European Journal for Philosophy of Science (2021) 11: 20
https://doi.org/10.1007/s13194-020-00341-6

PAPER IN THE PHILOSOPHY OF THE LIFE SCIENCES

The mechanistic stance

Jonny Lee[1] · Joe Dewhurst[2,3]

Received: 29 August 2019 / Accepted: 8 December 2020 / Published online: 5 January 2021

Abstract
It is generally acknowledged by proponents of 'new mechanism' that mechanistic explanation involves adopting a perspective, but there is less agreement on how we should understand this perspective-taking or what its implications are (if any) for practising science. This paper examines the perspectival nature of mechanistic explanation through the lens of the 'mechanistic stance', which falls somewhere between Dennett's more familiar physical and design stance. We argue this approach implies three distinct and significant ways in which mechanistic explanation can be interpreted as perspectival: 'phenomenon perspectivism', 'pattern perspectivism' and 'hierarchy perspectivism'. We evaluate the strength of the perspective-dependency implied by each of these, and along the way, discuss their significance for wider debates within the new mechanism literature, such as the nature of function attribution and an ontic vs epistemic understanding of explanation.

Keywords Mechanism · Function · Perspectivism · Biology · Cognitive neuroscience

1 Introduction

It is often noted by proponents of 'new mechanism' that mechanistic explanation in some sense involves adopting a perspective. However, little has been said about how to

---

This article belongs to the Topical Collection: Perspectivism in science: metaphysical and epistemological reflections
Guest Editor: Michela Massimi

✉ Jonny Lee
leejonathan.cw@gmail.com

Joe Dewhurst
joseph.e.dewhurst@gmail.com

[1] Temporarily Unaffiliated, 22 Bruntsfield Avenue, Edinburgh EH10 4EW, UK

[2] LMU Munich, Munich, Germany

[3] Munich Centre for Mathematical Philosophy, Ludwigstraße 31, 80539 Munchen, Germany

understand the role of perspective-taking in mechanistic explanation. For instance, many standard definitions of mechanistic explanation emphasise that there are only ever mechanisms 'for the production of some phenomenon', never simply mechanisms 'as such', suggesting the perspectival selection of an explanandum phenomenon might play an important role in understanding mechanistic explanation. More explicit support for perspectivism is found in Craver's (2013) theory of function attribution in mechanistic explanation and Glennan's (2017) emphasis on the alignment between new mechanism and existing perspectivist theories in philosophy of science, along with perspectival approaches to mechanistic explanation in specific cases such as physical computation (Dewhurst 2018a) and neuroscience (Kästner 2018). Despite these endorsements, the relationship between mechanistic explanation and perspectivism remains vague, leaving unanswered questions about the degree of observer-dependency involved in mechanistic explanation and the implications (if any) for scientific practice.

This paper discusses the role of perspective-taking in mechanistic explanation through the lens of the Dennettian stance framework. We present a novel 'mechanistic stance' that lies between Dennett's more familiar physical and design stance. This approach builds on Craver's (2013) brief suggestion of a "mechanistic design stance" (ibid.: 156) analogous to Dennett's design stance, but we go further in arguing that it constitutes a distinct interpretive perspective. We then examine three ways in which mechanistic explanation can be understood to be perspectival, via the mechanistic stance. First, we discuss 'phenomenon perspectivism', which refers to the way in which mechanistic explanation requires scientists to select the phenomenon they want to adopt a stance towards. We suggest this is a relatively weak sense of perspectivism, that accords with more pedestrian interpretations of perspective-taking in science. Second, we discuss 'pattern perspectivism', which refers to the way that the mechanistic stance's emphasis on organisational patterns implies a dual relationship between objective regularities in nature and recognition of those regularities by an observer. If the mechanistic stance approach is correct, then pattern perspectivism offers some insight into the nature of mechanistic explanation. However, the implied dependency on a perspective remains relatively moderate, because the regularities constituting the pattern exist prior to any perspective being taken. Finally, we consider 'hierarchy perspectivism', according to which the hierarchical structure of mechanisms also depends on our explanatory perspective, via the constitutive role played by function attribution. This has significant implications for the role of perspectives in mechanistic explanation, insofar as the structural organisation of mechanisms are themselves dependent on the epistemic considerations of agents. Along the way, we touch on how the mechanistic stance approach may inform debates over the nature of function ascription, whether mechanistic explanation is 'ontic' or 'epistemic', and the relationship between mechanistic explanation and other explanatory stances.

