

UCLA

Department of Statistics Papers

Title

Probabilistic Models of Cognition: Conceptual Foundations

Permalink

<https://escholarship.org/uc/item/1g84199d>

Authors

Nick Chater

Joshua B. Tenenbaum

Alan Yuille

Publication Date

2011-10-25

Probabilistic Models of Cognition: Conceptual FoundationsNick Chater¹Joshua B. Tenenbaum²& Alan Yuille³

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

²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Building 46-4015, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

³Departments of Statistics and Psychology, UCLA, 8967 Math Sciences Building, Los Angeles, CA 90095-1554, USA

Remarkable progress in the mathematics and computer science of probability has led to a revolution in the scope of probabilistic models. In particular, “sophisticated” probabilistic methods apply to structured relational systems such as graphs and grammars, of immediate relevance to the cognitive sciences. This Special Issue outlines progress in this rapidly developing field, which provides a potentially unifying perspective across a wide range of domains and levels of explanation. Here, we introduce the historical and conceptual foundations of the approach, explore how the approach relates to studies of explicit probabilistic reasoning, and give a brief overview of the field.

The history of probabilistic models of thought is, in a sense, as old as probability theory itself. Probability theory has always had a dual aspect., serving both as a *normative* theory for “correct” reasoning about chance events, but also as a *descriptive* theory of how people reason about uncertainty---as providing an analysis, for example, of the mental processes of an “intelligent” juror. The title of Bernouilli’s great book, *Ars Conjectandi* [1], “The Art of Conjecture,” nicely embodies this ambiguity, suggesting both a “how-to” guide for better reasoning, and a survey of how the “art” is actually practiced. That is, from its origins, probability theory was viewed as both mathematics and psychology.

From a modern perspective, this conflation seems anomalous. Mathematics has shaken free of its psychological roots, and become an autonomous, and highly formal, discipline. The philosophical thesis of “psychologism,” that mathematics (including probability) is a description of thought, fell from favour by the end of the nineteenth century. Moreover, the mathematics and psychology of probability have become divorced. The normative mathematical theory has seen spectacular developments in rigor, generality, and sophistication, going far beyond unaided intuition. Yet the descriptive study of how people judge probabilities has focussed on apparently systematic patterns of fallacious reasoning about chance [2].

This Special Issue is based on the premise that reconciliation is long overdue and that the mathematics of probability is a vital tool in building theories of cognition. The articles in this issue illustrate how probability provides a rich framework for vision and motor control, learning, language processing, reasoning, and beyond. Moreover, probabilistic models can be applied in various ways---ranging from analyzing a problem that the cognitive system faces, to explicating the function of the specific neural processes that solve it. Rather than advocating a monolithic and exclusively probabilistic view of the mind, we suggest instead that probabilistic methods have a range of valuable roles to play in understanding cognition. We hope that this Special Issue will help further inspire researchers in the cognitive and brain sciences to join the project of illuminating cognition from a probabilistic standpoint; and encourage mathematicians, statisticians and computer scientists to deploy the recent remarkable conceptual and computational armoury that they have developed to help understand cognition.

The ubiquity of probabilistic inference

The cognitive sciences view brain as an information processor; and information processing typically involves inferring new information from information that has been derived from the senses, from linguistic input, or from memory. This process of inference from old to new is, outside pure mathematics, typically uncertain. Probability theory is,

in essence, a calculus for uncertain inference, at least according to the subjective interpretation of probability (Box 1). Thus, *prima facie*, probabilistic methods have potentially broad application to uncertain inferences from sensory input to environmental layout; from speech signal to semantic interpretation; from goals to motor output; or from observation to inferring regularities of nature.

