Integrating Philosophy of Understanding with the Cognitive Sciences

Kareem Khalifa[1*], Farhan Islam[2], J.P. Gamboa[3], Daniel A. Wilkenfeld[4], Daniel Kostić[5]

[1]Department of Philosophy, Middlebury College, Middlebury, VT, USA

[2]Independent Scholar, Madison, WI, USA

[3]Department of History and Philosophy of Science, University of Pittsburgh, Pittsburgh, PA

[4]Department of Acute and Tertiary Care, University of Pittsburgh School of Nursing, Pittsburgh, PA

[5]Institute for Science in Society (ISiS), Radboud University, Nijmegen, The Netherlands

*** Correspondence:**
Kareem Khalifa
kkhalifa@middlebury.edu

Keywords: explanation, understanding, mechanism, computation, topology, dynamic systems, pluralism. (Min.5-Max. 8)

Abstract

We provide two programmatic frameworks for integrating philosophical research on understanding with complementary work in computer science, psychology, and neuroscience. First, philosophical theories of understanding have consequences about how agents should reason if they are to understand that can then be evaluated empirically by their concordance with findings in scientific studies of reasoning. Second, these studies use a multitude of explanations, and a philosophical theory of understanding is well suited to integrating these explanations in illuminating ways.

1 Introduction

Historically, before a discipline is recognized as a science, it is a branch of philosophy. Physicists and chemists began their careers as "natural philosophers" during the Scientific Revolution. Biology and psychology underwent similar transformations throughout the nineteenth and early twentieth centuries. So, one might think philosophical discussions of understanding will be superseded by a "science of understanding."

While we are no great forecasters of the future, we will suggest that philosophical accounts of understanding can make two important scientific contributions. First, they provide a useful repository of hypotheses that can be operationalized and tested by scientists. Second, philosophical accounts of understanding can provide templates for unifying a variety of scientific explanations.

We proceed as follows. Section 2 presents these two frameworks for integrating philosophical ideas about understanding with scientific research. Sections 3 discusses the first of these frameworks, in which philosophical theories of understanding propose hypotheses that are tested and refined by the cognitive sciences. Section 4 discusses the second, in which considerations of understanding

provide criteria for integrating different scientific explanations. Both of our proposals are intended to be programmatic. We hope that many of the relevant details will be developed in future work.

2 Two Frameworks for Integration

As several reviews attest (Baumberger, 2014; Baumberger, Beisbart, & Brun, 2016; Gordon, 2017; Grimm, 2021; Hannon, 2021), understanding has become a lively topic of philosophical research over the past two decades. While some work has been done to integrate these ideas with relevant findings from computer science, psychology, and neuroscience, these interdisciplinary pursuits are relatively nascent. While other frameworks are possible and should be developed, we propose two ways of effecting a more thoroughgoing synthesis between philosophy and these sciences (Figure 1). In the first framework for integrating philosophy with the cognitive sciences—what we call *naturalized epistemology of understanding* (Figure 1A)—the philosophy of understanding provides conjectures about reasoning that are tested and explained by the relevant sciences. In the second integrative framework—*understanding-based integration* (Figure 1B)—the philosophy of understanding provides broad methodological guidelines about how different kinds of scientific explanation complement each other. The two proposals are independent of each other: those unpersuaded by one may still pursue the other. We discuss each in turn.

A. Naturalized Epistemology of Understanding

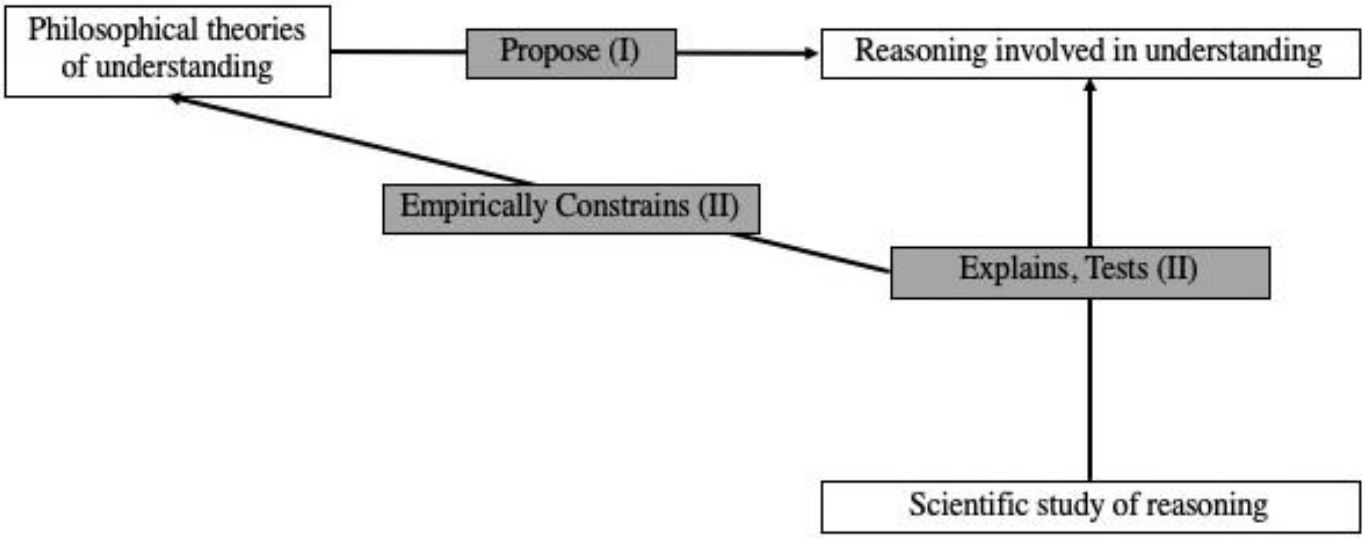

B. Understanding-Based Integration

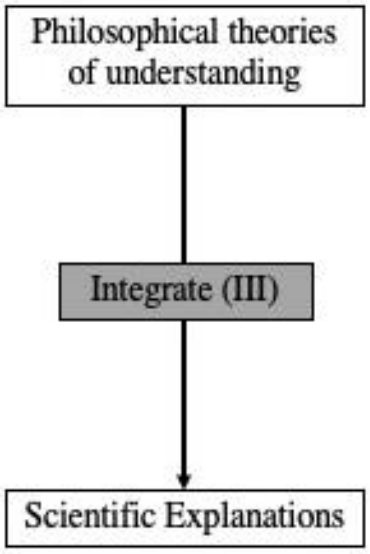

Figure 1. Two Ways to Integrate Philosophical Work on Understanding with Relevant Sciences.

3 Naturalized Epistemology of Understanding

In epistemology, naturalism is the position that philosophical analyses of knowledge, justification, and kindred concepts should be intimately connected with empirical science. Different naturalists specify this connection in different ways; see Rysiew (2020) for a review. Given that philosophical interest in understanding has only recently achieved critical mass, the more specific research program of a naturalized epistemology of understanding is nascent. We propose to organize much existing work according to the framework in Figure 1A. More precisely, philosophical theories of understanding propose how reasoning operates in understanding (Section 3.1), and these proposals are constrained by explanations and empirical tests found in sciences that study this kind of reasoning (Section 3.2).

3.1 Philosophical Theories Propose Reasoning in Understanding (I)

Two kinds of understanding have garnered significant philosophical attention: explanatory understanding (Greco, 2013; Grimm, 2010, 2014; Hills, 2015; Khalifa, 2012, 2013a, 2013b, 2017; Kuorikoski & Ylikoski, 2015; Potochnik, 2017; Strevens, 2013) and objectual understanding (Baumberger, 2019; Baumberger & Brun, 2017; Carter & Gordon, 2014; Dellsén, 2020; Elgin, 2004, 2017; Kelp, 2015; Kvanvig, 2003; Wilkenfeld, 2021). Explanatory understanding involves understanding why or how something is the case. (For terminological convenience, subsequent references to "understanding-why" are elliptical for "understanding-why or -how".) Examples include understanding why Caesar crossed the Rubicon and understanding how babies are made. Objectual understanding is most easily recognized by its grammar: it is the word "understanding" followed immediately by a noun phrase, e.g., understanding Roman history or understanding human reproduction. Depending on the author, the objects of objectual understanding are taken to be subject matters, phenomena, and for some authors (e.g., Wilkenfeld, 2013), physical objects and human behaviors. For instance, it is natural to think of Roman history as a subject matter but somewhat counterintuitive to think of it as a phenomenon. It is more natural to think of, e.g., the unemployment rate in February 2021 as a phenomenon than as a subject matter. Human reproduction, by contrast, can be comfortably glossed as either a subject matter or a phenomenon.

To clarify what they mean by explanatory and objectual understanding, philosophers have disambiguated many other senses of the English word "understanding." Frequently, these senses are briefly mentioned to avoid confusion but are not discussed at length. They are listed in Table 1. Scientists may find these distinctions useful when characterizing the kind of understanding they are studying. That said, we will focus on explanatory understanding hereafter. Thus, unless otherwise noted, all subsequent uses of "understanding" refer exclusively to explanatory understanding.

Kind of understanding	**Typical Complement**	**Examples**
Propositional	that + declarative sentence	I understand that you might not enjoy reading this book.
Broad Linguistic	name of a language	Schatzi understands German.
Narrow Linguistic	what + a linguistic expression + means	Schatzi understands what "Ich bin ein Berliner" means.
Procedural	how + infinitive	Miles understands how to play trumpet.
Non-explanatory Interrogative	embedded question that does not seek an explanation as its answer (most who, where, what, and when questions)	I understand who my friends are. I understand where my friends will be going. I understand what my friends are doing. I understand when my friends need a good laugh.

Table 1. Kinds of understanding that philosophers infrequently discuss (Khalifa 2017, 2)

Virtually all philosophers agree that one can possess an accurate explanation without understanding it, e.g., through rote memorization. In cases such as this, philosophers widely agree that the lack of understanding is due to the absence of significant *inferential* or *reasoning* abilities. However, philosophers disagree about *which* inferences characterize understanding. Three broad kinds of reasoning have emerged. First, some focus on the reasoning required to *construct* or *consider* explanatory models (De Regt, 2017; Newman, 2012, 2013, 2015). Second, others focus on the reasoning required to *evaluate* those explanatory models (Khalifa, 2017). On both these views, explanatory models serve as the *conclusions* of the relevant inferences. However, the third and most prominent kind of reasoning discussed takes explanatory information as *premises* of the relevant reasoning—paradigmatically the inferences about how counterfactual changes in the explanatory variable or *explanans* would result in changes to the dependent variable or *explanandum* ((Bokulich, 2011; Grimm, 2010, 2014; Hills, 2015; Hitchcock & Woodward, 2003; Kuorikoski & Ylikoski, 2015; Le Bihan, 2016; Potochnik, 2017; Rice, 2015; Verreault-Julien, 2017; Wilkenfeld, 2013;

Woodward, 2003). This is frequently referred to as the ability to answer "what-if-things-had-been-different questions." Many of these authors discuss all three of these kinds of reasoning—which we call *explanatory consideration, explanatory evaluation,* and *counterfactual reasoning*—often without explicitly distinguishing them in the ways we have here.

3.2 Scientific Studies of Reasoning's Contributions to the Philosophy of Understanding (II)

A naturalized epistemology of understanding begins with the recognition that philosophers do not have a monopoly on studying these kinds of reasoning. Computer scientists, psychologists, and neuroscientists take explanatory and counterfactual reasoning to be important topics of research. Undoubtedly, each discipline has important insights and contributions. Moreover, these scientific disciplines may raise interesting questions about understanding that are not on the current philosophical agenda.

