Wharton Department of Statistics and Data Science
Lawrence D. Brown Distinguished Lectures

Michael I. Jordan

University of California, Berkeley

- Wednesday, October 1, 12:00-1:00 PM
  Location: G60 Jon M. Huntsman Hall

A Collectivist, Economic, and Statistical Perspective on AI

Information technology is in the midst of a revolution in which omnipresent data collection and machine learning are impacting the human world as never before. The word "intelligence" is being used as a North Star for the development of this technology, with human cognition viewed as a baseline. This view neglects the fact that humans are social animals, and that much of our intelligence is social and cultural in origin. Thus, a broader framing is to consider the system level, where the agents in the system, be they computers or humans, are active, they are cooperative, and they wish to obtain value from their participation in learning-based systems. Agents may supply data and other resources to the system only if it is in their interest to do so, and they may be honest and cooperative only if it is in their interest to do so. Critically, intelligence inheres as much in the overall system as it does in individual agents. This is a perspective that is familiar in economics, although without the focus on learning algorithms. A key challenge is thus to bring (micro)economic concepts into contact with foundational issues in the computing and statistical sciences. I'll discuss some concrete examples of problems and solutions at this tripartite interface.

- Thursday, October 2, 1:45-2:45 PM
  Location: 345 Jon M. Huntsman Hall

Sequential Hypothesis Tests and Universal Log-Optimality for General Classes of E-processes

We consider the problem of sequential hypothesis testing by betting. For a general class of composite testing problems---which include bounded mean testing, equal mean testing for bounded random tuples, and some key ingredients of two-sample and independence testing as special cases---we show that any e-process satisfying a certain sublinear regret bound is adaptively, asymptotically, and almost surely log-optimal for a composite alternative. This is a strong notion of optimality that has not previously been established for the aforementioned problems and we provide explicit test supermartingales and e-processes satisfying this notion in the more general case. Furthermore, we derive matching lower and upper bounds on the expected rejection time for the resulting sequential tests in all of these cases. The proofs of these results make weak, algorithm-agnostic moment assumptions and rely on a general-purpose proof technique involving the aforementioned regret and a family of numeraire portfolios.

Joint work with Ian Waudby-Smith and Ricardo Sandoval.

- Friday, October 3, 10:00-11:00 AM
  Location: 370 Jon M. Huntsman Hall

Explorations in Conformal Prediction: Backward Conformal, Multivariate Conformal, Selection, and Decisions

Conformal prediction offers a distribution-free framework for constructing prediction sets with coverage guarantees. The basic theory is limited to marginal coverage, but the promise of conformal prediction hinges on finding approximations to conditional guarantees that are operational in practice, achieving robustness to assumptions, handling multivariate responses, selecting among multiple conformal predictions, and focusing coverage on decision-relevant inference. We discuss progress in all of these areas. Additionally, we present a notion of "backward conformal prediction" in which the size of the prediction set is held fixed and the coverage level is random, arguing that such an inferential concept is relevant in many real-world problems in which people have limited resources to follow up on putative discoveries.

Joint work with Anastasios Angelopoulos, Liviu Aolaritei, Francis Bach, Eugene Berta, Sacha Braun, Aymeric Dieuleveut, Mahmoud Hegazy, and Jordan Lekeufack.