Marr on Computational-Level Theories*

Oron Shagrir†‡

According to Marr, a computational-level theory consists of two elements, the what q1 and the why. This article highlights the distinct role of the Why element in the computational analysis of visual tasks. Three theses are advanced: (*a*) that the Why element plays a major explanatory role in computational-level theories, (*b*) that its goal is to explain why the computed function is appropriate for a given visual task, and (*c*) that the explanation consists in showing that the "inner" functional relations between the representing cells are similar to the "external" mathematical relations between the entities that are being represented. It is concluded that computational theories relate the mathematical input-output function that is being computed (specified by the What element) to the mathematical relations between the entities that the inputs and the q2 outputs represent (specified by the Why element).

1. Introduction. David Marr's *Vision* (1982) advances a computational approach to the study of visual processes. Although published almost 3 decades ago, its impact is still found in artificial intelligence, cognitive q3 science, and neuroscience.[1] Marr's work has also stimulated debates in philosophy about levels of explanation, top-down versus bottom-up methodologies, the nature of computation, externalism versus internalism in computational theories of mind, and the relations between content and computation.

Marr's conception of computational theory is somewhat unusual. In most computational approaches, the goal of a computational theory is to characterize the what and the how. The What element characterizes

*Received September 2009; revised January 2010.

†To contact the author, please write to: Departments of Philosophy and Cognitive Science, The Hebrew University, Jerusalem 91905, Israel; e-mail: shagrir@cc.huji.ac.il.

‡I am thankful to Merav Ahissar, Collin Allen, Carl Craver, Frances Egan, Yosef Grodzinsky, Amit Hagar, Amir Horowitz, Otto Lappi, Arnon Levy, David Kaplan, Gualtiero Piccinini, Itamar Pitowsky, Nick Shea, Anna-Mari Rusanen, and Shimon Ullman. This research was supported by the Israel Science Foundation grant 725/08.

1. According to the ISI Web of Knowledge index, *Vision* was cited over 150 times in 2008 alone.

Philosophy of Science, 77 (October 2010) pp. 000–000. 0031-8248/2010/7704-0001$10.00

the function that is being computed, and the How element specifies the algorithm by which the system computes this function. But, according to Marr, this How element belongs not to the computational level but to "the algorithmic level." In contrast, the two elements of Marr's computational theory are the what and the why: "The most abstract is the level of *what* the device does and *why*" (1982, 22).

But what exactly does Marr mean by "computational-level theory"? Some interpreters of Marr argue that computational-level theories specify the mathematical input-output function that is being computed (the algorithmic level then specifies the algorithm by which the system computes this function; see Egan 1992, 1995, 2010; Butler 1998). Other interpreters argue that computational theories also refer, in one way or another, to the visual content of these input and output states. This latter interpretation motivates the linkage between Marr's computational level and the "intentional" (Dennett 1994), "semantic" (Pylyshyn 1984), "knowledge" (Newell 1982), or "ecological" (Sterelny 1990) level (see Sterelny 1990, 44–46; Harnish 2002, 400ff.).

In what follows, I advance a different interpretation of Marr's conception of computational-level theories. This interpretation underscores the distinct character of the Why element in Marr's computational theories, an element that is often ignored.[2] According to my understanding, Marr's computational level can be seen as intentional, in the sense that it is concerned with the embedment of the visual system in the environment we happen to live in. But it is not intentional if "intentional" boils down to the specification of visual content. Computational theories, I argue, relate the mathematical function that is being computed (specified by the What element) to the mathematical relations between the things that are being represented (specified by the Why element).

This article proceeds as follows: in section 2, I review some of the well-known statements of Marr. I suggest that one role of the computational level—the role of the What element—is to provide a function-theoretic, mathematical description of what the system does. In section 3, I argue, on the basis of other statements, that there must be another crucial component in computational-level theories, namely, the Why element. In section 4, I clarify the goal of the Why element, which is to explain why the mathematical function that is being computed is appropriate for a given information-processing visual task. In section 5, I argue that this goal is accomplished when it is shown that the "inner" input-output functional relations between the representing cells are similar (in ways specified in the article) to the mathematical relations between the entities that these

2. A notable exception is Kroustallis (2006).

cells represent. In section 6, I show how the Why element combines with the What element to constitute a computational-level explanation. In section 7, I compare my view with other interpretations of Marr.

2. Marr's Computational Theories: The What. In *Vision*, Marr (1982) famously advances a three-level approach to the study of visual processes. The computational level specifies what is being computed and why. The algorithmic level characterizes the system of representations that is being used (e.g., decimal vs. binary) and the algorithm for the transformation from input to output. The implementation level specifies how the representations and algorithm are physically realized.

According to Marr, neuroscientists have been occupied with the working of vision at the implementation level, namely, with the behavior and properties of cells that implement the visual processes and states. This approach, Marr argues, is descriptive and does not explain behavior.[3] The algorithmic level, which was imported to the study of vision from computer science, advances our theoretical understanding of the structure of visual processes. But it too does not capture essential elements of the task as an information-processing problem.

q4 Marr thus concludes that "there must exist an additional level of understanding at which the character of the information-processing tasks carried out during perception are analyzed and understood in a way that is independent of the particular mechanisms and structures that implement them in our heads. This was what was missing—the analysis of the problem as an information processing task" (1982, 19). Marr emphasizes, here and elsewhere, two points about the study of vision. One is that the study of vision is primarily a study of an information-processing task: "Vision is therefore, first and foremost, an information-processing task" (3). The other is the prime explanatory role of the computational level in the analysis of visual, as information-processing, tasks: "Although algorithms and mechanisms are empirically more accessible, it is the top level, the level of computational theory, which is critically important from an information-processing point of view" (27).

What is an information-processing task? Marr characterizes it as "mapping from one kind of information to another" or as "mapping from one
q5 representation to another." He does not provide a detailed account of what is meant by representation or information. Like many researchers of vision, Marr apparently identifies these terms with some sort of selective response to stimuli or with some reliable causal correlation between the

3. Marr writes: "The key observation is that neurophysiology and psychophysics have as their business to *describe* the behavior of cells or of subjects but not to *explain* such behavior" (1982, 15).

Figure 1. Information-processing task that consists in a causal process (*solid line*) from brain state/event B_1 to brain state/event B_2, whereby B_1 represents (*dashed line*) an object/property W_1, in the visual scene, and B_2 represents object/property W_2.

activity of cells and certain types of stimuli. For example, he identifies "the apocryphal grandmother cell" (1982, 15) as "a cell that fires only when grandmother comes into view" (note on 15).

Characterizing an information-processing task somewhat schematically, we can present it as a process from one brain state (B_1) to another (B_2),
q6 whereby B_1 and B_2 are themselves representations. State B_1 represents objects/properties W_1, usually in the visual scene, and B_2 represents ob-
q7 jects/properties W_2 (fig. 1). Examples from the context of vision are shape from shading, depth from disparity, structure from apparent motion, sur-
q8 face orientation from optical flow, and edge detection.

Let us consider edge detection, the task of detecting oriented edges in the visual scene (fig. 2). Variable B_1, here, is the electrical activity of the photoreceptors, known as intensity values. This activity is sensitive to light intensities in the visual field, which consist, mainly, in light reflectance, geometry, illumination of the scene, and the viewpoint (W_1). These intensity values constitute what is known as the retinal image, wherein each "pixel" is sensitive to the light intensity in a certain location in the visual scene. Variable B_2 is the activity of cells in early visual cortex (V1); these are roughly the Hubel and Wiesel's cells that are sensitive to oriented lines. These "edges" are salient features in the physical scene, such as object boundaries (W_2).

