STAT 548: PhD Qualifying Course Papers

Overview. My research primarily concerns high-performance algorithms for Bayesian modeling. An important outcome of this research is the development of algorithms which support probabilistic programming languages such as Stan and TensorFlow Probability. These algorithms must be efficient, easy-to-use ("black box"), and general-purpose in order to accommodate the range of problems scientists work on and support a model development workflow. Two important classes of algorithms in this space are *Markov chain Monte Carlo* and *variational inference*.

This research is at the intersection of theory, computation, and application. The papers I have selected demonstrate how to work at this intersection. When considering the topic of each paper, you might ask yourself: would a scientist using Bayesian modeling benefit from this paper's contribution? Should the proposed method be implemented in a statistical software? Does the paper provide us with a sufficiently thorough analysis (theoretical and empircal) to answer these questions?

Paper list (chose 4)

- Nested $\widehat{R}$: Assessing the Convergence of Markov chain Monte Carlo when Running Many Short Chains. Margossian et al. *Bayesian Analysis* (2024)
- Incorporating Local Step-Size Adaptivity into the No-U-Turn Sampler using Gibbs Self Tuning. Bou Rabee et al. *arXiv:2408.08259* (2024)
- Pathfinder: Parallel quasi-Newton variational inference. Zhang et al. *Journal of Machine Learning Research* (2022)
- Advances in Black Box VI: Normalizing Flows, Importance Weighting, and Optimization. Agrawal et al. *Advances in Neural Information Processing Systems* (2020) (read appendix in arXiv version: arXiv:2006.10343)
- Variational Inference in Location-Scale Families: Exact Recovery of the Mean and Correlation Matrix. Margossian and Saul. *Artificial Intelligence and Statistics* (2025)

Evaluation.

- For each paper you will write a critical analysis which summarizes the contributions of the paper. You should contextualize the paper and discuss the existing literature. You can highlight the strengths and limitations of the paper, as a reviewer might do. If the paper has theoretical contributions, you should be able to reproduce the proofs step-by-step and fill in gaps which may have been glossed over in the paper. Finally, the report should discuss open questions and propose future research directions—ones which might actually lead to your own research project!
- In addition, we will come up with a small project for each paper. For example, you might extend one of the numerical experiments which evaluates the performance of a method across several models by adding a new model. In doing so, you will write code to implement the method, evaluate the performance and produce a figure describing the output—all important skills for research.

As you work through the paper, I will be particularly interested in your ability to identify which parts of the paper you understand thoroughly and which ones remain unclear to you. I might grill you on some of the details in the papers with the goal to help you develop the ability to grill yourself when you read papers. It's ok if some of the paper is unclear and it's something we should discuss together.