# Supplementary: On the Sample Complexity of Privately Learning Half-spaces

## 1. Proof of Lemma 12

**Lemma 1 (Restatement of Lemma 12)** *Algorithm 1 satisfies $(\epsilon, 0)$-differential privacy. Furthermore, there is at least one half-space (with the angle) $\phi^* \in S_{\mathcal{H}}$ with quality $q(S, \phi^*) = max_{\phi \in [0,2\pi)} q(S, \phi)$.*

**Proof** Given two neighboring samples $S, S' = S \cup \{(\mathbf{x}, y)\}$, we can observe that only the corresponding multiplicities of the point $(\mathbf{x}, y)$ in $S$ and $S'$ differs by 1. By the definition of $\mathcal{H}_\gamma$, there exists $\phi \in \mathcal{H}_\gamma$ such that $|\phi(\mathbf{x}) - \phi| < \gamma$ and $h_\phi(\mathbf{x}) = y$, which is the closest angle to $\phi(\mathbf{x})$ in $\mathcal{H}_\gamma$ (Line 3). The number of the angles, denoted as $n_\phi$ and $n'_\phi$, respectively, also differs by 1. So, we can consider $n_\phi$ as a sensitivity-1 function, and applying the Laplace Mechanism (Line 4) with noise distribution $Lap(\frac{1}{\epsilon})$ satisfies $(\epsilon, 0)$-differential privacy (Lemma 5). The max function can be considered as a post-processing step and introduce no privacy cost (Lemma 2). Therefore, the first statement holds.

We prove the second statement by contradiction. Suppose there is a point $(\mathbf{x}, y) \in S$ such that the corresponding half-space with the angle $\phi' \in S_{\mathcal{H}}$ and the highest quality misclassifies it. This implies that either $y = -1$ and $h_{\phi'}(\mathbf{x}) = 1$, or vice versa. Note that given the assumption of the dataset being realizable, $h_{\phi'}$ can only error on the points in $S$ that are between the target half-space and $h_{\phi'}$. Therefore, $\mathbf{x}$ is the only point in $S$ that sits in the area.

By our construction, a half-space (with the angle) $\phi'' \in \mathcal{H}_\gamma$ that correctly classifies $\mathbf{x}$ should be added to $S_{\mathcal{H}}$ (Line 3). Since there are no other points in $S$ positioned between $h_{\phi''}$ and the target half-space, $h_{\phi''}$ also correctly classifies all the points in $S$. Thus, $h_{\phi''}$ has a higher quality than $h_{\phi'}$.

Therefore, by contradiction, the second statement also holds. ■

## 2. Proof of Theorem 14

**Theorem 2 (Restatement of Theorem 14)** *For any $\epsilon, \delta, \alpha, \beta \in (0, 1)$, if there is an $(\epsilon, \delta)$-differentially private $(\alpha, \beta)$-empirical learner $A_{Thr}$ that learns thresholds on a finite domain $\mathcal{X}_{Thr}$ with $n_{Thr}(\mathcal{X}_{Thr}, \epsilon, \delta, \alpha, \beta)$ samples, then with sample complexity*

$$n = O(n_{Thr}(\mathcal{X}_{Thr}, \frac{\epsilon}{2}, \delta, \alpha, \beta)),$$

*Algorithm 3 is an $(\epsilon, \delta)$-differentially private $(\alpha, \beta)$-empirical learner for 2-dimensional half-spaces.*

**Proof** We analyze the privacy and accuracy as follows. First, for the privacy demand, our construction in the algorithm can be considered as a concatenation of three differentially private mechanisms $A_{Thr} \circ MakeThrData \circ MakeData$, which is $(2\epsilon, \delta)$-differentially private. To see that, we have MakeData satisfying $(\epsilon, 0)$-differentially private by Lemma 12. Concatenating it with MakeThrData does not introduce extra privacy cost by Dwork et al. (2014). Finally, further concatenating them with $A_{Thr}$ gives the result by the privacy guarantee of $A_{Thr}$ and the basic composition theorem in Dwork et al. (2006) and Dwork and Lei (2009). Therefore, by setting $\tilde{\epsilon} = \frac{\epsilon}{2}$ as the privacy parameter for MakeData and $A_{Thr}$, the privacy statement holds.

On the other hand, by the assumption that $S$ is realizable, our construction ensures that the half-space with quality $max_{\phi \in [0,2\pi)}(q(S, \phi))$ is in the dataset $S_{\mathcal{H}}$ by setting $\gamma$ as guaranteed by Lemma 11, along with MakeData (Lemma 12) and MakeThrData (Lemma 13). Hence, the accuracy is satisfied following the accuracy guarantee of $A_{Thr}$. Therefore, the statement holds. ■

## 3. Proof of Theorem 21

**Theorem 3** *For any $\epsilon, \delta, \alpha, \beta \in (0, 1)$ and $\delta' > 0$, if there is an $(\epsilon, \delta)$-differentially private $(\alpha, \beta)$-empirical learner $A_{Thr}$ that learns thresholds on a finite domain $\mathcal{X}_{Thr}$ with sample complexity $n_{Thr}(\mathcal{X}_{Thr}, \epsilon, \delta, \alpha, \beta)$, then with sample complexity*

$$n = O(n_{Thr}(\mathcal{X}_{Thr}, O(\frac{\epsilon}{log(\frac{1}{\delta'})}), \frac{\delta - \delta'}{2(d-1)}, \frac{\alpha}{d-1}, \frac{\beta}{d-1})),$$

*Algorithm 6 is an $(\epsilon, \delta)$-differentially private $(\alpha, \beta)$-empirical learner for d-dimensional half-spaces.*

**Proof** We analyze the privacy and accuracy as follows. First, for the privacy demand, we adapt the proof of Theorem 14 such that our construction ensures that for every two neighboring input samples, the concatenation of functions during each iteration

$$A_{Thr} \circ MakeHighDimThrData \circ MakeHighDimData$$

preserves $(2\epsilon, \delta)$-differential privacy by Corollary 18, the privacy guarantee of $A_{Thr}$ along with the basic composition theorem in Dwork et al. (2006) and Dwork and Lei (2009).