What is the job of computational theories in analyzing an information-processing task? One job is to specify the functional relations between

Friday Aug 13 2010 04:14 PM/PHOS10157/2010/77/4
/sji/lfw/ /editing started/use-graphics/narrow/default/

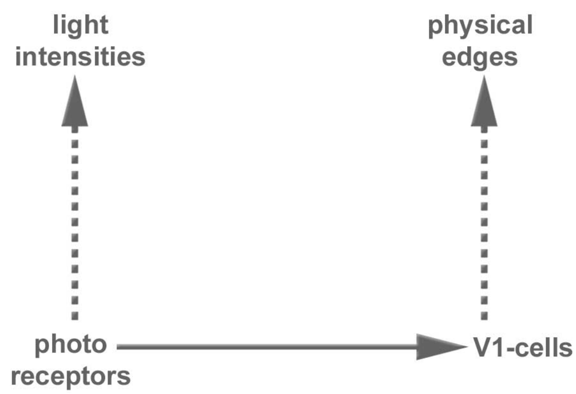

Figure 2. Edge detection as an information-processing task. Intensity values of the photoreceptors represent light intensities in the visual field, which consists, among other things, in light reflectance. Cells in the primary visual cortex (V1) represent oriented edges, such as object boundaries.

abstract properties of B_1 and B_2. Ideally, this mapping relation is defined in terms of mathematical properties abstracted from the electrical activity of the cells: "At one extreme, the top level, is the abstract computational theory of the device, in which the performance of the device is characterized as a mapping from one kind of information to another, the abstract properties of this mapping are defined precisely" (Marr 1982, 24).

In their "Theory of Edge Detection," Marr and Hildreth (1980; see also Marr 1982, chap. 2) describe the early visual task of edge detection in terms of colocated zero crossings of different size filters of the form
q9 $\nabla^2 G \times I(x, y)$ (fig. 3). The term $I(x, y)$ refers to an array of intensity values of the photoreceptors, in which each pixel covers a spatial point (or region) in the visual scene. The term $\nabla^2 G \times I$ describes the activity
q10 of retinal ganglion and LGN cells, which perform a certain filtering of
q11 the intensity values. Here the multiplication symbol is a convolution operator, and $\nabla^2 G$ is a filtering operator: G is a Gaussian that blurs the image, and ∇^2 is the Laplacian operator ($\partial^2/\partial x^2 + \partial^2/\partial y^2$) that is sensitive to sudden intensity changes in the image. The zero crossings of this formula are precisely those places in the image that have sharp intensity changes. This process takes place through several filters with different Gaussian distributions, each producing a different set of zero crossings (fig. 4). Detecting these zero crossings is the task of cells in the primary visual cortex (V1). The colocated zero crossings often signify edges, such as object boundaries, and are the basis for the zero-crossing (edge) segments in the raw primal sketch.

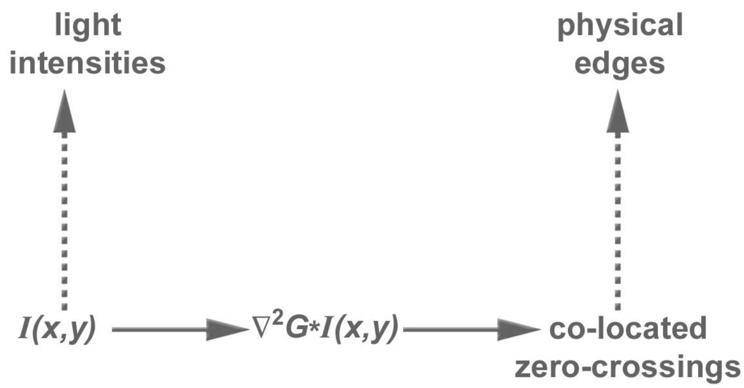

Figure 3. Formal description of edge detection: retinal image, $I(x, y)$, is convoluted through the filtering operator $\nabla^2 G$, where G is a Gaussian and ∇^2 is a second-derivative (Laplacian) operator. Early vision processes include several filters with different Gaussian distributions, and each produces a different set of zero crossings. Intensity values are interpreted as representing light intensities in the visual field, and the colocated zero crossings are interpreted as representing edges, such as object boundaries.

One might think that the job of computational theories stops here, namely, in the formal, function-theoretic specification of the input-output relations. Frances Egan, for one, associates Marr's "computational theory" with "the specification of *the function computed*" (1991, 196–97). She argues that computational theories provide no more than mathematical specifications: "The top level should be understood to provide a function-theoretic characterization," and "the theory of computation is a mathematical characterization of the function(s) computed" (1995, 185).[4] As Egan notes, Marr himself refers, in the epilogue of *Vision*, to the mathematical formula $\nabla^2 G \times I$ as the computational description of what the retina does: "Take the retina. I have argued that from a computational point of view, it signals $\nabla^2 G * I$ (the X channels) and its time derivative $\partial/\partial t(\nabla^2 G * I)$ (the Y channels). From a computational point of view, this is a precise specification of what the retina does" (1982, 337).

I think that Egan captures very well the way Marr characterizes the What element in computational theories. The job of this element is to provide a precise specification of what the system does, and the precise specification of what the retina does is provided by the formula $\nabla^2 G\ \times$

4. Egan's views are more sophisticated than that; I return to discuss her views in the last section.

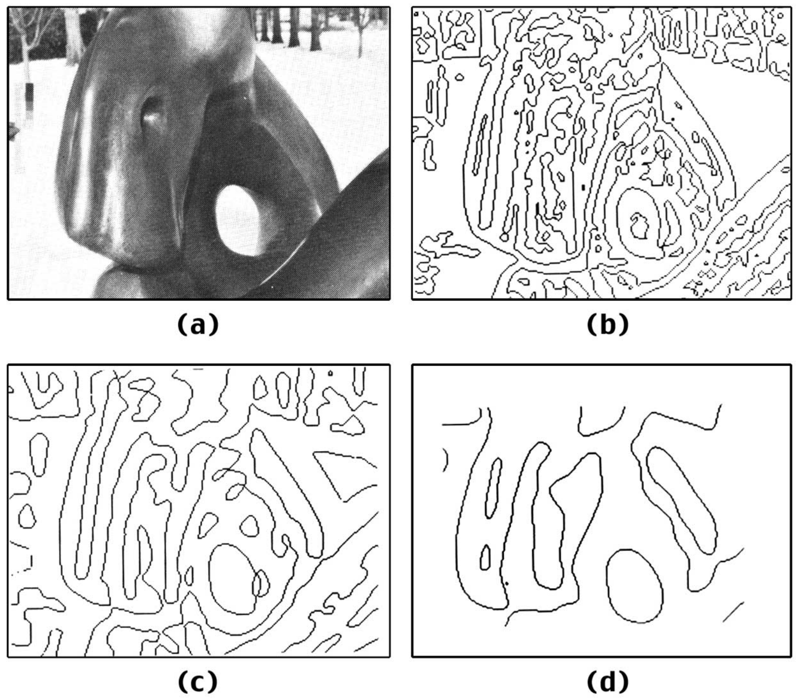

Figure 4. Different sets of zero crossings. Image (*a*) is convoluted with different-sized filters (*b*, *c*) and shows the zero crossings thus obtained (*d*). Many of the fine details obtained through the smaller-sized filter (*b*) are not obtained by the larger-sized filter (*c*), but some of the zero crossings obtained in *c* do not appear in the *b*. From Marr and Hildreth (1980), 201, fig. 6. Reprinted with permission from Royal Society Publishing.

