Simulation-Driven Parameter Estimation with a Focus on Control and Privacy

Braghadeesh Lakshminarayanan

Doctoral Thesis Defence
Department of Decision and Control Systems
KTH Royal Institute of Technology, Sweden

May 8, 2026

Outline

1. **Simulation-Driven Estimation** — motivation and the Two-Stage (TS) estimator
2. **Part 1: Analysis of the TS Estimator**
3. **Part 2: Fine-Tuning a Simulation-Driven Estimator**
4. **Part 3: Applications in Control**
5. **Part 4: Privacy in Point Estimation**

Simulation-Driven Estimation

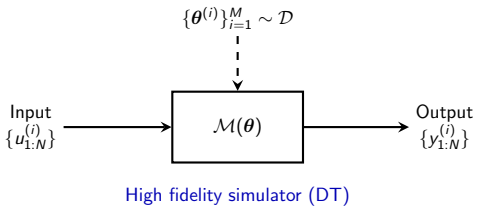

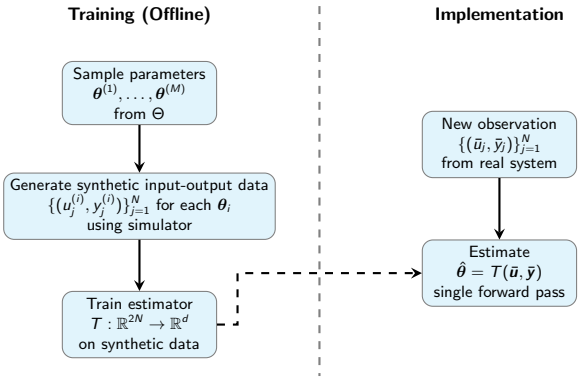

S. Garatti, S. Bittanti. "A new paradigm for parameter estimation in system modeling". *Int. J. Adapt. Control Sig. Proc.*, 2013.
T. Diskin et al. "Learning to estimate without bias". *Trans. Sig. Proc.*, 2022.
A. Ghosh et al. "DeepBayes—An estimator for parameter estimation in stochastic nonlinear dynamical models". *Automatica*, 2024.

Why Simulation-Driven Estimation?

Traditional methods

- PEM, EKF, MLE, ...
- High inference time

Simulation-driven (TS)

- Train offline using synthetic data
- Inference = single forward pass
- **Low inference time**

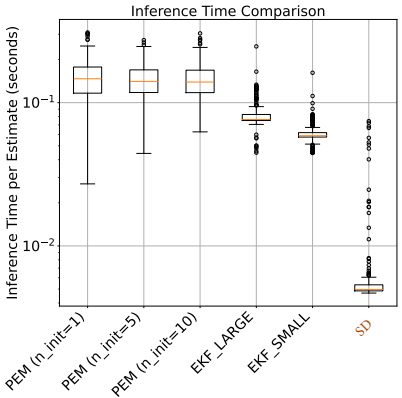

Part I

Analysis of a Simulation-Driven Estimator

The Two-Stage (TS) Estimator

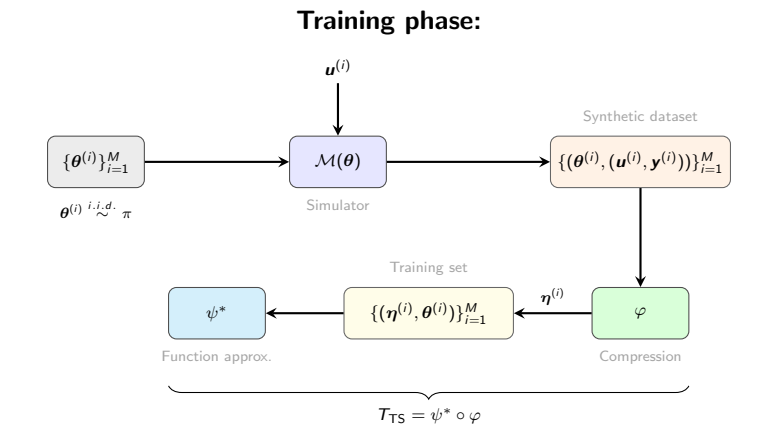

S. Garatti, S. Bittanti. "A new paradigm for parameter estimation in system modeling" *Int. J. Adapt. Control Sig. Proc.*, 2013.

Chapter 3

Statistical Framework of the TS Estimator

B.Lakshminarayanan, C.R. Rojas. "A Statistical Decision-Theoretical Perspective on the Two-Stage Approach to Parameter Estimation" *IEEE CDC*, pages 5369–5374, 2022.

