

Neural Machinery and Realization

Thomas W. Polger^{†‡}

The view that the relationship between minds and brains can be thought of on the model of software and hardware is pervasive. The most common versions of the view, known as functionalism in philosophy of mind, hold that minds are realized by brains. The question arises, What is the realization relation? I approach the question of realization through a case study: David Marr's (1982) computational account of early visual processing. Marr's work is instructive because it is the textbook case of the hierarchy of mechanisms that has seemed to bear out the arguments of functionalist philosophers and cognitive scientists. I argue that realization as employed by Marr has some but not all of the characteristics that it is usually taken to have.

1. Introduction. Alan Turing begins his groundbreaking 1950 paper, "Computing Machinery and Intelligence," by asking the question, "Can machines think?" Finding this question ambiguous, Turing proposes to substitute a question that he supposes to be "relatively unambiguous" (1950, 433). Through iteration, the question that Turing poses and revises has become transformed into the question of whether (or how) brains or machines realize mental states. Yet this question, like the original, demands clarification: We have no good explanation of the realization relation itself, and there have been few attempts to provide one. At the same time, the relevance of the realization relation has expanded beyond questions about minds and machines, giving the project of understanding realization some urgency. According to a widely held view, not only psychological states and properties but all those of the so-called special sciences are realized by the states and properties of basic (or more basic) sciences. Only recently has realization at last surfaced as a topic in its

[†]To contact the author, please write to: Department of Philosophy, University of Cincinnati, Cincinnati, OH, 45221-0374; e-mail: thomas.polger@uc.edu.

[‡]For valuable comments and discussion, I am indebted to the participants and audience of the PSA 2002 symposium "Realization and Explanation in Neuroscience," especially Carl Craver, Rob Wilson, Barbara Von Eckhardt, Jeffrey Poland, Peter Machamer, and Stuart Glennan. I am also grateful for many conversations about realization with Carl Gillett and Larry Shapiro.

Philosophy of Science, 71 (December 2004) pp. 997–1006. 0031-8248/2004/7105-0031\$10.00
Copyright 2004 by the Philosophy of Science Association. All rights reserved.

TABLE 1. MARR'S THREE LEVELS OF EXPLANATION FOR COGNITIVE CAPACITIES
(Marr 1982, 24).

Computational Theory	Representation and Algorithm	Hardware Implementation
What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out?	How can this computational theory be implemented? In particular, what is the representation for input and output, and what is the algorithm for the transformation?	How can the representation and algorithm be realized physically?

own right and a lively debate is erupting.¹ But deciding which account of realization to favor, and which dependencies are to be explained by the realization relation, is a matter of some dispute. Following Turing's tactic, I propose to set aside the question, "What is realization?", and substitute in its place a question that I suppose to be somewhat more manageable: What realization relation is invoked by David Marr's (1982) computational theory of vision?

2. Marr's Three Levels. I do not suppose that we should adopt Marr's theory of vision.² Nor do I argue that the theory provides the only or the best account of realization. Marr's theory is a useful starting point not because it is correct but because of its wide influence.

There are two primary reasons for turning to Marr to help us understand realization. One is the general impact that Marr's three-level analysis of explanation of vision has had on the cognitive sciences. According to Marr, explanation in the cognitive sciences has three hierarchically organized components: the computational theory, the representation or algorithm level explanation, and the hardware implementation level explanation (Table 1). Marr (1982, 22–24) explains these levels using the example of a cash register. The computational theory of the cash register tells what it does and why. The cash register performs arithmetic operations in order to combine the prices of various products and arrive at a total invoice. The representation and algorithm explanation tells us how, in the abstract, to satisfy the goals of the computation. Marr takes it that typical cash registers represent prices using Arabic numerals and perform arithmetic operations using elementary rules for adding decimal digits. Finally, the hardware implementation explanation tells how the algorithm is in fact carried out by some particular class of devices. Elementary adding is achieved by means of wires and transistors in the cash register,

1. See Heil 1992, Poland 1994, Horgan and Tienson 1996, Kim 1998, Shoemaker 2001, Wilson 2001, Gillett 2002, Polger 2004, Shapiro 2004, and Endicott forthcoming.

2. For one interesting alternative, see Purves and Lotto 2003.

but a “child who methodically adds two numbers from right to left, carrying a digit when necessary, may be using the same algorithm” (1982, 24).

