Leveling the Field: Talking Levels in Cognitive Science

by

Luke Kersten

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in partial fulfillment of the requirements for the degree of

Master of Cognitive Science

in

Cognitive Science

Carleton University
Ottawa, Ontario

Abstract

This thesis aims to advance the study of cognitive science by examining the "levels metaphor." The levels metaphor is defined as the application of levels talk to various aspects of scientific investigation. The thesis examines several applications of the levels metaphor within cognitive science and provides a conceptual framework for analyzing discussion. The thesis argues for a pluralistic approach to levels. The main claim is that different applications of the levels metaphor are justified insofar as attention is paid to how and why the metaphor is deployed. To show that my approach has practical applications, I discuss the role of levels within computational cognitive modeling.

Acknowledgments

There are several people I would like to acknowledge for helping to make this thesis possible. First, I would like to thank my co-supervisors Andrew Brook and Robert West. Their continued support and assistance during the preparation and writing of this thesis has been instrumental to its success. Second, I am grateful to the graduate students at the Institute of Cognitive Science at Carleton University. Conversations with a number of talented individuals helped to shape various aspects of this project – I am particularly appreciative of long discussions with Can Mekik and Gautam Sukumar. Finally, I would like to acknowledge my deep debt to Laura Davis, whose patient and encouraging support proved an invaluable resource throughout the entire project.

Table of Contents

List of Tables

List of Illustrations

Introduction

Early in the history of neuroscience it looked as if simply recording the cell activity of different neurons would reveal how the brain performed various cognitive functions. A string of discoveries revealed that particular neurons were responsive to specific sets of stimuli – for example, that the ganglion cells in the frogs' retina served as "bug detectors" (Barlow, 1953) or that cats and monkeys had "edge detector" cells in the visual cortex (Hubel & Wiesel, 1968). However, by the early nineteen seventies the cell-recording research programme had faltered. Part of the reason was that it turned out that knowing which cells in the brain were responsive to particular sensory information revealed little about why and how those neurons contributed to cognitive functions.

To some, the breakdown of neuroscience's early vision suggested an important lesson. Marr, for example, concluded that: "[t]here must exist an additional level of understanding at which the character of the information processing tasks carried out during perception are analyzed and understood in a way that is independent of the particular mechanisms and structures that implement them in our heads" (1982, p. 19). Profitable research required shifting to different "levels". Neurophysiology and psychophysics could not be the whole story when it came to investigating cognition.

Talk of levels has only grown in popularity since Marr. Even a cursory glance at the literature on levels reveals a variety of usages, particularly within cognitive science. To name just a few, authors speak of levels of abstraction (Floridi, 2008), analysis (Dawson, 1998), causation and explanation (Kim, 1998), processing (Craik & Lockhart, 1972), aggregation (Wimsatt, 1994), implementation (Marr, 1982), organization

(Churchland & Sejnowski, 1992) and realization (Gillett, 2002). Application of the "levels metaphor" is ubiquitous in cognitive science.

What is interesting is that in spite of the widespread application of the levels metaphor within cognitive science surprisingly little attention has been paid to the notion. 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 large litany of literature on the subject" (p.55). Or, as Craver (2015) puts it: "Despite the ubiquity of levels talk in contemporary science and philosophy, very little has been done to clarify the notion" (p.23). Though the notion of 'level' is part of the conceptual backdrop of cognitive science – similar to notions such as computation and representation, though perhaps less often and explicitly discussed – it has largely avoided sustained scrutiny.

This state of affairs calls out for improvement. Detailed investigation is needed to understand what levels are, how they function, and what role they should have within cognitive science. The purpose of this thesis is to provide an analysis of the levels metaphor within cognitive science, to explore and examine various applications of the metaphor, as it is commonly deployed within the field.

In this thesis, the "levels metaphor" is defined as application of levels talk to different aspects of scientific thinking. To take a simple example, when one says that sophomores are at a higher level than freshman, one is applying the levels metaphor. One is applying levels talk to the standing of students within the university system. The levels metaphor involves sorting objects or processes into different layers, usually on the basis of some principle or criterion – for example, in the above case the principle is the year of study. The levels metaphor is a tool or lens through which phenomena are organized and conceptualized.

Consider, as an analogy, the role of the ‘computer metaphor’ within cognitive science. Researchers have developed a number of complex and novel theories and models about cognition by making assumptions about the similarity between human and computer information processing, everything from computational theories of language acquisition to Bayesian theories of decision making. The computer metaphor has not only helped to structure and organize individual thinking, it has also acted as a larger paradigm from within which to conceptualize and conduct research.

The levels metaphor has often occupied a similar role within cognitive science. As we will see, similar to the computer metaphor, talk of levels has often figured centrally into how researchers have thought of and studied the cognitive world. To take one example, when used in functional analysis, levels orient movement through rounds of decomposition. Successive system decompositions are often interpreted as revealing additional levels of organization (Cummins, 1983).[1] The levels metaphor has often helped to guide and sustain scientific theorizing within cognitive science. As Craver points out: “These metaphors are too basic to how we organize the world to seriously recommend that they could or should be stricken from thought or expression” (2015, p. 2).

Why view the term ‘level’ as a metaphor? Consider what makes something a metaphor generally. Metaphors require mapping a primary subject (or topic) to a secondary subject (or vehicle) (Jackendoff, 1983). For example, in the metaphor “all the

---

[1] I will say more about the role of the levels metaphor within cognitive science in Chapters 1 and 2. For now, this description should provide at least a sense of how the idea of levels bears on cognition.

world is a stage", the primary subject is the world and the secondary subject is a stage. When the vehicle 'stage' is attached to the topic 'world', it conveys the sense that life is often similar to a theatrical performance. The primary subject is said to have similar properties to that of the secondary subject. Metaphoric thinking allows one to convey complex and novel ideas. Similarly, the term level often occupies a metaphoric role in everyday and scientific language. For example, when talking about buildings, level acts as a secondary subject, helping to describe the idea of being parallel to the ground. It conveys the idea that there is a spatial arrangement to landings within buildings.

In terms of substance, this thesis argues for a 'pluralistic approach' to understanding the levels metaphor. The main claim is that since it is unlikely that there is a single verdict when it comes to the conceptual soundness of the levels metaphor within cognitive science, different applications are justified insofar as they are evaluated and used with caution.[2] The central message is that no one application of levels metaphor is ideal for all of cognitive science.

The current thesis also looks to exemplify what Andrew Brook (2009) has dubbed "philosophy in cognitive science." In addition to providing conceptual analysis, philosophical work should also lend insight into the day-to-day issues preoccupying cognitive researchers. This thesis attempts to meet such a standard by (i) providing a general analysis of the notion of 'levels' and (ii) showing how a robust understanding of levels can actively contribute to the interpretation and construction of cognitive models. Part of the novelty of this thesis therefore lies in the descriptive and normative insight it provides to cognitive science.

[2] In this way, the central purpose of this thesis echoes that of Craver (2015).

Structurally, the thesis divides into three chapters. Chapter 1 introduces and outlines the levels metaphor. It explores several of the most prominent applications of the metaphor within cognitive science and offers a conceptual framework for organizing and analyzing the different applications. Chapter 2 investigates whether there is a preferred conception of levels for cognitive science. After outlining and contrasting several possibilities, it is argued that cognitive science is best served by a pluralistic approach to levels. Finally, Chapter 3 examines the potential contribution of the levels metaphor to cognitive modeling. The chapter introduces cognitive modeling, outlines a specific application of the levels metaphor and argues for adopting the pluralistic approach.

Chapter 1 - The Levels Metaphor and Cognitive Science

1.0 Chapter Overview

The purpose of this first chapter is to outline and clarify the various ways the levels metaphor has been applied within cognitive science. The chapter begins by introducing and motivating discussion of levels. It then turns to outlining and relating several specific applications of the levels metaphor. The result of this discussion is the creation of a conceptual framework that maps the logical space of levels talk.

1.1 Levels in Cognitive Science

An instructive place to begin discussion is by considering some general questions about levels. These will serve to both introduce and motivate the topic in a bit more detail. Three questions, in particular, can be considered.

1.1.1 The Question of Usage

First, how have people used levels in cognitive science? To what ends has the level metaphor generally been put? Two major applications stand out. First, levels have been used to motivate and sustain particular conceptions of cognition. Marr (1982), for example, uses levels to argue for an information-processing conception of cognition, writing: "The message was plain. There must exist an additional level of understanding at which the character of the information-processing tasks are carried out during perception" (1982, p.19). Because neurophysiology and psychophysics only describe the behaviour of cells, a full explanation of perception requires shifting to an additional level of analysis. At this additional level, cognition is best thought of as a form of information processing.

In a related vein, Pylyshyn claims that cognition is best thought of as a form of rule-governed symbol manipulation. Pylyshyn (1984) motivates this position by describing cognition at three different levels of description, writing: "Explaining cognitive phenomena requires that we advert to three distinct levels of this system…This tri-level nature of explanation in cognitive science is a basic feature of the computational view of mind" (p.xviii). In Pylyshyn's hands, levels support a computational view of cognition.

Consider another way levels have been used – this time in debates over connectionist and symbolicist models of cognition. Broadbent (1984) argues that cognitive models, particularly the kind of interest to cognitive psychologists, should be pitched at an "algorithmic level." This is in contrast to connectionist models, which address cognition at an "implementational level." For Broadbent, because connectionist models reside at an implementational level, they are ill suited to accurately describe and explain cognition.

In contrast, Dawson (2013) argues that the tension between connectionist and symbolic models has been overblown. When viewed from Marr's three levels, connectionist and symbolicist models turn out to be explanatorily equivalent. Using a mushroom classification problem, for example, Dawson et al. (2000) shows that the hidden units of an artificial neural network can be described in terms of executing 'productions'. These productions make the connectionist model closer in spirit to symbolic, classical models. The theoretical translatability of the two types of models at different levels shows their explanatory equivalence. Contrary to Broadbent, levels unify rather than undermine different approaches to cognition.

1.1.2 The Question of Value

Consider a second question: What is at stake when it comes to the levels metaphor? Why should one care about talk of levels in cognitive science? One answer is that levels talk provides a unifying framework for cognitive science. That cognitive science loses much of its supposed coherence without use of the levels metaphor.

Pylyshyn and Dawson maintain something like this position. Dawson, for example, writes: "[E]xplanations of information processors require working at four different levels of investigation, with each level involving a different vocabulary and being founded upon the methodologies of different disciplines" (2013, p.19). Levels talk offers a means to integrate, clarify, and order information-processing theories. Insofar as cognitive science assumes that cognition is a form of information processing it requires working at several different levels. Levels provide a conceptual home for integrating different theories, methods, and techniques.[3]

Another answer to the value question is that levels talk facilitates cognitive investigation. For example, construed ontologically – that is, as applying to structures in the world – the levels metaphor helps to describe the structure of cognition. It aids in figuring out how cognitive systems are hierarchically organized. Construed epistemically – that is, as applying to theories or explanations – the levels metaphor offers a tool for organizing cognitive investigation. It enhances the epistemic situation of researchers. As

[3] The interested reader might notice that whereas Dawson (2013) speaks of four levels, discussion in this thesis refers to only three – Dawson separates out the 'architectural' level from Marr's algorithmic level. The reason for this difference is that in this thesis, especially when speaking about Marr, the architecture of procedures (the symbolic building blocks and ways of combining them) and the procedures of cognition (how those blocks are combined to produce algorithms) are seen as parts of a single activity at one level. Furthermore, nothing major hinges on this delineation. The distinction is mostly cosmetic. Dawson's (1998) discussion of levels, for example, is structured in terms of three levels and yet still contains discussion of the architectural level.

we will see, this epistemic/ontic divide has an important role to play in discussions of levels.

1.1.3 The Question of Uniqueness

Finally, why should one think that there is anything unique to the application of the levels metaphor in cognitive science? Why is an analysis of the levels metaphor needed for cognitive science specifically? There are two points to consider here. One is to recall that cognitive science is inherently interdisciplinary. It draws from several disciplines, e.g., neuroscience, linguistics, philosophy, etc. Because of this, it would be surprising if the applications of the levels metaphor were in some way not unique to the interdisciplinary nature of the field, particularly given the likelihood of blending specific usages from its contributing disciplines.

The second point to notice is that the subject matter of cognitive science is at root an inherently "multilevel" phenomenon. Cognitive processes, states, and systems are both abstract in the sense that can be discussed without reference to material form, but also concrete in the sense that they are realized in a specific physical medium (Thagard, 2014). Any application of the levels metaphor in cognitive science will likely have to wrestle with the fact that cognitive phenomena live two lives as both biological and informational entities.[4]

1.1.4 Summary

What the preceding discussion has conveyed, I hope, is the sense that the notion of level cuts across many of the core elements of cognitive science. Rather than being an

[4] I am working here with the established assumption that cognitive states are information-bearing states (though of what kind I remain neutral) and that the processes that operate over those states are computational in character.

ancillary notion, use of the levels metaphor has been integral to many of cognitive science's central projects. Whether it is in terms of adjudicating longstanding debates or motivating foundational concepts, talk of levels has occupied an important role in the life of cognitive science.

1.2 Levels on Display

Although there has been excellent discussion in the wider philosophical literature about levels, the focus here, as mentioned, will center on how the notion is deployed within cognitive science.[5] For this reason, I outline four particular applications of the levels metaphor within cognitive science. These accounts represent several of the most influential and detailed positions that have been expressed within the field. This survey will help to highlight some of the details and larger points of discussion. It should also be mentioned that the intention here is not to outline each account in all its rich detail but, rather, to provide a basic sketch of each author's view.

1.2.1 The Tri-Level Hypothesis

The first account to consider is David Marr's. Marr (1977, 1982) has offered probably the most famous and influential account of levels. Marr identifies three levels of analysis.[6] Here is how Marr characterizes the view:

> At one extreme, the top level, is the abstract computational theory of the device, in which the performance of the device is characterized as a mapping from one kind of information to another, the abstract properties of this mapping are defined precisely, and its appropriateness and adequacy for the task at

---

[5] For a more general survey of levels talk see Craver (2015). Or, for earlier discussions of levels within philosophy of mind and cognitive science see McClamrock (1991) or Bechtel (1994). For excellent discussion of levels as it pertains to science as a whole see Wimsatt (1976, 1994).

[6] Some have claimed that a proper reading of Marr reveals that, "they're not levels and there aren't three" (Brook, personal communication).

> hand are demonstrated. In the center is the choice of representation for the input and output and the algorithm to be used to transform one into the other. And at the other extreme are the details of how the algorithm and representation are realized physically. (1982, p.25)

It will be helpful to unpack the view in a bit more detail.

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 here aims to translate general, everyday descriptions of cognitive phenomena into particular information-processing problems or tasks. For example, in his classical study of vision, Marr claimed that the function of vision was object recognition. Given this, the problem is to identity the particular constraints that allow the visual system to transform two-dimensional retinal images into three-dimensional representations of the environment. To explain the input-output function that turns distal information into complex, internal representations.

Next, 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 procedures that solve a particular information-processing problem. For instance, since photoreceptor cells detect only changes in light intensity, to extract useful information about geometric organization, the visual system must decide when changes in light intensity values reflect edges and when such changes reflect straight lines. To do this, the visual system employs particular mathematical operations. These operations constitute the algorithms or procedures by which the system carries out the larger input-output mapping defined at the computational level.

Finally, there is the implementational level. At the implementational level, the task is 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, in the case of the visual system, this might involve detailing how cells in the retina encode light intensities or how the optic nerve carries information downstream to the visual cortex.

On Marr's account, each of the three levels plays a different role within the cognitive investigation. Each satisfies different epistemic ends. The computational level explains what function is being computed, the algorithmic levels specifies how that function is carried out, and the implementational details what structures support the entire process. By providing different characterizations of a target system, each level provides answers to different types of questions. The resultant picture is one where the three levels promote different theories, invoke different methodologies, and employ different experimental techniques in the service of explaining some cognitive system.

For example, when investigators emphasize the spatiotemporal properties of a system, they operate at an implementational level. Using the techniques of neuroscience, such as neuroimaging or lesion studies, researchers identify the physical components that support some cognitive system. The interest is in investigating how the biophysical properties of neurons are suited to carrying out specific tasks – for example, how things like receptive fields and synaptic connections contribute to encoding and filtering information. When investigators focus on how components interact so as to produce a particular operation or procedure, they ascend to a higher, algorithmic level of analysis. Here researchers invoke methods and techniques from cognitive psychology – for example, randomized block designs or error rates, etc. In short, Marr's hierarchy of levels partitions or organizes investigation into its most explanatorily fruitful parts according to the types of questions different vocabularies address.

1.2.2 The Multiple Vocabularies of Cognition

The second account to consider is Zenon Pylyshyn's. Similar to Marr, Pylyshyn (1980, 1984) is also impressed by the explanatory power of levels, though Pylyshyn is more concerned with articulating the foundations of cognitive science than offering a unifying methodology. Pylyshyn similarly identifies three "levels of description." Here is Pylyshyn describing the view:

> There is a natural domain in psychology which corresponds roughly to what we pretheoretically call *cognition*, in which the principal generalizations covering behaviour occur at three autonomous levels of description, each conforming to different principles. These levels are referred to as the biological (or physical level, the symbolic (or syntactic or sometimes the functional level, and the semantic (or intentional) level. (1984, p.259)(original emphasis)

Again, like Marr, it will be worthwhile to unpack the view in a bit more detail.

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. Events at this level appeal to how an individual represents the world rather than to how the world actually is. It uses an "intentional" vocabulary.

To avoid confusion, it's also worth mentioning that Pylyshyn's semantic level should not be thought of as equivalent to Marr's computational level, though there is reason to see them as at least partially related (see Dawson, 1998). Roughly speaking, the difference amounts to the kinds of properties each level appeals to. Pylyshyn's semantic level deals with properties extrinsic to the system being described, while Marr's computational level deals with elements internal to the system. Intentional entities such as beliefs or desires refer to things outside a system – for example, beliefs about cars are about cars in the world. In contrast, the units relevant to computational analysis are those internal to a system – for example, the information added by a system to flesh out two-dimensional retinal images. Whereas Marr's computational level deals with abstract information-theoretic descriptions of cognitive systems, Pylyshyn's semantic level deals with representational content. The distinction is worth noting because the two levels are sometimes mistakenly taken to be equivalent.[7]

Second, there is the "symbolic or syntactic" level. Behaviour at the symbolic level is described in terms of the functional properties of an individual. For example, to explain why an individual does poorly on a memory-recall task, the symbolic level appeals to functional properties such as 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 this final level of description, individuals are described in terms of the familiar vocabulary of the physical sciences, e.g., neurology, chemistry, biology. Events at the biological level are explained under a physical description, referencing the neurological or chemical factors that are causally salient for a given behaviour. For example, to account for why someone is thinking of faces rather than sounds, explanations at the biological level might appeal to differences in brain activity or neural-chemistry.

