducible, and characterizes the non-deterministic relationship between the input and the desired output. Epistemic uncertainty is instead related to the trained model: it depends on the information available in the training data, the model's capacity to retain it, and the learning algorithm (Hüllermeier & Waegeman, 2021). There is currently no principled way to quantify the amount of task-related information present in the input, the training data, or the model. The state of the art for predictive uncertainty estimation instead relies on different sorts of proxies. These sometimes capture other elements, such as the noise of the labels, which we can use to our advantage. We focus here on the relevant methods to our work.

2.2.1 Sampled ensembles

Several networks independently train an ensemble of size N that can be interpreted as a distribution over predictions. The expected behavior is that different networks will only make similar predictions if they were sufficiently trained for a given input. The sampled variance of the networks' outputs is thus interpreted as the epistemic uncertainty. It is possible to include a random Bernoulli mask of probability p to each training sample, to ensure that each network undergoes different training. This method, used by Clements et al. (2020) and (Lee et al., 2020), has the same principle as single network Monte-Carlo dropout (Gal & Ghahramani, 2016). As the variance is sampled, the standard deviation is usually in the same scale as the prediction.

[2] We replaced ϵ in Mai et al. (2021) with ξ to avoid confusion with the reinforcement learning conventions.

When used at the very beginning of the training process, sampled ensembles present one particular challenge: as the networks are initialized, they all predict small values. The initial variance, instead of capturing the lack of knowledge, is then underestimated. To address this problem, Randomized Prior Functions (RPFs) enforce a prior in the variance by pairing each network with a fixed, untrained network which adds its predictions to the output (Osband et al., 2019). RPFs ensure a high variance at regions of the input space which are not well explored, and a lower variance when the trained networks have learned to compensate for their respective prior and converge to the same output. The scale of the prior is a hyper-parameter.[3]

2.2.2 Probabilistic networks

With probabilistic networks, the uncertainty is predicted using loss attenuation (Nix & Weigend, 1994; Kendall & Gal, 2017). A network outputs two values in its final layer given an input x: the predicted mean $\mu(x)$ and variance $\sigma^2(x)$. The network parameters θ are optimized by minimizing the negative log-likelihood of a heteroscedastic Gaussian distribution:

$$\mathcal{L}_{LA}(x, \theta_Q) = \frac{\left(\mu_\theta(x) - y(x)\right)^2}{\sigma^2_\theta(x)} + \log \sigma^2_\theta(x) \tag{3}$$

Probabilistic networks naturally down-weight the labels with high variance in the optimization process. The variance prediction is trained from the error between $\mu_\theta(x)$ and the label $y(x)$. Therefore, if there is noise on the labels or if the regression task changes over time, this will also be captured by loss attenuation. As the variance is predicted by a neural network, it may not be well calibrated, and may be over estimated (Kuleshov et al., 2018; Levi et al., 2020; Bhatt et al., 2021). This can (1) give wrong variance estimates but also (2) affect the learning process by ignoring a sample if the variance estimate is too high.

2.2.3 Probabilistic ensembles

Lakshminarayanan et al. (2017) combine sampled ensembles and probabilistic networks for probabilistic ensembles. The predictive variance is given by a Gaussian mixture over the variance predictions of each network in the ensemble. This method, when trained, is able to capture uncertainty more reliably than others (Ovadia et al., 2019), with N = 5 networks in the ensemble being sufficient. Probabilistic ensembles also suffer from underestimated early epistemic variance estimation. However, they seem empirically less prone to calibration issues (1) in the final variance estimation, because the mixture of Gaussians dampens single very high variances, and (2) in the learning process, because even if one network does not learn correctly, the others will.

2.3 Uncertainty and exploration in DRL

While our work focuses on using uncertainty estimates to mitigate the impact of unreliable supervision, we can take advantage of the structure in place to better drive the exploration/exploitation trade-off. In particular, we used BootstrapDQN (Osband et al., 2016) for exploration. In this method a single network is sampled from an ensemble at the beginning of each episode to select the action. This method is improved with the previously described RPFs (Osband et al., 2018). In continuous settings, we instead followed Lee et al. (2020) and added an Upper Confidence Bound (UCB) exploration bonus based on uncertainty prediction. As the variance in UCB is added to Q-values, it must be calibrated: we evaluate it with sampled ensembles.

3 Inverse-Variance Reinforcement Learning

3.1 Target uncertainty in Reinforcement Learning

Many model-free DRL algorithms use temporal difference updates. In methods such as DQN (Mnih et al., 2013), PPO (Schulman et al., 2017) and SAC (Haarnoja et al., 2018), a neural network is

[3]Using RPFs in the context of DRL can destabilize the exploration due to a significant initial bias added to the value of each state-action. It must be compensated for with uncertainty-aware exploration, such as Bootstrap (Osband et al., 2016), where it leads to a significant improvement in performance for exploration-based tasks.

trained to predict the Q-value of a given state-action pair $Q^\pi(s,a)$ [4] by minimizing the error between the target $T(s,a)$ and its prediction $\hat{Q}(s,a)$. $T(s,a)$ is computed according to Bellman's equation:

$$T(s,a) = r + \gamma \bar{Q}(s',a') \tag{4}$$

s' and r are sampled from the environment given (s,a), and a' is sampled from the current policy given s'. $\bar{Q}(s',a')$ is predicted by a copy of the Q-network (called the target network) which is updated less frequently to ensure training stability. The neural network's parameters θ are optimized using stochastic gradient descent to minimize the following loss function:

$$\mathcal{L}_\theta = \left|\left| T(s,a) - \hat{Q}_\theta(s,a) \right|\right|^2 \tag{5}$$

3.1.1 THE TARGET AS A RANDOM VARIABLE

The target $T(s,a)$ is a noisy approximation of $Q^\pi(s,a)$ that is distributed according to its distribution $p_T(T|s,a)$. The generative model used to produce samples of $T(s,a)$ is shown in Figure 1, and has the following components:

1. if the reward r is stochastic, it is sampled from $p_R(r|s,a)$ [5];
2. if the environment dynamics are stochastic, the next state s' is sampled from $p_{S'}(s'|s,a)$;
3. if the policy is stochastic a' is sampled from the policy $\pi(a'|s')$;
4. $\bar{Q}$ is a prediction from a probabilistic neural network $p_{\bar{Q}}(\bar{Q}|s',a')$;
5. T is deterministically generated from r and $\bar{Q}$ using equation (4).