The multilevel analysis has been important in at least two ways. First, a cognitive process can be explained at any of the levels; in particular, there may be a computational or algorithmic explanation of a cognitive phenomenon that does not mention or depend on any hardware implementation explanation. However, second, Marr himself seems to regard explanation at any one level as incomplete unless it is incorporated into a multilevel account. This has not stopped philosophers and cognitive scientists from regarding the levels as providing “autonomous” explanations for psychological phenomena—explanations that are each complete in their own right. Whether or not the levels are considered complete and autonomous, Marr’s account focuses our attention on the relationships between the levels of explanation and between the entities that they invoke. That relation is realization:

In order that a [computational] process shall actually run, however, one has to *realize* it in some way and therefore choose a representation for the entities that the process manipulates The choice, then, may depend on the type of hardware or machinery in which the algorithm is to be *embodied* physically.

This brings us to the third level, that of the device in which the process is to be *realized* physically. (Marr 1982, 23–24; emphasis added)

The importance of the claim that the relationship among the levels is realization should not be underestimated. Putnam introduced realization as the relationship that a machine has to a Turing machine program.³ If Marr is correct then this relation is fundamental to explanation in the cognitive sciences.

The second reason for turning our attention to Marr’s theory is the specific influence that his explanation of early visual processing has had for philosophical acceptance of realization-based approaches to cognition. Philosophical uses of the realization relation are sometimes directly derivative of Marr’s use. For example, Terence Horgan and John Tienson write,

The relationship between state types at Marr’s middle (algorithm) level and state types at the lowest (implementation) level, and *also*

3. Kim 1998, among others, credits Putnam with first using ‘realization’ in this way. See Lycan 1974 and Polger 2004 concerning the ambiguities in Putnam’s way of talking about Turing machines.

the relationship between state types at Marr's top (cognitive-transition [computational]) level and state types at his middle level, is the relation that philosophers call *realization* and cognitive scientists call *implementation*. (1996, 23).

Horgan and Tienson do not explore the nature of the realization relation itself, but they take Marr's account to exhibit some of its fundamental characteristics, e.g., that realization is a transitive relation, and that the realization relation is a one-many relation so that the computational and algorithmic levels do not have unique realizers (1996, 23).

Marr's theory is a good starting point because it has been so influential and because the relation of realization is basic to his account. If we can understand how realization operates in Marr's theory then we can understand how many philosophers and cognitive scientists employ the realization relation. Of course we might still have reasons for seeking an alternative explanation of realization; but at least we will be in a better position to understand the work that the relation is supposed to do.

3. Marr's Computational Theory of Vision. According to Marr, visual perception begins by taking as input the stimulation of retinal photoreceptors and deriving a representation of the three-dimensional scene before the perceiver. This process has three stages. First, the retinal response is processed to produce a primal sketch that makes explicit the two-dimensional information in the scene. Derivation of the primal sketch begins by locating the luminance boundaries ("edges") in the image projected onto the retina, forming the raw primal sketch. Then the luminance boundaries are grouped into objects and shapes, yielding the full primal sketch. Next, the full primal sketch is enriched to include information about the depth and orientation of contours in a viewer-centered framework. And, finally, a three-dimensional model is constructed to represent the objects and their spatial organization in an object-centered framework.

Let us examine just the very first stage of the process: production of the raw primal sketch. To explain the generation of the raw primal sketch we must provide computational, algorithmic, and hardware implementation explanations. The computational theory of the stage that generates the raw primal sketch is that it is the process of edge detection. The task is to locate the edges (luminance boundaries) in the image projected on the retina. The questions, then, are how that process is achieved (realized) at both the representation/algorithm and the hardware implementation levels of explanation. When we see how Marr answers those questions, we can determine the relation or relations of realization implicit in his theory.

Marr (1982, following Marr and Hildreth 1980) proposes that the al-




Figure 1. Marr uses response profiles to illustrate the idea of a zero-crossing, and to argue for the optimality of a second derivative filter: “intensity change (a) gives rise to a peak (b) in the first derivative and to a (steep) zero-crossing (c) in its second derivative” (Marr 1982, 54).

gorithm for accomplishing edge detection is the Laplacian of the Gaussian ($\nabla^2 G$).⁴ Because of the inevitable noise in the retinal representation, Marr proposes to apply a Gaussian (G) filter to the image; the purpose of this procedure is to wash out small noise in the image in favor of more coarse grained trends. Boundaries in the image (edges) will be located at the gradients from/to brighter to/from darker regions. (The two-dimensional luminance profile of such a boundary is depicted by Figure 1a.) The question then becomes how to find these luminance gradients. One suggestion is to apply a first order derivative operator to the Gaussian filtered image, which will yield peaks in the representation at the locations of the gradients (Figure 1b). But it is difficult to find local maxima and minima in a noisy image, even after the Gaussian filter. So Marr opts for a second derivative operator, which will produce a so-called zero-crossing in the representation at the location of the luminance boundary (Figure 1c). Since the algorithm is supposed to locate edges regardless of their orientation, both the Gaussian filter and second derivative operator should be orientation independent. Marr argues that the Laplacian (∇^2) should be used because it is the simplest second derivative isotropic operator. So the algorithm-level explanation for the process of generating the raw primal sketch is that we must find zero-crossings in the response profile of the $\nabla^2 G$ filter applied to the retinal image.

