

When differential privacy meets NLP: The devil is in the detail

Ivan Habernal

Trustworthy Human Language Technologies
Department of Computer Science
Technical University of Darmstadt
ivan.habernal@tu-darmstadt.de
www.trusthlt.org

Abstract

Differential privacy provides a formal approach to privacy of individuals. Applications of differential privacy in various scenarios, such as protecting users' original utterances, must satisfy certain mathematical properties. Our contribution is a formal analysis of ADePT, a differentially private auto-encoder for text rewriting (Krishna et al., 2021). ADePT achieves promising results on downstream tasks while providing tight privacy guarantees. Our proof reveals that ADePT is not differentially private, thus rendering the experimental results unsubstantiated. We also quantify the impact of the error in its private mechanism, showing that the true sensitivity is higher by at least factor 6 in an optimistic case of a very small encoder's dimension and that the amount of utterances that are not privatized could easily reach 100% of the entire dataset. Our intention is neither to criticize the authors, nor the peer-reviewing process, but rather point out that if differential privacy applications in NLP rely on formal guarantees, these should be outlined in full and put under detailed scrutiny.

1 Introduction

The need for NLP systems to protect individuals' privacy has led to the adoption of differential privacy (DP). DP methods formally guarantee that the output of the algorithm will be 'roughly the same' regardless of whether or not any single individual is present in the central dataset; this is achieved by employing randomized algorithms (Dwork and Roth, 2013). Local DP, a variant of DP, mitigates the need for a central dataset and applies randomization on each individual's datapoint. Local DP thus guarantees that its output for an individual A will be 'almost indistinguishable' from the output of any other individuals B or C.¹

¹See the *randomized response* for an easy explanation of local DP for a single bit (Warner, 1965).

This level of privacy protection makes local DP an ideal framework for NLP applications that operate on sensitive user input which should not be collected and processed globally by an untrusted party, e.g., users' verbatim utterances. When the utterances are 'privatized' by local DP, any future post-processing or adversarial attack cannot reveal more than allowed by the particular local DP algorithm's properties (namely the ϵ parameter; see later Sec. 2).

ADePT, a local DP algorithm recently published at EACL by Krishna et al. (2021) from Amazon Alexa, proposed a differentially private auto-encoder for text rewriting. In summary, ADePT takes an input textual utterance and re-writes it in a way such that the output satisfies local DP guarantees. Unfortunately, a thorough formal analysis reveals that ADePT is in fact not differentially private and the privatized data do not protect privacy of individuals as formally promised.

In this short paper, we shed light on ADePT's main argument, the privacy mechanism. We briefly introduce key concepts from differential privacy (DP) and present a detailed proof of the Laplace mechanism (Sec. 2). Section 3 introduces ADePT's (Krishna et al., 2021) architecture and its main privacy argument. We formally prove that the proposed ADePT's mechanism is in fact not differentially private (Sec. 4) and determine the actual sensitivity of its private mechanism (Sec. 5). We sketch to which extent ADePT breaches privacy as opposed to the formal DP guarantees (Sec. 6) and discuss a potential adversary attack (Appendix C).

2 Theoretical background

From a high-level perspective, DP works with the notion of *individuals* whose information is contained in a *database (dataset)*. Each individual's *datapoint (or record)*, which could be a single bit, a number, a vector, a structured record, a text document, or any arbitrary object, is considered private

and cannot be revealed. Moreover, even whether or not any particular individual A is in the database is considered private.

Definition 2.1. Let \mathcal{X} be a ‘universe’ of all records and $x, y \in \mathcal{X}$ be two datasets from this universe. We say that x and y are neighboring datasets if they differ in one record.

For example, let dataset x consist of $|x|$ documents where each document is associated with an individual whose privacy we want to preserve. Let y differ from x by one document, so either $|y| = |x| \pm 1$, or $|y| = |x|$ with i -th document replaced. Then by definition 2.1, x and y are neighboring datasets.

