An introduction to score-based generative models

Giovanni Conforti and Alain Durmus

January 11, 2024

Abstract

.

In simple words, generative modeling consists in learning a map capable of generating new data instances that resemble a given set of observations, starting from a simple prior distribution, most often a standard Gaussian distribution.

This course aims at providing a mathematical introduction to generative models and in particular to Score-based Generative Models (SGM). SGMs have gained prominence for their ability to generate realistic data across diverse domains, making them a popular tool for researchers and practitioners in machine learning. Participants will learn about the methodological and theoretical foundations, as well as some practical applications associated with these models.

The first two lectures motivate the use of generative models, introduce their formalism and present two simple though relevant examples: energy-based models and Generative Adversarial Networks. In the third and fourth lecture we present score-based diffusion models and explain how they provide an algorithmical framework to the basic idea that sampling from the time-reversal of a diffusion process converts noise into new data instances. We shall do so following two different approaches: a first elementary one that only relies on discrete transition probabilities, and a second one based on stochastic calculus. After this introduction, we derive sharp theoretical guarantees of convergence for score-based diffusion models assembling together ideas coming from stochastic control, functional inequalities and regularity theory for Hamilton-Jacobi-Bellman equations. The course ends with an overview of some of the most recent and sophisticated algorithms such as flow matching and diffusion Schödinger bridges (DSB), which bring an (entropic) optimal transport insight into generative modeling.

Syllabus

- Lecture 1: **Introduction to generative models**
  - Motivations/Applications: molecule generation, Large Language Model (LLM), inverse problem...

  - Simple generative models: Generative Adversarial Networks (GAN) and energy models.

- Lecture 2: **Probabilistic/statistical tools for generative models**
  - Diffusions and their discretization
  - Monte Carlo Markov Chains (MCMC)
  - Divergences on set of probability measures

- Lecture 3: **A "discrete" introduction to score-based diffusion models**
  - Backward/ancestral sampling
  - Score-based diffusion models through variational Inference

- Lecture 4: **A Itô-diffusion viewpoint on score-based diffusion models**
  - Time reversal of diffusion processes
  - Relative entropy on path space
  - Diffusion models as discretization of the time reversal

- Lecture 5: **Convergence bounds for score-based diffusion models via stochastic control**
  - Relative entropy bounds under a finite Fisher information condition,
  - Early stopping rules and bounds under a second moment condition.

- Lecture 6: **Optimal transport and entropic optimal transport for generative modeling**:
  - Entropic OT dynamic and static formulation
  - Sinkhorn algorithm
  - Diffusion Schrödinger bridges (DSB)
  - Flow matching, Diffusion Schrödinger bridge matching...