Local Differential Privacy: Refined Mechanism Design and Utility Analysis

Ye Zheng

Advisor: Dr. Yidan Hu

Committee: Dr. Sumita Mishra, Dr. Haibo Yang, Dr. Weijie Zhao

RIT | Rochester Institute of Technology

PDF & slides 👉 https://zhengyeah.com

Data Collection Everywhere

RIT

- Users' personal data are collected by companies for analysis or services

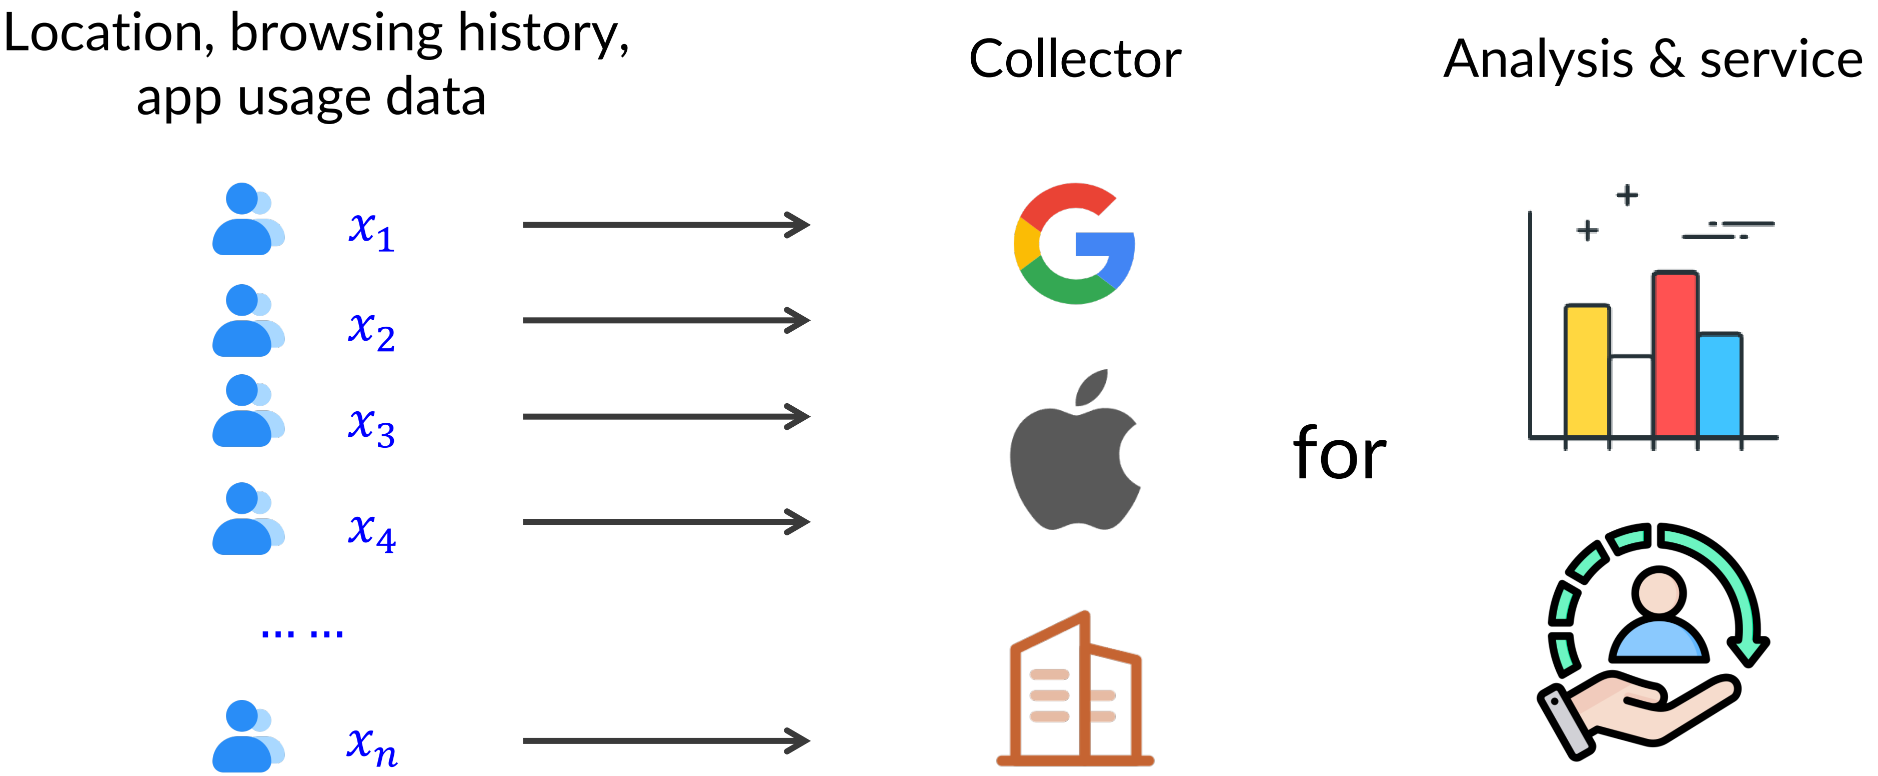

Users' Data Privacy

- Users' personal data are collected by companies for analysis or services
  - these companies may not be trusted to collect users' sensitive data

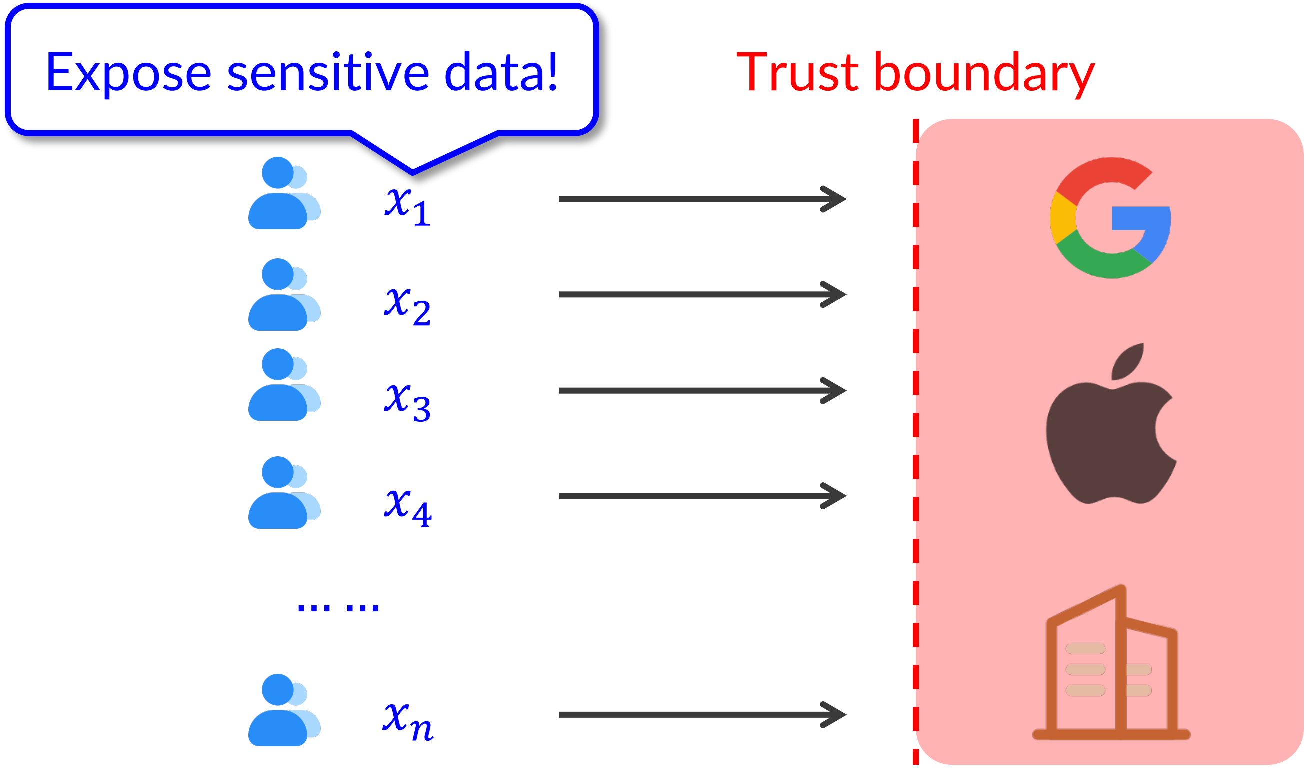

Users' Data Privacy

- Users' personal data are collected by companies for analysis or services
  - these companies may not be trusted to collect users' sensitive data

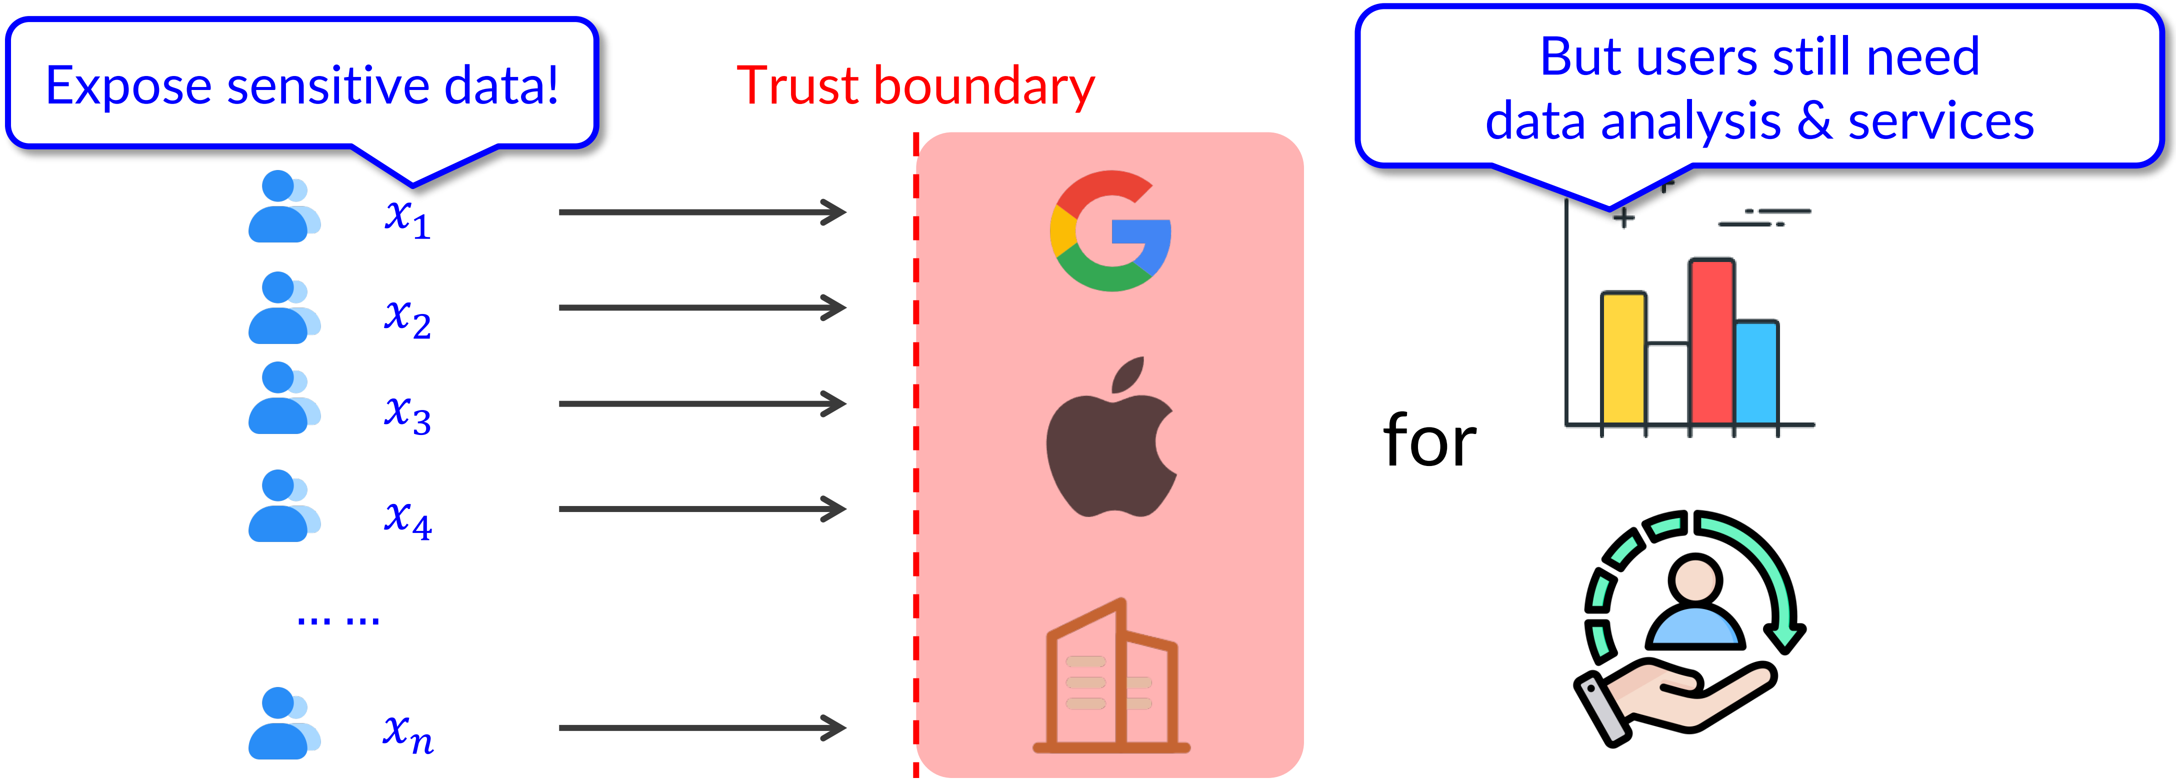

Users' Data Privacy

RIT

- Users' personal data are collected by companies for analysis or services
  - these companies may not be trusted to collect users' sensitive data

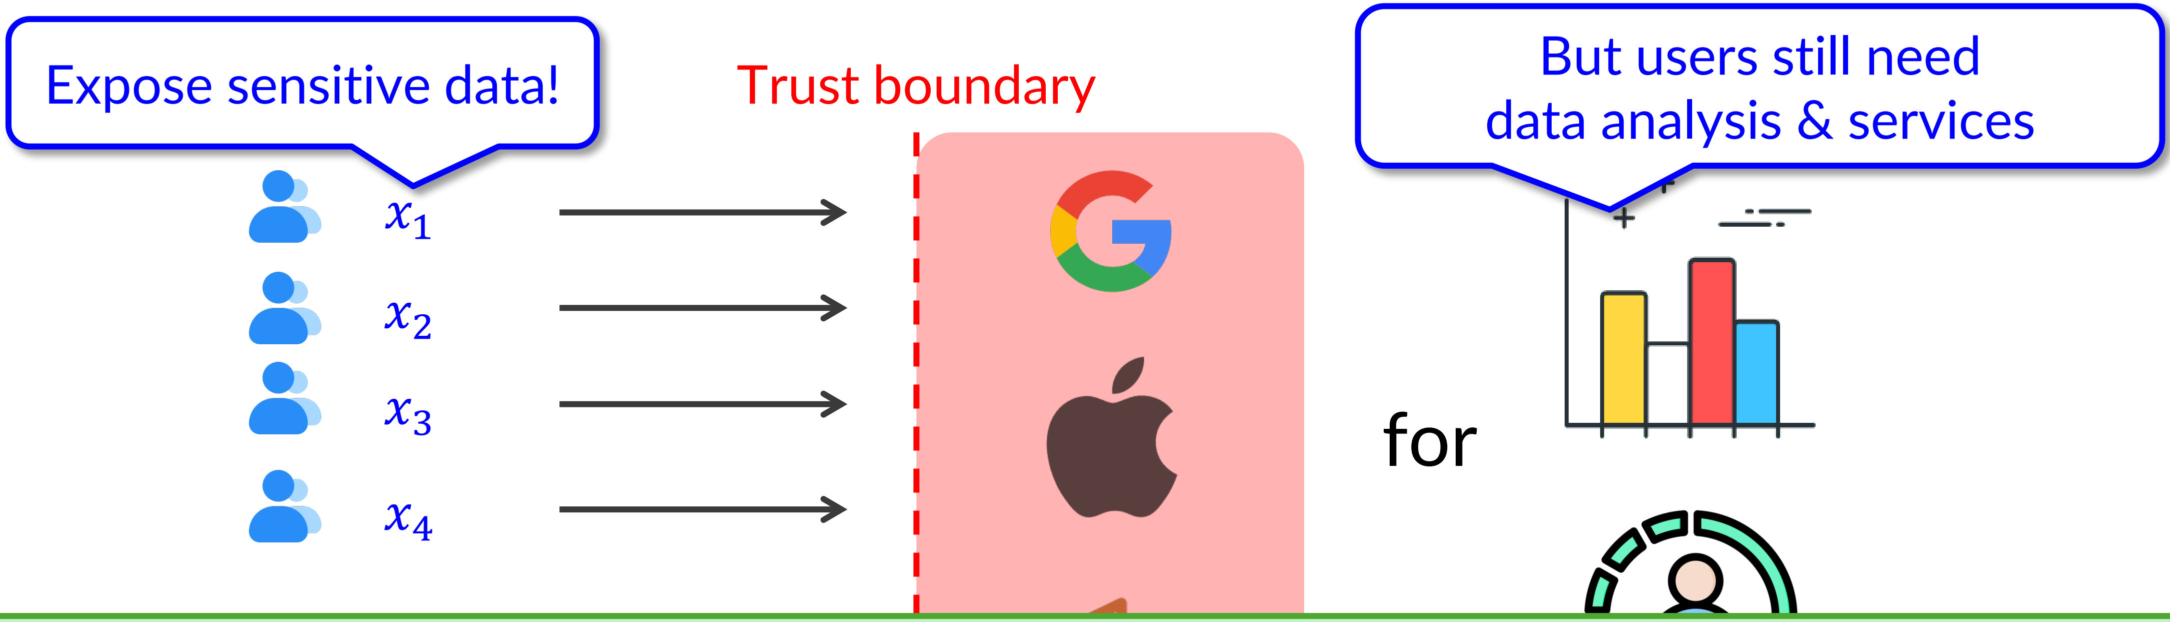

Q: How can we provide data analysis & services **while** protecting users' data privacy?