The paper proceeds as follows. In Section 2 we introduce 'new mechanism', outlining the contemporary understanding of mechanistic explanations. In Section 3 we present the mechanistic stance, emphasising its distinctive character in contrast to the physical and design stance. In Section 4 we use the mechanistic stance to discuss three different senses in which mechanistic explanation might be considered perspectival, and the further implications each of these would have for mechanisms, functions, and scientific explanation in general.

2 Mechanistic explanation

Despite some disagreement on the fringes of the literature, recent discussion has settled on a relatively uncontroversial definition of mechanistic explanation provided by Illari & Williamson (2012: 120): "A mechanism for a phenomenon consists of entities and activities organized in such a way that they are responsible for the phenomenon" (for other, similar definitions see Machamer et al. 2000: 3; Glennan 2002: S344; Bechtel and Abrahamsen 2005: 423). The idea is that a common form of explanation in at least the biological and cognitive sciences consists of first identifying a target phenomenon, and then describing a physical structure (a mechanism) whose component parts (entities) interact with one another (via their activities) to either produce or constitute the phenomenon to be explained. The production or constitution of the phenomenon by the mechanism is said to explain the phenomenon, insofar as we can now tell a satisfactory story about how the phenomenon is produced or constituted. Furthermore, mechanistic explanation is usually understood to be hierarchical: just as one mechanism can be used to explain the target phenomenon, so also can the components of that mechanism be themselves treated as smaller sub-mechanisms whose organisation and activities explain some aspect of the larger mechanism.[1] To illustrate, consider the phenomenon of blood being transported around the body. We can identify the cardiovascular system, taken as a whole, as the mechanism responsible for this phenomenon, and describe in some detail how the organised activities of its subparts (the heart, veins, arteries, and so on) contribute to this phenomenon. At a finer level of grain, the heart itself can also be described as a mechanism, one whose organised activities contribute to (and thus explain) the larger operation of the cardiovascular system as a whole. And so on for the component parts of the heart, and their own components, right down to a level at which mechanistic explanation gives way to some more fundamental, non-mechanistic kind of explanation.

One additional and more controversial aspect of mechanistic explanation is the idea that all (or at least some) mechanisms are *functional* (cf. Garson 2013; for discussion of non-functional mechanisms see Illari and Williamson 2012 and Glennan 2017). At first pass, the function of a mechanism or mechanism component is the contribution it makes to producing or constituting the phenomenon we are interested in explaining (although one's specific account of function attribution might complicate this). So, in the above example, the function of the cardiovascular system is to transport blood, and the function of the heart is to pump blood. Function attribution plays an important role in mechanistic explanation insofar as it both fixes the phenomenon to be explained and also serves as a heuristic guide to mechanistic decomposition. Once we know that the function of the heart is to pump blood, we have an approximate idea of the kind of mechanism it is (i.e. a pump), and know to begin looking for components such as valves when we begin to decompose it (in order to better explain the pumping activity). We will return to the topic of mechanistic functions in Sections 3 and 4.

[1] Although see Bechtel (2019) for some interesting discussion of 'hetararchical' (i.e. non-hierarchical) mechanisms. Bechtel's suggested move from hierarchical to hetararchical mechanisms might complicate this picture somewhat, but even hetararchical mechanisms must have some kind of organisational structure, and so much of what we have to say here will still apply.