Statistical Frameworks of the TS Estimator

Bayes TS
Minimize the average risk:

$$\psi^* = \arg\min_{\psi \in \mathcal{G}} \frac{1}{M} \sum_{i=1}^{M} \ell(\boldsymbol{\theta}^{(i)}, \psi(\boldsymbol{\eta}^{(i)}))$$

- $\boldsymbol{\theta}^{(i)} \sim \pi$ (prior)
- Standard supervised learning

Minimax TS
Minimize the worst-case risk:

$$\psi^* = \arg\min_{\psi \in \mathcal{G}} \max_{i=1,\dots,M} \ell(\boldsymbol{\theta}^{(i)}, \psi(\boldsymbol{\eta}^{(i)}))$$

- $\boldsymbol{\theta}^{(i)} \sim s$ (proposal)
- Robust: no prior needed

Both share the same structure $T_{\mathsf{TS}} = \psi^* \circ \varphi$; they differ only in how ψ^* is trained.

Minimax via epigraph formulation:

$$\min_{\boldsymbol{\beta},\, t} \; t \qquad \text{s.t.} \quad \ell(\boldsymbol{\theta}^{(i)}, \boldsymbol{\beta}^\top J(\boldsymbol{\eta}^{(i)})) \le t, \quad i = 1, \dots, M$$

Numerical Example

Model: $y_i \overset{i.i.d.}{\sim} \text{Weibull}(\lambda, k)$, λ: scale, k: shape

Bayes TS Estimator

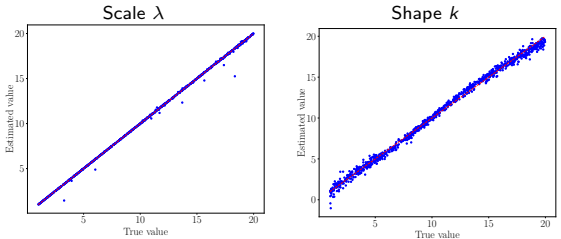

performance depends on prior π.

Minimax TS Estimator

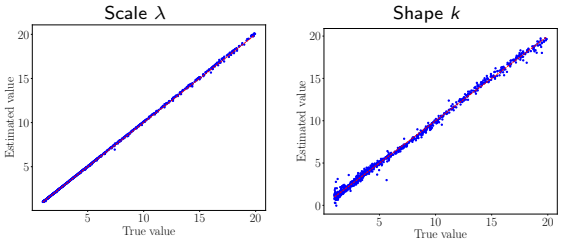

Robust (no prior); higher variance $\rightarrow$ motivates improved minimax (Ch. 4).

B.L., C.R. Rojas. *IEEE CDC*, 2022.

Chapter 4

Improved Minimax TS via Gradient Boosting

B. Lakshminarayanan and C. R. Rojas. "Minimax Two-Stage Gradient Boosting for Parameter Estimation". *IEEE CDC*, pages 1189–1194, 2023.

Improved Minimax TS via Gradient Boosting

Minimax objective:

$$\min_{\psi} \max_{i=1,\dots,M} \ell(\boldsymbol{\theta}^{(i)}, \psi(\boldsymbol{\eta}^{(i)}))$$

Problem: max is non-smooth.

↓ *soft-max approximation*

TSGBM objective:

$$\min_{\psi} \frac{1}{K} \log\left(\sum_{i=1}^{M} e^{K\, \ell(\boldsymbol{\theta}^{(i)}, \psi(\boldsymbol{\eta}^{(i)}))} \right)$$

GBM: iterative residual fitting

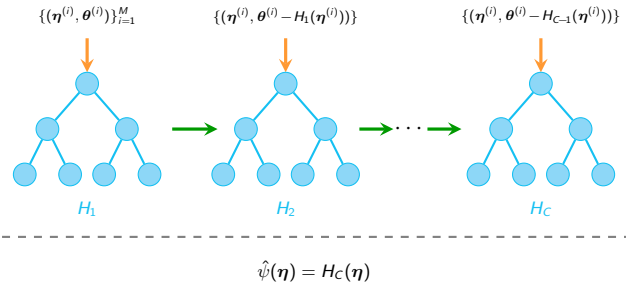

Smooth → usable as GBM loss | $K \to \infty$ recovers exact minimax

Numerical Example

Model: $y_i \overset{i.i.d.}{\sim} \text{Weibull}(\lambda, k)$, λ: scale, k: shape

Minimax TS (Ch. 3)

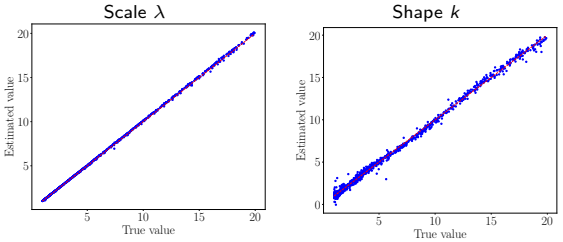

High variance for k; slow (epigraph + CVXPY per sample).