Privacy-Preserving Computation

RIT

- Users' personal data are collected by companies for analysis or services
  - these companies may be untrusted to collect users' sensitive data
  - how to compute $f(x_1, x_2, \ldots, x_n)$ without revealing $x_1, x_2, \ldots, x_n$?

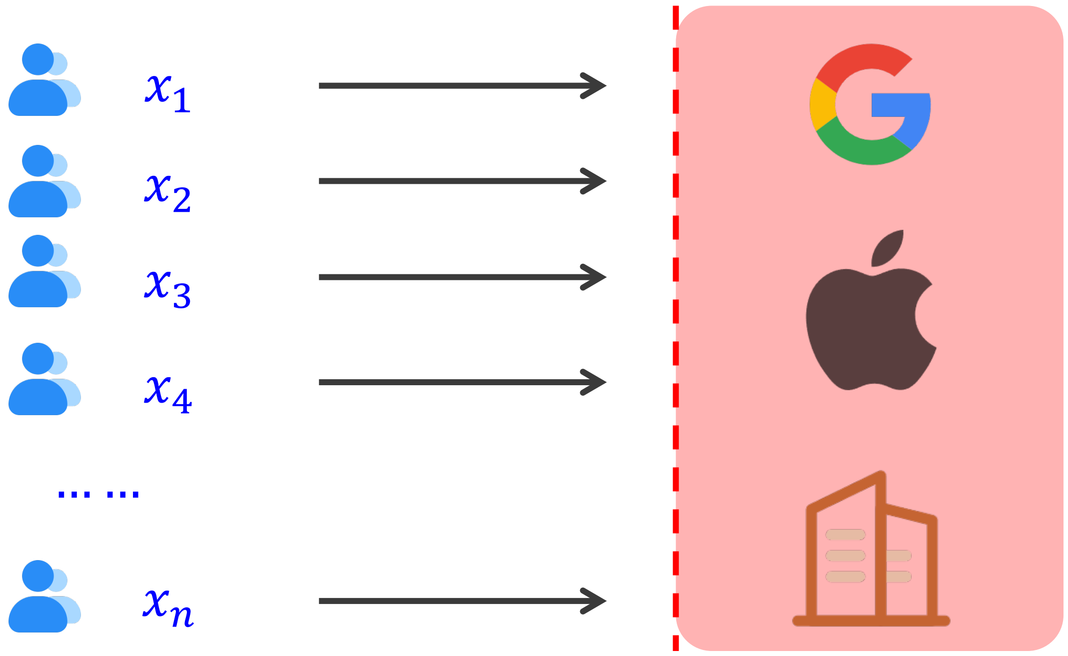

for $f(x_1, x_2, \ldots, x_n)$

Privacy-Preserving Computation - **Techniques**

- Homomorphic encryption (HE), multi-party computation (MPC), local differential privacy (LDP), etc

Privacy-Preserving Computation - **HE**

RIT

- Homomorphic encryption (HE):
  - "homomorphic": preserving structure
  - design algorithms $\{Enc, Dec\}$ $\rightarrow$ $Dec\big(f(Enc(x_1), Enc(x_2), \dots, Enc(x_n))\big) = f(x_1, x_2, \dots, x_n)$

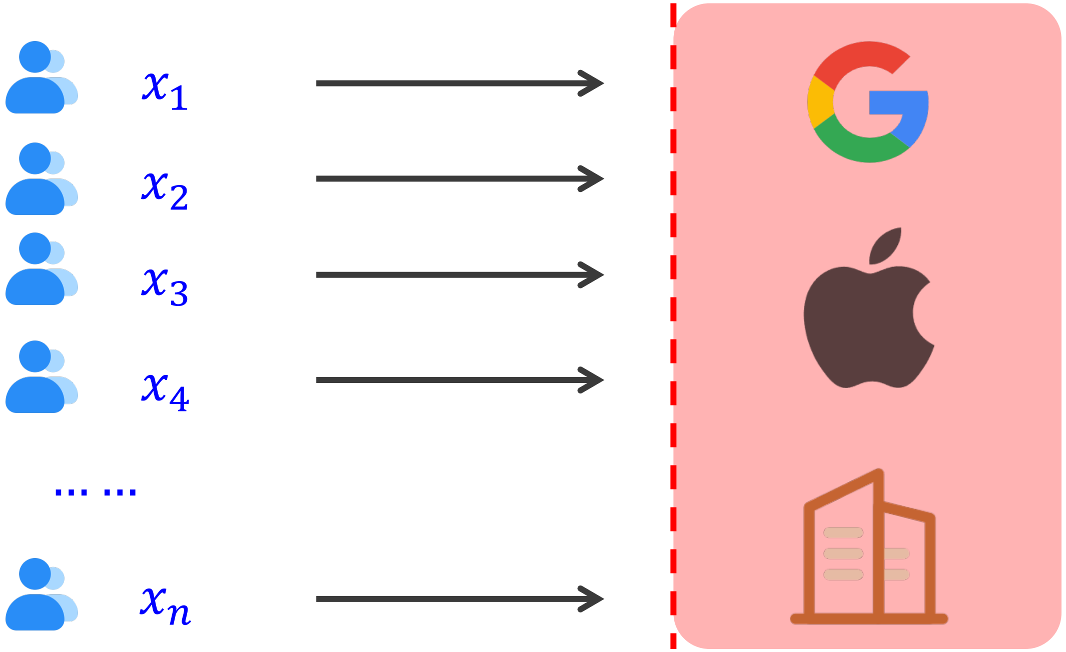

for $f(x_1, x_2, \dots, x_n)$

Privacy-Preserving Computation - **HE**

- Homomorphic encryption (HE):
  - "homomorphic": preserving structure
  - design algorithms $\{Enc, Dec\}$ $\rightarrow$ $Dec\big(f(Enc(x_1), Enc(x_2), \dots, Enc(x_n))\big) = f(x_1, x_2, \dots, x_n)$

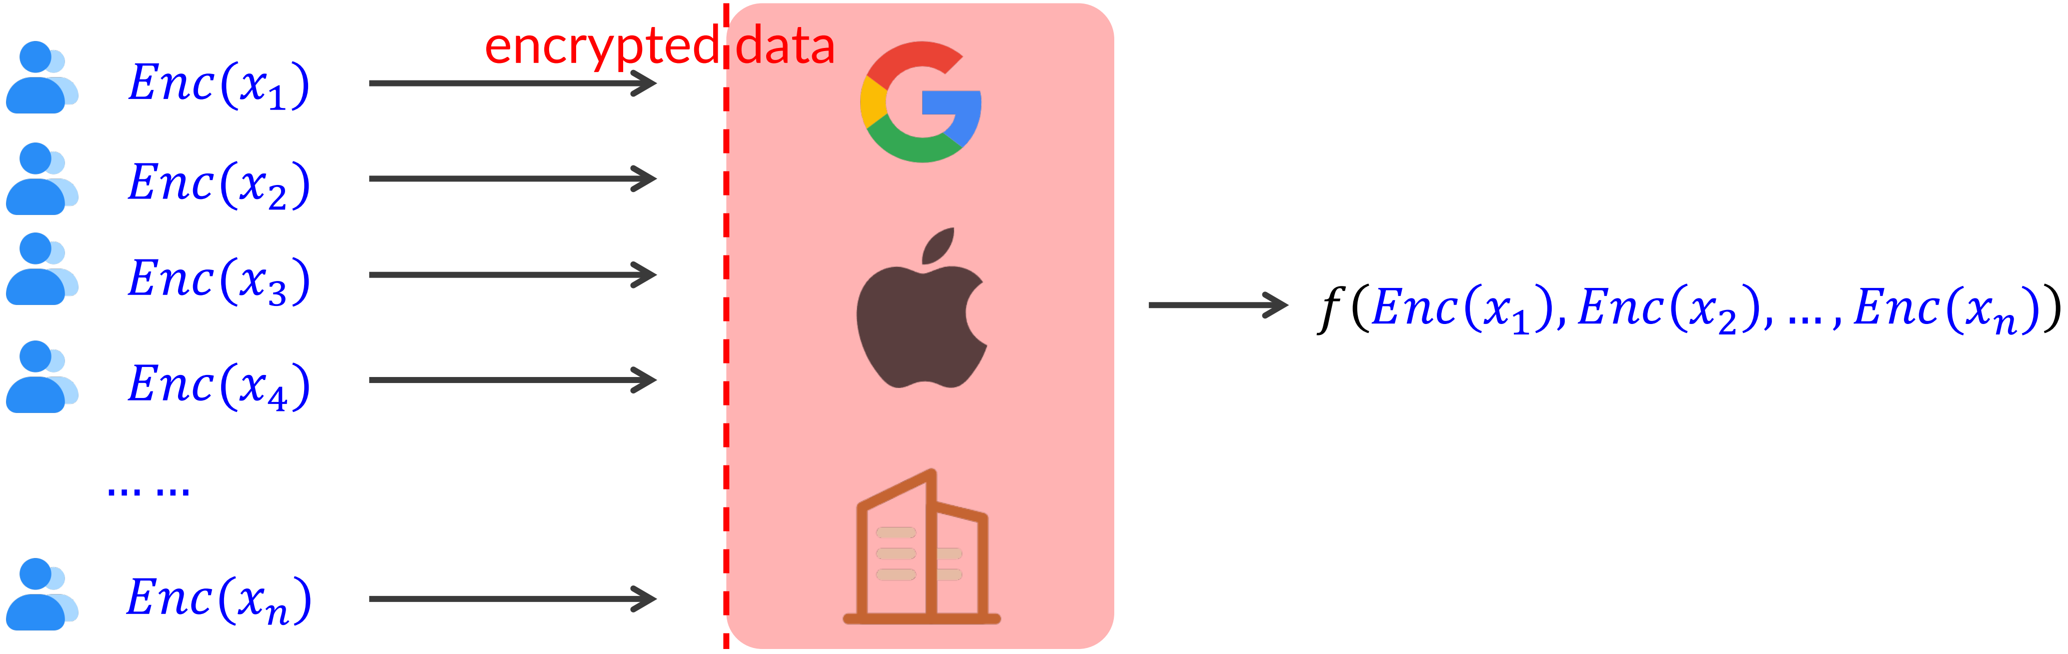

Privacy-Preserving Computation - **HE**

RIT

- Homomorphic encryption (HE):
  - "homomorphic": preserving structure
  - design algorithms $\{Enc, Dec\}$ $\rightarrow$ $Dec\left(f(Enc(x_1), Enc(x_2), \ldots, Enc(x_n))\right) = f(x_1, x_2, \ldots, x_n)$

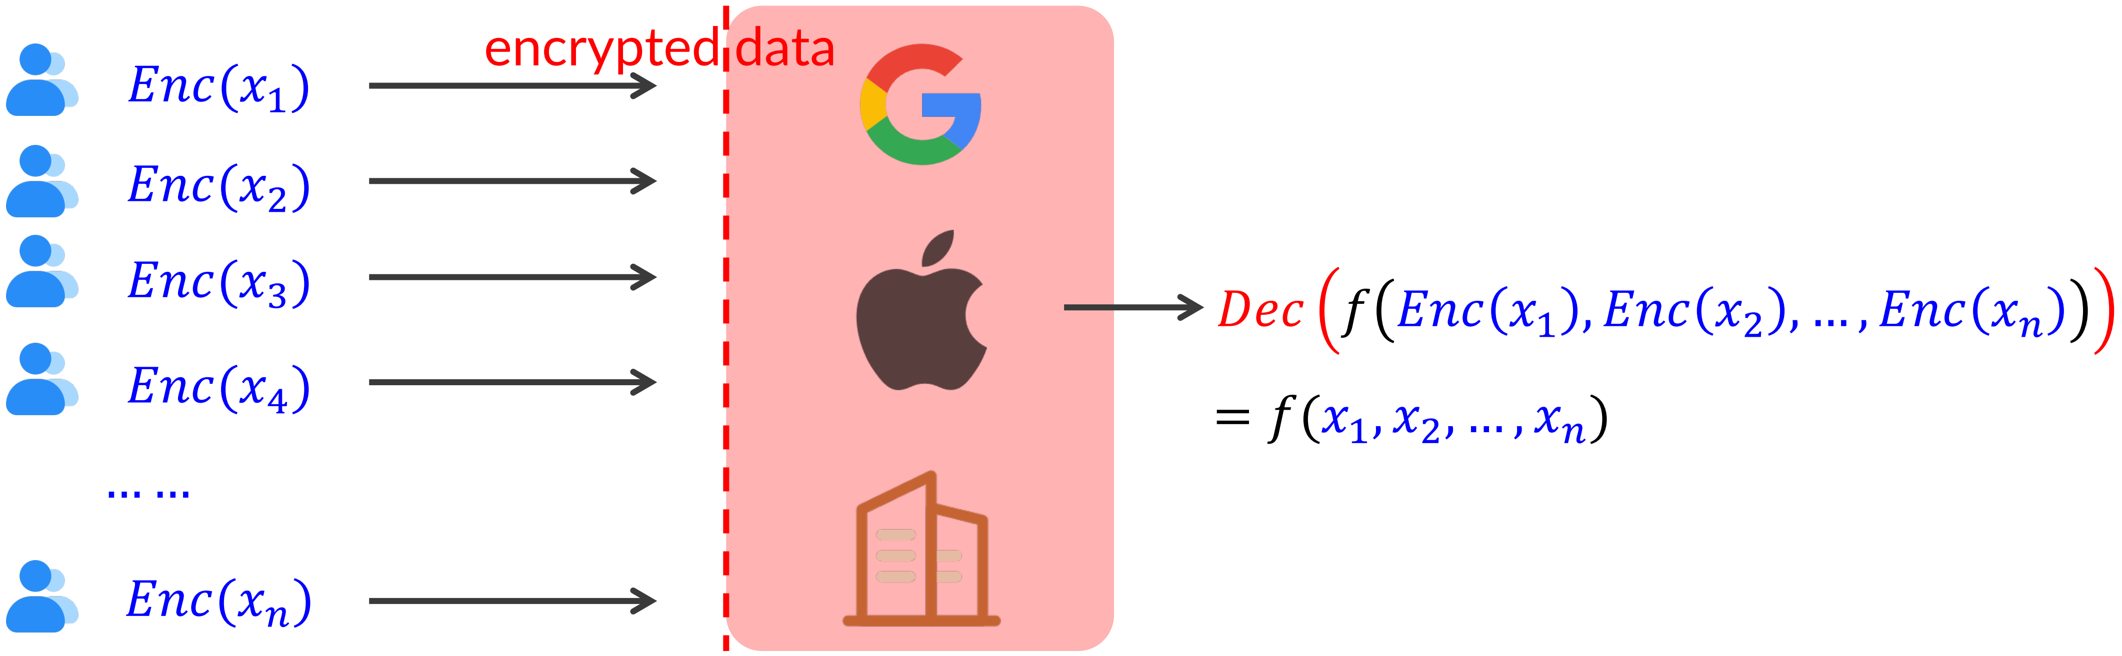

Privacy-Preserving Computation - **HE**

- Homomorphic encryption (HE):
  - "homomorphic": preserving structure
  - design algorithms $\{Enc, Dec\}$ $\rightarrow$ $Dec\big(f(Enc(x_1), Enc(x_2), \dots, Enc(x_n))\big) = f(x_1, x_2, \dots, x_n)$

