Locally Private Histograms in All Privacy Regimes

Clément Canonne
(University of Sydney)

Joint work with Abigail Gentle

SIMONS INSTITUTE for the Theory of Computing

The (distributed) setting

- Users have data (observations)
- Center wants to learn from this data

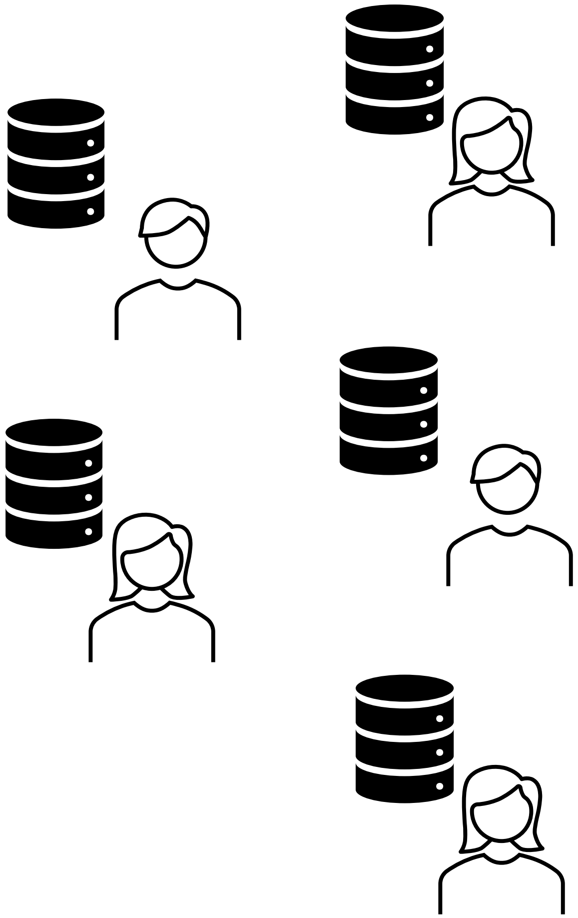

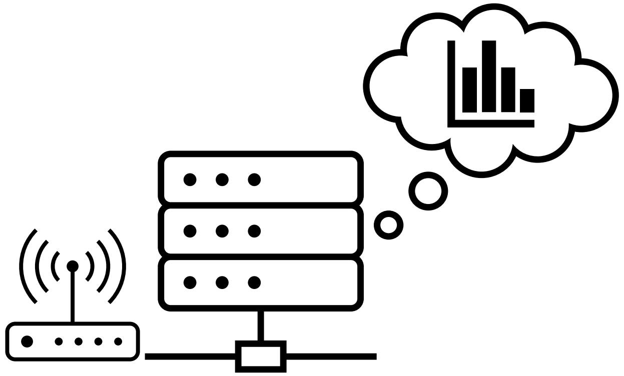

The (distributed) setting: example

- Users have data $X_1,\ldots,X_n$

  - Personal data
  - Observations i.i.d. from some distribution p

- Can we learn the counts/frequencies/histogram?
  - To learn about the users' preferences
  - To learn about the underlying data distribution

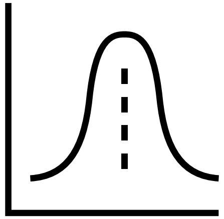

The (distributed) setting: example

- Users have data $X_1,\ldots,X_n$

  - Personal data
  - Observations i.i.d. from some distribution p

- Can we learn the counts/frequencies/histogram?

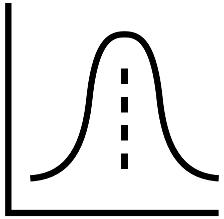

  - To learn about the users' preferences
  - To learn about the underlying data distribution

The (distributed) setting

- Users have data (observations) $X_1,\ldots,X_n$ (fixed or i.i.d.)
- Users have constraints: privacy, bandwidth, ...
- Center has goals: e.g., maximize utility given n users (minimize number)

The (distributed) setting

- Users have data (observations) $X_1,\ldots,X_n$ (fixed or i.i.d.)
- Users have constraints: **privacy**, bandwidth, ...
- Center has goals: e.g., maximize utility given n users (minimize number)

What is **private**?

- (Central) Privacy: Trust the Center
- Local Privacy: Trust Nobody
- Shuffle Privacy: Trust The Middle Box

Three variants of
Differential Privacy

What is **private**?

- (Central) Privacy: Trust the Center
- Local Privacy: Trust Nobody
- Shuffle Privacy: Trust The Middle Box

Three variants of
Differential Privacy

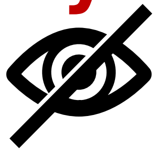

Local Differential privacy [DMNS06, KLNRS11]

$$\forall x, x', \ \forall S$$

$$\mathbb{P}\{R(x) \in S\} \le \underbrace{e^{\varepsilon}}_{\approx 1 + \varepsilon \ (?)} \mathbb{P}\{R(x') \in S\}$$

