Information, Computation & Quantization

Infinitely divisible privacy and beyond I
Resolution of the $s^2 = 2k$ conjecture

Aaradhya Pandey

Operations Research and Financial Engineering,
Princeton University

January 14, 2026

What is f-differential privacy?

What is Gaussian differential privacy?

What is Infinitely divisible privacy?

What is Poisson differential privacy?

Information, Computation and Quantization of differential privacy!

What is f-differential privacy?

A Lipschitz-type requirement on a Markov kernel or a Randomized algorithm.

Before differential privacy (Domain, Range and Examples of a Markov kernel)

- A **Markov kernel** is a map $M : \Omega_1 \times \mathscr{F}_2 \to [0,1]$ such that $\{M(w_1, \cdot)\}_{w_1 \in \Omega_1}$ is a measurable collection of probability measures on $(\Omega_2, \mathscr{F}_2)$.
- It is also known as a Channel, Cryptographic encryption, Randomized algorithm.
- Moreover, one can think of Markov kernels as a 'map' $(\Omega_1, \mathscr{F}_1) \xrightarrow{M} (\Omega_2, \mathscr{F}_2)$ that takes probability measures to probability measures, and functions to functions.

$$M(P_1)(A_2) := \int_{\Omega_1} M(w_1, A_2) dP_1(w_1) \quad M(f_2)(w_1) := \int_{\Omega_2} M(w_1, dw_2) f_2(w_2) \tag{1}$$

- **SRW:** $\Omega_1 = \Omega_2 = \mathbb{Z}$ with $\mathscr{F}_1 = \mathscr{F}_2 = 2^{\mathbb{Z}}$ and $M(w_1, \cdot) = p\delta_{w_1, w_1+1} + (1-p)\delta_{w_1, w_1-1}$.
- **Transition matrix:** Markov chain on a countable space $\Omega_1 = \Omega_2$, $\mathscr{F}_1 = \mathscr{F}_2 = 2^{\Omega_1}$ is given by a $\mathbb{N} \times \mathbb{N}$ transition 'matrix' $M(w_1, A_2) = \sum_{w_2 \in A_2} M(w_1, w_2)$.

Before differential privacy (Examples & a little algebra)

- **Normal:** $\Omega_1 = \Omega_2 = \mathbb{R}$, $\mathscr{F}_1 = \mathscr{F}_2 = \mathscr{B}_{\mathbb{R}}$, $M(\theta, A) := \int_{A-\theta} \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}} dx \equiv \{N(\theta, 1)\}_{\theta \in \mathbb{R}}$.
- **Convolve:** $\Omega_1 = \Omega_2 = \mathbb{G}$, $\mathscr{F}_1 = \mathscr{F}_2 = \mathscr{B}_{\mathbb{G}}$, $M(\theta, A) := \mu(A - \theta) \equiv \{\delta_\theta * \mu\}_{\theta \in \mathbb{G}}$.
- **Deterministic:** $M(w_1, A_2) = \mathbf{1}_{A_2}(f(w_1))$ for a function $(\Omega_1, \mathscr{F}_1) \xrightarrow{f} (\Omega_2, \mathscr{F}_2)$[1].
- One can do **algebra** as well as **analysis** on the 'class' of Markov kernels M.
- **Convex combination:** $\sum_{i=1}^k \lambda_i M_i$ for scalars $\lambda_i \geq 0, \sum_{i=1}^k \lambda_i = 1$ for $\Omega_1 \xrightarrow{M_i} \Omega_2$.
- **Composition:** $M_2 \circ M_1(w_1, A_3) := \int_{\Omega_2} M_1(w_1, dw_2) M_2(w_2, A_3)$ for $\Omega_1 \xrightarrow{M_1} \Omega_2 \xrightarrow{M_2} \Omega_3$.
- **Product:** $M_1 \otimes M_2(w_1, A_2 \times A_3) = M_1(w_1, A_2) M_2(w_1, A_3)$ for $\Omega_1 \xrightarrow{M_1} \Omega_2$, $\Omega_1 \xrightarrow{M_2} \Omega_3$.