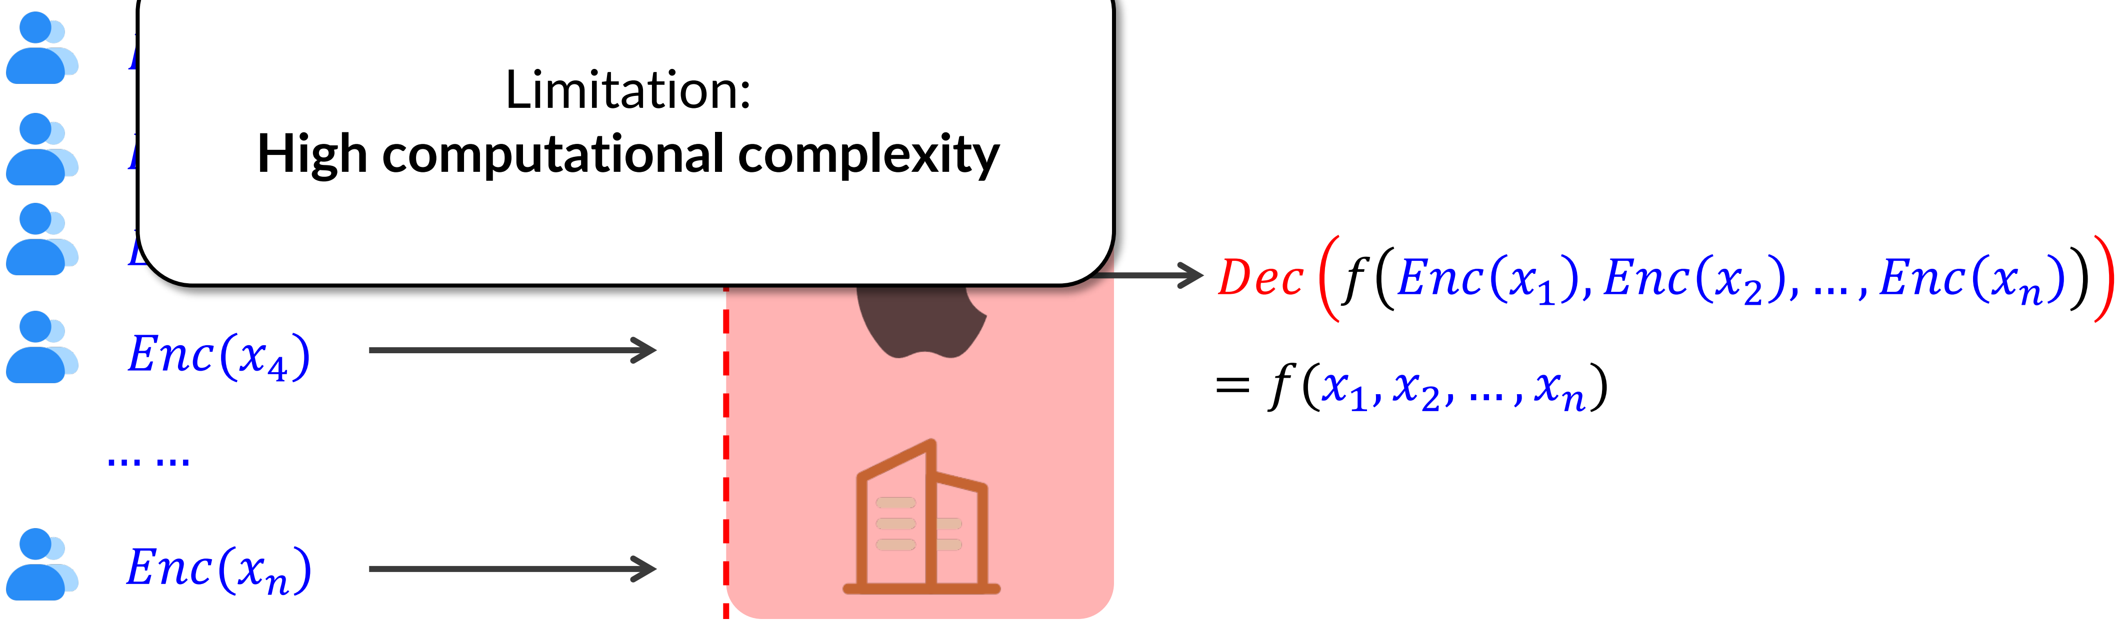

Limitation:
High computational complexity

Privacy-Preserving Computation - **MPC**

RIT

- Multi-party computation (MPC):
  - no central party
  - jointly compute f without revealing x_i
- Example: $f(x_1, x_2) = x_1 + x_2$

x_1

for $f(x_1, x_2)$

x_2

... ...

Privacy-Preserving Computation - **MPC**

RIT

- Multi-party computation (MPC):
  - no central party
  - jointly compute f without revealing x_i
- Example: $f(x_1, x_2) = x_1 + x_2$

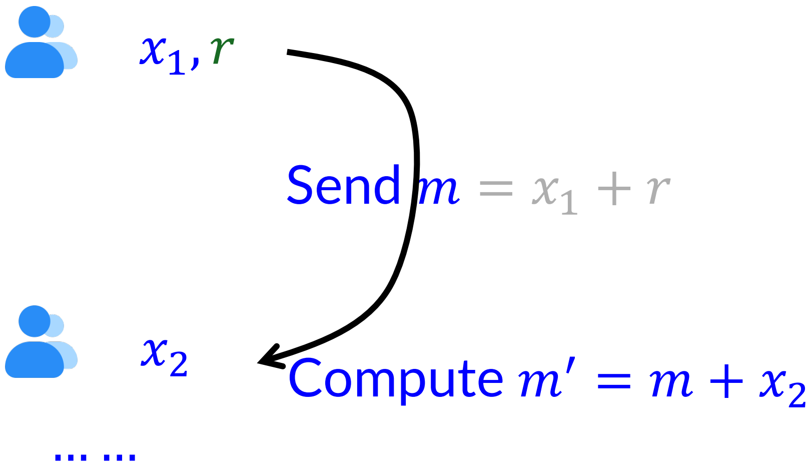

for $f(x_1, x_2)$

Privacy-Preserving Computation - **MPC**

- Multi-party computation (MPC):
  - no central party
  - jointly compute f without revealing x_i
- Example: $f(x_1, x_2) = x_1 + x_2$

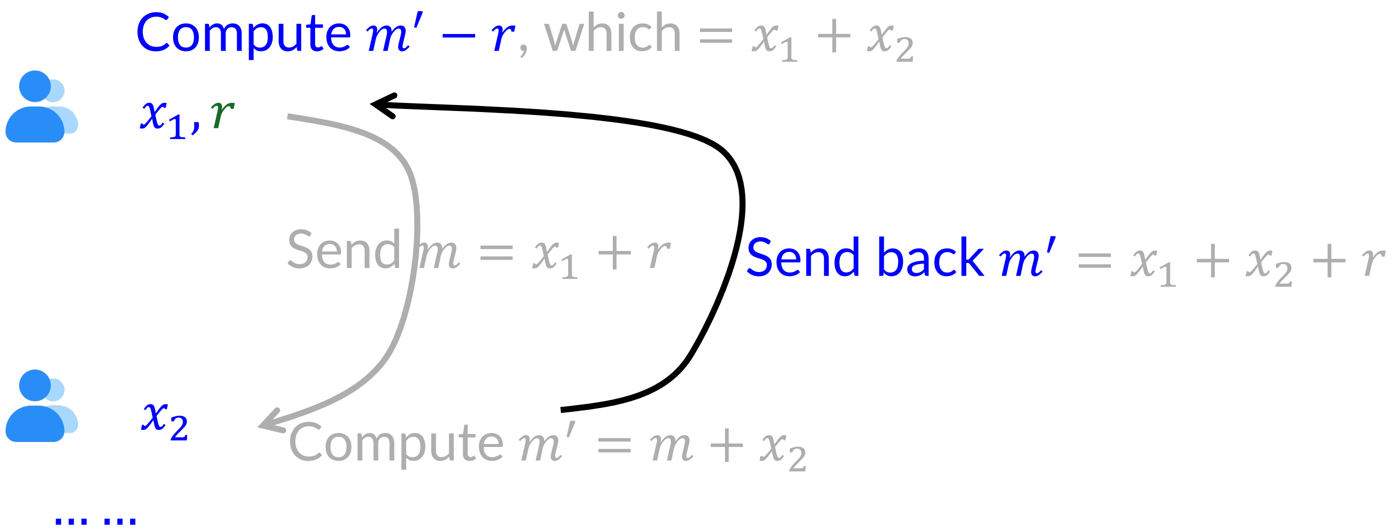

Privacy-Preserving Computation - **MPC**

RIT

- Multi-party computation (MPC):
  - no central party
  - jointly compute f without revealing x_i
- Example: $f(x_1, x_2) = x_1 + x_2$

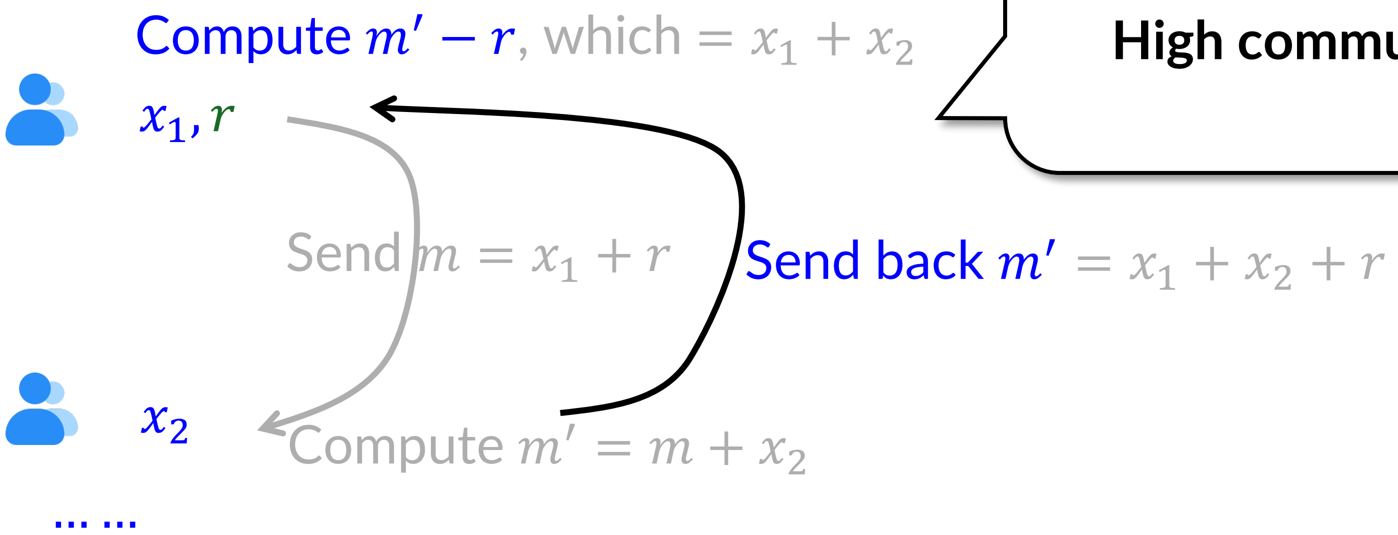

Limitation:
High communication complexity

Privacy-Preserving Computation - **LDP**

RIT

- Local differential privacy (LDP):
  - hard to differentiate the sensitive data from other data
  - each user locally perturbs x_i to $\tilde{x}_i$ $\rightarrow$ $f(\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_n) \approx f(x_1, x_2, \ldots, x_n)$

Privacy-Preserving Computation - **LDP**

RIT

- Local differential privacy (LDP):
  - hard to differentiate the sensitive data from other data
  - each user locally perturbs x_i to $\tilde{x}_i$ $\rightarrow$ $f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n) \approx f(x_1, x_2, \dots, x_n)$

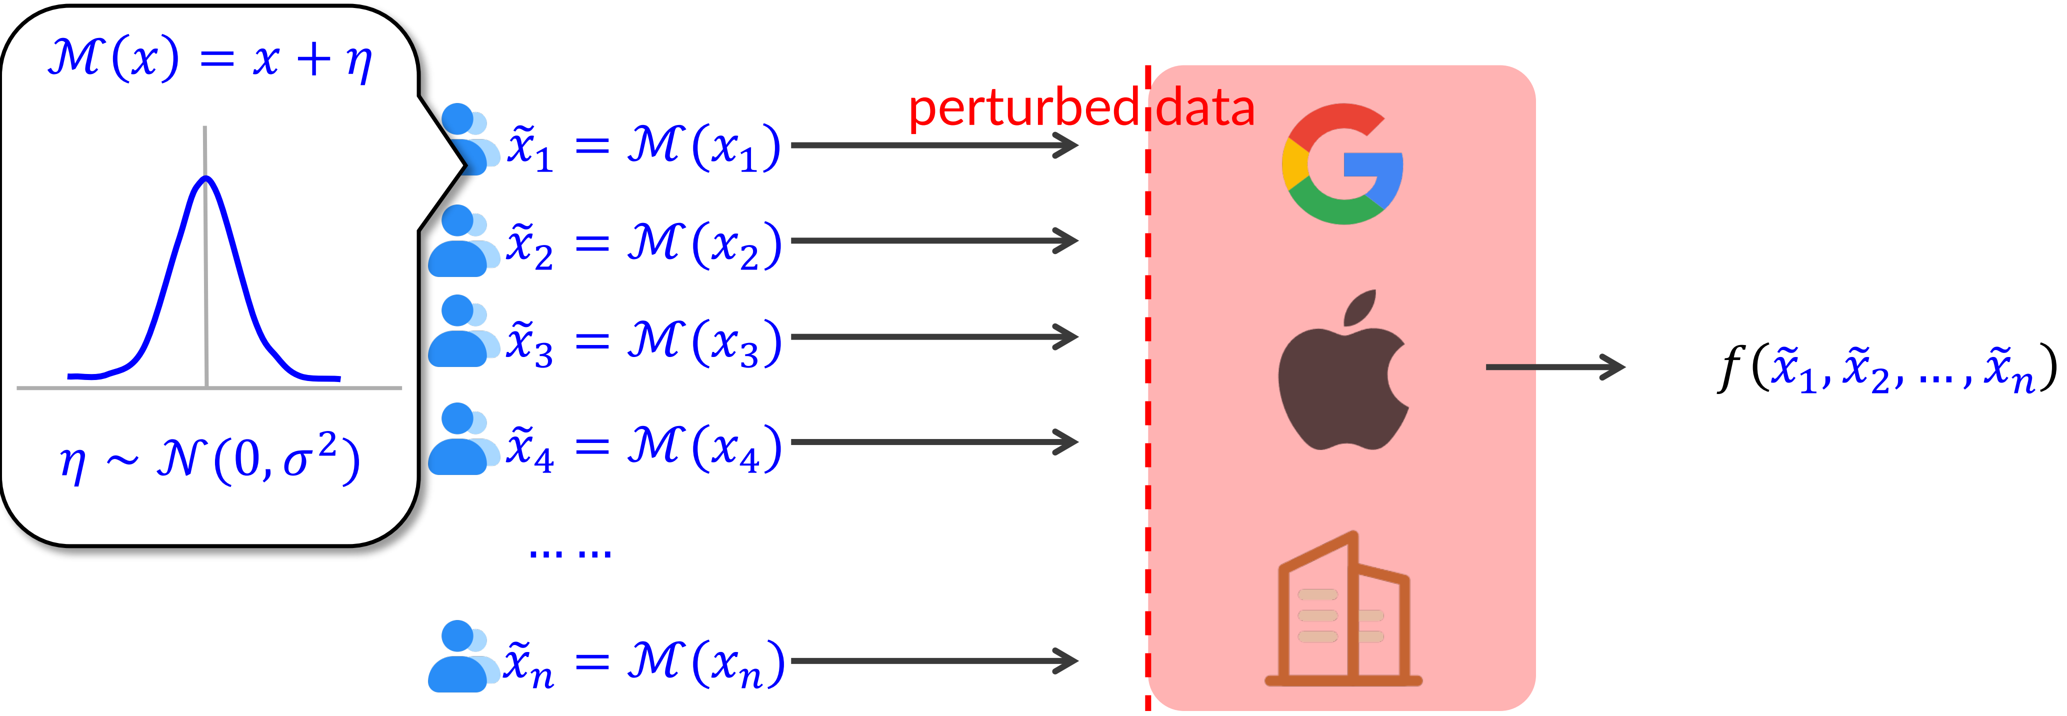

Privacy-Preserving Computation - **LDP**

- Local differential privacy (LDP):
  - hard to differentiate the sensitive data from other data
  - each user locally perturbs x_i to $\tilde{x}_i$ $\rightarrow$ $f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n) \approx f(x_1, x_2, \dots, x_n)$