Probability has, however, only recently become a major focus of attention in the cognitive sciences. One reason is that the field has often focussed on computational architecture (e.g., symbolic rule-based processing vs. connectionist networks), rather than the nature of the inferences, probabilistic or otherwise, implemented in that architecture. A second reason is that formal approaches to uncertain reasoning in psychology and artificial intelligence have often been studied using non-probabilistic methods, such as default logics, non-monotonic logics, or various heuristic techniques. A third reason is that probabilistic methods have typically been viewed as too restricted in scope to be relevant to cognitive processes defined over linguistic structural descriptions, logical representations, and networks of interconnected processing units. These restrictions have been substantially reduced by remarkable technical progress in the mathematics and computer science of probabilistic models (e.g., [3]).

The focus in this Special Issue is modelling cognitive abilities using such sophisticated forms of probabilistic inference. The term “sophisticated” is intended in at least two ways. First, the knowledge and beliefs of cognitive agents are modeled using probability distributions defined over structured systems of representation, such as graphs, generative grammars, or predicate logic. This development is crucial for making probabilistic models relevant to cognitive science, where structured representations are frequently viewed as theoretically central. Second, the learning and reasoning processes of cognitive agents are modeled using advanced mathematical techniques from statistical estimation, statistical physics, stochastic differential equations, and information theory.

Early examples of sophisticated probabilistic models include Grenander’s pattern theory [4] and Pearl’s work on Bayesian networks [5]. This approach has led to broad advances in the design of intelligent machines, with implications for computer vision, machine learning, speech and language processing, and planning and decision making (see [6], for a survey). Applying these ideas to modeling aspects of human cognition was not straightforward, despite pioneering work by Shepard [7] and Anderson [8]. Indeed, classic work in cognitive psychology by Kahneman, Tversky, and their colleagues suggested that human cognition might be non-rational, non-optimal, and non-probabilistic in fundamental ways (see Box 2).

Yet it seems increasingly plausible that human cognition may be explicable in rational probabilistic terms and that in core domains, human cognition approaches an optimal level of performance based on the statistical properties of the domain. Nevertheless, these new ideas remain unfamiliar to most cognitive scientists and it is only in the last five years that they have started making a significant impact on the field.

Vision is the subfield of cognitive science where these models are most advanced. Mumford [9], and Kersten and Yuille [10], give overviews of this approach from mathematical and cognitive science perspectives. Recent work [11, 12] has used these techniques to extend the classical ideal observer (who performs optimally) to complex stimuli, using Bayesian decision theory, and show that these models can account for many aspects of human visual perception. There have also been successes at formulating the classic Gestalt laws of perceptual organization in terms of probabilistic models [13] that relate to earlier psychological models of grouping and scene perception [14, 15].

Researchers have begun to explore grammatical models of vision using compositionality [16] and developed stochastic grammars for image parsing [17]. These provide links between vision and probabilistic approaches to language processing [18, 19] which are becoming increasingly successful at modeling experiments in psycholinguistics [20]. Indeed, whereas hierarchical symbolic representations have been viewed as problematic for probabilistic approaches, recent work in both vision and language has focussed instead on taking these as the structures over which sophisticated probabilistic models are defined. Specifically, determining which structure is most likely to underlie image or speech data requires using Bayes' theorem to combine a prior probability distribution over structures, and the probability of the data, given each structure. Computing this latter quantity amounts to 'synthesizing' the data, from candidate structures. Thus, in both language and vision, 'analysis-by-synthesis' becomes natural from a probabilistic viewpoint [3, 21].

An advantage of the probabilistic perspective is that it leads to techniques for coupling different sensory modes and for integrating perception with planning. Recent work [22-24] has built on theoretical studies [25] to model the integration of visual and haptic cues, yielding good fits with experimental data. Stankiewicz et al [26] have made use of modeling by Kaelbling et al [27] to design an ideal observer model for how humans navigate through mazes and demonstrated that this model fits data where people navigate mazes in virtual reality.

