Bayesian differential privacy for linear dynamical systems

Genki Sugiura, Kaito Ito, and Kenji Kashima
Kyoto University, Japan

Outline

- **Overview**
  - **Motivation, concept and example**
- Theoretical results
  - Problem setting and theorems

Motivation

- Utility of user data with **private information**
  - e.g., microgrid

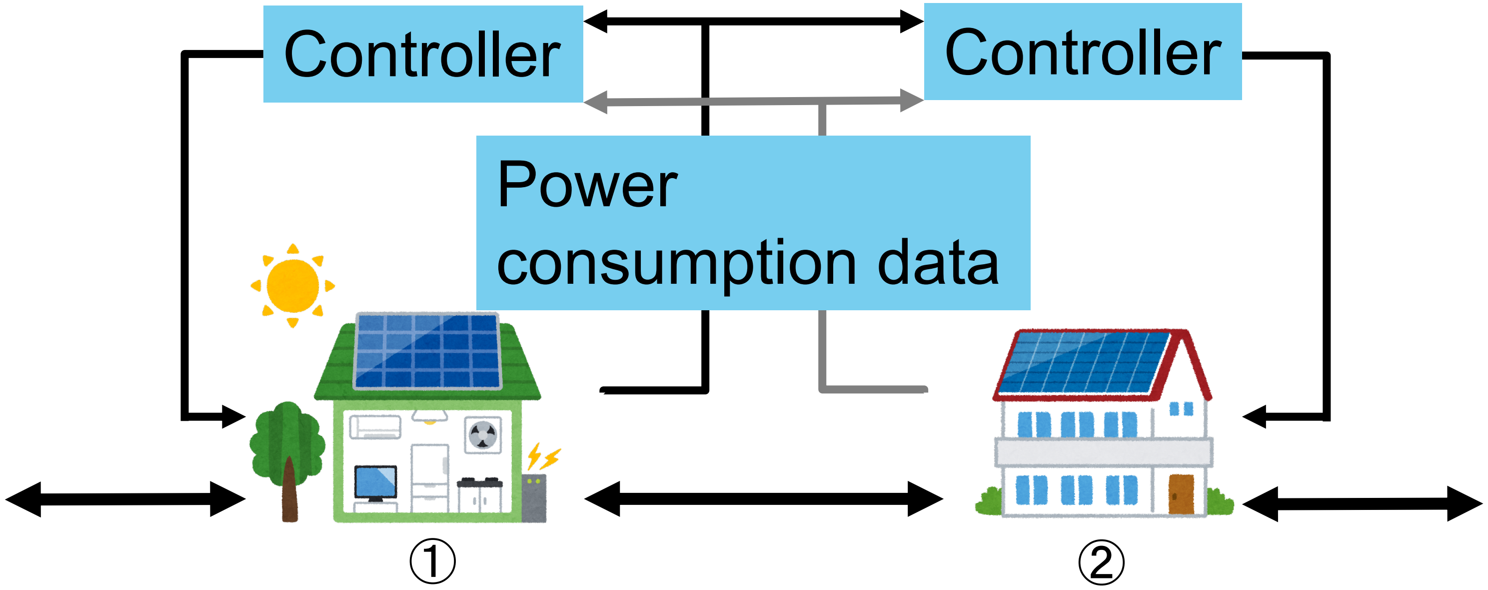

  - Input signal to ② is influenced by the consumption pattern of ①
    - **Risk of consumption data leakage**

→ Need for control method considering **privacy protection**

Privacy protection on dynamical systems

- Privacy protection by adding noise (differential privacy or DP)
  - Difficulty of distinguishing $u_1(t)$ and $u_2(t)$ = Privacy level

