Bayesian Design of Experiments and Parameter Recovery

Christian Frederiksen

Tulane University

Joint work with Nathan Glatt-Holtz, Justin Krometis, Victoria Sieck, and Laura Freeman

How Can We Better Estimate Unknowns?

Consider the problem of recovering θ given a model of the form

$$y = G(\theta, d) + \eta$$

where y are observations, d are observation conditions, and η is noise.

Theoretical Models	Empirical Models

Our success will depend on:

- What is measured, G
- Where it is measured, d
- How we use our measurements

Table of Contents

A DOE Case Study

A recent contribution[a] found sequential (classical) DOE generally improved testing efficiency in a simulation study. The study considered linear models both with and without the presence of interaction and quadratic terms.

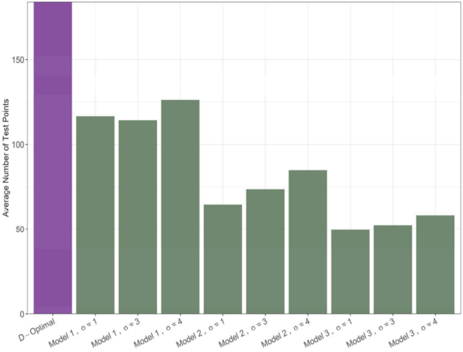

[a] Ahrens, M., Medlin, R., Pagán-Rivera, K., Dennis, J. "Case study on applying sequential analyses in operational testing". In: Quality Engineering (2022), pp. 1–12.

The Bayesian Ingredients

Bayesian DOE and parameter recovery both use the same ingredients:

- Prior - $p_0(\theta)$
- Likelihood - $L(y|\theta, d)$
- Posterior - $p(\theta|d, y)$

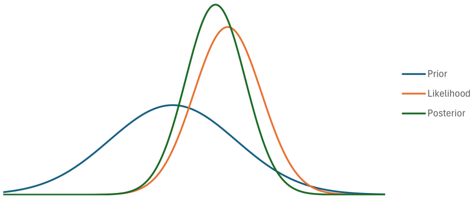

Using Bayes Rule we have

$$p(\theta|d, y) \propto L(y|\theta, d)p_0(\theta)$$

Sequential Bayesian DOE

Given previous test designs and observations $D_{N-1} = \{d_k, y_k\}_{k=1}^{N-1}$ we define

$$U_N(d) = \int u(d, y, \theta, D_{N-1}) L(y|\theta, d) p(\theta|D_{N-1}) d\theta dy$$

We then choose

$$d_N = \text{argmax } U_N(d)$$

- **Mutual Information**

$$u(d, y, \theta, D_{N-1}) = \log\left(\frac{p(\theta|d, y, D_{N-1})}{p(\theta|D_{N-1})}\right)$$

- **Bayesian D-Optimality**

$$u(d, y, \theta, D_{N-1}) = \frac{1}{\det(cov(\theta|y, d, D_{n-1}))}$$

Results

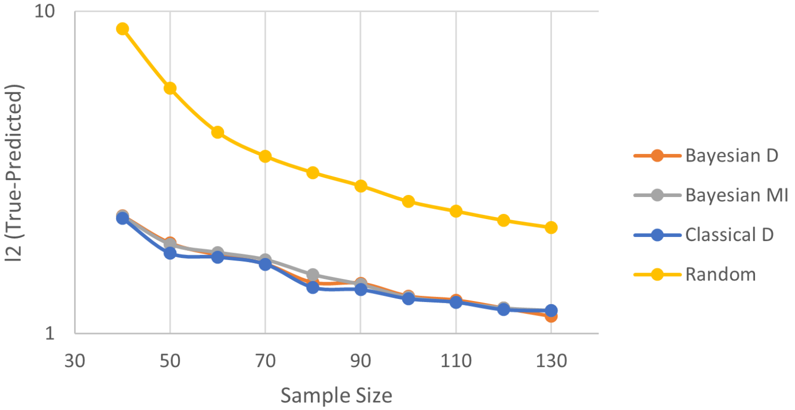
10
1
l2 (True-Predicted)
30
50
70
90
110
130
Sample Size
Bayesian D
Bayesian MI
Classical D
Random

Table of Contents

Advection Diffusion Inverse Problem

Consider a solute $\mathbf{u}$ diffusing and subject to advection by a flow θ

$$\frac{\partial \mathbf{u}}{\partial t}(t, \mathbf{x}, \theta, \mathbf{u}_0) + \theta(\mathbf{x}) \cdot \nabla \mathbf{u}(t, \mathbf{x}, \theta, \mathbf{u}_0) = \kappa \Delta \mathbf{u}(t, \mathbf{x}, \theta, \mathbf{u}_0), \qquad \mathbf{u}(0, \mathbf{x}, \theta) = \mathbf{u}_0(\mathbf{x})$$

Given a fixed parameter $\theta^\star$ and noisy observations $\mathbf{u}$ can we recover $\theta^\star$?

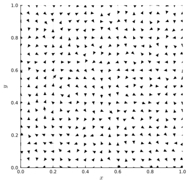

θ

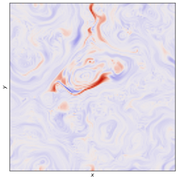

$\mathbf{u}$

Observation Model:

$$y_j = \mathbf{u}(t_j, \mathbf{x}_j, \theta^\star, \mathbf{u}_0) + \eta_j,$$
$$j = 1, \ldots, N$$

Posterior Consistency

Posterior Distribution: $p_N(\theta|d_N, y_N) \propto L(y_N|\theta, d_N)p_0(\theta)$

What happens as N grows?

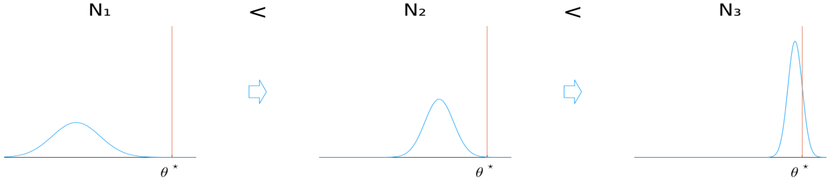

A novel contribution[1] establishes conditions under which posterior consistency holds not only for the advection diffusion problem but an entire class of PDE constrained problems.

[1] Frederiksen, C. (2024), "On Bayesian Recovery of Infinite Dimensional Parameters in Partial Differential Equations". [Doctoral Thesis, Tulane University].

Numerical Experiments

How can we quantify parameter recovery?

$$\mathrm{RMSE}(N) = \left[\int_H \|\theta - \theta^\star\|_{L^2}^2 \, p_N(d\theta) \right]^{1/2}$$

Does this actually work?

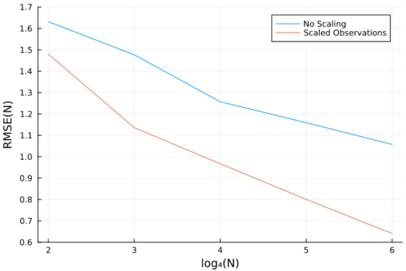

The same work[a] presents numerical experiments investigating where theory fails, characterizing the behavior of the posterior, and suggesting strategies which dramatically improve parameter recovery rates.

[a] Frederiksen, C. (2024), "On Bayesian Recovery of Infinite Dimensional Parameters in Partial Differential Equations". [Doctoral Thesis, Tulane University].

Table of Contents

Putting It All Together