BEACON: Bayesian Experimental Design for Adaptive and Continual Learning in Non-stationary Environments

Sara Pérez-Vieites

Aalto University
Finnish Center for Artificial Intelligence (FCAI)

7 October 2025

Index

Sequential Bayesian experimental design (BED)

⇒ **Goal:** choose design ξ_t that maximizes the **expected information gain (EIG)** about parameters θ given history $h_{t-1} = \{\xi_{1:t-1}, y_{1:t-1}\}$.

$$\xi_t^\star = \arg\max_{\xi_t \in \Omega} \mathcal{I}(\xi_t)$$

EIG definition (information gain about parameters):

$$\mathcal{I}(\xi_t) = \mathbb{E}_{p(\theta, y_t | \xi_t, h_{t-1})}\left[\log \frac{\overbrace{p(y_t | \theta, \xi_t)}^{\text{likelihood}}}{\underbrace{p(y_t | \xi_t)}_{\text{evidence}}} \right]$$

$$= \mathbb{E}_{p(\theta, y_t | \xi_t, h_{t-1})}\left[\log \frac{p(y_t | \theta, \xi_t)}{\mathbb{E}_{p(\theta | h_{t-1})} p(y_t | \theta, \xi_t)} \right]$$

⇒ The likelihood $p(y_t | \theta, \xi_t)$ is available in closed-form.

Sequential Bayesian experimental design (BED)

⇒ **Goal:** choose design ξ_t that maximizes the **expected information gain (EIG)** about parameters θ given history $h_{t-1} = \{\xi_{1:t-1}, \mathbf{y}_{1:t-1}\}$.

$$\xi_t^\star = \arg\max_{\xi_t \in \Omega} \mathcal{I}(\xi_t)$$

EIG definition (information gain about parameters):

$$\mathcal{I}(\xi_t) = \mathbb{E}_{p(\theta, \mathbf{y}_t | \xi_t, h_{t-1})}\left[\log \frac{\overbrace{p(\mathbf{y}_t | \theta, \xi_t)}^{\text{likelihood}}}{\underbrace{p(\mathbf{y}_t | \xi_t)}_{\text{evidence}}} \right]$$

$$= \mathbb{E}_{p(\theta, \mathbf{y}_t | \xi_t, h_{t-1})}\left[\log \frac{p(\mathbf{y}_t | \theta, \xi_t)}{\mathbb{E}_{p(\theta | h_{t-1})} p(\mathbf{y}_t | \theta, \xi_t)} \right]$$

⇒ The likelihood $p(\mathbf{y}_t | \theta, \xi_t)$ is available in closed-form.

Sequential Bayesian experimental design (BED)

⇒ **Goal:** choose design ξ_t that maximizes the **expected information gain (EIG)** about parameters θ given history $h_{t-1} = \{\xi_{1:t-1}, \mathbf{y}_{1:t-1}\}$.

$$\xi_t^\star = \arg\max_{\xi_t \in \Omega} \mathcal{I}(\xi_t)$$

EIG definition (information gain about parameters):

$$\mathcal{I}(\xi_t) = \mathbb{E}_{p(\theta, \mathbf{y}_t | \xi_t, h_{t-1})}\left[\log \frac{\overbrace{p(\mathbf{y}_t | \theta, \xi_t)}^{\text{likelihood}}}{\underbrace{p(\mathbf{y}_t | \xi_t)}_{\text{evidence}}} \right]$$

$$= \mathbb{E}_{p(\theta, \mathbf{y}_t | \xi_t, h_{t-1})}\left[\log \frac{p(\mathbf{y}_t | \theta, \xi_t)}{\mathbb{E}_{p(\theta | h_{t-1})} p(\mathbf{y}_t | \theta, \xi_t)} \right]$$

⇒ The likelihood $p(\mathbf{y}_t | \theta, \xi_t)$ is available in closed-form.

