

Computational Rationality for Human-Like Computing

Andrew Howes

University of Birmingham

If research on human-like computing is to be successful then it must be guided by a rigorous understanding of the phenomena that constitute human behaviour. Fortunately, some recent advances in cognitive science have been founded on the application of bounded optimality [9] – an idea borrowed from machine learning – to the challenge of developing computational models of human cognition [3, 5, 7]. *Computational rationality* develops the idea that people not only adapt to the structure of the environment, but also to that of the mind and brain itself. This work has offered transformative explanations of behavioural phenomena in at least three areas: decision making, dual tasking, and human-computer interaction [4–6]. By providing computational and empirically verified theories of human behaviour, these models define an agenda for human-like computing.

One recent application of computational rationality [6] was to the analysis of decision problems, including preference reversals, that were previously thought to provide evidence that people are irrational and that human behaviour is inconsistent with the axioms of rational choice. Howes et al. show, with computational modelling, that in fact people behave entirely rationally once the constraints, particularly uncertainty due to processing noise, imposed by the cognitive-neural system are taken into account in the definition of the optimisation problem. Howes et al.’s analysis is important, in part, because it suggests that heuristic models built to directly fit computer programs to apparent human biases are not on the path to building human-like computing. Rather, algorithms, that are grounded in rigorous theoretical analysis of decision problems are sufficient to explain a wide range of human behaviour.

Another application of computational rationality has been to the problem of understanding human visual search [1, 2, 8]. By defining visual search as a Partially Observable Markov Decision Problem (POMDP) we have shown that well known visual search phenomena are explained by deriving optimal strategies for eye movements given constraints on visual acuity. Human behaviour is shown to be a consequence of adaptation to the goal of making detection decisions given noise in the human visual system. This work marks a departure from accounts of human vision that fail to offer an optimal integration of bottom-up and top-down guidance. However, despite the success of our work, far more powerful machine learning techniques are required to find optimal solutions to a broad range of human vision problems. We are, therefore, currently exploring the potential role of deep reinforcement learning in these models.

Both of the examples above offer illustrations of the potential for computational rationality to frame the problem of how to use machine learning for human-like computing. Building on these examples my talk will seek to provide

an agenda of research in this area. In particular it will focus on (1) what are, and are not, important phenomena in the definition of ‘human-like’ and (2) how to deploy machine learning methods that generate these phenomena.

References

1. Butko, N.J., Movellan, J.R.: I-POMDP: An infomax model of eye movement. In: 2008 7th IEEE International Conference on Development and Learning. pp. 139–144. IEEE (2008)
2. Chen, X, Howes, A., Lewis, R.L., Myers, C.W., Houpt, J.: Discovering computationally rational eye movements in the distractor ratio task. In: The 1st Multidisciplinary Conference on Reinforcement Learning and Decision Making. Oct 25-27; Princeton University, Princeton, New Jersey, USA. (2013), <http://rldm.org/past-meetings/rldm2013/>
3. Griffiths, T.L., Lieder, F., Goodman, N.D.: Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in cognitive science* 7(2), 217–229 (2015)
4. Howes, A., Duggan, G.B., Kalidindi, K., Tseng, Y.C., Lewis, R.L.: Predicting short-term remembering as boundedly optimal strategy choice. *Cognitive Science* 40(5), 1192–1223 (2016)
5. Howes, A., Lewis, R.R.L., Vera, A.: Rational adaptation under task and processing constraints: implications for testing theories of cognition and action. *Psychological Review* 116(4), 717 (2009)
6. Howes, A., Warren, P., Farmer, G., El-Deredy, W., Lewis, R.: Why contextual preference reversals maximize expected value. *Psychological Review* 123(4), 368–391 (2016)
7. Lewis, R., Howes, A., Singh, S.: Computational rationality: linking mechanism and behavior through bounded utility maximization. *Topics in Cognitive Science* 6(2), 279–311 (2014)
8. Myers, C.W., Lewis, R.L. Howes, A.: Bounded optimal state estimation and control in visual search: explaining distractor ratio effects. In: M. Knauff, M. Pauen, N. Sebanz, & I.W.E. (ed.) *Proceedings of the 35th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society (2013)
9. Russell, S., Subramanian, D.: Provably bounded-optimal agents. *Journal of Artificial Intelligence Research* 2, 575–609 (1995)