Thirty years of Marr's Vision: Levels of Analysis in Cognitive Science

Participants

Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo,
200 University Avenue West, Waterloo, Ontario, Canada,
N2L 3G1

Tom Griffiths (tom_griffiths@berkeley.edu)
Institute of Cognitive and Brain Sciences, University of California, Berkeley,
132 Barker Hall, MC 3190, Berkeley, CA 94720, USA

Valerie Gray Hardcastle (valerie.hardcastle@uc.edu)
Departments of Philosophy and Psychology, University of Cincinnati,
Cincinnati, OH 45221, USA

Bradley C. Love (b.love@ucl.ac.uk)
Department of Cognitive, Perceptual and Brain Sciences,
University College London,
26 Bedford Way, London, UK, WC1H 0AP

Discussant

William Bechtel (bill@mechanism.ucsd.edu)
Department of Philosophy, University of California, San Diego,
9500 Gilman Drive, La Jolla, CA 92093, USA

Organisers

Richard P. Cooper (R.Cooper@bbk.ac.uk)
Department of Psychological Sciences,
Birkbeck, University of London

David Peebles (d.peebles@hud.ac.uk)
Department of Behavioural and Social Sciences
University of Huddersfield

Keywords: Marr; levels of description and explanation; computational level; algorithmic and representational level; implementational level.

Introduction

Thirty years after Marr's landmark posthumous book, *Vision* (Marr, 1982), the argument for which he is most cited remains the distinction between computational, algorithmic and representational, and the implementation levels. In the interim, many reformulations of this basic distinction have been proposed, but is it still relevant? This symposium will discuss whether there is still a place for the algorithmic and representational level, with its cognitive-level concepts, given the rise in reductionist neuroscience from below and Bayesian analysis from above.

Marr's Attacks: A Gentle Reminder
Chris Eliasmith

Marr's (1982) three levels can be seen as the result of two deep concerns he had about how brain theories were being constructed in his day. We can see these concerns giving rise to 1) an attack on reductionism; and 2) an attack on vagueness.

With respect to reductionism, Marr was interested in ensuring the centrality of not only mechanisms, but also of their function to our generation of brain theories (Marr and Poggio, 1977). With respect to vagueness, Marr wanted to ensure that our high-level descriptions of neural phenomena could be tested against empirical data (Marr, 1975).

Unfortunately, many researchers after Marr seem to have taken his purpose to be a divisive one. Some, such as Pylyshyn (1984), refer to the "three autonomous levels of description" (p. 259). In contrast, Marr (1982) seems to be suggesting that the intermediate, representational level, is a *bridge* between our more abstract characterizations, and more detailed characterizations (pp. 23-24).

I argue that Marr's levels should be understood in an integrative sense. I show that adopting this perspective provides critical constraints for building Marr-type brain models. I provide the details of one such model: a large-scale simulation of spiking neurons that reproduces detailed neural and behavioral results across a wide array of cognitive and non-cognitive tasks.

In short, adopting Marr's perspective on levels helps pave the way for the kind of unified models of brain function for which he, himself, was striving.

Bridging Levels of Analysis for Probabilistic Models of Cognition
Tom Griffiths

Most probabilistic models of cognition are intended to explain human behavior at the computational level, linking how people act to the solution to an abstract computational problem. This focus is quite different from that of other approaches to cognitive modeling, which tend to emphasize the algorithmic and implementational levels. This raises a number of important questions: When are theories at these different levels incompatible with one another? What are the implications of a computational-level analysis for theories at the other levels? How can we begin to draw connections across levels of analysis, for an integrated account of cognition? I will argue that we can only answer these questions by explicitly taking on the challenge of building a bridge between levels of analysis, considering how computational-level models can be translated to the algorithmic and implementational levels and how algorithmic- and implementational-level accounts might be cast at the computational level. I will illustrate this argument with examples drawn from recent work looking at Monte Carlo methods as a source of "rational process models" and analyses of the computational-level commitments of artificial neural networks.

