Research Statement: Gavin Kerrigan

gavin.k@uci.edu • gavinkerrigan.github.io

Introduction

My research aims to develop principled foundations for generative modeling and to build upon these foundations to solve challenging problems arising from the sciences and engineering. I am particularly excited by work in the intersection of applied mathematics and machine learning, and my previous work has brought together tools from optimal transport, differential equations, and functional analysis to design novel generative models. Applications of my work seek to employ generative models as a tool for scientific discovery and innovation, as highlighted by my work in climate science.

Previous Work

A key theme of my research has been the study of dynamics in the space of probability measures. Such dynamics allow one to model a complex, high-dimensional data distribution through a series of gradual transformations of a specified reference distribution. This framework unifies several classes of generative models, including flows and diffusion models.

One benefit of this perspective is *generality*. For instance, function-valued data, like solutions to partial differential equations (PDEs) or time series, are ubiquitous in the sciences and engineering. However, building generative models for functional data poses significant fundamental challenges due to the infinite-dimensional nature of such data. In [1], I developed the theory for building diffusion generative models in infinite-dimensional spaces, leveraging tools from measure theory and functional analysis. These techniques were further refined in [2], where I developed the theory for flow-based models in infinite-dimensional spaces. This work [2] was recognized with an Outstanding Student Paper award (top 0.3%) at AISTATS 2023.

A second benefit of this framework is *flexibility*. As there are many paths in the space of distributions which interpolate between the reference measure and the data measure, we are afforded great flexibility in this choice. One tool for designing paths is optimal transport, which has previously shown strong empirical success in unconditional generative modeling tasks. Motivated by this, I sought to employ similar techniques for conditional tasks, such as solving Bayesian inverse problems. However, prior to my work, the theory of conditional optimal transport (COT) was not sufficiently developed to justify this approach. Thus, in [3] I developed a dynamical theory of COT which is a conditional analogue of the renowned Benamou-Brenier theory. Equipped with these tools, I developed a conditional flow-based model for solving challenging amortized inference problems, including an infinite-dimensional Bayesian inverse problem arising from fluid flow PDEs.

I have additionally collaborated with domain experts in applying generative models to problems arising in climate science. In joint work with Stephan Mandt's group at UCI and Chris Bretherton's group at AI2, we developed a diffusion-based approach for super-resolving coarse-grid weather forecasts [4]. As numerical simulations are computationally expensive at high resolutions, this method enables practitioners to run cheaper low-resolution simulations which may be super-resolved to predict the high-resolution weather states. In ongoing work, I have also collaborated with Efi Foufoula-Georgiou's group at UCI to develop diffusion-based approaches for multi-satellite data fusion, with a focus on recovering precipitation intensities from microwave and infrared imaging.

Future Work

One tool I am interested in exploring in future work is that of gradient flows, which provides a framework for optimizing functionals in the space of probability measures. For example, diffusion models may be viewed as the gradient flow of the KL divergence. One potential avenue is to design functionals which regularize the behavior of generative models, such as enhancing fidelity in low-density regions, which is crucial for accurately sampling rare events (e.g. extreme precipitation events). Conversely, gradient flows could serve as a powerful technique for enhancing our theoretical understanding of existing models, such as convergence results for flow-based models. Gradient flows are intimately connected with curves in the space of measures, and thus this line of work naturally builds on my expertise in this domain.

I am also broadly interested in developing techniques which enhance the applicability of generative models to real-world problems. One pervasive limitation is the requirement for large-scale datasets, and alleviating this requirement would empower generative models in domains where data is sparse. A promising direction for overcoming this limitation is to design hybrid generative models which blend hand-crafted models with data-driven techniques. For instance, many problems in the physical sciences include known physical constraints, e.g. the dynamics of climate systems are often modeled through PDEs [5]. Building generative models which obey these known constraints can lead to more robust and interpretable models by leveraging the rich history of scientific modeling.

Second, a key advantage of generative models over discriminative models is their potential for accurate uncertainty quantification. Having a well-calibrated model not only enables informed decision making, but also enhances the trust of domain experts in model predictions. Moreover, accurate uncertainty quantification is critical in high-stakes applications, e.g. solving inverse problems arising in medical imaging. In future work, I am interested in developing techniques for understanding, evaluating, and improving the calibration of contemporary generative models. This is particularly challenging for generative models, as existing discriminative metrics for calibration do not naturally extend to high-dimensional unsupervised data.

Conclusion

Overall, my research spans the spectrum of generative modeling, ranging from theoretical and methodological advances to practical applications. I am particularly motivated by problems arising from the sciences and engineering, where the flow of ideas is bidirectional — practical challenges can engender theoretical developments, while theoretical advances can be leveraged to solve impactful problems. As our ability to collect and store vast amounts of data from our world grows, generative models will become increasingly important, and my work seeks to develop the tools necessary for leveraging this sea of unsupervised data to better understand our world.

[1] Diffusion Generative Models in Infinite Dimensions. Gavin Kerrigan, Justin Ley, Padhraic Smyth. *International Conference on AI and Statistics (AISTATS)*, 2023.
[2] Functional Flow Matching. Gavin Kerrigan, Giosue Migliorini, Padhraic Smyth.
International Conference on AI and Statistics (AISTATS), 2024.
[3] Dynamic Conditional Optimal Transport through Simulation-Free Flows. Gavin Kerrigan, Giosue Migliorini, Padhraic Smyth. *arXiv:2404.04240*, 2024.
[4] Precipitation Downscaling with Spatiotemporal Video Diffusion. Prakhar Srivastava, Ruihan Yang, Gavin Kerrigan, Gideon Dresdner, Jeremy McGibbon, Christopher Bretherton, Stephan Mandt. *arXiv:2312.06071*, 2023.
[5] Introduction to Climate Dynamics and Climate Modeling. H. Goosse et al. Centre de recherche sur la Terre et le climat Georges Lemaître-UCLouvain, 2010.