The Hierarchical Correspondence View of Levels:

A Case Study in Cognitive Science

(forthcoming in *Minds and Machines*)

Abstract. There is a general conception of levels in philosophy which says that the world is arrayed into a hierarchy of levels and that there are different modes of analysis that correspond to each level of this hierarchy, what can be labelled the 'Hierarchical Correspondence View of Levels" (or HCL). The trouble is that despite its considerable lineage and general status in philosophy of science and metaphysics the HCL has largely escaped analysis in specific domains of inquiry. The goal of this paper is to take up a recent call to domain-specificity by examining the role of the HCL in cognitive science. I argue that the HCL is, in fact, a conception of levels that has been employed in cognitive science and that cognitive scientists should avoid its use where possible. The argument is that the HCL is problematic when applied to cognitive science specifically because it fails to distinguish two important kinds of shifts used when analysing information processing systems: *shifts in grain* and *shifts in analysis*. I conclude by proposing a revised version of the HCL which accommodates the distinction.

Keywords: Levels, Hierarchy, Correspondence, Computational Modelling, Cognitive Systems

1. Introduction

There is a general conception of levels in philosophy which says that the world is arrayed into a hierarchy of levels, and that there are different modes of analysis that correspond to each level of this hierarchy, what I label the 'Hierarchical Correspondence View of Levels" (or HCL for short).[1] The HCL is a combination of three distinct ideas about levels: (i) that the world is

[1]The view has sometimes been called the 'correspondence views of levels.' However, this label does not capture the important role the notion of hierarchy plays in the view. What makes the view particularly attractive for many, as we will see, is not only the correspondence relation the view posits but also the conceptual tidiness of hierarchies.

organised into a hierarchy of ontological levels; (ii) that there are specific modes of analysis for describing these levels; and (iii) that a specific relationship obtains between the modes of analysis used to describe the world and the way the world is hierarchically organised. The HCL is a specific conception of how various ideas about levels fit together.

Oppenheim and Putnam's (1958) 'layer cake model' offers probably the most classic statement of the HCL. According to the layer cake model, there is a neat mapping between a given 'level' of science, such as chemistry, and a given 'level' of nature, such as the molecular level. The constituents comprising organised strata of phenomena in the world are neatly mapped to the predicates and theories linked with describing those constituents. For each mode of analysis in science, there is a distinct level of ontology to which it applies.[2]

However, despite its considerable lineage and general status within philosophy of science and metaphysics, the HCL has mostly escaped analysis within specific domains of inquiry. For instance, while Floridi (2008) provides an interesting and critical examination of the HCL, he does so with an eye to developing a more general concept of 'levels' for philosophy. However, as many are now concerned, there may be no such precise, general concept of levels for all of philosophy and science. The growing consensus is that any explication of the concept of levels has to be relative to a specific domain of inquiry (Brooks, 2017; Brooks & Eronen, 2018; Potochnik, 2022; Potochnik and Guilherme, 2020).

The goal of this paper is to take up this recent call to domain-specificity. My aim is to examine how the HCL applies particularly within cognitive science. I aim to show, first, that the HCL is a conception of levels that has been employed in cognitive science; and second, that cognitive scientists should avoid its use where possible. The HCL is problematic when applied to cognitive science, I suggest, because it fails to differentiate two important kinds of shifts crucial to analysing information processing systems: *shifts in grain* and *shifts in analysis*.

[2] For an extended analysis of the layer cake model, see Kim (2002), Baxendale (2016) or Brooks (2017).

In response, I propose a revised version of the HCL which accommodates the distinction. The task here is an extension of existing attempts to clarify and refine the concept of levels more broadly. However, unlike previous investigations, such as Craver's (2007), which focuses on the mechanisms of cognitive neuroscience, or Baxendale's (2016), which focuses on non-reductive physicalism in philosophy of mind, the current discussion centres specifically on the types of information processing systems that form the explanatory core of cognitive science.[3]

The paper is divided into four sections. First, in Section 2, I unpack the three main ideas of the HCL, alongside further precsifying the view via a formalisation. Then, in Section 3, I examine a detailed case study from cognitive science: namely, the work of Ron Sun (2005, 2008). I argue that Sun's work not only shows that the HCL exists in cognitive science, but that it also exerts an influence. Following this, in Section 4, I argue that the HCL fails to appreciate an important distinction between shifts in grain and shifts in analysis, and that this failure stems from an inability to appreciate the important role played by functional contextualisations within cognitive analysis. The result is an impoverished and restrictive conception of levels. Finally, in Section 5, I sketch a modified version of the HCL, one which accommodates the shift distinction and the role of functional contextualisation.

2. The Hierarchical Correspondence View of Levels

As mentioned, the HCL pulls together three distinct ideas about levels.

The first is that it assumes that natural phenomena, such as cognition, are organised into a hierarchy of ontological levels. Kim (2002) provides a clear statement of the idea: "what often seems implicit is a certain overarching ontological picture of the world according to which the entities of the natural world are organised in an ascending hierarchy of levels, from lower to higher, from simpler to more complex" (p.3). The rather intuitive, but often implicit, idea is

[3] For an early discussion, see Kersten, West, and Brook (2016).

that the world is arrayed into a hierarchy of levels. For example, on Oppenheim and Putnam's (1958) layer cake model, higher-order entities, such as cells, are assembled or composed out of lower-order entities, such as molecules, via part-whole relations.

As Potochnik (2017, 2022) points out, a hierarchy of ontological levels can be arranged according to a number of organisational schemes. One, for example, is in terms of compositional relations. On this scheme, higher levels are a stepwise function of the compositional relations obtaining between smaller and slower entities and processes (Wimsatt, 1976). Compositional relations are the most appealed to, but other schemes have been proposed, including part-whole relations (Oppenheim & Putnam, 1958), realisation (Fodor 1974), and temporal scales (DiFrisco, 2017). To clarify, I am not taking a stance on the plausibility of this ontological notion of levels, I am simply laying out its basic structure in order to show how it combines with other ideas about levels.[4]

The second idea is that the HCL assumes that there are various modes of analysis one can offer about the world, which we can assign to different levels. Floridi (2008) calls this the 'epistemological' conception of levels. For example, according to Marr's classic tripartite framework, cognitive investigations can be separated into three distinct modes of analysis or inquiry: the computational level, the algorithmic level, and the implementational level.[5] Each of these three levels attempt to answer a different question about a cognitive system – for example, the computational level answers questions about a cognitive system's function and constraints on the system, while the algorithmic level answers questions about the procedures and representations by which a cognitive system achieves its function.