A New Appreciation for Marr's Levels: Understanding How Brains Break
Valerie Gray Hardcastle and Kiah Hardcastle

Much work in the cognitive sciences, including computational neuroscience, now focuses on brains performing less than optimally. That is, while the original programs in artificial intelligence and the like aimed to articulate what thought was in ideal terms, much research now looks at how and why brains or other cognitive engines fail to function as they should. This focus on impairment affects how one can understand Marr's three levels. In this presentation, we use a method of exploring impulsivity and behavioral inhibition based on a neural network/ population activity model of the cortico-striatal circuitry as a case study to refine Marr's distinctions. In particular, we will show that the computational level should be redefined, for simply knowing the goal of a computation may not tell us much about why something has gone wrong and why the information-processing device is exhibiting abnormal behavior. We will also argue, as have many others, that the distinction between algorithm and hardware largely collapses when considering the brain.

The Primacy of Mechanism in Cognitive Science
Brad Love

Cognitive science is primarily concerned with the "how" questions of brain and behavior. These questions address mechanism, and therefore make contact with Marr's algorithmic level. From below, mechanistic accounts can be informed, constrained, and inspired by neuroscience. Rather than being reduced by neuroscience, cognitive models are proving valuable in interpreting fMRI data because these mechanisms help neuroscientists understand the function of brain regions. From above, despite many cognitive scientists professing a devotion to the computational level, very few are trained or focus their research on characterizing evolutionary environments, niches, and histories. I will argue that explanations formulated purely at the computational level are not sufficiently constrained, because rational Bayesian models are uninformed by a wide range of process-level data and their assumptions about the environment are generally not grounded in empirical measurement.

Given the recent surge of interest in computational-level theories of cognition, one question is whether integration across algorithmic and computational levels would be beneficial. One promising avenue for integration is to evaluate the representations on which Bayesian inference operates and the algorithms and heuristics that carry it out as psychological mechanisms. In other words, one means of integration is to evaluate Bayesian models at the algorithmic level. A number of researchers have adopted this strategy and have concluded that humans engage in forms of approximate Bayesian inference that are intended to reflect human capacity limitations. Although an improvement over purely rational approaches, approximate Bayesian models face significant challenges. One challenge is that people are suboptimal for reasons other than capacity limitations. In domains where people's behavior falls far short of that predicted by rational accounts, the layering of capacity limitations and suboptimalities onto the rational account may only serve as a lengthy detour to the algorithmic level.

Differentiating While Integrating Levels
William Bechtel

Are all three of Marr's levels needed? Should they be kept distinct? Symposiasts emphasize how cognitive science is or should integrate Marr's levels. This is important, but it is also important to emphasize the distinct contributions and methodologies of each level of inquiry. They represent three different perspectives required to understand mechanisms generally, but especially information processing mechanisms. Marr viewed neuroscience of his day as emphasizing the material implementation at the expense of the algorithmic-representational and computational levels, and that has been true of mechanistic science generally. But mechanisms only work insofar as they are organized, and this is especially true of information processing mechanisms that must insure that information is encoded appropriately within the mechanism and made available to the operations that require it. Moreover, it is crucial to understand how a mechanism functions in broader environments that determine the computations it needs to perform (and may fail to perform). Different modes of inquiry are required to examine each of these. This is especially true of the computational perspective, which requires looking outside the mechanism to the environment in which it operates and engaging in appropriate experimental and theoretical studies to understand what those demands really are.

References

Marr, D. (1975). Approaches to biological information processing. *Science*, 190:875-876.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco, CA: W. H. Freeman.

Marr, D. and Poggio, T. (1977). From understanding computation to understanding neural circuitry. *Neurosciences Research Program Bulletin*, *15*, 470-488.

Pylyshyn, Z. (1984), *Computation and Cognition*, Cambridge, MA: MIT Press.