[1] This is important in constructing Pseudorandom objects. See Luca Trevisan's lectures at Simons Institute.

Differential privacy– Step 0 (Domain and the Range)

- 'Every' new definition in statistics, CS is a requirement[2] on a Markov kernel M.
- Being Gaussian is also a requirement on an M! Can you guess why?
- **Differential privacy** is a specific statistical requirement on a Markov kernel M.
- '**Cryptographic security**' is a specific computational requirement on M.
- Now, we formalize the definition of **Differential privacy** $(\Omega_1, \mathscr{F}_1) \xrightarrow{M} (\Omega_2, \mathscr{F}_2)$.
- First, it requires us to define a symmetric relation $w \sim w'$ on Ω_1.
- In fact, the example of interest is a metric space (Ω_1, d_1) and $w \sim w' \leftrightarrow d_1(w, w') \leq 1$.
- We can further specialize to $\Omega_1 = \chi^k$ with $d_1(S, S') = \Delta(S, S') := \sum_{i=1}^{k} \mathbf{1}(S_i \neq S'_i)$.

[2] It also includes some more structure on the domain Ω_1 and range Ω_2 than just measurable spaces.

Differential privacy– Step 0 (The Lipschitz requirement, intuition)

- We can further specialize to $\Omega_1 = \chi^k$ with $d_1(S,S') = \Delta(S,S') := \sum_{i=1}^k \mathbf{1}(S_i \neq S'_i)$.

For $\Delta(S,S') \leq 1$ we want $M(S)$ and $M(S')$ to be statistically indistinguishable.

- It immediately suggests a computational analogue of differential privacy.
- We describe the notion of statistical indistinguishability through **Blackwell ordering**.
- Consider a hypothesis testing problem of two probability measures P, Q on $(\Omega_2, \mathscr{F}_2)$.
- For a test $(\Omega_2, \mathscr{F}_2) \xrightarrow{\varphi} ([0,1], \mathscr{B}_{[0,1]})$, **type I error=** $P(\varphi)$, **type II error=** $Q(1-\varphi)$.
- Consider the following trade-off function as a baseline notion of differential privacy.

$$f(\alpha) = T(P,Q)(\alpha) = \inf\{Q(1-\varphi) : P(\varphi) \leq \alpha, \varphi : (\Omega_2, \mathscr{F}_2) \xrightarrow{m} [0,1]\} \text{ for } \alpha \in [0,1].$$

Trade off functions and the definition of f-differential privacy

Theorem (Characterization of trade-off functions [2])

A function $f : [0,1] \to [0,1]$ is a trade-off function for some distributions P, Q on $(\Omega, \mathscr{F})$ if and only if f is convex, continuous , non-increasing, and $f(x) \leqslant 1 - x$ for $x \in [0,1]$.

Definition (f-differential privacy [2])

A Markov kernel $M : (\chi^k, \Delta) \to (\Omega_2, \mathscr{F}_2)$ is said to satisfy f-DP (for some TOF f) if

$$\inf_{\Delta(S,S') \leq 1} T(M(S), M(S')) \geq f \text{ pointwise everywhere on } [0,1]. \tag{2}$$

Theorem (Connections with the traditional definition [2])

A mechanism M is (ε, δ) DP if and only if M is $f_{\varepsilon,\delta} - DP$.

$$f_{\varepsilon,\delta}(\alpha) = \max\left\{0, 1 - \delta - e^{\varepsilon}\alpha, e^{-\varepsilon}(1 - \delta - \alpha)\right\} \tag{3}$$