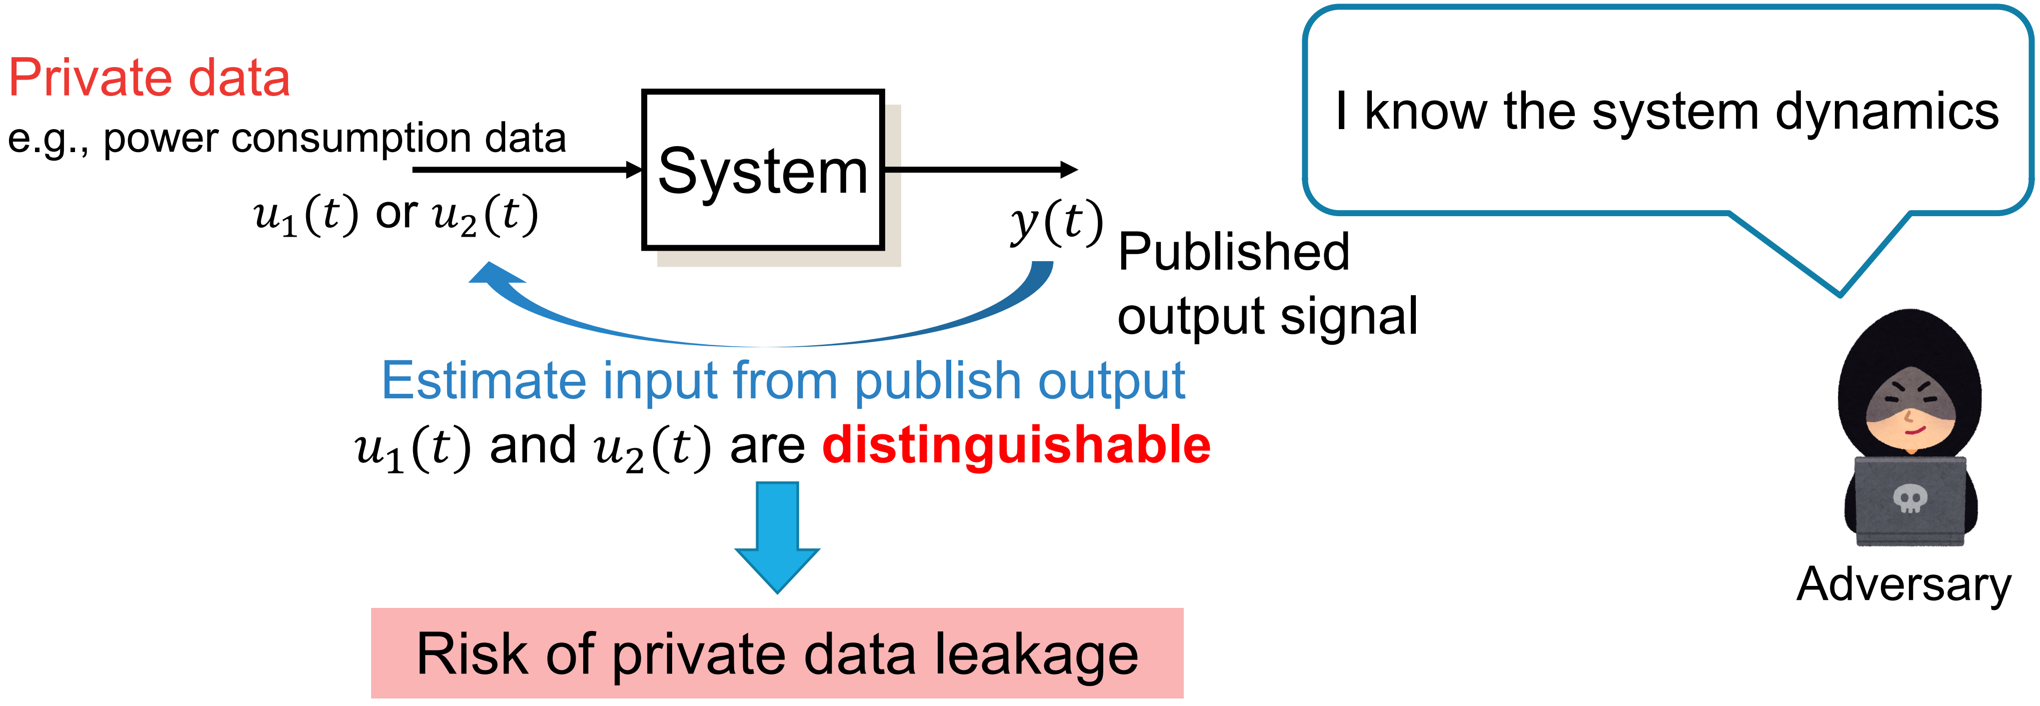

Privacy protection on dynamical systems

- Privacy protection by adding noise (differential privacy or DP)
  - Difficulty of distinguishing $u_1(t)$ and $u_2(t)$ = Privacy level