[7] See, for example, Dawson (1998, Ch.2).

Like Marr, Pylyshyn's three levels are also arranged hierarchically. At the bottom is the biological level, next comes the symbolic level, and at the top is the semantic level. However, for Pylyshyn, the purpose of talking at different levels is a bit different than for Marr. Whereas Marr is concerned with answering different types of questions at each level, Pylyshyn is focused on capturing distinct generalizations.

Pylyshyn provides an illustrative example. To explain why an individual might run to help someone in a car accident, two types of descriptions might be offered. Under an intentional description, one might explain the behaviour by referencing the individual's mental states – for example, that the individual *believed* that someone was in an accident, that they *desired* to help, and that they had the *goal* of getting help. Alternatively, under a physical or biological description, one might explain the behaviour by identifying the relevant biological factors – for example, how the contraction of arm muscles allows the individual to pick up a payphone or how the pavement provides resistance to the individual running, etc. At this biological-level, a descriptive story is told in terms of the underlying causal factors.

At each level, different descriptive vocabularies capture distinct sets of regularities. As Pylyshyn writes: "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.

1.2.3 The Leveled Architecture of Cognition

Newell (1980, 1990) offers a third application of the levels metaphor. 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. The architecture of human cognition is built up out of a hierarchy of multiple system levels. Each system level is realized by components at the next system level below. Here is how Newell describes it: "At level N, a collection of K components, say of characteristic size S, are put together to form a component at N + 1" (1990, p.119).

The hierarchical arrangement of system levels is dictated by the stability it affords. 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" (1990, p.117).

In terms of identifying system levels, Newell claims that as one ascends upward from the smallest components to the largest, 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.

These different operating speeds separate the human cognitive architecture into different "bands." Newell mentions four: the social band, the rational band, the cognitive band, and the biological band. Each band tracks different time scales of various system levels. They correspond to different groupings of the system levels. For example, neural systems operate at speeds between 100us and 10ms, while cognitive systems operate between 100ms and 10 sec. Qualitative shifts in time signatures reflect substantive shifts between different system levels.

1.2.4 A Middle Ground: Levels of Mechanisms

The final account to consider is William Bechtel's (1994, 2007). Bechtel's application of the levels metaphor is a bit different from the preceding three, as it is mostly 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. Here is Bechtel and Richardson defining the view:

> A machine is a composite of interrelated parts, each performing its own functions, that are combined in such a way that each contributes to producing a behaviour of the system. A mechanistic explanation identifies these parts and their organization, showing how the behaviour of the machine is a consequence of the parts and their organization. (1993, p.17)

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.

It is within the mechanistic context that Bechtel forms his account of levels. Rather than viewing levels as substantive divisions between different systems (for example, as Newell does), 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 a helpful 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. This is despite the fact that cell membranes are composed of sodium molecules. 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. 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.

One important consequence of Bechtel's analysis is that it is not to 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).

Contrasting Bechtel’s account with Newell’s might help to further bring out some of the view’s crucial elements. Consider, for example, that for Newell levels are constructed out of componentially arranged parts. This also holds true of Bechtel’s position. But also notice that whereas Newell slots these elements into a global arrangement based on considerations of size, Bechtel does not. Bechtel is careful to point that such a comprehensive organization of level is not possible. The components that figure in various levels are dependent on the wider explanatory investigation. Though both authors think that the cognition is constructed out of componentially arranged elements, they disagree on how those components are related as a whole.

1.3 Questioning Levels

What the preceding survey should make clear is that just what levels are, how they are identified, and what function they serve varies considerably from author to author. This is both an interesting and puzzling state of affairs. The task is to bring some structure to this wide-ranging discussion. I begin by exploring several questions that can be asked about levels. This gives way to an analysis of how each question applies to the four previously outlined accounts. The aim is to tease out the main points of difference between the four applications of the level metaphor.

1.3.1 Three Questions for Levels

Craver (2015) has recently provided an instructive discussion of levels within science and philosophy. One particularly interesting aspect of Craver's discussion is his identification of three separate "level questions" that can be asked about various applications of the levels metaphor.

First, there is the "Relata Question." The relata question looks to identify the different types of items that can figure into a given account of levels. It asks simply: What are the types of entities that are sorted into different levels? The relata question is interesting because it points to the "entity neutrality" of the levels metaphor. It highlights how the metaphor is indifferent to the types of entities it sorts. As Craver remarks: "The metaphor works in any context because it leaves open just what kinds of objects are to be arranged (2015, p.1). Consider, for example, the array of entities that might be sorted into different levels. The metaphor might sort abstract entities (such as numbers), concrete objects (such as neurons), types (such as pyramidal cells), tokens (such as individual neurons), or activities (such as action potentials). Part of what makes the metaphor so appealing is that it leaves open just which kinds of entities or activities slot into different levels.

Second, there is what Craver refers to as the "Relations Question." The relations question asks: In virtue of what is one item said to be at a higher or lower level than another item? The focus is on what distinguishes two items as being at different levels. For example, consider how one theory might be said to be at a higher or lower level than another. One way to do this would be to sort theories according to their "derivability." If one theory can be derived from a second (and not the reverse), then the first theory can be said to be at a "lower" level than the second, since the second is in some sense more "fundamental" than the first. By providing a criterion of derivability, the two theories might be arranged at higher and lower levels.[8]

Third, Craver points to a "Placement Question." Like the relations question, the placement question looks for the principle or criterion by which two items are sorted. The difference is that this time it focuses on when two items are sorted to the same level. For example, consider how puffins and porcupines might be sorted to the same level. One way to do this would be to appeal to scales of size. Since puffins and porcupines are roughly of the same size (between 50 and 90cm) they can be placed at the same level when taxonomizing species. Size scales offer one principle by which to arrange items at the same level. When used in conjunct with the relations question, the placement question helps to identify the global arrangement of levels, presumably because sorting relata will require identifying into which levels the items do and not slot.

In what follows, I fold the Relations Question and Placement Question into one. Call it the "Criterion Question." The reason for doing so is that in most cases the relations question and placement question track the same principle or criterion. In most cases, they are two sides of the same coin. For example, if two theories are sorted into different levels on the basis of one being derivable from another, then non-derivability is going to be the test for when two theories are at the same level. Of course, this will not always be the case. Economic theory is not derivable from string theory, but this does not mean that each resides at the same level. The point is simply that in most applications of the levels metaphor in cognitive science one criterion is used to answer both questions.

[8] This, in fact, is the standard usage in discussions of theory reduction in philosophy of science.

1.3.2 Applying the Level Questions

The relata and criterion questions provide useful tools for thinking about levels. Each helps to tease out key elements of the various applications of the metaphor. Given this, in what follows I consider how the level questions apply to each of the four previously discussed applications of the levels metaphor.

First, consider how the two level questions apply to Marr's tripartite account. In terms of the relata question, Marr's account sorts different "explanatory frameworks" or "theoretical languages." First, there are computational languages, such as Bayesian or information theory. These form the top level. Next, there are representational and processing languages, such as those provided by cognitive psychology and psychophysics. These form the second, intermediary level. Third, there are physical languages, such as neurology or biochemistry. These form the bottom level. In each case, the different languages slot into a specific place in the explanatory hierarchy.

Next, in terms of the criterion question, Marr's account stratifies different languages according to the different types of questions the languages address. Two languages are localizable to different levels when they answer different types of questions – for example, since computational theories answer what and why questions and algorithmic theories answer how questions, each language can be said to reside at a different level. This point gains further support when it is recalled that Marr's original motivation for ascending to a computational level was based on the idea that neurological theories alone were insufficient to answer questions about why and how certain brain regions instantiated particular cognitive functions (see Marr, 1982, p.14 or Ch.1, sec. 1.1.1).

The flip side, of course, is that when two languages answer the same question they reside at the same level. For example, although Marr does not explicitly discuss such cases, it can reasonably be assumed that because biochemical and electrophysiological theories each deal with questions about how cognitive processes are instantiated in the brain that both reside at the implementational level of analysis.

Next, consider Pylyshyn's account. First, notice that Pylyshyn deploys the levels metaphor with respect to different "descriptive vocabularies." For instance, the semantic level employs intentional terms, the symbolic level employs functional terms, and the biological level invokes physical terms. In each case, the three levels define the sets of linguistic terms that can be used for describing phenomena. Similar to Marr, Pylyshyn's account targets linguistic entities.

A more substantive difference emerges with respect to the Criterion Question. For here, unlike Marr, Pylyshyn focuses on capturing distinct behavioral regularities. Instead of relating languages based on the different types of questions they answer, Pylyshyn holds that two vocabularies reside at different levels when each captures distinct sets of regularities. Recall, for example, Pylyshyn's accident bystander. In this case, the semantic vocabulary, in referencing intentional entities, explains cross-situational regularities – for example, why other individuals might rush to offer assistance in the same situation. On the other hand, the biological vocabulary, in appealing to physical terms, captures individual specific regularities – for example, how the individual moved in the way that she did to help the accident victim in the specific case.

One might wonder how the regularities captured by Pylyshyn's vocabularies differ from those captured by Marr's questions. The difference is subtle but important. Whereas

Pylyshyn's vocabularies capture different sets of regularities in virtue of tracking different sets of components, Marr's questions re-describe the same set of component in different ways. Although both authors appeal to distinctive vocabularies, only one thinks that the different vocabularies pick out different sets of components within the same system.

What about Newell's account? How do the relata and criterion questions apply there? Recall that on Newell's account system levels represent functionally organized components. This means that unlike the previous two accounts the relata being sorted are not "theories" or "languages" but rather parts or properties of systems. Newell's system levels have components with distinct organizational features, such as location, size, shape, and motion. What's more, these components might be either of informational or biological character; they might be neurons or they might be cognitive structures. The point is that the system levels, as constituted by functionally organized components, are part of the ontological furniture of the world.

Turning to the criterion question, Newell's account suggests that system levels are distinguished on the basis of a criterion of prediction. One system level is distinguishable from another if the pattern of behaviour produced at the one system level fails to be predicted using organizing principles at one system level below. For example, Newell identifies one system level as the "deliberate act." This system level is responsible for agent behaviour operating at a 100ms time scale. Just below the deliberate act is the "neural circuit." Because the behaviour controlled by the deliberate act system level is not predictable on the basis of neurological principles, as they pertain to the neural circuit system level, it can be assumed that the deliberate act system level is at a higher level than the neural circuit system level. The organized components are arranged across levels according to the degree to which they can be used to predict the behaviour of other components one level above.

Finally, consider how the level questions apply to Bechtel's account. On the relata question, Bechtel's account is rather straightforward. What are sorted are different activities of mechanisms, along with the constitutive component parts; here it is important to recall that mechanisms are constituted by the interaction and organization of components as they realize some activity – for example, a rat learning to run a maze might constitute the activity supported by mechanisms in spatial memory.

How are the different mechanistic activities related? Bechtel's account appeals to componential relations. One activity and its mechanism are at a higher level than another if its constitutive components are part/whole related to components of another activity. For example, returning to the maze learning case, one component of spatial memory, the mechanism that realizes maze learning, is the hippocampus. The hippocampus "maps" orientations and location for the rat during learning. However, the ability of the hippocampus to map geometric features is, in turn, realized by changes in the synapses of pyramidal cells. Each component is related to a mechanism above as they realize different activities. Craver puts the point a bit more stringently: "Some component X's Ying, is at a lower mechanistic level than S's Ying if and only if X's Ying is a relevant spatiotemporal part of S's Ying" (2015, p.17). Simply put, for Bechtel and other mechanists, the relata of mechanistic levels are componentially related.

However, there is an important qualification to be inserted here, one that has already been alluded to. This is that levels of mechanisms are local in application. Unlike

Newell's system levels, Bechtel's levels do not form a well-delineated hierarchy. The reason for this is that part-whole relations do not neatly track organizational features such as location, size, shape, or motion. Whether or not one component is contained within a given mechanism depends on the nature of the investigation.

To return to the previous example, if one is interested in explaining maze learning, then one has to invoke properties of spatial memory, and then, in turn, decompose the activities of the spatial memory into properties of the hippocampus and pyramidal cells. However, this does not mean that pyramidal cells are always at a lower level than the hippocampus. It might be that a given pyramidal cell is a component in some hippocampal mechanisms but not others. Ordering mechanisms across various levels is relative to some explanatory endeavour; explanatory context is paramount.

An illustrative aid is provided in Table 1.1. The table shows how each of the four accounts answers the level questions. As can be seen, each account applies the levels metaphor slightly differently. Although there is some overlap, each answers the relata and criterion questions distinctively. This fact will prove interesting in the analysis to come, as confusion often arises when it is assumed that each application of the metaphor is the same.

Level Questions	Marr	Pylyshyn	Newell	Bechtel
Relata Question	Theoretical Languages	Descriptive Vocabularies	Components of Systems	Activities of Mechanisms
Criterion Question	Questions Answered	Generalizations Captured	Prediction	Part-whole Relations

Table 1.1 The two level questions as applied to four accounts of levels in cognitive science.

1.4 Analyzing Levels

Having outlined and related several of the major accounts of levels, it is time to provide a more general analysis. To this end, consider two general points that hold of the various applications of the levels metaphor.

1.4.1 Epistemic vs. Ontic Relata

The first is that there seem to be two general types of relata that the levels metaphor applies to. The first is ontic; the second is epistemic.

On the ontic application, the levels metaphor 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 of the metaphor, levels are part of the configuration of the world, picking out the various features constitutive of cognition.

On the epistemic application, the levels metaphor 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, the levels metaphor functions to provide perspectives for understanding cognition. It offers a means of enhancing an investigator's epistemic situation. Levels have more to do with the theoretical pre-occupations of investigators than with the organization of the world.

What is interesting about these two general applications is that when viewed as extremes they form a continuum on which various accounts can be located. On the one hand, there are those views, such as Marr's, that tend toward an epistemic application. On the other hand, there are also those views, such as Newell's, that fall closer to an ontic application. An account's positioning depends on the relata it implicates. When applied to linguistic items, the account has an epistemic bent. When applied to properties of the world, the account has an ontological character. Used as an organizing axis, these two opposing applications furnish an interesting point of comparison. Figure 1.1 provides an illustration.

Figure 1.1 Various accounts of levels arranged along a spectrum of epistemic vs. ontic relata.

Notice also that Bechtel's view provides something of an interesting middle ground. The account straddles the two poles. Recall that on Bechtel's view levels of mechanisms are both ontic structures and observer-dependent entities. Although mechanistic activities are real phenomenon in the world, the arrangement of mechanisms into different levels depends on the explanatory goals of the investigation. Levels gain partial ontic status via their application to components and activities, but also remain epistemically dependent in virtue of their relativeness to explanatory context.

One implication of applying the levels metaphor to epistemic versus ontic relata is that it changes how levels are taken to relate. 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. Consider each in turn.

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. 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.[9] For 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 about some entity rather than the entity itself.

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.[10] 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.

The point to note is that the supportive relation holds between different theories or linguistic items and not ontological structures. The relation is an explanatory one.

---

[9] For an instructive discussion of the many faces of realization in cognitive science, see Wilson & Craver (2007).

[10] There is obviously much to both the standard model and theory reduction in general, but this description should suffice to convey the basic point required here.

Whether or not the reduction is successful is beside the point. 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 emphasis 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.

The takeaway is that there are those views that take the relata to stand in metaphysical realization relations and there are those views take the relata to stand in explanatory realizations relations. That depending on the answer given to the relata question changes follow for how levels relate.

1.4.2 Two Uses for Levels

The second dimensions deals with how accounts use the levels metaphor, to what purpose or end the levels metaphor is generally put. Again, there are two general types of usage: one ontic, one epistemic.

On the ontic usage, the levels metaphor attempts to provide accurate descriptions of how cognition is organized. Levels are in the business of providing ontic descriptions. The metaphor seeks to furnish representative accounts of how cognition is 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. The levels metaphor functions to describe the organization and structure of cognition as described by the different system levels. Thus, according to the ontic usage, the levels metaphor has a distinctly ontic role to play in discussion of cognition.

Contrast this position with a second use of the levels metaphor that views levels as addressing how cognitive phenomena should be studied. On this epistemic usage, 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). The levels metaphor acts as an epistemic tool for researchers 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. The levels metaphor functions to expand understanding by simplifying and structuring discussion into three levels.

When viewed as extremes, these two general usages form a second continuum on which various accounts can be located. Figure 1.2 provides an illustration.

Figure 1.1 Various accounts of levels arranged along a spectrum of epistemic versus ontic usage.

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

1.4.3 A Conceptual Framework for Understanding Levels

The fact that levels talk can be sorted along two dimensions of analysis is not an insubstantial result. This is because it allows for the creation of a conceptual framework. A framework that can be used to the map the logical space of levels talk. Figure 3 provides an illustration.

Marr's account occupies the bottom left corner of the space. This is because it applies the levels metaphor 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 views in the middle. These views tow a more neutral line, conceiving of levels as at once ontic and epistemic in character to some degree. One might wonder why Pylyshyn is not positioned closer to Marr. After all, his relata are also linguistic entities.

The reason for this is that Pylyshyn also takes his levels to supervene one on top of the other in virtue of tracking distinct componentially arranged parts. In at least this way, Pylyshyn is a bit closer to more ontically oriented views in how it conceives of the using levels talk. Finally, there is Newell's view in the top right hand corner. Newell's position here is owed to his adherence to an almost entirely ontic application of the levels metaphor.

