Marr's Computational Level and Delineating Phenomena

Oron Shagrir and William Bechtel

Abstract: A key component of scientific inquiry, especially inquiry devoted to developing mechanistic explanations, is delineating the phenomenon to be explained. The task of delineating phenomena, however, has not been sufficiently analyzed, even by the new mechanistic philosophers of science. We contend that Marr's characterization of what he called the *computational level* (CL) provides a valuable resource for understanding what is involved in delineating phenomena. Unfortunately, the distinctive feature of Marr's computational level, his dual emphasis on both *what* is computed and *why* it is computed, has not been appreciated in philosophical discussions of Marr. Accordingly we offer a distinctive account of CL. This then allows us to develop two important points about delineating phenomena. First, the accounts of phenomena that figure in explanatory practice are typically not qualitative but precise, formal or mathematical, representations. Second, delineating phenomena requires consideration of the demands the environment places on the mechanism—identifying, as Marr put it, the basis of the computed function in the world. As valuable as Marr's account of CL is in characterizing phenomena, we contend that ultimately he did not go far enough. Determining the relevant demands of the environment on the mechanism often requires detailed empirical investigation. Moreover, often phenomena are *reconstituted* in the course of inquiry on the mechanism itself.

1. Introduction

Bogen and Woodward (1988) convincingly demonstrated that scientific explanations are directed at phenomena, not data. Phenomena are regular, repeatable types of events, processes, or states; Bogen and Woodward offer examples of what they mean by phenomena: "weak neutral currents, the decay of the proton, and chunking and recency effects in human memory" (p. 306). The new mechanistic philosophers of science have embraced Bogen and Woodward's focus on phenomena, holding that mechanisms are identified in terms of the phenomena they are to explain (Machamer, Darden, & Craver, 2000; Glennan, 2002; Bechtel & Abrahamsen, 2005). For the most part they, following the lead of Bogen and Woodward, have stayed with textbook accounts of phenomena, offering as examples the action potential, glycolysis, protein synthesis, and long-term potentiation. The specification of phenomena is generally treated as unproblematic—the challenge is explaining them. Kauffman (1971) noted the importance of selecting among the many things organisms do before attempting to explain how they do so as that selection will affect the explanation offered. Bechtel and Richardson (1993/2010) drew attention to the fact that often much research must be done to *delineate* phenomena and that sometimes in the course of developing a mechanistic account scientists end up recognizing that the phenomenon is different than they initially supposed. For example, research in biochemical genetics began by trying to account for the role of genes in generating phenotypic traits, but in the course of their research Beadle and Tatum (1941) recharacterized genes as involved in generating enzymes. Bechtel and Richardson refer to such revisions in the account of the phenomenon as *reconstituting* the phenomena. But even they do not develop the fact that the phenomena for which explanations are sought are typically characterized in a far more detailed, quantitative fashion, and that *saving* such quantitative features of phenomena is often a critical challenge in explanation and an important criterion in evaluating putative explanations.

Insofar as phenomena are the explananda for mechanistic explanation it is important to clarify what a phenomenon is. Although measuring phenomenon quantitatively is more important than mechanists have recognized, not everything that can be measured quantitatively is treated as a phenomenon to be explained by a mechanism, even if it is the effect of a mechanism and plays a role in evaluating proposed accounts of the mechanism. In the case of the action potential, the change over time of the electrical charge across the neuron membrane is part of the phenomenon, but the temporary increase in sodium concentration inside the neuron is not, although both can be characterized quantitatively. Likewise, the phenomenon of long-term potentiation is characterized by the increased number of action potentials generated by a neuron in response to a stimulus but not by how much ATP is consumed in the process. Given the multitude of items that can be measured quantitatively, it is important that we be able to differentiate those that do and those that do not count as phenomena for which a mechanism is sought.

We will argue that important insights into the role of phenomena in mechanistic explanations can be found in Marr's (1982) characterization of what he called *the computational level*. Marr introduces his well-known account of levels to counter what he took to be a shortcoming in the practice of neuroscience: the preoccupation with the components of the visual processing mechanism—the properties of cells and their behavior. Marr's objective was not to repudiate the search for mechanism but to recast it in terms of his tri-level framework of computational, algorithmic, and implementational levels. Marr contended that "Vision is . . . first and foremost, an information-processing task." Delineating this information processing-task – the phenomenon – is the job of what Marr called the computational level. The algorithmic level characterizes the system of representations that is being used, e.g., decimal vs. binary, and the algorithm employed for transforming representations of inputs into those of outputs. The implementation level specifies how the representations and algorithm are physically realized.

What is involved in characterizing vision as performing an information-processing task? Marr associates the computational level with two aspects, the *what* and the *why*. In the introductory, "philosophical outlook", chapter of *Vision*, Marr says that "the most abstract is the level of *what* the device does and *why*" (p. 22). The job of the what-aspect is to specify what is computed. The job of the why-aspect is to demonstrate the appropriateness and adequacy of what is being computed to the information-processing task (pp. 24-25). In "Artificial intelligence: a personal view", Marr states 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*" (Marr, 1977, p. 37).

But what exactly Marr means by these *what* and *why* aspects of CL? Marr never provided a systematic and detailed account of his notion of CL; what he does say about it is often brief and somewhat vague. Instead, Marr provided a set of computational theories of specific visual tasks. These impressive theories induced an enormous amount of research into computer and biological vision. The conceptual task of explicating the notion of a computational-level theory was left to philosophers, who provided, in turn, radically different interpretations.

Unfortunately, as we will argue in the next section, none of these interpretations is adequate to the distinctive role Marr envisaged for CL. We will review three of the most prominent accounts in part 2 and show how each fails short of what Marr seems to have had in mind. In part 3 we advance an alternative interpretation that we contend better captures what Marr saw as the importance of the *what* and *why* aspects of CL analysis. CL theory, as we see it, provides a formal or mathematical account of the task the visual system performs in the actual physical world in which it functions. Our goal, though, is not simply to engage in Marr exegesis. Rather, we contend that understanding what Marr had in mind by analysis of CL is extremely important for providing an adequate account of the role delineating phenomena plays in science, especially science devoted to the identification of mechanisms. As we argue in part 4, the phenomena for which mechanisms are sought require formal or mathematical characterizations that are grounded in the context in the world in which the mechanism functions. In part 5 we will argue that Marr did not go far enough in characterizing phenomena in CL terms. The tasks mechanisms are to perform are not simply givens to scientists, but typically discovered through empirical (observational or experimental), inquiry. Moreover, they are frequently revised in the course of developing explanations of them.

Following Marr, we will take visual perception as our primary exemplar. But the implications of Marr's approach extend to any phenomena that are appropriately characterized in computational terms, that is, information-processing terms. Marr's account was designed for neuroscience and, although some contest it, the computational metaphor is appropriate for brain function generally. The task for the brain and the various processes occurring in it is to extract and use information to control the functioning of an organism. Moreover, the reasons that justify reference to computation and information processing in the case of the brain apply far more broadly to control processes in living organisms. Cell signaling systems, for example, process information to control such activities as use of different metabolites for fuel, the repair of DNA, or the synthesis of proteins and researchers are increasingly employing information-processing language to characterize these processes (Shapiro, 2011).[1] But the activities thereby regulated—the transformation of energy into ATP, or the synthesis or degradation of proteins—are not appropriately characterized in information processing terms. Exploring what insight Marr's account of CL offers to characterizing phenomena in those cases goes beyond the scope of this paper.

