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

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

Location, browsing history, app usage data

Collector

Analysis & service

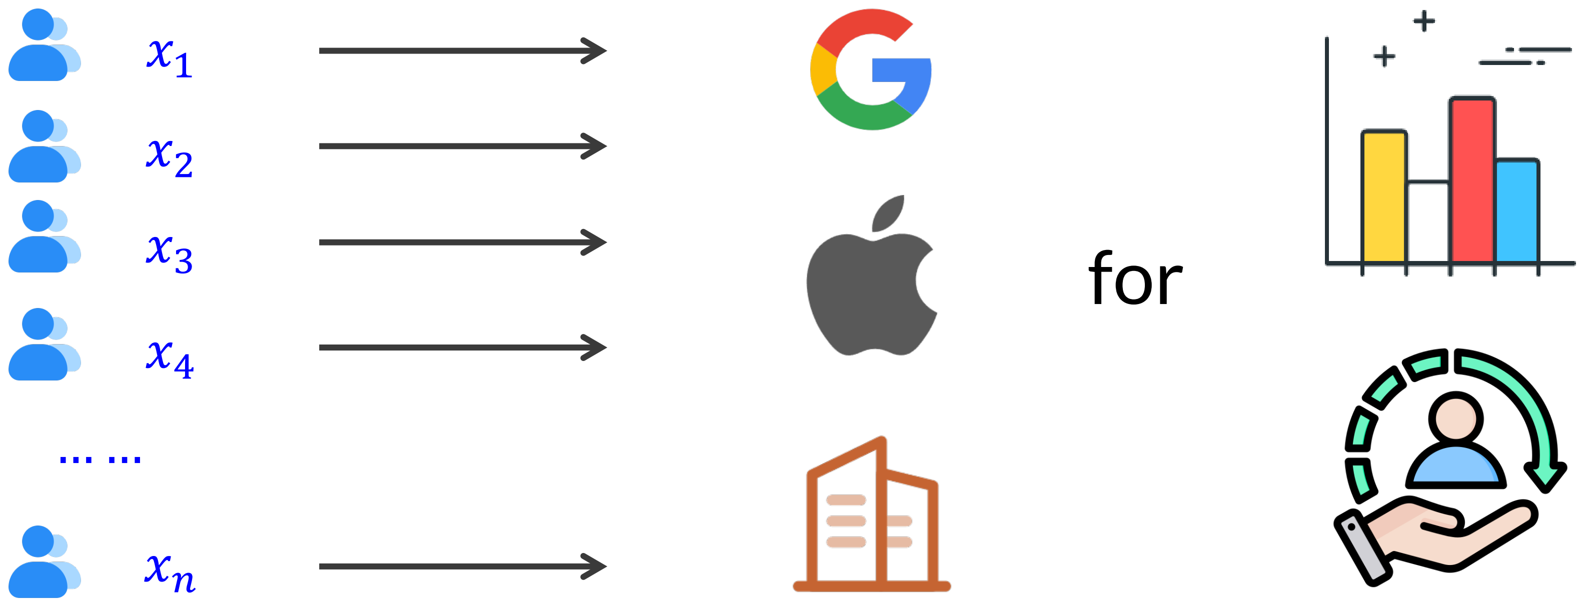

x_1

x_2

x_3

x_4

… …

x_n

for

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

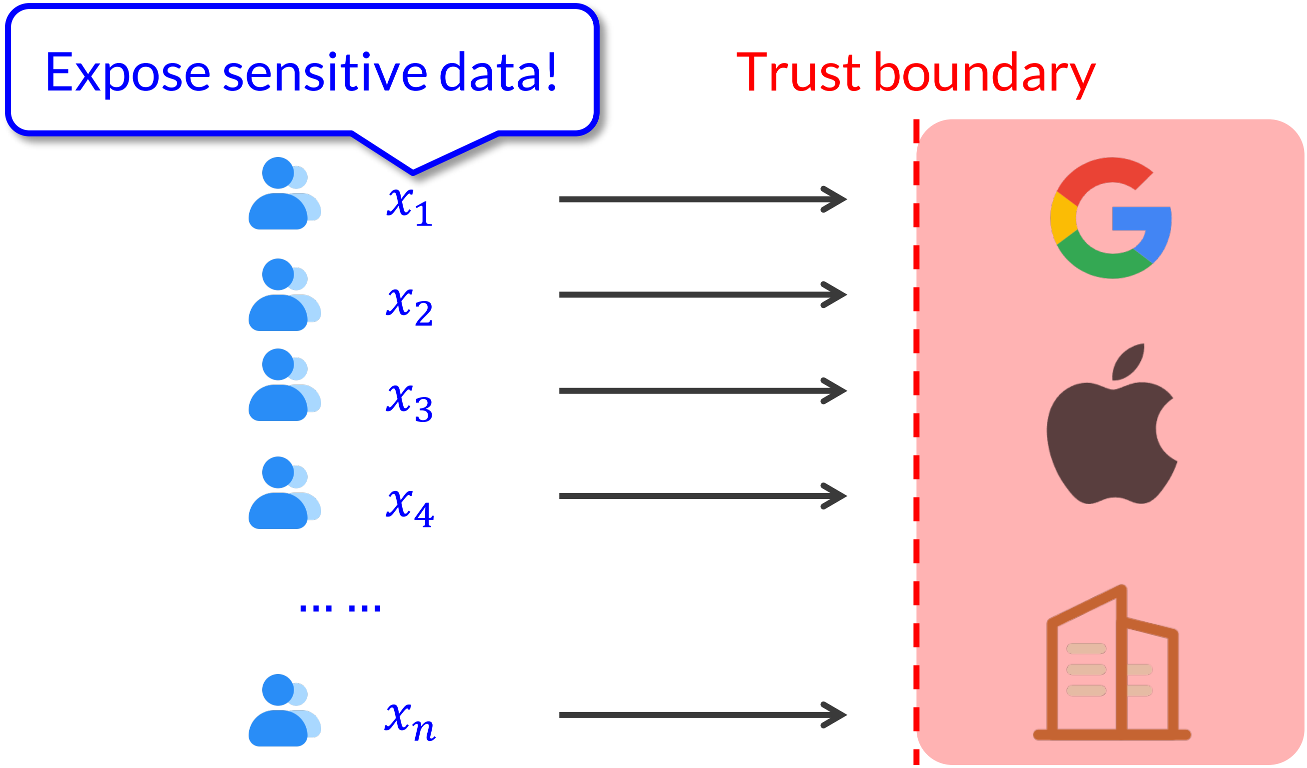

Users’ Data Privacy

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

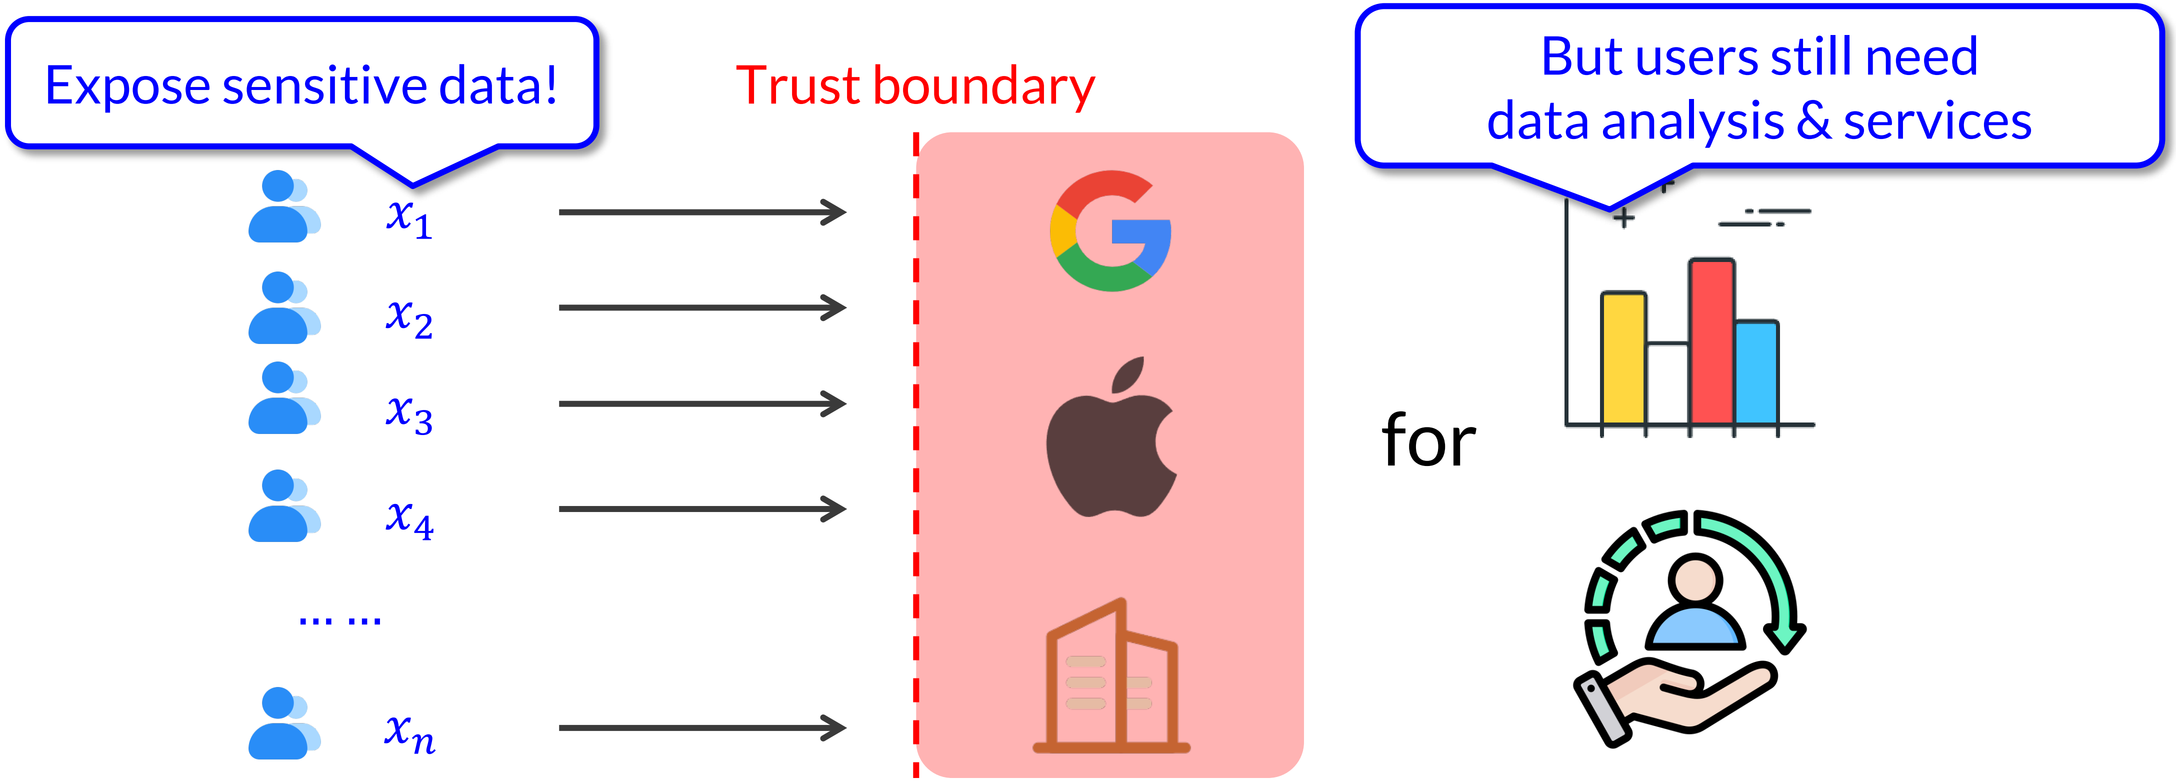

Expose sensitive data!