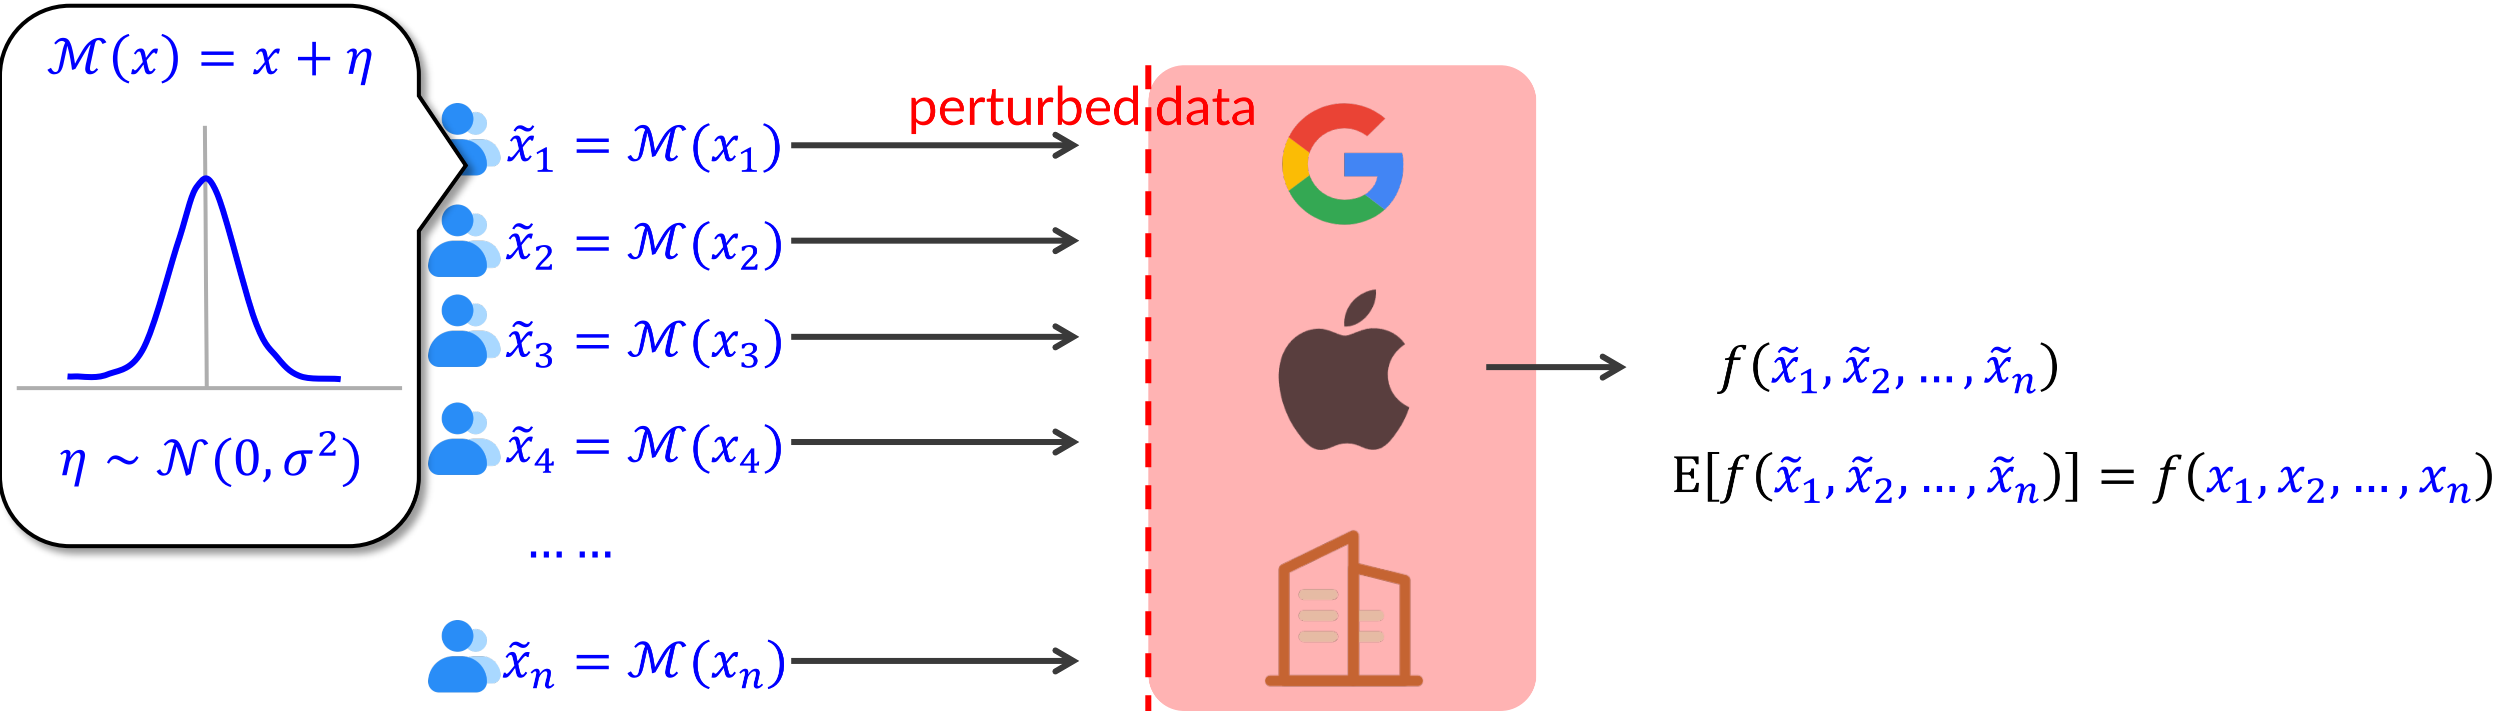

Privacy-Preserving Computation - **LDP**

- Local differential privacy (L
  - hard to differentiate the
  - each user locally perturb

  $\ldots, x_n)$

Advantages:
Negligible computational complexity
No communication between users

But approximated f

$\mathcal{M}(x) = x + \eta$

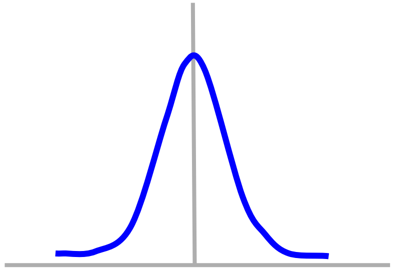

$\eta \sim \mathcal{N}(0, \sigma^2)$

$\tilde{x}_1 = \mathcal{M}(x_1)$

$\tilde{x}_2 = \mathcal{M}(x_2)$

$\tilde{x}_3 = \mathcal{M}(x_3)$

$\tilde{x}_4 = \mathcal{M}(x_4)$

… …

$\tilde{x}_n = \mathcal{M}(x_n)$

perturbed data

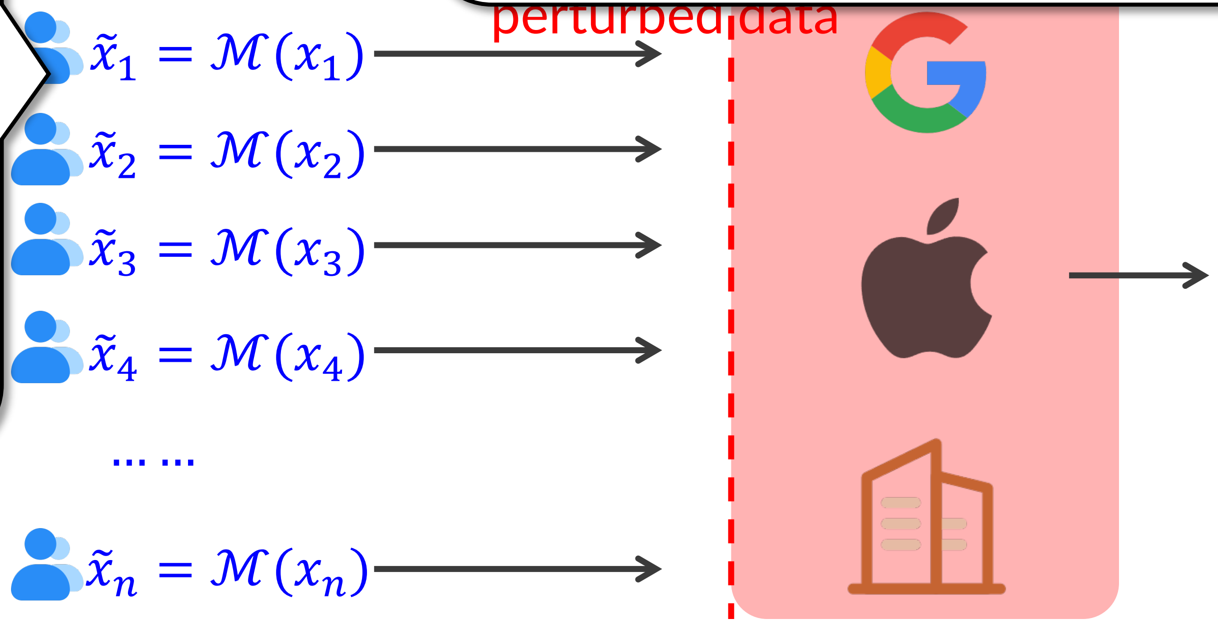

$f(\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_n)$

$\mathrm{E}[f(\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_n)] = f(x_1, x_2, \ldots, x_n)$