Several theorists within the new mechanism literature have suggested that mechanistic explanation is perspective dependent. Consider Craver's (2013) perspectival account of mechanistic functions, according to which the role of function attribution in mechanistic explanation is "ineliminably perspectival" (133; cf. Hardcastle 1999 for an earlier perspectival view of functions). Roughly, this is because a mechanism or component's function just is its causal role relative to a certain phenomenon, and any given system will have multiple causal roles; which we choose to focus on is a feature of our explanatory perspective (see next section for further discussion). Or take Glennan (2017), who has also embraced a form of mechanistic perspectivism, explicitly stating that his position

> falls within a family of philosophical positions which have sought at once to acknowledge the essential role of perspectives or stances within the complex sciences while taking seriously the reality or objectivity of entities and kinds seen from these perspectives and stances. (2017: 93)

He goes on to list several notable members of this family, including

> Wimsatt's (1994) account of perspectives and rainforest ontology, Dennett's (Dennett 1991a) account of real patterns, Dupre's (Dupré 1993) promiscuous realism, Giere's (2004) perspectival realism, Mitchell's (2009) integrative pluralism, Kellert et al. (2006) scientific pluralism, and Ladyman and Ross's (2007) rainforest realism. (*ibid*)

Glennan's perspectivism holds that there can be multiple correct ways to describe the world, in line with recent literature on pluralism in scientific modelling (cf. Kellert et al. 2006). Perspectivism is also discussed elsewhere in the literature on mechanistic explanation, beyond the general sense in which a mechanism is describable only relative to a phenomenon of interest. Kästner (2018) describes what she calls a "perspectival aspect" inherent in the mechanistic conception of levels, and analyses this in terms of the "epistemic perspectives" taken by different scientific researchers, and Dewhurst (2018a) has developed a perspectival approach to Piccinini's (2007, 2015) mechanistic account of physical computation (see also Coelho Mollo Forthcoming, for a critical response to this approach).

Despite these claims that mechanistic explanation is perspectival, there is little agreement on how exactly this 'mechanistic perspectivism' should be understood, and little systematic analysis of what it entails. In the remainder of this paper, we go some way to addressing this deficit by developing an approach to evaluating how, and to what extent, mechanistic explanation is perspectival based on the idea that mechanistic explanation involves adopting a distinctive kind of epistemic stance towards a system, one that picks up on underlying regularities, or 'real patterns', in the organisational structure and activities of the system. In this way, we mostly sidestep an evaluation of the many claims by different authors concerning perspectivism on a case-by-case basis, and instead present a general framework for understanding the perspective-dependency of mechanistic explanation. At the same time, we take one of the few sustained discussions of perspectivism in mechanistic explanation as a jumping-off point, Craver's (2013) discussion of mechanistic functions and perspectivism.

3 The mechanistic stance

In this section, we present the 'mechanistic stance' approach. At its core, the mechanistic stance approach suggests that mechanistic explanation involves adopting a particular epistemic perspective that utilises a set of explanatory tools to recognise distinct kinds of organisational pattern. Our focus is not on defending the mechanistic stance against alternative approaches but on outlining the positive story to understand the role of perspective-taking in mechanistic explanation through one promising lens.

The notion of the mechanistic stance builds on a suggestion made by Craver during a discussion of the nature of function ascriptions in mechanistic explanation. To appreciate the context in which a Dennettian stance approach to mechanistic explanation was first proposed, it is worth briefly revisiting Craver's discussion. For Craver, mechanistic functions come in three varieties: "as a way of tersely indicating an etiological explanation, as a way of framing constitutive explanations, and as a way of explaining the item by situating it within higher-level mechanisms" (Craver 2013, 133). Etiological functions are attributed in the context of giving an explanation in terms of the history of a system (cf. Craver 2013, 145–6). To say the (etiological) function of the heart is to pump blood is to appeal to the role it was selected for in the evolutionary history of the organism. Constitutive functions focus on the synchronic causal structure of a system, capturing how, in the here and now, a system might produce a phenomenon (cf. Craver 2013, 149–51). To say the heart (constitutively) functions as a pump is to say that it has the correct kind of physical structure to perform the pumping role. Finally, attributions of contextual function take into account a mechanism's wider environmental context (cf. Craver 2013, 151–4). Contextual functions are important when it comes to situating a component within a broader mechanism that it is part of. To say the heart (contextually) functions to pump *blood*, as opposed to just being a pump simpliciter, is to consider the heart's role within the broader context of the cardiovascular system, which is itself a mechanism for circulating oxygenated blood around the body.