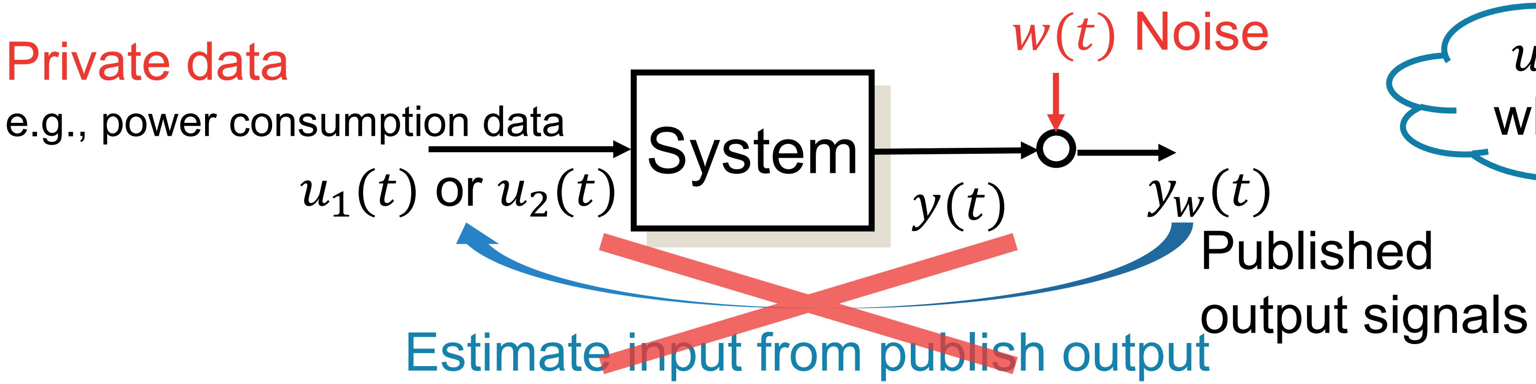

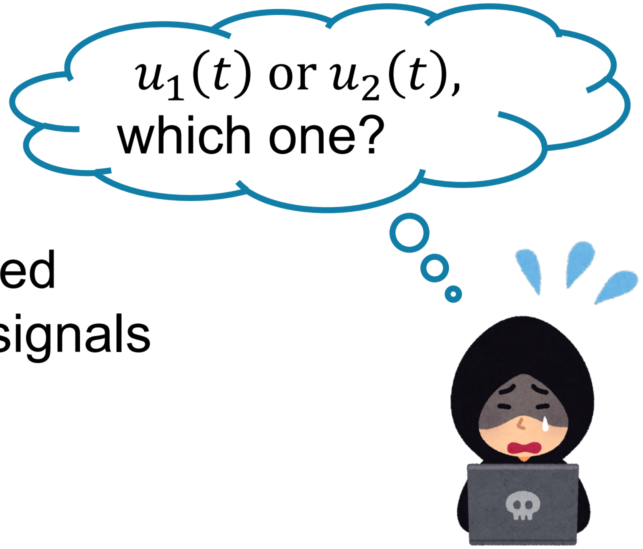

Adversary

$u_1(t)$ and $u_2(t)$ are **indistinguishable**

- To be considered
  - What types of input signal pairs $(u_1(t), u_2(t))$ can we protect?
  - Noise scale? → Large noise decreases the information usefulness

"Privacy protection level vs. information usefulness" [1]

[1] Y. Kawano, K. Kashima, and M. Cao, "Modular control under privacy protection: Fundamental trade-offs," Automatica, vol. 127, May 2021, Art. no. 109518.

Previous research

- Privacy protection by adding Gaussian noise[2, 3]

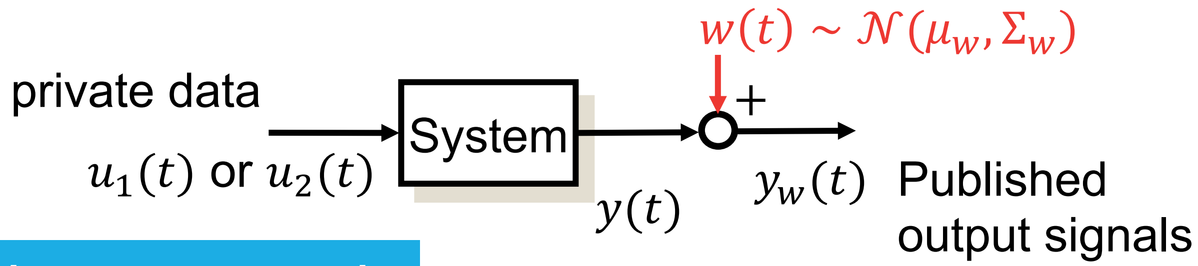

Previous research

Privacy is guaranteed only for similar data
Privacy protection of outlier data is **NOT** considered

u_1 and u_2 are similar

def

$||u_1 - u_2|| \leq c$ (predetermined)

u_1 and u_2 are far apart

This talk

Utilizing **the prior distribution of the data** for privacy protection of dynamical systems

- Providing privacy guarantees even for outlier data
- For static data, it has been studied as Bayesian DP[4]

[2] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Trans. Automat. Control*, vol. 59, no. 2, pp. 341–354, Feb. 2014.
[3] Y. Kawano and M. Cao, "Design of privacy-preserving dynamic controllers," *IEEE Trans. Automat. Control*, vol. 65, no. 9, pp. 3863–3878, Sep. 2020.
[4] A. Triastcyn and B. Faltings, "Bayesian differential privacy for machine learning," in Proc. Int. Conf. Mach. Learn., Nov. 2020, pp. 9583–9592.