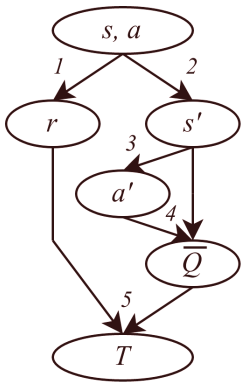

Figure 1: Bayesian network representing the target sampling process

As the variance of the noise of $T(s,a)$ is not constant, the training of the Q-network using $\mathcal{L}_\theta$ as in equation (5) is regression on heteroscedastic noisy labels.

3.1.2 VARIANCE OF THE TARGET

As seen in section 2.1, BIV can be used to reduce the impact of heteroscedastic noisy labels in regression, provided estimates of the label variances. We thus aim to evaluate $\sigma_T^2(T|s,a)$ based on the sampling process described in section 3.1.1. As r and Q-value estimation are independent given s and a, we have:

$$\sigma_T^2(T|s,a) = \sigma_R^2(r|s,a) + \gamma^2 \sigma_{S'A'\bar{Q}}^2\left(\bar{Q}|s,a\right) \tag{6}$$

where $p_{S'A'\bar{Q}}$ is the compound probability distribution based on components 2-4 in Figure 1:

$$p_{S'A'\bar{Q}}(\bar{Q}|s,a) = \iint p_{\bar{Q}}(\bar{Q}|s',a') p_{S'A'}(s',a'|s,a) da' ds' \tag{7}$$

where $p_{S'A'}(s',a'|s,a) = p_{A'}(a'|s')p_{S'}(s'|s,a)$. Using the law of total variance, the variance of $\bar{Q}$ is given by:

$$\sigma_{S'A'\bar{Q}}^2\left(\bar{Q}|s,a\right) = \mathbb{E}_{S'A'}\left[\sigma_{\bar{Q}}^2\left(\bar{Q}|s',a'\right)\right] + \sigma_{S'A'}^2\left(\mathbb{E}_{\bar{Q}}\left[\bar{Q}|s',a'\right]\right) \tag{8}$$

Plugging (8) into (6) gives:

$$\sigma_T^2(T|s,a) = \gamma^2 \underbrace{\left(\mathbb{E}_{S'A'}\left[\sigma_{\bar{Q}}^2\left(\bar{Q}|s',a'\right)\right]\right)}_{\text{Predictive variance of Q-network}} + \underbrace{\gamma^2\left(\sigma_{S'A'}^2\left(\mathbb{E}_{\bar{Q}}\left[\bar{Q}|s',a'\right]\right)\right) + \sigma_R^2(r|s,a)}_{\text{Policy and environment induced variance}} \tag{9}$$

[4] In this paper, we focus on Q-values, although the principle can also apply to state values.

[5] Stochastic reward can happen for example if the reward is corrupted or estimated (Romoff et al., 2018; Campbell et al., 2015)

We can identify two distinct components in equation 9 that contribute to the overall variance of the target. The first is the (expectation of the) variance that is due to the uncertainty in the neural network prediction of the value function, $\mathbb{E}_{S'A'}\left[\sigma^2_{\bar{Q}}\left(\bar{Q}|s',a'\right)\right]$. The second is the uncertainty due to the stochasticity of the environment and of the policy, $\sigma^2_R(r|s,a) + \gamma^2\sigma^2_{S'A'}\left(\mathbb{E}_{\bar{Q}}\left[\bar{Q}|s',a'\right]\right)$.

Uncertainty in neural network prediction For a given policy π and a given s', a', the agent may not have seen enough samples to have an accurate approximation of $Q^\pi(s',a')$. This corresponds to an epistemic source of uncertainty that should be captured by sampled ensembles. However, as π is updated, the regression target, $Q^\pi(s',a')$, is also changing. This can be interpreted as variability in the underlying process which will be captured by probabilistic networks. We can thus combine both sampling-based and probabilistic network methods and evaluate $\sigma^2_{\bar{Q}}\left(\bar{Q}(s',a')\right)$ with probabilistic ensembles.

We assume that the estimate of $\sigma^2_{\bar{Q}}(\bar{Q}|s',a')$ given a sampled (s',a') is unbiased and can therefore use it to directly approximate the expectation $\mathbb{E}_{S'A'}\left[\sigma^2_{\bar{Q}}\left(\bar{Q}|s',a'\right)\right]$. These values are used in the BIV loss $\mathcal{L}_{BIV}$ (equation 1 across a mini-batch sampled from the replay buffer. In this case, ξ is used to control the trust in the variance estimation, as explained in section 2.1.

Stochastic environment and policy The other potential source of variance in the target is the result of the stochasticity of the environment encapsulated by $p_R\left(r|s,a\right)$ and $p_{S'}(s'|s,a)$ and of the policy represented by $\pi(a'|s')$. Note that in model-free RL, we have no explicit representation of $p_R\left(r|s,a\right)$ or $p_{S'}(s'|s,a)$, which are necessary to estimate this source of uncertainty.

$Q^\pi(s,a)$ is defined as the expected value of the return. As a result, even in the case where where $\bar{Q}(s',a')$ = $Q^\pi(s,a)$ in (4), there is still noise over the value of the *target* that is being used as the label due to the stochasticity of the environment and policy that generate r, a' and s'.

If we assume that this noise is zero mean and normally distributed, this underlying stochasticity of the generating process is well-captured by a probabilistic network with a loss attenuation using the negative log-likelihood formulation described in equation 3.

3.1.3 Loss function for IV-RL

Finally, based on the motivation above, we propose our IV-RL loss:

$$\mathcal{L}_{\text{IVRL}} = \mathcal{L}_{\text{BIV}} + \lambda\mathcal{L}_{\text{LA}} \tag{10}$$

This loss is a simple linear combination of the LA and the BIV losses with a constant hyper-parameter λ. The result is that high-variance samples generated from the Q-value estimation will be down-weighted in the BIV loss, and while high-variance samples due to the stochasticity of the underlying environment and policy will be down-weighted in the LA loss. In the remainder of the paper we show how this loss can be applied to different architectures and algorithms, and demonstrate that in many cases it significantly improves sample efficiency.