Trust boundary

But users still need data analysis & services

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

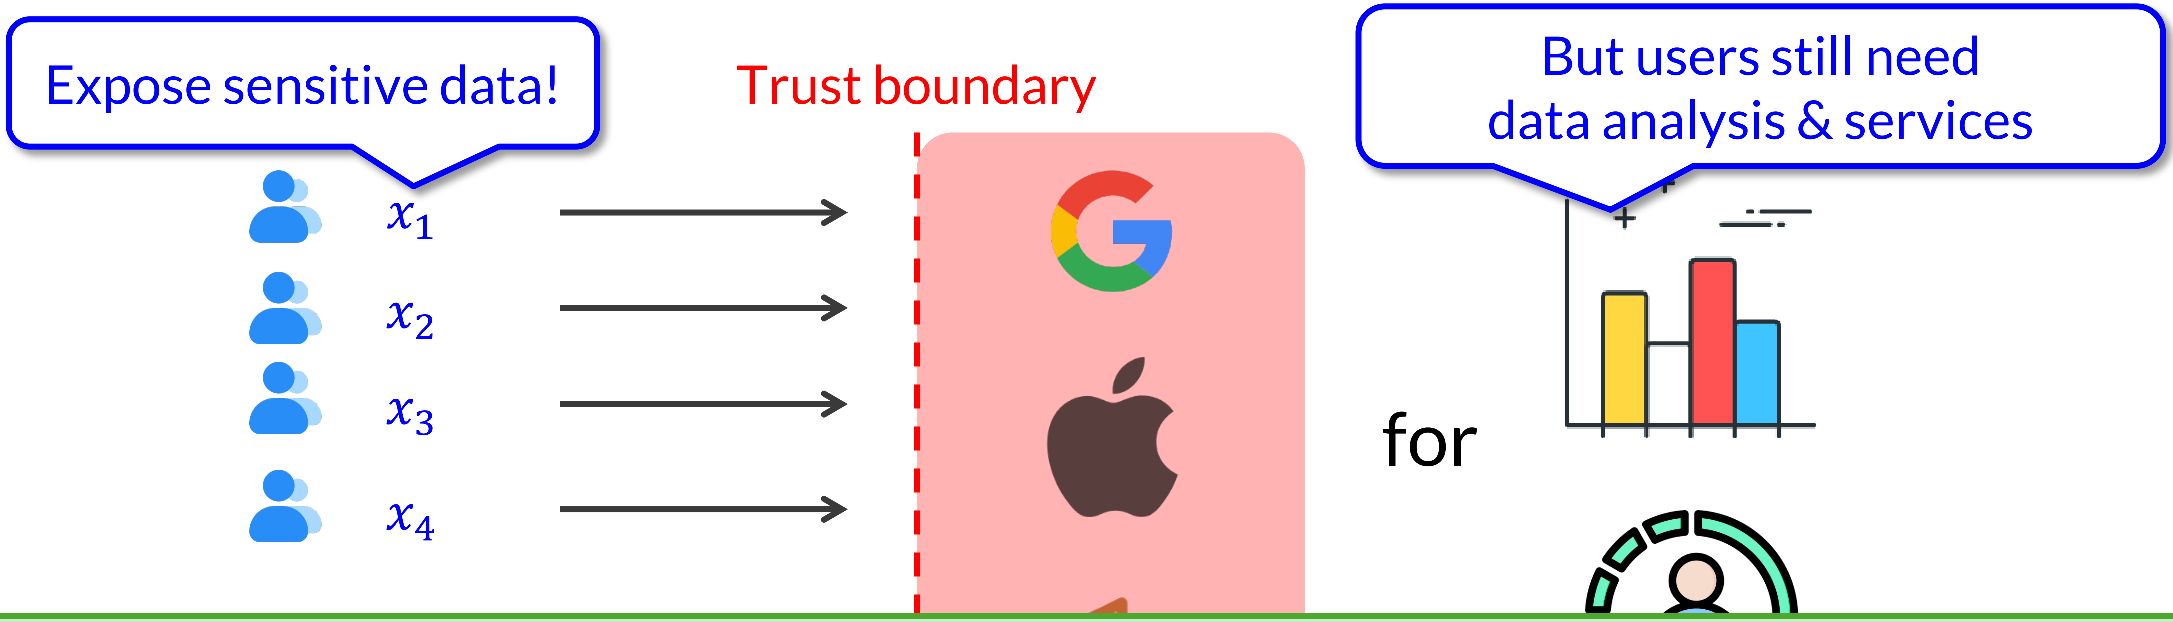

Expose sensitive data!

x_1

x_2

x_3

x_4

Trust boundary

But users still need data analysis & services

for

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

Privacy-Preserving Computation

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

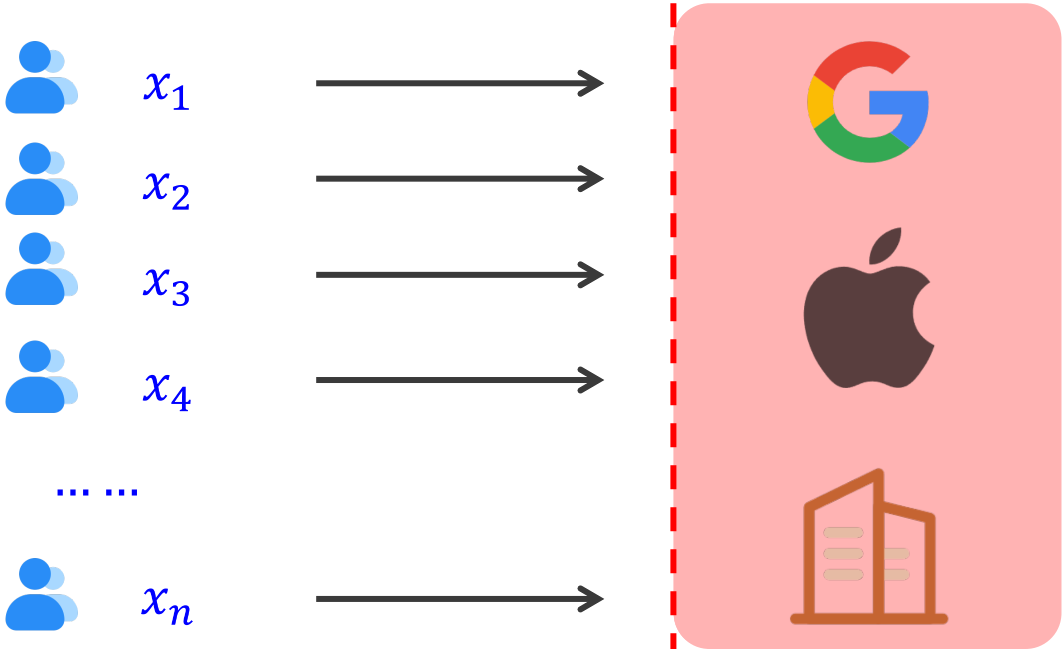

for $f(x_1, x_2, \dots, x_n)$

Privacy-Preserving Computation - **Techniques**

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

Privacy-Preserving Computation - **HE**

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

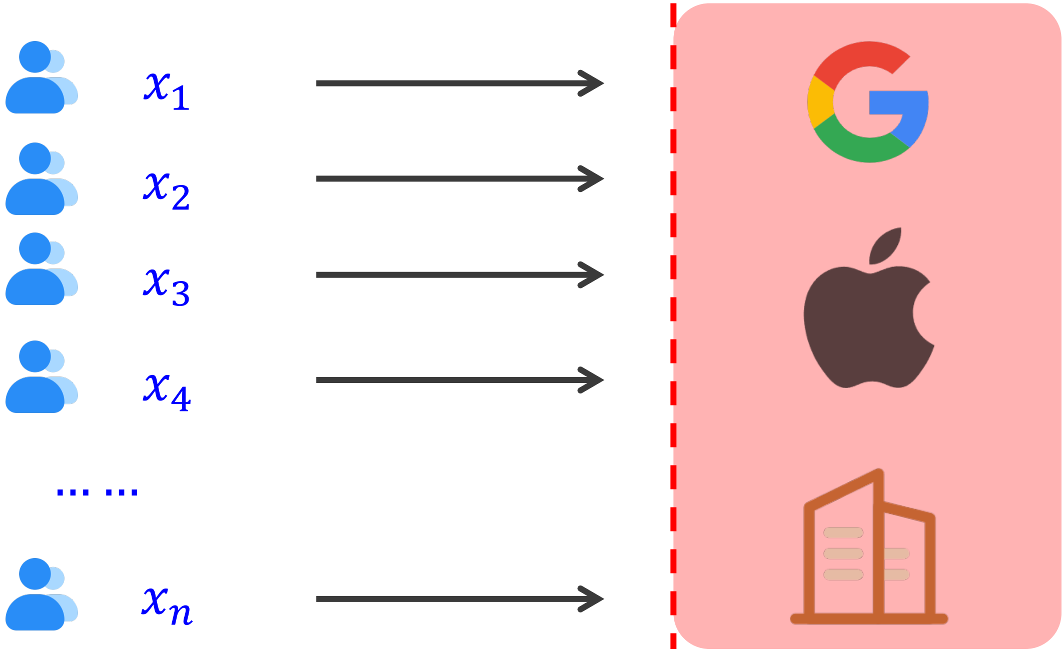

for $f(x_1, x_2, \ldots, x_n)$

Privacy-Preserving Computation - **HE**

- Homomorphic encryption (HE):
  - “homomorphic”: preserving structure
  - design algorithms $\{Enc, Dec\}$ → $Dec\big(f(Enc(x_1), Enc(x_2), \ldots, Enc(x_n))\big) = f(x_1, x_2, \ldots, x_n)$

