Henry's ASCM Dream

Adaptive Sampler for Complex Models

A tribute to Henry P. Wynn

Noha Youssef

nayoussef@aucegypt.edu

Department of Mathematics and Actuarial Science
The American University in Cairo

Henry Wynn

- My PhD supervisor at LSE. I started my PhD with him in 2006, and he taught me how to *think* about design.
- His dream: a design that is *truly adaptive*, learning from the observed outputs at every stage.

One question, three papers

Henry's question

Can a sequential design use the observed **outputs**, not just the inputs, and stay honestly Bayesian?

1. **The PhD:** K–L expansion + an adaptive entropy criterion. We tried a full hierarchical model with an inverse Wishart prior on the covariance, but the integration was intractable inside the design loop, so we used an empirical Bayes approximation instead.
   Youssef & Wynn (2019), *Sequential Analysis*
2. **Continued:** the K–L machinery in full
   Youssef (2023), *Comm. Statist. Theory Methods*
3. **The dream:** the fully Bayesian adaptive loop, tested
   Youssef (2026), *Statistical Papers*, mODa14

Maximum Entropy Sampling

- Expensive simulator $\Rightarrow$ Gaussian process surrogate.
- **Maximum Entropy Sampling:** sample where predictive uncertainty is largest.

The catch Henry would not let go of

In the Gaussian set-up, the criterion uses only the *inputs*. The observed outputs never enter, so the design is sequential, but not truly *adaptive*.

The PhD, part I: linearize the process

- The Karhunen–Loève expansion turns the GP into a linear model:

$$Y(\mathbf{x}) \;=\; \mathbf{h}^{\top}(\mathbf{x})\boldsymbol{\beta} + \boldsymbol{\phi}^{\top}(\mathbf{x})\boldsymbol{\gamma} + \epsilon(\mathbf{x})$$

- Covariance structure becomes **mean structure**: nonlinear covariance parameters are traded for linear coefficients $\boldsymbol{\gamma}$.

- No analytical eigen-solution? **Haar wavelets** solve it numerically, in one or many dimensions.

The PhD, part II: an adaptive entropy criterion

- Hierarchical Bayes on the linearized model; an empirical Bayes step keeps the design loop tractable.
- The criterion becomes **output-aware**: each new point depends on the updated estimates from the data.

Youssef & Wynn (2019)

The design concentrates points where the simulator has peaks and skips the flat regions. *Full adaptation is beneficial.*

The K–L expansion, continued

- Full numerical treatment of the K–L expansion: **Haar wavelets win** on accuracy and speed.
- Extension to **multivariate** processes.
- K–L designs predict better: the new mean parameters are updated by the data.
- Prior variances become a steering wheel for where the design points go.

Published in Youssef (2023), *Communications in Statistics, Theory and Methods*.

Henry's dream: do it properly

- One fixed entropy criterion, three ways of updating the GP hyperparameters inside the same sequential loop:
- **ML-II** (empirical Bayes). Estimates the hyperparameters by maximizing the marginal likelihood, then plugs the single best estimate into the predictions. Fast and stable, but it treats the estimate as the truth, so hyperparameter uncertainty is ignored.
- **INLA**. Approximates the full posterior of the hyperparameters with fast deterministic (Laplace) approximations on a mesh, so their uncertainty flows into the predictive variance. No sampling needed, but the mesh machinery is costly.
- **MAP + FullBayes**. Finds the posterior mode of the hyperparameters, then samples around it with MCMC and averages the predictions over those samples. More honest uncertainty than a plug-in, at the cost of sampling at every stage.

Benchmarks from 1D to 6D, identical budgets, Monte Carlo replication.

What we learned

- **Accuracy:** ML-II and MAP+FullBayes predict comparably well; the entropy rule finds the informative regions every time.
- **Uncertainty** is where they part ways: ML-II is fast but overconfident; MAP+FullBayes stays conservative; INLA depends on the configuration.

The message

The treatment of what we *do not know* governs the design.

Henry's result, in a real sense

- From 1987 to 2019 to 2026, this is one continuous line of work: information-based design works when we are honest about uncertainty.
- He did not see the final version, but the question behind it is his.

I like to think the design is, at last, the adaptive one he had in mind.

Carrying it forward

- Bayesian updating in high dimensions.
- Hybrid strategies: ML-II speed, full-Bayes calibration.
- Sequential design on constrained input spaces.

Shewry & Wynn (1987) · Youssef & Wynn (2019) · Youssef (2023) · Youssef (2026)

In memory of Henry P. Wynn

Teacher, collaborator, and the most adaptive mind I have known.

Thank you.

Noha Youssef