Global DP and queries In a typical setup, the database is not public but held by a *trusted curator*. Only the curator can fully access all datapoints and answer any *query* we might have, for example how many individuals are in the database, whether or not B is in there, what is the most common disease (if the database is medical), what is the average length of the documents (if the database contains texts), and so on. The types of queries are task-specific, and we can see them simply as functions with arbitrary domain \mathcal{X} and co-domain \mathcal{Z} . In this paper, we focus on a simple query type, the *numerical query*, that is a function with co-domain in \mathbb{R}^n .

For example, consider a dataset $x \in \mathcal{X}$ containing textual documents and a numerical query $f : \mathcal{X} \rightarrow \mathbb{R}$ that returns an average document length. Let’s assume that the length of each document is private, sensitive information. Let the dataset x contain a particular individual A whose privacy we want to breach. Say we also have some leaked background information, in particular a neighboring dataset $y \in \mathcal{X}$ that contains all datapoints from x except for A. Now, if the trusted curator returned the true value of $f(x)$, we could easily compute A’s document length, as we know $f(y)$, and thus we could breach A’s privacy. To protect A’s privacy, we will employ randomization.

Definition 2.2. Randomized algorithm $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Z}$ takes an input value $x \in \mathcal{X}$ and outputs a value $z \in \mathcal{Z}$ nondeterministically, e.g., by drawing from a certain probability distribution.

Typically, randomized algorithms are parameterized by a density (for $z \in \mathbb{R}^n$) or a discrete distribution (for categorical or binary z). The randomized algorithm ‘perturbs’ the input by drawing

from that distribution. We suggest to consult (Igamberdiev and Habernal, 2021) for yet another NLP introduction to differential privacy.

Definition 2.3. Randomized algorithm \mathcal{M} satisfies $(\varepsilon, 0)$ -differential privacy if and only if for any neighboring datasets $x, y \in \mathcal{X}$ from the domain of \mathcal{M} , and for any possible output $z \in \mathcal{Z}$ from the range of \mathcal{M} , it holds

$$\Pr[\mathcal{M}(x) = z] \leq \exp(\varepsilon) \cdot \Pr[\mathcal{M}(y) = z] \quad (1)$$

where $\Pr[\cdot]$ denotes probability² and $\varepsilon \in \mathbb{R}^+$ is the privacy budget. A smaller ε means stronger privacy protection, and vice versa (Wang et al., 2020; Dwork and Roth, 2013).

In words, to protect each individual’s privacy, DP adds randomness when answering queries such that the query results are ‘similar’ for any pair of neighboring datasets. For our example of the average document length, the true average length would be randomly ‘noisified’.

Another view on $(\varepsilon, 0)$ -DP is when we treat $\mathcal{M}(x)$ and $\mathcal{M}(y)$ as two probability distributions. Then $(\varepsilon, 0)$ -DP puts upper bound ε on Max Divergence $\mathbb{D}_\infty(\mathcal{M}(x) || \mathcal{M}(y))$, that is the maximum ‘difference’ of any output of two random variables.³

Differential privacy has also a Bayesian interpretation, which compares the adversary’s prior with the posterior after observing the values. The odds ratio is bounded by $\exp(\varepsilon)$, see (Mironov, 2017, p. 266).

Neighboring datasets and local DP The original definition of neighboring datasets (Def. 2.1) is usually adapted to a particular scenario; see (Desfontaines and Pej3, 2020) for a thorough overview. So far, we have shown the global DP scenario with a trusted curator holding a database of $|x|$ individuals. The size of the database can be arbitrary, even containing a single individual, that is $|x| = 1$. In this case, we say a dataset $y \in \mathcal{X}$ is neighboring if it contains another single individual ($y \in \mathcal{X}$, $|y| = 1$). This setup allows us to proceed without the trusted curator, as each individual queries its single record and returns differentially private output; this scenario is known as *local DP*.

In local differential privacy, where there is no central database of records, *any pair of data points*

²The definition holds both for densities p and probability mass functions P as \Pr .