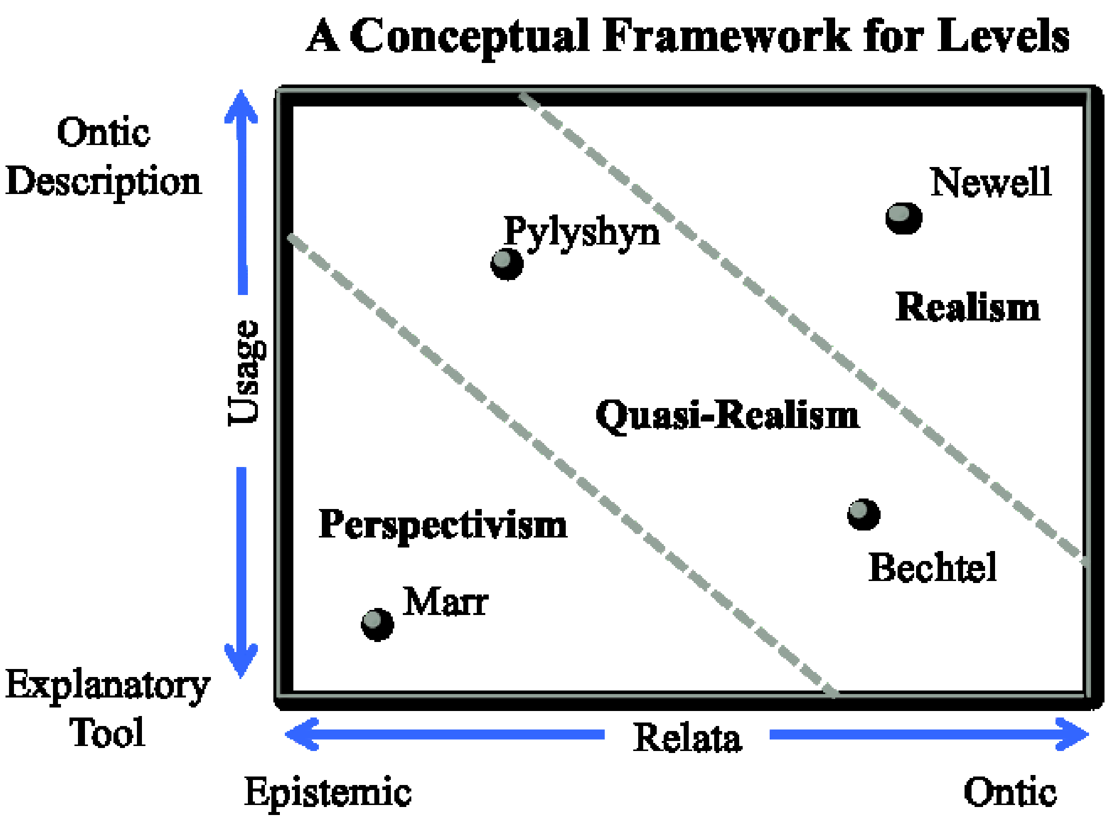

Figure 1.3 Various applications of the levels metaphor arranged according to commitment to epistemic versus ontic relata and usage.

What is interesting about this framework is that it maps many of the possible applications of levels talk in cognitive science. It is both descriptive and generative. It both positions various accounts of levels according to their application of the metaphor, while at the same time leaving it open as to how other views might apply the metaphor.

One might worry that the current framework uses two dimensions where only one is necessary, given that both dimensions include an ontic versus epistemic comparison. The reason two dimensions are more effective than one is that two dimensions allow the current framework to pull apart middle positions such as Bechtel or Pylyshyn's that deploy the levels metaphor both epistemically and ontically. Using only one dimension would fail to effectively convey the difference between positions such as these.

The framework also goes beyond mapping current and potential views. This is because it also reveals three more general, orienting positions that might be held with respective to levels. The first is what might be called the *perspectivist* or *perspectival* view. According to this view, levels denote epistemic activities; levels pick out different descriptive frameworks that separate investigation into its most explanatorily potent parts. Included here are views such as Marr's. The second is what might be called the *quasi-realist* view. Here levels straddle the epistemic/ontic divide, at once denoting ontic categorizations and observer-dependent perspectives. Bechtel and (possibly) Pylyshyn fall under this second category. Finally, there is the *realist* view. On this construal, levels reflect divisions in the organization of nature, or in this case cognition; they mark separations among layered entities or systems. Newell is included under this third heading.

Note that these three views are not actual positions any particular author holds, though, in fairness, some authors do get pretty close. Rather, each view represents a sort of center of gravity or idealization of many applications of the levels metaphor held with respect to levels in cognitive science. The positions can be used as general umbrella categories in which to explore different applications of the levels metaphor 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 under 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 is that the various views are just different sides of the same coin; that with additional work each can be distilled into a single 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. At this stage, it is just not clear.

What is clear is that there are three desirable features to the present framework. First, the framework helps to clarify relations between different accounts. 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. Marr fails 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 make for a poor fit.

Understood within the context of the present analysis, the issue becomes 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 brain organization. Churchland, Koch and Sejnowski have confused Marr's epistemological application of the levels metaphor 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. Thus, the framework can play a therapeutic role in discussions of levels.

Second, the conceptual framework provides interpretative benefits. 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. An 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, going back to Newell's account, if one is dubious about the appeal to evolution on the grounds that such arguments are underdetermined, then one can argue that the realist position is importantly under-motivated; that application of the level metaphor to cognitive structures is too quick. Or, to take the contrary position, if one is concerned that the perspectival view makes levels-talk arbitrary then the current analysis provides the direction in which to mount this challenge – for example, that perspectivism fails to get at the real structures and organization of cognition.[11] The framework helps to simplify discussion by cutting to the core elements of disagreement.

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

A quick qualification is necessary in closing. The preceding analysis is not meant to be the definitive guide to conceptualizing levels within cognitive science. It is entirely possible that other dimensions of analysis exist on which applications of the levels metaphor 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, I hope, is one concrete step toward making of further sense of the levels metaphor in cognitive science.

1.5 Conclusion

To recap, first, I introduced the topic of levels in cognitive science using three orienting questions. This helped to pull out some of the major motivations for talking about levels, along with providing a brief history of levels in cognitive science. Then, I outlined four influential applications of the levels metaphor. After this, I showed how two level questions could be applied to each account. This revealed two dimensions of analysis on which different accounts could be compared. When put together, these two

[11] I will say more about both these concerns in Chapter 2.

dimensions generated a conceptual framework, one that could be used to productively map the logical space of levels, along with clarifying and simplifying otherwise confusing discussion.

With all this in place, the task of the next chapter is to come to some consensus on which view of levels is ideal for cognitive science. To preview just a bit, the argument is that the preferred view of levels within cognitive science is a form of *pluralism*. The claim is that the realist, quasi-realist, perspectivist positions are all valid applications of the levels metaphor so long as it is clear how, when, and why each is being applied. There is no one true application of the levels metaphor for cognitive science.

Chapter 2 – Leveling the Field

2.0 Chapter Overview

The goal of Chapter 2 is to investigate whether there is a preferred conception of levels for cognitive science. After outlining and contrasting several possibilities, it is argued that cognitive science is best served by a 'pluralistic approach' to levels. The chapter concludes by expounding some of the advantages associated with adopting the pluralistic approach.

2.1 Three Conceptions of Levels

Near the end of Chapter 1, I introduced three general conceptions of levels: realism, perspectivism and quasi-realism. In what follows, I explicate each conception in a bit more detail, comparing and contrasting the respective views. After demonstrating the advantages and disadvantages of each conception, I argue that cognitive science is best served by adopting a 'pluralistic approach' to levels.

2.1.1 Realism, Perspectivism, and Quasi-Realism

In order to come to some consensus about how we should think about the levels metaphor in cognitive science, it will be informative to first begin with a general overview of the three main conceptions. This will help to tease apart some of the broad differences between the views (see also Ch.1, sec. 1.4.3).

First, consider realism. According to realism, the levels metaphor 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.

Second, consider perspectivism. On the perspectival view, the levels metaphor 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 (Ch. 1, sec. 1.4.2).

Third, on the quasi-realist view, the levels metaphor 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.

2.1.2 The Semantic Value of Levels

One way to further explicate the three conceptions is to focus on the semantic value of levels talk. As may have been noticed, each conception of levels conceives of what it means to say that 'x is at a higher level than y' slightly differently. The semantic value of statements about levels changes according to the conception being adopted.

First, consider that according to realism statements such as 'x is at a higher level than y' are literally true or false. Levels talk is truth apt. This reason for this is that realism conceives of levels in terms of a correspondence between how the world is and how statements describe it to be – I will say more about this in a bit. Insofar as there is a correspondence between statements about levels and the way cognition is actually organized, levels talk has a truth-value.

Contrast this view with perspectivism. On the perspectival view, the semantic value of levels is pragmatic. What is important is that discussion aids in investigation – for example, by facilitating prediction (see, e.g., Dennett, 1991). Perspectivism denies strict correspondence between the descriptions it offers and the structures it purports to describe. Talk of levels operates to enhance or augment inquiry more than it does to describe the structure of cognition. The semantic value of levels talk takes on a more instrumental favour on the perspectival view.

Finally, consider the quasi-realist view. According to quasi-realism, statements about levels have a truth-value but only relative to a specific explanatory context. For example, on Bechtel's version of the view, levels only relate to a given mechanism. Though it is true to say that some items are at a higher level than others, this is only relative to a given mechanistic investigation. The appropriateness of assigning truth-values to statements about levels extends only to particular mechanistic investigations. It is not possible to make global assessments of levels across mechanisms.[12]

To clarify the point further, consider an analogy with the semantic view of scientific theories (Savage, 1990; Mormann, 2007). The semantic view holds that scientific theories consist of the specification of a class of structures and the languages used to describe those structures. As van Fraassen describes it: "To present a theory is to specify a family of structures, its models; and second, to specify certain parts of those

[12] See also Craver (2015).

models (the empirical substructures) as candidates for the direct representation of observable phenomena (1980, p.64). Scientific theories involve specifying different ontic structures in a particular modeling language. Some authors opt for a realist interpretation, holding that the semantic relationship is truth-apt. Others opt for a more instrumental approach, claiming that what matters is empirical adequacy.

There is an analogous relationship in the structure of levels talk. Different applications of the levels metaphor take different stands on how theoretical languages apply to ontic structures. Each of the three main applications adopts a different stance on the semantic or truth-value relationship that obtains between levels of explanation and levels of organization: some realist, others more instrumentalist. Similar to the semantic view of scientific theories, the different conceptions of levels each hold to alternative pictures of how levels relate to the world, which colour the semantic value of statements about levels.

2.1.3 Levels of Explanation versus Levels of Organization

At this point one might ask: What is it that produces the differences in semantic value? What is it that makes it so that the realist conception differs with respect to the value of levels talk compared to the perspectival or quasi-realist views? The answer lies in how each view conceptualizes the relationship between the relata being sorted and the framework used to describe the relata. The point is best framed in terms of the relationship between levels of explanation and levels of organization.

One way to think about the difference between an ontic versus epistemic application of the levels metaphor is as tracking differences between levels of organization and levels of explanation. Levels of explanation involve the stratification of discussion into different explanatory or descriptive modes, while levels of organization involve the description of organizational features of nature; in this case, cognition. In identifying particular items for sorting (levels of organization) each conception of levels also makes certain assumptions about how to talk about the items being sorted (levels of explanation).

Consider the realist view first. On the realist view, the relationship between levels of explanation and organization is one of strict correspondence. Each level of explanation picks out or tracks a different level of organization. As each level of organization layers one on top of the other, so too do the different levels of explanation. Newell's account offers something pretty close to this view. Newell's "system levels" mark different levels of organization, while his "cognitive bands" identify different levels of explanation (see Ch.1, sec.1.2.3). The various cognitive bands identify different collections of system levels such that as one ascends up the different system levels one also ascends up the different cognitive bands.

Another example comes from Oppenheim-Putnam's (1958) layer-cake model of science. According to Oppenheim and Putnam's model, the various sciences map to different levels of mereology. Different collections of scientific theories are ordered so as to correspond to six ontological kinds of organizations: societies, organisms, cells, molecules, atoms, and elementary particles. Each kind of organization is defined by part/whole relationships among types, e.g., elementary particles are parts of atoms, atoms are parts of molecules, molecules are parts of cells, etc. Because each kind of organization is assigned a distinct science, each science is layered one on top the other. The layer-cake model maintains a neat hierarchical ordering between levels of organization and levels of explanation, though in Oppenheim and Putnam's case the metaphor applies to sciences as a whole rather than theories of cognition. The general picture is depicted in Figure 2.1.

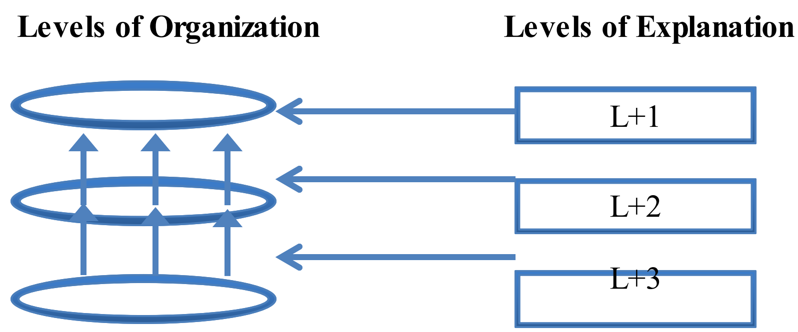

Figure 2.1 The relationship between levels of organization and levels of explanation according to realism.

Next, consider the perspectivist view. The perspectival view denies the neatness of the correspondence relationship. In contrast to realism, perspectivism views the relationship between levels of explanation and levels of organization as a bit looser. At times the relationship is one-to-one, but often it is one-to-many or many-many.[13] On at least some interpretations, Marr's discussion of levels holds to this view – another plausible example is McClamrock (1991).[14]

Consider, as an example, Marr's discussion of the visual system. Marr claims that the function of vision is object recognition. Studying vision requires analyzing how the visual system computes three-dimensional objects from two-dimensional retinal images. It involves identifying a particular input-output mapping. But whereas object recognition is one plausible characterization of the function of vision, another plausible characterization is that the visual system is for delivering motion detection; that the function of vision is to specify the position of objects in the environment. At this finer

[13] Craver makes a similar point in discussion of Oppenheim & Putnam's layer-cake model (2015, p.7)
[14] For an example from philosophy of biology see Wilson (2005).

grain of analysis, investigation is computational, but it applies to the components relevant to motion detection rather than object recognition.

The point is that there is no neat one-to-one mapping between levels of explanation and levels of organization on the perspectival view. Even granting that there may be some overlap between the two characterizations, the same level of explanation (the computational level) can track two different levels of organization (components of the visual system). The same level of explanation can be applied to different levels of organization. Another way to frame the point is to say that different levels of explanation are ‘grain sensitive’ (see Shapiro, 2000).

Craver’s (2015) discussion of neuroscience offers another example. Craver points out that the subject matter of neuroscience is not restricted to any one level of ontology. Neuroscience encompasses everything from ions and individual neurons to brain regions and behaviours of whole organisms. Like many other sciences, neuroscience often investigates phenomena that cross several distinct levels of organization. As Craver points out: “The multilevel structure of contemporary science emphasized in nearly every corner of the special sciences is not best understood as a hierarchy of theories” (2015, p.8). The various cognitive sciences are better viewed as operating at multiple levels of investigation, particularly in the case of neuroscience. Once again, Craver’s example helps to demonstrate the breakdown of any neat one-to-one mapping. Figure 2.2 provides an illustration.

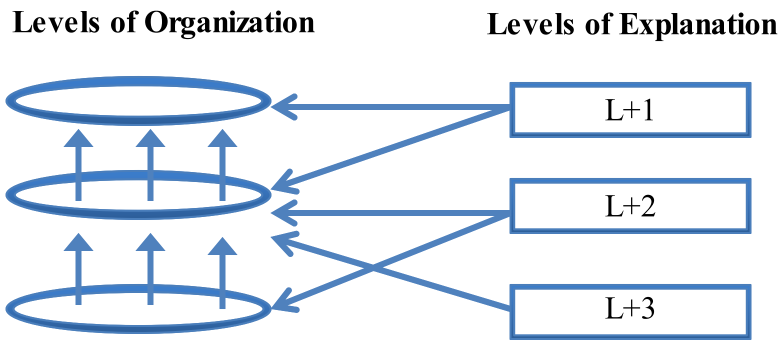

Figure 2.2 Levels of organization and levels of explanation according to the perspectivist view of levels.

Finally, consider the quasi-realist view. For authors such as Bechtel, various components, functionally arranged, form different mechanisms. Mechanisms are constituted at different levels of organization. However, this does not mean that mechanisms are simply products of nature. They are not read off the organization of the world. Rather, mechanisms are partly constituted by the various explanatory strategies used during investigation. Levels emerge in the movement between the objects being sorted and the way objects are discovered. This is what furnishes the ‘quasi’ element of the quasi-realism view. As Bechtel writes: “The multi-level nature of mechanisms can be couched in terms of a trichotomy of explanatory strategies” (2007, p.62). Levels emerge as a product of applying different kinds of explanatory strategies to mechanisms. The different explanatory strategies used during analysis stratify mechanisms into multiple layers. Exactly how many levels there are and how they are individuated is an empirical question, one whose answer ultimately depends on the phenomenon being investigated and the explanatory strategies being used.

Consider spatial memory as an illustrative example. Spatial memory can be explained by appealing to the activity of the hippocampus. The hippocampus is part of the mechanism that underlies spatial memory. Sodium channels are also part of the hippocampus, being essential building blocks in neurons. However, even though sodium channels are partly constitutive of hippocampal neurons, this does not entail that all sodium channels are part of the mechanism that supports spatial memory. Standing in mereological relations does not guarantee that a given component will be a higher or lower level. Instead, what determines a component's place in a given hierarchy of levels is the explanatory role it plays within the investigation. Because hippocampal neurons have a functional role to play in explaining spatial memory, they can be arranged at different levels. The leveled structure of mechanisms emerges via the explanatory relations between the functionally relevant parts. Figure 2.3 provides an illustration.

