
Privacy-preserving Q-Learning with Functional Noise in Continuous Spaces

Baoxiang Wang

The Chinese University of Hong Kong
Borealis AI, Edmonton
bxwang@cse.cuhk.edu.hk

Nidhi Hegde

Borealis AI, Edmonton
nidhi.hegde@borealisai.com

Abstract

We consider differentially private algorithms for reinforcement learning in continuous spaces, such that neighboring reward functions are indistinguishable. This protects the reward information from being exploited by methods such as inverse reinforcement learning. Existing studies that guarantee differential privacy are not extendable to infinite state spaces, as the noise level to ensure privacy will scale accordingly to infinity. Our aim is to protect the value function approximator, without regard to the number of states queried to the function. It is achieved by adding functional noise to the value function iteratively in the training. We show rigorous privacy guarantees by a series of analyses on the kernel of the noise space, the probabilistic bound of such noise samples, and the composition over the iterations. We gain insight into the utility analysis by proving the algorithm’s approximate optimality when the state space is discrete. Experiments corroborate our theoretical findings and show improvement over existing approaches.

1 Introduction

Increasing interest in reinforcement learning (RL) and deep reinforcement learning has led to recent advances in a wide range of algorithms [SB18]. While a large part of the advancement has been in the space of games, the applicability of RL extends to other practical cases such as recommendation systems [ZZZ⁺18, LSTS15] and search engines [RJG⁺18, HDZ⁺18]. With the popularity of the RL algorithms increasing, so have concerns about their privacy. Namely, the released value (or policy) function are trained based on the reward signal and other inputs, which commonly rely on sensitive data. For example, an RL recommendation system may use the reward signals simulated by users’ historical records. This historical information can thus be inferred by recursively querying the released functions. We consider differential privacy [DMNS06, DR14], a natural and standard privacy notion, to protect such information in the RL methods.

RL methods learn by carrying out actions, receiving rewards observed for that action in a given state, and transitioning to the next states. Observation of the learned value function can reveal sensitive information: *the reward function* is a succinct description of the task. It is also connected to the users’ preferences and the criteria of their decision-making; *the visited states* carry important contextual information on the users, such as age, gender, occupation, and etc.; *the transition function* includes the dynamics of the system and the impact of the actions on the environment. Among those, the reward function is the most vulnerable and valuable component, and studies have been conducted to infer this information [AN04, NR00]. In this paper, our aim is to design differentially private algorithms for RL, such that neighboring reward functions are indistinguishable.

There is a recent line of research on privacy-preserving algorithms by protecting the reward function. Balle et al. [BGP16] train the private value function using a fixed set of trajectories. However when a new state is queried this privacy guarantee will not hold. Similar results are also considered in

contextual multi-arm bandits [SS18, SS19, PFW⁺18], where the context vector is analogous to the state. The gap that these works leave lead us to design a private algorithm that is not dependent on the number of states queried to the value function.

In order to achieve this under continuous space settings, we investigate the Gaussian process mechanism proposed by Hall et al. [HRW13]. The mechanism adds functional noise to the value function approximation hence the function can be evaluated at arbitrarily many states while preserving privacy. We show that our choice of the reproducing kernel Hilbert space (RKHS) embeds common neural networks, hence a nonlinear value function can also be used. We therefore adapt Q-learning [MKS⁺15, WD92, Bai95] so that the value function is protected after each update, even when new states are visited.

We rigorously show differential privacy guarantees of our algorithm with a series of techniques. Notably, we derive a probabilistic bound of the sample paths thus ensuring that the RKHS norm of the noised function can be bounded. This bound is significantly better than a union bound of all noise samples. Further, we analyze the composition of the privacy costs of the mechanism. There is no known composition result of the functional mechanism, other than the general theorems that apply to any mechanism [KOV13, DRV10, BKN10]. Inspired by these theorems, we derive a privacy guarantee which is better than existing results. On the utility analysis, though there is no known performance analysis on deep reinforcement learning, we gain insights by proving the utility guarantee under the tractable discrete state space settings. Empirically, experiments corroborate our theoretical findings and show improvement over existing methods.

Related Works. There is a recent line of research that discusses privacy-preserving approaches on online learning and stochastic multi-armed bandit problems [SB18, Sze10]. The algorithms protect neighboring reward sequences from being distinguished, which is related to our definition of neighboring reward functions. In bandit problems, the algorithms preserve the privacy via mechanisms that add noise to the estimates of the reward distribution [TD17, TD16, MT15, TS13, KGGW15]. This line of work shares similar motivations as our work, but they do not scale to the continuous space because of the \sqrt{N} or $\sqrt{N \ln N}$ factor involved where N is the number of arms. Similarly, in the online learning settings, the algorithms preserve the privacy evaluated sequence of the oracle [GUK17, ALMT17, AS17, JKT12]. Their analyses are based on optimizing a fixed objective thus do not apply to our setting.

More closely related are privacy studies on contextual bandits [SS19, SS18], where there is a contextual vector that is analogous to the states in reinforcement learning. Equivalently, differentially private policy evaluation [BGP16] considers a similar setting where the value function is learned on a one-step MDP. It worth note that they also consider the privacy with respect to the state and the actions, though in this paper we will focus only on the rewards. The major challenge to extend these works is that reinforcement learning requires an iterative process of policy evaluation and policy improvement. The additional states that are queried to the value function are not guaranteed to be visited and protected by previous iterations. We propose an approach for both the evaluation and the improvement, while also extending the algorithm to nonlinear approximations like neural networks.

Differential privacy in a Markov decision process (MDP) has been discussed [Ven13] via the input perturbation technique. In the work, the reward is reformulated as a weighted sum of the utility and the privacy measure. With this formulation, it amounts to learn the MDP under this weighted reward. Essentially, input perturbation will cause relatively large utility loss and is therefore less preferred. Similarly, output perturbation can be used to preserve privacy, as shown in our analysis. It is though obvious that the necessary noise level is relatively larger and also depends on more factors than our algorithm does. Therefore, more subtle techniques will be required to improve the methods by input and output perturbation.

A general approach that can be applied to continuous spaces is the differentially private deep learning framework [ACG⁺16, CBK⁺19]. The method perturbs the gradient estimator in the updates of the neural network parameters to preserve privacy. In our problem, applying the method will require large noise levels. In fact, the algorithm considers neighboring inputs that at most one data point can be different, therefore benefits from a $1/B$ factor via privacy amplification [KLN⁺11, BKN10] where B is the batch size. This no longer holds in reinforcement learning, as all reward signals can be different for neighboring reward functions, causing the noise level to scale B times back.

2 Preliminaries

2.1 Markov Decision Process and Reinforcement Learning

Markov decision process (MDP) is a framework to model decisions in an environment. We use canonical settings of the discrete-time Markov decision process. An MDP is denoted by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, r, \rho_0, \gamma)$ which includes the state space \mathcal{S} , the action space $\mathcal{A} = \{1, \dots, m\}$, the stochastic transition kernel $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^+$, the reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, the initial state distribution $\rho_0 : \mathcal{S} \rightarrow \mathbb{R}^+$ and the discount factor $\gamma \in [0, 1)$. Denote m in the above as the number of actions in the action space. The objective is to maximize the expected discounted cumulative reward. Further define the policy function $\pi : \mathcal{S}, \mathcal{A} \rightarrow \mathbb{R}^+$ and the corresponding action-state value function as

$$Q^\pi(s, a) = \mathbb{E}_\pi \left[\sum_{t \geq 0} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a, \pi \right].$$

When the context is clear, we omit π and write $Q(s, a)$ instead.

We use the continuous state space setting for this paper, except in Appendix D. We investigate bounded and continuous state space $\mathcal{S} \subseteq \mathbb{R}$ and without loss of generality assume that $\mathcal{S} = [0, 1]$. The value function $Q(s, a)$ is treated as a set of m functions $Q_a(s)$, where each function is defined on $[0, 1]$. The reward function is similarly written as a set of m functions, each defined on $[0, 1]$. We do not impose any particular assumptions on the reward function.

Our algorithm is based on deep Q-learning [MKS⁺15, Bai95, WD92], which solves the Bellman equation. Our differential privacy guarantee can also be generalized to other Q-learning algorithms. The objective of deep Q-learning is to minimize the Bellman error

$$\frac{1}{2} (Q(s, a) - \mathbb{E}[r + \gamma \max_{a'} Q(s', a')])^2,$$

where $s' \sim \mathcal{T}(s, a, s')$ denotes the consecutive state after executing action a at state s . Similar to [MKS⁺15], we use a neural network to parametrize $Q(s, a)$. We will focus on output a learned value function where the reward function $r(\cdot)$ and $r'(\cdot)$ cannot be distinguished by observing $Q(s, a)$, as long as $\|r - r'\|_\infty \leq 1$. Here without ambiguity we write $r(\cdot), r'(\cdot)$ as r, r' , and the infinity norm $\|f(s)\|_\infty$ is defined as $\sup_s |f(s)|$.

2.2 Differential Privacy

Differential privacy [DKM⁺06, DMNS06] has developed into a strong standard for privacy guarantees in data analysis. It provides a rigorous framework for privacy guarantees under various adversarial attacks.

The definition of differential privacy is based on the notion that in order to preserve privacy, data analysis should not differ at the aggregate level whether any given user is present in the input or not. This latter condition on the presence of any user is formalized through the notion of neighboring inputs. The definition of neighboring inputs will vary according to the problem settings.

Let $d, d' \in \mathcal{D}$ be neighboring inputs.

Definition 1. A randomized mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{U}$ satisfies (ϵ, δ) -differential privacy if for any two neighboring inputs d and d' and for any subset of outputs $\mathcal{Z} \subseteq \mathcal{U}$ it holds that

$$\mathbb{P}(\mathcal{M}(d) \in \mathcal{Z}) \leq \exp(\epsilon) \mathbb{P}(\mathcal{M}(d') \in \mathcal{Z}) + \delta.$$

An important parameter of a mechanism is the (global) sensitivity of the output.

Definition 2. For all pairs $d, d' \in \mathcal{D}$ of neighboring inputs, the sensitivity of a mechanism \mathcal{M} is defined as

$$\Delta_{\mathcal{M}} = \sup_{d, d' \in \mathcal{D}} \|\mathcal{M}(d) - \mathcal{M}(d')\|, \quad (1)$$

where $\|\cdot\|$ is a norm function defined on \mathcal{U} .

Vector-output mechanisms. For converting vector-valued functions into a (ϵ, δ) -DP mechanism, one of the standard approaches is the Gaussian mechanism. This mechanism adds $\mathcal{N}(0, \sigma^2 \mathbb{I})$ to the output $\mathcal{M}(d)$. In this case $\mathcal{U} = \mathbb{R}^n$ and $\|\cdot\|$ in (1) is the ℓ^2 -norm $\|\cdot\|_2$.

Proposition 3 (Vector-output Gaussian mechanism; Theorem A.1 of [DR14]). *If $0 < \epsilon < 1$ and $\sigma \geq \sqrt{2\ln(1.25/\delta)}\Delta_{\mathcal{M}}/\epsilon$, then $\mathcal{M}(d) + y$ is (ϵ, δ) -differentially private, where y is drawn from $\mathcal{N}(0, \sigma^2 I)$.*

Function-output mechanisms. In this setting the output of the function is a function, which means the mechanism is a functional. We consider the case where \mathcal{U} is an RKHS and $\|\cdot\|$ in (1) is the RKHS norm $\|\cdot\|_{\mathcal{H}}$. Hall et al. [HRW13] have shown that adding a Gaussian process noise $\mathcal{G}(0, \sigma^2 K)$ to the output $\mathcal{M}(d)$ is differentially private, when K is the RKHS kernel of \mathcal{U} . Let \mathcal{G} denote the Gaussian process distribution.

Proposition 4 (Function-output Gaussian process mechanism [HRW13]). *If $0 < \epsilon < 1$ and $\sigma \geq \sqrt{2\ln(1.25/\delta)}\Delta_{\mathcal{M}}/\epsilon$, then $\mathcal{M}(d) + g$ is (ϵ, δ) -differentially private, where g is drawn from $\mathcal{G}(0, \sigma^2 K)$ and \mathcal{U} is an RKHS with kernel function K .*

Note that in [HRW13] the stated condition was $\sigma \geq \sqrt{2\ln(2/\delta)}\Delta_{\mathcal{M}}/\epsilon$. The improvement from constant 2 to 1.25 is natural but for the completeness we include a proof in Appendix B.

