Privacy and Geometry

Sasho Nikolov

University of Toronto

A Statistical Database and Linear Queries

- *Database*: $D \in \mathcal{U}^n$.
  - collection of n rows, one per individual
  - each row gives the *type* of the individual
  - universe $\mathcal{U}$: all possible types
- E.g. $\mathcal{U} = \{0,1\}^d$: each individual described by d binary attributes

A Statistical Database and Linear Queries

- *Database*: $D \in \mathcal{U}^n$.
  - collection of n rows, one per individual
  - each row gives the *type* of the individual
  - universe $\mathcal{U}$: all possible types
- E.g. $\mathcal{U} = \{0,1\}^d$: each individual described by d binary attributes
- *Linear query*: for $q : \mathcal{U} \to [0,1]$ and $D = \{r_1, \ldots, r_n\}$:

$$q(D) = \frac{1}{n} \sum_{i=1}^{n} q(r_i).$$

- E.g. *counting queries* when $q : \mathcal{U} \to \{0,1\}$.

A Statistical Database and Linear Queries

- *Database*: $D \in \mathcal{U}^n$.
  - collection of n rows, one per individual
  - each row gives the *type* of the individual
  - universe $\mathcal{U}$: all possible types
- E.g. $\mathcal{U} = \{0,1\}^d$: each individual described by d binary attributes
- *Linear query*: for $q : \mathcal{U} \to [0,1]$ and $D = \{r_1, \ldots, r_n\}$:

$$q(D) = \frac{1}{n} \sum_{i=1}^{n} q(r_i).$$

- E.g. *counting queries* when $q : \mathcal{U} \to \{0,1\}$.
- *Workload*: a collection of linear queries $\mathcal{Q} = \{q_1, \ldots, q_m\}$:

$$\mathcal{Q}(D) = \langle q_1(D), \ldots, q_m(D) \rangle.$$

Differential Privacy

Neighboring databases: D and D' that differ in at most one row.

Definition ([DMNS06])

An algorithm $\mathcal{A}$ is (ε, δ)-*differentially private* if for every two neighboring databases D, D', and every measurable subset S of the range of $\mathcal{A}$, $\mathcal{A}$ satisfies

$$\mathbb{P}[\mathcal{A}(D) \in S] \leq e^{\varepsilon}\mathbb{P}[\mathcal{A}(D') \in S] + \delta.$$

Measure of Error

Accuracy of algorithm $\mathcal{A}$ – *average (squared) error*:

$$\mathrm{Err}(\mathcal{A}, \mathcal{Q}, n) = \max_{D \in \mathcal{U}^n} \left(\mathbb{E} \frac{1}{m} \sum_{i=1}^{m} (\mathcal{A}(\mathcal{Q}, D)_i - q_i(D))^2 \right)^{1/2};$$

$$= \max_{D \in \mathcal{U}^n} \left(\mathbb{E} \frac{1}{m} \|\mathcal{A}(\mathcal{Q}, D) - \mathcal{Q}(D)\|_2^2 \right)^{1/2};$$

Measure of Error

Accuracy of algorithm $\mathcal{A}$ – *average (squared) error*:

$$\mathrm{Err}(\mathcal{A}, \mathcal{Q}, n) = \max_{D \in \mathcal{U}^n} \left(\mathbb{E} \frac{1}{m} \sum_{i=1}^{m} (\mathcal{A}(\mathcal{Q}, D)_i - q_i(D))^2 \right)^{1/2};$$

$$= \max_{D \in \mathcal{U}^n} \left(\mathbb{E} \frac{1}{m} \|\mathcal{A}(\mathcal{Q}, D) - \mathcal{Q}(D)\|_2^2 \right)^{1/2};$$

- Can define Err with worst case error, i.e. $\|\cdot\|_\infty$. But average error is natural for geometric techniques.

Sample Complexity

What is the smallest database size for which we can guarantee error at most α?

$$\mathsf{sc}(\mathcal{A}, \mathcal{Q}, \alpha) = \min\{n : \mathsf{Err}(\mathcal{A}, \mathcal{Q}, n) \leq \alpha\}$$
$$\mathsf{sc}_{\varepsilon,\delta}(\mathcal{Q}, \alpha) = \min\{\mathsf{sc}(\mathcal{A}, \mathcal{Q}, \alpha) : \mathcal{A} \text{ is } (\varepsilon, \delta) - \mathsf{DP}\}$$

Sample Complexity

What is the smallest database size for which we can guarantee error at most α?

$$\mathsf{sc}(\mathcal{A}, \mathcal{Q}, \alpha) = \min\{n : \mathsf{Err}(\mathcal{A}, \mathcal{Q}, n) \leq \alpha\}$$
$$\mathsf{sc}_{\varepsilon,\delta}(\mathcal{Q}, \alpha) = \min\{\mathsf{sc}(\mathcal{A}, \mathcal{Q}, \alpha) : \mathcal{A} \text{ is } (\varepsilon, \delta) - \mathsf{DP}\}$$

Goal:
Characterize the sample complexity in terms of natural properties of $\mathcal{Q}$.

- Understand the "hardness" of $\mathcal{Q}$.
- Understand "optimal" algorithms.