According to Marr's picture, the answer given to each mode of analysis constrains, to a greater or lesser extent, analysis at the level below. For example, depending on the nature of

[4] For critical discussion, see Wimsatt (1994), Kim (2002), Craver (2015), Eronen (2015), Potochnik (2010), or Baxendale (2016).
[5] For some, there is also an 'architectural' level that separates from the algorithmic level, but I leave such interpretative questions to one side for now (see Dawson 1998).

constraints and function identified at the computational level, the set of procedures and representations posited at the algorithmic level can change. If vision has the function of representing objects rather than facilitating movement, then a whole host of representations and algorithms are needed to explain this higher-level function.[6] There are different vocabularies or languages that can be used in analysing cognition (e.g., computational versus implementational), and these can be arranged into a mutually informing hierarchy. As Richardson (2009) makes the point:

> It is plausible to think that there are as many levels of description available as there are levels of organization [...] We may describe humans in terms of their physical make-up, their chemical constitution, their physiological structure, their gross anatomy, their cognitive capacities, their social role, and much more. All these give a distinctive perspective on human psychology. (p. 474).

In short, the HCL also assumes that different modes of analysis can be arranged into a hierarchy of explanatory levels.

Finally, the HCL takes a particular stance on the relation between the previous notions of levels: namely, it claims that for each mode of analysis there is a corresponding level of ontology. That is, there is systematic correspondence between a specific mode of analysis, i.e., a descriptive vocabulary or language, and specific level of the ontological hierarchy, i.e., a set of stratified structures in the world. This point is often assumed but not often made explicit. The claim is not that there is just *a* correspondence between hierarchically arranged modes of analysis and ontological levels, but, rather, that for every mode of analysis there is a *unique* level of the ontological hierarchy to which that mode of analysis applies. The HCL assumes that there is a one-to-one mapping between a given mode of analysis and a given level of

[6] The reverse case also holds for Marr. The format of representations and procedures used can similarly the set of computations that are implemented. Constraints flow top-down and bottom-up.

ontology, what I call the 'strict correspondence relation'. Potochnik and Guilherme (2020) indirectly critique this view when they write:

> [W]e grant the existence of part-whole and realization/implementation relationships, and this discussion is not meant to convince readers that there are no levels of organization. Rather, our point is that the entities and properties focal to different cognitive science explanations are not usually related to one another in any of the primary ways used to articulate levels.

I hasten to add two caveats. First, the HCL is not a conception of a specific notion of levels, such as the familiar concepts of 'levels of organisation' or 'levels of analysis'. Rather, it is a conception of how different notions about levels *relate*, i.e., the explanatory and ontological notions of levels. One consequence of this is that the HCL does not commit to how many levels there are, what they consist of, or what their target is. Instead, it simply spells out the relation between different notions of levels via a specific conception of correspondence and a hierarchical ordering. As we will see, this means that authors can put forward a number of different ontological and explanatory levels, while nonetheless remaining committed to the HCL.

Second, the HCL is not committed to a conception of how the different notions of levels relate internally. The HCL is compatible with a range of positions on inter-level relations, such as integration, reduction, or autonomy. One could, for example, maintain that Marr's computational level is insulated from details at the algorithmic level – i.e., that it is explanatorily autonomous, or vice-vera, while nonetheless maintaining that it addresses a distinct ontological level. As we will see, the only formal requirement on the HCL is that there is a weak mapping relation between elements of one level and those of another within either hierarchy. Such a constraint is compatible with stronger and weaker formulations, such as identity or multiple realisability (see List (2019, p.858) for discussion). Figure 1 provides an illustration.

2.1. A Formal Definition

One of the troubles with talk of levels, as several authors have pointed out, is that while its usage is fairly common across the natural and cognate sciences, its discussion, in both philosophical and scientific circles, often is often less than precise (Craver, 2015; Potochnik, 2017). To rectify this, following List (2019), I want to suggest that the HCL be further precisified using a pair of related concepts: 'systems of levels' and 'functor'. This formalisation enables us to generate a clearer picture of what exactly the HCL entails, and how to spot it in specific domains. A system of levels is a pair <L, S> defined as follows:

- L = a class of objects called *levels*.
- S = a class of mappings between levels, where each mapping σ has a *source level* l and a *target level* l' and is denoted σ: l → l'

While a system of levels can, in principle, capture any notion of levels, I apply the concept here to the ontological and explanatory notions of levels employed in the HCL.[7]

First, when applied to the ontological notion of levels, levels denote a set of possible *level-specific* worlds with a *total ordering* – i.e., S is reflexive, antisymmetric, transitive and connex.[8] Each level is a set of encoded facts that uniquely obtain at a world. For example, worlds at a physical level are those that encode the totality of physical facts – some levels are 'thicker' or 'thinner' depending on the number of facts they encode, e.g., physical versus psychological levels. The relations between level-specific worlds are captured by mappings. Each lower-level

[7] List (2019) uses the framework to several address debates in philosophy, such as those about consciousness and realisation.
[8] Note that this is specific to the HCL conception of levels, List allows for partial orders too in his system of levels).

world determines facts at a corresponding higher level world – for example, lower-level chemicals facts determine or settle higher-levels biological facts. For some non-empty set of level-specific worlds L ($\Omega_1 \ldots \Omega_n$), there must be a class of functions S of the form $\sigma: \Omega \rightarrow \Omega'$.

Second, when applied to the explanatory notion of levels, levels denote a set of languages – i.e., sets of formal expressions. A sentence is true at a given mode of analysis only if the object it denotes appears in the world. For example, the truth of a sentence, Φ, of a language, ℓ, is determined at a world, ω, if Φ is member of a maximally consistent subset of ℓ to which ω corresponds; otherwise, it is false. The relations between modes of analysis are also captured by mappings, only this time the elements are those of a language rather facts of a level-specific world – for example, sentences in a physical language and sentences in an intentional language. For some non-empty class of modes of analysis L ($\ell_1 \ldots \ell_n$), there is a class of functions S of the form $\sigma: \ell \rightarrow \ell'$.[9]

Complementing the notion of a system of levels is the concept of a 'functor'. A functor is a structure-preserving mapping between different systems of levels. A *functor*, F, from one system of levels, <L, S>, to another, <L', S'>, is a mapping which assigns to each level l in L a corresponding l' = F(l) in l', and assigns to each mapping σ: l → l' in S a corresponding mapping σ' = F (σ) in S', where F(σ): F(l) → (l').