TSGBM (Ch. 4)

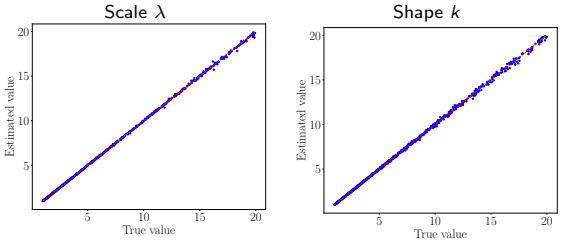

Low variance for k improved up to 10×; much faster via soft-max + LightGBM.

Chapter 5

Asymptotic Analysis of the TS Estimator

B. Lakshminarayanan and C. R. Rojas. "On Asymptotic Analysis of the Two-Stage Approach: Towards Data-Driven Parameter Estimation". Submitted to *Automatica*, 2024.

Asymptotic Analysis of the TS Estimator

Challenge: Unlike classical estimators (one limit: $N \to \infty$), the TS estimator involves *three* limits:

Training:	$M \to \infty$	number of parameter value draws
	$N \to \infty$	observation length per draw
Implementation:	$N \to \infty$	observation length at deployment

Result: $\hat{\theta}_{\text{TS}} \colon \mathbb{R}^N \times \mathbb{R}^M \to \mathcal{G} \circ h$ converges to $\bar{\theta}$ as $N, M \to \infty$

1. **Consistency:** $\hat{\theta}_0 = \bar{\theta}\Big((y_1^{(0)}, \ldots, y_N^{(0)})^\top\Big) \xrightarrow{\text{w.p.1}} \theta_0$ as $N \to \infty$
2. **Asymptotic Normality:** $\sqrt{N}(\hat{\theta}_0 - \theta_0) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ as $N \to \infty$

Key insight: TS is asymptotically consistent and normal under simplified assumptions.

Chapter 6

Non-Asymptotic Analysis of TS

B. Lakshminarayanan and C. R. Rojas. "Finite Sample Analysis of a Data-Driven Estimator". Submitted to *IEEE SPL*, 2026.

Finite-Sample Bounds for the TS Estimator

Goal: Non-asymptotic guarantee: $|\hat{\theta} - \theta_0| \leq \varepsilon$ with high probability.

Key idea: Use an IV regressor (two independent sequences per $\theta^{(i)}$) and decompose the error:

$$\hat{\boldsymbol{\beta}}_{\mathrm{IV}} - \boldsymbol{\beta}^* = \boldsymbol{B}_{N,M}^{-1} \boldsymbol{S}_{N,M}\, \boldsymbol{\beta}^* + \boldsymbol{B}_{N,M}^{-1}\, \boldsymbol{e}_{N,M}$$

where $\boldsymbol{B}_{N,M} \to \boldsymbol{B}_* := \mathbb{E}_\theta[g_*(\theta) g_*^\top(\theta)]$ as $M, N \to \infty$.

Result: For any $\varepsilon > 0$,

$$|\hat{\theta} - \theta_0| \leq \zeta + \varepsilon \|\boldsymbol{\beta}^*\|_2 + \alpha\, \frac{\varepsilon \|\boldsymbol{\beta}^*\|_2 + \alpha\zeta}{2\, \sigma_{\min}(\boldsymbol{B}_*)}$$

except with probability at most

$$4p^2 \exp\left[-\frac{M \sigma_{\min}^2(\boldsymbol{B}_*)}{128\, p^2 \alpha^4}\right] + 2p^2 \exp\left[-\frac{M\varepsilon^2}{32\alpha^4}\right] + \delta(N, \varepsilon)$$

Here ζ is the approximation bias: $\theta = {\boldsymbol{\beta}^*}^\top g_*(\theta) + r(\theta)$ with $|r(\theta)| \leq \zeta$.
Probability bound decays exponentially in M.

Chapter 7

Sampling from Jeffreys Prior

Y. Shi, **B. Lakshminarayanan**, and C. R. Rojas. "A Metropolis-Adjusted Langevin Algorithm for Sampling Jeffreys Prior". *IEEE CDC*, pages 2359–2364, 2025.

Diverse Training via Jeffreys Prior Sampling

Question: How to sample $\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(M)}$ to improve TS performance?

