Supplementary Material for Differentially Private Bayesian Optimization

Matt J. Kusner MKUSNER@WUSTL.EDU
Jacob R. Gardner GARDNER.JAKE@WUSTL.EDU
Roman Garnett GARNETT@WUSTL.EDU
Kilian Q. Weinberger KILIAN@WUSTL.EDU
Washington University in St. Louis, 1 Brookings Dr., St. Louis, MO 63130

Here we give the omitted proofs of intermediate results left out of the main paper.

With observation noise

Proof of Corollary 1. Let $\mathcal{V}, \mathcal{V}'$ be neighboring datasets. Let E denote the event that the global sensitivity bound of Theorem 1 holds. Thus, $\Pr[E] \geq 1-\delta$. If E holds, drawing $\tilde{\lambda}$ with probability proportional to $\exp(\epsilon\mu_T(\lambda)/(4\sqrt{\beta_{T+1}}+2c))$ is ϵ-differentially private by the privacy guarantee of the exponential mechanism (McSherry & Talwar, 2007). Specifically, the inequality holds: $\Pr[\mathcal{A}(\mathcal{V}) = \tilde{\lambda}^{(T)}|E] \leq e^{\epsilon}\Pr[\mathcal{A}(\mathcal{V}') = \tilde{\lambda}|E]$. We demonstrate this implies (ϵ, δ)-differential privacy,

$$\begin{aligned}
&\Pr[\mathcal{A}(\mathcal{V}) = \tilde{\lambda}] \\
&\leq \Pr[\mathcal{A}(\mathcal{V}) = \tilde{\lambda}|E]\Pr[E] + \big(1 - \Pr[E]\big) \\
&\leq e^{\epsilon}\Pr[\mathcal{A}(\mathcal{V}') = \tilde{\lambda}|E]\Pr[E] + \delta \\
&\leq e^{\epsilon}\Pr[\mathcal{A}(\mathcal{V}') = \tilde{\lambda}, E] + \delta \\
&\leq e^{\epsilon}\Pr[\mathcal{A}(\mathcal{V}') = \tilde{\lambda}] + \delta.
\end{aligned}$$

■

Proof of Corollary 2. Let $\mathcal{V}, \mathcal{V}'$ be neighboring datasets. Again let E denote the event that the global sensitivity of Theorem 3 holds (and thus $\Pr[E] \geq 1 - \delta$). If E holds, adding Laplacian noise as as described in Algorithm 1 to $\max_{t \leq T} v_t$ makes $\tilde{v}$ ϵ-differential private, by the guarantee of the Laplace mechanism. Specifically, the inequality holds: $\Pr[\mathcal{A}(\mathcal{V}) = \tilde{v}|E] \leq e^{\epsilon}\Pr[\mathcal{A}(\mathcal{V}') = \tilde{v}|E]$. Using the same technique as the proof of Corollary 1 it is straightforward to show that $\tilde{v}$ is (ϵ, δ)-differentially private. ■

Without observation noise

Proof of Theorem 5. Note that at time T the regret is $f(\lambda^*) - f(\lambda_T) \leq \Omega \triangleq Ae^{-\frac{T\tau}{(\log T)^{d/4}}}$ (de Freitas et al., 2012) with probability at least $1 - \frac{\delta}{2}$. Observe the similarity of the above expression to eq. (6) (with $\max_{t \leq T} f(\lambda_t)$ replaced with $f(\lambda_T)$). In fact, the remainder of this proof follows in nearly the same way as the proof of Theorem 3. The only differences are (a) we use $f(\lambda_T)$ instead of the max term, (b) we use the regret bound of de Freitas et al. (2012) and, (c) we need not bound the maximum v as there is no noise. ■

Proof of Corollary 3. Given the sensitivity bound of Theorem 5, the proof follows in the same way as the proof of Corollary 2, where E is the event that Theorem 5 holds. ■

Proof of Theorem 6. For a random variable $Z \sim Lap(b)$, recall that $\Pr[|Z| \leq ab] = 1 - e^{-a}$. Therefore, as defined in Algorithm 2, $|\tilde{f} - f(\lambda_T)| \leq ab$ for $b = \left(\frac{\Omega}{\epsilon} + \frac{c}{\epsilon}\right)$ with probability $1 - e^{-a}$. Note that, similar to eq. (7), we have for the noise-free setting,

$$ab \geq f(\lambda_T) - \tilde{f} \geq (f(\lambda^*) - \Omega) - \tilde{f}$$

where the second inequality follows from the regret bound of de Freitas et al. (2012) and holds w.p. at least $1-\delta$. This implies that $f(\lambda^*) - \tilde{f} \leq \Omega + ab$. We can use a similar analysis to eq. (8) to show that $f(\lambda^*) - \tilde{f} \geq -\Omega - ab$. Therefore $|\tilde{f} - f(\lambda^*)| \leq \Omega + ab$ w.p. greater than $1 - (\delta + e^{-a})$. ■

Without the GP assumption

Proof of Corollary 4. Given the total global sensitivity bound implied by Theorem 7, the proof is nearly identical to the proof of Corollary 2, where E is the event that the total global sensitivity holds. ■

References

de Freitas, Nando, Smola, Alex, and Zoghi, Masrour. Exponential regret bounds for gaussian process ban-

dits with deterministic observations. In *ICML*, 2012.

Dwork, Cynthia, McSherry, Frank, Nissim, Kobbi, and Smith, Adam. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pp. 265–284. Springer, 2006.

McSherry, Frank and Talwar, Kunal. Mechanism design via differential privacy. In *FOCS*, pp. 94–103. IEEE, 2007.

Srinivas, Niranjan, Krause, Andreas, Kakade, Sham M, and Seeger, Matthias. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, 2010.