x

$R(x)$

Histograms

private: user j holds x_j

$x_1, \ldots, x_n \in [k]$

$f_i = \frac{1}{n}\sum_{j=1}^{n} \mathbb{1}_{x_j = i}, \quad i \in [k]$

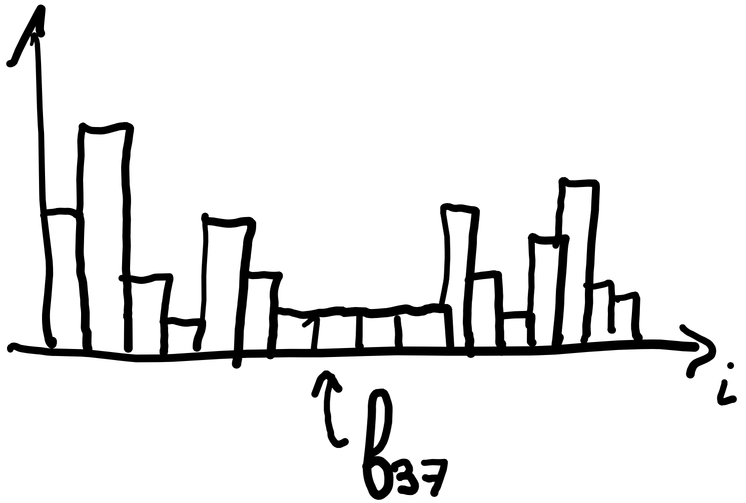

Goal: center outputs $\hat{f}_1, \ldots, \hat{f}_k$ minimizing

$$\mathbb{E}\left[\|\hat{f} - f\|_\infty\right]$$

over protocol's randomness

Histograms

Best $\mathbb{E}[\|\hat{f}-f\|_\infty] = \Theta\left(\sqrt{\frac{\log k}{n\varepsilon^2}}\right)$

Histograms

Best $\mathbb{E}[\|\hat{f}-f\|_\infty] = \Theta\left(\sqrt{\frac{\log k}{n\varepsilon^2}}\right)$

(and this is achieved by several protocols, including some* with $O(\log k)$ bits from each user)

*e.g., [Acharya-Sun '19]

Done!

Histograms

Best $\mathbb{E}\left[\|\hat{f}-f\|_\infty\right] = \Theta\left(\sqrt{\frac{\log k}{n\varepsilon^2}}\right)$

for $\varepsilon \ll 1$

("high privacy")

Histograms

Best $\mathbb{E}[\|\hat{f}-f\|_\infty] = ?$ for $\varepsilon \gg 1$

(and why should we care?)

Histograms: what do we need?

- Simple, efficient protocols
- Low communication
- Non-interactive
- No public randomness
- Good error…
- … including in the low-privacy regime

Histograms: what do we need?

For Lookup Hints, Apple uses a privacy budget with epsilon of 4, and limits user contributions to two donations per day. For emoji, Apple uses a privacy budget with epsilon of 4, and submits one donation per day. For QuickType, Apple uses a privacy budget with epsilon of 8, and submits two donations per day.

For Health types, Apple uses a privacy budget with epsilon of 2 and limits user contributions to one donation per day. The donations do not include health information itself, but rather which health data types are being edited by users.

For Safari, Apple limits user contributions to 2 donations per day. For Safari domains identified as causing high energy use or crashes, Apple uses a single privacy budget with epsilon of 4. For Safari Auto-play intent detection, Apple uses a privacy budget with epsilon of 8.

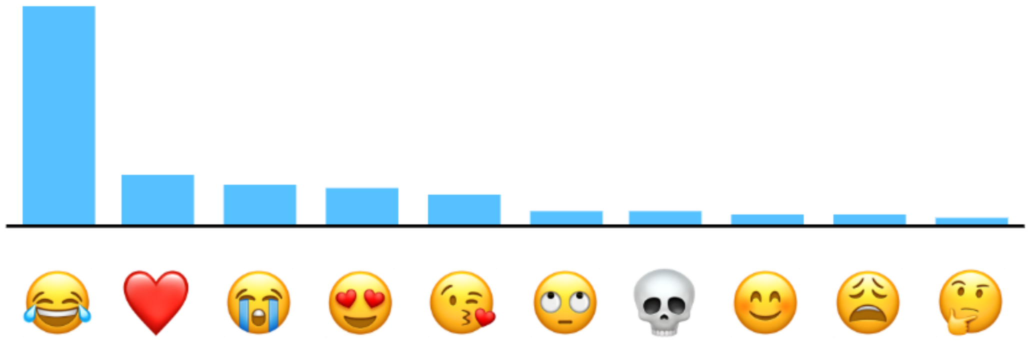

https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

A detour (?)

Let $X_1, \dots, X_k$ be σ^2-subgaussian r.v.s.

Then

$$\mathbb{E}\left[\max_{1 \leq i \leq k} |X_i|\right] \leq \sqrt{2\sigma^2 \log k}$$

A detour (?)

Let $X_1, \dots, X_k$ be σ^2-subgaussian r.v.s.

Then

$$\mathbb{E}\left[\max_{1\le i\le k} X_i\right] \le \sqrt{2\sigma^2 \log k}$$

<u>Proof.</u> $\forall t>0$, $\mathbb{E}\max_i X_i \le \frac{1}{t}\mathbb{E}\log e^{t\max_i X_i} \le \frac{1}{t}\log \mathbb{E}\max_i e^{tX_i}$

$\le \frac{1}{t}\log \sum_i \mathbb{E}e^{tX_i} \le \frac{1}{t}\log\left(n e^{\frac{t^2\sigma^2}{2}}\right)$

$= \frac{\log n}{t} + \frac{\sigma^2}{2}t.$ Choose $t = \sqrt{\frac{2\log n}{\sigma^2}}$. $\square$

A detour (?)

<u>Hoeffding's Lemma.</u>

If $X \in [0,1]$, then $\mathbb{E}\, e^{t(X - \mathbb{E}X)} \leq e^{\frac{t^2}{8}}$.

"Bernoulli's are $\frac{1}{4}$-subgaussian."

A detour (?)

<u>Hoeffding's Lemma.</u>

If $X \in [0,1]$, then $\mathbb{E}\, e^{t(X - \mathbb{E}X)} \leq e^{\frac{t^2}{8}}$.

"Bernoulli's are $\frac{1}{4}$-subgaussian."

<u>Corollary.</u> If $X = \sum_i X_i$, X is $\frac{n}{4}$-subgaussian.

