

# Leveling the Field: Talking Levels in Cognitive Science

**Luke Kersten (lukekersten@email.carleton.ca)**  
Institute of Cognitive Science, 1125 Colonel By Drive,  
Ottawa, ON K1S 5B6 CA

**Robert West (robert\_west@carleton.ca)**  
Institute of Cognitive Science, 1125 Colonel By Drive,  
Ottawa, ON K1S 5B6 CA

**Andrew Brook (andrew.brook@carleton.ca)**  
Institute of Cognitive Science, 3A40 Paterson Hall,  
Ottawa, ON K1S 5B6 CA

## Abstract

Talk of levels is everywhere in cognitive science. Whether it is in terms of adjudicating longstanding debates or motivating foundational concepts, one cannot go far without hearing about the need to talk at different ‘levels’. Yet in spite of its widespread application and use, the concept of levels has received little sustained attention within cognitive science. This paper provides an analysis of the various ways the notion of levels has been deployed within cognitive science. The paper begins by introducing and motivating discussion via four representative accounts of levels. It then turns to outlining and relating the four accounts using two dimensions of comparison. The result is the creation of a conceptual framework that maps the logical space of levels talk, which offers an important step toward making sense of levels talk within cognitive science

**Keywords:** levels; analysis; explanation; organization; cognitive science; conceptual framework

## Introduction

Levels are everywhere in cognitive science. One cannot go far without hearing about the need to talk at different levels, whether it is in terms of levels of organization, explanation, description or aggregation. Levels have been called upon to do everything from adjudicate longstanding debates in cognitive modeling (Broadbent, 1985; Dawson, 1998) to motivating computational theories of mind (Pylyshyn, 1984).

Yet in spite of its widespread application and use, the concept of levels has received little sustained attention in cognitive science. Although there has been excellent wider discussion in the philosophical literature about levels (see, e.g., Wimsatt, 1976; 1994), little has been done to develop the notion specifically within cognitive science – for early attempts see Bechtel (1994) and McClamrock (1991). As Wright and Bechtel (2007) point out: “levels-talk is virtually threadbare from overuse yet [the] various conceptions of levels are rarely analyzed in any sustained, substantive detail despite there being a litany of literature on the subject” (p.55). Or, as Craver puts it: “Despite the ubiquity of levels talk in contemporary science and philosophy, very

little has been done to clarify the notion” (2015, p.23).

This state of affairs calls out for improvement. In what follows, we aim to provide an analysis of the various ways levels have been deployed within cognitive science. The paper begins by introducing and motivating discussion via four representative accounts of levels. It then turns to outlining and relating the four accounts using two dimensions of comparison. The result is the creation of a conceptual framework that maps the logical space of levels talk. This conceptual framework offers one important step toward making sense of levels within cognitive science.

## Talking Levels

The first account to consider is David Marr’s (1977, 1982). Marr identifies three different “levels of analysis” for cognitive science.

First, there is the computational level. At the computational level, investigators look at what function a system performs, asking questions about what information-processing problem the system solves. Research at the computational level aims to translate general, everyday descriptions of cognitive phenomena into particular information-processing problems or tasks. Second, there is the algorithmic level. At the algorithmic level, researchers investigate by what steps a system solves an information-processing problem; they ask questions about the algorithms and representations used by the system. Research at this level attempts to specify in detail the set of information-processing procedures that solve a particular information-processing problem. Finally, there is the implementational level. At the implementational level, researchers attempt to determine what physical structures instantiate the algorithms used for solving the information-processing problem; what physical mechanisms realize or support the cognitive system under investigation.

For example, when investigators emphasize the spatiotemporal properties of cognitive systems using the techniques of neuroscience, such as neuroimaging or lesion studies, they operate at an implementational level. When investigators focus on how components interact so as to

produce the operations or procedures that carry out computations using the methods and techniques of cognitive psychology such as randomized block designs or error rates, they ascend to a higher, algorithmic level of analysis. Each level of analysis has a different role to play in the cognitive investigation, as each satisfies a different epistemic end.

