THE HONG KONG UNIVERSITY OF SCIENCE & TECHNOLOGY

Department of Mathematics

SEMINAR ON DATA SCIENCE AND APPLIED MATHEMATICS

The Emergence of Generalizability and Semantic Low-Dim Subspaces in Diffusion Models

By

Prof. Qing QU

University of Michigan

Abstract

Recent empirical studies have shown that diffusion models possess a unique reproducibility property, transiting from memorization to generalization as the number of training samples increases. This demonstrates that diffusion models can effectively learn image distributions and generate new samples. Remarkably, these models achieve this even with a small number of training samples, despite the challenge of large image dimensions, effectively circumventing the curse of dimensionality. In this work, we provide theoretical insights into this phenomenon by leveraging two key empirical observations: (i) the low intrinsic dimensionality of image datasets and (ii) the low-rank property of the denoising autoencoder in trained diffusion models. With these setups, we rigorously demonstrate that optimizing the training loss of diffusion models is equivalent to solving the canonical subspace clustering problem across the training samples. This insight has practical implications for training and controlling diffusion models. Specifically, it enables us to precisely characterize the minimal number of samples necessary for accurately learning the low-rank data support, shedding light on the phase transition from memorization to generalization. Additionally, we empirically establish a correspondence between the subspaces and the semantic representations of image data, which enables one-step, transferrable, efficient image editing. Moreover, our results have profound practical implications for training efficiency and model safety, and they also open up numerous intriguing theoretical questions for future research.

Speaker Bio: Qing Qu is an assistant professor in EECS department at the University of Michigan. Prior to that, he was a Moore-Sloan data science fellow at Center for Data Science, New York University, from 2018 to 2020. He received his Ph.D from Columbia University in Electrical Engineering in Oct. 2018. He received his B.Eng. from Tsinghua University in Jul. 2011, and a M.Sc.from the Johns Hopkins University in Dec. 2012, both in Electrical and Computer Engineering. His research interest lies at the intersection of foundation of data science, machine learning, numerical optimization, and signal/image processing. His current research interse focus on deep representation learning and diffusion models. He is the recipient of Best Student Paper Award at SPARS'15, and the recipient of Microsoft PhD Fellowship in machine learning in 2016, and best paper awards in NeurIPS Diffusion Model Workshop in 2023. He received the NSF Career Award in 2022, and Amazon Research Award (AWS AI) in 2023. He is the program chair of the new Conference on Parsimony & Learning, area chairs of NeurIPS and ICLR.

Date : 19 Dec. 2024 (Thursday)
Time : 10:30a.m.-11:30a.m.
Venue : Room 1409 (near lift 25/26)

All are welcome