
The Potential of the Return Distribution for Exploration in RL

Thomas M. Moerland¹ Joost Broekens¹ Catholijn M. Jonker¹

Abstract

This paper studies the potential of the return distribution for exploration in deterministic reinforcement learning (RL) environments. We study network losses and propagation mechanisms for Gaussian, Categorical and Gaussian mixture distributions. Combined with exploration policies that leverage this return distribution, we solve, for example, a randomized Chain task of length 100, which has not been reported before when learning with neural networks.

1. Introduction

Reinforcement learning (RL) is the dominant class of algorithms to learn sequential decision-making from data. Most RL approaches focus on learning the mean action value $Q(s, a)$. Recently, Bellemare et al. (2017) studied distributional RL, where one propagates the entire return distribution $p(Z|s, a)$ (of which $Q(s, a)$ is the expectation) through the Bellman equation. Bellemare et al. (2017) show increased performance in a variety of RL tasks.

However, Bellemare et al. (2017) did not yet leverage the return distribution for exploration. In the present paper, we identify the potential of the return distribution for informed exploration. The return distribution may be induced by two sources of stochasticity: 1) our stochastic policy and 2) a stochastic environment. For this work we assume a *deterministic* environment, which makes the return distribution entirely induced by the stochastic policy. Thereby, we may actually act optimistically with respect to this distribution.¹ The present paper explores this idea, in the context of neural networks and for different propagation distributions (Gaussian, Categorical and Gaussian mixture). Results show vastly improved learning in a challenging exploration task,

¹Department of Computer Science, Delft University of Technology, The Netherlands. Correspondence to: Thomas M. Moerland <T.M.Moerland@tudelft.nl>.

Published at the Exploration in Reinforcement Learning Workshop at the 35th International Conference on Machine Learning, Stockholm, Sweden. Copyright 2018 by the author(s).

¹In Section 7, we more thoroughly discuss the different types of uncertainty present in sequential decision making.

which had not been solved with neural networks before. We also provide extensive visual illustration of the process of return-based exploration, which shows a natural shift from exploration to exploitation.

2. Distributional Reinforcement Learning

We adopt a Markov Decision Process (MDP) (Sutton & Barto, 1998) given by the tuple $\{\mathcal{S}, \mathcal{A}, P, R, \gamma\}$. For this work, we assume a discrete action space and deterministic transition and reward functions. At every time-step t we observe a state $s_t \in \mathcal{S}$ and pick an action $a_t \in \mathcal{A} = \{1, 2, \dots, |\mathcal{A}|\}$. The MDP follows the transition dynamics $s_{t+1} = P(s_t, a_t) \in \mathcal{S}$ and returns rewards $r_t(s, a) = R(s_t, a_t) \in \mathbb{R}$. We act in the MDP according to a stochastic policy $\pi(\cdot|s) \in \mathcal{P}(\mathcal{A})$. The (discounted) return $Z^\pi(s, a)$ from a state-action pair (s, a) is a *random process* given by

$$Z^\pi(s, a) = \sum_{t=0}^{\infty} \gamma^t r_t \quad (1)$$

where $s_{t+1} = P(s_t, a_t)$, $a_{t+1} \sim \pi(\cdot|s_{t+1})$, $s_0 = s$, $a_0 = a$. The return Z^π is a random variable, where the distribution of Z^π is induced by the stochastic policy (as we assume a deterministic environment). Eq. 1 can be unwritten in recursive form, known as the *distributional Bellman equation* (Bellemare et al., 2017) (omitting the π superscript from now on):

$$Z(s, a) \stackrel{d}{=} r(s, a) + \gamma \mathbb{E}_{A'}[Z(s', A')] \quad (2)$$

where $\stackrel{d}{=}$ denotes distributional equality (Engel et al., 2005). The state action value $Q(s, a) = \mathbb{E}_Z[Z(s, a)]$ is the expectation of the return distribution. Applying this expectation to Eq. 2 gives

$$Q(s, a) = r(s, a) + \gamma \mathbb{E}_{A' \sim \pi(\cdot|s')} [Q(s', A')], \quad (3)$$

$s' = P(\cdot|s, a)$, which is known as the Bellman equation (Sutton & Barto, 1998). Most RL algorithms learn this mean action value $Q(s, a)$, and explore by some random perturbation of these means.

3. Distributional Perspective on Exploration

As mentioned in the introduction, the return distribution may be induced by two sources of stochasticity: 1) our stochastic policy and 2) a stochastic environment. Therefore, if we assume a deterministic environment, then the return distribution is entirely induced by our own policy. As we may modify our policy, it actually makes sense to act optimistically with respect to the return distribution.

As an illustration, consider a state-action pair with particular mean value estimate $Q(s, a)$. It matters whether this average originates from a highly varying return or from consistently the same return. It matters because our policy may *influence* the shape of this distribution, i.e. for the highly varying returns we may actively transform the distribution towards the good returns. In other words, what we really care about in deterministic domains is the best return, or the upper end of the return distribution, because it is an indication of what we may achieve once we have figured out how to act in the future. By starting from broad distribution initializations that gradually narrow when subpolicies converge, we observe a natural shift from exploration to exploitation.

4. Distributional Policy Evaluation

Following Bellemare et al. (2017), we introduce a neural network to model the return distribution $p_\phi(Z|s, a)$. For this work we will consider three parametric distributions $p_\phi(Z)$ to approximate the return distribution: Gaussian, Categorical (as previously studied by Bellemare et al. (2017)), and Gaussian mixture.