More recently, these ideas have started making an impact in causal learning and inference. This work has built on artificial intelligence approaches to probabilistic and causal reasoning using Bayesian networks [5, 28, 29]. Cheng's causal power model [30, 31] explains people's judgements about the strength of causal relations, as essentially a form of parameter estimation in a simple Bayesian network. Tenenbaum, Griffiths and colleagues [32-35] show that judgments about causal structure – which variables are causes of which other variables – could be explained using Bayesian model selection among a set of candidate Bayesian networks of the same class. Gopnik and colleagues [36] argue that children's causal learning could be modeled in this way.

Many other cognitive abilities may be explicable within this framework. How people learn the forms and meanings of words from linguistic and perceptual experience has been the subject of recent work [37-40], which draws on, and advances, state-of-the-art techniques developed in information retrieval, computational linguistics, and machine learning [18]. In earlier, related work, Anderson [8] and Shiffrin and Steyvers [41] considered how people form long-term memories, and prioritize the retrieval of memories, as a function of the statistics of their experience with the relevant events. Work on concept learning by Tenenbaum and Griffiths [42-44] showed that many phenomena of inductive generalization and similarity could be explained in terms of Bayesian inference over a hypothesis space of candidate concepts, on the assumption that the observed examples of a concept are a random sample from the concept's extension. In reasoning, work by Oaksford and Chater [45, 46], McKenzie [47, 48], and Krauss, Martignon, and Hoffrage [49] helps explain why people use simple heuristics for certain judgment and decision tasks as approximations to Bayesian inference. This approach relates to theoretical analyses [50-52] showing that simple heuristics can be sometimes serve as surprisingly good approximations.

Finally, studies of the temporal characteristics of human causal learning [53] suggest relationships to stochastic differential equations. Causal relationship models can

be learnt by variants of the Rescorla-Wagner associative learning model. The equilibria of this model have recently been classified [54] and the convergence rates analyzed using stochastic approximation theory [55]. Related techniques have been applied by Dayan and colleagues [56, 57] to analyzing the dynamics of animal learning behavior, and to understanding connections between basic human and animal learning processes. This approaches promises to provide a deeper understanding of learning phenomena that have typically been viewed in purely mechanistic, associative terms. More generally, the probabilistic viewpoint may help explain why the computational and neural mechanisms of the brain have the structure they do.

Levels of probabilistic explanation

Sophisticated probabilistic models can be related to cognitive processes in a variety of ways. This variety can usefully be understood in terms of Marr's [58] celebrated distinction between three levels of computational explanation: The *computational* level, which specifies the nature of the cognitive problem being solved, the information involved in solving it, and the logic by which it can be solved; the *algorithmic* level, which specifies the representations and processes by which solutions to the problem are computed; and the *implementational* level, which specifies how these representations and processes are realized in neural terms.

The probabilistic models and methods described in this Special Issue have potential relevance at each of these levels. As we have noted, the very fact that much cognitive processing is naturally interpreted as uncertain inference immediately highlights the relevance of probabilistic methods at the computational level. This level of analysis is focussed entirely on the nature of the problem being solved---there is no commitment concerning how the cognitive system actually attempts to solve (or approximately to solve) the problem. Thus, a probabilistic viewpoint on the problem of, say, perception or inference, is compatible with the belief that at the algorithmic level, the relevant cognitive processes operate via a set of heuristic tricks, rather than explicit probabilistic computations (e.g. [59, 60]).

One drawback of the heuristics approach, though, is that it is not easy to explain the remarkable generality and flexibility of human cognition [52]. Such flexibility seems to suggest that cognitive problems involving uncertainty may, in some cases at least, be solved by the application of probabilistic methods. Thus, we may take models such as stochastic grammars for language or vision, or Bayesian networks, as candidate hypotheses about cognitive representation. Yet, when scaled-up to real-world problems, full Bayesian computations are intractable, an issue that is routinely faced in engineering applications. From this perspective, the fields of machine learning, artificial intelligence, statistics, informational theory and control theory can be viewed as rich sources of hypotheses concerning tractable, approximate algorithms that might underlie probabilistic cognition.

