Fast, Self-Correcting Variational Inference for Bayesian Deep Learning

Dayta Dominic Bagus
Mathematical Informatics Laboratory
Supervisor: Prof. Kazushi Ikeda

Abstract (should be within 1st page)

This dissertation develops two novel algorithmic contributions to improve the stability and speed of Standard Monte Carlo Variational Inference (MCVI), or Black-Box Variational Inference (BBVI), with the goal of enabling Bayesian Deep Learning (BDL). First, we show that replacing the naive Monte Carlo average with the James-Stein (JS) estimator results in a guaranteed variance reduction by shrinking the noisy gradient, and does so without the model-specific analytical overhead of traditional methods like Rao-Blackwellization. Second, we address the computational bottleneck of requiring large sample sizes. We develop YOASOVI, a "self-correcting" algorithm that uses only a single Monte Carlo sample per iteration. It controls the extreme variance of single sample estimates using a Metropolis-type acceptance-rejection mechanism based on the ELBO value, ensuring a fast, greedy, and provably convergent optimization.

We demonstrate the power and flexibility of this framework by applying it to complex econometric and deep learning models for financial time series. We introduce the Transformer-Dynamic Factor Model (TransDFM), which replaces the DFM's linear components with a Bayesian neural network and a Bayesian SaMoVAR (Transformer) to model non-linear factor dynamics. Applied to the S&P 500 and PSE, the TransDFM achieves strong predictive accuracy and well-calibrated uncertainty. Despite its "black-box" nature, the model retains economic interpretability, as its primary latent factor demonstrates a significant alignment with systematic market risk (CAPM beta). Collectively, these contributions provide a suite of theoretically-grounded, efficient, and "black-box" VI algorithms that bridge the gap between high-performance deep learning and principled Bayesian inference.