Chez Pierre

Presents ...

Chez Pierre Seminar

Sitan Chen, Harvard University

"Theory for diffusion models."

Abstract: Diffusion models are now the de facto approach to generative modeling across many types of high-dimensional data, including images, audio, video, molecules, and control policies. Their empirical success has motivated efforts to understand their behavior at a more principled level, both mathematically and phenomenologically, leading to a lively exchange of ideas between machine learning and fields like physics and theoretical computer science.

In this talk I will survey some of our progress along these lines. The first part will overview a key theoretical result that initiated this line of work: given accurate estimation of the so-called "score" of a distribution, a diffusion model can provably generate samples from that distribution, even in the presence of a highly complex or rugged energy landscape. The second part will discuss two intriguing empirical phenomena of diffusion models: the behavior of diffusion models under "guidance," which can be viewed as a controlled perturbation of the sampling dynamics, and the sudden emergence of features in narrow "critical windows" of time during the generation process. In the last part, I will briefly touch upon recent efforts to extend these ideas to language modeling.