3.2 Q-value uncertainty and actor-critic structures

In section 3.1, we discussed how the target's uncertainty can be quantified and how the Bellman update can then interpreted as a heteroscedastic regression problem. This is applicable in most model-free DRL algorithms, whether they are based on Q-learning or policy optimization. In the special case of actor-critic algorithms, the state-values or Q-values predicted by the critic network are also used to train the policy π_ϕ's parameters by gradient ascent optimization. An estimate of their variance can also be used to improve the learning process.

The objective is to maximize the expected Q-value:

$$\text{E}_{s\sim D}[Q(s,\pi_\phi(s))] \tag{11}$$

where D is the state distribution over the probable agent trajectories. The expectation is computed by sampling a mini-batch B from a replay buffer. The Q-value is approximated by the critic as $\hat{Q}(s,a)$. The actor's parameters ϕ are then trained to maximize the unweighted average:

$$1/|B|\sum_{i\in B}\hat{Q}(s_i,\pi_\phi(s_i)) \tag{12}$$

If we instead consider the critic's estimation $\hat{Q}(s_i, \pi_\phi(s_i))$ as a random variable sampled from $p_{\hat{Q}}(\hat{Q}|s, a)$, with variance $\sigma^2_{\hat{Q}}(s_i, \pi_\phi(s_i))$, we can instead infer the expected value in equation (11) using Bayesian estimation (Murphy, 2012):

$$\left(\sum_{i \in B} 1/\sigma^2_{\hat{Q}}\left(s_i, \pi_\phi(s_i)\right)\right)^{-1} \sum_{i \in B} \frac{1}{\sigma^2_{\hat{Q}}\left(s_i, \pi_\phi(s_i)\right)} \hat{Q}(s_i, \pi_\phi(s_i)) \tag{13}$$

The normalized inverse-variance weights are a direct fit with the BIV loss in equation 1: we therefore also use BIV to train the actor. As the target and the Q-networks have the same structure, $\sigma^2_{\hat{Q}}(s_i, \pi_\phi(s_i))$ can be estimated using the same probabilistic ensembles used in section 3.1.2.

3.3 Algorithms

We have adapted IV-RL to DQN and SAC to produce IV-DQN and IV-SAC. The implementation details and specific algorithms can be found in appendix A.

4 Results

We have tested IV-RL across a variety of different environments, using IV-DQN for discrete tasks and IV-SAC for continuous tasks. To determine the effects of IV-RL, we propose as a baseline not only the original DQN and SAC, but also the non-IV version of the improvements used in IV-RL and which are not our contribution. We also include SUNRISE (Lee et al., 2020) as a baseline.

All ensemble-based methods use an ensemble size of $N = 5$. Unless specified otherwise, each result is the average of runs over 25 seeds: 5 for the environment $\times$ 5 for the network initialization. The hyperparameters are the result of a thorough fine-tuning process explained in appendix B. We discuss the computation time considerations in appendix C.

4.1 IV-DQN

We tested IV-DQN on different discrete control environments, including LunarLander and MountainCar from OpenAI Gym (Brockman et al., 2016), as well as Cartpole-Noise from BSuite[6] (Osband et al., 2020) where noise is applied on the reward. BootstrapDQN is an ensemble-based method which includes the bootstrap exploration from (Osband et al., 2016) and Randomized Prior Functions (Osband et al., 2018). The reported SUNRISE results are obtained by applying the SUNRISE weights to BootstrapDQN. The learning curves are shown in figure 2, where the results are averaged on a 100 episode window.

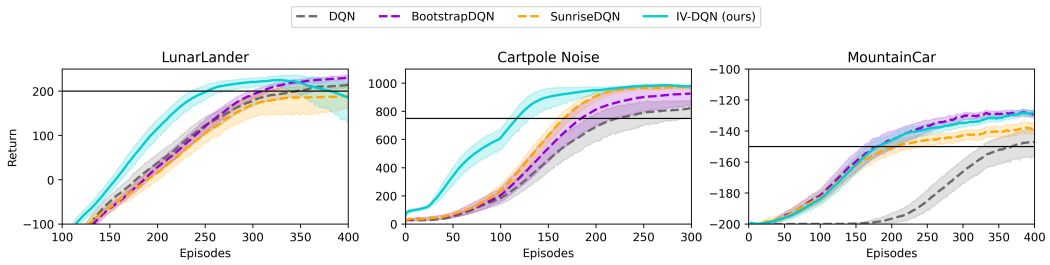

Figure 2: Using IV-RL shows improved performance over non-IV methods in Cartpole Noise and LunarLander, and does not impact MountainCar.

IV-DQN outperforms BootstrapDQN and SUNRISEDQN on LunarLander and Cartpole-Noise, which are control-based environments. However, in Mountain-Car, the exploration strategy is of crucial importance since the reward is sparse. As a result, we see that all weighting schemes perform comparably, as long as RPF and Bootstrap are present. We note that probabilistic ensembles

[6]The results from BSuite are averaged over 20 runs with different environment and network seeds

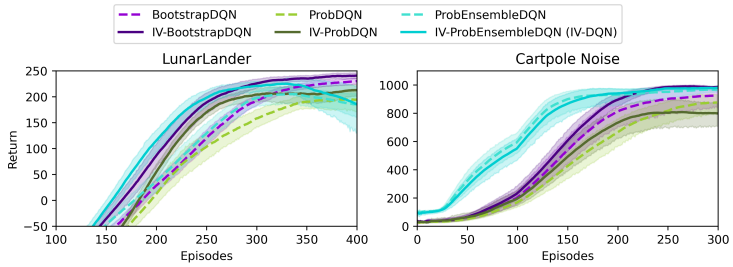

Figure 3: Ablation study: depending on the environment, the BIV or the probabilistic ensemble component is the most important factor of improvement.

on LunarLander lead to a drop in the return in the end, both in the IV and non-IV version: this may be due to catastrophic forgetting.

Table 1 shows the median number of episodes necessary to reach a given score for which the environment is considered as solved. IV-DQN shows a clear improvement over baselines in sample efficiency in both control-based environments, but fails to improve the exploration in Mountain Car.

Table 1: 25th - 50th - 75th percentiles of the number of episodes necessary for the return averaged with a 100-episode window to reach the solved score on different environments. IV-DQN shows significant improvements in sample efficiency when the environment is not exploration-based.