To perform policy evaluation, we need to discuss two topics:

1. How to propagate the distribution through the Bellman equation (based on newly observed data). We will denote the propagated distribution as $q(Z)$.
2. A loss function between the current network predictions and the new target: $L(p_\phi(Z), q(z))$.

Due to space restrictions, we will only show the propagation and loss for the Gaussian case. For the Categorical and Gaussian mixture outcome we specify the distributional loss in Appendix A.1 and the Bellman propagation in Appendix A.2.

Distribution propagation Define the *distributional Bellman back-up operator* \mathcal{T} as the recursive application of Eq. 2 to $Z(s, a)$. For a Gaussian network output, $p_\phi(Z|s, a) = \mathcal{N}(Z|\mu_\phi(s, a), \sigma_\phi(s, a))$, we need to propagate both the mean and variance through the Bellman equation.

$$\begin{aligned} \mu^q(s, a) &= \mathbb{E}_Z [\mathcal{T}Z(s, a)] \\ &= \mathbb{E}_Z [r(s, a) + \gamma \mathbb{E}_{A'} Z(s', A')] \\ &= r(s, a) + \gamma \mathbb{E}_{A'} \mu_\phi(s', A') \end{aligned} \quad (4)$$

$$\begin{aligned} \sigma^q(s, a) &= \text{Sd} [\mathcal{T}Z(s, a)] = \text{Sd} [r(s, a) + \gamma \mathbb{E}_{A'} Z(s', A')] \\ &= \gamma \text{Sd} [\mathbb{E}_{A'} Z(s', A')] = \gamma \mathbb{E}_{A'} \sigma_\phi(s', A') \end{aligned} \quad (5)$$

because $\gamma \geq 0$, $\pi(a'|s') \geq 0$, and we assume the next state distributions are independent so we may ignore the covariances.²

In practice, we approximate the expectation over the policy probabilities at the next state a' by sampling a next state once (either on- or off-policy). This is the most common solution in RL, and will be right in expectation over multiple traces. The 1-step bootstrap distribution estimate then becomes $q(Z|s, a) = \mathcal{N}(Z|r(s, a) + \gamma \mu_\phi(s', a'), \gamma \sigma_\phi(s', a'))$.

Loss Next, we want to move our current network predictions $p_\phi(Z|s, a)$ closer to the new target $q(Z|s, a)$, for which we will use a distributional distance. A well-known choice in machine learning is the *cross-entropy* $H(q, p)$:

$$L_{\text{CE}}(\phi) = H(q(Z), p_\phi(Z)) = \mathbb{E}_{q(Z)} [-\log p_\phi(Z)]. \quad (6)$$

For both the Gaussian and Categorical output distributions we can derive closed-form expressions for the cross-entropy H , see Appendix A.1. However, for the Gaussian mixture outcome we do not have a closed-form cross-entropy expression, and we instead minimize the L_2 distance. See Appendix A.1 for details as well.

In practice, we store a database of transition tuples $\{s, a, r, s', a'\}$, where a' can also be computed either on- or off-policy, and minimize:

$$L_{\text{CE}}(\phi) = \mathbb{E}_{\{s, a, r, s', a'\} \in \mathcal{D}} \left[\mathbb{E}_{q(Z|s, a)} [-\log p_\phi(Z|s, a)] \right] \quad (7)$$

where $q(Z) = f(r, s', a')$ is computed based on Bellman propagation. This completes the policy evaluation step for the Gaussian case.

²For random variables X, Y and scalar constants a, b, c we have: $\mathbb{E}[a + bX + cY] = a + b\mathbb{E}[X] + c\mathbb{E}[Y]$ and $\text{Var}[a + bX + cY] = b^2 \text{Var}[X] + c^2 \text{Var}[Y] + 2bc \text{Cov}[X, Y]$.

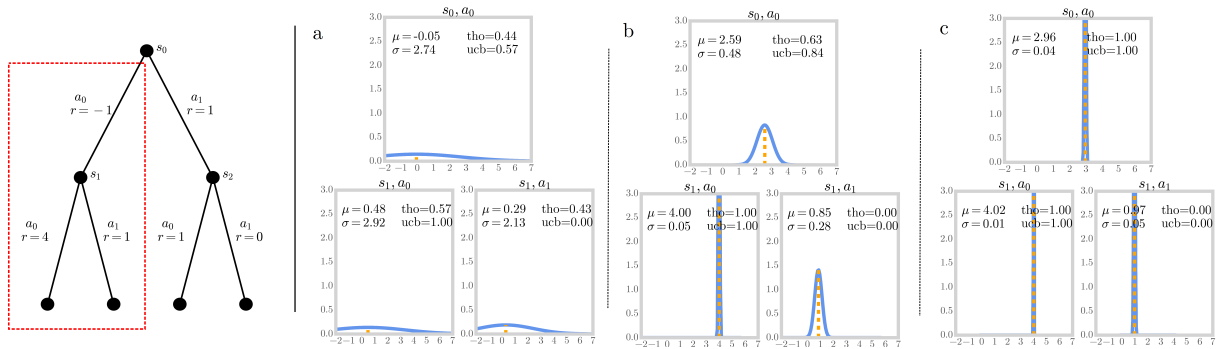


Figure 1. Gaussian distribution propagation. Left: Example 2-step MDP. Learning process for the state-action pairs in the dotted box is shown on the right. Right: Three (a-c) return exploration phases for the left half of the MDP. Each plot also displays the mean (μ), standard deviation (σ), and policy probabilities under Thompson (tho) sampling and UCB (ucb).