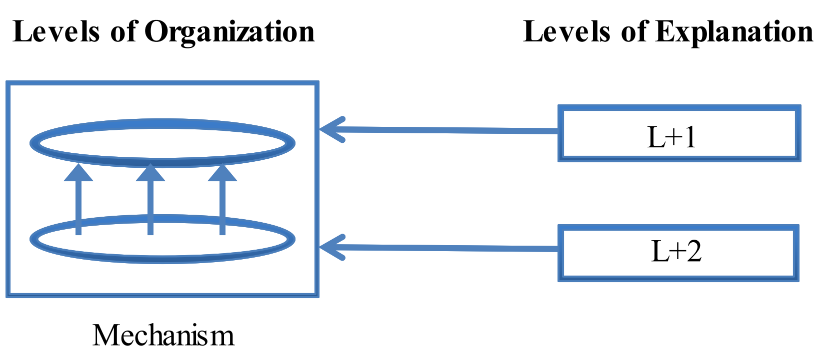

Figure 2.3 The relationship between levels of organization and levels of explanation according to the quasi-realist view.

Notice that what emerges on the quasi-realist view is probably best not viewed as 'levels', at least not when compared to the view espoused by realism. The notion of levels that emerges on the quasi-realist view has more affinity with different 'kinds of explanation' than with a mechanism whose structure is hierarchically arranged into layers. The levels constitutive of mechanisms on the quasi-realist view have more in common with perspectives than with stratified levels of mechanisms. Nonetheless, because authors such as Bechtel and Craver present their accounts as committed to the reality of levels as constitutive of mechanisms, my discussion will continue to reflect this presentation.

More generally, what this discussion highlights is that there are important assumptions at work behind the scenes for each conception of levels. Substantive assumptions lie just beneath the surface when it comes to how relata are discussed and how levels relate. Because each conception takes a different stand on the relationship between the mode of description and object of description, there are implications for what levels talk means and how levels function. The quasi-realist view orients the levels metaphor to mechanistic investigation, the realist view uses the levels metaphor to make claims about the structure of cognition, and the perspectival view uses the levels metaphor to unlock different explanatory insights.

2.2 The Question of Levels in Cognitive Science

So far I have shown how each conception of levels stakes a competing claim on what levels are, how levels should be thought to relate, and what function levels should have within cognitive science. The question to consider now is given the fact that levels play an important role in theorizing about cognition, which conception, if any, is best suited for cognitive science. The remainder of the chapter is occupied with providing an answer to this question.

2.2.1 Why the Debate Matters

Before delving too deeply into the debate, consider first why the debate might be important. One reason is that it may have methodological implications for how researchers study cognition. Consider, for example, how the different conceptions of levels might influence decompositional analysis. Functional decomposition is the process of explaining cognitive capacities via decomposition into sub-capacities. It involves breaking down a system into its component parts using task analysis. Some see functional decomposition as a defining part of scientific psychology (Cummins, 1983; Bechtel & Richardson, 1993).

One way to understand what happens in functional decomposition is that researchers move downward through different levels of organization as analysis progresses. That when a cognitive capacity is analyzed, researchers break down the capacity into sub-capacities that reside at lower levels of organization. This picture comports well with the realist conception of levels. Each round of functional decomposition analyzes systems or entities into further levels of organization, because cognition is stratified into different layers.

Notice how the picture changes on the perspectival view. If it is assumed that levels denote perspectives rather than hierarchically organized layers, then figuring out when analysis moves from one level of organization to another becomes more difficult. Though functional analysis is still involved in decomposing different-sized systems, it is less clear how those systems relate in a global hierarchy – for example, whether a system's sub-capacities reside at "lower" levels than the capacities from which they are decomposed. This, of course, does not make decompositional analysis futile. It simply means that interpretation of what decompositional analysis amounts to changes according to the conception of levels that is adopted.

The general point is that the centrality of the notion of levels entails effects for other areas of cognitive science. It means that depending on how the levels debate turns out consequences can emerge for how to think about other aspects of the field. Settling the question of levels may affect other decisions researchers need to make about cognition, such as in decompositional analysis. These considerations should begin to motivate the view that the debate is significant.

2.2.2 Realism

Time to return to determining which conception of levels is best suited for cognitive science. In what follows, the main arguments supporting each conception of levels are teased out and evaluated – in the interest of neutrality, I abstract away from the particularities of an author's view wherever possible.

To begin, consider the main line of support for realism, which seems to stem from largely evolutionary considerations (see also Ch.1, sec. 1.2.3 and Ch.1, sec. 1.4.3). Newell (1990) offers a representative example, writing: "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). On Newell's view, the realist conception of levels is vindicated on the basis of the hierarchical nature of cognition as a product of evolution. Since it is more likely that nature will build complex systems incrementally rather than all at once, it is correspondingly more likely that entities will be organized into layers of systems.[15] The argument looks something like the following:

[15] Herbert Simon (1962) also holds something similar. In fact, Newell credits Simon with first offering something like this argument.

1. Systems are arranged hierarchically for reasons of stability.
2. Thus, stability dictates the existence of a hierarchy of levels.

There are three points to note about this argument. First, the argument essentially boils down to a claim about evolution favouring systems constructed via the addition of modular parts. As entities grow in size and scale, assemblies and subassemblies combine to form larger and larger complex entities, ones whose productive powers are not strictly reducible to the level below. Second, support for premise (i) comes from the fact that different types of behaviour are controlled by different systems and that attempts to build complicated systems without first building stable subassemblies often fail (see Newell (1990, p.117-8) for details). Third, the argument's application is wide in scope. It holds not only of intelligent systems, but also of living systems as a whole. Since the evolutionary advantageousness of stability holds across all of nature, it is likely to be found in a whole host of systems in addition to intelligent ones.

Given the reasoning in favour of realism, consider two lines of reasoning that speak against the realist argument and the conception it supports. First, the argument's appeal to evolution is underdetermined. Evolutionary considerations are neutral between hierarchical and non-hierarchical architectures. Consider, for instance, the reverse version of the argument: Evolution favors massively distributed systems because such designs are stable in virtue of being more flexible and resistant to degradation – a massively connected neural network is a good example. Since evolution is more likely to select those systems that balance flexibility of function with resistance to damage, stability dedicates a non-leveled structure to intelligent systems.

Prima facie, this argument appears plausible enough. The problem is that the argument does not actually show what evolution favours. It is more than certainly possible, and even perhaps plausible, that evolution *might* favour non-hierarchically organized system. But this does not show that evolution actually *does* favour such organizations. Given the right circumstances, evolution might favour multi-level hierarchical systems, but it just as well might favour single-level massively distributed systems. In the abstract, evolutionary considerations cut both ways.

The realist might respond by arguing that empirical evidence decides in favour of a multi-level hierarchical structure. For example, perhaps ablation of parts shows that cognitive systems are leveled. As components are removed functionality is preserved. Does this not demonstrate the leveled structure of cognition? The problem, again, is that the evidence cuts both ways. The same set of facts can be used to support non-leveled designs.

Neural nets, for example, can lose discrete functionality without being hierarchically organized into levels – it would be odd to say that the nodes or connection weights of a connectionist network exist at higher or lower levels. If neural networks adopt coarse coding (where each node is responsible for coding a particular feature of the network input), then disruptions to relatively small numbers of connection weights or nodes can result in relatively specific functional deficits (see Dawson, 2004, Ch.3). What this shows is that when taken in isolation, evolution is insufficient to establish the kind of hierarchicality required to prop-up the realist conception of levels. Neural networks can exhibit the same kind of discrete functional losses without being hierarchically organized. Of course, this does not mean that evolutionary considerations are irrelevant. Rather, it means that such considerations are not decisive. Insofar as the realist conception relies on evolutionary considerations of stability, the position is under-motivated.

Second, there is an issue with how realism conceives of the relationship between levels. Or, to be a bit more specific, there is a problem with how realism conceptualizes the realization relation that holds between different levels of organization. The primary issue is that there is an emerging body of evidence from the biological domains that suggests that realization in complex systems is more complicated than the realist conception supposes.

Consider, for example, what properties determine an instance of fish camouflage. Some of the relevant properties include parts of the fish, such as those that determine the overall shape of the fish or the arrangement of pigmented cells; others include parts of the environment, such as the presence of leaf litter in the stream; and some further still include properties of predators, such as predators being aquatic or hunting by sight (Sazima et al. 2006).

The point to note is that the set of properties that determines the presence of camouflage involves a complex combination of properties, including everything from fish and environment parts to the predator population. A whole host of same or different level properties feature in the realization of higher-level entities or systems, such as being camouflaged. Higher-level properties are not exclusively realized by lower-level properties. The realization base of higher-level properties, such as camouflage, extends beyond properties one level below, such as genotypes. As Potochnik points out: "The properties that figure into lower-level explanations often fail to be the true supervenience bases of the properties that figure into higher-level explanations" (2010, p.129).

Potochnik (2010, 2012) calls this form of realization: "complex" realization – for a complimentary view, see Wilson's (2001, 2004) notion of "wide" realization. Complex realization threatens realism's appeal to hierarchicality. It entails that components at various levels of organization can form the realization base of a given entity. It shows that constitutive realization relations cross multiple levels. Or, to phrase things in the language of philosophy of mind, the subvenient realization base can include entities at levels of organization other than one level beneath the realized phenomenon. Insofar as the realist conception underestimates the complexity of the realization relation in hierarchically ordered systems, it offers an impoverished account of levels. The relationship between systems, cognitive or otherwise, is more complicated than the realist conception supposes. Figure 2.4 illustrates the difference in realization views.

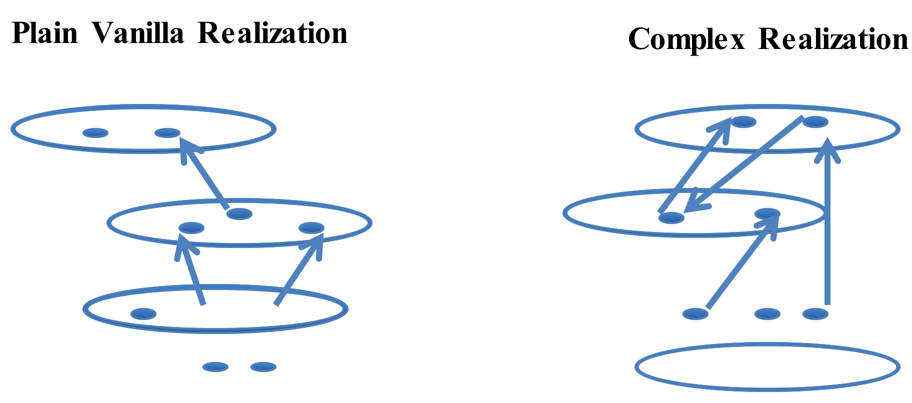

Figure 2.4. Complex versus regular realization.

It is important to be clear about what this discussion shows. Complex realization does not show that the entire realist conception of level is suspect. It does not show that components are not compositionally related or that the entities or activities are not "real" things in the world. Rather, what complex realization points out is that any global ordering of entities into levels is going to be difficult given the complex relations in which constitutive elements often stand. This result is not insubstantial. Given that realism is committed to hierarchicality, complex realization goes some way to casting doubt on the view's completeness.

2.2.3 Perspectivism

Consider perspectivism next. Support for perspectivism comes from a largely epistemic basis. The perspectivist view gains traction in virtue of what the levels metaphor offers cognitive investigation. This is usually cashed out in one of two ways. One is to appeal to what the metaphor offers specific discussions of cognitive phenomena. The other is to appeal to what the metaphor offers cognitive investigation as a whole.

Consider the specific case first. For some, the motivation for talking at different levels comes from the flexibility it provides in thinking about cognitive phenomena. Marr, for example, writes: "[i]t becomes possible, by separating explanations into different levels, to make explicit statements about what is being computed and why to construct theories" (1982, p.19). Different descriptive vocabularies or languages, such as neurology or information-processing theory, provide different insights into cognitive phenomena. With each new level comes a new way of conceptualizing cognitive phenomena. This is why authors such as Marr (1982) or Dawson (1998) talk about new "levels of understanding." Talk of levels is justified insofar as it provides novel epistemic insights and opens up new lines of investigation. A perspectival application of the levels metaphor allows researchers to "carve nature at her joints."

The other way the levels metaphor is sometimes justified is by appealing to levels as an organizational tool. The best example here again comes from Marr (1982).[16] On Marr's tri-level approach, research begins by characterizing cognitive systems in terms of information-processing problems. This involves specifying the function being computed and the environmental factors that constrain the choice of function. When one moves down a level (to the algorithmic level), focus shifts from the function as a whole to the steps and representations used by the system to solve the information-processing problems. The focus shifts to how the system carries out a given input-output mapping. Move another level down further still (to the implementational level) and focus shifts again. At this bottom level, researchers focus on what neurological mechanisms support the system's function; what physical states realize the computation under investigation. By stratifying investigation into different levels, the levels metaphor provides a road map for research. As investigation proceeds from one stage to the next, it moves from one level to another.[17]

One concern that speaks against perspectivism is how the view conceives of the movement between levels. The problem is that perspectivism seems to have a difficult time specifying what it is that constrains movement between different levels once a many-many relationship between levels of organization and explanation is conceded.

Consider, as an illustrative example, how one might explain a calculator. One approach would be to explain what was going on at a physical level – for instance, to explain the calculator's internal workings, such as its integrated circuit. A second approach would be to explain the calculator's basic architecture and procedures – for example, whether it worked in binary or used principles of Boolean algebra. A final

[16] See Dawson (1998, Ch.2) for another example.
[17] Dawson (1998, Ch.2, 2013, Ch.1) also offers something like this justification for levels.

approach would be to describe the relation between the calculator's input and outputs – that is, to describe what function the calculator computes, e.g., '4 - 2 = ' or '4 + 2 = '.[18]

Consider what might constrain these different types of explanations according to realist and perspectival views. On the realist view, movement between the three types of descriptions is constrained by different levels of organization. Implementational explanations target a low level of organization in the calculator (the level of transistors and resistors), while computational descriptions address an abstract level of organization (the level of digits and functions). Unique levels of organization determine the appropriateness of the different types of explanation.

Contrast this with the perspectival view. On the perspectival view, there is a many-many relationship between levels of organization and levels of explanation. A computational or algorithmic level description might address any given level of organization (recall Ch.2, sec. 2.1.3). As Shapiro puts it: "[O]ne investigator's computational level is another's algorithmic level; and one investigator's algorithmic level may be another's implementational level" (2000, p.441). There is a no unique level of organization determining the appropriateness of a given type of explanation.

The problem is that this characterization makes it unclear as to why different types of explanations are more or less appropriate when applied to different levels of organization. In spite of the fact that it seems natural and compelling to describe the calculator using multiple types of explanations, perspectivism fails to explain why this is the case. Whereas on the realist view movement between levels of explanation is constrained by the unique levels of organization each explanation tracks; on the perspectival view, no such similar constraints exist.

[18] These three approaches correspond roughly to Marr's tripartite approach.

Consider an analogous problem in debates over the computational theory of mind. Searle (1990), for instance, argues that computational theory is bankrupt when applied to cognition, since the conditions for ascribing computational descriptions are too weak – for example, he claims that a wall can implement a computer program because there is some pattern of molecular movements in the wall that is isomorphic to the formal structure of a computational description. Putnam (1988) defends a similar thesis, though somewhat less extreme. Computational descriptions are problematic because they are a purely interpretive or instrumental type of explanation.

A similar thought underwrites the concern about perspectivism. Even though we are naturally drawn to different types of descriptions – as in the case of the calculator – perspectivism does not provide an account of why we should find such explanations compelling. In spite of the naturalness of different types of explanation, perspectivism leaves matters unclear as to what constrains movement between different levels, particularly given its many-many mappings view between levels of organization and levels of explanation. Since one is able to move between different levels of explanation, especially higher-level ones, without restriction, there is a legitimate question as to why such explanations should be thought of as properly applying to cognitive systems in the first place.

One might worry that this concern points to a deep-seated problem in perspectivism. This is not the case. The argument is only directed specific instances of perspectivism, such as some interpretations of Marr (1982) or Dawson (1998), not perspectivism as a whole. The claim is simply that, as a systematic approach to understanding levels within cognitive science, certain embodiments or versions of the perspectival view are importantly incomplete. As will be shown in Chapter 3, perspectivism still has something important to add to the discussion of levels, it is just not adequate in every respect.

2.2.4 Quasi-Realism

Finally, consider the quasi-realist view. Two sets of considerations seem to favour quasi-realism. One is that it does not oversimplify the relationship between levels of organization and levels of explanation. The other is that it retains the 'realness' of levels talk.

First, recall that quasi-realism holds the relationship between levels of organization and levels of explanation to be more nuanced than the realist and perspectival views propose. Craver, for example, writes: "Levels of organization are, in a sense, levels of explanations, given that explanations for different topping off phenomena will often decompose the system into altogether different parts within parts" (2015, p.23). Levels of organization are, in part, constituted by explanatory strategies. Levels of organization are not distinct from the languages that describe them. Why this is important for authors such as Craver and Bechtel is that it suggests that quasi-realism may be better equipped to address some otherwise confusing examples.

Consider, for instance, the relationship between the nephra and the hippocampus. Nephrons are the basic structural and functional units of the kidney. They regulate the concentration of water and soluble substances in the blood, such as sodium salts (Pocock, 2006). The hippocampus is an elongated set of ridges on the floor of each lateral ventricle of the brain often associated with learning, memory and emotion.

On the realist view, the hippocampus and nephra are located at the same level because they are roughly of the same size or shape, for example. In contrast, according to the quasi-realist view, the nephra and hippocampus are non-localizable with respect to each other, because there is no overlap in properties or function. Whereas realism slots the two structures to the same level in virtue of considerations of location, size, shape, or motion, quasi-realism view reserves judgment in virtue of the fact that the two structures are not related part/whole. Craver (2015) views this non-monolithic element as an advantage of quasi-realism, as it allows quasi-realism to avoid complicating cases such as having to say that the nephra and hippocampus are at the same level despite having completely different functions, properties and constitutive material.

The point is more forcefully brought out in an example from Oppenheim and Putnam's discussion of the mereological hierarchy. According Oppenheim and Putnam, societies are the largest organizational kind in nature. Given this ordering, one might reasonably ask where hurricanes slot into the hierarchy. Are hurricanes, for example, at the same level as societies? None of the criterion that the realist might appeal will sort hurricanes and societies to the same level, whether location, size, shape, or motion. In contrast, the quasi-realist can withhold judgment. Because the two entities are not related part/whole, they cannot be said to reside at the same or different levels. Thus, the quasi-realist view seems to have the advantage of not forcing counter-intuitive conclusions on its adherents. To Craver (2015) this suggests that, "the mechanistic view seems to capture many of the intuitions that accompany the idea that world is organized into levels but without many of the objectionable elements of other applications of the levels metaphor" (2015, p.23). By not oversimplifying the relationship between levels of explanation and organization, quasi-realism purports to avoid complicating cases that may trip up other views.