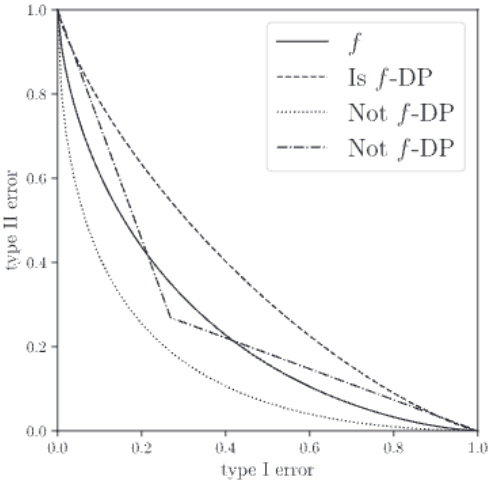
1.0
0.8
0.6
0.4
0.2
0.0
type II error
0.0
0.2
0.4
0.6
0.8
1.0
type I error
f
Is f-DP
Not f-DP
Not f-DP

Crash course on Hypothesis testing and ROC curves I

Definition (Trade-off function)

Given two probability distributions P, Q on a measurable space $(\Omega_2, \mathscr{F}_2)$, we define the *trade-off function* as the map $T(P,Q) : [0,1] \to [0,1]$ for $\alpha \in [0,1]$ as

$$T(P,Q)(\alpha) := \inf_{\varphi} \Big\{ Q(1-\varphi) \;\Big|\; P(\varphi) \leq \alpha, \varphi : (\Omega_2, \mathscr{F}_2) \xrightarrow{m} ([0,1], \mathscr{B}_{[0,1]}) \Big\}.$$

- For any **type I error** $P(\varphi)$, the trade-off function returns the smallest possible value of **type II error** $Q(1-\varphi)$ over all possible test functions $\varphi : (\Omega_2, \mathscr{F}_2) \xrightarrow{m} ([0,1], \mathscr{B}_{[0,1]})$.
- **Neyman-Pearson lemma:** Statistically optimal choice is the likelihood ratio test[3].

$$\varphi_\alpha(x) = \mathbf{1}\left(\log \frac{dQ}{dP} \geq z_\alpha \right) \text{ such that } P\left(\log \frac{dQ}{dP} \geq z_\alpha \right) = \alpha. \tag{4}$$

[3] Often we need to consider randomized likelihood ratio test. Think of polynomial threshold functions!

Blackwell ordering and ROC curves II

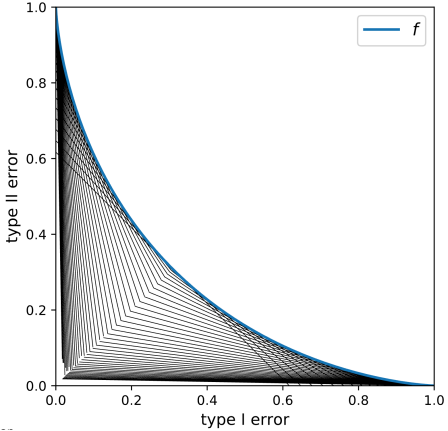

- **TOF:** A function $f : [0,1] \to [0,1]$ is a trade-off curve $T(P,Q)$ if and only if it is convex, continuous, non-increasing, and $f(\alpha) \leq 1 - \alpha$.
- **Blackwell ordering :** If $T_{P_0,Q_0}(\alpha) \geq T_{P_1,Q_1}(\alpha)$ for all α, then (P_1, Q_1) is easier to distinguish than (P_0, Q_0).
- **Complete indistinguishability:** $f(\alpha) = 1 - \alpha$ means random guess ROC with Bernoulli(α).

Blackwell's theorem and ROC curves III

Theorem (Equivalence of Blackwell informativeness and post-processing)

Let P_1, Q_1 be probability measures on Y_1 and P_0, Q_0 be probability measures on Y_0. The following two statements (Blackwell informativeness and post-processing) are equivalent:

1. ***Blackwell ordering:*** $T(P_0, Q_0)(\alpha) \geq T(P_1, Q_1)(\alpha)$ *for all* $\alpha \in [0,1]$.
2. ***Post-processing:*** $\exists$ *a Markov Kernel* $\mathrm{R}\colon Y_1 \to Y_0$ *such that* $(P_0, Q_0) = (\mathrm{R}(P_1), \mathrm{R}(Q_1))$.

