VISION

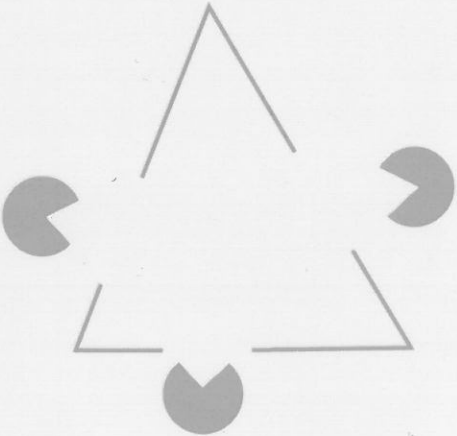

David Marr

FOREWORD BY
Shimon Ullman
AFTERWORD BY
Tomaso Poggio

Afterword

Marr's Vision and Computational Neuroscience

Tomaso Poggio

THIRTY YEARS AGO: A VISION OF COMPUTATIONAL NEUROSCIENCE

The link between computation and neuroscience—the realization that the brain is a computer—is old. Turing wrote about it. McCulloch, Pitts, and Lettvin followed the idea from the perspective of both computation and of neuroscience. Seeds of the idea can be found in centuries-old writings. Though it may not be true that this book started the field known as computational neuroscience, it is certainly true that it had a key role in its beginning and rapid growth. A few years ago at Cosyne, the main conference for computational neuroscience, I mentioned David Marr's work in my keynote talk. In the days afterward, a surprising number of well-known researchers came to me to recount how they entered the field after reading Marr's book and thought that their career was indeed due to *Vision*!

THIRTY YEARS LATER

Thirty years later we still do not understand the brain. Of course, this is not surprising. The problem of intelligence—of how intelligence is created by the brain and of how to make intelligent machines—is one of the greatest problems in science, possibly the most fundamental of all. In 1976 when I was working with Marr for a three-month period at MIT, we fully realized that a satisfactory understanding was far away because the problem was so deep and so difficult. At the time, we had hopes, however, that computational insights could help decrypt puzzles in neuroscience—in particular, in the neuroscience of the visual system. An example is the work we did with Francis Crick at the Salk Institute in La Jolla in 1979 (the discussions among the three of us plus Leslie Orgel eventually became the third part of *Vision*) trying to link the number and properties of cells in layer $4c\beta$ of primary cortex (V1) to computations based on the sampling theorem. Today, it is fair to say that, as a scientific community, computational neuroscience still disagrees about which computations are performed by the visual cortex. In a similar way, one could argue that system physiology has made little progress and, in fact, was not rescued at all by computational neuroscience, as we—and many others—had hoped. What has happened since the publication of *Vision*? What will happen next?

THE "LEVELS OF UNDERSTANDING" MANIFESTO

In trying to provide some answers, I start from one of the most enduring frameworks in Marr's *Vision*, which has been often cited and reformulated in many ways. The simple observation is that a complex system—like a computer and like the brain—should be understood at several different levels. For the purpose of this brief finale, let me list just three levels: the hardware, the algorithms, and the computations. In *Vision*, Marr emphasizes that explanations at different levels are largely independent of each other: a software engineer does not need to know the hardware in any great detail. The message was important at the time: the study of the problems to be solved—and of the associated computations—is relevant in its own right and is needed for a full understanding of the brain. I argue, however, that it is now time to reemphasize the connections between levels, if we want to make progress in computational neuroscience.

To explain let me recount the background of the argument. The section in *Vision* about levels of understanding is directly based on a paper (Marr and Poggio 1977) we wrote together for a booklet of NRP (the influential Neuroscience Research Program founded at MIT by Frank Schmitt) to which I had been invited to contribute. That paper was the original "manifesto" of our computational approach to the brain. Its content was a summary of a few long discussions Marr and I had in the spring of 1976 about levels of analysis of a complex system. We started with an argument described in a long paper (Reichardt and Poggio 1976) on the visual system of the fly. In this paper, Reichardt and I distinguished the three levels of single cells and circuits, of algorithms, and of behavior (of the organism). Marr insisted, correctly, on replacing the behavior level with the level of computation and of computational analysis. This was important for defining the approach of computational neuroscience. One key aspect of the original argument in Reichardt and Poggio 1976, however, almost disappeared in the process. In that paper we stressed that one ought to study the brain at different levels of organization, from the behavior of a whole animal to the signal flow—that is, the algorithms—to circuits and single cells. In particular, we expressed our belief—and Reichardt had written about it even earlier—that (1) insights gained on higher levels help ask the right questions and do the right experiments at lower levels, and (2) it is necessary to study nervous systems at all levels simultaneously. From this perspective, the importance of coupling experimental and theoretical work in the neurosciences follows directly: without close interaction with experiments, theory is very likely to be sterile.

