

How Private Is Your RL Policy? An Inverse RL Based Analysis Framework

Kritika Prakash, Fiza Husain, Praveen Paruchuri, Sujit Gujar

Machine Learning Lab, IIIT Hyderabad

kritika.prakash@research.iiit.ac.in, fiza.husain@students.iiit.ac.in, praveen.p@iiit.ac.in, sujit.gujar@iiit.ac.in

Abstract

Reinforcement Learning (RL) enables agents to learn how to perform various tasks from scratch. In domains like autonomous driving, recommendation systems and more, optimal RL policies learned could cause a privacy breach if the policies memorize any part of the private reward. We study the set of existing differentially-private RL policies derived from various RL algorithms such as Value Iteration, Deep-Q Networks, and Vanilla Proximal Policy Optimization. We propose a new Privacy-Aware Inverse RL analysis framework (PRIL), that involves performing reward reconstruction as an adversarial attack on private policies that the agents may deploy. For this, we introduce the reward reconstruction attack, wherein we seek to reconstruct the original reward from a privacy-preserving policy using the Inverse RL algorithm. An adversary must do poorly at reconstructing the original reward function if the agent uses a tightly private policy. Using this framework, we empirically test the effectiveness of the privacy guarantee offered by the private algorithms on instances of the FrozenLake domain of varying complexities. Based on the analysis performed, we infer a gap between the current standard of privacy offered and the standard of privacy needed to protect reward functions in RL. We do so by quantifying the extent to which each private policy protects the reward function by measuring distances between the original and reconstructed rewards.

Introduction

Recent advancements in reinforcement learning (RL) have found widespread application in many real-world domains. Often, these domains are built from rich data sources or real-world environments, which could contain sensitive information of many individuals. This is evident in domains such as autonomous driving, recommendation systems, trading, industrial assembly, and domestic service robots. For example, a recommendation system agent for an online shopping platform not only tracks the purchases made, but also how long the user hovers over an item that he/she did not purchase. Another example is when an autonomous driving agent not only learns the dynamics behind driving, but also identifies people and predicts their behaviours on the streets, and how to respond to such situations. The reward function in these

environments is often sensitive, as it is built on people's private information. RL agents trained in such environments should not expose the private information of individuals. We therefore, need to use privacy-preserving methods to protect the rewards from being memorized in the agent's policy in such a manner that the agent's utility is not compromised.

Recent works use *Differential Privacy* (DP) (Dwork et al. 2006) to make the agent's policy quantifiably private and to get a rigorous privacy guarantee. Like many other areas in AI, RL has also started adopting DP to establish a mathematical way of guaranteeing data privacy in RL environments. However, an important question is: does the privacy guarantee offered by the private policies translate well into protecting the reward function? If not, how can we understand the gap in achieving meaningful privacy? Does the type of reward function, algorithms (both training and DP algorithms), and environment play a role? In other words, does privacy in policy translate well to privacy in reward? To investigate this, we first build an autonomous agent whose aim is to learn a private policy using existing privacy techniques with the intent of reaching a goal state that helps to maximize its expected reward in discrete finite-state environments. We then investigate the true level of reward privacy offered by the existing state-of-the-art privacy techniques for RL algorithms.

We evaluate the private policies by estimating an adversary's ability to reconstruct the original reward function from the agent's learned policy. The field of inverse RL (Ng, Russell et al. 2000), arose to solve this exact problem of extracting a reward function given the observed, optimal behaviour. Given the "reverse engineering" nature of inverse RL, the reconstructed reward can be used as an adversarial attack on environments with protected rewards. Building on this key intuition, we propose the PRIL analysis framework that first performs the reward reconstruction attack, and then computes its similarity to the original private reward via various distance metrics. We apply this framework over a set of privacy preserving techniques:

1. Bellman update DP
2. Rényi-DP in deep learning (DL)
3. Functional noise DP

These are applied to a set of RL algorithms:

