LEAVE NO TRACE: LEARNING TO RESET FOR SAFE AND AUTONOMOUS REINFORCEMENT LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Deep reinforcement learning algorithms can learn complex behavioral skills, but real-world application of these methods requires a considerable amount of experience to be collected by the agent. In practical settings, such as robotics, this involves repeatedly attempting a task, resetting the environment between each attempt. However, not all tasks are easily or automatically reversible. In practice, this learning process requires considerable human intervention. In this work, we propose an autonomous method for safe and efficient reinforcement learning that simultaneously learns a forward and backward policy, with the backward policy resetting the environment for a subsequent attempt. By learning a value function for the backward policy, we can automatically determine when the forward policy is about to enter a non-reversible state, providing for uncertainty-aware safety aborts. Our experiments illustrate that proper use of the backward policy can greatly reduce the number of manual resets required to learn a task and can reduce the number of unsafe actions that lead to non-reversible states. [1]

1 INTRODUCTION

Deep reinforcement learning (RL) algorithms have the potential to automate acquisition of complex behaviors in a variety of real-world settings. Recent results have shown success on games (Mnih et al. (2013)), locomotion (Schulman et al. (2015)), and a variety of robotic manipulation skills (Pinto & Gupta (2017); Schulman et al. (2016); Gu et al. (2017)). However, the complexity of tasks achieved with deep RL in simulation still exceeds the complexity of the tasks learned in the real world. Why have real-world results lagged behind the simulated accomplishments of deep RL algorithms?

One challenge with real-world application of deep RL is the scaffolding required for learning: a bad policy, can easily put the system into an unrecoverable state from which no further learning is possible. For example, an autonomous car might collide at high speed, and a robot learning to clean glasses might break them. Even in cases where failures are not catastrophic, some degree of human intervention is often required to reset the environment between attempts (e.g. Chebotar et al. (2017)).

Most reinforcement learning algorithms require sampling from the initial state distribution at the start of each episode. On real-world tasks, this operation often corresponds to a human resetting the environment after every episode, an expensive solution for complex environments. Even when tasks are designed so that these resets are easy (e.g. Levine et al. (2016) and Gu et al. (2017)), manual resets are necessary when the robot or environment breaks (e.g. Gandhi et al. (2017)). The bottleneck for learning many real-world tasks is not that the agent collects data too slowly, but rather that data collection stops entirely when the agent is waiting for a manual reset. To avoid manual resets caused by the environment breaking, humans often add negative rewards to dangerous states and intervene to prevent agents from taking dangerous actions. While this works well for simple tasks, scaling to more complex environments requires writing large numbers of rules for types of actions the robot should avoid. For example, we might want a robot to avoid hitting itself, except when clapping. One interpretation of our method is as automatically learning these safety rules. Decreasing the

[1] Videos of our experiments: `https://sites.google.com/site/mlleavenotrace/`

number of manual resets required to learn to a task is important for scaling up reinforcement learning experiments outside simulation, allowing researchers to run longer experiments on more agents for more hours.

We propose to address these challenges by forcing our agent to "leave no trace." The goal is to learn not only how to do the task at hand, but also how to undo it. The intuition is that the sequences of actions that are reversible are safe; it is always possible to undo them to get back to the original state. This property is also desirable for continual learning of agents, as it removes the requirements for manual resets. In this work, we learn two policies that alternate between attempting the task and resetting the environment. By learning how to reset the environment at the end of each episode, the agent we learn requires significantly fewer manual resets. Critically, our value-based reset policy restricts the agent to only visit states from which it can return, intervening to prevent the forward policy from taking potentially irreversible actions. The set of states from which the agent knows how to return grows over time, allowing the agent to explore more parts of the environment as soon as it is safe to do so.

The main contribution of our work is a framework for continually and jointly learning a reset policy in concert with a forward task policy. We show that this reset policy not only automates resetting the environment between episodes, but also helps ensure safety by reducing how frequently the forward policy enters into unrecoverable states. Incorporating uncertainty into the value functions of both the forward and backward policy further allows us to make this process risk-aware, balancing exploration against safety. Our experiments illustrate that this approach reduces the number of "hard" manual resets required during learning of a variety of simulated robotic skills.

2 Related Work

Our works builds off of previous work in areas of safe exploration, multiple policies, and automatic curriculum generation. Previous work has examined safe exploration in small MDPs. Moldovan & Abbeel (2012a) examined risk-sensitive objectives for MDPs, and proposed a new objective of which minmax and expectation optimization are both special cases. Moldovan & Abbeel (2012b) consider safety using ergodicity, where an action is safe if it is still possible to reach every other state after having taken that action. These methods were limited to small, discrete MDPs where exact planning is straightforward. Our work includes a similar notion of safety, but can be applied to solve complex, high-dimensional tasks.