³ $\mathbb{D}_\infty(\mathcal{M}(x) || \mathcal{M}(y)) = \max_{z \in \mathcal{Z}} \left\{ \ln \frac{\Pr[\mathcal{M}(x)=z]}{\Pr[\mathcal{M}(y)=z]} \right\}$

(examples, input values, etc.) is considered neighboring (Wang et al., 2020). This also holds for ADePT: using the DP terminology, any two utterances x, y are neighboring datasets (Krishna et al., 2021).

Definition 2.4. Let $x, y \in \mathcal{X}$ be neighboring datasets. The ℓ_1 -sensitivity of a function $f : \mathcal{X} \rightarrow \mathbb{R}^n$ is defined as

$$\Delta f = \max_{x,y} \|f(x) - f(y)\|_1 \quad (2)$$

where $\|\cdot\|_1$ is a ℓ_1 -norm defined as $\|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|$ (Dwork and Roth, 2013, p. 31).

Definition 2.5. Laplace density with scale b centered at μ is defined as

$$\text{Lap}(t; \mu, b) = \frac{1}{2b} \exp\left(-\frac{|\mu - t|}{b}\right) \quad (3)$$

Definition 2.6. Laplace randomized algorithm (Dwork and Roth, 2013, p. 32). Given any function $f : \mathcal{X} \rightarrow \mathbb{R}^n$, the Laplace mechanism is defined as

$$\mathcal{M}_L(x, f, \varepsilon) = f(x) + (Y_1, \dots, Y_n) \quad (4)$$

where Y_i are i.i.d. random variables drawn from a Laplace distribution

$$Y_i \sim \text{Lap}(\mu = 0; b = \Delta f / \varepsilon) \quad (5)$$

An analogous definition centers the Laplace noise directly at the function’s output, that is

$$\begin{aligned} \mathcal{M}_L &= (Y_i \sim \text{Lap}(\mu = f(x)_1; b = \Delta f / \varepsilon), \\ &\dots, \\ &Y_n \sim \text{Lap}(\mu = f(x)_n; b = \Delta f / \varepsilon)) \end{aligned} \quad (6)$$

From Definition 2.6 also immediately follows that at point $z \in \mathbb{R}^n$, the density value of the Laplace mechanism $p(\mathcal{M}_L(x, f, \varepsilon) = z)$ is

$$\prod_{i=1}^n \frac{\varepsilon}{2\Delta f} \exp\left(-\frac{\varepsilon|f(x)_i - z_i|}{\Delta f}\right) \quad (7)$$

Theorem 2.1. The Laplace randomized algorithm preserves $(\varepsilon, 0)$ -DP (Dwork and Roth, 2013).

As ADePT relies on the proof of the Laplace mechanism, we show the full proof in Appendix A.

3 ADePT by Krishna et al. (2021)

Let \mathbf{u} be an input text (a sequence of words or a vector, for example; this is not key to the main argument). Enc is an encoder function from input \mathbf{u} to a latent representation vector $\mathbf{r} \in \mathbb{R}^n$ where n is the number of dimensions of that latent space. Dec is a decoder from the latent representation back to the original input space (again, a sequence of words or a vector). What we have so far is a standard auto-encoder, such that

$$\mathbf{r} = \text{Enc}(\mathbf{u}) \quad \text{and} \quad \mathbf{v} = \text{Dec}(\mathbf{r}). \quad (8)$$

Krishna et al. (2021) define ADePT as a randomized algorithm that, given an input \mathbf{u} , generates \mathbf{v} as $\mathbf{v} = \text{Dec}(\mathbf{r}')$, where $\mathbf{r}' \in \mathbb{R}^n$ is a clipped latent representation vector with added noise

$$\mathbf{r}' = \mathbf{r} \cdot \min\left(1, \frac{C}{\|\mathbf{r}\|_2}\right) + \boldsymbol{\eta} \quad (9)$$

where $\boldsymbol{\eta} \in \mathbb{R}^n$, $C \in \mathbb{R}$ is an arbitrary clipping constant, and $\|\cdot\|_2$ is an ℓ_2 (Euclidean) norm defined as $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$.