Cognitive psychological investigations into the nature of explanation and understanding frequently focus on the role of those states in our cognitive lives. To the extent that one can derive a general lesson from this literature, it is probably that both having and seeking explanations aid other crucial cognitive tasks such as prediction, control, and categorization. Developmental psychologists argue that having proper explanations promotes survival, and that at least the sense of understanding evolved to give us an immediate reward for gaining such abilities (Gopnik, 1998). In cognitive psychology, Koslowski, Marasia, Chelenza, and Dublin (2008) have argued that having an explanation better enables thinkers to incorporate evidence into a causal framework. Lombrozo and collaborators have done extensive empirical work investigating the epistemic advantages and occasional disadvantages of simply being prompted to explain new data. They find that under most normal circumstances trying to seek explanations enables finding richer and more useful patterns (Williams & Lombrozo, 2010). This work also has the interesting implication that the value of explanation and understanding depends on the extent to which there are genuine patterns in the world, with fully patterned worlds granting the most advantages from prompts to explain (ibid.), and more exception-laden worlds providing differential benefits (Kon & Lombrozo, 2019). It has also been demonstrated that attempts to explain can (perhaps counterintuitively) systematically mislead. For example, attempts to explain can lead to miscategorization and inaccurate predictions when there are no real patterns in the data (Williams, Lombrozo, & Rehder, 2013). Similarly, laypeople can be misguided by the appearance of irrelevant neuroscientific or otherwise reductive explanations (Hopkins, Weisberg, & Taylor, 2016; Weisberg, Keil, Goodstein, Rawson, & Gray, 2008). In more theoretical work, Lombrozo (2006; Lombrozo & Wilkenfeld, 2019) considers how different kinds of explanation can lead to understanding that is either more or less tied to specific causal pathways connecting explananda and explanantia versus understanding focused on how different pathways can lead to the same end result. Thagard (2012) has argued that explanatory reasoning is key to science's goals both intrinsically and as they contribute to truth and education.

One recent thread in the cognitive science and philosophy of understanding combines insights from information theory and computer science to characterize understanding in terms of data compression. Data compression (Grünwald, 2004) involves the ability to produce large amounts of information from relatively shorter hypotheses and explicitly encoded data sets—in computer science and model-centric physics, there is a burgeoning sense that understanding is tied to pattern recognition and data compression. Petersen (ms) helpfully documents an array of such instances. Li and Vitányi (2008) use compression and explanation almost interchangeably, and at some points even suggest a possible equivalence between compression and the scientific endeavor generally, as in Davies (1990). Tegmark (2014) likewise connects the notion of compression with the explanatory

goals of science. Wilkenfeld (2019) translates the importance of compression to good scientific (and non-scientific) understanding into the idiom of contemporary philosophy of science. While part of the inspiration characterizing understanding in terms of compression comes from the traditional "unificationist" philosophical position that understanding involves having to know fewer brute facts (Friedman, 1974) or argument patterns (Kitcher, 1989), the introduction of compression helps evade some objections to unificationist views, such as the fact that such views require explanations to be arguments (Woodward, 2003) and the fact that they allow for understanding via unification that no actual human agent can readily use (Humphreys, 1993). (Compression as a marker for intelligence has come under recent criticism (e.g., Chollet, 2019) as only accounting for past data and not future uncertainties; we believe Wilkenfeld's (2019) account evades this criticism by defining the relevant compression partially in terms of usefulness, but defending that claim is beyond the scope of this paper).

There has also been more direct work on leveraging insights from computer science in order to try to build explanatory schema and even to utilize those tools to reach conclusions about true explanations. Schank (1986) built a model of computerized explanations in terms of scripts and designed programs to look for the best explanations. Similarly, Thagard (1989, 1992, 2012)—who had previously (1978) done seminal philosophical work on good-making features of explanation and how they should guide theory choice—attempted to automate how computers could use considerations of explanatory coherence to make inferences about what actually occurred.

One underexplored area in the philosophy of understanding and computer science is the extent to which neural nets and deep learning machines can be taken to understand anything. While Turing (1950) famously argued that a machine that could behave sufficiently close to a person could thereby think (and thus, perhaps, understand), many argue that learning algorithms are concerned with prediction *as opposed to* understanding. The most extreme version of this position is Searle's (1980) claim that computers by their nature cannot achieve understanding because it requires semantic capacities when manipulating symbols (i.e., an ability to interpret symbols and operations, and to make further inferences based on those interpretations). Computers at best have merely syntactic capabilities (they can manipulate symbols using sets of instructions, without understanding the meaning of either symbols or operation upon them). However, at the point where deep learning machines have hidden representations (Korb, 2004), can generate new (seemingly theoretical) variables (ibid.), and can be trained to do virtually any task to which computer scientists have set their collective minds (including what looks from the outside like abstract reasoning in IBM's Watson and their Project Debater), it raises vital philosophical questions regarding on what basis we can continue to deny deep learning machines the appellation of "understander".

Elsewhere in cognitive science, early psychological studies of reasoning throughout the 1960s and 1970s focused on deductive reasoning and hypothesis testing (Osman, 2014). A major influence on this trajectory was Jean Piaget's (1952) theory of development, according to which children develop the capacity for hypothetico-deductive reasoning around age 12. The kinds of reasoning studied by psychologists then expanded beyond their logical roots to include more humanistic categories such as moral reasoning (Kohlberg, 1958). The psychology literature offers a rich body of evidence demonstrating how people reason under various conditions. For example, there is ample evidence that performance on reasoning tasks is sensitive to the semantic content of the problem being solved. One interpretation of this phenomenon is that in some contexts, people do not reason by applying content-free inference rules (Cheng & Holyoak, 1985; Cheng, Holyoak, Nisbett, &

Oliver, 1986; Holyoak & Cheng, 1995). This empirical possibility is of particular interest for philosophers. In virtue of their (sometimes extensive) training in formal logic, philosophers' reasoning practices may be atypical of the broader population. This in turn may bias their intuitions about how "people" or "we" reason in various situations, including when understanding. Another issue raised by sensitivity to semantic content is how reasoning shifts depending on the object of understanding. Although the distinctions explicated by philosophers (e.g., explanatory vs. objectual understanding) are clear enough, it is an open empirical question whether and how reasoning differs *within* these categories depending on the particular object and other contextual factors. As a final example, a further insight from psychology is that people may have multiple modes of reasoning that can be applied to the very same problem. Since Wason and Evans (1974) suggested the idea, dual-process theories have dominated the psychology of reasoning.[1] Although both terminology and precise hypotheses vary significantly among dual-process theories (Evans, 2011, 2012), the basic idea is that one system of reasoning is fast and intuitive, relying on prior knowledge, while another is slow and more cognitively demanding. Supposing two or more systems of reasoning can be deployed in the same situation, one important consideration is how they figure in theories about the reasoning involved in understanding. To the extent that philosophical accounts are not merely normative but also aim at describing how people actually reason when understanding, psychological studies provide valuable empirical constraints and theoretical considerations.

With the aid of techniques for imaging brains while subjects perform cognitive tasks, neuroscientists have also made great progress in recent decades on identifying regions of the brain involved in reasoning. While that is certainly a worthwhile goal, it may seem tangential to determining the kind of reasoning that characterizes understanding. Here, we suggest two ways in which findings from neuroscience may help with this endeavor. First, neuroscientific evidence can help resolve debates where behavioral data underdetermine which psychological theory is most plausible. More precisely, in cases where competing psychological models of reasoning make the same behavioral predictions, they can be further distinguished by the kinds of neural networks that would implement the processes they hypothesize (Operskalski & Barbey, 2017). For example, Goel, Buchel, Frith, and Dolan (2000) designed a functional magnetic resonance imaging (fMRI) experiment to test the predictions of dual mechanism theory vs. mental model theory. According to the former, people have distinct mechanisms for form- and content-based reasoning, and the latter should recruit language processing structures in the left hemisphere. Mental model theory, by contrast, claims that reasoning essentially involves iconic representations, i.e., non-linguistic representations whose structure corresponds to the structure of whatever they represent (Johnson-Laird, 2010). In early formulations of the theory, it was assumed that different kinds of reasoning problems depend on the same visuo-spatial mechanisms in the right hemisphere (Johnson-Laird, 1995). Goel et al. (2000) tested the theories against one another by giving subjects logically equivalent syllogisms with and without semantic content. As expected, behavioral performance was similar in both conditions. Neither theory predicts significant behavioral differences. Consistent with both theories, the content-free syllogisms engaged spatial processing regions in the right hemisphere. However, syllogisms with semantic content activated a left hemisphere ventral network that includes

[1] Though see Keren and Schul (2009), Osman (2004), and Stephens, Dunn, and Hayes (2018) for examples of criticisms.

language processing structures like Broca's area. Unsurprisingly, proponents of mental models have disputed the interpretation of the data (Kroger, Nystrom, Cohen, & Johnson-Laird, 2008). We do not take a stance on the issue here. We simply raise the case because it illustrates how neuroscience can contribute to debates between theories of reasoning pitched at the psychological level.

Neuroscientific evidence can also guide the revision of psychological models of understanding and reasoning. The broader point is about cognitive ontology. In the sense we mean here, a cognitive ontology is a set of standardized terms which refer to the entities postulated by a cognitive theory (Janssen, Klein, & Slors, 2017). The point of developing a cognitive ontology is to represent the structure of psychological processes and facilitate communication through a shared taxonomy. One role for neuroscience is to inform the construction of cognitive ontologies. Price and Friston (2005), for instance, defend a strong bottom-up approach. In their view, components in a cognitive model (e.g., a model of counterfactual reasoning) should be included or eliminated depending on our knowledge of functional neuroanatomy. Others agree that neuroscience has a crucial role to play in theorizing about cognitive architecture but reject that it has any special authority in this undertaking (Poldrack & Yarkoni, 2016; Sullivan, 2017). We take no position here on how exactly neuroscience should influence the construction of cognitive models and ontologies. Instead, we highlight this important interdisciplinary issue to motivate the potential value of neuroscience for models of understanding and the reasoning involved in it, including those developed by philosophers.

4 Philosophical Theories of Understanding Integrate Scientific Explanations (III)

Thus, there appear to be ample resources for a naturalized epistemology of understanding, in which explanations and empirical tests from the cognitive sciences empirically constrain philosophical proposals about the kinds of reasoning involved in understanding. However, we offer a second and distinct proposal for how the philosophy of understanding can inform scientific practice: as an account of how different explanations can be integrated (Figure 1B).

Such integration is needed when different explanations of a single phenomenon use markedly different vocabularies and concepts. This diversity of explanations is prevalent in several sciences—including the cognitive sciences. To that end, Section 4.1 presents different kinds of explanations frequently found in the cognitive sciences. Whether these different explanations are complements or competitors to each other raises several issues that are simultaneously methodological and philosophical. To address these issues, Section 4.2 presents a novel account of explanatory integration predicated on the idea that explanations are integrated to the extent that they collectively promote understanding. To illustrate the uniqueness of this account, Section 4.3 contrasts our account of integration with a prominent alternative in the philosophical literature.

Before proceeding, two caveats are in order. First, although we focus on the cognitive sciences, the account of explanatory integration proposed here is perfectly general. In principle, the same account could be used in domains ranging from particle physics to cultural anthropology. Second, our aim is simply to show that our account of integration enjoys some initial plausibility; a more thoroughgoing defense exceeds the current paper's scope.

4.1 A Variety of Scientific Explanations

Puzzles about explanatory integration arise only if there are explanations in need of integration, i.e., explanations whose fit with each other is not immediately obvious. In this section, we provide examples of four kinds of explanations found in the cognitive sciences: mechanistic, computational, topological, and dynamical.

4.1.1 Mechanistic Explanations

Mechanistic explanations are widespread in the cognitive sciences (Bechtel & Richardson, 1993; Craver, 2007; Craver & Tabery, 2019; Glennan, 2017; Illari & Williamson, 2010; Machamer, Darden, & Craver, 2000). Despite extensive discussion in the philosophical literature, there is no consensus on the proper characterization of mechanisms or how exactly they figure in mechanistic explanations.[2] For our purposes, we illustrate basic features of mechanistic explanations by focusing on Glennan's (2017, p. 17) minimal conception of mechanisms:

> A mechanism for a phenomenon consists of entities (or parts) whose activities and interactions are organized so as to be responsible for the phenomenon.

This intentionally broad proposal captures a widely held consensus among philosophers about conditions that are necessary for something to be a mechanism. Where they disagree is about further details, such as the nature and role of causation, regularities, and levels of analysis involved in mechanisms. At minimum, mechanistic explanations account for the phenomenon to be explained (the *explanandum*) by identifying the organized entities, activities, and interactions responsible for it.