Previous work has also used multiple policies for safety and for learning complex tasks. Han et al. (2015) learn a sequence of forward and backward policies to complete a complex manipulation task. Similar to Han et al. (2015), our work learns a reset policy to undo the actions of the forward policy. While Han et al. (2015) engage the reset policy when the forward policy fails, we preemptively predict whether the forward policy will fail, and engage the reset policy before allowing the forward policy to fail. Similar to our approach, Richter & Roy (2017) also proposes to use a safety policy that can trigger an "abort" to prevent a dangerous situation. However, in contrast to our approach, Richter & Roy (2017) use a heuristic, hand-engineered reset policy, while our reset policies are learned simultaneously with the forward policy. Kahn et al. (2017) uses uncertainty estimation via bootstrap to provide for safety. Our approach also uses bootstrap for uncertainty estimation, but unlike our method, Kahn et al. (2017) does not learn a reset or safety policy.

Learning a reset policy is related to curriculum generation: the reset controller is engaged in increasingly distant states. Prior methods have studied curriculum generation by maintaining a separate goal setting policy or network (Sukhbaatar et al., 2017; Matiisen et al., 2017; Held et al., 2017). In contrast to these methods, we do not set explicit goals, but only allow the reset policy to abort an episode. When learning the forward and reset policies jointly, the training dynamics of our reset policy resemble those of reverse curriculum generation (Florensa et al., 2017), but in reverse. In particular, reverse curriculum learning can be viewed as a special case of our method: our reset policy is analogous to the learner in the reverse curriculum, while the forward policy plays a role similar to the initial state selector. However, reverse curriculum generation requires that the agent can be reset to any state (e.g., in a simulator), while our method is specifically aimed at streamlining real-world learning, through the use of uncertainty estimation and early aborts.

3 PRELIMINARIES

In this section, we discuss the episodic RL problem setup, which motivates our proposed joint learning of forward and reset policies. RL considers decision making problems that consist of a state space $\mathcal{S}$, action space $\mathcal{A}$, transition dynamics $P(s' \mid s, a)$, an initial state distribution $p_0(s)$, and a scalar reward function $r(s, a)$. In episodic, finite horizon tasks, the objective is to find the optimal policy $\pi^*(a \mid s)$ that maximizes the expected sum of γ-discounted returns, $\mathbb{E}_\pi[\sum_{t=0}^{T} \gamma^t r(s_t, a_t)]$, where $s_0 \sim p_0$, $a_t \sim \pi(a_t|s_t)$, and $s_{t+1} \sim P(s' \mid s, a)$.

Typically, the RL training routines involve iteratively sampling new episodes, where at the end of each episode, a new starting state s_0 is sampled from a given initial state distribution p_0. In practical applications, such as robotics, this procedure effectively involves executing some hard-coded reset policy or human interventions to manually reset the agent. Our work is aimed at avoiding these manual resets by learning an additional reset policy that satisfies the following property: when the reset policy is executed from any state, the distribution over final states matches the initial state distribution p_0. If we learn such a reset policy, then the agent never requires querying the black-box distribution p_0 and can continually learning on its own.

4 CONTINUAL LEARNING WITH JOINT FORWARD-RESET POLICIES

Our method for continual learning relies on jointly learning a *forward policy* and *reset policy*, and using *early aborts* to avoid manual resets. The forward policy aims to maximize the task reward, while the reset policy takes actions to reset the environment. Both have the same state and action spaces, but are given different reward objectives. The forward policy reward $r_f(s, a)$ is the usual task reward given by the environment. The reset policy reward $r_r(s)$ is designed to be large for states with high density under the initial state distribution. For example, in locomotion experiments, the reset reward is large when the agent is standing upright. To make this set-up applicable for solving the task, we make the weak assumption on the task environment that there exists a policy that can reset from at least one of the reachable states with maximum reward in the environment. This assumption ensures that it is possible to solve the task without any manual resets.

We choose off-policy actor-critic methods as the base RL algorithm (Silver et al., 2014; Lillicrap et al., 2015), since its off-policy learning allows sharing of the experience between the forward and reset policies. Additionally, the Q-functions can be used to inform the performance of each policy during exploration. Our methods can also be used directly with any other Q-learning methods ((Watkins & Dayan, 1992; Mnih et al., 2013; Gu et al., 2017; Amos et al., 2016; Metz et al., 2017)).

