Tutorial: "Computational Cognitive Science through Probabilistic Programming"
Timothy J. O'Donnell, MIT

Time & Place
Monday, 14.12., 10:00 to 12:00 and 13:00 to 15:00
Tuesday, 15.12., 10:00 to 12:00 and 13:00 to 15:00
Zentrum für Allgemeine Sprachwissenschaft (ZAS), room 403,
Schützenstr. 18, 10117 Berlin

Registration: please register by email to uli@alum.mit.edu
(doesn't apply to members of FB IV of ZAS)

Abstract:
Psychologists, linguists, and other cognitive scientists are increasingly turning to intricately structured probabilistic, statistical, and information-theoretic models to capture the representation, acquisition, and use of human knowledge. While such approaches have proven powerful, their adoption has been limited by their complexity. Such models are tedious to implement, difficult to communicate precisely to other scientists, and challenging to integrate with one another. Recently a new paradigm has emerged to address these issues, known as Probabilistic Programming (http://probabilistic-programming.org/wiki/Home). Probabilistic programming unifies the idea of general purpose programming in familiar languages such as Python or Lisp with probabilistic modeling. This unification allows complex models to be expressed as programs for inherently stochastic virtual machines This perspective, in turn, makes possible inference algorithms that work "out of the box" for large classes of useful probabilistic models.

In this tutorial, we will provide an introduction to computational cognitive science from the perspective of probabilistic programming using the Church language (Goodman, Mansinghka, Roy, Bonawitz, and Tenenbaum, 2008; https://probmods.org/; http://dippl.org/). The tutorial will emphasize major concepts and principles of inference and the importance of highly structured representations such as graphs, grammars and logics. Following the general introduction specific topics of interest to attendees will be explored from a probabilistic programming perspective."

Timothy J. O'Donnell is a research scientist at the Massachusetts Institute of Technology. His work focuses on developing mathematical models of language generalization, learning, and processing, and draws on theoretical ideas from linguistics, experimental methods from psychology, and computational model- ing techniques from natural language processing, artificial intelligence, and machine learning. He is the author of the recent MIT Press monograph "Productivity and Reuse in Language: A Theory of Linguistic Computation and Storage."