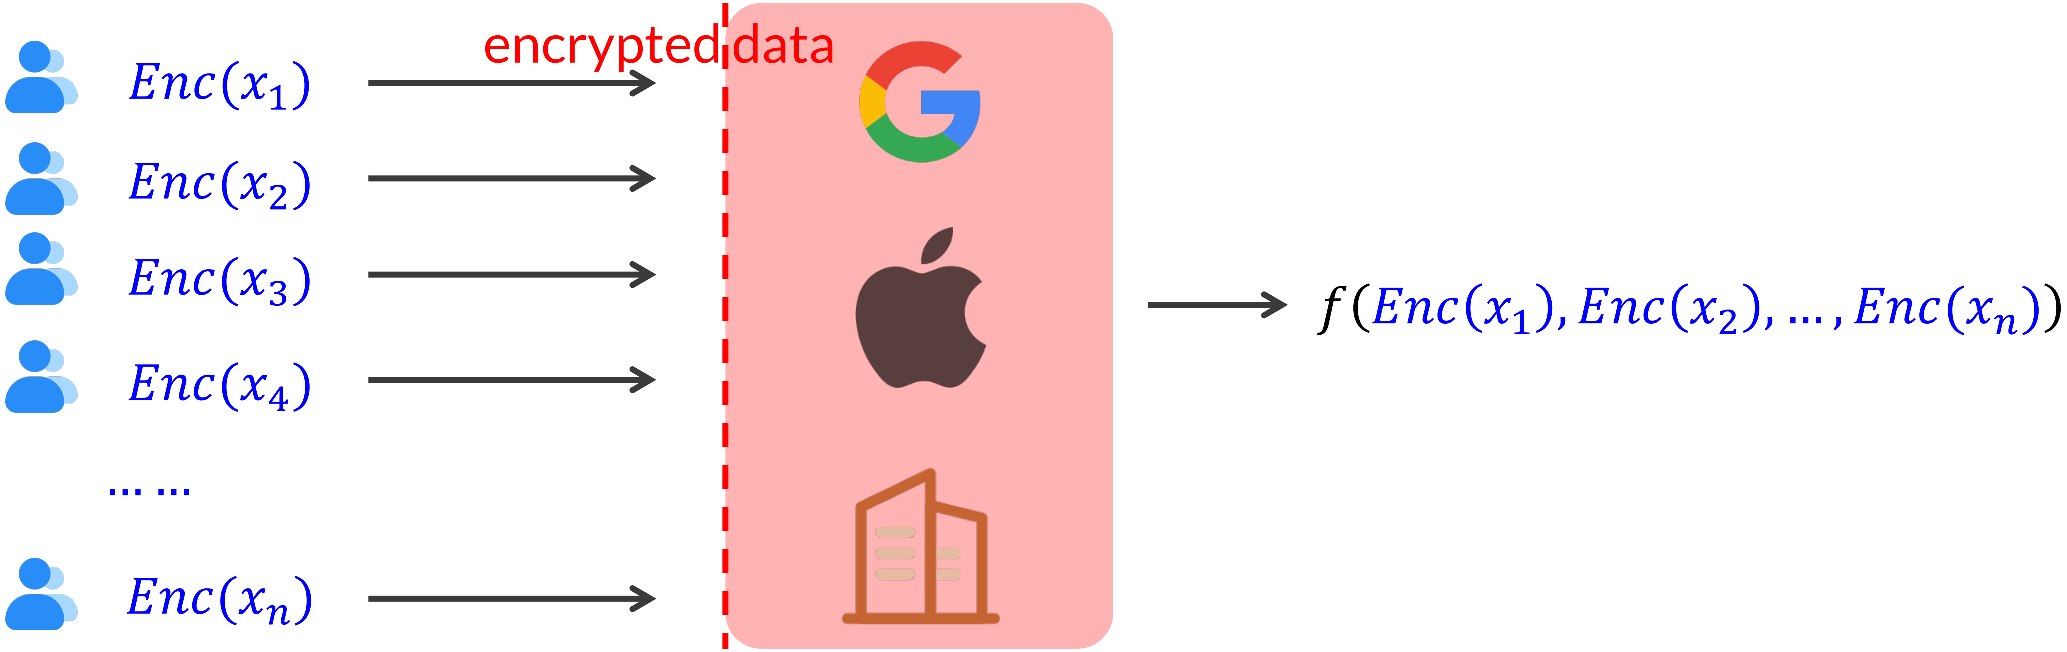

Privacy-Preserving Computation - **HE**

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

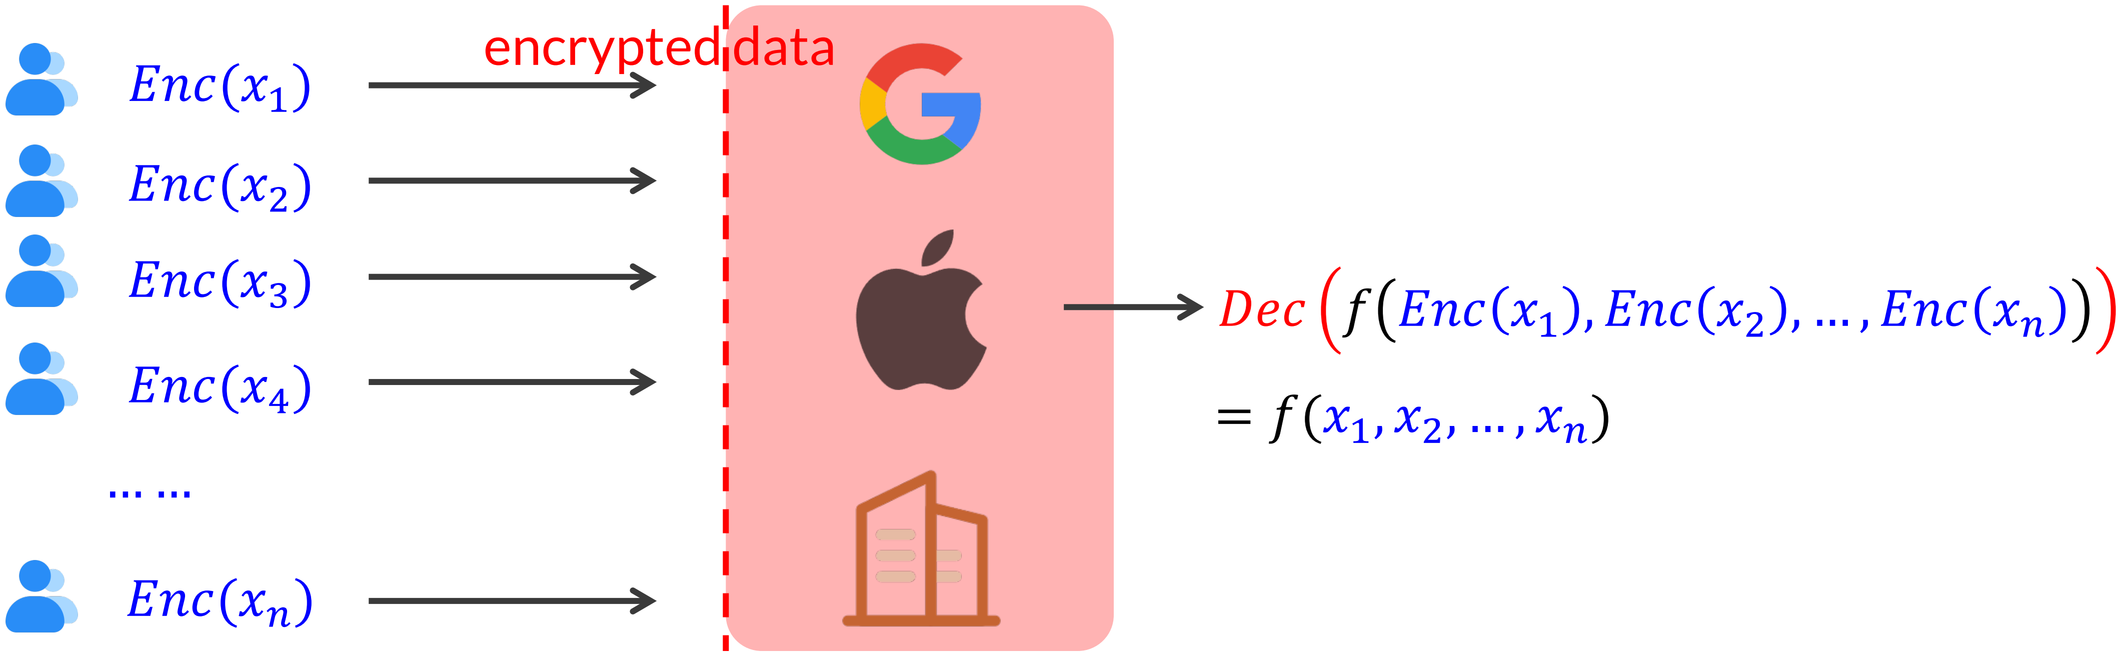

$Enc(x_1)$

$Enc(x_2)$

$Enc(x_3)$

$Enc(x_4)$

… …

$Enc(x_n)$

encrypted data

$Dec\left(f\left(Enc(x_1), Enc(x_2), \ldots, Enc(x_n)\right)\right)$

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

Privacy-Preserving Computation - **HE**

- Homomorphic encryption (HE):
  - “homomorphic”: preserving structure
  - design algorithms $\{Enc, Dec\}$ → $Dec\Big(f\big(Enc(x_1), Enc(x_2), \ldots, Enc(x_n)\big)\Big) = f(x_1, x_2, \ldots, x_n)$

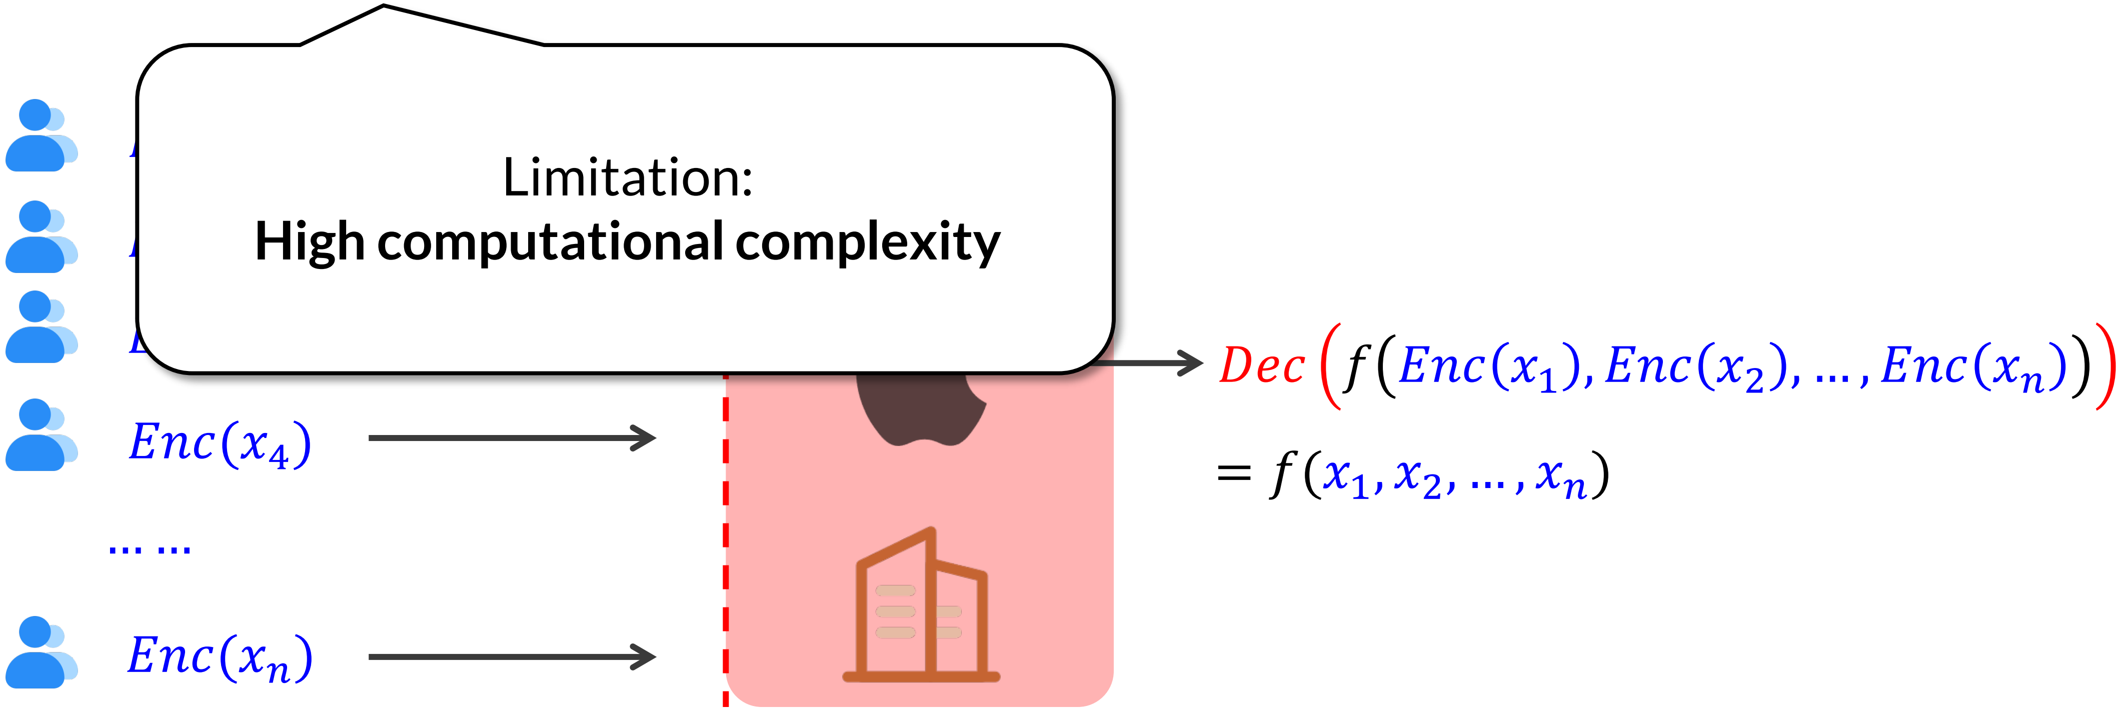

$$Dec\Big(f\big(Enc(x_1), Enc(x_2), \ldots, Enc(x_n)\big)\Big)$$
$$= f(x_1, x_2, \ldots, x_n)$$

