

NYU Stern School of Business
Department of Information, Operations & Management Sciences
IOMS Colloquium Series

TOPIC: Algorithms for Interpretable Machine Learning

SPEAKER: Cynthia Rudin (MIT)

DATE: Wednesday, December 3, 2014

TIME: 12:30 PM – 1:45 PM (Lunch at 12:15 PM)

PLACE: KMC 3-110

ABSTRACT

Possibly **the** most important obstacle in the deployment of predictive models is the fact that humans simply do not trust them. If it is known exactly which variables were important for the prediction and how they were combined, this information can be very powerful in helping to convince people to believe (or not believe) the prediction and make the right decision. In this talk I will discuss algorithms for making these non-black box predictions including:

- 1) "Bayesian Rule Lists" - This algorithm builds a decision list using a probabilistic model over permutations of IF-THEN rules. It competes with the CART algorithm.
- 2) "Supersparse Linear Integer Models" - This algorithm produces scoring systems, which are a type of model that is widely used in the medical community. It proposes an alternative to the Lasso method.

I will show applications to healthcare, including an alternative to the CHADS₂ score, which is one of the most widely used scoring systems in medicine. Our model was trained on over 1000 times the amount of data as CHADS₂ and is more accurate, but is just as interpretable. I will also show preliminary joint work with the MGH Sleep Lab on diagnosing forms of sleep apnea.

At the end of the talk, I will discuss random cool applications of machine learning being worked on in the Prediction Analysis Lab.

Links:

Building Interpretable Classifiers with Rules using Bayesian Analysis

<http://web.mit.edu/rudin/www/LethamRuMcMa14.pdf>

This is joint work with my student Ben Letham and colleagues Tyler McCormick and David Madigan.

Method and Models for Interpretable Linear Classification

<http://arxiv.org/abs/1405.4047>

This is joint work with my student Berk Ustun.

Other links are on my webpage under "Interpretable Predictive Models"

<http://web.mit.edu/rudin/www/PredicticsPapers.html>

BIO:

Cynthia Rudin is an associate professor of statistics at the Massachusetts Institute of Technology associated with the Computer Science and Artificial Intelligence Laboratory and the Sloan School of Management, and directs the Prediction Analysis Lab. Her interests are in machine learning, data mining, applied statistics, and knowledge discovery (Big Data). Her application areas are in energy grid reliability, healthcare, and computational criminology. Previously, Prof. Rudin was an associate research scientist at the Center for Computational Learning Systems at Columbia University, and prior to that, an NSF postdoctoral research fellow at NYU. She holds an undergraduate degree from the University at Buffalo where she received the College of Arts and Sciences Outstanding Senior Award in Sciences and Mathematics, and she received a PhD in applied and computational mathematics from Princeton University in 2004. She is the recipient of the 2013 INFORMS Innovative Applications in Analytics Award. She was given an NSF CAREER award in 2011. Her work has been featured in IEEE Computer, Businessweek, The Wall Street Journal, the Boston Globe, the Times of London, Fox News (Fox & Friends), the Toronto Star, WIRED Science, Yahoo! Shine, U.S. News and World Report, Slashdot, CIO magazine, and on Boston Public Radio. She is presently the chair-elect for the INFORMS Data Mining Section, and currently serves on committees for DARPA, the National Academy of Sciences, the US Department of Justice, and the American Statistical Association.