

Probabilistic models of cognition: where next?

Nick Chater¹, Joshua B. Tenenbaum² and Alan Yuille³

¹Department of Psychology, University College London, Gower Street, London, WC1E 6BT, UK

²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, USA

³Department of Statistics, University of California, Los Angeles, CA, USA

This Special Issue surveys the state of the art of probabilistic models across a broad range of topics in cognitive science. We suggest that the present shift towards probabilistic methods has deeper origins: viz., conceptual and technical developments in probability theory and statistics that provide the machinery to engage with cognitively relevant information-processing problems. These technical developments provide a rich range of models, tools and metaphors with which to reconceptualize cognition. Moreover, the application of these probabilistic ideas to relevant engineering problems, in speech and image processing, expert systems, robotics and machine learning, has provided a rich source of insights into some of the probabilistic reasoning problems solved by the brain. Here, we highlight some of the key developments that drive current work, and consider future technical and empirical challenges. We divide our discussion into three, interlinked, domains: representation, processing and learning, before drawing conclusions concerning the prospects for the field.

Representation

Classical probability theory focussed on the narrow domain of repeatable events and processes, where limiting relative frequencies are well-defined – a world of coins, dice and cards, of Markov chains and diffusion processes [1].

Yet, as is clear throughout this Special Issue, the focal domains of cognitive science require rich, compositional representations: the hierarchical structure of the motor system; multiple layers of complexity in visual images; the layers of phonological, syntactic and semantic regularities in language; and representation of rich causal theories of everyday events. Representing these domains requires rich, compositional data structures and processes: networks, grammars, feature-value matrices, schemas, and so on. Sophisticated probabilistic models, as applied in this issue, are now beginning to embrace such complex representations, and to show how they integrate with probabilistic methods. Hence, probabilistic models of vision or language need not be viewed, as formerly, as requiring radical representational simplification (e.g. to sets of binary features that might serve as input to a regression or a connectionist network) – rather, structure and probability can be integrated directly. Note, too, that this application of probability requires a conceptual shift, from viewing probability as modelling repeatable events, to viewing probability as a calculus for uncertain inference. Probability can thus be seen as a model of the

knowledge, and inferential potential, of the individual agent, faced with a stimulus and background information that will typically be unique, rather than repeatable. That is, we require a subjective, rather than frequentist, conception of probability [2].

Nevertheless, substantial representational challenges remain. For example, Bayesian networks [3] provide a powerful formalism for capturing dependency relations between discrete variables, but such variables correspond to ‘atomic’ states of the world (e.g. earthquake/no-earthquake; alarm on/off). Yet capturing inferential relations more fully will require treating states themselves as structured – the earthquake has a location, intensity, time; the alarm is a specific one; and a specific sounding of the alarm might be punctate, intermittent or continuous; the alarm has a range of further properties concerning its loudness, timbre, electrical properties. Both quake and alarm might be related in many, complex ways. It must be possible ultimately to represent such factors in an adequate model of knowledge representation. Similar issues arise whether considering the representation of sensory input or motor output. More than a century of logic and formal semantics has begun to deal with some of this complexity, using a variety of logical formalisms. Probabilistic methods need to enhance and integrate with, rather than by-pass, these powerful tools (see, e.g. [4]). More generally, integrating probability with rich data structures is crucially important in developing a viable framework for cognitive science.

Processing

People have difficulty with many kinds of explicit probabilistic reasoning. Yet important computational insights, especially in relation to Bayesian networks, have revealed how probabilistic inference can naturally map onto distributed, parallel networks of simple processors. How far should the cognitive system be viewed as a probabilistic engine? What architecture does it have? What processing limitations might it possess? And under what conditions might it be possible to harness the brain’s putative probabilistic machinery in explicit probabilistic reasoning tasks [5].

The view that world knowledge might be organized in some kind of probabilistic network raises interesting questions about how reasoning is controlled. One idea is that the structure of the network should express causal dependencies; then different modes of reasoning, such as supposing premises to be true, counterfactual reasoning, imagining, seeking explanations for particular facts, and so on, map on to different ways of intervening in this causal network [6]. It seems plausible that there are fundamental processing limitations at work in controlling

Corresponding author: Chater, N. (n.chater@ucl.ac.uk).

Available online 16 June 2006

these interventions. Can such a viewpoint be mapped onto existing data on human reasoning? (see, e.g. [7]).