We can use functors to explicate the notion of a one-to-one mapping for the HCL. For each level of <L, S>, the HCL claims that there is unique level in <L,' S'> to which it is assigned. Or to be more specific, for two systems of levels, one of which is an explanatory system of levels <L, S> and one of which is an ontological system of levels <L', S'>, there is a functor

[9] The following conditions hold under this definition:

S1) *Transitivity*: S is closed under *composition* of mappings, i.e., if S contains σ : l → l' and σ' : l' → l'', then it also contains the composite mapping $\sigma \cdot \sigma'$: l → l'' defined by first applying σ and then applying σ' (where composition is associative).

S2) *Reflexivity*: For each level L, there is an *identity mapping* 1_L : l → l in S, such that, for every mapping σ : l → l' , we have $1_L \cdot \sigma = \sigma = \sigma \cdot 1_L$

S3) *Uniqueness*: For any pair of levels l and l', there is at most one mapping from l to l' in S.

from <L, S> to <L’, S’>, and <L’, S’> to <L, S>. The two systems are structurally equivalent, i.e., each level and mapping of <L, S> uniquely corresponds to a level and mapping of <L’, S’>, and vice versa.

There is also a total ordering to levels according to the HCL; levels are *well defined*. That is, for any two levels, one can always say if one level is higher or lower than the other (e.g., the computational versus algorithmic). For example, for two levels Ω and Ω’, of a system of ontological levels, <L, S>, there is a function, S, such that either $\sigma: \Omega \rightarrow \Omega'$ or $\sigma: \Omega' \rightarrow \Omega$ (but at least one). In contrast, in a partial ordering, it could be that either Ω or Ω’ are higher (but at most one). Total orderings are more restrictive subset of partial orderings. Of course, a set of levels in L could be larger than a set of L’ such that the mappings between S and S’ are preserved, but the important point is that the mapping relations between L and L’ are in both directions (i.e. one-to-one).[10]

As an illustrative example, consider Marr’s (1982) tripartite framework. Framed in terms of the HCL, Marr’s computational, algorithmic and implementation levels constitute three languages for describing cognition $\{\ell_1, \ell_2, \ell_3\}$, as well as an ontology of levels $\{\Omega_{1\ldots}\ \Omega_n\}$.[11] Marr’s levels form a non-empty class of modes of analysis and mappings, <L, S>, and a corresponding non-empty class of levels-specific worlds and mappings, <L, S>, such that there is a mapping between each mode of analysis and mappings in <L, S> and a level-specific world and mappings in <L’, S’>. Thus, to say that a sentence Φ is true at the computational mode of analysis, ℓ_1, is to say that Φ denotes a set of facts or entities at a corresponding level-specific computational world, $\Omega_{L1,}$ such that $\Phi \in \ell_1$ and $\Phi \in \Omega_{L1}$.[12]

[10] For a view which importantly departs from the HCL in several respects, see Piccinini (2020).
[11] An 'ontology' here simply means a minimally rich set of worlds, Ω_L, such that each world in Ω_L settles everything that can be expressed in L.
[12] One possibility here, for instance, is Sterelny’s (1990) suggestion that the computational level maps to ecological level.

I should be clear, though. I am not claiming that this is actually what Marr holds. Rather, I am simply using it to illustrate how Marr's view could be interpreted along the lines of the HCL.

3. The HCL and Cognitive Science

Yet even granting the general existence and coherence of the HCL, one might nonetheless wonder whether the view has actually found its way into cognitive science. Potochnik and Guilherme (2020) gesture in this direction when they write: "All three positions [the reductionist, pluralist, mechanist] tend to share the presumption that the different phenomena involved in cognition, and the different investigations that target them, hold a hierarchical or leveled relationship to one another" (p.1310). However, to demonstrate that a particular domain is committed to a certain conception of levels, such as the HCL, one has to do a bit more. As Kim (2002, p.3) notes, different conceptions of levels are often hidden in the background of various domains. For this reason, in what follows, I consider a detailed case study. The aim is to demonstrate that, as a matter of fact, the HCL exists and has exerted an influence within cognitive science. The claim is not that all accounts in cognitive science subscribe to the HCL, but that at least some do, and this warrants caution.

3.1. A Case Study

In a series of papers, Sun (2005, 2008) lays out an ambitious yet detailed vision of computational cognitive modelling, what he labels the 'new hierarchy of four levels'.[13] As the moniker suggests, the account posits four distinct levels for computational cognitive modelling.

The first is the sociological level. Research here models large-scale objects, such as collective agent behaviour, inter-agent processes, socio-cultural processes, or environment-

[13] By 'computational cognitive modelling' I mean here the practice of providing detailed descriptions of the mechanisms and processes of cognition in terms of the science and technology of computing.

agent interactions (both physical and socio-cultural) (Vygotsky, 1986). One example is West et al.'s (2006) model of human game playing, which investigates how human game playing often deviates from optimal strategies. Such a model strives to capture the relationship *between* agents, rather than *within* an agent. Sun suggests that without understanding and incorporating socio-cultural factors researchers are in a difficult position when it comes to modeling cognitive processes; modeling is only partial or incomplete without the sociological level.

The second 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 here is Osherson et al.'s (1990) model of inductive reasoning, which models simple inductive inferences via the relation between superordinate categories and premises. Models at the psychological level incorporate data structures representing "knowledge" and "beliefs." Investigation involves collecting and interpreting behavioral data, usually via functional analysis.

The third is the componential level. The componential level addresses the intra-agent components responsible for behaviour at the psychological level. Sternberg's (1968) classic memory model offers an example, as it uses data collected from verbal protocol analysis to construct a model that mirrors the exhaustive search strategies used during digit span recall tasks. The componential level involves analysis both in terms of the structural and functional elements of cognitive processes, often using the resources of symbolic, connectionist or dynamicist approaches (Dawson, 2013).

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

More generally, Sun's hierarchy offers a procedure for investigating, developing, and interpreting computational cognitive models. Different models are constructed on the basis of different modes of analysis, each addressing distinct types of phenomena. Each level has (i) an object of analysis, (ii) a type of analysis, (iii) and an associated set of computational models. As Sun et al. describe the vision:

> [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 phenomenon […] 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).

There are a few points to note about Sun's account. The first is that 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. Another is that 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 are limited to one level at a time. The main constraint is that models address the appropriate process, entities, and causal relations. A third is that the 'phenomena' detailed or captured reside at distinct levels of the hierarchy. Human mentality is divided into different layers, each demarcating entities of varying size, complexity, and casual powers with a general cognitive architecture. Finally, 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. In the spirit of Marr (1982), higher and lower-level factors place constraints on modeling.