Theorem 3.1 (which is false). (Krishna et al., 2021) If $\boldsymbol{\eta}$ is a multidimensional noise, such that each element η_i is independently drawn from a distribution shown in equation 10, then the transformation from $\mathbf{u} \rightarrow \mathbf{v}'$ is $(\varepsilon, 0)$ -DP.

$$\text{Lap}(\eta_i) \sim \frac{\varepsilon}{4C} \exp\left(-\frac{\varepsilon|v_i|}{2C}\right) \quad (10)$$

Proof. Krishna et al. (2021) refers to the proof of Theorem 3.6 by Dwork and Roth (2013, p. 32), which is the proof of the Laplace mechanism. \square

First, v_i in Eq. 10 is ambiguous as it ‘semantically’ relates to \mathbf{v} which is the decoded vector that comes first *after* drawing a random value; moreover $\boldsymbol{\eta}$ and \mathbf{v} have different dimensions. Given that the authors employ Laplacian noise and base their proofs on Theorem 3.6 from Dwork and Roth (2013, p. 32), we believe that Eq. 10 is the standard Laplace mechanism

$$\eta_i \sim \text{Lap}(\mu = 0; b = \Delta f / \varepsilon), \quad (11)$$

such that each value η_i is drawn independently from a zero-centered Laplacian noise parametrized by scale b (Definition 2.6). Given the density from Eq. 3, we rewrite Eq. 11 as

$$\eta_i \sim \frac{\varepsilon}{2\Delta f} \exp\left(-\frac{\varepsilon|t|}{\Delta f}\right), \quad (12)$$

Krishna et al. (2021) set their clipped encoder output as the function f , that is⁴

$$f = \mathbf{r} \cdot \min\left(1, \frac{C}{\|\mathbf{r}\|_2}\right). \quad (13)$$

Theorem 3.2 (which is false). (Krishna et al., 2021) Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a function as defined in equation 13. The ℓ_1 -sensitivity Δf of this function is $2C$.

Proof. (Krishna et al., 2021) Maximum ℓ_1 norm difference between two points in a hyper-sphere of radius C is $2C$. \square

Thus by plugging the sensitivity Δf from Theorem 3.2 into Eq. 12, we obtain

$$\eta_i \sim \frac{\varepsilon}{4C} \exp\left(-\frac{\varepsilon|t|}{2C}\right), \quad (14)$$

which is what Krishna et al. (2021) express in Eq. 10. To sum up, the essential claim of Krishna et al. (2021) is that if each η_i is drawn from Laplacian distribution with scale $\frac{2C}{\varepsilon}$, their mechanism is $(\varepsilon, 0)$ differentially private.

4 ADePT with Laplace mechanism is not differentially private

Proof. Following the proof of Theorem 2.1, the following bound (Eq. 33) must hold for any x, y

$$\frac{p(M_L(x, f, \varepsilon) = z)}{p(M_L(y, f, \varepsilon) = z)} \leq \exp\left(\frac{\varepsilon}{\Delta f} \cdot \|f(y) - f(x)\|_1\right)$$

and thus this inequality must hold too

$$\exp\left(\frac{\varepsilon}{\Delta f} \cdot \|f(y) - f(x)\|_1\right) \leq \exp(\varepsilon) \quad (15)$$

Fix the clipping constant $C > 0$ arbitrarily ($C \in \mathbb{R}$), set dimensions to $n = 2$. Let $\mathbf{r}_y = (\frac{2}{3}C, \frac{2}{3}C)$ be the input y of the clipping function f from Eq. 13.