5. Distributional Exploration (Policy Improvement)

Our real interest is usually not in policy evaluation only, as we want to gradually improve our policy as well. The major benefit of probabilistic policy evaluation (previous section) is that we have additional information to balance exploration and exploitation. We will treat the return distribution as something against which we can act optimistically. Exploration under uncertainty has been extensively studied in the bandit literature. Two of the most successful algorithms, which we both consider in this work, are

1. Thompson sampling (Thompson, 1933), which takes a sample $z_a \sim p(Z|s, a)$ for each action and picks the action with the highest draw.
2. Upper Confidence Bounds (UCB) (Auer et al., 2002), which picks the action with the highest upper confidence bound $\mu_Z + c \cdot \sigma_Z$ for some constant $c \in \mathbb{R}^+$. Analytic expressions for σ for the different output distributions are provided in Appendix A.3.

6. Experiments

We now show several results of return-based exploration on a Toy example, Chain domain and OpenAI Gym task. Fig. 1, left shows an example 2-step MDP to illustrate the concept of return-based exploration. On the right of Fig. 1 we display three phases of learning in this MDP for a Gaussian $p_\phi(Z)$. Due to space constraints we only visualize the distributions for the left half of the MDP, the full learning process is shown in Figure 5 (Appendix). In Fig. 1a we just initialized the network, and both Thompson sampling and UCB follow almost uniform policies. After some training (Fig. 1b) the second state (s_1) distributions start converging, but the uncertainty at s_0 still remains broader (as it generalizes over the sometimes explored inferior a_1 in s_1). Thompson sampling and UCB gradually start to prefer a_0 in the root

state s_0 now. Finally, after some additional episodes (Fig. 1c) the distribution estimates have converged on the optimal state-action values, and both Thompson sampling and UCB have automatically converged on the optimal policy.

We next consider the Chain domain (Appendix C, Figure 4), which has been previously studied in RL literature (Osband et al., 2016). The domain consists of a chain of states of length N , with two available actions at each state. The only trace giving a positive, non-zero reward is to select the ‘correct’ action at every step, which is randomly picked at domain initialization. The domain has a strong exploration challenge, which grows exponentially with the length of the chain for undirected exploration methods (see Appendix C).

Figure 2 show the learning curves of return-based exploration for different types of output distributions, and compares the results to ϵ -greedy exploration. The plots progress row-wise to longer chain lengths. First of all, we note that ϵ -greedy learns slowly in the short chain, and does not learn at all in the longer chains. However, the methods with return uncertainty do learn, and consistently solve the domain even for the long chain of length 100. Note that this is a very challenging exploration problem, as we need to take 100 steps correctly while there is no structure in the domain at all (i.e., the correct action randomly changes at each depth in the chain, so a function approximator with local generalization is more harmful than beneficial). The mixture of Gaussians (mog) return distribution performs less stable than the Gaussian and categorical. This might be due to the L_2 -loss used for the Gaussian mixture, which is different from the cross-entropy losses used for the Gaussian and categorical distributional loss (see Appendix A). In Figure 6 we provide a full illustration of how exploration based on a Categorical return distribution performs. A detailed illustration of the learning process for the categorical return distribution is shown in Figure 6 (Appendix).

In Figure 3 we show the results of return-based exploration on a task from the OpenAI Gym repository: FrozenLake.




Figure 2. Learning curves on Chain domain for different types of return distributions. Plots progress row-wise for increased depth of the Chain, i.e. increased exploration difficulty. Exploration uses a UCB policy with constant $c_a \sim \text{Uniform}(1.7, 2.3)$ for each a (this induces slight randomness in the otherwise deterministic UCB decision). Results averaged over 10 repetitions.




Figure 3. Return-based exploration versus ϵ -greedy on FrozenLake-v1. The return-based exploration methods use Thompson sampling. Compared to the OpenAI Gym implementation, we modify the environment to be fully deterministic (by removing the random ‘slipping’ effect of the task). Results averaged over 5 repetitions.

Again, we observe that the return based exploration methods learn better than ϵ -greedy. These experiments use Thompson sampling for exploration, which shows that return-based exploration can be employed with both UCB or Thompson exploration.

7. Discussion

We shortly discuss why return-based exploration seems to work so well in the challenging exploration task of the Chain. It turns out that the return distributions really narrow when a certain action terminates the episode. In such cases, we bootstrap a very narrow next state distribution around 0 (because the reward function is assumed deterministic). On the Chain, we see that all the terminating actions very quickly narrow, while the trace along the full path keeps some additional uncertainty. It appears as if the return distribution in this implementation identifies a specific type of uncertainty related to the termination probabilities and asymmetry in the domain search tree, which relates our work to ideas from Monte Carlo Tree Search (Moerland et al., 2018) as well.

The benefit of exploration based on uncertainty is that policy improvement almost comes for free. The only thing that we propagate are full distributions, which we initialize wide,

and then gradually converge when the distributions behind it start converging. This creates a more natural transition from the exploration to the exploitation phase, a trait which most undirected methods (ϵ -greedy, Boltzmann) lack.

