

CROSS TALK

Cross Talk opposing view: Marr's three levels of analysis are not useful as a framework for neuroscienceJonathan W. Pillow 

Princeton Neuroscience Institute Princeton University, Princeton, New Jersey, USA

Email: pillow@princeton.edu

Handling Editors: Francisco Sepúlveda & Vaughan Macefield

The peer review history is available in the Supporting Information section of this article (<https://doi.org/10.1113/JP279550#support-information-section>).

Introduction

Let me begin by briefly describing how I came to this debate. Scientifically, I grew up revering the three levels. Among my wonderful mentors were many who venerated Marr's framework, and I readily embraced it as well. For many years I sought to describe my own research in Marrian terms, and cheerfully expounded on the Three Levels in lectures and summer courses.

However, as time went on I began to notice the framework's shortcomings. The Three Levels seemed inapplicable to many forms of inquiry. Marr's treatise itself felt overly didactic in prescribing a 'correct' approach to the nervous system. The rise of deep learning and failure of Marr's own program for understanding vision – outlined in the same 1982 book as the Three Levels of Analysis – tipped me further into opposition.

Despite these shortcomings, Marr's framework continues to enjoy widespread acclaim, particularly among theoretical, computational and cognitive neuroscientists¹. My current position 'against' the levels is primarily a reaction to this over-enthusiasm. So to be clear: I do not wish to suggest that the Three Levels have no value, or that we should set them aside and never speak of them again. I believe

Marr's book played a valuable role in promoting computational and theoretical approaches to the study of vision. And if you have never heard of the Three Levels, I would urge you to read Mátè Lengyel's companion piece with an open mind. However, I feel the Three Levels framework has, overall, proven less useful than hoped, and will outline three primary criticisms in the sections below.

'Computational' favouritism. My first criticism of Marr's framework is its undue favouritism for the 'computational' level, at the expense of the other two levels ('algorithmic' and 'implementation'). We can find level-1 chauvinism throughout the first chapter of *Vision* (Marr, 1982). For example, in a section entitled 'The Importance of Theory', Marr states:

[I]t is the top level, the level of computational theory, which is critically important from an information-processing point of view. The reason for this is that the nature of computations that underlie perception depends more upon the computational problems that have to be solved than upon the particular hardware in which their solutions are implemented. To phrase the matter another way, an algorithm is likely to be understood more readily by understanding the nature of the problem being solved than by examining the mechanism (and the hardware) in which it is embodied.

— Marr (1982), pg 27.

And of course there is the famous analogy to birds and feathers:

' [T]rying to understand perception by studying only neurons is like trying to understand bird flight by studying only feathers: It just cannot be done.' — Marr (1982), pg 27.

In my reading, the Three Levels of Analysis framework carries a clear **insinuation** that the computational level is superior to the other levels. Moreover, Marr's writing suggests that the 'correct' way to study the nervous system is to start with the computational analysis of a problem

before proceeding to the algorithm or its implementation in the brain.

As a sidebar, I would point out that Marr is most beloved by neuroscientists and cognitive scientists whose work sits at the computational level. It's only natural that these scientists should embrace Marr, since he affirms that their approach to studying the brain is the best! (**As an exercise, I urge you to notice that a talk that begins with a discussion of Marr's three levels is almost always a talk that seeks to tell you why the computational approach is the best approach for understanding a particular phenomenon**). Marr is rarely invoked in talks about anatomy or biophysics! Naturally, I object to this perspective, since my own work tends to focus on data-driven rather than theory-driven approaches.

More seriously, however, is the objection that Marr's prescription was simply wrong. His idea that we should start with a theory-driven understanding of the problem, and only later look to data-driven approaches for evidence of implementation, has not in general led to progress in understanding the difficult computational problems faced by the nervous system. Certainly, there are many cases of relatively simple problems where computation-first approaches have borne fruit, such as sound localization using interaural timing differences (see Fig. 1). However, Marr's computation-first approach to vision led him to formulate algorithms that relied on zero-crossings, 'blobs' and the so-called '2½D sketch', which proved to be dead ends of computer vision (see Fig. 1). More broadly, data-driven approaches have yielded a plethora of basic discoveries about the organization of the nervous system that were not anticipated by computational considerations of the problem (e.g. grid cells, 'up' and 'down states', and hippocampal replay, to name just a few). In my view, this adds weight to the position that many biological algorithms might **in many cases** be understood more readily by examining the hardware in which they are embodied.

Lack of separation between levels in biological systems. A second criticism of the Three Levels of Analysis framework is Marr's assertion that the levels could (and should) be considered independently:

¹'Drinking game: take a shot every time a presentation starts w Marr's levels of analysis #CCN2017.' [Collected on twitter, September 2017].