1. Deep Q Network (DQN)

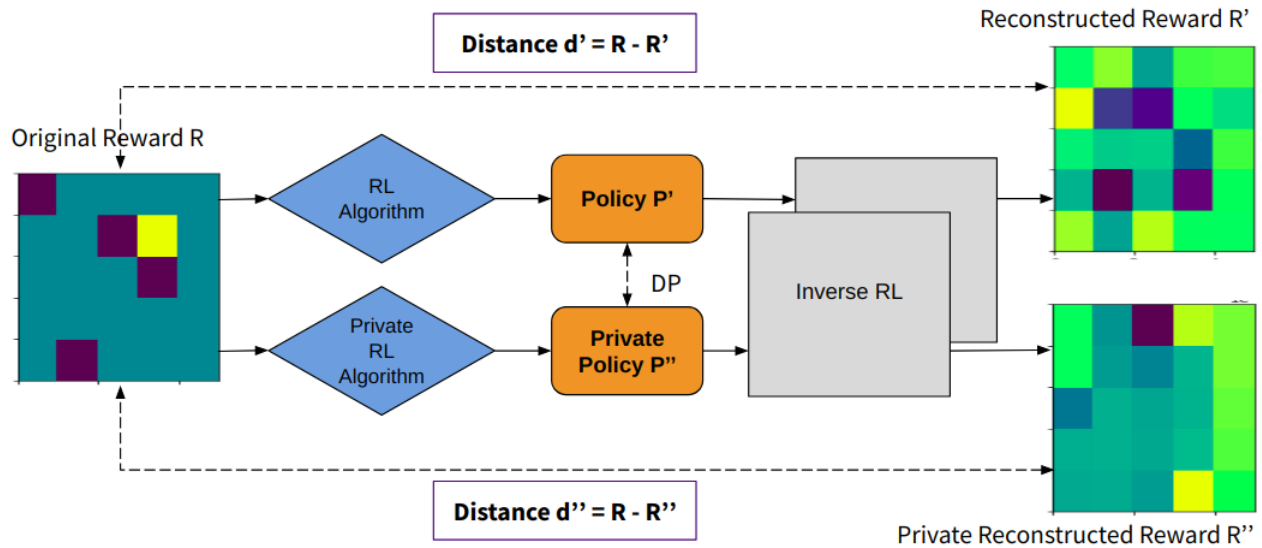


Figure 1: PRIL: Privacy-Aware Inverse RL analysis framework

2. Vanilla Proximal Policy Optimization (PPO)
3. Value Iteration (VI)

We build privacy into the agent from multiple perspectives - from a *DL perspective*, an *RL perspective*, and a *deep RL perspective*. Through experiments, we show that there exists a gap between the privacy offered via the current private RL methods, and the privacy needed to protect reward functions. We also present the privacy-utility trade-off achieved by each policy. We show that privacy in policy does not translate to privacy in reward, as the reconstruction error is independent of the DP budget. DP policies are unsuccessful at protecting the sensitive reward function due to a privacy gap. Our experiments demonstrate that there is a need to further inspect the effectiveness of DP policies to protect sensitive reward functions. It is a serious threat if an adversary is able to infer the rewards in spite of using a private policy.

In summary, our key contributions are:

1. We study and analyze the existing set of privacy techniques for RL.
2. We introduce a novel reward reconstruction attack and supporting PRIL framework.
3. We empirically evaluate the performance of various private deep RL policies within our framework.
4. We identify and quantify the gap between the privacy offered in policy and the privacy needed in reward.

The next section reviews related work, followed by a section that provides background on RL, DP, and inverse RL. We then explain the PRIL framework, followed by describing the experimental pipeline and setup. Then, we provide an empirical analysis and discussion. Finally, we conclude the work. All source code and experiments are made publicly available¹. Longer version of our paper with more detailed experiments and proofs is made available on ArXiv²

¹Link to code: <https://github.com/magnetar-iiith/PRIL>

²Link to ArXiv paper: <https://arxiv.org/abs/2112.05495>

(Prakash et al. 2021).

Related Work

Previous works on privacy-preserving RL make the use of DP. (Vietri et al. 2020) shows how to achieve joint-DP in episodic-RL via the upper confidence bound and Q-learning algorithms, where each training episode comes from a different environment. (Wang and Hegde 2019) makes the use of functional noise to make Q-learning private. (Vietri et al. 2020) gives us probably-approximately correct (PAC) and regret guarantees for private RL. (Balle, Gomrokchi, and Precup 2016) propose differentially private policy evaluation using the Monte-Carlo algorithm. (Hannun et al. 2019) introduces a private way of performing multi-party contextual bandits. For the actor-critic class of algorithms, (Lebensold et al. 2019) presents a differentially private critic, and (Seo and Yang 2020) presents a differentially private actor.

In RL, a significant amount of work has been done to perform adversarial attacks that target the quality of the learned policy (Gleave et al. 2019), (Huang et al. 2017), (Kos and Song 2017), (Chen et al. 2019). Building on this, there has been work on making RL policies robust to such attacks (Oikarinen, Weng, and Daniel 2020). However, very few works have looked into privacy attacks in RL. (Pan et al. 2019) shows that agents can memorize the environment and its private transition dynamics, by performing privacy attacks using genetic algorithms and candidate inference. (Fu, Luo, and Levine 2017) introduce an adversarial inverse RL algorithm based on an adversarial reward learning formulation to improve robustness.