An important direction for future work is to connect the policy-dependent return uncertainty, as studied in this paper, to the *statistical* (or epistemic) uncertainty of the mean action-value, which is a function of the local number of visits to a state. The return distribution mechanism in this paper clearly identifies a different aspect of (future policy) uncertainty, which may be related to the termination structure of subtrees, or to the fact that uncertainty in an MDP should propagate over steps as well (Dearden et al., 1998). In any case, due to the sequential nature of MDPs there appear to be more aspects to epistemic/reducible (Osband et al., 2018) uncertainty, and these distinctions have yet to be properly identified. Finally, another important extension is to stochastic environments (Depeweg et al., 2016; Moerland et al., 2017b), i.e., separating which part of the return distribution originates from our own policy uncertainty (for which we can be optimistic) and which part originates from the stochastic environment (for which we want to act on the expectation).

8. Conclusion

This paper identified the potential of the return distribution for targeted exploration. In deterministic domains, the return distribution is induced by our own policy, and since we may modify this policy ourselves, it makes sense to act optimistically with respect to this distribution. Exploration based on the return distribution, especially for the Gaussian and Categorical case, manages to solve the ‘randomized’ Chain of length 100 with function approximation, which we believe has not been reported before. Moreover, it also performs well in another task from the OpenAI Gym. Future work should expand these ideas to stochastic environments, and identify the connections to exploration based on the statistical uncertainty of the mean action value.

References

- Auer, Peter, Cesa-Bianchi, Nicolo, and Fischer, Paul. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Azizzadenesheli, Kamyar, Brunskill, Emma, and Anandkumar, Animashree. Efficient Exploration through Bayesian Deep Q-Networks. In *Deep Reinforcement Learning Symposium, NIPS*, 2017.
- Bellemare, Marc, Srinivasan, Sriram, Ostrovski, Georg, Schaul, Tom, Saxton, David, and Munos, Remi. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pp. 1471–1479, 2016.
- Bellemare, Marc G, Dabney, Will, and Munos, Rémi. A distributional perspective on reinforcement learning. *arXiv preprint arXiv:1707.06887*, 2017.
- Dearden, Richard, Friedman, Nir, and Russell, Stuart. Bayesian Q-learning. In *AAAI/IAAI*, pp. 761–768, 1998.
- Depeweg, Stefan, Hernández-Lobato, José Miguel, Doshi-Velez, Finale, and Udluft, Steffen. Learning and policy search in stochastic dynamical systems with bayesian neural networks. *arXiv preprint arXiv:1605.07127*, 2016.
- Engel, Yaakov, Mannor, Shie, and Meir, Ron. Reinforcement learning with Gaussian processes. In *Proceedings of the 22nd international conference on Machine learning*, pp. 201–208. ACM, 2005.
- Gal, Yarin, McAllister, Rowan Thomas, and Rasmussen, Carl Edward. Improving PILCO with bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop*, volume 951, pp. 2016, 2016.
- Guez, Arthur, Silver, David, and Dayan, Peter. Efficient Bayes-adaptive reinforcement learning using sample-based search. In *Advances in Neural Information Processing Systems*, pp. 1025–1033, 2012.
- Henderson, Peter, Doan, Thang, Islam, Riashat, and Meger, David. Bayesian Policy Gradients via Alpha Divergence Dropout Inference. *arXiv preprint arXiv:1712.02037*, 2017.
- Jeong, Heejin and Lee, Daniel D. Bayesian Q-learning with Assumed Density Filtering. *arXiv preprint arXiv:1712.03333*, 2017.
- Kaufmann, Emilie and Koolen, Wouter M. Monte-carlo tree search by best arm identification. In *Advances in Neural Information Processing Systems*, pp. 4897–4906, 2017.
- Mannor, Shie and Tsitsiklis, John. Mean-variance optimization in Markov decision processes. *arXiv preprint arXiv:1104.5601*, 2011.
- Moerland, Thomas M, Broekens, Joost, and Jonker, Catholijn M. Efficient exploration with Double Uncertain Value Networks. *arXiv preprint arXiv:1711.10789*, 2017a. Deep Reinforcement Learning Symposium @ NIPS 2017.
- Moerland, Thomas M, Broekens, Joost, and Jonker, Catholijn M. Learning Multimodal Transition Dynamics for Model-Based Reinforcement Learning. *arXiv preprint arXiv:1705.00470*, 2017b.
- Moerland, Thomas M, Broekens, Joost, Plaat, Aske, and Jonker, Catholijn M. Monte Carlo Tree Search for Asymmetric Trees. *arXiv preprint arXiv:1805.09218*, 2018.
- Morimura, Tetsuro, Sugiyama, Masashi, Kashima, Hisashi, Hachiya, Hirotaka, and Tanaka, Toshiyuki. Parametric return density estimation for reinforcement learning. *arXiv preprint arXiv:1203.3497*, 2012.
- Osband, Ian, Van Roy, Benjamin, and Wen, Zheng. Generalization and exploration via randomized value functions. *arXiv preprint arXiv:1402.0635*, 2014.
- Osband, Ian, Blundell, Charles, Pritzel, Alexander, and Van Roy, Benjamin. Deep exploration via bootstrapped DQN. In *Advances in Neural Information Processing Systems*, pp. 4026–4034, 2016.
- Osband, Ian, Aslanides, John, and Cassirer, Albin. Randomized Prior Functions for Deep Reinforcement Learning. *arXiv preprint arXiv:1806.03335*, 2018.
- Petersen, Kaare Brandt, Pedersen, Michael Syskind, et al. The matrix cookbook. *Technical University of Denmark*, 7(15):510, 2008.
- Russo, Daniel, Van Roy, Benjamin, Kazerouni, Abbas, and Osband, Ian. A Tutorial on Thompson Sampling. *arXiv preprint arXiv:1707.02038*, 2017.
- Sobel, Matthew J. The variance of discounted Markov decision processes. *Journal of Applied Probability*, 19(4):794–802, 1982.
- Sutton, Richard S and Barto, Andrew G. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- Tamar, Aviv, Di Castro, Dotan, and Mannor, Shie. Learning the variance of the reward-to-go. *Journal of Machine Learning Research*, 17(13):1–36, 2016.
- Tang, Yunhao and Agrawal, Shipra. Exploration by Distributional Reinforcement Learning. *arXiv preprint arXiv:1805.01907*, 2018.
- Tang, Yunhao and Kucukelbir, Alp. Variational Deep Q Network. *arXiv preprint arXiv:1711.11225*, 2017.