4.1 EARLY ABORTS

The reset policy learns how to transition the agent from a state reached by the forward policy back to an initial state. However, in challenging domains with irreversible states, the reset policy may be unable to reset from some states, and a costly "hard" reset may be required. The process of learning the reset policy offers us a natural mechanism for reducing these hard resets. We observe that, for states that are irrecoverable, the value function of the reset policy will be low. We can therefore use this value function (or, specifically, its Q-function) as a metric to determine when to terminate the forward policy, performing an *early abort*.

Before an action proposed by the forward policy is executed in the environment, it must be "approved" by the reset policy. In particular, if the reset policy's Q value for the proposed action is too small, then an early abort is performed: the proposed action is not taken and the reset policy takes control. Formally, early aborts restrict exploration to a 'safe' subspace of the MDP. Let $\mathcal{E} \subseteq \mathcal{S} \times \mathcal{A}$ be the set of (possible stochastic) transitions, and let $Q_{reset}(s, a)$ be the Q values of our reset policy at state s taking action a. The subset of transitions $\mathcal{E}^* \in \mathcal{E}$ allowed by our algorithm is

$$\mathcal{E}^* \triangleq \{(s, a) \in \mathcal{E} \mid Q_{reset}(s, a) > Q_{min}\} \tag{1}$$

Noting that $Q(s) \triangleq \max_{a \in \mathcal{A}} Q(s, a)$, we see that the states allowed under our algorithm $\mathcal{S}^* \subseteq \mathcal{S}$ are those states with at least one transition in $\mathcal{E}^*$:

$$\mathcal{S}^* \triangleq \{s \mid (s, a) \in \mathcal{E}^* \text{ for at least one } a \in \mathcal{A}\} \tag{2}$$

For intuition, consider tasks where the reset reward is 1 if the agent has successfully reset, and 0 otherwise. In these cases, the reset Q function indicates the probability that the reset will succeed. Early aborts occur when this probability for the forward policy's proposed action is too small. This can be interpreted as a learned, dynamic safety constraint, and a viable alternative for the manual constraints that are typically used for real-world robotic learning experiments. Early aborts promote safety by preventing the agent from taking actions from which it cannot recover. These aborts are dynamic because the states at which they occur change during throughout training, as more and more states are considered safe. This can make it easier to learn the forward policy, by preventing it from entering unsafe states. We analyze this experimentally in Section 6.

4.2 Hard Resets

Early aborts decrease the requirement for "hard" resets that move the agent back to the initial state, but do not eliminate them, since an imperfect reset policy might still miss a dangerous state early in the training process. However, it is challenging to identify whether it is possible for *any* policy to reset from the current state. Our approach is to approximate irreversible state identification with a necessary (but not sufficient) condition: we say we have reached an irreversible state if the reset policy fails to reset after N attempts, where N is a hyperparameter. Formally, we define a set of safe states $\mathcal{S}_{reset} \subseteq \mathcal{S}$, and say that we are in an irreversible state if the set of states visited by the reset policy over the past N episodes is disjoint from $\mathcal{S}_{reset}$. Increasing N decreases the number of hard resets. However, when we are in an irreversible state, increasing N means that we remain in that state (learning nothing) for more episodes. Section 6.3 empirically examines this trade-off. In practice, the setting of this parameter should depend on the cost of hard resets.

4.3 Value Function Ensembles

The accuracy of the Q-value estimates of our policies affect learning and reset performance through early aborts. Q-values of the reset policy may not be good estimates of the true value function for previously-unseen states. To address this, we train Q-functions for both the forward and reset policies that provide uncertainty estimates. Several prior works have explored how uncertainty estimates can be obtained in such settings (Gal & Ghahramani, 2016; Osband et al., 2016). We use the bootstrap ensemble in our method (Osband et al., 2016), though other techniques could be employed. In this approach, we train an ensemble of Q-functions, each with a different random initialization, which provides a distribution over Q-values at each state. Given this distribution over Q values, we can propose three strategies for early aborts:

> *Optimistic Aborts:* Perform an early abort only if all the q values are less than Q_{min}. Equivalently, do an early abort if $\max_\theta Q^\theta_{reset}(s, a) < Q_{min}$.
>
> *Realist Aborts:* Perform an early abort if the mean Q value is less than Q_{min}.
>
> *Pessimistic Aborts:* Perform an early abort if any of the Q values are less than Q_{min}. Equivalently, do an early abort if $\min_\theta Q^\theta_{reset}(s, a) < Q_{min}$.

We expect that optimistic aborts will provide better exploration at the cost of more hard resets, while pessimistic aborts should decrease hard resets, but may be unable to effectively explore. We empirically compare these strategies in Section 6.4.