While these previous works can tackle building privacy in policies, they do not investigate its impact on the underlying private data i.e., the reward function used to learn the policies. Many adversarial attacks such as the membership inference attack, the linkage attack, and the data-reconstruction attack have been used to evaluate the level of privacy attain-

Acronym List	
Abbreviation	Full Form
DP	Differential Privacy
RL	Reinforcement Learning
MDP	Markov Decision Process
DL	Deep Learning
LP	Linear Programming
RDP	Rènyi Differential Privacy
PAC	Probably Approximately Correct
PRIL	Privacy-Aware Inverse RL
VI	Value Iteration
DQN	Deep Q Network
PPO	Proximal Policy Optimization
DP-Bellman	Private Bellman update
DP-SGD	Private SGD optimizer + ReLU
DP-Shoe	Private SGD optimizer + tan-h
DP-Adam	Private Adam optimizer + ReLU
DP-FN	Private Functional Noise engine
VI-DP-Bellman	VI + DP-Bellman
DQN-DP-SGD	DQN + DP-SGD
DQN-DP-Shoe	DQN + DP-Shoe
DQN-DP-Adam	DQN + DP-Adam
DQN-DP-FN	DQN + DP-FN
PPO-DP-SGD	PPO + DP-SGD actor
PPO-DP-Shoe	PPO + DP-Shoe actor
PPO-DP-Adam	PPO + DP-Adam actor

Table 1: List of acronyms used

able by a data-analysis system. We introduce a new privacy attack that targets the private reward function. Our proposed attack - the reward-reconstruction attack is a special case of the data-reconstruction attack. We use inverse RL to learn the reward, and to assess the quality of private RL algorithms.

Background and Preliminaries

In this section, we introduce the basics of RL, inverse RL, and DP.

Reinforcement Learning

We focus on non-deterministic environments with discrete and finite - state and action spaces. Let $M = (S, A, P, R, \gamma, S_0)$ represent the MDP environment. Here, S is the set of finite discrete states, A is the set of finite actions, $P(s, a, s')$ is the transition probability of reaching state s' by taking an action $a \in A$ in state s , where $s, s' \in S$. $R(s)$ is the reward the agent receives in state $s \in S$, γ is the discount factor for future rewards, and S_0 is the initial state distribution over S . State value $V(s)$ is the value of expected return (sum of future discounted rewards) starting with state s . State-Action value $Q(s, a)$ denotes the value of taking an action in a state (following a policy π). The goal of an RL agent is to learn a policy that maximizes the expected cumulative reward. We use the following classes of RL algorithms to learn the optimal policy and value function for our experiments:

1. VI (Pashenkova, Rish, and Dechter 1996): VI computes an optimal state value function for an MDP. The method uses Bellman updates to converge to the optimal values.
2. DQN (Mnih et al. 2013): DQN is a model-free off-policy deep RL approach that uses a neural network approximator for the Q-function which uses a batch of past experiences (replay memory) to train the agent to learn the optimal policy.
3. PPO (Schulman et al. 2017): PPO is a first-order optimization based policy-gradient algorithm that uses the actor-critic approach to find the best policy. In this approach, the actor model learns to take an action in an observed state by improving upon the feedback given by the critic model - that takes state as an input, and finds a value function estimating future rewards with the help of an optimizer.

Inverse RL

Inverse RL (Ng, Russell et al. 2000) is a method of extracting a reward function, given the observed, optimal behaviour in an environment. We use the method of inverse RL in finite-state spaces to reconstruct the private reward function by solving a linear programming (LP) (Ignizio and Cavalier 1994) formulation that makes the given policy optimal by a large margin (as compared to other sub-optimal policies). Since this is an under-constrained problem, we choose the reward with the smallest L_1 -norm.

Differential Privacy

DP (Dwork et al. 2006) is considered to be the golden standard of computational privacy guarantees. It allows us to quantify the degree of privacy achievable by a mechanism. It is built on the concept of adjacent databases. In the context of our work, the RL agents learn optimal policies by exploring the environment and taking in rewards as a feedback for their actions. Since we care about the privacy of the reward signals, we say that two reward signals are adjacent if the maximum L_2 norm of their point-wise difference is upper bounded by 1.