Yuille and Kersten (this issue [8]) propose a probabilistic viewpoint suggesting that perception is a process of analysis-by-synthesis, and that probabilistic updating could involve some version of Markov Chain Monte Carlo (MCMC) processing. This viewpoint seems to make strong predictions about the seriality of cognitive processes – presumably, it is not possible to run multiple distributed computations over the same hardware. How might these ideas fit with research on attention, and anatomical localization of processing?

More generally, theories of processing naturally lead to speculations concerning how probabilistic inference might map onto neural hardware. Is MCMC, and faster variants, nonetheless too slow to be neurally plausible? Do reciprocal cortical connections send a top-down ‘synthesis’ signal, as the analysis-by-synthesis viewpoint suggests?

Learning

Learning has frequently been viewed in algorithmic terms – whether concerning principles of association, or principles of memory representation and storage. The probabilistic perspective views learning as an inferential task; for example, inferring the structure of the world from data, or inferring relationships between behaviour, environment and reward.

The fields of statistics and machine-learning have made considerable headway in inferring best-fit parameters, using hill-climbing (and related methods) over likelihood (or similar quantities). More challenging is inferring representational structures over which parameters are optimized. One problem is that the space of possible structures is often large and discontinuous; a second is that a direct application of probabilistic methods would involve assessing each structure by integrating a prior over its parameters, which seems computationally prohibitive; a third is that structures appear to be constrained in potentially highly abstract ways (e.g. the structure of theories in Tenenbaum, Griffiths and Kemp, this issue [9]).

The learnability of structure has three aspects. First, given available data, is enough information available in the input for particular structures to be learnable at all? Minimum description length methods, touched on in Chater and Manning (this issue [10]) might help address this question – structures will be unlearnable if the ‘cost’ of encoding them is greater than the saving in encoding the available data. Thus we can ask: in quantitative terms, what innate constraints are required to make learning possible in principle? Second, is the search through the space of possible structures feasible? What constraints on the search space might make it more feasible? Third, do models of structure-learning fit with neuroscientific and behavioural data concerning how people learn language, visual structure or motor control? Interesting recent work [11] has shown novel ways of

mapping parameters from a learning model into neural structure, which could provide an interesting future research direction.

Probabilistic learning methods typically aim to find a global quantitative fit between model and data; however, the philosophy of science suggests that scientific reasoning typically works by local, qualitative arguments connecting specific data and hypotheses. Should this piecemeal aspect of scientific reasoning be viewed as an inelegant approximation, necessary only in domains for which the cognitive system does not have relevant innate hardware or data availability? Or should qualitative scientific reasoning be viewed as giving clues to how qualitative conclusions can be drawn about the appropriate representational structures, while avoiding apparently intractable quantitative probabilistic computations?

Conclusion

We suggest that cognitive science and artificial intelligence have been undergoing a quiet probabilistic revolution, and the articles in this Special Issue provide evidence of that revolution. Probabilistic ideas provide a rich framework for building models of cognition; and powerful technical tools for building intelligent mechanisms that work. Future research will need to refine and elaborate these technical developments, fuse them with earlier theoretical insights, especially concerning representation, and connect probabilistic models more thoroughly with empirical data.

References

- 1 Grimmer, G.R. and Stirzaker, D.R. (2001) *Probability and Random Processes* (3rd edn), Oxford University Press
- 2 Hájek, Alan, (2003) Interpretations of probability. In *The Stanford Encyclopedia of Philosophy* (Zalta, E.N., ed.), <http://plato.stanford.edu/archives/sum2003/entries/probability-interpret/>
- 3 Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann
- 4 Milch, B. et al. (2005) BLOG: Probabilistic models with unknown objects. In *Proc. 19th Int. Joint Conf. Artif. Intell. (IJCAI)*, pp. 1352–1359, Morgan Kaufmann
- 5 Griffiths, T.L. and Tenenbaum, J.B. Optimal predictions in everyday cognition. *Psychol. Sci.* (in press)
- 6 Pearl, J. (2000) *Causality: Models, Reasoning, and Inference*, Cambridge University Press
- 7 Oaksford, M. and Chater, N. *Bayesian Rationality*, Oxford University Press (in press)
- 8 Yuille, A. and Kersten, D. (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* DOI:10.1016/j.tics.2006.05.002
- 9 Tenenbaum, J.B. et al. (2006) Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* DOI:10.1016/j.tics.2006.05.009
- 10 Chater, N. and Manning, C. (2006) Probabilistic models of language processing and acquisition. *Trends Cogn. Sci.* DOI:10.1016/j.tics.2006.05.006
- 11 Daw, N. et al. (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neurosci.* 8, 1704–1711