Likelihood-Free Inference in SSMs with Unknown Dynamics

Alexander Aushev, Thong Tran, Henri Pesonen, Andrew Howes, Samuel Kaski

Aalto University, University of Oslo, The University of Manchester, University of Birmingham

NeurIPS 2021 Workshop on Deep Learning and Inverse Problems

Introduction

We introduce a method capable of likelihood-free state approximation and state prediction in discrete-time state-space models (SSMs), where observed measurements $\boldsymbol{x}_t \in \mathbb{R}^n$ are emitted given a series of simulator parameters $\theta_t \in \mathbb{R}^m$. The simulator parameters follow unknown transition dynamics h_θ: $\theta_{t+1} \sim h_\theta(\theta_{t+1}|\theta_t)$, and generate observations according to a black-box g_θ: $\boldsymbol{x}_t \sim g_\theta(\boldsymbol{x}_t \,|\, \theta_t)$.

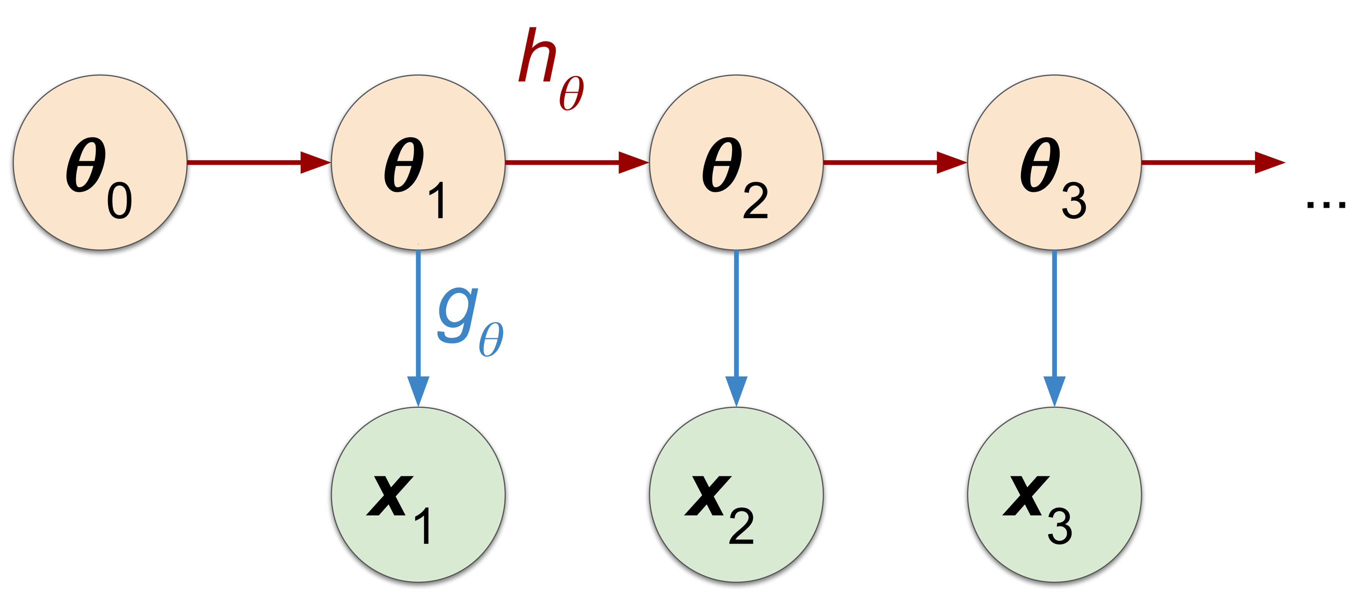

Methods

We follow a Bayesian optimization for LFI (BOLFI) approach [1], in which a Gaussian Process (GP) is used as a surrogate for discrepancy between observed measurements and simulated datasets. This way, the likelihood can be approximated through:

$$p(\mathbf{x}^*|\boldsymbol{\theta}) \propto F\left(\frac{\epsilon - \mu(\boldsymbol{\theta})}{\sqrt{\nu(\boldsymbol{\theta}) + \sigma^2}}\right) \quad (1)$$

where F is a normal CDF with mean 0 and variance 1, ε is a user-defined threshold, $\mu(.)$ and $\nu(.)$ are GP mean and variance respectively, and σ is a Gaussian likelihood noise.

To estimate state values θ_t, given $\boldsymbol{x}_t$, we employ a multi-objective surrogate $\tilde{\delta}_\theta$ for discrepancies, and then extract the LFI posterior with (1). At the same time, we form pairs of consecutive posterior samples and train a non-parametric surrogate for the state transition $\tilde{h}_\theta$, whose predictive posterior proposes candidates for future simulations:

$$p(\boldsymbol{\theta}_{T+1}|\mathbf{x}_T) = \int \hat{h}_{\boldsymbol{\theta}}(\boldsymbol{\theta}_{T+1}|\boldsymbol{\theta}_T)p(\boldsymbol{\theta}_T|\mathbf{x}_T)d\boldsymbol{\theta}_T \quad (2)$$

In the experiments, we use a Linear Model of Coregionalization (LMC) as $\tilde{\delta}_\theta$, and a Bayesian Neural Network (BNN) as $\tilde{h}_\theta$.