State-space models (SSMs)

Many real systems are **partially observable dynamical systems**, where data are generated via latent states $\boldsymbol{x}_t$:

$$\begin{aligned} \text{(state)} \quad & \boldsymbol{x}_t \sim f(\boldsymbol{x}_t | \boldsymbol{x}_{t-1}, \boldsymbol{\theta}, \boldsymbol{\xi}_t), \\ \text{(observation)} \quad & \boldsymbol{y}_t \sim g(\boldsymbol{y}_t | \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t). \end{aligned}$$

EIG objective:

$$\mathcal{I}(\xi_t) = \mathbb{E}_{p(\theta, y_t, x_{0:t} | \xi_t, h_{t-1})} \left[\log \frac{p(y_t | \theta, \xi_t)}{p(y_t | \xi_t)} \right] \quad (1)$$

$$\text{(likelihood)} \quad p(y_t | \theta, \xi_t) = \mathbb{E}_{p(x_{0:t} | \theta, \xi_t)} [g(y_t | x_t, \theta, \xi_t)], \quad (2)$$

$$\text{(evidence)} \quad p(y_t | \xi_t) = \mathbb{E}_{p(\theta | h_{t-1}) p(x_{0:t} | \theta, \xi_t)} [g(y_t | x_t, \theta, \xi_t)]. \quad (3)$$

⇒ Requires marginalization over $x_{0:t}$ → intractable likelihood.

State-space models (SSMs)

Many real systems are **partially observable dynamical systems**, where data are generated via latent states $\boldsymbol{x}_t$:

$$
\begin{aligned}
\text{(state)} \quad & \boldsymbol{x}_t \sim f(\boldsymbol{x}_t | \boldsymbol{x}_{t-1}, \boldsymbol{\theta}, \boldsymbol{\xi}_t), \\
\text{(observation)} \quad & \boldsymbol{y}_t \sim g(\boldsymbol{y}_t | \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t).
\end{aligned}
$$

EIG objective:

$$
\mathcal{I}(\boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{\theta}, \boldsymbol{y}_t, \boldsymbol{x}_{0:t} | \boldsymbol{\xi}_t, h_{t-1})} \left[\log \frac{p(\boldsymbol{y}_t | \boldsymbol{\theta}, \boldsymbol{\xi}_t)}{p(\boldsymbol{y}_t | \boldsymbol{\xi}_t)} \right] \tag{1}
$$

$$
\text{(likelihood)} \quad p(\boldsymbol{y}_t | \boldsymbol{\theta}, \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{x}_{0:t} | \boldsymbol{\theta}, \boldsymbol{\xi}_t)} [g(\boldsymbol{y}_t | \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t)], \tag{2}
$$

$$
\text{(evidence)} \quad p(\boldsymbol{y}_t | \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{\theta} | h_{t-1}) p(\boldsymbol{x}_{0:t} | \boldsymbol{\theta}, \boldsymbol{\xi}_t)} [g(\boldsymbol{y}_t | \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t)]. \tag{3}
$$

⇒ Requires marginalization over $\boldsymbol{x}_{0:t}$ → intractable likelihood.

State-space models (SSMs)

Many real systems are **partially observable dynamical systems**, where data are generated via latent states $\boldsymbol{x}_t$:

$$\begin{aligned} \text{(state)} \quad & \boldsymbol{x}_t \sim f(\boldsymbol{x}_t \mid \boldsymbol{x}_{t-1}, \boldsymbol{\theta}, \boldsymbol{\xi}_t), \\ \text{(observation)} \quad & \boldsymbol{y}_t \sim g(\boldsymbol{y}_t \mid \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t). \end{aligned}$$

EIG objective:

$$\mathcal{I}(\boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{\theta}, \boldsymbol{y}_t, \boldsymbol{x}_{0:t} \mid \boldsymbol{\xi}_t, h_{t-1})} \left[\log \frac{p(\boldsymbol{y}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t)}{p(\boldsymbol{y}_t \mid \boldsymbol{\xi}_t)} \right] \tag{1}$$

$$\text{(likelihood)} \quad p(\boldsymbol{y}_t \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{x}_{0:t} \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t)} \left[g(\boldsymbol{y}_t \mid \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t) \right], \tag{2}$$

$$\text{(evidence)} \quad p(\boldsymbol{y}_t \mid \boldsymbol{\xi}_t) = \mathbb{E}_{p(\boldsymbol{\theta} \mid h_{t-1}) p(\boldsymbol{x}_{0:t} \mid \boldsymbol{\theta}, \boldsymbol{\xi}_t)} \left[g(\boldsymbol{y}_t \mid \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t) \right]. \tag{3}$$

$\Rightarrow$ Requires marginalization over $\boldsymbol{x}_{0:t}$ $\rightarrow$ intractable likelihood.