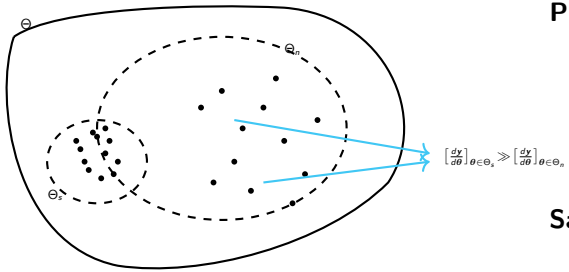

Θ_s: sensitive region (high FIM), needs *more* samples.
Θ_n: non-sensitive region (low FIM).
Uniform sampling ignores this geometry!

Proposed: Sample from Jeffreys prior

$$\pi_J(\boldsymbol{\theta}) \propto \sqrt{\det \boldsymbol{I}_{\boldsymbol{\theta}}}$$

- Concentrates tasks in Θ_s where data is most sensitive to $\boldsymbol{\theta}$

Sampling via Langevin SDE + MALA:

$$d\boldsymbol{\theta}_t = -\nabla_{\boldsymbol{\theta}} V(\boldsymbol{\theta}_t)\, dt + \sqrt{2}\, d\boldsymbol{W}_t$$

$V(\boldsymbol{\theta}) = -\frac{1}{2} \ln \det \boldsymbol{I}_{\boldsymbol{\theta}} \quad \Rightarrow$ stationary dist. $= \pi_J(\boldsymbol{\theta})$

When FIM is intractable (e.g. NLSS):

- Approximate $\boldsymbol{I}_{\boldsymbol{\theta}}$ via particle filter
- MH accept/reject corrects errors

Y. Shi, **B.L.**, C.R. Rojas. "A Metropolis-Adjusted Langevin Algorithm for Sampling Jeffreys Prior". *IEEE CDC*, 2025.

Numerical Example

Model: $y_i \overset{i.i.d.}{\sim} \text{Weibull}(\lambda, \kappa), \quad \lambda \in [1, 20], \quad \kappa \in [1, 20]$

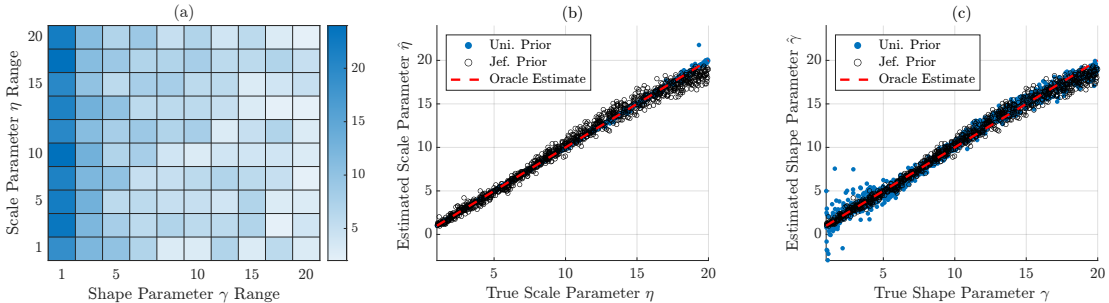

(a) Jeffreys prior concentrates samples at low κ (high FIM); (b) Both priors comparable for λ; (c) Jeffreys prior reduces variance and bias for $\kappa < 5$.

Part 2

Fine-Tuning a Simulation-Driven Estimator

The Out-of-Distribution Problem

TS is pretrained on $\Theta_p \subseteq \Theta$.

If $\boldsymbol{\theta}_0 \notin \Theta_p$:

- ψ has never seen features near $\boldsymbol{\theta}_0$
- Estimates would degrade significantly

Goal: Adapt the pretrained estimator using new observations.

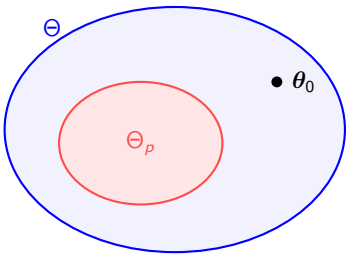

Chapter 8

Supervised Fine-Tuning of a Simulation-Driven Estimator

B. Lakshminarayanan, M. A. Guerrero, and C. R. Rojas. "Fine-Tuning a Simulation-Driven Estimator". *IEEE L-CSS*, vol. 9, pages 2975–2980, 2025.