The Sensitivity Polytope

Sensitivity Polytope $K_{\mathcal{Q}}$ [HT10]:

- convex hull of $\pm\mathcal{Q}(D)$ for all databases D of size $n = 1$
- D and D' neighboring $\Leftrightarrow |D| \cdot \mathcal{Q}(D) - |D'| \cdot \mathcal{Q}(D') \in K_{\mathcal{Q}}$.
- Describes how answers change between neighboring databases.

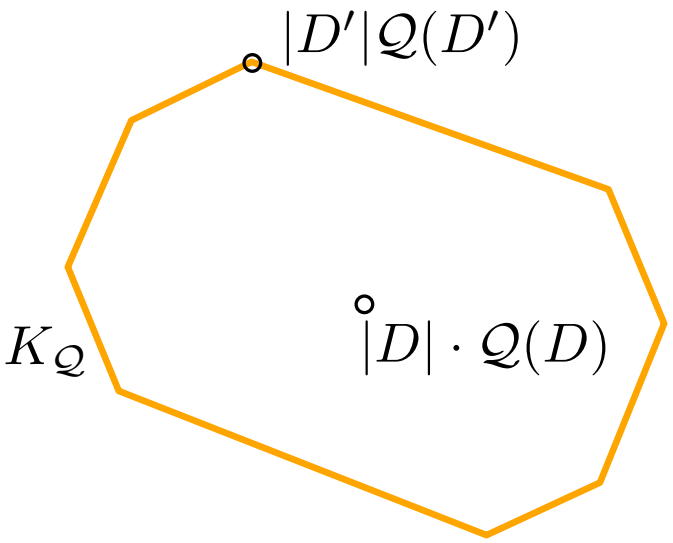

The Sensitivity Polytope

Sensitivity Polytope $K_{\mathcal{Q}}$ [HT10]:

- convex hull of $\pm\mathcal{Q}(D)$ for all databases D of size $n = 1$
- D and D' neighboring $\Leftrightarrow |D| \cdot \mathcal{Q}(D) - |D'| \cdot \mathcal{Q}(D') \in K_{\mathcal{Q}}$.
- Describes how answers change between neighboring databases.

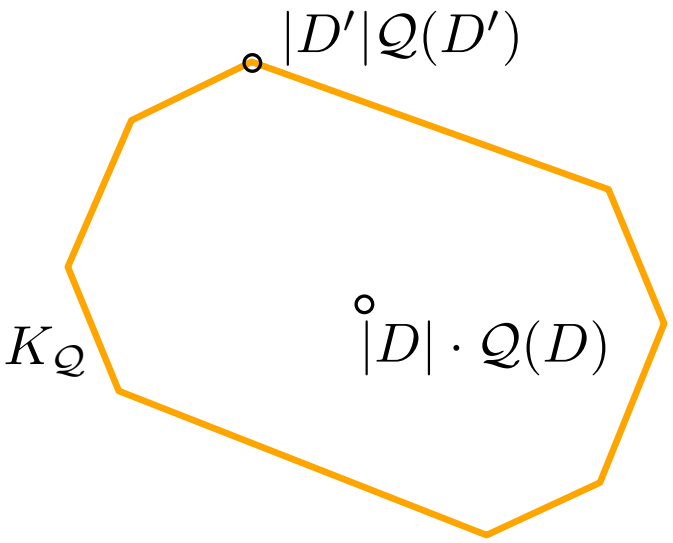

- Identify geometric measures of the **"size"** of $K_{\mathcal{Q}}$ that characterize sample complexity/error.

Mean Point Problem

- Can think of $D = \{r_1, \ldots, r_n\} \subseteq \mathcal{U}$ as $\{x_1, \ldots, x_n\} \subset K_{\mathcal{Q}}$:

$$x_i = \mathcal{Q}(\{r_i\}) \qquad\qquad \mathcal{Q}(D) = \frac{1}{n}\sum_{i=1}^{n} x_i \in K_{\mathcal{Q}}.$$

- From now on we treat the Mean Point Problem (MPP):
  - **Input**: $\{x_1, \ldots, x_n\} \subset K$
  - **Approximate**: $\frac{1}{n}\sum_{i=1}^{n} x_i$.
- We will assume $K \subseteq [0,1]^m$. Sample complexity is $\mathrm{sc}_{\varepsilon,\delta}(K, \alpha)$.

Mean Point Problem

- Can think of $D = \{r_1, \ldots, r_n\} \subseteq \mathcal{U}$ as $\{x_1, \ldots, x_n\} \subset K_{\mathcal{Q}}$:

$$x_i = \mathcal{Q}(\{r_i\}) \qquad\qquad \mathcal{Q}(D) = \frac{1}{n}\sum_{i=1}^{n} x_i \in K_{\mathcal{Q}}.$$

- From now on we treat the Mean Point Problem (MPP):
  - **Input**: $\{x_1, \ldots, x_n\} \subset K$
  - **Approximate**: $\frac{1}{n}\sum_{i=1}^{n} x_i$.
- We will assume $K \subseteq [0,1]^m$. Sample complexity is $\mathrm{sc}_{\varepsilon,\delta}(K, \alpha)$.
- Algorithms for the MPP imply algorithms for query release.