The quasi-realist view is also supported by the fact that it retains the “ontic” status of levels talk. Recall that one of the concerns for perspectivism is that levels of explanation are left untethered to cognitive structures. Quasi-realism seeks to avoid this concern by applying the levels metaphor to organizational features of mechanisms: properties and activities. Applying the metaphor in this way allows quasi-realism to side-step concerns over having an exclusively “theoretical” status. As Wright and Bechtel write: “[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 (2007, p.57).

In spite of its advantages, there are issues with the quasi-realist view. One of the primary issues is that even though quasi-realism is able to say when two items *are not* at the same level, it fails to provide an account of when two entities *are* at the same level. Recall that quasi-realism maintained that items are related across levels via part-whole relations: “Levels of mechanisms, like part-whole levels generally, are not monolithic divisions in the furniture of the world. Levels of mechanisms are defined only within a given part-whole hierarchy” (Craver, 2015, p.18). One item can only be said to be at different level from another item if each is related part to whole – for example, the hippocampus is at a different level from spatial memory because the hippocampus is part/whole related to spatial memory.

The problem with this appeal to part/whole relations is that it produces a purely negative principle for organizing items across levels. It only allows quasi-realism to say when two items are not at the same level, but not when two items are at the same level. This is a problem insofar as we think that a complete account of levels should be able to do both. An account should be able to not only say when two items are not at the same level, but it should also say when two items are at the same level – realism, for example, does not have this problem as it can use considerations of size to organize items to the same level. In making discussion of levels local and non-relatable, quasi-realism undermines its ability to provide a positive ordering of items across levels.

The quasi-realist might bite the bullet here. They might argue that there is nothing wrong with claiming that one cannot specify how items relate within level; it is just the logical consequence of the view.[19] However, there seems to be a legitimate question whether something so basic in how the world is understood should be so easily given up, i.e. the intuition that comparably large entities reside at the same level of organization. Theoretical equilibrium is important. Common sense cases should only be dropped when it comes at a corresponding benefit. However, it is not clear that quasi-realism has enough going for it to make this tradeoff – at least the realist and perspectivist views would certainly claim otherwise. I leave it to others to decide. The point here simply is that as desirable as quasi-realism is on some grounds, it is not without its drawbacks. Similar to realism and perspectivism, quasi-realism remains only partially adequate as an account of level within cognitive science.

2.2.5 Qualifications and Implications

In concluding this section, it will be worthwhile to qualify the preceding argumentation. The preceding discussion is not intended to undermine the various

[19] To some extent, this is the position adopted in Craver (2015)

conceptions of levels in their entirety. Rather, the discussion is meant to show that each type of application of the levels metaphor comes with its own conceptual baggage, that each view is importantly incomplete. Even though each conception adds something important to the discussion of levels, no one view is sufficient in every respect.

This conclusion provides some impetus for questioning an important assumption underlying the discussion: that there is only one answer to the question of levels in cognitive science. In showing that each conception is importantly incomplete, there emerges a substantive question as to whether there needs to be only one conception of levels for all of cognitive science. Thus, the discussion to follow speaks against a monistic and in favour of a pluralistic approach to levels.

2.3 The Pluralistic Approach to Levels

The position I intend to adopt with respect to levels is a form of pluralism. The claim is that there is no one single verdict when it comes to utility or conceptual soundness of the levels metaphor in cognitive science. The various applications of the levels metaphor are justified so long as each is carefully deployed.[20] In what follows, I lay out what I take to be the pluralistic approach to levels. This is achieved by (i) drawing an analogy to pluralism in philosophical logic and (ii) explicating the pluralistic account more substantively, outlining how the view seeks to accommodate each of the three general conceptions of levels. I conclude by outlining four general desirable features of the pluralistic view. These serve to motivate the view a bit more generally.

[20] The beginnings of this position can be found in Craver (2002, 2007, 2015)

2.3.1 Situating Pluralism

To introduce and motivate pluralism within discussion of levels, a helpful analogy can first be drawn to pluralism within philosophical logic. Recent years have seen a rise in the debate over whether there is a single correct logic for all of philosophy and science. Historically, the received view was that "classical logic" was the correct view to take with respect to understanding truth and natural language. It was the "only game in town", to borrow a phrase from Jerry Fodor. However, starting in the early 20th century a number of alternative logics began to develop – for example, intuitionist, paraconsistent, and quantum logics. These alternative logics prompted some to ask whether classical logic was indeed the only logic in town. The ensuing debate lasted for years; in many ways it is still going on.

One particularly interesting response to this debate, in some ways owed to the longevity of the debate itself, is that some have begun to question whether there is even a single correct logic to be had at all (Beall & Restall 2000, 2001, 2006; Restall 2002). These "pluralists" argue that as long as one is careful about how and when one is using different logics, it is perfectly acceptable to say that there are several "correct" logics. It is not that classical or paraconsistent logics are the correct choice for philosophy and science, it is that both are correct given the right circumstances. This is the contrast between *monist* and *pluralist* views of logic.

Consider the following simple case. The disjunctive syllogism P v Q, not Q, therefore P, might be interpreted in several ways. From the point of view of classical logic, the relation of logical consequence holds between the premises and conclusion. This is what makes the argument valid. P follows as logical consequence of the denial of

Q. However, on at least some logics, such as paraconsistent logics, since one can deny the law of excluded middle, which says that something cannot be both true and false at the same time, the argument is invalid. According to logics where both A and not A can be true, disjunctive syllogisms fail to hold as a valid argument form. From the point of view of paraconsistent logics, P does not follow from the denial of Q. The conclusion does not follow as a result of the premises.

On the monist view, only one of these two interpretations is correct. Either the classical view is right or the paraconsistent view is right, but not both. There is a substantive disagreement over the nature of logical consequence. In contrast, on the logical pluralist view, both types of interpretations are correct. It is reasonable to say that P does and does not follow from P v Q/not Q. Pluralists think that notions such as 'follows from' or 'valid' are inherently unsettled or vague. They may be cashed out in any number of ways.[21] As Beall puts it:

> What is important, here, is that for a pluralist it does not follow that A ∨ ~A is not true, or even, not *necessarily* true. It is consistent to maintain that all of the truths of classical logic hold, and that all of the arguments of classical logic are *valid* with the use of constructive mathematical reasoning, and the rejection of certain classical inferences. The crucial fact which makes this position consistent is the shift in context. (2000, p.13)(original emphasis).

There is no single correct usage. What determines the correctness of alternative interpretations is the utility of the formal system as a whole, not anything about the nature of logical consequence. What is important is that one can do things with paraconsistent or intuitionist logics that cannot be done with classical logics – for example, paraconsistent logics can help to make sense of inconsistency in a way that treats inconsistent

[21] Of course, I am vastly oversimplifying the debate, but for illustrative purposes such a description should suffice.

information as informative; something classical logics struggle with (see, e.g., Beall, 2006). Pluralism is motivated by the idea that different logics serve different purposes.

2.3.2 Pluralism and Levels

What I want to suggest is that an analogous case holds for the concept of levels within cognitive science. Because of the inherently unsettled and vague nature of levels, the preferable position to adopt is that different applications of the levels metaphor are equally appropriate given the right circumstances. Call this the pluralist approach to levels. In cases where researchers are more concerned about the theoretical power of various explanatory or descriptive frameworks, a perspectival view of levels is preferable. In cases where researchers are interested in a clear delineation of functional systems, a realist or quasi-realist view is preferable. It is the larger aims of research that are important. Some research programs may be more aptly suited to a realist application, others to a perspectival or quasi-realist application.

What I am suggesting is that instead of trying to patch up each view or fit every application of the metaphor into one usage or another, we instead acknowledge the limitations of each conception. We recognize that the utility of the levels metaphor lies not in any one single application but, rather, in being used to a multitude of ends.

Consider, as an example case, the study of vision. As we saw, on the perspectival view it makes sense to talk about the visual system under various descriptive or explanatory guises – for example information processing theory or neurophysiological accounts. The function of the visual system can be aptly described and explained via recourse to several different levels of understanding or explanation. Additionally, we saw that it makes sense to talk about the visual system as containing several different levels of organization. In both physical and informational terms, it makes sense to say that the visual system is arranged into distinct components that conspire to produce its various functions. Whether its neurons, optic nerves, function integrators, each of type of component fits together to sustain the system's functionality. It makes sense to taxonomize and stratify different sets of components according to functional type. Further still, we saw that it makes sense to treat different groupings of various components as discrete units, even as they crisscross different levels of organization. In line with quasi-realism, it makes sense to group different token components into distinct functional mechanisms.

What this shows is that each application of the levels metaphor captures an important part of the story about the visual system. Whether it is in terms of stratification by unit type, identification of functional mechanisms or unification of explanatory frameworks, each kind of levels talk brings something to the table. Each conception of levels offers something distinct to cognitive science.

The realist view emphasizes the centrality of capturing generalizations produced by different sets of components using distinct theoretical vocabularies (e.g., computation or neurology). The perspectival view emphasizes the advantages of applying the same level of explanation – for example, the computational level – to multiple levels of organization, which helps to informatively structure cognitive investigation. The quasi-realist view helps to direct attention away from the epistemic/ontic dichotomy and toward investigation of mechanisms, which allows it to ground discussion of levels in concrete structures while still appealing to different kinds of explanation. Each application of the levels metaphor adds something unique to the study of cognition.

Problems start when any one application is stretched too far, when any one application is thought to apply to all of cognition – for example, when the realist's one-to-one mapping between different kinds of explanation and organization is thought to apply wholesale. In some ways, the desire to find a unifying application of the levels metaphor is understandable. The flexibility and fecundity of the metaphor lends itself to such an effort. However, no single view puts all the pieces together perfectly. Perspectivism loses sight of the ontic status of cognitive structures; realism oversimplifies the explanatory and realization relations between levels; quasi-realism preserves level realness but at the expense of losing the notion of levels more generally – recall what was said about the notion of levels looking more like 'kinds of explanations' (see Ch.2, sec. 2.1.3).

The seductiveness of the levels metaphor is in some sense in its worst enemy. Because the metaphor is applied so effortlessly, several distinct but related usages naturally emerge. Each brings with it something important to discussion, but each does so at the expense of clarity and precision of one or more kinds. The solution is to drop the global requirement of the levels metaphor, to no longer think that the levels metaphor needs to apply uniformly in all cases. The realist, perspectival, and quasi-realist views are not engaged in an all-or-nothing struggle.

Notice also that this is a negative claim. The view is not that the levels metaphor in general is suspect but, rather, that cognitive science should move beyond treating levels talk homogenously. Only when the monist view is dropped as an overarching framework for understanding levels can each conception be given its proper place within the field.

2.3.4 Objections

To help explicate the current approach further, consider two objections that might be raised against the view. One is that pluralism does not leave enough room for disagreement between different views. That pluralism undermines the chance for substantive rivalry between different conceptions of levels. It problematically subscribes to an "anything goes" policy.

Consider Beall's response to this type of concern. Beall's answer is to suggest that disagreement among different logics exists only relative to a given discussion. He writes: "There is scope for rivalry and disagreement, when the meaning of the basic lexicon is settled. The moral of our pluralism goes as follows: Once you are specific what your logic is meant to do, there is scope for genuine disagreement" (2000, p.16).

A similar point holds for levels. Pluralism allows for disagreement, but only one relative to application of the metaphor within a specific context. For example, in Chapter 3, it is shown that the levels metaphor can be applied to cognitive modeling. Within that discussion, there emerges a substantive question as to whether a realist-based account can effectively accommodate several types of cognitive models. It is argued that a perspectival-based account can accommodate some types of cognitive models better than a realist-inspired account. What this discussion helps to show is that the substantive disagreement over which conception does the best job accommodating certain types of models. There is a real rivalry between the two views. Pluralism makes room for disagreement; it is just one that is contrastive and local in character.

Another worry is that pluralism makes it unclear how researchers are supposed to decide between different applications of the levels metaphor. There is a question of how one is to know which conception of levels is most appropriate within a given context. Here is the analogous problem in logic: "The objection, in short, is that pluralism has no non-arbitrary way of choosing which among its many (alleged) logics ought to govern our reasoning in the debate between monism vs pluralism. This, in turn, is supposed to be a mark against pluralism, given that monism, for obvious reasons, has no trouble at all in specifying which among its logics *it* chooses" (Beall, 2001, p.6).

Beall's answer is to suggest that logic and reason come apart. That so long as reasoning isn't exhausted by logical systems, one will still be able to decide between different logics principled. Nothing quite so lofty is required in the case of levels. Deciding appropriateness in the case of levels talk requires researchers to look for what is *not* accommodated by a given usage. If alternative applications of the levels metaphor in various contexts carry with them different constraints, then researchers need to be able to recognize and identify those constraints in order to determine whether what is excluded is something important to investigation.

To return to the previous example, if a realist-based approach to modeling excludes certain types of models such as those based on population-level data, then whether those models are desirable ones to include may depend upon further considerations about what is relevant to a framework. Are prototype models worth retaining, for example? Using the levels metaphor under a pluralistic vision entails that there is no a single procedure for recognizing and controlling different applications pre-investigation. An answer to appropriateness only comes after a conception has been deployed. The 'principledness' of deciding between different conceptions is (i) local to a given investigation and (ii) largely comes after the fact when contrasting different applications.

2.3.3 The Benefits of Pluralism

To conclude, consider some advantages associated with adopting pluralism. First, adopting the pluralistic approach allows researchers to accommodate many of the already-existing usages of levels. The pluralistic stance allows researchers to take seriously the fact that there are multiple accounts of levels actively being deployed within cognitive science. Instead of trying to fit various accounts into any one single model, the pluralistic stance allows acknowledgement of the diversity of views already operating with cognitive science.

Whether its Marr's tri-partite approach or Newell's system-level view, there is something to said for incorporating as many views as possible. In at least this way, the pluralistic conception accords more closely with general state of opinion in cognitive science. As the survey in Chapter 1 makes clear, application of the levels metaphor within cognitive science runs the gamut of usages. The pluralistic approach offers minimal resistance in accommodating much of what cognitive science already takes for granted.

One might respond by pointing out that accommodating peoples' views is secondary to establishing what is true. That it shouldn't matter how many views are accommodated but which view is the correct one to adopt. This is concern is certainly well founded, but misplaced in the present context. This is because it begs the question in favour of the monist view. It mistakenly assumes that there can only be one correct application of the levels metaphor. However, if pluralism is true, then the two desiradata (accommodation and truth) are not mutually exclusively. In having motivated the incompleteness of each conception of levels, the pluralism should appear a viable option over and above a purely monistic approach.

Second, a pluralistic conception of levels does not rule out particular kinds of levels talk *a priori*. One of the issues faced by the realist conception of levels, for example, is that it failed to accommodate the possibility of complex many-many relations between explanatory modes and ontic strata. This is only a problem if realism is taken to be the "correct" view of levels. If pluralism is adopted, then the conceptual boundaries of levels talk are opened up. Because pluralism eschews the single model fits all view, it retains the flexibility to move and readjust the conceptual boundaries of levels talk – for example, if in one context the perspectivist fails to deal the right kind of theoretical results, one can shift to the realist view. Movement is based on the specifics of inquiry rather than any narrow definition of what levels supposed to be or must relate.

Third, pluralism helps to make sense of stubborn debates concerning particular level accounts. For example, there is a longstanding issue over how to interpret Marr's tri-partite account of levels. Interpretations range from radically epistemic and internalist friendly accounts to distinctly metaphysical and staunchly externalist in character (e.g., Burge, 1986). In some sense the disagreement is owed to the vagueness of Marr's original presentation. However, more substantively, the disagreement is also owed to the fact that various authors often operate implicitly with a certain vision of levels in mind. Whether realist, perspectivist or quasi-realist in nature, background assumptions often do much to color and sustain interpretative debates.

On the pluralistic approach, this problem is ameliorated. If adopting a perspectival interpretation leads to greater methodological and explanatory insight, then that interpretation is desirable – arguably, this is the case when researchers study more modular systems, such as perception (see, e.g., Bermudez, 2005, Ch.1). Yet, in contrast, if adopting a realist interpretation yields stronger ontic claims (such as in the case of claims about neurological mechanisms), then that interpretation is desirable. Pluralism provides the flexibility to interpret levels talk along several lines; interpretative debates become less poignant once it is conceded that what matters more is how a view is put to use.

Fourth, and finally, pluralism provides a substantive and instructive guide to practicing cognitive scientists. Though I will say more about this in Chapter 3, it will suffice for now to point out that the pluralistic approach can assist in cognitive modeling. Consider recent discussion of levels and cognitive modeling. Sun (2005), for example, argues that arranging models into a hierarchy of levels helps to simplify and structure discussion. The problem, as we will see, is that Sun's framework limits the range of models that can be included, because it implicitly relies on a realistic conception of levels. In response, I construct a perspectival-inspired framework. This framework accommodates several of the models that strain Sun's account. This discussion acts as a further vindication of pluralism, as it helps to demonstrate the usefulness of alternative conceptions of levels. It shows that perspectivism, and by extension pluralism, adds something constructive to the discussion of modeling.

2.4 Conclusion

A few things should be clear by now, I hope. First, there are three general positions one may adopt with respect to levels: the realist, quasi-realist and perspectivist views. Each takes a unique stance on what levels are, how levels should function within cognitive science, and how different levels relate to one another. Second, the three conceptions of levels are each motivated by a distinct set of considerations: some rely on considerations based on evolution; others the utility of levels as epistemic lens. Third, each conception has different strengths and weaknesses based on the arguments and assumptions it employs. Fourth, and finally, the pluralistic approach to levels stands as an important and viable alternative to the realist, quasi-realist and perspectivist views. It is both a well-motivated and desirable approach, particularly in contrast to a monist approach.

I will close with this. One might have the nagging concern that the preceding discussion has been sparse on specifically "cognitive" content. It may not exactly be clear why all this discussion of levels should be properly thought of as part of cognitive science rather than something more akin to an exercise in philosophy of mind or science. There is some merit to this concern. In response, Chapter 3 shifts focus to more traditionally cognitive science topics. In particular, it looks to demonstrate the practical and theoretical role of the levels metaphor in the context of cognitive modeling. This discussion attempts to connect the preceding two chapters to topics closer to the heart of cognitive science.