	LunarLander (200)	Cartpole-Noise (750)	MountainCar (-150)
DQN	296 - 316 - 349	171 - 193 - max	304 - 333 - 403
BootstrapDQN	287 - 305 - 317	160 - 174 - 196	**134 - 149 - 206**
SunriseDQN	291 - 309 - 368	155 - 165 - 175	152 - 197 - 257
IV-DQN (ours)	**226 - 237 - 263**	**105 - 112 - 117**	142 - 163 - 200

In Figure 3 we perform an ablation of BIV and probabilistic ensembles for LunarLander and CartpoleNoise. We see different patterns: in the LunarLander, the BIV component is clearly responsible for the improved performance. In Cartpole-Noise, the probabilistic ensemble yields the more significant portion of the improvement. One hypothesis to explain this effect is that LunarLander has more states to explore than Cartpole-Noise, and thus is more prone to generate epistemic uncertainty, while Cartpole-Noise has more stochasticity in the environment due to the noise in the reward. In both cases, the use of BIV or probabilistic ensembles improves the performance, and their combination leads to the best sample efficiency.

We also show the results of using one single probabilistic network (ProbDQN and IV-ProbDQN) as opposed to probabilistic ensembles. In both cases, it is sub-optimal. This is likely due to the unstable nature of the single uncertainty estimate, which affects the learning process. More details are shown in appendix D and E. The use of probabilistic ensembles stabilizes the variance estimation, as explained in section 2.2.3.

4.2 IV-SAC

IV-SAC was applied to different continuous control environments from OpenAI Gym (Brockman et al., 2016) as implemented by MBBL (Wang et al., 2019). EnsembleSAC is a baseline using ensembles with an UCB exploration bonus (Lee et al., 2020). Table 2 shows the average return after 100k and 200k steps on 25 seeds. We note that the addition of an ensemble instead of a single network, along with the uncertainty-based exploration bonus, already allows the performance to increase compared to SAC. Except in Ant, the SUNRISE weights do not seem to lead to consistently better results. In comparison, IV-SAC leads to significant improvements in performance.

This improvement in performance after a fixed amount of training steps can be explained by an improved sample efficiency. This can be seen in figure 4. Even when IV-SAC's return is not significantly better than the baselines at 200k steps, such as in Ant or Hopper, it clearly is learning faster, which is also reflected by the scores at 100k steps.

Similarly to IV-DQN, we can separate the contribution from BIV and loss attenuation, as shown in the two first plots of figure 5. While both BIV and probabilistic ensembles alone have a significant

Table 2: Performance at 100K and 200K timesteps (100 and 200 episodes) for several robotics environments in OpenAI Gym. The results show the mean and standard error over 25 runs.

		Walker	HalfCheetah	Hopper	Ant
100k steps	SAC	-392 $\pm$ 187	3211 $\pm$ 136	322 $\pm$ 177	724 $\pm$ 29
	SampledSAC	-389 $\pm$ 209	3938 $\pm$ 112	1597 $\pm$ 152	852 $\pm$ 27
	SunriseSAC	46 $\pm$ 213	3879 $\pm$ 204	1618 $\pm$ 195	834 $\pm$ 130
	IV-SAC (ours)	**857** $\pm$ **231**	**4260** $\pm$ **118**	**2237** $\pm$ **116**	**948** $\pm$ **46**
200k steps	SAC	371 $\pm$ 189	3978 $\pm$ 148	1635 $\pm$ 162	985 $\pm$ 82
	SampledSAC	1337 $\pm$ 278	4757 $\pm$ 92	2652 $\pm$ 91	981 $\pm$ 65
	SunriseSAC	1423 $\pm$ 295	4785 $\pm$ 228	2572 $\pm$ 119	**1462** $\pm$ **154**
	IV-SAC (ours)	**3009** $\pm$ **193**	**5451** $\pm$ **151**	**2889** $\pm$ **50**	1281 $\pm$ 101

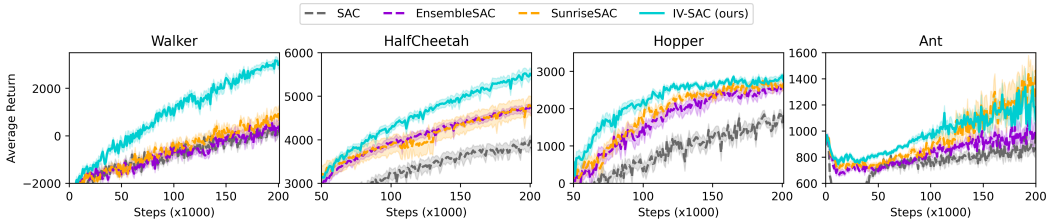

Figure 4: IV-SAC learns faster and leads to significantly better results than the baselines.

impact in Walker, it's only their combination that brings the most improvement in HalfCheetah. Similarly to the discrete control case, a simple probabilistic network with ProbSAC and IV-ProbSAC leads to a slight improvement over SAC, which is discussed in appendix E. We also show in figure 5 that the BIV weights provided with probabilistic ensemble variance estimations than the SUNRISE weights, or even the inverse variance weights of UWAC (Wu et al., 2021). The normalization of BIV allows it to better cope with uncalibrated variance predictions, as shown in appendix D.

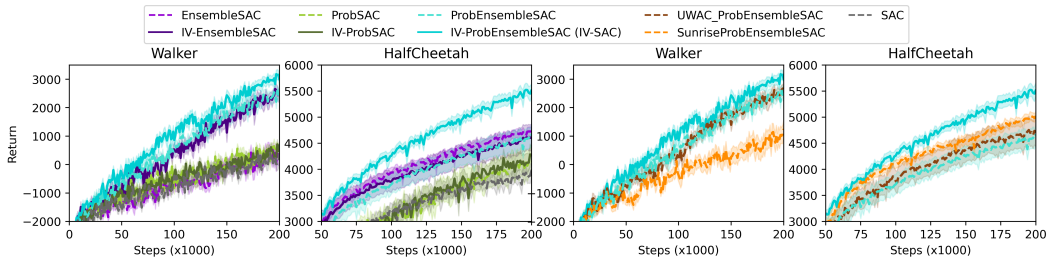

Figure 5: **Ablation Study:** (first two figures) Impact of using different uncertainty estimation methods (last two figures) Comparing different weighting schemes with probabilistic ensembles.