$$f(y) = \mathbf{r}_y \cdot \min\left(1, \frac{C}{\|\mathbf{r}_y\|_2}\right) \quad (\text{from Eq. 13})$$

$$= \mathbf{r}_y \cdot \min\left(1, \frac{C}{\frac{2\sqrt{2}}{3}C}\right) \quad (16)$$

$$= \mathbf{r}_y \cdot \min(1, 1.06066\dots) \quad (17)$$

$$= \mathbf{r}_y \cdot 1 = \left(\frac{2C}{3}, \frac{2C}{3}\right) \quad (18)$$

Similarly, let $\mathbf{r}_x = (-\frac{2}{3}C, -\frac{2}{3}C)$ be input x , for which we get analogically $f(x) = (-\frac{2}{3}C, -\frac{2}{3}C)$. Then

$$\|f(y) - f(x)\|_1 = \quad (19)$$

$$\left\|\left(\frac{2C}{3}, \frac{2C}{3}\right) - \left(-\frac{2C}{3}, -\frac{2C}{3}\right)\right\|_1 = \quad (20)$$

$$= \frac{8C}{3} \quad (21)$$

Plug Theorem 3.2 and Eq. 21 into Eq. 15

$$\exp\left(\frac{\varepsilon}{2C} \cdot \|f(y) - f(x)\|_1\right) \leq \exp(\varepsilon) \quad (22)$$

$$\exp\left(\frac{\varepsilon}{2C} \cdot \frac{8C}{3}\right) \leq \exp(\varepsilon) \quad (23)$$

$$\exp\left(\frac{4}{3} \cdot \varepsilon\right) \not\leq \exp(\varepsilon) \quad (24)$$

therefore Theorem 3.1 by Krishna et al. (2021) must be false. \square

In general, it is the inequality $\|\mathbf{x}\|_1 \geq \|\mathbf{x}\|_2$ that makes ADePT fail the DP proof.

5 Actual sensitivity of ADePT

Theorem 5.1. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a function as defined in Eq. 13. The sensitivity Δf of this function is $2C\sqrt{n}$.

Proof. See Appendix B. \square

Corollary 5.1. Since $2C\sqrt{n} = 2C$ only for $n = 1$, ADePT could be differentially private only if the encoder's latent representation $\mathbf{r} = \text{Enc}(\mathbf{u})$ were a single scalar.

Since Krishna et al. (2021) do not specify the dimensionality of their encoder's output, we can only assume some typical values in a range from 32 to 1024, so that the true sensitivity of ADePT is ≈ 6 to 32 times higher than reported.

6 Magnitude of non-protected data

How many data points actually violate the privacy guarantees? Without having access to the trained model and its hyper-parameters (C , in particular), it is hard to reason about properties of the latent space, where privatization occurs. We thus simulated the encoder's 'unclipped' vector outputs \mathbf{r} by sampling 10k vectors from two distributions: 1) uniform within $(-C, +C)$ for each dimension,

⁴We contacted the authors several times to double check that this formula is correct without a potential typo but got no response. However other parts of the paper give evidence it is correct, e.g., the authors use an analogy to a hyper-sphere which is considered euclidean by default.

and 2) zero-centered normal with $\sigma^2 = 0.1 \cdot C$. Especially the latter one is rather optimistic as it samples most vectors close to zero. In reality these latent space vectors are unbounded.

Each pair of such vectors in the latent space after clipping but before applying DP (Eq. 13) is ‘neighboring datasets’ so their ℓ_1 distance must be bound by sensitivity ($2C$ as claimed in Theorem 3.2) in order to satisfy DP with the Laplace mechanism.

We ran the simulation for an increasing dimensionality of the encoder’s output and measured how many pairs violate the sensitivity bound.⁵ Fig. 1 shows the ‘curse of dimensionality’ for norms. Even for a considerably small encoder’s vector size of 32 and unbounded encoder’s latent space, almost **none** of the data points would be protected by ADePT’s Laplace mechanism.




Figure 1: Simulation results. Percentage of ‘neighboring datasets’ that violate the distance bounds required by the Laplace mechanism with sensitivity $2C$.

7 Discussion

Local DP differs from centralized DP in such a way that there is no central database and once the privatized data item ‘leaves’ an individual, it stays so forever. This makes typical membership inference attacks unsuitable, as no matter what happens to the rest of the world, the probability of inferring the individual’s true value after observing their privatized data item is bounded by $\exp(\epsilon)$.