.

Chapter 3 – The Levels Metaphor and Cognitive Modeling

3.0 Chapter Overview

The aim of Chapter 3 is to examine the contribution of the levels metaphor to cognitive science. The chapter focuses on the role of the metaphor within cognitive modeling. It proceeds in three parts. First, cognitive modeling is introduced and examined. Second, Ron Sun's "Hierarchy of Four Levels" is outlined and related to the realist conception of levels. Third, after laying out problems with Sun's realist-based account, the pluralistic approach is re-introduced and its benefits demonstrated.

3.1 Cognitive Modeling

At the end of Chapter 2, I promised a concrete discussion of levels. I claimed that greater clarity with respect to levels would substantively contribute to the practice of cognitive science. I agreed, in other words, to engage in "philosophy in" rather than "philosophy of" cognitive science. Chapter 3 attempts to deliver on this promise. The goal is to show that the levels metaphor has a substantive role to play within cognitive modeling. After introducing the topic and outlining a specific example, it is argued that cognitive modeling can make productive strides if it adopts the pluralistic approach to levels. This discussion serves to once again motivate the importance of the levels metaphor within cognitive science and vindicate the pluralistic approach to levels.

3.1.1 Computational Cognitive Modeling

In its most general form, cognitive modeling attempts to approximate human or animal cognitive processes for the purpose of prediction and comprehension. The aim is to develop models that accurately represent the behaviours and interactions of cognitive agents. Traditionally, there are three varieties of cognitive models within psychology and cognitive science: mathematical, verbal-conceptual and computational. Mathematical models present relations between variables via complex mathematical equations. An example is Rescrola-Wagner's (1972) learning rule for classic conditioning. Verbal-conceptual models describe entities, relations and process in natural language. Baddeley's (2012) tri-component model of working-memory offers a good example. Computational models provide process-based accounts using algorithmic descriptions. Schank and Abelson's (1977) script model of planning provides an illustration. [22]

Each of these models has played an important role within psychology and cognitive science. This chapter focuses exclusively on computational cognitive models. The reason for this is that computational modeling is by a margin the most commonly used variety of model within cognitive science. It has been a mainstay of cognitive science since its earliest beginnings in the 1970s.

In computational cognitive modeling, researchers provide detailed descriptions of the mechanisms and processes of cognition in terms of the science and technology of computing (Luce, 1995; Turing, 1950). Computational models use detailed computer simulations to describe and specify the algorithms and programs that underlie cognitive processes. Computational theory provides a theoretically flexible and expressively powerful tool for exploring cognition.

There are three types of models that tend to dominate computational cognitive modeling. Probably the most well know variety is what many call "classic" or "symbolic" models (Minsky, 1981; Anderson, 1983). Symbolic models use rule-governed symbol manipulation to model complex cognitive processes, such as problem solving or concept

[22] Mathematical models can be seen as a subset of computational models.

formation (e.g., Newell & Simon, 1976). They rely on complex data structures and well-defined procedures to simulate higher-order or central cognitive processes, usually adhering to a strict distinction between structure and process.

A second popular type of model is what's called neural network or connectionist models. Gaining popularity in the mid-1980s, connectionist models move away from the symbol-based accounts in favour of a more distributed, massively parallel architecture (Rumelhart et al., 1986). Connectionist models often simulate complex information processing tasks using dozens of interconnected functional units.[23] These processor units usually perform simple functions, such as integration or step-wise functions. When connected in large quantities, the functional units are capable of solving extremely difficult problems, such as the traveling salesman problem (see Dawson, 2004). Unlike their symbolic counterparts, connectionist models do not adhere rigidly to the structure/process distinction; instead employing "distributed representations" (see, e.g., Clark, 1989). This has led some to suggest that connectionist networks offer a more biologically plausible model of cognition (Smolensky, 1988).[24]

A third more recent class of computational cognitive models is dynamical models (Beer, 2000; Chemero, 2008). Dynamical models represent non-linear interactions between components using complex differential equations. They emphasize the interactive and developmental character of cognitive processes. Because of this, dynamical models are sometimes thought to better capture the temporal dimensions of cognition. Furthermore, because the language of dynamical systems does not appeal to static structures, dynamical models are sometimes also thought to be more parsimonious

[23] For a contrast between symbolic and connectionist models see Dawson (1998).
[24] See Fodor & Pylyshyn (1988) for a different opinion.

than their representation-heavy cousins (van Geller, 1995). There are also hybrid models that blend each of the preceding three types of models together in varying degrees (see, e.g., Sun & Bookman, 1994).

Unifying these various types of models is the assumption that cognitive processes and system can be described and explained using computational simulation. Each model seeks to account for cognitive processes in terms instantiating computational programs. Though each disagrees with the others on the details, the consensus is that the language of the modeling is computation and that successful progress in studying cognition requires developing and implementing runnable computer programs.

3.1.2 The Explanatory Value of Modeling

In terms of explanatory value, the utility of cognitive modeling lies in its ability to produce computational simulations capable of making predictions that match human data. Sun and Ling (1998) lay out three types of correspondence for computational models.

First, there is correspondence to behavioral outcomes. When computational models match the same types of behaviours under roughly similar conditions as human subjects, models are behaviorally equivalent. An example is a model that makes roughly the same kinds of decisions as humans in different decision-making scenarios.[25] Second, when models capture the same deterioration of performance on a given task as human subjects, models correspond qualitatively. An example is matching primary and recency effects in memory task scenarios. Finally, when models match the same behaviours, as indicated by some performance measure, models can be said to correspond quantitatively.

[25] Pylyshyn (1984) calls this "weak equivalence."

Point-to-point matching of learning curves is an example of this last sort of correspondence.[26]

The development and comparison of computational models and empirical data offers an important way of exploring human cognition. By matching the "fit" of human data with cognitive models, researchers are able to establish systematic relationships that can uncover the computational underpinnings of cognition. As Sun (2008) points out: "Computational modeling contributes to the general, theoretical understanding of cognition through generating mechanic and process-based descriptions that match human data" (p.3). Finding a good fit between human data and computational models reveals important details about the mechanisms and processes underlying cognitive processes.[27] Successful cognitive modeling requires a healthy exchange with empirical data.

Cognitive models are also importantly different from cognitive architectures. Whereas cognitive models are computational simulations of specific cognitive processes, cognitive architectures are generally considered to be more domain-generic representations of cognition. Cognitive architectures carry with them structural assumption about the basic organizational features of cognition, assumptions which need not be included within a given cognitive model. They are the general schemas in which particular models are realized. Unlike more circumscribed cognitive models, cognitive architectures provide concrete frameworks for cross-domain analysis. Examples include ACT-R (Anderson, 1983), CLARION (Sun, 2002b), and NENGO (Eliasmith, 2002). This distinction is worth noting because, as will be important later, thinking about the relation

[26] Inferences made on the basis of quantitative and qualitative data can reveal the procedures by which a cognitive process is performed. When combined with behavioral equivalence, Pylyshyn (1984) calls this "strong equivalence."

[27] For discussion of how data can be used in cognitive science and modeling see Dawson (2013, Ch.3).

between models is a separate enterprise from thinking about the basic building blocks of cognition. Discussion in this chapter is restricted exclusively to the former. [28]

3.1.3 The Theoretical Status of Modeling

Computational cognitive modeling has largely occupied two roles within cognitive science. One is as a theoretical enterprise operating in close connection to cognitive psychology. The other is as an extension of AI – that is, the project of implementing cognitive processes in non-neural mediums, i.e. computers. Though both roles have been important to cognitive science, focus here centers on the more methodological and theoretical dimension of modeling.

One question that has generated substantial debate in recent years within modeling concerns the theoretical status of computational cognitive models (see, e.g., West & MacDougall, 2014). Some have wondered whether computational cognitive models are best viewed as tools for theory building, instantiations of theories, or full blown psychological or cognitive theories. Providing an answer to this question is important because it has important constraints on model interpretation.

For example, to take just one scenario, if models turn out to be instantiations of theories, then two models derived from the same cognitive theory cannot stand in direct conflict, even if they produce contradictory results. They are simply two different interpretations of how to implement a given theory computationally. One can be more faithful to a theory than another, but neither, strictly speaking, is more or less correct. It is right to say that theories are true or false, but not models. Given the importance of this

[28] For an interesting discussion of cognitive architecture see Sun (2004).

debate, exploring several possibilities will be instructive, particularly as it sets up the critical discussion to follow.

As mentioned, three options appear possible when it comes to the theoretical status of computational cognitive models. One is that models are tools for theory building. If a model confirms a specific hypothesis, then it corroborates the theory; if it does not, then it can lead to the formulation of a new verbal-conceptual theory. Since cognitive models are good at generating new phenomenon, such as data, they can be said to either falsify or verify existing cognitive theories. Models are important instruments for theory building.

A second view is that cognitive models are instantiations or embodiments of pre-existing cognitive theories. Computational cognitive models allow researchers to fill out verbal-conceptual theory by implementing them as a computer programs (Anderson & Lebiere, 1998). Computational cognitive models allow greater conceptual clarity and precision in how verbal-conceptual theories are understood. They allow researchers to fill in the gaps left by linguistic descriptions. In implementing computational cognitive models, researchers are able to identify and address problematic assumptions or gaps in a theory that might otherwise be ignored.

Baddeley's (2012) component model of working memory offers a good example. Conceptually, Baddeley's model is rather straightforward. Working memory consists of four elements: a central executive, a phonological loop, a visuo-spatial sketchpad and an episodic buffer. However, when implemented as a computer program, there are aspects of Baddeley's model that are underspecified – for example, how fast information is bound within the central executive or by what algorithms the central executive moves between different retrieval strategies. In the actual implementation of Baddeley's model as a computer program, important functional details are fleshed out, ones that might have otherwise been glossed over in a purely verbal-conceptual articulation of the theory. Models are not only simulations for corroborating or generating theories they also detailed elaborations and enhancements of theories.

Finally, there is the view that computational models are cognitive theories in their own right. Computational models are not only built on top of or for the sake of pre-existing cognitive theories they are also stand-alone cognitive theories. Here is Sun (2008) describing the view: "No verbal-conceptual theory completely specifies the computational mechanisms involved, let alone the dynamic processes that may emerge. Thus…the computational framework is, in essence, just another language for presenting a more rigorous and/or detailed theory" (p.7). The idea is that because cognitive models go beyond what is contained with the initial verbal-conceptual theories and generate novel predications and data they count as cognitive theories in their own right. Because there is no clear mapping between the elements of a given cognitive theory and the elements of a computational simulation, it is best to view cognitive models as inspired by but not reducible to pre-existing verbal-conceptual theories. In many ways, cognitive models exhibit an independent theoretical status.

3.2 The Levels Metaphor and Cognitive Modeling

One particularly interesting aspect of the preceding discussion is the question of how models relate once it is assumed that models are species of theories. Consider a contrast with the models-as-theory-instantiation view.

On the models-as-theory-instantiation view, models relate via the theories they implement. This means that models stand in secondary relations to one other, as the theories implemented by the models are the main units relating to truth. Models cannot be said to be true or false, only the theories they implement. Contrast this with the models-as-theories view. On this view, models are stand-alone theories. They have their own distinct theoretical status. No prior theoretical relations connect models to the structures in the world they purport to represent.

What is interesting is that if models have an independent theoretical status then there is a substantive question as to how they might relate. Unlike the models-as-theory-instantiation view, models can stand in competitive relations to the truth. This means that there is illegitimate question of how to test the theoretical adequacy of models. In some ways, this is the old question of how to judge empirical adequacy reframed in the context of modeling (see van Fraassen, 2001).

3.2.1 A New Hierarchy of Levels

One potential answer to the relation question is to say that models relate in virtue of being placed within levels of a hierarchy. This is the position adopted by Ron Sun (2005, 2008). Here is Sun describing the view:

> [A] scientific theory of cognition requires the construction of a hierarchy of different levels with consistent causal descriptions from a low level through a series of intermediate levels a high-level phenomena…Scientific understanding depends upon the selection of key elements of cognitive phenomena and the creation of models for such elements at appropriate levels (2005, p.634).

On Sun's view, organizing cognitive models into different levels resolves the issue of how to relate different models. As each model is slotted into a different level of the hierarchy, the relation between various models becomes visible. One reason Sun's account warrants sustained consideration is that it represents one of the most detailed and thought out applications of the levels metaphor within cognitive modeling. It offers an important first step toward applying the levels metaphor effectively and substantively within modeling.

Sun motivates his account by contrasting it with more traditional approaches: "Instead of focusing exclusively, as in the traditional theory of multiple levels of analysis, on what we consider to be minor differences in the computational tools that we use (e.g., algorithms, programs, and implementation), we want to focus on the very phenomena that we are supposed to study" (2005, p.619). Traditional approaches to modeling are too focused on the *tools* of modeling rather than the *phenomena* being modeled. Sun singles out Marr as an example of the former approach.

On Marr's account, computational cognitive modeling occurs at three levels: the computational, the algorithmic and the implementational. Each level addresses a different aspect of modeling: the computational level addresses function; the algorithmic level addresses representation and procedure, etc. The problem, according to Sun, is that a computational cognitive modeling based on Marr's account can gloss over important differences among the phenomenon being modeled. This is because the distinction between computational, algorithmic, implementation levels does not necessarily reflect a substantive division between phenomena in the world. It is a distinction largely drawn from computer science. Thus, it does not necessarily have purchase when applied to cognition. The result is that models founded on something such as Marr's tripartite account can fail to reflect substantive differences in cognitive phenomena.

A more successful approach is to focus on what phenomena or entities a model attempts to explain. Computational cognitive modeling needs to be first and foremost grounded in the phenomena it purports to represent. Once researchers are clear about the phenomenon being addressed, they are better positioned to see how and when different models relate. Sun argues that by reorienting focus to phenomena a hierarchy of levels can be constructed, one that can be used to guide and integrate various cognitive models. Sun dubs this "A New Hierarchy of Four Levels" – for brevity, the account is shortened to simply the "Hierarchy of Four Levels."

Consider each of Sun's four levels. First, at the top, there is the sociological level. Researchers at the sociological level model large-scale objects. These may include collective agent behaviour, inter-agent processes, socio-cultural processes or environment-agent interactions (both physical and socio-cultural) (Vygotsky, 1986). Sun argues that without understanding and incorporating socio-cultural factors, researchers are in an impoverished position when it comes to modeling cognitive processes. Modeling that ignores the sociological level is only partial or incomplete.

An example of modeling at the sociological is Robert West's model of human game playing (West & Lebiere, 2001; West et al., 2006). West's model investigates how human game playing often deviates from optimal strategies – for example, how people often attempt to detect sequential dependence in opponent's play in order to predict their next move. West's model embodies a sociological level model because it represents the relationship between agents rather than the processes within an individual.

Second, there is the psychological level. Phenomena modeled at the psychological level include the familiar cast of characters from Psychology (including folk psychology) and AI, such as individual behaviour, beliefs, concepts, and skills. An example case is Osherson et al. (1990). For Osherson et al. (1990), inductive reasoning is modeled by assessing propositional statements according to the similarity between premise and conclusion categories.

For example, consider two inferences: (i) Mice have property X/All mammals have property X and (ii) Horses have property X/All mammals have property X. The category mammal in the conclusion covers both mice and horses. The model can use this superordinate category to make judgments about similarity. This allows the models to handle various types of simple inferences. Models at the psychological level incorporate data structures representing "knowledge" and "beliefs." Investigation involves collecting and interpreting behavioral data, usually via functional analysis.

Third, there is the componential level. The componential level deals with intra-agent components. It addresses the underlying components responsible for behaviour at the psychological level. Sternberg's (1968) classic memory model is a prime example. Using data collected from verbal protocol analysis, Sternberg constructed a memory model that mirrored the exhaustive search strategies used during digit span recall tasks. Sternberg's model simulated the intra-agent processes of memory search. Componential level models simulate the mechanisms and processes of cognition by using the symbolic, connectionist or dynamicist approaches. Analysis is both structural in that it looks at components and functional in that it looks at cognitive capacities.

Finally, there is the physiological level. This is lowest level of the hierarchy. The physiological level models the biological substrate or implementational of intra-agent processes. Disciplines such as biology, physiology, cognitive neuroscience, etc., are relevant here. The physiological level supplies details about architectural and material elements. Leigh and Zee's (2006) model of eye movement offers an illustrative example. Leigh and Zee's model simulates how a neural network in the ocular-motor system calculates and controls eye movement based on eye-velocity and eye-position.

The importance of the physiological level is that it allows researchers to identify, in computational terms, the mechanistic primitives responsible for behaviour. Similar to Marr's implementational or Pylyshyn's biological level, the physiological level addresses the physical entities responsible for sustaining and realizing cognitive systems and processes.

In terms of operation, Sun's account has something like the following structure. First, a phenomenological description of social, cultural and cognitive processes is provided. This description provides the initial specification of the to-be-modeled phenomena. This is followed by data collection. A variety means (e.g., psychological experiments) are used to uncover data relevant to the modeling process – an example being error rates or response times. These data reveals the causal relationships obtaining among the entities under investigation. Once the causal relationships are identified, a process-based description is constructed. The process description constitutes the basic sketch of the computational model. After the process description, specification of the relevant algorithms and data structures is detailed. This further fills out the model. Finally, the model is implemented in a runnable computer program. After several simulations, data are generated and model adjusted. Figure 3.1 provides a depiction of Sun's four levels.

Level	Object of Analysis	Type of Analysis	Computational Model
1	Inter-agent processes	Social/cultural	Collections of agents
2	Agents	Psychological	Individual agents
3	Intra-agent processes	Componential	Modular construction of agents
4	Substrates	Physiological	Biological realization of modules

Table 3.1 Sun's Hierarchy of Four Levels

There are few points to note about Sun's account. First, model construction flows from the phenomenon to the model. Constructing models at each level of the hierarchy requires: (i) identifying a given phenomenon, (ii) providing a computational specification and (iii) implementing a computational model in a runnable computer program.

Second, models can be constructed downward and upward from the lowest to highest levels of the hierarchy. Any number of models can be constructed at a given level, though analysis and model implementation is largely limited to one level at a time. The main constraint is that computational models address the appropriate process, entities and causal relations.