Definition 1 (ϵ, δ) -DP: A randomized mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ with domain \mathcal{D} and range \mathcal{R} satisfies (ϵ, δ) -differential privacy if for any two adjacent inputs $d, d' \in \mathcal{D}$ and for any subset of outputs $S \subseteq \mathcal{R}$ it holds that

$$Pr[\mathcal{M}(d) \in S] \leq e^\epsilon Pr[\mathcal{M}(d') \in S] + \delta$$

Definition 2 α -Rènyi Divergence: For two probability distributions P and Q defined over \mathbf{R} , the Rènyi divergence of order $\alpha > 1$ is

$$D_\alpha(P||Q) = \frac{1}{\alpha - 1} \log_e E_{x \sim Q} \left(\frac{P(x)}{Q(x)} \right)^\alpha$$

where $P(x)$ and $Q(x)$ are the respective probability densities of P and Q at x .

An algorithm is said to have (α, ϵ) Rènyi DP (Mironov 2017) if for any two neighbouring databases, it holds that the Rènyi divergence of order α between outputs of the algorithm is less than e^ϵ .



Figure 2: All 24 FrozenLake environments used. Here, green: goal (G), red: frozen (F), yellow: start state, white: safe (S), and blue: high-reward (A).

Definition 3 (α, ϵ) -Rényi DP: A randomized mechanism $f : \mathcal{D} \rightarrow \mathcal{R}$ is said to have (ϵ) -Rényi differential privacy of order α , or (α, ϵ) -RDP for short, if for any adjacent $d, d' \in \mathcal{D}$ it holds that

$$D_\alpha(f(d) || f(d')) \leq e^\epsilon$$

Private RL Methods

We use the following private RL methods in our experiments:

1. *Bellman update DP*: In this method, noise is added locally to the Bellman update step of VI, such that it satisfies the definition of ϵ -DP (Dwork et al. 2006).
2. *Rényi-DP in DL*: This is a natural relaxation of DP, that we use for multiple DP methods (DQN-DP-SGD, PPO-DP-SGD, and more) in our work (Abadi et al. 2016), (Papernot et al. 2019).
3. *Functional noise DP in Q-Learning*: In this, functional noise is iteratively added to the value function in the training process. The aim is to protect the value function (Wang and Hegde 2019).

We share the entire training process - including every loss gradient update publicly (worst-case guarantee).

PRIL: Privacy-Aware Inverse RL Analysis Framework

We introduce a novel case of the data-reconstruction attack - the reward reconstruction attack for RL, as we wish to protect the reward function from adversaries. We assume that the adversary has knowledge of the environment and the learned private policy. Using this information, the adversary

tries to reverse engineer the reward function. While many methods can be used to do so, we focus on the LP based inverse RL technique in this paper, as it seems to be the best tool at our disposal. Using inverse RL, we perform the reward reconstruction attack, to determine how effective a private policy is at protecting the reward function. It does so by computing (a variety of) distances between the reconstructed reward and the original reward. The framework takes as input the original reward function R , an RL policy P' , and a private RL policy P'' trained using the same algorithm. Using the inverse RL algorithm, it predicts the reconstructed rewards, R' and R'' , from P' and P'' respectively. It then computes the distances $d'(R', R)$ and $d''(R'', R)$, and compares them. The larger the distance, the stronger the RL policy's privacy guarantee (in protecting the reward function). We use multiple distance metrics, such as - L_1 norm, L_2 norm, L_∞ norm, and number of sign changes.

Experimental Setup

We will now discuss our overall experimental pipeline and setup. We perform our experiments on 24 custom environments (as shown in Figure 2) in the *FrozenLake* domain - a discrete-state OpenAI Gym (Brockman et al. 2016) toolkit. In all these environments, the agent controls its movement and navigates in a grid-world. Additionally, the movement direction of the agent is uncertain and is only partially dependent on the direction chosen. The agent is rewarded for finding the most rewarding walkable path to the goal state. The grid-world environment has five possible states - safe (S), frozen (F), hole (H), high-reward (A) and goal (G). The agent has four possible actions - up, down, left and right. Half of the 24 environments are of a grid-size 5x5, and the




Figure 3: Original and reconstructed reward heatmaps for two 5x5 FrozenLake environments (1a, 2a) and for two 10x10 FrozenLake environments (3a, 4a)

remaining half are of a grid-size 10x10. The agent moves around the grid until it reaches the goal state. If it falls into the hole, it has to start from the beginning and is penalized. The process continues until it eventually reaches the goal state. We measure the performance of 8 private algorithms across three algorithm classes - VI, DQN, and PPO:

For VI, we evaluate the performance of VI-DP-Bellman (private Bellman update via local DP) as well as non-private VI. For DQN, we evaluate the performance of the following cases (with and without privacy):