Tesauro, Gerald, Rajan, VT, and Segal, Richard. Bayesian inference in monte-carlo tree search. *arXiv preprint arXiv:1203.3519*, 2012.

Thompson, William R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

White, DJ. Mean, variance, and probabilistic criteria in finite Markov decision processes: a review. *Journal of Optimization Theory and Applications*, 56(1):1–29, 1988.

A. Distributional Details

The current network distributions are denote by $p_\phi(Z|s, a)$, which we want to update with a newly calculated target distribution $q(Z|s, a)$. For readability, we will omit the dependency on s, a in the remainder of this section. We study three types of network output distributions:

1. Gaussian: $p(Z) = \mathcal{N}(Z|\mu, \sigma)$
2. Categorical: $p(Z)$ parametrized by the number of bins $N \in \mathbb{N}$ and edges $Z_{min}, Z_{max} \in \mathbb{R}$. Define the set of bins as $\{z_i = Z_{min} + 0.5\Delta z + i\Delta z : 0 \leq i < N\}$, for $\Delta z := \frac{Z_{max} - Z_{min}}{N}$. Each bin has associated density $p(z_i)$, with $\sum_{i=1}^N p(z_i) = 1$.
3. Gaussian mixture: $p(Z) = \sum_{i=1}^M p_i \mathcal{N}(Z|\mu_i, \sigma_i)$, for M mixtures. Here p_i denotes the weight of the i -th mixture component, $\sum_i p_i = 1$.

We now detail the loss, Bellman propagation and analytic standard deviation (as used in the UCB policy) for each of these output distributions.

A.1. Loss

Gaussian The main text already introduced the cross-entropy loss $L_{CE} = H(q, p)$ used for Gaussian $p_\phi(Z)$. Here we derive the analytical expression of this cross-entropy:

$$\begin{aligned} H(q, p) &= \mathbb{E}_{q(Z)}[-\log p(Z)] \\ &= - \int q(Z) \log \frac{1}{\sqrt{2\pi\sigma_p^2}} \exp\left(-\frac{(Z - \mu_p)^2}{2\sigma_p^2}\right) dZ. \end{aligned} \quad (8)$$

Bringing everything that does not depend on Z out of the integral and taking the logarithm:

$$H(q, p) = \frac{1}{2} \log(2\pi\sigma_p^2) + \frac{1}{2\sigma_p^2} \int q(Z)(Z - \mu_p)^2 dZ \quad (9)$$

Which can be rewritten as

$$\begin{aligned} H(q, p) &= \frac{1}{2} \log(2\pi\sigma_p^2) + \frac{1}{2\sigma_p^2} \left[\mathbb{E}_{q(Z)}[Z^2] \right. \\ &\quad \left. - \mathbb{E}_{q(Z)}[Z\mu_p] + \mathbb{E}_{q(Z)}[\mu_p^2] \right] \end{aligned} \quad (10)$$

We can rewrite the second moment $\mathbb{E}_{q(Z)}[Z^2]$ in terms of the mean and variance of Z :

$$\mathbb{E}_{q(Z)}[Z^2] = \sigma_q^2 + \mu_q^2. \quad (11)$$

Therefore, we can simplify the full expression to

$$\begin{aligned} H(q, p) &= \frac{1}{2} \log(2\pi\sigma_p^2) + \frac{1}{2\sigma_p^2} (\sigma_q^2 + \mu_q^2 - 2\mu_q\mu_p + \mu_p^2) \\ &= \frac{1}{2} \log(2\pi\sigma_p^2) + \frac{\sigma_q^2 + (\mu_q - \mu_p)^2}{2\sigma_p^2} \end{aligned} \quad (12)$$

which is used as the closed-form loss for the Gaussian experiments (Eq. 6) in this paper. Note that we also experimented with (other) closed form distributional losses for Gaussians, such as the Bhattacharyya distance and Hellinger distance, but these did not significantly improve performance.