5 Conclusion

We present Inverse Variance Reinforcement Learning (IV-RL), a framework for model-free deep reinforcement learning which leverages uncertainty estimation to enhance sample efficiency and performance. Motivated by a thorough analysis of the sources of noise that contribute to errors in the target, we use a combination of Batch Inverse Variance (BIV) weighting and probabilistic ensembles to estimate the variance of the target and down weight the uncertain samples in two complementary ways. Our results show that these two components are beneficial, and that their combination significantly improves the state of the art in terms of learning efficiency.

We have adapted our method to both discrete (DQN) and continuous (SAC) reinforcement learning problems. IV-RL can be easily adapted to other model-free algorithms such as PPO or TRPO. Future research could be made to apply the ideas proposed in IV-RL to model-based RL or algorithms solving similar tasks such as imitation learning, active learning, curriculum and continual learning, or even in the sim2real process.

ETHICS STATEMENT

In this work, we present a method to enhance the sample efficiency of deep reinforcement learning algorithms. As such, it is agnostic to the applications, and per se does not raise any particular ethical issue. We however strongly encourage the user of our algorithm to ensure they have carefully thought about the ethical issues related to the particular field of application, such as medicine, robotics, communication, finance, etc.

As environmental sustainability can also be considered an ethical issue (Universite de Montreal, 2018), we publish the carbon footprint of our work.

Experiments were conducted using a private infrastructure, which has a carbon efficiency of 0.028 kgCO_2eq/kWh. A cumulative of 12367 days, or 296808 hours, of computation was mainly performed on hardware of type RTX 8000 (TDP of 260W). We assume full power usage of the GPUs, although this was not always the case.

Total emissions are estimated to be 2160.76 kgCO_2eq of which 0 percents were directly offset. This is equivalent to 8730 km driven by an average car, or 1.08 metric ton of burned coal.

Estimations were conducted using the MachineLearning Impact calculator presented in Lacoste et al. (2019).

REPRODUCIBILITY STATEMENT

To allow reproducibility of our results, we submitted a link to download an anonymous version of source code we used to produce them. This also includes a configuration file with the hyperparameters used to produce each result presented in this paper.

The environments and implementaion we used (OpenAI Gym, BSuite, MBBL) are all publicly accessible, although a Mujoco license is needed to run some of them.

REFERENCES

Siddharth Aravindan and Wee Sun Lee. State-aware variational thompson sampling for deep q-networks. *arXiv:2102.03719 [cs]*, Feb 2021. URL http://arxiv.org/abs/2102.03719. arXiv: 2102.03719.

Marc G. Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. *arXiv:1707.06887 [cs, stat]*, Jul 2017. URL http://arxiv.org/abs/1707.06887. arXiv: 1707.06887.

Dhaivat Bhatt, Kaustubh Mani, Dishank Bansal, Krishna Murthy, Hanju Lee, and Liam Paull. f-cal: Calibrated aleatoric uncertainty estimation from neural networks for robot perception. *arXiv:2109.13913 [cs]*, Sep 2021. URL http://arxiv.org/abs/2109.13913. arXiv: 2109.13913.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv:1606.01540 [cs]*, Jun 2016. URL http://arxiv.org/abs/1606.01540. arXiv: 1606.01540.

J.S. (Jeffrey S.) Campbell, S.N. (Sidney N.) Givigi, and Howard M. Schwartz. Handling stochastic reward delays in machine reinforcement learning, January 2015.

Richard Y. Chen, Szymon Sidor, Pieter Abbeel, and John Schulman. Ucb exploration via q-ensembles. *arXiv:1706.01502 [cs, stat]*, Nov 2017. URL http://arxiv.org/abs/1706.01502. arXiv: 1706.01502.

William R. Clements, Bastien Van Delft, Benoît-Marie Robaglia, Reda Bahi Slaoui, and Sébastien Toth. Estimating risk and uncertainty in deep reinforcement learning. *arXiv:1905.09638 [cs, stat]*, Sep 2020. URL http://arxiv.org/abs/1905.09638. arXiv: 1905.09638.

Richard Dearden, Nir Friedman, and Stuart J. Russel. Bayesian q-learning. 1998. URL http://ai.stanford.edu/~nir/Papers/DFR1.pdf.

Sebastian Flennerhag, Jane X. Wang, Pablo Sprechmann, Francesco Visin, Alexandre Galashov, Steven Kapturowski, Diana L. Borsa, Nicolas Heess, Andre Barreto, and Razvan Pascanu. Temporal difference uncertainties as a signal for exploration. *arXiv:2010.02255 [cs, stat]*, Oct 2020. URL http://arxiv.org/abs/2010.02255. arXiv: 2010.02255.

Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, and et al. Noisy networks for exploration. *arXiv:1706.10295 [cs, stat]*, Jul 2019. URL http://arxiv.org/abs/1706.10295. arXiv: 1706.10295.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR. URL http://proceedings.mlr.press/v48/gal16.html.

Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. 2018. ISSN 0264-410X (Print)r0264-410X (Linking). URL http://arxiv.org/abs/1801.01290.

Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. *arXiv:1709.06560 [cs, stat]*, Jan 2019. URL http://arxiv.org/abs/1709.06560. arXiv: 1709.06560.

Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. *Machine Learning*, 110(3):457–506, Mar 2021. ISSN 0885-6125, 1573-0565. doi: 10.1007/s10994-021-05946-3.

Moksh Jain, Salem Lahlou, Hadi Nekoei, Victor Butoi, Paul Bertin, Jarrid Rector-Brooks, Maksym Korablyov, and Yoshua Bengio. Deup: Direct epistemic uncertainty prediction. *arXiv:2102.08501 [cs, stat]*, Feb 2021. URL http://arxiv.org/abs/2102.08501. arXiv: 2102.08501.

Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *arXiv:1703.04977 [cs]*, Oct 2017. URL http://arxiv.org/abs/1703.04977. arXiv: 1703.04977.

Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. *arXiv:1807.00263 [cs, stat]*, Jun 2018. URL http://arxiv.org/abs/1807.00263. arXiv: 1807.00263.

Aviral Kumar, Justin Fu, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. *arXiv:1906.00949 [cs, stat]*, Nov 2019. URL http://arxiv.org/abs/1906.00949. arXiv: 1906.00949.