1. *DQN-DP-SGD*: DP-SGD optimizer + ReLU activations
2. *DQN-DP-Adam*: DP-Adam optimizer + ReLU activations
3. *DQN-DP-Shoe*: DP-SGD optimizer + tan-h activations
4. *DQN-DP-FN*: DQN + functional noise

For PPO, we evaluate the performance of the following cases (with and without privacy in the actor network):

1. *PPO-DP-SGD*: DP-SGD optimizer + ReLU activations
2. *PPO-DP-Adam*: DP-Adam optimizer + ReLU activations
3. *PPO-DP-Shoe*: DP-SGD optimizer + tan-h activations

All the experiments were performed across the following set of ϵ privacy budgets: 0.1, 0.105, 0.2, 0.5, 1.0, 2.0, 5.0, 10.0, and ∞ (no privacy).

We evaluate the *privacy-utility trade-off* by simultaneously measuring the average returns of the the private policy over multiple sample trajectories during test time (as shown in Figure 5). For each private RL algorithm, we consider the non-private version of the RL policy (with an indefinite privacy budget) as the baseline for reward reconstruction. The more private the policy is, the larger the reward distance (between the original reward and the reconstructed reward).

Reward distance as a measure of privacy guarantee: We used four reward distance metrics for our experiments - L_1 distance, L_2 distance, L_∞ distance, and the sign change

count. The idea is to measure the similarity between the original reward function and the recovered reward function. This is a measure of the degree of privacy of a policy - the larger the reward distance, the more private the policy is. The metrics are calculated as follows:

- L_1 distance: Normalize the rewards R, R' using L_1 -norm, and then take the L_1 distance across the 2 vectors.
- L_2 distance: Normalize the rewards R, R' using L_2 -norm, and then take the L_2 distance across the 2 vectors.
- L_∞ distance: Normalize the rewards R, R' using L_∞ -norm, and then take the L_∞ distance across the 2 vectors.
- Sign change count: Measure the number of sign changes from R to R' .

Since each distance metric is in a different space, all the distances evaluated together allow us to get a deeper insight into the reward reconstruction mechanism, and the optimality and privacy of policies.

Policy return as a measure of agent utility: We measure how much utility the learned private policies achieve by observing how they perform during test-time, by calculating the average discounted returns over multiple trajectories played by the agent following the policy.

We build 24 custom FrozenLake environments using the Open AI gym toolkit. We use LP solvers to help solve the objective functions of Inverse RL. We build the Deep RL experiments using TensorFlow 2.4, and add privacy using TensorFlow Privacy. We build VI-DP-Bellman private algorithm from scratch, and use the publicly available code provided for the DQN-DP-FN strategy (Wang and Hegde 2019). We use a reference implementation (Alger 2016) for finite-state space inverse RL that makes the use of cvxopt (Andersen et al. 2013) to solve the LP formulation. We use Linux OS based servers for training all the RL agents with a total of 8 GPUs and 8 CPUs. All experiments spanned across 9 privacy budgets, 24 environments, 8 algorithms, repeating each experiment 10 times (to account for the randomness

