UNIVERSITY AT ALBANY
State University of New York

College of Arts and Sciences
Department of Mathematics and Statistics

Data Science and Analysis Seminar

Wei Deng
Morgan Stanley

From Fast Generation to Reasoning: Training Diffusion Language Models With Distillation and RL

Tuesday, April 28, 2026
3:00 P.M. via ZOOM
ID: 951 8640 0117/Pass: 012887

Abstract. Diffusion language models (DLMs) offer a parallel, iterative alternative to autoregressive generation, but remain limited by slow inference and underdeveloped post-training techniques. In this talk, we present two recent advances in ICLR'26 that address these challenges from complementary perspectives. First, DiDi-Instruct introduces a principled distillation framework based on integral KL divergence, enabling few-step generation with up to 64× speedup while maintaining or improving generation quality. Second, GDPO proposes a low-variance reinforcement learning approach for DLMs, improving reasoning performance through efficient sequence-level training. Together, these results highlight a unified view: with proper distillation and optimization, diffusion language models can become both efficient and reasoning-capable.

Wei Deng is a Machine Learning Researcher at Morgan Stanley. He received his Ph.D. from Purdue University in 2021. His current research focuses on fast sampling, reasoning, and the evolution of (diffusion) language models. He also has a strong interest in generative methods for financial applications, including generative hedging. He has served as an Area Chair for both NeurIPS and ICLR.