Categorical For the categorical target distribution $q(Z)$ we again minimize the cross-entropy H with $p_\phi(Z)$:

$$L_{CE} = H(q(Z), p_\phi(Z)) = - \sum_i q(Z_i) \log p_\phi(Z_i). \quad (13)$$

Gaussian Mixture There is no closed form expression for the KL-divergence or cross-entropy between two Gaussian mixtures. We could of course approximate such a loss by repeated sampling, but this will strongly increase the computational burden. Therefore, we instead searched for a distance measure between Gaussian mixtures that does have a closed form expression, which is the L_2 -distance:

$$\begin{aligned}
 L_{L2}(q(Z), p_\phi(Z)) &= \int \left(q(Z) - p(Z; \phi) \right)^2 dZ \\
 &= \sum_{i,i'} q_i q_{i'} \int q_i(Z) q_{i'}(Z) dZ \\
 &+ \sum_{j,j'} p_j p_{j'} \int p_i(Z) p_{i'}(Z) dZ \\
 &- 2 \sum_{i,j} q_i p_j \int q_i(Z) p_j(Z) dZ. \quad (14)
 \end{aligned}$$

We may simplify the remaining integrals in this expression, since for any two Gaussians \mathcal{N}_1 and \mathcal{N}_2 we have (Petersen et al., 2008):

$$\int \mathcal{N}_1(Z|\mu_1, \sigma_1) \mathcal{N}_2(Z|\mu_2, \sigma_2) dZ = \mathcal{N}(\mu_1|\mu_2, \sigma_1 + \sigma_2), \quad (15)$$

Therefore, Eq. 14 simplifies to

$$\begin{aligned}
 L_{L2}(q(Z), p_\phi(Z)) &= \sum_{i,i'} q_i q_{i'} \mathcal{N}(\mu_i^q | \mu_{i'}^q, \sigma_i^q + \sigma_{i'}^q) \\
 &+ \sum_{j,j'} p_j p_{j'} \mathcal{N}(\mu_j^p | \mu_{j'}^p, \sigma_j^p + \sigma_{j'}^p) \\
 &- 2 \sum_{i,j} q_i p_j \mathcal{N}(\mu_i^q | \mu_j^p, \sigma_i^q + \sigma_j^p), \quad (16)
 \end{aligned}$$

which can be evaluated in $O(M^2)$ time for M mixture components.

Sample-based loss For some output distributions we either do not have a density (like some deep generative models) or the available analytic distributional loss performs suboptimal. However, we can always sample from our model. For example, for a 1-step Q-learning update, we can (repeatedly) sample from our network at the next timestep $Z'_k \sim p(Z|s', a')$, transform these through the Bellman equation, and then train our model on a negative log-likelihood loss:

$$\begin{aligned}
 L_{NLL} &= \mathbb{E}_{q(Z)} [-\log p_\phi(Z)] \\
 &\approx - \sum_k \log p_\phi(r + \gamma \cdot Z'_k), \quad Z'_k \sim p(Z|s', a')
 \end{aligned} \quad (17)$$

Results of this approach are not shown, but were comparable to the results with approximate return propagation shown in

the Results section. However, this approach is clearly more computationally expensive.

A.2. Bellman Propagation

Given a data tuple $\{s, a, r, s', a'\}$, where a' may either be on- or off-policy, and a bootstrapped distribution $p_\phi(z'|s', a')$, we want to calculate the one-step Bellman transformed distribution $\mathcal{T}p(z|s, a)$.

Categorical For the categorical distribution, we may Bellman transform each individual atom/bin, and then project the probabilities of the transformed means back on the atoms (denoted by operator Ψ). This procedure follows Bellemare et al. (2017):

$$\begin{aligned}
 q(Z_i) &= \left(\Psi \mathcal{T} Z \right)_i \\
 &= \sum_{j=1}^N p_{\phi^\dagger}(Z'_j) \cdot \left[1 - \frac{[r + \gamma Z'_j]_{Z_{min}}^{Z_{max}} - Z_i}{\Delta Z} \right]_0^1 \quad (18)
 \end{aligned}$$

Gaussian mixture For the Gaussian mixture case, we have

$$\begin{aligned}
 q(Z) &= \mathcal{T}Z(s, a) \\
 &= r(s, a) + \gamma \mathbb{E}_{a' \sim \pi(\cdot|s')} [Z'] \\
 &= r(s, a) + \gamma \mathbb{E}_{a' \sim \pi(\cdot|s')} \left[\sum_i p_i Z'_i \right] \\
 &= \sum_i p_i \left[r(s, a) + \gamma \mathbb{E}_{a' \sim \pi(\cdot|s')} [Z'_i] \right] \quad (19)
 \end{aligned}$$

This implies that we may propagate each Gaussian mixture component individually, as discussed in Section 4, keeping each mixture weight the same.

A.3. Standard Deviation

For UCB exploration, we require fast (i.e., analytic) access to the distribution standard deviation, to prevent repeatedly having to sample. Clearly, for the Gaussian output we directly have the standard deviation available.

Categorical For a categorical output distribution with categories z_i and associated probabilities $p(z_i)$, we have the standard deviation as:

$$\text{Sd}[Z] = \sum_i p(Z_i) (Z_i - \mathbb{E}[Z])^2 \quad (20)$$

where $\mathbb{E}[Z] = \sum_i Z_i \cdot p(Z_i)$.

Gaussian mixture For a Gaussian mixture model with mixture weights p_i , mixture means μ_i and mixture standard deviation σ_i , we start from:

$$\begin{aligned} \text{Var}[y] &= \mathbb{E}_{p(Z)}[Z^2] - (\mathbb{E}_{p(Z)}[Z])^2 \\ &= \sum_i p_i \cdot \mathbb{E}_{p(Z_i)}[Z_i^2] - \left(\sum_i p_i \cdot \mu_i\right)^2 \end{aligned} \quad (21)$$

Now we may again use Eq. 11 to rewrite the second moments of the mixture components in terms of their means and variances, i.e. $\mathbb{E}_{p(Z_i)}[Z_i^2] = \sigma_i^2 + (\mu_i)^2$. Plugging this expression into Eq. 21 gives:

$$\begin{aligned} \text{Var}[y] &= \sum_i p_i \cdot \left(\sigma_i^2 + (\mu_i)^2\right) - \left(\sum_i p_i \cdot \mu_i\right)^2 \\ &= \sum_i p_i \cdot \sigma_i^2 + \sum_i p_i (\mu_i)^2 - \left(\sum_i p_i \cdot \mu_i\right)^2 \end{aligned} \quad (22)$$

This last expression gives the variance of the mixture in terms of the weight, mean and variance of the mixture components.