Third, the 'phenomena' detailed in Sun's account reside at different levels of organization. Though not explicit, the types of phenomena Sun appeals to are situated at different levels of nature. Cognition is divided into distinct layers, each demarcating entities of varying size, complexity and casual powers. As we will see, this aspect of Sun's account pushes the view closer to that of the realist conception of levels.[29]

---

[29] The realist may, of course, use other organizational features beside size, such as location, shape or motion – size is just the most the commonly appealed to usually.

Fourth, and finally, computational cognitive modeling is constrained by factors and processes at different levels of the hierarchy. Though in many cases models can be constructed without reference to lower levels, consideration of lower and higher level factors is critical to successful modeling. Higher and lower level factors place constraints on modeling.

There are a number of advantages to adopting the Hierarchy of Four Levels, according to Sun. One is that the hierarchy is able to integrate models across levels. This is because the hierarchy allows researchers to engage in cross- and mixed-level analysis. In cross-level analysis, researchers integrate details from lower levels into higher-level models – for example, a researcher might include details about cognitive biases from the psychological or componential level into a model about religious rituals at the sociological level. In mixed level analysis, modelers build lower level principles or rules into higher-level models. Because causal relationships at a higher level sometimes generate predictions inconsistent with observations at a lower level, a model must occasionally include rules that indicate the conditions under which a more detailed model can supersede a higher-level model – see Sun et al. (2005, p.628) for details. The upshot is that models situated at different levels can be integrated to greater or lesser degrees. The hierarchy allows researchers to integrate otherwise disconnected models depending on the amount of multi-level constraints each model accommodates.

A second advantage is that the hierarchy places important constraints on model viability. By reorienting model construction to phenomena, investigation is made more principled. It places model testing on solid theoretical footing. For example, since it is inadequate to look for only the neurological correlates of cognitive processes, any model that tries to do so will have difficulty finding a home within the hierarchy. Models that try to integrate phenomena across levels without incorporating the relevant intermediate levels, such as the componential level, are likely to be unsuccessful; or at least less so than those that do. Because the hierarchy arranges models into different levels based on considerations of size and complexity, the hierarchy is effective at sorting models according to their theoretical viability. To find a place within the hierarchy, models must track appropriate kinds of phenomena – that is, the models must track phenomena that can be identified at one of the four levels of the hierarchy.

3.2.2 Realism and Modeling

With Sun's framework outlined, I turn now to detailing how the account relates to the realist conception of levels. I begin by outlining the core features of the realist view.

Recall from Chapter 1 that the levels questions revealed two core features of the realist view: (i) that levels sort ontic structures standing in part-whole relations and (ii) that sorting ontic structures into different levels often uses a criterion of prediction or causality (see also Ch.2, sec. 2.1.1). Consider what happens when (i) and (ii) are used as conditions for realism.

With respect to (i), Sun's accounts sorts models according to the differently sized processes and entities they track, along with their attendant causal relations. As one moves up the hierarchy of levels, one moves from smaller to larger phenomena, each with its own distinct ontological standing. Sociological phenomena are larger than psychological or componential phenomena and have their own distinct ontic standing. Notice also that the objects and causal relationships at higher levels are defined in terms of combinations of more detailed objects and processes at lower levels. Sun writes, for instance: “Higher-level entities would be made up of sets of more detailed entities, and causal relationships at higher level would be generated by the casual relationships amongst the equivalent entities at more detailed levels” (2005, p.624). Different entities and processes relate via part-whole relations on the basis of considerations of size and complexity. This is again quite similar to the realist view. Recall, as a representative example, Newell’s view of system levels: “At level N, a collection of K components, say of characteristic size S, are put together to form a component at N + 1” (1990, p.119). On both accounts, the existence of levels is dependent on the mereological stratification of different sized entities.

Next, consider (ii). Similar to authors like Newell, Sun’s hierarchy stratifies phenomena on the basis of causality and prediction. He writes, for example: “Causal relationships between detailed (deeper-level) entities must explain (or “generate”) the causal relationships that exist between higher-level entities. Higher-level relationships may be generated deterministically or probabilistically from lower-level (more detailed) relationships (2005, p.622). Phenomena sort to different levels based on identifying the distinct causal relationships that obtain between their constitutive entities and processes. By investigating the causal structures and relations that obtain between different entities and processes, phenomena can be sorted into different levels of the hierarchy. Because sociological-level phenomena support different causal relations than componential level phenomena, the models it constructs reside at a different level than componential level models. As Sun puts it: “[I]n cognitive science we can get a handle on the essential casual structures and processes of cognition by looking for what is invariant under different mappings from one domain to another or from one level to another” (2005, p.630).

Sorting phenomena across levels requires understanding causal relations and structures. Sun's account also seems to subscribe to condition (ii).

Interestingly, if Sun's account is committed to both (i) and (ii), then it embodies the core elements of the realist conception of levels, at least in broad outline. This is interesting not least of all because it shows the enduring appeal and influence of three main conceptions of levels. It shows that realism, as a general approach to applying the levels metaphor, is used even in the context of modeling. The realist, quasi-realist and perspectivist views provide useful tools for assessing and conceptualizing not only theories within cognitive science generally but cognitive modeling specifically – the views act as centers of gravity not only in theorizing but also in the practice of cognitive science.

The conclusion is also interesting because it highlights the potential practical import of the levels metaphor. As an orienting framework for cognitive modeling, the Hierarchy of Four Levels offers an important procedure for investigating, constructing and interpreting models. It is not only a way to think about what models attempt to explain but also a procedure for how to implement those computational explanations. The fact that Sun's account subscribes to something like the realist conception goes some way to demonstrating the concrete value of the levels metaphor to cognitive science. It shows that the three general applications of the levels metaphor not only have a role in organizing theoretical discussions about cognition, but also that they have a concrete role in how researchers think about, construct, and interpret specific models of cognition. It goes to show that the levels metaphor can have direct impact on the research efforts of cognitive science.

There are some differences between Sun's model and the realist conception worth mentioning. One is that Sun's framework grafts model levels onto pre-identified levels of explanation and organization. This is why modeling follows from phenomenological descriptions. The different model levels track different levels of explanation as they track different levels of organization. Levels of description or explanation (whether sociological, psychological or componential) provide the initial identification and specification of the to-be-modeled phenomena. The model levels are parasitic on levels of explanation and organization. Though never explicitly stated, this aspect of Sun's account follows in virtue of the account's emphasis on phenomena and the order in which models are constructed. Figure 3.2 provides an illustration.

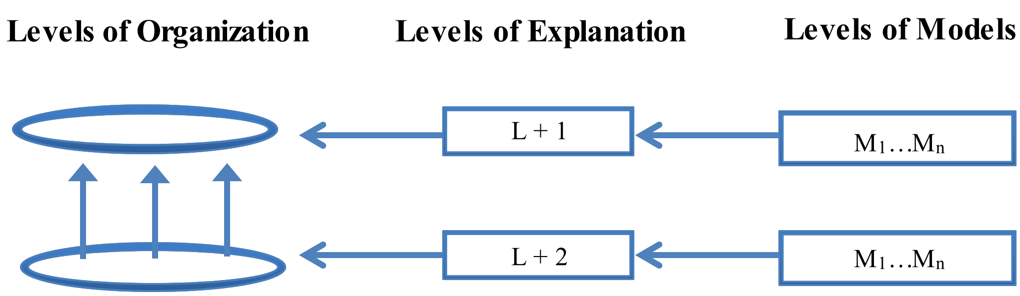

Figure 3.1 The relation between levels of organization, explanations, and models according to Sun's Hierarchy of Four Levels.

Another difference is that movement from one level to another is partly based on considerations of information compression – in addition to considerations of causal structure. Because the Hierarchy of Four Levels deals with phenomena of varying size and complexity, different levels of description are necessary to address informational overload. "To keep the information content of description on every level within human capacities, entities need to be arranged in hierarchies and defined in such a way that the interaction between two entities at a high level is a small proportion of the interactions with the population of more detailed entities that make up the two higher-level entities (Sun et al., 2005, p.622). Levels of modeling reflect a need to compress information, particularly as each level deals with larger and larger phenomena. The move between levels is often motivated by a need to abstract away from details in favour of more general information. This allows investigation to proceed in a scientifically tractable way.

To summarize the main point of this section, Sun's account models (no pun intended) the Hierarchy of Four Levels on something very close to the realist application of the levels metaphor. It is in virtue of correspondence with levels of organization and explanation that models find a home within the hierarchy – as further evidenced by the fact that Sun labels one of his commitments "The Dictum of Causal Correspondence." In both form and detail, Sun's account subscribes to something very close the realist view.

3.2.2 The Hierarchy of Four Levels Reconsidered

Having outlined the allegiance of Sun's framework to the realist conception of levels, I now turn to the more specific task of evaluating Sun's account. In bringing a critical eye to the Hierarchy of Four Levels, I argue that Sun's account is importantly incomplete. That it fails to incorporate several important and viable types of computational cognitive models. There are two examples I want to point to in underscoring the incomplete nature of Sun's account. Each takes issue with the kind of phenomena Sun's account views as central to organizing and constructing models.

The first is cognitive models based on population data. Consider prototype versus exemplar models. Both prototype and exemplar attempt to describe and explain categorization – the cognitive process by which we organize and arrange stimuli into different categories. Each model does so through a different means. Prototype models compare new stimuli to a single prototype in a category; while exemplar models compare new stimuli to multiple known exemplars in a category (see Minda & Smith, 2002). Under Sun's framework, prototype and exemplar models are most likely situated at the componential level. Each model attempts to explain categorization by appealing to intra-agent processes. Prototype models posit one set of mechanisms and processes for categorization; exemplar models another.

Here is the problem. Prototype and exemplar models are often based on and better fitted to two different types of human data (Minda & Smith, 2002). Prototype models are better suited to accommodate population-level data; exemplar models are better equipped to deal with individual-level data. Population data are abstract entities calculated using measures such as mean ratings; they are usually used to reduce the noise caused by variability among individual-level data points. Individual data are the data collected from observing behaviours using some performance measure, such as error rates or response times.

The question is whether there is room in Sun's hierarchy for models based on population data. At what level of organization do population-level data reside? Though inferences are made on the basis of population data, it does not appear to neatly fit within any of Sun's four levels. Individual-level data certainly seem to, presumably slotting into the componential or psychological level. Population data, on the other hand, which is the basis on which at least some computational cognitive models are constructed, seem to fail to map neatly to any one level of organization. Population data are certainly larger than the phenomena picked out by individual data, but they do not seem to be reducible to such phenomena. The problem is that in spite of the fact that population data play a crucial role in modeling, they do not seem to neatly fit anyway within Sun's hierarchy. Even though, as a theoretical construct, population data plays a crucial role in generating and building cognitive models, it does not have any sort of clear ontological standing within the Hierarchy of Four Levels.

The issue, then, is one of over-restriction. Sun's framework problematically limits the range of entities and processes that can figure in model construction. It excludes entities that can play an important role in modeling, but which find no clear basis within the Hierarchy of Four Levels. In particular, there seems to be a good question of why, given the active and prevalent use of population data, computational cognitive models based on abstract entities should be excluded a priori from Sun's framework. The argument, in other words, is that models constructed on the basis of population data, such as prototype models of categorizations, are counterexamples to the sufficiency of Sun's hierarchy. Models based on population data demonstrate the incompleteness of Sun's account.

Consider a second example – this time coming from the other direction. Chris Eliasmith and colleagues (2003, 2005, 2012) have recently developed a large-scale model of brain functioning. They call it "Spaun" (Semantic Pointer Architecture: Neural Network). Spaun is similar to other artificial neural networks in that it consists of groups of neurons and connections weights. These neurons form distributed representations of vectors and the connections specify the computations to be performed on those vectors – these are the two basic principle of the Neural Engineering Framework underwriting the Spaun model. However, unlike traditional neural models, Spaun does not use a learning rule during simulation. Instead, Spaun uses the connection weights to approximate the computation necessary to solve a given task. These approximations are pre-set before simulation.

More specifically, the model consists of three hierarchies, an action selection mechanism and five systems. The modules are considered to be cortical and subcortical areas of the brain, each implementing different computational operations – for details see Eliasmith (2013, Ch.7). In concert with the action selector, Spaun is able to seamlessly move between several cognitive tasks without changing any of its basic elements. The action selector allows the model to activate different spiking neuron chains that instantiate the computations used to solve everything from digit recognition and working memory tasks to 'question and counting tasks.

For present purposes, what is interesting about Spaun is how it interprets its explanatory target. Spaun's central purpose is to simulate in one unified model a set of neural mechanisms that perform several cognitive tasks. But notice what the model is not. It is not a neural model *of* cognitive processes, for it is purports to be an implementation of cognitive processes *in* a computer-simulated brain. Here is Eliasmith describing the model: "[W]e present here a spiking neuron model of 2.5 million neurons that is centrally directed to bridging the brain-behavior gap. Our model embodies neuroanatomical and neurophysiological constraints, making it directly comparable to neural data" (2012, p.1202). If Spaun is a simulation of the implementation of cognitive processes in brain-ware, then this raises the question of where the model slots into Sun's hierarchy of levels.

Two options appear possible: either Spaun resides at the componential or physiological level. Recall that componential-level models deal with intra-agent components and processes, while neurological level models deal with the physical implementing structures of intra-agent processes. Is Spaun better thought of as a physiological or componential model? At first glance, the answer is perhaps not immediately obvious. It seems to make sense to treat Spaun as either a componential or physiological level model. The model deals with the implementational of cognitive processes in neuroanatomical structures. It purports to embody both neuronal and functional processes as they relate to psychological behaviour.

The ambiguity, I suggest, is generated by Sun's appeal to strict correspondence between hierarchical levels. Because phenomena are arranged according to size and causal structure, Sun's model places constraints on what types of entities are properly suited to supply targets for computational cognitive modeling. Sun's emphasis on hierarchically arranged entities creates problems for how it understands models that straddle different levels. Models in computational neuroscience, of which Spaun is an instance, eschew any clear distinction between process and form. This makes it difficult to assign Spaun to a given level of the hierarchy. Sun's hierarchy is unable to accommodate Spaun in any sort of non-awkward or post-hoc way. Though the hierarchy is well suited to accommodate several styles of models, it falters with respect to other models that blur the lines between levels of the hierarchy.

Here is the argument in short: Because of its allegiance to the realist conception of levels, Sun's framework places constraints on the kinds of phenomena that can be used in the construction of cognitive models. However, given the viability of the models excluded, Sun's account is importantly incomplete. It unnecessarily excludes otherwise useful and well-established models. Informally, the argument looks something like the following:

(i) Frameworks that purport to serve cognitive modeling should incorporate the majority of models

(ii) Viable models are excluded on Sun's account

(iii) Thus, Sun's framework is incomplete.

A defense of premise (i) and (ii) will help to explicate the argument in a bit more detail. In terms of (i), Sun's framework is set to apply to all of cognitive modeling insofar as the levels described within the model are thought to range over all the phenomena relevant to the study of cognition: "[A] scientific theory of cognition requires the construction of a hierarchy of different levels…and the creation of models for such elements at appropriate levels" (Sun et al., 2005, p.634). The framework stands to apply to cognitive modeling as a whole, since models are supposed to track the range of cognitive phenomena at different levels of the hierarchy.

In terms of premise (ii), population-level models and Spaun function as a demonstration of the fact that at least some viable cognitive models are not included within Sun's account – by viable I mean the fact that types of models are taken seriously and have an established research program around them. The phenomena on which prototype and Spaun models are based are not accommodated within Sun's hierarchy. This shows the gaps in Sun's framework. Thus, insofar as Sun's framework is taken to apply to all of cognitive modeling, it is at least partially inadequate.

A qualification is necessary at this point. The concern just raised should not be read as a full-scale indictment of Sun's framework. It is not an attempt to undermine Sun's overarching account. Rather, the concern is simply that insofar as Sun's account aligns itself with the realist conception of levels, it falls into the monistic trap of thinking that there is only one true application of the levels metaphor. As mentioned, it does so insofar as Sun's framework purports to incorporate cognitive modeling as a whole, and not only some portion of it.[30] In highlighting cases that fall outside the purview of Sun's framework, I am attempting to draw attention to embracing alternative approaches to thinking about the role of the levels metaphor in cognitive modeling. The counterexamples used should provide some *prime facie* motivation for considering alternative approaches to modeling. Sun's framework is only one among several applications of the levels metaphor within computational cognitive modeling.

3.3 An Alternative Approach to Modeling

If Sun's Hierarchy of Four Levels embodies the realist application of the levels metaphor within cognitive modeling, then advancing discussion requires providing an alternative application of the metaphor. Thus, in line with the discussion of Chapter 2, this section provides an alternative framework for modeling, one based on the perspectivist conception of levels. This alternative framework provides an important supplement to the realist account, demonstrating the fecundity of other applications of the levels metaphor. The choice of perspectivism in this context is simply based on the fact that it offers the starkest contrast to the realist framework – in principle a similar framework can be constructed for quasi-realism.

3.3.1 Perspectivism and Modeling

If realism purports to provide a picture of how nature is organized into different levels, perspectivism offers a picture of levels as tools for cognitive investigation. Levels

[30] Had the view been more modest in reach and claimed to only apply to subset of models, then the argument could be made that Sun's account should not be construed monistically.

talk provides answers to different types of questions about cognitive phenomena. The implication for cognitive modeling is that models constructed at different levels within a perspectivist framework aim at facilitating answers to different types of questions. The value of levels to cognitive modeling therefore lies in the theoretical and explanatory utility it affords talking about information processing devices.

McClamrock's discussion of levels is useful here. He writes:

> [T]here are three perspectives that we can take toward it [cognition] – or if you prefer, three general kinds of questions that we can pose about it, or three kinds of explanations of it we might try to give: questions about that structure itself; questions about the functional, context-dependent properties of the parts and relations in that structure and their contribution to the functioning of the system as a whole; and questions about the implementation of the primitive parts of [the] algorithmic structure (1991, p.193).

Talk at different levels involves answering different questions about cognition – for example, questions about the syntactic structure of representations, the relational or global function of different processes or the primitives used to implement different computational processes. Shifts between different levels are shifts in the kinds of question one wants to answer.

