Value Gradient Sampling

Sangwoong Yoon

AI Research Fellow

Korea Institute for Advanced Study (KIAS)

We propose the Value Gradient Sampler (VGS), a trainable sampler inspired by optimal control. VGS generates samples from a given unnormalized density (i.e., energy) by drifting and diffusing randomly initialized particles, similar to Langevin Monte Carlo. Minimizing the KL divergence between the target density and the samples, VGS can synthesize accurate samples in fewer steps. We formulate this KL divergence minimization as an optimal control problem and apply value-based dynamic programming to obtain the optimal drift and diffusion at each sampling step. During sampling, particles drift along the gradient of the learned value function, which is learned using standard reinforcement learning techniques like temporal difference learning. Being a fast, adaptive, and accurate sampling method, VGS can be applied to generate negative samples in contrastive divergence training of energy-based models. We demonstrate the effectiveness of VGS in training energy-based models for industrial anomaly detection.