Sampling challenge and nested particle filters (NPFs)

Problem: Sample full trajectories $\boldsymbol{x}_{0:t}$ at each new time step—computational cost grows quadratically, $\mathcal{O}(t^2)$.

Goal: Maintain a joint posterior $p(\theta, \boldsymbol{x}_{0:t} | h_t)$ that can be updated recursively as new data arrive.

Approach: nested particle filters (NPFs)[1]

- Two-layer structure ($M \times N$ particles) to approximate $p(\mathrm{d}\boldsymbol{\theta}, \mathrm{d}\boldsymbol{x}_{0:t} | h_t)$.
- Updates one step forward — **no need to replay past data**, linear cost $\mathcal{O}(t)$.
- **Asymptotic convergence guarantees** as number of particles $M, N \to \infty$.

⇒ Recursive and consistent estimator of EIG.

[1] Crisan & Míguez (2018). *Nested particle filters for online parameter estimation in discrete-time state-space Markov models.* Bernoulli.

Sampling challenge and nested particle filters (NPFs)

Problem: Sample full trajectories $\boldsymbol{x}_{0:t}$ at each new time step—computational cost grows quadratically, $\mathcal{O}(t^2)$.

Goal: Maintain a joint posterior $p(\boldsymbol{\theta}, \boldsymbol{x}_{0:t} | h_t)$ that can be updated recursively as new data arrive.

Approach: nested particle filters (NPFs)[1]

- Two-layer structure ($M \times N$ particles) to approximate $p(\mathrm{d}\boldsymbol{\theta}, \mathrm{d}\boldsymbol{x}_{0:t} | h_t)$.
- Updates one step forward — **no need to replay past data**, linear cost $\mathcal{O}(t)$.
- **Asymptotic convergence guarantees** as number of particles $M, N \to \infty$.

$\Rightarrow$ Recursive and consistent estimator of EIG.

[1] Crisan & Míguez (2018). *Nested particle filters for online parameter estimation in discrete-time state-space Markov models.* Bernoulli.

Sampling challenge and nested particle filters (NPFs)

Problem: Sample full trajectories $\boldsymbol{x}_{0:t}$ at each new time step—computational cost grows quadratically, $\mathcal{O}(t^2)$.

Goal: Maintain a joint posterior $p(\boldsymbol{\theta}, \boldsymbol{x}_{0:t} \mid h_t)$ that can be updated recursively as new data arrive.

Approach: nested particle filters (NPFs)[1]

- Two-layer structure ($M \times N$ particles) to approximate $p(\mathrm{d}\boldsymbol{\theta}, \mathrm{d}\boldsymbol{x}_{0:t} \mid h_t)$.
- Updates one step forward — **no need to replay past data**, linear cost $\mathcal{O}(t)$.
- **Asymptotic convergence guarantees** as number of particles $M, N \to \infty$.

$\Rightarrow$ Recursive and consistent estimator of EIG.

[1] Crisan & Míguez (2018). *Nested particle filters for online parameter estimation in discrete-time state-space Markov models*. Bernoulli.

Algorithm for partially observable systems

Key idea: Combine **EIG optimization** with **online inference via NPFs**.

At each time t:

1. **Optimize design ξ_t** using stochastic gradient ascent on $\widehat{\mathcal{I}}(\xi_t)$.
2. **Collect data y_t** under optimized design.
3. **Update posterior** via nested particle filter (jitter, propagate, resample).

$\Rightarrow$ **Sequential design + inference with linear cost in T.**

Algorithm for partially observable systems

Key idea: Combine **EIG optimization** with **online inference via NPFs**.

At each time t:

1. **Optimize design $\boldsymbol{\xi}_t$** using stochastic gradient ascent on $\widehat{\mathcal{I}}(\boldsymbol{\xi}_t)$.
2. **Collect data $\boldsymbol{y}_t$** under optimized design.
3. **Update posterior** via nested particle filter (jitter, propagate, resample).

⇒ Sequential design + inference with linear cost in T.

Algorithm for partially observable systems

Key idea: Combine **EIG optimization** with **online inference via NPFs**.

