

Local Differential Privacy: Refined Mechanism Design and Utility Analysis

Ye Zheng

Advisor: Dr. Yidan Hu

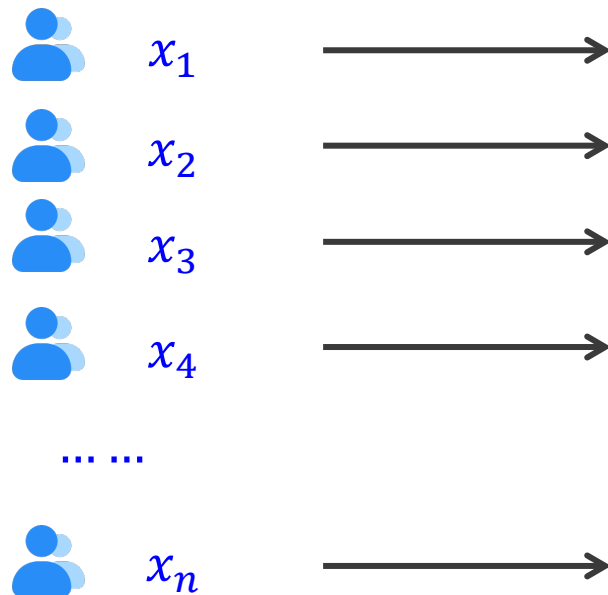
Committee: Dr. Sumita Mishra, Dr. Haibo Yang, Dr. Weijie Zhao

RIT | Rochester Institute of Technology

PDF & slides  <https://zhengyeah.com>

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

Location, browsing history,
app usage data




Collector

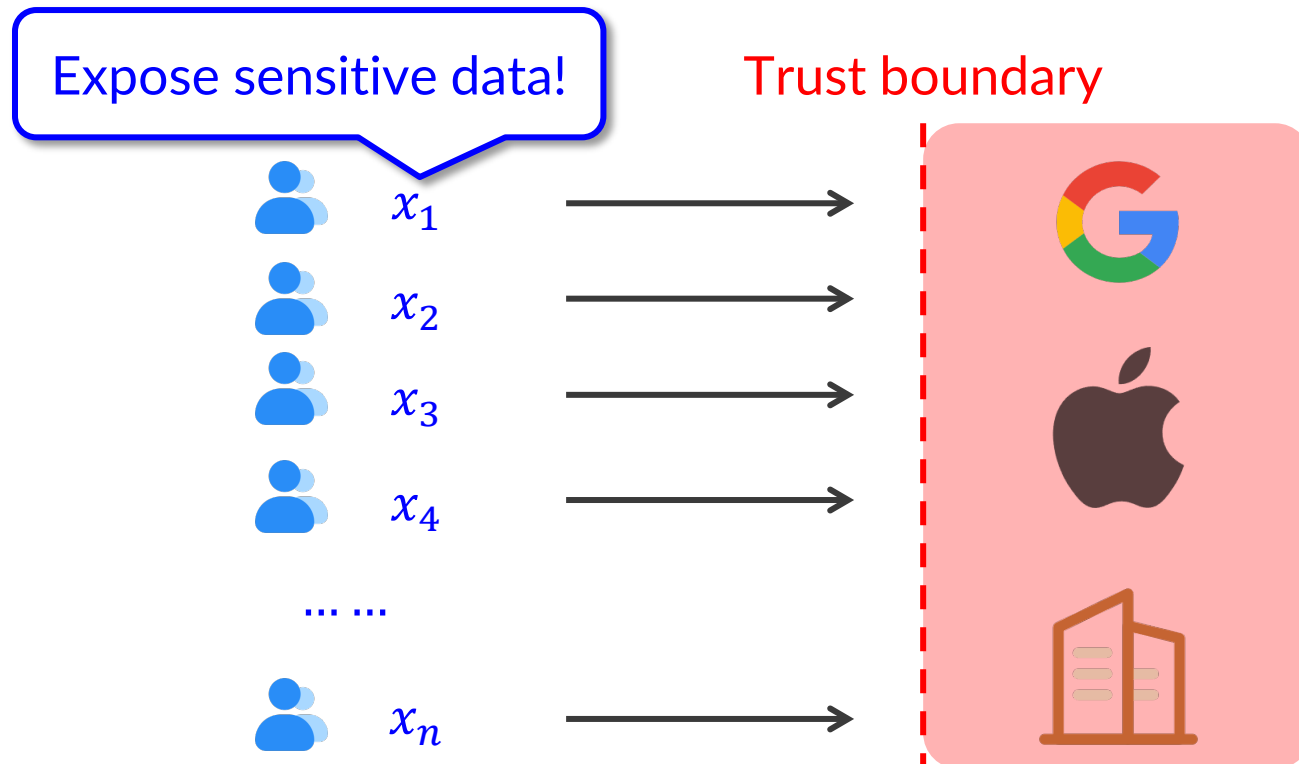


Analysis & service

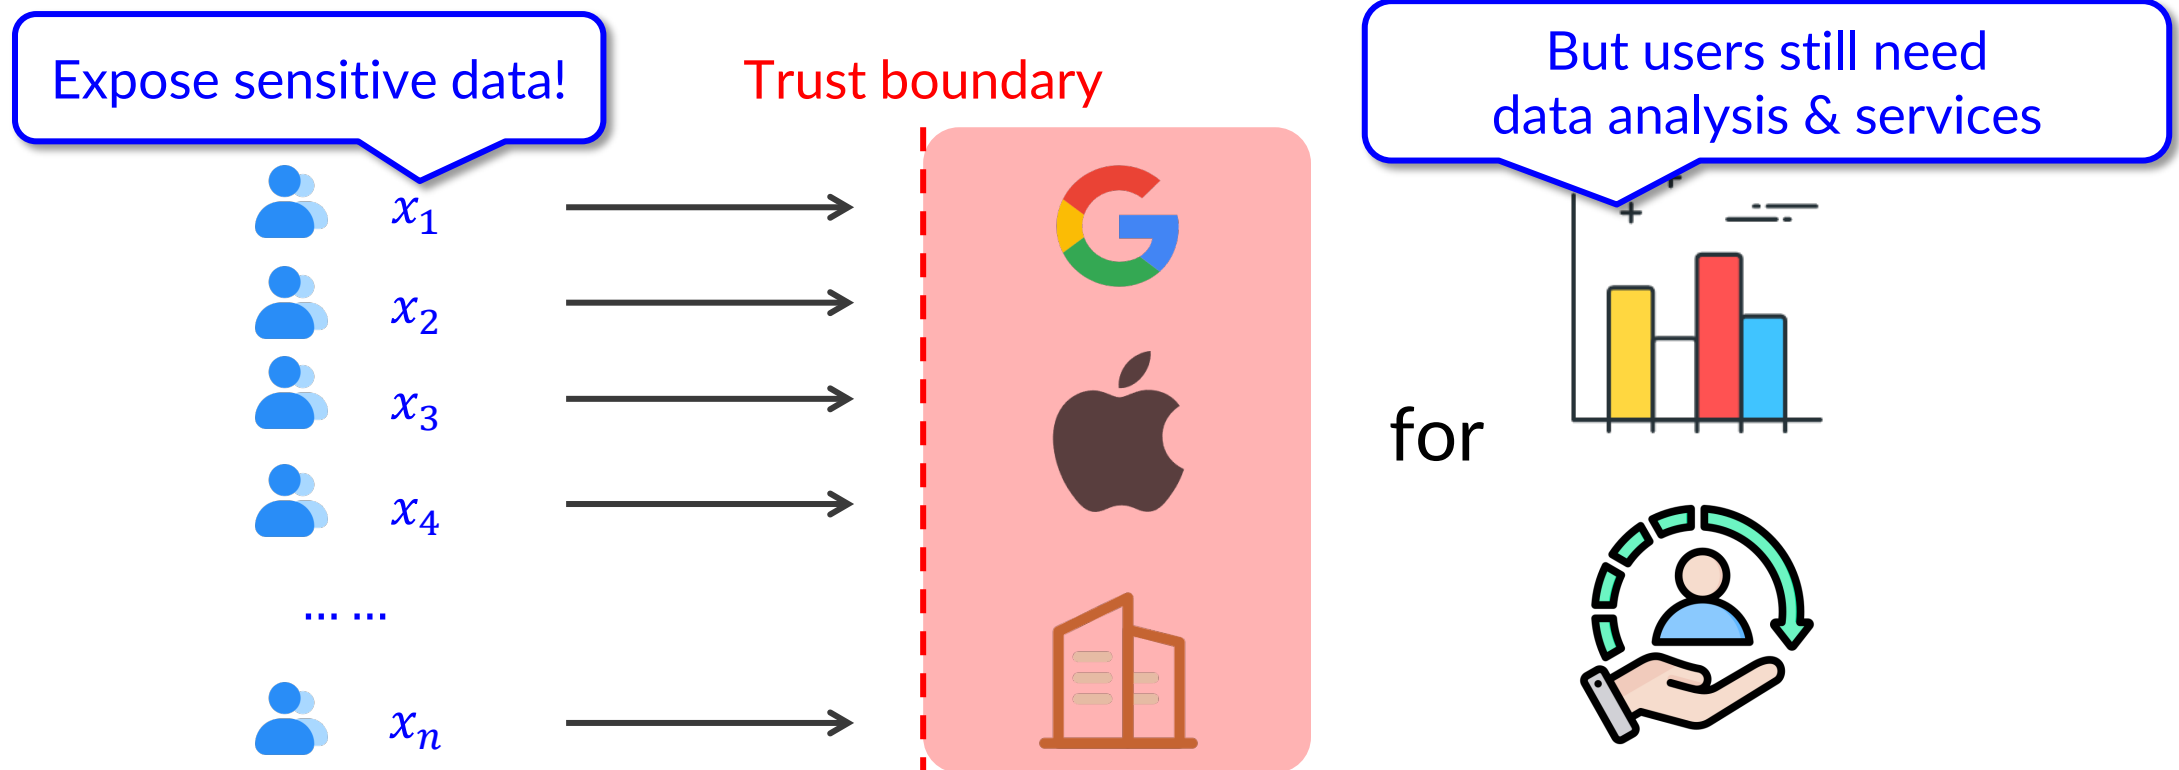
for



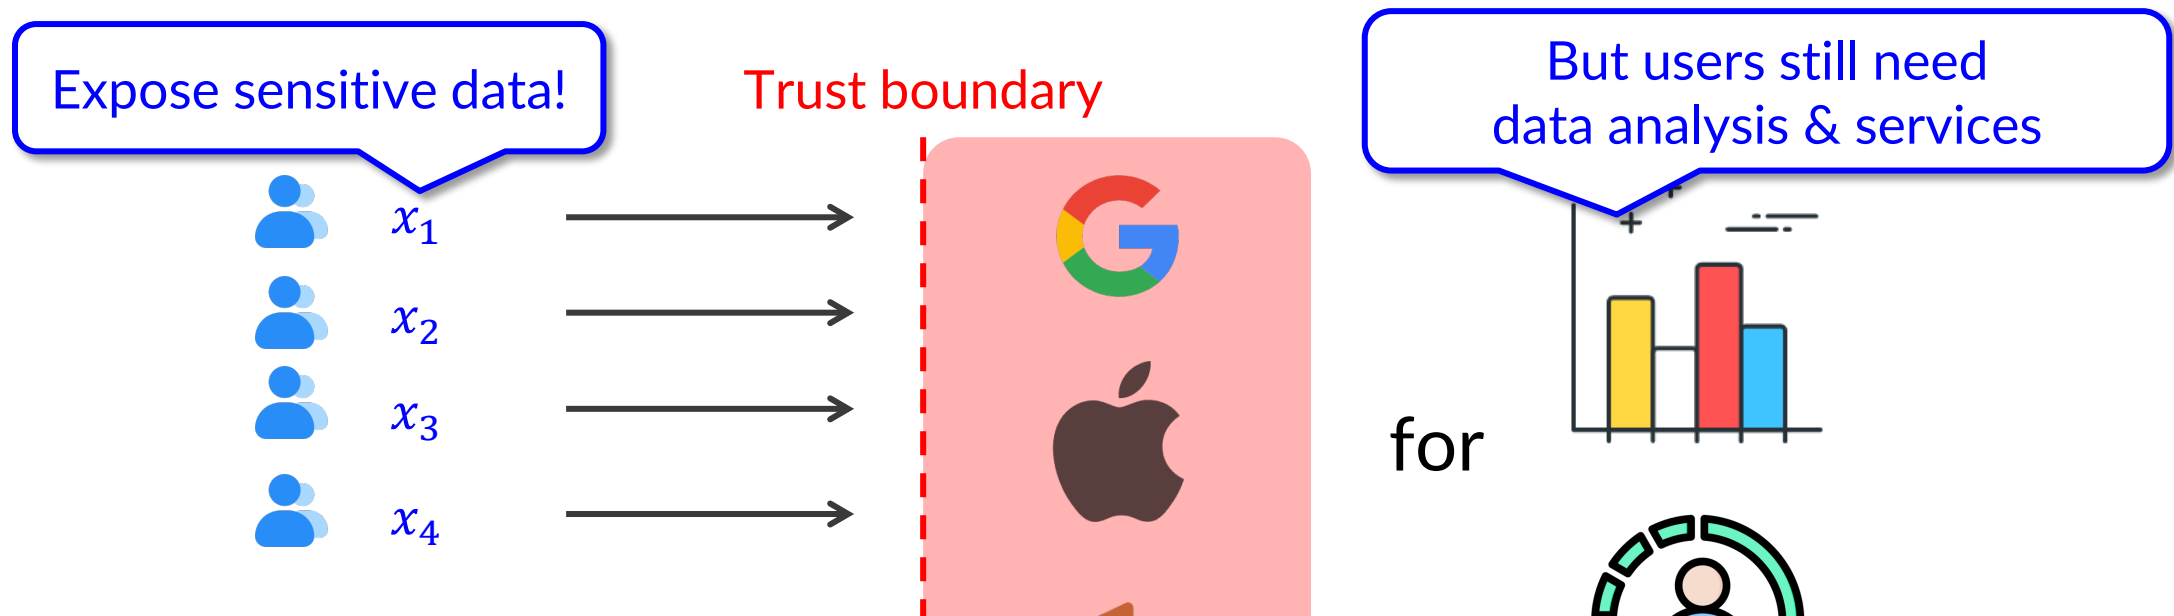
- Users' personal data are collected by companies for analysis or services
 - these companies may **not be trusted** to collect users' sensitive data



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



- 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?

- 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 ?



Privacy-Preserving Computation - Techniques

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

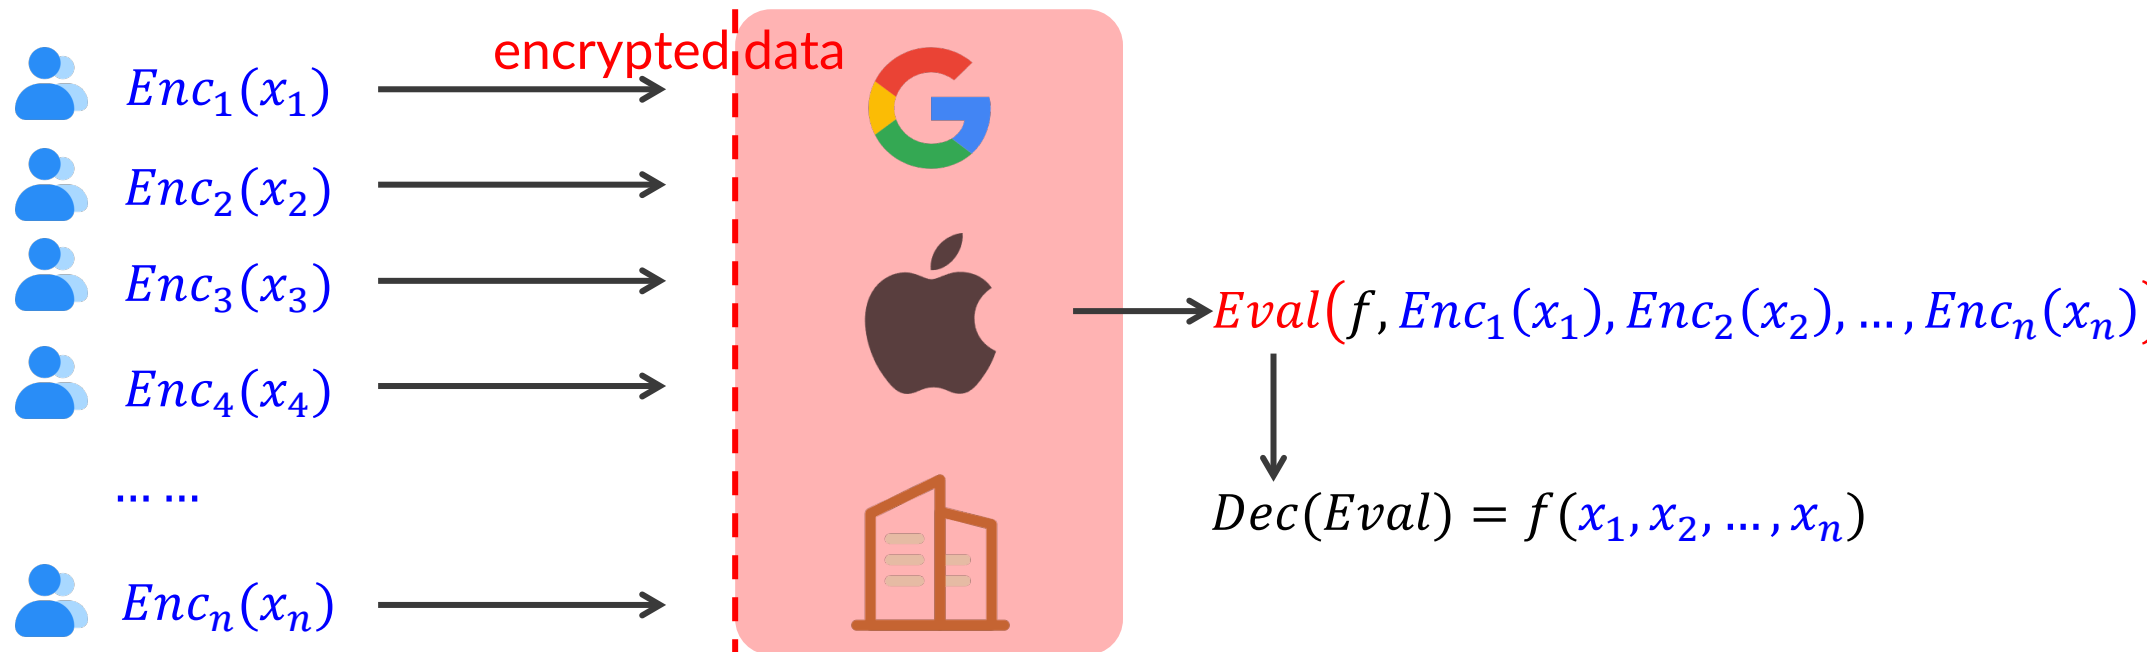
- Homomorphic encryption (HE):
 - “homomorphic”: preserving structure
 - design $\{Enc, Dec, Eval\} \rightarrow Dec(Eval(f, Enc_1(x_1), Enc_2(x_2), \dots, Enc_n(x_n))) = f(x_1, x_2, \dots, x_n)$



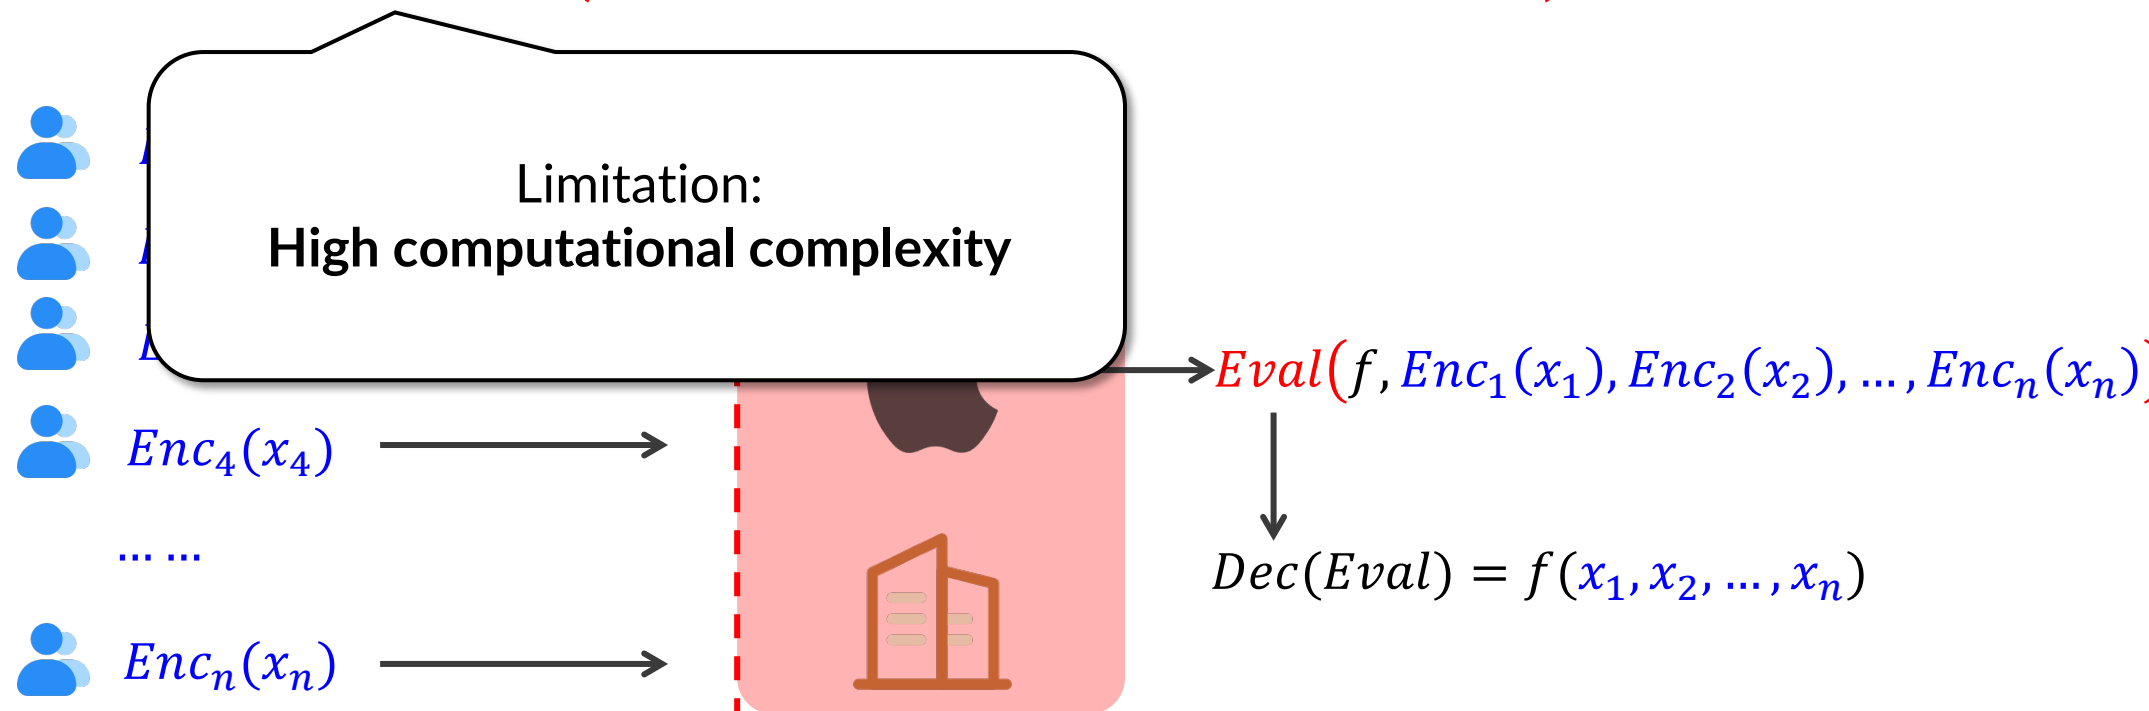
- Homomorphic encryption (HE):

- “homomorphic”: preserving structure

- design $\{Enc, Dec, Eval\} \rightarrow Dec(Eval(f, Enc_1(x_1), Enc_2(x_2), \dots, Enc_n(x_n))) = f(x_1, x_2, \dots, x_n)$



- Homomorphic encryption (HE):
 - “homomorphic”: preserving structure
 - design $\{Enc, Dec, Eval\} \rightarrow Dec(Eval(f, Enc_1(x_1), Enc_2(x_2), \dots, Enc_n(x_n))) = f(x_1, x_2, \dots, x_n)$



- 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

....