Mean Point Problem

- Can think of $D = \{r_1, \ldots, r_n\} \subseteq \mathcal{U}$ as $\{x_1, \ldots, x_n\} \subset K_{\mathcal{Q}}$:

$$x_i = \mathcal{Q}(\{r_i\}) \qquad\qquad \mathcal{Q}(D) = \frac{1}{n}\sum_{i=1}^{n} x_i \in K_{\mathcal{Q}}.$$

- From now on we treat the Mean Point Problem (MPP):
  - **Input**: $\{x_1, \ldots, x_n\} \subset K$
  - **Approximate**: $\frac{1}{n}\sum_{i=1}^{n} x_i$.
- We will assume $K \subseteq [0,1]^m$. Sample complexity is $\mathrm{sc}_{\varepsilon,\delta}(K, \alpha)$.
- Algorithms for the MPP imply algorithms for query release.
- Sample complexity lower bounds for MPP imply lower bounds for query release, up to losing a $\mathrm{poly}(\frac{1}{\alpha})$ factor.

Packing

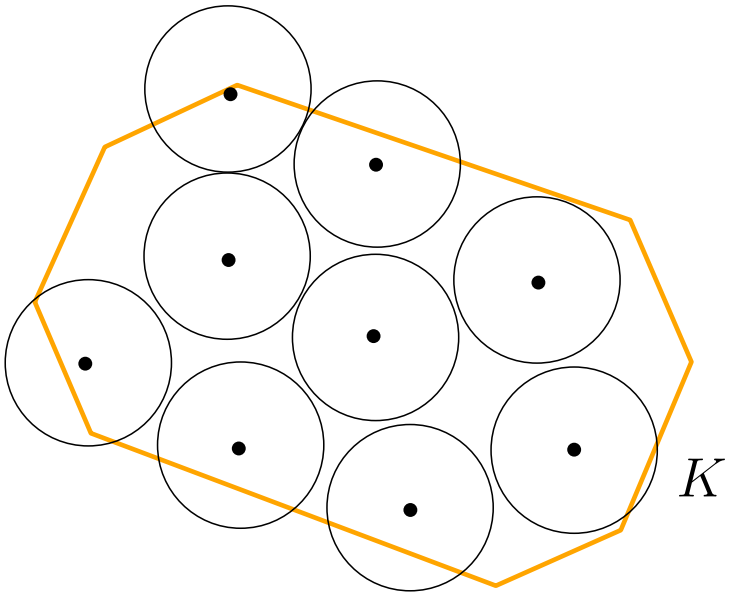

- α-*Packing*: $y_1, \ldots, y_N \in K$ s.t. $i \neq j \implies \|y_i - y_j\|_2 \geq 2\alpha$.
- $N(K, \alpha) =$ size of the largest α-packing.

Packing

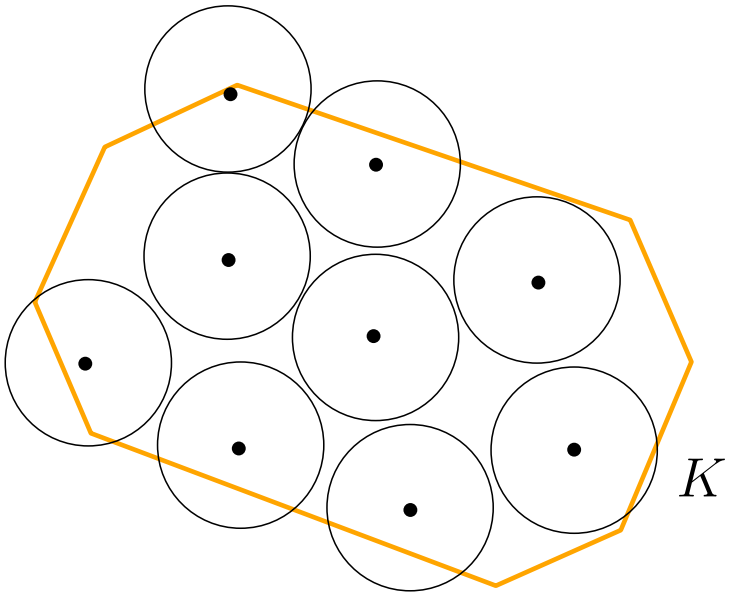

- α-*Packing*: $y_1, \ldots, y_N \in K$ s.t. $i \neq j \implies \|y_i - y_j\|_2 \geq 2\alpha$.
- $N(K, \alpha) =$ size of the largest α-packing.
- *Claim*: $\mathrm{sc}_{\varepsilon,0}(K, \alpha) \geq \frac{1}{\varepsilon} \log(N(K, 2\alpha\sqrt{m})/2)$

Packing Lower Bound

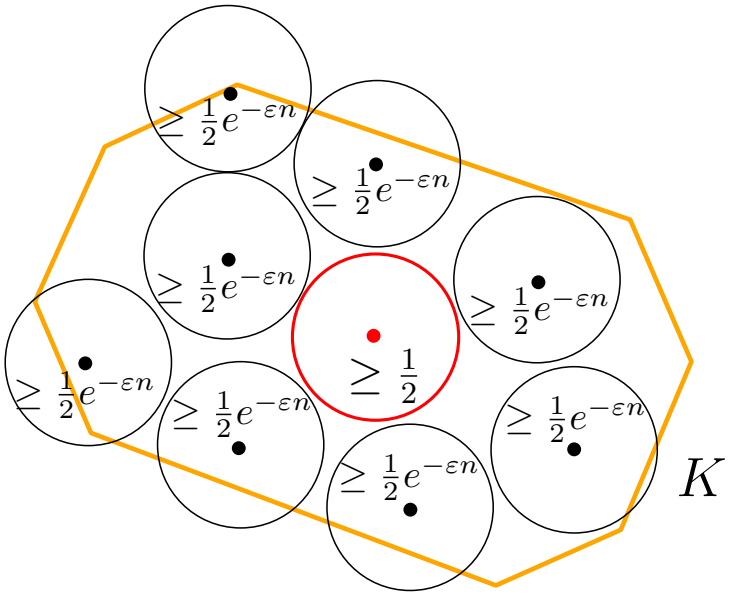

- $D_i = \{y_i, y_i \ldots, y_i\}$, where $\{y_1, \ldots, y_N\}$ is a $2\alpha\sqrt{m}$-packing.