For example, the ATIS dataset used in ADePT contains 5,473 utterances of lengths 1 to 46 tokens, with a quite limited vocabulary of 941 words. In theory, the search space of all possible utterances would be of size $941^{46} \approx 6 \times 10^{136}$, and under ϵ -DP all of them are multiplicatively indistinguishable – for example, after observing “on april first i need a ticket from tacoma to san jose departing

before 7 am” from ADePT’s autoencoder privatized output, the true input might well have been “on april first i need a flight going from phoenix to san diego” or “monday morning i would like to fly from columbus to indianapolis” and our posterior certainty of any of those is limited by the privacy bound. However, since outputs of ADePT are leaking privacy, attacks are possible. We sketch a potential scenario in Appendix C.

There are two possible remedies for ADePT. Either the latent vector clipping in Eq. 9 could use ℓ_1 -norm, or the Laplacian noise in Eq. 10 could use the correct sensitivity as determined in Theorem 5.1. In either case, the utility in the downstream tasks as presented by Krishna et al. (2021) are expected to be worse due to a much larger amount of required noise.

8 Conclusion

This paper revealed a potential trap for NLP researchers when adopting a local DP approach. We believe it contributes to a better understanding of the exact modeling choices involved in determining the sensitivity of local DP algorithms. We hope that DP will become a widely accessible and well-understood framework within the NLP community.

Acknowledgements

The independent research group TrustHLT is supported by the Hessian Ministry of Higher Education, Research, Science and the Arts. Thanks to Max Glockner, Timour Igamberdiev, Jorge Cordona, Jan-Christoph Klie, and the anonymous reviewers for their helpful feedback.

References

- Damien Desfontaines and Balázs Pejó. 2020. [SoK: Differential privacies](#). *Proceedings on Privacy Enhancing Technologies*, 2020(2):288–313.
- Cynthia Dwork and Aaron Roth. 2013. [The Algorithmic Foundations of Differential Privacy](#). *Foundations and Trends® in Theoretical Computer Science*, 9(3-4):211–407.
- Andrea Gadotti, Florimond Houssiau, Luc Rocher, Benjamin Livshits, and Yves-Alexandre de Montjoye. 2019. When the Signal is in the Noise: Exploiting Diffix’s Sticky Noise. In *28th USENIX Security Symposium*, pages 1081–1098, Santa Clara, CA, USA. USENIX Association.
- Timour Igamberdiev and Ivan Habernal. 2021. [Privacy-preserving graph convolutional networks for text classification](#). *arXiv preprint*.

⁵Code available at <https://github.com/habernal/emnlp2021-differential-privacy-nlp>

Satyapriya Krishna, Rahul Gupta, and Christophe Dupuy. 2021. **ADePT: Auto-encoder based Differentially Private Text Transformation**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2435–2439, Online. Association for Computational Linguistics.

Ilya Mironov. 2017. **Rényi Differential Privacy**. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275, Santa Barbara, CA, USA. IEEE.

Teng Wang, Xuefeng Zhang, Jingyu Feng, and Xinyu Yang. 2020. **A Comprehensive Survey on Local Differential Privacy toward Data Statistics and Analysis**. *Sensors*, 20(24):7030.

Stanley L. Warner. 1965. **Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias**. *Journal of the American Statistical Association*, 60(309):63–69.

A Proof of Laplace mechanism

Theorem A.1. *Negative triangle inequality for absolute values. For $a, x, y \in \mathbb{R}$,*

$$|x - a| - |y - a| \leq |x - y|. \quad (25)$$

Proof is directly based on the triangle inequality.