- **Gaussian trade-off function:** Let Φ be the Gaussian CDF. Then we have $G_\theta(\alpha) := T(N(0,1), N(\theta,1))(\alpha) = \Phi(\Phi^{-1}(1-\alpha) - \theta)$ for $\theta \geq 0$.
- **Gaussian comparison:** $T(P,Q) \geq G_\theta$ means it is at least as difficult to distinguish the pair (P,Q) than it is to a pair of Normals with one having a shifted mean.

What is Gaussian differential privacy?

Gaussian baseline: $G_\theta(\alpha) := T(N(0,1), N(\theta,1))(\alpha) = \Phi(\Phi^{-1}(1-\alpha) - \theta)$ for $\theta \geq 0$.

Gaussian differential privacy, and a Composition theorem

Definition (Gaussian differential privacy [2])

A Markov kernel $M : (\chi^k, \Delta) \to (\Omega_2, \mathscr{F}_2)$ is said to satisfy θ-GDP (for some $\theta \geq 0$) if

$$\inf_{\Delta(S,S')\leq 1} T(M(S), M(S')) \geq G_\theta \text{ pointwise everywhere on } [0,1]. \tag{5}$$

Theorem (A general composition theorem [2])

If $M_1 : X^N \to Y$ is f-DP, $M_2 : X^N \times Y \to Z$ is g-DP for all $y \in Y$, then the joint mechanism

$$M = (M_1, M_2) : X^N \to Y \times Z \text{ defined as } M(x) = (M_1(x), M_2(x, M_1(x))) \text{ is } f \otimes g\text{-DP}. \tag{6}$$

- $T(P_1, Q_1) \otimes T(P_2, Q_2) := T(P_1 \otimes P_2, Q_1 \otimes Q_2)$ for $f = T(P_1, Q_1), g = T(P_2, Q_2)$.
- $f \otimes g$ is well-defined, commutative, associative, $f \otimes \mathrm{Id} = \mathrm{Id} \otimes f = f$, $\mathrm{Id}(\alpha) = 1 - \alpha$.

A central limit theorem for Gaussian differential privacy

$$\mathrm{kl}(f) := -\int_0^1 \log|f'(x)|\,dx, \quad \kappa_2(f) := \int_0^1 \log^2|f'(x)|\,dx, \quad \kappa_3(f) := \int_0^1 \left|\log|f'(x)|\right|^3 dx$$

Theorem (Central Limit theorem[2])

Let $\{f_{ni} : 1 \le i \le n\}_{n=1}^{\infty}$ be a triangular array of symmetric trade-off functions and assume the following limits for some constants $k \ge 0$ and $s > 0$ as $n \to \infty$:

1. $\sum_{i=1}^n \mathrm{kl}(f_{ni}) \to k$ *and* $\sum_{i=1}^n \kappa_2(f_{ni}) \to s^2$,
2. $\max_{1\le i\le n} \mathrm{kl}(f_{ni}) \to 0$ *and* $\sum_{i=1}^n \kappa_3(f_{ni}) \to 0$.

Then, we have the following convergence pointwise for all $\alpha \in [0,1]$.

$$\lim_{n\to\infty} f_{n1} \otimes f_{n2} \otimes \cdots \otimes f_{nn}(\alpha) = T(N(-k,s^2), N(k,s^2))(\alpha) \tag{7}$$

- Miraculously, 'in all examples' we always have $s^2 = 2k$. 'A proof of this interesting observation or the construction of a counterexample is left for future work'.

What is Infinitely divisible privacy?