Fine-Tuning TS: Supervised Approach

Architecture:
Trunk $\hat{\boldsymbol{w}}_T$: frozen
Heads $\hat{\boldsymbol{w}}_H$: retrained on $\mathcal{D}_{\text{FT}}$

Steps:

1. **Detect** OOD via whitened feature-space test
2. **Sample** cloud around $\hat{\boldsymbol{\theta}}_{\text{init}}$, simulate DT $\rightarrow \mathcal{D}_{\text{FT}}$
3. **Retrain** heads with a weighted loss

$$T_{\text{TS}}^{\text{FT}} = \psi(\varphi(\cdot);\, \hat{\boldsymbol{w}}_T,\, \hat{\boldsymbol{w}}_H^{\text{FT}})$$

Only *forward simulation* needed.

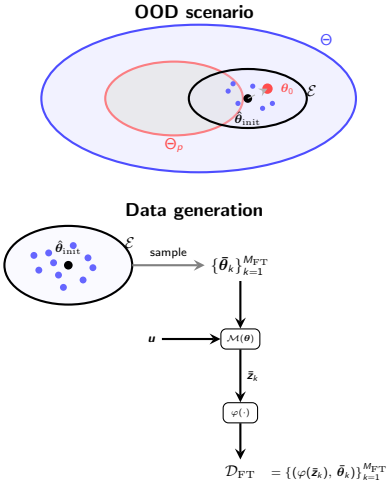

Chapter 9

Unsupervised Fine-Tuning of a Simulation-Driven Estimator

B. Lakshminarayanan and C. R. Rojas. "Trajectory-Level Self-Supervision for Simulation Driven Estimators". *IFAC Journal of Systems and Control*, vol. 35, page 100374, 2026.

Unsupervised Fine-Tuning via Trajectory Matching

Idea: Directly minimize simulation error: no synthetic dataset needed for retraining.

$$\ell_{\text{sim}}(\boldsymbol{\Omega}) = \frac{1}{K'N} \sum_{k=1}^{K'} \sum_{t=1}^{N} (\tilde{y}_{t,k} - y_t)^2$$

- Differentiable simulator needed
- Updates *all* weights $\boldsymbol{\Omega}$

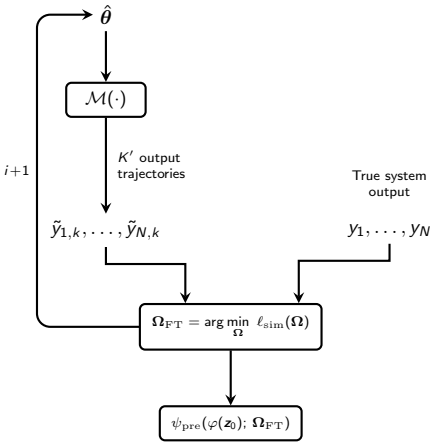

Numerical example

VDP

Pretrained: $\mu \in [0, 2.5]$. OOD 2: $\mu = 5.0$.

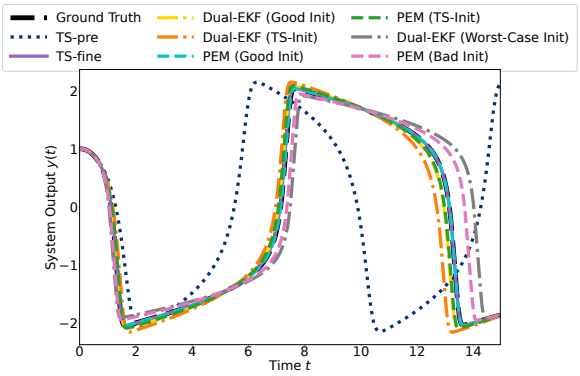

- TS-FT closely tracks ground truth
- TS-pre and poorly initialized baselines diverge

TS-pre: < 2 ms. Fine-tuning overhead incurred *only* upon OOD detection.

Part 3

Applications in Control

Chapter 10: The Two-Stage Approach for Controller Tuning

F. Dettù, **B. Lakshminarayanan**, S. Formentin, and C. R. Rojas. "From Data to Control: A Two-Stage Simulation-Based Approach". *ECC*, pages 3428–3433, 2024.

Chapter 11: Direct Learning of Controller Tuning Rules

B. Lakshminarayanan, F. Dettù, C. R. Rojas, and S. Formentin. "Inverse Supervised Learning of Controller Tuning Rules". *Automatica*, vol. 178, p. 112356, 2025.