At each time t:

1. **Optimize design** $\boldsymbol{\xi}_t$ using stochastic gradient ascent on $\widehat{\mathcal{I}}(\boldsymbol{\xi}_t)$.
2. **Collect data** $\boldsymbol{y}_t$ under optimized design.
3. **Update posterior** via nested particle filter (jitter, propagate, resample).

$\Rightarrow$ **Sequential design + inference with linear cost in T.**

Example: moving source model

⇒ **State dynamics.**

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \Delta t \begin{bmatrix} v_x \cos\phi_{t-1} \\ v_y \sin\phi_{t-1} \\ v_\phi \end{bmatrix} + \boldsymbol{\epsilon}_t$$

$\mathbf{x}_t = (p_{x,t}, p_{y,t}, \phi_t)^\top$, $\boldsymbol{\theta} = (v_x, v_y)^\top$, and $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$.

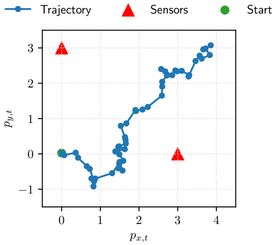

⇒ **Observation model.** J fixed sensors at positions (s_x^j, s_y^j):

$$\log y_{t,j} \,|\, \mathbf{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t \sim \mathcal{N}\big(\log \mu_{t,j},\, \sigma^2\big),$$

$$\mu_{t,j} = b + \underbrace{\frac{\alpha_j}{m + \|(p_{x,t}, p_{y,t}) - (s_x^j, s_y^j)\|^2}}_{\text{distance to source}} \underbrace{\left(\frac{1 + d\cos\Delta_{t,j}}{1+d}\right)^k}_{\text{directional sensitivity}},$$

- $\Delta_{t,j}$ is angular mismatch between sensor orientation and source direction.
- Design: sensor orientations $\boldsymbol{\xi}_t = (\xi_{t,1}, \ldots, \xi_{t,J})$.

Example: moving source model

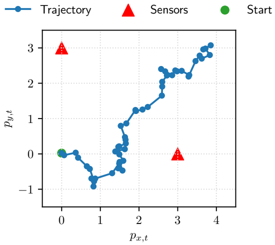

⇒ **State dynamics.**

$$\boldsymbol{x}_t = \boldsymbol{x}_{t-1} + \Delta t \begin{bmatrix} v_x \cos\phi_{t-1} \\ v_y \sin\phi_{t-1} \\ v_\phi \end{bmatrix} + \boldsymbol{\epsilon}_t$$

$\boldsymbol{x}_t = (p_{x,t}, p_{y,t}, \phi_t)^\top$, $\boldsymbol{\theta} = (v_x, v_y)^\top$, and $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$.

⇒ **Observation model.** J fixed sensors at positions (s_x^j, s_y^j):

$$\log y_{t,j} \,|\, \boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\xi}_t \sim \mathcal{N}\left(\log \mu_{t,j},\, \sigma^2\right),$$

$$\mu_{t,j} = b + \frac{\alpha_j}{\underbrace{m + \|(p_{x,t}, p_{y,t}) - (s_x^j, s_y^j)\|^2}_{\text{distance to source}}} \underbrace{\left(\frac{1 + d\cos\Delta_{t,j}}{1 + d}\right)^k}_{\text{directional sensitivity}},$$

- $\Delta_{t,j}$ is angular mismatch between sensor orientation and source direction.
- Design: sensor orientations $\boldsymbol{\xi}_t = (\xi_{t,1}, \ldots, \xi_{t,J})$.

Example: moving source model

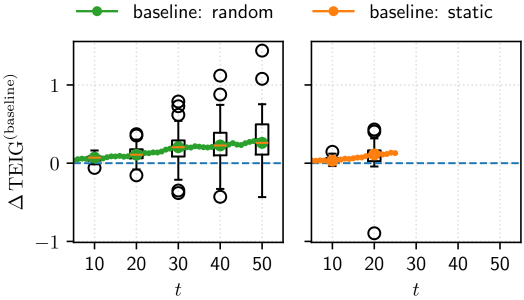

- $\Delta\textbf{TEIG}^{(\text{baseline})} = \sum_{\tau=1}^{t}\left(\widehat{\mathcal{I}}(\boldsymbol{\xi}_\tau^\star) - \widehat{\mathcal{I}}(\boldsymbol{\xi}_\tau^{(\text{baseline})})\right)$
- Average over 50 seeds.
- Random = random designs.
- Static = static BED version of our approach.

