LEARNED REFERENCE-BASED DIFFUSION SAMPLING FOR MULTI-MODAL DISTRIBUTIONS

LOUIS GRENIOUX

Abstract: In recent years, score-based diffusion methods have emerged as powerful tools for sampling from complex probability distributions using only evaluations of unnormalized densities. These methods approximate the time-reversal of a diffusion process, gradually transforming an easy-to-sample base distribution into the target distribution. However, their performance critically depends on hyperparameters that typically require ground truth samples for tuning—an issue that becomes particularly problematic for multi-modal distributions.

In this talk, I will discuss this fundamental challenge and introduce Learned Reference-based Diffusion Sampling (LRDS), a novel approach that incorporates prior knowledge of the target distribution's structure to improve sampling efficiency. LRDS consists of two key steps: first, it learns a reference diffusion model (which we see as an energy-based model) tailored to high-density regions of the target distribution; second, it leverages this reference to guide the training of a more effective diffusion-based sampler. I will present experimental results demonstrating that LRDS significantly improves sampling quality compared to existing methods, particularly for challenging multi-modal distributions.