2. Shortcomings of Extant Accounts of Marr's Computational Level (CL)

We cannot review all attempts to explicate Marr's notion of CL, but will focus on three whose shortcomings are illuminating. According to the first (the "standard" interpretation), CL characterizes the information-processing task, mainly in

[1] The concept of information has been employed in many different ways in biology, where it took on special significance after Watson and Crick (1953) used it to characterize the function of the genetic code. Some, inspired by Shannon (1948) have treated information in causal terms (effects carry information about their causes). Others such as Maynard Smith (2000) have defended a teleosemantic notion in which the content of a signal is fixed by natural selection. Yet others have rejected the application of the concept of information to genes as metaphorical (Griffiths, 2001). See (Levy, 2011) for a valuable discussion that elucidates the roles different accounts of information play in biological theorizing.

intentional terms. According to the second, the aim of CL is to provide a mathematical or a formal theory, and according to the third, CL provides a sketch of mechanism.

2.1. The "standard" interpretation: Specifying an information-processing task

Most interpreters of Marr assume that the role of the computational level is specifying an information-processing visual or cognitive task: "At the highest level was a specification of what task a system was designed to perform: e.g., in the case of vision, to construct a three-dimensional representation of distal stimuli on the basis of inputs to the retina" (Horst, 2009).This information-processing task is often described in terms of the contents of the input and the output representations: "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)" (Bermúdez, 2005, p. 18). Thus edge-detection is the mapping from representations of light intensities to representations of physical edges (e.g., object boundaries). Shape-from-shading is the mapping from representations of shading to representations of shape, and so on. When put in the context of the *what* and *why* aspects, the standard interpretation apparently associates the *what* with the mapping of inputs representations to output representations, and the *why* with the informational (or "intentional") content of these representations. Thus the computational level specifies, for example, that early visual processes map representations of light intensities to representations of oriented lines ("edges").

Another claim made by the standard interpretation is that these specified visual information-processing tasks are the phenomena to be explained. In other words, the specification of the information-processing task 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, p. 41). *De facto*, most interpreters think that the real explanatory level is the algorithmic level where it is shown "how the brain performs this representational conversion" (Ramsey, p. 41). Ramsey continues: "In this three-tiered framework, it is the middle, algorithmic level where the CCTC theories attempt to explain the kinds of processes that account for mentality" (p. 42). In the last sentence Ramsey mentions classical theories (CCTC), but he adds: "This is the general form of cognitive science explananda, even for non-CCTC accounts like connectionism" (p. 41).

We agree that the phenomena to be explained are visual information-processing tasks, couched in intentional terms of input and output representations (i.e., edge-detection, shape-from-shading and so on). We also think that this specification itself is often not trivial and requires lengthy scientific investigation. We contend, however, that this intentional specification is not the job, or at least the main job, of CL. It is often made, at least to some extent, before we invoke CL at all. Using techniques such as single-cell recording, neuroscientists had discovered 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 long before Marr invoked his computational theories. We see no reason to call a specification of a task in terms of informational content of the inputs and outputs a “computational theory”. This would trivialize Marr’s notion of CL-level theory Indeed, those who hold the standard interpretation refrain from Marr’s label of computational theory. Thus Dennett who associates Marr's computational level with his intentional level, says that "this specification was at what he [Marr] called, misleadingly, the computational level" (Dennett, 1994, p. 681).[2] But, of course, the labeling would be misleading only if the job of computational level theories is providing such intentional descriptions of cognitive tasks. We will argue, however, that the job of CL goes far above and beyond that and that the standard interpretation misses what makes CL-level analysis distinctive.

[2] Sterelny (1990, p. 46), Ramsey (2007, pp. 41, note 43), and Horst (2009) make similar comments.

2.2. Frances Egan: Providing a mathematical or formal theory

Frances Egan associates Marr's CL with "the specification of *the function computed*" (Egan, 1991, pp. 196-197). She argues that CL provides 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" (Egan, 1995, p. 185). The aim of CL, on this view, is to specify the input-output mathematical function that the system computes (then the algorithmic levels specifies the algorithm by means of which the system computes this function, and the implementation level specifies how this algorithm is implemented in the brain). Thus, for example, the computational theory of early vision provides the mathematical formula $\nabla^2 G * I$ as the computational description of what the retina does. As Marr put it: "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).[3]

Proponents of the standard interpretation might agree with Egan that CL also provides a mathematical description of the computed function. Egan departs from the standard interpretation in two ways. One is her insistence that CL does *not* provide an intentional, information-processing, characterization of the input-output function. Egan (2010) cites Chomsky, who writes that when Marr talks about 'representation', it "is not to be understood relationally, as 'representation of'" (Chomsky, 1995, p. 53). What is being represented, according to Egan, is immaterial from a computational point of view: "*Qua* computational device, it does not matter that input values represent *light intensities* and output values the rate of change of

[3] The term I stands for a two-dimensional array ("the retinal image") of intensity values detected by the photoreceptors (which is the input). This image is convoluted (here signified by '$*$') through a filter $\nabla^2 G$, where G is a Gaussian and ∇^2 is a second-derivative (Laplacian) operator. This operation is arguably performed in the retinal ganglion cells.

light intensity. The computational theory characterizes the visual filter as a member of a well understood class of mathematical devices that have nothing essentially to do with the transduction of light" (Egan, 2010, p. 255). We invoke the representational content only *after* the computational-level theory has accomplished its task of specifying the mathematical function. The cognitive, intentional, characterization is what Egan terms a *gloss* on the mathematical characterization provided by the computational theory. This intentional characterization "forms a bridge between the abstract, mathematical characterization that constitutes the explanatory core of the theory and the intentionally characterized pre-theoretic explananda that define the theory's cognitive domain" (pp. 256-257).[4]

The other departure from the standard interpretation is mentioned in the last sentence cited. According to Egan, CL is a mathematical *theory* whose aim is *explanatory*. What it explains is the intentional, information-processing, characterization of the function that the visual system performs. Thus, Egan agrees with the standard interpretation as to the need for such an intentional, information-processing account. She contends, however, that this characterization is pre-theoretic and so does not constitute part of the computational theory. The computational theory, which consists solely of mathematical descriptions, aims to *explain* this pre-theoretic explananda. That the early visual system computes the $\nabla^2 G * I$ operations explains how it performs edge-detection. The explanation (presumably) is that the system detects edges by detecting the zero-crossings generated by the second-derivative filters $\nabla^2 G * I$ (where Gaussians are used at different scales).

We think that Egan captures very well the way Marr characterizes the *what* aspect of CL. The job of this element is to provide a precise specification of *what* the system

[4] Egan's main motivation here is avoiding a context-dependent individuation of computational states; see Shagrir (2001) for discussion.

does, and the precise specification of what the retina does is provided by the formula $\nabla^2 G * I$. However, Egan downplays the fact that there is another component to CL, namely, the *why* aspect. When Marr says "from a computational point of view, this is a precise specification of what the retina does," he refers to *what* the retina does, not to the *why*. After characterizing *what* early visual processes do, Marr says that "the term *edge* has a partly physical meaning – it makes us think of a real physical boundary, for example" (p. 68). And, he adds, "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" (p. 68). So it seems that Marr thinks that CL has to cover another aspect, beyond providing mathematical characterizations.