⇒ **Advantage over baselines grows with t.**

Example: moving source model

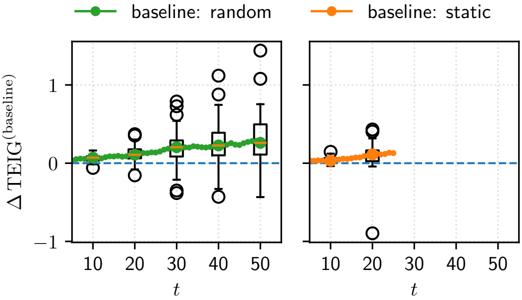

- $\Delta\textbf{TEIG}^{(\text{baseline})} = \sum_{\tau=1}^{t}\left(\widehat{\mathcal{I}}(\boldsymbol{\xi}_\tau^\star) - \widehat{\mathcal{I}}(\boldsymbol{\xi}_\tau^{(\text{baseline})})\right)$
- Average over 50 seeds.
- Random = random designs.
- Static = static BED version of our approach.

⇒ **Advantage over baselines grows with t.**

Summary

- Introduced a Bayesian experimental design framework for **partially observable dynamical systems**.
- Derived **recursive EIG and gradient estimators** using **nested particle filters** for online optimization and inference.

- Pérez-Vieites, S., Iqbal, S., Särkkä, S., & Baumann, D. *Online Bayesian experimental design for partially observable dynamical systems.* Submitted.

Summary

- Introduced a Bayesian experimental design framework for **partially observable dynamical systems**.
- Derived **recursive EIG and gradient estimators** using **nested particle filters** for online optimization and inference.

- Pérez-Vieites, S., Iqbal, S., Särkkä, S., & Baumann, D. *Online Bayesian experimental design for partially observable dynamical systems*. Submitted.

Index

Vision: Adaptive and Robust BED

Vision: Extend BED beyond short, controlled experiments to realistic deployments.

Three complementary directions:

- **Objective 1:** Continual adaptation.
- **Objective 2:** Non-ergodic dynamics.
- **Objective 3:** Non-stationary dynamics.

Vision: Adaptive and Robust BED

Vision: Extend BED beyond short, controlled experiments to realistic deployments.

Three complementary directions:

- **Objective 1:** Continual adaptation.
- **Objective 2:** Non-ergodic dynamics.
- **Objective 3:** Non-stationary dynamics.

Objective 1: Continual adaptation

Most BED formulations assume:

- **Finite horizon:** small, fixed number of experiments ($T \ll \infty$).

Problem: In long runs, design policies degrade[2]; static BED infeasible due to high dimension in the design space.

- **Changing environments:** new conditions emerge that fixed policies cannot adapt to.
- **Complex "big worlds":** even stationary systems can appear non-stationary when high-dimensional or heavy-tailed.

[2] Ivanova et al., (2024). *Step-dad: Semi-amortized policy-based Bayesian experimental design.* ICLR Workshop on Data-centric Machine Learning Research (DMLR).

Objective 1: Continual adaptation

Most BED formulations assume:

- **Finite horizon:** small, fixed number of experiments ($T \ll \infty$).

Problem: In long runs, design policies degrade[2]; static BED infeasible due to high dimension in the design space.

- **Changing environments:** new conditions emerge that fixed policies cannot adapt to.
- **Complex "big worlds":** even stationary systems can appear non-stationary when high-dimensional or heavy-tailed.

[2] Ivanova et al., (2024). *Step-dad: Semi-amortized policy-based Bayesian experimental design.* ICLR Workshop on Data-centric Machine Learning Research (DMLR).

Objective 1: Continual adaptation

Most BED formulations assume:

- **Finite horizon:** small, fixed number of experiments ($T \ll \infty$).

Problem: In long runs, design policies degrade[2]; static BED infeasible due to high dimension in the design space.

- **Changing environments:** new conditions emerge that fixed policies cannot adapt to.
- **Complex "big worlds":** even stationary systems can appear non-stationary when high-dimensional or heavy-tailed.