I. However, Egan downplays the fact that there is another component to the computational level, namely, the Why element. Marr refers in the passage to what the retina does but not to why. As we shall see in the next section, Marr stresses that a computational-level theory goes above and beyond the characterization of the what, namely, of the mathematical function that is being computed.

3. Beyond the What: Asking Why. Marr repeatedly associates the computational level with two components, the what and the why. In the introductory chapter of *Vision*, where Marr presents his "philosophical outlook," he states that "the most abstract is the level of *what* the device does and *why*" (1982, 22). The role of the What element is to specify "what is computed," whereas the role of the Why element is to demonstrate the appropriateness and adequacy of what is being computed to the information-processing task (24–25). In "Artificial Intelligence: A Personal View," Marr says that at the computational level, "the underlying nature of a particular computation is characterized, and its basis in the physical world is understood. One can think of this part as an abstract formulation of *what* is being computed and *why*" (1977, 37).

When discussing the cash register example (1982, 22–24), Marr says that what is being computed by the device is addition. But he then goes on to state that this characterization is only one-half of the computational explanation: "The other half of this level of explanation has to do with the question of *why* the cash register performs addition and not, for instance, multiplication when combining the prices of the purchased items to arrive at a final bill" (22). I return to analyze this example below, but it is important to see at this point that Marr refers to the computational level as a level of explanation and that the Why element is an essential part of this computational explanation.

Marr's statements about the why are not merely declarative. The quest for why occurs in his computational theories of vision. Take his theory of edge detection, for example. After characterizing what as being computed in early vision—namely, the colocated zero crossings of $\nabla^2 G \times I$—Marr emphasizes that we still have to confront another critical question: "Up to now I have studiously avoided using the word *edge*, preferring instead to discuss the detection of intensity changes and their representation by using oriented zero-crossing segments. The reason is that the term *edge* has a partly physical meaning—it makes us think of a real physical boundary, for example—and all we have discussed so far are the zero values of a set of roughly band-pass second-derivative filters. We have no right to call these edges, or, if we do have a right, then we must say so and why" (1982, 68).

It is not easy to tell what Marr means by this statement. But it is clear enough that he is concerned with the veridical relation between the visual system and the visual scene, for example, between colocated zero crossings that are the basis of edge segments and physical edges such as object boundaries. This concern is emphasized, perhaps more explicitly, by Marr and Hildreth, who say that "the concept of 'edge' has a partly visual and partly physical meaning. One of our main purposes . . . is to make explicit this dual dependence" (1980, 211).

The specific questions that Marr raises in this respect are (*a*) why are the zero crossings that result from different-sized filters related to the same feature in the physical environment (see fig. 4)? After all, as Marr puts it, "there is no [a] priori reason why the zero-crossings obtained from the different-sized filters are related" (1982, 68). And (*b*) why is this feature often an edge in the physical sense, for example, an object boundary? After all, even if these zero crossings are related to a single feature, there is no a priori reason as to why this feature is a physical edge.

Another, and even more striking, example of the quest for why appears in the theory of stereo vision. As we occasionally experience, there is an angular discrepancy in the position of an object in the two retinal images. This discrepancy is termed disparity. The disparity is usually larger when the object is closer to the eyes (a finger touching your nose) and smaller when it is farther away. The visual system deploys disparity to compute several features, of which the most significant is depth. The first step of this process is matching up elements from the visual scene—that is, finding the two elements, one from the left retinal image and the other from the right retinal image—that correspond to the same object. The difficulty of the task stems, among other things, from the ambiguity of elements in
q12 the images and the multiple possibilities of matching the dots.[5]

q13 An example is illustrated in figure 5. The four projections in the left eye's view (L_1, . . . , L_4) can be paired in various ways with the four projections (R_1, . . . , R_4) in the right eye's view. Of all the pairing possibilities, only the four filled circles are correct, whereas the other 12 open circles are false targets. Still, the visual system solves the correspondence problem, even in highly ambiguous scenes. What explains this remarkable ability?

Marr and Poggio (1976, 1979) argue that the pairing function that provides matching is the one that satisfies two constraints: (*a*) Uniqueness, a black dot from one image can match no more than one black dot from the other image (this constraint rules out, e.g., the function that matches L_1 to R_1 and also L_1 to R_2), and (*b*) Continuity, disparity varies smoothly almost everywhere. This constraint rules out functions that match up pairs with very different disparities. Marr (1982, 111–16) demonstrates that these constraints are efficient; they provide necessary and sufficient conditions for matching in most natural scenes.[6]

So far, we have dealt with the What element of the computational-level theory. This element states very precisely what the visual system computes, namely, a function that satisfies Uniqueness and Continuity (UC) pairing. This characterization, however, does not exhaust the role of the computational level. Marr argues that a computational level also has to explain why computing this function is appropriate for the information-processing task. The computational level has to explain why the UC-pairing func-

5. For recent work on the neural basis of these processes, see Cumming (2002) and Durand, Celebrini, and Trotter (2007).

6. The demonstration involves a third constraint, that of compatibility, which asserts that black dots can match only black dots.

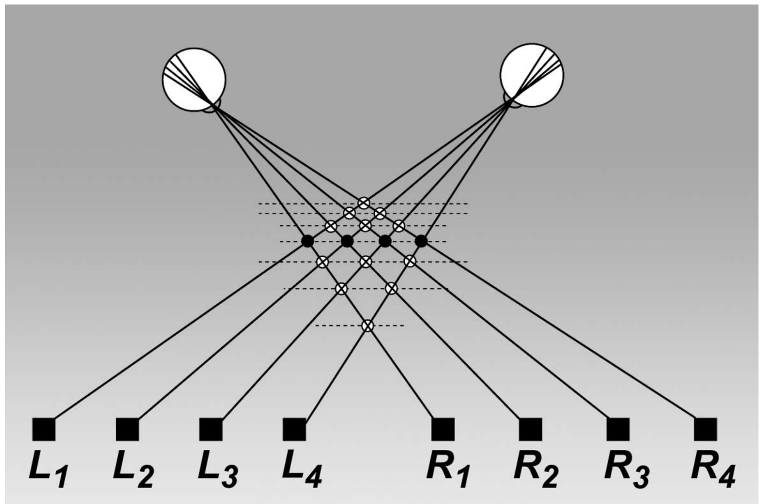

Figure 5. Ambiguity in correspondence. Four projections on the left eye's image (L_1, L_2, L_3, L_4) could match any of the four projections (R_1, R_2, R_3, R_4) on the right eye's image. Of the 16 possible pairings, only four are correct (*filled circles*), while the remaining 12 (*open circles*) are false targets. Dashed lines signify the amount of (horizontal) disparity; circles (pairs) that are on the same line have the same disparity. From Marr and Poggio (1976), 285, fig. 1. Reprinted with permission from American Association for the Advancement of Science.

tion—and not another pairing function—provides matching. As Marr puts it, "The real question to ask is *Why* might something like that work? For the plain fact is that if we look just at the pair of images, there is no reason whatever why L_1 should not match R_3; L_2 match R_1, and even L_3 match R_1" (1982, 112).

