Probability, programs, and the mind: Building structured Bayesian models of cognition

Noah D. Goodman (ngoodman@stanford.edu)
Department of Psychology,
Stanford University, Stanford CA 94305 USA

Joshua B. Tenenbaum (jbt@mit.edu)
Department of Brain and Cognitive Sciences,
Massachusetts Institute of Technology, Cambridge MA 02139 USA

Objectives and scope

Human thought is remarkably flexible: we can think about infinitely many different situations despite uncertainty and novelty. Probabilistic models of cognition (Chater, Tenenbaum, & Yuille, 2006) have been successful at explaining a wide variety of learning and reasoning under uncertainty. They have borrowed tools from statistics and machine learning to explain phenomena from perception (Yuille & Kersten, 2006) to language (Chater & Manning, 2006). Traditional symbolic models (e.g. Newell, Shaw, & Simon, 1958; Anderson & Lebiere, 1998), by contrast, excel at explaining the productivity of thought, which follows from compositionality of symbolic representations. Indeed, there has been a gradual move toward more structured probabilistic models (Tenenbaum, Kemp, Griffiths, & Goodman, 2011) that incorporate aspects of symbolic methods into probabilistic modeling. Unfortunately this movement has resulted in a complex "zoo" of Bayesian models. We have recently introduced the idea that using programs, and particularly *probabilistic programs*, as the representational substrate for probabilistic modeling tames this unruly zoo, fully unifies probabilistic with symbolic approaches, and opens new possibilities in cognitive modeling. The goal of this tutorial is to introduce probabilistic models of cognition from the point of view of probabilistic programming, both as a unifying idea for cognitive modeling and as a practical tool.

The probabilistic programming language Church (Goodman, Mansinghka, Roy, Bonawitz, & Tenenbaum, 2008), mathematically grounded on the stochastic λ-calculus, provides a universal language for representing probabilistic models. We will use Church to introduce key ideas and examples of probabilistic modeling. A Church program represents a probabilistic model, and hence inferences that can be drawn from this model, without committing to a process level implementation of inference. This will allow us to focus the tutorial on structured representations and probabilistic inference phenomena without worrying about the details of inference algorithms (such as Markov chain Monte Carlo) that tutorials on Bayesian modeling often become bogged down in. On the other hand, because there are existing inference tools for Church (e.g. Wingate, Stuhlmüller, & Goodman, 2011), students will get hands-on experience with performing inference over different probabilistic models.

The tutorial will include several in-depth case studies where the probabilistic programming viewpoint is particularly useful. These include intuitive theories, such as naive physics and theory of mind, and models of inductive learning that exhibit learning-to-learn and structured abstraction.

Tutorial format

This full-day tutorial aims to introduce students to key ideas of, and new tools for constructing, structured probabilistic models. We will assume only basic familiarity with probability and with programming (i.e. minimal mathematical or statistical background). The tutorial will thus be appropriate for a general Cognitive Science audience, as well for practitioners of bayesian modeling who want to learn about probabilistic programming.

We will teach this tutorial drawing on a combination of infrastructure and materials that we've have used to teach graduate-level classes at Stanford and MIT (and which has been used by others at UCSD and University College Dublin). In particular, students will use the on-line ChurchServ interface to Church, in order to explore these tools without the need to install special software. This interface has been integrated into a Wiki document on Probabilistic Models of Cognition (`http://projects.csail.mit.edu/church/wiki/Probabilistic_Models_of_Cognition`) that we will use for portions of the tutorial.

In addition, we will create new examples focussed on aspects of the approach that we expect to be both new and interesting to a Cognitive Science audience. These include models of physics and vision, based on forward-simulation with standard graphics and vision simulators, and models of language understanding that predict detailed, quantitative human data.

We will use the morning session to introduce the ideas of probabilistic modeling and the Church language, to illustrate basic ideas (such as explaining away, and hierarchical models), and to provide hands-on exercises using Church to create models. The afternoon session will be devoted to case studies of more sophisticated applications of these ideas to cognition, including studies from vision, language, and reasoning.

We, the instructors, have extensive experience in probabilistic modeling of cognition and extensive experience teaching courses and tutorials on these techniques. In addition we are active at the forefront of developing probabilistic programming languages, both conceptually and as practical tools. Both of the instructors have extensive experience teaching tutorials on probabilistic models of cognition specifically from the viewpoint of Church, including courses to graduate students, high-school students, linguists, and psychologists.

Tutorials on Bayesian Models of Inductive Learning have been taught at the Annual Conference of the Cognitive Science Society in 2006, 2008, and 2010 (all co-taught by JBT). This tutorial will be similar in covering the ideas of recent work in Bayesian modeling, but will do so from a different viewpoint and will introduce a different skill set (Church and probabilistic programming). We have presented similar tutorials at the European Summer School For Logic Language and Information 2010 (NDG), the Institute for Pure and Applied Mathematics (NDG and JBT), and several other venues. We will adjust the tutorial based on feedback from those experience as well as the particular audience we expect at Cognitive Science.

References

Anderson, J., & Lebiere, C. (1998). *The atomic components of thought*. Lawrence Erlbaum.

Chater, N., & Manning, C. D. (2006). Probabilistic models of language processing and acquisition. *TRENDS in Cognitive Sciences*, *10*, 335–344.

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006, July). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, *10*(7), 287–291.

Goodman, N. D., Mansinghka, V. K., Roy, D. M., Bonawitz, K., & Tenenbaum, J. B. (2008). Church: a language for generative models. *Uncertainty in Artificial Intelligence*.

Newell, A., Shaw, J., & Simon, H. (1958). Elements of a theory of human problem solving. *Psychological Review*, *65*(3), 151.

Tenenbaum, J., Kemp, C., Griffiths, T., & Goodman, N. (2011). How to grow a mind: Statistics, structure, and abstraction. *science*, *331*(6022), 1279.

Wingate, D., Stuhlmüller, A., & Goodman, N. (2011). Lightweight implementations of probabilistic programming languages via transformational compilation. In *Proceedings of the 14th international conference on artificial intelligence and statistics* (p. 131).

Yuille, A., & Kersten, D. (2006). Vision as bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, *10*, 301–308.