Consider the case of the action potential. A mechanistic explanation of this phenomenon specifies parts such as voltage-gated sodium and potassium channels. It describes how activities of the parts, like influx and efflux of ions through the channels, underlie the rapid changes in membrane potential. It shows how these activities are organized such that they are responsible for the characteristic phases of action potentials. For example, the fact that depolarization precedes hyperpolarization is explained in part by the fact that sodium channels open faster than potassium channels. In short, mechanistic explanations spell out the relevant physical details.

Importantly, not all theoretical achievements in neuroscience are mechanistic explanations. As a point of contrast, compare Hodgkin and Huxley's (1952) groundbreaking model of the action potential. With their mathematical model worked out, they were able to predict properties of action potentials and neatly summarize empirical data from their voltage clamp experiments. However, as Hodgkin and Huxley (1952) explicitly pointed out, their equations lacked a physical basis. There is some disagreement among philosophers about how we should interpret the explanatory merits of the model (Craver & Kaplan, 2020; Favela, 2020a; Levy, 2014), but what is clear is that the Hodgkin and Huxley model is a major achievement that is *not* a mechanistic explanation of the action potential. We return to issues such as these in Section 4.3.

4.1.2 Computational Explanations

Mechanistic explanations are sometimes contrasted with other kinds of explanation. In the philosophical literature, computational explanations are perhaps the most prominent alternative. Computational explanations are frequently considered a subset of *functional explanations*. The latter

[2] See Craver (2014) for an overview of the latter issue.

explain phenomena by appealing to their function and the functional organization of their parts (Cummins, 1975, 1983, 2000; Fodor, 1968). Insofar as computational explanations are distinct from other kinds of functional explanations, it is because the functions to which they appeal involve information processing. Hereafter, we focus on computational explanations.

In computational explanations, a phenomenon is explained in terms of a system performing a computation. A computation involves the processing of input information according to a series of specified operations that results in output information. While many computational explanations describe the object of computation as having representational content, some challenge this as a universal constraint on computational explanations (Dewhurst, 2018; Fresco & Miłkowski, 2021; Piccinini, 2015). We will use "information" broadly, such that we remain silent on this issue. Here, "operations" refer to logical or mathematical manipulations on information such as addition, subtraction, equation (setting a value equal to something), "AND", etc. For example, calculating *n!* involves taking in input *n* and calculating the product of all natural numbers less than or equal to *n* and then outputting said product. Thus, we can explain why pressing "5", "!", "=", in sequence on a calculator results in the display reading "120"; the calculator *computes* the factorial.

More detailed computational explanations of this procedure are possible. For example, the calculator performs this computation by storing *n* and iteratively multiplying the stored variable by one less than the previous iteration from *n* to 1. In this case, the operations being used are equation, multiplication, and subtraction. The information upon which those operations are being performed are the inputted value for *n* and the stored variable for the value of the factorial at that iteration.

4.1.3 Topological Explanations

In topological or "network" explanations, a phenomenon is explained by appeal to graph-theoretic properties. Scientists infer a network's structure from data, and then apply various graph-theoretic algorithms to measure its topological properties. For instance, clustering coefficients measure degrees of interconnectedness among nodes in the same neighborhood. Here, a node's *neighborhood* is defined as the set of nodes to which it is directly connected. An individual node's *local* clustering coefficient is the proportion of edges within its neighborhood divided by the number of edges that could possibly exist between the members of its neighborhood. By contrast, a network's *global* clustering coefficient is the ratio of closed triplets to the total number of triplets in a graph. A triplet of nodes is any three nodes that are connected by at least two edges. An *open* triplet is connected by exactly two edges; a *closed* triplet, by three. Another topological property, average (or "characteristic") path length, measures the mean number of edges needed to connect any two nodes in the network.

In their seminal paper, Watts and Strogatz (1998) applied these concepts to a family of graphs and showed how a network's topological structure determines its dynamics. First, *regular graphs* have both high global clustering coefficients and high average path length. By contrast, *random graphs* have low global clustering coefficients and low average path length. Finally, they introduced a third type of *small-world graph* with high clustering coefficient but low average path length.

Highlighting differences between these three types of graphs yields a powerful explanatory strategy. For example, because regular networks have larger average path lengths than small-world networks, things will "diffuse" throughout the former more slowly than the latter, largely due to the greater number of edges to be traversed. Similarly, because random networks have smaller clustering coefficients than small-world networks, things will also spread throughout the former more slowly than the latter, largely due to sparse interconnections within neighborhoods of nodes. Hence, *ceteris*

paribus, propagation/diffusion is faster in small-world networks. This is because the fewer long-range connections between highly interconnected neighborhoods of nodes shorten the distance between neighborhoods of nodes that are otherwise very distant and enables them to behave as if they were first neighbors. For example, Watts and Strogatz showed that the nervous system of *C. elegans* is a small-world network, and subsequent researchers argued that this system's small-world topology explains its relatively efficient information propagation (Bullmore & Sporns, 2012; Latora & Marchiori, 2001).

4.1.4 Dynamical Explanations

In dynamical explanations, phenomena are accounted for using the resources of dynamic systems theory. At root, a system is dynamical if its state space can be described using differential equations, paradigmatically of the following form:

$$\dot{\boldsymbol{x}}(\boldsymbol{t}) = \boldsymbol{f}(\boldsymbol{x}(\boldsymbol{t}); \boldsymbol{p}, \boldsymbol{t})$$

Here, *x* is a vector (often describing the position of the system of interest), ***f*** is a function, *t* is time, and ***p*** is a fixed parameter. Thus, the equation describes the evolution of a system over time. In dynamical explanations, these equations are used to show how values of a quantity at a given time and place would uniquely determine the phenomenon of interest, which is typically treated as values of the same quantity at a subsequent time.

For example, consider dynamical explanations of why bimanual coordination—defined roughly as wagging the index fingers of both hands at the same time—is done either in- or anti-phase. Haken, Kelso, and Bunz (1985) use the following differential equation to model this phenomenon:

$$\frac{d\phi}{dt} = -asin\phi - 2bsin2\phi$$

Here ϕ is relative phase, having a value of either 0 degrees or 180 degrees (representing in- and anti-phase conditions respectively) and *b/a* is the coupling ratio inversely related to the oscillations' frequency. The explanation rests on the fact that only the in- and anti-phase oscillations of the index fingers are basins of attraction.

4.2 Understanding-Based Integration

Thus far, we have surveyed four different kinds of explanation—mechanistic, computational, topological, and dynamical. Moreover, each seems to have some explanatory power for some phenomena. This raises the question as to how these seemingly disparate kinds of explanation can be integrated. We propose a new account of "understanding-based integration" (UBI) to answer this question. A clear account of understanding is needed if it is to integrate explanations. To that end, Section 4.2.1 presents Khalifa's (2017) model of understanding. Section 4.2.2 then extends this account of understanding to provide a framework for explanatory integration.

4.2.1 An Account of Understanding

We highlight two reasons to think that Khalifa's account of understanding is especially promising as a basis for explanatory integration. First, as Khalifa (2019) argues, his is among the most demanding philosophical accounts of understanding. Consequently, it serves as a useful ideal to which scientists should aspire. Second, this ideal is not utopian. This is especially clear with Khalifa's requirement that scientists evaluate their explanations relative to the best available methods and evidence. Indeed, among philosophical accounts of understanding, Khalifa's account is uniquely sensitive to the

centrality of hypothesis testing and experimental design in advancing scientific understanding (Khalifa, 2017, forthcoming), and thus makes contact with workaday scientific practices. In this section, we present its three core principles.

Khalifa's first central principle is the *Explanatory Floor*:

> Understanding why *Y* requires possession of a correct explanation of why *Y*.

The Explanatory Floor's underlying intuition is simple. It seems odd to understand why *Y* while lacking a correct answer to the question, "Why *Y*?" For instance, the person who lacks a correct answer to the question "Why do apples fall from trees?" doesn't understand why apples fall from trees. Since explanations are answers to why-questions, the Explanatory Floor appears platitudinous. Section 4.3.2.1 provides further details about correct explanation.

The Explanatory Floor is only one of three principles comprising Khalifa's account and imposes only a necessary condition on understanding. By contrast, the second principle, the *Nexus Principle*, describes how understanding can improve:

> Understanding why *Y* improves in proportion to the amount of correct explanatory information about *Y* (= *Y's* explanatory nexus) in one's possession.

To motivate the Nexus Principle, suppose that one person can correctly identify two causes of a fire, and another person can only identify one of those causes. *Ceteris paribus*, the former understands why the fire occurred better than the latter. Crucially in what follows, however, "correct explanatory information" is not limited to correct explanations. The explanatory nexus also includes the *relationships* between correct explanations. We return to these "inter-explanatory relationships" below.

Furthermore, recall our earlier remark that gaps in understanding arise when one simply has an accurate representation of an explanation (or explanatory nexus) without significant cognitive ability. This leads to the last principle, the *Scientific Knowledge Principle*:

> Understanding why *Y* improves as one's possession of explanatory information about *Y* bears greater resemblance to scientific knowledge of *Y*'s explanatory nexus.

Once again, we may motivate this with a simple example. Consider two agents who possess the same explanatory information that nevertheless differ in understanding because of their abilities to relate that information to relevant theories, models, methods, and observations. The Scientific Knowledge Principle is intended to capture this idea. Khalifa provides a detailed account of scientific knowledge of an explanation:

> An agent *S* has scientific knowledge of how/why *Y* if and only if there is some *X* such that *S*'s belief that *X explains Y* is the safe result of *S's* scientific explanatory evaluation (SEEing).

The core notions here are safety and SEEing. Safety is an epistemological concept that requires an agent's belief to not easily have been false given the way in which it was formed (Pritchard, 2009). SEEing then describes the way a belief in an explanation should be formed to promote understanding. SEEing consists of three phases:

1. *Considering* plausible potential explanations of how/why *Y*;

2. *Comparing* those explanations using the best available methods and evidence; and
3. Undertaking *commitments* to these explanations on the basis these comparisons. Paradigmatically, commitment entails that one believes only those plausible potential explanations that are decisive "winners" at the phase of comparison.

Thus, scientific knowledge of an explanation is achieved when one's commitment to an explanation could not easily have been false given the way that one considered and compared that explanation to plausible alternative explanations of the same phenomenon.

4.2.2 Understanding-Based Integration

With our account of understanding in hand, we now argue that it provides a fruitful account of how different explanations, such as the ones discussed in Section 4.1, can be integrated. The Nexus Principle is the key engine of integration. As noted above, this principle states that understanding improves in proportion to the amount of explanatory information possessed. In the cognitive sciences, a multitude of factors explain a single phenomenon. According to the Nexus Principle, understanding improves not only when more of these factors are identified, but when the "inter-explanatory relationships" between these factors are also identified.

One "inter-explanatory relationship" is that of *relative goodness*. Some explanations are *better* than others, even if both are correct. For instance, the presence of oxygen is explanatorily relevant to any fire's occurrence. However, oxygen is rarely judged as the *best* explanation of a fire. Per the Nexus Principle, grasping facts such as these enhances one's understanding. Parallel points apply in the cognitive sciences. For example, it has been observed that mental simulations that involve episodic memory engage the default network significantly more than mental simulations that involve semantic memory (Parikh, Ruzic, Stewart, Spreng, & De Brigard, 2018). Hence, episodic memory better explains cases in which the default network was more active during a mental simulation than does semantic memory.

However, correct explanations can stand in other relations than superiority and inferiority. There are also "structural" relationships between different correct explanations. For instance, the aforementioned explanation involving the default network contributes to a more encompassing computational explanation of counterfactual reasoning involving three core stages of counterfactual thought (Van Hoeck, Watson, & Barbey, 2015). First, alternative possibilities to the actual course of events are mentally simulated. Second, consequences are inferred from these simulations. Third, adaptive behavior and learning geared toward future planning and problem-solving occurs. The default network figures prominently in the explanation of (at least) the first of these processes (Figure 2).