The final question is how the process of finding zero-crossings in $\nabla^2 G$ can be implemented by the visual system. This process has two stages: calculating $\nabla^2 G$, and then locating the zero-crossings in the resulting representation. According to Marr, the calculation of $\nabla^2 G$ is accomplished by certain retinal ganglion cells and so-called X -cells in the lateral geniculate. Detection of the zero-crossings is accomplished by “simple cells”

4. The complete argument for using $\nabla^2 G$ is advanced in Marr and Hildreth (1980).

in the cortex, of the sort identified by Hubel and Wiesel (1962). So the ganglion and X -cells realize the $\nabla^2 G$ filter, and the simple cells realize the zero-crossing detector. Together, these cells realize a system for producing the raw primal sketch.

If this is correct then we are in a position to ask: What is the realization relation appealed to by Marr's analysis? How does Marr argue that particular neural structures realize the particular representations and algorithms that are identified as optimal?

4. Realization for Marr. Begin with the retinal ganglion and X -cells. These cells have center-surround receptive fields. That is, they are stimulated by retinal cells which have a spatial organization of concentric circles, such that (i) when the retinal cells in the center of the receptive field are stimulated they stimulate the ganglion cell, and (ii) when the retinal cells in the surround portion of the receptive field are stimulated they inhibit the ganglion cell. This describes a so-called "on-center" cell; in an "off-center" cell, the spatial relation of the excitatory and inhibitory retinal cells is inside out. The first reason for thinking that ganglion and X -cells realize $\nabla^2 G$ is that their center-surround physiology has the same spatial organization as the Laplacian filter. Second, like the Laplacian, the center-surround cells are orientation insensitive, thus isotropic. Third, the cells are insensitive to the exact location of a stimulus within the subportions of the receptive field, that is, they only distinguish between center and surround but not exact locations in center and surround. As such, they "blur" the exact location of the stimuli as the Gaussian filter does, and the spatial characteristics of the receptive field govern the amount of blur; moreover, Marr notes that the sizes of the receptive fields approximate the spatial characteristics of the ideal Gaussian filter for edge detection (1982; Marr and Hildreth 1980). Indeed, Marr reports, the center-surround cells are known to have a response profile that can be described as the difference of two Gaussians (DOG), which approximates the $\nabla^2 G$ function.⁵ These observations can be summarized by saying that the re-

5. This simplified story is misleading in certain ways, but the complication does not assist our present discussion. It is worth noting that in fact Marr does not propose that a single center-surround cell realizes the $\nabla^2 G$ function, but that it is realized by two such cells (one on-center and one off-center) with adjacent receptive fields—one contributing the positive portion of the signal and one the negative. That means the cells do not have the same spatial organization as the Laplacian filter. (Instead they have the spatial configuration of two overlapping Laplacian filters, as illustrated in Figure 2.) The reasoning that Marr offers (1982, 64–65, especially figure 2-17) involves observing that if we bisect the $\nabla^2 G$ curve and look at each part separately, we notice that each is approximated by the activity of a center-surround cell. But even if we suppose that this curve-fitting argument can be defended, there is a problem: It is clear