Corollary A.1. *Definition 2.4 implies that Δf is an upper bound value on the ℓ_1 norm of the function output for any neighboring x and y . In other words*

$$\|f(x) - f(y)\|_1 \leq \Delta f \quad (26)$$

The actual proof (Dwork and Roth, 2013). We will prove that for any x, y the following ratio

$$\frac{p(M_L(x, f, \varepsilon) = z)}{p(M_L(y, f, \varepsilon) = z)} \quad (27)$$

is bounded by $\exp(\varepsilon)$ and thus satisfies Definition 2.3. Fix $z \in \mathbb{R}^n$ arbitrarily. By plugging Eq. 7 into Eq. 27, we get

$$= \prod_{i=1}^n \frac{\frac{\varepsilon}{2\Delta f} \exp\left(-\frac{\varepsilon|f(x)_i - z_i|}{\Delta f}\right)}{\frac{\varepsilon}{2\Delta f} \exp\left(-\frac{\varepsilon|f(y)_i - z_i|}{\Delta f}\right)} \quad (28)$$

$$= \prod_{i=1}^n \frac{\exp\left(-\frac{\varepsilon}{\Delta f}|f(x)_i - z_i|\right)}{\exp\left(-\frac{\varepsilon}{\Delta f}|f(y)_i - z_i|\right)} \quad (29)$$

$$= \prod_{i=1}^n \exp\left(\frac{\varepsilon}{\Delta f} \cdot \underbrace{|f(y)_i - z_i| - |f(x)_i - z_i|}_{\text{Apply Theorem A.1}}\right) \quad (30)$$

$$\leq \prod_{i=1}^n \exp\left(\frac{\varepsilon}{\Delta f} \cdot |f(y)_i - f(x)_i|\right) \quad (31)$$

$$= \exp\left(\frac{\varepsilon}{\Delta f} \cdot \underbrace{\sum_{i=1}^n |f(y)_i - f(x)_i|}_{\text{Def. of } \ell_1 \text{ norm}}\right) \quad (32)$$

$$= \exp\left(\frac{\varepsilon}{\Delta f} \cdot \underbrace{\|f(y) - f(x)\|_1}_{\leq \Delta f \text{ Corollary A.1}}\right) \quad (33)$$

$$\leq \exp\left(\frac{\varepsilon}{\Delta f} \cdot \Delta f\right) \quad (34)$$

$$= \exp(\varepsilon) \quad (35)$$

which is what we wanted. By symmetry we get the proof for $\frac{p(M_L(y, f, \varepsilon) = z)}{p(M_L(x, f, \varepsilon) = z)} \leq \exp(\varepsilon)$. \square

B Proof of Theorem 5.1

Proof. The definition of sensitivity corresponds to the maximum ℓ_1 distance of any two vectors \mathbb{R}^n from the range of f . As Eq. 13 bounds all vectors to their ℓ_2 (Euclidean) norm, we want to find the distance between two opposing points on an n -dimensional sphere that have maximal ℓ_1 distance.

Let $n \in \mathbb{N} > 0$ be the number of dimension and $C \in \mathbb{R}$ a positive constant. We solve the following optimization problem

$$\begin{aligned} \max_{x_1, \dots, x_n} \quad & f(x_1, \dots, x_n) = |x_1| + \dots + |x_n| \\ \text{s.t.} \quad & \sqrt{x_1^2 + \dots + x_n^2} = C \end{aligned}$$

First, we can get rid of the absolute values in $f(x_1, \dots, x_n)$ as the maximums will be symmetric, i.e. $\max(|a| + |b|) = \max(-a + |b|)$.

Using Lagrange multipliers, we define the constraints as

$$g(x_1, \dots, x_n) = \sqrt{x_1^2 + \dots + x_n^2} - C = 0,$$

hence

$$\begin{aligned} \mathcal{L}(x_1, \dots, x_n, \lambda) &= f(x_1, \dots, x_n) \\ &\quad + \lambda \cdot g(x_1, \dots, x_n) \\ &= x_1 + \dots + x_n + \\ &\quad \lambda \sqrt{x_1^2 + \dots + x_n^2} - \lambda C \end{aligned}$$