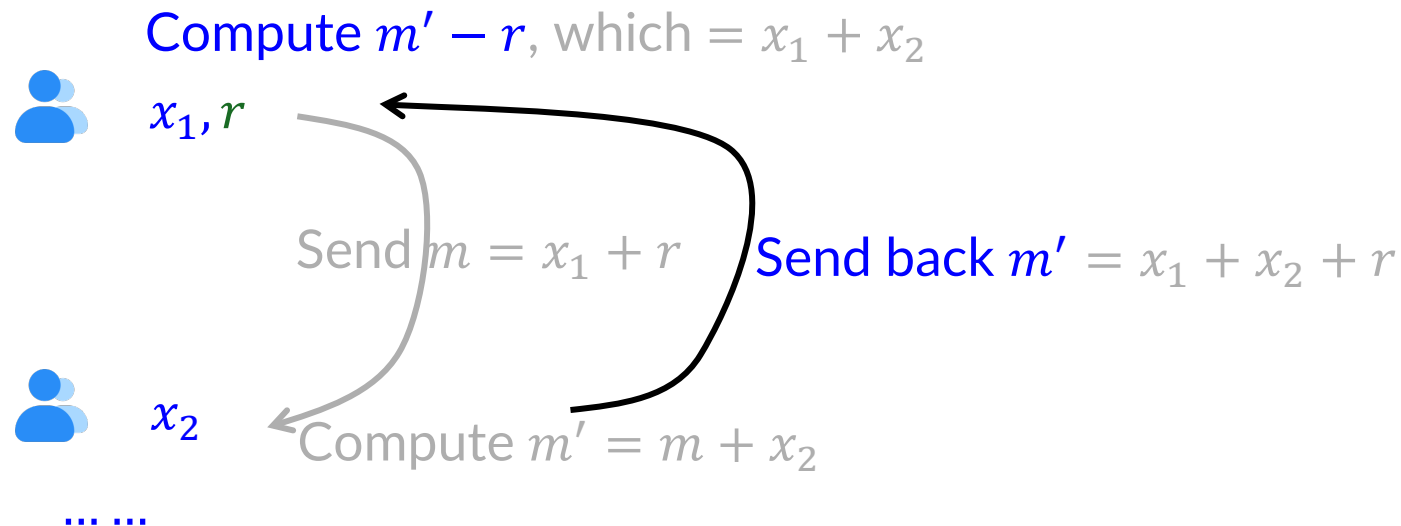
for $f(x_1, x_2)$

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

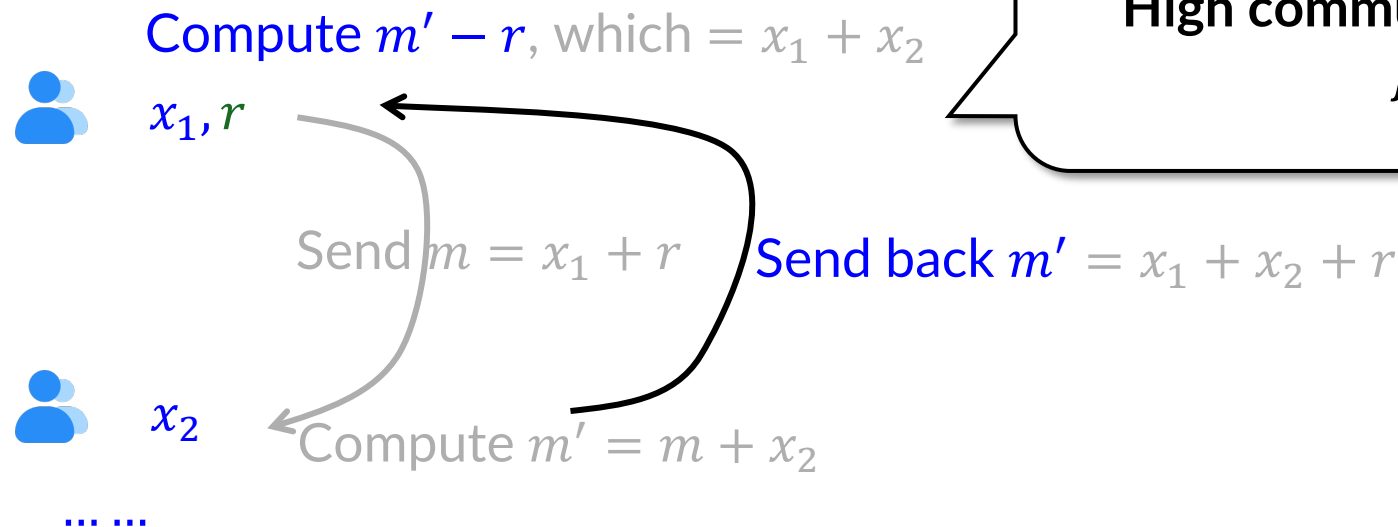


for $f(x_1, x_2)$

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



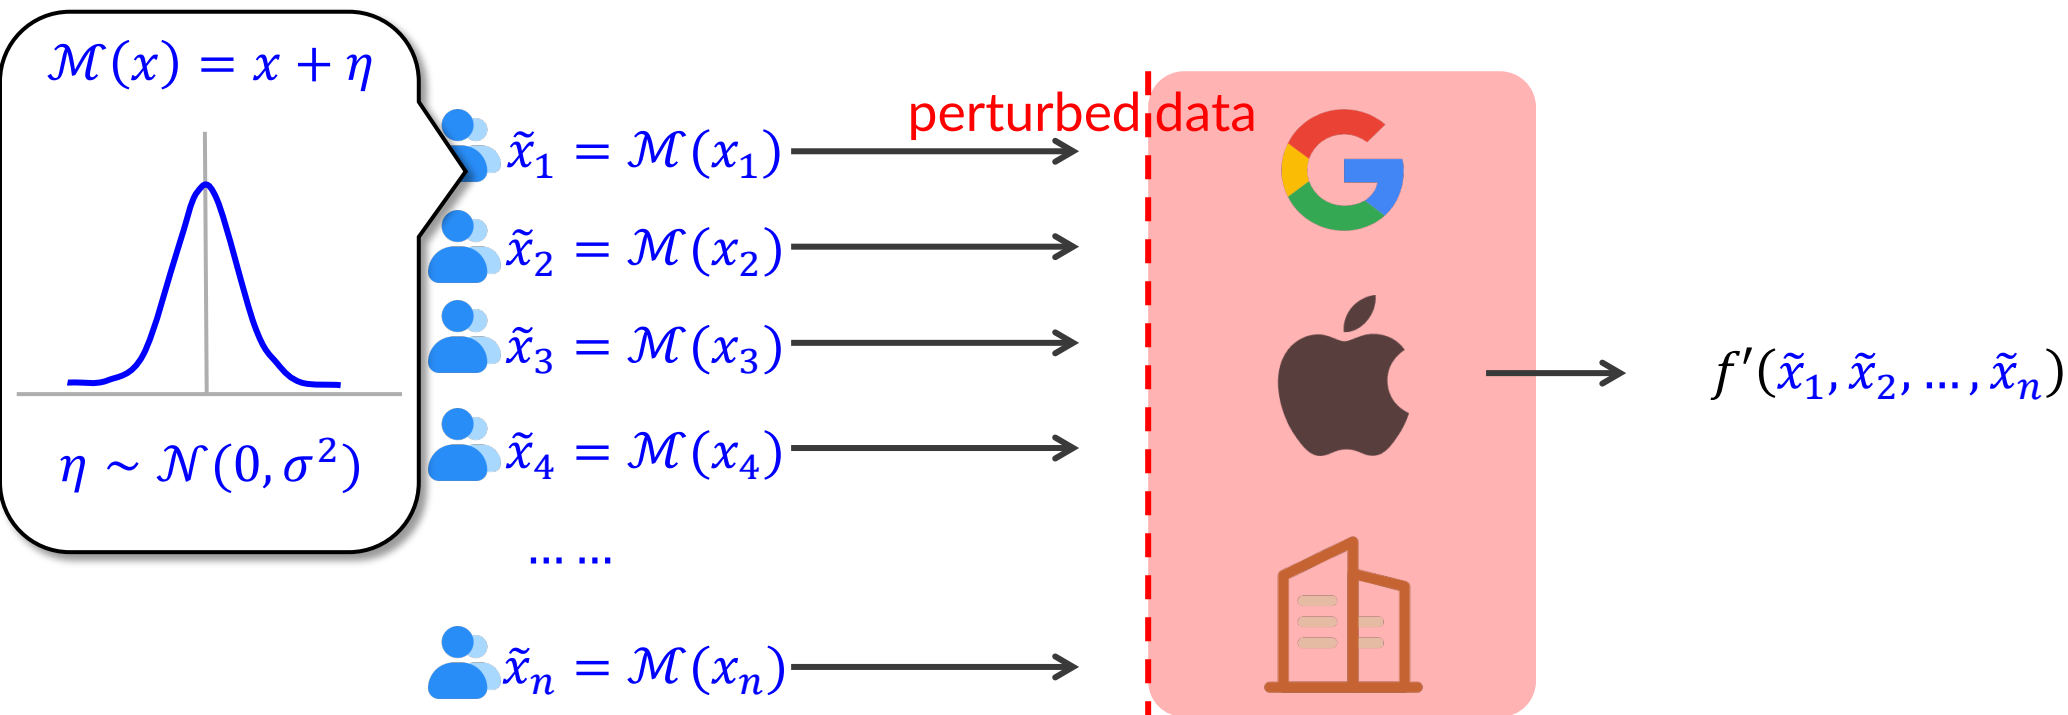
- Multi-party computation (MPC):
 - no central party
 - jointly compute f without revealing x_i
- Example: $f(x_1, x_2) = x_1 + x_2$



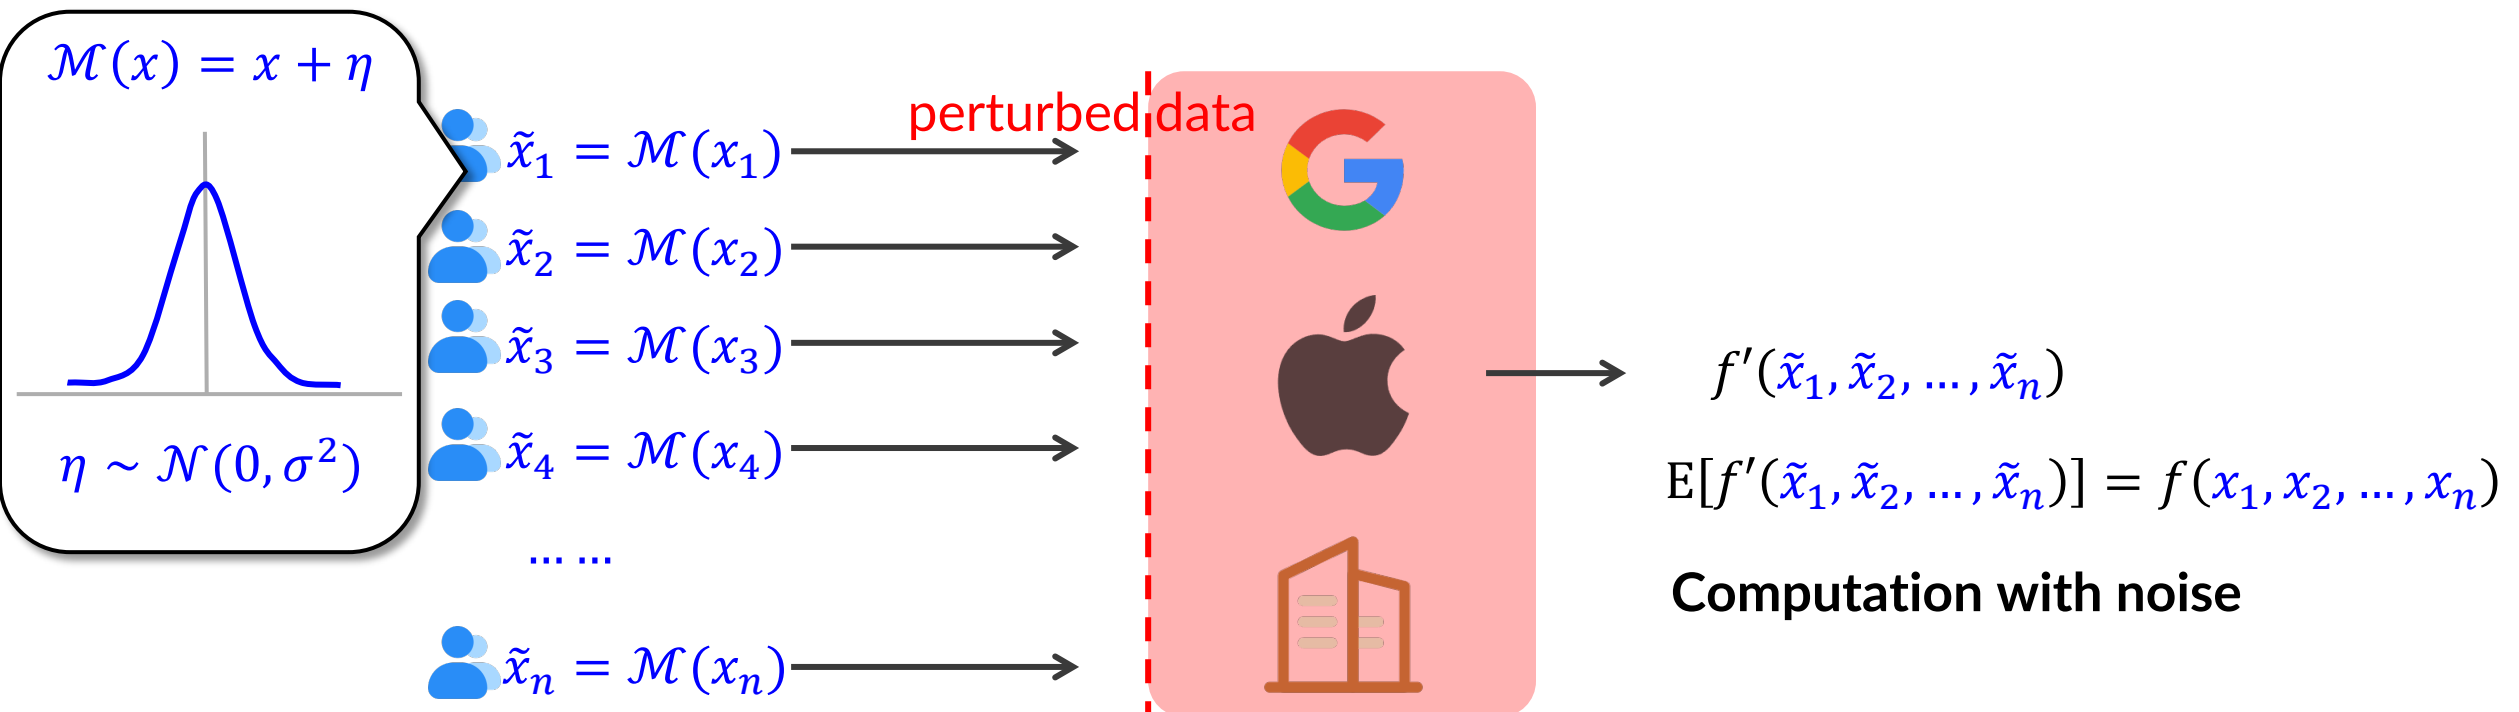
Limitations:
High communication complexity
 f -specific

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

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



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




- Local differential privacy (LDP)
 - **hard to differentiate** the
 - each user locally perturb

Advantages:
Negligible computational complexity
No communication between users


But approximated f

\dots, x_n)


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



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



perturbed data



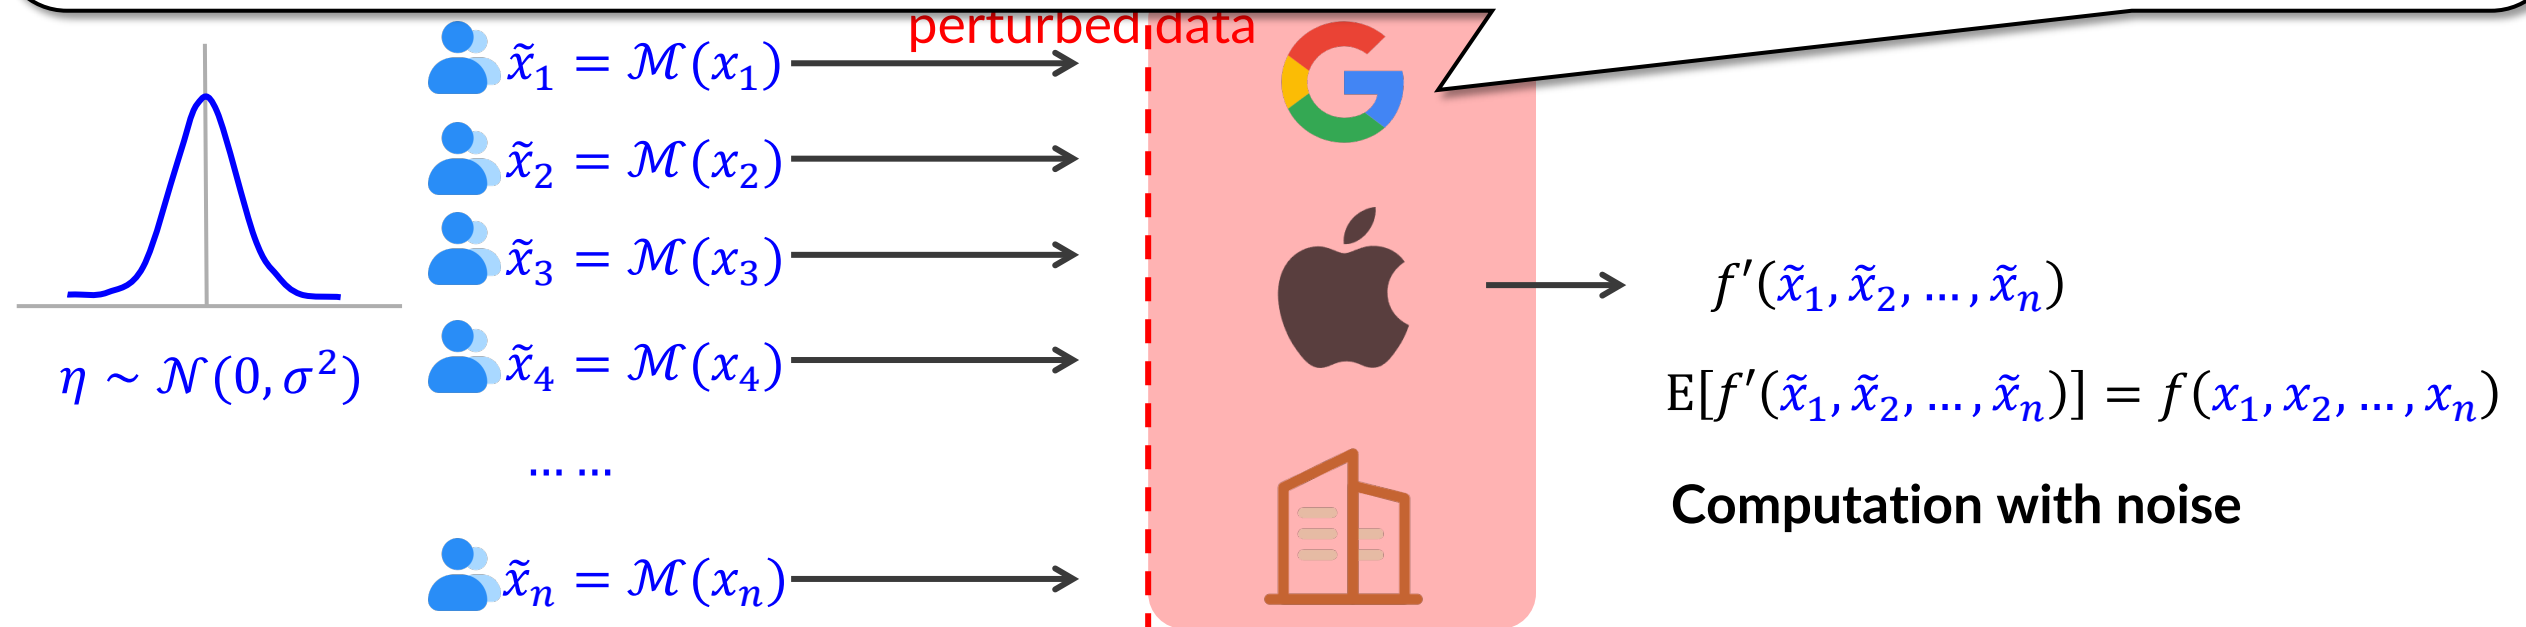
$$f'(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$$

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

Computation with noise



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






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



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



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


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

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

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

... ..

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



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

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

Computation with noise

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




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

Distinguishability


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

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

- 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	

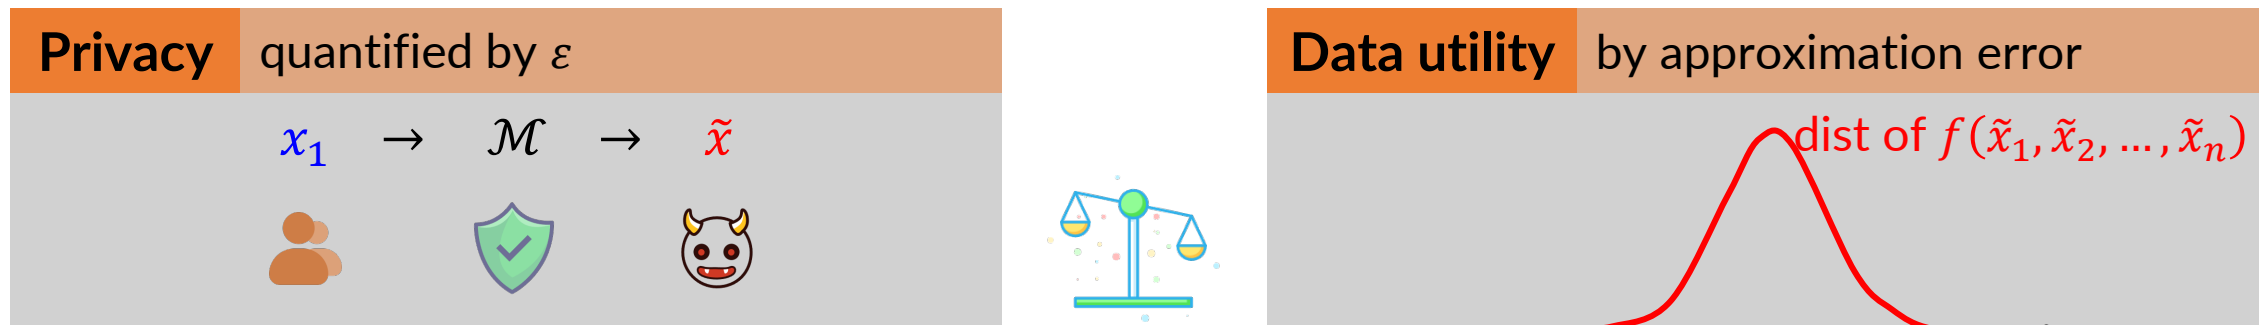


Data utility	by approximation error
	dist of $f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$
	
$f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n) \approx f(x_1, x_2, \dots, x_n)$	

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

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

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



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

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




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

$f(x_1, x_2, \dots, x_n) := h: \mathbb{R}^n \rightarrow \{1, 2, \dots, K\}$ 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 $f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$

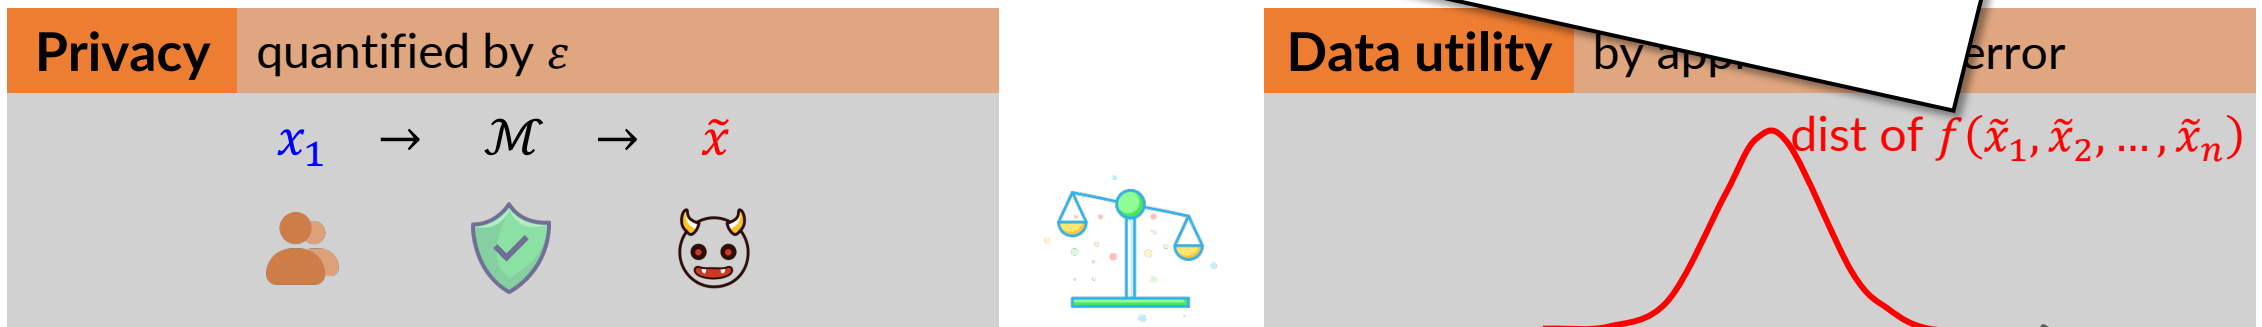


$f(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n) \approx f(x_1, x_2, \dots, x_n)$

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

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

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



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


- Advancing LDP's **mechanism design** and **utility analysis**


- 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


 $\tilde{x}_1 = \mathcal{M}(x_1)$ →

 $\tilde{x}_2 = \mathcal{M}(x_2)$ →

 $\tilde{x}_3 = \mathcal{M}(x_3)$ →

 $\tilde{x}_4 = \mathcal{M}(x_4)$ →

...