B. Related Work

Return Uncertainty While the distributional Bellman equation (Eq. 2) is certainly not new (Sobel, 1982; White, 1988), nearly all RL research has focussed on the mean action-value. Most papers that do study the underlying return distribution study the ‘variance of the return’. Engel et al. (2005) learned the distribution of the return with Gaussian Processes, but did not use it for exploration. Tamar et al. (2016) studied the variance of the return with linear function approximation. Mannor & Tsitsiklis (2011) theoretically studies policies that bound the variance of the return.

The variance of the return has actually primarily been in the context of *risk-sensitive RL*. In several scenarios we may want to avoid incidental large negative pay-offs, which can e.g. be disastrous for a real-world robot, or in a financial portfolio. Morimura et al. (2012) studied parametric return distribution propagation as well. They do risk-sensitive exploration by softmax exploration over *quantile* Q-functions (also known as the *Value-at-Risk* (VaR) in financial management literature). Their distribution losses are based on KL-divergences (including Normal, Laplace and skewed Laplace distributions), but their implementations do remain in the tabular setting.

Bellemare et al. (2017) was the first to theoretically study the distributional Bellman operator, and also implement a distributional policy evaluation algorithm in the context of neural networks. Thereby, there work can be considered

the basis of our work, where we present an extension that uses the return distribution for exploration. Concurrently with our work, Tang & Agrawal (2018); Tang & Kucukelbir (2017) interpreted the return distribution from a variational perspective and leveraged it for exploration as well. Moerland et al. (2017a) also provided initial work on the return distribution for exploration. Our present paper is more extensive on the theoretical side, for example specifying full distributional loss functions and comparing different types of network output distributions. However, Moerland et al. (2017a) does try to connect the concept of return distribution to the statistical uncertainty of the mean action value as well, which both seem plausible quantities for exploration.

Another branch of related work is from the Tree Search community. Various papers have focussed on propagating distributions within the tree, e.g. Tesauro et al. (2012) and Kaufmann & Koolen (2017). The tree search approach by Moerland et al. (2018) does not explicitly propagate distributions (only σ -like estimates), but their idea (the remaining uncertainty should also incorporate the remaining uncertainty in the subtree below an action) is observable in the return-based exploration and learning visualizations of this paper as well.

Other Uncertainty-based exploration methods There exists a long history of work on the statistical uncertainty of the mean action value for exploration, in the context of function approximation for example by Osband et al. (2016), Gal et al. (2016) and more recently Azizzadenesheli et al. (2017), Henderson et al. (2017) and Jeong & Lee (2017). Moreover, the uncertainty theme for exploration also appears in count-based exploration approaches (Bellemare et al., 2016) and model-based RL (Guez et al., 2012; Moerland et al., 2017b).

C. Randomized Chain




Figure 4. Chain domain. Example MDP where undirected exploration is highly inefficient. Based on Osband et al. (2014).

We here present the randomized Chain, which we believe is the correct implementation of a well-known RL task known as the Chain (Osband et al., 2014) (Fig. 4). The domain illustrates the difficulty of exploration with sparse rewards. The MDP consists of a chain of states $\mathcal{S} = \{1, 2, \dots, N\}$. At each time step the agent has two available actions: a_1 (‘left’) and a_2 (‘right’). At every step, one of both actions is the ‘correct’ one, which deterministically moves the agent one step further in the chain. The wrong action terminates the

episode. All states have zero reward except the final chain state N , which has $r = 1$.

Variants of these problem have been studied more frequently in RL (Osband et al., 2014). In the ‘ordered’ implementation, the correct action is always the same (e.g. a_2), and the optimal policy is to always walk right. This is the variant illustrated in Fig. 4 as well. However, in our ‘randomized’ Chain implementation the correct action is randomly picked at domain initialization. The problem with the ordered version is that it introduced a systematic bias which is easily exploited when learning with neural networks. Due to the generalization of neural networks, it relatively easily predicts to always take action a_2 , and then suddenly solves the entire chain. With the randomized version, there is actually no structure in the domain at all, and learning with a neural network only makes the domain more complicated. The ‘randomized’ version therefore gives the true exponential complexity, as reported before (Osband et al., 2014; Moerland et al., 2017a), when learning with neural networks.

D. Implementation Details

Network architecture consists of a 3 layer neural network *per discrete action* with 256 nodes in each hidden layer and ELU activations. Learning rates were 0.0005 on all experiments. Optimization is performed with stochastic gradient descent on minibatches of size 32 using Adam updates in Tensorflow. We use a replay database of size 50,000. After collecting a new (set of) roll-outs, we randomly sample an equal amount of data from the replay for processing. All new collected data is processed on-policy, while all replay data is processed off-policy. The maximum length per episode is 200. We use discount factor $\gamma = 0.995$. All ϵ -greedy experiments have ϵ fixed at 0.05 throughout learning.