What I want to suggest is that the Sun's account is (i) implicitly but concretely committed to the HCL, and (ii) that such a commitment has important methodological/explanatory implications.

With respect to endorsement, consider first that the objects and causal relations at higher levels of Sun's hierarchy are defined in terms of combinations of objects and processes at lower levels. Sun et al. write, 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 compositional relations on the basis of considerations of size and complexity. The objects of the componential level, for example, which include intra-agent cognitive systems and processes, are composed of physiological level entities and processes, such as action potentials.[14] As one moves up the hierarchy, one moves from smaller to larger phenomena, each with its own distinct ontological standing. There is an ontological system of levels <L', S'>, consisting of several level-specific worlds L' ($\Omega_1 \ldots \Omega_n$) and a class of functions S' ($\sigma: \Omega \rightarrow \Omega'$), such that each element of one level can be mapped to an element of another level in the hierarchy.

Second, notice that computational models are assigned to different levels of the hierarchy in virtue of the complexity of the processes or entities they describe. The hierarchy stratifies models on the basis of causality and prediction. 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"

[14] Gallagher, et al. (2015) offers a similar view, emphasising the aim of "multi-scale explanations involving factors at various scales (neuroscientific, psychological, phenomenological, social, and so on) all contributing to an integrated explanation" (p. 156-157).

(2005, p.630). For instance, because sociological-level phenomena support different causal relations than componential level phenomena, the models it constructs reside at a different level. As one moves up Sun's hierarchy, one moves up in descriptions of increasingly complex and causally dense phenomena. Or, to put the point again in more formal terms, there is an explanatory system of levels <L, S>, consisting of several languages L ($\ell_1 \ldots \ell_n$), and a class of functions S ($\sigma: \ell \rightarrow \ell'$), such each element of one level can be mapped to an element of another level in the hierarchy.

Finally, notice that the relation proposed between the different aspects of the hierarchy is one of systematic correspondence. As Sun et al. write:

> A theory on any level creates entities at that level of descriptive detail as well as causal relationship between those entities that correspond with a range of data. Entities at a higher level often tend to package sets of lower-level entities in such a way that the higher-level causal relationships can be specified without reference to the internal structure of the higher-level entities (2005, p.622).

The phenomena addressed by the different types of analysis, and thereby computational models, are based on what unique causal relationships they capture. Each type of analysis addresses only certain types of entities and process at specific organisational levels. For instance, because sociological phenomena support different causal relations than componential level phenomena, models constructed targeting sociological phenomena reside at a different level of the hierarchy than componential ones. There is unique mapping from the explanatory system of levels to the ontological system of levels, and vice-versa.[15]

As is hopefully clear, Sun's account exhibits all the hallmarks of the HCL. There is a set of explanatory levels consisting of distinct modes of analysis (e.g., sociological, psychological,

[15] I am here bundling the computational models and modes of analysis together within the explanatory system of levels for ease of exposition, but the point still holds.

componential, physiological), a set of ontological levels describing various types of phenomena (e.g., collective behaviour, beliefs, sub-personal states, and neural processes), a total ordering to the hierarchy (e.g., the sociological level is at the top of the hierarchy and the physiological level is at the bottom), and there is a unique correspondence relation which holds between levels (i.e., each explanatory level describes distinct types of entities and processes at a given ontological level). The account envisages two distinct explanatory and ontological hierarchies linked by a systematic mapping relation.

With respect to the explanatory and methodological implications, consider first that Sun's modeling methodology makes particular sense if premised on a systematic relation between explanatory and ontological hierarchies. Recall that a researcher must initially characterise a phenomenon within a specific mode of analysis, such as the cognitive or social. Then, after collecting experimental data, certain causal relationships between entities at a particular level become conspicuous. Using these casual relationships, a researcher can develop a fleshed out computational model, implementing and testing the model against empirical data. But notice that it is only when there is an ordered, systematic hierarchy of levels, such as the HCL entails, that there is good reason to think modelling can unfold via the identification of phenomena at particular levels of ontology, descriptions in terms of specific modes of analysis, and implementations as computational models. Without the HCL, it is unclear why what is being modelled at a given mode of analysis should correspond to an appropriate ontological level. As Sun et al. (2005) expresses the point: "such tying-together is necessary for developing a deep theory in cognitive science."

Second, notice that if ontological levels are hierarchically organised, each mapping to a corresponding mode of analysis (e.g., sociological, psychological, etc.), then this makes sense of why Sun is committed to the prospects of mixed- and cross-level modeling: Such a vision requires a total ordering to a hierarchy – a key feature of the HCL. That is, it requires the ability to identify where a given phenomenon stands in the hierarchy. Without being able to clearly identify the target phenomenon, it will be unclear which facts about which entities and processes (sociological, psychological, etc.) should be incorporated into a model. So long as researchers are clear about which phenomenon are being targeted, which constraints from higher and lower levels are being included, and which types of analysis are being used, they can provide not only precise, detailed models at a specific mode of analysis, but also interesting cross- and mixed-level analyses. The very idea of cross- and mixed-level modeling finds natural support within the HCL.

To be clear on this point, the claim is not that the HCL entails, or excludes, mixed- or cross-level modeling. This is a further assumption specific to Sun's account. Rather, the point is that in implicitly adopting the HCL, Sun's account envisions a particular relationship between the various conceptions of levels (e.g., the explanatory and ontological), and this orients and constrains the methodological lessons it draws in thinking about how to organise and construct computational models. To use a metaphor, the HCL offers a loose scaffold onto which researchers can hang their more specific visions of levels.

Taking a step back, the fact that Sun's account both exhibits the hallmarks of the HCL and orients its thinking about computational modeling as a result should prove interesting. It shows that the HCL does, as a matter of fact, direct the construction and interpretation of work within cognitive science. In siting at the cross-roads of experimental, computational, and philosophical work in cognitive science, Sun's account offers an illustrative example of the HCL in action, showing how a general yet intuitive picture of levels can influence the practice of cognitive science.[16]

[16] While I do not think that Sun's is the only example of the HCL in cognitive science, with Fodor (1978) and Danks (2014) also likely offering plausible candidates, it offers a particularly clear and developed case, and so I have chosen to focus on it here.

4. The Shifting Nature of Levels

The dangers of the HCL are well documented. For example, in philosophy of biology, Potochnik and McGill (2012) worry that a hierarchical notion of levels fails to make room for complex forms of realisation, Craver (2007) worries that it fails to make sense of the rich multi-layered structure of the brain, while Potochnik and Guilherme (2020) question the value of talking about levels at all.[17] These are important criticisms, and there is much to be said for them.

