

Implicit Diffusion: Efficient Optimization through Stochastic Sampling

We present a new algorithm to optimize distributions defined implicitly by parameterized stochastic diffusions. Doing so allows us to modify the outcome distribution of sampling processes by optimizing over their parameters. We introduce a general framework for first-order optimization of these processes, that performs jointly, in a single loop, optimization and sampling steps. This approach is inspired by recent advances in bilevel optimization and automatic implicit differentiation, leveraging the point of view of sampling as optimization over the space of probability distributions. We provide theoretical guarantees on the performance of our method, as well as experimental results demonstrating its effectiveness. We apply it to training energy-based models and finetuning denoising diffusions.

Orateur: Dr BERTHET, Quentin (Google)

Classification de Session: Mini-symposium