3 Differentially Private Q-Learning

3.1 Our Algorithm

We present our algorithm for privacy-preserving Q-learning under the setting of continuous state space in Algorithm 1. The algorithm is based on deep Q-learning proposed by Mnih et al. [MKS⁺15]. We achieve privacy by perturbing the learned value function at each iteration, by adding a Gaussian process noise. The noise-adding is described by line 19-20 of the algorithm, where \hat{g} is the noise. This noise is a discrete estimate of the continuous sample path, evaluated at the states s_t visited in the trajectories. Intuitively, when (s, z) is an element of the list \hat{g} , it implies $g(s) = z$ for the sample path g . Line 14-18 describes the necessary maintenance of \hat{g} to simulate the Gaussian process. Line 7-9 samples a new Gaussian process sample path for every J iterations, which controls the balance between the approximation factor of privacy and the utility. The other steps are similar to [MKS⁺15].

Algorithm 1 Differentially Private Q-Learning with Functional Noise

```

1: Input: the environment and the reward function  $r(\cdot)$ 
2: Parameters: target privacy  $(\epsilon, \delta)$ , time horizon  $T$ , batch size  $B$ , action space size  $m$ , learning rate  $\alpha$ , reset factor  $J$ 
3: Output: trained value function  $Q_{\theta}(s, a)$ 
4: Initialization:  $s_0 \in [0, 1]$  uniformly,  $Q_{\theta}(s, a)$  for each  $a \in [m]$ , linked list  $\hat{g}_k[B][2] = \{\}$ 
5: Compute noise level  $\sigma = \sqrt{2(T/B) \ln(e + \epsilon/\delta)C(\alpha, k, L, B)}/\epsilon$ ;
6: for  $j$  in  $[T/B]$  do
7:   if  $j \equiv 0 \pmod{T/JB}$  then
8:      $\hat{g}_k[B][2] \leftarrow \{\}$ ;
9:   end if
10:  for  $b$  in  $[B]$  do
11:     $t \leftarrow jT/B + b$ ;
12:    Execute  $a_t = \arg \max_a Q_{\theta}(s_t, a) + \hat{g}_a(s_t)$ ;
13:    Receive  $r_t$  and  $s_{t+1}$ ,  $s \leftarrow s_{t+1}$ ;
14:    for  $a \in [m]$  do
15:      Insert  $s$  to  $\hat{g}_a[:, 1]$  such that the list remains monotonically increasing;
16:      Sample  $z_{at} \sim \mathcal{N}(\mu_{at}, \sigma d_{at})$ , according to Equation (2), Appendix A;
17:      Update the list  $\hat{g}_a(s) \leftarrow z_{at}$ ;
18:    end for
19:     $y_t \leftarrow r_t + \gamma \max_a Q_{\theta}(s_{t+1}, a) + \hat{g}_a(s_{t+1})$ ;
20:     $l_t \leftarrow \frac{1}{2}(Q_{\theta}(s_t, a_t) + \hat{g}_a(s_t) - y_t)^2$ ;
21:  end for
22:  Run one step SGD  $\theta \leftarrow \theta + \alpha \frac{1}{B} \nabla_{\theta} \sum_{t=jB}^{(j+1)B} l_t$ ;
23: end for
24: Return the trained  $Q_{\theta}(s, a)$  function;

```

Insight into the algorithm design. To satisfy differential privacy guarantees, we require two reward functions r and r' to be indistinguishable upon observation of the learned functions, as long as $\|r - r'\|_\infty \leq 1$. The major difficulty is that the reward signal $r(s, a)$ can appear at any s , and all the reward signals can be different under r and r' . Therefore, we will need a stronger mechanism of privacy that does not rely on the finite setting where at most one data point in a (finite) dataset is different, like in [ACG⁺16] and [BGP16]. This is also the major challenge in extending [BGP16] from policy evaluation to policy improvement. The natural approach to address the challenge is to treat a function as one “data point”, which leads to our utilization of the techniques studied by Hall et al. [HRW13].

3.2 Privacy, Efficiency, and Utility of the Algorithm

Privacy analysis. There are three main components in the privacy analysis. First, we have to define the RKHS to invoke the Gaussian process mechanism in Proposition 4. This RKHS should also include the value function approximation we used in the algorithm, namely, neural networks. Second, we give a privacy guarantee of composing the mechanism for T/B iterations. There is not a known composition result of such a functional mechanism, other than the general theorems that apply to any mechanism [KOV13, DRV10, BKN10]. But we derive such a privacy guarantee which is better than existing results. Third, as the sample path is evaluated on multiple different states, the updated value function can be unbounded, which subsequently induces the RKHS norm to be unbounded. This is addressed by showing a probabilistic uniform bound of the sample path over the state space.

Our privacy guarantee is shown in the following theorem.

Theorem 5. *The Q -learning algorithm in Algorithm 1 is $(\epsilon, \delta + J \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2))$ -DP with respect to two neighboring reward functions $\|r - r'\|_\infty \leq 1$, provided that $2k > 8.68\sqrt{\beta}\sigma$, and*

$$\sigma \geq \sqrt{2(T/B) \ln(e + \epsilon/\delta)} C(\alpha, k, L, B)/\epsilon,$$

where $C(\alpha, k, L, B) = ((4\alpha(k+1)/B)^2 + 4\alpha(k+1)/B)L^2$, $\beta = (4\alpha(k+1)/B)^{-1}$, L is the Lipschitz constant of the value function approximation, B is the batch size, T is the number of iterations, and α is the learning rate.

Theorem 5 provides a rigorous guarantee on the privacy of the reward function. We now present three statements to address the challenges mentioned above and support the theorem.

Lemma 6 and its corollary, informally stated below and formally stated in Appendix C, describe the RKHS that is necessary to both embed the function approximators we use and invoke the mechanism in Proposition 4.

Lemma 6 (Informal statement). *The Sobolev space H^1 with order 1 and the ℓ^2 -norm is defined as*

$$H^1 = \{f \in C[0, 1] : \partial f(x) \text{ exists; } \int_0^1 (\partial f(x))^2 dx < \infty\},$$

where $\partial f(x)$ denotes weak derivatives and the RKHS kernel is $K(x, y) = \exp(-\beta|x - y|)$.

Immediately following Lemma 6, we show that the common neural networks are in the Sobolev space. That includes neural networks with nonlinear activation layers such as a ReLU function, a sigmoid function, or the tanh function. The proof of the following corollary is also in Appendix C.

Corollary 7. *Let $\hat{f}_W(x)$ denote the neural network with finitely many finite parameters W . For $\hat{f}_W(x)$ with finitely many layers, if the gradient of the activation function is bounded, then $\hat{f}_W(x) \in H^1$.*

By the corollary $\hat{f}_W(x)$ is Lipschitz continuous. Denote L as the Lipschitz constant which only depends on the network architecture. It follows from Lemma 6 immediately that, in the algorithm for any $Q(s, a)$ and $Q'(s, a)$, $\|Q(\cdot, a) - Q'(\cdot, a)\|_{\mathcal{H}}^2 \leq 2r_0^2(1 + \beta/2)/(1 - \gamma)^2 + L^2/\beta$, for each a , where it assumes bounded reward $|r(s, a)| \leq r_0$. This will lead to an alternative privacy guarantee, but less preferred than in Theorem 5 due to the $1/(1 - \gamma)^2$ and the r_0 factor.

Line 19 and 20 use \hat{g} , which is the list of Gaussian random variables evaluated at the Gaussian process sample paths. Using a union tail bound we can derive a probabilistic bound of these variables, but it will cause the approximation factor to be $\delta + \mathcal{O}(1 - (1 - \exp(2k - \sqrt{\beta}\sigma))^T)$, which is unrealistically

large. We show in the lemma below that with high probability the entire sample path is uniformly bounded over any state s_t . We can then calibrate the δ to cover the exponentially small tail bound $\mathcal{O}(\exp(-u^2))$ of the noise. The proof is in Appendix A.

Lemma 8. *Let \mathbb{P} the probability measure over H^1 of the sample path f generated by $\mathcal{G}(0, \sigma^2 K)$. Then almost surely $\max_{x \in [0,1]} f(x)$ exists, and for any $u > 0$*

$$\mathbb{P}\left(\max_{x \in [0,1]} f(x) \geq 8.68\sqrt{\beta}\sigma + u\right) \leq \exp(-u^2/2).$$

Proof of Theorem 5. Let Q and Q' denote the learned value function of the algorithm given r and r' , respectively, where $\|r - r'\|_\infty \leq 1$. To make Q and Q' indistinguishable, we inspect the update step in line 21. Let Q_0 denote the value function after and before the update, we have

$$\|Q - Q_0\|_\infty \leq \alpha L(2 + \hat{g}_a(s_{t+1}) - \hat{g}_a(s_t))/B.$$

As per Lemma 8, with probability at least $1 - \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2)$, we have $|Q - Q_0| \leq 2\alpha L(k+1)/B$. By the triangle inequality, for any $\|r - r'\|_\infty \leq 1$, the corresponding Q and Q' satisfies $\|Q - Q'\|_\infty \leq 4\alpha L(k+1)/B$, given that Q_0 is fixed by the previous step. Let $f = Q - Q'$, we have

$$\|f\|_{\mathcal{H}}^2 \leq (1 + \beta/2)(4\alpha L(k+1)/B)^2 + L^2/2\beta$$

by the formal statement Lemma 6, Appendix C. We choose $1/\beta = 4\alpha(k+1)/B$ and have $\|f\|_{\mathcal{H}}^2 \leq ((4\alpha(k+1)/B)^2 + 4\alpha(k+1)/B)L^2$. Now by Proposition 4, adding $g \sim \mathcal{G}(0, \sigma^2 K)$ to Q will make the update step $(\epsilon', \delta' + \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2))$ -differentially private, given that $\sigma \geq \sqrt{2 \ln(1.25/\delta')} \|f\|_{\mathcal{H}}/\epsilon'$, where $K(x, y) = \exp(-4\alpha L(k+1)|x - y|/B)$ is our choice of the kernel function. Thus each iteration of update has a privacy guarantee.

It amounts to analyze the composition of T/B many iterations. It is shown by the composition theorem [KOV13, Mir17] that any $\sigma \geq \sqrt{2(T/B) \ln(1.25/\delta) \ln(e + \epsilon/\delta)} \|f\|_{\mathcal{H}}/\epsilon$ is sufficient. This is the best known bound, but we continue to derive the specific bound for our algorithm. Let z (a function, either Q or Q') be the output of a single update of the algorithm. Denote $v = 4\alpha(k+1)/B$ and $T' = T/B$ for simplicity. By Lemma 11, Appendix A, we have

$$\mathbb{E}_0[(\mathbb{P}_1(z \in S)/\mathbb{P}_0(z \in S))^\lambda] \leq \exp\left(\frac{(\lambda^2 + \lambda)(v^2 + v)L^2}{2\sigma^2}\right),$$

where \mathbb{P}_0 and \mathbb{P}_1 are the probability distribution of z given r and r' , respectively. This moment generating function will scale exponentially if multiple independent instances of z are drawn. Namely, let \mathbf{z} be the vector of T' many independent z , and $\mathbb{P}_0^{T'}$ and $\mathbb{P}_1^{T'}$ be its probability distribution under r and r' . We have for $\lambda > 0$, $\mathbb{E}_0[(\mathbb{P}_1^{T'}(\mathbf{z} \in S)/\mathbb{P}_0^{T'}(\mathbf{z} \in S))^\lambda] \leq \exp(\frac{(\lambda^2 + \lambda)(v^2 + v)L^2}{2\sigma^2} T')$. Thus,

$$\exp(\lambda(\ln(\mathbb{P}_1^{T'}(\mathbf{z})/\mathbb{P}_0^{T'}(\mathbf{z})) - \epsilon)) = \exp\left(\frac{T'\|f\|_{\mathcal{H}}^2}{2\sigma^2}(\lambda + \frac{1}{2}(1 - \frac{2\epsilon\sigma^2}{T'\|f\|_{\mathcal{H}}^2}))^2 - \frac{1}{4}(1 - \frac{2\epsilon\sigma^2}{T'\|f\|_{\mathcal{H}}^2})^2\right).$$