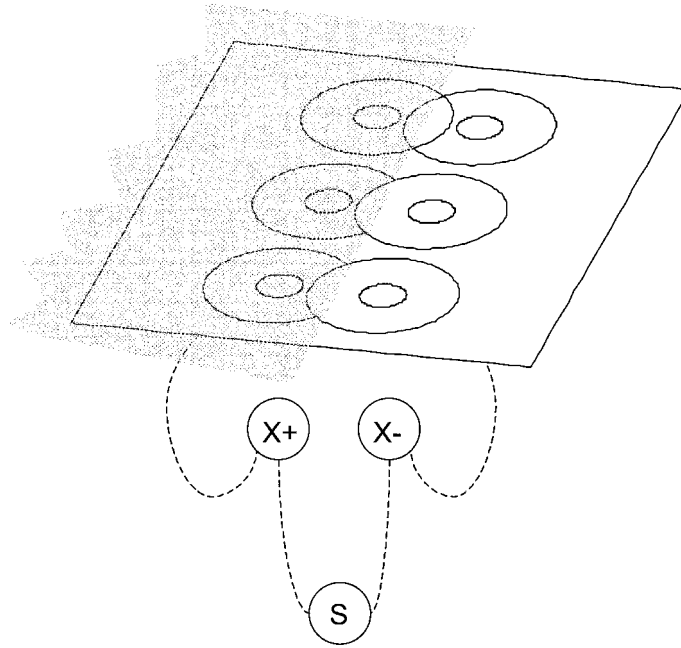


Figure 2. On Marr's model, the ganglion or X -cells with on-center or off-center surround receptive fields ($X+$, $X-$), stimulate the simple cells (S) in visual cortex. The X -cells realize $\nabla^2 G$, and the simple cells act as zero-crossing detectors on the resulting signal.

sponse profile of retinal ganglion and X -cells to stimuli that fall in their receptive fields matches the response profile of the $\nabla^2 G$ operator and thereby realizes that operator. Furthermore, the ganglion and X -cells stimulate and inhibit so-called simple cells in the visual cortex, which are known to be sensitive to oriented lines that fall in the receptive fields of their upstream ganglion and X -cells (Figure 2). In short, the ganglion and X -cells have the response profile of $\nabla^2 G$, and the simple cells have the response profile of a zero-crossing detector.

Realization for Marr, then, is having the same response profile, relative to some specification of input-output mappings.⁶ This understanding of

that the resulting signal has only a positive value and therefore that no manipulation on it amount to identifying a zero crossing—for it has none. (See Section 5, below.)

6. Here I am ignoring questions about how to determine those mappings, but that is clearly an important issue in its own right.

Marr matches the notion of realization that is used by many philosophers. For example, Robert Cummins writes:

What an adding machine does is *instantiate* the plus function. It *instantiates* addition by *satisfying* the function g whose arguments and values represent arguments and values of the addition function, or in other words, have those arguments and values as interpretations. (1989, 89)

According to Cummins a physical system realizes an abstract operation by having inputs and outputs that can be mapped onto the inputs and outputs of the abstract function. And this is exactly the basis for Marr's defense of his account of the realization of edge detection. Realization is some sort of mapping relation of inputs and outputs.

This is a boring conclusion, but it has at least one interesting consequence: Realization of the computational level need not proceed via any intermediate representation or algorithmic level, for the hardware can directly realize the computational level. That is, the inputs and outputs of the system composed of ganglion cells, X-cells, simple cells, and their connections realizes the function of edge detection regardless of whether it also realizes zero-crossing detection on $\nabla^2 G$, and it does so directly rather than indirectly by realizing the representational or algorithmic level. Contrary to Horgan and Tienson, Marr's analysis does not require that realization be transitive. In fact, since the mapping functions from hardware to representation and from hardware to computation are almost certainly different—specifically, because the higher levels each have unique outputs to be mapped, one a response profile and the other a representation of edges—it is difficult to see how the relation could be transitive. Since realization is widely thought to be transitive, this is an interesting result.

5. Objections to the Input-Output Interpretation of Marr. The boring conclusion is that realization for Marr is simply input-output mapping relative to some specification of inputs and outputs. The more interesting conclusion is that this does not support the claim that realization must be transitive. Each of these conclusions may seem hasty. There are two related objections. First, one might argue that by looking independently at realization of the representation/algorithm level and of the computational level I have neglected the importance that Marr places on the integration of the levels into a single explanation. Thus, the visual system does not realize an edge detector merely because it has inputs and outputs that an edge detector would have, a claim that borders on triviality. Rather, the system has an internal organization that realizes the algorithmic mechanism that Marr hypothesizes; it is only said to realize that computational

function because it has internal states that realize the algorithm. And, one might continue, there are cells that calculate $\nabla^2 G$, and those cells are wired to “zero-crossing detectors” that respond to zero-crossings in $\nabla^2 G$. So Marr’s notion of realization is more nuanced than I have admitted. Second, one might worry that my conclusion that realization for Marr is not transitive is simply a symptom of my neglect of the complexity of his account. The middle level of representation/algorithm is crucial to Marr. After all, one of his most important insights is that $\nabla^2 G$ is the optimal function for edge detection. An account that discounts the importance of that function to realizing edge detection has obviously not appreciated Marr’s theory.