 $\tilde{x}_n = \mathcal{M}(x_n)$ →



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


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


- Advancing LDP's mechanism design and utility analysis


binary $x \rightarrow$ numerical x


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

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


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

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

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

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

...

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



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


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


- Advancing LDP's mechanism design and utility analysis


binary $x \rightarrow$ numerical x


Part 1: correlated \mathcal{M} \rightarrow Part 2: optimal piecewise-based \mathcal{M} \downarrow 1D $x \rightarrow$ 2D x

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


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


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

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

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

...

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



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

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

- Advancing LDP's mechanism design and utility analysis

binary $x \rightarrow$ numerical x


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


$1D\ x \rightarrow 2D\ x$


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


Mechanism-level

Task-level


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


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

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

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

...

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



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

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

- **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

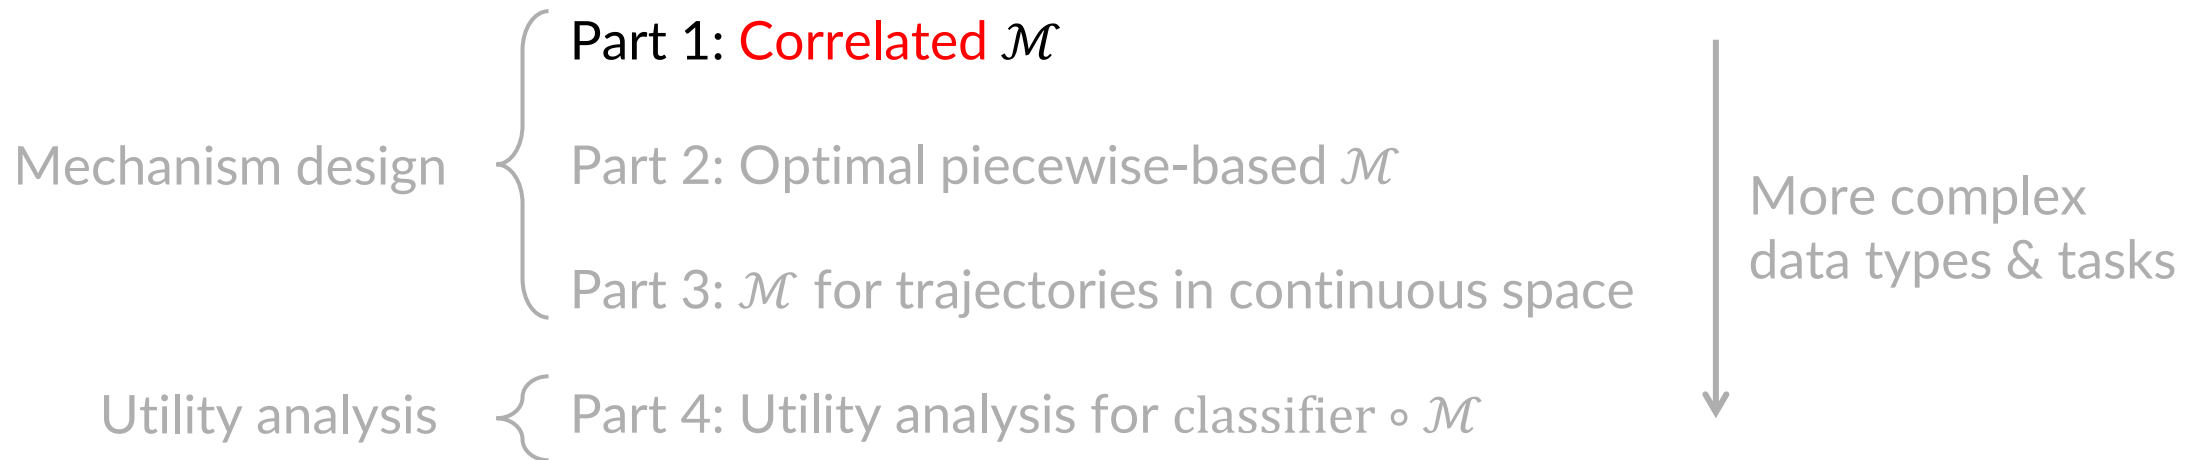

Local Differential Privacy: Refined Mechanism Design and Utility Analysis



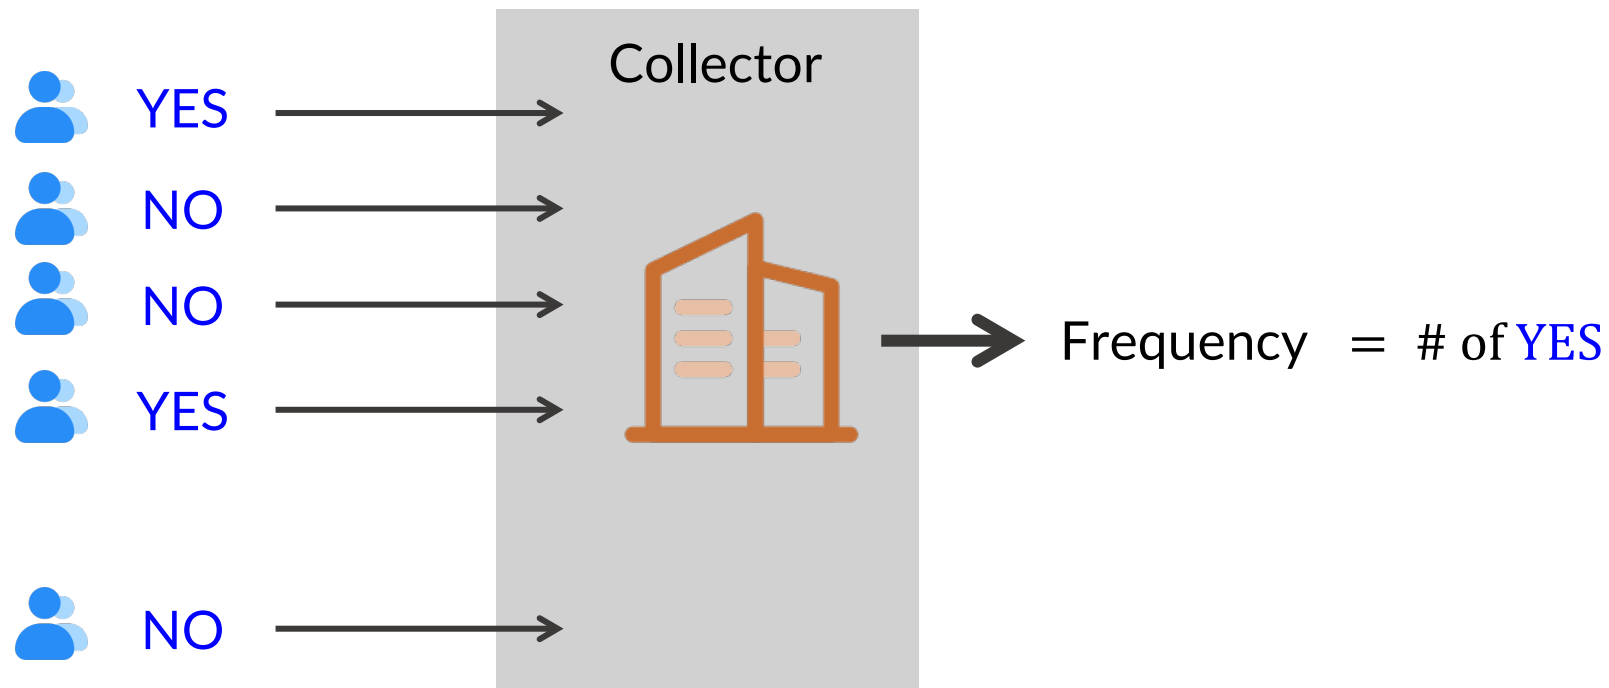
- Mechanism design { Part 1: **Correlated** \mathcal{M}
Part 2: **Optimal** piecewise-based \mathcal{M}
Part 3: \mathcal{M} for trajectories in **continuous space**
- Utility analysis { Part 4: Utility analysis for **classifier** $\circ \mathcal{M}$

More complex
data types & tasks

Local Differential Privacy: Refined Mechanism Design and Utility Analysis



- Social science: How to know how many people engage in tax evasion?
 - ask one person if they had evaded tax
 - the person answers YES or NO



- People have privacy concerns on **sensitive/embarassing question**
 - i.e. **don't want to answer** → how to collect answers in a privacy-preserving way?
- A privacy mechanism \mathcal{M} satisfies LDP if

For any truth x_1, x_2 ,
and randomized answer y :

$$\max \frac{\Pr[\mathcal{M}(x_1) = y]}{\Pr[\mathcal{M}(x_2) = y]} \leq e^\epsilon$$

Distinguishability of x_1 (YES) and x_2 (NO)
from y (randomized answer)

- People have privacy concerns on **sensitive/embarassing question**
 - i.e. **don't want to answer** → how to collect answers in a privacy-preserving way?
- A privacy mechanism \mathcal{M} satisfies LDP if

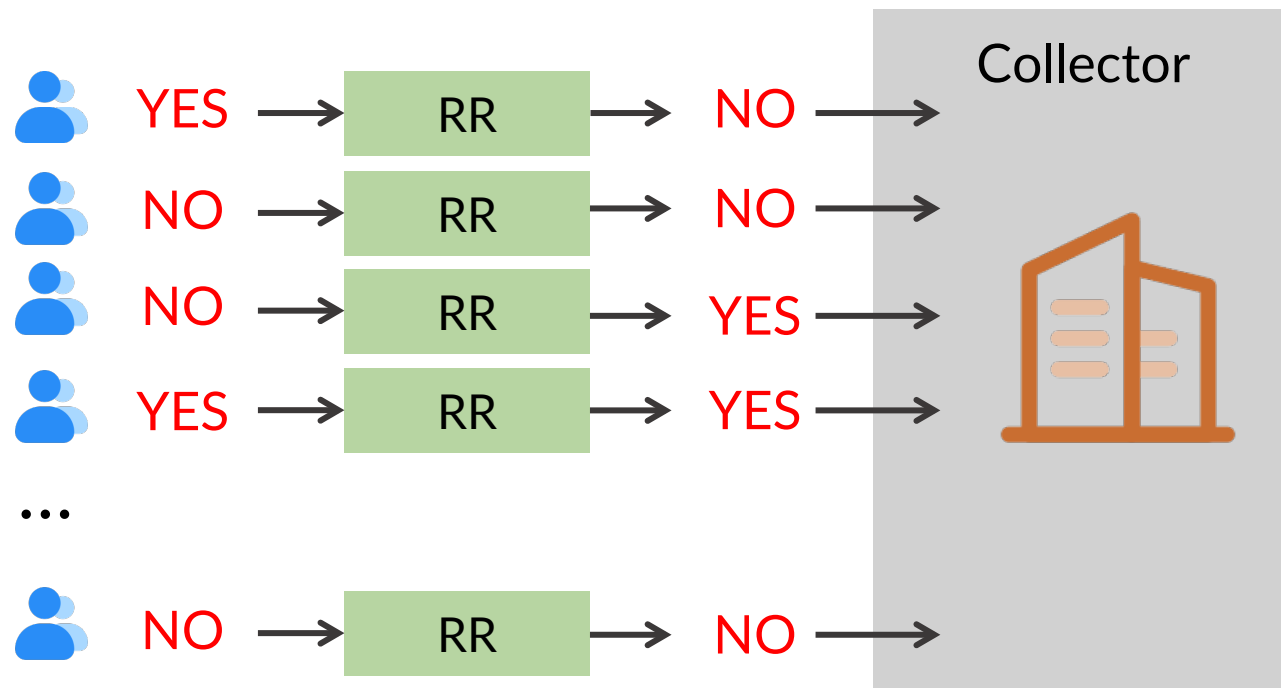
For any truth x_1, x_2 ,
and randomized answer y :

$$\max \frac{\Pr[\mathcal{M}(x_1) = y]}{\Pr[\mathcal{M}(x_2) = y]} \leq e^\epsilon$$

Distinguishability of x_1 (YES) and x_2 (NO)
from y (randomized answer)

- **quantifiable hardness** to distinguish x_1 (YES) and x_2 (NO) from the randomized answer y
- defense against inference from data collectors  or adversaries 

- People have privacy concerns on sensitive/embarrassing question
 - i.e. don't want to answer → how to collect answers in a privacy-preserving way?
- Randomized Response: Randomize the truth before answering the collector



Randomized Response for Privacy

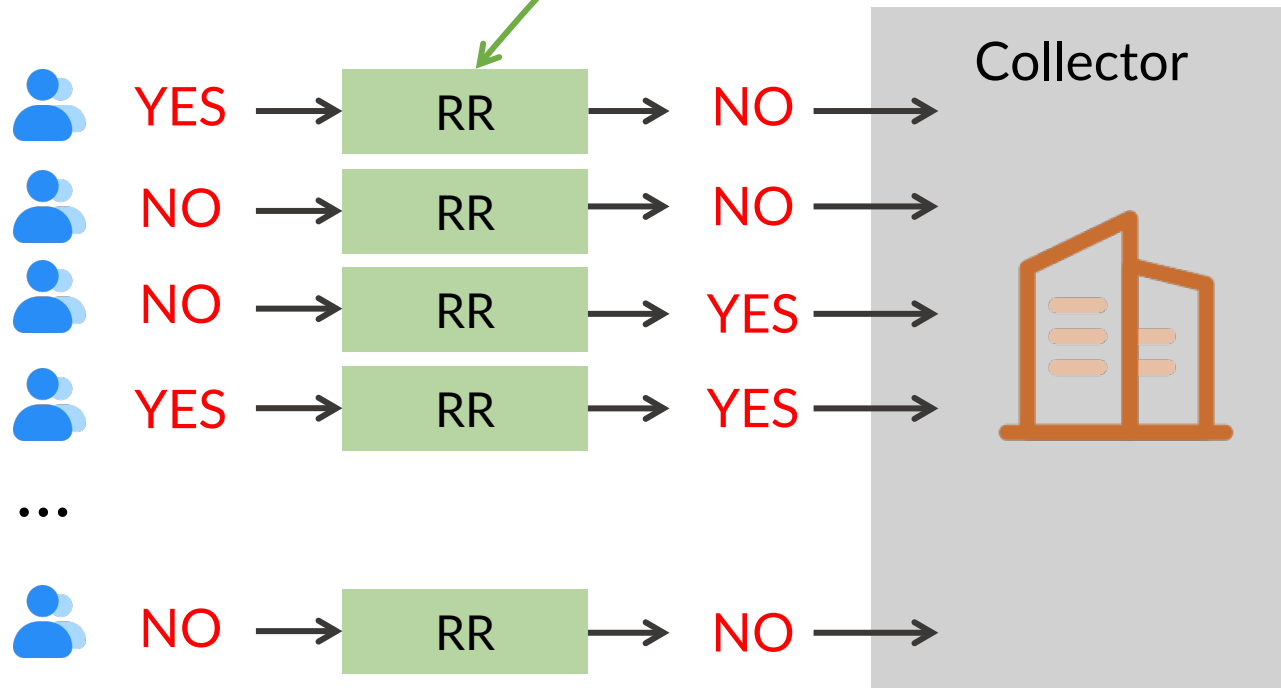
$$\max \frac{\Pr[\mathbf{RR}(x_1) = y]}{\Pr[\mathbf{RR}(x_2) = y]} \leq e^{\ln \frac{p}{1-p}}$$

- People have privacy concerns on sensitive/embarrassing questions - i.e. don't want to answer → how to collect answers in a private way
- Randomized Response: Randomize the truth before answering

RR: [Warner, 1965]
answer truth with probability p

$$\mathbf{RR}(x) = \begin{cases} x & \text{w. p. } p \\ \neg x & \text{w. p. } 1 - p \end{cases}$$

Private



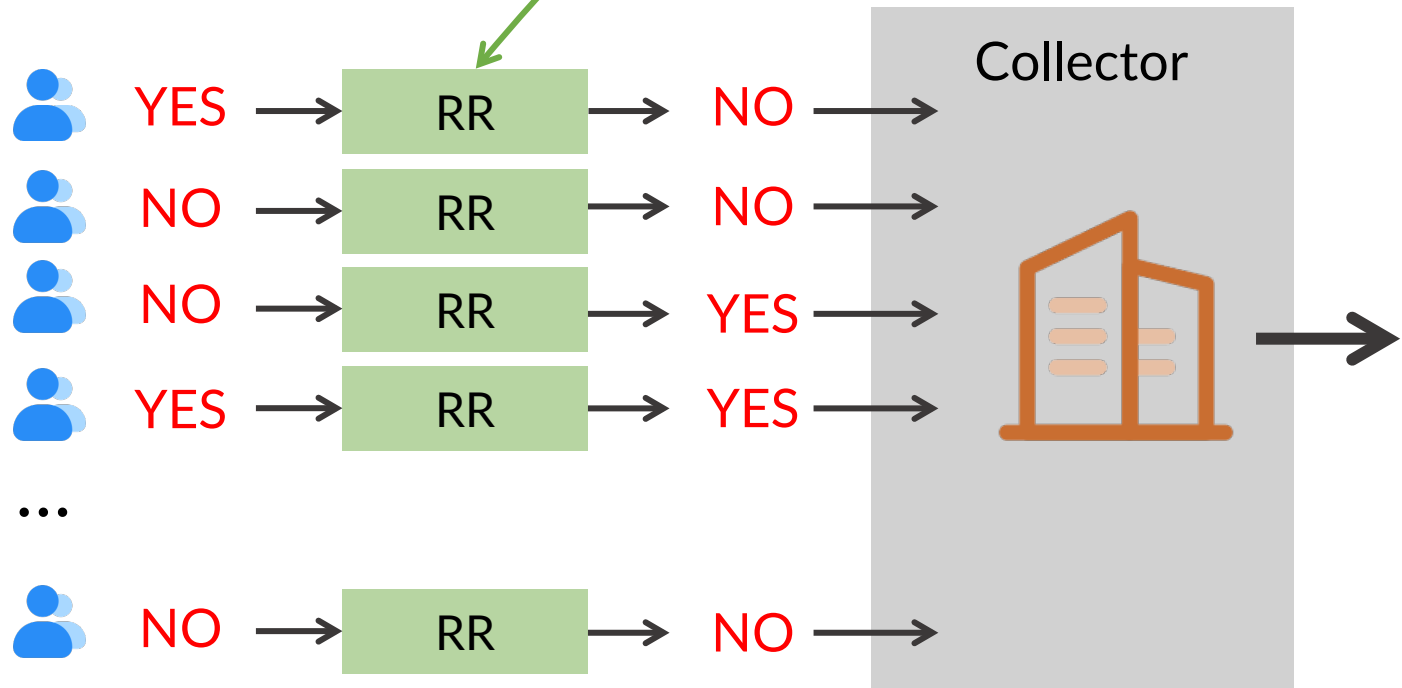
Randomized Response for Privacy

$$\max \frac{\Pr[\mathbf{RR}(x_1) = y]}{\Pr[\mathbf{RR}(x_2) = y]} \leq e^{\ln \frac{p}{1-p}}$$

- People have privacy concerns on sensitive/embarrassing questions - i.e. don't want to answer → how to collect answers in a private way
- Randomized Response: Randomize the truth before answering

Private

RR: [Warner, 1965]
answer truth with probability p

$$\mathbf{RR}(x) = \begin{cases} x & \text{w. p. } p \\ \neg x & \text{w. p. } 1 - p \end{cases}$$


estimated frequency

$$= \frac{\# \text{ of YES} - \# \text{ of people} \times q}{p - q}$$

Unbiased:
expectation = truth

- Randomization reduces data utility

$$\text{Var}\left[\frac{\# \text{ of YES} - \# \text{ of people} \times q}{p - q}\right] = \frac{\text{Var}[\# \text{ of YES}]}{(p - q)^2} = \frac{npq}{(p - q)^2}$$

- **classical result:** summation of variance from all n independent randomization

- Randomization reduces data utility

$$\text{Var}\left[\frac{\# \text{ of YES} - \# \text{ of } \text{👤} \times q}{p - q}\right] = \frac{\text{Var}[\# \text{ of YES}]}{(p - q)^2} = \frac{npq}{(p - q)^2}$$

- **classical result:** summation of variance from all n independent randomization

- Q: Can we do better?


- Randomization reduces data utility

$$\text{Var}\left[\frac{\# \text{ of YES} - \# \text{ people} \times q}{p - q}\right] = \frac{\text{Var}[\# \text{ of YES}]}{(p - q)^2} = \frac{npq}{(p - q)^2}$$


- **classical result:** summation of variance from all n **independent** randomization

- Q: Can we do better?
 - yes, by correlated (joint) randomization

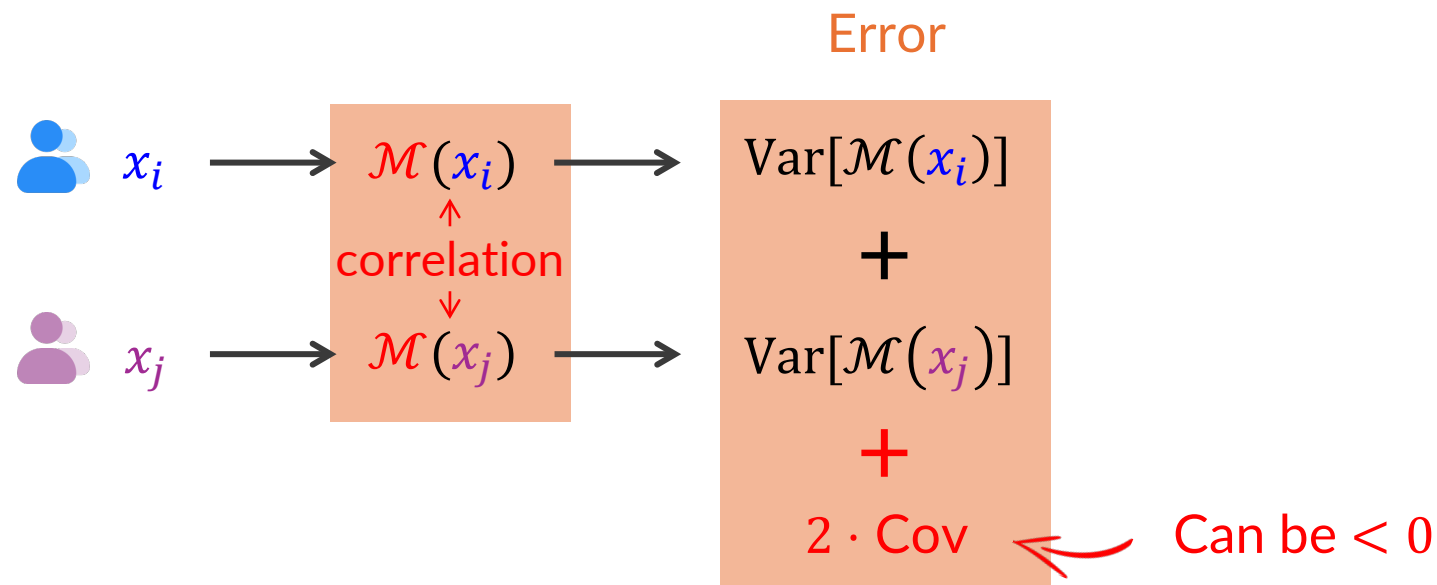
- Existing LDP mechanisms: Each user perturbs their data **independently**



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



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



- JRR: Better data utility by joint/correlated randomization
- **Example:** 2-person ($x_1 = \text{YES}$ and $x_2 = \text{YES}$) with $p = 0.8$ ($P[T = 1] = 0.8$)

RR: Joint distribution

	$T_1 = 1$	$T_1 = 0$	Truthfulness of x_1
$T_2 = 1$	0.64 (= p^2)	0.16 (= pq)	
$T_2 = 0$	0.16 (= pq)	0.04 (= q^2)	

Truthfulness
of x_2

Joint/Correlated RR (JRR)

- JRR: Better data utility by joint/correlated randomization
- **Example:** 2-person ($x_1 = \text{YES}$ and $x_2 = \text{YES}$) with $p = 0.8$ ($P[T = 1] = 0.8$)

RR: Joint

$P[T_1 = 1] = 0.8$

$P[T_2 = 1] = 0.8$

	$T_1 = 1$	$T_1 = 0$	Truthfulness of x_1
$T_2 = 1$	0.64 (= p^2)	0.16 (= pq)	
$T_2 = 0$	0.16 (= pq)	0.04 (= q^2)	