Finally, turning to the implementational level, one may ask whether the brain itself should be viewed in probabilistic terms. Intriguingly, many of the sophisticated probabilistic models that have been developed with cognitive processes in mind map naturally onto highly distributed, autonomous, and parallel computational architectures, which seem to capture the qualitative features of neural architecture. Indeed, computational neuroscience [61] has attempted to understand the nervous system as implementing probabilistic calculations; and neurophysiological findings, ranging from spike trains in the blow-fly visual system [62], to cells apparently involved in decision

making in monkeys [63] have been interpreted as conveying probabilistic information. How far it is possible to tell an integrated probabilistic story across levels of explanation, or whether the picture is more complex, remains to be determined by future research.

Box 1. Subjective probability in a nutshell

The mathematical properties of probability are relatively uncontroversial. But the *interpretation* of probability is not. Most scientists are familiar with the “frequentist” interpretation [64]: that probabilities are limiting relative frequencies of repeated identical “experiments,” such as coin flips or dice rolls.

Crucially, this interpretation is *not* in play here---in cognitive science applications, probabilities refer to “degrees of belief” [65]. Thus, a person’s degree of belief that a coin that has rolled under the table has come up heads may be around 1/2; this degree of belief may increase rapidly to 1, as she moves her head bringing it into view. Her friend observing the same event, may have different prior assumptions and obtain a different stream of sensory evidence. Thus the two people are viewing the same event, but their belief states and hence their *subjective* probabilities may differ. Moreover, the relevant information is defined by the specific details of the situation. This particular pattern of prior information and evidence will never be repeated; and hence cannot define a limiting frequency.

Probabilistic analyses of perceptual, linguistic, learning or motor tasks typically follow this pattern—the issue is to understand what is believed, and what can be inferred, about the objects in the environment [3], the future state of the motor system [66], the message being conveyed [21], or the regularities linking cause and effect [67, 68].

Why should degrees of belief follow the laws of probability? There are a range of convergent justifications, but two of the more notable are Cox’s axioms and the “Dutch book” argument. Cox proposes several qualitative axioms that any reasonable measure of degree of belief should satisfy, and it can then be proven that only probability measures satisfy those axioms [69]. The “Dutch book” argument suggests that any violation of the laws of probability leads to trouble, e.g., combinations of gambles that each appear fair but which, together, guarantee a loss [65].

The subjective interpretation of probability generally aims to evaluate conditional probabilities, $\Pr(h_j|d)$, i.e., probabilities of alternative hypotheses, h_j (about the state of reality), given certain data, d (e.g., available to the senses). By the definition of conditional probability, for any propositions, A and B , the probability that both are true, $\Pr(A, B)$, is by definition the probability that A is true, $\Pr(A)$, multiplied by the probability that B is true, given that A is true, $\Pr(B|A)$. Applying this identity, simple algebra gives Bayes’ theorem:

$$\Pr(h_j | d) = \frac{\Pr(d | h_j) \Pr(h_j)}{\Pr(d)}$$

The centrality of Bayes theorem to the subjective approach to probability has led to the approach commonly being known as the Bayesian approach. But the real content of the approach is the subjective interpretation of probability; Bayes’ theorem itself is just an elementary, if spectacularly productive, identity in probability theory.

Box 2. How can probability theory be hard for a probabilistic mind?

Termining probabilities as degrees of *belief* invites comparison with the folk psychological notion of belief, in which our everyday accounts of each other's behaviour are formed. This in turn suggests that people might reasonably be expected to introspect about the probabilities associated with their beliefs. In practice, people often appear poor at making such numerical judgments; and poor, too, at numerical probabilistic reasoning problems, where they appear to fall victim to a range of probabilistic fallacies [2]. The fact that people can appear to be such poor probabilists may seem to conflict with the thesis that many aspects of cognition can fruitfully be modelled in probabilistic terms.

Yet this conflict is only apparent. People struggle not just with probability, but with all branches of mathematics. Yet the fact that, e.g., Fourier analysis, is hard to understand does not imply that it, and its generalizations, are not fundamental to audition and vision. The ability to introspect about the operations of the cognitive system are the exception rather than the rule---hence, probabilistic models of cognition do not imply the cognitive naturalness of learning and applying probability theory.