The second account to consider is Zenon Pylyshyn's. Pylyshyn (1980, 1984) similarly identifies three "levels of description" for cognitive science, though where Marr is concerned with offering a unifying methodology for cognitive investigation Pylyshyn is concerned with articulating the foundations of cognitive science.

First, there is what Pylyshyn identifies as the "semantic" level. At this level, psychological behaviour is described and explained using the representational or semantic content of an individual's mental states; what the individual believes, desires, etc. Second, there is the "symbolic or syntactic" level. Behaviour at the symbolic level is described in terms of functional properties. For example, to explain why an individual does poorly on a memory-recall task, the symbolic level appeals to control structures or memory storage capacities. This is in contrast to explaining performance by referencing what information the individual fails to recall, for example. Third, there is the "biological level". At the biological level, individuals are described in terms of the familiar vocabulary of the physical sciences, e.g., neurology, chemistry, biology. To account for why someone is thinking of faces rather than sounds, for example, explanations at the biological level might appeal to differences in brain activity or neural-chemistry.

Similar to Marr, Pylyshyn (1980, 1984) is also impressed by the explanatory power of levels. However, whereas Marr's levels are concerned with answering different types of questions, Pylyshyn's levels are focused on capturing distinct generalizations. For Pylyshyn, different descriptive vocabularies capture distinct sets of regularities. He writes, for instance: "When we have principles of operation that cannot be stated within a certain vocabulary – but can be captured in another, more abstract (here, functional) vocabulary to which the terms of the first vocabulary stand in multiple relations – we have a *prima facie* case for the existence of an independent level" (1984, p.33). Some valid generalizations are only expressible at particular levels of description.

Newell (1980, 1990) offers a third account. Newell defines what he calls "systems levels". System levels are collections of components that, in virtue of their organization and interaction, produce particular functions or behaviours. System levels mark functional divisions between different sets of organized components. They define the basic "technology" by which the human cognitive architecture is constructed. Each system level is realized by components at the next system level below.

Furthermore, since complicated systems are more likely to be resistant to degradation if they are built out of assemblies of stable subcomponents, it is more likely that system levels will be successively layered. As Newell puts

it: "If stable subassemblies are created, layer upon layer, then each one has a reasonable probability of being constructed out of a few parts. Thus, there exists a general argument that stability dictates the existence of levels" (1990, p.117).

In terms of identifying system levels, Newell claims that as one ascends upward from the smallest components to the largest components, system levels are revealed by their unique time signatures; as one moves up the hierarchy of system levels the size of each level will increase while speed will decrease. This follows in virtue of the aggregative size of levels. If components of one system level are of a characteristic size and those components are put together to form components at the next system level above, then it follows that the higher-level components will take longer to operate than their constitutive elements.

Methodologically, the implication is that the time required to perform different tasks exposes different system levels. Different system levels are composed of increasingly larger components that operate at successively slower speeds. Qualitative shifts in time signatures reflect substantive shifts between different system levels.

The final account to consider is William Bechtel's (1994, 2007). Bechtel's use of levels is a bit different from the preceding three, as it is formed within a larger discussion of mechanistic explanation. An explanation qualifies as mechanistic when it identifies a hierarchical system whose components, in virtue of their organization, produce some activity or behaviour. The goal of mechanistic explanation is to decompose a given mechanism, and its constitutive activities, into its underlying component parts, showing how those components conspire to produce the activity of the composite whole (see, e.g., Bechtel & Richardson, 1993; Craver, 2007). Rather than viewing levels as substantive divisions between different systems (as Newell does, for example), Bechtel prefers to view levels as local fields of analysis. Levels pick out the various explanatory strategies that can be used during mechanistic decomposition. As Bechtel writes: "[m]ultiple cycles of analysis thus give rise to a hierarchy of levels that is confined to a given mechanism" (2007, p.56).