As this example illustrates, grasping the relationships between different kinds of explanations can advance scientists' understanding. In Figure 2, a computational account of mental simulation explains certain aspects of counterfactual reasoning, but mental simulation is then explained mechanistically: the default network consists of parts (e.g., ventral medial prefrontal cortex, posterior cingulate cortex) whose activities and interactions (anatomical connections) are organized so as to be responsible for various phenomena related to mental simulations. Quite plausibly, scientific understanding increases when the relationship between these two explanations is discovered.

Importantly, this is but an instance of an indefinite number of other structures consisting of inter-explanatory relationships (see Figure 3 for examples). In all of these structures, we assume that for all i, X_i is a correct explanation of its respective explanandum. Intuitively, a person who could not

distinguish these different explanatory structures would not understand *Y* as well as someone who did. For instance, a person who knew that X_1 only explains *Y* through X_2 in Figure 3A, or that X_1 and X_2 are independent of each other in Figure 3B, or that X_3 is a common explanation or "deep determinant" of both X_1 and X_2 in Figure 3D, etc. seems to have a better understanding than a person who did not grasp these relationships. Undoubtedly, explanations can stand in other relationships that figure in the nexus.

Thus, the Nexus Principle provides useful guidelines for how different kinds of explanations should be integrated. Moreover, we have already seen that different kinds of explanations can stand in fruitful inter-explanatory relationships, and that these relationships enhance our understanding. In some cases, we may find that one and the same phenomenon is explained both mechanistically and non-mechanistically, but one of these explanations will be better than another. As noted above, "better than" and "worse than" are also inter-explanatory relationships. So, the Nexus Principle implies that knowing the relative strengths and weaknesses of different explanations enhances understanding.

The Scientific Knowledge Principle also plays a role in UBI. Suppose that X_1 and X_2 are competing explanations of *Y*. SEEing would largely be achieved when, through empirical testing, X_1 was found to explain significantly more of *Y*'s variance than X_2. This gives scientists grounds for thinking X_1 better explains *Y* than X_2 and thereby bolsters our understanding of *Y*. Importantly, SEEing is also how scientists discover other inter-explanatory relationships. An example is the aforementioned study that identified the inter-explanatory relationships between episodic memory, semantic memory, the default network, and mental simulation (Parikh et al., 2018).

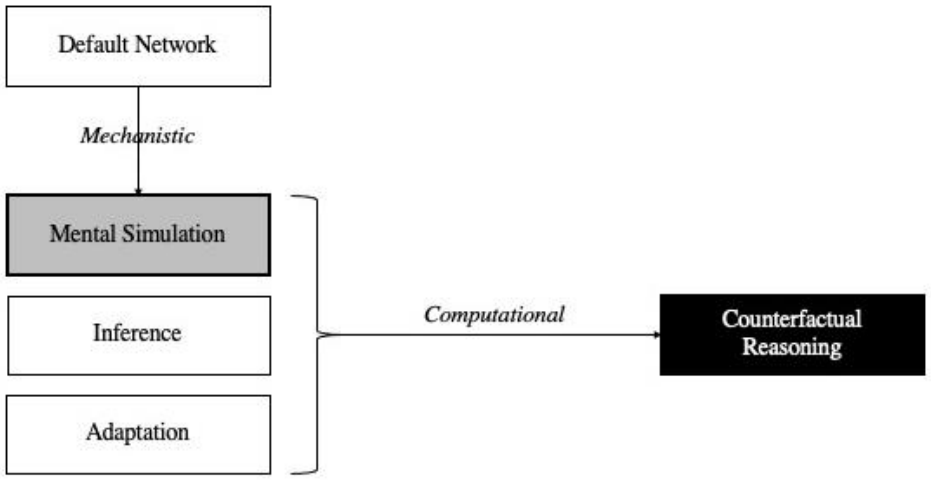

Figure 2. Computational and Mechanistic Explanations Involved in Counterfactual Reasoning

Mental simulation (gray box) both contributes to the computational explanation of counterfactual reasoning (black box) and is mechanistically explained by the activation of the default network.

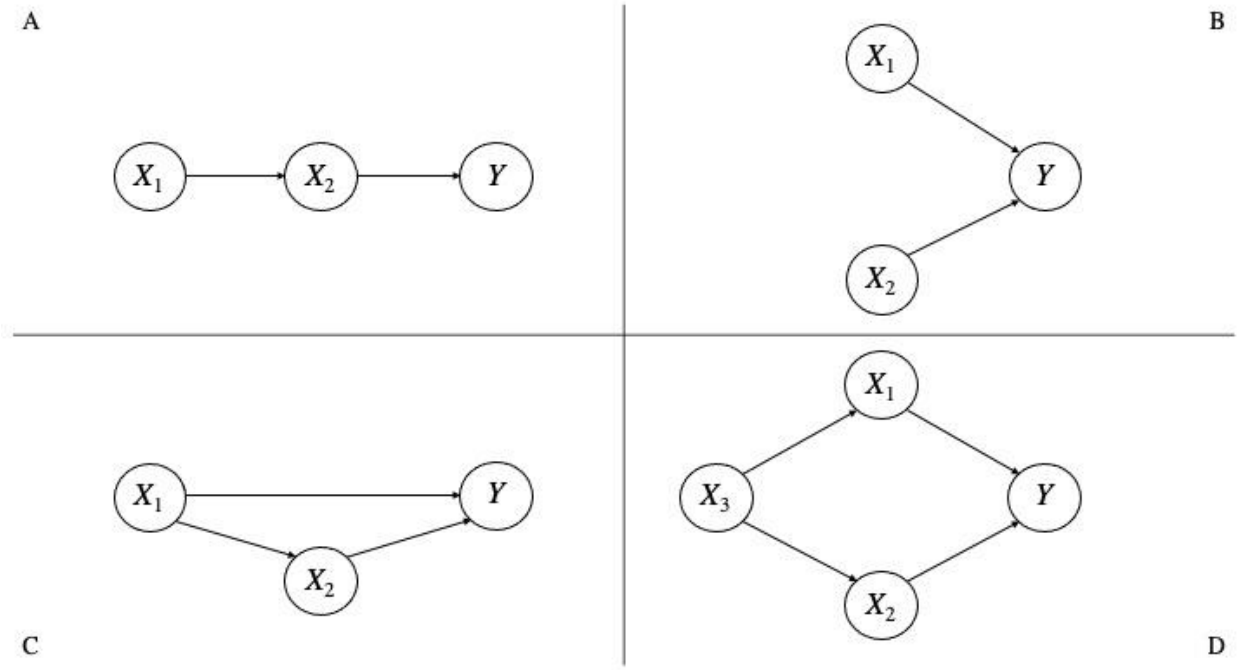

Figure 3. Different Inter-Explanatory Relationships.

Letters at the head of an arrow denote phenomena to be explained; those at the tail, factors that do the explaining. For example, X_1 explains X_2 and X_2 explains Y in Figure 3A.

4.3 Mechanism-Based Integration

Aside from UBI, several other philosophical accounts of explanatory integration in the cognitive sciences are available (Kaplan, 2017; Miłkowski & Hohol, 2020). We provide some preliminary comparisons with the most prominent of these accounts, which we call *mechanism-based integration* (MBI). According to *strong* MBI, all models in the cognitive sciences are explanatory only insofar as they provide information about mechanistic explanations. In response, several critics of MBI—whom we call *pluralists*—have provided examples of putatively non-mechanistic explanation (see Table 2). When presented with putatively non-mechanistic explanations, e.g., of the computational, topological, and dynamical varieties, mechanists (i.e., MBI's proponents) have two strategies available. First, the negative strategy argues that closer scrutiny of the relevant sciences reveals the putatively non-mechanistic explanation to be no explanation at all (Kaplan, 2011; Kaplan & Craver, 2011). The assimilation strategy argues that closer analysis of the relevant sciences reveals the putatively non-mechanistic explanation to be a mechanistic explanation, often of an elliptical nature (Hochstein, 2016; Miłkowski, 2013; Piccinini, 2006, 2015; Piccinini & Craver, 2011; Povich, 2015; Zednik, 2011). Mechanists inclined toward strong MBI frequently use the negative and assimilation strategies in a divide-and-conquer-like manner: the negative strategy applies to some putatively non-mechanistic explanations and the assimilation strategy applies to the rest. However, more prevalent is a *modest* form of MBI that simply applies these strategies to *some* putatively non-mechanistic explanations.

Modest MBI diverges from pluralism on a case-by-case basis. Such cases consist of an explanation where the negative or assimilation strategy seems apt but stands in tension with other considerations that suggest the model is both explanatory and non-mechanistic. On this latter front,

several pluralists argue that computational, topological, and dynamical explanations' formal and mathematical properties are not merely abstract representations of mechanisms (Chirimuuta, 2018; Darrason, 2018; Egan, 2017; Huneman, 2018; Lange, 2017; Rusanen & Lappi, 2016; Serban, 2015; van Rooij & Baggio, 2021; Weiskopf, 2011). Others argue that these explanations cannot (Chemero, 2009; Rathkopf, 2018; Silberstein & Chemero, 2013) or need not (Shapiro, 2019) be decomposed into mechanistic components or that they cannot be intervened upon in the same way that mechanisms are intervened upon (Meyer, 2018; Ross, 2020; Woodward, 2013). Some argue that these putatively non-mechanistic explanations are non-mechanistic because they apply to several different kinds of systems that have markedly different mechanistic structures (Chirimuuta, 2014; Ross, 2015). Pluralist challenges specific to different kinds of explanations can also be found (e.g., Kostić, 2018; Kostić & Khalifa, manuscript).

In what follows, we will show how UBI is deserving of further consideration because it suggests several plausible alternatives to the assimilation and negative strategies. As such, it contrasts with both strong and modest MBI. While we are partial to pluralism, our discussion here is only meant to point to different ways in which mechanists and pluralists can explore the issues that divide them. Future research would determine whether UBI outperforms MBI.

Explanans	**Explanandum**	**Scientific Example**	**Philosophical Work Discussing Example**
Computational Explanations			
Difference of Gaussians	Stereoscopic Vision	Marr (1982); Rodieck (1965)	Bechtel and Shagrir (2015); Egan (2017); Kaplan (2011*); Kaplan and Craver (2011); Rusanen and Lappi (2016); Shagrir (2010); Shapiro (2019)
Exhaustive Search	Recall (Memory)	Sternberg (1969)	Shapiro (2017, 2019)
Gain Field Encoding	Hand-Eye Coordination	Pouget, Deneve, and Duhamel (2002); Pouget and Sejnowski (1997); Shadmehr and Wise (2005); Zipser and Andersen (1988)	Egan (2017); Kaplan (2011*); Rusanen and Lappi (2016); Serban (2015); Shagrir (2006*)

Geon Composition	Object Recognition	Hummel and Biederman (1992)	Buckner (2015*); Povich (2015*); Weiskopf (2011)
Hybrid Computation	Efficiency of Brain	Sarpeshkar (1998)	Chirimuuta (2018)
Inhibitory Feedback	Normalization	Carandini and Heeger (2012)	Chirimuuta (2014); Serban (2015)
Internal Integration	Eye Movement	Seung, Lee, Reis, and Tank (2000)	Egan (2017)
Line Attractor of Choice Axis, Stimuli's Selection Vector	Context-Dependent Decision Making	Mante, Sussillo, Shenoy, and Newsome (2013)	Chirimuuta (2018)
Mapping Non-Coplanar Points to Unique Rigid Configuration	Three-Dimensional Visual Structure of Moving Objects	Ullman (1979)	Egan (2017); Shagrir and Bechtel (2014*)
Optimization of Spatial and Spectral Information Recovery (Gabor Function)	V1 Receptive Fields	Daugman (1985)	Chirimuuta (2014, 2018)
Similarity of Stimulus to Stored Exemplars	Categorization	Kruschke (2008); Love, Medin, and Gureckis (2004)	Buckner (2015*); Povich (2015*); Weiskopf (2011)
Topological Explanations			
Closeness Centrality	Speech and Tonal Processing	Mišić et al. (2018)	Kostić (2020)
Mean Connectivity	Ictogenicity	Helling, Petkov, and Kalitzin (2019)	Kostić and Khalifa (2021)