In the context of modeling, this means that models are constructions that aim at providing answers to different types of questions. They are computational embodiments of theoretical answers to different questions. Some models answer questions about representations and algorithms, others questions about content, function or interpretation. On the perspectivist view, computational cognitive modeling organizes and constructs models on a case-by-case in accordance with the different types of information processing questions modelers ask.

3.3.2 Filling the Gaps

To get a better picture of the perspectival framework, return to previous two examples that caused trouble for Sun's framework.

On the perspectival account, one way to understand prototype versus exemplar models is as offering two differing answers to the same question. Though both models aim at detailing the means by which individuals categorize objects, prototype models simply do so through a more abstract characterization of individuals than exemplar models. This is why prototype models capture generalizations based on population data better than exemplar models and vice-versa. Prototype and exemplar reside at the same level not because they address the same sized phenomenon, but because they address the same type of question. Contextual factors, such as the experimental evidence and methodological design, produce the differences between the models. The models themselves, however, are localizable to the same level in virtue of answering the same type of question – in this case, a question about the procedures by which categorization is made possible.

A similar point holds in the case of Spaun. Spaun can be understood as answering two types of questions simultaneously. One is what are the functional primitives and processes that facilitate cognitive processes; the other is how in the brain cognitive processes and primitives are implemented. Spaun attempts to bridge the neural/cognitive gap not by providing a link between different-sized entities, but by answering different types of questions. Of course, it does so by modeling something in the world. But that something is not neatly organized into a hierarchy of componentially arranged entities and processes. The success with which it does so I leave for others to decide.

The point is that the perspectivist framework provides the resources to accommodate and incorporate various cognitive models that strain more realist-based approaches, such as Sun's. This is in virtue of its commitment to the many-many mappings possible between different levels of organization, explanations and models. Because modeling on the perspectival framework flows from questions to phenomena, it is more flexible in the relationship it envisions between different types of models, levels of explanation and levels of organization. Figure 3.2 provides an illustration.

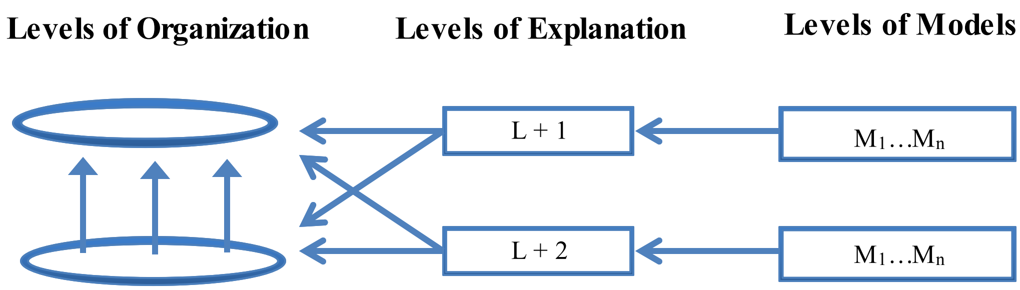

Figure 3.2 The relation between levels of organization, explanations, and models according to the perspectival view.

Contrasting the perspectival framework with the realist framework might help to further explicate the view. The main point of contrast between the realist and perspectival frameworks is that under the perspectivist rubric, models do not slot into any sort of clear global arrangement. Though models can be constructed into different levels, these levels do not reflect a corresponding tiered structure in the nature of the phenomenon being modeled. Understanding complex cognitive system needs to first and foremost be sensitive to function and form. That is not to say that structure is not important to modeling, but only that structure should not be envisaged in a hierarchical fashion – recall the discussion of complex realization in Chapter 2.

The perspectivist approach flips the script, maintaining that model construction follows the questions and tools being used rather than the phenomenon. Understanding cognitive phenomena is part-and-parcel of the activity of modeling. On the perspectivist framework, computational cognitive modeling is seen as a leveled enterprise that moves between several different perspectives, each adding something important to the mix.

3.3.3 A Problem and a Solution

At this point one might worry that without an overarching structure, such as the one provided by Sun's hierarchy, modeling devolves into a series of interesting but disconnected models. Fortunately, this is not the case. Principled work can still be done within the perspectivist framework largely due to the interconnected nature of the questions being asked. For example, consider the shift between questions about algorithms and implementation. On Marr's account, this shift reflects a move between computational and biological entities. It reflects a move between talking about representations and procedures and physical structures.

On the perspectival account, although this kind of shift is possible, it does not exhaust the relationship. The shift between algorithm and implementation is better understood as a move between processes and primitives. Primitives are the basic functional units operative within a given cognitive system at a specific level of organization. There is no limit to the number of times a system can be broken down into different primitives. The primitives of a virtual machine, for example, might be the basic building blocks in a programming language like LISP, but they might be the algorithmic units used in performing computations when compared to the units constitutive of programming in JAVA. As McClamrock points out:

> [I]f we were to take the "three levels" view as making a claim

> about the actual number of levels of organization (or stages of natural decomposition) in cognitive systems, it would be a very substantive (and I think false) empirical claim. It would be claiming that cognitive systems will not have any kind of multiple nesting of levels of organization. Why should our explanatory framework for cognition have this kind of limiting assumption built into it?" (1991, p.191).

Modeling can still be a principled activity under the perspectival framework so long as modelers are clear about when, why and how they move from answering one type of question to another. Because of the deep connection between the questions one asks and the types of models that are constructed, principled relationships can be still be established between different models.

3.3.4 General Remarks

What I have been trying to show is that a perspectival application of the levels metaphor within cognitive modeling brings something important to the table. It addresses important gaps in the realist conception of modeling, such as accommodating mixed- and cross-level models such as Spaun. However, as the previous discussion of pluralism tried to show, this does not mean that perspectivism should be taken as the superior approach to modeling. Instead, it should be taken to simply show that neglecting different applications of the levels metaphor produces anemic approaches to modeling. The perspectivist framework offers an important supplement to Sun's hierarchical approach. Researchers are well founded in adopting either, but cautioned against conceiving of either as the be all and end all of cognitive modeling.

More generally, what the above discussion shows is that computational cognitive modeling is best served not by adopting a perspectivist, realist, or quasi-realist application of the levels metaphor but rather by recognizing the advantages of moving between different applications. Much like the case for pluralism in philosophical logic, the issue is not whether there is some deep nature to the levels in modeling but whether different applications of the metaphor yield different benefits when investigating cognition. One way to see the preceding discussion is as a sort of proof of principle that each application of the levels metaphor adds something important to cognitive science. That being a pluralist is an important and viable position to adopt when it comes to the levels metaphor.

3.3.5 Summary

So, here is what I have done in this final chapter. First, I outlined the field of computational cognitive modeling as whole, detailing several of the major models currently in use and outlining one of the important debates within the field.

Second, I outlined in some detail Ron Sun's Hierarchy of Four Levels. The purpose of this discussion was twofold: (i) it served to illustrate the theoretical and practical importance of the levels metaphor and (ii) it served to show the pervasive influence of the three main conceptions of levels. The latter point was demonstrated by showing that Sun's account of modeling was committed to the realist conception of levels. Discussion then turned critical. Several shortcomings in Sun's framework were highlighted, particularly with respect to its ability to incorporate different types of cognitive models.

Third, and finally, in service of demonstrating the desirability of the pluralistic approach to levels, I outlined both what a perspectivist approach to modeling might look like and how it can accommodate several of the issues faced by a realist-inspired, monist approach. In so doing, I argued that perspectivism, and by extension quasi-realism, offered an important supplement to a realist approach to modeling. This discussion further served to motivate avoiding a monistic approach to levels.

3.4 Conclusion

A few general concluding remarks about the thesis are in order. First, I have tried to show that despite the ubiquity and intuitiveness of the levels metaphor, little work in cognitive science has been devoted to the topic. I have tried to show not only that discussion of levels is important to cognitive science, but that it also has practical implications for conducting research.

Second, in attempting to fill the gap in the literature, I offered an examination and analysis of various applications of the metaphor within cognitive science. This resulted in the creation of a conceptual framework for analyzing levels. Not only is this framework an important tool for philosophers, but it is also an important tool for researchers who want to flesh out their philosophical assumptions when conducting research.

Third, the thesis has also tried to highlight the vigilance required when it comes to thinking about and using the levels metaphor. It was argued that problems emerge when monistic assumptions are made about whether there needs to be only a single verdict when it comes to the conceptual soundness of the levels metaphor. In this way, the present thesis acts as a sort of cautionary tale for those who wish to utilize the levels metaphor within cognitive science.

In sum, this thesis attempted to take one small but concrete step toward making sense of the levels in cognitive science, advancing and clarifying how and when the notion can be successfully deployed. More work is obviously required. But I hope the thesis has at least outlined the contours of the problems and some possible solutions. In this way, whatever enduring value the thesis has lies more with the problems and positions it lays out than the solutions it offers.

Bibliography

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Anderson, J., & Liebere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Lawrence Erlbaum Associates.

Baddeley, A. D. (2012). Working memory: theories, models, and controversies. *Annual Review of Psychology*, 63, 1-29.

Barlow, H. B. (1953). Summation and inhibition in the frog's retina. *Journal of Physiology*, 119, 69-88.

Beall, J, & Restall, G. (2000). Logical pluralism. *Australasian Journal of Philosophy*, 78, 475–493.

Beall, J. & Restall, G. (2001). Defending logical pluralism. In *Logical Consequence: Rival Approaches Proceedings of the 1999 Conference of the Society of Exact Philosophy*, Stanmore: Hermes, pp. 1–22.

Beall, J. and G. Restall. (2006). *Logical Pluralism*. Oxford: Oxford University Press.

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.

Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4, 91-99.

Bermudez, J. (2005). *Philosophy of Psychology: A Contemporary Introduction*. New

York, NY: Routledge.

Broadbent, D. (1985). A question of levels: Comments on McCelland and Rumelheart. *Journal of Experimental Psychology: General*, 114, 189-192.

Brook, A. (2009). Introduction: Philosophy in and Philosophy of Cognitive Science. *Topics in Cognitive Science, 1*(2), 216–230.

Burge, T. (1986). Individualism and psychology. *Philosophical Review*, 95, 3-45.

Carruthers, P. (2006). *The Architecture of the Mind*. Oxford: Oxford University Press.

Churchland, P.S., Koch, C., & Sejnowski, T.J. (1990). What is computational neuroscience? In Schwartz, E.L. (ed.), *Computational Neuroscience* (pp. 46-55). Cambridge, MIT Press.

Clark, A. (1989). *Microcognition*. Cambridge, MA: MIT Press.

Craik, F. I. M. & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning & Verbal Behavior*, *11*(6), 671-684.

Craver, C. (2002). Interlevel experiments and multilevel mechanisms in the neuroscience of memory. *Philosophy of Science*, 69, 83-97.

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.

Cummins, R. (1983). *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.

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., Boechler, 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.

Dawson, M. R. W. (2002). Computer modeling of cognition: Levels of analysis. In L. Nadel (ED.), *Encyclopedia of cognitive science* (pp.635-638). London: Macmillan.

Dennett, D. (1991). Real Patterns. *The Journal of Philosophy, 88*(1), 27-51.

Egan, M. F. (1992). Individualism, Computationalism and Perceptual Content. *Mind*, 101, 443-59.

Eliasmith, C. (2005). A unified approach to building and controlling spiking attractor networks. *Neural computation*, 7, 1276-1314.

Eliasmith, C., Anderson, C.H. (2003). *Neural engineering: Computation, Representation and dynamics in neurological systems*. Cambridge, MA: MIT Press.

Eliasmith, C., Stewart T. C., Choo X., Bekolay T., DeWolf T., Tang Y., Rasmussen, D. (2012). A large-scale model of the functioning brain. Science. 338(6111), 1202-1205.

Eronen, M. (2015). Levels of Organization: A deflationary Account. Biology and Philosophy, *30*(1), 39-58.

Floridi, L. (2008). The method of levels of abstraction. *Minds and Machines*, *18*(3), 303-329.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3–71.

Gillett, C. (2002). The dimensions of realization: A critique of the standard view.

Analysis, *62*(276), 316-323.

Hubel, D., & Wiesel, T. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215–243.

Jackendoff, R. (1983). Semantics and Cognition. Cambridge, Mass.: MIT Press.

Kim, J. (1998). *Mind in a physical world*. Cambridge, MA: MIT Press.

Leigh R., J., & Zee, D. (2006). *The neurology of eye movements* (4th ed.). New York: Oxford University Press.

Luce, R. D. (1995). Four tensions concerning mathematical modeling in psychology. *Annual Review of Psychology*, 46, 1-26.

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.

Minda, J. P., & Smith, J. D. (2002). Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 28, 275-292.

Minsky, M. (1981). A framework for representing knowledge. In J. Haugeland (ed.), *Mind design* (pp.95-128). Cambrigde, MA: MIT Press.

Mormann, T. (2007). The Structure of Scientific Theories in Logical Empiricism. In A. Richardson and T. Uebel (eds.), *The Cambridge Companion to Logical Empiricism* (pp. 136–162). Cambridge: Cambridge University Press.

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.

Newell, A., & Simon, H. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of ACM*, 19, 113-126.

Osherson, D. N., Smith, E.E., Wilkie, O., Lopez, A., Shafir, E. (1990). Category-based induction. *Psychological Review*, 97, 185-200.

Piccinini, G. (2015). *Physical Computation: A Mechanistic Account*. Published to Oxford Scholarship Online. DOI: 10.1093/acprof:oso/9780199658855.001.0001

Pocock, G. (2006). *Human Physiology* (3rd Edition). Oxford, UK: Oxford University Press.

Potochnik, A. (2010). Levels of Explanation Reconceived. *Philosophy of Science*, 77(1), 59-72.

Potochnik, A., & McGill, B. (2012). The Limits of Hierarchical Organization. *Philosophy of Science 79*(1), 120-140.

Putnam, H. (1960). Minds and Machines. In (eds.) S. Hook *Dimensions of Mind*. New York: New York University Press.

Putnam, H. (1988). *Representation and Reality*. Cambridge, MA: MIT Press.

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.

Rescorla, R.A. & Wagner, A.R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W.F. Prokasy (eds.), *Classical Conditioning II* (pp. 64–99). Appleton-Century-Crofts.

Rumelhart, D., McClelland, J., & the PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Savage, C.W. (1990). Preface. In C.W. Savage (ed.), *Scientific Theories. Minnesota Stuides in the Philosophy of Science* (vol.4) (pp. vii-ix). Minneapolis: Univesrity of Minnesota Press.

Sazima, I., Carvalho, L., Mendonca, F., & Zuanon, J. (2006). Fallen Leaves on the Water-Bed: Diurnal Camouflage of Three Night Active Fish Species in an Amazonian Streamlet. *Neotropical Ichthyology*, 4(1), 119–22.

Searle, J. (1980). Minds, brains and programs. *The Behavioral and Brain Sciences*, 3(3), 417-24.

Searle, J. (1990). Is the Brain a Digital Computer?. *Proceedings and Addresses of the American Philosophical Association*, 64, 21–37.

Schank, R., & Abelson, R. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.

Shapiro, L. (2000). Michael R.W. Dawson, Understanding Cognitive Science. *Minds and Machines, 10*, 440-444.

Smolensky. R. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, 1-74.

Simon, A. (1962). The Architecture of Complexity. *Proceedings of the American Philosophical Society* 106, 467-482.

Sternberg, S. (1968). Memory-scanning: Mental processes revealed by reaction-time experiments. *American scientist*, *57*(4), 421-457.

Sun, R. (2002a). Hybrid connectionist symbolic systms. In M. Arbib (Ed.) *Handbook of brain theories and neural networks* (2nd ed.) (pp.543-547). Cambridge, MA: MIT Press.

Sun, R. (2002b). *Duality of Mind*. Mahwah, NJ: Lawrence Erlbaum.

Sun, R. (2004). Desiderata for cognitive architectures. *Philosophical Psychology*, *17*(3), 341-373.

Sun, R. (2008). Theoretical status of computational cognitive modeling. *Cognitive Systems Research*, 1-17.

Sun, R., & Bookman, L. (1994). *Computational architectures integrating neural and symbolic processes: A perspective on the state of the state*. Norwell, MA: Kluwer Academic Publishers.

Sun, R., Coward, A. L., & Zenzen, M. J. (2005). On levels of Cognitive Modeling. *Philosophical Psychology,* 18(5), 613–637.

Thagard, P. (2014). Cognitive Science. In Edward N. Zalta (ed.), The Stanford *Encyclopedia of Philosophy*: 1-15. DOI: http://plato.stanford.edu/entries/cognitive-science.

Turing, A. (1950). Computing machinery and intelligence. *Mind*, *59*(236), 433-460.

van Fraassen, B. (1980). *The Scientific Image*. Oxford: Oxford University Press.

van Fraassen, B. (2001). Constructive Empiricism Now. *Philosophical Studies*, 106, 151–170.

van Gelder, T., & Port, R. F. (1995) It's about time: An overview of dhymanical approach to cognition. In R. F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics in cognition*. Cambridge, MA: MIT Press.

Vygotsky, L. (1986). *Mind in society*. Hillsdale, NJ: Lawrence Erlabaum Associates.

West, R., & Lebiere, C. (2001). Simple games as dynamic coupled systems: Randomness and other emergent properties. *Cognitive Systems Research*, 1, 221–239.

West, R., Lebiere, C., & Bothell, D. (2006). Cognitive architectures, playing games, and human evolution. In (ed.) R. Sun, *Cognition and multi-agent interaction: From cognitive modeling to social simulation* (pp.103-124). Cambridge, MA: Cambridge University Press.

West, R. & MacDougall, K. (2014). The macro-architecture hypothesis: Modifying Newell's system levels to include macro-cognition. *Biologically Inspired Cognitive Architectures*, *8*, 140–149.

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. (2001). Two Views of Realization. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 104(1): 1-31.

Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences*. Cambridge: Cambridge University Press.

Wilson, R. (2005). *Genes and the Agents of Life: The Individual in the Fragile Sciences: Biology*. Cambridge University Press.

Wilson, R., & Craver, C. (2007). Realization: Metaphysical and Scientific Perspectives. In P. Thagard (ed.), *Philosophy of Psychology and Cognitive Science*, Amsterdam: Elsvier, pp. 81-104.

Wright, C.D., and Bechtel, W. (2007). Mechanisms and Psychological Explanation. In P. Thagard (ed.), *Philosophy of Psychology and Cognitive Science*, Amsterdam: Elsvier, pp. 31-77.