Problem setting

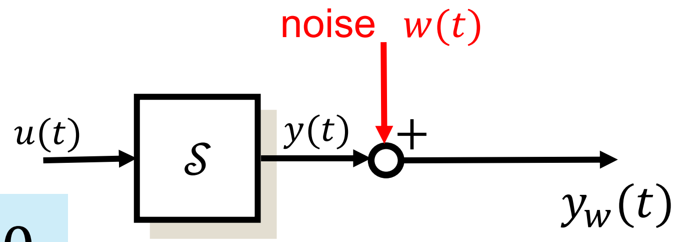

$$\mathcal{S} : \begin{cases} x(t+1) = Ax(t) + Bu(t), & x(0) = 0 \\ y(t) = Cx(t) + Du(t) \end{cases}$$

$x \in \mathbb{R}^n$: state, $u \in \mathbb{R}^m$: input, $y \in \mathbb{R}^q$: output

I know $\mathcal{S}$ and prior distr. of U_T

Adversary

Public output

- $$Y_{w,T} := \begin{bmatrix} y_w(0) \\ \vdots \\ y_w(T) \end{bmatrix} = N_T U_T + W_T$$

$$N_T := \begin{bmatrix} D & 0 & \cdots & \cdots & 0 \\ CB & D & \ddots & \ddots & \vdots \\ CAB & CB & D & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ CA^{T-1}B & CA^{T-2}B & \cdots & CB & D \end{bmatrix}$$

Private data

- $$U_T := \begin{bmatrix} u(0) \\ \vdots \\ u(T) \end{bmatrix} \sim \mathcal{N}(0, \Sigma_u)$$

Prior distr.

Noise

- $$W_T := \begin{bmatrix} w(0) \\ \vdots \\ w(T) \end{bmatrix} \sim \mathcal{N}(0, \Sigma_w)$$

Example: Private data with prior distribution

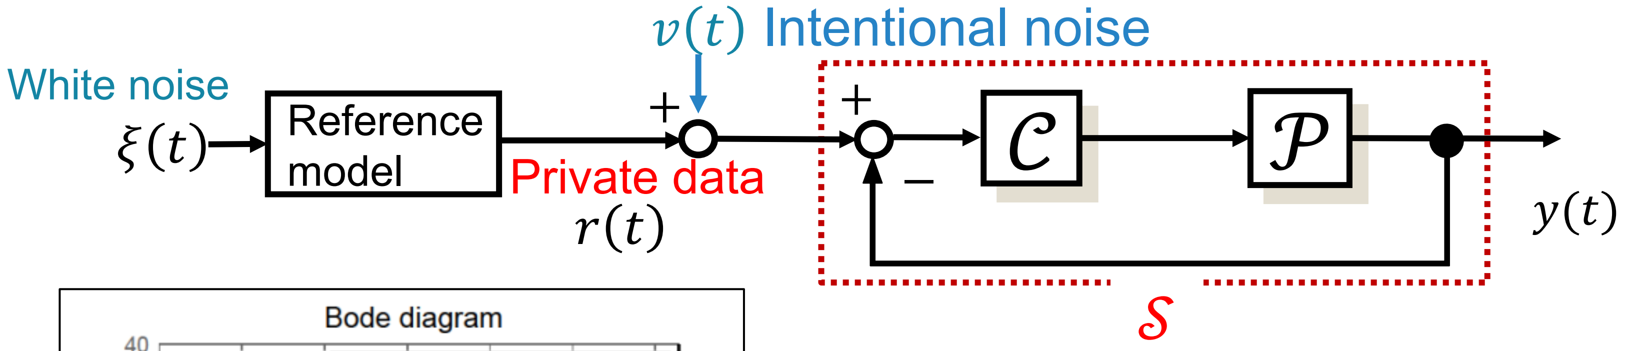

$$\mathcal{S}: \begin{cases} x(t+1) = Ax(t) + B(r(t) + v(t)), \\ y(t) = Cx(t) \end{cases}$$