Packing Lower Bound

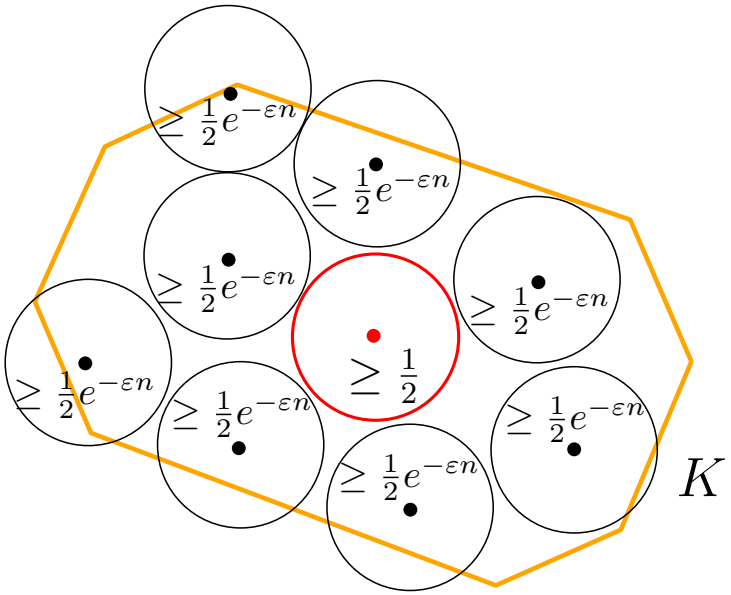

- $D_i = \{y_i, y_i \ldots, y_i\}$, where $\{y_1, \ldots, y_N\}$ is a $2\alpha\sqrt{m}$-packing.
- Assume $\mathbb{P}(\frac{\|\mathcal{A}(D_i) - y_i\|_2}{\sqrt{m}} \leq 2\alpha) \geq \frac{1}{2}$ for all i.

Packing Lower Bound

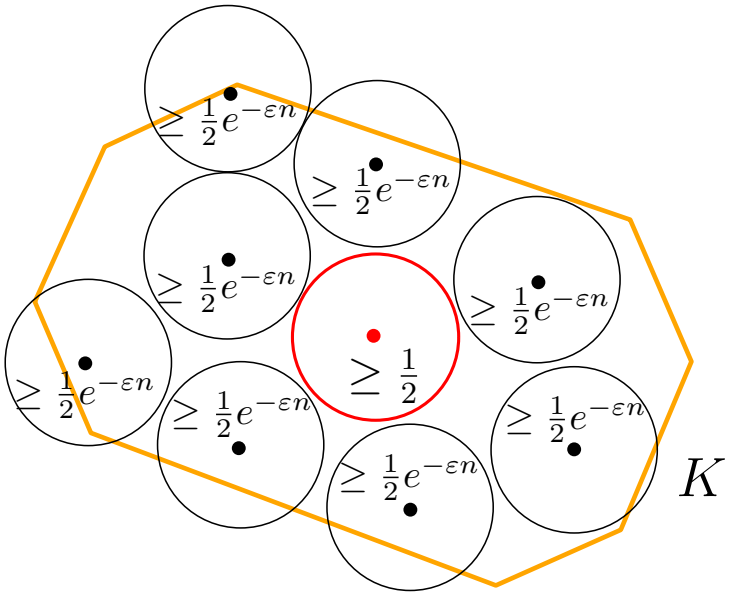

- $D_i = \{y_i, y_i \ldots, y_i\}$, where $\{y_1, \ldots, y_N\}$ is a $2\alpha\sqrt{m}$-packing.
- Assume $\mathbb{P}(\frac{\|\mathcal{A}(D_i)-y_i\|_2}{\sqrt{m}} \leq 2\alpha) \geq \frac{1}{2}$ for all i. By *group privacy*,

$$\mathbb{P}\left(\frac{\|\mathcal{A}(D_1)-y_i\|_2}{\sqrt{m}} \leq 2\alpha\right) \geq e^{-\varepsilon n}\mathbb{P}\left(\frac{\|\mathcal{A}(D_i)-y_i\|_2}{\sqrt{m}} \leq 2\alpha\right) \geq \frac{1}{2}e^{-\varepsilon n}.$$