Bechtel (2007) offers the following example. Suppose a biological mechanism involves sodium molecules crossing over a cell membrane. On the mechanistic account, the sodium molecules and cell membrane can be said to be at the same level if both are implicated in the operation of the biological mechanism in question. The status of standing at higher or lower level depends, in part, on the explanatory role occupied by the component within the larger investigation – in contrast to, for example, slotting into a global organization of size, shape or motion. If the investigator pursues another cycle of decomposition, the components might be further decomposed, but this will still be local to the analysis being offered.

On Bechtel's analysis, it is not possible to say that one mechanism exists at a higher or lower level than another.

There is no a global ordering available for mechanisms. Levels simply define the various explanatory strategies employed at any one cycle of decomposition. The existence of levels is mechanism dependent. As Bechtel puts it: “[L]evels on the mechanistic account are real in that they deal with the particularities of actual components and their operations, but they are perspectival in that they are defined with respect to specific foci on mechanistic activities (Wright & Bechtel, 2007, p.57).

These four accounts should provide a flavor of the kind of levels talk that has prevailed in much of cognitive science. As one can see, the range is quite large. Just what a level is, how it is to be used, and why such talk is important varies notably from author to author. The question to address is how to make sense of these various uses of levels.

### Analyzing Levels

There are two general points of comparison or dimensions of analysis that can be used to examine levels talk. The first deals with the types of items that figure into an account of levels; what ‘relata’ the view implicates. The second deals with how an account uses levels talk; to what end or purpose an account deploys levels. Consider each in turn.

First, there seem to be two general types of relata used in levels talk. The first is ontic relata. On this view, levels talk is taken to hold of activities, structures, and properties. It applies to cognitive entities as they appear *in the world*. Newell offers something like this application with his system levels account: “[A] system level is a property of nature, and not just something in the head of the observer” (1990, p.118). On the ontic application, levels are part of the configuration of the world, picking out the various features constitutive of cognition.

The second type of relata is epistemic relata. On this view, levels talk is taken to hold of ‘linguistic’ or ‘theoretical’ entities. Marr is reasonably interpreted as subscribing to something like this view. When applied to epistemic relata, levels talk functions to provide perspectives for understanding cognition. It offers a means of enhancing an investigator’s epistemic situation.

When viewed as extremes, these two general applications form a continuum on which various accounts can be located. Figure 1 provides an illustration.

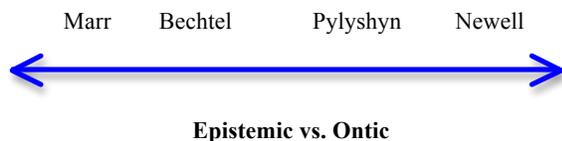


Figure 1: Levels accounts arranged along a spectrum of commitment to epistemic vs. ontic relata.

On the one hand, there are those views such as Marr’s that apply levels talk to epistemic relata. On the other hand, there are also those views such as Newell’s that apply levels

talk to ontic relata. An account’s positioning depends on the relata it implicates.

Notice that Bechtel’s view provides something of an interesting middle ground, as it straddles the two poles. On Bechtel’s view, levels of mechanisms are both ontic structures and observer-dependent entities. Mechanistic activities are real phenomenon in the world, but their arrangement into different levels depends on the explanatory goals of the investigation.

An additional implication of applying levels talk to epistemic versus ontic relata is that it changes how levels relate within a given account. If, on the one hand, an account applies to ontic relata, then its levels relate via a metaphysical realization relation. If, on the other hand, the account applies to explanatory or theoretical entities, then its levels relate in virtue of an explanatory realization relation.

Metaphysical realization is the familiar notion from philosophy of mind that says that two entities stand in a realization relation if and only if when there is change in the higher-level entity there is a corresponding change in the underlying entity change, but not the reverse (Kim, 1998). The standard example is the relation between mental and physical states. Since there cannot be change among mental states without a corresponding change among physical states, physical states are said to realize mental states. Metaphysical realization involves a determinative, asymmetric dependence relation between different properties, states or activities.