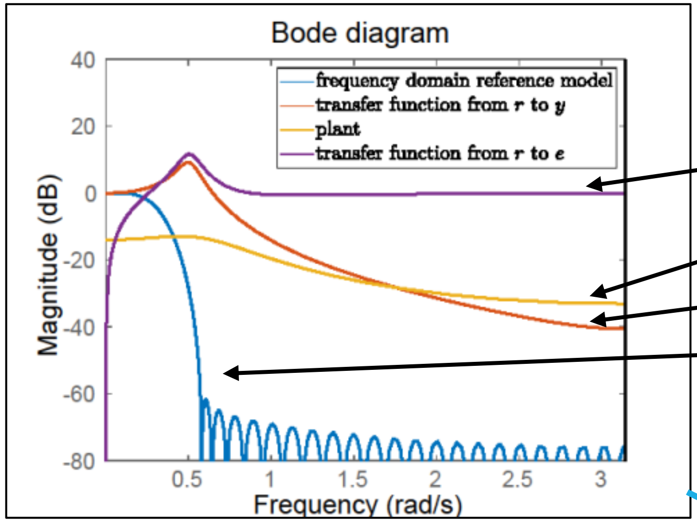

$G_{r \to e}$

G_{plant}

$G_{r \to y}$

Frequency domain reference model (Low pass filter)

I know both system dynamics and reference model.

- With prior information, $r(t)$ is **easily estimated** from $y(t)$
  - e.g., $r(t)$ concentrates on the low frequency range
  - Need for larger noise → "Privacy protection level vs. information usefulness"

Example:
Private data with prior distribution

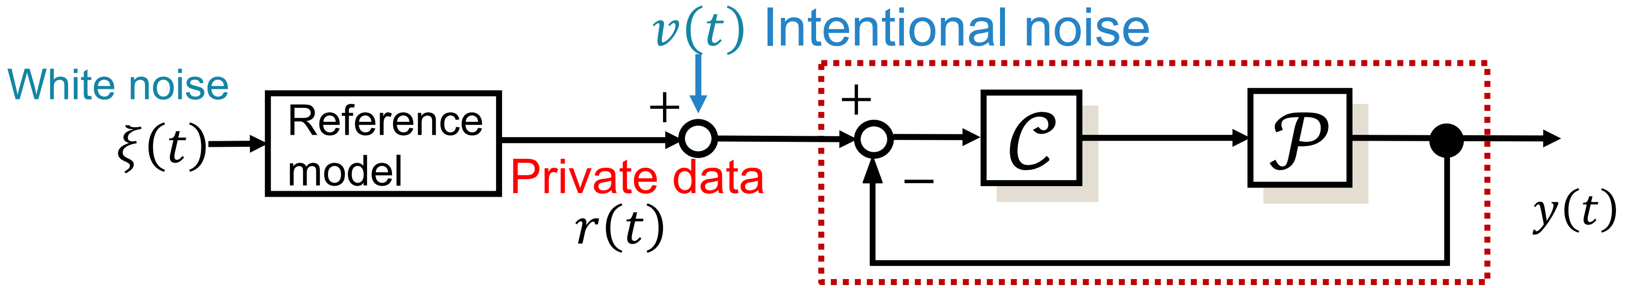

$$\mathcal{S}:\begin{cases} x(t+1) = Ax(t) + B(r(t)+v(t)), \\ y(t) = Cx(t) \end{cases}$$

Reference

Noise-free
(good tracking, no pricacy)

