



Data Science Institute COLUMBIA UNIVERSITY

COLLOQUIUM SERIES EVENT FACULTY SPOTLIGHT



John Paisley
Assistant Professor
of Electrical Engineering

THURSDAY, MARCH 23, 2017
5PM-6PM | Davis Auditorium
412 CEPSR, 530 West 120th Street

LIGHT FARE AVAILABLE | OPEN TO THE PUBLIC | REGISTRATION NOT REQUIRED

STRUCTURED AND SCALABLE PROBABILISTIC TOPIC MODELS

ABSTRACT:

Advances in scalable machine learning have made it possible to learn highly structured models on large data sets. In this talk, I will discuss some of our recent work in this direction. I will first briefly review scalable probabilistic topic modeling with stochastic variational inference. I will then then discuss two structured developments of the LDA model in the form of tree-structured topic models and graph-structured topic models. I will present our recent work in each of these areas.

John Paisley joined the Department of Electrical Engineering at Columbia University in Fall 2013 and is an affiliated faculty member of the Data Science Institute at Columbia University. He received the B.S., M.S. and Ph.D. degrees in Electrical and Computer Engineering from Duke University in 2004, 2007 and 2010. He was then a postdoctoral researcher in the Computer Science departments at Princeton University and UC Berkeley, where he worked on developing probabilistic models for large-scale text and image processing applications. He is particularly interested in developing Bayesian models and posterior inference techniques that address the Big Data problem, with applications to data analysis and exploration, recommendation systems, information retrieval, and compressed sensing.

