Mixture Proportion Estimation for Weakly Supervised Learning*

Clayton Scott

Abstract: Weakly supervised learning has come to refer to a spectrum of problems in machine learning, lying somewhere between supervised and unsupervised learning, where partial or noisy label information is available in some sense. In this talk I will overview many examples of weakly supervised learning problems, and argue that several of them can be solved in terms of a fundamental problem called mixture proportion estimation (MPE). MPE is the problem of inferring the maximum proportion of one distribution that is present in another, given random samples from each distribution. I will discuss several approaches to MPE, including one based on the kernel mean embedding, which exhibits state-of-the-art performance and comes with theoretical guarantees.

Bio: Clayton Scott is Associate Professor in the Department of Electrical Engineering and Computer Science, with a courtesy appointment in the Department of Statistics, at the University of Michigan. He received his undergraduate degree from Harvard University in mathematics, and his master's and doctoral degrees in electrical engineering from Rice University. His research interests include the theory, methods, and applications of statistical machine learning.

*Machine Learning External Seminar, Gatsby Unit, October 12, 2016.