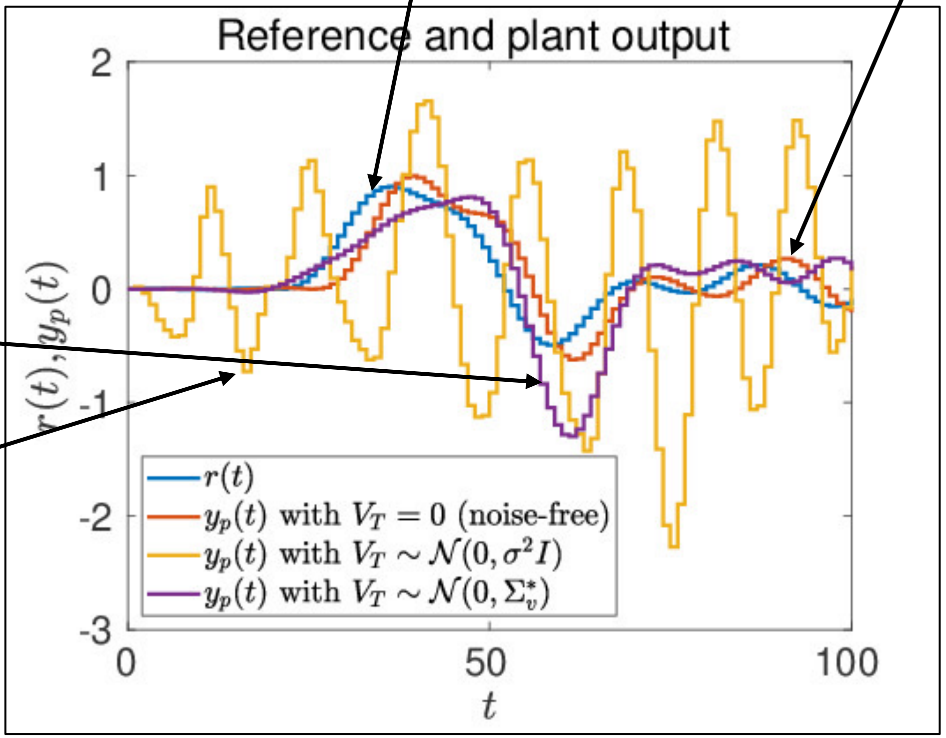

Guaranteeing
the same privacy level

Optimal noise
(Minor deterioration)

Optimal i.i.d. noise
(not tracking at all)

Outline

- Overview
  - Motivation, concept and example
- **Theoretical results**
  - **Problem setting and theorems**

Differential privacy for dynamical systems

$$Y_{w,T} = N_T U_T + W_T \cdots (\bigstar)$$

$$U_T \longrightarrow \boxed{(\bigstar)} \longrightarrow Y_{w,T}$$

(ε, δ)-differential privacy*

For given $\varepsilon > 0$, $(\bigstar)$ satisfies
ε - Differential privacy (ε - DP)

$\Updownarrow$ def

$$\frac{\mathbb{P}[Y^1_{w,T} \in S]}{\mathbb{P}[\underbrace{Y^2_{w,T}}_{N_T U_T^2 + W_T} \in S]} \leq e^{\varepsilon} \quad \forall S \subset \mathbb{R}^{(T+1)q}$$

for similar (U_T^1, U_T^2)

$||U_T^1 - U_T^2|| \leq c, \quad c > 0$

$$\frac{\mathbb{P}[Y^1_{w,T} \in S]}{\mathbb{P}[Y^2_{w,T} \in S]} \leq e^{\varepsilon}$$

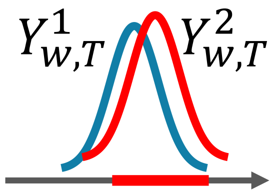

ε:**small**

$\downarrow$

distr. of $Y^1_{w,T}$ and $Y^2_{w,T}$ are **close**

$\downarrow$

difficult to distinguish U_T^1, U_T^2

* For the simplicity of the presentation, we take $\delta = 0$.

Main Problem: Conventional differential privacy

Problem

Provide the privacy notion which guarantees difficulty of distinguishing (U_T^1, U_T^2) even if $||U_T^1 - U_T^2|| > c$