Explanatory realization, on the other hand, has received a little less attention than its ontological cousin. In explanatory realization, the dependence relation tracks items of particular theories, languages, or models. It addresses explanatory rather than metaphysical entities. The realization relation holds between statements rather than entities.

An illustrative example comes from theory reduction in the philosophy of science. On the standard model of theory reduction, one theory reduces to a second if and only if the first theory is derivable from the second given certain “bridge laws” (Nagel, 1961). Reduction occurs when all of the items of one theory can be translated into items of the second theory via specific law-like statements or logical connectives. For example, the theory of thermal conductivity reduces to electrical conductivity because (most) of the terms in thermal conductivity can be translated into terms of electrical conductivity via covering generalizations such as the Wiedemann-Franz Law.

Whether or not the reduction is successful is beside the point. The point to note is that the supportive relation holds between different theories or linguistic items and not ontological structures. Explanatory realization shifts attention to how descriptions of entities within theories sustain or determine the theoretical fruitfulness of entities described at other levels.

On the epistemic application, hierarchies of levels stand in similar sorts of explanatory realization relations. This is

why both Pylyshyn and Marr emphasize the collective use of several levels. On Marr's account, for example, the computational level identifies a given input-output mapping, while the algorithmic level describes the procedures by which the function is computed. The algorithmic level supplies crucial constraints on satisfying the input-output mapping. The algorithmic level makes it possible for the computational level to track real computations in the world.

Second, different accounts of levels can also be measured to the extent to which they use levels talk for epistemic versus ontic or metaphysical purposes.

On the ontic usage, levels talk attempts to provide descriptions of how cognition is organized. Levels talk is in the business of providing accurate descriptions of how the world is organized. Discussion attempts to furnish representative accounts of how cognitive systems and capacities are structured into levels.

Consider Newell's account, for example. On Newell's view, different "cognitive bands" address different sets of "system levels" – for example, the biological band addresses the neural circuit system level, while the cognitive band deals with the deliberate act system level. On Newell's view, levels talk functions to describe the organization and structure of cognition as described by the different system levels. Thus, according to the ontic usage, levels talk has a distinctly ontic role to play in discussion of cognition.

Contrast this position with a second, epistemic type of usage. On this usage, levels talk deals with how cognitive phenomena should be studied. Levels function to provide perspectives or viewpoints from which to investigate cognition. As Dawson puts it: "levels do not attempt to explain the nature of information processing devices, but instead provide an epistemology – a way to inquire about the nature of the world" (2013, p.53). Levels talk acts as a tool for inquiry rather than as an ontic description of cognition.

Consider Marr's account, for example. On Marr's view, the computational, algorithmic and implementational levels help to organize cognitive investigation into its most explanatory fruitful parts. The computational level deals with function of the cognitive system, the algorithmic level deals with the procedures used to carry out said function, and the implementational level deals with physical material in which the computation is instantiated. Each level partitions investigation into distinct units of analysis such that it structures and simplifies discussion.

When viewed as extremes, these two general usages form a second continuum on which various accounts can be located. Figure 2 provides an illustration.

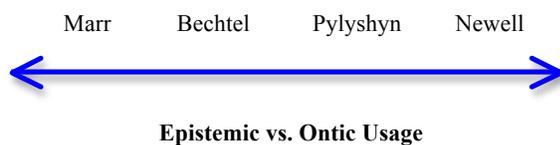


Figure 2: Levels accounts arranged along a spectrum of commitment to explanatory versus ontic usage.

On the one hand, there are those views such as Marr's that deploy levels talk for explanatory purposes – these accounts view levels talk as a tool for inquiry. On the other hand, there are those views such as Newell and Pylyshyn's that use levels talk for metaphysical or ontic purposes – these accounts view levels talk as a means to describe cognitive structures as they appear in the world. Different accounts employ levels talk to different ends.