Aviral Kumar, Abhishek Gupta, and Sergey Levine. Discor: Corrective feedback in reinforcement learning via distribution correction. *arXiv:2003.07305 [cs, stat]*, Mar 2020. URL http://arxiv.org/abs/2003.07305. arXiv: 2003.07305.

Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*, 2019.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv:1612.01474 [cs, stat]*, Nov 2017. URL http://arxiv.org/abs/1612.01474. arXiv: 1612.01474.

Kimin Lee, Michael Laskin, Aravind Srinivas, and Pieter Abbeel. Sunrise: A simple unified framework for ensemble learning in deep reinforcement learning. *arXiv:2007.04938 [cs, stat]*, Jul 2020. URL http://arxiv.org/abs/2007.04938. arXiv: 2007.04938.

Dan Levi, Liran Gispan, Niv Giladi, and Ethan Fetaya. Evaluating and calibrating uncertainty prediction in regression tasks. *arXiv:1905.11659 [cs, stat]*, Feb 2020. URL http://arxiv.org/abs/1905.11659. arXiv: 1905.11659.

Vincent Mai, Waleed Khamies, and Liam Paull. Batch inverse-variance weighting: Deep heteroscedastic regression. *arXiv:2107.04497 [cs, stat]*, Jul 2021. URL http://arxiv.org/abs/2107.04497. arXiv: 2107.04497.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *NIPS*, pp. 9, 2013.

Kevin P. Murphy. *Machine learning: a probabilistic perspective*. Adaptive computation and machine learning series. MIT Press, 2012. ISBN 978-0-262-01802-9.

D.A. Nix and A.S. Weigend. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, pp. 55–60 vol.1. IEEE, 1994. ISBN 978-0-7803-1901-1. doi: 10.1109/ICNN.1994.374138. URL http://ieeexplore.ieee.org/document/374138/.

OpenAI, Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, and et al. Dota 2 with large scale deep reinforcement learning. *arXiv:1912.06680 [cs, stat]*, Dec 2019. URL http://arxiv.org/abs/1912.06680. arXiv: 1912.06680.

Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. *arXiv:1602.04621 [cs, stat]*, Jul 2016. URL http://arxiv.org/abs/1602.04621. arXiv: 1602.04621.

Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper/2018/file/5a7b238ba0f6502e5d6be14424b20ded-Paper.pdf.

Ian Osband, Benjamin Van Roy, Daniel J. Russo, and Zheng Wen. Deep exploration via randomized value functions. *Journal of Machine Learning Research*, 20(124):1–62, 2019.

Ian Osband, Yotam Doron, Matteo Hessel, John Aslanides, Eren Sezener, Andre Saraiva, Katrina McKinney, Tor Lattimore, Csaba Szepesvari, Satinder Singh, and et al. Behaviour suite for reinforcement learning. *arXiv:1908.03568 [cs, stat]*, Feb 2020. URL http://arxiv.org/abs/1908.03568. arXiv: 1908.03568.

Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua V. Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. *arXiv:1906.02530 [cs, stat]*, Dec 2019. URL http://arxiv.org/abs/1906.02530. arXiv: 1906.02530.

Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. *arXiv:1705.05363 [cs, stat]*, May 2017. URL http://arxiv.org/abs/1705.05363. arXiv: 1705.05363.

Joshua Romoff, Peter Henderson, Alexandre Piché, Vincent François-Lavet, and Joelle Pineau. Reward estimation for variance reduction in deep reinforcement learning. (CoRL):1–26, 2018.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. pp. 1–12, 2017. ISSN 0008-5286 (Print)r0008-5286 (Linking). doi: 10.1007/s00038-010-0125-8.

David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, and et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, Jan 2016. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature16961.

Malcolm Strens. A bayesian framework for reinforcement learning. *Proceedings of the Seventeenth International Conference on Machine Learning*, 02 2001.

Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: an introduction*. Adaptive computation and machine learning series. The MIT Press, second edition edition, 2018. ISBN 978-0-262-03924-6.

William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285, Dec 1933. ISSN 00063444. doi: 10.2307/2332286.

Universite de Montreal. Montreal declaration for a responsible development of artificial intelligence 2018, 2018.

Tingwu Wang, Xuchan Bao, Ignasi Clavera, Jerrick Hoang, Yeming Wen, Eric Langlois, Shunshi Zhang, Guodong Zhang, Pieter Abbeel, and Jimmy Ba. Benchmarking model-based reinforcement learning. *arXiv:1907.02057 [cs, stat]*, Jul 2019. URL http://arxiv.org/abs/1907.02057. arXiv: 1907.02057.

Yue Wu, Shuangfei Zhai, Nitish Srivastava, Joshua Susskind, Jian Zhang, Ruslan Salakhutdinov, and Hanlin Goh. Uncertainty weighted actor-critic for offline reinforcement learning. *arXiv:2105.08140 [cs]*, May 2021. URL http://arxiv.org/abs/2105.08140. arXiv: 2105.08140.

A Algorithms in details

In this section, we are detailing the different algorithms described in the paper, and explaining the design choices that we made.

A.1 EnsembleDQN

EnsembleDQN is based on DQN (Mnih et al., 2013). Instead of a Q- and a target network, we use a Q-ensemble and a target ensemble. It is the base algorithm upon which the others are built.

Network architecture The Q-ensemble and the target ensemble are initialized identically with N neural networks. We used multi-layer perceptrons with 2 ReLU-activated hidden layers of 64 nodes.

TD update Following the implementation of Osband et al. (2016), the networks are paired one to one between the Q- and the target networks during the TD update. The system can be considered as several DQNs working in parallel, whose only interaction is during the computation of the variance.

The best action (s', a') is selected based on the mean value across the ensemble of target networks. Voting was also tested, with slightly lower performances.

Exploration During exploration, an $e-$greedy selection is applied over the highest evaluated action of the Q-network, once again based on the mean value across the ensemble.

In the Bootstrap variant, during exploration, one network of the ensemble is used to predict the best action for a whole episode. This also follows Osband et al. (2016), and is an approximation of Thompson sampling (Thompson, 1933). We added RPFs (Osband et al., 2018) to tackle the variance underestimation at early training.

Masked replay buffer To mask the replay buffer, each step saved in the replay buffer is associated to a N-sized boolean vector m which is used as a mask during training. Each element of m indicates if the corresponding network in the Q-ensemble should be trained using this sample. These elements of m are independently generated when the sample is saved, using a Bernoulli distribution with probability p_m. They are kept fixed during the training. The role of the mask is to ensure variability in the training of the different networks of the Q-ensemble, and thus maintain a significant variance prediction.