Indeed, probabilistic models may be most applicable to cognitive process that are particularly well-optimized, and which solve the probabilistic problem of interest especially effectively. Thus, vision or motor control may be especially tractable to a probabilistic approach; and our explicit attempts to reason about chance might often, ironically, be poorly modelled by probability theory [70]. Nonetheless, some conscious judgments have proven amenable to probabilistic analyses, such as assessments of covariation or causal efficacy [33, 47], uncertain reasoning over causal models [71, 72], or predicting the extent of everyday events [73]. But unlike textbook probability problems, these are exactly the sorts of critical real-world judgments for which human cognition should be expected to be optimized.

References (main body)

1. Bernoulli, J. (1713) *Ars conjectandi*, Basel: Thurnisiorum.
2. Kahneman, D., Slovic, P., and Tversky, A., Eds (1982) *Judgment under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press
3. Yuille, A. & Kersten, D. (this issue). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences*.
4. Grenander, U. (1993). *General Pattern Theory: A Mathematical Study of Regular Structures*. Oxford University Press Inc., New York, NY.
5. Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufman Publishers.
6. Russell, S. and Norvig, P. (2003) *Artificial Intelligence: A Modern Approach; Second Edition*. Prentice Hall, New Jersey.
7. Shepard, R. N. (1987). Towards a universal law of generalization for psychological science. *Science*, 237:1317-1323
8. Anderson, J. R. (1990). *The adaptive character of thought*. Erlbaum.
9. Mumford, (2002) Pattern Theory: The Mathematics of Perception, in "Proceedings of the International Congress of Mathematicians, Beijing, 2002", vol. 1, Higher Educ. Press, Beijing, 2002.
10. Kersten, D., and Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, 13(2), 1-9
11. Knill, L., and Kersten, D. (1995) Object classification for human and ideal observers. *Vision Research*, 35: 549- 568, 1995.
12. Weiss Y., Simoncelli E.P. and Adelson E.H., (2002) Motion Illusions as Optimal Percepts, *Nature Neuroscience* 5(6):598 – 604.
13. Zhu, S. C. (1999). Embedding Gestalt Laws in Markov Random Fields. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 21(11)
14. Shipley, T.F. and Kellman, P. J., eds. (2001) *From Fragments to Objects: Segmentation and Grouping in Vision*. Amsterdam, The Netherlands: Elsevier Science.
15. Chater, N. (1996) Reconciling simplicity and likelihood principles in perceptual organization. *Psychological Review* 103, 566-581.

16. S. Geman, D., Potter, and Chi, Z. (2002) Composition systems. *Quarterly of Applied Mathematics*, LX, 2002, 707-736
17. Tu, Z., Chen, X., Yuille, A., and Zhu, SC. (2005) Image Parsing: Unifying Segmentation, Detection, and Object Recognition, *International Journal of Computer Vision*, 63(2), 113–140
18. Manning, C., & Schütze, H. (2000) *Foundations of Statistical Natural Language Processing*, MIT Press
19. Geman, S & Johnson, M (2003) Probability and statistics in computational linguistics, a brief review. In: *Mathematical foundations of speech and language processing*. (Johnson, M., Khudanpur, S., Ostendorf, M., and Rosenfeld, R., eds) *Volumes in Mathematics and its Applications*, Vol. 138, Springer-Verlag, New York, 2003, 1-26
20. Jurafsky, D. (2003) Probabilistic Modelling in Psycholinguistics: Linguistic Comprehension and Production. In *Probabilistic Linguistics* (Bod, R., Hay, J., and Jannedy, S., Eds), MIT Press.
21. Chater, N. & Manning, C. (this issue). Probabilistic models of language processing and acquisition. *Trends in Cognitive Science*.
22. Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415, 429-433.
23. Schrater, P. R., and Kersten, D. (2000). How Optimal Depth Cue Integration Depends on the Task. *International Journal of Computer Vision*, 40: 71-89.
24. Schrater, P. R., and Kersten, D. (2000). Vision, Psychophysics, and Bayes. In: (R. Rao, B. Olshausen, and M. Lewicki, Ed.) *Statistical Theories of the Brain*, MIT Press.
25. Clark, J.J., and Yuille, A.L., (1990) *Data Fusion for Sensory Information Processing Systems*, Kluwer Academic Publishers, Boston, MA.
26. Stankiewicz, B.J., Legge, G.E., Mansfield, J.S. and Schlicht, E.J. (in press). Lost in Virtual Space: Human and Ideal Wayfinding Behavior. *Journal of Experimental Psychology: Human Perception and Performance*.
27. Kaelbling, L., Littman, M., and Cassandra, A. (1998). Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence*. 101.
28. Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge: Cambridge University Press.
29. Spirtes, P., Glymour, C., and Scheines, R (2000) *Causation, Prediction and Search*. MIT Press.