Required level of difficulty of distinguishing data

Uniform over similar data

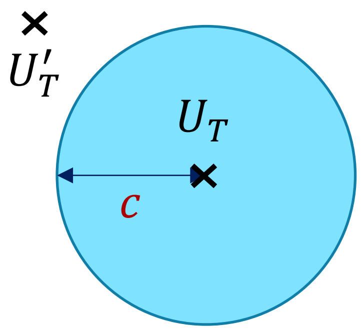

Conventional DP :
Protecting only “similar” data

Weighted in terms of prior distribution

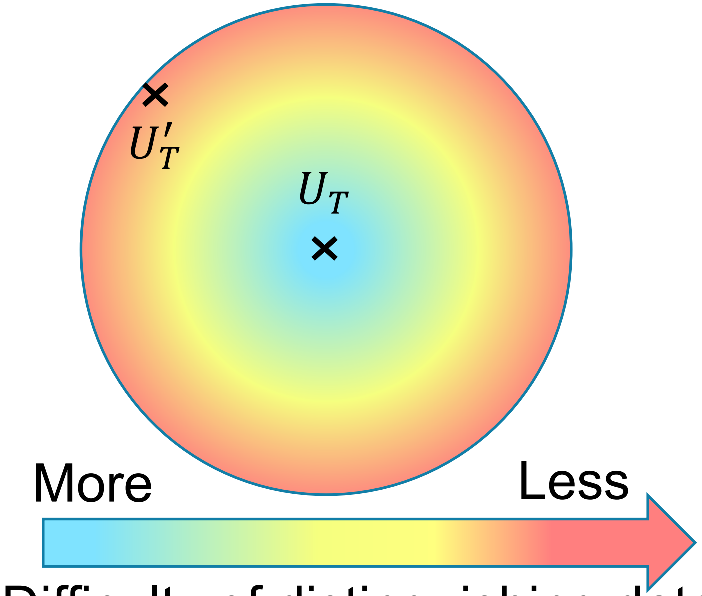

Difficulty of distinguishing data

Main result 1
Bayesian differential privacy for dynamical systems

$$Y_{w,T} = N_T U_T + W_T \cdots (\bigstar)$$

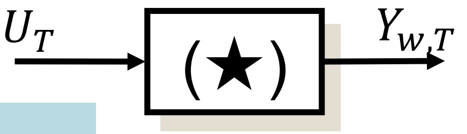

$(\mathbb{P}_{U_T}, \gamma, \varepsilon, \delta)$-Bayesian differential privacy*

U_T^1, U_T^2: i. i. d. with the distr. $\mathbb{P}_{U_T}$

For given $1 \geq \gamma \geq 0, \varepsilon > 0$, $(\bigstar)$ satisfies

$(\mathbb{P}_{U_T}, \gamma, \varepsilon)$-Bayesian differential privacy ($(\mathbb{P}_{U_T}, \gamma, \varepsilon)$ - BDP)

$\Updownarrow$ def

$$\mathbb{P}_{U_T}\left[\frac{\mathbb{P}_W\left[Y_{w,T}^1 \in S \mid U_T^1\right]}{\mathbb{P}_W\left[Y_{w,T}^2 \in S \mid U_T^2\right]} \leq e^{\varepsilon}\right] \geq \gamma \quad \forall S \subset \mathbb{R}^{(T+1)q}$$

c.f.) ε – DP

$$\frac{\mathbb{P}\left[Y_{w,T}^1 \in S\right]}{\mathbb{P}\left[Y_{w,T}^2 \in S\right]} \leq e^{\varepsilon} \quad \forall S \subset \mathbb{R}^{(T+1)q}$$

Small ε = High privacy level
Large γ = High privacy level even for far apart data pair U_T, U_T'

* For the simplicity of the presentation, we take $\delta = 0$.

Main result 2
Gaussian noise guaranteeing BDP

$$Y_{w,T} = N_T U_T + W_T \cdots (\bigstar)$$

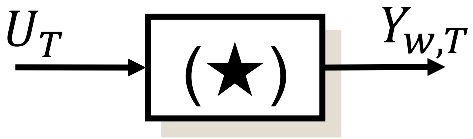

$U_T \sim \mathcal{N}(0, \Sigma_u)$

Prior distr.

$W_T \sim \mathcal{N}(\mu_w, \Sigma_w)$

Design parameter

$$N_T := \begin{bmatrix} D & 0 & \cdots & \cdots & 0 \\ CB & D & \ddots & \ddots & \vdots \\ CAB & CB & D & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ CA^{T-1}B & CA^{T-2}B & \cdots & CB & D \end{bmatrix}$$

$(\mathbb{P}_{U_T}, \gamma, \varepsilon, \delta)$ -BDP is satisfied if $\Sigma_w \succ 0$ is chosen such that

$$\lambda_{\max}\left(\Sigma_u^{1/2} N_T^\top \Sigma_w^{-1} N_T \Sigma_u^{1/2}\right)^{1/2} \leq \frac{1}{c(\gamma, T)} R^{-1}(\varepsilon)$$

become smaller
when Σ_w is large
(low information utility)

become smaller
when ε is small
γ is large
(High privacy level)

How to maximize information usefulness with privacy guarantee?

* For the simplicity of the presentation, hereafter we omit the argument δ of the function R

Main result 3
Optimal Gaussian noise guaranteeing BDP

$$Y_{w,T} = N_T U_T + W_T \cdots (\bigstar)$$

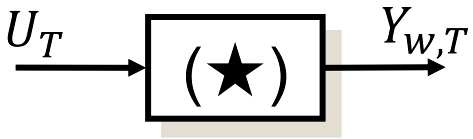

$U_T \sim \mathcal{N}(0, \Sigma_u)$ Prior distr.

$W_T \sim \mathcal{N}(\mu_w, \Sigma_w)$ Design parameter

$$N_T := \begin{bmatrix} D & 0 & \cdots & \cdots & 0 \\ CB & D & \ddots & \ddots & \vdots \\ CAB & CB & D & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ CA^{T-1}B & CA^{T-2}B & \cdots & CB & D \end{bmatrix}$$

$$\begin{cases} \min\limits_{\Sigma_w \succ 0} & \mathrm{Tr}(\Sigma_w) \\ \mathrm{s.t.} & \text{(sufficient condition for BDP)} \end{cases}$$