- We determine all baseline notions of privacy, proposed in an HDSR Review paper.
- We resolve the $s^2 = 2k$ conjecture, proposed in the celebrated JRSSB paper on GDP.
- We connect Differential privacy with Le Cam's theory of convergence of experiments.
- We propose Poisson differential privacy, a new baseline suited for random graphs
- We provide an optimal mechanism for a given discretized statistic, such as degrees.

Infinitely divisible privacy– Step I (What about Poisson differential privacy)

- **DRS** showed Gaussian limits of composing (triangular) $f_i^{(n)}$-DP mechanisms.

$$\lim_{n\to\infty} f_1^{(n)} \otimes \cdots \otimes f_n^{(n)} = T\left(N\left(-\frac{s^2}{2}, s^2\right), N\left(\frac{s^2}{2}, s^2\right)\right) \text{ for some } s \in \mathbb{R}.$$

- $f_i^{(n)} = T\left(\mathrm{B}\left(\frac{\lambda_1}{n}\right), \mathrm{B}\left(\frac{\lambda_2}{n}\right)\right)$, then $\lim_{n\to\infty} f_1^{(n)} \otimes \cdots \otimes f_n^{(n)} = T(P(\lambda_1), P(\lambda_2))$!

What are all limit baseline TOFs?

Infinitely divisible privacy– Step II (combining Poisson and Gaussian baselines)

Theorem (An infinitely divisible extension [3])

Consider a sequence of trade-off functions (TOFs) $\{f_n = T(P_n, Q_n)\}_{n\in\mathbb{Z}_{\geq 1}}$ *with* $f_n(0) = 1$ *for all* $n \in \mathbb{Z}_{\geq 1}$ *and If* $f_n^{\otimes n}(\alpha) \to f_\infty(\alpha)$ *pointwise on* $[0,1]$ *for some TOF* $f_\infty = T(P_\infty, Q_\infty)$ *then* $f_\infty \in \mathscr{I}_T$ *(class of infinitely divisible trade-off functions).*

$$\mathscr{I}_T = \{f = T(P,Q) : P \text{ is infinitely divisible on } \mathbb{R} \text{ and } dQ(x) = e^x dP(x)\}$$

- P infinitely divisible $\leftrightarrow \log \hat{P}(t) = itm - \frac{1}{2}t^2\sigma^2 + \int_{\mathbb{R}}(e^{itx} - 1 - itx\mathbf{1}_{|x|\leq 1})d\nu(x)$, where ν a positive Borel measure on $\mathbb{R}\setminus 0$ satisfying $\int_{\mathbb{R}}(x^2 \wedge 1)d\nu(x) < \infty$.

$$P \overset{d}{=} \textbf{Gaussian part} + \textbf{Poisson part}$$

Combining Gaussian and Poisson– Resolution of the $s^2 = 2k$ conjecture of **DRS** [2]

- **Gaussians:** Let $Z \sim N(0,1)$. Then $f_\infty = G_\mu = T(N(0,1), N(\mu,1)) \in \mathscr{I}_T$ by taking

$$P_\infty = N\left(-\frac{\mu^2}{2}, \mu^2\right) \stackrel{d}{=} -\frac{\mu^2}{2} + |\mu| Z, \text{ and } dQ_\infty(x) = e^x dP_\infty(x) = N\left(\frac{\mu^2}{2}, \mu^2\right) \quad (8)$$

- **Poissons:** $T(c_1 X + c_2, c_1 Y + c_2) = T(X,Y)$. $f_\infty = T(P(\lambda_1), P(\lambda_2)) \in \mathscr{I}_T$ by taking

$$P_\infty \stackrel{d}{=} \lambda_1 - \lambda_2 + N \log(\lambda_2/\lambda_1), \text{ and } dQ_\infty(x) = e^x dP_\infty(x). \quad (9)$$

- Consider the limiting Gaussian case with $P_\infty \stackrel{d}{=} -k + |s| Z$ (no Poisson part).
- **DRS** asked whether it is always the case that $s^2 = 2k$? The answer is Yes, always.
- Because, we need $dQ_\infty(x) = e^x dP_\infty(x)$ to be a probability measure too, since $f_\infty = T(P_\infty, Q_\infty) \in \mathscr{I}_T$. So, $Q(\mathbb{R}) = P(\exp x) = \exp(-k + s^2/2) = 1 \implies s^2 = 2k$.