Limitation:
High computational complexity

Privacy-Preserving Computation - **MPC**

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

👥 x_1

👥 x_2

… …

for $f(x_1, x_2)$

Privacy-Preserving Computation - **MPC**

RIT

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

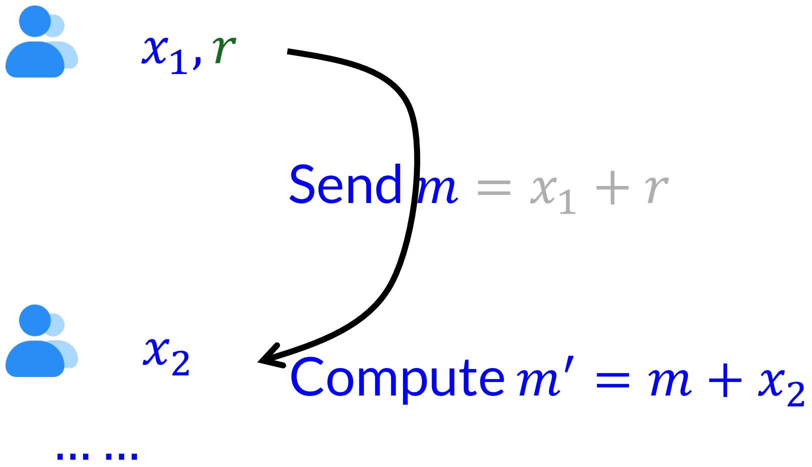

x_1, r

Send $m = x_1 + r$

x_2

Compute $m' = m + x_2$

… …

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: $f(x_1, x_2) = x_1 + x_2$

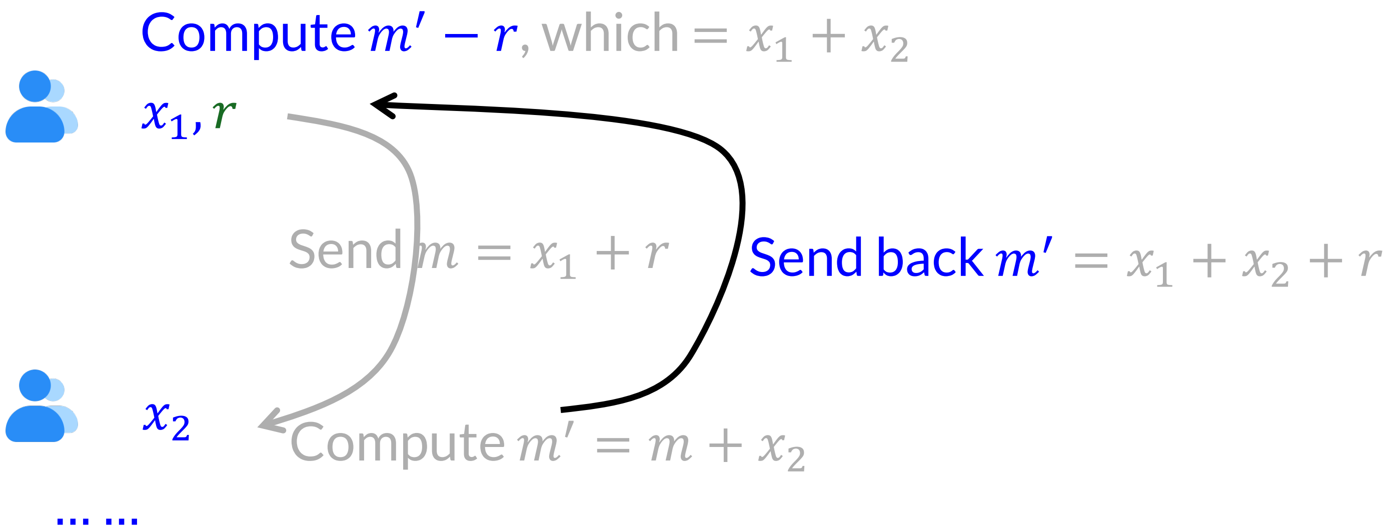

Compute $m' - r$, which $= x_1 + x_2$

x_1, r

Send $m = x_1 + r$

Send back $m' = x_1 + x_2 + r$

x_2

Compute $m' = m + x_2$

… …

Privacy-Preserving Computation - **MPC**

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

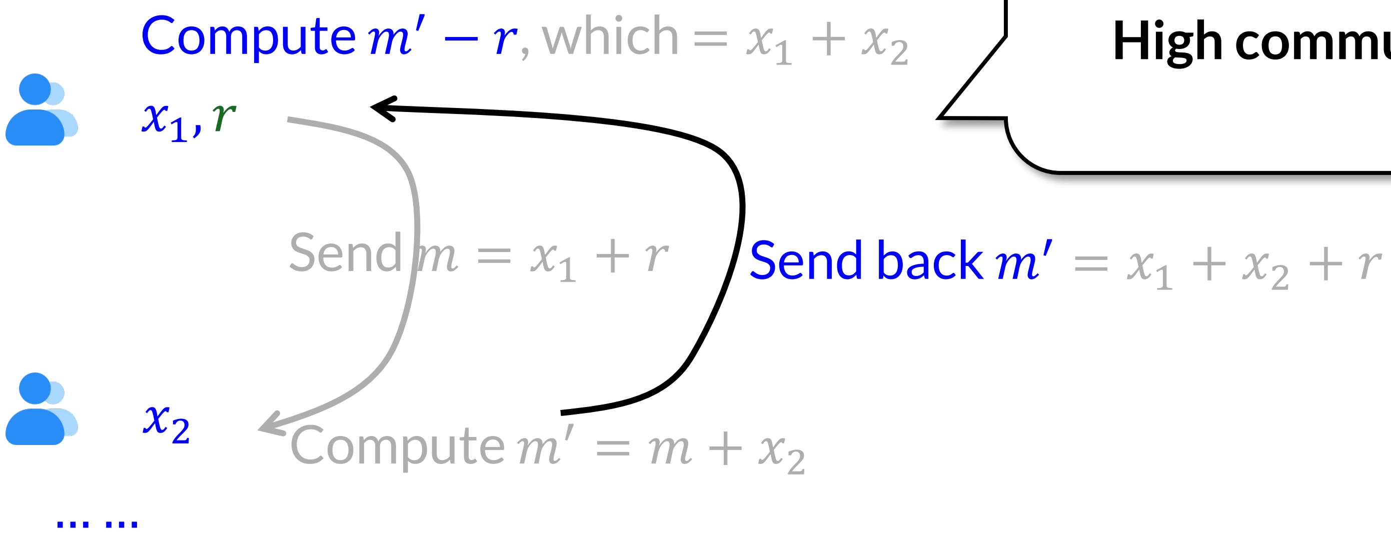

Compute $m' - r$, which $= x_1 + x_2$

x_1, r

Send $m = x_1 + r$

Send back $m' = x_1 + x_2 + r$

x_2

Compute $m' = m + x_2$

… …

Limitation:
High communication complexity

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 (LDP):
  - <span></span>hard to differentiate the sensitive data from other data
  - each user locally perturbs x_i to $\tilde{x}_i$ → $f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n) \approx f(x_1, x_2, \dots, x_n)$