[2] Ivanova et al., (2024). *Step-dad: Semi-amortized policy-based Bayesian experimental design.* ICLR Workshop on Data-centric Machine Learning Research (DMLR).

Objective 1: Continual adaptation

Most BED formulations assume:

- **Finite horizon:** small, fixed number of experiments ($T \ll \infty$).

Problem: In long runs, design policies degrade[2]; static BED infeasible due to high dimension in the design space.

- **Changing environments:** new conditions emerge that fixed policies cannot adapt to.
- **Complex "big worlds":** even stationary systems can appear non-stationary when high-dimensional or heavy-tailed.

[2] Ivanova et al., (2024). *Step-dad: Semi-amortized policy-based Bayesian experimental design.* ICLR Workshop on Data-centric Machine Learning Research (DMLR).

Objective 1: Continual adaptation

Goal: Given a design policy, π_ϕ, adapt policy parameters ϕ over time while preserving critical knowledge.

Challenge: *stability–plasticity dilemma*[3]
too stable → no adaptation; too adaptive → forgetting.

Idea: regularisation-based CL

- **Elastic weight consolidation (EWC)**[4]: weight regularisation using Fisher information,

$$\frac{\lambda}{2}\sum_i F_i(\phi_{i,t} - \phi_{i,t-1})^2$$

- **Variational continual learning (VCL)**[5]: variational inference penalty via

$$\mathrm{KL}\big(q_t(\phi) \,\|\, q_{t-1}(\phi)\big)$$

⇒ Avoiding high memory cost of replay buffers.

[3] Wang et al., (2024). *A comprehensive survey of continual learning: Theory, method and application.* TPAMI.
[4] Kirkpatrick et al., (2017). *Overcoming catastrophic forgetting in neural networks.* Proc. Nat. Acad. of Sci.
[5] Nguyen et al., (2018). *Variational continual learning.* ICLR.

Objective 1: Continual adaptation

Goal: Given a design policy, π_ϕ, adapt policy parameters ϕ over time while preserving critical knowledge.

Challenge: *stability–plasticity dilemma*[3]
too stable → no adaptation; too adaptive → forgetting.

Idea: regularisation-based CL

- **Elastic weight consolidation (EWC)**[4]: weight regularisation using Fisher information,

$$\frac{\lambda}{2}\sum_i F_i(\phi_{i,t} - \phi_{i,t-1})^2$$

- **Variational continual learning (VCL)**[5]: variational inference penalty via

$$\mathrm{KL}\big(q_t(\phi)\,\|\,q_{t-1}(\phi)\big)$$

⇒ Avoiding high memory cost of replay buffers.

[3] Wang et al., (2024). *A comprehensive survey of continual learning: Theory, method and application.* TPAMI.
[4] Kirkpatrick et al., (2017). *Overcoming catastrophic forgetting in neural networks.* Proc. Nat. Acad. of Sci.
[5] Nguyen et al., (2018). *Variational continual learning.* ICLR.

Objective 1: Continual adaptation

Goal: Given a design policy, π_ϕ, adapt policy parameters ϕ over time while preserving critical knowledge.

Challenge: *stability–plasticity dilemma*[3]
too stable → no adaptation;   too adaptive → forgetting.

Idea: regularisation-based CL

- **Elastic weight consolidation (EWC)**[4]: weight regularisation using Fisher information,

$$\frac{\lambda}{2}\sum_i F_i(\phi_{i,t} - \phi_{i,t-1})^2$$

- **Variational continual learning (VCL)**[5]: variational inference penalty via

$$\mathrm{KL}\big(q_t(\phi)\,\|\,q_{t-1}(\phi)\big)$$

⇒ Avoiding high memory cost of replay buffers.

[3] Wang et al., (2024). *A comprehensive survey of continual learning: Theory, method and application.* TPAMI.
[4] Kirkpatrick et al., (2017). *Overcoming catastrophic forgetting in neural networks.* Proc. Nat. Acad. of Sci.
[5] Nguyen et al., (2018). *Variational continual learning.* ICLR.

Objective 1: Continual adaptation

Goal: Given a design policy, π_ϕ, adapt policy parameters ϕ over time while preserving critical knowledge.

Challenge: *stability–plasticity dilemma*[3]
too stable → no adaptation; too adaptive → forgetting.