X_i ↑ indep^t Bernoulli

Back in track

RAPPOR: - user j encodes $x_j \in [k]$ as $e_{x_j} \in \{0,1\}^k$

- flips each bit independently w.p. $\frac{1}{e^{\varepsilon/2}+1}$
- sends $Y_j \in \{0,1\}^k$ to the center

Center: Set $\hat{p}_i = \left(\frac{1}{n} \sum_{j=1}^{n} Y_j - \frac{1}{e^{\varepsilon/2}+1} \right) \frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}, \quad i \in [k]$

Back in track

RAPPOR:

- user j encodes $x_j \in [k]$ as $e_{x_j} \in \{0,1\}^k$
- flips each bit independently w.p. $\frac{1}{e^{\varepsilon/2}+1}$
- sends $Y_j \in \{0,1\}^k$ to the center

Center: Set $$\hat{p}_i = \Big(\underbrace{\frac{1}{n}\sum_{j=1}^{n} Y_j}_{\frac{n}{4}\text{-subgaussian!}} - \frac{1}{e^{\varepsilon/2}+1} \Big) \frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}, \quad i \in [k]$$

Back in track

RAPPOR: - user j encodes $x_j \in [k]$ as $e_{x_j} \in \{0,1\}^k$

- flips each bit independently w.p. $\frac{1}{e^{\varepsilon/2}+1}$

- sends $Y_j \in \{0,1\}^k$ to the center

Center: Set $\hat{p}_i = \left(\frac{1}{n}\sum_{j=1}^{n} Y_j - \frac{1}{e^{\varepsilon/2}+1}\right)\frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}$, $i \in [k]$

$$\mathbb{E}\|\hat{p}-p\|_\infty \lesssim \sqrt{\frac{\log k}{n}} \cdot \frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}$$

Back in track

RAPPOR: - user j encodes $x_j \in [k]$ as $e_{x_j} \in \{0,1\}^k$

- flips each bit independently w.p. $\frac{1}{e^{\varepsilon/2}+1}$

- sends $Y_j \in \{0,1\}^k$ to the center

Center: Set $\hat{p}_i = \left(\frac{1}{n} \sum_{j=1}^{n} Y_j - \frac{1}{e^{\varepsilon/2}+1} \right) \frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}$, $i \in [k]$

$$\mathbb{E}\|\hat{p}-p\|_\infty \lesssim \sqrt{\frac{\log k}{n}} \; \frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}$$

$\frac{1}{\varepsilon}$ when $\varepsilon \ll 1$

1 when $\varepsilon \gg 1$

It doesn't go to 0...