2.3. Piccinini and Craver: CL as a sketch of mechanism

In a recent paper Piccinini and Craver (2011) argue that it is best to conceive Marr's computational and algorithmic levels as *sketches* of mechanism. On the one hand, the two levels are *not* levels of mechanisms "because they do not describe component/sub-component relations" (p. 303). On the other hand, the two levels "constrain the range of components that can be in play and are constrained in turn by the available components" (p. 303). In this sense, of constraining, the computational and algorithmic levels are sketches. They are placeholders for structural components or sub-capacities in a mechanism. At the beginning of their paper, Piccinini and Craver say that a sketch of mechanism is a description in which some structural aspects (of the mechanism) are omitted. Once the missing aspects are filled in, the description turn into "a full-blown mechanistic explanation"; the sketches themselves can be thus seen as "elliptical or incomplete mechanistic explanations" (p. 284). They are, in a way, a guide or a first step towards the structural components that constitute the full-blown mechanistic explanations.[5]

---

[5] Kaplan (2011) advances a somewhat similar view arguing that computational models in neuroscience are explanatory to the extent that they are tied to the norms

We agree with Piccinini and Craver that CL puts constraints on the mechanistic explanation of the phenomenon. This seems to align with Marr's methodological approach (to be discussed below). But we reject their attempt to collapse Marr's three levels into two by closely intertwining the computational and algorithmic levels. Theirs is not unique among philosophical and theoretical approaches to cognitive science in attempting to collapse Marr's levels (see, e.g., Pylyshyn, 1984; Newell, 1980, both of whom collapse Marr's computational and algorithmic level before adding an addition semantic (Pylyshyn) or knowledge (Newell) level.) But this approach is foreign to Marr. If anything, it is the algorithmic and implementational levels that belong together as both look inside the mechanism to the operations that enable it to compute a function.[6] Piccinini and Craver are right to observe that both the computational and algorithmic levels are abstract, in that they omit certain structural aspects of the mechanism (both levels are also abstract in the sense that they provide mathematical or formal descriptions). But Marr is far keener to point to a fundamental difference between the computational and algorithmic levels. The algorithmic level (much like the implementation level) is directed to the *inner working* of the mechanism, i.e. to causal relations (signified by arrows) between sub-components.[7] The computational level looks *outside*, to identifying the

---

of mechanistic explanations. When referring to Marr, Kaplan argues that "according to Marr, the ultimate adequacy of these computational and algorithmic specifications as explanations of human vision is to be assessed in terms of how well they can be brought into registration with known details from neuroscience about their biological implementation" (p. 343).

[6]Thus Marr (1982) writes that "there must exist an additional level of understanding [i.e., CL] 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" (p. 19), and that "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" (p. 27).

[7] There are reasons to reject as well Piccinini and Craver's contention that the algorithmic level offers only a sketch of a mechanism. An algorithm can provide a complete account of the operations in a mechanism. In doing so it will not specify the parts of the mechanism, as that is the task of the implementation account, but

function computed and relating it to the environment in which the mechanism operates. Marr's former student and colleague, Shimon Ullman, puts this point about CL succinctly in his manuscript on visual motion: "In formulating the computational theory, a major portion concerns the discovery of the implicit assumptions utilized by the visual system. Briefly, these are valid assumptions about the environment that are incorporated into the computation" (Ullman, 1979)pp. 3-4). We will elaborate on this point below.

3. Recognizing what is distinctive about CL

We offer an alternative interpretation of Marr's CL that keeps equally in focus the *what* and *why* questions associated with it. Accordingly, we emphasize two aspects of Marr's CL. One is the quantitative characterization of the phenomena (associated with the *what*). The other is the role of contextual or environmental constraints (associated with the *why*). To make things more concrete we focus on one specific information-processing task – the correspondence problem in stereo vision. As we proceed, we identify respects in which our interpretation agrees and differs with the three interpretations above.

3.1.The correspondence problem

There is an angular discrepancy in the position of an object in the two retinal images. This discrepancy is known as disparity. The disparity is usually larger when the object is closer to the eyes (as in looking at a finger touching your nose) and smaller when it is further away. The visual system deploys disparity to compute several features such as 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

---

then the implementation account is also incomplete insofar as it fails to specify the operations the parts perform. Moreover, as Levy and Bechtel (2013) argue, it is often a virtue in explanation to abstract from details of the mechanism to reveal what is actually responsible for the phenomenon of interest.

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 the images and the multiple possibilities of matching elements.

Marr illustrates the ambiguity of elements in Figure 1. The four projections in the left eye's view, $L_1,\ldots,L_4$, can be paired in 16 possible ways with the four projections, $R_1,\ldots,R_4$, in the right eye's view, but only 4 are correct (filled circles). The remaining 12 (open circles) are false targets. The horizontal dashed lines signify the amount of (horizontal) disparity; circles (pairs) that are on the same line have the same disparity. Strikingly, the visual system solves the correspondence problem even in highly ambiguous scenes.

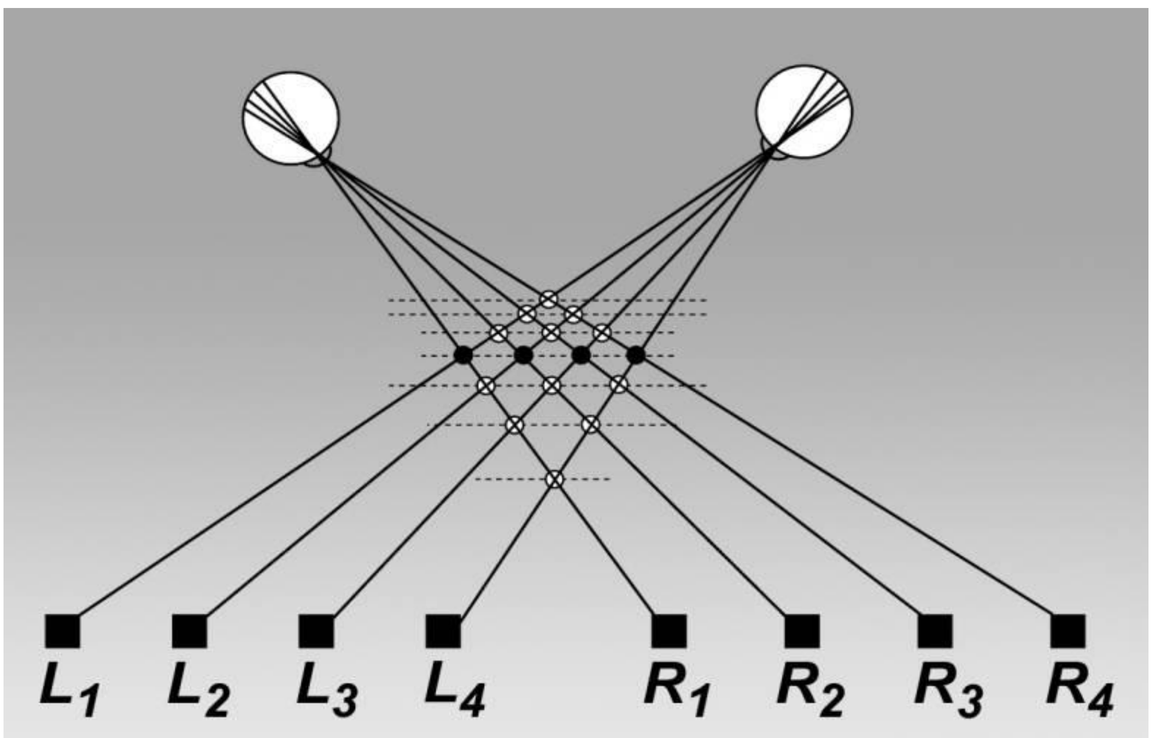

Figure 1. Marr's portrayal of the ambiguity in matching elements to determine the depth of an object.

According to the standard interpretation, characterizing the correspondence problem provides an intentional characterization of the input-output description of the task and exhausts the role of the computational level. The computational level states that the task at hand is the cognitive function whose input are elements from the left and right retinal images (say, edges, bars and so forth), and its output is some array of pairs of elements from the left and right images that correspond to the same worldly feature. With this characterization of CL, the standard interpretation would have researchers turn to the mechanistic levels of algorithms and implementations for the explanation. This, however, is not Marr's view. His computational level – both its *what* and *why* aspects – advance beyond this intentional description.