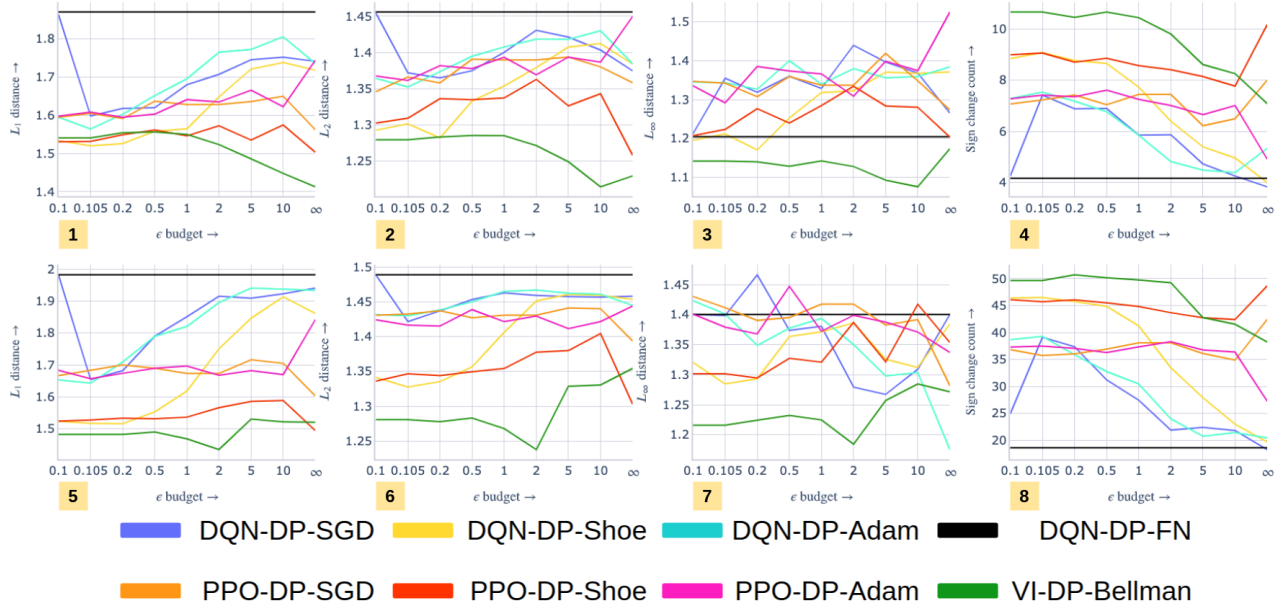


Figure 4: Reward distance vs privacy budget graphs for all strategies: 1,5: L_1 distance, 2,6: L_2 distance, 3,7: L_∞ distance, 4,8: Sign change counts. 1,2,3,4: averaged over 5x5 grid sized environments, 5,6,7,8: averaged over 10x10 grid sized environments

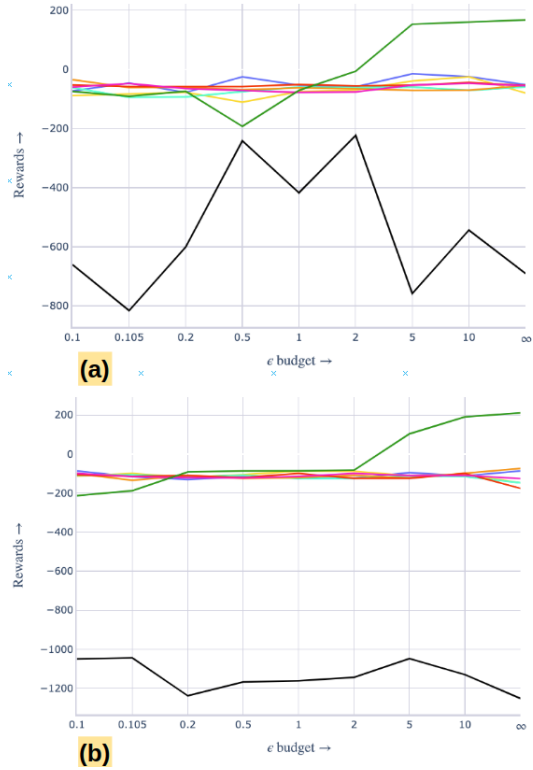


Figure 5: Utility (test-time return) vs privacy trade-off for all policies averaged across grid-sizes 5x5 (a) and 10x10 (b). Legend is the same as that in figure 4.

stemming from private noise mechanisms and DL optimization). The total runtime for the entire set of experiments was 3 weeks. In the longer version of the paper on ArXiv (link included), we provide a detailed record of all DL and RL hyperparameters (such as learning rate, discount factor, convergence criteria, and more). We also list the l_2 sensitivities, and Gaussian noise standard deviations used corresponding to each privacy budget ϵ , along with the formal privacy guarantees offered.

Analysis and Discussion

Figure 4 presents the variation in reward distances (y-axis) (of each type: L_1 , L_2 , L_∞ , and sign change counts) with an increase in the ϵ privacy budget (x-axis) (9 discrete values) including the no-privacy case at the very end ($\epsilon = \infty$). The first row shows results averaged over the 12 Frozen-Lake environments of grid size 5x5, whereas the second row shows the results corresponding to the environments of grid size 10x10. Each graph shows this relationship for all 8 private algorithms: DQN-DP-SGD, DQN-DP-Shoe, DQN-DP-Adam, DQN-DP-FN, PPO-DP-SGD, PPO-DP-Shoe, PPO-DP-Adam, and VI-DP-Bellman. The graphs show that there is no clear indication of any private strategy improving at reconstructing the reward (w.r.t. all distances) with a relaxation in the privacy budget - thus, rendering all strategies ineffective at being a truly meaningful private strategy. We observe the same lack of trend across both rows: for the 5x5 results in row 1 and 10x10 results in row 2.

Figure 5 presents the trade-off between the amount of utility (expected return: y-axis) and the degree of privacy (ϵ budget: x-axis) achieved by a private RL privacy. Graph 1 gives us the average trade-off for 5x5 environments, and graph 2 - for the 10x10 environments. Almost all algorithms exhibit




Figure 6: Aggregated L_2 distances across all environments and policy variants of the three main classes of RL algorithms: DQN, PPO, and VI.

comparable performance - with the exception of DQN-DP-FN, which performs significantly worse.