LDP – Applications

Chrome uses LDP to collect homepage settings, extension usage, etc

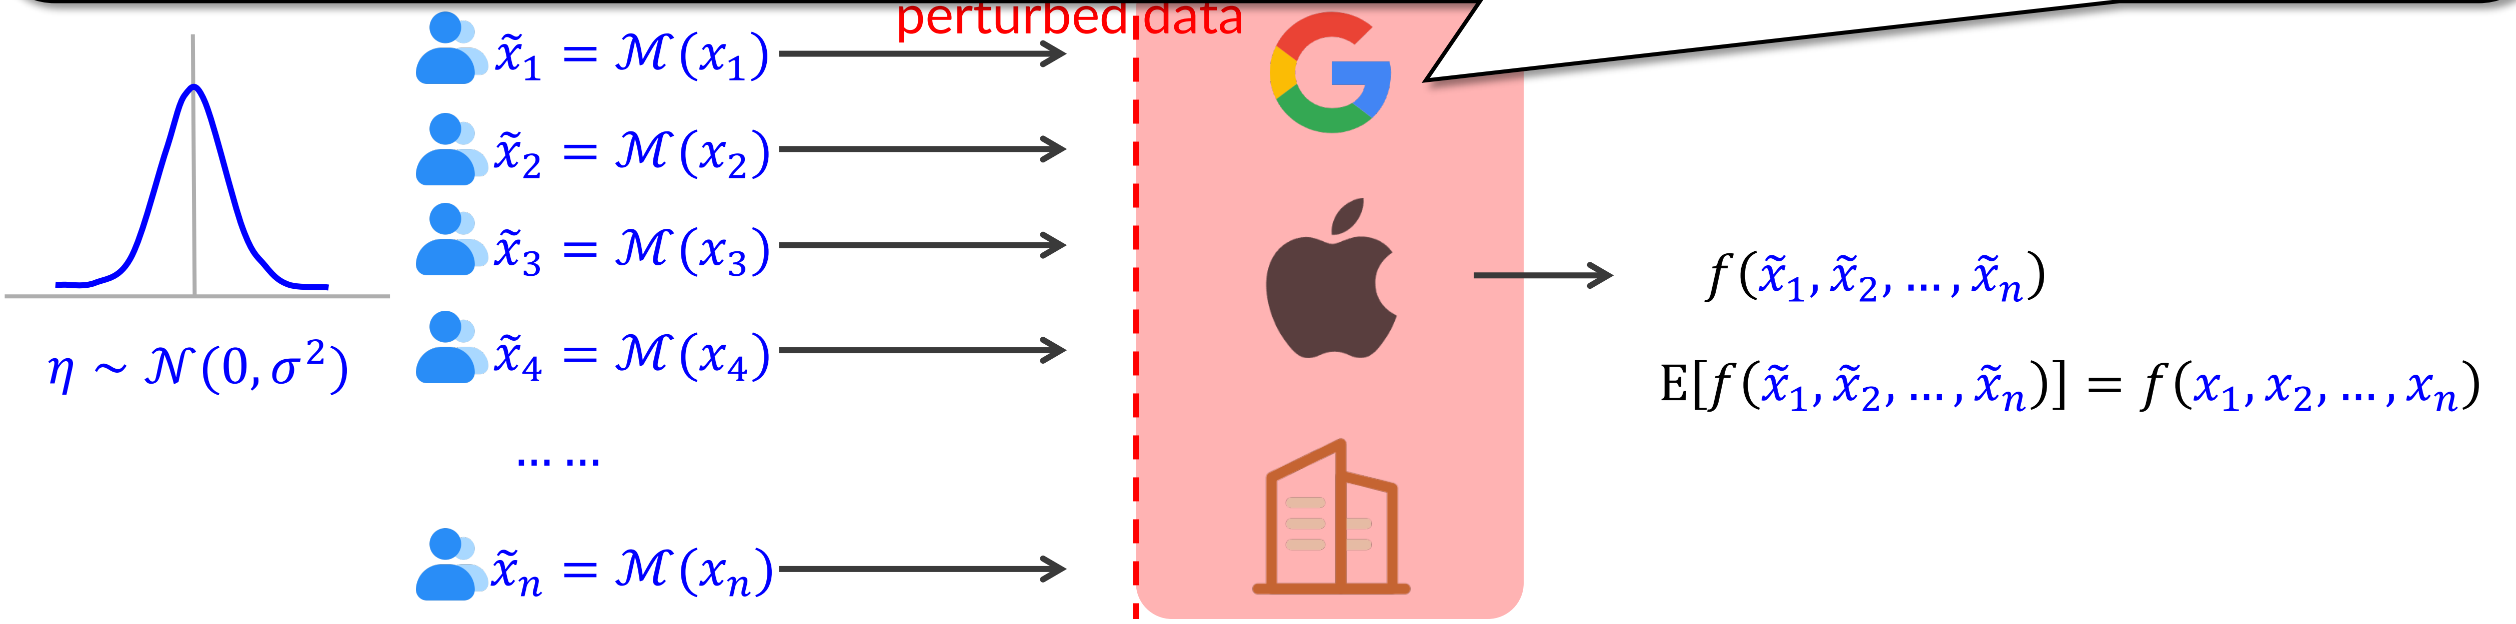

LDP – Applications

Chrome uses LDP to collect homepage settings, extension usage, etc

Emoji usage, new keyboard words, Safari URL statistics, health analytics

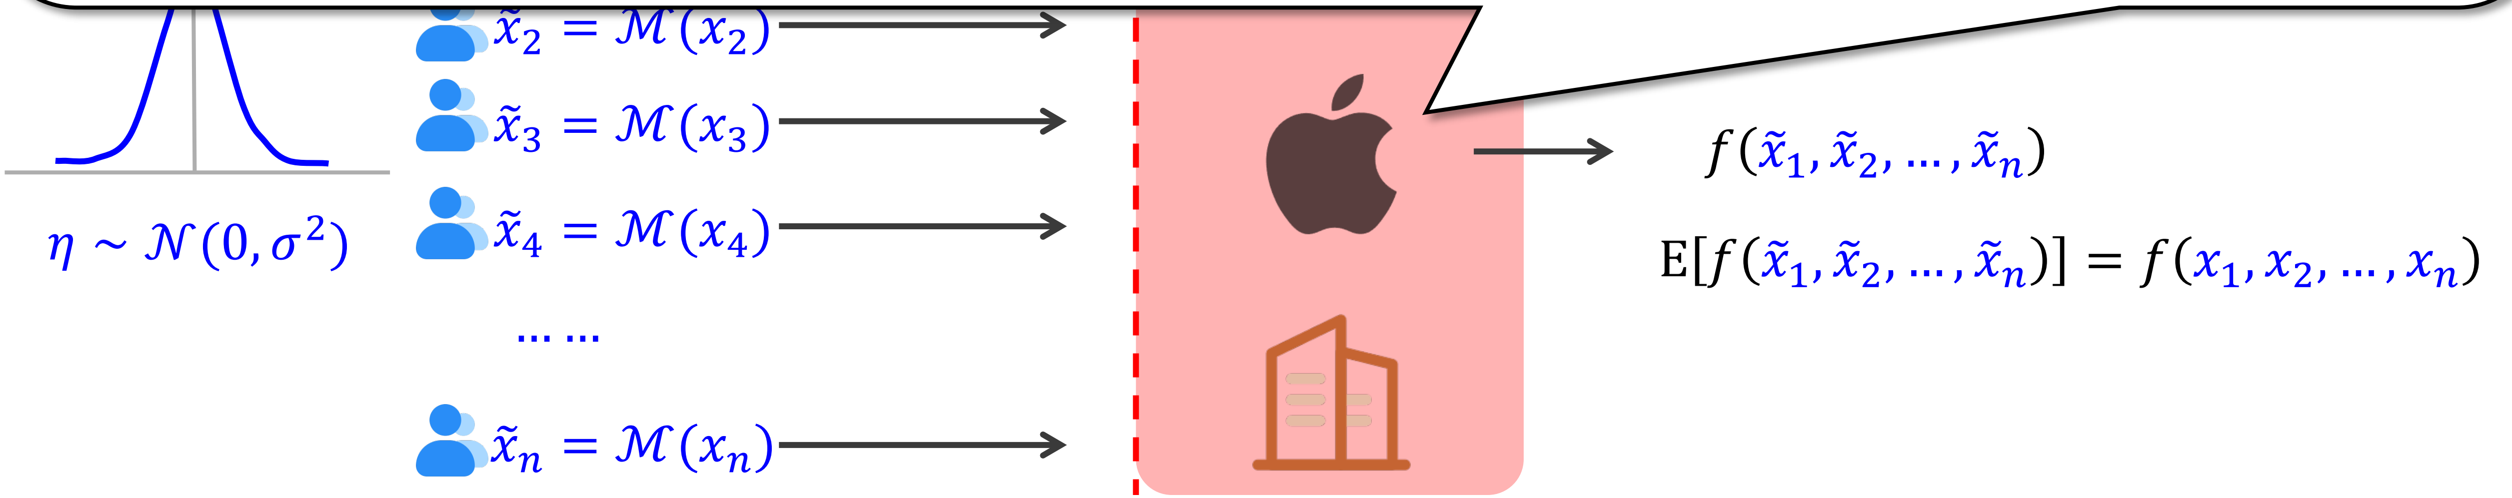

LDP – Formal Privacy

RIT

- After applying $\mathcal{M}$, the confidence of distinguishing sensitive x_1 and x_2 from observation $\tilde{x}$:

$$\forall x_1, x_2 \in \mathcal{D}, \forall y \in \widetilde{\mathcal{D}} \qquad \max \frac{\Pr[\mathcal{M}(x_1) = \tilde{x}]}{\Pr[\mathcal{M}(x_2) = \tilde{x}]} \leq e^{\varepsilon}$$

LDP – Formal Privacy

- After applying $\mathcal{M}$, the confidence of distinguishing sensitive x_1 and x_2 from observation $\tilde{x}$:

$$\forall x_1, x_2 \in \mathcal{D}, \forall y \in \widetilde{\mathcal{D}} \qquad \max \frac{\Pr[\mathcal{M}(x_1) = \tilde{x}]}{\Pr[\mathcal{M}(x_2) = \tilde{x}]} \leq e^{\varepsilon}$$

- The collector's / adversary's view: hard to infer the sensitive data

Privacy quantified by ε

$x_1 \rightarrow \mathcal{M} \rightarrow \tilde{x}$

Provable defense against data inference attacks

Data utility by approximated error

$$f(\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_n) \approx f(x_1, x_2, \ldots, x_n)$$

LDP – Formal Privacy

RIT

- After applying $\mathcal{M}$, the confidence of distinguishing sensitive x_1 and x_2 from observation $\tilde{x}$:

$$\forall x_1, x_2 \in \mathcal{D}, \forall y \in \widetilde{\mathcal{D}} \qquad \max \frac{\Pr[\mathcal{M}(x_1) = \tilde{x}]}{\Pr[\mathcal{M}(x_2) = \tilde{x}]} \leq e^{\varepsilon}$$

- The collector's / adversary's view: hard to infer the sensitive data

Privacy quantified by ε

$x_1 \rightarrow \mathcal{M} \rightarrow \tilde{x}$

Data utility by approximated error

Fundamental direction: **Design** of $\mathcal{M}$ to **optimize** the privacy—utility tradeoff

LDP – Formal Privacy

Utility analysis of $f \circ \mathcal{M}$

$$f(x_1, x_2, \dots, x_n) \coloneqq \sum_{i=1}^{n} x_i \quad \text{or} \quad f(x_1, x_2, \dots, x_n) \coloneqq \{x_1, x_2, \dots, x_n\} \rightarrow$$ Variance, MSE

$$f(x_1, x_2, \dots, x_n) \coloneqq h: \mathbb{R}^n \rightarrow \{1,2,\dots,K\}$$ is a classifier $\rightarrow$?

Privacy quantified by ε

$x_1 \rightarrow \mathcal{M} \rightarrow \tilde{x}$

Provable defense against data inference attacks

Data utility by app... error

dist of $f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$

$f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n) \approx f(x_1, x_2, \dots, x_n)$

LDP – Formal Privacy

Utility analysis of $f \circ \mathcal{M}$

$f(x_1, x_2, \dots, x_n) \coloneqq \sum_{i=1}^{n} x_i$ or $f(x_1, x_2, \dots, x_n) \coloneqq \{x_1, x_2, \dots, x_n\}$ → Variance, MSE

$f(x_1, x_2, \dots, x_n) \coloneqq h: \mathbb{R}^n \to \{1,2, \dots, K\}$ is a classifier → ?

Privacy quantified by ε

$x_1 \to \mathcal{M} \to \tilde{x}$

Data utility by app … error

Fundamental direction: **Utility analysis** of complex task f

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

Part 1: correlated $\mathcal{M}$

Part 2: optimal piecewise-based $\mathcal{M}$

Part 3: $\mathcal{M}$ for trajectories in continuous space

Part 4: utility analysis for classifier $\circ\, \mathcal{M}$

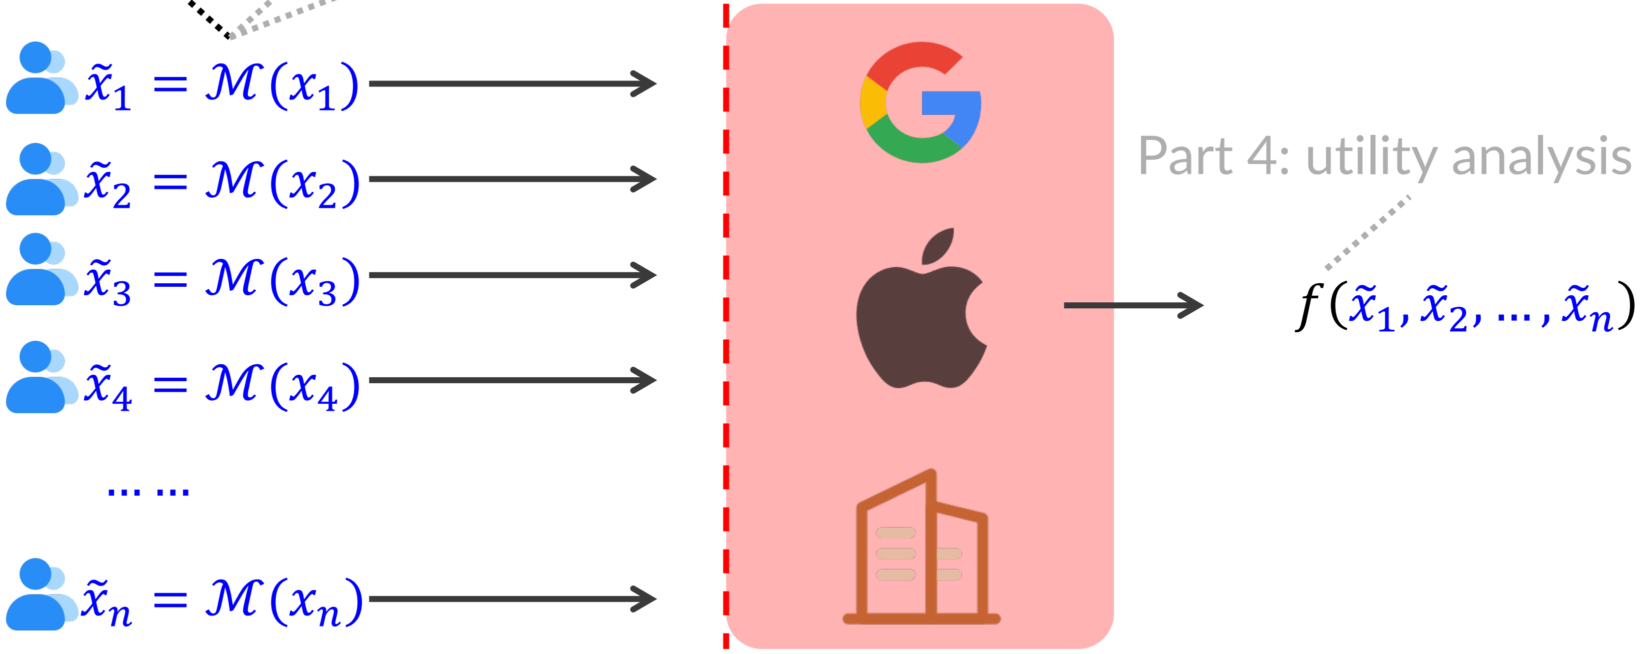

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

binary x → numerical x

Part 1: correlated $\mathcal{M}$ → **Part 2: optimal piecewise-based $\mathcal{M}$**