3.2. Specifying the task in quantitative terms (the *what*)

Let us start with the *what* aspect. Marr and Poggio (1976, 1979) provide a quantitative, mathematical description of the function solving the correspondence problem. This is a pairing function that satisfies two conditions: (a) Uniqueness: a black dot from one image can match no more than one black dot from the other image. This constraint rules out, for example, 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.

We see, then, that CL provides more than an intentional description of the phenomenon to be explained, i.e., matching elements from the left and right images. CL provides a quantitative characterization of this matching function: It specifies the (input-output) mathematical function that the system computes in order to reach matching. CL shows that the visual system solves the correspondence problem by computing a pairing function that satisfies the Uniqueness and Continuity constraints (in short: UC-pairing function). This role of CL is consistent with Egan's interpretation that highlights the centrality of a mathematical or formal theory. It is also consistent with Piccinini and Craver's claim that CL is a sketch of a mechanism. The computed, mathematical, function constrains the possible algorithms that the system might use, which are just the algorithms for a UC-pairing function (Marr and Poggio, 1979, propose a constraint-satisfaction attractor neural network). And the computational and algorithmic levels constrain the possible “full blown” mechanistic explanations that can be provided. However, both Egan, on the one hand, and Piccinini and Craver, on the other, do not notice that this quantitative characterization of the task is associated with the *what* aspect of CL: What is being computed is a UC-pairing function. This aspect, however, does not exhaust the role of the computational level. CL is also involved with embedding this function in the environment of the perceiving subject.

3.3.The role of the environment (the *why*)

Marr often emphasizes CL is involved with what he calls physical or natural constraints. As his students once put it, CL includes "an analysis of how properties of the physical world constrain how problems in vision are solved" (Hildreth & Ullman, 1989, p. 582). These physical constraints are features in the physical environment of the perceiving individual (1982, pp. 22-23); they are not features of the mechanism described abstractly. To avoid ambiguities with physical features of the inner implementing mechanisms we call them *contextual constraints*. It should be noticed that these constraints are *not* the informational contents of the representations, but facts about the physical environment we happen to live in.

In our case, Marr and Poggio relate the Uniqueness and Continuity conditions to contextual, environmental physical features. 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 specifies that "a given point on a physical surface has a unique position in space at any one time" (Marr and Poggio 1976, p. 284; see also Marr 1982, pp. 112-113). Continuity ("disparity varies smoothly almost everywhere") is motivated by the cohesiveness of matter constraint, 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, p. 284; see also Marr 1982, pp. 112-113).

What is the role of the contextual constraints in the analysis of vision, and of cognition more generally? We identify two related but different roles, one methodological and another explanatory. The methodological role has to do with the *discovery* of the computed function. The claim here is that we appeal to physical external factors in order to discover the mathematical function that is being computed. Thus, for example, we derive continuity ("contextual constraint") from the fact that the world around us consists of objects whose surfaces are by and large smooth; only a small fraction of the image is composed of features such as object boundaries that result in changes in depth. Thus overall disparity is mostly continuous. Or, returning to the example of edge-detection, the discovery that early visual processes compute derivation (either of first or second degree) is made through the observation that in our perceived environment sharp changes in light reflectance occur along physical edges such as boundaries of objects. This contextual feature puts substantial constraints on the mathematical function that is being computed, i.e., that it has to do with some form of derivation.

The methodological role of the physical constraints is related to a top-down methodology that is often associated with Marr's framework (that the scientific investigation should proceed from the top, computational, level, down to the algorithmic and implementation levels). A central claim of this approach is that it would be practically impossible to extract the computed mathematical function by abstracting from neural mechanisms. The way to go is to extract what the system computes from relevant cues in the physical world that constrain the computed function. The contextual constraints play a central role in this top-down approach.

The other role of the contextual constraints is explanatory (we note that on p. 22 Marr refers to CL as a "level of explanation"). This explanatory role of constraints is tied to the *why* aspect: The contextual constraints play the role of answering the question of why the computed mathematical function is appropriate for the given information-processing, visual, task. Thus consider again the correspondence problem. After characterizing the *what* (what is being computed is the UC-pairing function), Marr asks why the UC-pairing function – 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, p. 112; emphasis original). Marr is asking why should computing UC-pairing, and not any of the other functions, provide a solution for the correspondence problem. The algorithms and the neural mechanisms that underlie this function cannot answer to this question. These mechanisms specify *how* the system computes the UC-function, but they do not explain why computing this function, and not another function, lead to matching.

Marr explains why the UC-pairing function leads to matching by relating the conditions of uniqueness and continuity to facts about the physical world we happen to live in. Computing a UC-pairing function leads to matching because the UC-pairing function corresponds to spatial localization and the cohesiveness of matter in our world. Imagine a world consisting of objects with spiky surfaces that give rise to a reflection function that is almost never smooth. This will mean that the disparity between the images changes almost everywhere. In our example (fig. 1), the disparity between L1 and R1 is very different from the disparity between the between L2 and R2, and so on. In this world it might be impossible to find a function that satisfies continuity, and end even if there is such function there is no reason to assume that computing it will lead to matching. Had we lived in such a world, then computing this function would not lead to matching, but, if anything, to something else. Computing UC-pairing function is appropriate for matching in our case due to certain contingent facts about the physical environment in which we are embedded.

The methodological and explanatory roles of the constraints are related, of course. On the one hand, the contextual constraints explain, at least partly, the fact that the visual system computes the UC-function and not another function. On the other hand, Marr's methodological moral is that we can deploy these constraints in order to discover that the computed function is one satisfying the conditions of uniqueness and continuity.

4. Insights from Marr's CL for Mechanistic Explanation

Having articulated our account of Marr's CL level that sharply distinguishes it from the algorithmic and implementational levels and takes seriously his construal as involving both *what* and *why* aspects, we can return to mechanistic explanation. As we discussed above, Piccinini and Craver treated CL as offering a mechanism sketch. On our construal, CL is not providing a sketch of a mechanism but something quite different—it is characterizing the phenomenon for which a mechanism is sought as explanation. There is an important role for mechanism sketches in developing mechanistic explanations, but insofar as the sketch identifies operations in the mechanism it is an account at Marr's algorithmic level and insofar as it identifies these operations with component parts of the mechanism, it is at the implementational level. With respect to the mechanism, CL only specifies the task the mechanism performs and offers no suggestions to how it does it. Thus, it characterizes the phenomenon without trying to explain it mechanistically (although, as we have noted, it does figure in a different type of explanation, that concerned with why the mechanism is appropriate for the task).

Egan is correct to draw out the fact that CL offers mathematical characterizations of the task the mechanism is to perform. This is a crucial aspect of the way phenomena are delineated in scientific inquiries. If they weren't delineated mathematically, the quest for mechanistic explanation would often be unmanageably underdetermined. Many mechanisms can perform in qualitatively the same way, but quantitatively their performance differs. The challenge is to explain the actual phenomenon characterized quantitatively. This quantitative detail is also important to researchers as it provides a major tool for evaluating proposed mechanistic explanations. Of course the mechanistic explanation must also appeal to parts and operations that are known to constitute the mechanism. Yet, even when this condition is met, researchers find it important to assess whether the proposed mechanism could account for the quantitative details of the phenomenon. This is where computational modeling plays an increasingly central role in neuroscience and biology more generally (Bechtel, 2011).