Since the argument holds for any $\lambda > 0$, let $\lambda = -\frac{1}{2}(1 - \frac{2\epsilon\sigma^2}{T'\|f\|_{\mathcal{H}}^2}) > 0$, then

$$\begin{aligned} \mathbb{P}_1^{T'}(\mathbf{z}) - \exp(\epsilon)\mathbb{P}_0^{T'}(\mathbf{z}) &\leq \mathbb{E}_0[\exp(\lambda(\ln(\mathbb{P}_1^{T'}(\mathbf{z})/\mathbb{P}_0^{T'}(\mathbf{z})) - \epsilon)) + \lambda \ln \lambda - (\lambda + 1) \ln(\lambda + 1))] \\ &= \exp\left(-\frac{T'\|f\|_{\mathcal{H}}^2}{2\sigma^2}(1 - \frac{2\sigma^2}{T'\|f\|_{\mathcal{H}}^2}\epsilon)^2 + \lambda \ln \lambda - (\lambda + 1) \ln(\lambda + 1)\right) \\ &\leq \exp\left(-\frac{T'\|f\|_{\mathcal{H}}^2}{2\sigma^2}(1 - \frac{2\sigma^2}{T'\|f\|_{\mathcal{H}}^2}\epsilon)^2\right)(\lambda + 1) \\ &= \exp\left(-\frac{\sigma^2}{2T'\|f\|_{\mathcal{H}}^2}(\epsilon - \frac{T'\|f\|_{\mathcal{H}}^2}{2\sigma^2})^2\right) \frac{1}{1 + \frac{\sigma^2}{T'\Delta^2}(\epsilon - \frac{T'\|f\|_{\mathcal{H}}^2}{2\sigma^2})}. \end{aligned}$$

We desire to find ϵ and δ so that this difference $\mathbb{P}_1^{T'}(\mathbf{z}) - \exp(\epsilon)\mathbb{P}_0^{T'}(\mathbf{z})$ is less than δ . We use similar techniques as in the proof Theorem 4.3 of [KOV13]. We choose

$$\epsilon = \frac{T'\|f\|_{\mathcal{H}}^2}{2\sigma^2} + \sqrt{\frac{2T'\|f\|_{\mathcal{H}}^2 w}{\sigma^2}},$$

where $w = \ln(e + \sqrt{T'\|f\|_{\mathcal{H}}^2/\sigma^2}/\delta)$. Thus the first term is e^{-w} and the second term is $\frac{1}{1 + \sqrt{\frac{2\sigma^2 w}{T'\|f\|_{\mathcal{H}}^2}}}$.

This ensures that $e^{-w} \leq \delta/\sqrt{T'\|f\|_{\mathcal{H}}^2/\sigma^2}$ and $\frac{1}{1 + \sqrt{\frac{2\sigma^2 w}{T'\|f\|_{\mathcal{H}}^2}}} \leq \frac{1}{1 + \sqrt{\frac{2\sigma^2}{T'\|f\|_{\mathcal{H}}^2}}}$, thereby guaranteeing

that $\mathbb{P}_1^{T'}(z) - \exp(\epsilon)\mathbb{P}_0^{T'}(z) \leq \delta$ for differential privacy. We solve $\epsilon = \frac{T'\|f\|_{\mathcal{H}}^2}{2\sigma^2} + \sqrt{\frac{2T'\|f\|_{\mathcal{H}}^2 w}{\sigma^2}}$ and find the sufficient condition that

$$\begin{aligned}\sigma &= \sqrt{(2T'\|f\|_{\mathcal{H}}^2/\epsilon^2) \ln(e + \epsilon/\delta)} \\ &\leq \sqrt{2(T/B)((4\alpha(k+1)/B)^2 + 4\alpha(k+1)/B)L^2 \ln(e + \epsilon/\delta)}/\epsilon,\end{aligned}$$

as desired. When this sufficient condition is satisfied, the approximation factor will be no larger than δ plus J times the uniform bound derived above by Lemma 8. Namely, it achieves $(\epsilon, \delta + J \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2))$ -DP. \square

Time complexity. We show that the noise adding mechanism in our algorithm is efficient. In fact, the most complex step introduced by the noise-adding is the insertion in line 15. This can be negligible compared with the steps such as computing gradients and executing actions. A complete proof of the below proposition is given in Appendix A.

Proposition 9. *The noised value function (during either training or released) in Algorithm 1 can respond to N_q queries in $\mathcal{O}(N_q \ln(N_q))$ time.*

Utility analysis. To the best of our knowledge, there has not been a study to rigorously analyze the utility of deep reinforcement learning. In fact, in the continuous state space setting, the solution of the Bellman equation is not unique in general. Hence, it is unlikely for Q-learning to achieve a guaranteed performance, even if it converges. However, we gain insights by analyzing the algorithm's learning error in the discrete state space setting. The learning error is defined by the discrepancy between the learned state value function and the ground truth of the optimal state value function. We consider the worst case $J = 1$ for utility (which is the best case for $(\epsilon, \delta + \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2))$ -differential privacy) where the noise is the most correlated through the iterations. We show the upper bound of the utility loss, which has a limit of zero as the number of states approaches infinity. The proof involves the linear program formulation of MDP, which is given in Appendix D.

Proposition 10. *Let v' and v^* be the value function learned by our algorithm and the optimal value function, respectively. In the case $J = 1$, $|S| = n < \infty$, and $\gamma < 1$, the utility loss of the algorithm satisfies*

$$\mathbb{E}\left[\frac{1}{n}\|v' - v^*\|_1\right] \leq \frac{2\sqrt{2}\sigma}{\sqrt{n\pi}(1-\gamma)}.$$

3.3 Discussion

Extending to other RL algorithms. Our algorithm can be extended to the case where it learns both a policy function and a value function coordinately, like the actor-critic method [MBM⁺16]. If in the updates of the policy function, only the Q function is used in the policy gradient estimation, for example $\nabla_{\theta_\pi} \ln \pi(a|s)Q(s, a)$, then the algorithm has the same privacy guarantee. Also, any post-processing of the private Q function will not break the privacy guarantee. This also includes experience replay and ϵ -greedy policies [MKS⁺15].

However, consider the case where the reward is directly accessed in the policy gradient estimation, for example $\nabla_{\theta_\pi} \ln \pi(a|s)A(s, a)$ in [DWS12, SML⁺15] where $A(s, a) = \sum_{t=0}^T (\lambda\gamma)^t (r_t + v(s_{t+1}) - v(s_t))$. In this setting, the privacy guarantee no longer holds. To extend the privacy guarantee to this case, one should add noise to the policy function as well.

Extending to high-dimensional tasks. We assumed in our analysis $\mathcal{S} = [0, 1]$ for simplicity. The setting extends to any bounded $\mathcal{S} \subseteq \mathbb{R}$ by scaling the interval. Our approach can also be extended to high-dimensional spaces by choosing a high-dimensional RKHS and kernel. For example, the kernel function $\exp(-\beta|x - y|)$ where x and y now belongs to \mathbb{R}^n and $|\cdot|$ is the Manhattan distance. It is also possible to use other RKHS and kernels for the Gaussian process noise, such as the space of band-limited functions. Other than the re-calibration of the noise level to the new kernel, the privacy guarantee in the theorem holds in general for the respective definition of $\|\cdot\|_{\mathcal{H}}$. We note that the time

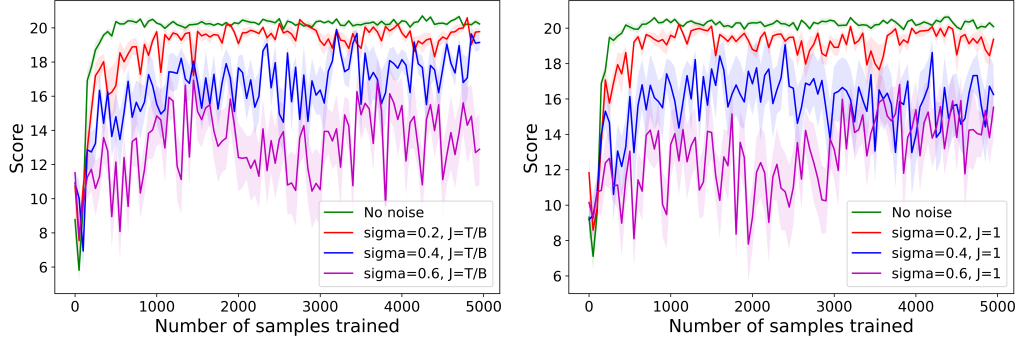


Figure 1: Empirical results of Algorithm 1 on different noise levels. The y-axis denotes the return. The x-axis is the number of samples the agent has trained on. Each episode has 50 samples. The shadow denotes 1-std. The learning curves are averaged over 10 random seeds. The curves are generated without smoothing.

complexity derived in Proposition 9 does not extend to other kernel functions, which requires the algorithm to take $\mathcal{O}(N_q^2)$ in the noise generating process.

4 Experiments

We present empirical results to corroborate our theoretical guarantees and to demonstrate the performance of the proposed algorithm on a small example. The exact MDP we use is described in Appendix E.1. The implementation is attached along with the manuscript submission.

We first plot the learning curve with a variety of noise levels and J values in Figure 1. With the increase of the noise level, the algorithm requires more samples to achieve the same return than the non-private version. This demonstrates the empirical privacy-utility tradeoff. We observe that with the noise being reset every round ($J = T/B$), the algorithm is likely to converge with limited sub-optimality as desired, especially when $\sigma < 0.4$. Therefore, as $J \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2)$ is exponentially small, we suggest using $J = T/B$ in practice to achieve a better utility.

The algorithm is then compared with a variety of baseline methods where they target the same (ϵ, δ) privacy guarantee, as shown in Figure 2(a) and 2(b). The policy evaluation method proposed by Balle, Gomrokchi and Precup [BGP16] is not differentially private under our context (while it is (ϵ, δ) -DP with respect to the reward sequences). We compare with it to illustrative the utility, where it is observed that their approach shares similar performance with ours. Note that studies on contextual bandits by Sajed and Sheffet [SS19] and by Shariff and Sheffet [SS18] consider an equivalently one-step MDP as [BGP16] and thus will yield the same method. We also compared our approach with the input perturbation method proposed by Venkatasubramanian [Ven13] and the differentially private deep learning framework by Abadi et al. [ACG⁺16]. Both the approaches are differentially private under our setting, while our algorithm significantly outperforms them. Especially, on the higher privacy regime $\epsilon = 0.45$, both the baseline methods do not improve over the training due to the large noise level needed. The baseline implementations and the exact calculation of the parameters are detailed in Appendix E.2 and E.3, respectively.

5 Conclusion

We have developed a rigorous and efficient algorithm for differentially private Q-learning in continuous state space settings. Releasing and querying the algorithm’s output value function will not distinguish two neighboring reward functions. To achieve this, our method applies functional noise taken from sample paths of a Gaussian process calibrated appropriately according to sensitivity calculated under the RKHS measure. Theoretically, we show the privacy guarantee and insights into the utility analysis. Empirically, experiments corroborate our theoretical findings and show improvement over existing methods. Our approach is general enough to be extended to other domains beyond reinforcement learning.

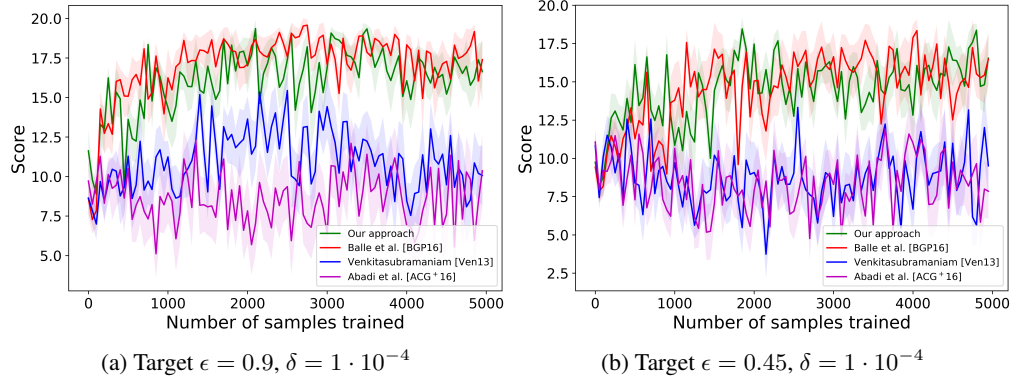


Figure 2: Empirical comparisons with other methods. Same configurations as Figure 1.