These three levels are coupled, but only loosely. ... [T]here is a wide choice available at each level, and the explication of each level involves issues that are rather independent of the other two.

— Marr (1982), pg 25.

Marr's proposed separation of levels makes sense when thinking about digital computers. For example, introductory courses in computer science will discuss the sorting problem (i.e. the problem of putting the elements of a list into an order), and then cover specific algorithms for sorting ('quicksort', 'bubble sort', 'insertion sort', etc). It is also straightforward to separate the consideration of an algorithm and its implementation in a particular hardware in digital computers.

However, it is much less obvious that this form of separation is possible or even desirable in the study of biological systems. As Tony Movshon put it:

David Marr said that people who study vision should separate their thoughts about the biology from their thoughts about the

computation. He wanted you to think about the algorithm and the implementation as two different problems... For a connectionist, that's anathema, because if you believe that the connection is the computation, then what Marr is saying is absolutely wrong.

— Tony Movshon (*Seung & Movshon, 2012*).

The proposed separation is belied by fields such as neuromorphic computing. Moreover, it is interesting to note that Marr proposed the Three Levels of Analysis framework before the rise of deep learning, in which the architecture of a neural network and the learning rule used to train it are often fundamental to the network's computational abilities.

Arguments against the separation of levels have also been put forth by proponents of probabilistic approaches, who have noted that 'resource-rational' accounts of cognition often require consideration of approximations and algorithmic constraints, which do not arise at the computational level (Griffiths et al., 2015).

An arbitrary and incomplete taxonomy. A third criticism is that the three levels proposed by Marr do not cleanly or neatly divide the space of approaches to studying the nervous system. The attempt to characterize neuroscience research in terms of these levels is therefore often unproductive or meaningless.

To give just one example, consider the case of recurrent neural networks (RNNs), which have surged in popularity for studying neural computation. To which Marr level of analysis do RNN-based approaches belong? Some have argued that RNNs sit at the computational level (e.g. Ritter et al., 2017). In my reading, Marr himself would disagree with this assignment; Marr wrote that 'for far too long, a heuristic program for carrying out some task was held to be a theory of that task' and he expressed dismay about the absence of computational theory in settings for which there is 'no way to determine whether a program would deal with a particular case other than by running the program' (Marr, 1982, pg 28). It seems to me one could plausibly make the case that RNNs sit at either the computational, algorithmic, or the implementational level, depending on one's framing of the problem. Ultimately, however, my experience has been that such debates do not produce meaningful answers because the three levels simply do not provide a useful taxonomy of approaches.

As additional evidence for this conclusion, the literature is rife with alternative schemes, most of which have either fewer or more than three levels. Earlier work from Marr himself, with his long-time collaborator Tomaso Poggio (Marr & Poggio, 1976), proposed *four* distinct levels of analysis: (1) computation; (2) algorithms; (3) mechanisms; and (4) hardware. Poggio later proposed 'learning' as an even higher level of analysis ('I propose – and I am sure David would agree – that learning should be added to the list of levels of understanding, above the computational level'), along with 'Evolution', 'Wetware', 'Hardware' and 'Circuits and Components', bringing the proposed number of levels to seven (Poggio, 2012). Others have argued that – in the case of analogue computation – the algorithm and implementation collapse into a single level, so there are really only two levels of analysis to consider in the case of neural systems (Maley, 2021).

Despite these and a wide variety of other proposals, the three levels of

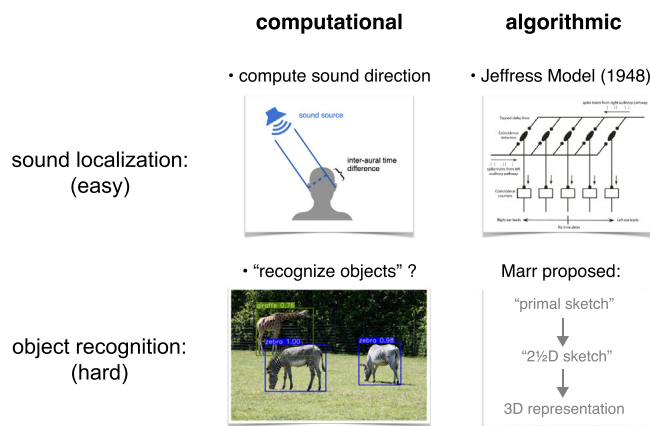


Figure 1. Marr's "computation-first" approach has succeeded with 'easy' but not 'hard' problems