As Egan correctly observes, the mathematical, quantitative, characterization (what she calls a mathematical or a formal theory) plays an explanatory role with respect to the pre-theoretic, intentionally characterized, explananda phenomenon. What Egan disregards, however, is that the mathematical theory has this explanatory role only if we embed the mechanism in the physical environment of the perceiving individual. The mathematical operation $\nabla^2 G * I$ is explanatory with respect to the phenomenon of edge-detection only when we relate this mathematical function with the relation that holds between magnitudes existing in the world. As Egan notes, correctly again (!), the informational content of the cells in the retina and in the primary visual cortex have no explanatory role in CL. They are, perhaps, only a *gloss* on the mathematical characterization that the computational theory provides. But this does not entail that there are no *other* contextual features that play an explanatory role. Indeed, according to Marr the relevant contextual features are physical ("contextual") constraints that indicate intensity changes in the image result from "surface discontinuities or from reflectance or illumination boundaries" (Marr & Hildreth, 1980, p. 187). The upshot is that the formal theory constitutes only one part of the explanation (associated with the *what*). "The other half of this level of explanation" (1982, p. 22), as Marr put it, has to do with the *why*, namely with why the visual system performs the mathematical operation $\nabla^2 G * I$, and not, say, exponentiation or factorization when detecting edges.

What makes CL explanatory with respect to edge-detection – so that the *what* and the *why* conspire to provide an explanation – is an intriguing question. One proposal is that the visual system works much like scientific models (for a survey, see Frigg &

Hartmann, 2006). It models its environment by preserving certain relations in the environment. CL describes this modeling relation, explaining its role in the visual task.[8] This is shown in Figure 2 in which the top portion identifies the relation in the world and the bottom portion the operations occurring on representations in the visual system. The dotted arrows indicate that the representations in the brain stand in for features in the world itself. The detection of visual edges (say, zero-crossing segments) mirrors a pertinent relation in the visual field in the sense that there is an isomorphism (or some other morphism) between this visual process and the visual field. This morphism is exemplified by the (alleged) fact that the visual system and the visual field have a shared mathematical description (or structure). On the one hand, the visual system computes the zero-crossings of second-derivative operations (over the retinal pixels) to detect edges; this is shown in the bottom span of figure 2. On the other hand, the reflection function in the visual field changes sharply along physical edges such as object boundaries. These changes can be described in terms of extreme points of first-derivatives or zero-crossing of second derivatives.

Figure 2. Edge-detection. Early visual processes (bottom span) detect "visual edges" in the retinal image by computing the zero-crossings of $\nabla^2 G*I$ (see note 3); the second-derivative operations $\nabla^2 G*I$ are performed by the ganglion and LGN cells. The intensity values encode (dashed arrow) "light intensities" of the visual field that combine different factors such as the reflectance of the visible surfaces. The visual edges (e.g., segments of co-located zero-crossings) encode physical edges such as object boundaries. .

Figure 2 makes clear how the CL accounts for edge detection: It is important to compute the function $\nabla^2 G*I$ because that is the relation that holds between

[8] This modeling idea is discussed in some detail by Shagrir (2010a).

magnitudes existing in the world: a mechanism that computes it will identify edges. This match between the task and the mechanism shows why the mechanism succeeds. The *what* aspect provides a description of the mathematical function that is being computed. The *why* aspect employs the contextual constraints in order to show how this function matches with the environment.

There are debates about whether the matching relation in models is similarity, isomorphism, partial isomorphism or homomorphism.[9] And, of course, not all mechanisms are perfectly adapted to their environments. There is a long tradition of showing that cognitive systems with limited resources employ heuristics that succeeded well enough in the actual world, but which can be expected to fail under specifiable conditions (Simon, 1996). Our proposal, though, works for heuristics as well as optimal procedures—heuristics work as well as they do because they capture real relations in the world (between cues and outcomes). The why-aspect of CL accounts does not require showing that the computational performed is optimal, only that it is grounded in the world in which the visual system is operating.[10]

What are the relations between CL explanations and mechanistic explanations? On the one hand, it is important to recognize that the task to be performed is conceptually independent of any mechanism that performs it, including the particular inputs the organism receives or the specific outputs it produces in solving it. While Marr viewed the algorithm he took to be operative in our brains as computing $\nabla^2 G * I$, computing that function would still be a task for a perceptual

[9] Swoyer (1991) talks about isomorphism, but others about partial isomorphism (French & Ladyman, 1999; Da Costa & French, 2003; Bueno & French, 2011), homomorphism (Bartels, 2006) and similarity (Giere, 2004; Weisberg, 2013).
[10] Edge-detection is by no means an isolated example of this kind of CL explanation. Shagrir (2010b) discusses the case of stereovision. Another example is Ullman's (1979) structure-from-motion algorithm in which the 3-D structure and motion of objects is inferred from the 2-D transformations of their projected images. Here the mathematical function computed reflects spatial relations in the target, assuming the constraint of rigidity.

system even if our brains failed to do so. By actually computing it, our brains solve a problem that is specified by the relation between light intensities and physical edges occurring in the world, as it is clearly shown in Figure 2.

On the other hand, Marr does not offer CL as an alternative explanation to mechanistic explanations, but as a *complementary* one. Mechanistic explanations describe the mechanisms by means of which the nervous system changes from one neural state to another. It describes, for example, how certain activity in the photoreceptors (that represent light intensities) lead, through the activity of the retinal ganglion cells, to the activation of cells in V1 (that are sensitive to oriented lines). This mechanistic description is surely an explanation at the level of neural circuitry. But it does not by itself explain the information-processing task of edge detection (This is perhaps what Marr means when he says: "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, p. 15)). This mechanistic description does not explain why this particular neural mechanism has to do with the detection of edges and not, say, with the detection of color. The CL provides the answer to this question: The mechanism implements a certain mathematical function (of the zero-crossings of $\nabla^2 G * I$) and this function matches the relations in the world, e.g., sharp changes in light intensities that typically occur along object boundaries. When the CL explanation is in place, the mechanistic – algorithmic and implementational – descriptions *explain* how exactly the visual system computes the mathematical function.

While one might accept our contention that Marr's CL accounts require turning to the world to address both the *what* and *why* aspects, one might still question whether there is a similar need to look outside a mechanism to its context in delineating the phenomenon it explains. Isn't it sufficient to show that the targeted mechanism exhibits regular behavior? We offer two responses to this question.

First, as we noted at the beginning, not all regularities that can be stated mathematically are appropriate targets for explanation. This applies both to naturally occurring ones and to ones that can be detected experimentally. Looking to the task that needs to be performed by the mechanism given the structure of the world provides a way of identifying which regularities require mechanistic explanation. Contrasting examples illustrates this. Although the heat generated by animals can be quantified, the hundred-year effort to explain animal heat terminated quickly when around 1930 it was recognized that heat was a waste product, not a source of energy that animals could use to perform work. The identification that instead, adenosine triphosphate (ATP) was the molecule in which energy released through metabolism was stored, resulted in extensive research to explain how, for example oxidative metabolism could result in synthesis of three ATPs. As these examples make clear, looking to the environment is important is important in mechanistic research in general, but it is especially relevant in the context of information-processing mechanisms where the task being performed is an important guide to what operations carry information needed for the mechanism to perform its task.

Second, it is the world that both sets the task and determines the resources available to the mechanism in performing the task. Part of the investigatory strategy researchers employ in developing mechanistic explanations is to identify these resources and their utilization within the mechanism. Mechanisms are typically not closed systems but consist of parts and operations that interact with their environment in generating phenomena. The visual system is an example. Although Marr and many other vision researchers focused only on the steps in processing stimuli and not the activities of the organism that determines what stimuli impact its retina, perceivers are often active—they move their eyes, heads, or whole bodies in the course of seeing. As they do so, the projections onto their retina change. Moreover, some of these movements are directed by the visual system as it actively samples the visual array to procure information (Ballard, 1991). Since many mechanisms actively engage their environment as they operate, it is important to capture these interactions in characterizing the phenomenon itself. Otherwise, researchers end up trying to explain a phenomenon that does not actually occur and may require resources that are not available in the mechanism. This concern is what lay behind calls for ecological validity in psychology research by Brunswik (1943), J. J. Gibson (1979), Neisser (1976), and others. (We discuss Gibson and Marr's response to Gibson further in the following section.)