Again, it is not immediately apparent what Marr means by this Why question. But I think it is safe to say at this point that Marr thinks that the business of a computational-level theory is not exhausted by the characterization of the mathematical mapping function. Computational theories should do more than that. They should also explain the basis of this mathematical function in the physical world. In what follows, I attempt to explicate what Marr means by a computational-level theory. Specifically, I focus on the following issues: (*a*) What exactly is the problem that the Why element aims to address (sec. 4)? (*b*) How does the Why element address this problem (sec. 5)? And (*c*) how does the Why element combine with the What element to constitute a computational-level theory (sec. 6)?

4. Why Ask Why: The Appropriateness Problem. The picture emerging so far is that the role of computational theories is to analyze information-processing tasks such as edge detection. The role of the What element is in characterizing, in mathematical terms, what is computed. In the case of edge detection, the characterization is in terms of the zero crossings of different-sized Laplacian filters of the image. The role of the Why element is to demonstrate the basis of this mapping function in the physical world (Marr 1977, 37) or the appropriateness and the adequacy of this mapping to the information-processing task (Marr 1982, 24). In the case of edge detection, it is to demonstrate why detecting zero crossings is appropriate for detecting edges. But what does this appropriateness come down to?

I want to suggest that appropriateness comes down to explaining why computing zero crossings is related to edge detection. After all, the mapping process that starts from the retinal image, *I*, and is described by the colocated zero crossings of $\nabla^2 G \times I$ is an "internal" mechanical process that takes place in our brain, whereas the information-processing task is defined, at least partly, by "external" features in the visual field, such as light reflectance, illumination, and object boundaries. We thus should wonder why on earth computing these zero crossings ends up with representations of physical edges: Why does it not end up with a representation of color? Why does it end up with meaningful representations at all?

A similar question arises in the context of stereo vision. The question is why computing the UC-pairing function has anything to do with matching: Why does the process that pairs dots from the two images, and is governed by the Uniqueness and Continuity constraints, end up with matching, namely, a representation of the same physical feature? Why is it that this UC pairing, and none of the other pairing functions, is appropriate for the task of stereo vision, leading to matching?

The problem, in its most general form, is that of explaining why a process that starts from a representation of W_1 (i.e., B_1) ends up with a representation of W_2 (i.e., B_2) and not with a representation of something else. How is it that the B_1–B_2 relations track or mirror the external W_1–W_2 relations? After all, we start the process with B_1, which represents W_1, and then proceed from B_1 to B_2 through a mechanical (causal) process that takes place in our brain. We thus should wonder what is it about the causal relation between B_1 and B_2 that guarantees that the neural representation of W_1 will lead to a neural representation of W_2.

The What element specifies the mathematical relation (function), *f*, between B_1 and B_2; in the case of edge detection, it is the colocated zero crossings of $\nabla^2 G \times I$. But this does not solve the appropriate problem; it simply moves the problem one level up. For now, we ask why the *f*-relation between B_1 and B_2 is appropriate for the information-processing task that is defined in terms of W_1 and W_2: Why does the process that starts from a representation of W_1 (i.e., B_1) and computes a mathematical function *f* end up with a representation of W_2? Why does the neural process that starts with a representation of light intensities, that is, the retinal image *I*, and computes the zero crossings of $\nabla^2 G \times I$ end up with a representation of physical edges such as object boundaries and not, say, with a representation of the object's colors? Why is it that the *f*-relations between B_1 and B_2 mirror, as it were, the W_1–W_2 relations?[7]

5. Explaining Why: The Role of Physical Constraints. Why does detecting zero crossings help in finding the edges of objects? Why are the zero crossings obtained from different-sized filters related to the same physical feature, and why is it that this feature is (often) a physical edge, for example, the boundary of an object? To answer these questions, Marr famously appeals to physical features in the environment known as natural or physical constraints: "The business of isolating constraints that are both powerful enough to allow a process to be defined and generally true of the world is a central theme of our inquiry" (1982, 23).

In the case of edge detection, Marr (1982, 68ff.) refers to the constraint of spatial localization, which means (in this context) that the things in the world that give rise to intensity changes are spatially localized. Generally speaking, intensity changes should be present under different-sized filters (this will not be the case if two or more local intensity changes are averaged together or two different physical phenomena operate on the same region). Another pertinent physical fact is that intensity changes in the image result from "surface discontinuities or from reflectance or illumination boundaries" (Marr and Hildreth 1980, 187). Taken together, it follows that a set of independent zero-crossing segments that have the same position and orientation indicates "the presence of an intensity change in the image that is due to a single physical phenomenon (a change in reflectance, illumination, depth, or surface orientation)" (Marr 1982, 70). Thus, the explanation goes, colocated zero crossings usually refer to a single physical feature because the (different) zero crossings point to the same location in the visual field, and in this location we (often) find only

7. There are affinities between Marr's appropriateness question and Fodor's question of how psychological processes "mirror" certain semantic relations (Fodor 1994, 9). But there are also differences. Fodor asks a how question about mechanisms. He is concerned with the mechanisms that support truth-preserving relations, and he argues that these mechanisms must be "Classical" (1994, 8–10; Fodor and Pylyshyn 1988). Marr argues that the question about mechanisms is the business of the algorithmic level, not the computational level. The computational level is concerned with why these (Classical) relations are "truth preserving" in the first place, reflecting the relevant external ("semantic") relations.

one feature that gives rise to intensity changes. This feature is (often) a physical boundary because sharp intensity changes occur along surface discontinuities, and they do not occur along the solid faces of surfaces.

What can we say about this sort of explanation? First, the explanation relates the zero crossings to contingent facts about the actual world, which comes in the form of physical constraints, that is, spatial location. It indicates that we happen to live in a world in which sudden changes in the retinal image (which the zero values of $\nabla^2 G \times I$ measure) strongly correlate with object boundaries. We could have lived in a world that consists of surfaces that sharply change reflectance across their solid faces. In such a world, the zero crossings of $\nabla^2 G \times I$ would be a poor way of detecting boundaries of objects. Zero crossings would correlate, if anything, with something else. In this sense, the explanation is ecological (Sterelny 1990).

Second, the explanation does not appeal to adaptive, that is, learning or evolutionary, processes. Marr obviously assumes that the correlation between zero crossings and object boundaries was established by some adaptive-based process. Still, he apparently does not think that the reference to learning and evolution explains why some things can correlate and some not, just as learning and evolution do not explain why the simple perceptron can correlate with an AND function but not with an
q14 XOR function.

Third, a careful examination reveals that Marr's explanation appeals to similarity between the internal mapping relations and external relations between the features that are being represented. The similarity is not at the level of physical properties. After all, the physiological properties of the brain are quite different from the physical and optical properties that make up our visual field. The similarity is at a more abstract level of mathematical properties. Put differently, the description in terms of zero crossings of $\nabla^2 G \times I$ refers to relations at two different levels. It describes internal mapping relations between the retinal and V1 cells. But it also describes relations between light intensities in the visual field (signified by I) and sudden changes in light intensity; these changes happen to occur, at least in our world, along object boundaries. The mathematical relation between the features that are being represented—for example, light reflectance in the visual field and light reflectance along object boundaries—is also that of (extreme points of) derivation.