Idea: regularisation-based CL

- **Elastic weight consolidation (EWC)**[4]: weight regularisation using Fisher information,

$$\frac{\lambda}{2}\sum_i F_i(\phi_{i,t} - \phi_{i,t-1})^2$$

- **Variational continual learning (VCL)**[5]: variational inference penalty via

$$\mathrm{KL}\big(q_t(\phi) \,\|\, q_{t-1}(\phi)\big)$$

⇒ Avoiding high memory cost of replay buffers.

[3] Wang et al., (2024). *A comprehensive survey of continual learning: Theory, method and application.* TPAMI.
[4] Kirkpatrick et al., (2017). *Overcoming catastrophic forgetting in neural networks.* Proc. Nat. Acad. of Sci.
[5] Nguyen et al., (2018). *Variational continual learning.* ICLR.

Objective 2: Non-ergodic dynamics

Most BED formulations assume:

- **Ergodicity:** time averages ≈ ensemble averages.

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} y_{\cdot,t} = \lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} y_{i,\cdot}.$$

Problem: When ergodicity breaks, incremental utilities misalign with total information gain. Standard BED objectives become unreliable.

- Multimodal or heavy-tailed observations → trajectories get trapped in one mode.
- Irreversible or "dead-end" states (e.g. stuck robots, terminated experiments).

Objective 2: Non-ergodic dynamics

Most BED formulations assume:

- **Ergodicity:** time averages ≈ ensemble averages.

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} y_{\cdot,t} = \lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} y_{i,\cdot}.$$

Problem: When ergodicity breaks, incremental utilities misalign with total information gain. Standard BED objectives become unreliable.

- Multimodal or heavy-tailed observations → trajectories get trapped in one mode.
- Irreversible or "dead-end" states (e.g. stuck robots, terminated experiments).

Objective 2: Non-ergodic dynamics

Most BED formulations assume:

- **Ergodicity:** time averages ≈ ensemble averages.

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} y_{\cdot,t} = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} y_{i,\cdot}.$$

Problem: When ergodicity breaks, incremental utilities misalign with total information gain. Standard BED objectives become unreliable.

- Multimodal or heavy-tailed observations → trajectories get trapped in one mode.
- Irreversible or "dead-end" states (e.g. stuck robots, terminated experiments).

Objective 2: Non-ergodic dynamics

Goal: Design utilities that remain reliable when ergodicity fails.

Idea: Learn transformations of incremental utilities that restore alignment between expected and time-averaged values.[6]

- **Detect** and diagnose loss of ergodicity during operation.
- **Learn transformations** $\mathcal{T}(U_t)$ that make incremental EIG ergodic again:

$$\mathbb{E}[\mathcal{T}(U_t)] \approx \frac{1}{T}\sum_t \mathcal{T}(U_t)$$

[6] Baumann et al. (2025). *Reinforcement learning with non-ergodic reward increments.* TMLR.

Objective 2: Non-ergodic dynamics

Goal: Design utilities that remain reliable when ergodicity fails.

Idea: Learn transformations of incremental utilities that restore alignment between expected and time-averaged values.[6]

- **Detect** and diagnose loss of ergodicity during operation.
- **Learn transformations** $\mathcal{T}(U_t)$ that make incremental EIG ergodic again:

$$\mathbb{E}[\mathcal{T}(U_t)] \approx \frac{1}{T}\sum_t \mathcal{T}(U_t)$$

[6] Baumann et al. (2025). *Reinforcement learning with non-ergodic reward increments.* TMLR.

Objective 3: Non-stationary dynamics

Most BED formulations assume:

- **Stationary model:** known and fixed likelihood $p(\boldsymbol{y}_t | \boldsymbol{\theta}, \boldsymbol{\xi}_t)$.

Problem: Static models become misspecified[7] as the environment evolves.

- **Parameter drift:** gradual changes in system behaviour (e.g. component wear, patient response evolution).
- **Regime switching:** abrupt transitions between modes (e.g. equipment faults, environment changes).

Example: *Industrial prognostics* — from slow degradation to sudden faults.

[7] Forster et al. (2025). *Improving Robustness to Model Misspecification in Bayesian Experimental Design.* Symp. Advances in Approximate Bayesian Inference.