Acknowledgement

We would like to thank Ruitong Huang, who provides helpful insight on the composition analysis and the algorithm design, and Kry Yik Chau Lui, who points out the idea to extend our approach to high-dimensional Sobolev spaces.

References

- [ACG⁺16] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318. ACM, 2016.
- [ALMT17] Jacob Abernethy, Chansoo Lee, Audra McMillan, and Ambuj Tewari. Online linear optimization through the differential privacy lens. *arXiv preprint arXiv:1711.10019*, 2017.
- [AN04] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *International Conference on Machine Learning*, page 1. ACM, 2004.
- [AS17] Naman Agarwal and Karan Singh. The price of differential privacy for online learning. In *International Conference on Machine Learning*, pages 32–40, 2017.
- [Bai95] Leemon Baird. Residual algorithms: Reinforcement learning with function approximation. In *Machine Learning Proceedings 1995*, pages 30–37. Elsevier, 1995.
- [BGP16] Borja Balle, Maziar Gomrokchi, and Doina Precup. Differentially private policy evaluation. In *International Conference on Machine Learning*, pages 2130–2138, 2016.
- [BKN10] Amos Beimel, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. In *Theory of Cryptography Conference*, pages 437–454. Springer, 2010.
- [CBK⁺19] MAP Chamikara, P Bertok, I Khalil, D Liu, and S Camtepe. Local differential privacy for deep learning. *arXiv preprint arXiv:1908.02997*, 2019.
- [CW16] Yichen Chen and Mengdi Wang. Stochastic primal-dual methods and sample complexity of reinforcement learning. *arXiv preprint arXiv:1612.02516*, 2016.
- [DFVR02] Daniela Pucci De Farias and Benjamin Van Roy. *The linear programming approach to approximate dynamic programming: Theory and application*. PhD thesis, Stanford University, 2002.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006.

- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- [DR14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [DRV10] Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 51–60. IEEE, 2010.
- [DWS12] Thomas Degris, Martha White, and Richard S Sutton. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*, 2012.
- [GUK17] Pratik Gajane, Tanguy Urvoy, and Emilie Kaufmann. Corrupt bandits for preserving local privacy. *arXiv preprint arXiv:1708.05033*, 2017.
- [HDZ⁺18] Yujing Hu, Qing Da, Anxiang Zeng, Yang Yu, and Yinghui Xu. Reinforcement learning to rank in e-commerce search engine: Formalization, analysis, and application. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 368–377. ACM, 2018.
- [HRRU14] Justin Hsu, Aaron Roth, Tim Roughgarden, and Jonathan Ullman. Privately solving linear programs. In *International Colloquium on Automata, Languages, and Programming*, pages 612–624. Springer, 2014.
- [HRW13] Rob Hall, Alessandro Rinaldo, and Larry Wasserman. Differential privacy for functions and functional data. *Journal of Machine Learning Research*, 14(Feb):703–727, 2013.
- [JKT12] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Conference on Learning Theory*, pages 24–1, 2012.
- [KGGW15] Matt Kusner, Jacob Gardner, Roman Garnett, and Kilian Weinberger. Differentially private bayesian optimization. In *International Conference on Machine Learning*, pages 918–927, 2015.
- [KLN⁺11] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- [KOV13] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. *arXiv preprint arXiv:1311.0776*, 2013.
- [Lal13] Steven P. Lalley. Introduction to gaussian processes. <https://galton.uchicago.edu/~lalley/Courses/386/GaussianProcesses.pdf>, 2013. Accessed: 2018-12-27.
- [LSTS15] Elad Liebman, Maytal Saar-Tsechansky, and Peter Stone. Dj-mc: A reinforcement-learning agent for music playlist recommendation. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 591–599, 2015.
- [MBM⁺16] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937, 2016.
- [Mir17] Ilya Mironov. Renyi differential privacy. In *Computer Security Foundations Symposium (CSF), 2017 IEEE 30th*, pages 263–275. IEEE, 2017.
- [MKS⁺15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [MT15] Nikita Mishra and Abhradeep Thakurta. (nearly) optimal differentially private stochastic multi-arm bandits. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 592–601. AUAI Press, 2015.
- [NR00] Andrew Y Ng and Stuart J Russell. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, pages 663–670. Morgan Kaufmann Publishers Inc., 2000.

- [PFW⁺18] Yangchen Pan, Amir-massoud Farahmand, Martha White, Saleh Nabi, Piyush Grover, and Daniel Nikovski. Reinforcement learning with function-valued action spaces for partial differential equation control. *arXiv preprint arXiv:1806.06931*, 2018.
- [RJG⁺18] Corby Rosset, Damien Jose, Gargi Ghosh, Bhaskar Mitra, and Saurabh Tiwary. Optimizing query evaluations using reinforcement learning for web search. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1193–1196. ACM, 2018.
- [SB18] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [SML⁺15] 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.
- [SS18] Roshan Shariff and Or Sheffet. Differentially private contextual linear bandits. In *Advances in Neural Information Processing Systems*, pages 4301–4311, 2018.
- [SS19] Touqir Sajed and Or Sheffet. An optimal private stochastic-mab algorithm based on an optimal private stopping rule. In *International Conference on Machine Learning*, 2019.
- [Sze10] Csaba Szepesvári. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.
- [TD16] Aristide CY Tossou and Christos Dimitrakakis. Algorithms for differentially private multi-armed bandits. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [TD17] Aristide Charles Yedia Tossou and Christos Dimitrakakis. Achieving privacy in the adversarial multi-armed bandit. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [TS13] Abhradeep Guha Thakurta and Adam Smith. (nearly) optimal algorithms for private online learning in full-information and bandit settings. In *Advances in Neural Information Processing Systems*, pages 2733–2741, 2013.
- [Ven13] Parv Venkatasubramanian. Privacy in stochastic control: A markov decision process perspective. In *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*, pages 381–388. IEEE, 2013.
- [WD92] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- [ZZZ⁺18] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. Drn: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, pages 167–176. International World Wide Web Conferences Steering Committee, 2018.

A Proof of Lemma 8, Lemma 11 and Proposition 9

In the proofs, we will refer to some properties of the Gaussian process and its sample paths. We put those properties in Claim 12, at the end of this section. The claim and the notation used in the claim are reused in the proof of Lemma 8, Lemma 11, and Proposition 9. For simplicity, when the result can be verified immediately, such as calculating an inverse matrix, we will omit the steps.

Notations for this section. We investigate a dyadic rational. For a sample path f and $n \geq 1$, define $f_{n0} = \{f(x_0), f(x_2), \dots, f(x_{2n})\}$ and $f_{n1} = \{f(x_1), f(x_3), \dots, f(x_{2n-1})\}$, where $x_i = i/2n$, $i = 0, \dots, 2n$. For a deterministic function g defined on $[0, 1]$, define g_{n0} and g_{n1} similarly as $g_{n0} = (\lim_{x \rightarrow x_0} g(x), \lim_{x \rightarrow x_2} g(x), \dots, \lim_{x \rightarrow x_{2n}} g(x))^T$ and $g_{n1} = (\lim_{x \rightarrow x_1} g(x), \lim_{x \rightarrow x_3} g(x), \dots, \lim_{x \rightarrow x_{2n-1}} g(x))^T$. Also define $\beta_n = \beta/2n$. Our goal is to investigate the desired properties when $\lim_{n \rightarrow \infty}$.

By the definition of the Gaussian process,

$$\begin{pmatrix} f_{n1} \\ f_{n0} \end{pmatrix} \sim \mathcal{N}\left(0, \sigma^2 \begin{bmatrix} K_{11} & K_{10} \\ K_{10}^T & K_{00} \end{bmatrix}\right),$$

where K_{11} , K_{10} , and K_{00} are depending on n . If $f \sim \mathcal{G}(0, \sigma^2 K)$, the conditional distribution $f_{n1}|f_{n0} \sim \mathcal{N}(K_{10}K_{00}^{-1}f_{n0}, \sigma^2(K_{11} - K_{10}K_{00}^{-1}K_{10}^T))$. If $f \sim \mathcal{G}(g, \sigma^2 K)$, $f_{n1}|f_{n0} \sim \mathcal{N}(g_{n1} + K_{10}K_{00}^{-1}(f_{n0} - g_{n0}), \sigma^2(K_{11} - K_{10}K_{00}^{-1}K_{10}^T))$.

We restate Lemma 8 and prove it. Recall that the lemma proves a tail bound for the maximum of a GP sample path.

Lemma 8. *Let \mathbb{P} the probability measure over H^1 of the sample path f generated by $\mathcal{G}(0, \sigma^2 K)$. Then almost surely $\max_{x \in [0, 1]} f(x)$ exists, and for any $u > 0$*

$$\mathbb{P}\left(\max_{x \in [0, 1]} f(x) \geq 8.68\sqrt{\beta}\sigma + u\right) \leq \exp(-u^2/2).$$

Proof. We start with the base case that $f_{10} = \{f(0), f(1)\}$. $\mathbb{E}[\max f_{10}] = \frac{1}{2}\mathbb{E}[|f(0) - f(1)|] = \sqrt{(1 - \exp(-\beta))/\pi}\sigma \leq \sqrt{\beta/\pi}\sigma$, where the second equation is due to $f(0) - f(1) \sim \mathcal{N}(0, 2(1 - \exp(-\beta))\sigma^2)$.

We desire to bound the expectation $\mathbb{E}[\max f_{n0}]$ for all $n > 1$. To complete this, we first prove a general result on the expectation of the maximum of the Gaussian variables. Let $z_1, \dots, z_n \sim \mathcal{N}(0, \sigma_z^2)$ be n independent Gaussian random variables, by the Chernoff bound we have

$$\exp(t\mathbb{E}[\max_i z_i]) \leq \mathbb{E}[\exp(t \max_i z_i)] = \mathbb{E}[\max_i \exp(tz_i)] \leq n \exp(t^2 \sigma_z^2/2).$$

By choosing $t = \sqrt{2 \ln n}/\sigma_z$, we conclude that $\mathbb{E}[\max_i z_i] \leq \sqrt{2 \ln n} \sigma_z$. Meanwhile, it is obvious that $\text{Var}[\max_i z_i] \leq \sigma_z$.

Denote $\mu_n = \mathbb{E}[\max f_{2^n 0}]$. As $f_{2^n 0} \subset f_{2^{n+1} 0}$, the series μ_n is non-decreasing. We derive an upperbound of $\mu_{n+1} - \mu_n$. Let $x_i = i/2^{n+1}$ and $\xi_{i,n} = f(x_{2i-1}) - \frac{\exp(-\beta_{2^n})}{1 + \exp(-2\beta_{2^n})}(f(x_{2i-2}) + f(x_{2i}))$, we have that $\xi_{i,n} \sim \mathcal{N}(0, \sigma^2 \frac{1 - \exp(-2\beta_{2^n})}{1 + \exp(-2\beta_{2^n})})$. Further, $\xi_{i,n}$ and $\xi_{j,n}$ are independent. It is true as $K_{10}K_{00}^{-1}$ is nonzero only at its two diagonals, which indicates that $f(x_{2j-1})$ is not depending on other point if $f(x_{2j-2}) + f(x_{2j})$ is given. Thus, we have $\xi_{i,n}$ i.i.d. for $i = 1, \dots, 2^n$.

As we shown before the upper bound $\mathbb{E}[\max_i z_i] \leq 2\sqrt{\ln n} \sigma_z$, the expectation is monotonically increasing on the variance σ_z . In general, we have $(1 - \exp(cx))/(1 + \exp(cx)) < x$ for all positive x , if and only if $c \leq 2$. Thus $\frac{1 - \exp(-2\beta_{2^n})}{1 + \exp(-2\beta_{2^n})} < \beta_{2^n}$, and consequently we have $\mathbb{E}[\max_i \xi_{i,n}] \leq \mathbb{E}[\max_i \xi'_{i,n}]$, for $\xi'_{i,n} \sim \mathcal{N}(0, \sigma^2 \beta_{2^n})$.