Truthfulness of x_2

Independent T_1 and T_2 ($P[T_1 \cap T_2] = P[T_1] \cdot P[T_2]$)

Joint probability = Π of marginal probabilities

- JRR: Better data utility by joint/correlated randomization
- Example:** 2-person ($x_1 = \text{YES}$ and $x_2 = \text{YES}$) with $p = 0.8$ ($P[T = 1] = 0.8$)

RR: Joint

$P[T_1 = 1] = 0.8$

$P[T_2 = 1] = 0.8$

	$T_1 = 1$	$T_1 = 0$
$T_2 = 1$	0.64 (= p^2)	0.16 (= pq)
$T_2 = 0$	0.16	0.04 (= q^2)

$\text{Cov}[y_1, y_2] = 0$

JRR: Joint distribution

	$T_1 = 1$	$T_1 = 0$
$T_2 = 1$	0.6 (= $p^2 + \rho pq$)	0.2 (= $pq - \rho pq$)
$T_2 = 0$	0.2 (= $pq - \rho pq$)	0 (= $q^2 + \rho pq$)

Independent T_1 and T_2 ($P[T_1 \cap T_2] = P[T_1] \cdot P[T_2]$)

Joint probability = Π of marginal probabilities

Joint/Correlated RR (JRR)

- JRR: Better data utility by joint/correlated randomization
- Example:** 2-person ($x_1 = \text{YES}$ and $x_2 = \text{YES}$) with $p = 0.8$ ($P[T = 1] = 0.8$)

Same marginal prob for each person (same ϵ)

RR: Joint

$P[T_1 = 1] = 0.8$

$P[T_2 = 1] = 0.8$	$T_1 = 1$	$T_1 = 0$
$T_2 = 1$	0.64 (= p^2)	0.16 (= pq)
$T_2 = 0$	0.16	0.04 (= q^2)

$\text{Cov}[y_1, y_2] = 0$

JRR: Joint

$P[T_1 = 1] = 0.8$

$P[T_2 = 1] = 0.8$	$T_1 = 1$	$T_1 = 0$
$T_2 = 1$	0.6 (= $p^2 + \rho pq$)	0.2 (= $pq - \rho pq$)
$T_2 = 0$	0.2 (= $pq - \rho pq$)	0 (= $q^2 + \rho pq$)

Independent T_1 and T_2 ($P[T_1 \cap T_2] = P[T_1] \cdot P[T_2]$)

Joint probability = Π of marginal probabilities

Joint/Correlated RR (JRR)

- JRR: Better data utility by joint/correlated randomization
- Example:** 2-person ($x_1 = \text{YES}$ and $x_2 = \text{YES}$) with $p = 0.8$ ($P[T = 1] = 0.8$)

Same marginal prob for each person (same ϵ)

RR: Joint

$P[T_1 = 1] = 0.8$

$P[T_2 = 1] = 0.8$

	$T_1 = 1$	$T_1 = 0$
$T_2 = 1$	0.64 (= p^2)	0.16 (= pq)
$T_2 = 0$	0.16	0.04 (= q^2)

$\text{Cov}[y_1, y_2] = 0$

Independent T_1 and T_2 ($P[T_1 \cap T_2] = P[T_1] \cdot P[T_2]$)

Joint probability = Π of marginal probabilities

JRR: Joint

$P[T_1 = 1] = 0.8$

$P[T_2 = 1] = 0.8$

	$T_1 = 1$	$T_1 = 0$
$T_2 = 1$	0.6 (= $p^2 + \rho pq$)	0.2 (= $pq - \rho pq$)
$T_2 = 0$	0.2	0 (= $q^2 + \rho pq$)

$\text{Cov}[y_1, y_2] = -0.04$

$P[T_1 = 0 \cap T_2 = 0] = 0 \neq P[T_1 = 0] \cdot P[T_2 = 0] = 0.04$

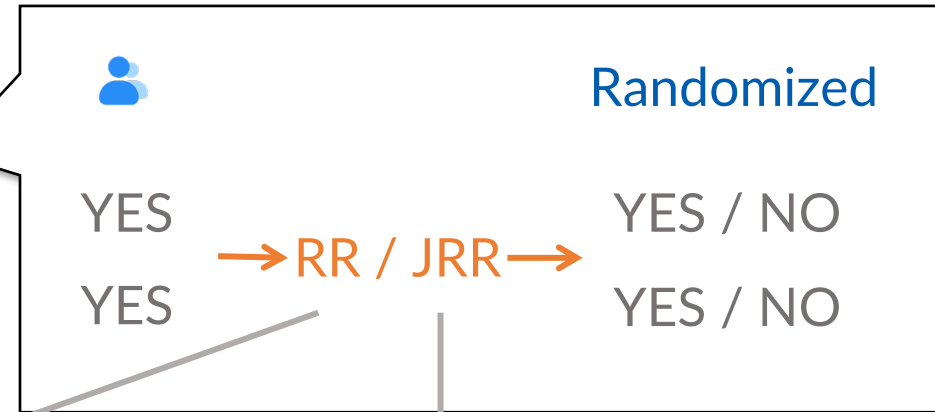
NOT independent T_1 and T_2

Joint probability \neq Π of marginal probabilities

Utility: JRR's Variance

- Variance: (# 👤 = 2, $p = 0.8$)

$$\text{Var}[\hat{n}_{\text{YES}}] = \frac{\text{Var}[\# \text{ of YES}]}{(0.8 - 0.2)^2}$$



RR

$$\text{Var}[\# \text{ of YES}] = 2 \times pq = \mathbf{0.32}$$

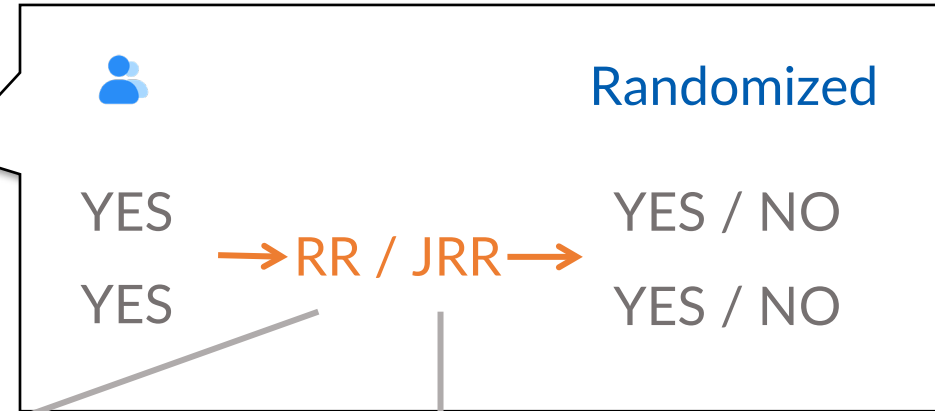
JRR

$$\begin{aligned} \text{Var}[\# \text{ of YES}] &= 0.32 + 2 \cdot \mathbf{Cov}[y_1, y_2] \\ &= 0.32 + 2 \cdot (-0.04) \\ &= \mathbf{0.24} \end{aligned}$$

Utility: JRR's Variance

- Variance: (# 👤 = 2, $p = 0.8$)

$$\text{Var}[\hat{n}_{\text{YES}}] = \frac{\text{Var}[\# \text{ of YES}]}{(0.8 - 0.2)^2}$$



RR

$$\text{Var}[\# \text{ of YES}] = 2 \times pq = \mathbf{0.32}$$

JRR

$$\begin{aligned} \text{Var}[\# \text{ of YES}] &= 0.32 + 2 \cdot \mathbf{Cov}[y_1, y_2] \\ &= 0.32 + 2 \cdot (-0.04) \\ &= \mathbf{0.24} \end{aligned}$$

Better data utility

- Correlated randomization with 2 persons x_{2i-1} and x_{2i}

JRR: Joint distribution


	$T_{2i-1} = 1$	$T_{2i-1} = 0$
$T_{2i} = 1$	$p^2 + \rho pq$	$(1 - \rho)pq$
$T_{2i} = 0$	$(1 - \rho)pq$	$q^2 + \rho pq$

$\rho \in [-1,1]$:
correlation coefficient

- RR is a special case of JRR with $\rho = 0$ (no correlation)


JRR – Variance Heatmap

- Effect of ρ and p (when $\varepsilon = 1, n = 10^4, n_{\text{Yes}} = 200$, and malicious users $m = 0$ & 500)




JRR – Variance Heatmap


- Effect of ρ and p (when $\varepsilon = 1, n = 10^4, n_{\text{Yes}} = 200$, and malicious users $m = 0$ & 500)



- Comparison with RR under the same privacy level - JRR: $\epsilon(n, m, \rho, p)$, RR: $\epsilon(p)$

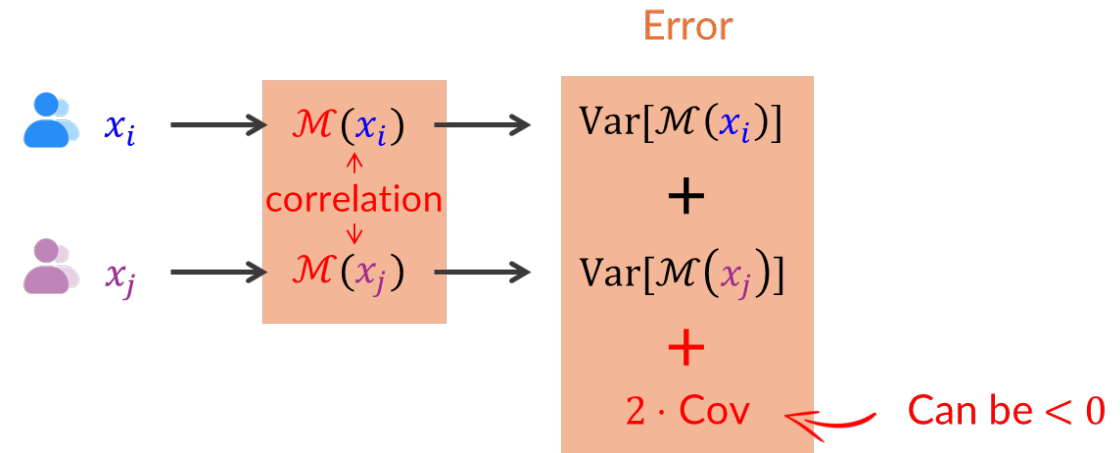


Real-world datasets ($\epsilon = 0.1$)



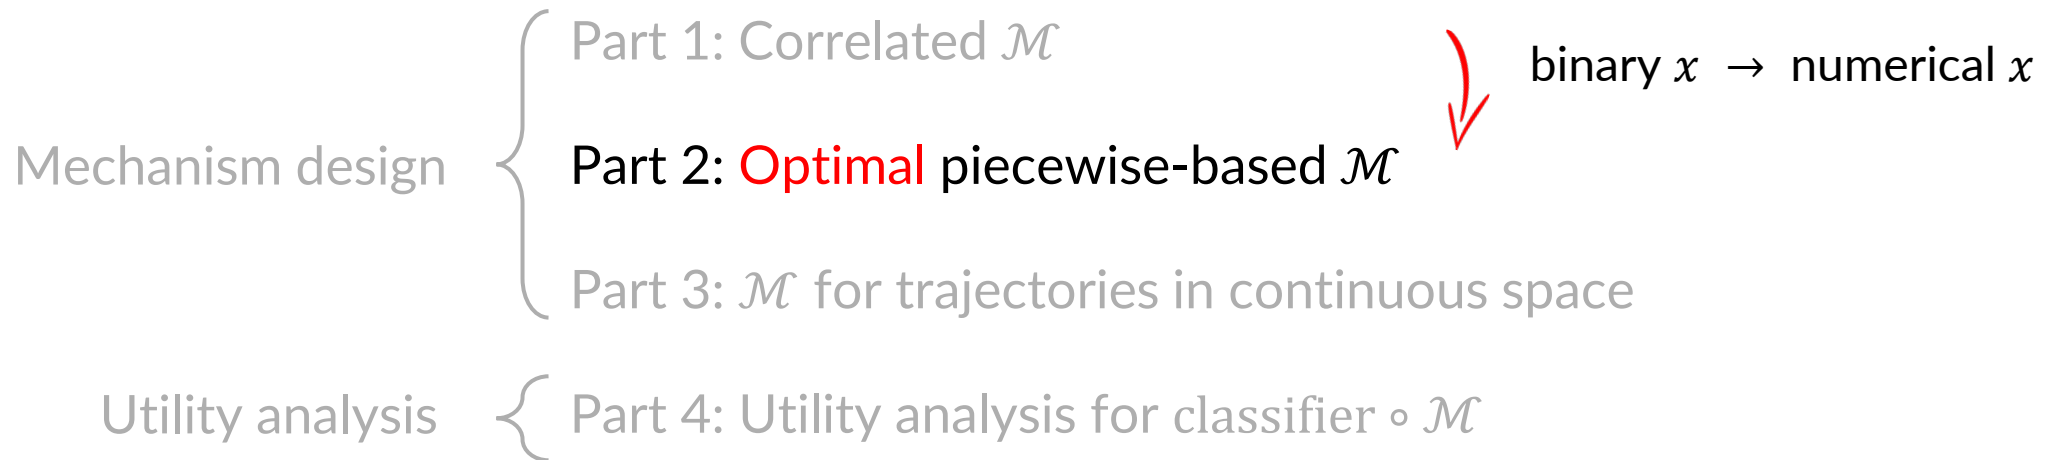

Synthetic datasets ($n = 10^4$)

- RQ: Can the classical RR be better?
- Contributions:^{*}
 - introducing correlation into RR
 - extensible to other LDP mechanisms
 - RR as a special case of $\rho = 0$
 - analytical privacy & data utility analysis



^{*} Locally Differentially Private Frequency Estimation via Joint Randomized Response, PETS'25

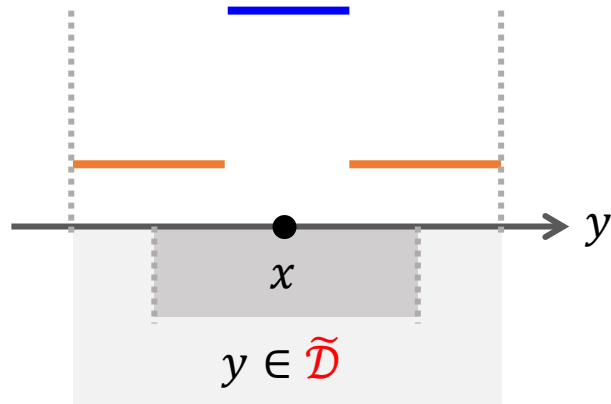
Local Differential Privacy: Refined Mechanism Design and Utility Analysis



3-Piecewise Mechanism



- 3-piecewise distributions on bounded numerical domain $\mathcal{D} \rightarrow \tilde{\mathcal{D}}$
 - given input x , sample output y from a distribution




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

3-Piecewise Mechanism

- 3-piecewise distributions on bounded numerical domain $\mathcal{D} \rightarrow \tilde{\mathcal{D}}$

- given input x , sample output y from a distribution



Sampling probability depends on ε

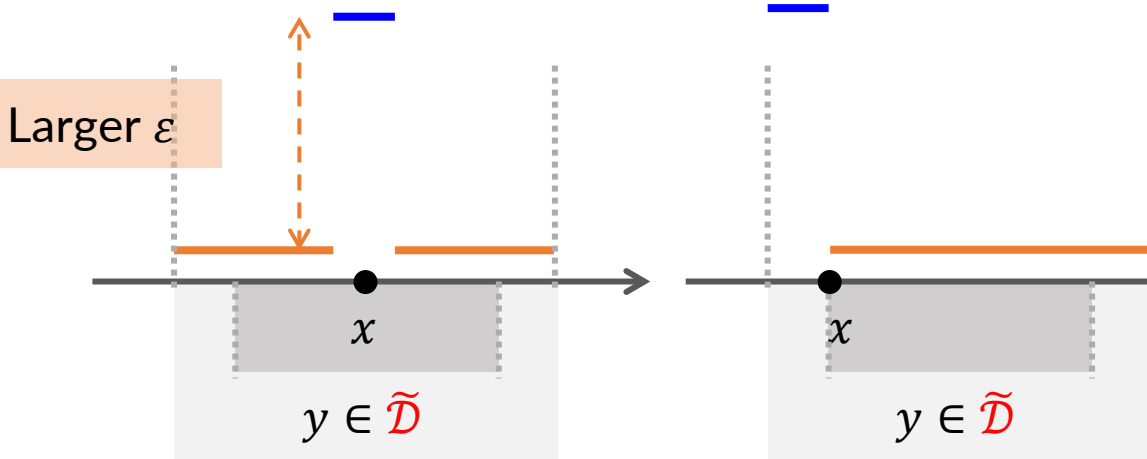
Sampling interval depends on x and ε

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

3-Piecewise Mechanism

- 3-piecewise distributions on bounded numerical domain $\mathcal{D} \rightarrow \tilde{\mathcal{D}}$

- given input x , sample output y from a distribution



Sampling probability depends on ϵ

Sampling interval depends on x and ϵ

$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_\epsilon & \text{if } y \in [l_{x,\epsilon}, r_{x,\epsilon}] \\ \frac{p_\epsilon}{\exp(\epsilon)} & \text{if } y \in \tilde{\mathcal{D}} \setminus [l_{x,\epsilon}, r_{x,\epsilon}] \end{cases}$$

- Instantiations: PM [ICDE'19], SW [SIGMOD'20], PTT [TMC'24] (design different $p_\epsilon, l_{x,\epsilon}, r_{x,\epsilon}$)
 - different errors, but **without optimality**, which is important for building blocks

3-Piecewise Mechanism

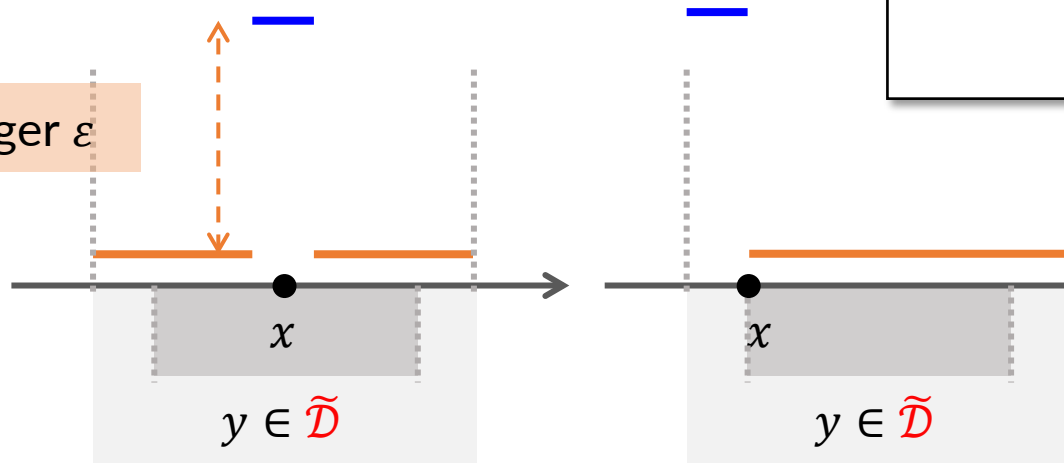
- 3-piecewise distributions on **bound**
 - given input x , sample output y



NOT enough to study **optimality** of piecewise-based mechanism

- only 3 pieces, two probabilities

Larger ε



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

- Instantiations: PM [ICDE'19], SW [SIGMOD'20], PTT [TMC'24] (design different $p_\varepsilon, l_{x,\varepsilon}, r_{x,\varepsilon}$)
 - different errors, but **without optimality**, which is important for building blocks

3-Piecewise Mechanism


- 3-piecewise distributions on **bound**
 - given input x , sample output y



NOT enough to study optimality of piecewise-based mechanism

- only 3 pieces, two probabilities
- piecewise distribution can have more pieces and probabilities

Larger ϵ




$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_\epsilon & \text{if } y \in [l_{x,\epsilon}, r_{x,\epsilon}] \\ \frac{p_\epsilon}{\exp(\epsilon)} & \text{if } y \in \tilde{\mathcal{D}} \setminus [l_{x,\epsilon}, r_{x,\epsilon}] \end{cases}$$

- Instantiations: PM [ICDE'19], SW [SIGMOD'20], PTT [TMC'24] (design different $p_\epsilon, l_{x,\epsilon}, r_{x,\epsilon}$)
 - different errors, but **without optimality**, which is important for building blocks

Generalized Piecewise-based Mechanism

- Most generalized version: m -piecewise distribution




$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_{1,\varepsilon} & \text{if } y \in [l_{1,x,\varepsilon}, r_{1,x,\varepsilon}] \\ p_{2,\varepsilon} & \text{if } y \in [l_{2,x,\varepsilon}, r_{2,x,\varepsilon}] \\ \dots & \dots \\ p_{m,\varepsilon} & \text{if } y \in [l_{m,x,\varepsilon}, r_{m,x,\varepsilon}] \end{cases}$$

$$\frac{p_{i,\varepsilon}}{p_{j,\varepsilon}} \leq e^\varepsilon \text{ (LDP constraint)}$$

Generalized Piecewise-based Mechanism

- Most generalized version: m -piecewise distribution



$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_{1,\varepsilon} & \text{if } y \in [l_{1,x,\varepsilon}, r_{1,x,\varepsilon}] \\ p_{2,\varepsilon} & \text{if } y \in [l_{2,x,\varepsilon}, r_{2,x,\varepsilon}] \\ \dots & \dots \\ p_{m,\varepsilon} & \text{if } y \in [l_{m,x,\varepsilon}, r_{m,x,\varepsilon}] \end{cases}$$

$$\frac{p_{i,\varepsilon}}{p_{j,\varepsilon}} \leq e^\varepsilon \text{ (LDP constraint)}$$

- Error (data utility):


$$\mathcal{L}(y, x)$$



$$\mathcal{L}(y, x) := |y - x|^p$$

Generalized Piecewise-based Mechanism

- Most generalized version: m -piecewise distribution




$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_{1,\varepsilon} & \text{if } y \in [l_{1,x,\varepsilon}, r_{1,x,\varepsilon}] \\ p_{2,\varepsilon} & \text{if } y \in [l_{2,x,\varepsilon}, r_{2,x,\varepsilon}] \\ \dots & \dots \\ p_{m,\varepsilon} & \text{if } y \in [l_{m,x,\varepsilon}, r_{m,x,\varepsilon}] \end{cases}$$

$$\frac{p_{i,\varepsilon}}{p_{j,\varepsilon}} \leq e^\varepsilon \text{ (LDP constraint)}$$

- Expected error:

$$\int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

- Most generalized version: m -piecewise distribution



$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_{1,\varepsilon} & \text{if } y \in [l_{1,x,\varepsilon}, r_{1,x,\varepsilon}] \\ p_{2,\varepsilon} & \text{if } y \in [l_{2,x,\varepsilon}, r_{2,x,\varepsilon}] \\ \dots & \dots \\ p_{m,\varepsilon} & \text{if } y \in [l_{m,x,\varepsilon}, r_{m,x,\varepsilon}] \end{cases}$$


$$\frac{p_{i,\varepsilon}}{p_{j,\varepsilon}} \leq e^\varepsilon \text{ (LDP constraint)}$$

- Expected error:

$$\min_{\mathcal{M}: p_i, l_i, r_i} \int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

Find \mathcal{M} to minimize the error at x

- Most generalized version: m -piecewise distribution



$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_{1,\varepsilon} & \text{if } y \in [l_{1,x,\varepsilon}, r_{1,x,\varepsilon}] \\ p_{2,\varepsilon} & \text{if } y \in [l_{2,x,\varepsilon}, r_{2,x,\varepsilon}] \\ \dots & \dots \\ p_{m,\varepsilon} & \text{if } y \in [l_{m,x,\varepsilon}, r_{m,x,\varepsilon}] \end{cases}$$


$$\frac{p_{i,\varepsilon}}{p_{j,\varepsilon}} \leq e^\varepsilon \text{ (LDP constraint)}$$

- Expected error:

$$\min_{\mathcal{M}: p_i, l_i, r_i} \max_{x \in \mathcal{D}} \int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

Find \mathcal{M} to minimize the worst-case error

- Most generalized version: m -piecewise distribution



$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_{1,\varepsilon} & \text{if } y \in [l_{1,x,\varepsilon}, r_{1,x,\varepsilon}] \\ p_{2,\varepsilon} & \text{if } y \in [l_{2,x,\varepsilon}, r_{2,x,\varepsilon}] \\ \dots & \dots \\ p_{m,\varepsilon} & \text{if } y \in [l_{m,x,\varepsilon}, r_{m,x,\varepsilon}] \end{cases}$$

$$\frac{p_{i,\varepsilon}}{p_{j,\varepsilon}} \leq e^\varepsilon \text{ (LDP constraint)}$$

Solved \mathcal{M} is
the optimal piecewise-based mechanism

Mathematically \equiv to find the optimal
piecewise distribution under the LDP constraint

$$\min_{\mathcal{M}: p_i, l_i, r_i} \max_{x \in \mathcal{D}} \int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