Part 3: $\mathcal{M}$ for trajectories in continuous space

Part 4: utility analysis for classifier $\circ\, \mathcal{M}$

f is a classifier

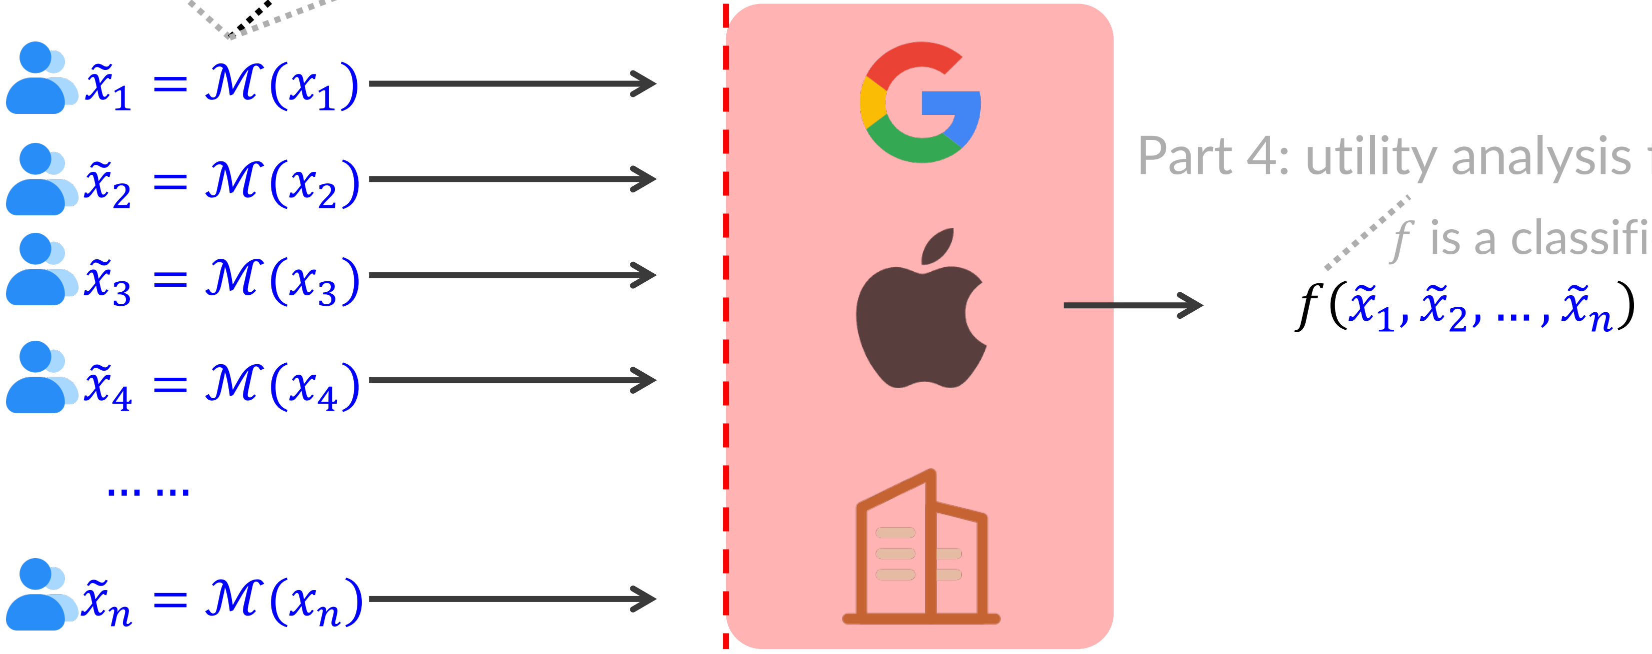

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

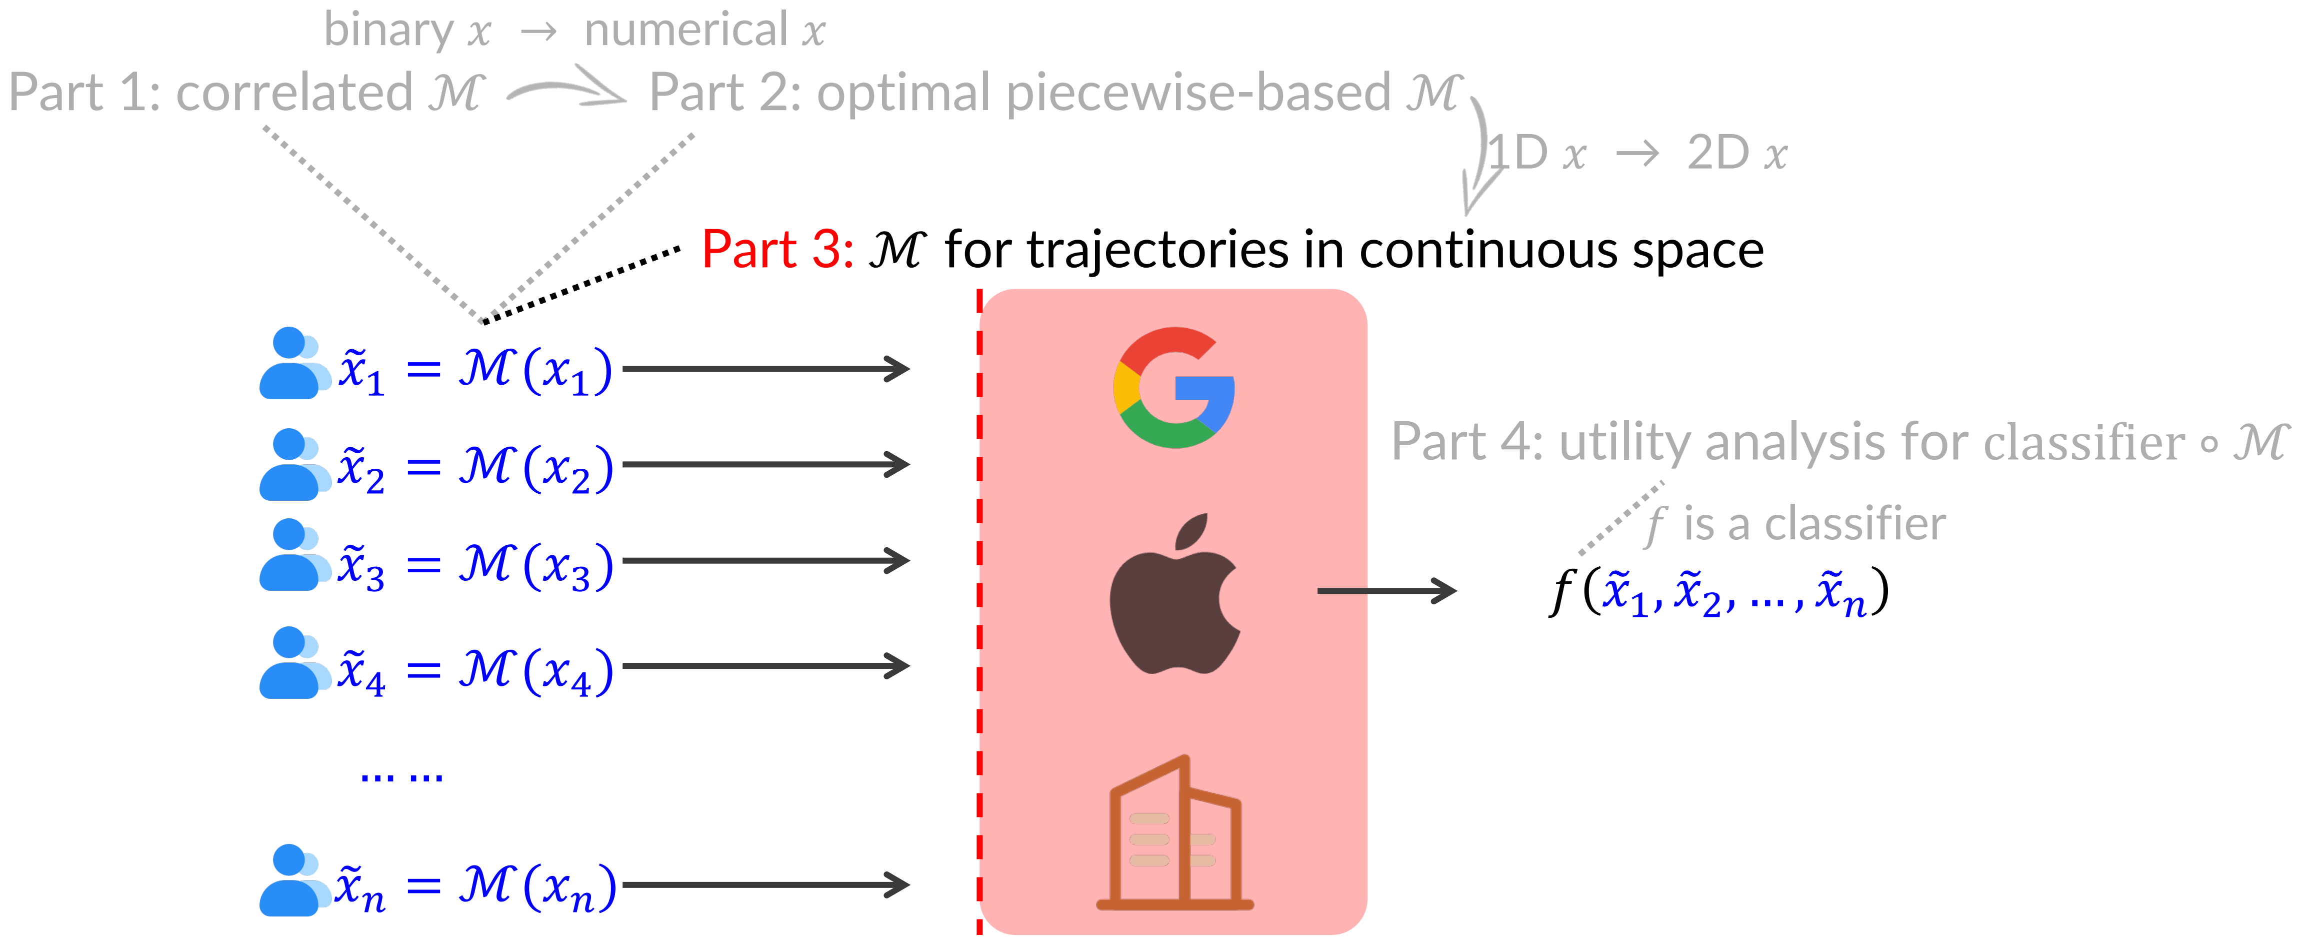

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

binary x → numerical x

Part 1: correlated $\mathcal{M}$ → Part 2: optimal piecewise-based $\mathcal{M}$

1D x → 2D x

Part 3: $\mathcal{M}$ for trajectories in continuous space

Mechanism-level
↓
Task-level

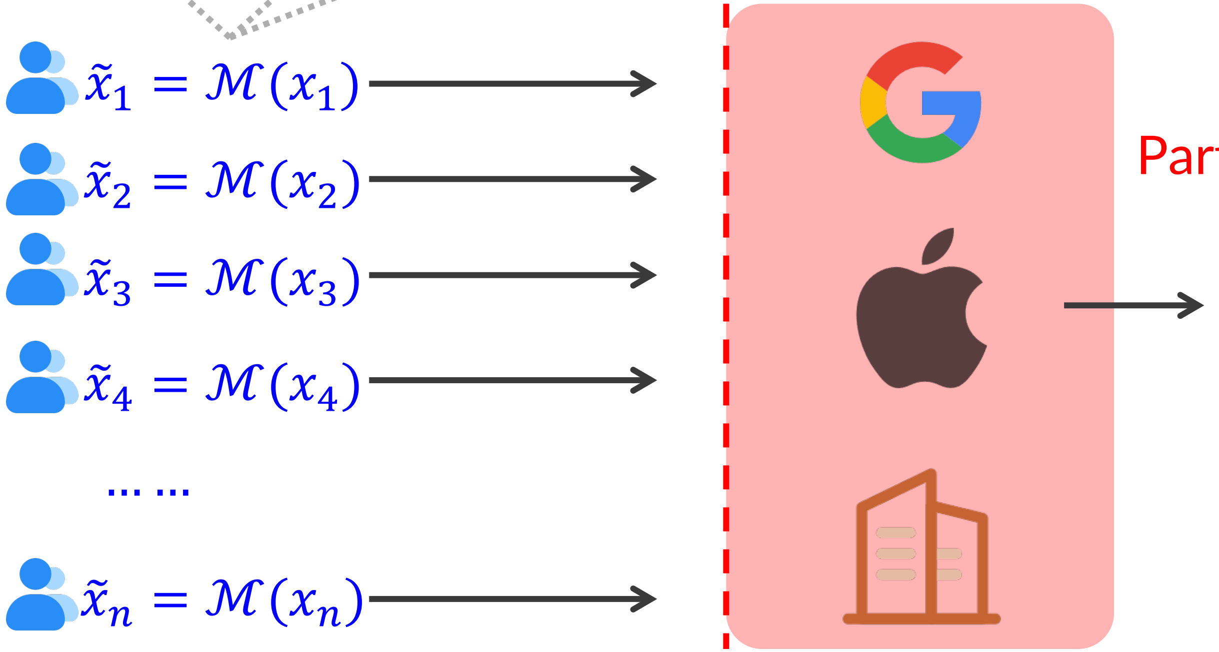

Part 4: utility analysis for classifier $\circ\ \mathcal{M}$

f is a classifier

$f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$

Part 1 – Correlated $\mathcal{M}$*

- Existing LDP mechanisms: Each user perturbs their data independently

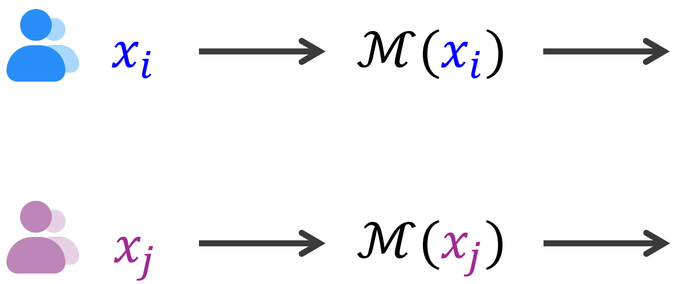

* [PETS'25] Locally Differentially Private Frequency Estimation via Joint Randomized Response

Part 1 – Correlated $\mathcal{M}$*

- Existing LDP mechanisms: Each user perturbs their data independently

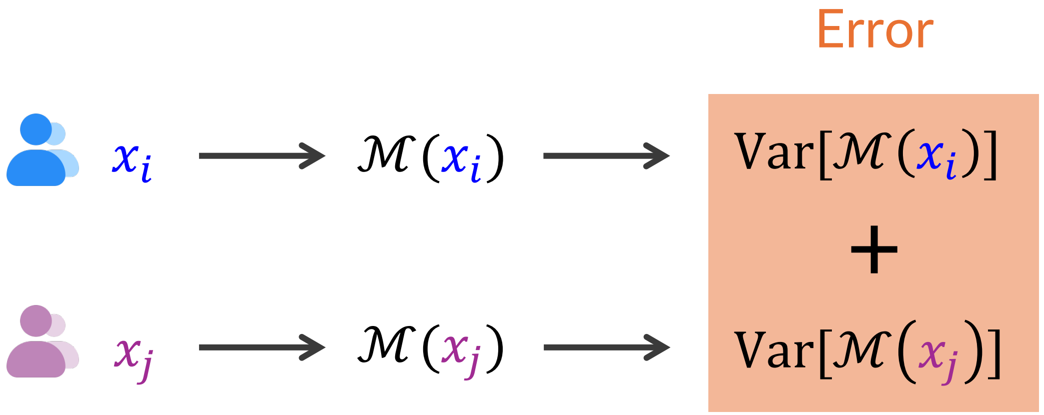

* [PETS'25] Locally Differentially Private Frequency Estimation via Joint Randomized Response

Part 1 – Correlated $\mathcal{M}$*

RIT

- Existing LDP mechanisms: Each user perturbs their data independently
- Correlated LDP mechanisms: Users' data are perturbed by correlated $\mathcal{M}$

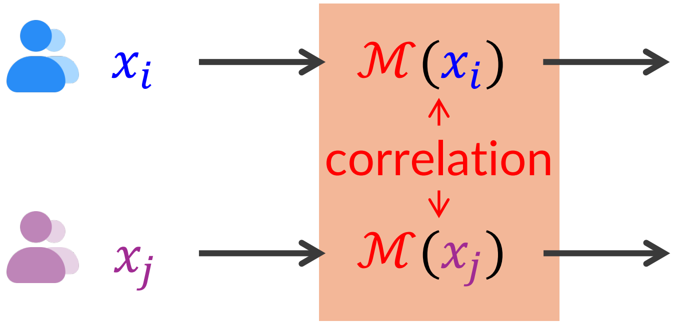

* [PETS'25] Locally Differentially Private Frequency Estimation via Joint Randomized Response

Part 1 – Correlated $\mathcal{M}$*

RIT

- Existing LDP mechanisms: Each user perturbs their data independently
- Correlated LDP mechanisms: Users' data are perturbed by correlated $\mathcal{M}$

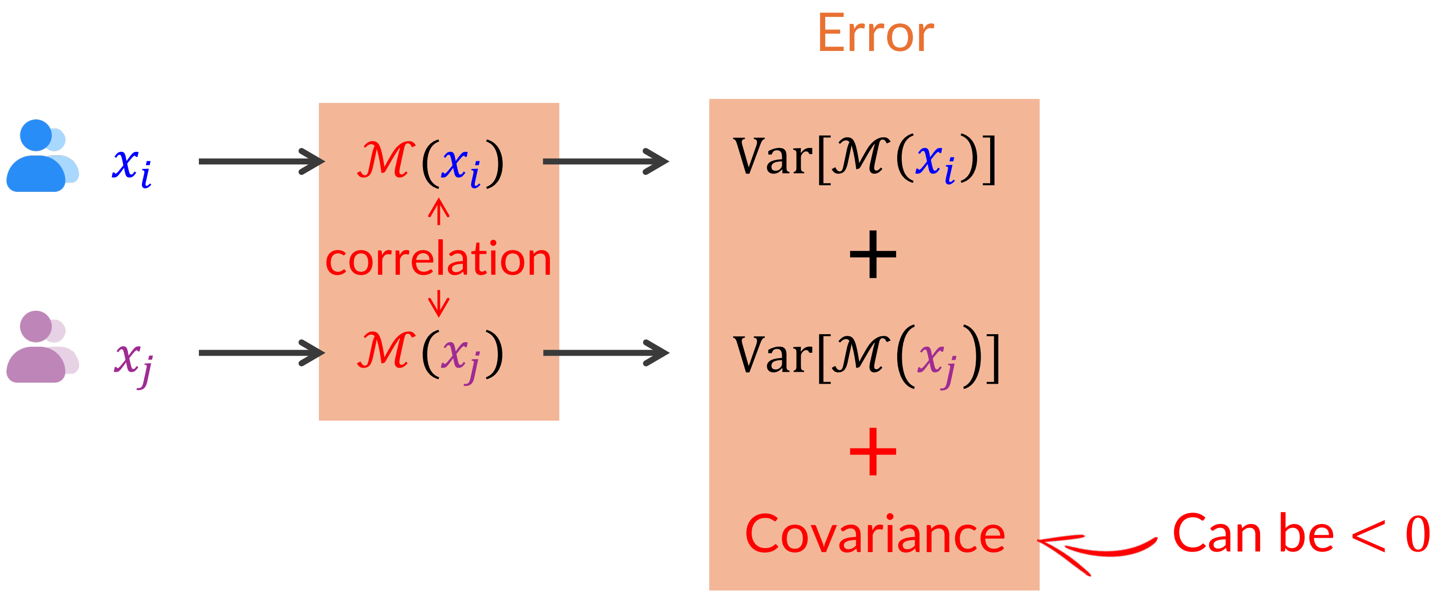

* [PETS'25] Locally Differentially Private Frequency Estimation via Joint Randomized Response

Part 2 – Optimal Piecewise-based $\mathcal{M}^\dagger$

RIT

binary x → numerical x

- SOTA for bounded numerical data: Piecewise-based mechanisms (3-piece heuristic PDF)

$\tilde{x} \leftarrow \mathcal{M}(x)$: sampling PDFs

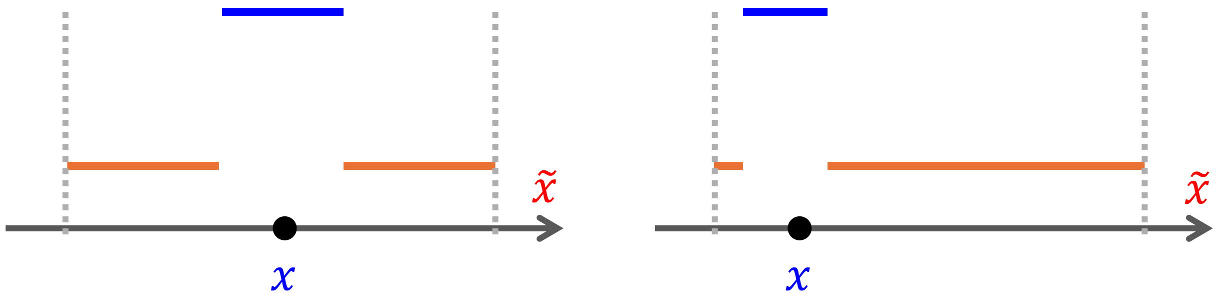

† [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

Part 2 – Optimal Piecewise-based $\mathcal{M}^{\dagger}$

RIT

- SOTA for bounded numerical data: Piecewise-based mechanisms (3-piece heuristic PDF)

