

RUTGERS UNIVERSITY
Department of Statistics, Department of Computer Science
www.stat.rutgers.edu

*Workshop on Big Data:
Theoretical Foundation of Big Data*

Speaker: **Dr. Robert Schapire**
Microsoft Research and Princeton University

Title: **A fast and simple algorithm for the contextual bandits problem**

Time: **4:10pm – 5:05pm, Thursday, October 16, 2014**

Place: ***CoRE Auditorium***

Abstract

In the contextual bandit learning problem, the learner must repeatedly decide what action to take in response to an observed context, and is then permitted to observe the received reward, but only for the chosen action. The goal is to learn through experience to behave nearly as well as the best policy (or decision rule) in some possibly very large space of possible policies. We assume that the learner can only access this policy space using an oracle for solving empirical cost-sensitive classification problems; in practice, most off-the-shelf classification algorithms could be used for this purpose. In this very general setting, we present a fast and simple algorithm that achieves a regret guarantee that is statistically optimal. Moreover, this algorithm makes very modest use of the oracle, which it calls far less than once per round, on average. These properties suggest this may be the most practical contextual bandit learning algorithm among all existing approaches that are provably effective for general policy classes.

This is joint work with Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford and Lihong Li.

Bio:

Robert Schapire is a Principal Researcher at Microsoft Research in New York City, currently on leave from Princeton University. He received his ScB in math and computer science from Brown University in 1986, and his SM (1988) and PhD (1991) from MIT under the supervision of Ronald Rivest. After a short post-doc at Harvard, he joined the technical staff at AT&T Labs (formerly AT&T Bell Laboratories) in 1991. Since 2002, he has been with the Computer Science Department at Princeton University, and was named the David M. Siegel '83 Professor in Computer Science in 2013. He joined Microsoft Research in 2014. His awards include the

1991 ACM Doctoral Dissertation Award, the 2003 Gödel Prize, and the 2004 Kanelakkis Theory and Practice Award (both of the last two with Yoav Freund). He is a fellow of the AAAI, and a member of the National Academy of Engineering. He mainly studies machine learning especially theoretically well-grounded algorithms, with a particular focus on a constellation of closely related methods and topics that includes boosting, online learning, game theory, and maximum entropy.