Online Bayesian Experimental Design for Partially Observable Dynamical Systems

Sara Pérez-Vieites[1], Sahel Iqbal[2], Simo Särkkä[3,4], Dominik Baumann[4]

[1]University of Helsinki, [2]University of Oxford, [3]ELLIS Institute Finland, [4]Aalto University.

sara.perezvieites@helsinki.fi*, sahel.iqbal@stats.ox.ac.uk, simo.sarkka@aalto.fi, dominik.baumann@aalto.fi

Problem Statement

We consider **sequential Bayesian experimental design (BED)** for partially observable state-space models (SSMs) with a *design* input $\boldsymbol{\xi}_t$ that shapes data acquisition:

$$\text{(transition)} \quad \boldsymbol{x}_t \sim f(\boldsymbol{x}_t \mid \boldsymbol{x}_{t-1}, \boldsymbol{\theta}, \boldsymbol{\xi}_t), \tag{1}$$

$$\text{(observation)} \quad \boldsymbol{y}_t \sim g(\boldsymbol{y}_t \mid \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t). \tag{2}$$

- $\boldsymbol{\theta}$: static parameters – to learn
- $\boldsymbol{x}_t$: latent states – to learn (partially observed)
- $\boldsymbol{\xi}_t$: design variables – to optimize (online)

Optimization objective: *expected information gain* (EIG). We choose $\boldsymbol{\xi}_t$ to maximize the expected reduction in uncertainty about $\boldsymbol{\theta}$:

$$\mathrm{EIG}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{y}_t \mid \boldsymbol{\xi}_t, h_{t-1})}\Big[\mathcal{H}[p(\boldsymbol{\theta} \mid h_{t-1})] - \mathcal{H}[p(\boldsymbol{\theta} \mid \boldsymbol{y}_t, \boldsymbol{\xi}_t, h_{t-1})]\Big] \tag{3}$$

$$= \mathbb{E}_{p(\boldsymbol{\theta} \mid h_{t-1})\, p(\boldsymbol{y}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t)}\Big[\log \underbrace{p(\boldsymbol{y}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t)}_{\text{likelihood}} - \log \underbrace{p(\boldsymbol{y}_t \mid \boldsymbol{\xi}_t)}_{\text{evidence}}\Big]. \tag{4}$$

- $\mathcal{H}[\cdot]$ is the Shannon entropy,
- $h_{t-1} = \{\boldsymbol{\xi}_{1:t-1}, \boldsymbol{y}_{1:t-1}\}$ is the history up to time $t-1$,

$\longrightarrow$ The **optimal design** $\boldsymbol{\xi}_t^\star$ is the one that maximizes $\mathrm{EIG}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_t)$:

$$\boldsymbol{\xi}_t^\star = \arg\max_{\boldsymbol{\xi}_t \in \Omega} \; \mathrm{EIG}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_t) \tag{5}$$

Method: Online BED approach

Challenges:

(i) **Intractability and latent-state marginalizations:**

$$\text{(likelihood)} \quad p(\boldsymbol{y}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{x}_{0:t} \mid \boldsymbol{\theta}, h_{t-1})}\big[g(\boldsymbol{y}_t \mid \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t)\big],$$

$$\text{(evidence)} \quad p(\boldsymbol{y}_t \mid \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{\theta} \mid h_{t-1})\, p(\boldsymbol{x}_{0:t} \mid \boldsymbol{\theta}, h_{t-1})}\big[g(\boldsymbol{y}_t \mid \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t)\big].$$

(ii) **Sequential inference bottleneck:** $p(\boldsymbol{x}_{0:t}, \boldsymbol{\theta} \mid h_t)$ changes every step; naïve recomputation would replay the entire history each time.

Key idea:

Extend BED to partial observability by:

- deriving *new* Monte Carlo estimators of $\mathrm{EIG}_{\boldsymbol{\theta}}$ (and its gradient) that treat the latent-state integrals explicitly; and
- leveraging nested particle filters (NPFs) (Crisan and Míguez, 2018), that is online Bayesian inference methods, to *reuse* state–parameter particles and *avoid replaying past data*.

TL;DR

- **Problem:** identifying the system parameters $\boldsymbol{\theta}$ *online*, where latent states $\boldsymbol{x}_t$ evolve over time, observed via noisy measurements $\boldsymbol{y}_t$. These measurement (data acquisition) processes can be influenced by *designs* $\boldsymbol{\xi}_t$.
- **Proposed solution:** choose designs $\boldsymbol{\xi}_t$ by maximizing the *expected information gain* (EIG) about $\boldsymbol{\theta}$, marginalizing out latent states $\boldsymbol{x}_t$. Leverage a nested particle filters (NPF) to approximate the intractable posterior $p(\boldsymbol{x}_t, \boldsymbol{\theta} \mid h_t)$ *online*.