For Craver, all ascriptions of functions to mechanisms and mechanism components are perspectival. For instance, there are many possible ways in which one could describe the physical structure of a system, only some of which will be relevant to our current explanation, so the appropriate constitutive function to attribute to a system is dependent on a particular explanatory context. Meanwhile, contextual functions appeal to features beyond the boundaries of a system, and so must involve judgements about *which* boundaries to focus on, as a system is typically embedded in a nested hierarchy of environments, each of which might be relevant in a different explanatory context.[2] The underlying point, for Craver, is that functions are not part of the causal structure of the world, and so we should not try to give a naturalistic account of them.

[2] See Shagrir and Bechtel (2017) for a recent discussion on the importance of environmental context in fixing the target phenomenon of a mechanistic explanation.

Rather, function ascriptions result from adopting a perspective on the phenomenon a mechanism is 'supposed' to constitute or produce. For Craver, the central goal of function attribution "is to make the busy and buzzing confusion of complex systems intelligible and, in some cases, usable" (Craver 2013: 140). All this leads him to suggest that attributions of function within mechanistic explanation involve adopting

> […] a kind of mechanistic design stance, liberated from Dennettian associations with adaptationism and optimality: a stance that there is a behavior that the mechanism as a whole exhibits (that it is the mechanism *of* a behavior) and that the components of the mechanism are organized and interact such that they exhibit its overall behavior. (Craver 2013: 156)

Craver thus offers the beginnings of an answer to whether and in what sense mechanistic explanation is perspectival. The details remain vague, however, and the extent to which mechanistic explanation is really perspectival in any significant sense remains unclear. In what follows, we will further develop what we call the 'mechanistic stance' (dropping 'design' from Craver's original proposal for reasons that will soon become clear). We will then evaluate in what sense this implies any significant perspectivism in mechanistic explanation.

To understand and develop Craver's suggestion that mechanistic explanation involves adopting a sort of Dennettian 'stance', we first need to understand Dennett's broader stance framework. According to Dennett, different kinds of explanation involve adopting different stances from which to interpret a system. These stances are defined by the approach they take towards explaining the behaviour of a system, and the different set of explanatory tools they adopt as a result. The tools best suited to explaining one kind of phenomenon may be ill-suited to explaining another, and so we adopt different explanatory stances to explain different kinds of phenomenon. Explaining human behaviour using the tools of particle physics would be extremely complicated, and also unnecessary given access to different kinds of explanatory tools that are better suited for this purpose. Each stance, we will later suggest, can be seen as a different kind of explanatory perspective that is more or less appropriate for different contexts.

Important to the stance framework is the notion that different stances operate by tracking different kinds of patterns in nature. We take patterns to be discernible regularities that are in-principle detectable by an observer (see Dennett 1991a; cf. Ross 2000; Andersen 2017; Suñé and Martinez 2019; Millhouse forthcoming). They can be defined formally in terms of string compression in algorithmic information theory: a pattern is present in some dataset if there is a compressed way of describing that dataset, or equivalently, a shorter program that outputs the same data in information theoretic terms (see Li and Vitányi 2008 for an overview, and for the original definitions, Solomonoff 1964a, 1964b; Kolmogorov 1965; Chaitin 1966). Patterns allow us to make sense of how many theoretical terms in scientific and ordinary explanation predict and explain. For example, propositional attitudes, such as beliefs and desires, refer not to discrete, concrete kinds in the brain, but rather to behavioural patterns of a whole organism, composed of myriad physical constituents. According to Dennett, such patterns are abstract but real insofar as they consist of objective regularities in the information theoretic sense. In his original presentation of what he terms 'real patterns', Dennett describes this initially in terms of image compression, where the discovery of a pattern in an image allows for a more efficient encoding than the raw bit-map, by making use of regularities to produce a shorter description of the image. Similarly, by describing and explaining behaviour in terms of patterns like beliefs and desires we avoid having to keep track of the microstructural details of our conspecifics, and by discovering regularities in nature we can produce scientific theories and models that leave out many fine-grained physical details. For example, provided with a coarse-grained description of someone who "believes it is going to rain and desires to stay dry", plus some background information about what kind of equipment is available to them, we can predict they will put on a raincoat or carry an umbrella. The same prediction could, in principle, be made with a fine-grained description of their neurological and physiological states, but this description would be far more complex and the process of generating the prediction would be far more arduous. In this way, taking this explanatory stance on a complex system like a human being can allow us to make compressed descriptions and predictions of their behaviour.