Thus, Marr not only demonstrates that the internal mathematical function correlates with the contingent world that we live in. He also underscores the basis of this correlation, which is a similarity of mathematical structures. This mathematically based similarity, I maintain, is the key in addressing the appropriateness problem and, hence, in computational-level explanations. In the remainder of this section, I show that this idea

Friday Aug 13 2010 04:14 PM/PHOS10157/2010/77/4
/sji/lfw/ /editing started/use-graphics/narrow/default/

of mathematically based similarity is also central in other computational-level theories. In the next section, I spell out the structure of the explanation.

Let us start with the cash register example. The What element characterizes the mapping function from input digits to output digits as addition. We arrive at this characterization when noticing that the machine maps digits to digits and that this mapping satisfies the rules of commutativity, associativity, zero, and inverses. We then ask why computing addition is appropriate for the information-processing task, which is combining the prices of purchased items to arrive at the final bill. Or, as Marr puts it, we ask "why the cash register performs addition and not, for instance, multiplication when combining the prices of the purchased items to arrive at the final bill" (1982, 22). After all, there is no a priori reason for the final bill to be the sum of the purchased items and not their product. We can certainly think of stores in which the cashier executes multiplication and not addition.

In answering this question, Marr demonstrates that the external relation between the final bill and the purchased items in this particular case is also that of addition. Spending enough time around the store, we see that the relation between the purchased items and the final bill is that of addition and not of multiplication. The reason, according to Marr, is that the rules ("constraints") of purchasing in this store define addition. These are the rules of zero ("if you buy nothing, it should cost you nothing; and buying nothing and something should cost you the same as buying just the something"; 1982, 22), commutativity ("the order in which goods are presented to the cashier should not affect the total"), associativity ("arranging the goods into two piles separately should not affect the total amount you pay"), and inverses ("if you buy an item and then return it for a refund, your total expenditure should be zero"; 23).

Establishing that the relation between the purchased items and the final bill is that of addition, Marr draws the conclusion that the input-output addition mapping in the cash register is appropriate for the task in this particular store. This explanation appeals to the fact that this internal mapping (of addition), defined over digits, corresponds to an external relation between the represented items (in the abstract), namely, between the prices of purchased items and the final bill.

Let us now turn to stereo vision. We recall that Marr argues that the visual system computes a UC-pairing function. But we also recall that he asks why computing this function is appropriate for matching: why is it that the UC pairing, and not any of the other pairings, provides a solution for the correspondence problem? In answering this question, Marr motivates Uniqueness and Continuity by appealing to physical constraints. Uniqueness ("a black dot from one image can match no more than one black dot from the other image") is motivated by the spatial localization constraint, which says that "a given point on a physical surface has a unique position in space at any one time" (Marr and Poggio 1976, 284; see also Marr 1982, 112–13). Continuity ("disparity varies smoothly almost everywhere") is motivated by cohesiveness of matter, which says that "matter is cohesive, it is separated into objects, and the surfaces of objects are generally smooth compared with their distance from the viewer" (Marr and Poggio 1976, 284; see also Marr 1982, 112–13).

We can note that here, too, the reference to physical constraints shows that the internal UC pairing corresponds to an external UC pairing. The similarity is not easily seen, but it is there: what is being represented by each image is an array of light intensities in the visual scene. Recall that these light intensities partially depend on the viewpoint and that the viewpoints of the two eyes are slightly different. What the physical constraints show is that when you apply a UC pairing to these external arrays of light intensities (not to the images), you get matching: the output of the function consists of pairs of elements from the visual scene, whereas the physical constraints dictate that in our world the elements of each pair are really one physical feature.

Marr is unique in making the idea of mathematically based similarity a key feature of computational-level theories. The idea itself, however, is not unique to Marr's theories of vision. It is found in Classical (see discussion in Ramsey 2007, 77–92) and Connectionist (see discussion in Churchland 2007) models of cognition. It is also found in many computational models in neuroscience, for example, the Zipser-Andersen model (Zipser and Andersen 1988; see discussion in Grush 2001). Another example is the oculomotor integrator, whose task is to hold the eyes still between the saccadic movements (Robinson 1989; Seung 1996, 1998; see discussion in Shagrir 2010). In this example, both the internal and the external relations are described in terms of integration. The internal mapping relation, from one stable state of the neural memory, is described as integration over the pulse saccadic inputs. The external relation between the eye positions (represented by the pertinent stable states) is described as integration over the eye velocity with respect to time (which is represented by the pulse inputs).

6. The Structure of Computational-Level Explanations. According to Marr, the job of computational-level theories is to analyze vision as an information-processing task "in a way that is independent of the particular mechanisms and structures that implement them in our heads" (1982, 19). It is time to clarify how the What element and the Why element join forces in providing such an analysis. The gist of the account is this: an information-processing task is mapping from one representation to an-

Friday Aug 13 2010 04:14 PM/PHOS10157/2010/77/4
/sji/lfw/ /editing started/use-graphics/narrow/default/

other (fig. 1), such as in edge detection (fig. 2). The goal of a computational analysis is to explain why the process that starts from a representation of W_1 (i.e., B_1) ends up with a representation of W_2 (i.e., B_2) and not with a representation of something else. A computational-level theory answers this question by pointing to a similarity at the more abstract level between the internal B_1–B_2 relations and the external W_1–W_2 relations. It states that the mathematical B_1–B_2 relations are "isomorphic" (in a sense specified below) to the mathematical W_1–W_2 relations. The role of the What element is to specify the mathematical function that is being computed, namely, the *f*-relation between B_1 and B_2. The role of the Why element is to demonstrate that the W_1–W_2 relation is also an *f*-relation.

Putting it schematically, the structure of a computational-level explanation is, ideally, as follows:

$$i(B_1) = W_1, \quad \text{where } i = \text{represents},$$

$$f(B_1) = B_2,$$

$$f(W_1) = W_2,$$

$$f(i(x)) = i(f(x)).$$

Therefore, $i(B_2) = W_2$.

The first premise states that B_1 represents W_1. In our example, it states that the intensity values, $I(x, y)$, represent light intensities in the visual scene, which result from such features as light reflectance, illumination, geometry, and viewpoint. The second premise states that the mathematical relation between B_1 and B_2 is *f*. Characterizing this function is the business of the What element. The What element provides a description of the internal mapping relations between the abstract properties of B_1 and B_2. In our example, the What element specifies the mapping relations between abstract properties of the photoreceptors, ganglion cells, and cells in V1 (abstracted from the electric activity of these cells) in terms of zero crossings of $\nabla^2 G \times I$.

The job of the Why element is to establish the third premise, namely, to show that the W_1–W_2 relations are also *f*-relations. Establishing this premise is often trivial, but it sometime takes considerable sophistication and effort. Edge detection is an example. As Marr says, it does not immediately follow that if the term $I(x, y)$ refers to the array of light intensities in the visual field, then the colocated zero crossings of the different-scale filtering of $\nabla^2 G \times I$ stand for physical edges. That they do is a contingent fact about our visual environment and should be argued for. The argument, we saw, appeals to the physical constraints, which are facts and assumptions about our physical world.