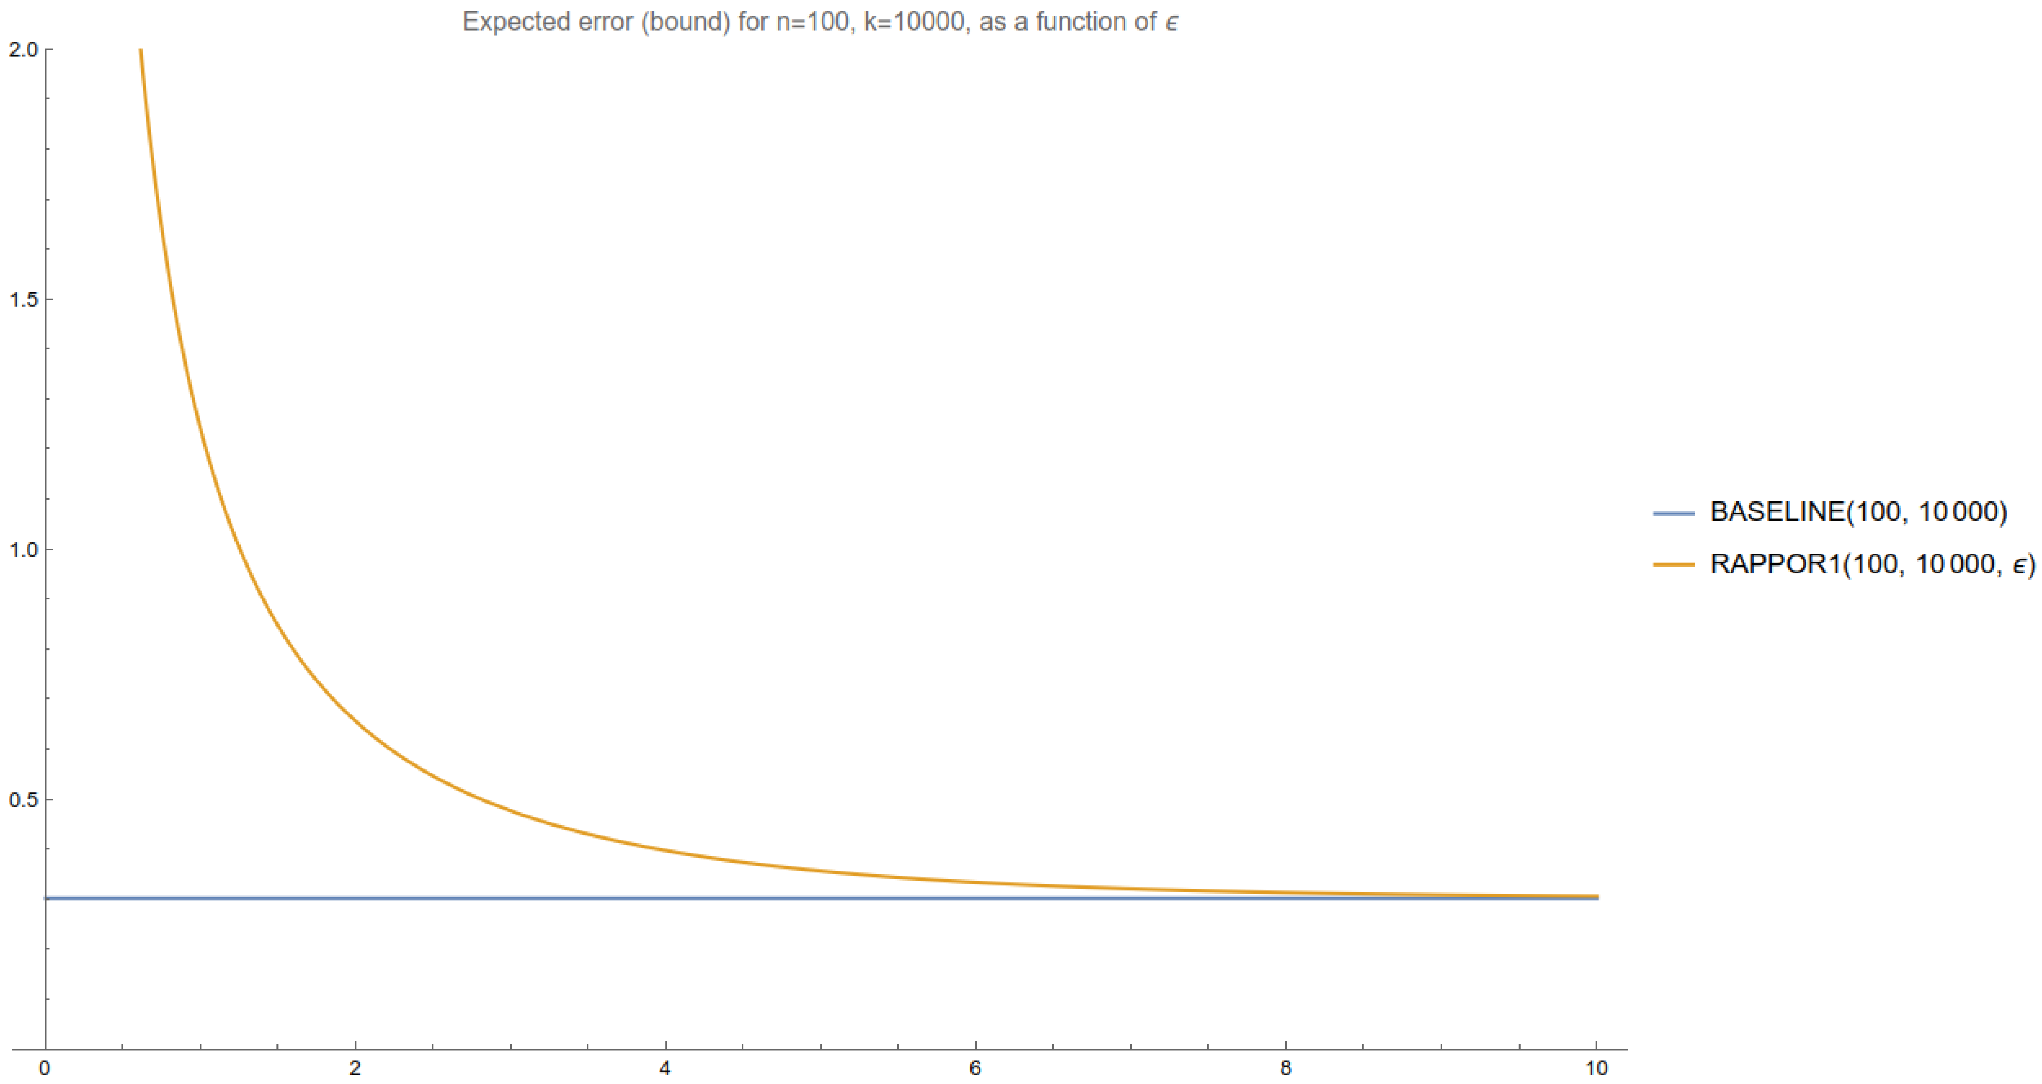

What can we do?

What can we do?

Modify the algorithm: can get

$$O\left(\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}\right)$$

via a general "repetition scheme". (Blows up communication by $\lceil \varepsilon \rceil$)

It goes to 0!