Algorithm Design selection at time t

Input: samples $\{\boldsymbol{\theta}_{t-1}^{(m)}, \boldsymbol{x}_{t-1}^{(m,n)}\}_{m=1,n=1}^{M,N} \sim p(\boldsymbol{x}_{t-1}, \boldsymbol{\theta} \mid h_{t-1})$;
initial design $\boldsymbol{\xi}_t^{(0)}$; stepsizes $\{\eta_k\}_{k=0}^{K-1}$
Output: samples $\{\boldsymbol{\theta}_t^{(m)}, \boldsymbol{x}_t^{(m,n)}\}_{m=1,n=1}^{M,N} \sim p(\boldsymbol{x}_t, \boldsymbol{\theta} \mid h_t)$; design $\boldsymbol{\xi}_t^\star$

for $k = 0, \dots, K-1$ **do** ▷ design optimization loop
  (1) Draw $(\tilde{\boldsymbol{y}}_t^{(\ell)}, \tilde{\boldsymbol{x}}_t^{(\ell)}, \boldsymbol{\theta}_{t-1}^{(\ell)})$ from
  $g(\cdot \mid \boldsymbol{x}_{t-1}, \boldsymbol{\theta}_{t-1}, \boldsymbol{\xi}_t^{(k)}) f(\cdot \mid \boldsymbol{x}_{t-1}, \boldsymbol{\theta}_{t-1}, \boldsymbol{\xi}_t^{(k)}) p(\boldsymbol{x}_{0:t-1}, \boldsymbol{\theta} \mid h_{t-1})$
  (2) Approximate $p(\tilde{\boldsymbol{y}}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t^{(k)})$, $p(\tilde{\boldsymbol{y}}_t \mid \boldsymbol{\xi}_t^{(k)})$, and $\widehat{\nabla}_{\boldsymbol{\xi}_t} \mathrm{EIG}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_t^{(k)})$
  (3) $\boldsymbol{\xi}_t^{(k+1)} \leftarrow \big(\boldsymbol{\xi}_t^{(k)} + \eta_k \widehat{\nabla}_{\boldsymbol{\xi}_t} \mathrm{EIG}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_t^{(k)})\big)$
end for
Set $\boldsymbol{\xi}_t^\star \leftarrow \boldsymbol{\xi}_t^{(K)}$; collect $\boldsymbol{y}_t$ ▷ observe new data

Jittering: $\boldsymbol{\theta}_{t|t-1}^{(m)} \sim \kappa_M(\cdot \mid \boldsymbol{\theta}_{t-1}^{(m)})$ ▷ update beliefs (NPF step)
for $m = 1$ to M **do**
  Propagate states: $\boldsymbol{x}_{t|t-1}^{(m,n)} \sim f(\cdot \mid \boldsymbol{x}_{t-1}^{(m,n)}, \boldsymbol{\theta}_{t|t-1}^{(m)}, \boldsymbol{\xi}_t^\star)$
  Update state weights $w_{\boldsymbol{x},t}^{(m,n)}$ and resample to obtain $\boldsymbol{x}_t^{(m,n)}$
end for
Update parameter weights $w_{\boldsymbol{\theta},t}^{(m)} = w_{\boldsymbol{\theta},t-1}^{(m)} \sum_{n=1}^{N} w_{\boldsymbol{x},t}^{(m,n)}$
Resample to obtain $\boldsymbol{\theta}_t^{(m)}$

Consistency of the EIG estimator

Theorem 1. Let $\hat{\mathcal{I}}(\boldsymbol{\xi}_t)$ denote the NMC estimator of the EIG, constructed using M parameter particles, N state particles per parameter, and L pseudo-observations. Under mild assumptions, for any t and $\boldsymbol{\xi}_t$

$$\hat{\mathcal{I}}(\boldsymbol{\xi}_t) \xrightarrow[L,M,N \to \infty]{\text{a.s.}} \mathcal{I}(\boldsymbol{\xi}_t).$$

Assumptions: small/regular jittering kernels, Lipschitz state posterior in $\boldsymbol{\theta}$, bounded positive likelihood, integrand of EIG is Lipschitz and square-integrable.

Example: Moving source location

State. A single source moves in the plane such that

$$\boldsymbol{x}_t = \begin{bmatrix} p_{x,t} \\ p_{y,t} \\ \phi_t \end{bmatrix} = \begin{bmatrix} p_{x,t-1} + v_x \cos(\phi_{t-1}) \\ p_{y,t-1} + v_y \sin(\phi_{t-1}) \\ \phi_{t-1} \end{bmatrix} + \boldsymbol{w}_t$$

where $\boldsymbol{w}_t \sim \mathcal{N}(\boldsymbol{0}, Q)$ and $\boldsymbol{\theta} = (v_x, v_y)$ are unknown parameters.

