Learning the Physical World with Energy-Based Model

Jiahe Huang

My research interests center on artificial intelligence for scientific applications and computer vision, with a particular focus on leveraging **energy-based models** to bridge the **physical world and machine learning insights**. I am driven by the potential of diffusion models to capture complex physical phenomena and generate high-fidelity data, which can significantly advance scientific understanding. By utilizing these neural methods, I aim to address scientific challenges such as **physics simulations** and **decision-making** processes in real-world applications. Additionally, I am interested in incorporating **3D-aware techniques** to enhance data augmentation and improve model generalization. I will elaborate in the following sections.

Learning to simulate physics with generative prior

Predicting fluid systems governed by partial differential equations (PDEs) from sparse sensor data is a significant challenge in computational physics. PDEs are fundamental to modeling a wide range of physical phenomena, but their analytical solutions are often intractable, especially in complex, real-world scenarios like turbulent flows described by the Navier-Stokes equations. These challenges are compounded by the difficulty of reconstructing high-dimensional solutions from sparse or noisy observations. Since 2023, I have focused on addressing this challenge by integrating AI techniques into PDE solving, specifically leveraging diffusion models as powerful generative models suited to the nature of PDEs and capable of learning the physical distribution patterns.

Diffusion models excel in modeling the stochastic processes inherent in fluid dynamics, making them particularly effective for capturing the chaotic behavior of turbulent flows. Their ability to learn iterative PDE priors through energy-based modeling allows them to approximate complex PDE solutions even with limited data. By applying physics-informed constraints, diffusion models can iteratively solve inverse problems while ensuring that their gradually converged solutions respect the underlying laws of physics, bridging the gap between traditional PDE solving and modern AI methods in science. This approach not only enables accurate and robust predictions in scenarios involving turbulence or noisy data but also highlights the potential of AI in advancing scientific understanding of PDE-governed systems.

As the project leader, I developed a framework using physics-informed guided sampling, which combines observation loss and PDE function loss to enforce physical constraints, enabling the reconstruction of both material properties (coefficients) and flow properties (solutions) for static PDEs. For dynamic PDEs, the framework reconstructs the flow properties at key time steps, such as initial and final states, even with extremely sparse observations.

Through extensive experiments on various types of PDEs, I demonstrated that DiffusionPDE has several advantages: 1) it can solve both the solution (or final state) prediction and parameter (or initial state) estimation tasks simultaneously, 2) it accurately recovers missing data even with very limited ($\approx 3\%$) observations, which is crucial for the real-world application, and 3) it shows the potential of using a single generative model to effectively solve complex mathematical equations. My first-authored work [1] was accepted for **oral presentation** at the **AI for Science Workshop at ICML 2024**.

Building on this work, I further evaluated the performance of our guided sampling approach by comparing it with Classifier-Free Guidance (CFG). Our results indicate that guided sampling outperforms CFG, as it applies physical constraints more directly. This research [2] has been accepted to **NeurIPS 2024**.

3D world learning with 2D multi-view diffusion model

Building on my physics simulation experience, I have focused on modeling the three-dimensional physical world. Diffusion models effectively generate new 2D images, enhancing dataset diversity. To ensure high image quality and label accuracy in multi-view scenarios, I incorporated depth-based 3D scene reconstruction, bridging 3D structures with their 2D representations. Since Summer 2024, I have developed data augmentation techniques by integrating depth information with diffusion models. This automated pipeline generates high-quality augmented data through transformations like rotation and translation without manual prompts. Our project leverages 3D methods to produce 2D images, showcasing diffusion models' potential in understanding complex physical phenomena.

Decision making with graph diffusion model

Based on my work in physics simulation and 3D modeling where I utilized diffusion models to learn and represent complex physical systems, I became interested in how these models could be applied to decision-making tasks that also involve dynamic structures. Recognizing that many real-world optimization problems mirror the complexity of physical phenomena, I saw an opportunity to extend the application of diffusion models to new domains. Graph diffusion models have demonstrated promising potential for combinatorial optimization problems, achieving solutions with low optimality gaps. Since Fall 2024, I have been working on this topic. While current simple graph diffusion models struggle with solving highly complex combinatorial optimization problems [4], I implemented CFG techniques to increase its flexibility and generalization capabilities. I also integrated reinforcement learning into graph diffusion models to boost performance in tasks like Directed Acyclic Graph optimization, where rewards are measured incrementally.

Future Research Agenda

Academia provides the intellectual freedom to pursue innovative research, contributing to a future where artificial intelligence enhances human science in safe and productive ways. I look forward to contributing to the "AI+" community in the future. Specifically, I aim to further explore the application of energy-based models and other neural methods in natural science and the real world.

- **Simulating more dynamic and 3D physics systems with diffusion model:** My prior work [2] focuses on mapping initial to final states in dynamic PDEs and I am extending this research to the reconstruction of dynamic solutions over time with stochastic observations via video diffusion, aligning the approach more closely with real-world scenarios. Moreover, 3D simulations are often crucial for practical applications, such as pressure simulations on object surfaces using the 3D Navier-Stokes equation [5]. For example, GINO [3] applies neural operators to GNNs to learn 3D geometries. Given the limited availability of sensors in real-world situations, however, reconstructing 3D spaces from sparse observations is a significant yet underexplored challenge. My prior research on data augmentation has provided me with a deeper understanding of 3D generative models. Combining 3D graph-based methods with diffusion models offers a promising approach to address this challenge effectively.
- **Real-world decision-making with diffusion model:** My previous work on conditional graph diffusion points to a promising direction for improvement. While existing methods primarily condition on total cost, exploring conditions based on the remaining steps roughly predicted by basic heuristic methods could offer additional insights and enhance solution quality and model efficiency.

Related Work

[1] **Jiahe Huang**, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-solving under partial observation. In *ICML 2024 AI for Science Workshop*, 2024.

[2] **Jiahe Huang**, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-solving under partial observation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

References

[3] Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36, 2024.

[4] Zhiqing Sun and Yiming Yang. DIFUSCO: Graph-based diffusion solvers for combinatorial optimization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[5] Nobuyuki Umetani and Bernd Bickel. Learning three-dimensional flow for interactive aerodynamic design. *ACM Transactions on Graphics (TOG)*, 37(4):1–10, 2018.