Inverse Supervised Learning for Controller Tuning

Parameter	Trajectory	Design rule
$\tilde{\boldsymbol{\theta}}_1$	$\mathcal{D}_1^N = \{(u_k^{(1)}, y_k^{(1)})\}_{k=1}^N$	$C(\boldsymbol{\Psi}_1, \mathcal{M}(\tilde{\boldsymbol{\theta}}_1))$
$\tilde{\boldsymbol{\theta}}_2$	$\mathcal{D}_2^N = \{(u_k^{(2)}, y_k^{(2)})\}_{k=1}^N$	$C(\boldsymbol{\Psi}_2, \mathcal{M}(\tilde{\boldsymbol{\theta}}_2))$
$\vdots$	$\vdots$	$\vdots$
$\tilde{\boldsymbol{\theta}}_M$	$\mathcal{D}_M^N = \{(u_k^{(M)}, y_k^{(M)})\}_{k=1}^N$	$C(\boldsymbol{\Psi}_M, \mathcal{M}(\tilde{\boldsymbol{\theta}}_M))$

$$f^* = \arg\min_f \frac{1}{M} \sum_{i=1}^{M} \left\| \boldsymbol{\Psi}_i - f\left(y_1^{(i)}, u_1^{(i)}, \ldots, y_N^{(i)}, u_N^{(i)}\right) \right\|_2^2$$

TS-ISL: $f^* = \psi^* \circ \varphi$, φ: ARX compression, ψ^*: FFNN or GBM
Direct: f^*: CNN or WaveNet (end-to-end, no compression)

A Simple Case Study: Yaw-Rate Tracking

Objective: Achieve desired yaw-rate r under uncertain vehicle parameters.

Vehicle dynamics:

$$\dot{\beta} = -r - \frac{C_f \alpha_f}{M_{\text{veh}} v_x} - \frac{C_r \alpha_r}{M_{\text{veh}} v_x}$$

$$\dot{r} = -\frac{l_f C_f \alpha_f}{J_z} + \frac{l_r C_r \alpha_r}{J_z}$$

State: $\boldsymbol{x} = (\beta, r, \alpha_f, \alpha_r, z, \delta_f)^\top$
Input: $u = \delta_f^{\text{cmd}}$, Output: $y = r$

Uncertain parameters:

$$M_{\text{veh}} = M_0 + M_\delta, \quad M_\delta \in [0, 300] \text{ kg}$$
$$C_{f,r} = C_{f,r}^{\text{nom}} \mu_s, \quad \mu_s \in [0.3, 1.1]$$

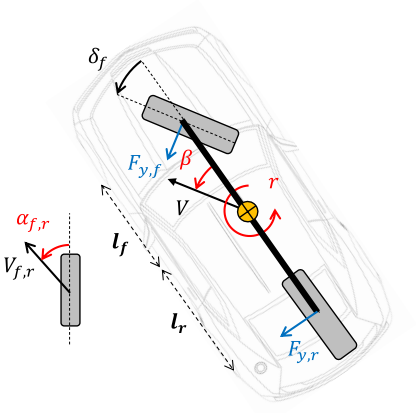

Control Design

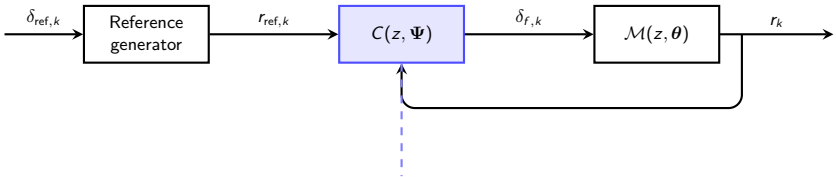

Controller parameters $\Psi = (K_p, K_i)$ are tuned to meet the loop-shaping specifications for each uncertain θ.

Results

Performance index

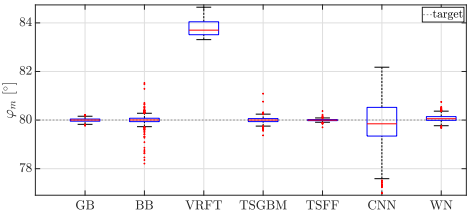

Computation (inference) time

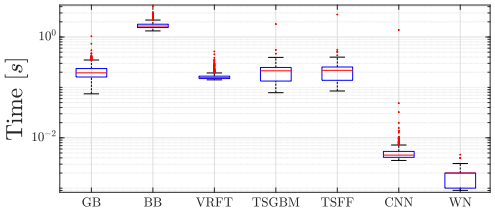

GB: Grey Box; BB: Black Box; VRFT: Virtual Reference Feedback Tuning.

Our methods: TSFF, TSGBM, CNN, WN.

For the same control performance, re-tuning time is significantly lower for our approach.

Part 4

Privacy in Point Estimation

Chapter 12: Differentially Private Point Estimation

B. Lakshminarayanan and C. R. Rojas. "A Unified Approach to Differentially Private Bayes Point Estimation". *IFAC WC*, pp. 8375–8380, 2023.