The usual additive noise mechanism achieving Gaussian Differential privacy

Theorem (Gaussian additive mechanism [2])

Given a metric space (Ω_1, d_1) of datasets and a real-valued 'continuous' statistic $g : (\Omega_1, d_1) \to \mathbb{R}$, with the privacy baseline $G_\mu = T(N(0,1), N(\mu,1))$ for some $\mu > 0$, consider the noise-adding mechanism (Markov Kernel) $M : (\mathbb{R}, \mathscr{B}_\mathbb{R}) \to (\mathbb{R}, \mathscr{B}_\mathbb{R})$ as $M(g(x)) \overset{d}{=} g(x) + cZ$ for $Z \sim N(0,1)$, so that $c^{-1} w_g(1) = \mu$.

$$\inf_{d(x,y)\leq 1} T(M(g(x)), M(g(y))) \geq T(Z, Z+\mu) = f_\infty \tag{10}$$

- Consequently, M satisfies μ GDP with $w_g(1) := \sup_{d(x,y)\leq 1} |g(x) - g(y)|$.
- But what if $g : (\Omega_1, d_1) \to \mathbb{Z}$ is 'discrete' and privacy baseline $f_\infty = T(P(\lambda_1), P(\lambda_2))$?
- Construct a mechanism $M : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ that satisfies Poisson differential privacy!

An Optimal Transport-based mechanism achieving Poisson Differential Privacy

Theorem ([3])

Given a metric space (Ω_1, d_1) of datasets and a statistic $g : (\Omega_1, d_1) \to \mathbb{Z}_{\geq 0}$ with the baseline Poisson trade-off function $f_\infty = T(P(\mu_1), P(\mu_2))$ for some $\mu_2 > \mu_1 > 0$. Consider the mechanism $M : (\mathbb{Z}_{\geq 0}, 2^{\mathbb{Z}_{\geq 0}}) \to (\mathbb{Z}_{\geq 0}, 2^{\mathbb{Z}_{\geq 0}})$ defined as

$$M(g(x)) \sim P(N_2 e^{N_1 g(x)}), \text{ with } N_1 = \frac{\log\left(\frac{\mu_2}{\mu_1}\right)}{w_g(1)} \text{ and } N_2 = \frac{|\mu_2 - \mu_1|}{w_{h\circ g}(1))}, \text{ where} \tag{11}$$

$$h : (\mathbb{R}, |\cdot|) \to (\mathbb{R}, |\cdot|) \quad h(y) = e^{N_1 y} = \left(\frac{\mu_2}{\mu_1}\right)^{\frac{y}{w_g(1)}}. \tag{12}$$

Then, M satisfy (symmetric) $f_{\mu_1,\mu_2} = \min(f_\infty, f_\infty^{-1})$ -DP. More precisely, we have

$$\inf_{d_1(x_1,x_2)\leq 1} T(M(g(x_1)), M(g(x_2))) \geq f_{\mu_1,\mu_2} = \min(f_\infty, f_\infty^{-1}). \tag{13}$$

Beyond Infinite divisibility

Theorem ([2] Every 'divergence' satisfying DPI is a functional of the TOF)

If $D(R(P), R(Q)) \leq D(P, Q)$ for probability distributions P, Q and Markov kernels R, then there exists a functional $l_D : \mathscr{F} \to \mathbb{R}$ such that $D(P, Q) = l_D(T(P, Q))$.

- One can go beyond infinitely divisible limits under compositions of a random number of nearly perfect differentially private operations.
- How do we incorporate computation into our framework?
- How do we incorporate quantization into our framework?