Consequently, we can bound the total privacy cost by verifying that $A_{HighH}$ satisfies the Reorder-Slice-Compute paradigm: we first set $\tau = d - 1, m = n_{Thr}$ and the sorters according to the lexicographical order of $(Q_{\phi_1^*,\dots,\phi_{i-1}^*}(S, \phi_i), \phi_i)$ in the $i^{th}$ iteration; we actually perform the steps 3, 4 of Algorith 7 in Line 2 of MakeHighDimThrData for the noisy selection of elements, and we perform the steps 5, 6 of Algorithm 7 in Line 14, 15 of $A_{HighH}$ to ensure that the points are used only once, along with the corresponding half-spaces. Consequently, the total privacy cost preserves $(\epsilon, \delta)$-differential privacy following Lemma 20, by setting privacy parameters of $\tilde{\epsilon} = O(\frac{\epsilon}{log(\frac{1}{\delta'})}), \tilde{\delta} = \frac{\delta - \delta'}{2(d-1)}$ for some $\delta' > 0$.

Next, for the accuracy, the proof follows similarly to Beimel et al. (2019) and Kaplan et al. (2020) by leveraging the accuracy guarantee of $A_{Thr}$. We aim to prove by induction that by setting accuracy parameters $\tilde{\alpha} = \frac{\alpha}{d-1}, \tilde{\beta} = \frac{\beta}{d-1}$ after the $i^{th}$ iteration, with probability at least $1 - i \cdot \frac{\beta}{d-1}$, the output values $\phi_1^*, ...\phi_i^*$ satisfy

$$Q_{\phi_1^*,...,\phi_{i-1}^*}(S, \phi_i^*) \geq (1 - \frac{\alpha}{d-1})^i \cdot OPT,$$

where $OPT = max_{\phi_1,...,\phi_{d-1} \in [0,2\pi)} q(S, (\phi_1, ..., \phi_{d-1}))$.

For the base case $i = 1$, by the guarantee of *Discretize* (Lemma 11), along with *MakeHighDimData* (Corollary 18) and *MakeHighDimThrData* (Corollary 19), there exists at least one half-space in $S_{Thr}$ that maximizes $Q(S, \cdot)$. Therefore, with the accuracy guarantee of $A_{Thr}$, with probability at least $1 - \frac{\beta}{d-1}$, the output value $\phi_1^*$ satisfies

$$Q(S, \phi_1^*) \geq (1 - \frac{\alpha}{d-1}) \cdot OPT.$$

Next, assume the statement holds for $i = k - 1$, such that with probability at least $1 - (k - 1) \cdot \frac{\beta}{d-1}$, the output values $\phi_1^*, ...\phi_{k-1}^*$ satisfy

$$Q_{\phi_1^*,...,\phi_{k-2}^*}(S, \phi_{k-1}^*) \geq (1 - \frac{\alpha}{d-1})^{k-1} \cdot OPT.$$

Therefore, by the guarantee of Lemma 11, Corollary 18, and Corollary 19, there exists $\phi_k' \in S_{Thr}$ such that $Q_{\phi_1^*,...,\phi_{k-1}^*}(S, \phi_k') = max_{\phi_k \in \mathcal{H}}(Q_{\phi_1^*,...,\phi_{k-1}^*}(S, \phi_k))$, along with the accuracy guarantee of $A_{Thr}$, with probability at least $(1 - \frac{\beta}{d-1})(1 - (k-1) \cdot \frac{\beta}{d-1}) \geq 1 - k \cdot \frac{\beta}{d-1}$, the output values $\phi_1^*, ..., \phi_k^*$ satisfy

$$\begin{aligned} Q_{\phi_1^*,...,\phi_{k-1}^*}(S, \phi_k^*) &\geq (1 - \frac{\alpha}{d-1})(1 - \frac{\alpha}{d-1})^{k-1} \cdot OPT \\ &= (1 - \frac{\alpha}{d-1})^k \cdot OPT. \end{aligned}$$

This concludes the accuracy statement such that after the $d-1$ iterations, with probability at least $1 - \beta$, $A_{HighH}$ outputs a half-space $h_{(\phi_1^*,...,\phi_{d-1}^*)}$ such that

$$\begin{aligned} Q_{\phi_1^*,...,\phi_{d-2}^*}(S, \phi_{d-1}^*) &\geq (1 - \frac{\alpha}{d-1})^{d-1} \cdot OPT \\ &\geq (1 - (d-1) \cdot \frac{\alpha}{d-1}) \cdot OPT = (1 - \alpha) \cdot OPT. \end{aligned}$$

■

## References

Amos Beimel, Shay Moran, Kobbi Nissim, and Uri Stemmer. Private center points and learning of halfspaces. In *Conference on Learning Theory*, pages 269–282. PMLR, 2019.

Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 371–380, 2009.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer, 2006.

Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Haim Kaplan, Yishay Mansour, Uri Stemmer, and Eliad Tsfadia. Private learning of halfspaces: Simplifying the construction and reducing the sample complexity. *Advances in Neural Information Processing Systems*, 33:13976–13985, 2020.