### A Conceptual Framework for Levels

The fact that levels talk can be sorted along two dimensions of analysis allows for the creation of a conceptual framework that can be used to map the logical space of levels talk. Figure 3 provides an illustration.

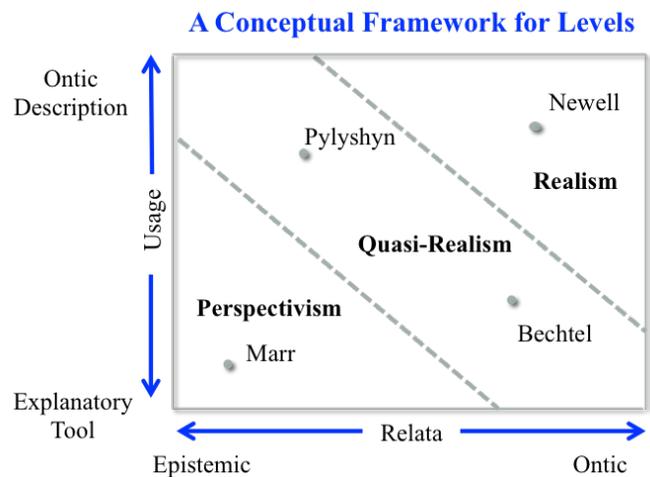


Figure 3: Levels views compared along commitment to ontic versus epistemic relata and usage.

Marr's account occupies the bottom left corner of the space. This is because it applies levels talk to exclusively theoretical or linguistic entities. Both the relata and usage are taken to be explanatory in nature. Marr's account is followed by Pylyshyn and Bechtel's accounts in the middle. These views tow a more neutral line, conceiving of levels as at once ontic and epistemic in character to some degree – Pylyshyn's view, for example, applies to epistemic relata but is used to make claims about the leveled structure of cognition. Finally, there is Newell's view in the top right hand corner of the space. Newell's position here is owed to his adherence to almost entirely ontic application of levels talk – levels talk applies to ontic relata (system levels) and is used to describe the structure of cognitive architecture in terms of the system levels.

What is interesting about this framework is that it maps many of the possible applications of levels talk. It is both descriptive and generative. Various accounts are positioned with respect to how they talk about levels, but at the same time the framework leaves it open as to how other accounts might use levels.

Now one might worry that the framework uses two dimensions of analysis where only one is necessary, particularly given that both points of comparison include an ontic versus epistemic element. The reason two dimensions are more effective than one is that two dimensions allow the current framework to distinguish positions, such as Bechtel or Pylyshyn's, that deploy level talk both epistemically and ontically. Using only one dimension would fail effectively convey the difference between positions such as these, as Pylyshyn's view implicates epistemic relata but makes ontic claims and Bechtel's view implicates ontic relata but uses levels as an epistemic tool.

One might further worry that the present analysis focuses too much on "first generation" cognitive science. That it does not pay sufficient attention to neurological or biological levels. Though this objection is well motivated, it is misplaced in present context. Nothing in the current analysis precludes discussion of neurological or biological levels. It is simply that neurological or biological levels have not featured prominently in discussion of levels – for discussion, see Craver (2015). Although it is certainly right to say that cognitive processes, states, and systems are abstract in the sense that they can be discussed without reference to material form but concrete in the sense that they are realized in a specific physical medium (Thagard, 2014), the current analysis is neutral with respect to the importance of this fact. The current analysis simply attempts to map the space of levels talk as it has been most prominently developed within cognitive science. Whether or not neurological levels are crucial to understanding cognition is of secondary concern to the present investigation.

The framework also reveals three more general, orienting positions that might be held with respect to levels. The first is what might be called the "perspectivist or perspectival" view. On the perspectival view, levels talk applies to the products or units of scientific investigation. Levels deal with theories, explanations or research programs rather than ontic structures. Different epistemic structures are sorted according to the different questions they address. Items on the perspectival view relate in virtue of standing in explanatory rather than ontic relations – recall, for example, what was said about explanatory versus metaphysical realization. Included here are views such as Marr's.

