Final Scientific/Technical Report for "Practical and Optimal Sequential Bayesian Experimental Design for Complex Systems Incorporating Human Experimenter Preferences"

Report Number: DOE-UMICH-0021398

Department of Energy
Office of Advanced Scientific Computing Research
Award Number: DE-SC0021398
Program Manager: William Spotz

Project Title:
Practical and Optimal Sequential Bayesian Experimental Design for Complex Systems Incorporating Human Experimenter Preferences

Period of Performance:
09/01/2020 – 08/31/2024

Institution Recipient:
The Reagents of the University of Michigan

Team Members:
Xun Huan, University of Michigan (Principal Investigator)
Nikola Banovic, University of Michigan (Co-Principal Investigator)

Executive Summary

Experiments are indispensable for developing models of complex systems. Carefully designed experiments can provide substantial savings for these expensive data-acquisition opportunities. However, designs based on heuristics are often suboptimal for systems with multiphysics, nonlinear dynamics, and uncertain and noisy environments. Optimal experimental design, while leveraging predictive models, seeks to systematically quantify and maximize the value of experiments.

In this project, we focused on the design of multiple experiments, where current approaches are largely suboptimal: batch-design does not adapt to new data acquired during the experiment campaign (no feedback), and greedy/myopic design ignores future dynamics and consequences (no lookahead). We developed the mathematical framework and computational methods for sequential optimal experimental design (sOED) for complex systems. We enabled tractable model-based sOED in a rigorous manner through novel algorithms based on reinforcement learning, and investigated the effects of human experimenters on the design process.

Our methods are fully Bayesian, able to quantify and update uncertainty in a principled manner. The traits aimed by our approach—mathematical rigor and optimality, human effects and uncertainty quantification, computational practicality—are crucial for elevating the standards of artificial intelligence (AI) to support decision-making in scientific domains, and contribute toward trust and realistic adoption of AI in experimental design practice.

Research Contributions

Engineering and science revolve around the interplay between data and models: leveraging data to develop, calibrate, and improve models, and using models to predict outcomes, control processes, and guide decision-making. When data is expensive to acquire, carefully designing experiments becomes crucial. Optimal experimental design (OED) is a field dedicated to systematically quantifying the value of experiments and identifying the best conditions to conduct them. In particular, sequential experimental design involves planning multiple experiments conducted in sequence, where the results of earlier experiments can inform the design of subsequent ones.

The major goals of this project are to develop a mathematical framework and computational methods for sOED for complex systems. We have two main objectives:

1. Enable tractable model-based sequential optimal experimental design that designs multiple experiments under uncertainty, in a rigorous and provably optimal manner integrating feedback and lookahead, while leveraging predictive science models.
2. Incorporate quantified effects of the human experimenter into the experimental design utility that can filter out human heuristics bias, and conceive designs that properly reflect what practitioners value in real-life experiments beyond information gain.

Objective 1: Enabling Tractable Sequential Optimal Experimental Design

We formulated the sequential optimal experimental design (sOED) problem under a finite-horizon partially observable Markov decision process (POMDP) in a Bayesian setting with information-theoretic utilities. This general framework accommodates continuous random variables, non-Gaussian posteriors, and nonlinear forward models. The sOED design policy incorporates elements of feedback and lookahead simultaneously, and we show it to generalize the commonly-used batch and greedy design strategies. We solved for the sOED policy using policy gradient (PG) techniques from reinforcement learning, and provided a derivation for the PG expression in the sOED context. Adopting an actor-critic approach, the policy and value functions are parameterized using deep neural networks and improved via PG estimates produced from simulated episodes of designs and observations. The new PG-sOED algorithm was validated on a linear-Gaussian benchmark and compared against other design baselines on a sensor movement problem for contaminant source inversion in a convection-diffusion field. Related publications: [5, 4].

