THE UNIVERSITY OF CHICAGO | DEPARTMENT OF STATISTICS

Master's Thesis Presentation

Qi Chen
Department of Statistics
The University of Chicago

"Graphic model Geometry-Aware Hamiltonian Variational Auto-Encoder"

December 9, 2025, at 10:30 AM
Jones 111, 5747 S. Ellis Avenue

Abstract

Variational Autoencoders (VAEs) provide a powerful probabilistic framework for generative modeling, yet their assumption of an isotropic Gaussian latent space often limits their ability to capture the true geometric structure of complex data. This thesis investigates the Riemannian Hamiltonian Variational Autoencoder (RHVAE), a geometry-aware extension of the standard VAE that endows the latent space with a learned Riemannian metric and employs Riemannian Hamiltonian Monte Carlo (RHMC) for efficient sampling. The learned metric adapts to the underlying data manifold, enabling more expressive posterior representations and improved sample quality.

To enhance sampling stability, a tempering scheme is introduced within the RHMC integration steps, and a centroid-based approximation is proposed to accelerate metric computation without sacrificing performance. The model is evaluated on both synthetic datasets and the MNIST benchmark through qualitative visualization and quantitative analysis using the FrÅLechet

Inception Distance (FID). Experimental results demonstrate that RHVAE achieves competitive performance with flow-based models such as the Generative Latent Flow (GLF) while substantially

outperforming Gaussian baselines. Moreover, the analysis shows that geometry-aware sampling stabilizes generation, that moderate latent dimensionalities provide the best tradeoff between flexibility and regularization, and that clustering-based metric compression significantly reduces computational cost while preserving generation quality in higher-dimensional latent spaces.

In addition, a classification experiment using class-specific HVAE models shows that the

---

For information about building access for persons with disabilities, please contact Keisha Prowoznik at 773.702-0541 or send an email to kprowoznik@uchicago.edu. If you wish to subscribe to our email list, please visit the following website: https://lists.uchicago.edu/web/info/statseminars.

learned Riemannian geometry also provides strong discriminative ability: evaluating the ELBO under class-specific models yields high accuracy (up to 95% for $d = 8$), indicating that the geometry-aware latent structure captures class-dependent features effectively.

Overall, the RHVAE framework offers a principled and scalable approach to generative modeling,

effectively combining geometric representation learning with Hamiltonian dynamics to capture complex latent structures.