Packing Lower Bound

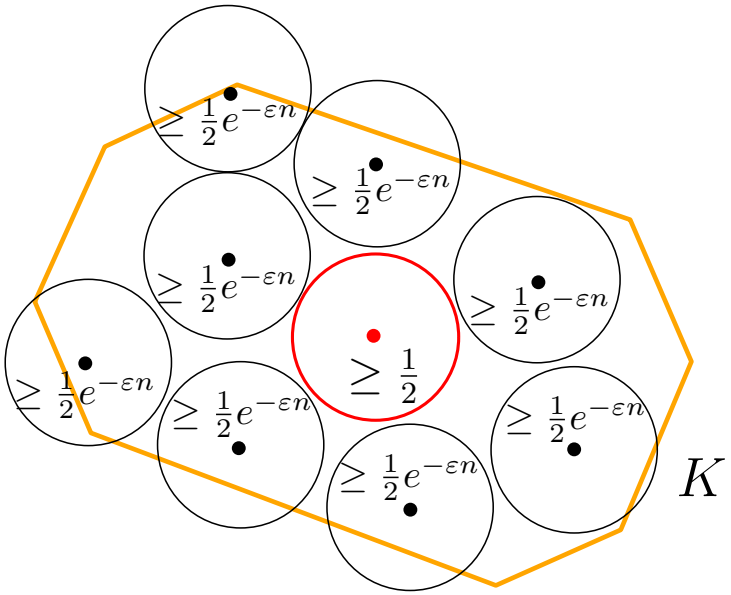

- $D_i = \{y_i, y_i \ldots, y_i\}$, where $\{y_1, \ldots, y_N\}$ is a $2\alpha\sqrt{m}$-packing.
- Assume $\mathbb{P}(\frac{\|\mathcal{A}(D_i)-y_i\|_2}{\sqrt{m}} \leq 2\alpha) \geq \frac{1}{2}$ for all i. By *group privacy*,

$$\mathbb{P}\left(\frac{\|\mathcal{A}(D_1)-y_i\|_2}{\sqrt{m}} \leq 2\alpha\right) \geq e^{-\varepsilon n}\mathbb{P}\left(\frac{\|\mathcal{A}(D_i)-y_i\|_2}{\sqrt{m}} \leq 2\alpha\right) \geq \frac{1}{2}e^{-\varepsilon n}.$$

- Events are disjoint: $\frac{1}{2}Ne^{-\varepsilon n} \leq 1$.

Tightness of the Lower Bound

Using the exponential mechanism:

$$\frac{1}{\varepsilon}\log N(K, 2\alpha\sqrt{m}) \lesssim \mathrm{sc}_{\varepsilon,0}(K,\alpha) \lesssim \frac{1}{\varepsilon}\log N(K, \alpha\sqrt{m}/4)$$

Tightness of the Lower Bound

Using the exponential mechanism:

$$\frac{1}{\varepsilon}\log N(K, 2\alpha\sqrt{m}) \lesssim \mathrm{sc}_{\varepsilon,0}(K,\alpha) \lesssim \frac{1}{\varepsilon}\log N(K, \alpha\sqrt{m}/4)$$

- The lower bounds hold for MPP, not query release. We lose a $\frac{1}{\alpha^2}$ factor for query release.
- The upper bound is certified by an exponential time algorithm. A polynomial time algorithm [HT10, BDKT12, NTZ13] achieves the same result up to factors $\mathrm{poly}(\log m, \log|\mathcal{U}|)$.)

Tightness of the Lower Bound

Using the exponential mechanism:

$$\frac{1}{\varepsilon}\log N(K, 2\alpha\sqrt{m}) \lesssim \mathrm{sc}_{\varepsilon,0}(K,\alpha) \lesssim \frac{1}{\varepsilon}\log N(K, \alpha\sqrt{m}/4)$$

- The lower bounds hold for MPP, not query release. We lose a $\frac{1}{\alpha^2}$ factor for query release.
- The upper bound is certified by an exponential time algorithm. A polynomial time algorithm [HT10, BDKT12, NTZ13] achieves the same result up to factors $\mathrm{poly}(\log m, \log|\mathcal{U}|)$.)
- *Recent*: Analogous characterization using packing numbers for Concentrated Differential Privacy. (with Blasiok, Bun, Kattis, Steinke)

Kolmogorov width

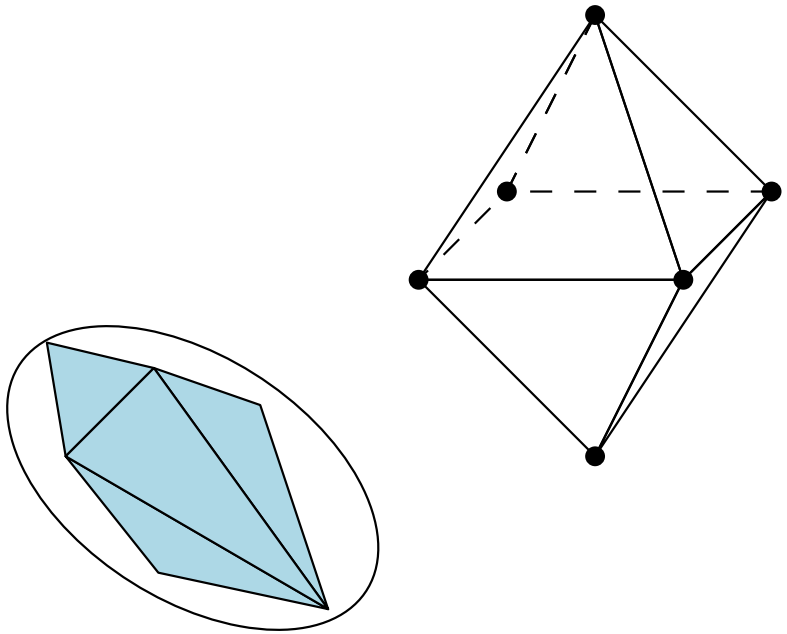

$$d_k(K) = \inf\{\mathrm{radius}(P_W K) : \mathrm{co}-\dim(W) < k\},$$

where P_W is the orhtogonal projection onto W.