4.4 Algorithm Summary

Our full algorithm (Algorithm 1) consists of alternately running a forward policy and reset policy. When running the forward policy, we perform an *early abort* if the Q value for the reset policy is less than Q_{min}. Only if the reset policy fails to reset after N episodes do we do a manual reset.

5 Small-Scale Didactic Example

We first present a small didactic example to illustrate how our forward and reset policies interact and how cautious exploration reduces the number of hard resets. We first discuss the gridworld in Figure 1(a). The states with red borders are absorbing, meaning that the agent cannot leave them and must use a hard reset. The agent receives a reward of 1 for reaching the goal state, and 0 otherwise.

Algorithm 1 Joint Training

repeat
  for $j \leftarrow 1 \cdots max_steps_per_episode$ **do**
    $a \leftarrow$ FORWARD_AGENT.ACT(s)
    if RESET_AGENT.Q(s, a) $< Q_{min}$ **then**
      Switch to reset policy. ▷ Early Abort
    $(s, r) \leftarrow$ ENVIRONMENT.STEP(a)
    Update the forward policy.
  for $i \leftarrow 1 \cdots max_steps_per_episode$ **do**
    $a \leftarrow$ RESET_AGENT.ACT(s)
    $(s, r) \leftarrow$ ENVIRONMENT.STEP(a)
    Update the reset policy.
  Let $\mathcal{S}^i_{reset}$ be the final states from the last N reset episodes.
  if $\mathcal{S}^i_{reset} \cap \mathcal{S}_{reset} = \emptyset$ **then**
    $s \leftarrow$ ENVIRONMENT.RESET() ▷ Hard Reset

(a) Early aborts for the gridworld

(b) Early aborts with an absorbing goal

Figure 1: *Didactic example:* The agent must reach a goal state without entering one of the red absorbing states. The background color indicates how frequently an early abort occurred in that state. Note the high frequency of early aborts next to the absorbing goal in (b).

The states are colored based on the number of early aborts triggered in each state. Note that most aborts occur next to the initial state, when the forward policy attempts to enter the top-left absorbing state but is blocked by the backward policy. In Figure 1(b), we present a harder example, where the task can be successfully completed by reaching one of the two goals, exactly one of which is reversible. The forward policy has no preference for which goal is better, but the backward policy successfully prevents the forward policy from entering the absorbing goal state, as indicated by the much larger early abort count in the blue-colored state next to the absorbing goal.

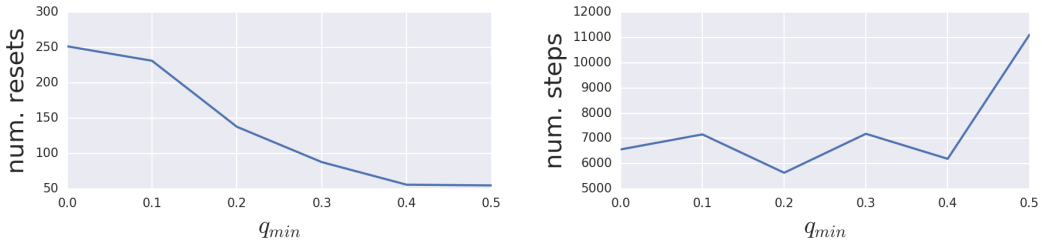

Figure 2: Early abort threshold: In our didactic example, increasing the early abort threshold causes more cautious exploration (left) without severely increasing the number of steps to solve (right).

Figure 2 shows how changing the early abort threshold to explore more cautiously reduces the number of failures. Increasing Q_{min} from 0 to 0.4 reduced the number of hard resets by a 78% without increasing the number of steps to solve the task. In a real-world setting, this might produce a substantial gain in efficiency, as time spend waiting for a hard reset could be better spent collecting

more experience. Thus, for some real-world experiments, increasing Q_{min} could decrease training time even if it requires more steps to learn.

6 Continuous Environment Experiments

In this section, we use three complex continuous control environments to answer questions about our approach. Figure 3 shows *pusher*, *cliff cheetah* and *cliff walker*. In each environemnt, agents can reach irreversible states: the pusher can knock the puck outside its workspace and the cheetah and walker and the cheetah can jump off a cliff. Crucially, reaching these irreversible states do not terminate the episode, so the agent remains in the irreversible state until it calls for a hard reset. Becaues of space constraints, plots for cliff walker, as well as full experimental details, are in the Appendix.

