

Deep Generative Models for Scientific Applications

Jeffrey Regier, University of Michigan

Deep Generative Models (DGMs) combine Bayesian inference and deep learning. Conditional distributions are encoded with neural networks, giving these models the flexibility to reflect the complex structure of scientific data. In this talk, I discuss DGMs for a variety of applications, including genome-wide association studies, electronic health records, single-cell RNA-sequencing, spatial transcriptomics, and astronomical images. I report on the final application at length. DGMs, fitted with variational inference and stochastic optimization, let us achieve better results than previous MCMC-based approaches, nearly 10,000 times faster.