30. Novick, L. R. & Cheng, P. W. (2004). Assessing interactive causal influence. *Psychological Review*, 111, 455-485.
31. Cheng, P.W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367-405.
32. Tenenbaum, J. B., and Griffiths, T. L. (2001) Structure learning in human causal induction, in *Advances in Neural Information Processing Systems 13*. (Leen, T., Dietterich, T., and Tresp, V., eds). Cambridge, MIT Press, pp 59-65.
33. Griffiths, T. L., and Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology* 51(4), 285-386.
34. Steyvers, M., Tenenbaum, J. B., Wagenmakers, E., and Blum, B. (2003). Inferring causal networks through observations and interventions. *Cognitive Science*, 27, 453-489.
35. Tenenbaum, J. B. and Griffiths, T. L. (2003). Theory-based causal inference. In S. Becker, S. Thrun, and K. Obermayer (eds.), *Advances in Neural Information Processing Systems 15*. Cambridge, MA: MIT Press.
36. Gopnik, A. Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T, and Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111, 1-30
37. Brent, M. R. (1999). Speech segmentation and word discovery: a computational perspective. *Trends in Cognitive Sciences*, 3, 294-301.
38. Tenenbaum, J. B. and Xu, F. (2000). Word learning as Bayesian inference. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, August 2000.
39. Griffiths, T.L., and Steyvers, M. (2003) Prediction and semantic association. *Advances in Neural Information Processing Systems 15*.
40. Blei, D., Griffiths, T. L., Jordan, M. I., and Tenenbaum J. B., (2004). Hierarchical topic models and the nested Chinese restaurant process. In *Advances in Neural Information Processing Systems 16*.
41. Shiffrin, R. M., and Steyvers, M. (1997). A model for recognition memory: REM: Retrieving effectively from memory. *Psychonomic Bulletin and Review*, 4, 145-166.
42. Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. In *Advances in Neural Information Processing Systems 11*. (M. Kearns, S. Solla, and D. Cohn, eds.), Cambridge, MA: MIT Press.

