# Statistics Seminar

Department of Mathematics and Statistics

| | |
|---|---|
| **DATE:** | Thursday, April 9, 2026 |
| **TIME:** | 1:30pm – 2:30pm |
| **LOCATION:** | WH 100E |
| **SPEAKER:** | Yiyi Cao, Binghamton University |
| **TITLE:** | Bridging empirical auditing and theoretical privacy guarantees - toward robust and interpretable privacy evaluation in modern machine learning |


**Abstract**

Privacy auditing has emerged as a practical approach for evaluating privacy risks in machine learning models, particularly in black-box or limited-access settings where the training process is unavailable. At the same time, differential privacy provides rigorous theoretical guarantees through frameworks such as Rényi differential privacy, but these guarantees are often difficult to interpret and may not directly reflect observable privacy risks.

This talk first reviews the foundations of differential privacy, including composition techniques and modern privacy accounting, as well as recent advances in empirical privacy auditing, such as membership inference. Despite progress, current approaches exhibit several limitations: empirical auditing methods often produce weak or sample-inefficient signals and are sensitive to probe design and distribution mismatch, while theoretical guarantees are not directly aligned with observable privacy risks. Moreover, translating empirical results into formal guarantees may lead to substantial information loss. Motivated by these gaps, a research agenda is proposed toward bridging empirical auditing and formal privacy guarantees. The goal is to develop more robust auditing methods, more interpretable representations of privacy loss, and principled connections between empirical evidence and theoretical guarantees.