
RICARDO BAPTISTA, California Institute of Technology
Dynamics and Memorization Behaviour of Score-Based Diffusion Models

Diffusion models have emerged as a powerful framework for generative modeling that relies on score matching to learn gradients of the data distribution's log-density. A key element for the success of diffusion models is that the optimal score function is not identified when solving the denoising score matching problem. In fact, the optimal score in both unconditioned and conditioned settings leads to a diffusion model that returns to the training samples and effectively memorizes the data distribution. In this presentation, we study the dynamical system associated with the optimal score and describe its long-term behavior relative to the training samples. Lastly, we show the effect of two forms of score function regularization on avoiding memorization: restricting the score's approximation space and early stopping of the training process. These results are numerically validated using distributions with and without densities including image-based problems.