Motif Frequency	Functional Connectivity	Adachi et al. (2011)	Kostić and Khalifa (2021, manuscript)
Navigation Efficiency, Diffusion Efficiency	Efficiency of Neuronal Communication	Seguin, Razi, and Zalesky (2019)	Kostić (2020)
Network Communicability	Cognitive Control	Gu et al. (2015)	Kostić (2020)
Small-Worldness	Information Propagation	Watts and Strogatz (1998)	Kostić and Khalifa (manuscript)
Dynamical Explanations			
Coupling of Eye and Bodily Movements	Onset of Motor Control	Kelso et al. (1998); Shenoy, Sahani, and Churchland (2013)	Chemero and Silberstein (2008); Favela (2020b); Vernazzani (2019*)
Coupling Ratio	Bimanual Coordination (Relative Phase)	Haken et al. (1985)	Chemero (2000, 2001); Golonka and Wilson (2019*); Kaplan and Craver (2011*); Lamb and Chemero (2014); Meyer (2018); Stepp, Chemero, and Turvey (2011); Zednik (2011*)
Strength of Memory Trace, Salience of Target, Waiting Time, Stance	Infant Reaching (A-not-B Error)	Thelen, Schöner, Scheier, and Smith (2001)	Gervais (2015); Meyer (2018); Povich (forthcoming*); van Eck (2018*); Venturelli (2016);

			Verdejo (2015); Zednik (2011*)
Potassium and Sodium Ion Flows	Neural Excitability	FitzHugh (1961); Hodgkin and Huxley (1952); Nagumo, Arimoto, and Yoshizawa (1962)	Craver and Kaplan (2011*); Favela (2020a, 2020b); Hochstein (2017*); Kaplan and Bechtel (2011*); Kaplan and Craver (2011*); Ross (2015)

Table 2. Putatively non-mechanistic explanations discussed by philosophers.

The *explanans* (first column) is the factor that explains. The *explanandum* (second column) is the phenomenon to be explained. An asterisk indicates that the author takes the explanation to be mechanistic.

4.3.1 Assimilation Strategy

According to mechanists' assimilation strategy, many putatively non-mechanistic explanations are in fact elliptical mechanistic explanations or "mechanism sketches" (Miłkowski, 2013; Piccinini, 2015; Piccinini & Craver, 2011; Povich, 2015, forthcoming; Zednik, 2011). Thus, when deploying the assimilation strategy, mechanists take computational, topological, and dynamical models to fall short of a (complete) mechanistic explanation, but to nevertheless provide important information about such mechanistic explanations. Mechanists have proposed two ways that putatively non-mechanistic explanations can provide mechanistic information, and thereby serve as mechanism sketches. First, putatively non-mechanistic explanations can be *heuristics* for discovering mechanistic explanations. Second, putatively non-mechanistic explanations can *constrain* the space of acceptable mechanistic explanations.

An alternative interpretation is possible. The fact that non-mechanistic models assist in the identification of mechanistic explanations does not entail that the former is a species of the latter. Consequently, putatively non-mechanistic explanations can play these two roles with respect to mechanistic explanations without being mere mechanism sketches. In other words, "genuinely" *non-mechanistic* explanations can guide or constrain the discovery of *mechanistic* explanations. Earlier explanatory pluralists (McCauley, 1986, 1996) already anticipated precursors to this idea, but did not tie it explicitly as a response to mechanists' assimilation strategy.

Moreover, this fits comfortably with our account of scientific explanatory evaluation (SEEing) and hence with UBI. Heuristics of discovery are naturally seen as advancing SEEing's first stage of considering plausible potential explanations. Similarly, since the goal of SEEing is to identify correct explanations and their relationships, it is a consequence of UBI that different kinds of explanations of the related phenomena constrain each other. For instance, suppose that we have two computational explanations of the same phenomenon, and that the key difference between them is that only the first of these is probable given the best mechanistic explanations of that phenomenon.

Then that counts as a reason to treat the first computational explanation as better than the second. Hence, SEEing entails mechanistic explanations can constrain computational explanations.

More generally, UBI can capture the same key inter-explanatory relationships that mechanists prize without assimilating putatively non-mechanistic explanations to mechanistic explanation. Indeed, like many mechanists, UBI suggests that not only do putatively non-mechanistic explanations guide and constrain the discovery of mechanistic explanations, but that the converse is also true. (Section 4.3.2.2 provides an example of this.) Parity of reasoning entails that mechanistic explanations should thereby be relegated to mere "computational, topological, and dynamical sketches" in these cases, but mechanists must resist this conclusion on pain of contradiction. Since UBI captures these important inter-explanatory relationships without broaching the more controversial question of assimilation, it need not determine which models are mere sketches of adequate explanations. Future research would evaluate whether this is a virtue or a vice.

4.3.2 Negative Strategy

Mechanists' assimilation strategy becomes more plausible than the UBI-inspired alternative if there are good grounds for thinking that the criteria that pluralists use to establish putatively non-mechanistic explanations as genuine explanations are insufficient. This is the crux of the mechanists' negative strategy. As with the assimilation strategy, we suggest that UBI provides a suggestive foil to the negative strategy.

The negative strategy's key move is to identify a set of non-explanatory models that pluralists' criteria would wrongly label as explanatory. Two kinds of non-explanatory models—how-possibly and phenomenological models—exemplify this mechanist argument. How-possibly models describe factors that *could* but do not *actually* produce the phenomenon to be explained. For instance, most explanations begin as conjectures or untested hypotheses. Those that turn out to be false will be how-possibly explanations. Phenomenological models, which accurately describe or predict the target phenomenon without explaining it, provide a second basis for the negative strategy. Paradigmatically, phenomenological models correctly represent non-explanatory correlations between two or more variables. Mechanists claim that pluralist criteria of explanation will wrongly classify some how-possibly and some phenomenological models as correct explanations. By contrast, since models that accurately represent mechanisms are "how-actually models," i.e., models that cite explanatory factors responsible for the phenomenon of interest, MBI appears well-positioned to distinguish correct explanations from how-possibly and phenomenological models.

However, UBI can distinguish correct explanations from how-possibly and phenomenological models. Moreover, it can do so in two distinct ways that do not appeal to mechanisms. First, it can do so on what we call *structural* grounds, i.e., by identifying non-mechanistic criteria of explanation that are sufficient for funding the distinction. It can also defuse the negative strategy on what we call *procedural* grounds, i.e., by showing that the procedures and methods that promote understanding also distinguish correct explanations from these non-explanatory models.

4.3.2.1.Structural Defenses

We suggest that the following provides a structural defense against the negative strategy:

If *X* correctly explains *Y,* then the following are true:

(1) *Accuracy Condition: X* is an accurate representation, and

(2) *Counterfactual Condition*: Had the objects, processes, etc. represented by *X* been different, then *Y* would have been different.

These are only necessary conditions for correct explanations. They are also sufficient for distinguishing correct explanations from how-possibly and phenomenological models but are likely insufficient for distinguishing correct explanations from every other kind of non-explanatory model. Identifying these other models is a useful avenue for future iterations of the negative strategy and responses thereto.

Situating this within UBI, these conditions are naturally seen as elaborating the Explanatory Floor, which claims that understanding a phenomenon requires possession of a correct explanation. Crucially, mechanists and pluralists alike widely accept these as requirements on correct explanations, though we discuss some exceptions below. Reasons for their widespread acceptance becomes clear with a simple example. Consider a case in which it is hypothesized that taking a certain medication (*X*) explains recovery from an illness (*Y*). If it were discovered that patients had not taken the medication, then this hypothesis would violate the accuracy condition. Intuitively, it would not be a correct explanation, but it would be a how-possibly model.

More generally, how-possibly models are correct explanations *modulo* satisfaction of the accuracy condition. Consequently, pluralists can easily preserve this distinction without appealing to mechanisms; accuracy is sufficient. Just as mechanisms can be either accurately or inaccurately represented, so too can computations, topological structures, and system dynamics be either accurately or inaccurately represented. Similarly, just as inaccurate mechanistic models can be how-possibly models but cannot be correct explanations, so too can inaccurate computational, topological, and dynamical models be how-possibly models but cannot be how-actually models.

Analogously, the counterfactual condition preserves the distinction between correct explanations and phenomenological models. Suppose that our hypothesis about recovery is confounded by the fact that patients' recovery occurred two weeks after the first symptoms, and that this is the typical recovery time for anyone with the illness in question, regardless of whether they take medication. Barring extenuating circumstances, e.g., that the patients are immunocompromised, these facts would seem to cast doubt upon the claim that the medication made a difference to their recovery. In other words, they cast doubt on the following counterfactual: had a patient not taken the medication, then that patient would not have recovered when she did. Consequently, the hypothesis about the medication explaining recovery violates the counterfactual condition. Moreover, the hypothesis does not appear to be correct, but would nevertheless describe the patients' situation, i.e., it would be a phenomenological model.

More generally, phenomenological models are correct explanations *modulo* satisfaction of the counterfactual condition. Just as a mechanistic model may accurately identify interacting parts of a system that correlate with but do not explain its behavior, a non-mechanistic model may accurately identify computational processes, topological structures, and dynamical properties of a system that correlate with but do not explain its behavior. In both cases, the counterfactual condition accounts for the models' explanatory shortcomings; no appeal to mechanisms is needed.

4.3.2.2.Procedural Defenses

Admittedly, structural defenses against the negative strategy are not unique to UBI; other pluralists who are agnostic about UBI have invoked them in different ways. By contrast, our second *procedural* defense against the negative strategy is part and parcel to UBI. Procedural defenses show

that the procedures that promote understanding also distinguish correct explanations from how-possibly and phenomenological models.

The Scientific Knowledge Principle characterizes the key procedures that simultaneously promote understanding and distinguish correct explanations from these non-explanatory models. Recall that SEEing consists of three stages: *considering* plausible potential explanations of a phenomenon, *comparing* them using the best available methods, and forming *commitments* to explanatory models based on these comparisons. This provides a procedural defense against the negative strategy. How-possibly and phenomenological models will only be acceptable in the first stage of SEEing: prior to their deficiencies being discovered, they frequently deserve *consideration* as possible explanations of a phenomenon. By contrast, correct explanations must "survive" the remaining stages of SEEing: they must pass certain empirical tests at the stage of comparison such that they are acceptable at the stage of commitment. Indeed, it is often through SEEing that scientists come to distinguish correct explanations from how-possibly and phenomenological models.

Crucially, consideration is most effective when it does not prejudge what makes something genuinely explanatory. This minimizes the possibility of missing out on a fruitful hypothesis. Consequently, both mechanistic and non-mechanistic explanations should be included at this initial stage of SEEing. However, our procedural defense supports pluralism only if some computational, topological, or dynamical explanations are acceptable in light of rigorous explanatory comparisons. As we see it, this is a strength of our procedural defense, for it uses the empirical resources of our best science to adjudicate debates between mechanists and pluralists that often appear intractable from the philosophical armchair.

Nevertheless, we can point to an important kind of explanatory comparison—which we call *control-and-contrast*—that deserves greater philosophical and scientific attention when considering explanatory integration in the cognitive sciences. Control-and-contrast proceeds as follows. Let X_1 and X_2 be two potential explanations of Y under consideration. Next, run two controlled experiments: one in which the explanatory factors in X_1 are absent but those in X_2 are present and the second in which the explanatory factors in X_1 are present but those in X_2 are absent. If Y is only present in the first experiment, then the pair of experiments suggests that X_2 is a better explanation of Y than X_1. Conversely, if Y is only present in the second experiment, the pair of experiments suggests that X_1 is a better explanation of Y than X_2. If Y is present in both experiments, the experiments are inconclusive. If Y is absent in both experiments, then the experiments suggest that the combination of X_1 and X_2 better explains Y than either X_1 or X_2 does in isolation. Since we suggest that both mechanistic and non-mechanistic explanations should be considered and thereby play the roles of X_1 and X_2, we also suggest that which of these different kinds of explanations is correct for a given phenomenon Y should frequently be determined by control-and-contrast.