Kolmogorov width and Approximate DP

[NTZ13, Nik15]:

$$\max_{k \le \varepsilon n} \frac{\sqrt{k} d_k(K)}{\varepsilon \sqrt{mn}} \lesssim \inf\{ \mathrm{Err}(\mathcal{A}, K, n) : \mathcal{A} \text{ is } (\varepsilon, \delta) - DP\}$$

$$\lesssim (\log n)(\log 1/\delta)^{1/4}(\log |\mathcal{U}|)^{1/4} \cdot \max_{k \le \varepsilon n} \frac{\sqrt{k} d_k(K)}{\varepsilon \sqrt{mn}}$$

Kolmogorov width and Approximate DP

[NTZ13, Nik15]:

$$\max_{k \le \varepsilon n} \frac{\sqrt{k} d_k(K)}{\varepsilon \sqrt{mn}} \lesssim \inf\{ \mathrm{Err}(\mathcal{A}, K, n) : \mathcal{A} \text{ is } (\varepsilon, \delta) - DP\}$$

$$\lesssim (\log n)(\log 1/\delta)^{1/4}(\log |\mathcal{U}|)^{1/4} \cdot \max_{k \le \varepsilon n} \frac{\sqrt{k} d_k(K)}{\varepsilon \sqrt{mn}}$$

- The lower bound is a reconstruction attack.
- The upper bound is an efficient algorithm. (Running time polynomial in $n, m, |\mathcal{U}|$.

Kolmogorov width and Approximate DP

[NTZ13, Nik15]:

$$\max_{k \le \varepsilon n} \frac{\sqrt{k} d_k(K)}{\varepsilon \sqrt{mn}} \lesssim \inf\{ \mathrm{Err}(\mathcal{A}, K, n) : \mathcal{A} \text{ is } (\varepsilon, \delta) - DP\}$$

$$\lesssim (\log n)(\log 1/\delta)^{1/4}(\log |\mathcal{U}|)^{1/4} \cdot \max_{k \le \varepsilon n} \frac{\sqrt{k} d_k(K)}{\varepsilon \sqrt{mn}}$$

- The lower bound is a reconstruction attack.
- The upper bound is an efficient algorithm. (Running time polynomial in $n, m, |\mathcal{U}|$.
- The $(\log |\mathcal{U}|)^{1/4}$ gap is unavoidable if using Kolmogorov width. A geometric lower bound based on fingerprinting code attacks may be stronger for constant α [KN17].

The K-norm Mechanism[HT10]

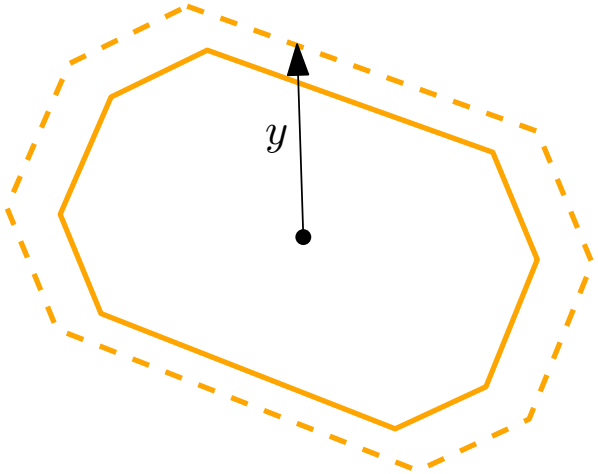

Output $\tilde{x}$ with density $p(z) \propto \exp(\varepsilon \|z - x\|_K)$, where

$$\|y\|_K = \inf\{t : y \in tK\},$$

is the smallest scaling of K that contains y.

The K-norm Mechanism[HT10]

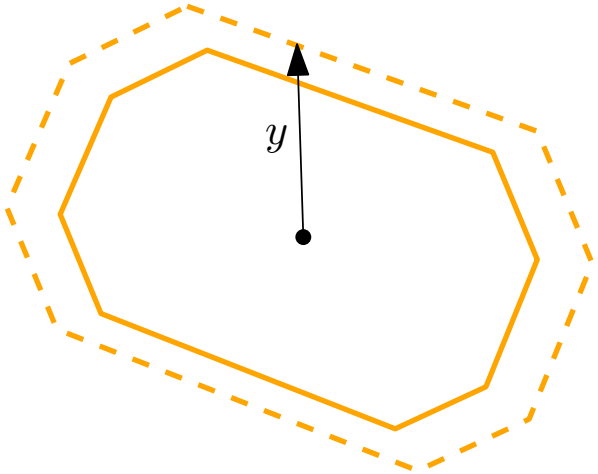

Output $\tilde{x}$ with density $p(z) \propto \exp(\varepsilon \|z - x\|_K)$, where

$$\|y\|_K = \inf\{t : y \in tK\},$$

is the smallest scaling of K that contains y.

- The error is controlled by $\left(\mathbb{E}_{y \sim K} \|y\|_2^2\right)^{1/2}$.
- Deep results in convex geometry (the *slicing problem*) relate this quantity to the volume of K and packing numbers.