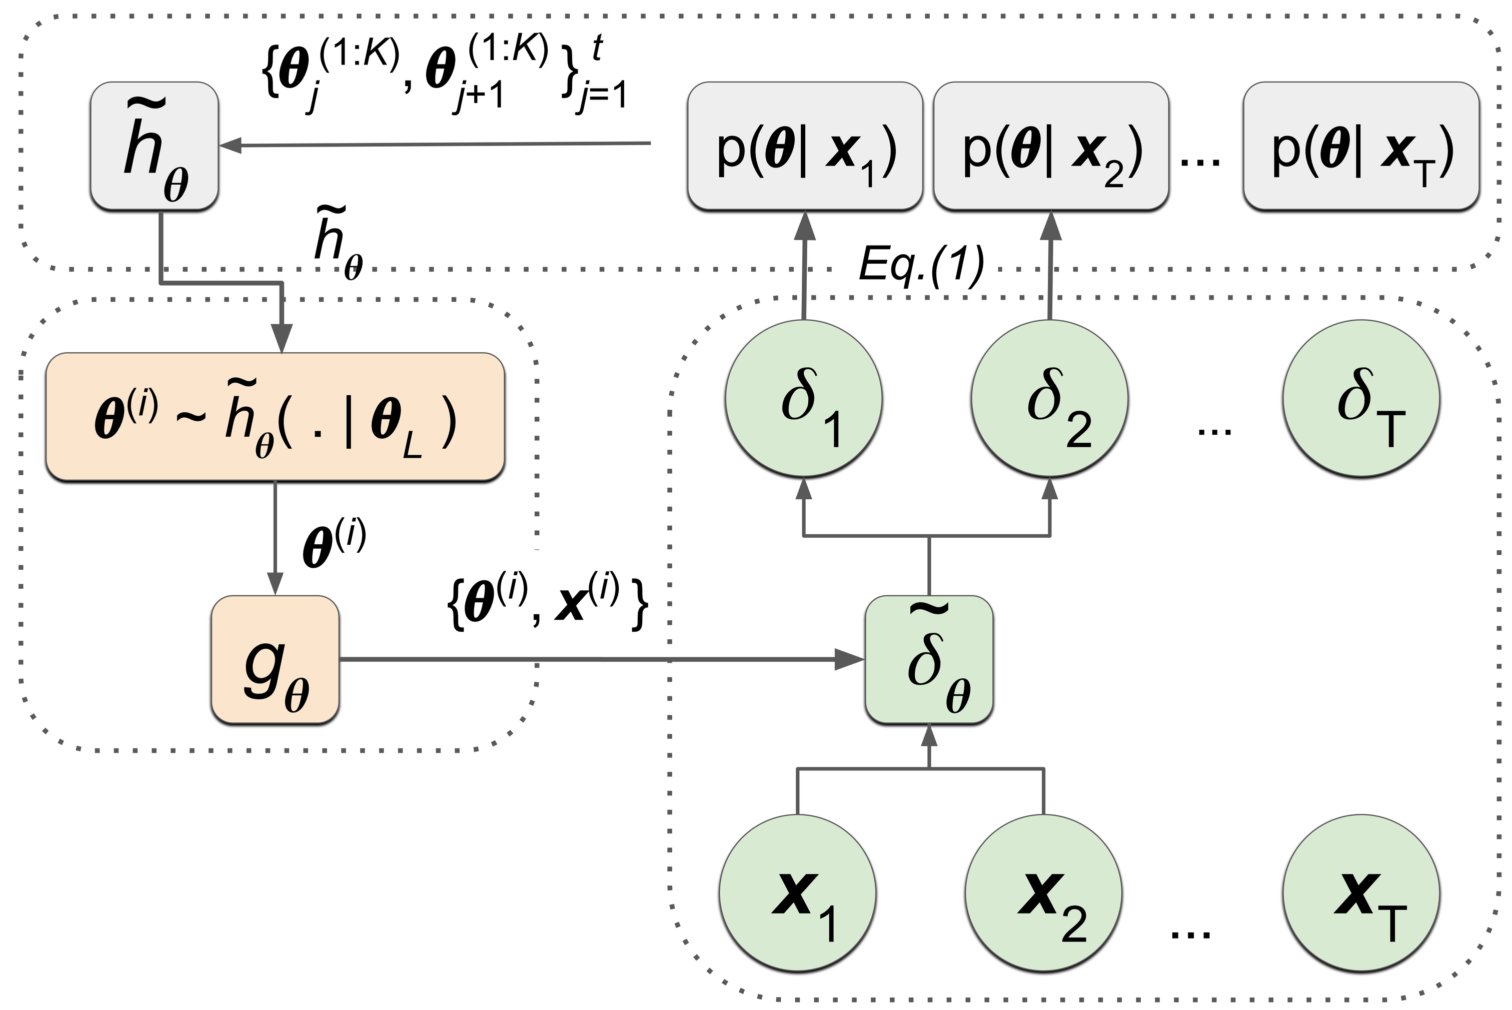

Experiments

We simulated behavioural data from humans that completed a certain task. For the UMAP task, the user evaluated dataset embeddings for a classification problem. During the Gaze task, the user searched for a target on a display. Our goal in the experiments was to track the changing parameters of user models and learn their dynamics.