Observation. Sensors fixed at positions $\{\boldsymbol{s}_j\}_{j=1}^J \subset \mathbb{R}^2$ return a noisy log–intensity with distance attenuation and cardioid directivity,

$$\log y_{t,j} \mid \boldsymbol{x}_t, \boldsymbol{\xi}_t \sim \mathcal{N}\left(\log\left[b + \frac{\alpha_j}{m + \|\boldsymbol{p}_t - \boldsymbol{s}_j\|^2}\left(\frac{1 + \cos \Delta_{t,j}(\xi_{t,j})}{2}\right)\right], \sigma^2\right),$$

$$\Delta_{t,j}(\xi_{t,j}) = \xi_{t,j} - \operatorname{atan2}\big((\boldsymbol{p}_t - \boldsymbol{s}_j)_y, (\boldsymbol{p}_t - \boldsymbol{s}_j)_x\big).$$

Design. The *design* $\boldsymbol{\xi}_t = (\xi_{t,1}, \dots, \xi_{t,J})$ sets **sensor *orientations***.

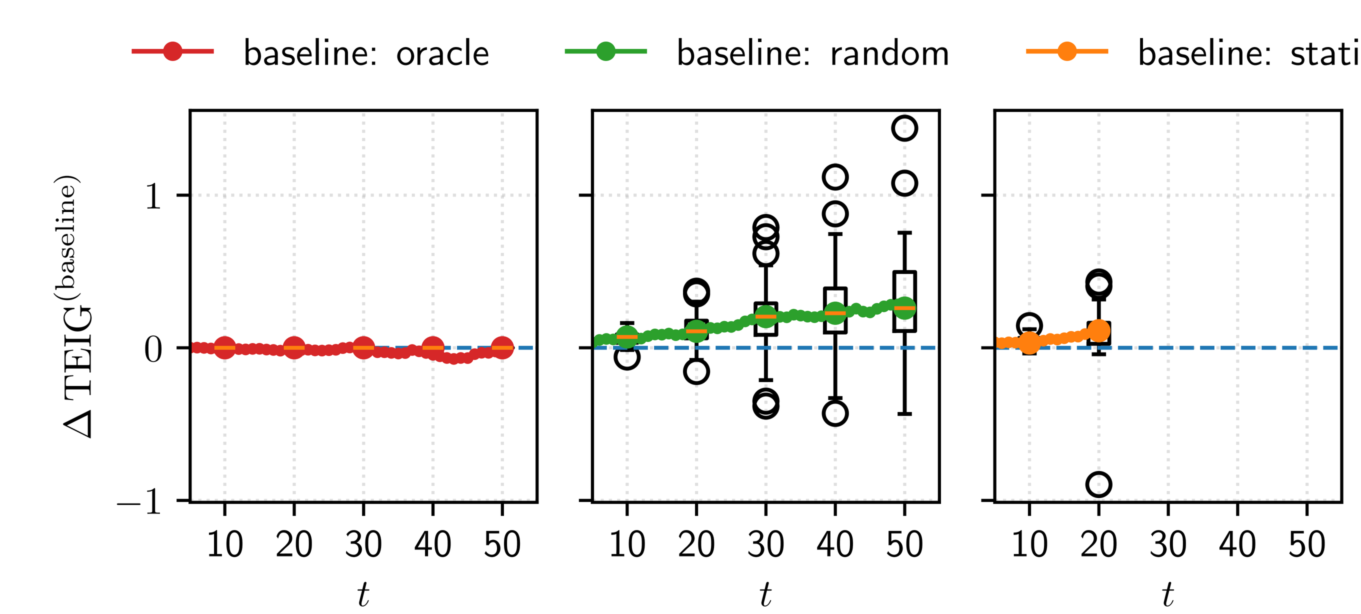

Assessment. We report *relative improvements*,

$$\Delta \mathrm{TEIG}^{(\text{baseline})} = \sum_{s=1}^{t}\left(\widehat{\mathrm{EIG}}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_s^{\text{BAD-PODS}}) - \widehat{\mathrm{EIG}}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_s^{\text{baseline}})\right),$$

against two baselines to highlight adaptive gains:

- *Random* baseline: random designs sampled uniformly from Ω.
- *Static* baseline: a non-adaptive version of our method where the full sequence $\boldsymbol{\xi}_{1:T}$ is optimized offline.
- *Oracle* baseline: evaluates EIG on a dense discretisation of the space (instead of gradient-based optimisation).

References

[1] D. Crisan and J. Míguez, "Nested particle filters for online parameter estimation in discrete-time state-space Markov models," *Bernoulli*, vol. 24, no. 4A, pp. 3039–3086, 2018.

[2] S. Pérez-Vieites, S. Iqbal, S. Särkkä, and D. Baumann, "Online Bayesian experimental design for partially observed dynamical systems," *arXiv preprint arXiv:2511.04403*, 2025.