The Projection Algorithm

Gaussian Noise Algorithm [DN03, DN04, DKM+06].

1. $x = \frac{1}{n}\sum_{i=1}^{n} x_i$
2. Sample $w \sim N(0, \frac{m}{n^2}\sigma^2_{\varepsilon,\delta})^m$, $\sigma^2_{\varepsilon,\delta} = O\left(\frac{\log(1/\delta)}{\varepsilon^2}\right)$.
3. Output $\tilde{x} = x + w$.

Bad when $n \ll m$!

Idea: We know $x \in K$. Use that!

The Projection Algorithm

Projection Algorithm $\mathcal{A}_{\text{proj}}$ [NTZ13]

1. $x = \frac{1}{n}\sum_{i=1}^{n} x_i$
2. Compute $\tilde{x} = x + w$, $w \sim N(0, \frac{m}{n^2}\sigma^2_{\varepsilon,\delta})^m$.
3. Output $\hat{x} = \arg\min\{\|x - \tilde{x}\|_2 : x \in K\}$.

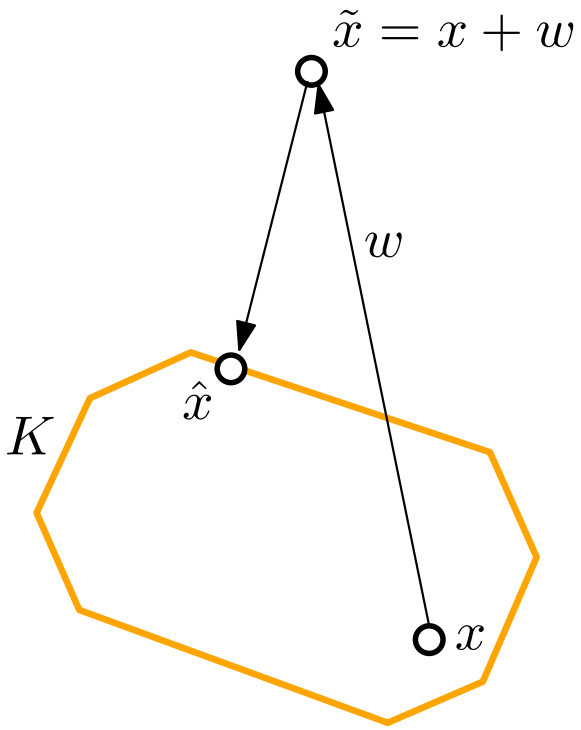

Bounding the Error

Support function of $K \subseteq \mathbb{R}^m$: $h_K(y) = \max\{\langle x, y\rangle : x \in K\}$.

- measures width of K in the direction of y

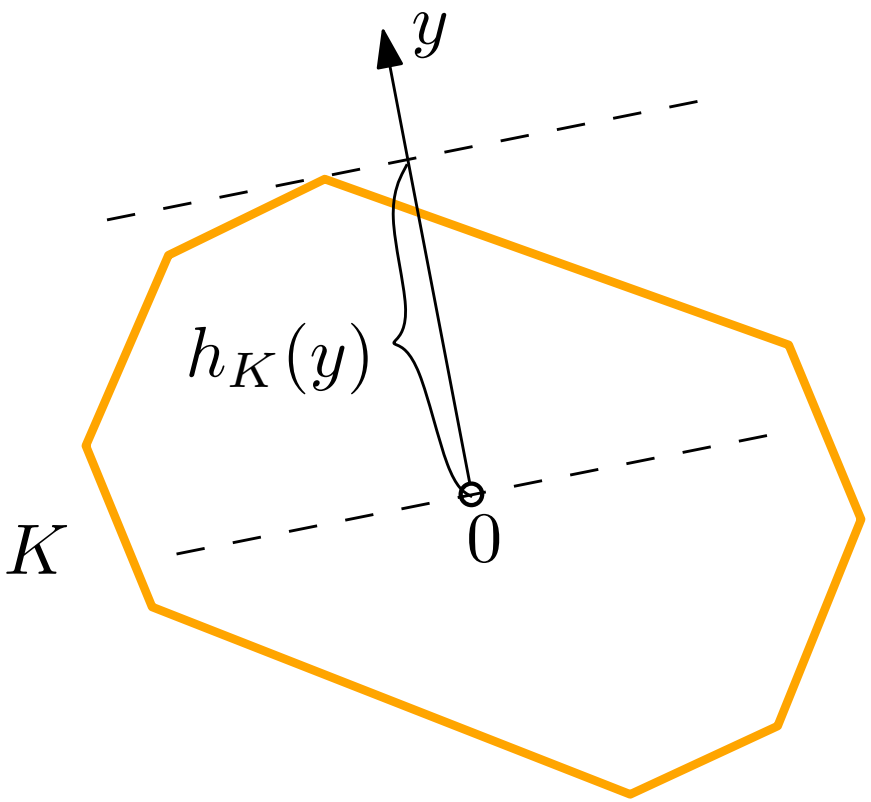

Bounding the Error

Support function of $K \subseteq \mathbb{R}^m$: $h_K(y) = \max\{\langle x, y\rangle : x \in K\}$.