In some cases, scientists are only interested in controlling-and-contrasting explanations of the same kind. However, even in these cases, the controls are often best described in terms of other kinds of explanation. For instance, as discussed above, the default mode network mechanistically explains mental simulations involved in episodic memory. By contrast, when mental simulations involve semantic memory, inferior temporal and lateral occipital regions play a more pronounced role (Parikh et al., 2018). Both episodic and semantic memory are functional or computational concepts that can figures as controls in different experiments designed to discover which of these mechanisms explains a particular kind of mental simulation. Less common is controlling-and-contrasting explanations of different kinds. Perhaps this is a lacuna in current research. Alternatively, it may turn out that

different kinds of explanation rarely compete and are more amenable to integration in the ways outlined above.

The procedural defense complements the structural defense in two ways. First, not all pluralists accept the accuracy condition. Their motivations for this are twofold. First, given that science is a fallible enterprise, our best explanations today are likely to be refuted. Second, many explanations invoke idealizations, i.e., known inaccuracies that nevertheless enhance understanding. The procedural defense does not require the accuracy condition but can still preserve the distinction between correct explanations and non-explanatory models. Instead, the procedural defense only requires that correct explanations be acceptable on the basis of the best available scientific methods and evidence.

Second, tests such as control-and-contrast regiment the subjunctive conditionals that characterize the counterfactual condition. In evaluating counterfactuals, it is notoriously difficult to identify what must be held constant, what can freely vary without altering the truth-value of the conditional, and what must vary in order to determine the truth-value of the conditional. Our account of explanatory evaluation points to important constraints on this process. Suppose that we are considering two potential explanations X_i and X_j of some phenomenon Y. To compare these models, we will be especially interested in counterfactuals such as, "Had the value of X_i been different (but the value of X_j had remained the same), then the value of Y would have been different," and also, "Had the value of X_i been different (but the value of X_j had remained the same), then the value of Y would have been the same." These are precisely the kinds of counterfactuals that will be empirically supported or refuted by control-and-contrast.

5 Conclusion

Fruitful connections between the philosophy and science of understanding can be forged. In a naturalized epistemology of understanding, philosophical claims about various forms of explanatory and counterfactual reasoning are empirically constrained by scientific tests and explanations. By contrast, in understanding-based integration, the philosophy of understanding contributes to the science of understanding by providing broad methodological prescriptions as to how diverse explanations can be woven together. Specifically, understanding-based integration includes identification of inter-explanatory relationships, consideration of different kinds of explanations, and evaluation of these explanations using methods such as control-and-contrast. As our suggestions have been of a preliminary character, we hope that future collaborations between philosophers and scientists will advance our understanding of understanding.

Adachi, Y., Osada, T., Sporns, O., Watanabe, T., Matsui, T., Miyamoto, K., & Miyashita, Y. (2011). Functional Connectivity between Anatomically Unconnected Areas Is Shaped by Collective Network-Level Effects in the Macaque Cortex. *Cerebral Cortex, 22*(7), 1586-1592. doi:10.1093/cercor/bhr234

Baumberger, C. (2014). Types of understanding: Their nature and their relation to knowledge. *Conceptus, 40*(98), 67–88.

Baumberger, C. (2019). Explicating Objectual Understanding: Taking Degrees Seriously. *Journal for General Philosophy of Science, 50*(3), 367-388. doi:10.1007/s10838-019-09474-6

Baumberger, C., Beisbart, C., & Brun, G. (2016). What is understanding? An overview of recent debates in epistemology and philosophy of science. In S. R. Grimm, C. Baumberger, & S.

741 Ammon (Eds.), *Explaining understanding: New perspectives from epistemology and*
742 *philosophy of science* (pp. 1-34). New York: Routledge.

743 Baumberger, C., & Brun, G. (2017). Dimensions of Objectual Understanding. In S. G. Christoph
744 Baumberger & S. Ammon (Eds.), *Explaining understanding: New perspectives from*
745 *epistemology and philosophy of science* (pp. 165-189): Routledge.

746 Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity : decomposition and localization*
747 *as strategies in scientific research*. Princeton, N.J.: Princeton University Press.

748 Bechtel, W., & Shagrir, O. (2015). The Non-Redundant Contributions of Marr's Three Levels of
749 Analysis for Explaining Information-Processing Mechanisms. *Topics in Cognitive Science,*
750 *7*(2), 312-322. doi:https://doi.org/10.1111/tops.12141

751 Bokulich, A. (2011). How scientific models can explain. *Synthese, 180*(1), 33-45.
752 doi:10.1007/s11229-009-9565-1

753 Buckner, C. (2015). Functional kinds: a skeptical look. *Synthese, 192*(12), 3915-3942.

754 Bullmore, E., & Sporns, O. (2012). The economy of brain network organization. *Nature Reviews*
755 *Neuroscience, 13*(5), 336-349. doi:10.1038/nrn3214

756 Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature*
757 *Reviews Neuroscience, 13*(1), 51-62. doi:10.1038/nrn3136

758 Carter, J. A., & Gordon, E. C. (2014). Objectual understanding and the value problem. *American*
759 *philosophical quarterly, 51*(1), 1-13.

760 Chemero, A. (2000). Anti-Representationalism and the Dynamical Stance. *Philosophy of Science,*
761 *67*(4), 625-647. doi:10.1086/392858

762 Chemero, A. (2001). Dynamical explanation and mental representations. *Trends in cognitive*
763 *sciences, 5*(4), 141-142.

764 Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, Mass.: MIT Press.

765 Chemero, A., & Silberstein, M. (2008). After the Philosophy of Mind: Replacing Scholasticism with
766 Science. *Philosophy of Science, 75*(1), 1-27. doi:10.1086/587820

767 Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive psychology, 17*(4),
768 391-416. doi:https://doi.org/10.1016/0010-0285(85)90014-3

769 Cheng, P. W., Holyoak, K. J., Nisbett, R. E., & Oliver, L. M. (1986). Pragmatic versus syntactic
770 approaches to training deductive reasoning. *Cognitive psychology, 18*(3), 293-328.
771 doi:https://doi.org/10.1016/0010-0285(86)90002-2

772 Chirimuuta, M. (2014). Minimal models and canonical neural computations: the distinctness of
773 computational explanation in neuroscience. *Synthese, 191*(2), 127-153. doi:10.1007/s11229-
774 013-0369-y

775 Chirimuuta, M. (2018). Explanation in Computational Neuroscience: Causal and Non-causal. *The*
776 *British Journal for the Philosophy of Science, 69*(3), 849-880. doi:10.1093/bjps/axw034

777 Chollet, F. (2019). On the measure of intelligence. *arXiv preprint arXiv:1911.01547*.

778 Craver, C. F. (2007). *Explaining the brain: mechanisms and the mosaic unity of neuroscience*.
779 Oxford: Clarendon Press.

Craver, C. F. (2014). The Ontic Account of Scientific Explanation. In I. M. Kaiser, R. O. Scholz, D. Plenge, & A. Hüttemann (Eds.), *Explanation in the Special Sciences: The Case of Biology and History* (pp. 27-52). Dordrecht: Springer Netherlands.

Craver, C. F., & Kaplan, D. M. (2011). Towards a Mechanistic Philosophy of Neuroscience. In S. French & J. Saatsi (Eds.), *Continuum Companion to the Philosophy of Science* (pp. 268): Continuum.

Craver, C. F., & Kaplan, D. M. (2020). Are More Details Better? On the Norms of Completeness for Mechanistic Explanations. *The British Journal for the Philosophy of Science, 71*(1), 287-319. doi:10.1093/bjps/axy015

Craver, C. F., & Tabery, J. (2019). Mechanisms in Science. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 2019 Edition)*. URL = <https://plato.stanford.edu/archives/sum2019/entries/science-mechanisms/>.

Cummins, R. C. (1975). Functional analysis. *Journal of philosophy, 72*(20), 741-765.

Cummins, R. C. (1983). *The nature of psychological explanation*. Cambridge, Mass.: MIT Press.

Cummins, R. C. (2000). How does it work?" versus" what are the laws?": Two conceptions of psychological explanation. *Explanation and cognition*, 117-144.

Darrason, M. (2018). Mechanistic and topological explanations in medicine: the case of medical genetics and network medicine. *Synthese, 195*(1), 147-173. doi:10.1007/s11229-015-0983-y

Daugman, J. G. (1985). Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America A, 2*(7), 1160-1169. doi:10.1364/JOSAA.2.001160

Davies, P. C. (1990). Why is the physical world so comprehensible. In W. Zurek (Ed.), *Complexity, entropy and the physics of information* (pp. 61-70): Addison-Wesley Publishing Company.

De Regt, H. W. (2017). *Understanding scientific understanding*. New York: Oxford University Press.

Dellsén, F. (2020). Beyond Explanation: Understanding as Dependency Modelling. *The British Journal for the Philosophy of Science, 71*(4), 1261-1286. doi:10.1093/bjps/axy058

Dewhurst, J. (2018). Individuation without Representation. *The British Journal for the Philosophy of Science, 69*(1), 103-116. doi:10.1093/bjps/axw018

Egan, F. (2017). Function-Theoretic Explanation and the Search for Neural Mechanisms. In D. M. Kaplan (Ed.), *Explanation and Integration in Mind and Brain Science* (pp. 145-163). Oxford: Oxford University Press.

Elgin, C. Z. (2004). True enough. *Philosophical issues, 14*(1), 113-131. Retrieved from http://dx.doi.org/10.1111/j.1533-6077.2004.00023.x

Elgin, C. Z. (2017). *True enough*. Cambridge, MA: MIT Press.

Evans, J. S. B. T. (2011). Dual-process theories of reasoning: Contemporary issues and developmental applications. *Developmental Review, 31*(2), 86-102. doi:https://doi.org/10.1016/j.dr.2011.07.007

Evans, J. S. B. T. (2012). Dual-process theories of deductive reasoning: Facts and fallacies. In *The Oxford handbook of thinking and reasoning.* (pp. 115-133). New York, NY, US: Oxford University Press.

820 Favela, L. H. (2020a). The dynamical renaissance in neuroscience. *Synthese*. doi:10.1007/s11229-
821 020-02874-y

822 Favela, L. H. (2020b). Dynamical systems theory in cognitive science and neuroscience. *Philosophy*
823 *compass, 15*(8), e12695. doi:https://doi.org/10.1111/phc3.12695

824 FitzHugh, R. (1961). Impulses and Physiological States in Theoretical Models of Nerve Membrane.
825 *Biophysical Journal, 1*(6), 445-466. doi:https://doi.org/10.1016/S0006-3495(61)86902-6

826 Fodor, J. A. (1968). *Psychological Explanation: An Introduction To The Philosophy Of Psychology*.
827 New York: Random House.

828 Fresco, N., & Miłkowski, M. (2021). Mechanistic Computational Individuation without Biting the
829 Bullet. *The British Journal for the Philosophy of Science, 72*(2), 431-438.
830 doi:10.1093/bjps/axz005

831 Friedman, M. (1974). Explanation and scientific understanding. *Journal of philosophy, 71*(1), 5-19.

832 Gervais, R. (2015). Mechanistic and non-mechanistic varieties of dynamical models in cognitive
833 science: explanatory power, understanding, and the 'mere description' worry. *Synthese,*
834 *192*(1), 43-66. doi:10.1007/s11229-014-0548-5

835 Glennan, S. (2017). *The new mechanical philosophy* (First edition. ed.). Oxford: Oxford University
836 Press.

837 Goel, V., Buchel, C., Frith, C., & Dolan, R. J. (2000). Dissociation of Mechanisms Underlying
838 Syllogistic Reasoning. *NeuroImage, 12*(5), 504-514.
839 doi:https://doi.org/10.1006/nimg.2000.0636

840 Golonka, S., & Wilson, A. D. (2019). Ecological mechanisms in cognitive science. *Theory &*
841 *Psychology, 29*(5), 676-696. doi:10.1177/0959354319877686

842 Gopnik, A. (1998). Explanation as Orgasm. *Minds and Machines, 8*(1), 101-118.
843 doi:10.1023/A:1008290415597

844 Gordon, E. C. (2017). Understanding in Epistemology. In *Internet Encyclopedia of Philosophy*.
845 https://iep.utm.edu/understa/.