However, in what follows, I develop a different worry. I argue that the HCL is uniquely problematic for cognitive science because it conflates two important kinds of explanatory shifts: (i) shifts in grain and (ii) shifts in analysis. In so doing, it fails to capture an important part of explanatory practice within cognitive science: functional contextualisation. While I previously argued that the HCL is active within cognitive science as a matter of descriptive fact, I now suggest that cognitive scientists *should* avoid its use where possible. To be clear, I take this argument to apply specifically to cognitive science. It might be that the HCL fails in other domains for different reasons, but here I am concerned solely with cognitive science. I begin by laying out the two kinds of shifts, and then say why the HCL neglects functional contextualisation.[18]

First, consider *shifts in analysis*. These are shifts in the theoretical vocabulary or language used to describe a complex system. When one switches from one mode of analysis to another, such as moving from the language of biology or neurology to computation or information processing, one is shifting to different descriptive or explanatory categories.[19] Shifts in the mode of analysis mark movements between distinct types of descriptive or explanatory

[17] See also Kim (2002), Rueger & McGivern (2010), and Eronen (2013).
[18] For alternative, pluralist worry, see Potochnik and Guilherme (2020).
[19] Here I am taking 'theoretical vocabulary' to simply mean the explanatory tools one uses plus the language in which those tools are couched. So, for example, computational analysis not only uses the language of computing but also couches its explanations in terms of how some phenomenon can be said to compute or not.

vocabularies. Second, consider *shifts in grain*. These are shifts in the generality of one's description. For example, when one omits properties relevant to the aerodynamics of an aeroplane in building a model (to instead, say, focus on geometric proportions), one shifts the coarseness of one's description. Shifts in grain decrease or increase the amount of detail provided within a given description (Kersten, 2020).

A helpful way to think about the distinction is in relation to the notion of *abstraction*. On at least one plausible account, abstraction involves selectively attending to a subset of features within a target system relative to the parameters of a background theory (Portides, 2018). For example, in mechanics, the material composition and internal structure of a body are often ignored in setting-up equations for motion. Relative to a background theory, such as physics, a mechanical model of a target system (e.g., a body in motion) is created by focusing on certain features, such as weight or height, at the expense of others. The process of abstraction results in a simplified description of a target system via the selective effects of a constraining background theory. Framed in terms of abstraction, the thought is that shifts in grain reflect the process of omitting or adding detail within a description (e.g., selectively attending to target features within a model), whereas shifts in analysis involve changes to the constraining background theory (e.g., the language of physics).

The problem is that the HCL runs together shifts of grain and analysis. The HCL assumes that shifts in grain systematically coincide at certain points with shifts in analysis. Pylyshyn (1984, p.33) expresses the idea nicely when he writes: "It is the fact that certain arbitrary sets of physical properties have something in common at the functional level, which cannot be captured in terms of a finite description at the physical, that leads us to postulate the existence of a new level with a distinct taxonomy or vocabulary." To speak of moving from higher- to lower-level organisationally layered components and processes (e.g., spatial maps and LTP mechanisms) within a memory system, for instance, involves, at certain fixed points, moving from higher- to lower-level modes of analysis (e.g., componential to physiological levels). Because of the systematic mapping between modes of analysis and levels of ontology, shifts among the latter precipitate changes amongst the former.

To be clear on this point, the claim is not that every shift in grain reflects a shift in mode of analysis; it is more specific. Rather, the idea is that at certain crucial points in analysis enough shifts in grain will precipitate or require a shift in mode of analysis. Such a view is indeed understandable enough to an extent. In many cases shifts in analysis do appear to reflect shifts in grain – certain causal generalisations, for example, such as those concerning intentionality, do seem naturally captured by some vocabularies and not others, as Pylyshyn (1984) and others point out. Here the HCL does well. However, the trouble is that there also appear to be a number of cases in which the relationship breaks down.

Consider, for example, Changeux's (2017) 'dynamical nesting model'. For Changeux, the brain is a nested assembly of functional structures at multiple levels of organisation, including: the level of genome, the TF–gene network level, the level of epigenetic action on synapse formation, and the level of long-range connectivity. Each level is reciprocally inter-regulated and in constant dynamic evolution with the others. But notice that while the brain has a dynamic multi-level structural organisation, all of the levels are governed by explanations using one mode of analysis: namely, the level of physical chemistry. Contrasting his approach with others Changeux (2017) writes, for example: "the models aimed at representing and/or simulating a process and/or behaviour on the basis of minimal, yet realistic, architectures and activity patterns most often use a single level of organisation. To attempt a type of modelling that spans several levels, as proposed here, is in itself a theoretical position" (p.169). While shifts in grain of analysis occur across multiple temporal and spatial scales, the actual mode of analysis remains constant (see Green and Batterman (2017) and Hauies and Burnston (2021) for further examples).

The reverse case also holds. For example, consider that at the level of primitive computing components, a transistor (a physical device) can be interpreted as a logic gate (a computational device) when its states are assigned the values 1 and 0. Such an interpretation is made possible by the fact that the transistor supports two stable but different states. To arrive at a computational description, one must selectively attend to a subset of the features of the transistor (its degrees of freedom), as they accord with computational theory; only then does the physical description of the transistor transform into a mathematical description of a device that processes 1s and 0s, i.e., a logic gate (Kersten, 2024). Importantly, though, this shift in mode of analysis (computational versus physical) does not involve a corresponding change in the grain of the structures (i.e., transistors versus logic gates). While the structure is re-interpreted relative to a new theoretical vocabulary (i.e., the language of computation versus physics), the actual structure (the transistor) remains at the same grain of analysis. There is a change in the mode of analysis without a corresponding change in the grain of analysis.

It seems, then, that there are two types of shifts available for complex systems, which dissociate in various ways. There are cases where shifts in grain and shifts in analysis decouple, i.e., the same kind of analysis can be used across different grains, or vice versa, and there are cases where different kinds of analysis are used on the same grain. This fact leads to three possible options in conceptual space. If shifts in grain and shifts in analysis systematically come together (e.g., as Sun proposes between componential and psychological levels), then something like the HCL follows. But if the two shifts robustly come apart, then there are two further possibilities: either there are shifts in analysis without shifts in grain (e.g., physical versus computer vocabulary), or there are shifts in grain without shifts in the analysis (e.g., brain model cases). As we have seen, the HCL, in conflating the two kinds of shifts, restricts the conceptual space so as to only include the first set of cases but this excludes other important cases. This has two problematic consequences I want to suggest.