By the inequality $\exp(-x)/(1 + \exp(-2x)) < \frac{1}{2}$ for $x > 0$, we relax that $\frac{\exp(-\beta_n)}{1 + \exp(-2\beta_n)} < \frac{1}{2}$. Thus, $f(x_{2i-1}) \leq \xi_{i,n} + \frac{1}{2}(f(x_{2i-2}) + f(x_{2i}))$. Taking maximum and expectation on both sides, we have

$$\mathbb{E}[\max_{1 \leq i \leq 2^n} f(x_{2i-1})] \leq \mathbb{E}[\max_{1 \leq i \leq 2^n} \xi_{i,n}] + \frac{1}{2}(f(x_{2i-2}) + f(x_{2i}))$$

$$\begin{aligned}
&\leq \mathbb{E}[\max_{1 \leq i \leq 2^n} \xi'_{i,n}] + \frac{1}{2} \mathbb{E}[\max_{1 \leq i \leq 2^n} f(x_{2i-2})] + \frac{1}{2} \mathbb{E}[\max_{1 \leq i \leq 2^n} f(x_{2i})] \\
&\leq \sqrt{2 \ln 2^n} \sqrt{\sigma^2 \beta_{2^n}} + \frac{1}{2} \mu_n + \frac{1}{2} \mu_n \\
&= \mu_n + \sqrt{n \beta / 2^n} \sigma.
\end{aligned}$$

Thus, $\mathbb{E}[\max f_{2^{n_1}} - \max f_{2^{n_0}}] \leq \sqrt{n \beta / 2^n} \sigma$. Meanwhile, we have that $\text{Var}[\max f_{2^{n_1}} - \max f_{2^{n_0}}] \leq \text{Var}[\max f_{2^{n_1}}] + \text{Var}[\max f_{2^{n_0}}] \leq 2(\beta/2^n)\sigma^2$. Let z be a random variable subject to $\mathbb{E}[z] = \sqrt{n \beta / 2^n} \sigma$, $\text{Var}[z] = 2(\beta/2^n)\sigma^2$, then there exists a z such that $\mathbb{E}[\max(0, \max f_{2^{n_1}} - \max f_{2^{n_0}})] \leq \mathbb{E}[\max(0, z)]$. We will bound this value. Denote $c = \sqrt{\beta/2^n} \sigma$ for simplicity. Then we have

$$\begin{aligned}
\exp\left(\frac{1}{c} \mathbb{E}[\max(z, 0)]\right) &\leq \mathbb{E}[\exp\left(\frac{1}{c} \max(z, 0)\right)] \\
&\leq \mathbb{E}[\max(\exp\left(\frac{z}{c}\right), \exp(0))] \\
&\leq \mathbb{E}[\exp\left(\frac{z}{c} + 1\right)] \\
&\leq \exp(\sqrt{n} + 1) + 1.
\end{aligned}$$

Consequently,

$$\mathbb{E}[\max(z, 0)] \leq (\sqrt{n} + 1 + \frac{1}{\exp(\sqrt{n} + 1)}) \sqrt{\beta/2^n} \sigma.$$

Hence we have

$$\mu_{n+1} - \mu_n = \mathbb{E}[\max(0, \max f_{2^{n_1}} - \max f_{2^{n_0}})] \leq (\sqrt{n} + 1 + \frac{1}{\exp(\sqrt{n} + 1)}) \sqrt{\beta/2^n} \sigma.$$

By induction, we have $\mu_n \leq \sqrt{\beta/\pi} \sigma + \sum_{i=0}^{\infty} (\sqrt{i} + 1 + \frac{1}{\exp(\sqrt{i} + 1)}) \sqrt{\beta/2^i} \sigma < 8.68 \sqrt{\beta} \sigma$, for any integer n . Since the dyadic rational $\cup_{i=0}^{\infty} f_{2^i 0}$ is dense and compact on $[0, 1]$ and f is continuous with probability one, $\mathbb{E}[\max f]$ shares the same upper bound of μ_n almost surely. It is shown in Theorem 3 of [Lal13], that if the expectation $\mathbb{E}[\max f]$ is bounded then $\max f$ is sub-Gaussian. The lemma follows. \square

The following Lemma 11 shows an upper bound of the moment generating function of the Gaussian process with the kernel introduced in Lemma 6. The lemma will be used to compose the Gaussian process mechanisms in Theorem 5.

Lemma 11. *Let $g \in H^1$ be continuous almost everywhere. Denote $\mathbb{P}_0(f)$ and $\mathbb{P}_1(f)$ be the probability measure over H^1 of the sample path generated by $\mathcal{G}(0, \sigma^2 K)$ and $\mathcal{G}(g, \sigma^2 K)$, respectively. The sample path $f \sim \mathbb{P}_0$ satisfies, for any $\lambda > 0$ and any set S of sample paths,*

$$\mathbb{E}_0[(\mathbb{P}_1(f \in S)/\mathbb{P}_0(f \in S))^\lambda] \leq \exp((\lambda^2 + \lambda)\|g\|_{\mathcal{H}}^2/2\sigma^2).$$

Proof. By the conditional distribution $f_{n1}|f_{n0}$ in Claim 12, we have the ratio of the probability density

$$\begin{aligned}
\ln \frac{\mathbb{P}_1(f_{n1}|f_{n0})}{\mathbb{P}_0(f_{n1}|f_{n0})} &= -\frac{(g_{n1} - K_{10}K_{00}^{-1}g_{n0})^T (K_{11} - K_{10}K_{00}^{-1}K_{10}^T)^{-1} (g_{n1} - K_{10}K_{00}^{-1}g_{n0})}{2\sigma^2} \\
&\quad + \frac{2(g_{n1} - K_{10}K_{00}^{-1}g_{n0})^T (K_{11} - K_{10}K_{00}^{-1}K_{10}^T)^{-1} (f_{n1} - K_{10}K_{00}^{-1}f_{n0})}{2\sigma^2} \\
&= -\frac{1 + \exp(-2\beta_n)}{2\sigma^2(1 - \exp(-2\beta_n))} \sum_{i=1}^n (g(x_{2i-1}) - \frac{\exp(-\beta_n)}{1 + \exp(-2\beta_n)} (g(x_{2i-2}) + g(x_{2i})))^2 \\
&\quad + \frac{1 + \exp(-2\beta_n)}{\sigma^2(1 - \exp(-2\beta_n))} \sum_{i=1}^n (g(x_{2i-1}) - \frac{\exp(-\beta_n)}{1 + \exp(-2\beta_n)} (g(x_{2i-2}) + g(x_{2i}))) \\
&\quad \times (f(x_{2i-1}) - \frac{\exp(-\beta_n)}{1 + \exp(-2\beta_n)} (f(x_{2i-2}) + f(x_{2i}))).
\end{aligned}$$

Note that in the equation above $f(x_{2i-1}) - \frac{\exp(-\beta_n)}{1+\exp(-2\beta_n)}(f(x_{2i-2}) + f(x_{2i}))$ and $f(x_{2j-1}) - \frac{\exp(-\beta_n)}{1+\exp(-2\beta_n)}(f(x_{2j-2}) + f(x_{2j}))$ are i.i.d. with the distribution $\mathcal{N}(0, \sigma^2 \frac{1-\exp(-2\beta_n)}{1+\exp(-2\beta_n)})$. Hence, We have

$$\begin{aligned}
& \mathbb{E}_0[\exp(\lambda \ln \frac{\mathbb{P}_1(f_{n1}|f_{n0})}{\mathbb{P}_0(f_{n1}|f_{n0})})] \\
&= \exp(-\lambda \frac{1+\exp(-2\beta_n)}{2\sigma^2(1-\exp(-2\beta_n))} \sum_{i=1}^n (g(x_{2i-1}) - \frac{\exp(-\beta_n)}{1+\exp(-2\beta_n)}(g(x_{2i-2}) + g(x_{2i})))^2) \\
&+ \sum_{i=1}^n (\frac{\lambda^2(1+\exp(-2\beta_n))^2}{2\sigma^4(1-\exp(-2\beta_n))^2} (g(x_{2i-1}) - \frac{\exp(-\beta_n)}{1+\exp(-2\beta_n)}(g(x_{2i-2}) + g(x_{2i})))^2 \\
&\times \sigma^2 \frac{1-\exp(-2\beta_n)}{1+\exp(-2\beta_n)}) \\
&= \exp((\lambda^2 - \lambda) \frac{1+\exp(-2\beta_n)}{2\sigma^2(1-\exp(-2\beta_n))} \sum_{i=1}^n (g(x_{2i-1}) - \frac{\exp(-\beta_n)}{1+\exp(-2\beta_n)}(g(x_{2i-2}) + g(x_{2i})))^2).
\end{aligned}$$

Meanwhile,

$$\begin{aligned}
\mathbb{E}_0[\exp(\lambda \ln \frac{\mathbb{P}_1(f_{10})}{\mathbb{P}_0(f_{10})})] &= \mathbb{E}_0[\exp(2\lambda \frac{f(0)g(0) + f(1)g(1) - \exp(-\beta)(f(0)g(1) + f(1)g(0))}{2\sigma^2(1-\exp(-2\beta))} \\
&- \lambda \frac{g(0)^2 + g(1)^2 - 2\exp(-\beta)g(0)g(1)}{2\sigma^2(1-\exp(-2\beta))})] \\
&= \mathbb{E}_0[\exp(2\lambda \frac{(g(0) - \exp(-\beta)g(1))f(0) + (g(1) - \exp(-\beta)g(0))f(1)}{2\sigma^2(1-\exp(-2\beta))} \\
&- \lambda \frac{g(0)^2 + g(1)^2 - 2\exp(-\beta)g(0)g(1)}{2\sigma^2(1-\exp(-2\beta))})] \\
&= \exp(\frac{4\lambda^2}{8\sigma^4(1-\exp(-2\beta))^2} \text{Var}((g(0) - \exp(-\beta)g(1))f(0) + (g(1) - \exp(-\beta)g(0))f(1))) \\
&- \lambda \frac{g(0)^2 + g(1)^2 - 2\exp(-\beta)g(0)g(1)}{2\sigma^2(1-\exp(-2\beta))}) \\
&= \exp(\frac{4\lambda^2}{8\sigma^4(1-\exp(-2\beta))^2} (\sigma^2(g(0) - \exp(-\beta)g(1))^2 + \sigma^2(g(1) - \exp(-\beta)g(0))^2 \\
&+ 2\exp(-\beta)\sigma^2(g(0) - \exp(-\beta)g(1))(g(1) - \exp(-\beta)g(0))) \\
&- \lambda \frac{g(0)^2 + g(1)^2 - 2\exp(-\beta)g(0)g(1)}{2\sigma^2(1-\exp(-2\beta))}) \\
&= \exp((\lambda^2 - \lambda) \frac{g(0)^2 + g(1)^2 - 2\exp(-\beta)g(0)g(1)}{2\sigma^2(1-\exp(-2\beta))}).
\end{aligned}$$