In this section we have focused on two important insights that can be gleaned for the task of delineating phenomena for mechanistic explanation. The first is that phenomena are typically characterized not just qualitatively, as they typically are in the accounts of the new mechanistic philosophers of science, but also in quantitative or formal terms (for recent exceptions, see Bechtel, 2013; Bechtel & Abrahamsen, 2010; Brigandt, 2013; Kaplan, 2011). In describing the talk of edge detection in his CL account, Marr identified the mathematical function that needed to be computed. Second, in delineating phenomena researchers often, as Marr did at the CL level, focus outwards on the context in which the mechanism operates. Among other things, this allows researchers to identify the resources available to the mechanism in producing the phenomenon. We will return to show how Marr's account generalizes to other phenomena beyond vision in the concluding section, but first point to two limitations of the account Marr offered.

5. Delineating Phenomena: Going Beyond Marr's Account of CL

As much as Marr emphasized the importance of developing an analysis of CL that showed both quantitative rigor and addressed the context in the world in which the visual mechanism operated, it is noteworthy that he did not develop two other aspects of the CL account that are critical in delineating phenomena—that empirical, even experimental, research is required to identify the quantitative relations that constitute the phenomena and that characterizations of phenomena are often revised in the course of developing mechanistic explanations of them.

5.1. Empirical Inquiry to Delineate Phenomena

Despite the attention Marr paid to CL, he pursued CL accounts with an intuitive, almost arm-chair approach. Poggio (1981), in articulating and defending Marr's approach and bringing out clearly how CL analysis is directed at the world outside the visual systems, nonetheless also claims: "No high-level specific preunderstanding is required, but only general knowledge about the physical world. An example of such general knowledge is that the world is constituted mainly of solid, non-deformable objects of which only one can occupy a given point in space and time." (p. 259). He also notes "It is probably fair to say that most physiologists and students of psychophysics have often approached a specific problem in visual perception with their personal 'computational' prejudices about the goal of the system and why it does what it does." This almost trivializes the importance of CL analysis. But we contend that Marr did, or should have, intended something more radical.

We take a cue as to what CL analysis ought to involve from Gibson, of whom Marr said: "In perception, perhaps the nearest anyone came to the level of computational theory was Gibson." (1982, p. 29). The basis for this comment is that Gibson more than most psychologists took seriously the importance of the environment in which perception occurs. Although he adopted the biological term *ecological*, his principle focus was on the physical features of the environment (specifically, those physical features about which information is available in the light). Much of Marr's discussion of Gibson is critical, focusing on Gibson's repudiation of representations and internal processing (Gibson claimed that vision was *direct*—we directly see objects in the world by *picking up* information available in the light). At the same period as Marr was writing *Vision*, Ullman published a detailed criticism of Gibson's account of direct perception (Ullman, 1980). We focus, however, on why Marr saw Gibson as the person who came closest to offering a computational theory.[11]

---

[11] Gibson would have bristled at being associated with anything called a *computational theory* and even more to Marr's advocacy of analyzing vision in terms

What an ecological approach to perception meant for Gibson and many who have subsequently pursued his project is that psychologists should study the perceiving organism in the context of the world in which it functions, considering both what the organism uses vision for and the resources the world provides for performing those tasks. Both require empirical inquiry. Studying perceiving organisms reveals that they use vision to accomplish biological needs—detect resources and threats in their environments and safely navigate through it. Often these tasks can be performed by picking up on information in the environment without having to build up a complete representation of the world (by converting 2D representations into 3D representations).[12]

Gibson referred to what an organism picks up through vision or other senses as affordances: "The affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or ill" (J. J. Gibson, 1979). In particular, they are possibilities for action in the world that are relative to the organism and its capacities for acting. An example he used is that a surface that is nearly horizontal and flat, sufficiently extended, rigid relative to the weight of the animal, affords support that can be walked or run on. Moreover, he stressed that these potentials exist regardless of whether the organism has learned to pick up on them. (Gibson was a pioneer in treating perception as a skill to be learned; see J. J. Gibson & Gibson, 1955. ; This topic that became the focus of Eleanor Gibson's research; see E. J.

---

of algorithms. It is possible, however, to view Gibson's arguments for direct perception and his eschewal of internal processing as methodological—as a strategy for focusing on the richness of what he called "information in the light" that was neglected by most psychologists who jumped too quickly to address how organisms process stimuli that they have designed to probe the visual system with less attention to how such stimuli reflect the inputs the visual system typically confronts. In this he is allied with Marr's contention of the importance of CL analysis.

[12] Such inquiry has been pursued subsequently by, for example, Turvey (1992).

Gibson, 1969.) When the objects of vision are other agents, vision captures emotional information and presents others as entities to engage, fight, flee from, etc. Gibson maintained that these affordances were not in the organism but in the world, although they might only be relevant to organisms with particular capacities for action and so “picked up” by them.

In identifying affordances the perceiver is typically not passive but moves about in the world, and even when not moving physically, moves its eyes to focus on different parts of the visual field. As we noted in the previous section, once one recognizes that perceivers move to acquire information, it is not sufficient to characterize the input they use when functioning in their environment in terms of retinal images. Rather, it is better to focus on what Gibson termed the “optic array”—the pattern found in the light that changes as either the perceiver changes vantage points or objects in the world move. Among other things, the optic array provides information as to how the perceiver and perceived objects are situated vis-á-vis each other.

Gibson initiated a research program that has provided substantial information about the information in the optic array. Lee and Reddish (1981), for example showed that a parameter τ, easily calculated from the rate of expansion in the optic array, specifies time to impact even for accelerating agents such as gannets diving into the ocean. By experimentally manipulating the size of doorways, Warren and Whang (1987) showed that the optic array carried information about whether a person could simply walk through or whether they would have to turn sideways. An important offshoot of Gibson’s research are investigations such as those of Findlay and Gilchrist (2003) into how agents determine appropriate eye movements (saccades) to secure useful information.

From the perspective of attempts to explain the information processing involved in vision, these inquiries are all CL inquiries. But, contrary to Poggio, they reveal information about how the visual system is situated in the body and world that was not part of general knowledge but stemmed from empirical investigations. Although we have not emphasized it, the results of these inquiries into the world in which vision operates can be stated in a precise, quantitative manner.

5.2. Reconstituting Computational Level Accounts

A standard picture of scientific inquiry is that researchers begin with a problem to be solved such as a phenomenon to be explained and their efforts are then directed at solving the problem or explaining the phenomenon. But as we are all aware, attempts at solving a problem often leads to recognition that the problem was somewhat different from what it was initially taken to be. Likewise, efforts at explaining a phenomenon by studying the mechanism can lead scientists to recognize that the phenomenon is different than they took it to be (Craver, 2007, p. 261). One of the most important developments in the analysis of vision since Marr has been the discovery that there are two streams of visual processing beyond V1: the ventral stream projects to areas in the medial temporal lobe while the dorsal stream projects to areas in the parietal lobe. In their paper identifying these pathways, Mishkin, Ungerleider, and Macko (1983) characterized them as involved in respectively determining the identity of an object and its location. Subsequently, Milner and Goodale (1995) offered evidence to support the claim that the dorsal stream serves to identify possibilities for action. These two streams, however, are not fully independent as there are connections at several points between them (van Essen & Gallant, 1994) and, as Findlay and Gilchrist (2003, Chapter 1) discuss, areas such as the frontal eye fields, critical in regulating saccades, receives inputs from both. These discoveries revealed that there are at least two components of the phenomenon of vision that were not differentiated prior to research on the responsible mechanism.