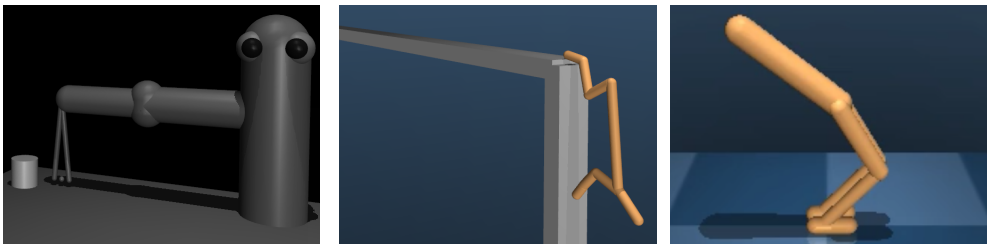

Pusher Cliff Cheetah Cliff Walker

Figure 3: *Continuous control environments*

6.1 Learning with Fewer Hard Resets

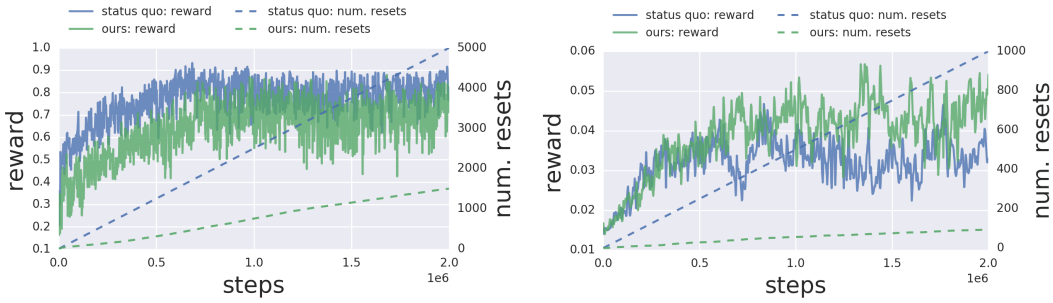

Pusher Cliff Cheetah

Figure 4: Few Hard Resets: We compare our method with the status quo on pusher (left) and cliff cheetah (right). The solid lines show the rewards, while the dotted lines show the number of resets, each over the sample steps. Our method achieves equal or better rewards with fewer manual resets.

Our first goal is to reduce the number of hard resets during learning. In this section, we compare our algorithm to the standard episodic learning setup with just a forward policy ("status quo"). As shown in Figure 4 (left), the conventional approach learns the pusher task somewhat faster than ours, but our approach eventually achieves the same reward threshold with only half the number of hard resets. In the cliff cheetah task, not only does our approach use an order of magnitude fewer hard resets, but the final reward of our method is substantially higher. This suggests that, besides reducing the number of resets, the early aborts can actually aid learning by preventing the forward policy from wasting exploration time, since the reset policy will trigger an abort when the forward policy is abort to enter an irreversible state (e.g. falling off the cliff).

6.2 Early Aborts avoid Hard Resets

Algorithm 1 was designed with the hypothesis that the early aborts would help avoid hard resets. To test whether early aborts prevent hard resets, we can see if the number of hard resets increases

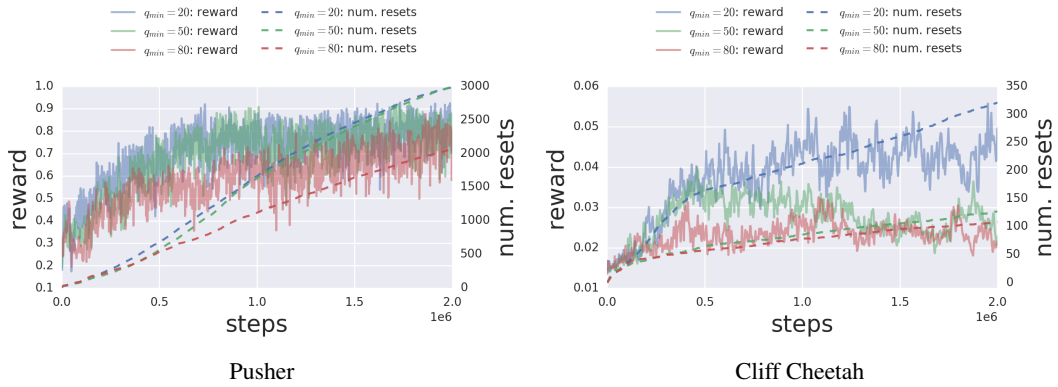

Pusher

Cliff Cheetah

Figure 5: Early Abort Threshold: Increasing the early abort threshold to act more cautiously helps avoid hard resets.

when we lower the early abort threshold. Figure 5 shows the effect of three values for Q_{min} while learning the pusher and cliff cheetah. In both environments, decreasing the early abort threshold increased the number of hard resets, supporting our hypothesis that early aborts prevent hard resets. On pusher, increasing Q_{min} to 80 allowed the agent to learn a policy that achieved nearly the same reward using 33% fewer hard resets. The cliff cheetah task has lower rewards than pusher, even an early abort threshold of 10 is enough to prevent 69% of the total early aborts that the status quo would have performed.