Objective 3: Non-stationary dynamics

Most BED formulations assume:

- **Stationary model:** known and fixed likelihood $p(\boldsymbol{y}_t|\boldsymbol{\theta}, \boldsymbol{\xi}_t)$.

Problem: Static models become misspecified[7] as the environment evolves.

- **Parameter drift:** gradual changes in system behaviour (e.g. component wear, patient response evolution).
- **Regime switching:** abrupt transitions between modes (e.g. equipment faults, environment changes).

Example: *Industrial prognostics* — from slow degradation to sudden faults.

[7] Forster et al. (2025). *Improving Robustness to Model Misspecification in Bayesian Experimental Design.* Symp. Advances in Approximate Bayesian Inference.

Objective 3: Non-stationary dynamics

Most BED formulations assume:

- **Stationary model:** known and fixed likelihood $p(\boldsymbol{y}_t | \boldsymbol{\theta}, \boldsymbol{\xi}_t)$.

Problem: Static models become misspecified[7] as the environment evolves.

- **Parameter drift:** gradual changes in system behaviour (e.g. component wear, patient response evolution).
- **Regime switching:** abrupt transitions between modes (e.g. equipment faults, environment changes).

Example: *Industrial prognostics* — from slow degradation to sudden faults.

[7] Forster et al. (2025). *Improving Robustness to Model Misspecification in Bayesian Experimental Design.* Symp. Advances in Approximate Bayesian Inference.

Objective 3: Non-stationary dynamics

Goal: Enable BED under evolving dynamics, maintaining model validity and design relevance

Approach: Incorporate ideas from Bayesian filtering and changepoint detection[8].

- **Gradual drift:** latent parameters θ_t evolve via transition $p(\theta_t | \theta_{t-1})$ (online filtering).
- **Regime switching:** latent mode ψ_t with transition $p(\psi_t | \psi_{t-1})$ enables changepoint-aware designs.

[8] Duran-Martin (2025). *Adaptive, robust and scalable Bayesian filtering for online learning.* PhD Thesis, Queen Mary University of London.

Objective 3: Non-stationary dynamics

Goal: Enable BED under evolving dynamics, maintaining model validity and design relevance

Approach: Incorporate ideas from Bayesian filtering and changepoint detection[8].

- **Gradual drift:** latent parameters $\boldsymbol{\theta}_t$ evolve via transition $p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_{t-1})$ (online filtering).
- **Regime switching:** latent mode ψ_t with transition $p(\psi_t | \psi_{t-1})$ enables changepoint-aware designs.

[8] Duran-Martin (2025). *Adaptive, robust and scalable Bayesian filtering for online learning.* PhD Thesis, Queen Mary University of London.

Summary

Towards adaptive, robust, and realistic Bayesian experimental design (BED).

- **Objective 1: Continual adaptation** — adapt policies over long deployments without retraining.
- **Objective 2: Non-ergodic dynamics** — reliable/robust objectives under non-ergodic dynamics.
- **Objective 3: Non-stationary dynamics** — maintain validity under evolving environments.

Thank you!

References

- **Crisan & Míguez (2018).** *Nested particle filters for online parameter estimation in discrete-time state-space Markov models.* Bernoulli, 24(4A), 3039–3086.
- **Ivanova et al., (2024).** *Step-dad: Semi-amortized policy-based Bayesian experimental design.* ICLR Workshop on Data-centric Machine Learning Research (DMLR).
- **Wang et al., (2024).** *A comprehensive survey of continual learning: Theory, method and application.* IEEE TPAMI
- **Kirkpatrick et al. (2017).** *Overcoming catastrophic forgetting in neural networks.* Proc. Natl. Acad. Sci. (PNAS), 114(13), 3521–3526.
- **Nguyen et al. (2018).** *Variational continual learning.* Int. Conf. on Learning Representations (ICLR).
- **Baumann et al. (2025).** *Reinforcement learning with non-ergodic reward increments: robustness via ergodicity transformations.* Trans. Machine Learning Research (TMLR).
- **Forster et al., (2025).** *Improving Robustness to Model Misspecification in Bayesian Experimental Design.* 7th Symposium on Advances in Approximate Bayesian Inference – Workshop Track.
- **Duran-Martin (2025).** *Adaptive, robust and scalable Bayesian filtering for online learning.* PhD Thesis, Queen Mary University of London.