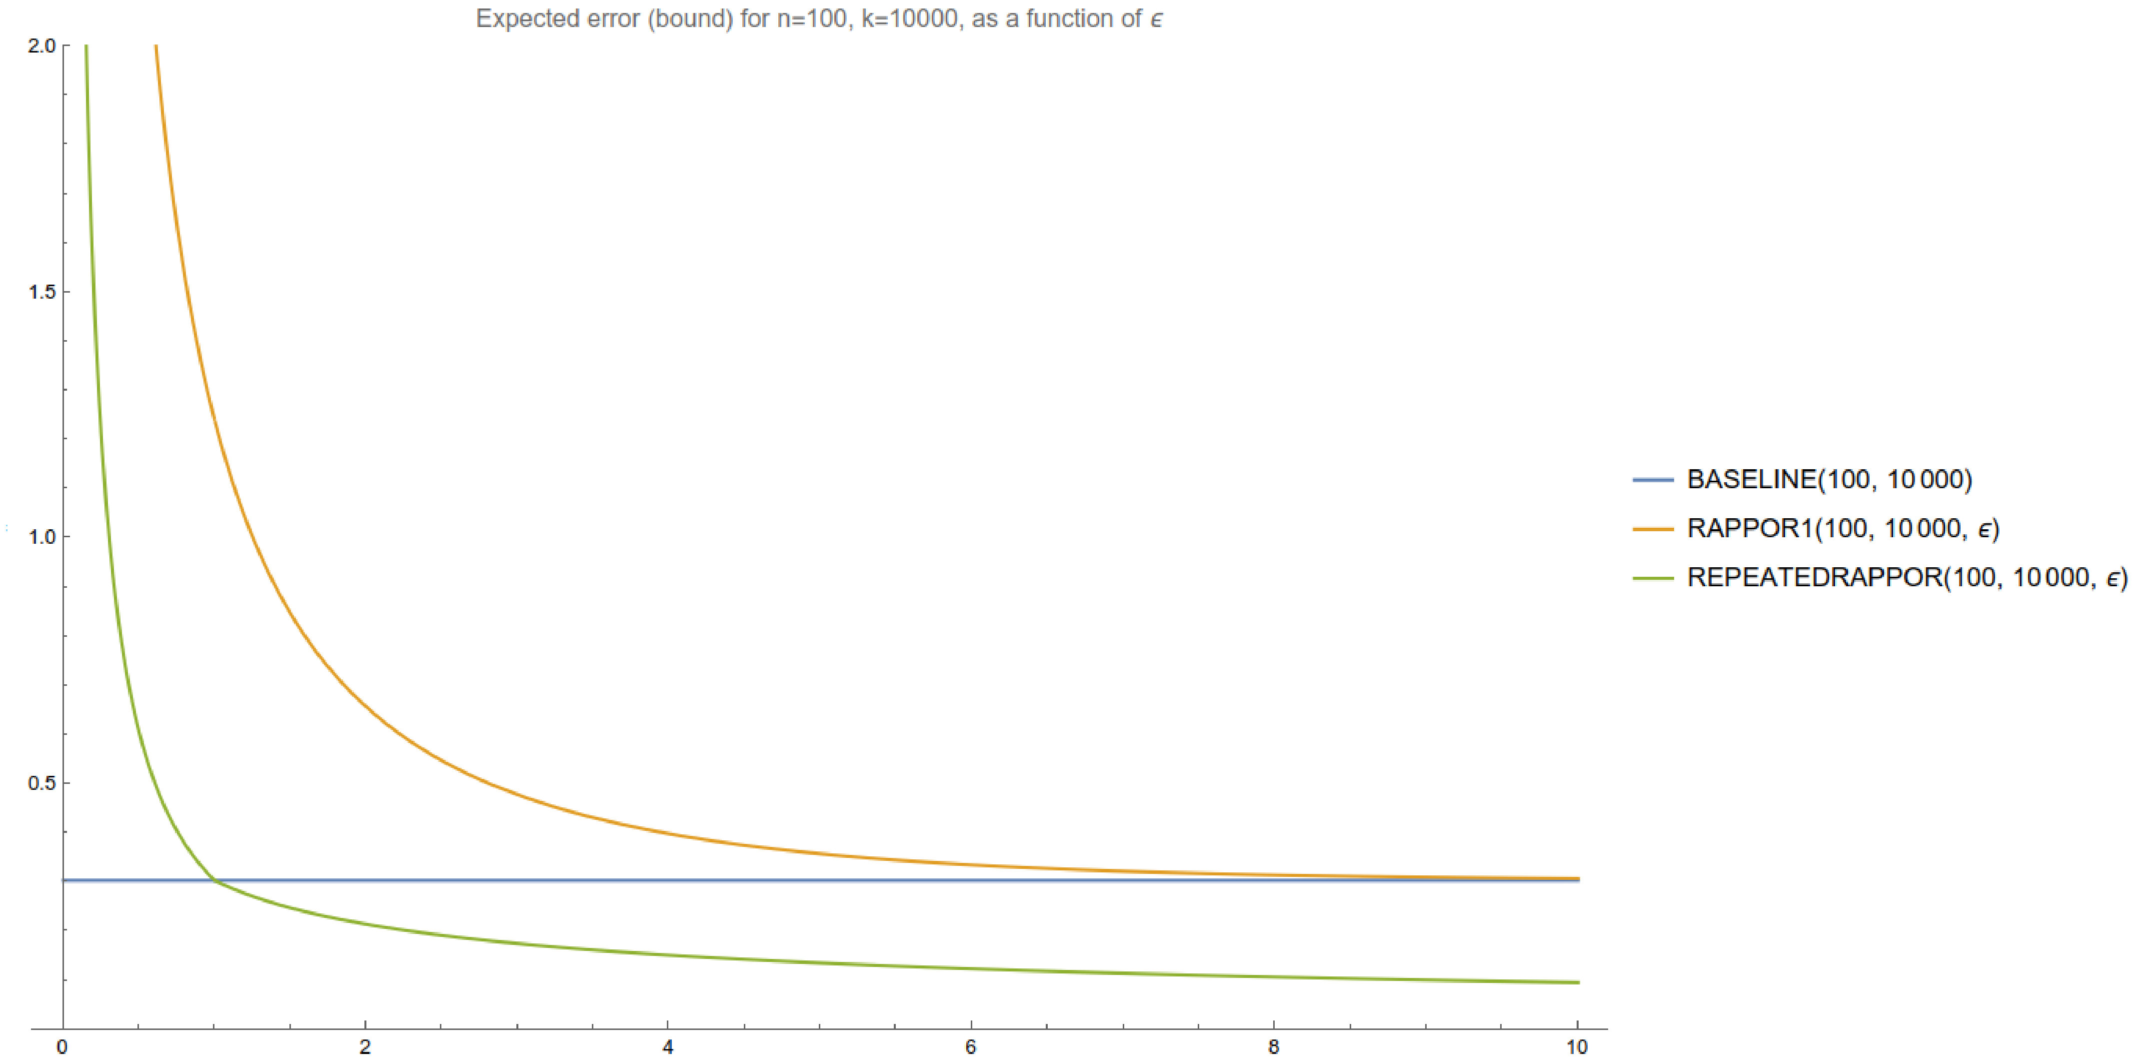

Expected error (bound) for n=100, k=10000, as a function of ϵ
2.0
1.5
1.0
0.5
0
2
4
6
8
10
BASELINE(100, 10 000)
RAPPOR1(100, 10 000, ϵ)
REPEATEDRAPPOR(100, 10 000, ϵ)

What else can we do?

Don't Modify the algorithm: can get

$$O\left(\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}\right)$$

just by analyzing RAPPOR better!