Explicating the structure of explanation reveals that the first three premises do not suffice to derive the conclusion. There is another, implicit, assumption here, which is the fourth premise. This premise asserts that the visual system is a representational system in some very strong sense: it respects, as it were, mathematical relations, or at least some mathematical relations (i.e., f). Put differently, the premise states that the representation function, i, is an isomorphism over the (mathematical) relation f. Marr never makes this assumption explicit, nor does he argue for it.[8] But the assumption is essential for the derivation of the conclusion.[9]

This sketch of computational explanations requires further elaboration and explication. Here, I will make two brief comments. One is about abstraction: the function f is mathematical in the sense that its domain and range are mathematical entities such as real numbers, geometrical relations, set-theoretic structures, and so forth. Thus, the function f relates numbers (or other mathematical entities) and not physical entities. These entities are mathematical values, as magnitudes that abstract from pertinent physical properties. At one level, the function relates numbers that abstract from the representations (e.g., electrical cellular activity). At another, it relates magnitudes that abstract representational contents (e.g., distal objects in the environment). The other comment is about approximation: a visual system, as a biological system, seldom computes the function f that is stated by the computational-level theory; it only approximates it. Likewise, the function f is, at best, an approximation of the distal relation. A computational-level theory thus relates the ideally computed function f, with an (ideal) external f-relation.[10]

7. Computation and Intentionality. In closing, I compare my interpretation of Marr with other interpretations. My aim, however, is neither to provide an exhaustive survey of all other interpretations nor to criticize them. It is to locate my interpretation in the wider philosophical landscape

8. This assumption is in accord with comments (private communication) by Frances Egan, David Kaplan, and Shimon Ullman to the effect that the notion of representation that Marr assumes, at least implicitly, is more than simple covariation. Ullman specifically pointed out that this condition is part of what he sees as a representational system.

9. The derivation can be presented as follows: according to the fourth premise, $f(i(B_1)) = i(f(B_1))$, and so, according to the second premise, $f(i(B_1)) = i(B_2)$. According to the first premise, $f(i(B_1)) = f(W_1)$, and so, according to the third premise, $f(i(B_1)) = W_2$. Taken together, we get the conclusion: $i(B_2) = W_2$.

10. Marr himself (1982, 28–29) presents the computational-level theory as a level of competence, stating that the distinction between it and the How levels of mechanisms (see, e.g., Craver 2007) is "roughly his [Chomsky's] distinction between competence and performance"; see also Horgan and Tienson (1994), Gillman (1996), and Polger (2004).

and to clarify the ways in which Marr's computational-level theories can be seen as intentional.

Interpretations of Marr have focused on issues of explanation, individuation, and methodology. Let us start with explanation. It is widely agreed that, according to Marr, a theory of vision should analyze a visual, information-processing task. The debates concern the role of computational-level theories in this analysis. Most philosophers take it that the role of all computational-level theories is to provide an information-processing description of the task, in terms of the visual (representational) content of neural states. In other words, the claim is that the main job of computational-level theories is "the specification of the explanandum—the cognitive task that we are attempting to explain. Marr calls this the 'computational' level, where the specification is typically an input-output function" (Ramsey 2007, 41). Bermúdez also writes that "a computational analysis will identify the information with which the cognitive system has to begin (the *input* to that system) and the information with which it needs to end up (the *output* from that system)" (2005, 18). And Horst (2009) concludes that "at the highest level was a specification of what task a system was designed to perform: for example, in the case of vision, to construct a three-dimensional representation of distal stimuli on the basis of inputs to the retina."

It should be noted that interpreters of Marr who adopt this line of thinking believe that Marr confuses the "computational" with the "representational." Sterelny, for example, says that "Marr, very confusingly, calls it the 'computational' level" (1990, 46). Dennett, after associating Marr's computational level with his intentional level, says that "this specification was at what he [Marr] called, misleadingly, the computational level" (1994, 681). Ramsey notes that "this [computational] label is somewhat misleading" (2007, 41 n. 3), and Horst (2009) says that "this level Marr (somewhat unfortunately) called the 'computational level.'" But, of course, Marr is confused only if he takes his computational-level theories to provide such information-processing descriptions of the task.

My view is that computational-level theories do not aim to provide such intentional descriptions. The characterization of the task in terms of informational content is often made before we invoke the computational-level theory; in the case of edge detection it is being made at the level of neurophysiology. That photoreceptors are sensitive to light reflectance, that information from the retina arrives to V1, and that cells in V1 are sensitive to oriented lines were all discovered by neurophysiologists, using techniques such as single-cell recording, long before Marr q15 invoked his computational theories.[11] Computational theories, according to Marr, aim to explain something about these information-processing tasks. They aim to explain why the neural process that starts with the representation of light intensities ends up with a representation of physical edges.

Other interpreters also think that the computational level plays an explanatory role. Shapiro (1997), for example, writes that "at the computational level of theory the theorist describes what I shall call *chief* tasks and *service* tasks. . . . The chief task of the visual system is the derivation of 3-D shape representations from information encoded in 2-D images. Service tasks are those tasks the completion of which contribute to the achievement of the chief task" (134). In particular, he argues, the information-processing description of the service tasks, in terms of informational content, contributes to the understanding of the chief task. This is in accord with the functional picture of computational explanations, according to which the capacity of a system (chief task) is explained in terms of the capacities of the components (service tasks) of which it is composed.[12]

I agree with Shapiro that the information-processing descriptions of service tasks (e.g., stereo disparity) account for the chief task (e.g., stereo vision). But I insist that the paradigm cases of computational theories that Marr and his students advance—edge detection, stereo disparity, and structure from motion—are all (service) tasks that do not comfortably break down to further service tasks. The explanation of these tasks is very different and refers to two kinds of formal relations. One is the inner input-output mathematical relation, and the other is the outer mathematical relation between what is being represented by the inputs and the outputs. The first sort of relation, of the input-output function, is not intentional. The latter sort can be seen as some kind of mathematical or formal content, in that these mathematical relations are abstracted from

11. This practice is customary in computational neuroscience. Theoreticians invoke the computational approach to analyze tasks whose information-processing descriptions are determined beforehand via electrophysiological, single-cell-recording experiments. To take two other examples: the Zipser-Andersen model simulates the activity of PPC cells that combine information about eye position and stimulus location in retinotopic coordinates (Andersen, Essick, and Siegel 1985). And Canon, Robinson, and Shamma (1983) propose a neural network that models the activity of the oculomotor system, on the basis of prior experimental results; these results show that some cells encode eye velocity and others eye position (see Robinson 1989).

12. This picture is attributed to Haugeland (1978) and Cummins (1983). We can note, however, that the functional strategy fits better with Marr's algorithmic level than with his computational level.

representational contents, for example, physical edges.[13] Thus, if computational explanations are intentional, it is by virtue of referring to these features—that is, the formal content—and not to specific content.

Egan presents a very different picture according to which the aim of computational-level theories is to provide a mathematical characterization of the (input-output) computed function. The intentional content is invoked to explain how the computed mathematical function can serve a cognitive function in a particular context. Thus, Egan argues that "an explanation of how the visual system detects the depth of the scene from information contained in two-dimensional images is forthcoming only when the states characterized in formal terms by the theory are construed as *representations of distal properties*" (1995, 190).[14]

I agree that one aim of computational-level theories is to characterize the (mathematical) computed function; this is the job of the What element. I also agree that computational-level theories are concerned with the embedment of the (computing) system in its environment; indeed, the role of the Why element is to explain the physical basis of the (computed) mathematical function in the physical world. Still, I resist the idea that the computational explanations answer this embedment question by specifying the intentional content of visual states. The embedment question is not "What is represented by the zero crossings?" For one thing, we know the answer to this before we invoke the computational-level theory, from the electrophysiological experiments; for another, Marr had no reason to make his dramatic statements about the why, if the role of the Why element is just to specify the content of the representations. The issue, rather, is explaining why the outputs of the process represent what they represent: why computing zero crossings leads to representations of edges and not to something else. And this explanation is provided not by referring to the intentional content of the visual states but by revealing the structural similarities between the internal input-output relations and the pertinent external relations.[15]

13. The term "formal content" is coined by Sher (1996) and "mathematical content" by Egan (1995). Sher explicitly argues that formal content is an abstraction from concrete properties and objects in the world. Egan seems to associate mathematical content with the mathematical function that is being "implemented."