43. Tenenbaum, J. B. (2000). Rules and similarity in concept learning. In *Advances in Neural Information Processing Systems 12*. (S. Solla, T. Leen, and K. R. Mueller, eds.), Cambridge, MA: MIT Press.
44. Tenenbaum, J. B. and Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24, 629-641.
45. Oaksford, M. and Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 608-631.
46. Chater, N., and Oaksford, M. (1999). The probability heuristics model of syllogistic reasoning. *Cognitive Psychology*, 38, 191-258
47. McKenzie, C. R. M. (2003). Rational models as theories -- not standards -- of behavior. *Trends in Cognitive Sciences*, 7, 403-406
48. McKenzie, C. R. M., and Mikkelsen, L. A. (2000). The psychological side of Hempel's paradox of confirmation. *Psychonomic Bulletin and Review*, 7, 360-366
49. Krauss, S., Martignon, L. and Hoffrage U. (1999). Simplifying Bayesian inference: The general case. In *Model-based reasoning in scientific discovery* (L. Magnani, N. J. Nersessian and P. Thagard, Eds.) (pp. 165–179). New York: Kluwer Academic/Plenum Publishers
50. Pearl, J. (1984) *Heuristics: Intelligent Search Strategies for Computer Problem Solving* Addison-Wesley, Reading, MA..
51. Coughlan, J. M. & Yuille, A. L. (2002). Bayesian A* Tree Search with Expected O(N) Node Expansions: Applications to Road Tracking. *Neural Computation*, 14, 1929-1958.
52. Chater, N., Oaksford, M., Nakisa, R. and Redington, M. (2003). Fast, frugal and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes*, 90, 63-86.
53. Danks, D., Griffiths, T.L., and Tenenbaum, J.B. (2003). Dynamical causal learning, *Advances in Neural Information Processing Systems 15*. (Becker, S., Thrun, S., and Obermayer, K. (eds). (pp. 67-74). Cambridge, MIT Press.
54. Danks, D (2003) Equilibria of the Rescorla-Wagner Model *Journal of Mathematical Psychology*, 47: 109-121
55. Yuille, A. L. (2004). The Rescorla-Wagner Algorithm and Maximum Likelihood Estimation of Causal Parameters. *Advances in Neural Information Processing Systems*, 17.
56. Kakade, S and Dayan, P (2002). Dopamine: Generalization and bonuses. *Neural Networks* 15, 549-559.

57. Dayan, P., Kakade, S. & Montague, P. R. (2000). Learning and selective attention. *Nature Neuroscience*, 3 , 1218-1223.
58. Marr, David (1982) *Vision*. San Francisco: W.H. Freeman.
59. Churchland, P. S., Ramachandran, V. S., & Sejnowski, T. J. (1994). A Critique of Pure Vision. In C. Koch & J. L. Davis (Eds.), *Large Scale Neuronal Theories of the Brain* (pp. 23-60). Cambridge, MA: MIT Press.
60. Gigerenzer, G./Todd, P.M./ABC Group, eds. (1999). *Simple Heuristics that Make Us Smart*. New York: Oxford University Press
61. Dayan, P, and Abbott, L.F., (2001) *Theoretical Neuroscience: Computational and Mathematical Modelling of Neural Systems*, MIT Press.
62. Rieke, F. et al. (1999). *Spikes*. MIT Press.
63. Shadlen, MN, Gold JI (2004) The neurophysiology of decision-making as a window on cognition. In *The Cognitive Neurosciences*, 3rd edition. (Gazzaniga MS, ed): MIT Press.

References (box 1)

64. von Mises, R (1981) *Probability, Statistics and Truth*, 2nd rev. English ed., New York, Dover, 1981.
65. Ramsey, F. P. (1931) *The Foundations of Mathematics*. London: Kegan Paul, Trench, Trubner
66. Koerding, K. & Wolpert, D. (this issue). TITLE. *Trends in Cognitive Sciences*.
67. Courville, A. C., Daw, N. D. & Touretzky, D. S. (this issue). Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences*.
68. Tenenbaum, J. B., Griffiths, T. & Kemp, C. (this issue). TITLE. *Trends in Cognitive Sciences*.
69. Cox, R. T. (1961). *The Algebra of Probable Inference*, Johns Hopkins University Press.

References (box 2)

70. Stewart, N., Chater, N. & Brown, G. D. A. (in press). Decision by sampling. *Cognitive Psychology*.
71. Krynski, T. R. and Tenenbaum, J. B. (2003). The role of causal models in reasoning under uncertainty. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*.

72. Sloman, S. and Lagnado, D. (2004). Do we “do”? *Cognitive Science*, 29, 5-39.
73. Griffiths, T. L. and Tenenbaum, J. B. (in press). Optimal predictions in everyday cognition. *Psychological Science*.