Marr's prescription that computational nature of the problem to be solved should be considered before the algorithm or implementation has had noteworthy successes in 'easy' problem settings, such as estimating the azimuthal angle of a sound source using two ears (top row). Jeffress (1948) proposed an algorithm to solve this problem using delay lines, long before the appearance of any experimental evidence for such circuits (Ashida & Carr, 2011; Carr & Konishi, 1988). However, for 'hard' problems such as object recognition (bottom row), it is difficult to formulate a meaningful computational-level description of the problem to be solved. Moreover, computer vision approaches based on a Marrian 'computational' understanding of the problem (including those proposed by Marr, shown on the right), have been overtaken by approaches relying on computationally intensive learning algorithms applied to large datasets (an approach that I would argue does not fit neatly within the three levels). Figure panels adapted from Cariani (2011) (upper right) and Zhao & Li (2020) (lower left).

analysis proposed in Marr's book have somehow maintained their position at the forefront of the field's awareness. My goal is not to advocate for one of these alternative frameworks, but to point out the arbitrariness and incompleteness of the three particular levels from Marr (1982). Personally, I find the distinction between *descriptive* (seeking to capture statistical structure in an empirical dataset) and *normative* (seeking to achieve optimality according to some theoretical objective) modelling approaches to be a more useful taxonomy for computational neuroscience (Dayan & Abbott, 2001).

Conclusion

The computational problems solved by our brains are in general so difficult that, in contrast to Marr's prognosis, it often makes sense to look carefully at data before attempting to formulate a computational theory. Moreover, the brain is not a general-purpose computer, but a specialized computer that evolved to solve particular problems using a specific organization and architecture. Rigid separation between computations and algorithms, or between algorithms and implementations, which make clear sense in digital computers, has less to offer for understanding biological computations. Finally, the failure of Marr's own approach to vision, inspired by the computer vision ideas of his day and laid out in the remaining chapters of his book (See Table 1-1, Marr, 1982) give further reason to approach the Three Levels with skepticism. Therefore, let us continue to celebrate the many scientific contributions of David Marr, and to seek computational understanding of the nervous system, but let us move away from the Three Levels of Analysis as the default framework doing so.

References

- Ashida, G. O., & Carr, C. E. (2011). Sound localization: Jeffress and beyond. *Current Opinion In Neurobiology*, *21*(5), 745–751.
- Cariani, P. (2011). Jeffress model. *Scholarpedia*, *6*(7), 2920.
- Carr, C. E., & Konishi, M. (1988). Axonal delay lines for time measurement in the owl's brain stem. *Proceedings of the National Academy of Sciences*, *85*(21), 8311–8315.
- Dayan, P., & Abbott, L. (2001). *Theoretical Neuroscience*. MIT Press.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, *7*(2), 217–229.
- Jeffress, L. A. (1948). A place theory of sound localization. *Journal of Comparative and Physiological Psychology*, *41*(1), 35.
- Maley, C. J. (2021). The physicality of representation. *Synthese*, 1–26.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information, henry holt and co personal*, volume 2. MIT Press.
- Marr, D., & Poggio, T. (1976). *From understanding computation to understanding neural circuitry*. Technical report, MIT Artificial Intelligence Laboratory.
- Poggio, T. (2012). The levels of understanding framework, revised. *Perception*, *41*(9), 1017–1023.
- Ritter, S., Barrett, D. G. T., Santoro, A., & Botvinick, M. M. (2017). Cognitive psychology for deep neural networks: A shape bias case study. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, volume 70 of Proceedings of Machine Learning Research, pages 2940–2949*. PMLR.
- Seung, H., & Movshon, J. (2012). *Connectomics: Sebastian Seung vs. Tony Movshon*. <https://www.youtube.com/watch?v=q4KrhDZQ088>. Debate at Columbia University.
- Zhao, L., & Li, S. (2020). Object detection algorithm based on improved yolov3. *Electronics*, *9*(3), 537.

Additional information

Competing interests

The author declares no conflict of interest.

Author contributions

Sole author.

Funding

Simons Collaboration for the Global Brain: Jonathan Pillow, SCGB AWD543027; McKnight Foundation (McKnight): Jonathan Pillow, McKnight Scholar's Award; NIH BRAIN Initiative: Jonathan Pillow, 9R01DA056404-04.

Acknowledgements

The author would like to thank the University of Cambridge for hosting the debate, and to Elisa Galliano, Mátè Lengyel and the other members of the Programme Committee ('David Marr – 50 years on') for organizing the debate and inviting me to participate. The author would also like to especially thank Mátè in particular for his good-natured, enthusiastic, open minded engagement with the topic. Finally, the author would like to thank the many mentors who inspired him with an early love for Marr's Three Levels, including Eero Simoncelli, Tony Movshon, Peter Latham and Peter Dayan; The author hopes they will not be too disappointed in **his renunciation!**

Keywords

computational neuroscience, David Marr, three levels of analysis

Supporting information

Additional supporting information can be found online in the Supporting Information section at the end of the HTML view of the article. Supporting information files available:

Peer Review History