14. Egan does not think, however, that the specification of content is part of computational-level theory; it is made after the computational-level theory has accomplished its task of specifying the mathematical function.

15. It is true that computational explanations refer to physical facts (constraints) about the physical environment. But even if we want to count these facts as intentional (Burge 1986), they are not, in any obvious sense, the things that are being represented by the inputs and the outputs. After all, spatial location is a fact about our physical world,

Much of the philosophical debate about Marr concerns individuation. Most of Marr's interpreters argue that his computational theories are intentional, in the sense that they make an essential reference to specific visual content. The claim, in other words, is that content affects the computational (type) identity of visual states.[16] Egan, in contrast, argues that computational types are affected by mathematical content but not by specific visual content: "For the purpose of individuation, the precise mathematical description given by the theory of computation is the description that counts" (1995, 186). She thus concludes that computational theories are individualistic (see also Egan 2010).

I agree with Egan that specific content does not always make a computational difference, namely, that two systems can be computationally identical even if they carry different content. I also agree with Egan that computational ("formal") types are affected by mathematical content in the sense that they refer to mathematical relations that are abstracted from representational content, for example, physical edges. Still, it does not follow that specific content does not affect computational individuation (and, hence, that computational theories cannot be externalistic). When you abstract from different contents, you sometimes get the same mathematical content, in which case the systems are computationally identical, but you sometimes get different mathematical content, in which case they are not.[17]

Finally, I discuss methodology. Marr famously advocates a top-down methodology, wherein the top level is the computational level. Most interpreters of Marr hold that the top level specifies an intentional description of the task; they think that the methodology is top down in the sense that the intentional description restricts the algorithm (see Shapiro 1997, 137). I agree that the top computational level provides guidance to the specification algorithms, although I think that the specification is function theoretic (mathematical) and not information processing (intentional). The main difficulty with the suggested interpretation, however, lies elsewhere: it is that the intentional (information-processing) description itself is very often provided through electrophysiological single-cell-recording experiments that combine "implementation," that is, action potentials,

but it is not a representational content of nerve cells, nor does Marr say or imply that our visual system represents it.

16. See Burge (1986), Kitcher (1988), Segal (1989, 1991), Sterelny (1990), Davies (1991), Morton (1993), Shapiro (1993, 1997), Peacocke (1994), and Silverberg (2006). They debate among themselves as to whether this content is "wide" or "narrow."

17. For further discussion of this argument, see Wilson (1994, 2004, 162ff.), Bontly (1998), Shagrir (2001), Horowitz (2007), Piccinini (2008), and Sprevak (2010).

with environmental stimuli. This makes the methodology no more top down than bottom up.[18]

According to my interpretation, Marr's contentious claim about methodology refers to the specification of the mathematical function that is being computed. The obvious method to arrive at this mathematical function is through abstraction from neurological properties (bottom up). But Marr argues that this method is not effective and that the more effective way is to start with meaningful (physical) constraints on the relations between the features that are being represented.[19] In the case of edge detection, we start with the fact that there are sudden changes in light reflectance along physical edges, and we infer from this that the pertinent mathematical relation is that of a derivative. The next step is to look for this mathematical relation inside the brain. In the case of edge detection, we look for derivation in early visual processes.

When properly understood, we can see that Marr's methodological framework is still widespread in computational neuroscience. It is found, for example, in the study of the oculomotor integrator (mentioned above). The first stage in the investigation was the discovery, through single-cell experiments, that the external relation between the features represented by inputs (i.e., transient eye velocity) and outputs (i.e., eye position) is that of integration (in fact, this is the reason the system was called integrator in the first place).[20] Realizing this, theoreticians moved to the next step of looking for an integration relation in the cellular activity of
q16 the representing cells (see Robinson 1989; Seung 1996, 1999). Such studies not only materialize Marr's conceptual and methodological maxims. They also combine experimental studies, alongside sophisticated theoretical work, which are exemplars of Marr's vision.

REFERENCES

Andersen, Richard A., Greg K. Essick, and Ralph M. Siegel. 1985. "Encoding of Spatial Location by Posterior Parietal Neurons." *Science* 230:456–58.

Bermúdez, José Luis. 2005. *Philosophy of Psychology: A Contemporary Introduction*. New York: Routledge.

Bontly, Thomas D. 1998. "Individualism and the Nature of Syntactic States." *British Journal for the Philosophy of Science* 49:557–74.

Burge, Tyler. 1986. "Individualism and Psychology." *Philosophical Review* 95:3–45.

Butler, Keith. 1998. "Content, Computation, and Individuation." *Synthese* 114:277–92.

18. For further discussion, see Bermúdez (2005), esp. 21 n. 3; see also Churchland and Sejnowski (1992).

19. As Marr's students put it, a computational theory includes "an analysis of how properties of the physical world constrain how problems in vision are solved" (Hildreth and Ullman 1989, 582); for more details, see Ullman (1979, 3–4).