Find \mathcal{M} to minimize the worst-case error

- Challenges
 1. min-max problem & multiple variables
 2. optimal results only for a specific m

$$\min_{\mathcal{M}: p_i, l_i, r_i} \max_{x \in \mathcal{D}} \int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

↑
2. $i \in [m]$

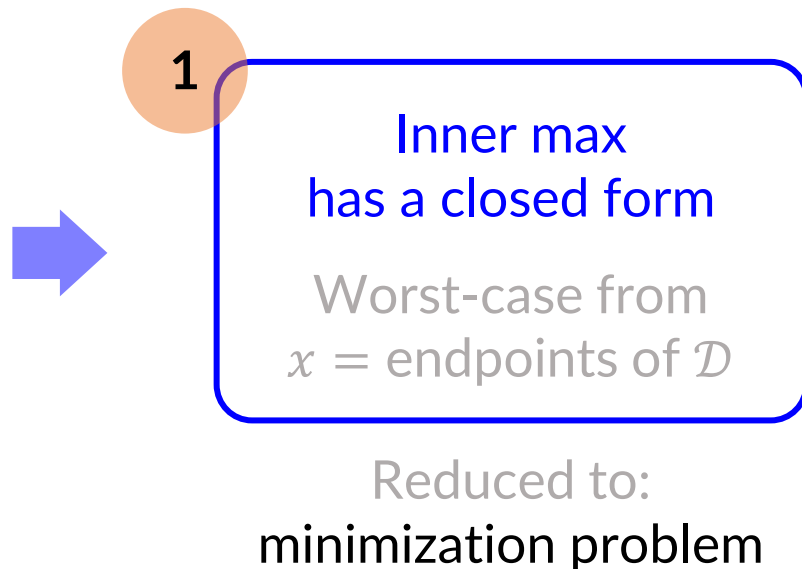
↑
1. variables p_i, l_i, r_i

- Challenges
 - min-max problem & multiple variables
 - optimal results only for a specific m

$$\min_{\mathcal{M}: p_i, l_i, r_i} \max_{x \in \mathcal{D}} \int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

↑
2. $i \in [m]$

↑
1. variables p_i, l_i, r_i



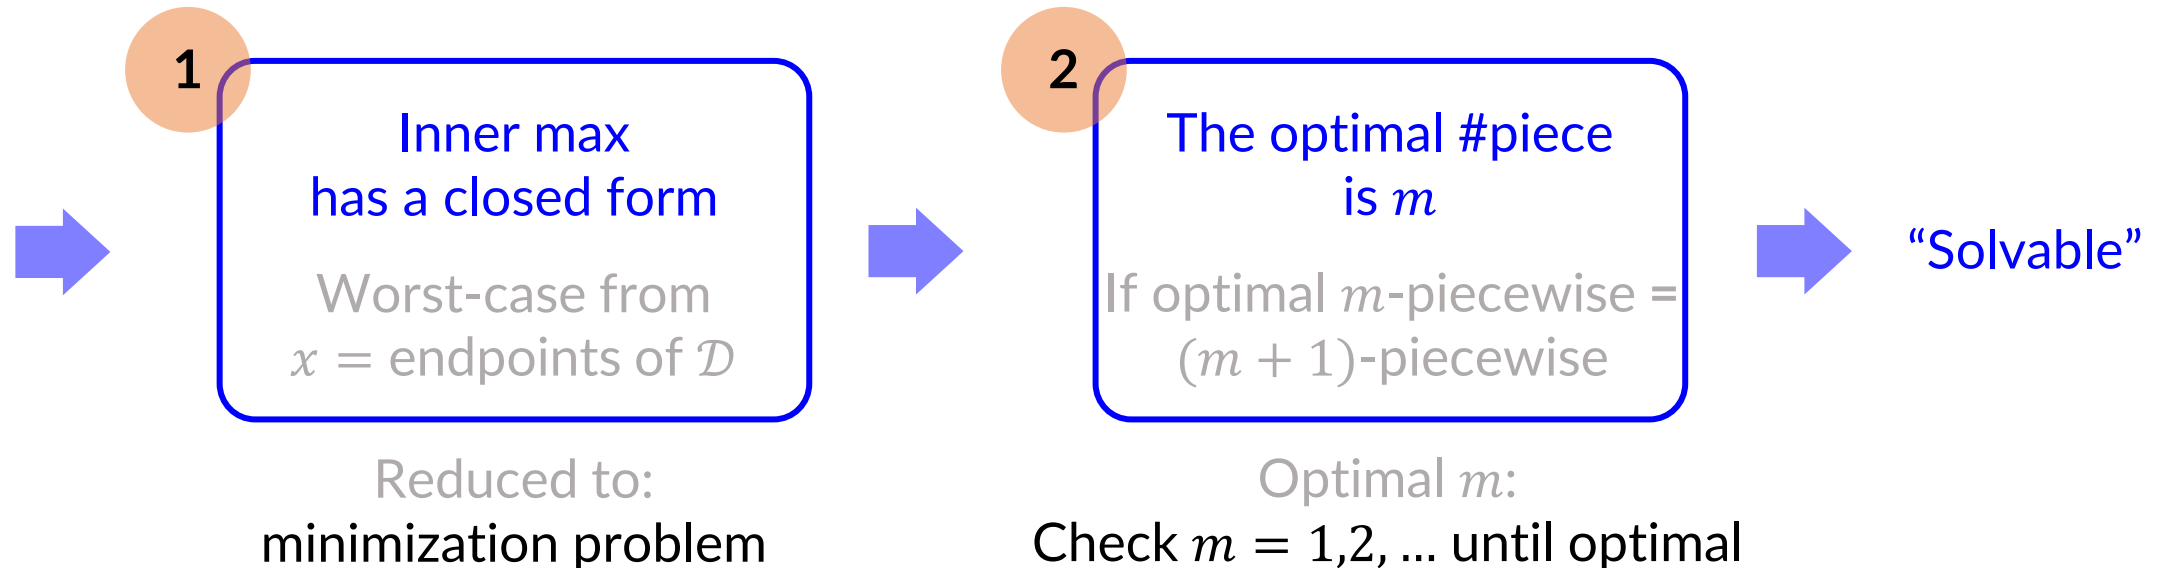
- Challenges

1. min-max problem & multiple variables
2. optimal results only for a specific m

$$\min_{\mathcal{M}: p_i, l_i, r_i} \max_{x \in \mathcal{D}} \int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

↑
2. $i \in [m]$

↑
1. variables p_i, l_i, r_i

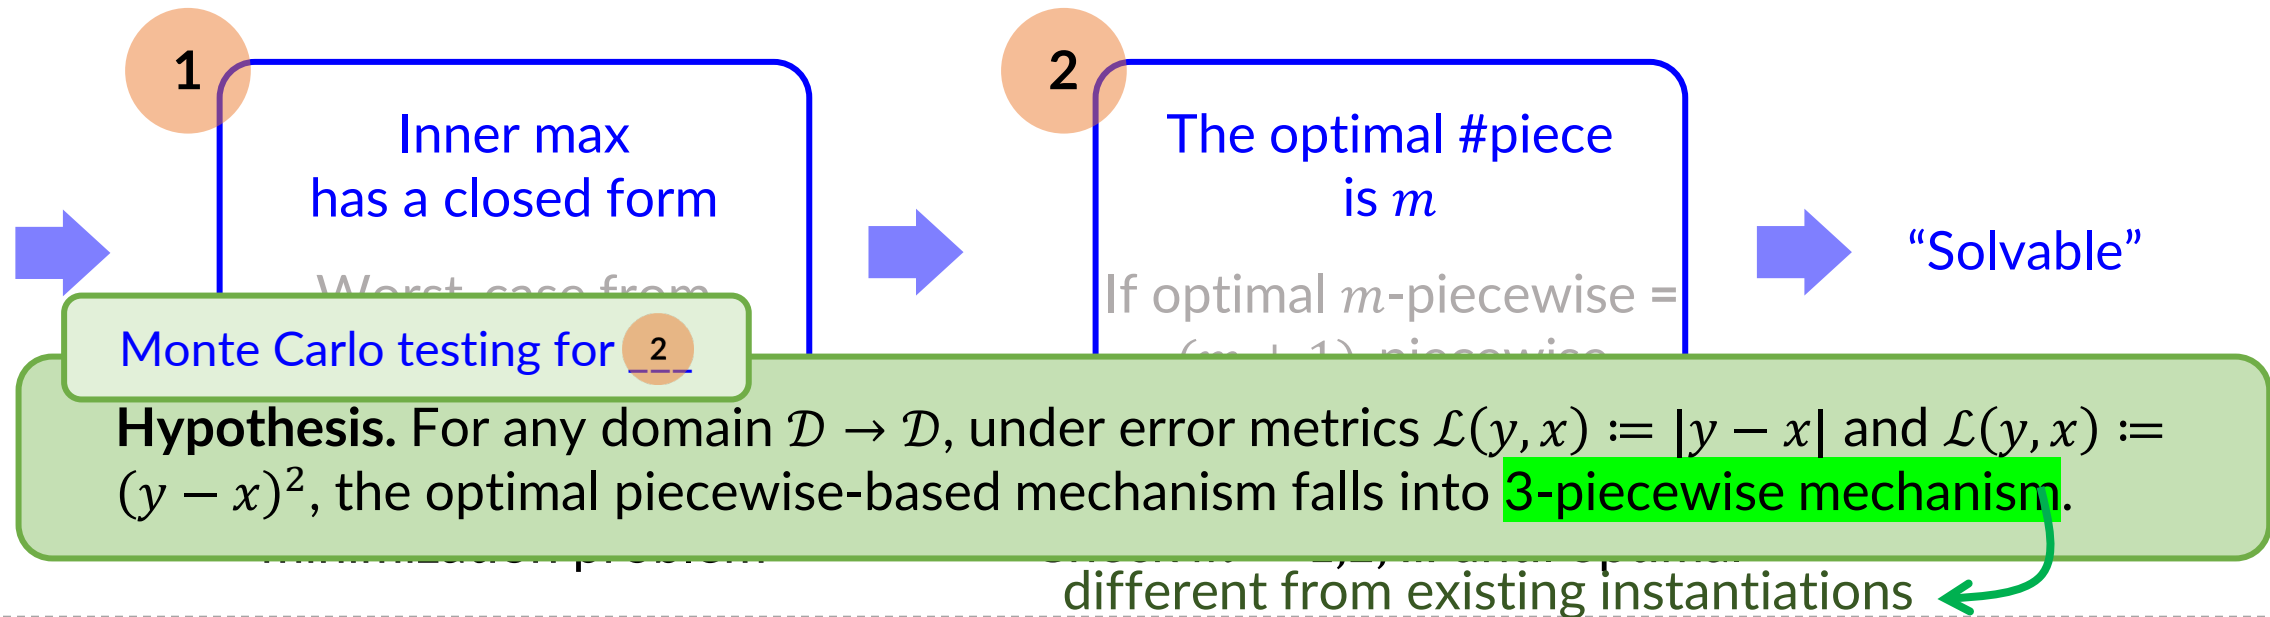


Challenges

1. min-max problem & multiple variables
2. optimal results only for a specific m

$$\min_{\mathcal{M}: p_i, l_i, r_i} \max_{x \in \mathcal{D}} \int_{\tilde{\mathcal{D}}} \mathcal{L}(y, x) \cdot pdf[\mathcal{M}(x) = y] dy$$

2. $i \in [m]$
1. variables p_i, l_i, r_i




- Optimal $\mathcal{M}: [0,1) \rightarrow [0,1)$ under $\mathcal{L} := |y - x|$

$$pdf[\mathcal{M}(x) = y] = \begin{cases} \exp\left(\frac{\varepsilon}{2}\right) & \text{if } y \in [l_{x,\varepsilon}, r_{x,\varepsilon}) \\ \exp\left(-\frac{\varepsilon}{2}\right) & \text{if } y \in [0,1) \setminus [l_{x,\varepsilon}, r_{x,\varepsilon}) \end{cases}$$


$$[l_{x,\varepsilon}, r_{x,\varepsilon}) = \begin{cases} [0, 2C) & \text{if } x \in [0, C) \\ x + [-C, C) & \text{if } x \in [C, 1 - C) \\ [1 - 2C, 1) & \text{otherwise} \end{cases}$$

$$C = \frac{\exp(\varepsilon/2) - 1}{2(\exp(\varepsilon) - 1)}$$

- When $\varepsilon = 1$:



(a) $x = 0$.




(b) $x = 0.5$.

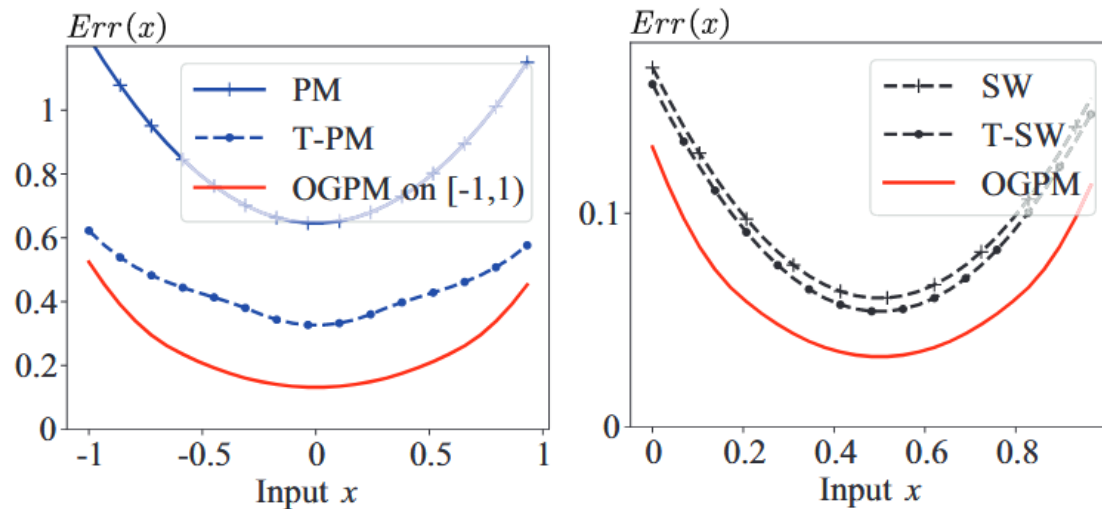
Comparison of Expected Errors

Whole-domain error (i.e. each point in \mathcal{D}) ($\epsilon = 2$)

Compressed PM, SW



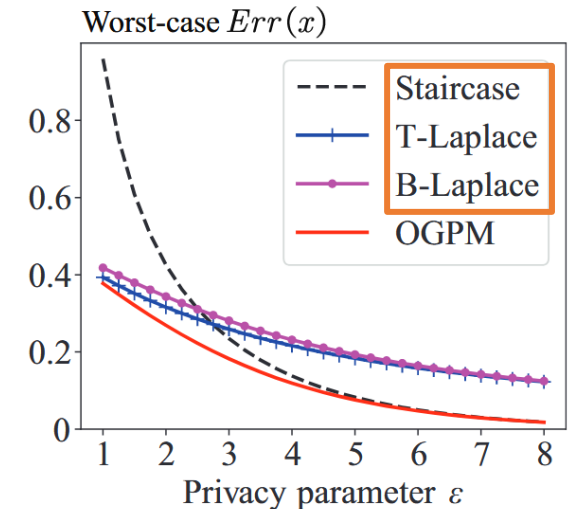
Original and truncated PM, SW



Lowest error

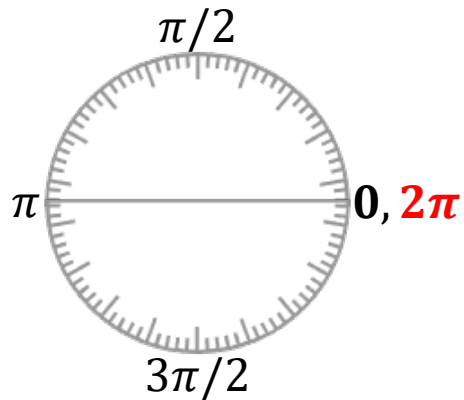
Worst-case error

Non-piecewise-based



Lowest error


- Different meaning of distance, e.g. $\text{distance}(0, 2\pi) = 0 \rightarrow$ requires different mechanisms



$$\mathcal{L} \rightarrow \mathcal{L}_{\text{mod}}$$

$$\mathcal{L}_{\text{mod}}(y, x) := \min(\mathcal{L}(y, x), \mathcal{L}(y, 2\pi - x))$$

- Different meaning of distance, e.g. $\text{distance}(0, 2\pi) = 0 \rightarrow$ requires different mechanisms




$$\mathcal{L} \rightarrow \mathcal{L}_{\text{mod}}$$

$$\mathcal{L}_{\text{mod}}(y, x) := \min(\mathcal{L}(y, x), \mathcal{L}(y, 2\pi - x))$$

$$\Rightarrow \min_{\mathcal{M}: p_i, l_i, r_i} \max_{x \in [0, 2\pi]} \int_{\tilde{\mathcal{D}}} \mathcal{L}_{\text{mod}}(y, x) \cdot \text{pdf}[\mathcal{M}(x) = y] dy$$

- Linking** to problems in the classical domain



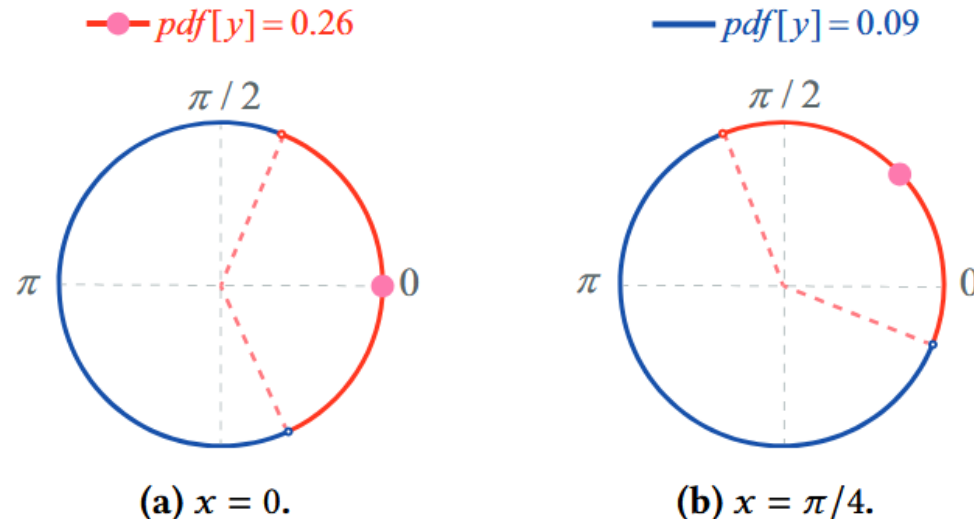
- Optimal $\mathcal{M}: [0, 2\pi) \rightarrow [0, 2\pi)$ under \mathcal{L}_{mod}

$$pdf[\mathcal{M}(x) = y] = \begin{cases} \frac{1}{2\pi} \exp\left(\frac{\varepsilon}{2}\right) & \text{if } y \in [l_{x,\varepsilon}^{\text{mod}}, r_{x,\varepsilon}^{\text{mod}}) \\ \frac{1}{2\pi} \exp\left(-\frac{\varepsilon}{2}\right) & \text{if } y \in [0, 2\pi) \setminus [l_{x,\varepsilon}^{\text{mod}}, r_{x,\varepsilon}^{\text{mod}}) \end{cases}$$

$$[l_{x,\varepsilon}^{\text{mod}}, r_{x,\varepsilon}^{\text{mod}}) = [x - C, x + C) \text{ mod } 2\pi$$


$$C = \pi \frac{\exp(\varepsilon/2) - 1}{\exp(\varepsilon) - 1}$$

- When $\varepsilon = 1$:




Whole-domain error ($\epsilon = 2$)

PM, SW on the **flattened** domain



Worst-case error


Purkayastha mechanism [CCS'21]*



Lowest error

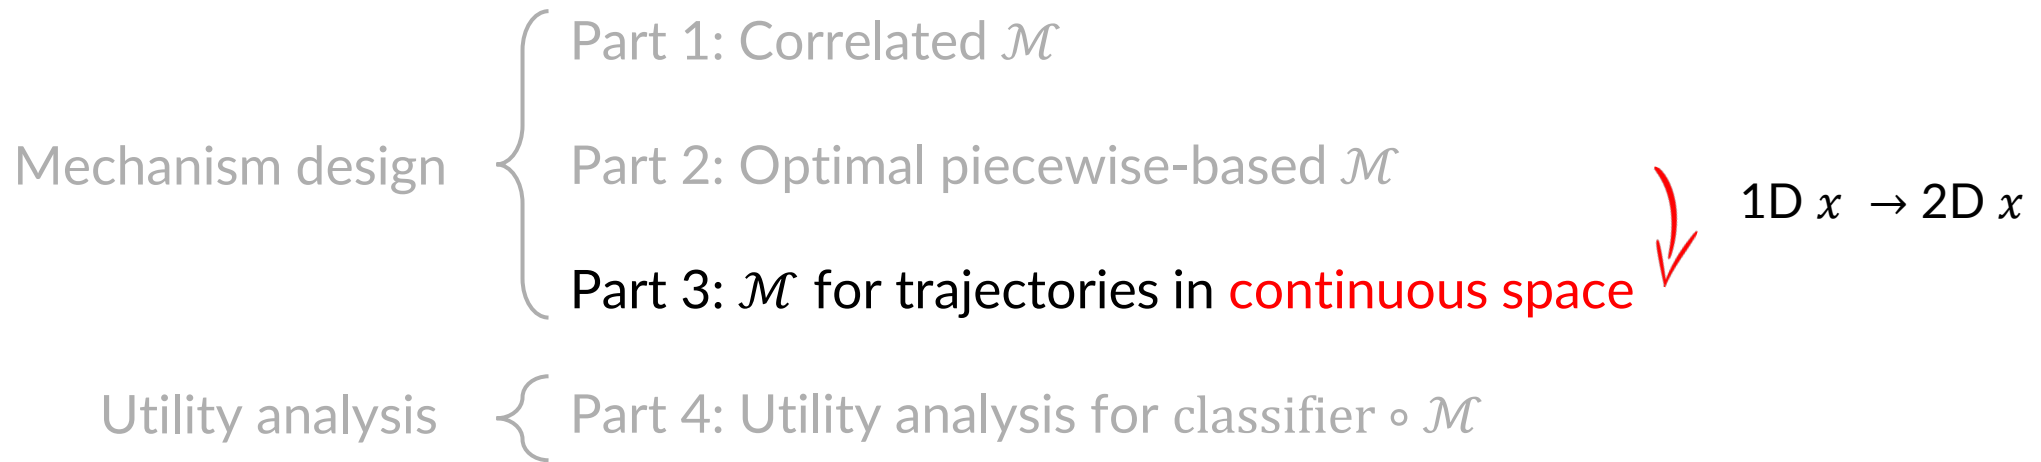

* Differential Privacy for Directional Data, CCS'21

- RQ: What is the optimal piecewise-based mechanism?
- Contributions:^{*}
 - solving framework for the optimality
 - closed-form mechanisms for classical domains & circular domains
 - comparison with non-piecewise-based mechanisms