846 Greco, J. (2013). Episteme: knowledge and understanding. In K. Timpe & C. A. Boyd (Eds.), *Virtues*
847 *and their vices* (pp. 285-301). Oxford: Oxford University Press.

848 Grimm, S. R. (2010). The goal of understanding. *Studies in the history and philosophy of science,*
849 *41*(4), 337-344. doi:10.1016/j.shpsa.2010.10.006

850 Grimm, S. R. (2014). Understanding as knowledge of causes. In A. Fairweather (Ed.), *Virtue*
851 *epistemology naturalized* (Vol. 366, pp. 329-345). Dordecht: Springer International
852 Publishing.

853 Grimm, S. R. (2021). Understanding. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*.

854 Grünwald, P. (2004). A tutorial introduction to the minimum description length principle. *arXiv*
855 *preprint math/0406077*.

856 Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., . . . Bassett, D. S.
857 (2015). Controllability of structural brain networks. *Nature Communications, 6*(1), 8414.
858 doi:10.1038/ncomms9414

859 Haken, H., Kelso, J. A. S., & Bunz, H. (1985). A theoretical model of phase transitions in human
860 hand movements. *Biological Cybernetics, 51*(5), 347-356. doi:10.1007/BF00336922

Hannon, M. (2021). RECENT WORK IN THE EPISTEMOLOGY OF UNDERSTANDING. *American philosophical quarterly, 58*(3), 269-290. doi:10.2307/48616060

Helling, R. M., Petkov, G. H., & Kalitzin, S. N. (2019). *Expert system for pharmacological epilepsy treatment prognosis and optimal medication dose prescription: computational model and clinical application*. Paper presented at the Proceedings of the 2nd International Conference on Applications of Intelligent Systems, https://doi.org/10.1145/3309772.3309775. https://doi.org/10.1145/3309772.3309775

Hills, A. (2015). Understanding why. *Noûs, 49*(2), 661-688. doi:10.1111/nous.12092

Hitchcock, C. R., & Woodward, J. (2003). Explanatory generalizations, part II: plumbing explanatory depth. *Noûs, 37*(2), 181-199.

Hochstein, E. (2016). One mechanism, many models: a distributed theory of mechanistic explanation. *Synthese, 193*(5), 1387-1407.

Hochstein, E. (2017). Why one model is never enough: a defense of explanatory holism. *Biology & Philosophy, 32*(6), 1105-1125. doi:10.1007/s10539-017-9595-x

Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology, 117*(4), 500-544. doi:https://doi.org/10.1113/jphysiol.1952.sp004764

Holyoak, K. J., & Cheng, P. W. (1995). Pragmatic reasoning with a point of view. *Thinking & Reasoning, 1*(4), 289-313. doi:10.1080/13546789508251504

Hopkins, E. J., Weisberg, D. S., & Taylor, J. C. V. (2016). The seductive allure is a reductive allure: People prefer scientific explanations that contain logically irrelevant reductive information. *Cognition, 155*, 67-76. doi:https://doi.org/10.1016/j.cognition.2016.06.011

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99*(3), 480-517. doi:10.1037/0033-295X.99.3.480

Humphreys, P. (1993). Greater Unification Equals Greater Understanding? *Analysis, 53*(3), 183-188. doi:10.2307/3328470

Huneman, P. (2018). Outlines of a theory of structural explanations. *Philosophical studies, 175*(3), 665-702. doi:10.1007/s11098-017-0887-4

Illari, P. M., & Williamson, J. (2010). Function and organization: comparing the mechanisms of protein synthesis and natural selection. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences, 41*(3), 279-291. doi:https://doi.org/10.1016/j.shpsc.2010.07.001

Janssen, A., Klein, C., & Slors, M. (2017). What is a cognitive ontology, anyway? *Philosophical Explorations, 20*(2), 123-128. doi:10.1080/13869795.2017.1312496

Johnson-Laird, P. N. (1995). Mental models, deductive reasoning, and the brain. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences* (pp. 999--1008): MIT Press.

Johnson-Laird, P. N. (2010). Mental models and human reasoning. *Proceedings of the National Academy of Sciences, 107*(43), 18243-18250. doi:10.1073/pnas.1012933107

Kaplan, D. M. (2011). Explanation and description in computational neuroscience. *Synthese, 183*(3), 339-373.

Kaplan, D. M. (2017). *Explanation and integration in mind and brain science* (First edition. ed.). Oxford, United Kingdom: Oxford University Press.

Kaplan, D. M., & Bechtel, W. (2011). Dynamical Models: An Alternative or Complement to Mechanistic Explanations? *Topics in Cognitive Science, 3*(2), 438-444. doi:https://doi.org/10.1111/j.1756-8765.2011.01147.x

Kaplan, D. M., & Craver, C. F. (2011). The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective. *Philosophy of Science, 78*(4), 601-627. doi:10.1086/661755

Kelp, C. (2015). Understanding phenomena. *Synthese, 192*(12), 3799–3816. doi:10.1007/s11229-014-0616-x

Kelso, J. A. S., Fuchs, A., Lancaster, R., Holroyd, T., Cheyne, D., & Weinberg, H. (1998). Dynamic cortical activity in the human brain reveals motor equivalence. *Nature, 392*(6678), 814-818. doi:10.1038/33922

Keren, G., & Schul, Y. (2009). Two Is Not Always Better Than One:A Critical Evaluation of Two-System Theories. *Perspectives on Psychological Science, 4*(6), 533-550. doi:10.1111/j.1745-6924.2009.01164.x

Khalifa, K. (2012). Inaugurating understanding or repackaging explanation? *Philosophy of Science, 79*(1), 15-37. Retrieved from http://www.jstor.org/stable/10.1086/663235

Khalifa, K. (2013a). Is understanding explanatory or objectual? *Synthese, 190*(6), 1153-1171. doi:10.1007/s11229-011-9886-8

Khalifa, K. (2013b). The role of explanation in understanding. *British Journal for the Philosophy of Science, 64*(1), 161-187. doi:10.1093/bjps/axr057

Khalifa, K. (2017). *Understanding, Explanation, and Scientific Knowledge*. Cambridge: Cambridge University Press.

Khalifa, K. (2019). Is *Verstehen* Scientific Understanding? *Philosophy of the Social Sciences, 49*(4), 282-306. doi:10.1177/0048393119847104

Khalifa, K. (forthcoming). Should Friends and Frenemies of Understanding be Friends? Discussing de Regt. In K. Khalifa, I. Lawler, & E. Shech (Eds.), *Scientific Understanding and Representation: Modeling in the Physical Sciences*. London: Routledge.

Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In P. Kitcher & W. C. Salmon (Eds.), *Scientific explanation* (Vol. XIII, pp. 410-506). Minneapolis: University of Minnesota Press.

Kohlberg, L. (1958). *The development of modes of moral thinking and choice in the years 10 to 16.* (PhD). University of Chicago, Chicago. /z-wcorg/ database.

Kon, E., & Lombrozo, T. (2019). Scientific Discovery and the Human Drive to Explain. In D. A. Wilkenfeld & R. Samuels (Eds.), *Advances in Experimental Philosophy of Science* (pp. 15). London: Routledge.

Korb, K. B. (2004). Introduction: Machine Learning as Philosophy of Science. *Minds and Machines, 14*(4), 433-440. doi:10.1023/B:MIND.0000045986.90956.7f

Koslowski, B., Marasia, J., Chelenza, M., & Dublin, R. (2008). Information becomes evidence when an explanation can incorporate it into a causal framework. *Cognitive Development, 23*(4), 472-487. doi:https://doi.org/10.1016/j.cogdev.2008.09.007

Kostić, D. (2018). The topological realization. *Synthese, 195*(1), 79-98. doi:10.1007/s11229-016-1248-0

Kostić, D. (2020). General theory of topological explanations and explanatory asymmetry. *Philosophical Transactions of the Royal Society B: Biological Sciences, 375*(1796), 20190321. doi:doi:10.1098/rstb.2019.0321

Kostić, D., & Khalifa, K. (2021). The directionality of topological explanations. *Synthese*. doi:10.1007/s11229-021-03414-y

Kostić, D., & Khalifa, K. (manuscript). Decoupling Topological Explanation from Mechanisms.

Kroger, J. K., Nystrom, L. E., Cohen, J. D., & Johnson-Laird, P. N. (2008). Distinct neural substrates for deductive and mathematical processing. *Brain Research, 1243*, 86-103. doi:https://doi.org/10.1016/j.brainres.2008.07.128

Kruschke, J. K. (2008). Models of categorization. In R. Sun (Ed.), *The Cambridge Handbook of Computational Psychology* (pp. 267--301): Cambridge University Press.

Kuorikoski, J., & Ylikoski, P. (2015). External representations and scientific understanding. *Synthese, 192*(12), 3817-3837. doi:10.1007/s11229-014-0591-2

Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.

Lamb, M., & Chemero, A. (2014). *Structure and application of dynamical models in cognitive science.* Paper presented at the Proceedings of the annual meeting of the cognitive science society.

Lange, M. (2017). *Because without cause : non-causal explanation in science and mathematics*. New York: Oxford University Press.

Latora, V., & Marchiori, M. (2001). Efficient Behavior of Small-World Networks. *Physical Review Letters, 87*(19), 198701. doi:10.1103/PhysRevLett.87.198701

Le Bihan, S. (2016). Enlightening Falsehoods: A modal view of scientific understanding. In S. R. Grimm, C. Baumberger, & S. Ammon (Eds.), *Explaining understanding: New perspectives from epistemology and philosophy of science* (pp. 111-136). London: Routledge.

Levy, A. (2014). What was Hodgkin and Huxley's Achievement? *The British Journal for the Philosophy of Science, 65*(3), 469-492. doi:10.1093/bjps/axs043

Li, M., & Vitányi, P. (2008). *An introduction to Kolmogorov complexity and its applications* (Vol. 3): Springer.

Lombrozo, T. (2006). The structure and function of explanations. *Trends in cognitive sciences, 10*(10), 464-470. doi:https://doi.org/10.1016/j.tics.2006.08.004

Lombrozo, T., & Wilkenfeld, D. (2019). Mechanistic versus Functional Understanding. In S. R. Grimm (Ed.), *Varieties of Understanding* (pp. 209-230): Oxford University Press.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A Network Model of Category Learning. *Psychological Review, 111*(2), 309-332.

Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science, 67*, 1-25.

Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature, 503*(7474), 78-84. doi:10.1038/nature12742

Marr, D. (1982). *Vision : a computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.

McCauley, R. N. (1986). Intertheoretic Relations and the Future of Psychology. *Philosophy of Science, 53*(2), 179-199. Retrieved from http://www.jstor.org/stable/187691

McCauley, R. N. (1996). Explanatory pluralism and the coevolution of theories in science. In R. N. McCauley (Ed.), *The Churchlands and Their Critics* (pp. 17-47): Blackwell Publishers.

Meyer, R. (2018). The Nonmechanistic Option: Defending Dynamical Explanation. *British Journal for the Philosophy of Science*, 0-0.

Miłkowski, M. (2013). *Explaining the Computational Mind*: MIT Press.

Miłkowski, M., & Hohol, M. (2020). Explanations in cognitive science: unification versus pluralism. *Synthese*. doi:10.1007/s11229-020-02777-y

Mišić, B., Betzel, R. F., Griffa, A., de Reus, M. A., He, Y., Zuo, X.-N., . . . Zatorre, R. J. (2018). Network-Based Asymmetry of the Human Auditory System. *Cerebral Cortex, 28*(7), 2655-2664. doi:10.1093/cercor/bhy101

Nagumo, J., Arimoto, S., & Yoshizawa, S. (1962). An Active Pulse Transmission Line Simulating Nerve Axon. *Proceedings of the IRE, 50*(10), 2061-2070. doi:10.1109/JRPROC.1962.288235

Newman, M. (2012). An inferential model of scientific understanding. *International studies in the philosophy of science, 26*(1), 1-26. doi:10.1080/02698595.2012.653118

Newman, M. (2013). Refining the inferential model of scientific understanding. *International studies in the philosophy of science, 27*(2), 173-197. doi:10.1080/02698595.2013.813253

Newman, M. (2015). Theoretical Understanding in Science. *The British Journal for the Philosophy of Science, 68*(2), 571-595. doi:10.1093/bjps/axv041

Operskalski, J. T., & Barbey, A. K. (2017). Cognitive neuroscience of causal reasoning. In *The Oxford handbook of causal reasoning.* (pp. 217-242). New York, NY, US: Oxford University Press.