Need for Privacy in Point Estimates

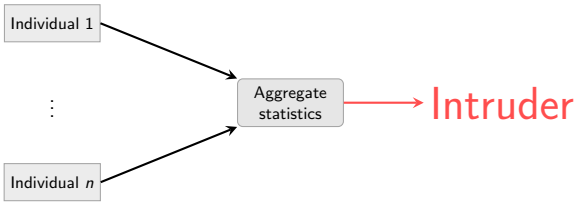

- Aggregate statistics: sample mean, sample covariance
- **Possible to infer an individual**[1]
- Differential privacy (DP) overcomes this challenge
  - Still "optimality" in the sense of balancing accuracy-privacy trade-off is challenging!

[1] Homer, N. et al. "Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays". *PLOS Genetics*, 2008.

Differential Privacy (DP)

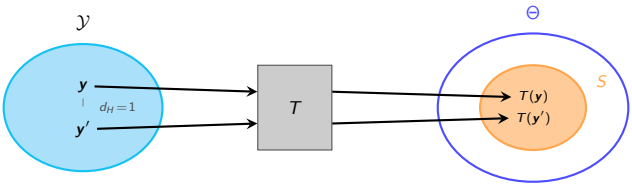

$$\mathbb{P}[T(\boldsymbol{y}) \in S] \leq e^{\varepsilon}\, \mathbb{P}[T(\boldsymbol{y}') \in S]$$

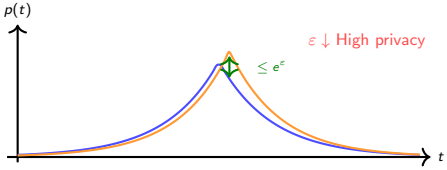

Smaller $\varepsilon \rightarrow$ output distributions become indistinguishable $\rightarrow$ stronger privacy.

Laplace Mechanism

Standard approach to achieve ε-DP: Add calibrated noise to the estimate.

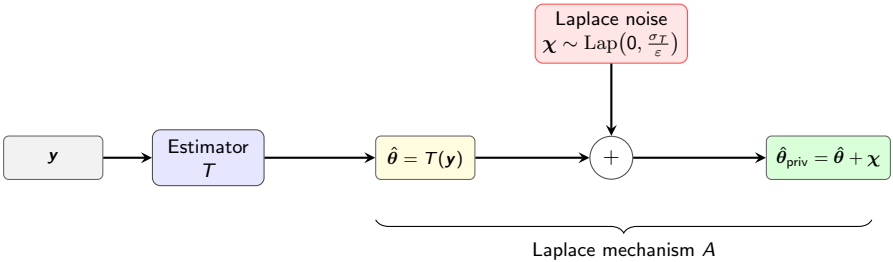

$\sigma_T = \sup_{d_H(\boldsymbol{y},\boldsymbol{y}')=1} \|T(\boldsymbol{y}) - T(\boldsymbol{y}')\|_1$ (ℓ_1-sensitivity)

Two questions:

1. Is the Laplace noise addition *optimal* for Bayes estimation under DP?
2. What if T has no closed-form expression to compute σ_T?

Unified Bayes Private Point (UBaPP) Estimator

Key idea: Integrate DP constraint directly into Bayes risk minimization.

$$\min_{\delta_{p,\varepsilon}} \quad \int_\Theta \int_{\mathcal{T}} \int_\Theta \ell(\theta, \tilde{\theta})\, \delta_{p,\varepsilon}(\tilde{\theta}|\boldsymbol{t})\, \nu_\theta(\boldsymbol{t})\, \pi(\theta)\, d\theta\, d\boldsymbol{t}\, d\tilde{\theta}$$

$$\text{s.t.} \quad \delta_{p,\varepsilon}(\tilde{\theta}|S(\boldsymbol{y})) \leq e^{\varepsilon}\, \delta_{p,\varepsilon}(\tilde{\theta}|S(\boldsymbol{y}'))\,, \quad \forall\, \tilde{\theta} \in \Theta, \quad \boldsymbol{y}, \boldsymbol{y}' \text{ s.t. } d_H(\boldsymbol{y}, \boldsymbol{y}') = 1 \quad \leftarrow \text{ DP}$$

$$\int_\Theta \delta_{p,\varepsilon}(\tilde{\theta}|\boldsymbol{t})\, d\tilde{\theta} = 1, \qquad \delta_{p,\varepsilon}(\tilde{\theta}|\boldsymbol{t}) \geq 0$$

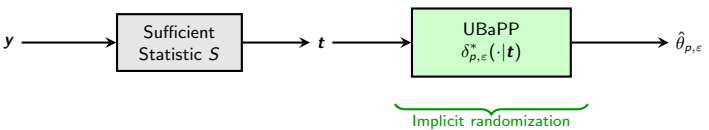

UBaPP is optimal by construction!