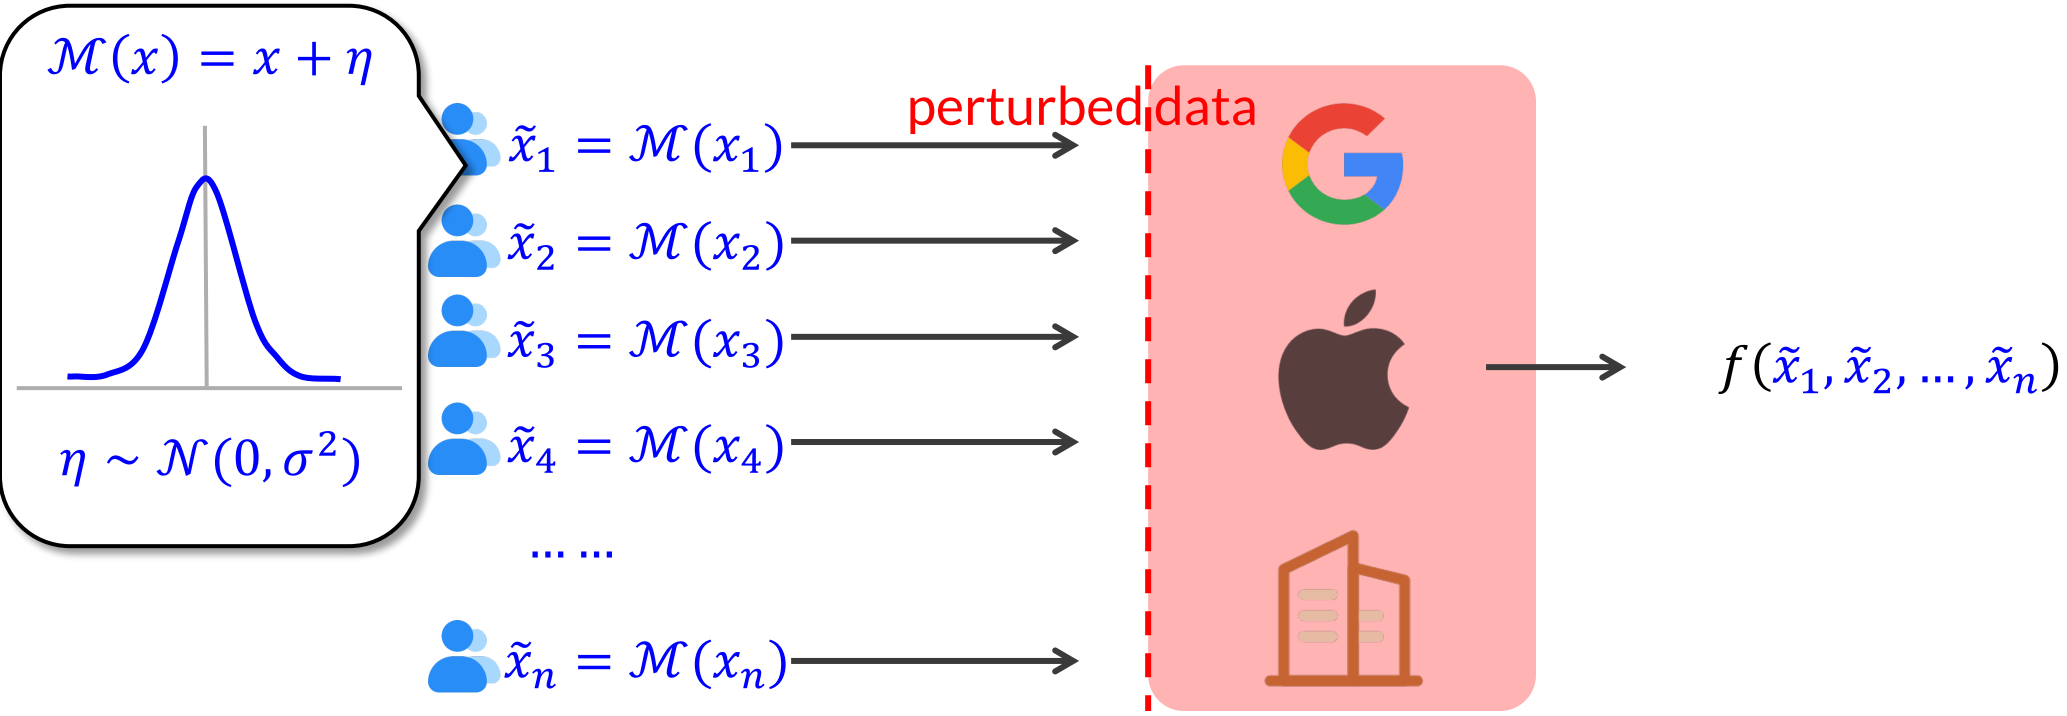

$\mathcal{M}(x) = x + \eta$

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

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

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

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

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

… …

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

perturbed data

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

Privacy-Preserving Computation - **LDP**

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

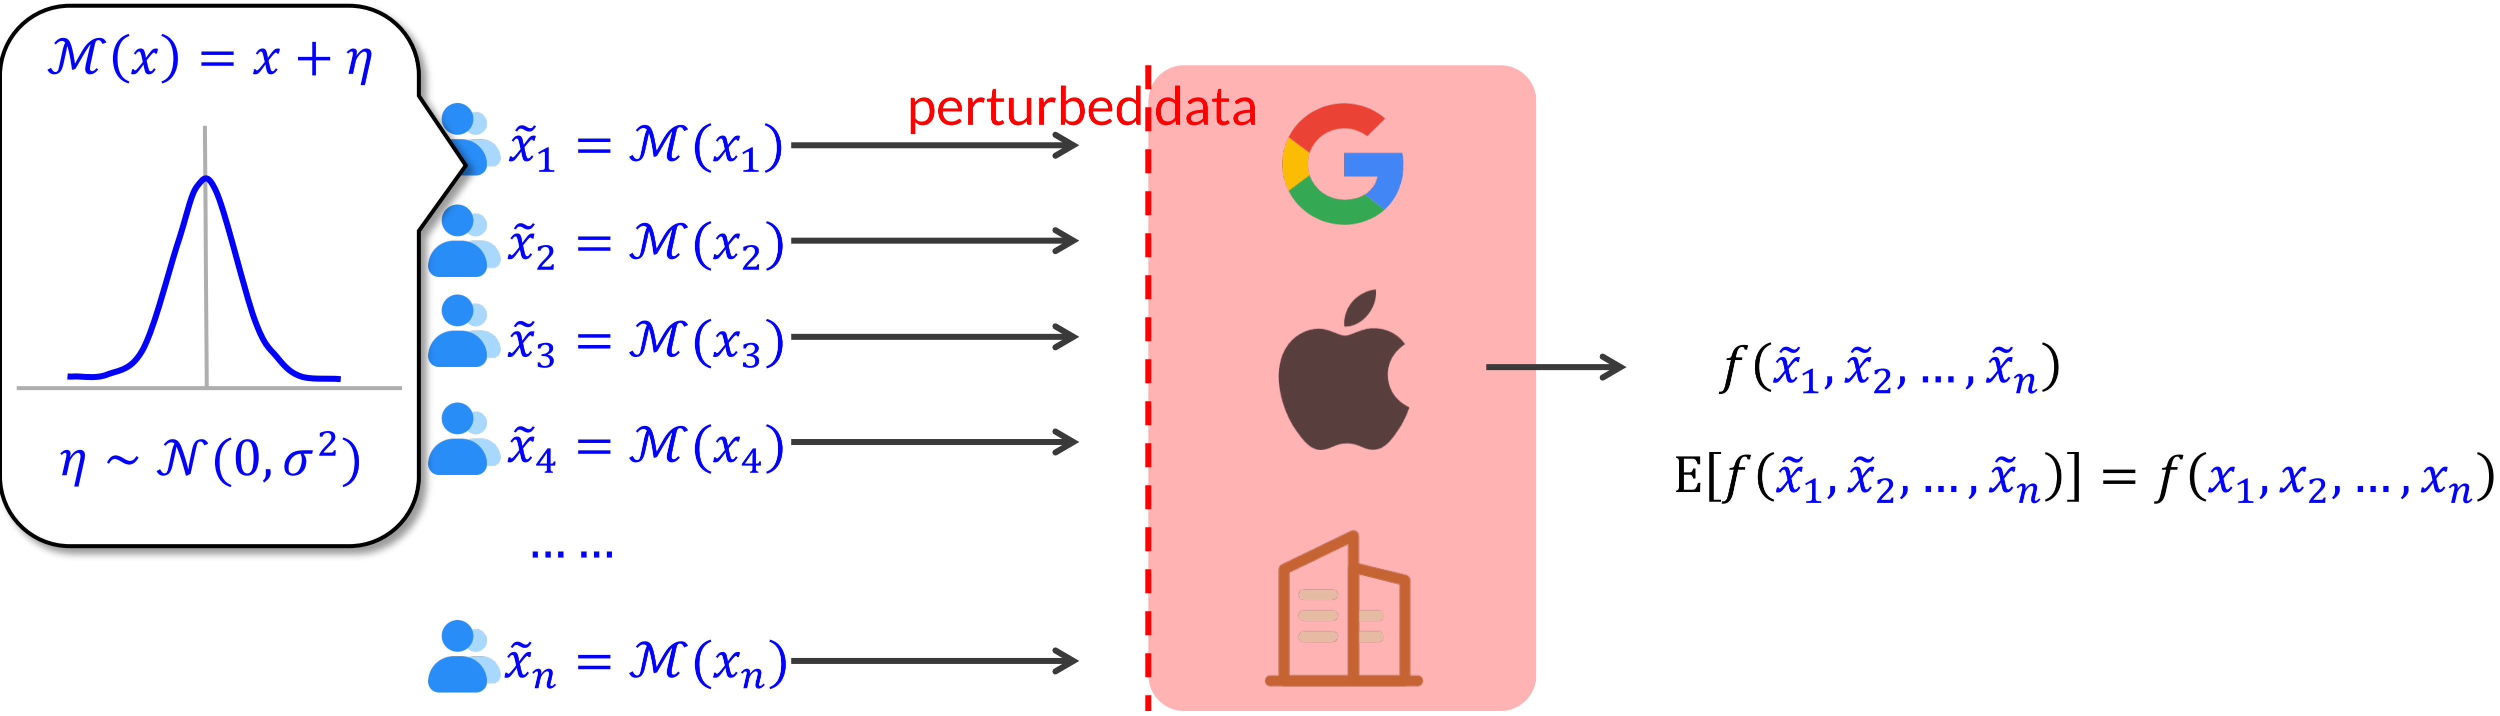

$\mathcal{M}(x) = x + \eta$

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

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

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

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

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

… …

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

perturbed data

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

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

Privacy-Preserving Computation - **LDP**

- Local differential privacy (L
  - hard to differentiate the s
  - each user locally perturb $x_n)$