We developed variational sequential optimal experimental design (vsOED), a novel method that employs one-point reward formulation and variational posterior approximations to form a provable lower bound to the expected information gain. Numerical methods were developed following an actor-critic reinforcement learning approach, including derivation and estimation of variational and policy gradients to optimize the design policy, and posterior approximation using Gaussian mixture models and normalizing flows. vsOED accommodates nuisance parameters, implicit likelihoods, and multiple candidate models, while supporting flexible design criteria that can target designs for model discrimination, parameter inference, goal-oriented prediction, and their weighted combinations. We demonstrated vsOED across various engineering and science applications, illustrating its superior sample efficiency compared to existing sequential experimental design algorithms. Related publications: [3, 4]

This project also contributed to a review paper [2], published in *Acta Numerica*, which surveyed modern OED from its foundations in classical design theory to current research involving OED for complex models. Key highlights include selection of design criteria, flexibility of the Bayesian and decision-theoretic approach well-suited to nonlinear and non-Gaussian statistical models, methods for estimating or bounding the values of these design criteria, computational optimization techniques for finding a design, and the latest advances in sOED. Additionally, this project resulted in a PhD thesis [3] that focused on sOED and vsOED, as well as robust OED and robust sOED that incorporate mean-plus-variance risk measures into the information-theoretic utility framework.

Objective 2: Quantifying Human Experimenter Preferences and Effects

A key aspect of this objective is to understand and incorporate human preferences into OED. This naturally involves personalization of the design criteria, to tailor the goals and other interactions of the design process relevant to each user's background and preferences. However, personalization requires information about users that platforms often collect without their awareness or their enthusiastic consent. We studied how the transparency of AI inferences on users' personal data affects their privacy decisions and sentiments when sharing data for personalization. We conducted two experiments where participants answered questions about themselves for personalized public arts recommendations. Participants indicated their consent to let the system use their inferred data and explicitly provided data after awareness of inferences. Our results showed that participants chose restrictive consent decisions for sensitive and incorrect inferences about them and for their answers that led to such inferences. Our findings expanded existing privacy discourse to inferences and informed future directions for shaping existing consent mechanisms in light of increasingly pervasive AI inferences. Related publication: [1].

We investigated the effects of Bayesian uncertainty quantification (UQ) to human practitioners in a healthcare application setting. This understanding is useful towards tailoring OED to important clinical uses and to improve transparency and trustworthiness of often opaque AI tools. DNN-based AI tools are being used to provide clinicians with fast and accurate predictions that are highly valuable for high-stakes medical decision-making, such as in brain tumor segmentation and treatment planning. However, these models largely lack transparency about the uncertainty in their predictions, potentially giving clinicians a false sense of reliability that may lead to grave consequences in patient care. Growing calls for Transparent and Responsible AI have promoted UQ to capture and communicate uncertainty in a systematic and principled manner. However, traditional Bayesian UQ methods remain prohibitively costly for large, million-dimensional tumor segmentation DNNs such as the U-Net. We presented a computationally-efficient UQ approach via the partially Bayesian neural networks (pBNN). In pBNN, only a single layer, strategically selected based on gradient-based sensitivity analysis, is targeted for Bayesian inference. We illustrated the effectiveness of pBNN in capturing the full uncertainty for a 7.8-million parameter U-Net. We also demonstrated how practitioners and model developers can use the pBNN's predictions to better understand the model's capabilities and behavior. Related publication: [7].

We proposed a sample size determination method to aid OED where the number of experiments is not fixed a priori, and where data collection is tedious and resource intensive. We studied this in the context of experiments collecting data for a specific inverse reinforcement learning (IRL) model of human behavior. This was approached in two cases: (1) pre-hoc experiment design—conducted in the planning stage before any data is collected, to guide the estimation of how many samples to collect; and (2) post-hoc dataset analysis—performed after data is collected, to decide if the existing dataset has sufficient samples and whether more data is needed. We validated our approach in experiments with a realistic model of behaviors of people with Multiple Sclerosis (MS) and illustrated how to pick a reasonable sample size target. Our work enabled model designers to perform a deeper, principled investigation of the effects of dataset size on IRL model parameters. Related publication: [6].