For the categorical outcome we put the bin edges slightly above and below the highest and lowest expected reward in the domain. In the chain we use $N = 7$ bins, on the other domains we use $N = 31$ bins. For the Gaussian output we add an initialization bias of 1 to the standard deviation at initialization. For the Gaussian mixture output we use $M = 5$ mixture components, where we spread out the mixture means upon initialization. Due to the logarithm appearing in the Gaussian cross-entropy loss we see that the gradient may explode when the standard deviation strongly narrows. We mitigate this problem by clipping gradients.

Thompson sampling is best implemented in an ‘episode-wise’ fashion, where we sample from a posterior distribution over parameters once at the beginning of a new episode (Russo et al., 2017). This ensures deep exploration. However, for the return based uncertainty we directly sample in the network output space per action, and we cannot implement this correlated form of Thompson sampling.

Full code is available from https://github.com/tmoer/return_distribution_exploration.git.

Double Uncertain Exploration

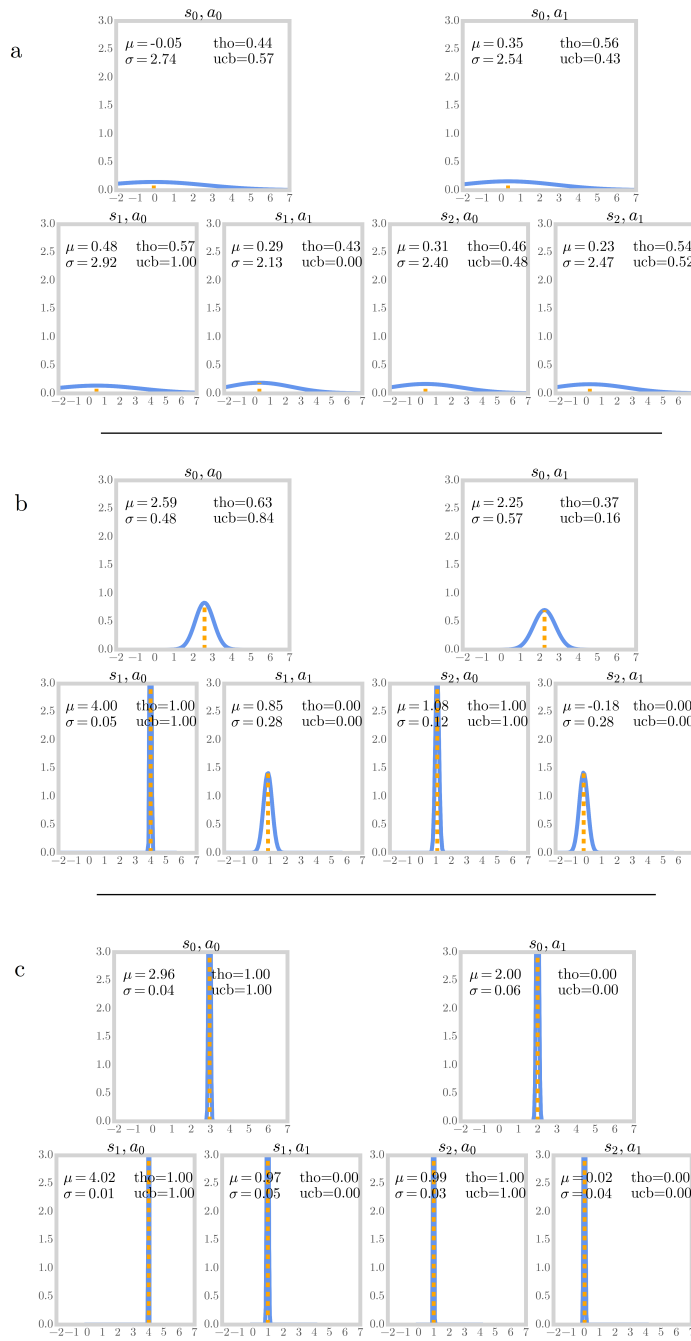


Figure 5. Example of Gaussian distribution propagation on the Toy example of Fig.1. **a)** Return distributions at initialization. Both Thompson sampling and UCB have largely uniform policies. **b)** Distributions after training for 64 episodes. The terminal state distributions start gradually converging, while the distributions at state s_0 remain broader. The terminal node decisions are already greedy, while the first node already starts to assign higher probability to the optimal action a_0 . **c)** Converged distributions after training for some additional time. Both UCB and Thompson sampling now deterministically sample the optimal policy.

Double Uncertain Exploration




Figure 6. Example of return distribution-based exploration on the Chain of length 7. Each plot (a-c) shows successive states (left-to-right) and both actions (up and down). The correct action at each step (randomly drawn at domain initialization) is indicated by a green box around the plot. We use a categorical $p(Z)$ with 7 atoms between -0.2 and 1.2 . **a).** Return distributions after 2 episodes. The distributions are almost uniform, which makes the policy fully exploratory. **b).** Return distributions after 28 episodes. Both the correct and wrong action have propagated mass towards 0. However, the distributions of the wrong actions converge faster, because the correct actions propagate the remaining uncertainty at the next timestep. The correct action in the last state already started to move towards a value of 1. **c).** Converged return distributions after 68 episodes. All the correct actions have now backpropagated the return from the end of the chain. The policy now consistently exploits. Note that some of the wrong actions (red boxes) put some mass at a return of 1 as well. This is due to the generalization from neighboring states (which are treated as continuous in the network input), but disappears with enough data.