MATH/STAT 567A: Theoretical Statistics I

Fall 2026, University of Arizona

Course Description

This course introduces a selection of modern machine learning methods, with a strong emphasis on *statistical foundations and insights*. The course will cover the following main topics as much as possible:

1. Deep Learning Theory
   - Approximation: What types of functions or tasks can neural networks represent, and how well can they approximate them?
   - Optimization: Methods for efficiently finding (near-)optimal parameters that enable neural networks to approximate target functions.
   - Generalization: How models trained on observed data perform on unseen data.
2. Foundation Models
   - Evolution of Generative Models: Energy-Based Models (EBM), Variational Autoencoders (VAE), Generative Adversarial Networks (GAN), Denoising Diffusion Probabilistic Models (DDPM), and Score-based SDEs.
   - Large Language Models: Building a Transformer from scratch, covering the full lifecycle including pre-training, fine-tuning, alignment, and trustworthiness.

Instructor & Contact Information

Instructor: Ganghua Wang **Email:** ganghua@uchicago.edu
Office Hours & Location: TBD **Class Meetings:** TBD

Textbook & References

1. *Deep learning theory lecture notes* by Dr. Matus Telgarsky. Electronic copy available here.

2. *Generative AI* by Dr. Jie Ding. Interactive course website at: `https://genai-course.jding.org/index.html`

3. *The Principles of Diffusion Models: From Origins to Advances*

Prerequisites

Undergraduate calculus, linear algebra, and probability theory; familiarity with Python.

Course Policies & Additional Resources

Please find general university policies here.

Disclaimer: This is a tentative syllabus and all information contained herein is subject to change. Any changes will be announced in class and/or via email.