UBaPP: Finite Case

Finite case: When $|\Theta|, |\mathcal{T}|$ are finite, $\delta_{p,\varepsilon}$ is described by a matrix $\boldsymbol{P} \in \mathbb{R}^{|\Theta| \times |\mathcal{T}|}$.

UBaPP reduces to a **linear program**:

$$\begin{aligned} \min_{\boldsymbol{P}} \quad & \operatorname{tr}(\boldsymbol{V} \operatorname{diag}(\boldsymbol{\pi})\, \boldsymbol{L}\, \boldsymbol{P}) \\ \text{s.t.} \quad & P_{k,i} \le e^{\varepsilon} P_{k,i'}, \quad \forall\, k, \quad i, i' \text{ s.t. } d_H(\boldsymbol{y}_i, \boldsymbol{y}_{i'}) = 1 \\ & \mathbf{1}^\top \boldsymbol{P} = \mathbf{1}^\top, \qquad \boldsymbol{P} \ge 0 \end{aligned}$$

$\boldsymbol{L}$: loss matrix, $\boldsymbol{V}$: likelihood matrix, $\boldsymbol{\pi}$: prior. Solvable via CVXPY.

Privacy Results: MSE vs. ε

High privacy regime

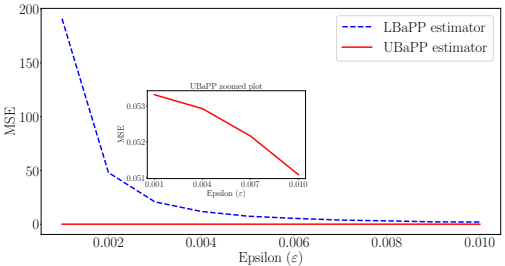

LBaPP MSE $\to \infty$ as $\varepsilon \to 0$.
UBaPP MSE stays bounded.

Low privacy regime

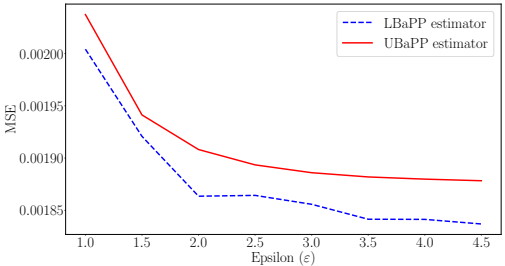

Both estimators perform similarly.
Small gap due to discretization of Θ.

High accuracy is achieved by our approach in the high privacy regime!

Privacy Results: Heat Maps of $\delta^*_{p,\varepsilon}$

$\varepsilon = 0.0001$

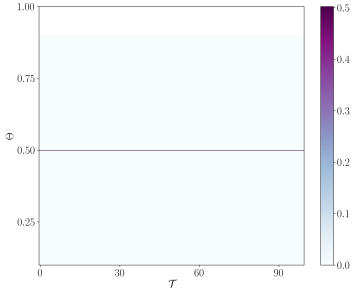

Deterministic at prior mean $\theta = 0.5$ for all t.

$\varepsilon = 0.01$

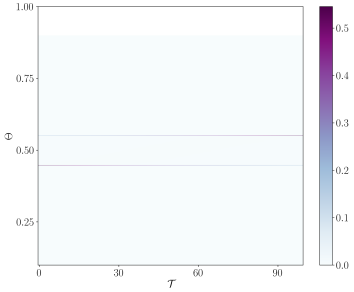

Implicit randomization: probability spread over Θ.

$\varepsilon = 5$

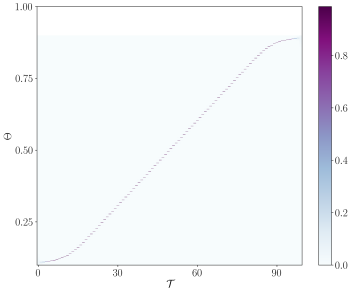

Converges to non-private Bayes (deterministic, data-dependent).

$\varepsilon \to 0$: data-independent $\longrightarrow$ $\varepsilon \to \infty$: data-dependent (non-private)

Summary

1. **Theoretical analysis of TS estimator**
   Statistical frameworks, consistency, asymptotic normality
2. **Fine-tuning for OOD adaptation**
   Supervised (OOD detection + retraining) and unsupervised (trajectory matching via differentiable simulator)
3. **Data-driven controller tuning**
   Meta-learning design rules via TS-ISL, CNN, and WaveNet
4. **Differentially private Bayes estimation**
   UBaPP: unified risk–privacy optimization, tractable LP for finite case

Thank You for Your Attention!

Braghadeesh Lakshminarayanan

blak@kth.se