A sneak peak into computation I

- What happens when we shrink the $\sigma-$algebra $\mathscr{G} \subset \mathscr{F}$ for (P,Q) on $(\Omega,\mathscr{F})$?

$$f(\alpha) = T(P,Q)(\alpha) = \inf\{Q(1-\varphi) : P(\varphi) \leq \alpha, \varphi : (\Omega,\mathscr{F}) \xrightarrow{m} [0,1]\} \text{ for } \alpha \in [0,1].$$

Proposition 2. *Let* $P_0 = N(0,1), P_\mu = N(\mu,1)$ *for* $\mu > 0$ *on* $(\Omega,\mathcal{F}) \equiv (\mathbb{R},\mathcal{B}_{\mathbb{R}})$ *with a sub-*σ *algebra* $\mathcal{G} := \sigma([n,n+1) : n \in \mathbb{Z})$. *Then,* $T(P_0,P_\mu)(\alpha) = \Phi(\Phi^{-1}(1-\alpha)-\mu) \forall \alpha \in [0,1]$, *and*

$$T_{\mathcal{G}}(P_0,P_\mu)(\alpha_k) = \Phi(\Phi^{-1}(1-\alpha_k)-\mu) \text{ for } \alpha_k = 1-\Phi(k) \text{ for all } k \in \mathbb{Z}, \text{ and} \quad (111)$$

$$T_{\mathcal{G}}(P_0,P_\mu)(\alpha) = T_{\mathcal{G}}(P_0,P_\mu)(\alpha_{k+1}) - \lambda^* q_k, \text{ for } \alpha \in [\alpha_{k+1},\alpha_k] \text{ for all } k \in \mathbb{Z} \quad (112)$$

$$\text{and } \lambda^* = \frac{\alpha-\alpha_{k+1}}{\alpha_k-\alpha_{k+1}}, \text{ with } q_k = \Phi(k+1-\mu)-\Phi(k-\mu) \quad (113)$$

A sneak peak into computation II

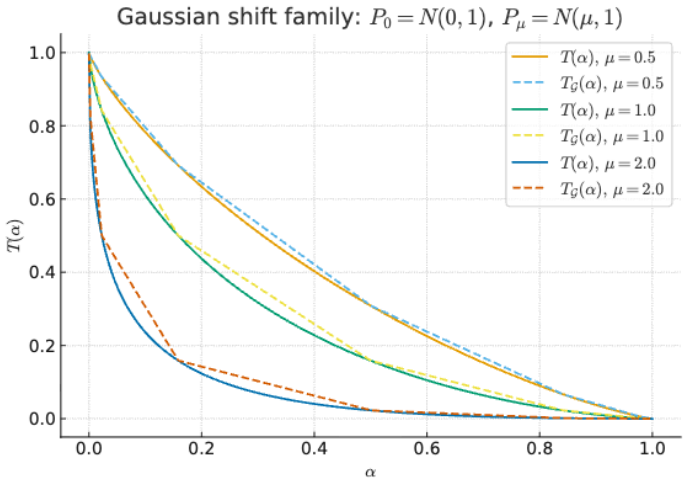

[1] Rachel Cummings and collaborators. “Advancing Differential Privacy: Where We Are Now and Future Directions for Real-World Deployment”. In: *Harvard Data Science Review* (2024). DOI: 10.1162/99608f92.d3197524. URL: https://doi.org/10.1162/99608f92.d3197524.

[2] Jinshuo Dong, Aaron Roth, and Weijie J. Su. “Gaussian Differential Privacy”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 84.1 (2022), pp. 3–37. DOI: 10.1111/rssb.12401.

[3] Aaradhya Pandey, Arian Maleki, and Sanjeev Kulkarni. *Infinitely divisible privacy and beyond I: resolution of the $s^2 = 2k$ conjecture*. Nov. 2025. DOI: 10.48550/arXiv.2512.00734. arXiv: 2512.00734 [math.ST]. URL: https://arxiv.org/abs/2512.00734.