Osman, M. (2004). An evaluation of dual-process theories of reasoning. *Psychonomic Bulletin & Review, 11*(6), 988-1010. doi:10.3758/BF03196730

Osman, M. (2014). Reasoning research: Where was it going? Where is it now? Where will it be going? In *New approaches in reasoning research.* (pp. 104-123). New York, NY, US: Psychology Press.

Parikh, N., Ruzic, L., Stewart, G. W., Spreng, R. N., & De Brigard, F. (2018). What if? Neural activity underlying semantic and episodic counterfactual thinking. *NeuroImage, 178*, 332-345. doi:https://doi.org/10.1016/j.neuroimage.2018.05.053

Petersen, S. (ms). Explanation as compression.

Piaget, J. (1952). *The origins of intelligence in children* (M. Cook, Trans.). New York: W. W. Norton & Co.

Piccinini, G. (2006). Computational explanation in neuroscience. *Synthese, 153*(3), 343-353.

Piccinini, G. (2015). *Physical computation: a mechanistic account*. Oxford: Oxford University Press.

Piccinini, G., & Craver, C. (2011). Integrating psychology and neuroscience: functional analyses as mechanism sketches. *Synthese, 183*(3), 283-311. doi:10.1007/s11229-011-9898-4

Poldrack, R. A., & Yarkoni, T. (2016). From Brain Maps to Cognitive Ontologies: Informatics and the Search for Mental Structure. *Annual Review of Psychology, 67*(1), 587-612. doi:10.1146/annurev-psych-122414-033729

Potochnik, A. (2017). *Idealization and the aims of science*. Chicago: The University of Chicago Press.

Pouget, A., Deneve, S., & Duhamel, J.-R. (2002). A computational perspective on the neural basis of multisensory spatial representations. *Nature Reviews Neuroscience, 3*(9), 741-747. doi:10.1038/nrn914

Pouget, A., & Sejnowski, T. J. (1997). Spatial Transformations in the Parietal Cortex Using Basis Functions. *Journal of Cognitive Neuroscience, 9*(2), 222-237. doi:10.1162/jocn.1997.9.2.222

Povich, M. (2015). Mechanisms and Model-Based Functional Magnetic Resonance Imaging. *Philosophy of Science, 82*(5), 1035-1046. doi:10.1086/683438

Povich, M. (forthcoming). Mechanistic Explanation in Psychology. In H. Stam & H. L. De Jong (Eds.), *The SAGE Handbook of Theoretical Psychology.* London: Sage.

Price, C. J., & Friston, K. J. (2005). Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology, 22*(3-4), 262-275. doi:10.1080/02643290442000095

Pritchard, D. (2009). Safety-based epistemology: whither now? *Journal of philosophical research, 34*, 33-45.

Rathkopf, C. (2018). Network representation and complex systems. *Synthese, 195*(1), 55-78. doi:10.1007/s11229-015-0726-0

Rice, C. (2015). Moving Beyond Causes: Optimality Models and Scientific Explanation. *Noûs, 49*(3), 589-615. doi:10.1111/nous.12042

Rodieck, R. W. (1965). Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vision Research, 5*(12), 583-601. doi:https://doi.org/10.1016/0042-6989(65)90033-7

Ross, L. N. (2015). Dynamical Models and Explanation in Neuroscience. *Philosophy of Science, 82*(1), 32-54.

Ross, L. N. (2020). Distinguishing topological and causal explanation. *Synthese*. doi:10.1007/s11229-020-02685-1

Rusanen, A.-M., & Lappi, O. (2016). On computational explanations. *Synthese, 193*(12), 3931-3949.

Rysiew, P. (2020). Naturalism in Epistemology. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Sarpeshkar, R. (1998). Analog Versus Digital: Extrapolating from Electronics to Neurobiology. *Neural Computation, 10*(7), 1601-1638. doi:10.1162/089976698300017052

Schank, R. C. (1986). *Explanation patterns : understanding mechanically and creatively*. Hillsdale, N.J.: L. Erlbaum Associates.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences, 3*(3), 417-424. doi:10.1017/S0140525X00005756

Seguin, C., Razi, A., & Zalesky, A. (2019). Inferring neural signalling directionality from undirected structural connectomes. *Nature Communications, 10*(1), 4289. doi:10.1038/s41467-019-12201-w

Serban, M. (2015). The scope and limits of a mechanistic view of computational explanation. *Synthese, 192*(10), 3371-3396.

Seung, H. S., Lee, D. D., Reis, B. Y., & Tank, D. W. (2000). Stability of the Memory of Eye Position in a Recurrent Network of Conductance-Based Model Neurons. *Neuron, 26*(1), 259-271. doi:https://doi.org/10.1016/S0896-6273(00)81155-1

Shadmehr, R., & Wise, S. P. (2005). *The computational neurobiology of reaching and pointing : a foundation for motor learning*. Cambridge: MIT Press.

Shagrir, O. (2006). Why we view the brain as a computer. *Synthese, 153*(3), 393-416. doi:10.1007/s11229-006-9099-8

Shagrir, O. (2010). Marr on Computational-Level Theories. *Philosophy of Science, 77*(4), 477-500. doi:10.1086/656005

Shagrir, O., & Bechtel, W. (2014). Marr's Computational Level and Delineating Phenomena.

Shapiro, L. (2017). Mechanism or Bust? Explanation in Psychology. *The British Journal for the Philosophy of Science, 68*(4), 1037-1059. doi:10.1093/bjps/axv062

Shapiro, L. (2019). A tale of two explanatory styles in cognitive psychology. *Theory & Psychology, 29*(5), 719-735. doi:10.1177/0959354319866921

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience, 36*(1), 337-359. doi:10.1146/annurev-neuro-062111-150509

Silberstein, M., & Chemero, A. (2013). Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences. *Philosophy of Science, 80*(5), 958-970. doi:10.1086/674533

Stephens, R. G., Dunn, J. C., & Hayes, B. K. (2018). Are there two processes in reasoning? The dimensionality of inductive and deductive inferences. *Psychological Review, 125*(2), 218-244. doi:10.1037/rev0000088

Stepp, N., Chemero, A., & Turvey, M. T. (2011). Philosophy for the Rest of Cognitive Science. *Topics in Cognitive Science, 3*(2), 425-437. doi:https://doi.org/10.1111/j.1756-8765.2011.01143.x

Sternberg, S. (1969). Memory scanning: mental processes revealed by reaction-time experiments. *American Scientist, 57*(4), 421-457. Retrieved from http://www.jstor.org/stable/27828738

Strevens, M. (2013). No understanding without explanation. *Studies in history and philosophy of science part A, 44*(3), 510-515. Retrieved from http://www.sciencedirect.com/science/article/pii/S003936811200115X

Sullivan, J. A. (2017). Coordinated pluralism as a means to facilitate integrative taxonomies of cognition. *Philosophical Explorations, 20*(2), 129-145. doi:10.1080/13869795.2017.1312497

Tegmark, M. (2014). *Our mathematical universe: My quest for the ultimate nature of reality*: Vintage.

Thagard, P. (1978). The best explanation: criteria for theory choice. *Journal of philosophy, 75*, 76-92.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences, 12*(3), 435-502.

Thagard, P. (1992). *Conceptual revolutions*. Princeton: Princeton University Press.

Thagard, P. (2012). *The cognitive science of science: Explanation, discovery, and conceptual change*. Cambridge: MIT Press.

Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences, 24*(1), 1-34.

Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind, LIX*(236), 433-460. doi:10.1093/mind/LIX.236.433

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

van Eck, D. (2018). Rethinking the explanatory power of dynamical models in cognitive science. *Philosophical Psychology, 31*(8), 1131-1161. doi:10.1080/09515089.2018.1480755

Van Hoeck, N., Watson, P. D., & Barbey, A. K. (2015). Cognitive neuroscience of human counterfactual reasoning. *Frontiers in Human Neuroscience, 9*(420). doi:10.3389/fnhum.2015.00420

van Rooij, I., & Baggio, G. (2021). Theory Before the Test: How to Build High-Verisimilitude Explanatory Theories in Psychological Science. *Perspectives on Psychological Science, 16*(4), 682-697. doi:10.1177/1745691620970604

Venturelli, A. N. (2016). A Cautionary Contribution to the Philosophy of Explanation in the Cognitive Neurosciences. *Minds and Machines, 26*(3), 259-285. doi:10.1007/s11023-016-9395-0

Verdejo, V. M. (2015). The systematicity challenge to anti-representational dynamicism. *Synthese, 192*(3), 701-722. doi:10.1007/s11229-014-0597-9

Vernazzani, A. (2019). The structure of sensorimotor explanation. *Synthese, 196*(11), 4527-4553. doi:10.1007/s11229-017-1664-9

Verreault-Julien, P. (2017). Non-causal understanding with economic models: the case of general equilibrium. *Journal of Economic Methodology, 24*(3), 297-317. doi:10.1080/1350178X.2017.1335424

Wason, P. C., & Evans, J. S. B. T. (1974). Dual processes in reasoning? *Cognition, 3*(2), 141-154. doi:https://doi.org/10.1016/0010-0277(74)90017-1

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature, 393*(6684), 440-442. doi:10.1038/30918

Weisberg, D. S., Keil, F. C., Goodstein, J., Rawson, E., & Gray, J. R. (2008). The Seductive Allure of Neuroscience Explanations. *Journal of Cognitive Neuroscience, 20*(3), 470-477. doi:10.1162/jocn.2008.20040

Weiskopf, D. A. (2011). Models and mechanisms in psychological explanation. *Synthese, 183*(3), 313-338. doi:10.1007/s11229-011-9958-9

1140 Wilkenfeld, D. A. (2013). Understanding as representation manipulability. *Synthese*, *190*(6), 997-
1141 1016. doi:10.1007/s11229-011-0055-x

1142 Wilkenfeld, D. A. (2019). Understanding as compression. *Philosophical studies*, *176*(10), 2807-2831.
1143 doi:10.1007/s11098-018-1152-1

1144 Wilkenfeld, D. A. (2021). Objectually Understanding Informed Consent. *Analytic Philosophy*, *62*(1),
1145 33-56.

1146 Williams, J. J., & Lombrozo, T. (2010). The Role of Explanation in Discovery and Generalization:
1147 Evidence From Category Learning. *Cognitive Science*, *34*(5), 776-806.
1148 doi:https://doi.org/10.1111/j.1551-6709.2010.01113.x

1149 Williams, J. J., Lombrozo, T., & Rehder, B. (2013). The hazards of explanation: Overgeneralization
1150 in the face of exceptions. *Journal of Experimental Psychology: General*, *142*(4), 1006-1014.
1151 doi:10.1037/a0030996

1152 Woodward, J. (2003). *Making things happen: a theory of causal explanation*. New York: Oxford
1153 University Press.

1154 Woodward, J. (2013). Mechanistic Explanation: Its Scope and Limits. *Proceedings of the Aristotelian*
1155 *Society, Supplementary Volumes*, *87*, 39-65. Retrieved from
1156 http://www.jstor.org/stable/23482050

1157 Zednik, C. (2011). The Nature of Dynamical Explanation. *Philosophy of Science*, *78*(2), 238-263.
1158 doi:10.1086/659221

1159 Zipser, D., & Andersen, R. A. (1988). A back-propagation programmed network that simulates
1160 response properties of a subset of posterior parietal neurons. *Nature*, *331*(6158), 679-684.
1161 doi:10.1038/331679a0

1162