$\tilde{x} \leftarrow \mathcal{M}(x)$: sampling PDFs

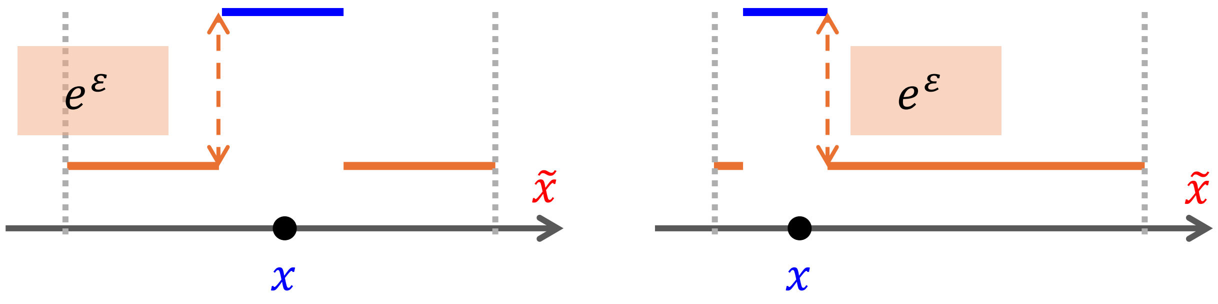

$$\text{pdf}[\mathcal{M}(x) = \tilde{x}] = \begin{cases} p_\varepsilon & \text{if } \tilde{x} \in [l_{x,\varepsilon}, r_{x,\varepsilon}] \\ \frac{p_\varepsilon}{e^\varepsilon} & \text{if } \tilde{x} \in \tilde{\mathcal{D}} \backslash [l_{x,\varepsilon}, r_{x,\varepsilon}] \end{cases}$$

† [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

Part 2 – Optimal Piecewise-based $\mathcal{M}^{\dagger}$

RIT

binary x → numerical x

- SOTA for bounded numerical data: Piecewise-based mechanisms (3-piece heuristic PDF)
- Too heuristic → More generalized version

$\tilde{x} \leftarrow \mathcal{M}(x)$: general sampling PDFs

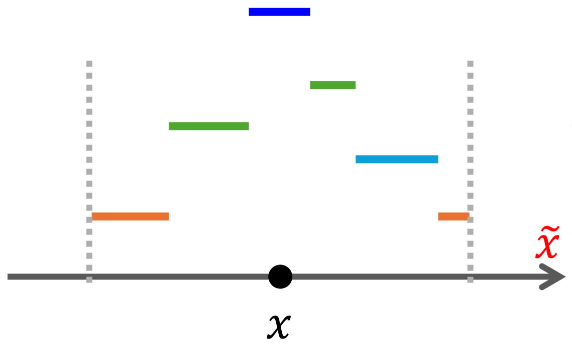

Potentially has lower error

† [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

Part 2 – Optimal Piecewise-based $\mathcal{M}^{\dagger}$

RIT

binary x → numerical x

- SOTA for bounded numerical data: Piecewise-based mechanisms (3-piece heuristic PDF)
- Too heuristic → More generalized version

$\tilde{x} \leftarrow \mathcal{M}(x)$: general sampling PDFs

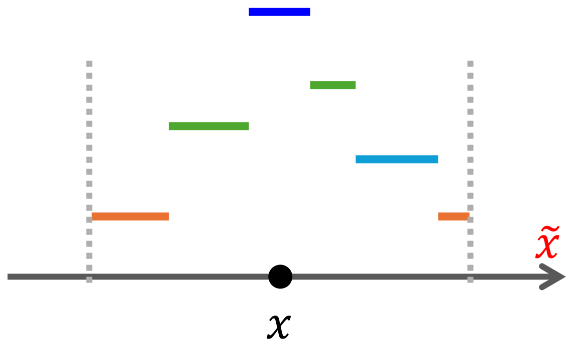

Potentially has lower error

- What is the optimal piecewise-based mechanism?

† [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

Part 2 – Optimal Piecewise-based $\mathcal{M}^{\dagger}$

RIT

binary x → numerical x

- Linear data domain → Circular data domain

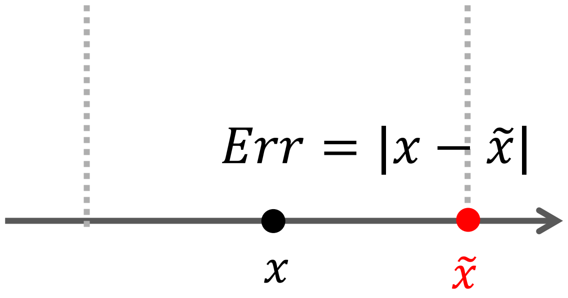

† [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

Part 2 – Optimal Piecewise-based $\mathcal{M}^\dagger$

RIT

- Linear data domain → Circular data domain

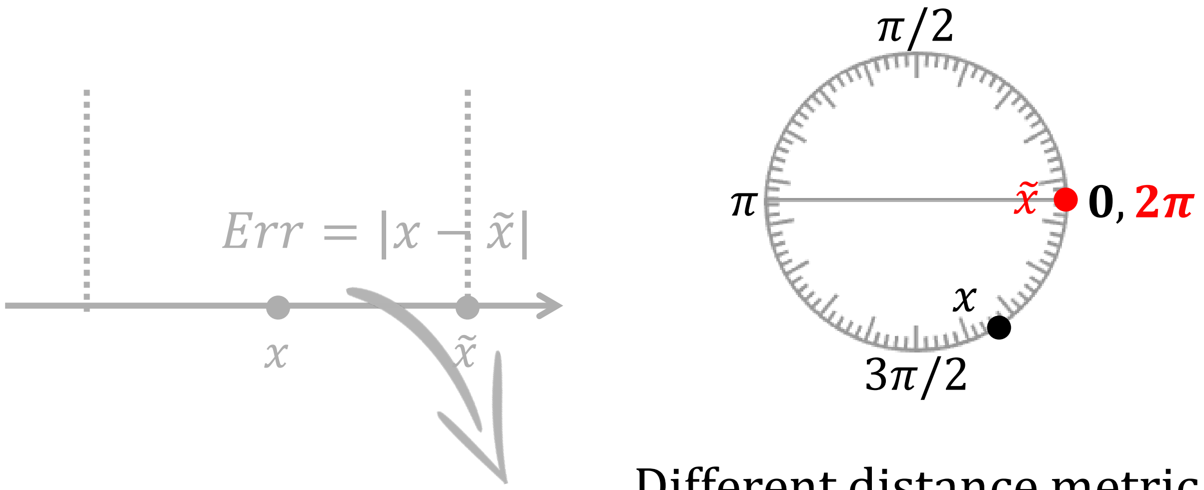

Different distance metric

$$Err = \min(|x - \tilde{x}|, |2\pi - x - \tilde{x}|)$$

† [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

Part 2 – Optimal Piecewise-based $\mathcal{M}^{\dagger}$

RIT

- Linear data domain → Circular data domain

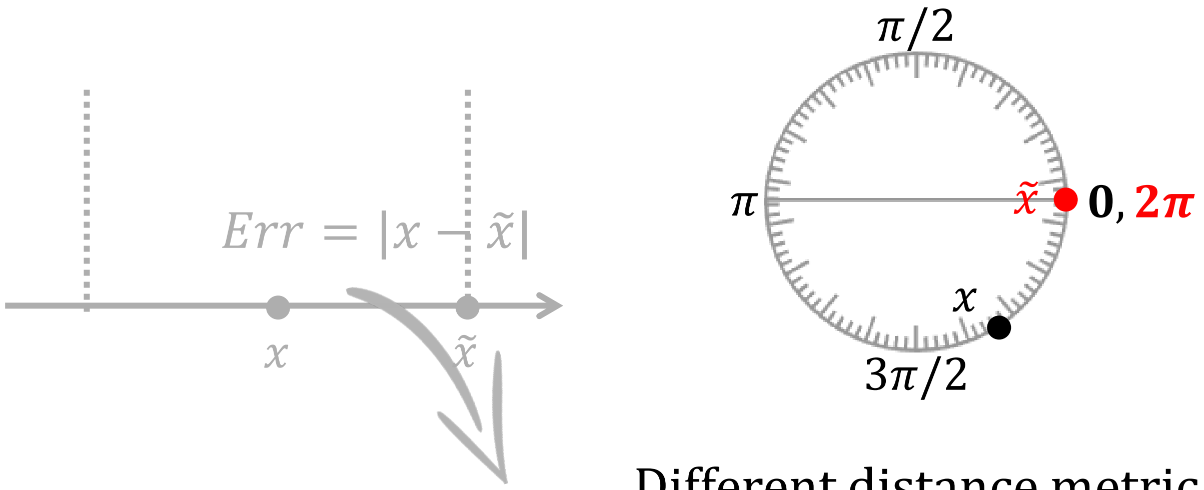

Different distance metric

$$Err = \min(|x - \tilde{x}|, |2\pi - x - \tilde{x}|)$$

- What is the optimal piecewise-based mechanism for circular domain?

† [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

Part 3 – $\mathcal{M}$ for Trajectories in Continuous Space‡

- Existing $\mathcal{M}$ for trajectory collection: assuming discrete location space

$$\mathcal{S} = \{p_1, \ldots, p_n\}$$

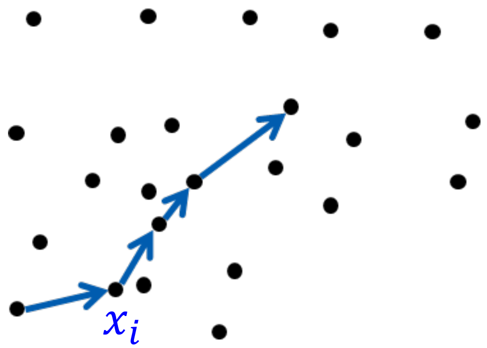

‡ [PETS'26] TraCS: Trajectory Collection in Continuous Space under Local Differential Privacy

Part 3 – $\mathcal{M}$ for Trajectories in Continuous Space‡

1D x → 2D x

- Existing $\mathcal{M}$ for trajectory collection: assuming discrete location space

$\mathcal{S} = \{p_1, \dots, p_n\}$

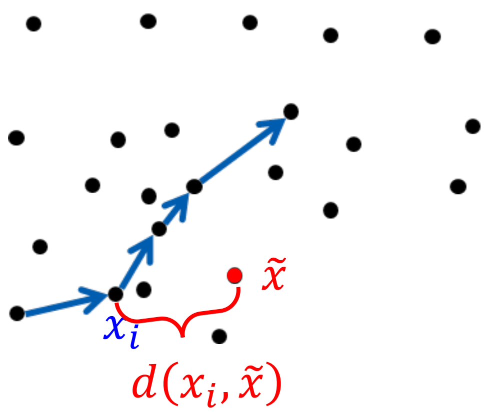

Make each x_i satisfying LDP:

For $\tilde{x} \in S$

$\Pr[\mathcal{M}(x_i) = \tilde{x}] \propto e^{d(x_i, \tilde{x})}$

‡ [PETS'26] TraCS: Trajectory Collection in Continuous Space under Local Differential Privacy

Part 3 – $\mathcal{M}$ for Trajectories in Continuous Space‡

1D x → 2D x

- Existing $\mathcal{M}$ for trajectory collection: assuming discrete location space

$\mathcal{S} = \{p_1, \dots, p_n\}$

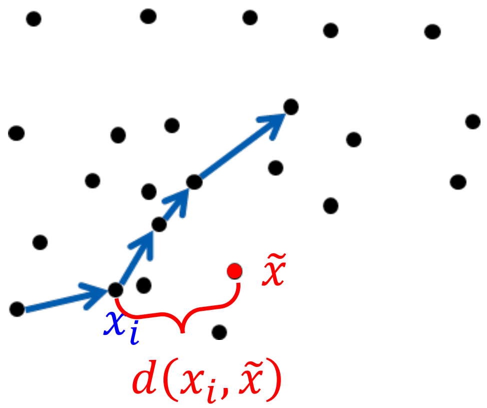

Make each x_i satisfying LD

For $\tilde{x} \in S$

$\Pr[\mathcal{M}(x_i) = \tilde{x}] \propto e^{d(x_i, \tilde{x})}$

Limitations

- expensive to sample
- only applicable to discrete S

‡ [PETS'26] TraCS: Trajectory Collection in Continuous Space under Local Differential Privacy

Part 3 – $\mathcal{M}$ for Trajectories in Continuous Space‡

1D x → 2D x

- Discrete location space → Continuous location space

$$\mathcal{S}_{\text{dis}} \subset \mathcal{S} = [0,1.5] \times [0,1]$$

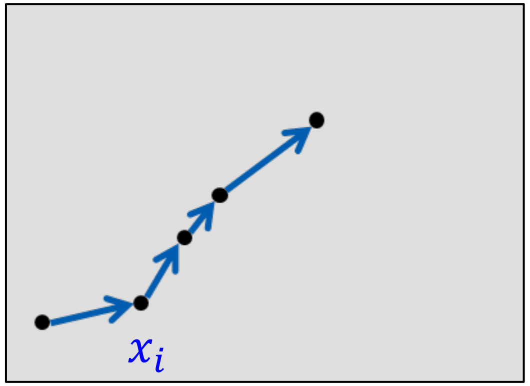

‡ [PETS'26] TraCS: Trajectory Collection in Continuous Space under Local Differential Privacy

Part 3 – $\mathcal{M}$ for Trajectories in Continuous Space‡

1D x → 2D x

- Discrete location space → Continuous location space

$\mathcal{S}_{\text{dis}} \subset \mathcal{S} = [0,1.5] \times [0,1]$

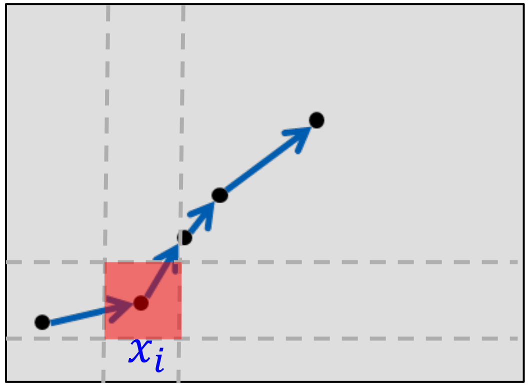

Make each x_i satisfying LDP in $\mathcal{S}$

$$\Pr[\mathcal{M}(x_i) = \tilde{x} \in \blacksquare] = p_\varepsilon$$

$$\Pr[\mathcal{M}(x_i) = \tilde{x} \in \blacksquare] = p_\varepsilon / e^\varepsilon$$

‡ [PETS'26] TraCS: Trajectory Collection in Continuous Space under Local Differential Privacy

Part 3 – $\mathcal{M}$ for Trajectories in Continuous Space[‡]

1D x → 2D x

- Discrete location space → Continuous location space

$\mathcal{S}_{\text{dis}} \subset \mathcal{S} = [0,1.5] \times [0,1]$

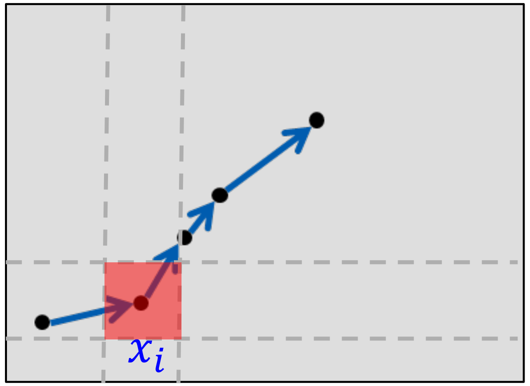

Make each x_i satisfying LD

$\Pr[\mathcal{M}(x_i) = \tilde{x} \in ■] = p_\varepsilon$

$\Pr[\mathcal{M}(x_i) = \tilde{x} \in ■] = p_\varepsilon / e^{\varepsilon}$

Benefits

- negligible sampling complexity
- applicable to discrete space by post-processing