Finally, with (♥) follows by cancelling the $g(i-1)g(i)$ terms, (♦) by relaxing the exponential terms, and $z(i)$ indicating number of bits before the i 's last 1-bit, we have

$$\begin{aligned}
\mathbb{E}_0[\exp(\lambda \ln \frac{\mathbb{P}_1(f_{2^n0})}{\mathbb{P}_0(f_{2^n0})})] &= \mathbb{E}_0(\exp(\lambda \ln \frac{\mathbb{P}_1(f_{2^{n-1}1}|f_{2^{n-1}0}) \cdots \times \mathbb{P}_1(f_{2^{01}}|f_{2^{00}})\mathbb{P}_1(f_{2^{00}})}{\mathbb{P}_0(f_{2^{n-1}1}|f_{2^{n-1}0}) \cdots \times \mathbb{P}_0(f_{2^{01}}|f_{2^{00}})\mathbb{P}_0(f_{2^{00}})})) \\
&= \mathbb{E}_0[\exp(\lambda \ln \frac{\mathbb{P}_1(f_{10})}{\mathbb{P}_0(f_{10})})] \prod_{k=0}^{n-1} \mathbb{E}_0[\exp(\lambda \ln \frac{\mathbb{P}_1(f_{2^{k1}}|f_{2^{k0}})}{\mathbb{P}_0(f_{2^{k1}}|f_{2^{k0}})})] \\
&= \exp(\frac{\lambda^2 - \lambda}{2\sigma^2} \sum_{k=0}^{n-1} \sum_{i=1}^{2^k} \frac{1+\exp(-2\beta/2^k)}{1-\exp(-2\beta/2^k)} (g((2i-1)2^{-(k+1)}))
\end{aligned}$$

$$\begin{aligned}
& - \frac{\exp(-\beta/2^k)}{1 + \exp(-2\beta/2^k)} (g((2i-2)2^{-(k+1)}) + g((2i)2^{-(k+1)}))^2 \\
& + \frac{(\lambda^2 - \lambda)}{2\sigma^2} \frac{1}{1 - \exp(-2\beta)} (g(0)^2 + g(1)^2 - 2\exp(-\beta)g(0)g(1)) \\
& \stackrel{(\heartsuit)}{=} \exp\left(\frac{\lambda^2 - \lambda}{2\sigma^2} \left(\frac{g(0)^2 + g(1)^2}{1 - \exp(-2\beta)} + \sum_{k=0}^{n-1} \sum_{i=1}^{2^k} \frac{1 + \exp(-2\beta/2^k)}{1 - \exp(-2\beta/2^k)} (g((2i-1)2^{-(k+1)})^2 \right. \right. \\
& \quad \left. \left. + \frac{\exp(-2\beta/2^k)}{(1 + \exp(-2\beta/2^k))^2} (g((2i-2)2^{-(k+1)})^2 + g((2i)2^{-(k+1)})^2) \right) \right. \\
& \quad \left. - \sum_{i=1}^{2^n} \frac{2\exp(-\beta/2^{n-1})}{1 - \exp(-2\beta/2^{n-1})} g((i-1)2^{-n})g(i2^{-n}) \right) \\
& = \exp\left(\frac{\lambda^2 - \lambda}{2\sigma^2} \left(\sum_{i=0}^{2^n} \left(\frac{1 + \exp(-2\beta/2^{z(i)})}{1 - \exp(-2\beta/2^{z(i)})} + 2 \sum_{k=z(i)+1}^{n-1} \frac{\exp(-2\beta/2^k)}{1 - \exp(-4\beta/2^k)} - \frac{\exp(-\beta/2^{n-1})}{1 - \exp(-2\beta/2^{n-1})} \right) g(i2^{-n})^2 \right. \right. \\
& \quad \left. \left. + \sum_{i=1}^{2^n} \frac{\exp(-\beta/2^{n-1})}{1 - \exp(-2\beta/2^{n-1})} g(((i-1)2^{-n}) - g(i2^{-n}))^2 + \frac{g(0)^2 + g(1)^2}{1 - \exp(-2\beta)} \right) \right) \\
& \stackrel{(\spadesuit)}{\leq} \exp\left(\frac{\lambda^2 - \lambda}{2\sigma^2} \left(\sum_{i=0}^{2^n} \left(\frac{1 + \exp(-2\beta/2^{z(i)})}{1 - \exp(-2\beta/2^{z(i)})} + \sum_{k=z(i)+1}^{n-1} 2^k/2\beta - 2^{n-1}/2\beta \right) g(i2^{-n})^2 \right. \right. \\
& \quad \left. \left. + 2^n/2\beta \times (g(((i-1)2^{-n}) - g(i2^{-n}))^2) + \frac{1}{2}(g(0)^2 + g(1)^2) \right) \right) \\
& \leq \exp\left(\frac{\lambda^2 - \lambda}{2\sigma^2} \left(\sum_{i=0}^{2^n} \beta^2 2^{-n} g(i2^{-n})^2 + \frac{1}{2\beta} 2^{-n} ((g((i-1)2^{-n}) - g(i2^{-n}))/2^{-n})^2 + \frac{1}{2}(g(0)^2 + g(1)^2) \right) \right).
\end{aligned}$$

The lemma follows immediately by letting $\lim n \rightarrow \infty$. \square

We restate Proposition 9 and prove it. Recall that g is the sample path and \hat{g} the linked list to estimate its evaluations.

Proposition 9. *The noised value function (during either training or released) in Algorithm 1 can respond to N_q queries in $\mathcal{O}(N_q \ln(N_q))$ time.*

Proof. The value function $Q(\cdot)$ is deterministic. It amounts to show that there is an approach to efficiently estimate the Gaussian process sample path g .

We consider the n -th query, where the previous $n-1$ queries have been computed and stored. Let x_1, \dots, x_{n-1} be the previous queries and $g(x_1), \dots, g(x_{n-1})$ the known value on the sample path. When $x_1 \leq \dots \leq x_{n-1}$, $K_{10}K_{00}^{-1}$ has only two non-zero elements in each row, where the two elements are consecutive. This property holds per the computation in Claim 12, even if $x_i \neq i/2n$. In this case, the two elements need not to be one, but other elements must be zero.

Therefore, the mean $\mu_{at} = K_{10}K_{00}^{-1}$ and the variance $d_{at} = K_{11} - K_{10}K_{00}^{-1}K_{10}^T$ can be computed using these two elements in constant time. The exact calculation is shown below. Thus the noised value function can be calculated in $\mathcal{O}(\ln(N_q))$ time, which is the time complexity of inserting x_n into sorted list $x_1 \leq \dots \leq x_{n-1}$. The proposition follows.

The exact value of μ_{at} and d_{at} can be verified immediately so we omit the steps of the derivation. Denote, in the linked list \hat{g} in the algorithm, s^+ as the element s links to and s^- as the element that links to s . Treat $s^+ = 1$ and $s^- = 0$ for non-existence. When $\hat{g}_k[b] = (s, z)$, denote $\hat{g}_k(s) = z$.

Using the arguments above, we have

$$\begin{aligned}\mu_{at} &= \frac{(\exp(\beta(s-s^-)) - \exp(-\beta(s-s^-)))\hat{g}(s^-)}{\exp(\beta(s^+-s^-)) - \exp(-\beta(s^+-s^-))} + \frac{(\exp(\beta(s^+-s)) - \exp(-\beta(s^+-s)))\hat{g}(s^+)}{\exp(\beta(s^+-s^-)) - \exp(-\beta(s^+-s^-))} \\ d_{at} &= -\frac{(\exp(\beta(s-s^-)) - \exp(-\beta(s-s^-)))\exp(\beta(s-s^-))}{\exp(\beta(s^+-s^-)) - \exp(-\beta(s^+-s^-))} \\ &\quad - \frac{(\exp(\beta(s^+-s)) - \exp(-\beta(s^+-s)))\exp(\beta(s^+-s))}{\exp(\beta(s^+-s^-)) - \exp(-\beta(s^+-s^-))} + 1.\end{aligned}\tag{2}$$

□

Claim 12. *The following equations hold.*

$$\begin{aligned}K_{11} &= \begin{bmatrix} 1 & \exp(-2\beta_n) & \cdots & \exp(-(2n-2)\beta_n) \\ \exp(-2\beta_n) & 1 & \cdots & \exp(-(2n-4)\beta_n) \\ \vdots & \vdots & \ddots & \vdots \\ \exp(-(2n-2)\beta_n) & \exp(-(2n-4)\beta_n) & \cdots & 1 \end{bmatrix}, \\ K_{10} &= \begin{bmatrix} \exp(-\beta_n) & \exp(-\beta_n) & \cdots & \exp(-(2n-1)\beta_n) \\ \exp(-3\beta_n) & \exp(-\beta_n) & \cdots & \exp(-(2n-3)\beta_n) \\ \vdots & \vdots & \ddots & \vdots \\ \exp(-(2n-1)\beta_n) & \exp(-(2n-3)\beta_n) & \cdots & \exp(-\beta_n) \end{bmatrix}, \\ K_{00} &= \begin{bmatrix} 1 & \exp(-2\beta_n) & \cdots & \exp(-2n\beta_n) \\ \exp(-2\beta_n) & 1 & \cdots & \exp(-(2n-2)\beta_n) \\ \vdots & \vdots & \ddots & \vdots \\ \exp(-2n\beta_n) & \exp(-(2n-2)\beta_n) & \cdots & 1 \end{bmatrix}, \\ K_{00}^{-1} &= \frac{1}{1 - \exp(-4\beta_n)} \begin{bmatrix} 1 & -\exp(-2\beta_n) & 0 & \cdots & 0 \\ -\exp(-2\beta_n) & 1 + \exp(-4\beta_n) & -\exp(-2\beta_n) & \cdots & 0 \\ 0 & -\exp(-2\beta_n) & 1 + \exp(-4\beta_n) & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -\exp(-2\beta_n) \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}, \\ K_{10}K_{00}^{-1} &= \frac{\exp(-\beta_n)}{1 + \exp(-2\beta_n)} \begin{bmatrix} 1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 1 & 1 \end{bmatrix}, \\ K_{10}K_{00}^{-1}K_{10}^T &= \frac{\exp(-\beta_n)}{1 + \exp(-2\beta_n)} \begin{bmatrix} 2z(1) & z(1) + z(3) & \cdots & z(2n-1) + z(2n-3) \\ z(1) + z(3) & 2z(1) & \cdots & z(2n-3) + z(2n-5) \\ \vdots & \vdots & \ddots & \vdots \\ z(2n-3) + z(2n-5) & z(2n-5) + z(2n-7) & \cdots & z(1) + z(3) \\ z(2n-1) + z(2n-3) & z(2n-3) + z(2n-5) & \cdots & 2z(1) \end{bmatrix},\end{aligned}$$

where we write $z(x) = \exp(-x\beta_n)$ for simplicity.

$$K_{11} - K_{10}K_{00}^{-1}K_{10}^T = \frac{1 - \exp(-2\beta_n)}{1 + \exp(-2\beta_n)}\mathbb{I}.$$

Proof. The claim may not be obvious, but it can be verified immediately. □

B Proof of Proposition 4

The following claim improves Proposition 3 of [HRW13], from the constant 2 to constant 1.25. The claim investigates the anisotropic Gaussian noise mechanism, which can be regarded as a discrete version of the Gaussian process mechanism.