Together these objections amount to the claim that Marr’s edge detection computation is realized in the visual system via the realization of the representation/algorithm Marr specifies, and that realization is therefore both complex and transitive after all. We must therefore ask: Does the visual system in fact detect edges by locating zero-crossings in the Laplacian of the Gaussian function?

Do some cells realize the $\nabla^2 G$ function? Yes. The receptive field structure and response profiles of some retinal ganglion and X-cells can be said to realize $\nabla^2 G$. Specifically and as reported by Marr, their response characteristics fit the DOG function which approximates $\nabla^2 G$. However, whether we think of these cells as realizing DOG or $\nabla^2 G$, it is clear that the realization relation is simple input-output mapping.

Perhaps more importantly, do some cells realize zero-crossing detectors on the $\nabla^2 G$ function? No. No zero values in $\nabla^2 G$ are identified. In fact, Marr himself is candid about this point:

From a physiological point of view, zero-crossing segments are easy to detect without relying on the detection of zero values, which would be a physiologically implausible idea. The reason is that just to one side of the zero-crossing will lie a peak positive value of the filtered image $\nabla^2 G * I$, and just to the other side, a peak negative value. (1982, 64).

The visual system, in other words, may be said to detect the locations at which zero-crossings would be located, but it does not do so by locating zero-crossings. Yet to say that the system identifies the locations where zero-crossings would occur is just to say that it identifies the locations of edges. Simple cells may be edge detectors, in which case their realization relation is a simple input-output mapping. But simple cells do not locate zero-crossings in $\nabla^2 G$, so they do not realize Marr’s representation/algorithm level.

Once more we have one boring and one interesting conclusion. The interesting conclusion is that the visual system does not realize the zero-

crossing edge detection algorithm. In fact, the mechanism that Marr describes is more like the first-order peak and trough detector that Marr and Hildreth dismiss (1980) than it is like the second derivative $\nabla^2 G$ operator. The boring conclusion is that center-surround and simple cells probably have something to do with edge detection, and so may (partially) realize (something like) Marr's computational level theory. But that amounts to little more than saying that they play a part in detecting edges.

6. Conclusion. Marr's realization relation is a simple input-output mapping relation. The input-output conception of realization has virtues, but it is not suited to all purposes.⁷ In particular, it has proven unhelpful if we are interested in questions about the causal properties of realized states and processes, e.g., mental causation. Recognition of the limits of the input-output mapping account of realization has led some philosophers to take a new look at the relation. Now we see that those investigations will not benefit from looking back to Marr.

REFERENCES

- Cummins, Robert (1989), *Meaning and Mental Representation*. Cambridge, MA: MIT Press.
- Endicott, Ronald (forthcoming), *The Book of Realization: Analyzing a Cluster of Concepts for the Philosophy of Mind, Solving a Computational Conundrum for Cognitive Science*.
- Gillett, Carl (2002), "The Dimensions of Realization: A Critique of the Standard View", *Analysis* 64(4): 316–323.
- Heil, John (1992), *The Nature of True Minds*. New York: Cambridge University Press.
- Horgan, Terence, and John Tienson (1996), *Connectionism and the Philosophy of Psychology*. Cambridge, MA: MIT Press.
- Hubel, David, and Torsten Wiesel (1962), "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex", *Journal of Physiology (London)* 160: 106–154.
- Kim, Jaegwon (1998), *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, MA: MIT Press.
- Lycan, William (1974), "Mental States and Putnam's Functionalist Hypothesis", *Australasian Journal of Philosophy* 52: 48–62.
- Marr, David (1982), *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman.
- Marr, David, and Ellen Hildreth (1980), "Theory of Edge Detection", *Proceedings of the Royal Society of London B* 207: 187–217.
- Poland, Jeffrey (1994), *Physicalism: The Philosophical Foundations*. New York: Oxford University Press.
- Polger, Thomas (2004), *Natural Minds*. Cambridge, MA: MIT Press.
- Purves, Dale, and R. Beau Lotto (2003), *Why We See What We Do: An Empirical Theory of Vision*. Sunderland, MA: Sinauer.
- Shapiro, Lawrence (2004), *The Mind Incarnate*. Cambridge, MA: MIT Press.
- Shoemaker, Sydney (2001), "Realization and Mental Causation", in Carl Gillett and Barry Loewer (eds.), *Physicalism and Its Discontents*. Cambridge: Cambridge University Press.
- Turing, Alan (1950), "Computing Machinery and Intelligence", *Mind* 59(236): 433–460.
- Wilson, Robert (2001), "Two Views of Realization", *Philosophical Studies* 104: 1–30.

7. See my 2004.