The gradient $\nabla_{x_1, \dots, x_n, \lambda} \mathcal{L}(x_1, \dots, x_n, \lambda)$ is

$$\begin{aligned} &\left(\frac{\partial \mathcal{L}}{\partial x_1}, \dots, \frac{\partial \mathcal{L}}{\partial x_n}, \frac{\partial \mathcal{L}}{\partial \lambda} \right) = \\ &\left(\frac{x_1 \lambda}{\sqrt{x_1^2 + \dots + x_n^2}} + 1, \dots, \right. \\ &\left. \frac{x_n \lambda}{\sqrt{x_1^2 + \dots + x_n^2}} + 1, \sqrt{x_1^2 + \dots + x_n^2} - C \right) \end{aligned}$$

Solve $\nabla_{x_1, \dots, x_n, \lambda} \mathcal{L}(x_1, \dots, x_n, \lambda) = 0$ by the following system of $n + 1$ equations

$$\begin{aligned} \frac{x_1 \lambda}{\sqrt{x_1^2 + \dots + x_n^2}} + 1 &= 0 \\ \dots &= 0 \\ \frac{x_n \lambda}{\sqrt{x_1^2 + \dots + x_n^2}} + 1 &= 0 \\ \sqrt{x_1^2 + \dots + x_n^2} - C &= 0 \end{aligned}$$

From the first n expressions we get

$$\begin{aligned} \lambda &= -\frac{\sqrt{x_1^2 + \dots + x_n^2}}{x_1} = \dots = \\ &= -\frac{\sqrt{x_1^2 + \dots + x_n^2}}{x_n}, \end{aligned}$$

hence $x_1 = x_2 = \dots = x_n$. Plugging into the last term we obtain

$$x_1 = x_2 = \dots = x_n = \frac{C}{\sqrt{n}} \quad (36)$$

Geometrically, x_i corresponds to the size of an edge of a hypercube embedded into a hypersphere of radius C .

Now let $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$ such that they have maximum ℓ_1 norm (Eq. 36) and their ℓ_2 norm is C (that is the output of function f after clipping in Eq. 13)

$$\begin{aligned} \mathbf{x} &= \left(-\frac{C}{\sqrt{n}}, \dots, -\frac{C}{\sqrt{n}} \right), \\ \mathbf{x}' &= \left(\frac{C}{\sqrt{n}}, \dots, \frac{C}{\sqrt{n}} \right) \end{aligned}$$

Then their ℓ_1 distance is

$$\begin{aligned} \|\mathbf{x} - \mathbf{x}'\|_1 &= \sum_{i=1}^n \left| -\frac{C}{\sqrt{n}} - \frac{C}{\sqrt{n}} \right| \\ &= n \cdot \left(\frac{2C}{\sqrt{n}} \right) = 2C\sqrt{n} \end{aligned} \quad (37)$$

□

C Potential attacks

Here we only sketch a potential attack on a single individual's privatized output \mathbf{v} . We do not speculate on the actual feasibility as differential privacy operates with the worst case scenario, that is the theoretical possibility that the adversary has unlimited compute power and unlimited background knowledge. However, real life examples show that anything less protective than DP can be attacked and it is mostly a matter of resources.⁶

We expect to have access to the trained ADePT autoencoder as well as the ATIS corpus (without the single individual whose value we try to infer, to be fair). We would need to find the privatized latent vector of \mathbf{v} , that is \mathbf{r}' , which could be possible by exploiting and probing the model. Second, by employing a brute-force attack, we can train a LM on ATIS to generate a feasible search space of input utterances, project them to the latent space, and explore the neighborhood of \mathbf{r}' . This would drastically reduce the search space. Then, depending on the geometric properties of that latent space, it might be the case that 'similar' utterances are closer to each other, increasing the probability of finding a similar utterance which might be a 'just good enough' approximation for the adversary.

⁶Diffix, a EU-based company, claimed their system is a better alternative to DP but did not provide formal guarantees for such claims. A paper from [Gadotti et al. \(2019\)](#) was a bitter lesson for Diffix, as it shows a successful attack. The bottom line is that without formal guarantees, it is impossible to prevent any future attacks.