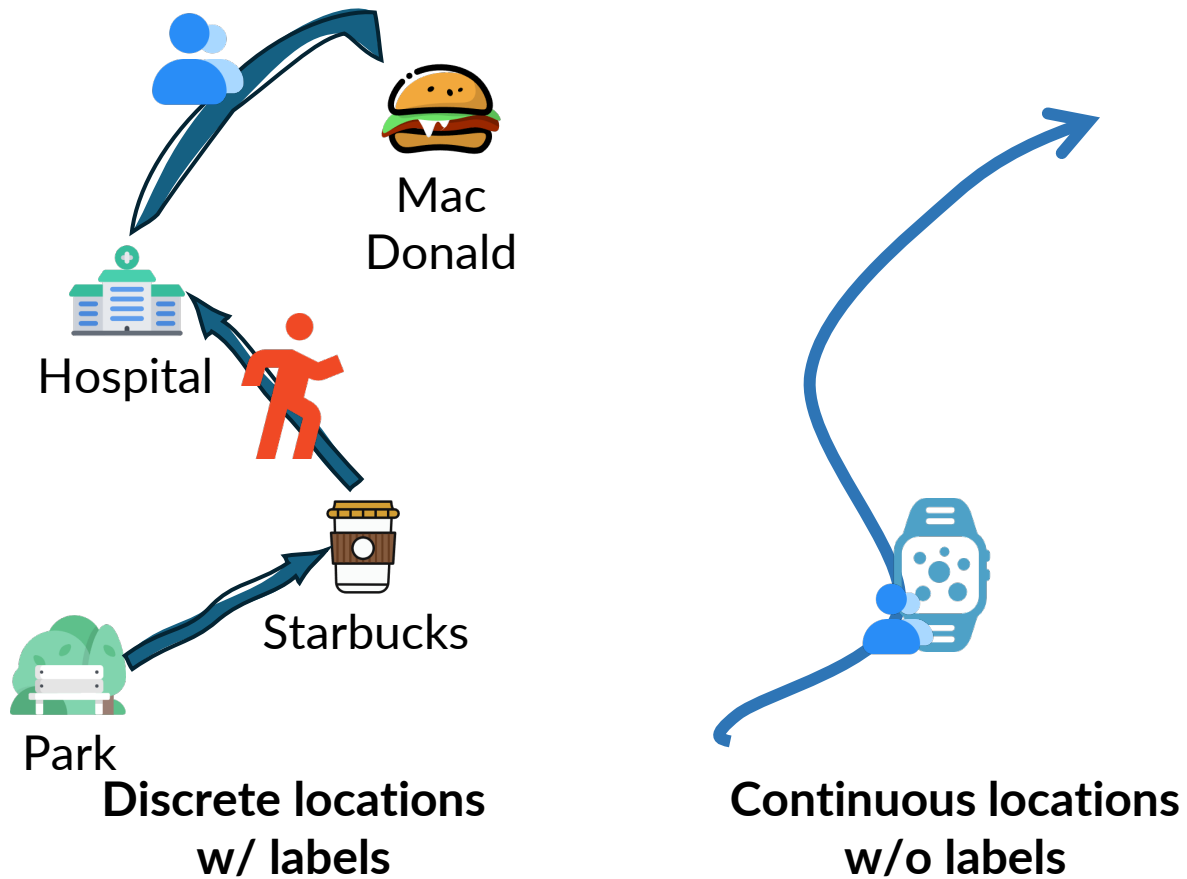


^{*} Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy, PETS'25

Local Differential Privacy: Refined Mechanism Design and Utility Analysis

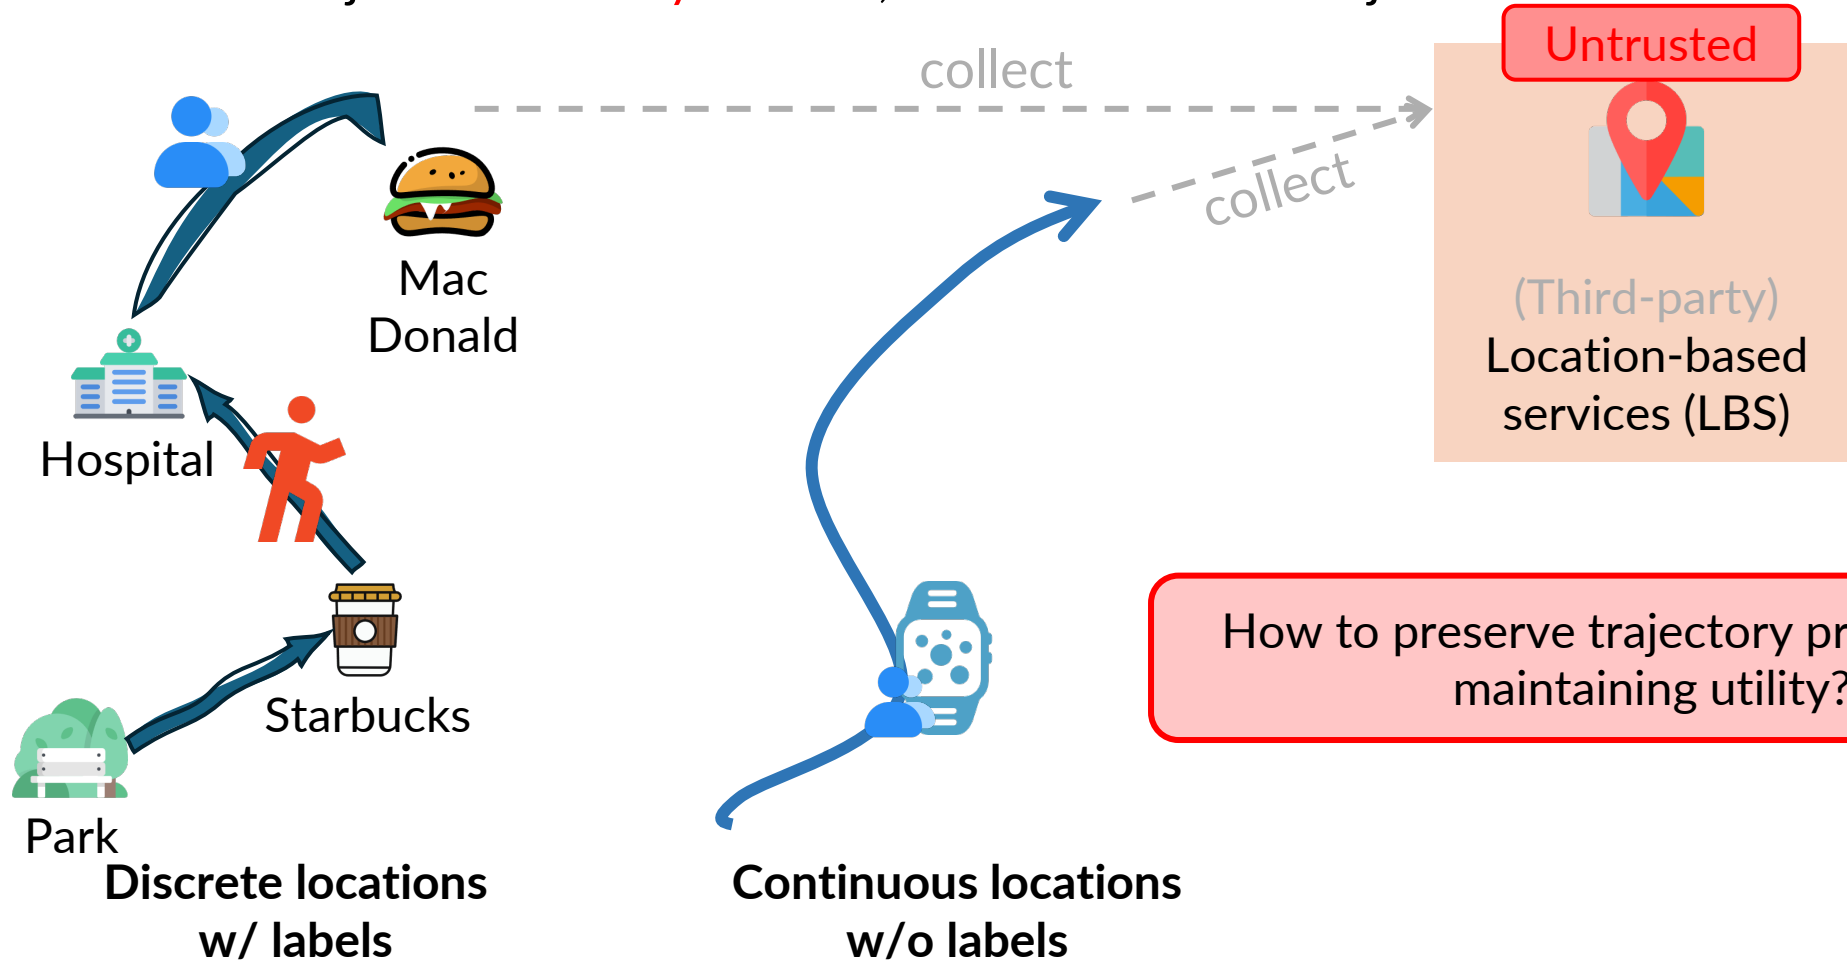


- Sensitive trajectories: **daily routine**, **wearable-sensor** trajectories



Trajectory Collection

- Sensitive trajectories: **daily routine**, **wearable-sensor** trajectories



- LDP-fy: perturb a trajectory with LDP guarantee (provable privacy)
 - cannot distinguish location τ_1 from τ_2 with confidence quantified by e^ϵ

$$\forall \tau_1, \tau_2, \tau_* \in \mathcal{S}: \frac{\Pr[\mathcal{M}(\tau_1) = \tau_*]}{\Pr[\mathcal{M}(\tau_2) = \tau_*]} \leq e^\epsilon$$

location space

distinguishability of τ_1 from τ_2 when observing τ_*

- LDP-fy: perturb a trajectory with LDP guarantee (provable privacy)
 - cannot distinguish location τ_1 from τ_2 with confidence quantified by e^ϵ

$$\forall \tau_1, \tau_2, \tau_* \in \mathcal{S}: \frac{\Pr[\mathcal{M}(\tau_1) = \tau_*]}{\Pr[\mathcal{M}(\tau_2) = \tau_*]} \leq e^\epsilon$$


location space

distinguishability of τ_1 from τ_2 when observing τ_*


- Target: LDP mechanism \mathcal{M} (for the location space)
 - provable privacy for users' trajectories, in **continuous spaces** and discrete spaces
 - as high trajectory utility as possible

LDP-fy a Trajectory in A Discrete Space

- Discrete location spaces: **Point of interests** or **discretized cells**




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



$$\mathcal{S} = \{c_1, \dots, c_n\}$$

LDP-fy a Trajectory in A Discrete Space

- Discrete location spaces: **Point of interests** or **discretized cells**



- Exponential mechanism:


$$\Pr[\mathcal{M}_{\text{exp}}(\tau) = \tau^*] = \frac{\exp(\varepsilon d(\tau, \tau^*))}{\sum_{\tau' \in \mathcal{S}} \exp(\varepsilon d(\tau, \tau'))}$$

← Pairwise distance

← Sum of distance

LDP-fy a Trajectory in A Discrete Space

- Discrete location spaces: **Point of interests** or **discretized cells**




- Exponential mechanism:

$$\Pr[\mathcal{M}_{\text{exp}}(\tau) = \tau^*] = \frac{\exp(\varepsilon d(\tau, \tau^*))}{\sum_{\tau' \in \mathcal{S}} \exp(\varepsilon d(\tau, \tau'))}$$

← Pairwise distance
← Sum of distance

LDP-fy a Trajectory in A Discrete Space

- Discrete location spaces: **Point of interests** or **discretized c**



- Exponential mechanism:

$$\Pr[\mathcal{M}_{\text{exp}}(\tau) = \tau^*] = \frac{\exp(\varepsilon d(\tau, \tau^*))}{\sum_{\tau' \in \mathcal{S}} \exp(\varepsilon d(\tau, \tau'))}$$

← Pairwise distance
← Sum of distance

Limitations

1. Efficiency


- each sample costs $\mathcal{O}(n)$

2. Trajectory Utility

- decreases as n increases


3. Applicability

- cannot apply to continuous \mathcal{S}




Continuous Spaces: Better & Universal

- Operate on a continuous space (covering the discrete space)



Continuous Spaces: Better & Universal

- Operate on a continuous space (covering the discrete space)




Simple sampling: (satisfying LDP for \mathcal{S})

$$\Pr[\mathcal{M}(x_i) = \tilde{x} \in \text{red square}] = p_{\text{high}}$$

$$\Pr[\mathcal{M}(x_i) = \tilde{x} \in \text{gray square}] = p_{\text{low}}$$

Continuous Spaces: Better & Universal

- Operate on a continuous space (covering the discrete space)



Simple sampling: (satisfying LDP for \mathcal{S})


$$\Pr[\mathcal{M}(x_i) = \tilde{x} \in \text{red square}] = p_{\text{high}}$$

$$\Pr[\mathcal{M}(x_i) = \tilde{x} \in \text{gray square}] = p_{\text{low}}$$

- Benefits:
 - Efficiency:** $\mathcal{O}(1)$ sampling complexity
 - Trajectory utility:** “ n -independent”
 - Applicability:** Both continuous spaces & discrete spaces

Continuous Spaces: Better & Universal

- Operate on a continuous space (covering the discrete space)



Simple sampling: (satisfying LDP for \mathcal{S})


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

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

- Benefits:
 - Efficiency:** $\mathcal{O}(1)$ sampling complexity
 - Trajectory utility:** “ n -independent”
 - Applicability:** Both continuous spaces & discrete spaces

Continuous Spaces: Better & Universal

- Operate on a continuous space (covering the discrete space)



Simple sampling: (satisfying LDP for \mathcal{S})

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

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

- Benefits:

1. **Efficiency:** $\mathcal{O}(1)$ sampling complexity

2. **Trajectory utility:** “ n -independent”


3. **Applicability:** Both continuous spaces & discrete spaces

Rounding-to-the-nearest

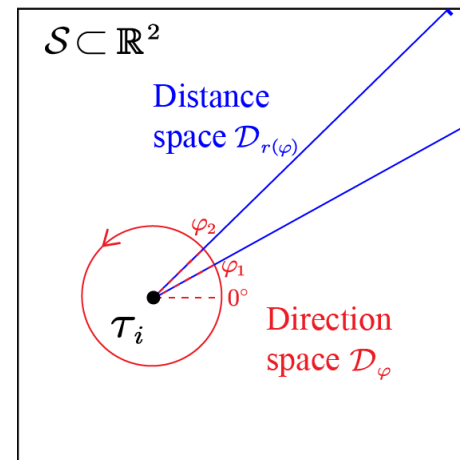
(post-processing)

TraCS: Trajectory Collection in Continuous Spaces

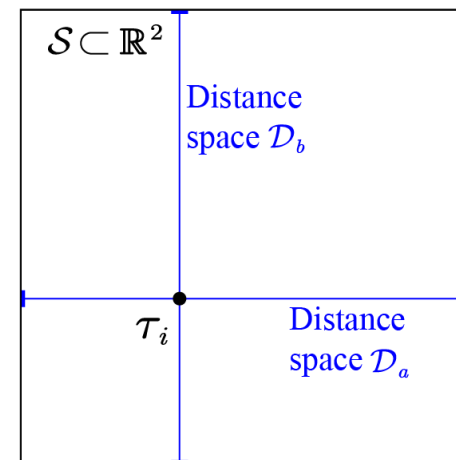
- TraCS-D: **direction-distance** perturbation
- TraCS-C: **coordinates** perturbation
- **Key idea:** decomposes \mathcal{S} into two subspaces



TraCS
decomp.
at \mathcal{T}_i




TraCS-D: $\mathcal{D}_\varphi \otimes \mathcal{D}_{r(\varphi)}$



TraCS-C: $\mathcal{D}_a \times \mathcal{D}_b$

TraCS: Trajectory Collection in Continuous Spaces

- TraCS-D: **direction-distance** perturbation
- TraCS-C: **coordinates** perturbation
- **Key idea:** decomposes \mathcal{S} into two subspaces \rightarrow design \mathcal{M} for **each subspace**



- Leverage piecewise-based mechanisms \mathcal{M}_0 and \mathcal{M}_-

$$\begin{array}{ccc} & \nearrow & \nwarrow \\ [0,2\pi) \rightarrow [0,2\pi) & & [0,1) \rightarrow [0,1) \end{array}$$


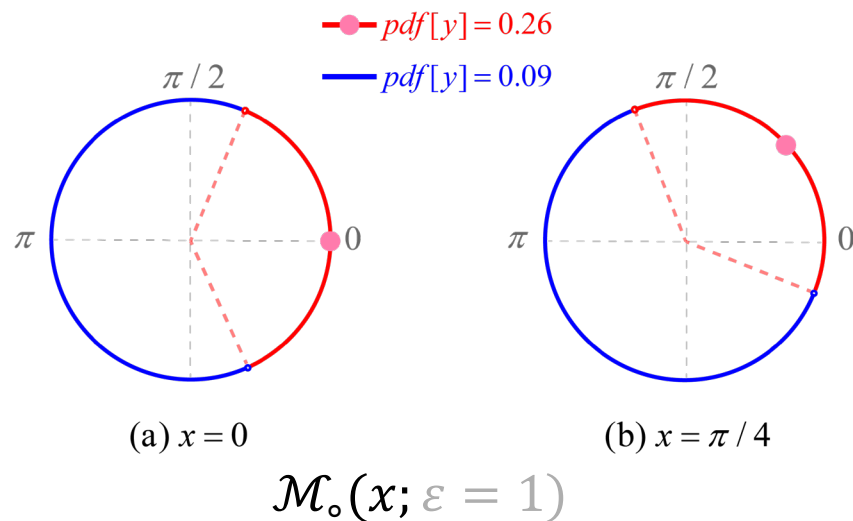
- other LDP mechanisms for bounded numerical domain also applicable

- Leverage piecewise-based mechanisms \mathcal{M}_0 and \mathcal{M}_-

$[0, 2\pi) \rightarrow [0, 2\pi)$ $[0, 1) \rightarrow [0, 1)$

- other LDP mechanisms for bounded numerical domain also applicable

- Examples of \mathcal{M}_0 and \mathcal{M}_-^*



* Optimal Piecewise-based Mechanism for Collecting Bounded Numerical Data under Local Differential Privacy, PETS'25

- Claims

- continuous spaces: better **utility** than naïve baselines, e.g. planar Laplace + truncation
- discrete spaces: better **efficiency and utility** than existing methods, ATP*, NGram**, L-SRR***

* Trajectory Data Collection with Local Differential Privacy, VLDB'23

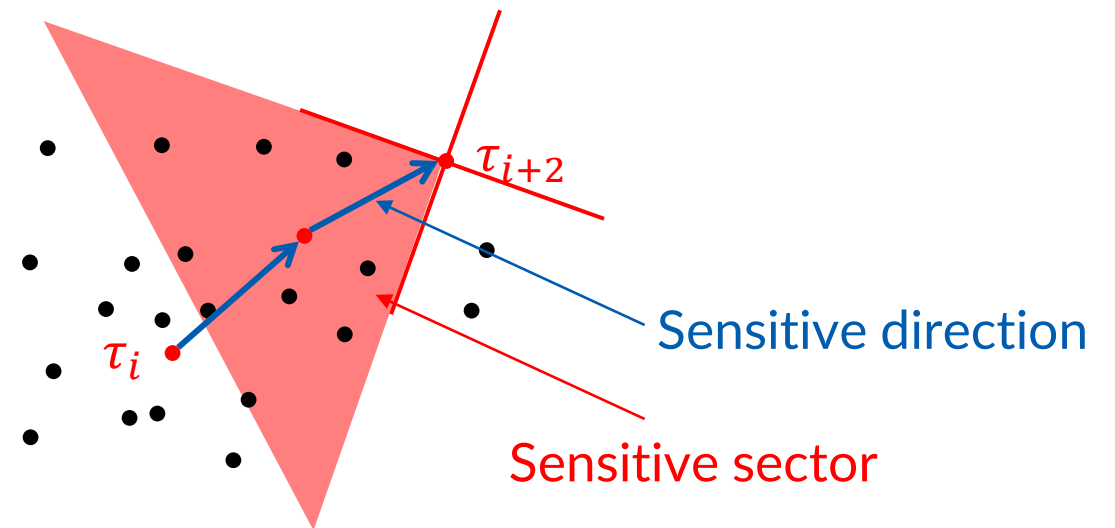
** Real-World Trajectory Sharing with Local Differential Privacy, VLDB'21

*** L-SRR: Local Differential Privacy for Location-Based Services with Staircase Randomized Response, CCS'22

- Claims
 - continuous spaces: better utility than naïve baselines, e.g. planar Laplace + truncation
 - discrete spaces: better efficiency and utility than existing methods, ATP*, NGram**, L-SRR***

- **ATP: direction perturbation**

1. divide direction sectors, e.g. $k = 4$



* Trajectory Data Collection with Local Differential Privacy, VLDB'23

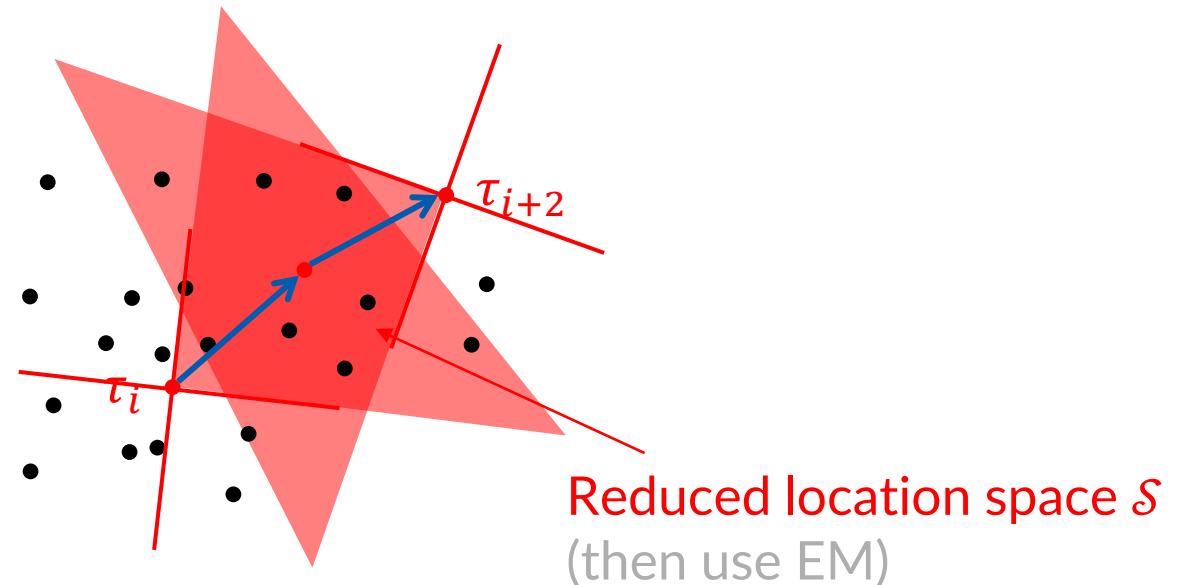
** Real-World Trajectory Sharing with Local Differential Privacy, VLDB'21

*** L-SRR: Local Differential Privacy for Location-Based Services with Staircase Randomized Response, CCS'22

- Claims
 - continuous spaces: better utility than naïve baselines, e.g. planar Laplace + truncation
 - discrete spaces: better efficiency and utility than existing methods, ATP*, NGram**, L-SRR***

- ATP: direction perturbation**

1. divide direction sectors, e.g. $k = 4$
2. perturb sector
3. $\mathcal{M}_{\text{exp}}(\tau)$ on reduced \mathcal{S}




* Trajectory Data Collection with Local Differential Privacy, VLDB'23

** Real-World Trajectory Sharing with Local Differential Privacy, VLDB'21

*** L-SRR: Local Differential Privacy for Location-Based Services with Staircase Randomized Response, CCS'22

- Claims
 - continuous spaces: better utility than naïve baselines, e.g. planar Laplace + truncation
 - discrete spaces: better efficiency and utility than existing methods, ATP*, NGram**, L-SRR***
- **NGram**: reachability constraint from public knowledge
 1. distance reachability, i.e. the next location cannot be too far
 2. $\mathcal{M}_{\text{exp}}(\tau)$ on reduced \mathcal{S}




* Trajectory Data Collection with Local Differential Privacy, VLDB'23

** Real-World Trajectory Sharing with Local Differential Privacy, VLDB'21

*** L-SRR: Local Differential Privacy for Location-Based Services with Staircase Randomized Response, CCS'22

- Claims
 - continuous spaces: better utility than naïve baselines, e.g. planar Laplace + truncation
 - discrete spaces: better efficiency and utility than existing methods, ATP*, NGram**, L-SRR***
- **L-SRR: Staircase Randomized Response**
 1. group the locations to k groups
 2. $\mathcal{M}_{\text{SRR}}(\tau)$ on k groups




* Trajectory Data Collection with Local Differential Privacy, VLDB'23

** Real-World Trajectory Sharing with Local Differential Privacy, VLDB'21


*** L-SRR: Local Differential Privacy for Location-Based Services with Staircase Randomized Response, CCS'22

Evaluations – Continuous Spaces

- Trajectory utility metric: Average error (AE)
 - T-Laplace: Truncated Laplace; Strawman: k -RR + uniform sampling for direction perturbation (similar to SRR)



(a) $S = [0, 1) \times [0, 1)$




(b) $S = [0, 2) \times [0, 10)$


Lower error when ϵ becomes larger

Evaluations – Discrete Spaces


- TKY and CHI datasets
 - TraCS for the rectangular area + rounding-to-the-nearest




(a) TKY location space



(b) CHI location space



(a) TKY dataset





(b) CHI dataset

Lower error when ϵ becomes larger

Evaluations – Discrete Spaces


- TKY and CHI datasets
 - TraCS for the rectangular area + rounding



(a) TKY dataset

Table 3: Time cost comparison (in milliseconds). averaged


	ATP	NGram	L-SRR	TraCS-D	TraCS-C
Total	145.7	100.9	6.2	0.06	0.05
Perturb	125.8	92.8	0.003	0.018	0.003



(b) CHI dataset

Negligible time cost

- RQ: Trajectory collection under LDP
- Our results:
 - **operating in continuous spaces can do better**
 - better efficiency: $\mathcal{O}(1)$ sampling complexity
 - better trajectory utility, especially for larger ϵ
 - better applicability, for both continuous and discrete \mathcal{S}




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

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

* TraCS: Trajectory Collection in Continuous Space under Local Differential Privacy, PETS'26