The first is that it means that the HCL imposes an unnecessary methodological constraint on investigation. In envisaging the shift between modes of analysis as a shift between levels of organisation or ontology, the HCL restricts how many times a given mode of analysis might apply. For example, as Bechtel (1994) notes, one can just as easily apply Marr's modes of analysis to low-level intracellular processes, such as oxidative phosphorylation, as one can to high-level cognitive processes, such as vision. The reason for this is that, at core, Marr's levels answer different types of questions (i.e., via changes in the background theory). We can ask what purpose oxidative phosphorylation serves at one level, or what metabolites contribute to the process at another. Computational and algorithmic analysis do not in every case mark shifts in different sized properties. There is no reason a conception of levels should place such a constraint on investigation ahead of time.

The second is that it means that the HCL makes for a questionable empirical hypothesis. In over-stating the fit between explanatory and ontological levels, the HCL precludes the possibility of cognitive systems having a multi-nested structure. McClamrock (1991) foreshadows this worry in discussing Marr's account: "[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" (p.191). The HCL assumes that complex systems will not have nest structured within a given level of ontology.[20] However, again, as we saw, this seems like an overreach. There is no reason a conception of levels should make such a specific assumption about the structure of complex systems such as cognition.

[20] To be clear, by 'nested structure' I mean the fact that within a level of ontology there might exist several temporal and spatial grains.

4.1. Functional Contextualisations

So why does the HCL run the two types of shifts together, and so become overly restrictive? It does this because it fails to appreciate an important part of good explanatory practice about complex information processing systems: namely, it neglects the role played by 'functional contextualisations'.

Simply put, functional contextualisations are analyses of the way in which a physical structure's functional role is often determined by situating or embedding it within a wider system. As Cummins (1983) puts the idea: "to ascribe a function to something is to ascribe a capacity to it that is singled out by its role in an analysis of some capacity of a containing system" (p.99). Functional contextualisations are not simply capacities picked out by their place in a system but, rather, analysis of the activity of some item in terms of how it is organised into the workings of a system (Craver, 2013). Figure 2 provides an illustration.

For example, to say that the heart distributes oxygen or pumps blood, one has to not only describe the heart's local/intrinsic properties, such as its constricting and releasing activities, but also its position within the circulatory system. Without reference to the external objects and activities of the circulatory system, the functional role of the heart is unclear; the object's local/intrinsic properties fail to identify its functional role. If the contextual description changes, then so too does the heart's function. So, if, for example, the heart is described as a 'noisemaker', then it is embedded in a new system (a new causal nexus), which thereby implicates a new set of components and activities, such as the resonance frequency of the heart's chamber walls or the capacity to use the heart in a three-piece band. Characterised one way, the heart forms part of the mechanism for blood circulation; characterised another way, it forms part of the mechanism for sound generation.[21]

[21] Craver (2013) frames these upward and downward changes in terms of 'contextual' and 'constitutive' explanations but the idea is the same.

Functional contextualisations are equally important in cognitive investigations. This is because cognitive systems, in being information processing systems, are, in principle, decomposable any number of times. As McClamrock notes: “the degree of functional decomposition (i.e., the true level of organization) might be done any number of times for some particular information processing system” (1991, p.191). The functional decomposability of a given cognitive or perceptual system is not only tied to its organisational structure, but also to how the system is initially functionally contextualised.

Consider vision, for example. If characterised in terms of facilitating object recognition, a whole suite of properties and activities are salient to the investigation. These include fine-grained properties, such as those in the primary visual cortex involved in edge detection, and coarse-grained properties, such as those involved in categorisation. However, if, in contrast, vision is characterised in terms of facilitating action-guidance, then an entirely new set of properties and activities are implicated, such as those involved in maintaining body images or reaching behaviour. The functional decomposability of vision depends not only on the structural organisation of the phenomenon but also on how the activity is initially characterised.

To further clarify, though, the claim is not that the properties *representing* objects in the visual system are more or less coarse or fine-grained. Rather, the point is that those models which purport to explain the visual system propose differently-grained properties depending on the functional contextualisation. Which properties are structurally relevant to decomposition is a product of, in part, the initial functional contextualisation within a wider embedding system. This is what leads to more or less fine- or coarse-grained properties being invoked in the explanans.

Once it is appreciated that functional characterisations help fix organisational relations, it starts to make sense why shifts in grain and analysis come apart. Shifts in grain are only possible relative a given functional contextualisation relative to a specified vocabulary (embodied in a background theory); functional contextualisations set the grain of analysis. Without either implicitly or explicitly contextualising a property or object, the levels of functional decomposability available for the system remain unclear, particularly for systems with nested structure.

For example, to return to the cell case, once a function has been specified, such as oxidative phosphorylation, a number of different forms of analysis can be applied. One might, in the spirit of Marr, ask algorithmic questions about what metabolites contribute to the phosphorylation process, or, alternatively, one might ask questions about what intracellular components and activities are responsible for implementing phosphorylation. Once the grain of analysis has been set shifts among modes of analysis can occur, not the other way round. In running the two kinds of shifts together, the HCL fails to appreciate the role of functional contextualisation in analysis. It mistakenly assumes that the grain of analysis is set by the mode of analysis at certain points.

Mechanists have long appreciated this point. Toulin (1975), for example, writes:

> Indeed, the very organization of organisms – the organization that is sometimes described as though in simply involved a 'hierarchy' of progressive larger structure – can be better viewed as involving a 'ladder' of progressively more complex systems. All these systems, whatever their level of complexity, need to be analysed and understood in terms of the functions they serve, and also of the mechanisms they call into play (p.53).

The question of whether two properties are localisable to the same level is only answerable relative to a particular functional contextualisation (Craver, 2015). It is only once a property has been provided with a contextual function that we can begin to speak of parts being at a 'higher' or 'lower' level. This does not mean we cannot say whether one part is larger or smaller, faster, or slower than another (shifts in grain are still possible). Rather, it means that things like spatial or temporal relations do not by themselves mark organisational divides in the world.

In short, the HCL is problematic for cognitive science because it fails to account for key aspects of cognitive systems. It neglects the role of functional contextualisations, and so fails to capture the iterative applications of modes of analysis and the possibility of nested structure within information processing systems.