6.3 Multiple Reset Attempts

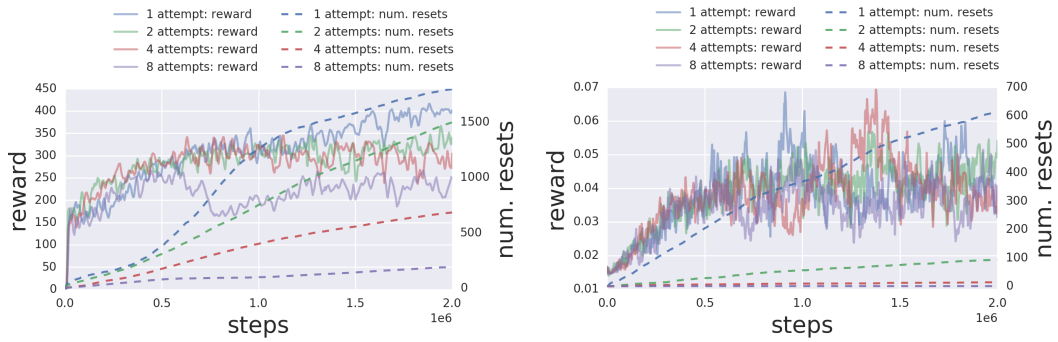

Pusher

Cliff Cheetah

Figure 6: Reset Attempts: Increasing the number of reset attempts reduces hard resets, but sometimes stalls learning.

While early aborts help avoid hard resets, our algorithm includes a mechanism for requesting a manual reset if the agent reaches an unresettable state. As described in Section 4.2, the reset agent is given a number of attempts at resetting; we only perform a hard reset if the reset agent fails to reset too many times.

Figure 6 illustrates the effect of the number of reset attempts on hard resets and reward. On the pusher task, when our algorithm was given a single reset attempt, it used 64% fewer hard resets than the status quo approach would have. Increasing the number of reset attempts to 4 resulted in another more than 2.5x reduction in hard resets, while decreasing the reward by less than 25%. Further increasing the number of reset attempts to 8 caused a larger drop in reward. On the cliff cheetah task, increasing the number of reset attempts brought the number of resets down to nearly 0, without changing the reward. Surprisingly, these results indicate that for some tasks, it is possible to learn an equally good policy with significantly fewer hard resets.

6.4 ENSEMBLES ARE SAFER

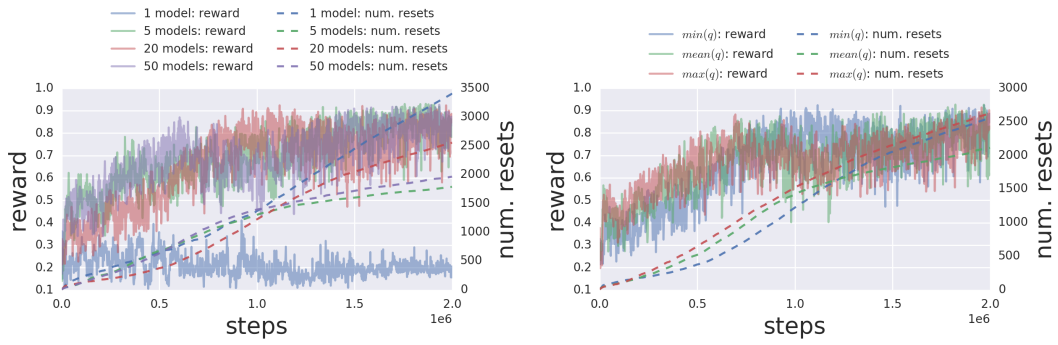

Pusher

Cliff Cheetah

Figure 7: Value function ensemble: (left) Increasing the ensemble size increases the policy reward and helps decrease the rate of hard resets throughout learning. (right) The method for combing value functions has little effect.

Our approach uses an ensemble of value functions to trigger early aborts. Our hypothesis was that our algorithm would be sensitive to bias in the value function if we used a single Q network. To test this hypothesis, we varied the ensemble size from 1 to 50. Figure 6 (left) shows the effect on learning the pushing task. An ensemble with one network failed to learn, but still required many hard resets. Increasing the ensemble size slightly decreased the number of hard resets without affecting the reward.