‡ [PETS'26] TraCS: Trajectory Collection in Continuous Space under Local Differential Privacy

Part 4 – Utility Analysis for classifier $\circ\ \mathcal{M}$¶

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$

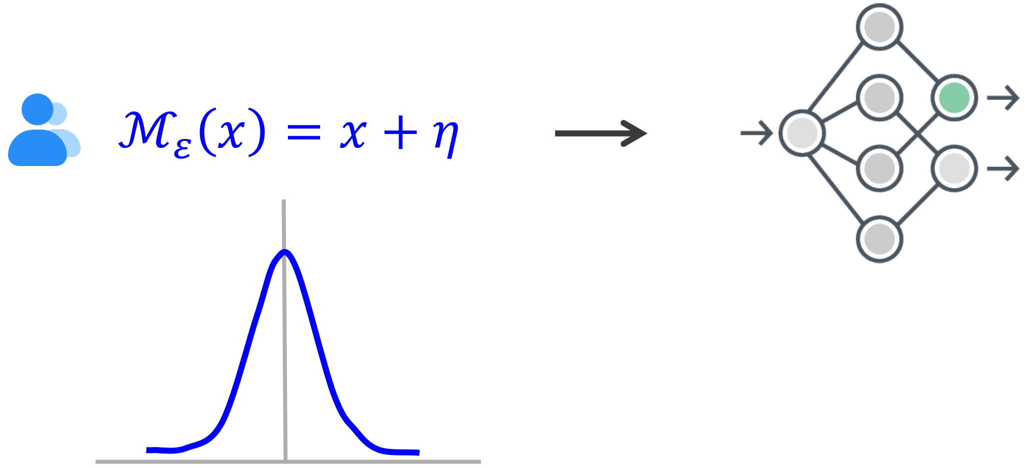

$\mathcal{M}_\varepsilon(x) = x + \eta$ →

$\eta \sim \mathcal{N}(0, \sigma^2)$

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier $\circ\ \mathcal{M}$¶

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$

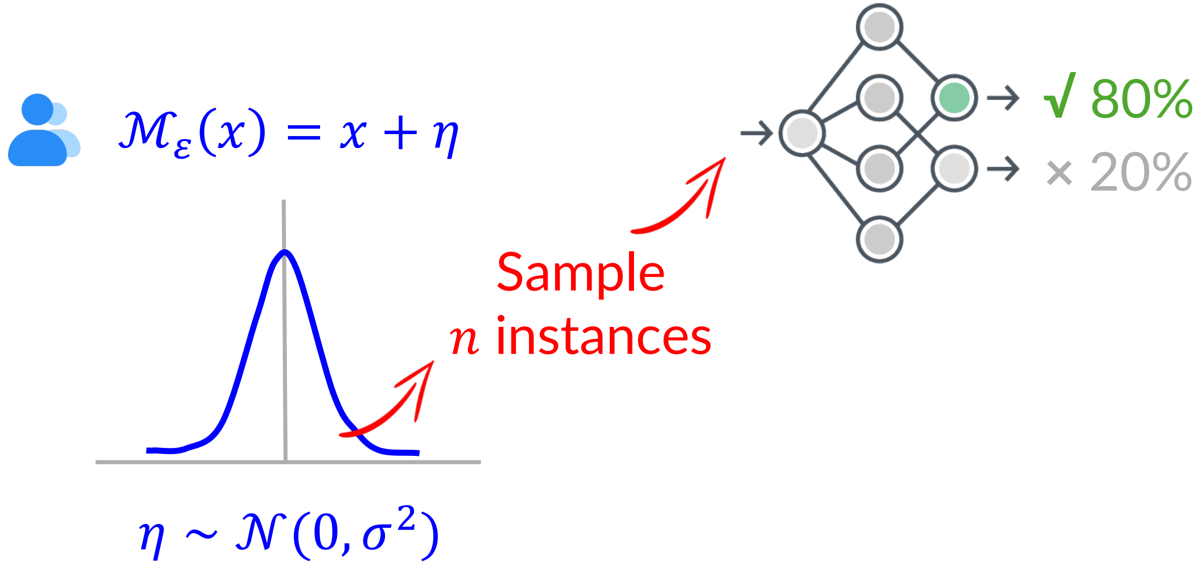

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier $\circ\ \mathcal{M}$[¶]

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$

$$\mathcal{M}_{\varepsilon'}(x) = x + \eta \longrightarrow$$

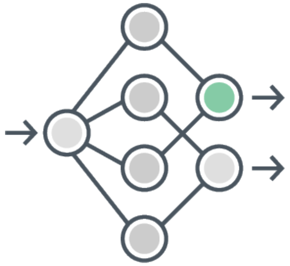

Change ε'

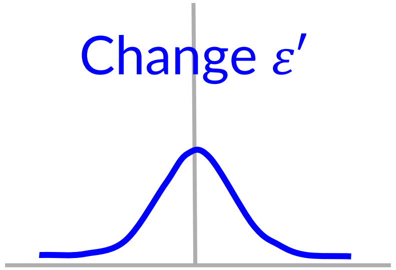

$$\eta \sim \mathcal{N}(0, \sigma'^2)$$

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier $\circ\ \mathcal{M}$¶

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$

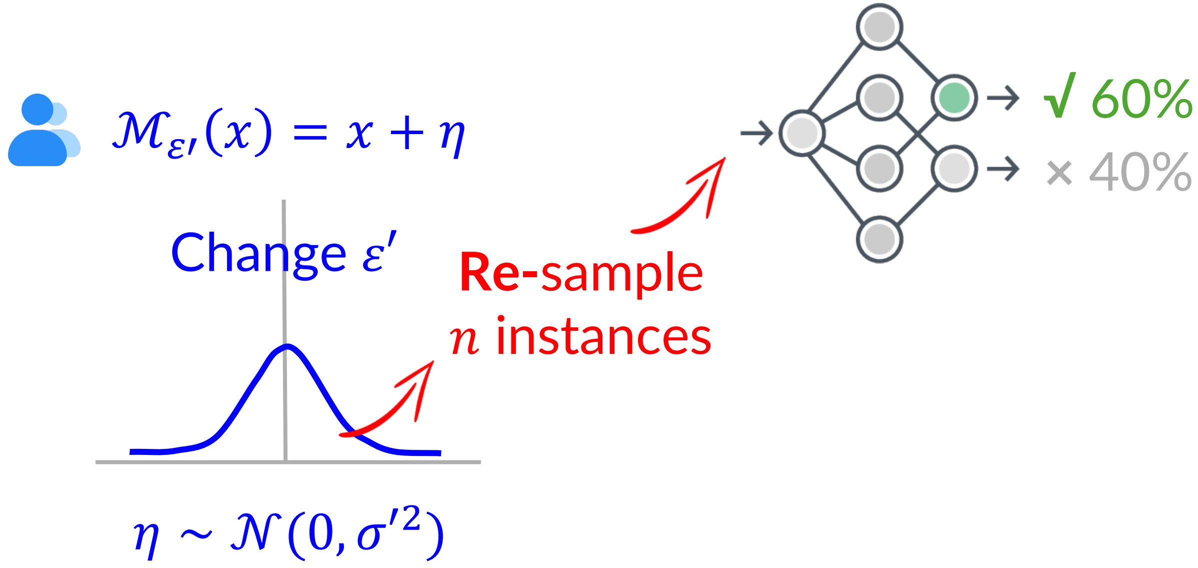

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier $\circ\ \mathcal{M}$[¶]

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$

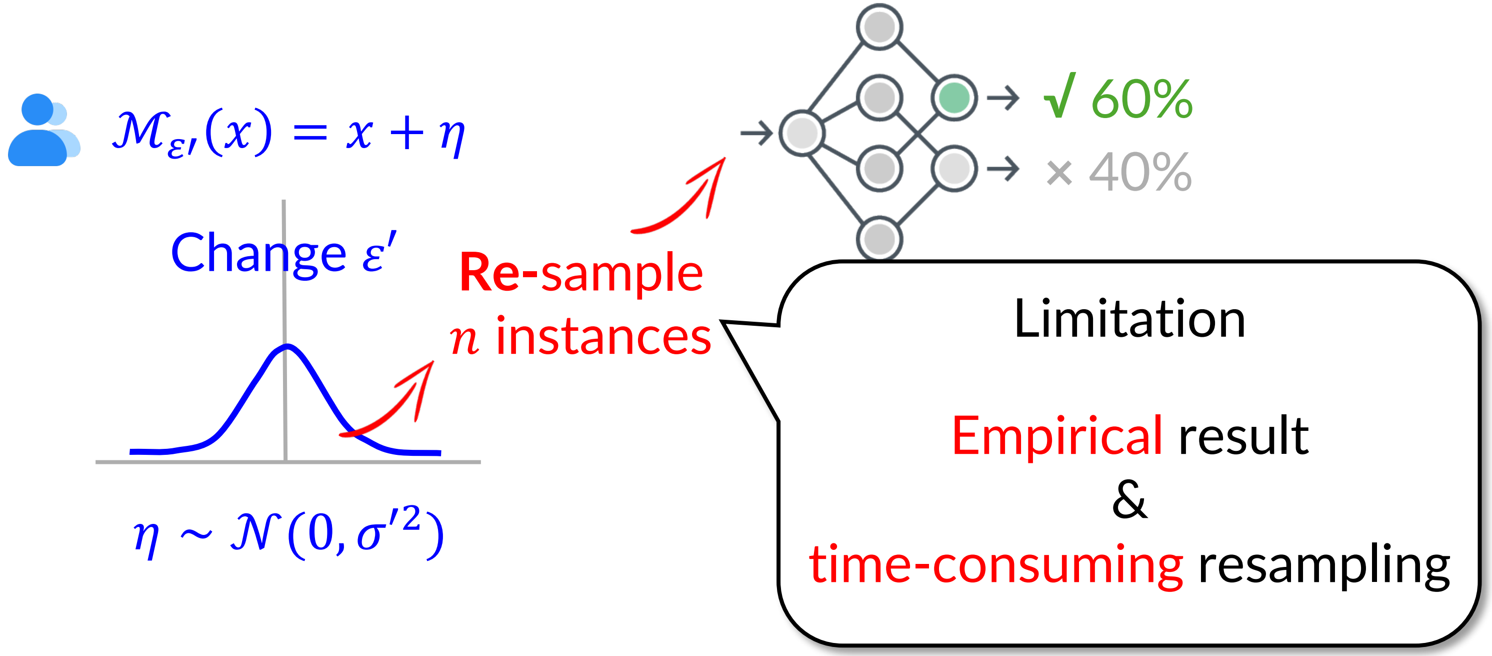

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier $\circ\ \mathcal{M}$¶

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$ → Analytical classifier utility

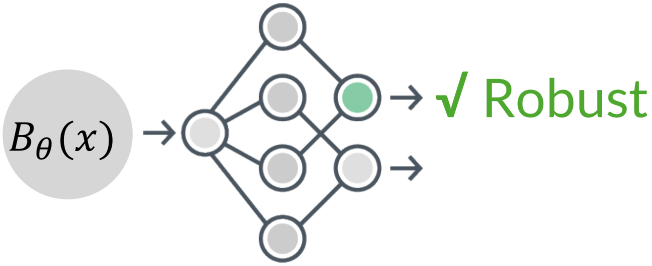

Robustness:
$B_\theta(x)$ is robust

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier ∘ $\mathcal{M}$¶

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$ → Analytical classifier utility

$\mathcal{M}_\varepsilon(x) = x + \eta$

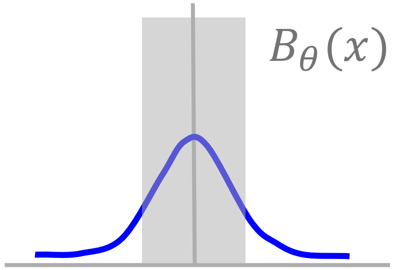

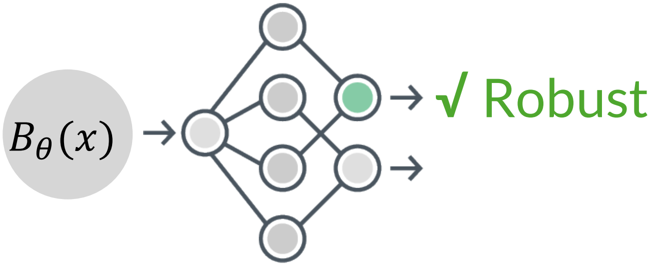

Concentration:

$\Pr[\mathcal{M}_\varepsilon(x) \in B_\theta(x)]$

$:= p(\varepsilon, \theta)$

Robustness:

$B_\theta(x)$ is robust

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier ∘ $\mathcal{M}$ ¶

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$ → Analytical classifier utility

$\mathcal{M}_\varepsilon(x) = x + \eta$

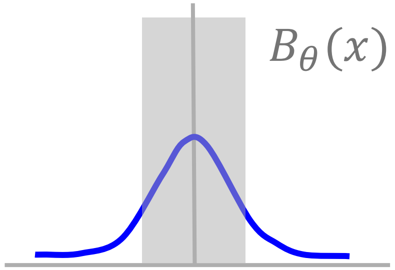

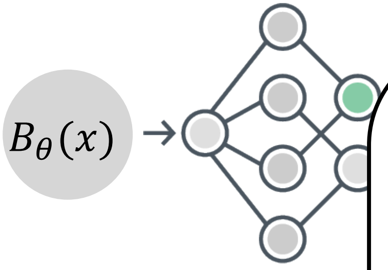

Concentration:
$\Pr[\mathcal{M}_\varepsilon(x) \in B_\theta(x)]$
$:= p(\varepsilon, \theta)$

Robustr
$B_\theta(x)$ is

Connected by θ

Analytical classifier utility

> *With probability at least $p(\varepsilon, \theta)$, the classifier preserves its correct classification under input $\mathcal{M}_\varepsilon(x)$.*

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Part 4 – Utility Analysis for classifier $\circ\, \mathcal{M}$¶

mechanism-level → task level

- Empirical classifier utility under $\mathcal{M}$ → Analytical classifier utility

$\mathcal{M}_{\varepsilon'}(x) = x + \eta$

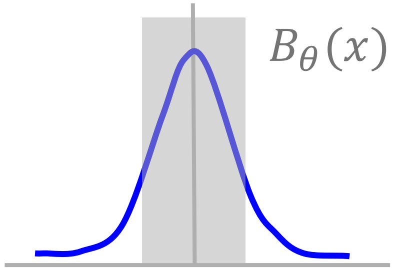

Concentration:
$\Pr[\mathcal{M}_{\varepsilon'}(x) \in B_\theta(x)]$
$:= p(\varepsilon', \theta)$

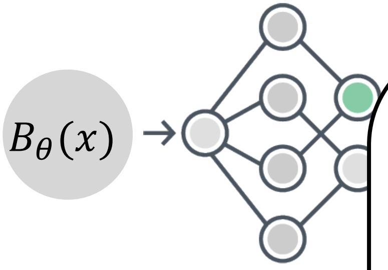

Robustr

$B_\theta(x)$ is

Connected by θ

Analytical classifier utility

With probability at least $p(\varepsilon', \theta)$, the classifier preserves its correct classification under input $\mathcal{M}_\varepsilon(x)$.

¶ [PETS'26] Quantifying Classifier Utility under Local Differential Privacy

Societal Impact

- **New LDP building blocks**
  - correlated LDP mechanisms
  - optimal piecewise-based mechanisms

Sensor networks & Federated learning, etc

- **Universal trajectory collection mechanisms**
  - applicable to both continuous / discrete space

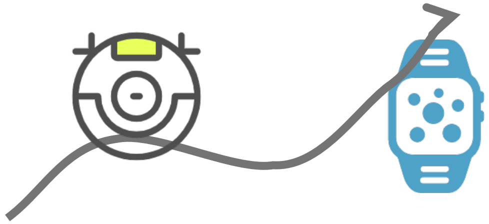

Smart home & wearable devices' trajectories, etc

- **Analytical view of classifier utility under LDP-perturbed inputs**
  - choosing best ε when using classifiers