A.2 IV-SAC

Network architecture The networks used in SAC are multilayer perceptrons with flattened inputs and 2 ReLU-activated hidden layers of 256 nodes. Our implementation is based on rlkit[7].

[7]https://github.com/rail-berkeley/rlkit

IV-RL on the critic SAC is a complex algorithm with several details which have no impact on IV-RL, for example the accounting of policy entropy in the definition of the Q-value. On the critic side, the implementation of IV-RL is mostly similar to IV-DQN's TD-update, except for one detail: "clipped double-Q learning".

SAC uses the "clipped double-Q learning": it uses two Q-networks and target networks, and always chooses the smallest value to avoid overestimation in the Q-function. To compute both the value and the variance, we always replace each Q- and target network with a pair of networks which returns the smallest value of the two.

IV-RL on the actor Following the example of Lee et al. (2020), we use an ensemble for the actor network when using ensembles for the critic. Each network in the actor is paired to two networks in the critic. When training, the Q-value and its variance for each state-action are computed similarly as when training the critic.

Action selection The use of an ensemble for the policy allows for different methods to select the next action given a state. In this respect, we follow again the example of Lee et al. (2020).

During the training of the Q-network, as we pair each Q-network with its corresponding actor, the action is directly output by this single actor network. In SAC, the action is sampled from Gaussian stochastic policy, with a predicted mean and variance.

During exploration too, the action is sampled from a Gaussian policy. We additionally take advantage of the ensemble to implement UCB. Each policy network outputs an action. The mean Q-value $\mu_Q(s,a)$ and variance $\sigma_Q^2(s,a)$ for each action a are computed by sampling over all networks of the critic, including both Q-networks. The chosen action has the highest UCB score: $a^* = \arg\max_a[\mu_Q(s,a) + \lambda_{\text{UCB}}\sigma_Q(s,a)]$, where λ_{UCB} is a hyperparameter.

The variance estimation method for UCB may be different than the one used for the training the networks: this is because, for UCB, the variance estimation must calibrated, to be scaled with the Q-values. Probabilistic networks are therefore disqualified.

During testing, the agent takes the average action, as sampled from the mean predictions output by the policy networks.

B HYPERPARAMETER TUNING

Deep reinforcement learning algorithms are known to be particularly difficult to evaluate due to the high stochasticity of the learning process and the dependence on hyperparameter. To thoroughly compare different algorithms, it is important to carefully tune the hyperparameters, based on several environments and network initialization seeds (Henderson et al., 2019).

For every result presented in this paper, the hyperparameters for each algorithm were tuned using grid search. Each combination of hyperparameters was ran five times, with different seeds on both initialization and environment. The best 3 or 4 configurations were then selected to be ran 25 times - this time, the combinations of 5 environment and 5 initialization seeds. The configuration selected to be shown in the paper is the best of these configurations based on the 25 runs.

Table 3 describes the type and sets of parameters that were optimized for the relevant algorithms. Note that the learning rate lr, discount rate γ and target new update rate τ are also relevant for SAC. Because there was too many hyperparameters to tune, we kept for all SAC experiments the default values used in *rlkit*[8] implementation : $actorlr = 0.0003$, $criticlr = 0.0003$, $\gamma = 0.99$ and $\tau = 0.005$.

The hyperparameters used for each curve can be found in the configuration file of the code submitted as additional material.

[8]https://github.com/rail-berkeley/rlkit

Table 3: List and set of hyperparameters that were tuned when testing the algorithm.

Hyperparameter	Relevant algorithms	Set
Effective minibatch size	All	$\{32, 64, 128, 256\}$ for DQN+ $\{256, 512, 1024\}$ for SAC+
Learning rate lr	DQN+	$\{0.0005, 0.001, 0.005, 0.01\}$
Discount rate γ	DQN+	$\{0.98, 0.99, 0.995\}$
Target net update rate τ	DQN+	$\{0.001, 0.005, 0.01, 0.05\}$
e-decay	DQN+ with e-greedy exploration	$\{0.97, 0.98, 0.99, 0.991, 0.995\}$
Mask probability	All with masked ensembles	$\{0.5, 0.7, 0.8, 0.9, 1\}$
RPF prior scale	All with RPF	$\{0.1, 1, 10\}$
Loss attenuation weight λ	All with loss attenuation	$\{0.01, 0.1, 0.5, 1, 2, 5\}$
Minimal EBS ratio	All with minimal EBS	$\{16/32, 24/32, 28/32, 30/32, 31/32\}$ for DQN+ $\{0.5, 0.9, 0.95, 0.99\}$ for SAC+
Minimal variance ϵ	All with constant ϵ	$\{0.5, 1, 5, 100, 1000, 10000\}$
UCB weight λ_{UCB}	All SAC+ with UCB (0 means Bootstrap)	$\{0, 1, 10\}$ as suggested in Lee et al. (2020)
Sunrise temperature	All with Sunrise	$\{10, 20, 50\}$ as suggested in Lee et al. (2020)
UWAC β	All with UWAC	$\{0.8, 1.6, 2.5\}$ as suggested in Wu et al. (2021)

C COMPUTATION TIME

Computation time Replacing a single neural network by an ensemble significantly increases the computation time necessary for the training an agent, as shown in table 4. While this can be a limitation of our work, we consider the cases where computation time is secondary to sample efficiency when samples are the more expensive element, such as in the real world.

Table 4: The table shows average time taken (in seconds) for each algorithm to train for 1000 steps on walker2d environment. We can see that all Ensemble based environments are significantly more computationally expensive than SAC.

	SAC	EnsembleSAC	SUNRISE	IV_SAC
Runtime/1000 steps (s)	8.34	62.16	62.1	74.56

D VARIANCE ESTIMATION

The difference of performance between the used of BIV with sampled ensembles and probabilistic ensembles is sometimes very significant. Looking at the variance prediction can allow us to shed a insightful light on this phenomenon.

D.1 IV-DQN

In figure 6, we show the mean of the sample variance estimations per mini-batch, depending on the variance estimation method.

