Guanting Chen, UNC

Title: Lean-Nexus: Scaling AI-for-Math to Real-World Mathematical Verification

Abstract: Recent AI-for-Math systems perform well on clean benchmarks but struggle with real-world mathematical documents containing implicit assumptions, heterogeneous notation, long dependency chains, and subtle errors. We introduce Lean-Nexus, an agentic framework for understanding, formalizing, and verifying research-level mathematics directly from LaTeX or PDF. Lean-Nexus iteratively performs interpretation, decomposition, formalization, verification, diagnosis, and repair, allowing it to distinguish formalization mismatches from genuine logical issues. It also produces structured reasoning artifacts, including proof dependencies, error annotations, repair trajectories, and verifier feedback. Experiments on research papers show that Lean-Nexus can identify subtle issues and suggest meaningful corrections, demonstrating a practical path toward scaling AI-for-Math from curated benchmarks to real-world mathematical research.

Bio: Guanting Chen is an Assistant Professor in the Department of Statistics and Operations Research at UNC–Chapel Hill. His research focuses on AI-for-Math, large language models, and sequential decision-making. He is also the founder of RigorisAI, which develops AI-for-Math tools for formalization, verification, and scientific reasoning.

Michael Kay, NCSU (to co-present with **Jake Benhart**)

Title: Pull-Based Engineering Education at Scale Using Agentic AI

Abstract: AI has made producing an engineering result cheap, moving the scarce skill from production to judgment: proposing the right model, and deciding whether the result can be trusted. ISE 754, Logistics Engineering, has been rebuilt for Fall 2026 around that shift. Students read the lecture before class alongside an agent that answers their own questions rather than delivering a prepared explanation, and the class hour is then built from what they could not resolve. Explanation is pulled at the moment of difficulty rather than pushed on a schedule. Assessment follows the same logic: take-home work is done with AI and counts for almost nothing, while most of the grade is earned in class, by hand, without it. Projects go further, casting the student as a consultant whose client releases only what their questions earn, so knowing what to ask is part of the work. Agents are what make this affordable, since each student can be answered on their own difficulty at no marginal cost. What carries it is a private repository per student in NC State's GitHub Enterprise organization: they pull what the course publishes, push their work as the submission, and receive feedback as commits. This semester is the pilot, and the goal is to scale it to much larger classes.

Eric Laber, Duke

Title: Model-assisted Reinforcement Learning

Abstract: Reinforcement learning approaches are categorized as model-based or model-free by their dependence on a system dynamics model. We propose a model-assisted estimator that combines the advantages of model-based estimators with the model-free framework of V-learning. We show that this estimator is doubly-robust, being consistent when either (but not necessarily both) a model-free estimator or a model-based estimator is correctly specified. Furthermore, it is efficient when both estimators are correctly specified. The model-assisted estimator exhibits considerable improvement over model-based and model-free estimators in experiments.

Bio: Laber is the James B. Duke Distinguished Professor of Statistical Sciences at Duke University. His work focuses on the development of methodologies for statistical reinforcement learning with applications in precision medicine, public health, e-sports, and retail. He is also passionate about K-12 STEM outreach. More information about his lab is available at https://laber-labs.com/

Fred Livingston, NCSU

Title: TBD

Shawn Mankad

Title: Confidently Wrong: The Illusion of Precision in LLM-Generated Tabular Data

Abstract: Large language models are increasingly used to generate synthetic tabular data, but this talk shows the resulting data carries a systematic, non-vanishing statistical distortion. We model LLM output as a contamination mixture of memorized and model-generated rows. This provably compresses residual variance and biases regression coefficients, producing an illusion of precision: synthetic-data intervals grow more confident in the wrong answer, rather than widening to reveal it. We then show the same diagnosis motivates a fix: the Covariance-Aligned Refit (CAR) corrects the observable distortion directly, and Prediction-Powered CAR (PP-CAR) adaptively decides how much to trust the corrected signal. Across five real datasets, this yields confidence intervals that are both narrower than reference-only and as valid.