20. The story is told in Robinson (1989).

Cannon, Stephen C., David A. Robinson, and Shihab A. Shamma. 1983. "A Proposed Neural Network for the Integrator of the Oculomotor System." *Biological Cybernetics* 49:127–36.
Churchland, Patricia S., and Terrence J. Sejnowski. 1992. *The Computational Brain*. Cambridge, MA: MIT Press.
Churchland, Paul M. 2007. *Neurophilosophy at Work*. New York: Cambridge University Press.
Craver, Carl F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. New York: Oxford University Press.
Cumming, Bruce. 2002. "An Unexpected Specialization for Horizontal Disparity in Primate Primary Visual Cortex." *Nature* 418:633–36.
Cummins, Robert E. 1983. *The Nature of Psychological Explanation*. Cambridge MA: MIT Press.
Davies, Martin. 1991. "Individualism and Perceptual Content." *Mind* 100:461–84.
Dennett, Daniel. 1994. "Cognitive Science as Reverse Engineering: Several Meanings of 'Top-Down' and 'Bottom-Up.'" In *Logic, Methodology and Philosophy of Science IX*, ed. Dag Prawitz, Brian Skyrms, and Dag Westerståhl, 679–89. Amsterdam: Elsevier Science.
Durand, Jean-Baptiste, Simona Celebrini, and Yves Trotter. 2007. "Neural Bases of Stereopsis across Visual Field of the Alert Macaque Monkey." *Cerebral Cortex* 17:1260–73.
Egan, Frances. 1991. "Must Psychology Be Individualistic?" *Philosophical Review* 100:179–203.
———. 1992. "Individualism, Computation and Perceptual Content." *Mind* 101:443–59.
———. 1995. "Computation and Content." *Philosophical Review* 104:181–204.
———. 2010. "A Modest Role for Content." *Studies in History and Philosophy of Science*, forthcoming.
Fodor, Jerry A. 1994. *The Elm and the Expert*. Cambridge, MA: MIT Press.
Fodor, Jerry A., and Zenon W. Pylyshyn. 1988: "Connectionism and Cognitive Architecture: A Critical Analysis." *Cognition* 28:3–71.
Gillman, Daniel. 1996. "Optimization and Simplicity: Computational Vision and Biological Explanation." *Synthese* 107:293–323.
Grush, Rick. 2001. "The Semantic Challenge to Computational Neuroscience." In *Theory and Method in the Neurosciences*, ed. Peter Machamer, Rick Grush, and Peter McLaughlin, 155–72. Pittsburgh: University of Pittsburgh Press.
Harnish, Robert M. 2002. *Minds, Brains, Computers: An Historical Introduction to the Foundations of Cognitive Science*. Malden, MA: Blackwell.
Haugeland, John. 1978. "The Nature and Plausibility of Cognitivism." *Behavioral and Brain Sciences* 2:215–26.
Hildreth, Ellen, and Shimon Ullman. 1989. "The Computational Study of Vision." In *Foundations of Cognitive Science*, ed. Michael Posner, 581–630. Cambridge, MA: MIT Press.
Horgan, Terence E., and John L. Tienson. 1994. "A Nonclassical Framework for Cognitive Science." *Synthese* 101:305–45.
Horowitz, Amir. 2007. "Computation, External Factors, and Cognitive Explanations." *Philosophical Psychology* 20:65–80.
Horst, Steven W. 2009. "The Computational Theory of Mind." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University, http://plato.stanford.edu/entries/computational-mind/.
Kitcher, Patricia S. 1988. "Marr's Computational Theory of Vision." *Philosophy of Science* 55:1–24.
Kroustallis, Basileios. 2006. "Content Individuation in Marr's Theory of Vision." *Journal of Mind and Behavior* 27:57–72.
Marr, David C. 1977. "Artificial Intelligence: A Personal View." *Artificial Intelligence* 9:37–48.
———. 1982. *Vision*. San Francisco: Freeman.
Marr, David C., and Ellen Hildreth. 1980. "Theory of Edge Detection." *Proceedings of the Royal Society of London* B 207:187–217.

Marr, David C., and Tomaso A. Poggio. 1976. "Cooperative Computation of Stereo Disparity." *Science* 194:283–87.
———. 1979. "A Computational Theory of Human Stereo Vision." *Proceedings of the Royal Society of London* B 204:301–28.
Morton, Peter A. 1993. "Supervenience and Computational Explanation in Vision Theory." *Philosophy of Science* 60:86–99.
Newell, Allen. 1982. "The Knowledge Level." *Artificial Intelligence* 18:87–127.
Peacocke, Christopher. 1994. "Content, Computation and Externalism." *Mind and Language* 9:303–35.
Piccinini, Gualtiero. 2008. "Computation without Representation." *Philosophical Studies* 137:205–41.
Polger, Thomas W. 2004. "Neural Machinery and Realization." *Philosophy of Science* 7: 997–1006.
Pylyshyn, Zenon W. 1984. *Computation and Cognition.* Cambridge, MA: MIT Press.
Ramsey, William. 2007. *Representation Reconsidered*. Cambridge: Cambridge University Press.
Robinson, David A. 1989. "Integrating with Neurons." *Annual Review of Neuroscience* 12: 33–45.
Segal, Gabriel. 1989. "Seeing What Is Not There." *Philosophical Review* 98:189–214.
———. 1991. "Defense of a Reasonable Individualism." *Mind* 100:485–94.
Seung, Sebastian H. 1996. "How the Brain Keeps the Eyes Still." *Proceedings of the National Academy of Sciences* 93:13339–44.
———. 1998. "Continuous Attractors and Oculomotor Control." *Neural Networks* 11:1253–58.
Shagrir, Oron. 2001. "Content, Computation and Externalism." *Mind* 110:369–400.
———. 2010. "Computation, San Diego Style." *Philosophy of Science*, forthcoming.
Shapiro, Lawrence A. 1993. "Content, Kinds, and Individualism in Marr's Theory of Vision." *Philosophical Review* 102:489–513.
———. 1997. "A Clearer Vision." *Philosophy of Science* 64:131–53.
Sher, Gila. 1996. "Did Tarski Commit 'Tarski's Fallacy'?" *Journal of Symbolic Logic* 61: 653–86.
Silverberg, Arnold. 2006. "Chomsky and Egan on Computational Theories of Vision." *Minds and Machines* 16:495–524.
Sprevak, Mark D. 2010. "Computation, Individuation, and the Representation Condition." *Studies in History and Philosophy of Science*, forthcoming.
Sterelny, Kim. 1990. *The Representational Theory of Mind: An Introduction*. Cambridge, MA: Blackwell.
Ullman, Shimon. 1979. *The Interpretation of Visual Motion*. Cambridge, MA: MIT Press.
Wilson, Robert A. 1994. "Wide Computationalism." *Mind* 103:351–72.
———. 2004. *Boundaries of the Mind: The Individual in the Fragile Sciences; Cognition*. Cambridge: Cambridge University Press.
Zipser, David, and Richard A. Andersen. 1988. "A Back-Propagation Programmed Network that Simulates Response Properties of a Subset of Posterior Parietal Neurons." *Nature* 331:679–84.

QUERIES TO THE AUTHOR

q1. Au: Per preferred journal style, italics omitted here and throughout when used for emphasis yet meaning seems clear without them.

q2. Au: Please shorten abstract (currently 136 words) to no more than 100 words.

q3. Au: Footnote 1: Added "index" here, OK? Or do you prefer "Web site," or "database," or something else to describe the ISI Web of Knowledge?

q4. Au: Here and throughout, quotations were run into the text when they did not meet the minimum 100 word requirement for block quotations.

q5. Au: Please provide page number(s) for two quotes here in second sentence of this paragraph.

q6. Au: Sentences cannot begin with mathematical symbols; therefore, "state" inserted here. Similar changes made throughout.

q7. Au: Figure 1: Replaced "event/property" with "object/property" to match all other uses in text and figure note.

q8. Au: For parallelism with remainder of sentence, do you wish to add a "from" clause to "edge detection" (e.g., "edge detection from ...")?

q9. Au: Explicit multiplication asterisks and dots replaced with crosses here and throughout for consistency and because crosses are preferred over asterisks or dots. (See also query 11.)

q10. Au: Please define LGN here.

q11. Au: Use of an operator (vs. variable) in run of text is not preferred (for clarity), thus "the multiplication symbol" inserted here to replace "*" (or " × "; see query 9).

q12. Au: Are you using "dots" (here and below) interchangeably with "circles" (below and in figure legend)? If yes, should "circles" be used exclusively for consistency (i.e., should we replace all uses of "dots" with "circles")? If no, can additional explanation of what "dots" means and what it refers to be incorporated here or possibly in n. 5? Are they dots viewed on a wall?

q13. Au: Figure 5: Defined AAAS as American Association for the Advancement of Science here, OK? If not, please define.

q14. Au: Please define AND and XOR here.

q15. Au: Footnote 11 (formerly n. 13): Please define PPC here.

q16. Au: Seung 1999 not in reference list; please add, or omit citation here. Do you intend Seung 1998 instead? NB: These citations were originally in n. 25 (and moved here to conform to preferred journal style).

Friday Aug 13 2010 04:14 PM/PHOS10157/2010/77/4
/sji/lfw/ /editing started/use-graphics/narrow/default/