Remember Hoeffding?

<u>Kearns – Saul inequality.</u>

If $X \sim \mathrm{Bern}(p)$, then $\mathbb{E}\, e^{t(X-\mathbb{E}X)} \leq e^{\frac{1-2p}{4\log\frac{1-p}{p}} t^2}$

"Bernoullis are $\frac{1-2p}{2\log\frac{1-p}{p}}$ – subgaussian."

Remember Hoeffding?

Kearns–Saul inequality.

If $X \sim \mathrm{Bern}(p)$, then $\mathbb{E}\, e^{t(X-\mathbb{E}X)} \leq e^{\frac{1-2p}{4\log\frac{1-p}{p}} t^2}$

"Bernoullis are $\underbrace{\frac{1-2}{2\log\frac{1-p}{p}}}_{\sigma^2(p)}$ - subgaussian."

Remember RAPPOR?

In RAPPOR, each user flips each of its k bits with proba. $\underbrace{\frac{1}{e^{\varepsilon/2}+1}}_{=p}$

Each bit is either $\mathrm{Bern}(p)$ or $\mathrm{Bern}(1-p)$

Each bit is σ^2-subgaussian for

$$\sigma^2 = \sigma^2(p) = \sigma^2(1-p)$$

Remember RAPPOR?

In RAPPOR, each user flips each of its k bits with proba. $\underbrace{\frac{1}{e^{\varepsilon/2}+1}}_{=p}$

Each bit is either $\mathrm{Bern}(p)$ or $\mathrm{Bern}(1-p)$

Each bit is σ^2-subgaussian for

$$\sigma^2 = \sigma^2(p) = \sigma^2(1-p) = \frac{e^{\varepsilon/2}-1}{e^{\varepsilon/2}+1} \cdot \frac{1}{\varepsilon}$$

Remember RAPPOR?

In RAPPOR, each user flips each of its k bits with proba. $\frac{1}{e^{\varepsilon/2}+1}$

This gives error bound

$$\mathbb{E}\,\|\hat{f}-f\|_\infty \lesssim \sqrt{\frac{e^{\varepsilon/2}+1}{(e^{\varepsilon/2}-1)\varepsilon} \cdot \frac{\log k}{n}}$$

Remember RAPPOR?

In RAPPOR, each user flips each of its k bits with proba. $\frac{1}{e^{\varepsilon/2}+1}$

This gives error bound

$$\mathbb{E}\,\|\hat{f}-f\|_\infty \lesssim \sqrt{\frac{e^{\varepsilon/2}+1}{(e^{\varepsilon/2}-1)\varepsilon}\cdot\frac{\log k}{n}}$$

$\sqrt{\frac{\log k}{n\varepsilon^2}}$ when $\varepsilon \ll 1$

$\sqrt{\frac{\log k}{n\varepsilon}}$ when $\varepsilon \gg 1$

So...


RAPPOR ACHIEVES THE SAME ERROR OUT-OF-THE-BOX!
SO, WE'RE DONE, RIGHT
SO, WE'RE DONE, RIGHT?
imgflip.com

We couldn't prove a matching lower bound...

We couldn't prove a matching lower bound...

So we tried to improve the upper bound.

Local Glivenko-Cantelli [CK23]

Let μ be a product distribution over $\{0,1\}^d$ with mean vector $p \in [0,1]^d$. Define $\tilde{p}_i = \min(p_i, 1-p_i)$, and assume wlog $\tilde{p}_1 \geq \dots \geq \tilde{p}_d$.

Given $X_1, \dots, X_n \sim \mu$, setting

$$\hat{p} = \frac{1}{n}\sum_{j=1}^{n} X_j$$

we have

$$\mathbb{E}\left[\|\hat{p}-p\|_\infty\right] \lesssim \max_{1\leq i\leq d} \sqrt{\frac{\tilde{p}_i \log(i+1)}{n}} + \frac{\log n}{n} \max_{1\leq i\leq d} \frac{\log(i+1)}{\log\frac{1}{\tilde{p}_i}}$$

Local Glivenko-Cantelli [CK23] + RAPPOR

For us, $p_i \in \{p, 1-p\}$ $\forall i$, so $\tilde{p}_i = p = \frac{1}{e^{\varepsilon/2}+1}$. We immediately* get

$$\mathbb{E}\|\hat{f}-f\|_\infty \lesssim \frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}\left(\sqrt{\frac{\log k}{n(e^{\varepsilon/2}+1)}} + \frac{\log k}{n\max(\varepsilon,1)}\right)$$

$$\underset{\varepsilon \gg 1}{\asymp} \sqrt{\frac{\log k}{n e^{\varepsilon/2}}} + \frac{\log k}{n\varepsilon}$$

Local Glivenko-Cantelli [CK23] + RAPPOR

For us, $p_i \in \{p, 1-p\}$ $\forall i$, so $\tilde{p}_i = p = \frac{1}{e^{\varepsilon/2}+1}$. We immediately* get

$$\mathbb{E}\|\hat{f}-f\|_\infty \lesssim \frac{e^{\varepsilon/2}+1}{e^{\varepsilon/2}-1}\left(\sqrt{\frac{\log k}{n(e^{\varepsilon/2}+1)}} + \frac{\log k}{n\max(\varepsilon,1)}\right)$$

$$\underset{\varepsilon \gg 1}{\lesssim} \sqrt{\frac{\log k}{n e^{\varepsilon/2}}} + \frac{\log k}{n\varepsilon}$$

*This is a lie.

We did nothing, and it got better!