The first thing to notice is the different profiles that can be seen in LunarLander and Cartpole-Noise. In the first environment, the scale of the variance predicted by sampled ensembles and probabilistic ensembles do not differ significantly. This is completely different in Cartpole-Noise, where the probabilistic ensembles quickly capture much more uncertainty. This explains the difference in dynamics that can be seen in Figure 3.

Although we do not have results for ProbDQN in Cartpole-Noise, we can see in LunarLander why single probabilistic networks are less reliable than probabilistic ensembles. The mean variance they predict is very high, but, as can be seen in Figure 7, the median is a lot lower. The mean is thus driven by extremely high variance predictions. Instead, sampled and probabilistic ensembles show similar profiles for mean and median variance. The low median variance confirms that single probabilistic networks fail to capture epistemic uncertainty, but also that they have very high volatility.

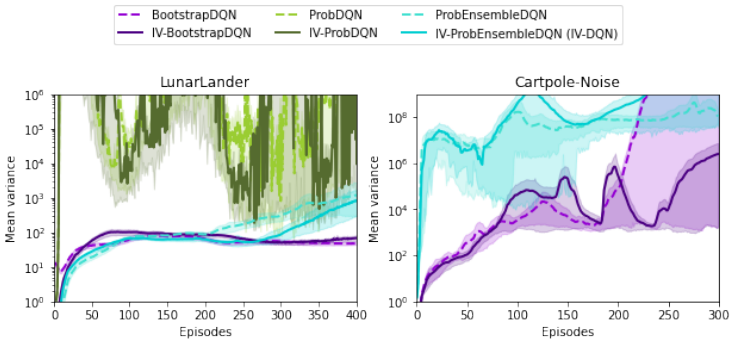

Figure 6: Mean target variance predictions for discrete environments

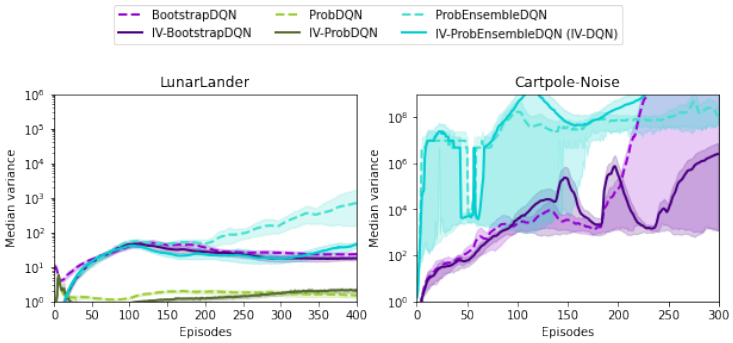

Figure 7: Median target variance predictions on discrete environments

D.2 IV-SAC

In the continuous control environments, the story is different, as can be seen in Figures 8 and 9.

The first thing to notice is the volatility of the mean of the variance produced by probabilistic ensembles. Although the scale is not completely unreasonable and is still close to the median, (on the opposit of probabilistic networks in DQN), it still suffers from some calibration issues. This explains why weighing schemes such as UWAC, which are not normalized like BIV, under-perform with probabilistic ensembles.

Second, we see that the epistemic variance in EnsembleSAC is higher than the one in IV-EnsembleSAC. One hypothesis to explain this is that using IV reduces the impact of very noisy supervision which may confuse the networks and increase the variance in the ensemble. Probabilistic ensembles do not seem to suffer from this.

One aspect that could be investigated more in details is the fact that in some cases, the uncertainty grows. This is particularly interesting in LunarLander with probabilistic ensembles, where it coincides with a drop in the performance after having reached the best score. This could be linked to catastrophic forgetting.

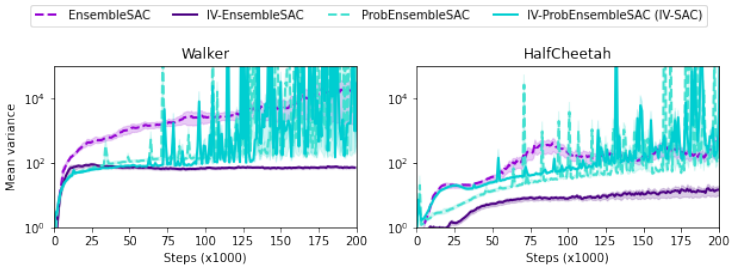

Figure 8: Mean target variance predictions on continuous environments

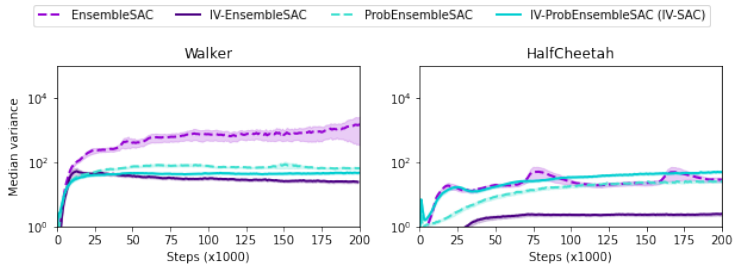

Figure 9: Median target variance predictions on continuous environments

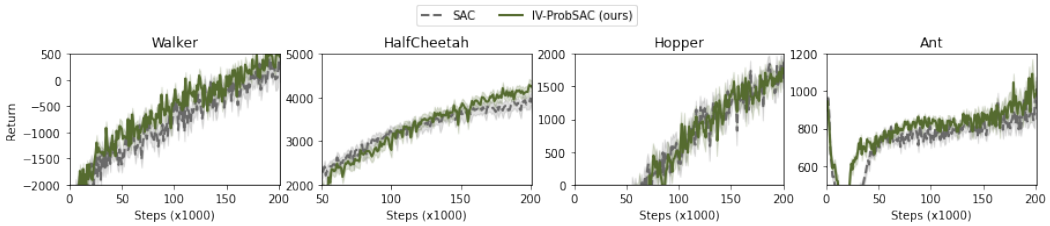

Figure 10: Single probabilistic networks combined with BIV lead to a very slight improvement.

E PROBABILISTIC NETWORKS

What if, instead of using predictive ensembles, we simply replaced the Q-network by a probabilistic network and trained it using the IV-RL loss?

Figure 10 shows the performances on 4 continuous environments, compared to SAC. We can see a slight improvement, but it is not statistically significant. Based on the insights given in section D, we make the hypothesis that it is due to the volatility of the variance estimation with a single probabilistic network.