Designing and controlling complex systems via diffusion generative models

Tailin Wu

In science and engineering, controlling the evolution of complex physical systems and designing their high-dimensional parameters are fundamental tasks, widely applied in disciplines such as condensed matter physics and fluid dynamics. Traditional design and control methods often require extensive computation due to complex physical dynamics (e.g., adjoint methods); or they struggle to adapt to scenarios with strong nonlinearity and strong coupling (e.g., PID). On the other hand, recent methods based on deep learning (e.g., deep reinforcement learning) also face problems such as adversarial examples and difficulties in optimizing long-term control sequences. In this work, we propose a design and control method for complex physical systems based on diffusion generative models. We consider the state trajectory of the physical system, design parameters, and control sequences as a joint variable and learn their joint probability distribution from data via a diffusion model (represented by an energy function). During inference, new samples are generated by simultaneously minimizing the learned energy function and a given design or control objective. Thus, the generated samples (including design parameters, control sequences, and system state trajectories) are both physically plausible and closely approximate the optimal design and control objectives. Additionally, we propose further methods that enable the generation of more complex design parameters than those used during training and control sequences that break free from the training set distribution constraints. We conduct design and control experiments in one-dimensional and two-dimensional partial differential equation (PDE) systems. In the experiment of designing aircraft wing shapes, although our model was trained only with data on the interaction between a single 2D wing shape and fluid, it was able to design multiple wing shapes and their formations to improve the lift-to-drag ratio, more complex than those in the training. In the control experiments, our method revealed that the "quick closing-slow opening" is an efficient movement pattern for jellyfish, consistent with existing findings in fluid dynamics, and showed unique advantages over deep reinforcement learning.