Local Differential Privacy: Refined Mechanism Design and Utility Analysis




- Mechanism design {
- Part 1: Correlated \mathcal{M}
 - Part 2: Optimal piecewise-based \mathcal{M}
 - Part 3: \mathcal{M} for trajectories in continuous space
- Utility analysis {
- Part 4: Utility analysis for **classifier** $\circ \mathcal{M}$



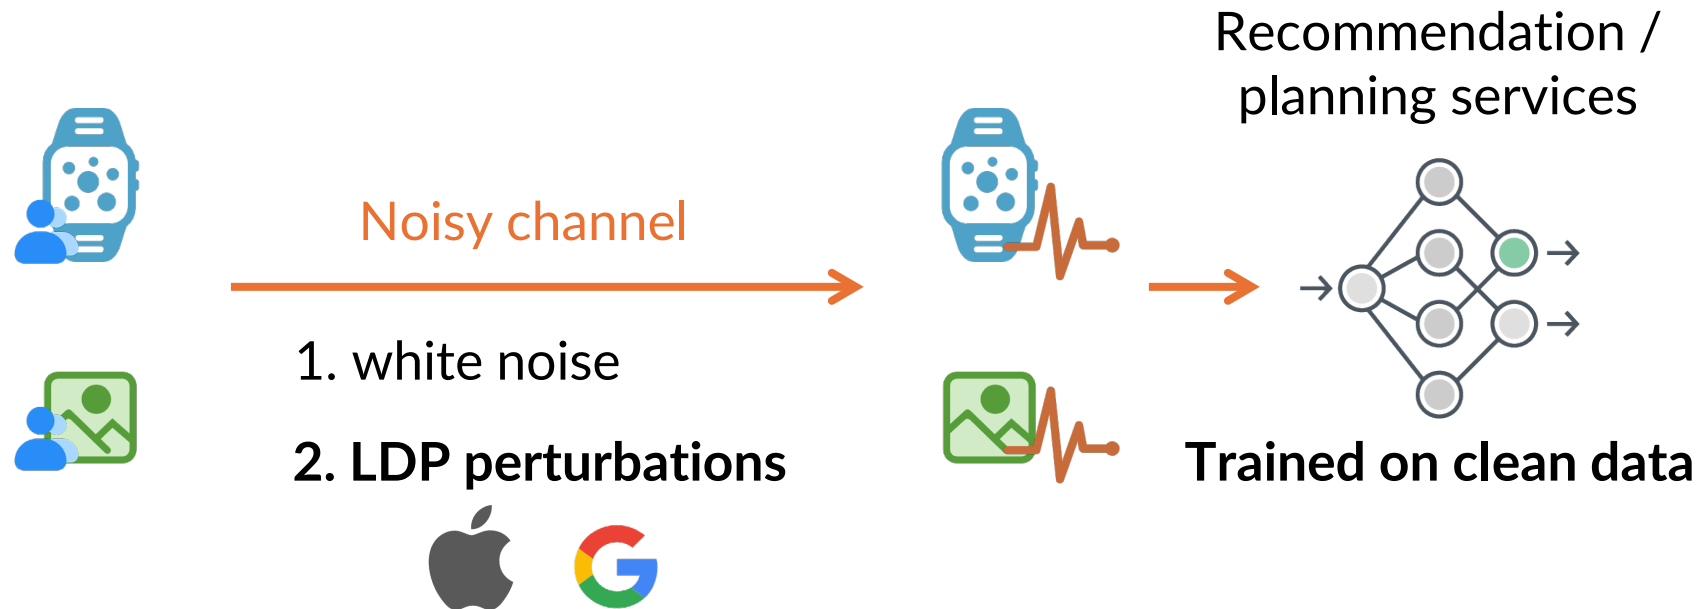
Mechanism-level utility
↓
Task-level utility

- Input data for classifiers may be noisy




Classifier Utility under Noisy Inputs

- Input data for classifiers may be noisy



Classifier Utility under Noisy Inputs

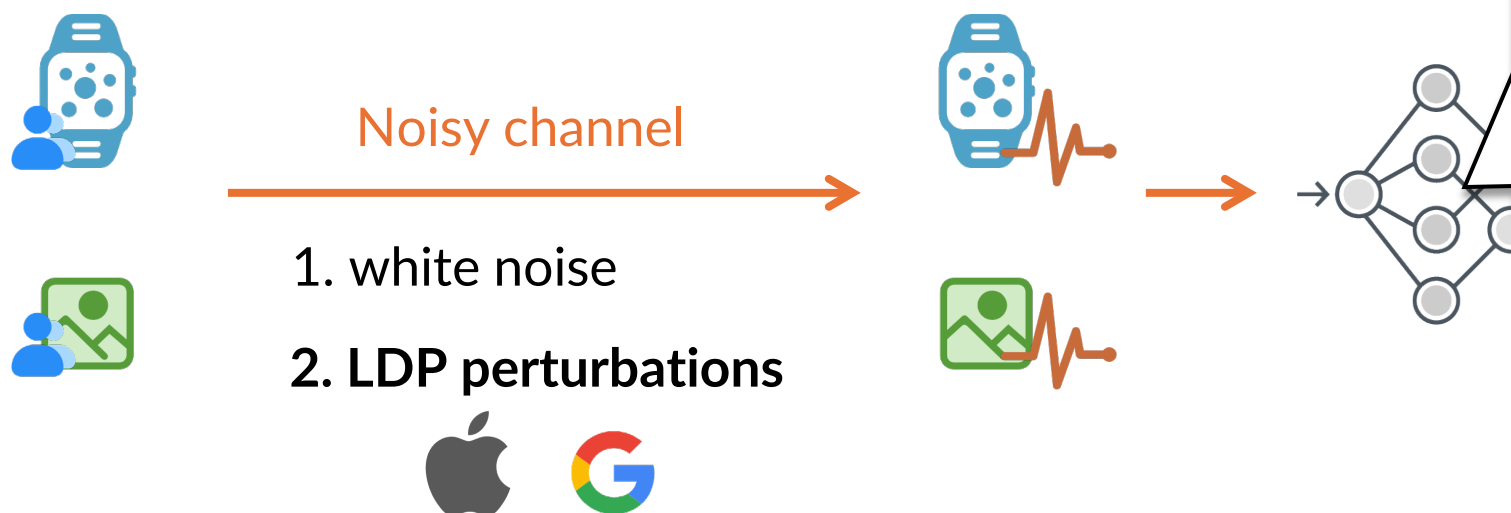
- Input data for classifiers may be noisy



“
What’s the utility of the trained classifier *under a noisy channel, e.g. LDP perturbation?*
”

Classifier Utility under Noisy Inputs

- Input data for classifiers may be noisy



“
What’s the utility of the trained classifier *under a noisy channel, e.g. LDP perturbation?*
”

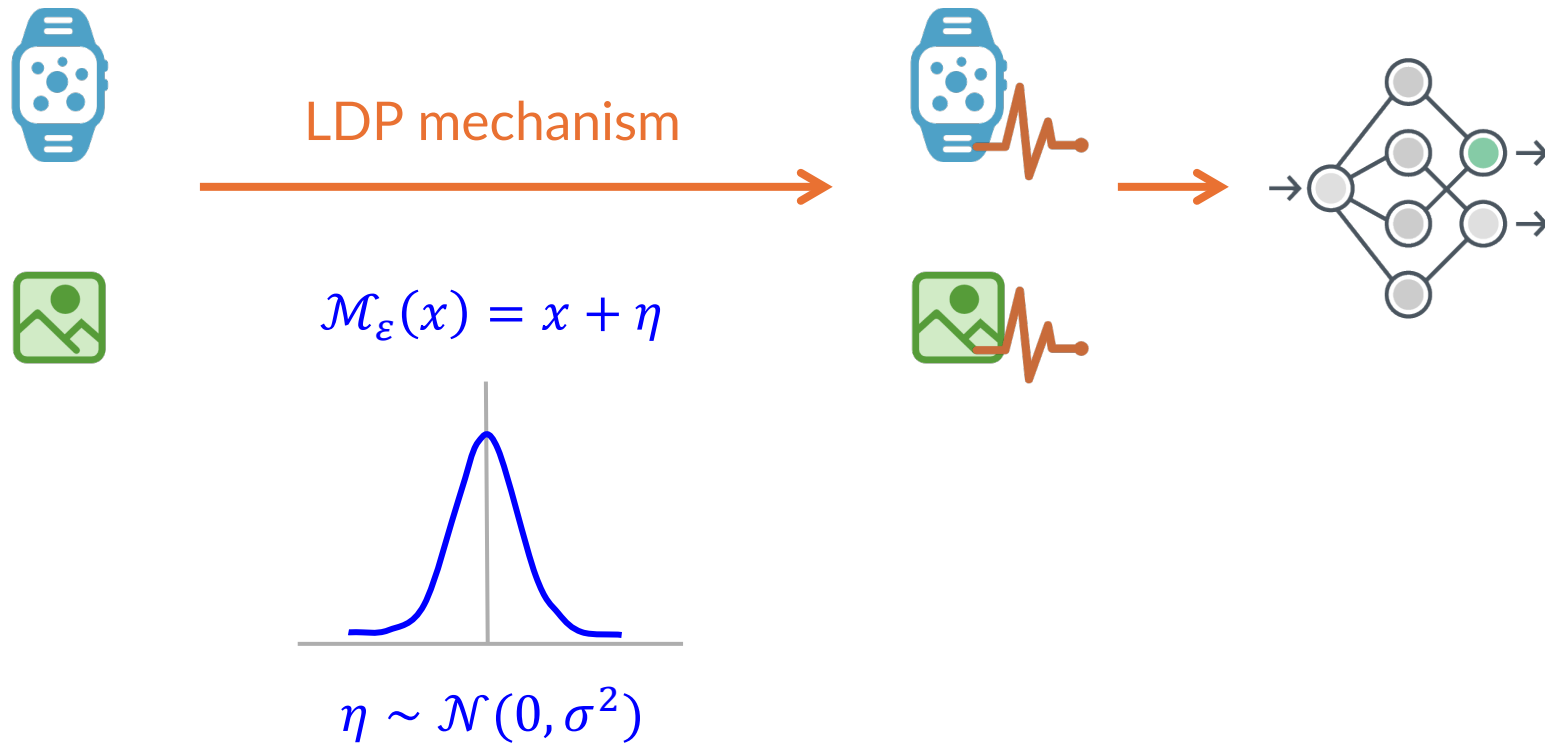
- Q: How can classifier designers/users know the classifier’s accuracy under LDP-perturbed data?

For an LDP-friendly classifier

For a better privacy-utility balance

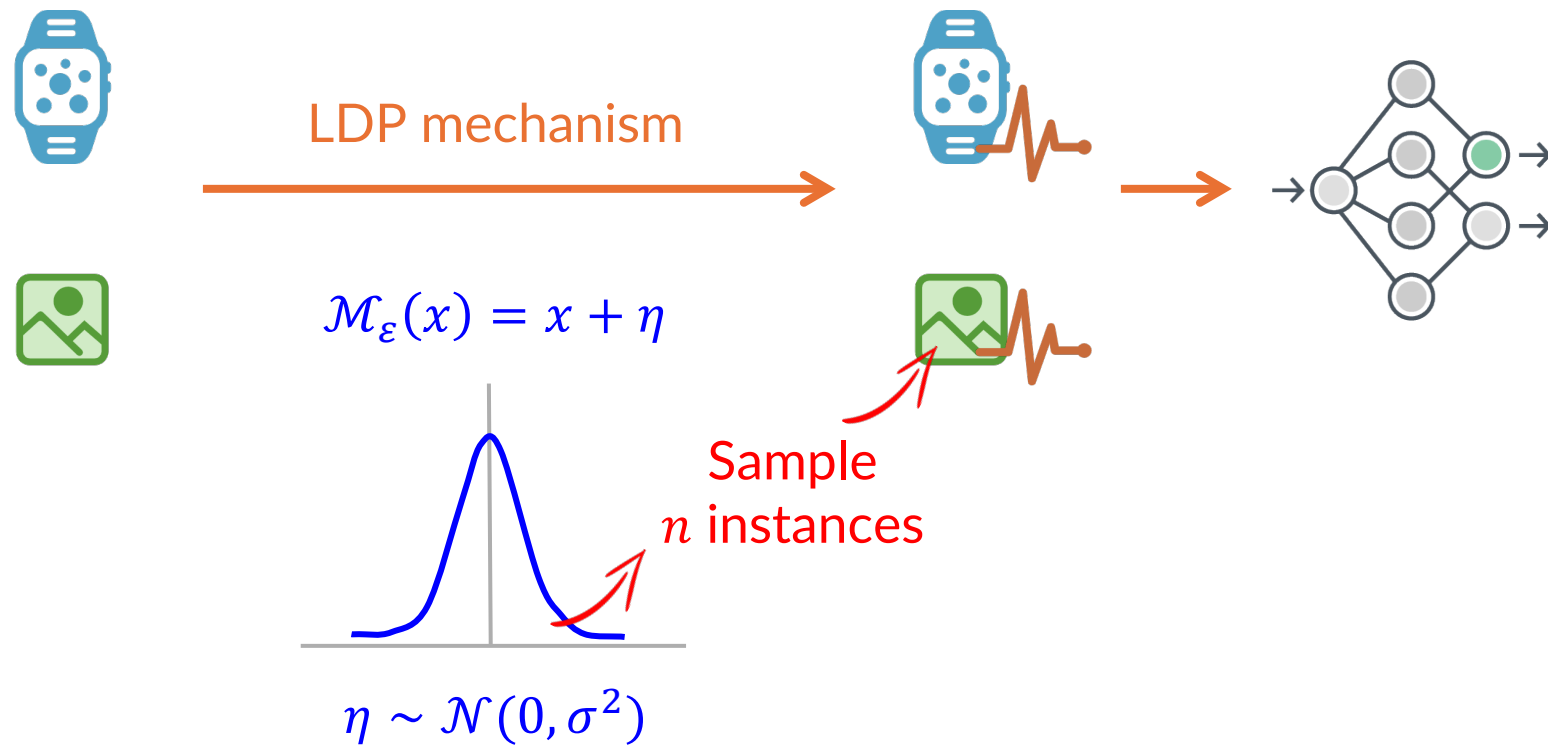
Empirical Classifier Utility under LDP-Data

- Empirical approach: Sample and then test
 - \mathcal{M} 's variance or MSE doesn't help – cannot provide a classifier accuracy



Empirical Classifier Utility under LDP-Data

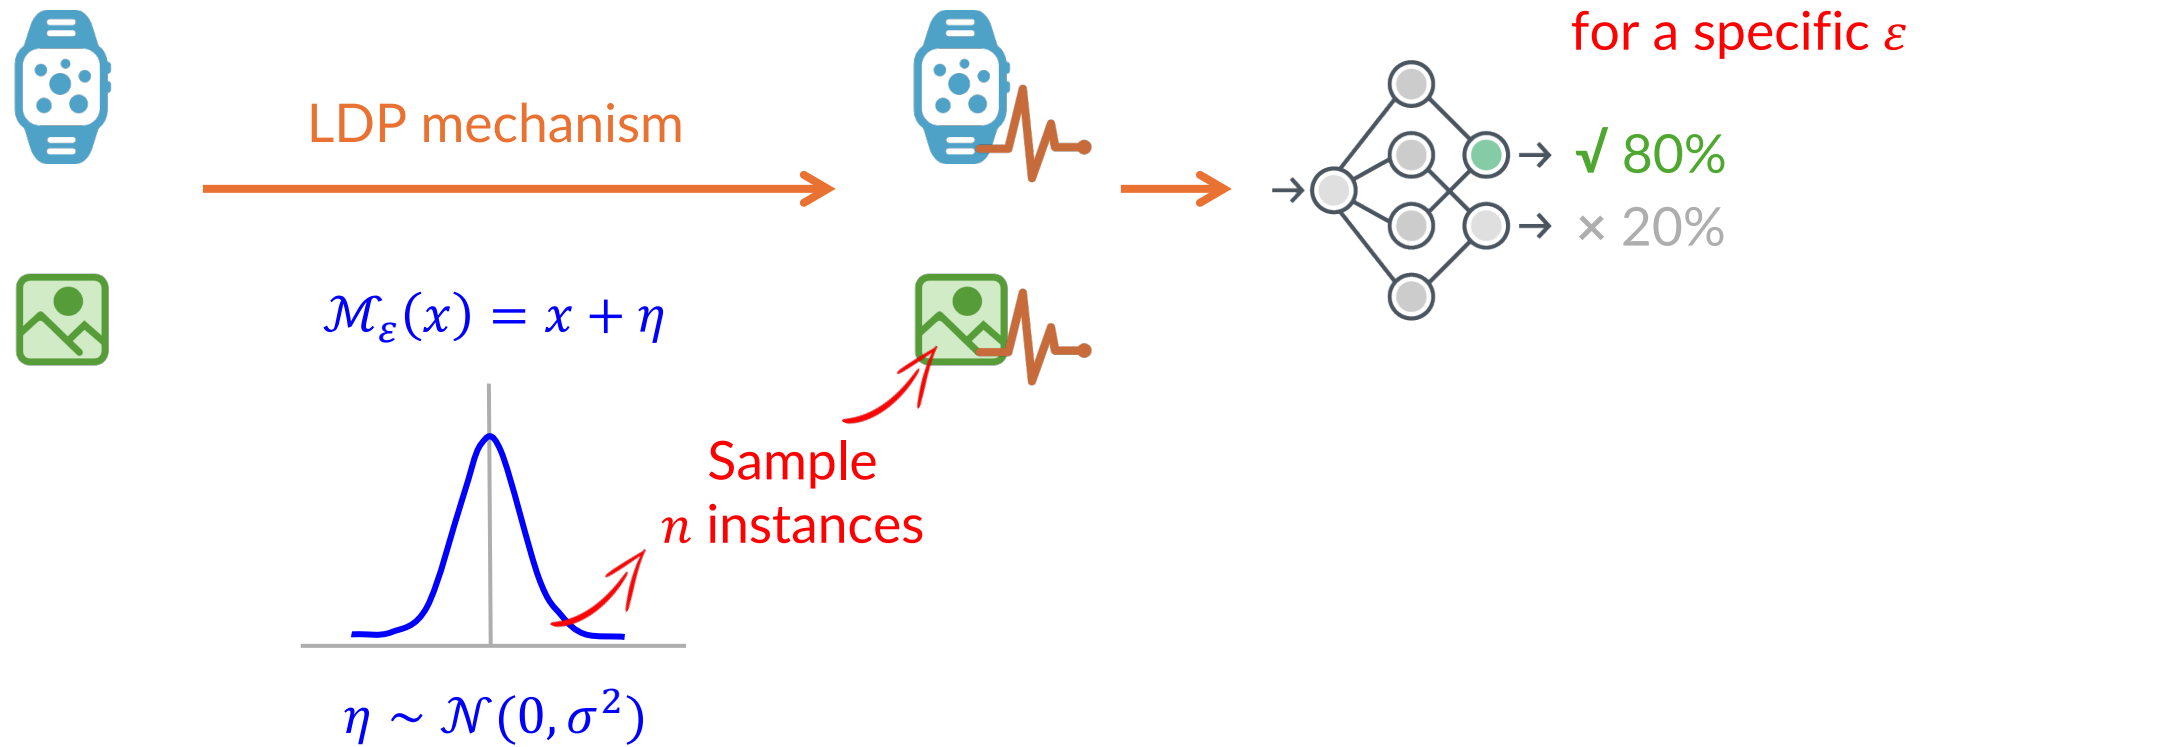
- Empirical approach: Sample and then test
 - \mathcal{M} 's variance or MSE doesn't help – cannot provide a classifier accuracy



Empirical Classifier Utility under LDP-Data

- Empirical approach: Sample and then test

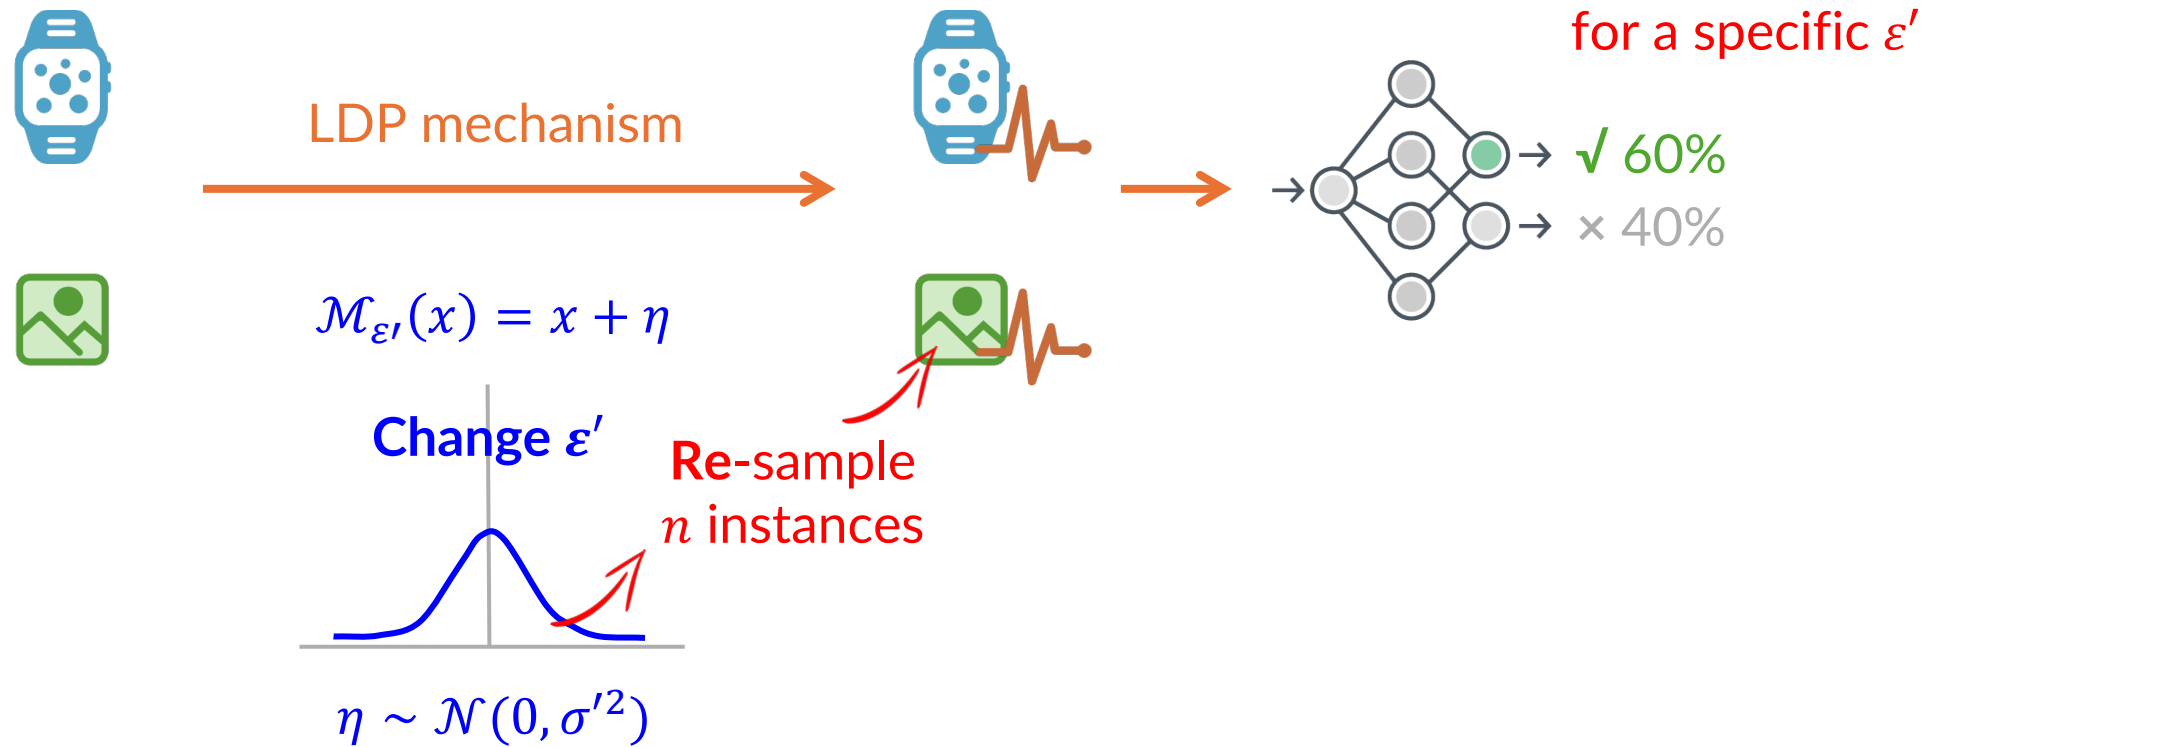
- \mathcal{M} 's variance or MSE doesn't help – cannot provide a classifier accuracy