5. The HCL Reconsidered

As the preceding argument demonstrates, the strict correspondence relation and the total ordering of the ontological and explanatory hierarchies offer overly restrictive assumptions when applied to complex systems such as information processing systems. Given this, what I want to do now is briefly sketch what an alternative version of the HCL might look like without these problematic elements.

First, consider what a version of the HCL might look like without the strict correspondence relation. If the HCL no longer claims that there is an isomorphic mapping from a system of levels <L, S> to <L', S'>, then while each level and mapping of <L, S> can be assigned to a corresponding level and mapping in <L', S'>, there is a no mapping in the other direction – <L, S> and <L', S'> are no longer structurally equivalent. So, while <L, S> and <L', S'> still denote distinct notions of levels – for example, levels of ontology and modes of analysis – there is now only a mapping from <L, S> to <L', S'>.

To illustrate, return to Marr's (1982) tripartite framework. According to the original version of the HCL, Marr's computational, algorithmic and implementation levels denoted three languages for describing cognition $\{\ell_1, \ell_2, \ell_3\}$ and an ontology of levels $\{\Omega_{1\ldots}\ \Omega_n\}$. Marr's levels formed a non-empty class of modes of analysis, L, and a corresponding non-empty class of levels-specific worlds, Ω_L, such that there was a mapping between each mode of analysis in L and a level-specific world in Ω_L.

However, according to this weaker version of the HCL, while Marr's levels still denote a non-empty class of modes of analysis, L, and a corresponding non-empty class of levels-specific worlds, any level ℓ of <L, S> can now be mapped to one or more level-specific worlds Ω_L in <L', S'>. To say that a sentence Φ is true at the computational mode of analysis, ℓ_1, for instance, is still to say that Φ denotes a set of facts or entities at a corresponding level-specific world, $\Omega_{L,}$ such that $\Phi \in \ell_1$ and $\Phi \in \Omega_{L,}$ but it is no longer the case that ℓ_1 needs to only be mapped to Ω_{L1}. It can be mapped to any level-specific world, as long as the appropriate orderings are preserved.

Second, in terms of total ordering of the hierarchy, while there would still be a mapping of levels *between* <L, S> and <L', S'>, it would no longer be the case that the levels *within* <L, S> or <L', S'> could be given a total ordering. This does not mean that levels cannot still be organised – one could still, for example, map elements of ℓ_1 to elements to level ℓ_2 – but one can no longer tell whether for any two levels that one is necessarily higher or lower than the other. For any two levels in a system of levels, there is function S from either from $\sigma: \ell \rightarrow \ell'$ or from $\sigma: \ell' \rightarrow \ell$ but not both.

So, again, returning to Marr's view as an example, while one can say that a particular set of descriptions at the computational level ℓ_1, ones that map to a level-specific world Ω_{L1}, are 'above' another set of descriptions at the algorithmic level ℓ_2, one cannot say that another set of computational descriptions ℓ_3, ones mapped to a different level-specific world Ω_{L2}, are necessarily 'higher' or 'lower' than ℓ_2. So, in the modified version of the HCL, there are two distinct systems of levels, but there is no strict correspondence or total ordering of levels.[22]

The real question is: can this weaker version of the HCL accommodate the previous worries?

---

[22] Formally, this revised version of the HCL respects the previous three conditions S1- S3. The important different is that drops the total ordering property.

First, with respect to the question of iteratively applying modes of analysis, if the mapping relation between systems of levels is no longer one-to-one, then a given mode of analysis can, in principle, apply to more than one ontological level. A particular language in <L, S>, ℓ_1, no longer needs to be mapped uniquely to a level-specific world in <L', S'> such as Ω_{L1}. It could just as easily map to Ω_{L2} or Ω_{L3}. This modified version of the HCL can now accommodate applying the same mode of analysis, such as Marr's computational level, to both lower-level processes, such as oxidative phosphorylation, as well as higher-level processes, such as object recognition. Without the strict correspondence relation, different modes of analysis are no longer tied to specific grains of analysis. Investigation at a given mode of analysis can span several shifts in grain without a corresponding shift in mode of description.

Second, with respect to the question of nested structure, if systems of levels are no longer locked into a mapped one-to-one, then there are multiple mappings possible different between modes of analysis and levels of ontology. For some complex system, S, composed of ontological levels, Ω_L, a language, ℓ_1, can apply to some subset of Ω_L, while another language, ℓ_2, can apply to a subset of that subset. Unlike Newell or Sun's views, in eschewing the strict correspondence relation, the modified version of the HCL can accommodate the possibility of cognitive systems with nested structure. A given mode of analysis can now address entities and process at spanning multiple grains of analysis within shifts in the mode of description.

Finally, with respect to functional contextualisation, if a system of ontological levels is only partially rather than totally ordered, then how a given levels gets initially fixed is undetermined. Since one can only say on a partial ordering that a level is locally higher or lower, the HCL leaves the question open as to how a particular entity or activity at a given ontological level is initially assigned. In a dropping the commitment to total orderings the same entity or activity can be situated or embedded in different systems depending on the explanatory need. As Potochnik and Guilherme (2020, pp.1317-8) ably put the point: "the question is not which levels are explanatory, but which explanatory style is called for given the potential causal pattern of interest.” We end with a shift toward explanatory pluralism.

6. Conclusion

So, what have I shown and why is it important? First, I have shown that the HCL can, in fact, be found in the cognitive science, and that it has it driven certain explanatory and methodological work. This reveals what many have already suspected: that talk of levels is not just a harmless metaphor but one which can actively drive and structure how thinking proceeds on philosophical and scientific topics (Craver, 2015; Potochnik, 2012, 2017). Second, I have shown that the HCL is specifically problematic when applied to the types of systems studied in cognitive science. Unlike previous worries, which often stem from more general concerns about accommodating outlier cases, the present argument derives its force from the subject matter of cognitive science itself: namely, information processing systems. Third, I have shown that while the HCL is generally unsuitable for cognitive science, there is a way to make the view serviceable. Following a number of other authors, such as Kim (2002), I have shown that a local approach can deliver a conception of levels that is responsive to the needs of a particular domain of inquiry – in the present case, accounting for the role of functional contextualisations, iterative applications of modes of analysis, and the possibility of nested structure. Finally, in applying and extending List’s (2019) formal framework, I have hopefully been able to bring an additional level of precision to the discussion of levels in cognitive science.

What lessons are to be drawn? On the one hand, the lesson is quite general. It is one similar to others drawn elsewhere about levels: we do not need a monolithic and all-inclusive hierarchy of levels to make progress on scientific and philosophical questions. A local, domain specific approach is preferable. On the other hand, the lesson is quite specific: cognitive systems not only require their own unique conception of levels, but such a conception has to be pay special attention to various properties unique to the types of information processing systems studied in cognitive science.