- measures width of K in the direction of y

Main Fact:

$$\begin{aligned}\|\hat{x} - x\|_2^2 &\leq |\langle \hat{x} - x, w\rangle| \\ &\leq 2h_K(w).\end{aligned}$$

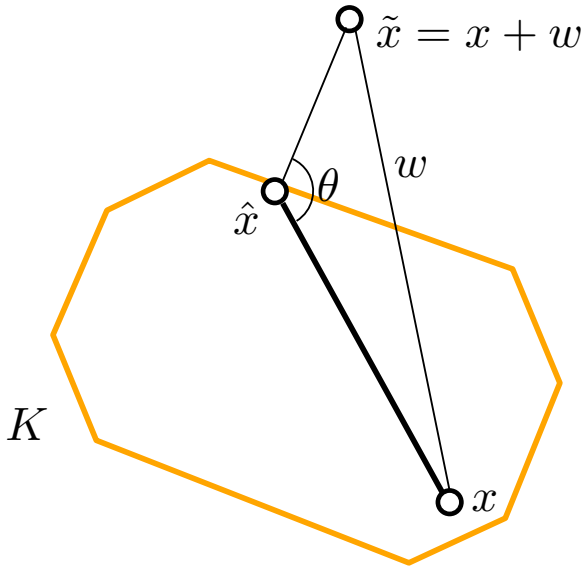

Key observation: θ is obtuse.

Bounding the Error

Support function of $K \subseteq \mathbb{R}^m$: $h_K(y) = \max\{\langle x, y\rangle : x \in K\}$.

- measures width of K in the direction of y

Main Fact:

$$\|\hat{x} - x\|_2^2 \leq |\langle \hat{x} - x, w\rangle| \leq 2h_K(w).$$

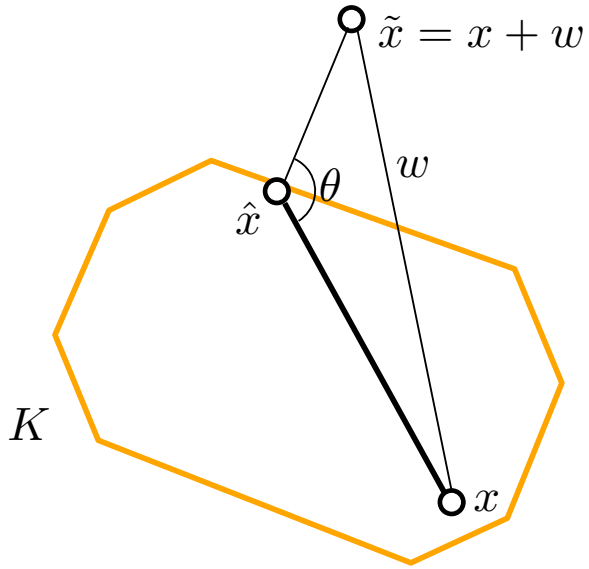

Key observation: θ is obtuse.

Error controlled by *mean width* $M^*(K) = \mathbb{E}_z h_K(z)$ where z is a random unit vector:

$$\mathrm{sc}_{\varepsilon,\delta}(\mathcal{A}_{\mathrm{proj}}, K, \alpha) \lesssim \frac{M^*(K)\sqrt{\log(1/\delta)}}{\varepsilon\alpha^2}.$$

Conclusion

Geometric viewpoint on differential privacy:

- the sensitivity polytope $K_{\mathcal{Q}}$ gives a geometric picture of the sensitivity of the queries;
- geometric measures of the "size" of $K_{\mathcal{Q}}$ that characterize sample complexity;
- geometric tools to design and analyze *algorithms*.

Conclusion

Geometric viewpoint on differential privacy:

- the sensitivity polytope $K_{\mathcal{Q}}$ gives a geometric picture of the sensitivity of the queries;
- geometric measures of the "size" of $K_{\mathcal{Q}}$ that characterize sample complexity;
- geometric tools to design and analyze *algorithms*.

Many open problems:

- Close gaps between upper and lower bounds.
- Data dependent results.
- Extend the theory to *non-linear* queries (convex optimization).
- Extend the theory to *interactive* mechanisms.

Aditya Bhaskara, Daniel Dadush, Ravishankar Krishnaswamy, and Kunal Talwar.
Unconditional differentially private mechanisms for linear queries.
In *Proceedings of the 44th symposium on Theory of Computing*, STOC '12, pages 1269–1284, New York, NY, USA, 2012. ACM.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor.
Our data, ourselves: Privacy via distributed noise generation.
4004:486–503, 2006.

C. Dwork, F. Mcsherry, K. Nissim, and A. Smith.
Calibrating noise to sensitivity in private data analysis.
In *TCC*, 2006.

Irit Dinur and Kobbi Nissim.
Revealing information while preserving privacy.
pages 202–210, 2003.

Cynthia Dwork and Kobbi Nissim.

Privacy-preserving datamining on vertically partitioned databases.
In *Advances in Cryptology–CRYPTO 2004*, pages 528–544. Springer, 2004.

Moritz Hardt and Kunal Talwar.
On the geometry of differential privacy.
In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 705–714, New York, NY, USA, 2010. ACM.

Assimakis Kattis and Aleksandar Nikolov.
Lower bounds for differential privacy from gaussian width.
To appear in SoCG 2017, 2017.

Aleksandar Nikolov.
An improved private mechanism for small databases.
In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 1010–1021. Springer, 2015.

Aleksandar Nikolov, Kunal Talwar, and Li Zhang.
The geometry of differential privacy: the sparse and approximate cases.
In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 351–360. ACM, 2013.