I believe that computational neuroscience over the past thirty years can be described—of course to a first approximation—as mostly exploring each level of understanding independent of the others. To illustrate the point, let me sketch some of the past research trends in computational neuroscience.

RECENT TRENDS IN COMPUTATIONAL NEUROSCIENCE

Much interesting work has been done at different levels. The most basic is probably at the level of biophysics and elementary circuits. A good example is provided by present models of realistic cortical networks, yielding oscillations in the gamma band and analyzing circuits and channel dynamics capable of generating them. Another example is the experimental and theoretical work on circuits with balanced excitation and inhibition and their properties in terms of transmitting and gating information. The analysis of their properties is developing in parallel with the experimental characterization of balanced excitation and inhibition in many brain areas. Another instance of computational work at a somewhat higher level is the analysis, in simulated networks, of operations such as normalization to explain properties of shape recognition, motion estimation, and attention.

It is thus not surprising to find quite a few exemplars of this type of work in conferences such as Computational and Systems Neuroscience, or Cosyne. At the same time, one also finds at Cosyne papers at the computational level, where neurons are not mentioned at all. An obvious recent trend is the emphasis on Bayesian inference as a framework to model the brain: graphical models and hierarchical Bayesian models have been presented as appropriate languages to describe the computations and the algorithms used by nervous systems. An obvious feature of this trend today is that the connection with neurons is missing (though some interesting efforts have been made).

Overall, this body of work at the lowest and the highest levels is what Marr prescribed and provides good foundations for the field. Naturally, Marr's message is sometimes lost. For instance, an explanation of the biophysics of oscillations in the neural activity of cortical areas appears to be regarded in several papers as a full explanation in itself, whereas, in the spirit of computational neuroscience, one must also eventually understand what is the computational role of oscillations and what is the algorithm that controls them. In other words, oscillations may be a symptom or the mechanism of attention, but which computation is actually performed by oscillations?

The level of understanding philosophy also suggests that attempts to understand the brain exclusively within a bottom-up approach are unlikely to succeed. The Blue Brain project—which of course is worthwhile for a number of reasons—could be misinterpreted from this point of view as an attempt to reconstruct every detail of a cortical column along with the belief that one would then be able to infer the computations performed by the cortex. At the other end of the spectrum, Bayesian explanations of psychophysical data—though intriguing and successful—cannot be accepted yet as a complete explanation since in most cases the connection with the underlying neural circuits is missing.

HOW TO UNDERSTAND INTELLIGENCE

This brief look at computational neuroscience in the last decades suggests that, though the problem is far from being solved, significant progress has been made at each of the levels of understanding—in a sense following Marr's prescription. After thirty years it is now time to go beyond it and shift gears. I feel that Marr would also think it is time to look again at the levels of understanding framework—now *emphasizing the connections between levels and their synergies*. In particular, I believe that neuroscience can help computational theory and even computer science as suggested by recent models of visual cortex, which are leading to interesting approaches in computer vision. In 1989, when Marr wrote *Vision*, our belief was that computational theories would help neuroscientists. The rise of computational neuroscience during the last several years has shown that to some extent this has indeed happened. Importantly, the table is now turning: in the near future, neuroscience may well be providing new ideas and approaches to artificial intelligence.

Emphasizing the connections between levels is also a recognition that the problem of explaining the brain is very difficult and that we need to use every bit of information, every approach, every technique we have. It is important also to recognize, as I mentioned, that the emphasis on coupling the different levels de facto implies an emphasis on a very close interaction between experiments and models, as a necessary condition for fruitful future work in computational neuroscience.

Finally, let me comment on the problem of learning, which is an intriguing and interesting omission in Marr's *Vision* quest to understand intelligence and the brain, especially because learning was the focus of his famous papers (1969, 1970) on the cerebellum and the neocortex. I am sure that this omission would have been corrected had Marr had the time. Of course it is important to understand the computations and the representations used by the brain—this is the main objective of the book—but it is also important to understand how an individual organism, and in fact a whole species, *learns* and develops them from experience of the natural world. One could even argue that a description of the learning algorithms and their a priori assumptions is deeper and more useful than a description of the details of *what* is actually learned. I have been arguing for the last two decades that the problem of learning is at the core of the problem of intelligence and of understanding the brain. Learning, I think, should have been included explicitly in Turing's operational definition of intelligence—his famous Turing test. Not surprisingly, the language of modern statistical learning, including regularization, SVMs, graphical models, hierarchical Bayesian models, is permeating various areas of computer science and is also a key component of today's computational neuroscience. I am not sure that Marr would agree, but I am tempted to add learning as the very top level of understanding, above the computational level. We need to understand not only what are the goals and the constraints of a computation are but also how a child could learn it and what the role of nature and nurture is in its development. Only then may we be able to build intelligent machines that could learn to see—and think—without the need to be programmed to do it.