Publications

[1] S. Asthana, J. Im, Z. Chen and N. Banovic (2024), "I know even if you don't tell me": Understanding users' privacy preferences regarding AI-based inferences of sensitive information for personalization, in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*, ACM, article 782, pp. 1–21.

[2] X. Huan, J. Jagalur and Y. Marzouk (2024), Optimal experimental design: Formulations and computations, *Acta Numerica* **33**, 715–840.
[3] W. Shen, J. Dong and X. Huan (2023), Variational sequential optimal experimental design using reinforcement learning. Available at arXiv:2306.10430.
[4] W. Shen (2023), Reinforcement learning based sequential and robust Bayesian optimal experimental design, PhD thesis, University of Michigan.
[5] W. Shen and X. Huan (2023), Bayesian sequential optimal experimental design for nonlinear models using policy gradient reinforcement learning, *Computer Methods in Applied Mechanics and Engineering* **416**, 116304.
[6] T. Hossain, W. Shen, A. D. Antar, S. Prabhudesai, S. Inoue, X. Huan and N. Banovic (2023), A Bayesian approach for quantifying data scarcity when modeling human behavior via inverse reinforcement learning, *ACM Transactions on Computer-Human Interaction* **30**(1), 8:1–8:27.
[7] S. Prabhudesai, J. Hauth, D. Guo, A. Rao, N. Banovic and X. Huan (2023), Lowering the computational barrier: Partially Bayesian neural networks for transparency in medical imaging AI, *Frontiers in Computer Science* **5**, 1071174.

Presentations/Posters

- (2024/09) Xun Huan gave a presentation "Bayesian Optimal Experimental Design" at the University of Michigan Applied Physics Seminar.
- (2024/06) Xun Huan gave a presentation "Sequential optimal experimental design for digital twins" at the IMSI Workshop on Mathematical and Statistical Foundations of Digital Twins.
- (2024/02) Jiayuan Dong gave a presentation "A Reinforcement Learning and Variational Approach to Sequential Optimal Experimental Design" at SIAM UQ.
- (2024/02) Snehal Prabhudesai gave a presentation "Computing and Communicating Uncertainty in AI-based Medical Imaging Applications for Healthcare Decision Making" at SIAM UQ.
- (2023/06) Xun Huan gave a presentation "Bayesian Sequential Optimal Experimental Design" at the University of Toronto Institute for Aerospace Studies Seminar Series.
- (2023/05) Jiayuan Dong gave a presentation "Variational Bayesian Optimal Experimental Design with Normalizing Flows" at ICODOE.
- (2023/05) Xun Huan gave a presentation "Variational Bayesian Sequential Optimal Experimental Design via Policy Gradient" at ICODOE.
- (2023/04) Tahera Hossain and Wanggang Shen gave a presentation "A Bayesian Approach for Quantifying Data Scarcity when Modeling Human Behavior via Inverse Reinforcement Learning" at CHI 2023.
- (2023/03) Xun Huan gave a presentation "A Reinforcement Learning Framework for Bayesian Sequential Experimental Design" at the National University of Singapore Centre for Data Science and Machine Learning (CDSML) Seminar.
- (2023/02) Xun Huan gave a presentation "Reinforcement Learning Framework for Bayesian Sequential Optimal Experimental Design" at SIAM CSE.
- (2023/02) Snehal Prabhudesai gave a presentation "Lowering The Computational Barrier : Selective Bayesian Uncertainty Quantification for Transparency in Medical Imaging AI" at SIAM CSE.
- (2022/09) Xun Huan gave a presentation "Bayesian Reinforcement Learning for Optimal Sensor Relocation in Convection-Diffusion Fields" at SIAM MDS.
- (2022/06) Xun Huan presented a poster "Goal-Oriented Optimal Experimental Design for Nonlinear Systems" at the ISBA World Meeting.
- (2022/04) Snehal Prabhudesai gave a presentation "Partially Bayesian Neural Networks: Low-Cost Bayesian Uncertainty Quantification for Deep Learning in Medical Image Segmentation" at SIAM UQ.