LMI constraint

Assumption: N_T has full row rank

Minimum energy Gaussian noise guaranteeing BDP is

$$\Sigma_w^* := c(\gamma, T)^2 R(\varepsilon)^2 N_T \Sigma_u N_T^\top$$

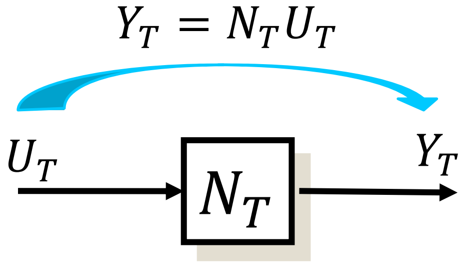

W_T having the same shape of distr. as Y_T can efficiently protect U_T

Main result 4
Input noise mechanism

$$N_T := \begin{bmatrix} D & 0 & \cdots & \cdots & 0 \\ CB & D & \ddots & \ddots & \vdots \\ CAB & CB & D & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ CA^{T-1}B & CA^{T-2}B & \cdots & CB & D \end{bmatrix}$$

Assumption: N_T is regular

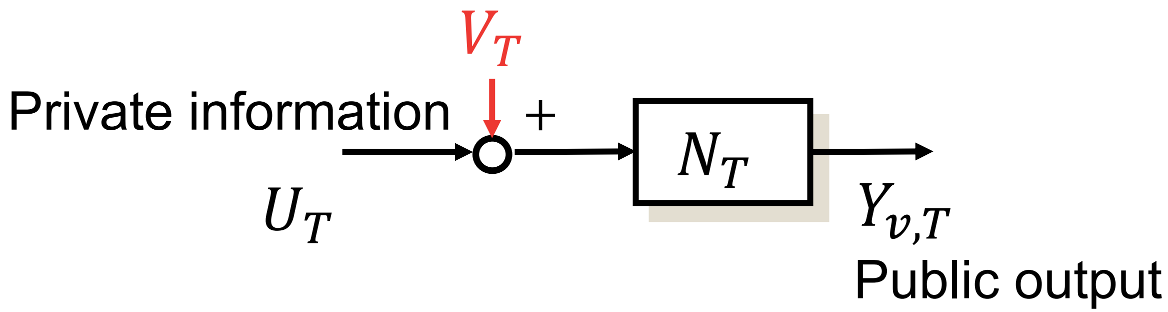

$$Y_{v,T} = N_T U_T + N_T V_T$$
$$V_T \sim \mathcal{N}(0, \Sigma_v)$$

Design parameter

In input noise case,
- sufficient condition for BDP guarantee
- optimal Gaussian noise

are **independent of system parameters**

BDP Condition: $\lambda_{\min}\left(\Sigma_u^{-1/2}\Sigma_v\Sigma_u^{-1/2}\right)^{1/2} \geq c(\gamma, T)R(\varepsilon)$

Opt. noise: $\Sigma_v^* = c(\gamma, T)^2 R(\varepsilon)^2 \Sigma_u$

Summary

- Objective
  - To protect input data to control systems
    - Previous research does not consider distant data sets
- Results
  - Introduced Bayesian differential privacy (BDP) to linear dynamical systems
    - Provided privacy guarantees even for distant data sets
    - Derived the minimum energy Gaussian noise guaranteeing BDP
      - Privacy VS. information utility
- Future work
  - Privacy protection in the infinite horizon
    - Difficulty: our privacy parameter $c(\gamma, T)$ is increasing function of T
      - Noise $\to \infty$ as $T \to \infty$

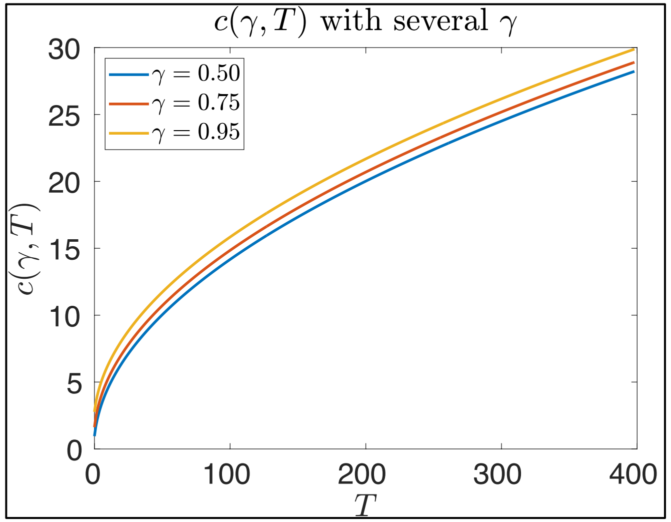

BDP condition:

$$\lambda_{\min}\left(\Sigma_u^{-1/2}\Sigma_v\Sigma_u{}^{-1/2}\right)^{1/2} \geq c(\gamma, T)R(\varepsilon)$$