Even the characterization of the object recognition process on which Marr focused has been significantly revised in recent years. Although the fact that there are at least as many and likely many more recurrent as feed-forward projection through cortex has been known since the pioneering research of Lorente de Nó (1938), there was little understanding of what function these might serve. References to top-down processing were frequent, especially in cognitive science, during the period in which Marr was working, but he was highly skeptical of them since they seemed incompatible with the fact that we often see what we don't expect to see. But evidence of the prevalence of recurrent activity in the brain has continued to grow and recently a number of researchers have developed accounts that accommodate it (Dayan, Hinton, Neal, & Zemel, 1995; Rao & Ballard, 1999; Llinás, 2001; Hawkins & Blakeslee, 2004; Hohwy, Roepstorff, & Friston, 2008; Huang & Rao, 2011; Clark, in press). They have recast the phenomenon of vision as starting with the brain predicting what it will next encounter through its senses and only engaging in further processing of input information when it contravenes what it predicted. Through the combination of empirical and conceptual research, the phenomenon of vision on which Marr focused is being reconstituted.

Marr was right to emphasize both the *what* and *why* elements of CL, but he did not go far enough in exploring how these are to be identified. Empirical investigations conducted at the point at which the mechanism engages its environment are required to determine what are the stimuli to which the perceiver is responding and, although we have not addressed it, the uses to which the perceiver puts the information. Moreover, the CL account is not final when investigation of the mechanism begins but often must be revised in light of what is discovered by the mechanism itself.

6. Conclusion

Our goal in this paper has been to develop a characterization of CL that is more adequate to Marr's insistence that it involves both a *what* and a *why* aspect than extant interpretations. The *what* aspect requires developing a mathematical description of the task for vision. The *why* aspect forces researchers to look to the structures in the world that the organism engages through its visual system. It shows that the function computed by the visual system is effective because it matches a mathematical relation that exists in the world (e.g., between light intensities and physical edges). We argued, however, that Marr did not go far enough either in recognizing that empirical inquiry such as Gibson pursued is often required to identify the task confronted by the visual system or that the characterization of the task must often be revised as research on the mechanism proceeds. The CL analysis, so construed, identifies the phenomena of vision—the visual system processes information provided by light so as to compute functions that correspond to those realized in the physical world, thereby enabling organisms to perform their activities.

Following Marr, we have focused on the visual system and thus discussed CL analyses of visual information available to organisms. But as we indicated at the outset, this perspective can be extended to other brain systems. The most straightforward extensions are to other sensory systems and motor systems that compute functions that relate directly to structures in the environment. Motor systems must compute commands that enable the body to operate in the environment, including changes in the environment that result from the execution of the motor processes. It is by looking to the environment that researchers can identify the function that the motor system must compute. A nice example involves the oculomotor system that controls eye movements. One of its tasks (performed by the *vestibulo-ocular reflex* or VOR) is keeping the visual world stable on the retina when the head is moving. Experimental studies show that the system converts transient eye-velocity-encoded inputs into persistent eye-position-encoded outputs. It was thus concluded that the system network is a *neural integrator*.[13] In this case

[13] It is hypothesized that the neural integrator also serves for other eye-movement operations such as saccadic and pursuit movements (Robinson, 1989; Goldman, Kaneko, Major, Aksay, Tank, & Seung, 2002).

the researchers infer from contextual cues ("contextual constraint") that the relations between the *encoded* velocity and position are that of integration to the claim mathematical integration is what is computed (Shagrir, 2012).

The challenge in characterizing CL analysis is somewhat greater for more central cognitive activities such as episodic memory. Following Ebbinghaus, memory has often been studied using laboratory tasks such as learning lists of words that are relatively far removed from those humans typically confront. Inspired by Gibson, researchers such as Neisser (1982; Neisser & Winograd, 1988) investigated real stimuli and real tasks (e.g., providing testimony in legal proceedings). One of the upshots of this endeavor was to demonstrate how reconstructive memory is (a claim that has been pursued by other researchers as well; see, e.g., Schacter, 1996). What makes it reconstructive is that, in the process of recall, pieces of information that are retrieved are organized together in ways that are at least partly responsive to the context in which retrieval is required. This points to the retrieval context as partly shaping the task of memory recall. It is much more challenging to characterize memory retrieval in terms of a mathematical function, and this may be one of the reasons why research on the mechanisms of episodic memory is less advanced than the research on the mechanisms of vision.

The information processing perspective applies more generally than just to brain function. Biological systems often employ systems that control other systems. At the cellular level, this is carried out chemically through cell signaling system. In single-celled organisms, which are the most prevalent life forms on the planet, molecular systems pick up information about the internal state or conditions in the environment of the cell, and regulate such activities as the synthesis of new proteins. In characterizing these phenomena, both the *what* and *why* aspects of Marr's CL level are appropriate: researchers both specify the relationship between the signal picked up and the response generated mathematically, and relate this to conditions external to the control system. This outward focus is important, as it is in vision, to specifying which mathematical relations constitute the phenomena to be explained and the resources available to the system in generating the phenomena.

We have limited our discussion in this paper to information processing contexts. But we think that Marr's account of CL provides insights into the tasks confronted in delineating phenomena and can help fill a lacuna in the accounts the new mechanists in philosophy of science have offered of the task of delineating phenomena. For example, mechanistic explanations are also advanced for phenomena such as protein synthesis and the generation of action potentials that do not themselves serve to process information. Developing detailed accounts of the phenomena and the contexts in which they are performed is also vitally important in those endeavors. Hence, some of the lessons derived from the CL analysis may extend to these explanations. However, since these explanations do not involve processing information, the distinctive why feature of CL analysis which we have emphasized does not apply. Our contention is that Marr's valuable insight is that with information processing mechanisms, the CL level plays a crucial role in identifying the relation in the world that the information processing system must compute in order to succeed. Moreover, we have argued that without a CL analysis, the quest for mechanism would be impaired and a crucial part of the explanation would be unavailable.

References

Ballard, D. H. (1991). Animate vision. *Artificial Intelligence, 48*, 57-86.

Bartels, A. (2006). Defending the structural concept of representation. *Theoria-Revista De Teoria Historia Y Fundamentos De La Ciencia, 21*, 7-19.

Beadle, G. W., & Tatum, E. L. (1941). Genetic control of biochemical reactions in *Neurospora*. *Proc Natl Acad Sci U S A, 27*, 499-506.

Bechtel, W. (2011). Mechanism and biological explanation. *Philosophy of Science, 78*, 533-557.

Bechtel, W. (2013). Understanding biological mechanisms: Using illustrations from circadian rhythm research. In K. Kampourakis (Ed.), *The Philosophy of Biology* (Vol. 1, pp. 487-510): Springer Netherlands.

Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences, 36*, 421-441.

Bechtel, W., & Abrahamsen, A. (2010). Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part A, 41*, 321-333.

Bechtel, W., & Richardson, R. C. (1993/2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge, MA: MIT Press. 1993 edition published by Princeton University Press.

Bermúdez, J. L. (2005). *Philosophy of psychology : a contemporary introduction*. New York: Routledge.

Bogen, J., & Woodward, J. (1988). Saving the phenomena. *Philosophical Review, 97*, 303-352.

Brigandt, I. (2013). Systems biology and the integration of mechanistic explanation and mathematical explanation. *Stud Hist Philos Biol Biomed Sci*.

Brunswik, E. (1943). Organismc achievement and enviornmental probability. *The Psychological Review, 50*, 255-272.

Bueno, O., & French, S. (2011). How theories represent. *British Journal for the Philosophy of Science, 62*, 857-894.

Chomsky, N. (1995). Language and nature. *Mind, 104*, 1-61.

Clark, A. (in press). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*.

Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. New York: Oxford University Press.