- (2022/04) Wanggang Shen gave a presentation "Optimal Bayesian Design of Sequential Experiments Using Deep Deterministic Policy Gradient" at SIAM UQ.
- (2022/01) Xun Huan gave a presentation "Optimal Bayesian Design of Finitely Sequential Experiments with Deep Deterministic Policy Gradient" at the Royal Statistical Society (RSS) Computational Statistics and Machine Learning session on Automatic Experimentation.
- (2021/12) Xun Huan gave a presentation "Bayesian Sequential Optimal Experimental Design for Nonlinear Systems via Policy Gradient" at the Applied Reinforcement Learning Seminar.
- (2021/11) Xun Huan gave a presentation "Reinforcement Learning for Sequential Optimal Experimental Design" at the New England Statistical Society (NESS) NextGen: Data Science Day.
- (2021/10) Xun Huan gave a presentation "Bayesian Optimal Experimental Design for Batch and Sequential Experiments" at the Dartmouth College Applied and Computational Mathematics Seminar.
- (2021/08) Xun Huan gave a presentation "Closed-Loop Bayesian Design of Sequential Experiments via Dynamic Programming and Reinforcement Learning" at the IFIP TC7 Conference on System Modelling and Optimization.
- (2021/07) Wanggang Shen gave a presentation "Sequential Optimal Experimental Design Using Reinforcement Learning with Policy Gradient" at USNCCM.
- (2021/06) Xun Huan gave a presentation "Optimal Bayesian Sequential Design Using Reinforcement Learning with Policy Gradient Methods" at ISBA World Meeting.
- (2021/05) Xun Huan gave a presentation "Optimal Sequential Bayesian Design of Experiments Using Reinforcement Learning with Policy Gradient" at EMI/PMR Conference.
- (2021/03) Wanggang Shen gave a presentation "Optimal Bayesian Design of Sequential Experiments using Reinforcement Learning with Policy Gradient Methods" at SIAM CSE.
- (2021/01) Wanggang Shen gave a presentation "Sequential Optimal Experimental Design Using Reinforcement Learning with Policy Gradient" at WCCM-ECCOMAS Congress.
- (2020/12) Xun Huan gave a presentation "Optimal Sequential Bayesian Design of Experiments Using Reinforcement Learning with Policy Gradient" at the Machine Learning in Science & Engineering (MLSE) Conference.

Conference Sessions Organized

- (2023/05) Xun Huan organized a session "Variational Methods for Optimal Experimental Design" at ICODOE.
- (2023/02) Xun Huan co-organized a two-part minisymposium "Model-Based Optimal Experimental Design" (parts I, II) at SIAM CSE.
- (2022/04) Xun Huan co-organized a three-part minisymposium "Model-Based Optimal Experimental Design" (parts I, II, III) at SIAM UQ.
- (2021/07) Xun Huan co-organized a minisymposium "Optimal Experimental Design in Computational Science and Engineering" at USNCCM.
- (2021/01) Xun Huan co-organized a minisymposium "Optimal Experimental Design in Computational Science and Engineering" at WCCM-ECCOMAS Congress.

People Supported

- Xun Huan, Associate Professor, Department of Mechanical Engineering, University of Michigan
- Nikola Banovic, Associate Professor, Department of Computer Science and Engineering, University of Michigan
- Wanggang Shen, PhD Student (graduated 2023), Department of Mechanical Engineering, University of Michigan

- Snehal Prabhudesai, PhD Student, Department of Computer Science and Engineering, University of Michigan
- Sumit Asthana, PhD Student, Department of Computer Science and Engineering, University of Michigan
- Anindya Das Antar, PhD Student, Department of Computer Science and Engineering, University of Michigan
- Divya Ramesh, PhD Student, Department of Computer Science and Engineering, University of Michigan

Acknowledgment

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under Award Number DE-SC0021398.

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.