Advantages:
Negligible computational complexity
No communication between users

But approximated f

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

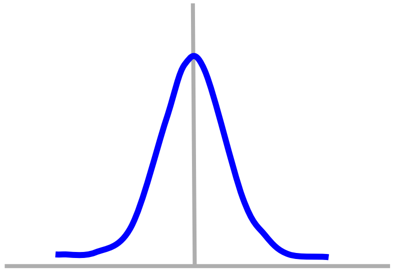

$\eta \sim \mathcal{N}(0, \sigma^2)$

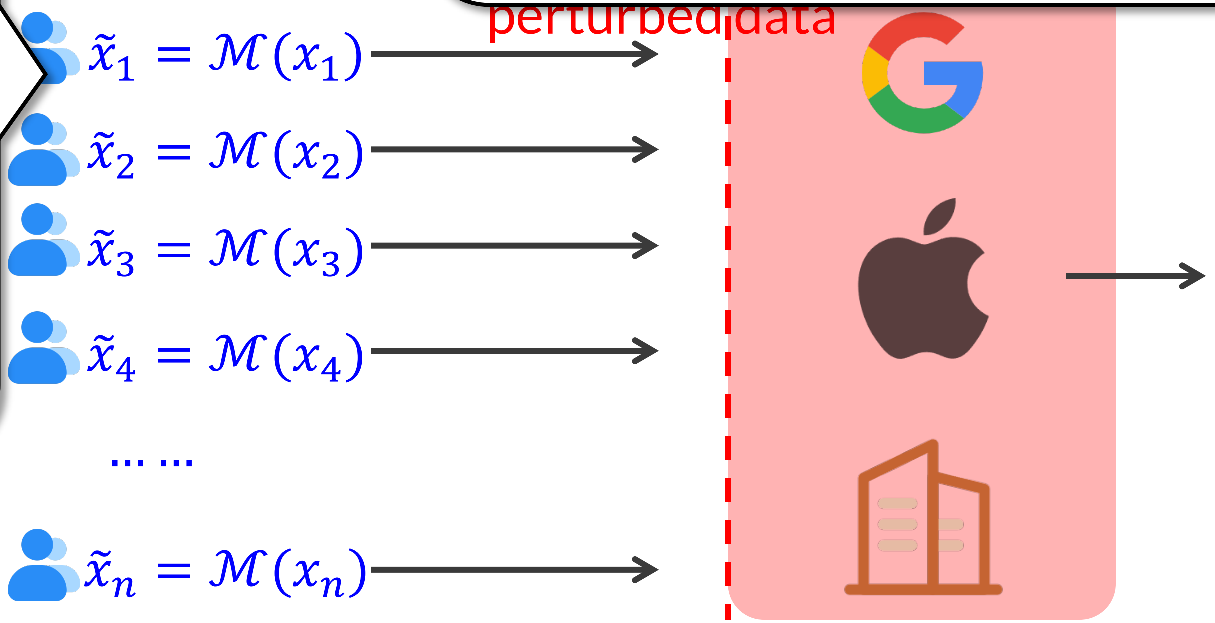

perturbed data

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

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

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

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

… …

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

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

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

LDP – Applications

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

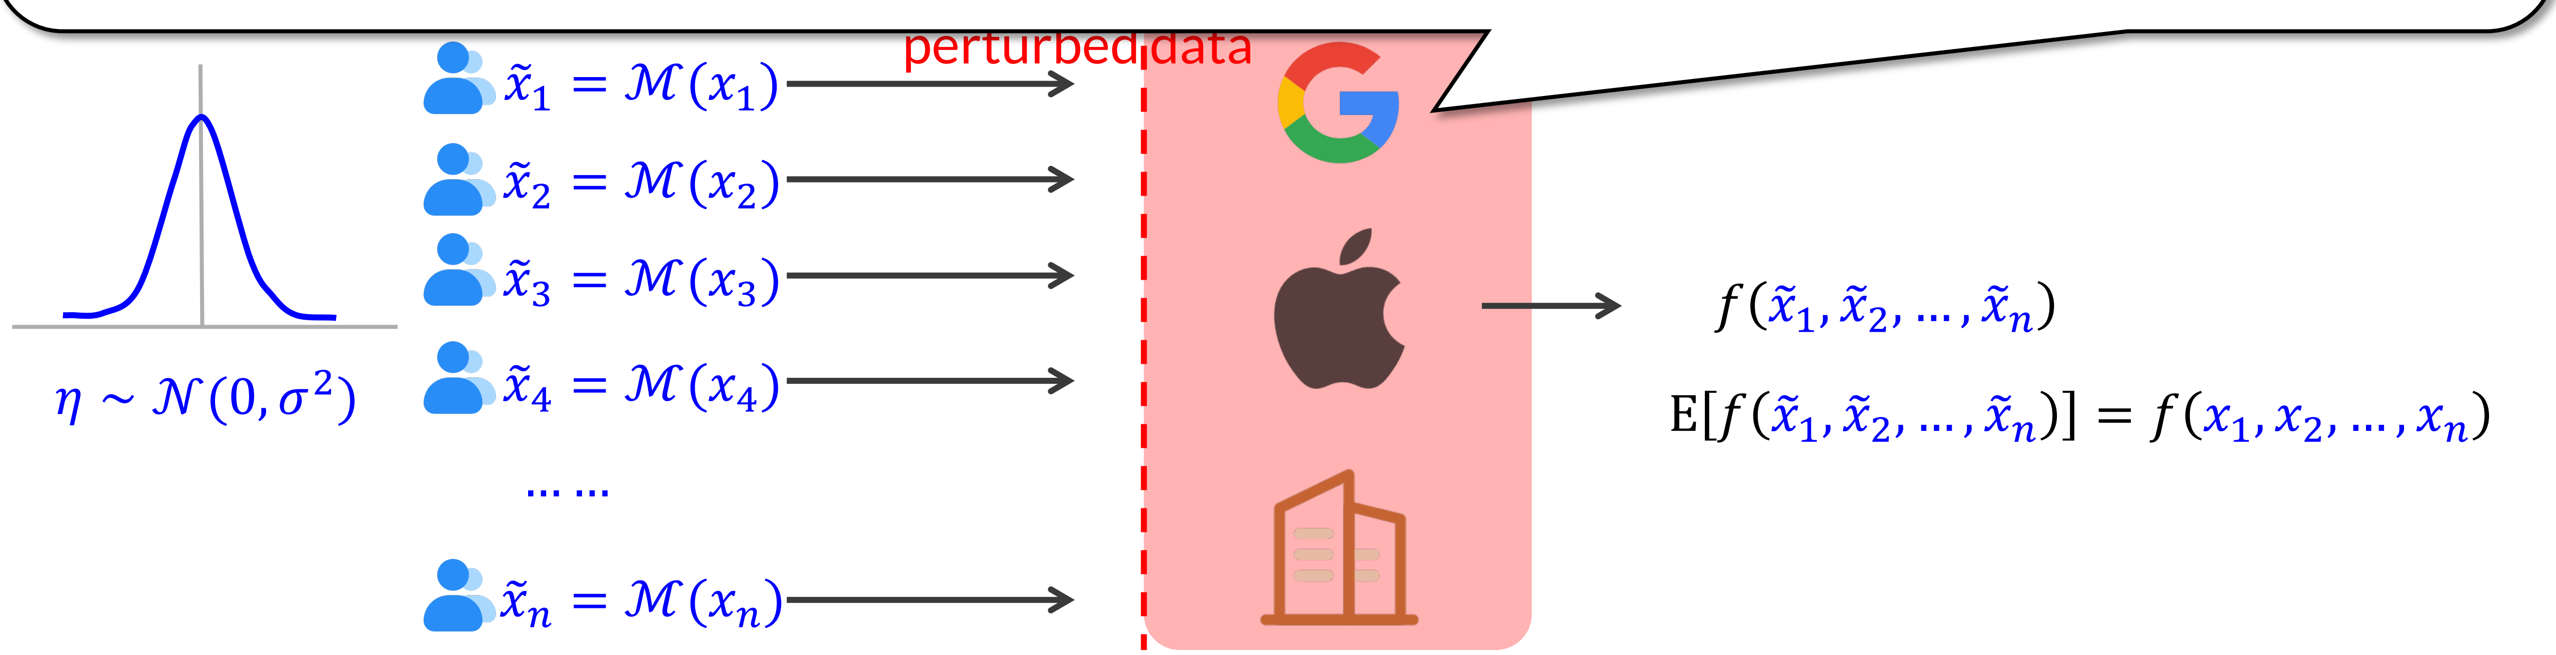

$\eta \sim \mathcal{N}(0, \sigma^2)$

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

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

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

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

… …

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

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

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

LDP – Applications

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

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

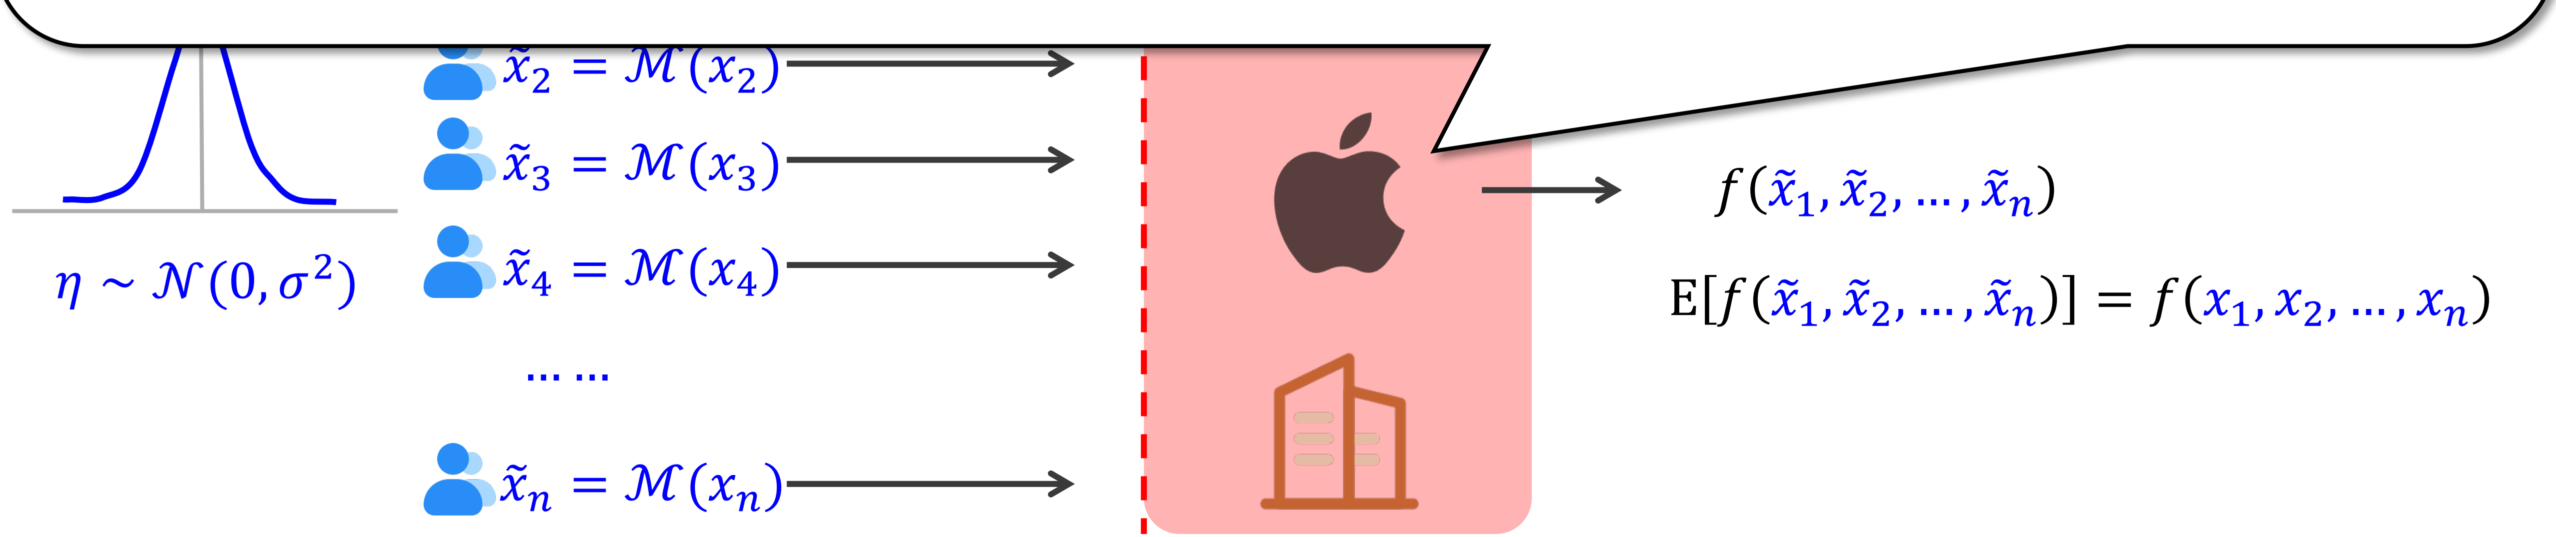
$\eta \sim \mathcal{N}(0, \sigma^2)$
$\tilde{x}_2 = \mathcal{M}(x_2)$
$\tilde{x}_3 = \mathcal{M}(x_3)$
$\tilde{x}_4 = \mathcal{M}(x_4)$
… …
$\tilde{x}_n = \mathcal{M}(x_n)$
$f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$
$\mathrm{E}[f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)] = f(x_1, x_2, \dots, x_n)$

LDP – Formal Privacy

RIT

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

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

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

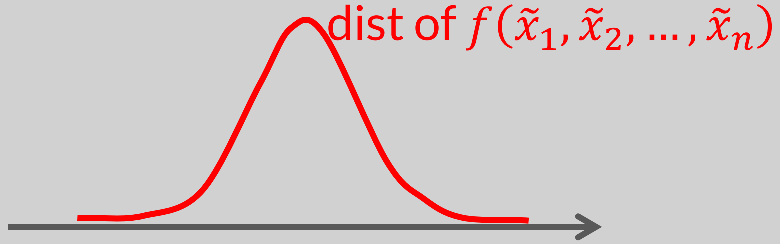

$$f(\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_n) \approx f(x_1, x_2, \ldots, x_n)$$

LDP – Formal Privacy

- 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

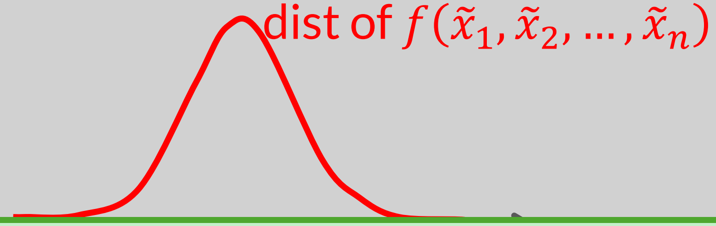

Fundamental direction: **Design** of $\mathcal{M}$ to **optimize** the privacy—utility tradeoff

LDP – Formal Privacy

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

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

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

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

RIT

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

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

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

Privacy quantified by ε

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

Data utility by app error

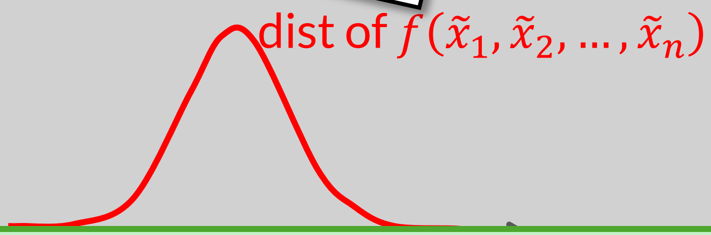

Fundamental direction: **Utility analysis** of complex task f

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

This Proposal: LDP Theory

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy

  - *correlated LDP mechanisms

Existing LDP mechanisms:
Each user perturbs their data independently

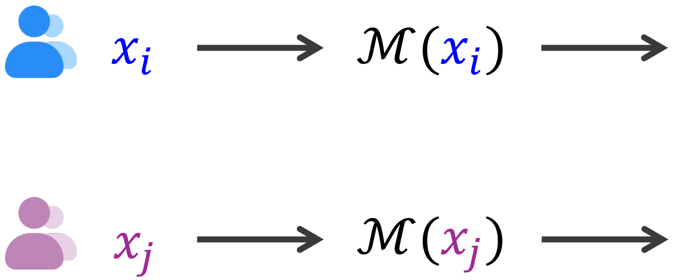

$$x_i \longrightarrow \mathcal{M}(x_i) \longrightarrow$$

$$x_j \longrightarrow \mathcal{M}(x_j) \longrightarrow$$

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

This Proposal: LDP Theory

RIT

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy

  - *correlated LDP mechanisms

Existing LDP mechanisms:
Each user perturbs their data independently

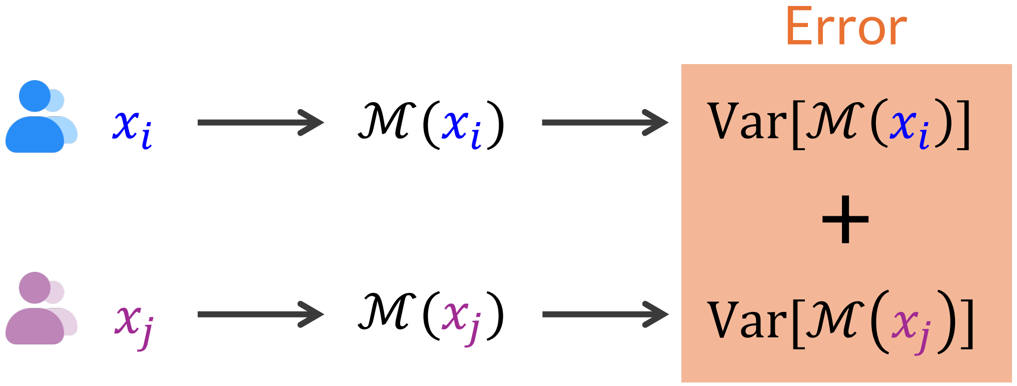
Error
$x_i \longrightarrow \mathcal{M}(x_i) \longrightarrow \mathrm{Var}[\mathcal{M}(x_i)]$
+
$x_j \longrightarrow \mathcal{M}(x_j) \longrightarrow \mathrm{Var}[\mathcal{M}(x_j)]$

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis
- New LDP mechanisms: (for better privacy

  - *correlated LDP mechanisms

Correlated LDP mechanisms:
Users’ data are perturbed by correlated $\mathcal{M}$

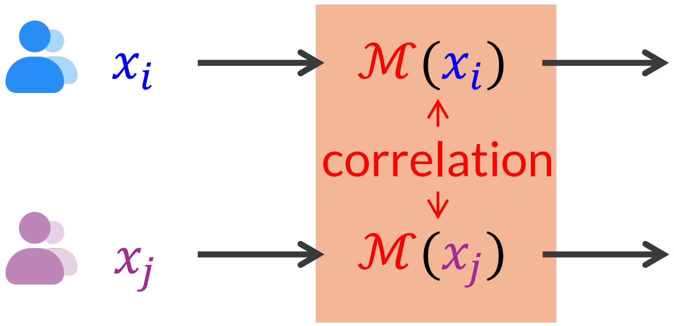

* [PETS’25] Locally Differentially Private Frequency Estimation via Joint Randomized Response

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy

  - *correlated LDP mechanisms

Correlated LDP mechanisms:
Users’ data are perturbed by correlated $\mathcal{M}$

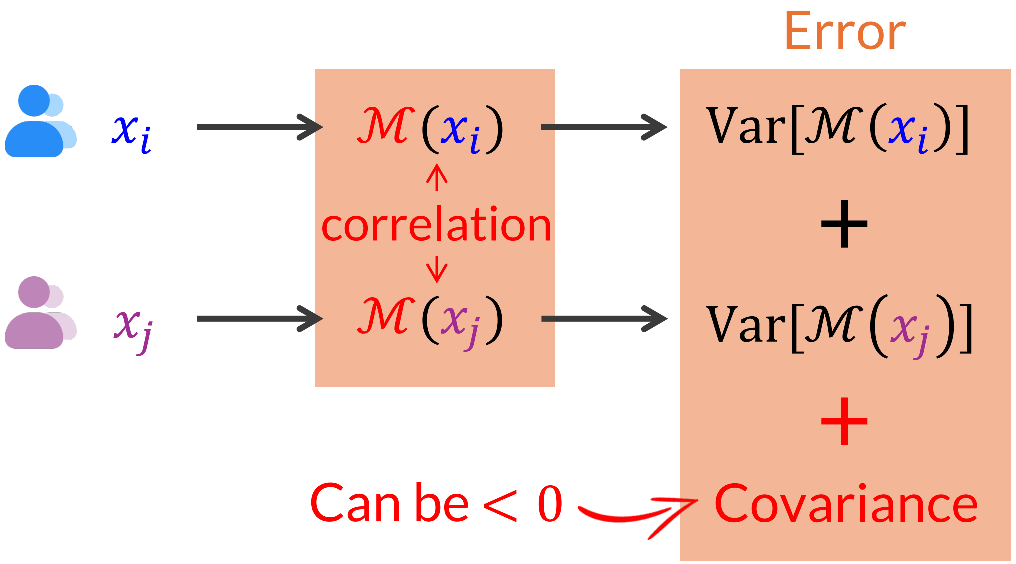

* [PETS’25] Locally Differentially Private Frequency Estimation via Joint Randomized Response

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy

  - correlated LDP mechanisms

  - [†]optimal piecewise-based mechanis

From binary
to numerical

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

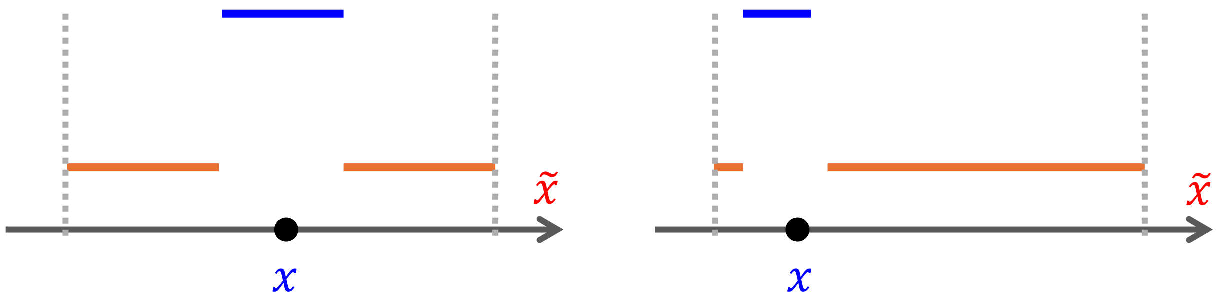

[†] [PETS'25] Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy

This Proposal: LDP Theory

RIT

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy

  - correlated LDP mechanisms

  - [†]optimal piecewise-based mechanis

From binary
to numerical

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

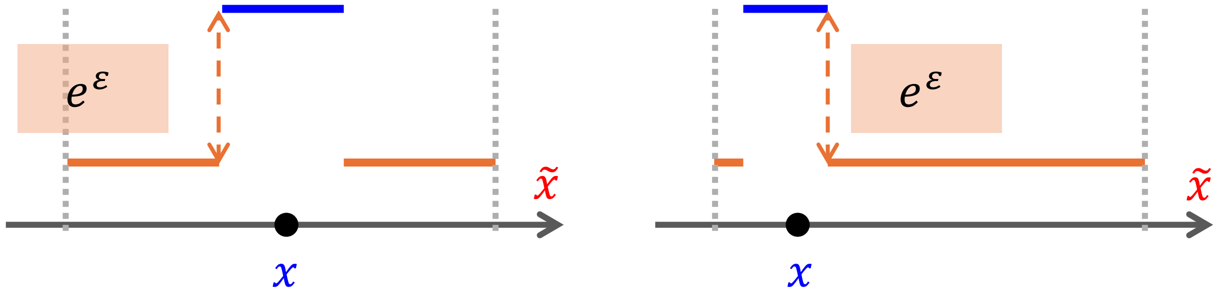

$$\mathrm{pdf}[\mathcal{M}(x) = \tilde{x}] = \begin{cases} p_\varepsilon & \text{if } \tilde{x} \in [l_{x,\varepsilon}, r_{x,\varepsilon}] \\ \dfrac{p_\varepsilon}{e^\varepsilon} & \text{if } \tilde{x} \in \tilde{\mathcal{D}} \setminus [l_{x,\varepsilon}, r_{x,\varepsilon}] \end{cases}$$