Bio: Shawn Mankad is an Associate Professor of IT & Analytics and the Dean's Scholar at NC State's Poole College of Management. His research develops and applies advanced statistical and machine-learning methods to study business, economic, and policy issues. Before joining NC State, Professor Mankad held faculty positions at Cornell University and the University of Maryland and was a visiting scholar at the Federal Reserve Bank of Philadelphia. He earned a bachelor's degree in mathematics from Carnegie Mellon University and a PhD in statistics from the University of Michigan. Personal Website: http://mankad-research.github.io

Pelin Pekgun

Title: Analytics for Good: Fighting Hunger and Reducing Waste with Data and Analytics

Abstract: Food insecurity remains a persistent humanitarian challenge, with approximately one in seven people in the United States experiencing food insecurity even as up to 40% of the nation's food supply goes uneaten. Addressing this disconnect requires better decisions about how food and other resources are accessed, allocated, and utilized within a complex supply chain. Yet nonprofit food-assistance organizations often operate with limited resources, heterogeneous capacities, incomplete data, fragmented IT systems, and limited analytical and technical expertise. Drawing on research and collaborations with food banks, this talk will explore how data, analytics, and AI can support decision making across food recovery and assistance systems, and what these mission-driven settings reveal about designing analytics solutions that are practical, interpretable, sustainable, and human-centered with appropriate privacy, governance, and ethical safeguards.

Bio: Pelin Pekgün is the Thomas H. Davis Professor in Business, Professor of Analytics, and Department Chair at Wake Forest University School of Business. Her research interests include applications of management science and operations research in supply chain management, pricing and revenue management, marketing/operations interface, and humanitarian and non-profit operations. Her research has appeared in top journals including Manufacturing & Service Operations Management, Production and Operations Management, Journal of Operations Management, and Journal of Marketing. Prior to joining academia, she led the operations research team in North America at JDA Software's Pricing and Revenue Management Group, with a particular focus on hospitality and travel.

Professor Pekgün currently serves as an Associate Editor for Manufacturing & Service Operations Management and a Department Editor for Decision Sciences journals. She also serves as the Secretary on the INFORMS Board of Directors and Executive Committee and as the Vice President, Industry on the Production and Operations Management Society (POMS) Board. She holds M.S. and Ph.D. degrees in industrial and systems engineering from Georgia Institute of Technology, and B.S. and M.S. degrees in industrial engineering from Bogazici University in Istanbul, Turkey.

Munindar Singh, NCSU

Title: Multiagent Systems for Agentic AI

Abstract: This talk will provide a brief survey of some of my recent and ongoing research on multiagent systems that is geared toward LLM-based agents as seen in agentic AI. Key points include (1) how to specify interactions among agents in a way that does not curtail their flexibility yet provides precise guard rails, (2) how to specify and use governance structures such as norms, and (3) how to support human-agent collaboration by endowing agents with the ability to model the mental states of humans and thereby to better understand incomplete and ambiguous instructions.

Bio: Dr. Munindar P. Singh is the SAS Institute Distinguished Professor and an Alumni Distinguished Graduate Professor in the Department of Computer Science at North Carolina State University. Munindar's research addresses how to develop artificial intelligence agents that are responsive to human needs, cooperate effectively with each other, and behave ethically and safely singly and in ensembles.

Munindar is a Fellow of AAAI, AAAS, ACM, and IEEE, a Foreign Member of Academia Europaea (honoris causa), and a recipient of the ACM/SIGAI Autonomous Agents Research Award.

Hong Wan, NCSU

Title: TBD

Serhan Ziya

Title: From Individual Predictions to Aggregate Decisions: Proactive Inpatient Bed Requests from the Emergency Department

Abstract: Emergency department (ED) boarding is a persistent operational challenge and an important contributor to patient length of stay. Prior work has explored using early predictions of hospital admission to initiate inpatient bed requests for individual ED patients before their disposition decisions are finalized. In this work, we consider an alternative approach in which bed-request decisions are made at the aggregate level. At regular decision epochs, information about patients currently in the ED is combined with information about previously requested beds to determine how many additional inpatient beds to request. We formulate this dynamic decision problem as a Markov decision process that balances ED boarding, idling of proactively prepared inpatient beds, and variability in bed requests. Because direct solution of the resulting model is computationally challenging, we develop several decision policies, including an MDP-based heuristic, a newsvendor-based heuristic, and a reinforcement learning approach using deep Q-networks. We evaluate these policies using a simulation model based on data from a large academic medical center. Our results suggest that aggregate proactive bed requests can substantially reduce ED boarding and length of stay while maintaining reasonable levels of inpatient-bed idling. (Joint work with Qian Cheng, Nilay T. Argon, and Aniruddhan Ganesaraman.)

Bio: Serhan Ziya is a professor in the Department of Statistics and Operations Research at the University of North Carolina. He holds a BS degree in Industrial Engineering (1997) from Bogazici University, Istanbul, Turkey and MS and Ph.D. degrees in Industrial and Systems Engineering (1999, 2003) from Georgia Institute of Technology. His research interests are in service operations, with a focus on healthcare operations, queueing systems, revenue management, and pricing. More information can be found from his website at http://ziya.web.unc.edu/.