Online Bayesian Experimental Design for Partially Observable Dynamical Systems

Sara Pérez-Vieites*, Sahel Iqbal, Simo Särkkä, Dominik Baumann

Department of Electrical Engineering and Automation, Aalto University, Finland.
Finnish Center for Artificial Intelligence (FCAI).
sara.perezvieites@aalto.fi*, sahel.iqbal@aalto.fi, 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 $\boldsymbol{\theta}$:

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

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

- $\mathcal{H}[\cdot]$ is the Shannon entropy,
- $h_{t-1} = \{\boldsymbol{\xi}_{1:t-1}, \boldsymbol{y}_{1:t-1}\}$ is the history up to time $t-1$,
- $p(\boldsymbol{y}_t\,|\,\boldsymbol{\theta},\boldsymbol{\xi}_t)$ is the likelihood, and $p(\boldsymbol{y}_t\,|\,\boldsymbol{\xi}_t)$ is the evidence.

Notation: To reduce clutter, all distributions are understood to condition on h_{t-1} unless shown explicitly.

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

Challenges & key idea

Challenges:

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

$$\textit{(likelihood)} \quad p(\boldsymbol{y}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{x}_{0:t}|\boldsymbol{\theta},h_{t-1})}\big[\, g(\boldsymbol{y}_t\,|\,\boldsymbol{x}_t,\boldsymbol{\theta},\boldsymbol{\xi}_t)\big], \tag{5}$$

$$\textit{(evidence)} \quad p(\boldsymbol{y}_t \mid \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{\theta}|h_{t-1})\,p(\boldsymbol{x}_{0:t}|\boldsymbol{\theta},h_{t-1})}\big[\, g(\boldsymbol{y}_t\,|\,\boldsymbol{x}_t,\boldsymbol{\theta},\boldsymbol{\xi}_t)\big]. \tag{6}$$

(ii) **Sequential inference bottleneck:** $p(\boldsymbol{x}_{0:t}, \boldsymbol{\theta}\,|\,h_t)$ 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) [1], that is **online Bayesian inference methods**, to *reuse* state–parameter particles and *avoid replaying past data*.

Assumption: f and g are differentiable w.r.t. $\boldsymbol{\xi}_t$ (for gradient-based design).

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] T. Rainforth, R. Cornish, H. Yang, A. Warrington, and F. Wood, "Tighter variational bounds are not necessarily better," in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.

[3] S. Iqbal, A. Corenflos, S. Särkkä, and H. Abdulsamad, "Nesting particle filters for experimental design in dynamical systems," in *International Conference on Machine Learning*, pp. 21047–21068, PMLR, 2024.

[4] S. Iqbal, H. Abdulsamad, S. Pérez-Vieites, S. Särkkä, and A. Corenflos, "Recursive nested filtering for efficient amortized Bayesian experimental design," *arXiv preprint arXiv:2409.05354*, 2024.

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 $p(\boldsymbol{x}_t, \boldsymbol{\theta}\,|\,h_t)$ *online.*
- **Control view:** the design $\boldsymbol{\xi}_t$ acts as a control input that influences the evolution of the latent states $\boldsymbol{x}_t$ and the observations $\boldsymbol{y}_t$, thereby affecting the information gained about the parameters $\boldsymbol{\theta}$.

Method: Online BED approach

Algorithm Per-step design selection at time t

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

Prediction (simulator): for each (m,n), propagate one step forward through f and g to get states and pseudo-observations $\{\tilde{\boldsymbol{x}}_t^{(m,n)}, \tilde{\boldsymbol{y}}_t^{(m,n)}\}$
for $k = 0, \ldots, K-1$ **do** ▷ design optimization loop
  Inner expectations: approximate $p(\tilde{\boldsymbol{y}}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t^{(k)})$ and $p(\tilde{\boldsymbol{y}}_t \mid \boldsymbol{\xi}_t^{(k)})$ (and their gradients)
  Gradient estimate: compute $\widehat{\nabla}_{\boldsymbol{\xi}_t}\mathrm{EIG}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_t^{(k)})$
  Ascent step: $\boldsymbol{\xi}_t^{(k+1)} \leftarrow \big(\boldsymbol{\xi}_t^{(k)} + \eta_k\, \widehat{\nabla}_{\boldsymbol{\xi}_t}\mathrm{EIG}_{\boldsymbol{\theta}}(\boldsymbol{\xi}_t^{(k)})\big)$
end for
Set $\boldsymbol{\xi}_t \leftarrow \boldsymbol{\xi}_t^{(K)}$; collect $\boldsymbol{y}_t$ ▷ observe new data
Update (one step of NPF): $p(\boldsymbol{x}_t, \boldsymbol{\theta}\,|\,h_t)$ ▷ update beliefs

Example: Moving source location

State. A single source moves in the plane with state $\boldsymbol{x}_t = (p_{x,t},\, p_{y,t},\, \phi_t)^\top$, such that

$$p_{x,t} = p_{x,t-1} + v_x \cos(\phi_{t-1}) + w_{x,t}, \tag{7}$$

$$p_{y,t} = p_{y,t-1} + v_y \sin(\phi_{t-1}) + w_{y,t}, \tag{8}$$

$$\phi_t = \phi_{t-1} + w_{\phi,t}, \tag{9}$$

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,

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

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

Design. The *design* $\xi_t = (\xi_{t,1}, \ldots, \xi_{t,J})$ sets **sensor *orientations***.