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

This Proposal: LDP Theory

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy

  - correlated LDP mechanisms

  - [†]optimal piecewise-based mechanis

From binary
to numerical

Too heuristic → More generalized version
Potentially has lower error

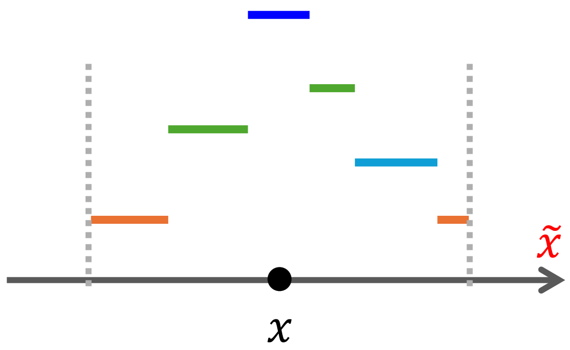

What is the optimal piecewise-based mechanism?

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

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy
  - correlated LDP mechanisms
  - [†]optimal piecewise-based mechanis

From binary to numerical

Linear data domain → Circular data domain

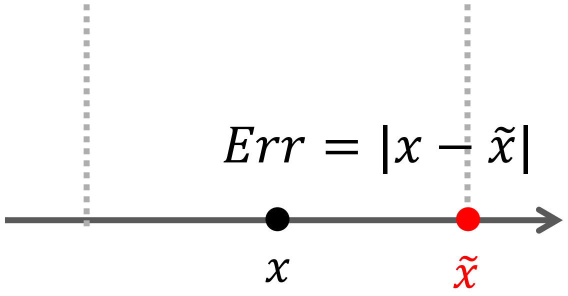

What is the optimal piecewise-based mechanism?

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

This Proposal: LDP Theory

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy
  - correlated LDP mechanisms
  - †optimal piecewise-based mechanis

From binary to numerical

Linear data domain → Circular data domain

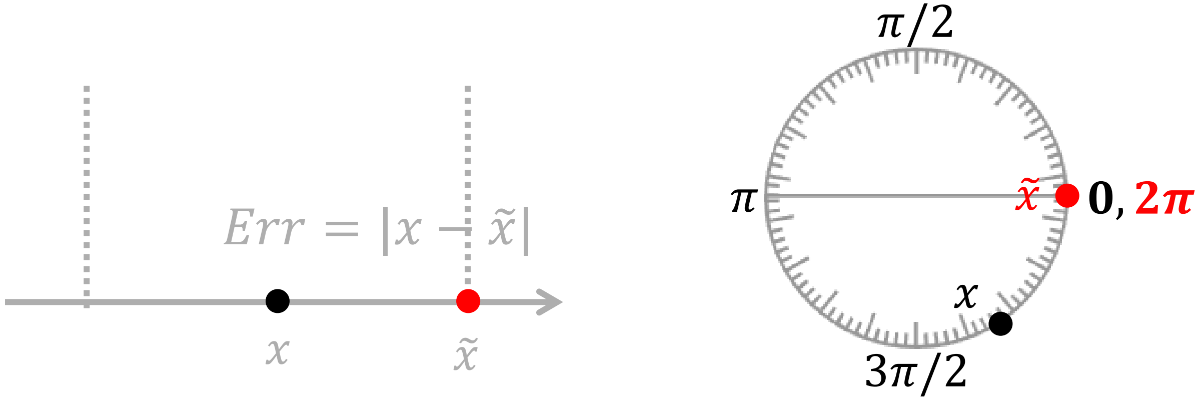

$$Err = \min(|x - \tilde{x}|, |2\pi - x - \tilde{x}|)$$

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

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy

  - correlated LDP mechanisms

  - [†]optimal piecewise-based mechanis

From binary to numerical

Linear data domain → Circular data domain

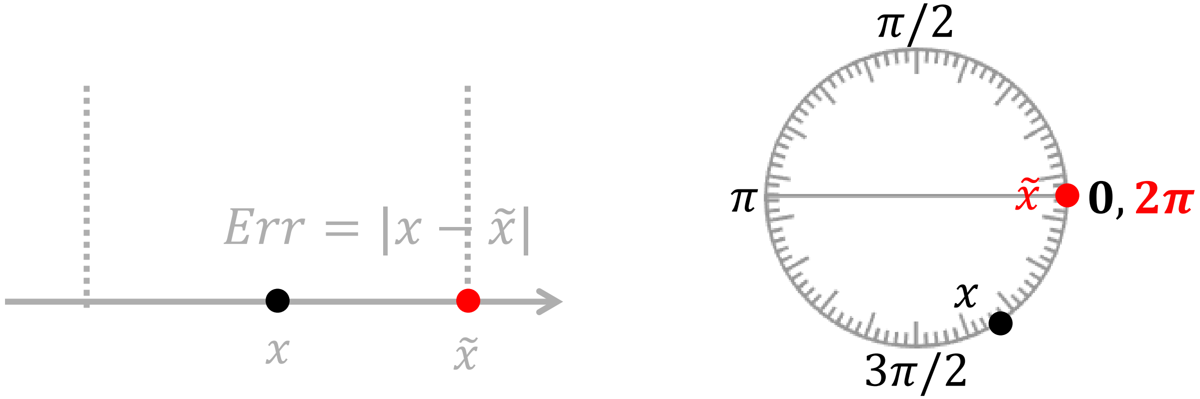

$$Err = \min(|x - \tilde{x}|, |2\pi - x - \tilde{x}|)$$

Optimal piecewise-based mechanism for circular domain?

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

This Proposal: LDP Theory

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—u

  - correlated LDP mechanisms

  - optimal piecewise-based mechanism

  - [‡]trajectory collection in continuous space

From 1D to 2D

Discrete location space

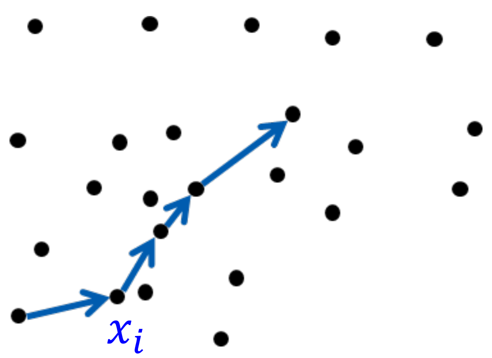

For $\tilde{x} \in S$

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

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

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

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—

  - correlated LDP mechanisms

  - optimal piecewise-based mechanism

  - [‡]trajectory collection in continuous space

From 1D to 2D

Discrete location space

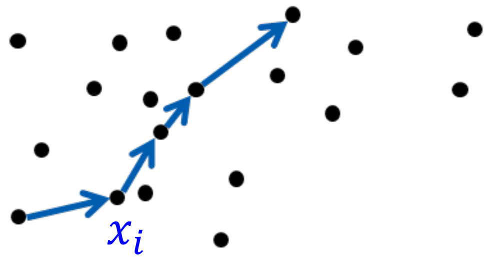

$$\mathcal{S} = \{p_1, \dots, p_n\}$$

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

This Proposal: LDP Theory

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—

  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - [‡]trajectory collection in continuous space

From 1D to 2D

Discrete location space → Continuous location space

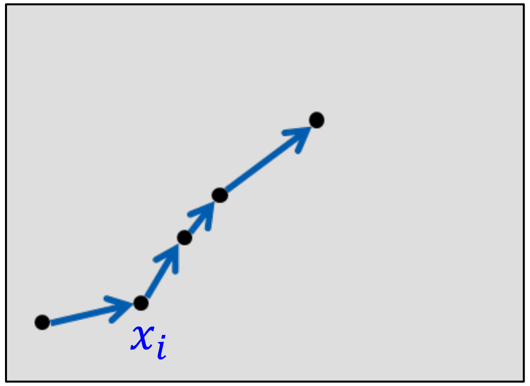

$$\mathcal{S} = [0,1.5] \times [0,1]$$

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - [‡]trajectory collection in continuous space

From 1D to 2D

Discrete location space → Continuous location space

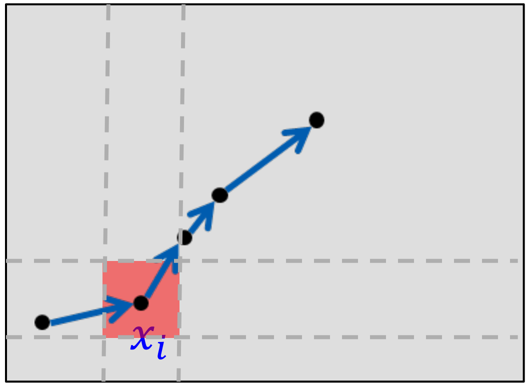

$\mathcal{S} = [0{,}1.5] \times [0{,}1]$

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

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

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—u
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - [‡]trajectory collection in continuous space

From 1D to 2D

Discrete location space → Continuous location space

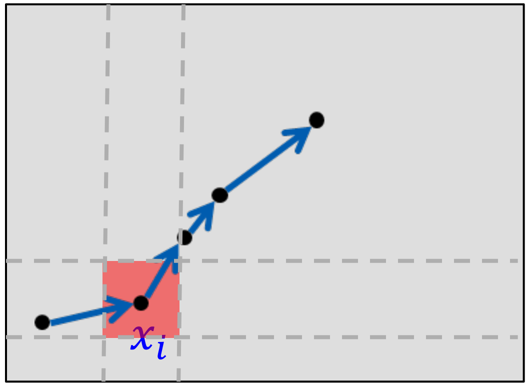

$$\mathcal{S} = [0,1.5] \times [0,1]$$

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

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

LDP for continuous S

Benefits:

- negligible sample complexity
- embedded discrete spaces

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

This Proposal: LDP Theory

- Advancing LDP's mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—u

  - correlated LDP mechanisms

  - optimal piecewise-based mechanism

  - [‡]trajectory collection in continuous space

From 1D to 2D

Rounding to the nearest

Does not affect privacy

Discrete location space → Continuous location space

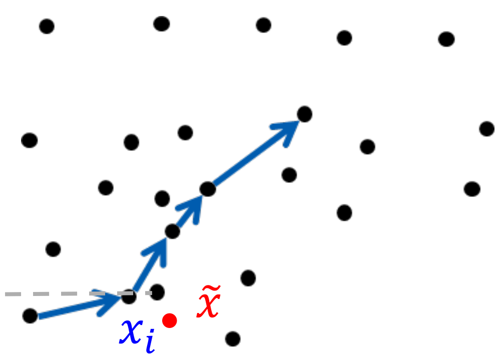

$$\mathcal{S} = [0,1.5] \times [0,1]$$

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

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

LDP for continuous S

Benefits:

- negligible sample complexity
- embedded discrete spaces

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—u
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - [‡]trajectory collection in continuous space

From 1D to 2D

Rounding to the nearest

Does not affect privacy

Discrete location space → Continuous location space

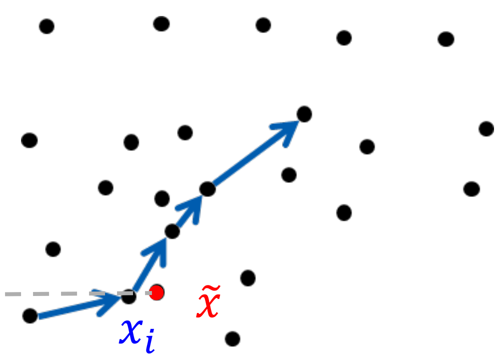

$\mathcal{S} = [0,1.5] \times [0,1]$

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

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

LDP for continuous S

Benefits:

- negligible sample complexity
- embedded discrete spaces

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

This Proposal: LDP Theory

RIT

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—u
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:
  - ¶classifier utility under LDP-inputs