Interestingly, the particular method for combining the value functions (maximum, minimum, or mean, as discussed in Section 4.3), makes relatively little difference for the number of resets or final performance, as shown in Figure 7 (right). This suggests that much of the benefit of ensemble comes from its ability to produce less biased abort predictions in novel states, rather than the particular risk-sensitive rule that is used.

7 CONCLUSION

In this paper, we presented a framework for automating reinforcement learning based on two principles: automated resets between trials, and early aborts for the learner policy to avoid unrecoverable states. Our method simultaneously learns a forward and backward policy, with the value functions of the two policies used to balance exploration against recoverability. Experiments in this paper demonstrate that our algorithm not only reduces the number of manual resets required to learn a task, but also learns to avoid unsafe states.

While our algorithm can be applied to a wide range of tasks, it requires a small number of manual resets to learn some tasks. During the early stages of learning we cannot accurately predict the consequences of our actions. We cannot learn to avoid a dangerous state until we have visited that state (or a similar state) and experienced a manual reset. Second, we treat all manual resets as equally bad. In practice, some manual resets are worse than others. For example, it is worse for a grasping robot to break a wine glass than to push a block out of its workspace. An approach not studied in this paper for handling these cases would be to specify costs associated with each type of manual reset, and incorporate these reset costs into the learning algorithm.

While the experiments for this paper were done in simulation, where manual resets are inexpensive, the next step is to apply our algorithm to real robots, where manual resets are costly. A challenge introduced when switching to the real world is automatically identifying when the agent has reset. In simulation we can access the state of the environment directly to compute the distance between the current state and initial state. In the real world, we must infer states from noisey sensor observations to deduce if they are the same. If we cannot distinguish between the state where the forward policy started and the state where the reset policy ended, then we have succeeded in Leaving No Trace!

REFERENCES

Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. *arXiv preprint arXiv:1609.07152*, 2016.

Yevgen Chebotar, Mrinal Kalakrishnan, Ali Yahya, Adrian Li, Stefan Schaal, and Sergey Levine. Path integral guided policy search. *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3381–3388, 2017.

Carlos Florensa, David Held, Markus Wulfmeier, and Pieter Abbeel. Reverse curriculum generation for reinforcement learning. *arXiv preprint arXiv:1707.05300*, 2017.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059, 2016.

Dhiraj Gandhi, Lerrel Pinto, and Abhinav Gupta. Learning to fly by crashing. *arXiv preprint arXiv:1704.05588*, 2017.

Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 3389–3396. IEEE, 2017.

Weiqiao Han, Sergey Levine, and Pieter Abbeel. Learning compound multi-step controllers under unknown dynamics. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pp. 6435–6442. IEEE, 2015.

David Held, Xinyang Geng, Carlos Florensa, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. *arXiv preprint arXiv:1705.06366*, 2017.

Gregory Kahn, Adam Villaflor, Vitchyr Pong, Pieter Abbeel, and Sergey Levine. Uncertainty-aware reinforcement learning for collision avoidance. *arXiv preprint arXiv:1702.01182*, 2017.

Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, pp. 0278364917710318, 2016.

Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Tambet Matiisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher-student curriculum learning. *arXiv preprint arXiv:1707.00183*, 2017.

Luke Metz, Julian Ibarz, Navdeep Jaitly, and James Davidson. Discrete sequential prediction of continuous actions for deep rl. *arXiv preprint arXiv:1705.05035*, 2017.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

Teodor M Moldovan and Pieter Abbeel. Risk aversion in markov decision processes via near optimal chernoff bounds. In *Advances in neural information processing systems*, pp. 3131–3139, 2012a.

Teodor Mihai Moldovan and Pieter Abbeel. Safe exploration in markov decision processes. *arXiv preprint arXiv:1205.4810*, 2012b.

Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. In *Advances in Neural Information Processing Systems*, pp. 4026–4034, 2016.

Lerrel Pinto and Abhinav Gupta. Learning to push by grasping: Using multiple tasks for effective learning. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 2161–2168. IEEE, 2017.

Charles Richter and Nicholas Roy. Safe visual navigation via deep learning and novelty detection. In *Proc. of the Robotics: Science and Systems Conference*, 2017.

John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.

John Schulman, Jonathan Ho, Cameron Lee, and Pieter Abbeel. Learning from demonstrations through the use of non-rigid registration. In *Robotics Research*, pp. 339–354. Springer, 2016.

David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 387–395, 2014.

Sainbayar Sukhbaatar, Ilya Kostrikov, Arthur Szlam, and Rob Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play. *arXiv preprint arXiv:1703.05407*, 2017.

Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.

A ADDITIONAL FIGURES

For each experiment in the main paper, we choose one or two demonstrative environments. Below, we show all experiments run on all environments.

(a) Experiment in Figure 4 for cliff walker

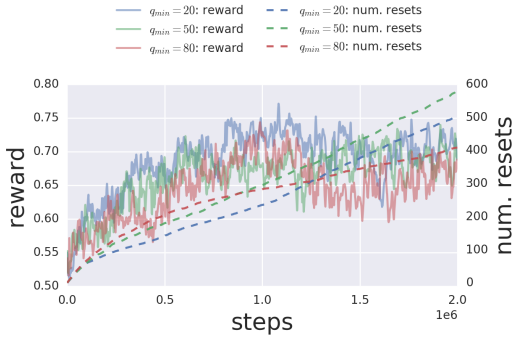

(b) Experiment in Figure 5 for cliff walker

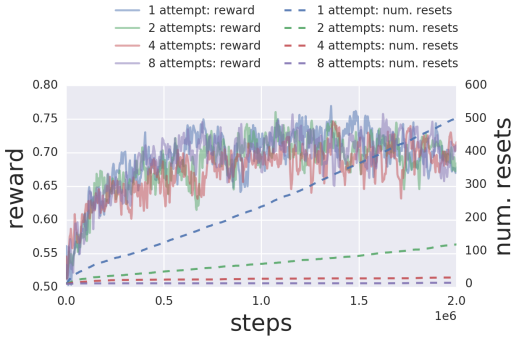

(a) Experiment in Figure 6 for cliff walker

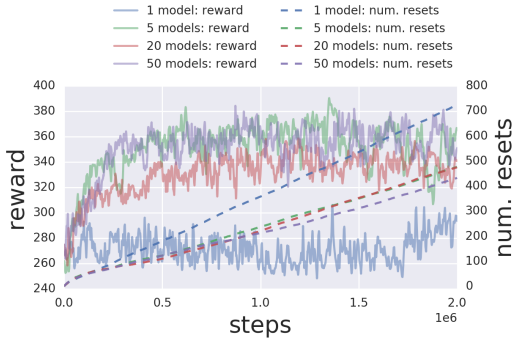

Figure 10: Experiment in Figure 7 for cliff walker

B EXPERIMENTAL DETAILS

B.1 GRIDWORLD ENVIRONMENT

To generate Figure 1, we averaged early abort counts across across 10 random seeds. For Figure 2 we took the median result across 10 random seeds. Both gridworld experiments used 5 models in the ensemble.

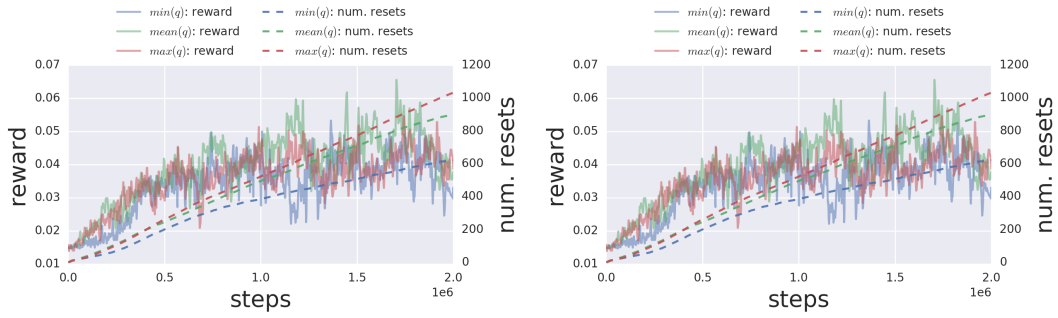

Figure 11: Experiment in Figure 7 for cliff cheetah

B.2 Continuous Control Environments

We did not do hyperparameter optimization for our experiments, but did run with multiple random seeds. To aggregate results, we took the median number across random seeds that solved the task. For most experiments, all random seeds solved the task.

We used the same DDPG hyperparameters for all continuous control environments:

> *Actor Network:* Two fully connected layers of sizes 400 and 300, with tanh nonlinearities throughout.
>
> *Critic Network:* We apply a 400-dimensional fully connected layer to states, then concatenate the actions and apply another 300-dimensional fully connected layer. Again, we use tanh nonlinearities.

Unless otherwise noted, experiments used an ensemble of size 20, $Q_{min} = 10$, 1 reset attempt, and early aborts using $min(q)$. The experiments in Section 6.1, our model used 2 reset attempts to better illustrate the potential for our approach to reduce hard resets.