The second is what might be called the "realist" view. According to this view, levels talk applies to items *in the world*. This means that levels deal with entities, activities and properties. These items exist as part of the furniture of cognition. These structures relate via compositional relations. Each item is constituted or realized by items residing at lower levels. Often, a principle of prediction or causality is used to sort items into different levels. Newell is included under this second heading.

Finally, there is the "quasi-realist" view. On this view, levels talk applies to the components and activities of mechanisms. These are features of the world. However, in contrast to realism and closer to perspectivism, levels also have an epistemic bent. Though levels count as real parts of

cognition, they only emerge during or as part of investigation. As Bechtel writes: "Constitutive strategies describe the mechanism's component parts, their operations, and their organization, showing how the mechanism's constituency is responsible for its activity" (2007, p.62). Levels are theoretical impositions, being explanatory strategies, but they are also part of the mechanisms constitutive of cognition. Bechtel falls under this third category.

Note that these three views are probably best regarded as representative positions rather than actual positions individual authors' hold. Each view represents a sort of general umbrella category or idealized position that can be used to explore different types of levels talk without getting into all the messiness of interpreting individual authors' viewpoints.

What is it about the realist, quasi-realist, and perspectivist positions that make other views cluster around them? One possibility is that the divisions reflect larger differences in overarching conceptions of cognition or science (see, e.g., van Fraassen, 1980). That the realism/anti-realism split reflects a larger divide that runs through most theoretical discussions in the sciences. Another possibility might be that the various views are just different sides of the same coin; that with additional work each can be distilled into one underlying view. For example, perhaps Marr's more perspectivist view collapses into the realist view if it is cashed out in specifically causal terms. Or, perhaps the claims of the realist can be deflated and shown to be subsumable under the quasi-realist's account.

There are three desirable features to the present framework. First, it helps to clarify confusions between different accounts of levels. For example, in their review of Marr's tripartite account of levels, Churchland, Koch, and Sejnowski write: "when we measure Marr's three levels of analysis against levels of organization in the nervous system, the fit is poor and confusing" (1990, p.38). For Churchland, Koch and Sejnowski, Marr's levels are either inadequate or patently false. They fail to offer accurate characterizations of the organization of cognition as it realized in the brain. Since there are a great many spatial levels of organization in the brain, Marr's levels makes for a poor fit.

Understood within the context of the present analysis, the issue is immediately clear: Churchland, Koch and Sejnowski have mistaken the relata of Marr's account. Marr's levels are not a poor fit for brain organization, because they are not meant to fit with brain organization. Churchland, Koch and Sejnowski have confused Marr's epistemological application of levels talk with an ontological one. In cases such as these, knowing what elements to look for can help identify key areas where potential misunderstandings might arise. The framework can play a therapeutic role in discussion.

Second, the conceptual framework furnishes interpretative benefits. This is because it allows the underlying reasoning of each account to be more easily

understood. For example, on Newell's view, evolution is thought to favor systems constructed through addition of modular parts, since it is more likely that nature will build complex systems incrementally rather than all at once. Understood in realist terms, this reasoning makes sense. Attention to the spatial and temporal features of levels weds naturally to a focus on evolutionary considerations. Once it is clear at what end each account aims, the underlying reasoning is more easily explicated.

Third, in the case of more substantive disagreements over levels, the framework helps to clarify different lines of support, and thus chart potential argumentative paths through the levels space. For example, returning to Newell's account, if one is dubious of the appeal to evolution on the grounds that such arguments are underdetermined, then one can argue that the realist position is importantly undermotivated; that application of levels talk to cognitive structures is too quick. Or, to take the contrary position, if one is concerned that the perspectival view relegates levels talk to an arbitrary status, then the current analysis provides a direction in which to mount this challenge – for example, that it fails to get at the real structures and organization of cognition. The framework helps to simply discussion by cutting to the core elements of disagreement.