Empirical classifier utility

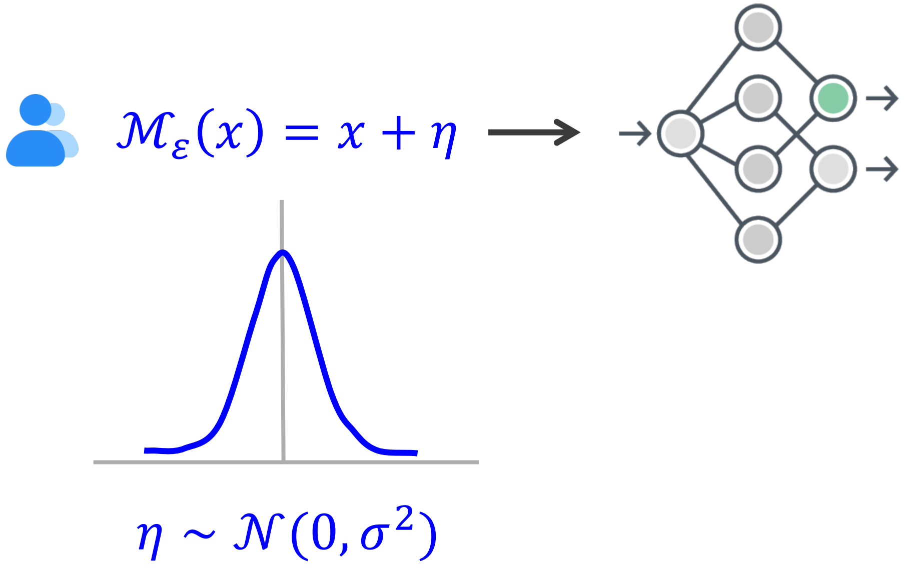

$\mathcal{M}_{\varepsilon}(x) = x + \eta$

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

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis
- New LDP mechanisms: (for better privacy—
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:
  - [¶]classifier utility under LDP-inputs

Empirical classifier utility

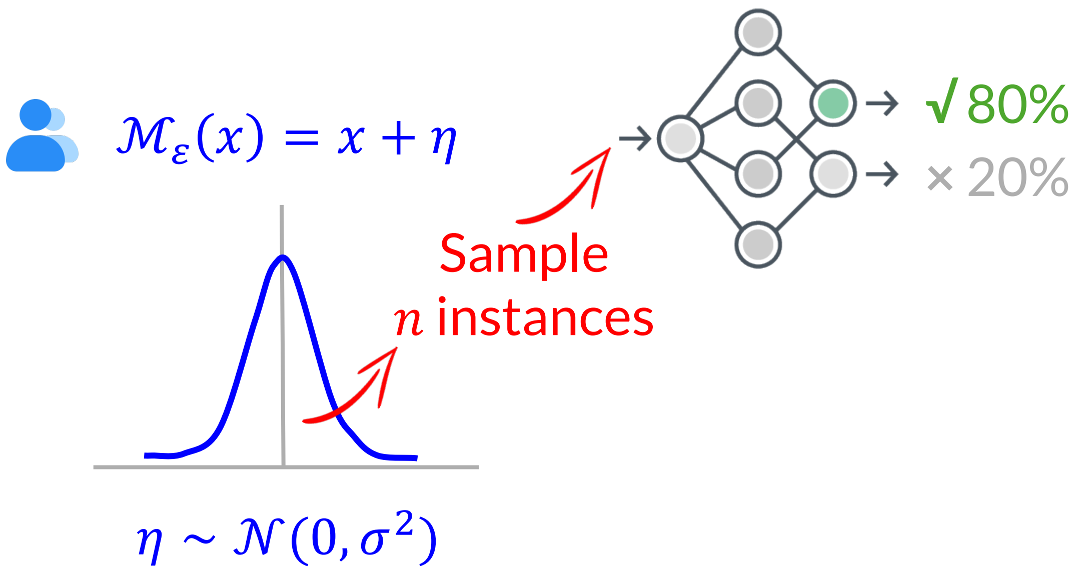
$\mathcal{M}_\varepsilon(x) = x + \eta$
Sample
n instances
√ 80%
× 20%
$\eta \sim \mathcal{N}(0, \sigma^2)$

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

This Proposal: LDP Theory

RIT

- Advancing LDP’s mechanism design and utility analysis
- New LDP mechanisms: (for better privacy—u
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:
  - [¶]classifier utility under LDP-inputs

Empirical classifier utility

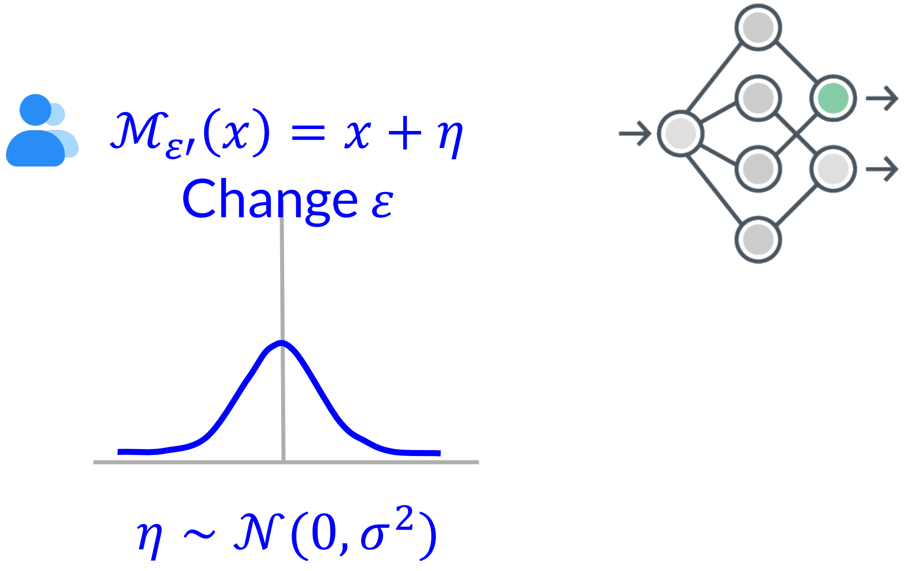

$\mathcal{M}_{\varepsilon'}(x) = x + \eta$

Change ε

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

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—

  - correlated LDP mechanisms

  - optimal piecewise-based mechanism

  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:

  - [¶]classifier utility under LDP-inputs

Empirical classifier utility

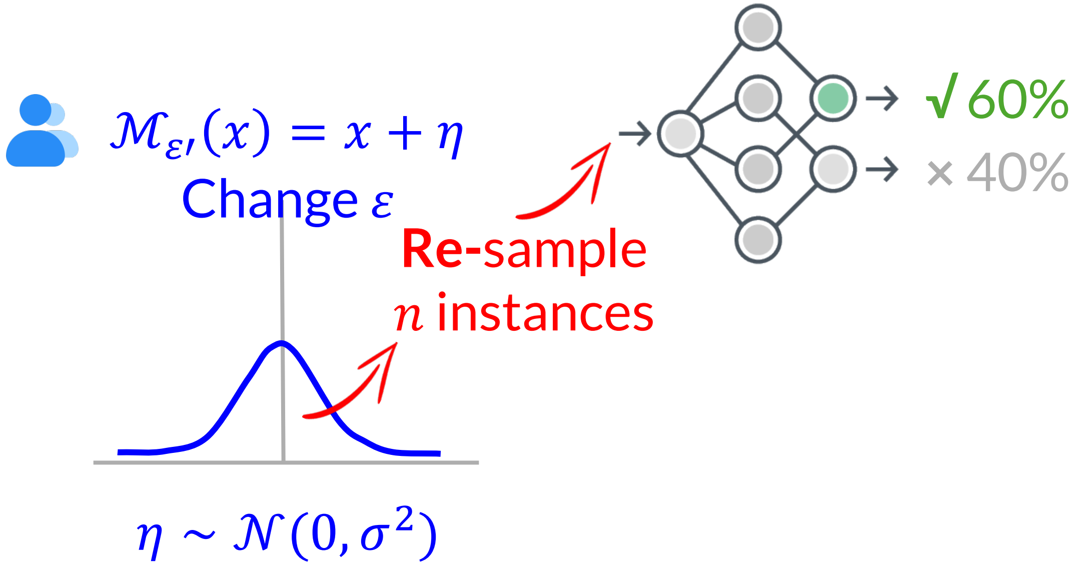

$\mathcal{M}_{\varepsilon'}(x) = x + \eta$
Change ε
Re-sample
n instances
√ 60%
× 40%
$\eta \sim \mathcal{N}(0, \sigma^2)$

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:
  - ¶classifier utility under LDP-inputs

Empirical classifier utility → Analytical classifier utility

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

This Proposal: LDP Theory

RIT

- Advancing LDP’s mechanism design and utility analysis
- New LDP mechanisms: (for better privacy—
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:
  - [¶]classifier utility under LDP-inputs

Empirical classifier utility → Analytical classifier utility

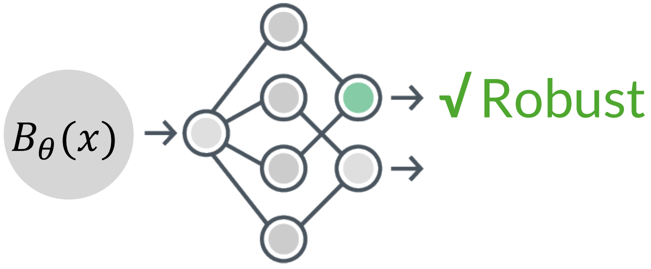

Robustness:
$B_\theta(x)$ is robust

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis

- New LDP mechanisms: (for better privacy—u
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:
  - [¶]classifier utility under LDP-inputs

Empirical classifier utility → Analytical classifier utility

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

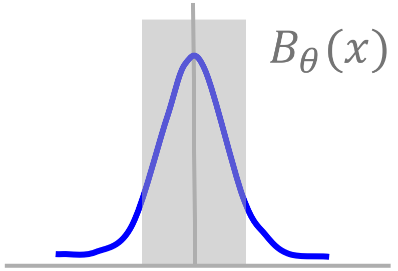

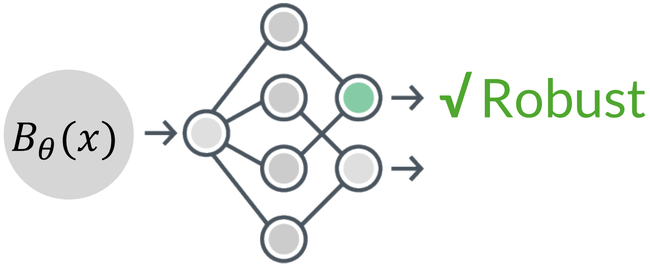

Concentration:
$\Pr[\mathcal{M}_\varepsilon(x) \in B_\theta(x)]$
$:= p(\varepsilon, \theta)$

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

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

This Proposal: LDP Theory

- Advancing LDP’s mechanism design and utility analysis
- New LDP mechanisms: (for better privacy—u
  - correlated LDP mechanisms
  - optimal piecewise-based mechanism
  - trajectory collection in continuous space

mechanism-level
to task level

- New utility quantification:
  - [¶]classifier utility under LDP-inputs

Empirical classifier utility → Analytical classifier utility

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

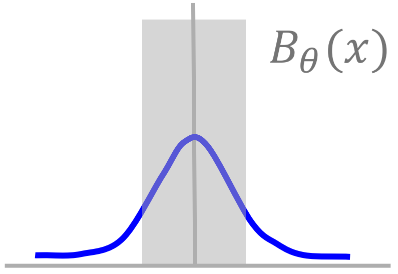

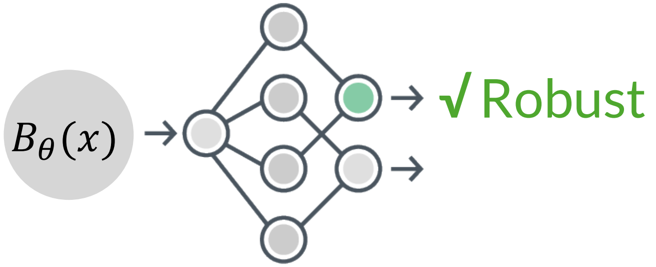

Concentration:
$\Pr[\mathcal{M}_\varepsilon(x) \in B_\theta(x)]$
$\coloneqq p(\varepsilon, \theta)$

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

Connected by θ

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

Social Contributions

RIT

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

Sensor networks &
Federated learning

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

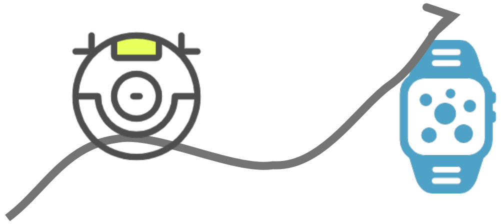

Smart home & wearable
devices' trajectories

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