Claim 13. *Let f be an vector-input vector-output function. Define the sensitivity under the Mahalanobis distance as $\Delta = \max_{x, x'} \|M^{-1/2}(f(x) - f(x'))\|_2$ where M is positive definite symmetric. Then if $0 < \epsilon < 1$ and $\sigma \geq \sqrt{2 \ln(1.25/\delta)} \Delta / \epsilon$, $f(x) + \sigma M^{1/2} y$ is (ϵ, δ) -DP, where y is drawn from $\mathcal{N}(0, \sigma^2 I)$.*

Proof. Let $z \in \mathbb{R}^d$ and $c = \ln(\mathbb{P}(f(x) + \sigma M^{1/2} y = z) / \mathbb{P}(f(x') + \sigma M^{1/2} y = z))$,

$$\begin{aligned} c(z) &= \ln \frac{\mathbb{P}(f(x) = z)}{\mathbb{P}(f(x') = z)} \\ &= \frac{(z - f(x))^T M^{-1} (z - f(x))}{2\sigma^2} - \frac{(z - f(x'))^T M^{-1} (z - f(x'))}{2\sigma^2} \\ &= \frac{\|f(x) - f(x')\|_M^2 - 2y^T M^{-1/2} (f(x) - f(x'))}{2\sigma^2}. \end{aligned}$$

Hence, when $y \sim \mathcal{N}(0, 1)$,

$$c(z) \sim \mathcal{N}(\|f(x) - f(x')\|_M^2 / 2\sigma^2, 2\|f(x) - f(x')\|_M^2 / 2\sigma^2).$$

The rest of the argument follows the approach in [DR14] page 261, which is described in the setting of one-dimensional random variables and isotropic M . We show the argument in our setting for completeness. For δ -approximation privacy we would like to have $\mathbb{P}(c < \epsilon) > 1 - \delta/2$. We consider the following tail bound of the Gaussian distribution: $\forall t$,

$$\mathbb{P}(c \geq \mathbb{E}[c] + t) \leq \exp(-t^2 / 2 \text{Var}(c)) \sqrt{\text{Var}(c) / 2t^2 \pi},$$

which indicates that it is sufficient if both $\ln(\sqrt{2/\pi\delta^2}) \leq \ln(t\sigma/\|f(x) - f(x')\|_M) + t^2\sigma^2/2\|f(x) - f(x')\|_M^2$ and $\|f(x) - f(x')\|_M^2/2\sigma^2 + t \leq \epsilon$ are satisfied for some t . The conditions are further reduced to $\ln(\sqrt{2/\pi\delta^2}) \leq \ln(t\sigma/\Delta) + t^2\sigma^2/2\Delta^2$ and $t \leq \epsilon - \Delta^2/2\sigma^2$, respectively. We insert $t = \epsilon - \Delta^2/2\sigma^2$ to the first inequality and derive:

$$\ln\left(\frac{\epsilon\sigma}{\Delta} - \frac{\Delta}{2\sigma}\right) + \left(\frac{\epsilon^2\sigma^2}{2\Delta^2} + \frac{\Delta^2}{8\sigma^2} - \frac{\epsilon}{2}\right) \geq \ln(\sqrt{2/\pi\delta^2}).$$

With $\epsilon \leq 1$ we have $\frac{\epsilon\sigma}{\Delta} - \frac{\Delta}{2\sigma} \geq 1$ whenever $\sigma\epsilon/\Delta \geq 3/2$. With $\sigma\epsilon/\Delta \geq 3/2$ we have $\frac{\epsilon^2\sigma^2}{2\Delta^2} + \frac{\Delta^2}{8\sigma^2} - \frac{\epsilon}{2} \geq \sigma^2\epsilon^2/2\Delta^2 - 4/9$ per the monotonicity with respect to $\sigma\epsilon/\Delta$. Hence, it is sufficient that both $\sigma\epsilon/\Delta \geq 3/2$ and $\sigma^2\epsilon^2/2\Delta^2 - 4/9$ are satisfied. The choice $\sigma \geq \sqrt{2 \ln(1.25/\delta)} \Delta / \epsilon$ the immediately follows, as desired. \square

Now Proposition 4 of this paper follows the Claim 13 above, and Proposition 7 and Proposition 8 of [HRW13].

C Proofs of Lemma 6 and Corollary 7

Recall that Lemma 6 finds the desired RKHS and kernel and Corollary 7 is an immediate result that neural networks belong to this RKHS. The key observation is that we cannot restrict the value of $f(0)$ and $f(1)$ to be zero (which is common in some analysis of RKHS), as the value function should not be assumed to have zero value at the boundary.

Lemma 6. *We consider the one-dimensional function with bounded variable $x \in \mathbb{R}$, which, without loss of generality, can be treated as $x \in [0, 1]$. The Sobolev space H^1 with order 1 and the ℓ^2 -norm (also written as $W^{1,2}$ conventionally) is defined as*

$$H^1 = \{f \in C[0, 1] : \partial f(x) \text{ exists; } \int_0^1 (\partial f(x))^2 dx < \infty\},$$

where $\partial f(x)$ denotes weak derivatives and $\int(\cdot)dx$ denotes the Lebesgue integration. If H^1 is equipped with inner product

$$\langle f, g \rangle = \frac{1}{2}(f(0)g(0) + f(1)g(1)) + \frac{1}{2\beta} \int_0^1 \partial f(x) \partial g(x) + \beta^2 f(x)g(x) dx$$

where $\beta > 0$, it is an RKHS with kernel $K(x, y) = \exp(-\beta|x - y|)$.

Note that a function being differentiable almost everywhere does not imply that the function has a weak derivative. A counterexample is the Cantor function, which is thus excluded from H^1 . Any function equal to $\partial f(x)$ almost everywhere is considered identical to $\partial f(x)$ in H^1 .

Proof. The Sobolev space defined in the lemma does not constrain the value to be zero on its border $\{0, 1\}$, hence standard arguments do not apply. However the arguments will be similar. It suffices to show that $f(x)^2 \leq c\langle f, f \rangle$ for some c and that H^1 is complete. The former can be seen by showing that for any nonzero f , $\langle f, g \rangle = 0$ if and only if $g = 0$. By the Cauchy-Schwartz inequality,

$$\begin{aligned} \langle f, f \rangle &\geq \frac{1}{2\beta} \int_0^1 (\partial f(x))^2 dx \geq \frac{1}{2\beta} \int_0^x (\partial f(z))^2 dz \\ &\geq \frac{1}{2\beta} \left(\int_0^x \partial f(z) dz \right)^2 = \frac{1}{2\beta} f(x)^2. \end{aligned}$$

We set $c = 2\beta$ as desired. For the completeness of H^1 , we show that for any sequence $\{f_n\}$ with $\langle f_n - f_{n+1}, f_n - f_{n+1} \rangle$ converging to zero, the limit of the sequence is in H^1 . In fact, $\langle f_n - f_{n+1}, f_n - f_{n+1} \rangle$ converging to zero indicates that $\int_0^1 (f_n - f_{n+1}) dx$ converges to zero, which then indicates that $\{f_n(x)\}$ converges pointwise for any x . The first part of the lemma follows. We then verify that $f(y) = \langle f, K_y \rangle$ for any function $f(y) \in H^1$. In fact,

$$\begin{aligned} \langle f, K_y \rangle &= \frac{1}{2}(f(0)K_y(0) + f(1)K_y(1)) + \frac{\beta}{2} \int_0^1 f(x)K_y(x) dx \\ &\quad + \frac{1}{2\beta} (f(x)\partial K_y(x)|_0^1 - \int_0^1 f(x)\partial^2 K_y(x) dx) \\ &= \frac{1}{2}(f(0)K_y(0) + f(1)K_y(1)) + \frac{\beta}{2} \int_0^1 f(x)K_y(x) dx + \frac{1}{2\beta} f(x)(-\beta u(x-y)K_y(x))|_0^1 \\ &\quad - \frac{1}{2\beta} \int_0^1 f(x)((-\beta u(x-y))^2 K_y(x) - 2\beta \delta(x)K_y(x)) dx \\ &= -\frac{1}{2\beta} \int_0^1 -f(x)2\beta \delta(x-y)K_y(x) dx = f(y), \end{aligned}$$

where $u(x)$ is the sign function and $\delta(x)$ is the impulse function. By the Riesz representation theorem, $K(x, y)$ is the unique kernel of H^1 equipped with the inner product $\langle f, g \rangle$ defined above. \square

Corollary 7. Let $\hat{f}_W(x)$ denote the neural network with finite many finite parameters W . For $\hat{f}_W(x)$ with finite many layers, if the gradient of the activation function is bounded, then $\hat{f}_W(x) \in H^1$.

Proof. Let $\psi_i(\cdot)$ be the activation function so the i -th layer is represented by the function $\hat{f}_i(x) = \psi(w_i x + b_i)$. Let the gradient of $\psi_i(\cdot)$ be bounded by c . As per the chain rule we have

$$\int_0^1 (\partial \hat{f}_W(x))^2 dx \leq \prod_{i=1}^{N_c} (f_1(\dots f_{i-1}(x)) \partial f_i|_{x=x_i})^2 < \infty$$

for some x_i . Hence $\hat{f}_W(x) \in H^1$. \square

D Proof of Proposition 10

As the proposition is under the finite state space setting, the context in this section will be different from the rest of the paper. We first restate the proposition and define the notations and preliminaries needed in the analysis. We then show some intermediate claims before we give the final proof.

In the discrete state space setting, we have $\mathcal{S} = \{1, \dots, n\}$. The stochastic transition kernel is the probability distribution $\mathbb{P}(s'|s, a)$, denoted as the matrices $P_a \in \mathbb{R}^{n \times n}$ (each row sums up to one), $a = 1, \dots, m$. We write the reward function as $r_a \in \mathbb{R}^n$, $a = 1, \dots, m$. In this setting, finding the optimal action-state value function is equivalent to finding the optimal state value function $v(s)$, denoted as a vector $v \in \mathbb{R}^n$. The Bellman equation for the optimal value function is given by

$$v \geq \gamma P_i v + r_i, \quad (3)$$

for each $i = 1, \dots, m$.

Proposition 10. *Let v' and v^* be the value function learned by our algorithm and the optimal value function, respectively. In the case $J = 1$, $|\mathcal{S}| = n < \infty$, and $\gamma < 1$, the utility loss of the algorithm satisfies*

$$\mathbb{E}\left[\frac{1}{n} \|v' - v^*\|_1\right] \leq \frac{2\sqrt{2}\sigma}{\sqrt{n\pi}(1-\gamma)}.$$

Our utility analysis is based on the linear program formulation under discrete state spaces [DFVR02, CW16] and the sensitivity of linear programs [HRRU14]. In the discrete setting, v is optimal if and only if the Bellman equation (3) is satisfied. In fact, the ‘if’ relation is immediate, and the ‘only if’ relation is shown in [SB18] page 64. By exhausting the action set under the max operator and numbering the actions from 1 to m , the Bellman equation is formulated into the below linear program:

$$\begin{aligned} & \underset{v}{\text{minimize}} && \mathbf{e}^T v \\ & \text{subject to} && (\mathbb{I} - \gamma P_i)v - r_i \geq 0, \quad i = 1, \dots, m, \end{aligned} \quad (4)$$

where \mathbf{e} is the all-one vector and $\mathbf{e}^T v$ is the dummy objective. The dual of the linear program (4) is

$$\begin{aligned} & \underset{\lambda_1, \dots, \lambda_m}{\text{maximize}} && \sum_i \lambda_i^T r_i \\ & \text{subject to} && \sum_i (\mathbb{I} - \gamma P_i^T) \lambda_i = \mathbf{e}, \\ & && \lambda_i \geq 0, \quad i = 1, \dots, m. \end{aligned}$$

We consider the discrete version of Algorithm 1. The Gaussian process noise degenerates to multi-variate Gaussian noise. It is also observed that adding noise to the value function is equivalent to adding noise to the reward function, as they are additive in the update. With $J = 1$, it uses the same sample of noise through the training process.

The convergence of the algorithm is guaranteed. In fact, per [SB18] Section 4.4, the value iteration algorithm will converge to the optimal value function of the noised reward function. Formally, given the transition matrices and the noised reward signal $r'_i = r_i + z_i$ for $i = 1, \dots, m$ where $z_i \sim \mathcal{N}(0, \sigma^2 \mathbb{I})$, Algorithm 1 is guaranteed to converge to a value function v' . We desire to show that

$$\mathbb{E}\left[\frac{1}{n} \|v' - v^*\|_1\right] \leq \frac{2\sqrt{2}\sigma}{\sqrt{n\pi}(1-\gamma)},$$

where v' and v^* are the optimal value function under the reward r' and r , respectively. v' and v^* are therefore the solution of the system (4) under the reward signal r' and r , respectively.