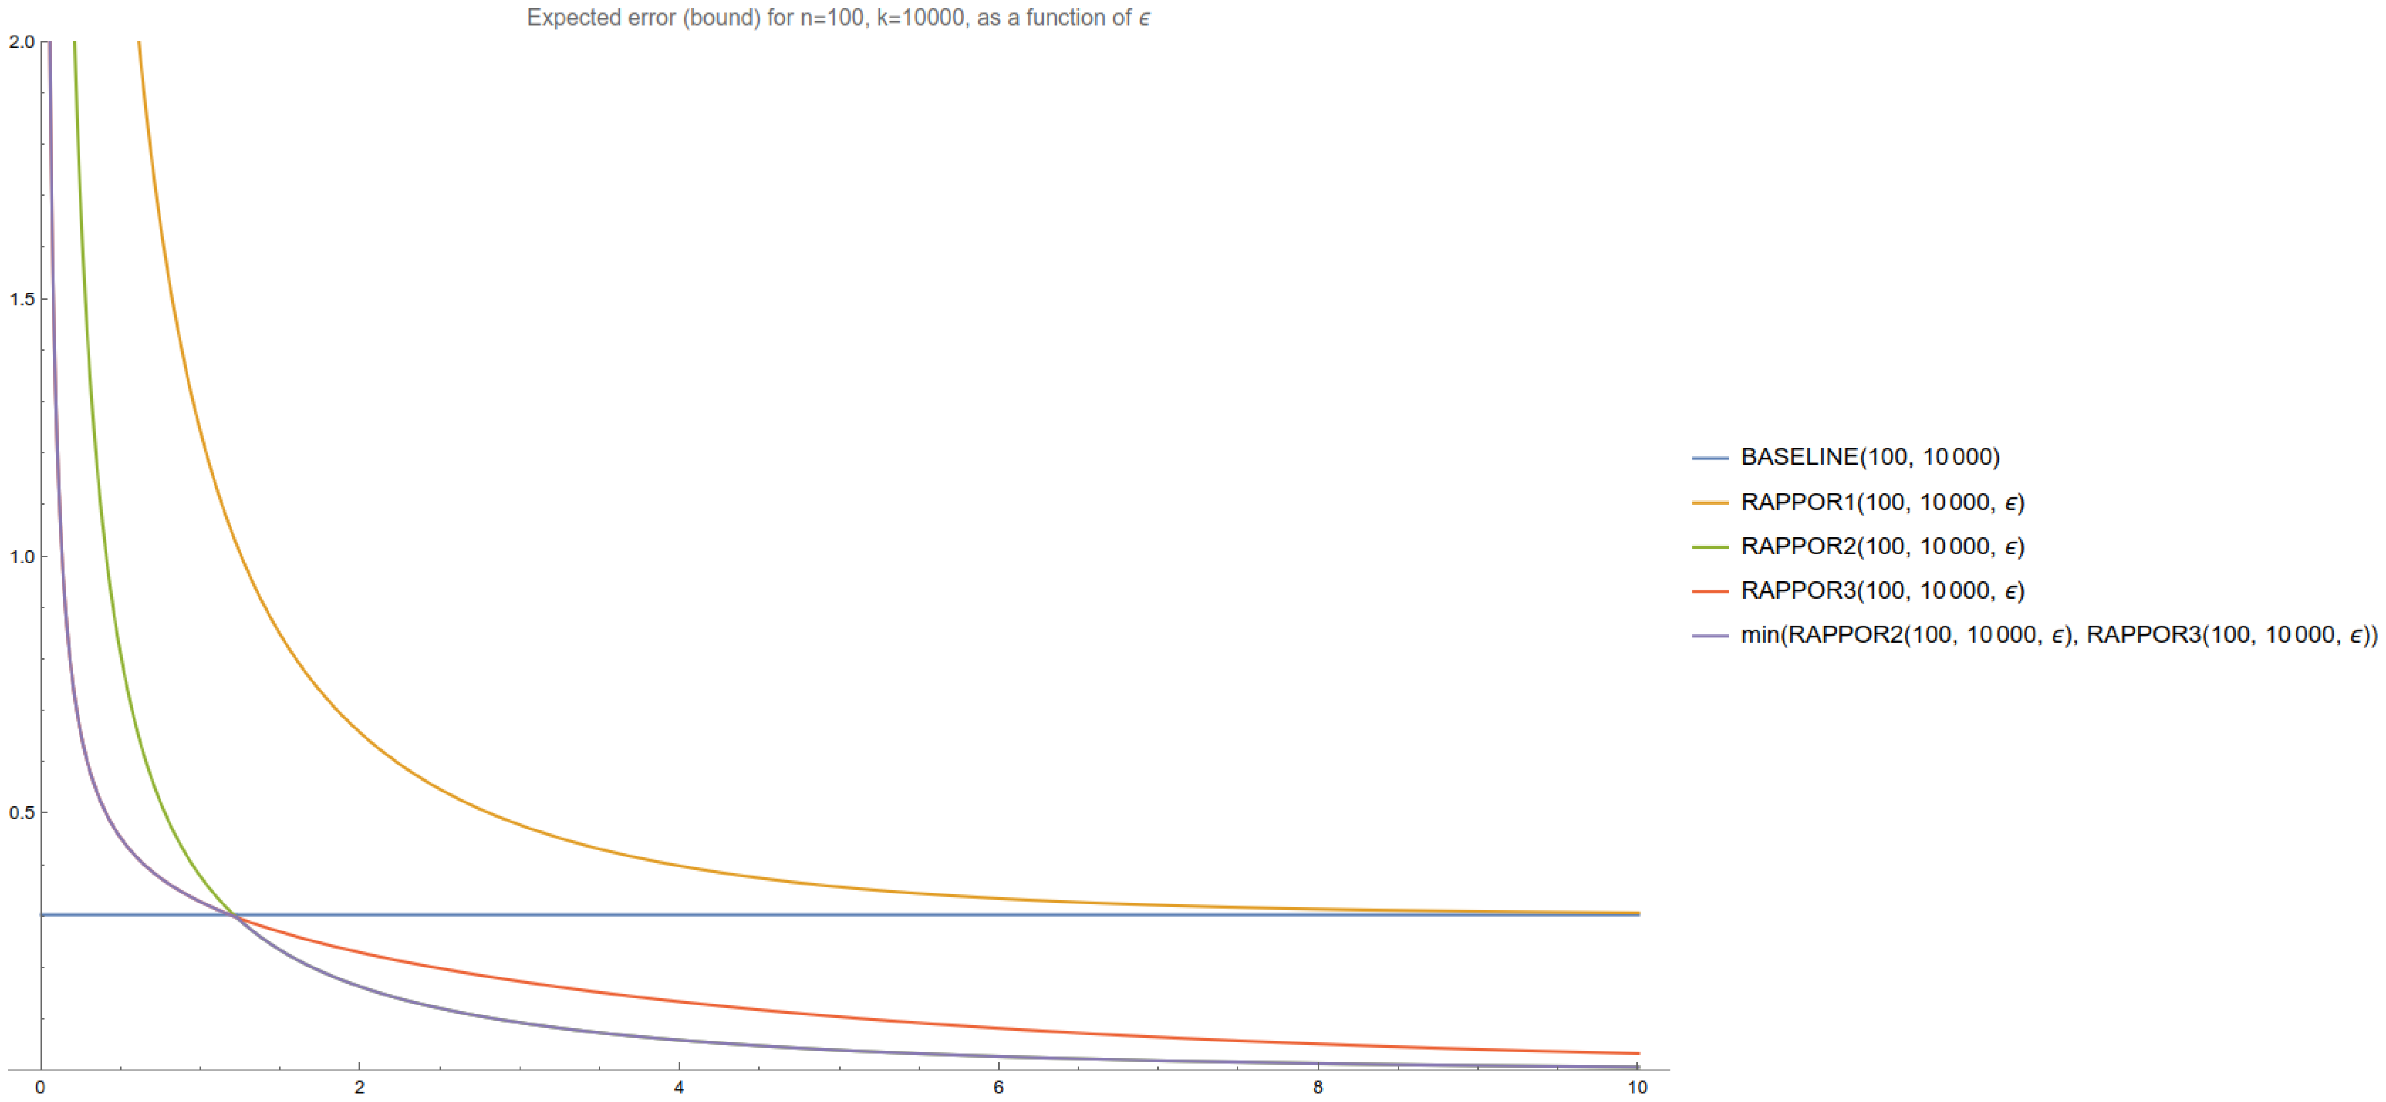

Also, this is optimal*

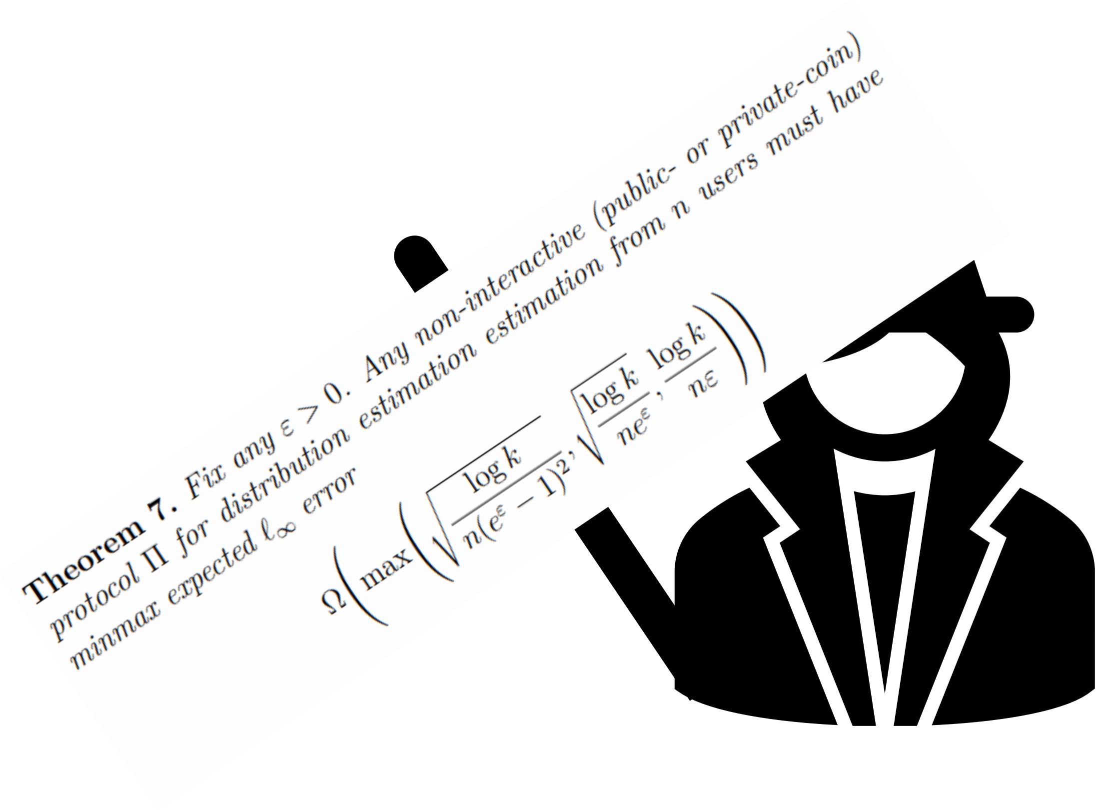

*This is a lie.

Summary

- Algorithms can "improve" without any modification
- Can we go beyond RAPPOR? (yes: work in progress)
- Even for "solved" problems, there is so much we don't know!
- Trying to prove lower bounds is useful
- Not mentioned here: beyond worst-case, and other protocols

Thank you.

The plot thickens