Da Costa, N. C. A., & French, S. (2003). *Science and partial truth: A unitary approach to models and scientific reasoning*. Oxford: Oxford University Press.

Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz machine. *Neural Computation, 7*, 889-904.

Dennett, D. C. (1994). Cognitive science as reverse engineering: Several meanings of 'top-down' and 'bottom-up'. In D. Prawitz, B. Skyrms & D. Westerstahl (Eds.), *Logic, metholodology and philospohy of science IX*. Amsterdam: Elsevier.

Egan, F. (1991). Must psychology be individualistic? *Philosophical Review, 100*, 179-203.

Egan, F. (1995). Computation and content. *The Philosophical Review, 104*, 181-203.

Egan, F. (2010). Computational models: a modest role for content. *Studies in History and Philosophy of Science, 41*, 253-259.

Findlay, J. M., & Gilchrist, I. D. (2003). *Active vision : the psychology of looking and seeing*. Oxford ; New York: Oxford University Press.

French, S., & Ladyman, J. (1999). Reinflating the semantic approach. *International Studies in the Philosophy of Science, 13*, 103-121.

Frigg, R., & Hartmann, S. (2006). Models in science. In E. N. Zalta (Ed.), *Encyclopedia of philosophy*: <plato.stanford.edu/archives/spr2006/entries/models-science/>.

Gibson, E. J. (1969). *Principles of perceptual learning and development*. New York,: Appleton-Century-Crofts.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Gibson, J. J., & Gibson, E. J. (1955). Perceptual learning; differentiation or enrichment? *Psychological Review, 62*, 32-41.
Giere, R. N. (2004). How models are used to represent reality. *Philosophy of SCience, 71*, 742-752.
Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science, 69*, S342-S353.
Goldman, M. S., Kaneko, C. R., Major, G., Aksay, E., Tank, D. W., & Seung, H. S. (2002). Linear regression of eye velocity on eye position and head velocity suggests a common oculomotor neural integrator. *Journal of Neurophysiology, 88*, 659-665.
Griffiths, P. E. (2001). Genetic Information: A Metaphor in Search of a Theory. *Philosophy of Science, 68*, 394-412.
Hawkins, J., & Blakeslee, S. (2004). *On intelligence*. New York: Times Books.
Hildreth, E. C., & Ullman, S. (1989). The computational study of vision. *Foundations of cognitive science* (pp. 581-630): MIT Press.
Hohwy, J., Roepstorff, A., & Friston, K. (2008). Predictive coding explains binocular rivalry: An epistemological review. *Cognition, 108*, 687-701.
Horst, S. W. (2009). The computational theory of mind. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*: <plato.stanford.edu/entries/comuptational-mind/>.
Huang, Y., & Rao, R. P. N. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science, 2*, 580-593.
Kaplan, D. M. (2011). Explanation and description in computational neuroscience. *Synthese, 183*, 339-373.
Kauffman, S. A. (1971). Articulation of parts explanation in biology and the rational search for them. In R. C. Bluck & R. S. Cohen (Eds.), *PSA 1970* (pp. 257-272). Dordrecht: Reidel.
Lee, D., & Redish, P. E. (1981). Plummeting gannets: A paradigm of ecological optics. *Nature, 293*, 293-294.
Levy, A. (2011). Information in biology: A fictionalist account1. *Noûs, 45*, 640-657.
Levy, A., & Bechtel, W. (2013). Abstraction and the organization of mechanisms. *Philosophy of Science, 80*, 241-261.
Llinás, R. R. (2001). *I of the vortex: From neurons to self*. Cambridge, MA: MIT Press.
Lorente de Nó, R. (1938). Analysis of the activity of the chains of internuncial neurons. *Journal of Neurophysiology, 1*, 207-244.
Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science, 67*, 1-25.
Marr, D. C. (1977). Artificial intelligence - personal view. *Artificial Intelligence, 9*, 37-48.
Marr, D. C. (1982). *Vision: A computation investigation into the human representational system and processing of visual information*. San Francisco: Freeman.
Marr, D. C., & Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London. Series B. Biological Sciences, 207*, 187-217.
Marr, D. C., & Poggio, T. (1976). Cooperative computation of stereo disparity. *Science, 194*, 283-287.

Marr, D. C., & Poggio, T. (1979). A Computational Theory of Human Stereo Vision. *Proceedings of the Royal Society of London. Series B. Biological Sciences, 204*, 301-328.
Maynard Smith, J. (2000). The concept of information in biology. *Philosophy of Science, 67*, 177-194.
Milner, A. D., & Goodale, M. G. (1995). *The visual brain in action*. Oxford: Oxford University Press.
Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in Neurosciences, 6*, 414-417.
Neisser, U. (1976). *Cognition and reality: Principles and implications of cognitive psychology*. San Francisco: W. H. Freeman.
Neisser, U. (1982). *Memory observed: Remembering in natural contexts*. San Francisco: W. H. Freeman.
Neisser, U., & Winograd, E. (1988). *Remembering reconsidered : ecological and traditional approaches to the study of memory*. Cambridge England ; New York: Cambridge University Press.
Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4*, 135-183.
Piccinini, G., & Craver, C. (2011). Integrating psychology and neuroscience: functional analyses as mechanism sketches. *Synthese, 183*, 283-311.
Poggio, T. (1981). Marr's computational approach to vision. *Trends in Neurosciences, 4*, 258-262.
Pylyshyn, Z. W. (1984). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press.
Ramsey, W. (2007). *Representation reconsidered*. Cambridge, UK New York: Cambridge University Press.
Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci, 2*, 79-87.
Robinson, D. A. (1989). Integrating with neurons. *Annual Review of Neuroscience, 12*, 33-45.
Schacter, D. L. (1996). *Searching for memory: The brain, the mind, and the past*. New York: Basic Books.
Shagrir, O. (2001). Content, computation and externalism. *Mind, 110*, 369-400.
Shagrir, O. (2010a). Brains as analog-model computers. *Studies In History and Philosophy of Science Part A, 41*, 271-279.
Shagrir, O. (2010b). Marr on computational-level theories. *Philosophy of Science, 77*, 477-500.
Shagrir, O. (2012). Structural representations and the brain. *The British Journal for the Philosophy of Science*.
Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*, 379-423, 623-656.
Shapiro, J. A. (2011). *Evolution: A view from the 21st century*. Upper Saddle River, NJ: FT Press Science.
Simon, H. A. (1996). *The sciences of the artificial* (Third ed.). Cambridge, MA: MIT Press.

Sterelny, K. (1990). *The representational theory of mind : an introduction*. Oxford, OX, UK ; Cambridge, Mass., USA: B. Blackwell.

Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese, 87*, 449-508.

Turvey, M. T. (1992). Ecological foundations of cognition: Invariants of perception and action. In H. L. Pick, P. van den Broek & D. C. Knill (Eds.), *Cognition: Conceptual and methodological issues* (pp. 85-117). Washington, DC: American Psychological Association.

Ullman, S. (1979). *The interpretation of visual motion*. Cambridge, Mass.: MIT Press.

Ullman, S. (1980). Against direct perception. *The Behavioral and Brain Sciences, 3*, 373-415.

van Essen, D. C., & Gallant, J. L. (1994). Neural mechanisms of form and motion processing in the primate visual system. *Neuron, 13*, 1-10.

Warren, W. H., Jr., & Whang, S. (1987). Visual guidance of walking through apertures: body-scaled information for affordances. *Journal of experimental psychology. Human perception and performance, 13*, 371-383.

Watson, J. D., & Crick, F. H. C. (1953). Genetical implications of the structure of deoxyribonucleic acid. *Nature, 171*, 964-967.

Weisberg, M. (2013). *Simulation and similarity: Using models to understand the world*. New York: Oxford.