Lemma 14 ([DFVR02] and [CW16]). *There exists an optimal dual solution λ_i^* , $i = 1, \dots, m$, an optimal deterministic policy $\pi^*(\cdot)$, and the corresponding transition matrix P^* , such that*

$$\sum_i \lambda_i^* = (\mathbb{I} - \gamma P^{*T})^{-1} \mathbf{e},$$

and the k -th entry of λ_i^ equals to the k -th entry of $\sum_i \lambda_i^*$ if $\pi^*(k) = i$, and zero otherwise.*

Proof. Similar proofs are presented in [DFVR02] and [CW16]. For the completeness of our paper we prove the claim under our context and notations. Denote the superscript (k) as the k -th element for a vector and as the k -th row for a matrix. Specify ξ_i^* to be any dual optimal solution and construct the policy $\pi^*(k) = \arg \max_i \xi_i^{*(k)}$. Then let

$$\lambda^* = (\mathbb{I} - \gamma P^{*T})^{-1} \mathbf{e},$$

where P^* is the transition matrix of $\pi^*(\cdot)$. The inversion exists since all the eigenvalues of the Markov matrix P^* are smaller than one. Define $\lambda_i^*, i = 1, \dots, m$, such that $\lambda_i^{*(k)} = \lambda^{*(k)}$ whenever $\pi^*(k) = i$ and zero otherwise. We have for λ_i^* that

$$\sum_k \sum_i \lambda_i^{*(k)} (\mathbb{I} - \gamma P_i)^{(k)} = \mathbf{e},$$

which is a rewrite of the dual feasibility by summing over k . We also have $\lambda_i^{*(k)} = 0$ whenever $\xi_i^{*(k)} = 0$ for any i and k , and together with the slackness

$$\xi_i^{*T} ((\mathbb{I} - \gamma P_i)v - r_i) = 0,$$

we have $\lambda_i^{*T} ((\mathbb{I} - \gamma P_i)v - r_i) = 0$. The optimality of $\lambda_i^*, i = 1, \dots, m$ follows. \square

Claim 15. *The ℓ^1 -norm of the dual optima $\|\sum_i \lambda_i^*\|_1$ is exactly $n/(1 - \gamma)$.*

Proof. By definition we have $\|\sum_i \lambda_i^*\|_1 = \|\lambda^*\|_1$ and $(\mathbb{I} - \gamma P^{*T})\lambda^* = \mathbf{e}$. Since P^* is a Markov matrix, we have $\|P^{*T}\lambda^*\|_1 = \|\lambda^*\|_1$. Taking ℓ^1 -norm and we have $\|\lambda^*\|_1 - \gamma\|\lambda^*\|_1 = \|\mathbf{e}\|_1$. The claim follows. \square

The following lemma justifies that there exists an algorithm to attain bounded suboptimality, given only the noised reward signal. We regard this property as the robustness of the dual system.

Lemma 16. *Let $\lambda'_i, i = 1, \dots, m$, be the optimal solution of the system*

$$\begin{aligned} & \underset{\lambda_1, \dots, \lambda_m}{\text{maximize}} && \sum_i \lambda_i'^T r'_i \\ & \text{subject to} && \sum_i (\mathbb{I} - \gamma P_i^T) \lambda_i = \mathbf{e}, \\ & && \sum_i \mathbf{e}^T \lambda_i \leq \frac{n}{1 - \gamma}, \\ & && \lambda_i \geq 0, \quad i = 1, \dots, m, \end{aligned} \tag{5}$$

we have

$$\mathbb{E}[\sum_i \lambda_i'^T r_i] \geq \sum_i \lambda_i^{*T} r_i - \frac{2\sqrt{2}n\sigma}{\sqrt{\pi}(1 - \gamma)}.$$

Proof. With (\heartsuit) follows the strong duality and (\spadesuit) follows the non-negativity and the convexity, we have

$$\begin{aligned} \mathbb{E}[\sum_i \lambda_i'^T r_i] &= \mathbb{E}[\sum_i \lambda_i'^T (r'_i - z_i)] \\ &\geq \mathbb{E}[\sum_i \lambda_i^{*T} r'_i - \sum_i \lambda_i'^T z_i] \\ &= \mathbb{E}[\sum_i \lambda_i^{*T} (r_i + z_i) - \sum_i \lambda_i'^T z_i] \\ &\stackrel{(\heartsuit)}{=} \sum_i \lambda_i^{*T} r_i + \mathbb{E}[\sum_i (\lambda_i^* - \lambda'_i)^T z_i] \\ &\stackrel{(\spadesuit)}{\geq} \sum_i \lambda_i^{*T} r_i - \frac{2}{m(1 - \gamma)} \mathbb{E}[\sum_i \|z_i\|_1] \\ &= \sum_i \lambda_i^{*T} r_i - \frac{2\sqrt{2}n\sigma}{\sqrt{\pi}(1 - \gamma)}. \end{aligned} \quad \square$$

It suffices to discuss the connection between the robustness of the primal and the robustness of the dual, which will help us to give a rigorous bound of the utility loss.

Intuitively, if we replace the maximum over λ_i by the fixed λ'_i in the below derivation of the slackness equation, the subsequent equations will yield $\mathbb{E}[\sum_i \lambda'_i r_i]$ which is desired. We observe that relaxing policy optimization (primal) side of the system at the saddle point results in an infeasible point at the value learning (dual) system. It amounts to show that this infeasible point can be mapped to the set of suboptimal values functions. Let A and B be the optimal value of the primal and the dual, the derivation of the slackness equation can be written as

$$\begin{aligned} A &= \min_v \max_{\lambda_1, \dots, \lambda_m \geq 0} \mathbf{e}^T v - (\lambda_1^T ((\mathbb{I} - \gamma P_1)v - r_1) + \dots + \lambda_m^T ((\mathbb{I} - \gamma P_m)v - r_m)) \\ &\geq \max_{\lambda_1, \dots, \lambda_m \geq 0} \min_v \mathbf{e}^T v - (\lambda_1^T ((\mathbb{I} - \gamma P_1)v - r_1) + \dots + \lambda_m^T ((\mathbb{I} - \gamma P_m)v - r_m)) \\ &= \max_{\lambda_1, \dots, \lambda_m \geq 0} \min_v (\lambda_1^T r_1 + \dots + \lambda_m^T r_m) - (-e^T + \lambda_1^T (\mathbb{I} - \gamma P_1) + \dots + \lambda_m^T (\mathbb{I} - \gamma P_m))v = B. \end{aligned}$$

Claim 17. *The stochastic policy $\pi'(i|k) = \lambda_i^{(k)} / \sum_{i'} \lambda_{i'}^{(k)}$ achieves the value v' such that $\mathbf{e}^T v' = \sum_i \lambda_i^{(k)} r_i$.*

Proof. With Lemma 14 showing the existence, specify $\lambda'' = (\mathbb{I} - \gamma P'')^{-1} \mathbf{e}$ and λ_i'' to be the optimal solution of (5) where P'' is the corresponding transition matrix. The Bellman equation indicates that $((\mathbb{I} - \gamma P_i)v' - r_i)^{(k)} = 0$ whenever $\lambda_i^{(k)} > 0$. It is equivalent to $(\mathbb{I} - \gamma P'')v' - \tilde{r} = 0$ where $\tilde{r}^{(k)} = r_{\pi(k)}^{(k)}$, $k = 1, \dots, n$. Hence,

$$\mathbf{e}^T v' = \mathbf{e}^T (\mathbb{I} - \gamma P'')^{-1} \tilde{r} = \tilde{r}^T (\mathbb{I} - \gamma P'')^{-1} \mathbf{e} = \tilde{r}^T \lambda'' = \sum_i \lambda_i^{(k)} r_i. \quad \square$$

Armed with the above results, we prove the proposition of the utility guarantee.

Proof of Proposition 10. By Lemma 16, our algorithm finds λ'_i by solving (5) which satisfies that $\sum_i \lambda_i^{*T} r_i - \mathbb{E}[\sum_i \lambda_i^T r_i] \leq \frac{2\sqrt{2}n\sigma}{\sqrt{\pi(1-\gamma)}}$. By Claim 17 we have $\mathbb{E}[\sum_i \lambda_i^T r_i] = \mathbb{E}[\mathbf{e}^T v']$. The strong duality then suggests $\sum_i \lambda_i^{*T} r_i = \mathbf{e}^T v^*$. As $\mathbb{E}[\|v' - v^*\|_1] = \mathbf{e}^T v^* - \mathbb{E}[\mathbf{e}^T v']$, the proposition follows. \square

E Details of the Experiments

E.1 The Environment

The MDP environment is defined as follows: $\mathcal{S} = [0, 1]$ and the state s denotes the location of the agent. s_0 is uniformly distributed on \mathcal{S} . $\mathcal{A} = \{0, 1\}$. If the agent chooses action 1, the agent will randomly move towards the right by a random amount sampled uniformly from $[0, 0.25]$. If after the move s is greater than 1, it will be reset to 1. Respectively, if the agent chooses action 0, the agent will randomly move towards the left by a random amount sampled uniformly from $[0, 0.25]$. If after the move s is less than 0, it will be reset to 0. The reward $0.5 - |s - 0.5|$ is given at each step, which encourages the agent to move close to the middle of the state space. Each episode of the MDP terminates at the 50th step. The algorithms are trained on 100 episodes or equivalently 5000 samples. The code is available with this manuscript submission.

E.2 The Baseline Approaches

Balle, Gomrokchi and Precup [BGP16] consider differentially private policy evaluation, where the value function is learned on a one-step MDP using a linear function approximator. This work protects the reward sequence from being distinguishable, but does not ensure the privacy of newly visited states when the value function is released. Thus we do not consider the work as differentially private under our aim of protecting the reward function. Studies on differentially private contextual bandits by Sajed and Sheffet [SS19] and by Shariff and Sheffet [SS18] are considering the equivalent problem, while we use [BGP16] to represent these works.

We also compare with the algorithm proposed by Venkatasubramanian [Ven13], via input perturbation. In the work, every reward signal is protected by a Gaussian noise thus making the algorithm differentially private. The privacy guarantee is straightforwardly derived by the composition theory by Kairouz, Oh, and Viswanath [KOV13].

We finally compare our algorithm with the differentially private deep learning by Abadi et al. [ACG⁺16]. As we use a neural network, we can perturb the gradient estimator in the updates such that all inputs are indistinguishable. We use the derived bound in Theorem 1 of [ACG⁺16]. The c_2 constant in that theorem is assigned by $\sqrt{2}$ by following the proof of the theorem.

E.3 Parameters of Our Approach

We have demonstrated the algorithms on the target of $\epsilon = 0.9$, $\delta = 1 \cdot 10^{-4}$ for Figure 2(a) and $\epsilon = 0.45$, $\delta = 1 \cdot 10^{-4}$ for Figure fig:empirical-compare(b), respectively. In this section we show how exactly these privacy targets are achieved.

Theorem 5 indicates that our algorithm is $(\epsilon, \delta + J \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2))$ -DP when

$$\sigma \geq \sqrt{2(T/B) \ln(e + \epsilon/\delta) C(\alpha, k, L, B) / \epsilon},$$

where $C(\alpha, k, L, B) = ((4\alpha(k+1)/B)^2 + 4\alpha(k+1)/B)L^2$, $\beta = (4\alpha(k+1)/B)^{-1}$. We reset the noise on every iteration, namely, let $J = \lfloor T/B \rfloor$. We rewrite the term $J \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2)$ as a tight bound $1 - (1 - \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2))^J$, which is the probability that all J sample paths are bounded by $2k$. Now we derive the set of parameters.

Let $\delta_g = 1 - (1 - \exp(-(2k - 8.68\sqrt{\beta}\sigma)^2/2))^J$ and $v = (4\alpha(k+1)/B)$. Then $\beta = 1/v$ and $C \approx vL^2$. Plugging in both the values and $T = 5000$, $B = 64$ we have $2k - 8.68\sqrt{\beta}\sigma \approx 2k - 8.6\sqrt{k+1}$. Similar to [ACG⁺16] we target 1×10^{-4} -approximation, where it is sufficient if $\delta \leq 5 \times 10^{-5}$ and $\delta_g \leq 5 \times 10^{-5}$. To satisfy $\delta_g \leq 5 \times 10^{-5}$ we need $2k - 8.6\sqrt{k+1} = \ln(1 - \exp(\ln(1 - \delta_g)/J)) \approx 3.5$, when $J = 78$. Thus $k = 23$ will be sufficient. Plugging this k value back to v we have $v = 6.19 \times 10^{-5}$, when $\alpha = 3 \times 10^{-4}$. Finally we target high-privacy regime $\epsilon = 0.9$ and plug in $L^2 = 16$ and have $\sigma \approx \sqrt{2(T/B) \ln(\epsilon/\delta) v L^2 / \epsilon} \approx 0.313$.

Approximations are made in the above arguments, but it is immediate to verify that when $\alpha = 3 \times 10^{-4}$, $k = 23$, $L^2 = 16$, $B = 64$, $T = 5 \times 10^3$, and $\sigma = 0.32$ the algorithm is $(0.9, 1 \times 10^{-4})$ -differentially private. When $\sigma = 0.74$ the algorithm is $(0.45, 1 \times 10^{-4})$ -differentially private. The above parameters correspond to Figure 2(a) and 2(b), respectively.