Disclosure Statement

The authors report there are no competing interests to declare.

References

Baxendale, M. (2016). The Layer Cake Model of the World and Non-Reductive Physicalism. *Kriterion – Journal of Philosophy*, 30(1), 39-60.

Bechtel, W. (1994). Levels of Description and Explanation in Cognitive Science. *Minds and Machines*, 4, 1-25.

Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. London: Routledge.

Bermudez, J. (2005). *Philosophy of psychology: A contemporary introduction*. New York: Routledge.

Brooks, D. (2017). In Defense of Levels: Layer Cakes and Guilt by Association. *Biological Theory*, 12(3), 142–156.

Brooks, D, and Eronen, M. (2018). The significance of levels of organization for scientific research: A heuristic approach. *Studies in History and Philosophy of Biological and Biomedical Sciences*. Online first.

Changeux, J. (2017). Climbing Brain Levels of Organisation from Genes to Consciousness. *Trends in Cognitive Sciences*, 21(3), 168-181.

Craver, C. (2007). Constitutive explanatory relevance. *Journal of Philosophical Research*, 32, 1-20.

Craver, C. (2013). Functions and Mechanisms: A Perspectivalist View. In: Huneman P. (eds) Functions: selection and mechanisms. Synthese Library (Studies in Epistemology, Logic, Methodology, and Philosophy of Science), vol 363 (pp 133-158). Springer, Dordrecht.

Craver, C. (2015). Levels. In T. Metzinger & J. M. Windt (Eds). *Open MIND:* 8(T). Frankfurt am Main: MIND Group. 10.15502/9783958570498.

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

Danks, D. (2014). *Unifying the mind: Cognitive representations as graphical models*. Cambridge, MA: MIT Press.

Dawson, M. (1998). *Understanding Cognitive Science*. Blackwell Publisher.

Eronen, M. (2013). No Levels, No Problems: Downward Causation in Neuroscience. *Philosophy of Science*, 80(5), 1042–1052.

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

Fodor, J. (1974). Special Sciences: The Disunity of Science as a Working Hypothesis. *Synthese*, 28: 97-115.

Gallagher, S., Janz, B., Reinerman, L., Trempler, J., & Bockelman, P. (2015). *A Neurophenomenology of Awe and Wonder: Towards a Non-Reductionist Cognitive Science*. Springer.

Green, S., and Batterman, R. (2017). Biology meets physics: Reductionism and multi-scale modeling of morphogenesis. *Studies in the History of Philosophy of Science Part 1*, 61, 20-34. doi: 10.1016/j.shpsc.2016.12.003.

Haueis, P., and Burnston, D. (2021). Evolving Concepts of 'Hierarchy' in Systems Neuroscience. In Fabrizio Calzavarini & Marco Viola (eds.), *Neural Mechanisms: New Challenges in the Philosophy of Neuroscience*.

Kersten, L. (2020). How to be concrete: Mechanistic computation and the abstraction problem. *Philosophical Explorations*, *23*(3), 251-266.

Kersten, L. (2024). An idealised account of mechanistic computation. *Synthese* 203(99), 1-24. https://doi.org/10.1007/s11229-024-04526-x.

Kersten, L., West, R., and Brook, A. (2016). Leveling the Field: Talking Levels in Cognitive Science. In A. Papafragou, D. Grodner, D. Mirman, & J.C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of Cognitive Science Society* (pp.2399-2404). Austin, TX: Cognitive Science Society.

Kim, J. (2002). The Layered Model: Metaphysical Consideration. *Philosophical Explorations*, 5(1), 2–20.

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

List, C. (2019), Levels: Descriptive, Explanatory, and Ontological. *Noûs*, 53: 852-883. https://doi.org/10.1111/nous.12241

Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.

McClamrock, R. (1991). Marr's Three Levels: A Re-Evaluation. *Minds and Machines*, 1, 185-196.

Miłkowski, M., Clowes, R.W., Rucińska, Z., Przegalińska, A., Zawidzki, T., Gies, A., Krueger, J., McGann, M., Afeltowicz, Ł., Wachowski, W.M. and Stjernberg (2018). From Wide Cognition to Mechanisms: A Silent Revolution. *Frontiers in psychology*, 9, 2393.

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

Oppenheim, P., and Putnam, H. (1958). Unity of Science as a Working Hypothesis. In *Minnesota Studies in the Philosophy of Science*, vol. 2, ed. Herbert Feigl, Michael Scriven, and Grover Maxwell, 3–36. Minneapolis: University of Minnesota Press.

Piccinini, G. (2020). *Neurocognitive mechanisms: Explaining Biological Cognition*. Oxford University Press. https://doi.org/10.1093/oso/9780198866282.001.0001.

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

Potochnik, A. (2017). *Idealization and the Aims of Science*. University of Chicago Press.

Potochnik, A., (2022). Our World Isn't Organized into Levels. In Dan Brooks Brooks, James DiFrisco & William C. Wimsatt (eds.), *Levels of Organization in Biology*. Cambridge, USA: MIT Press.

Potochnik, A., and McGill, Brian. (2012). The Limitations of Hierarchical Organization. *Philosophy of Science*, 79(1), 120-140.

Potochnik, A., and Guilherme, S. (2020). Patterns in Cognitive Phenomena and Pluralism of Explanatory Styles. *Topics in Cognitive Science*, 12(4), 1306-1320.

Portides, D. (2018). Idealization and abstraction in scientific modeling. *Synthese*. https://doi.org/10.1007/s11229-018-01919-7.

Pylyshyn, Z. (1984). *Computation and Cognition: Towards a Foundation for Cognitive Science.* Cambridge, MA: MIT Press.

Richardson, R. C. (2009). Multiple realization and methodological pluralism. *Synthese*, 167(3), 473-492.

Rueger, A, and McGivern, P. (2010). Hierarchies and Levels of Reality. *Synthese*, 176(3), 379–97.

Sterelny, K. (1990). *The representational theory of mind: An introduction.* Basil Blackwell.

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

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

Sun, R., Coward, A., and Zenzen, M. (2005). On Levels of Cognitive Modeling. *Philosophical Psychology,* 18(5), 613–637.

Thagard, P. (2005). *Mind: Introduction to Cognitive Science*. MIT Press.

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

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.

Wimsatt, W. (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, W. (1994). The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets. *Canadian Journal of Philosophy*, 20, 207-274.