Empirical Classifier Utility under LDP-Data

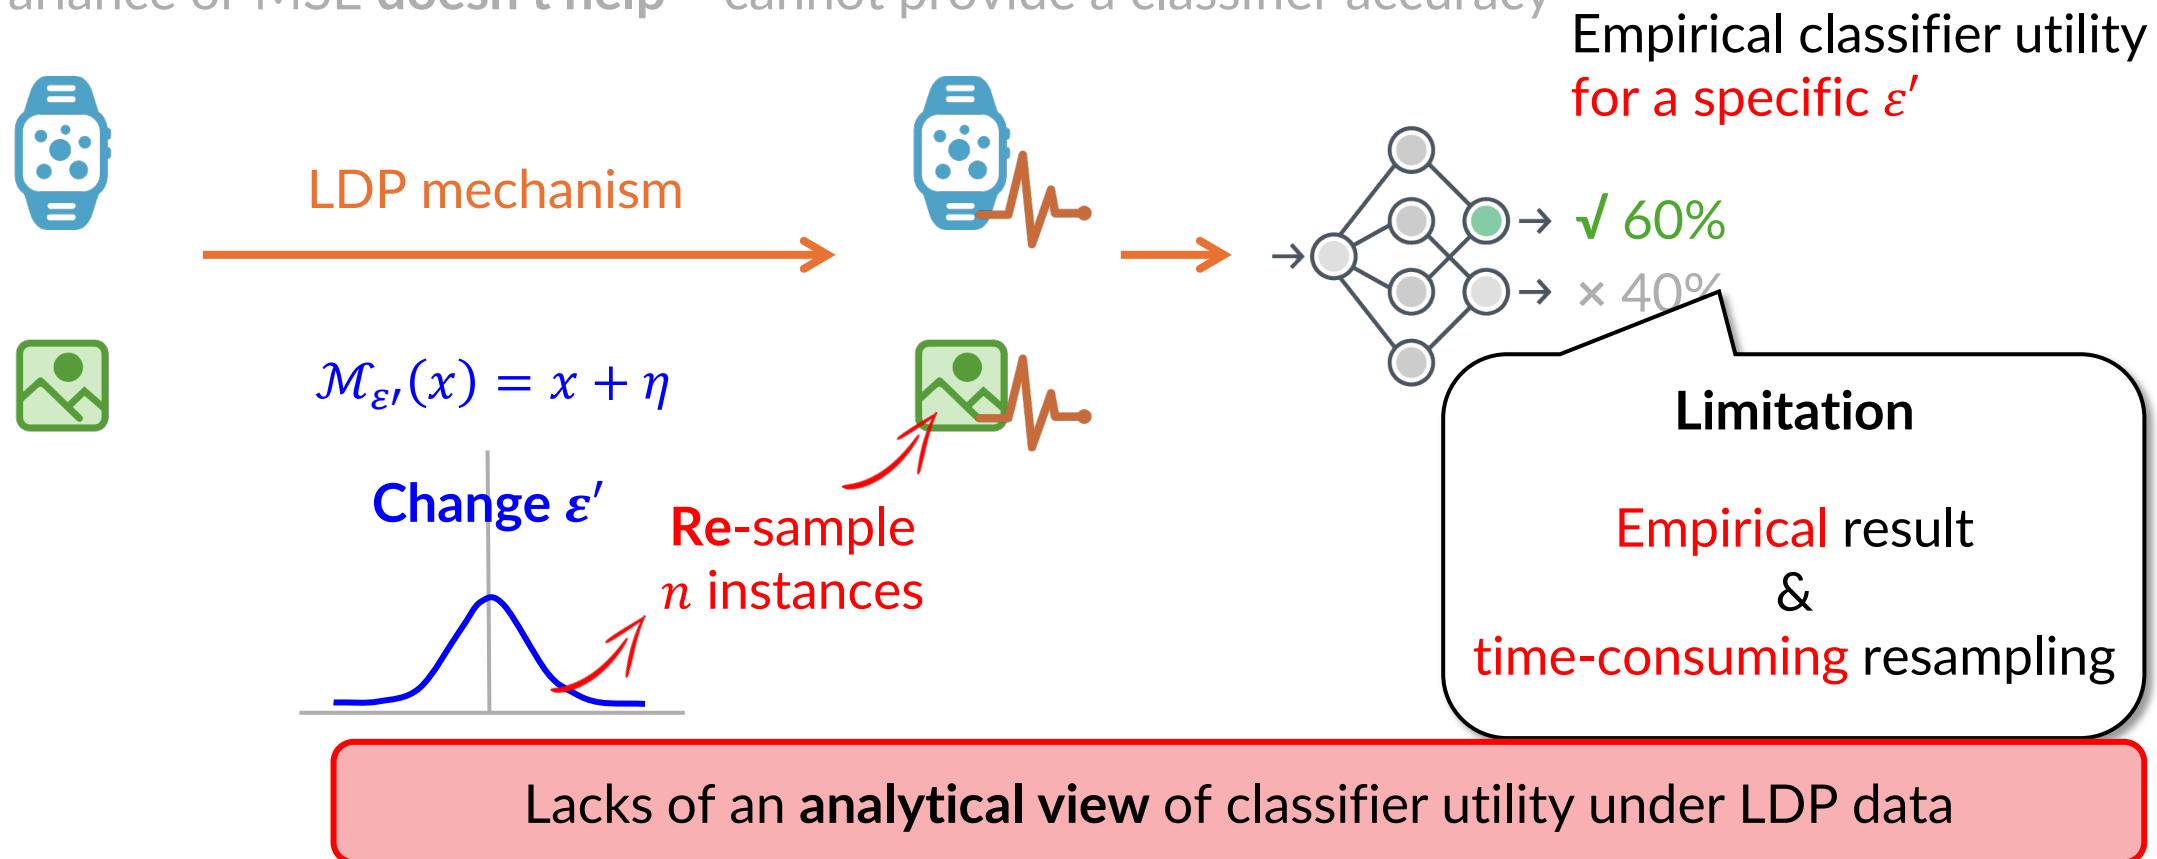
- Empirical approach: Sample and then test

- \mathcal{M} 's variance or MSE doesn't help – cannot provide a classifier accuracy




Empirical Classifier Utility under LDP-Data

- Empirical approach: Sample and then test
 - \mathcal{M} 's variance or MSE doesn't help – cannot provide a classifier accuracy



- Analytical approach: **connecting LDP with robustness**

$B_\theta(x)$ is robust




Robustness:

A θ -ball around x that doesn't change classification

- Analytical approach: **connecting LDP with robustness**

$\mathcal{M}_\varepsilon(x)$ concentrates in $B_\theta(x)$




Concentration:

Probability of falling into a region

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


$B_\theta(x)$ is robust



Robustness:


A θ -ball around x that doesn't change classification

- Analytical approach: **connecting LDP with robustness**



- Analytical approach: **connecting LDP with robustness**

$\mathcal{M}_\varepsilon(x)$ concentrates in $B_\theta(x)$




Concentration:

Probability of falling into a region

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

$B_\theta(x)$ is robust



Analytical classifier utility

“


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

”

Empirical Utility → Analytical Utility

- Analytical approach: **connecting LDP with robustness**

$\mathcal{M}_\varepsilon(x)$ concentrates in $B_\theta(x)$




Concentration:

Probability of falling into a region

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

$B_\theta(x)$ is robust



Analytical classifier utility

“

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


”

- Counterpart to the empirical “80%” accuracy

Empirical Utility \rightarrow Analytical Utility

- Analytical approach: **connecting LDP with robustness**

$\mathcal{M}_{\varepsilon'}(x)$ concentrates in $B_{\theta}(x)$




Concentration:

Probability of falling into a region

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

$B_{\theta}(x)$ is robust



Analytical classifier utility

“

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

”


Analytical & systematic view for any ε without resampling

One-Dimension Example

- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$

$B_{0.3}(x)$ is robust

Classifier:
(robustness analysis)




One-Dimension Example


- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$

$B_{0.3}(x)$ is robust

Classifier:
(robustness analysis)



LDP mechanism $\mathcal{M}_{\text{Lap}}(x)$:




One-Dimension Example

- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$

$B_{0.3}(x)$ is robust


Classifier:
(robustness analysis)

LDP mechanism $\mathcal{M}_{\text{Lap}}(x)$:
(concentration analysis)



One-Dimension Example


- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$



One-Dimension Example

- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$

Concentration analysis

$$p(\varepsilon, \theta) = \text{cdf}_{\text{Lap}}[\theta] - \text{cdf}_{\text{Lap}}[-\theta]$$
$$= 1 - e^{-\theta\varepsilon}$$
$$p(\varepsilon = 2, \theta = 0.3) = 0.46$$



One-Dimension Example

- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$

Concentration analysis

$$p(\varepsilon, \theta) = \text{cdf}_{\text{Lap}}[\theta] - \text{cdf}_{\text{Lap}}[-\theta]$$
$$= 1 - e^{-\theta\varepsilon}$$
$$p(\varepsilon = 2, \theta = 0.3) = 0.46$$

“ With probability at least $p(2,0.3) = 0.46$, the classifier *preserves its correct classification* under input $\mathcal{M}_{\varepsilon=2}(0.5)$. ”



One-Dimension Example


- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$

Concentration analysis

$$p(\varepsilon, \theta) = \text{cdf}_{\text{Lap}}[\theta] - \text{cdf}_{\text{Lap}}[-\theta]$$
$$= 1 - e^{-\theta\varepsilon}$$

For any ε :


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




Fixed Classifier (θ)

- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism $\mathcal{M}_{\text{Lap}}(x)$

This classifier's utility $p(\epsilon, 0.3)$




$B_{0.3}(x)$ is robust




Benefit 1: Choose the **best** ϵ for a desired classifier utility

- No universally optimal LDP mechanism for all ϵ and θ





(a) $\epsilon = 2$



(b) $\epsilon = 4$

- No universally optimal LDP mechanism for all ϵ and θ



Benefit 2: Systematic comparison of different LDP mechanisms


(b) $\epsilon = 4$

- Concentration analysis on the robustness region S

$$p(\varepsilon, S) = \int_S \text{pdf}[\mathcal{M}_\varepsilon(x) = \tilde{x}] d\tilde{x}$$

LDP mechanism $\mathcal{M}_\varepsilon(x)$ →

Robustness region S →




Beyond Classifier – Robustness in Other Systems

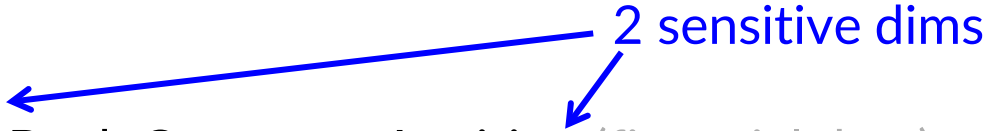
- **Communication systems**




- **Control systems**



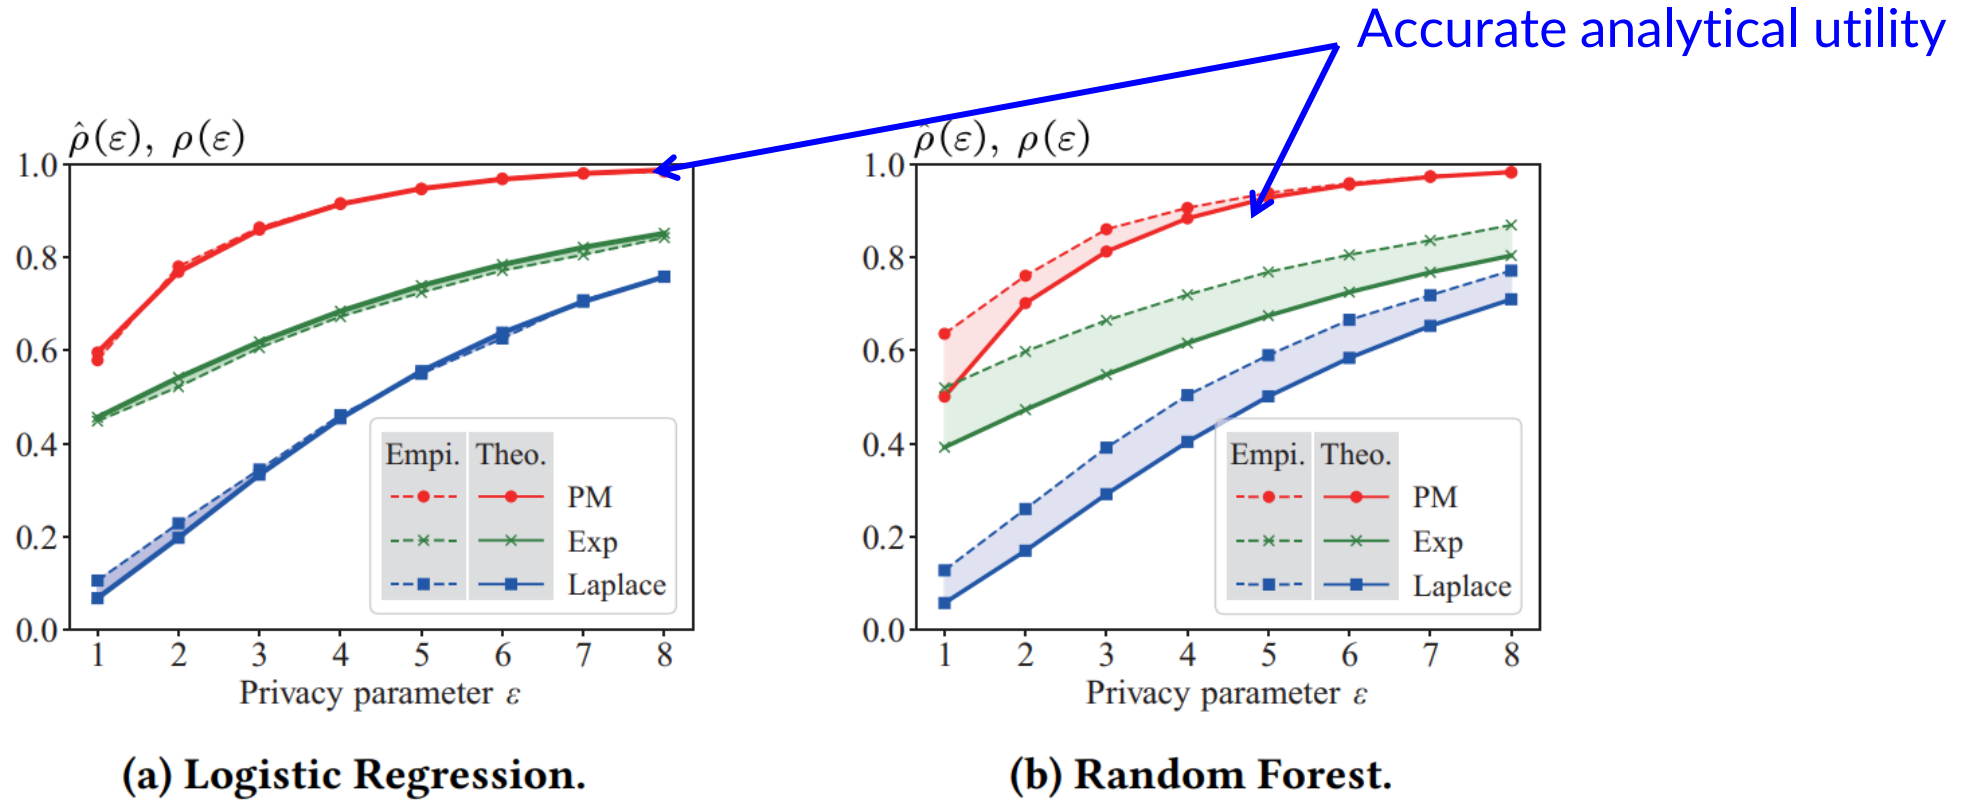
Case Studies – Quality of The Analytical Utility

- Claim: Accurate (compared with the empirical utility) & efficient
 - Classifiers:
 - Logistic Regression, Random Forest
 - Datasets: Stroke Prediction (medical data), Bank Customer Attrition (financial data),
- 
- 2 sensitive dims

Case Studies – Quality of The Analytical Utility

- Claim: Accurate (compared with the empirical utility) & efficient
- Classifiers:
 - Logistic Regression, Random Forest
- Datasets: Stroke Prediction (medical data), Bank Customer Attrition (financial data),
- Classifier utility:
 - analytical utility $p(\varepsilon, S)$: approximated S for black-box classifiers
 - empirical utility $\hat{p}(\varepsilon)$: 2000 samples from \mathcal{M}_ε and then used for testing

- x = the first record, with noisy “Age” and “BMI”



- x = the first record, with noisy “Age” and “BMI”

Time cost comparison (ms)

	PM	Exponential	Laplace
Empirical^a	6.56 + 1.38	859.83 + 1.53	11.29 + 1.53
Theoretical^b	0.24	0.94	0.30

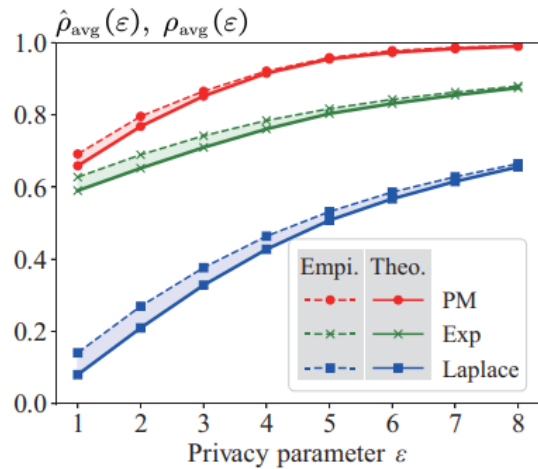
← Negligible time cost

^a Time of 2000 samples + inference.


^b Time to compute $\rho(\varepsilon, S)$ only; computing S takes 5.80 ms but is a one-time cost amortized across all ε values.

- Average-case and worst-case utility


$$x \sim P_x \rightarrow \begin{cases} p_{\text{avg}}(\epsilon) = \mathbb{E}_{x \sim P_x} [p_x(\epsilon, S)] \\ p_{\text{wor}}(\epsilon) = \min_{x \sim P_x} [p_x(\epsilon, S)] \end{cases}$$



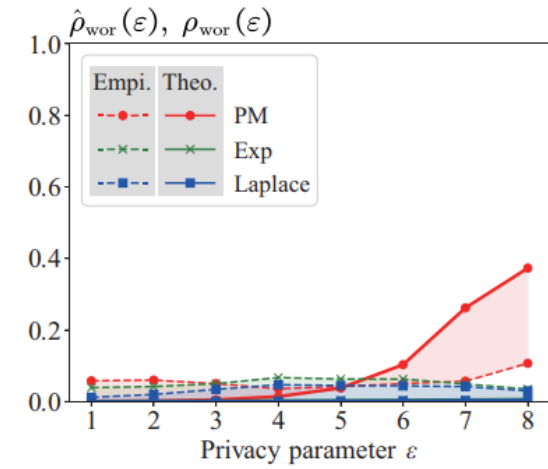
(a) LR: Average-case utility.



(b) LR: Worst-case utility.



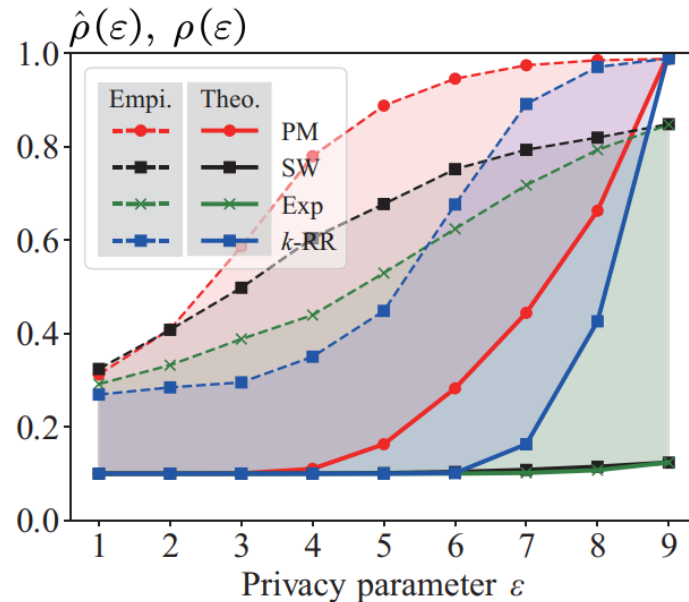
(a) RF: Average-case utility.



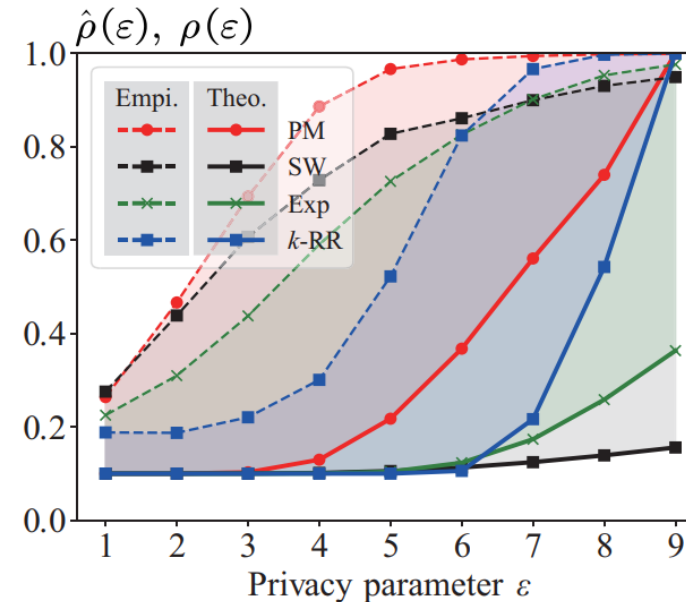
(b) RF: Worst-case utility.

Case Studies – Neural Networks

- Conservative analytical utility for high-dim classifiers (49-dim)
 - but high analytical utility \rightarrow high empirical utility




1st image



2nd image

- **RQ: Classifier utility under LDP-perturbed inputs**
- **Our approach:**
 - **connects LDP with robustness**
 - **provides analytical classifier utility**
 - the analytical utility is accurate for low-dim classifiers



* Quantifying Classifier Utility under Local Differential Privacy, PETS'26

LDP: Refined Mechanism Design and Utility Analysis

Privacy-Preserving Computation - LDP

RIT

- Local differential privacy (LDP)
 - hard to differentiate the
 - each user locally perturbs

Advantages: Negligible computational complexity, No communication between users

But approximated f

$\mathcal{M}(x) = x + \eta$, $\eta \sim \mathcal{N}(0, \sigma^2)$

$\tilde{x}_1 = \mathcal{M}(x_1)$, $\tilde{x}_2 = \mathcal{M}(x_2)$, $\tilde{x}_3 = \mathcal{M}(x_3)$, $\tilde{x}_4 = \mathcal{M}(x_4)$, ..., $\tilde{x}_n = \mathcal{M}(x_n)$

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

This Proposal: LDP Theory

RIT

- Advancing LDP's mechanism design and utility analysis
 - binary $x \rightarrow$ 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 $\circ \mathcal{M}$

$1D\ x \rightarrow 2D\ x$

Mechanism-level \downarrow Task-level

f is a classifier

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

This Paper: Joint/Correlated RR (JRR)

RIT

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

Error

$\text{Var}[\mathcal{M}(x_i)] + \text{Var}[\mathcal{M}(x_j)] + 2 \cdot \text{Cov}$ (Can be < 0)

Optimal Piecewise-based Mechanism

RIT

- Most generalized version: m -piecewise distribution

$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_{1,\epsilon} & \text{if } y \in [1, x, \epsilon] \\ p_{2,\epsilon} & \text{if } y \in [2, x, \epsilon] \\ \dots & \dots \\ p_{m,\epsilon} & \text{if } y \in [m, x, \epsilon] \end{cases}$$

$\frac{p_{i,\epsilon}}{p_{j,\epsilon}} \leq e^\epsilon$ (LDP constraint)

Find \mathcal{M} to minimize the worst-case error

Solved \mathcal{M} is the optimal piecewise-based mechanism

Mathematically \equiv to find the optimal piecewise distribution under the LDP constraint

Continuous Spaces: Better & Universal

One-Dimension Example

RIT

- Operate on a continuous space (covering the discrete space)

Simple sampling: (satisfying LDP for \mathcal{S})

$\Pr[\mathcal{M}(x_i) = \tilde{x} \in \square] = p_{\text{high}}$

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

$\mathcal{S} = [a_{\text{sta}}, a_{\text{end}}] \times [b_{\text{sta}}, b_{\text{end}}]$

- Benefits:
 - Efficiency: $\mathcal{O}(1)$ sampling complexity
 - Trajectory utility: "n-independent"
 - Applicability: Both continuous spaces & discrete spaces

- Classifier $h: [0,1] \rightarrow \{1, 2\}$ under Laplace mechanism \mathcal{M}_{Lap}

Concentration analysis

$p(\epsilon, \theta) = \text{cdf}_{\text{Lap}}[\theta] - \text{cdf}_{\text{Lap}}[-\theta] = 1 - e^{-\theta\epsilon}$

For any ϵ : With probability at least $p(\epsilon, \theta)$, the classifier preserves its correct classification under input $\mathcal{M}_\epsilon(x)$.

$B_{0.3}(x)$ is robust

The probability of having $B_{0.3}(x)$

Thank you!