Figure 3 shows the heatmaps and reward structures of 4 different MDPs (row 1) and their corresponding reconstructed rewards (row 2) using the VI-DP-Bellman algorithm. From 1a and 1b we can observe that the agent is able to clearly detect the obstacle bar on the right. From 2a and 2b we can infer that the agent finds a straight line path along the edges that is rewarding. But it also wrongly identifies some states as rewarding, even though they might be dead-ends (bright yellow bottom-right corner). Perhaps, this could be because the agent is able to survive the cost of the state directly above it, if it means that the agent can easily reach the goal state in a short amount of time. In 3a and 3b we learn that even in such a rich environment, the agent is able to reconstruct the reward structure from an implementation of the policy. It can clearly identify rewarding blocks within the entire maze. And finally, in 4a and 4b, the agent learns to stick to the edges to maximize reward early-on - which is why its evaluation of the central region is poor, and privacy preserving. Figure 6 shows that DQN algorithms give us the strongest reward privacy, followed by PPO and VI algorithms - that do a very poor job at protecting the reward - despite having very similar utilities (from figure 5).

Based on the experiments performed, we can say that there is a considerable gap between the privacy provided by the existing private methods, and the level of privacy needed to protect the reward function from the inverse RL attack. We can also infer that techniques using Deep RL methods are able to learn the policy in a more general manner, as compared to non-deep methods. We address the need for better privacy techniques for RL algorithms that can effectively protect the reward function. We hope that our work inspires a deeper theoretical understanding of the limits to minimizing the gap, as well as its consequences in real-world applications. Besides thoroughly testing our code, we perform an extensive span of experiments. Contrasting our results with the baseline (no privacy), we find the reward distances to be quite similar. We therefore, believe that the source of the pri-

vacy gap is not experimental error. Our survey of papers that experiment on FrozenLake shows that the commonly used grid sizes are $\{4 \times 4, 8 \times 8\}$ while we experimented with slightly larger grid sizes - $\{5 \times 5, 10 \times 10\}$. We expect to observe a similar (or worse) privacy gap upon further increase of grid size since the reward would be richer in information, and the DP sensitivity is independent of the grid size.

While we demonstrate our work on a grid-world domain, we believe it is extendible to real-world domains with sensitive data. Our work is the first in this direction and serves as evidence that there is a need to inspect further. Deep RL is increasingly being used for recommendation systems in dynamic environments. Consider the case when the recommendation engine for every user is a unique private RL policy whose job is to recommend items to users and learn their preferences in an online fashion (given the user’s historical data). The reward is the user’s feedback (ratings) to the recommended action. While the policy provides privacy guarantees for its training process, it can leak the user’s feedback when subject to the re-identification attack via reward reconstruction. PRIL can help assess this threat better.

We surveyed a range of Inverse RL (IRL) algorithms: finite state space LP, sample trajectories ((Ng, Russell et al. 2000)), deep IRL ((Wulfmeier, Ondruska, and Posner 2015a)), and maximum entropy IRL ((Ziebart et al. 2008), (Wulfmeier, Ondruska, and Posner 2015b)). Despite starting with the simplest case - LP for finite state spaces, we observe a significant privacy gap. The LP method acts as a baseline for other IRL methods. With increased complexity, the reward function would be represented parametrically which would allow the system to evaluate performance on much larger and richer (and maybe continuous) environments. As the performance of IRL attack improves, we expect the issue of privacy gap to become even more important to address.

Our work introduces a novel direction of evaluating the privacy guarantees of RL systems. In the future, we hope to build on our work in multiple ways: extending to the multi-agent scenario, extending to a diverse set of domains, assessing the effect of generalization and exploration on privacy in RL, testing the performance of other RL algorithms such as PPO-Clip and PPO-KL, and evaluating the effect of using other complex inverse RL algorithms such as (Wulfmeier, Ondruska, and Posner 2015a), (Ziebart et al. 2008) and (Wulfmeier, Ondruska, and Posner 2015b).

Conclusion

This paper introduces a new Privacy-Aware Inverse RL analysis framework (PRIL) for enhancing reward privacy in reinforcement learning (RL) that performs a novel reward reconstruction attack and demonstrated its ability to fairly assess the level of privacy achieved in protecting the reward structure from adversarial attacks. We studied the set of existing privacy techniques for RL, performed a detailed evaluation of their effectiveness and identified that there is a significant gap between the current standard of privacy offered and the standard of privacy needed to protect reward functions in RL. We quantify this gap by measuring distances between the original and reconstructed rewards via the reward reconstruction attack.