For all these reasons, the present framework offers an important tool for cognitive science. It provides a means to thinking about and sorting through various types of levels talk. It also reveals key elements of the various accounts while staying neutral with respect to alternative ways of applying levels talk.

In closing, a quick qualification needs to be made. The preceding analysis is not meant to be the definitive guide to conceptualizing levels. It is entirely possible that other dimensions of analysis exist on which applications of levels talk can be analyzed. However, the present account does identify several important features that cut across a large swath of levels talk. What the present discussion has provided, we hope, is one concrete step toward making of further sense of levels in cognitive science.

## References

- Bechtel, W., & Richardson, R. (1993). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.
- Bechtel, W. (1994). Levels of Description and Explanation in Cognitive Science. *Minds and Machines*, 4, 1-25.
- Broadbent, D. (1985). A question of levels: Comments on McClelland and Rumelhart. *Journal of Experimental Psychology: General*, 114, 189-192.
- Churchland, P.S., Koch, C., & Sejnowski, T.J. (1990). What is computational neuroscience? In Schwartz, E.L. (ed.), *Computational Neuroscience*. Cambridge, MIT Press.
- Craver, C. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford, UK: Clarendon Press.
- Craver, C. (2015). Levels. In T. Metzinger & J. M. Windt (Eds). *Open MIND*: 8(T). Frankfurt am Main: MIND Group. doi: [10.15502/9783958570498](https://doi.org/10.15502/9783958570498)
- Dawson, M. R. W. (1998). *Understanding Cognitive Science*. Blackwell Publisher.
- Dawson, M. R. W. (2013). *Mind, Body, World: Foundations of Cognitive Science*. Alberta, Edmonton: Athabasca Press.
- Dawson, M. R. W., Boehler, P. M., & Valsangkar-Smyth, M. (2000). Representing space in a PDP network: Coarse allocentric coding can mediate metric and nonmetric spatial judgements. *Spatial Cognition and Computation*, 2, 181-218.
- Kim, J. (1998). *Mind in a physical world*. Cambridge, MA: MIT Press.
- Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.
- Marr, D., & Poggio, T. (1977). From understanding computation to understanding neural circuitry. *Neurosciences Research Program Bulletin*, 15, 470-488.
- Mcclamrock, R. (1991). Marr's Three Levels: A Re-Evaluation. *Minds and Machines*, 1, 185-196.
- Nagel, T. (1961). *The Structure of Science. Problems in the Logic of Explanation*, New York: Harcourt, Brace & World, Inc.
- Newell, A. (1980). Physical symbol systems. *Cognitive Science*, 4, 66-83.
- Newell, A. (1990). *Theories of Unified Cognition*. Cambridge, MA: Harvard University Press. Pinker, S. (1997). *How the Mind Works*, New York: W. W. Norton & Company.
- Pylyshyn, Z. (1980). Cognition and computation: Issues in the foundation of cognitive science. *Behavioral and Brain Sciences*, 3(1), 154-169.
- Pylyshyn, Z. (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science*. Cambridge, MA: MIT Press.
- van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Oxford University Press.
- Wimsatt, C. (1976). Reductionism, Levels of Organization, and the Mind-Body Problem. In G. Globus, G. Maxwell & I. Savodnik, *Consciousness and the Brain* (pp. 205-267), Spring Us.
- Wimsatt, C. (1994). The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets. *Canadian Journal of Philosophy*, 20, 207-274.
- Wilson, R., & Craver, C. (2007). Realization: Metaphysical and Scientific Perspectives. In P. Thagard (ed.), *Philosophy of Psychology and Cognitive Science*, Amsterdam: Elsevier, pp. 81-104.
- Wright, C.D., and Bechtel, W. (2007). Mechanisms and Psychological Explanation. In P. Thagard (ed.), *Philosophy of Psychology and Cognitive Science* (pp. 31-77). Amsterdam: Elsevier.