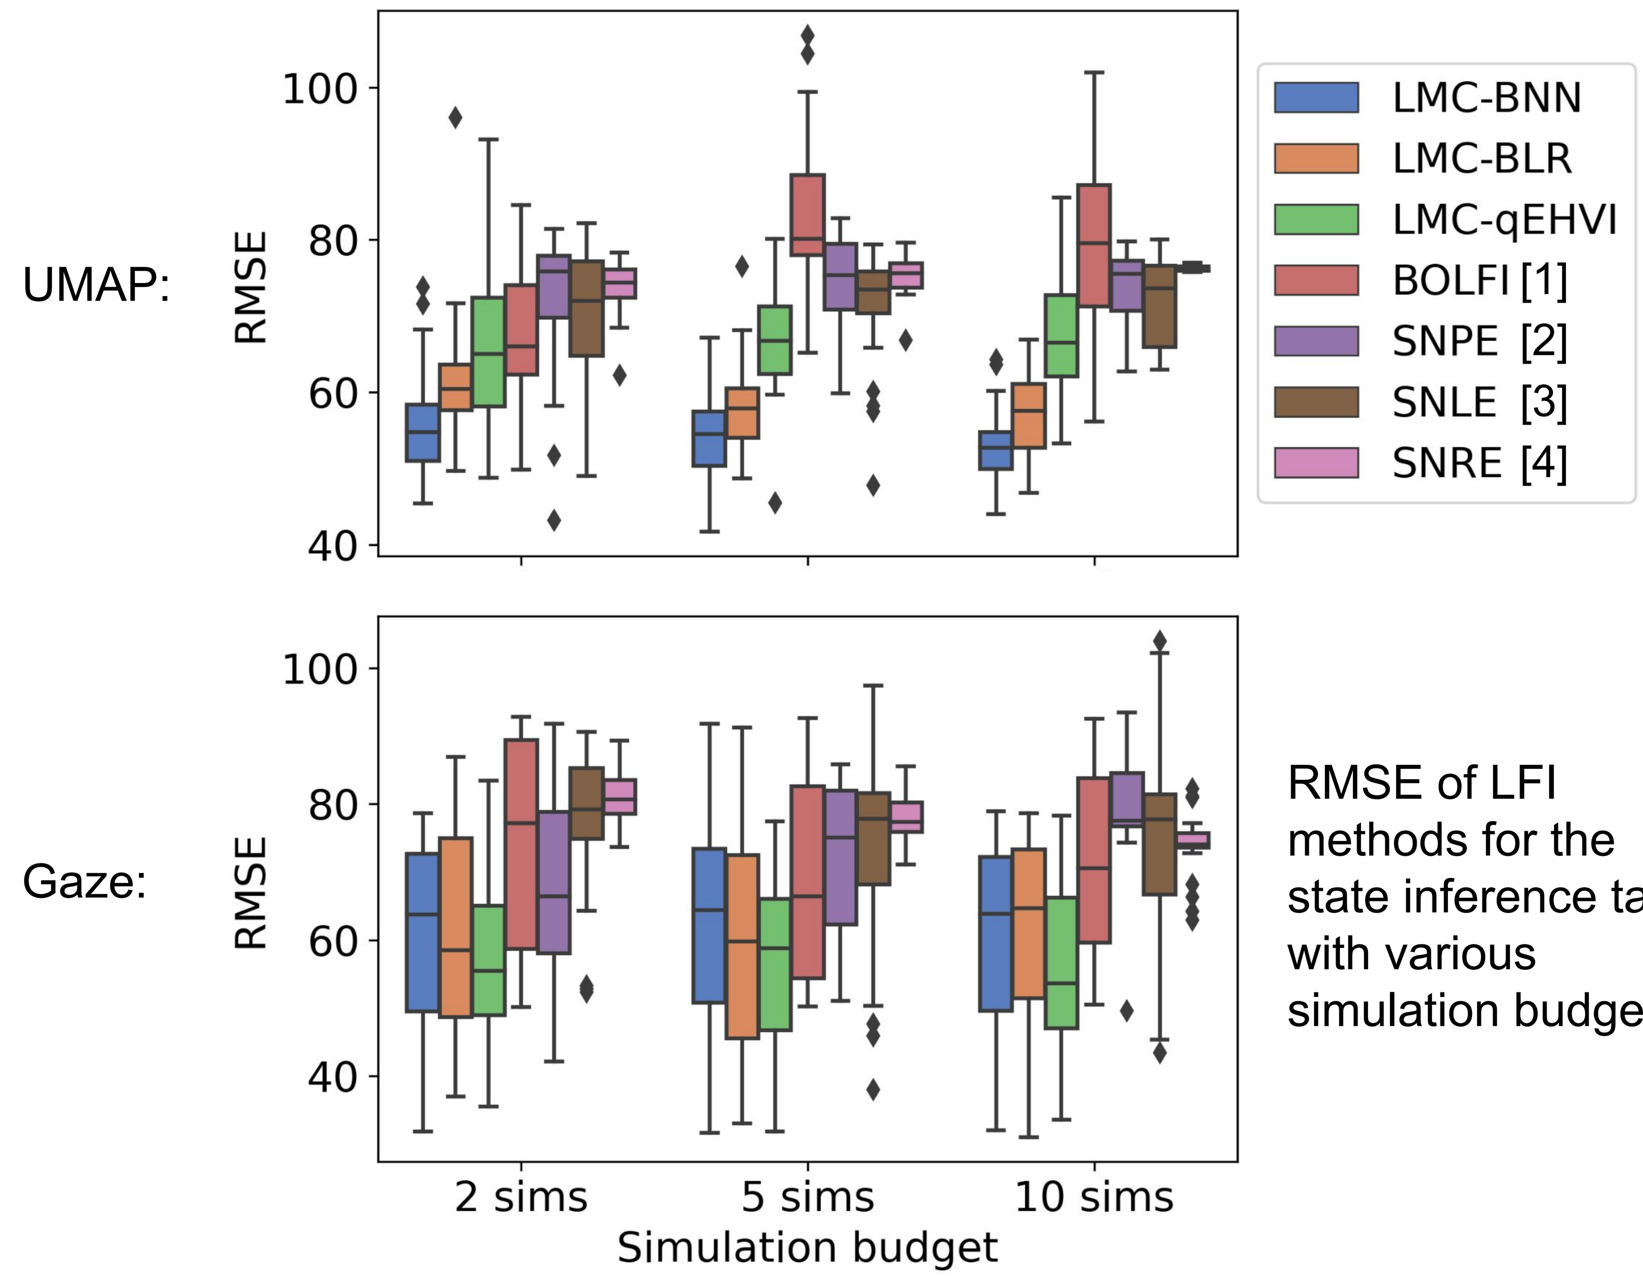

RMSE of LFI methods for the state inference task with various simulation budgets.

We also compare transition dynamics models (rows) in SSMs with tractable likelihoods (columns). The performance was measured with 95% confidence interval (CI) of the RMSE between sampled vs ground truth trajectories of length 50.

Methods	Linear Gaussian	Non-linear Non-Gaussian	Stochastic Volatility
LMC-BNN	205 ± 9	**165 ± 15**	135 ± 22
LMC-BLR	**64 ± 7**	258 ± 37	**100 ± 37**
GP-SSM	284 ± 71	2204 ± 1111	3206 ± 1175
PR-SSM	253 ± 68	610 ± 510	1378 ± 740

Take-Home Messages

- Current LFI methods assume simplified transition dynamics, which do not work with non-linear and non-Gaussian SSMs;
- We use samples from LFI posterior approximations to learn a deep learning model of transition dynamics, and use it to determine where to run simulations next.
- By leveraging time-series information, we improve upon current LFI methods for the state inference task.
- We demonstrate the proposed method is needed in a crucial case of user modelling, where user models are non-stationary because user's preferences and abilities change over time.

References

1. Gutmann MU, Corander J. Bayesian optimization for likelihood-free inference of simulator-based statistical models. Journal of Machine Learning Research. 2016.
2. Greenberg D, Nonnenmacher M, Macke J. Automatic posterior transformation for likelihood-free inference. InInternational Conference on Machine Learning 2019 May 24 (pp. 2404-2414). PMLR.
3. Papamakarios G, Sterratt D, Murray I. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In The 22nd International Conference on Artificial Intelligence and Statistics 2019 Apr 11 (pp. 837-848). PMLR.
4. Durkan C, Murray I, Papamakarios G. On contrastive learning for likelihood-free inference. International Conference on Machine Learning 2020 Nov 21 (pp. 2771-2781). PMLR.