References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 308–318.
- Alger, M. 2016. Inverse Reinforcement Learning. <https://zenodo.org/record/555999#.YkwFGnUzaV6>. Accessed: 2022-04-05.
- Andersen, M. S.; Dahl, J.; Vandenberghe, L.; et al. 2013. CVXOPT: A Python package for convex optimization. Available at cvxopt.org, 54.
- Balle, B.; Gomrokchi, M.; and Precup, D. 2016. Differentially private policy evaluation. In *International Conference on Machine Learning*, 2130–2138. PMLR.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym. *arXiv preprint arXiv:1606.01540*.
- Chen, T.; Liu, J.; Xiang, Y.; Niu, W.; Tong, E.; and Han, Z. 2019. Adversarial attack and defense in reinforcement learning-from AI security view. *Cybersecurity*, 2(1): 1–22.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In Halevi, S.; and Rabin, T., eds., *Theory of Cryptography*, 265–284. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-540-32732-5.
- Fu, J.; Luo, K.; and Levine, S. 2017. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*.
- Gleave, A.; Dennis, M.; Wild, C.; Kant, N.; Levine, S.; and Russell, S. 2019. Adversarial policies: Attacking deep reinforcement learning. *arXiv preprint arXiv:1905.10615*.
- Hannun, A.; Knott, B.; Sengupta, S.; and van der Maaten, L. 2019. Privacy-Preserving Multi-Party Contextual Bandits. *arXiv preprint arXiv:1910.05299*.
- Huang, S.; Papernot, N.; Goodfellow, I.; Duan, Y.; and Abbeel, P. 2017. Adversarial attacks on neural network policies. *arXiv preprint arXiv:1702.02284*.
- Ignizio, J. P.; and Cavalier, T. M. 1994. *Linear programming*. Prentice-Hall, Inc.
- Kos, J.; and Song, D. 2017. Delving into adversarial attacks on deep policies. *arXiv preprint arXiv:1705.06452*.
- Lebensold, J.; Hamilton, W.; Balle, B.; and Precup, D. 2019. Actor Critic with Differentially Private Critic. *arXiv preprint arXiv:1910.05876*.
- Mironov, I. 2017. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, 263–275. IEEE.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Ng, A. Y.; Russell, S. J.; et al. 2000. Algorithms for inverse reinforcement learning. In *ICML*, volume 1, 2.
- Oikarinen, T.; Weng, T.-W.; and Daniel, L. 2020. Robust Deep Reinforcement Learning through Adversarial Loss. *arXiv preprint arXiv:2008.01976*.
- Pan, X.; Wang, W.; Zhang, X.; Li, B.; Yi, J.; and Song, D. 2019. How you act tells a lot: Privacy-leakage attack on deep reinforcement learning. *arXiv preprint arXiv:1904.11082*.
- Papernot, N.; Chien, S.; Song, S.; Thakurta, A.; and Erlingsson, U. 2019. Making the shoe fit: Architectures, initializations, and tuning for learning with privacy. <https://openreview.net/pdf?id=rJg851rYwH>. Accessed: 2021-09-09.
- Pashenkova, E.; Rish, I.; and Dechter, R. 1996. Value iteration and policy iteration algorithms for Markov decision problem. In *AAAI'96: Workshop on Structural Issues in Planning and Temporal Reasoning*. Citeseer.
- Prakash, K.; Husain, F.; Paruchuri, P.; and Gujar, S. P. 2021. How Private Is Your RL Policy? An Inverse RL Based Analysis Framework. *arXiv preprint arXiv:2112.05495*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Seo, K.; and Yang, J. 2020. Differentially Private Actor and Its Eligibility Trace. *Electronics*, 9(9): 1486.
- Vietri, G.; Balle, B.; Krishnamurthy, A.; and Wu, S. 2020. Private reinforcement learning with pac and regret guarantees. In *International Conference on Machine Learning*, 9754–9764. PMLR.
- Wang, B.; and Hegde, N. 2019. Privacy-preserving Q-Learning with Functional Noise in Continuous State Spaces. *arXiv preprint arXiv:1901.10634*.
- Wulfmeier, M.; Ondruska, P.; and Posner, I. 2015a. Deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888*.
- Wulfmeier, M.; Ondruska, P.; and Posner, I. 2015b. Maximum entropy deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888*.
- Ziebart, B. D.; Maas, A. L.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, 1433–1438. Chicago, IL, USA.