We should emphasise at this point that the appropriateness of a stance, and the kinds of details it leaves out, will be relative to the phenomenon that we want to explain.[3] So, while attributions of mental states might be most appropriate for explaining coarse-grained human behaviour, they are less well-suited to explaining the tendency of a human to fall to the ground when dropped from a height, and an explanation in terms of gravitational force is probably more appropriate in such contexts. This is not to say the patterns themselves are observer-relative (algorithmic information theory gives a clear answer to the question of whether there is a compressible pattern in any dataset), but rather that the patterns that are most relevant to an explanation will depend on the phenomenon we want to explain.

In articulating the stance framework, Dennett presents three stances we regularly adopt in both our everyday and scientific explanations: the physical stance, the design stance, and the intentional stance.[4] When one adopts the physical stance, one decides to interpret the behaviour of a system by determining the physical properties of a system's constituents, and how those constituents interact with each other and the physical constituents of other systems, via physical laws. This stance typically offers fine-grained explanations and predictions about a system's behaviour, which can be computationally costly to keep track of but afford a high degree of precision. In cases where such precision is necessary or useful the physical stance might offer the best explanatory perspective, but in many other cases we can adopt stances that offer less fine-grained understandings but which nonetheless explain and predict the behaviour of a system to an adequate degree. Of course, some simplified physical models leave out many fine-grained details, so it is not straightforwardly true that the physical stance is *necessarily* more computationally costly than other stances. However, as a general rule, where another stance offers sufficient understanding of some phenomenon, for a

[3] We thank an anonymous reviewer for pressing us on this point.

[4] Dennett and others have occasionally mentioned other useful stances, so we do not take this list to be exclusive. Robbins and Jack (2006), for instance, discuss a 'phenomenal stance', and Vermass et al. (2011) draw a further distinction between an 'intentional design stance' and a 'teleological design stance'.

particular explanatory or predictive purpose, the physical stance will offer a comparatively more fine-grained (and hence more computationally costly) model.

One can often gain sufficient understanding of a system by adopting the design stance, and asking "what is this system *for?*". The design stance ignores the internal machinery of a system, instead treating it as a system that has been designed by an agent (such as a person), or quasi-agent (such as evolution by natural selection). The behaviour of the system is then predicted on the assumption that the artefact and its parts were designed to optimally perform what they were designed to do, i.e., their 'function'. Importantly, the design stance works for biological systems as well as artificial ones; we can predict the behaviour of a lion, and even lion parts, by treating them as 'designed' to survive in a certain environment and by further assuming that the parts of the system are suited to this end. For instance, we might predict that at midday on the savannah, a lion will seek shade, because it is designed to keep itself at a regular temperature.

Sometimes we can interpret the behaviour of a system without even considering its design or function, and instead simply considering its intentions. When one adopts the intentional stance towards a system, one treats the system as a rational, goal-seeking agent in possession of beliefs, desires, and other mental states. One can then predict the behaviour of the system by working out what the agent should do, given the conjunction of its beliefs and desires with principles of rationality. For example, upon seeing you open an empty fridge on a hot summer's day, I might predict that you will journey to the supermarket, attributing a belief about the fridge's empty contents, a desire for a refreshing drink, an additional belief about where to acquire groceries. Likewise, when we observe two different people reach towards a fridge door in two different environments we can quickly generate a host of similar predictions about what happens next, despite great variation in the physical constituents (all the way down to the microphysical details) of those persons and their environments, and subsequent variation in the particular details of their actions. Tracking intentional stance patterns is far more efficient for explaining behaviour of this kind, even if it leaves out many microphysical details.

Each stance allows us to focus on patterns that might not be apparent from other perspectives. To be clear, the patterns themselves are always 'out there' in the data, and in this sense objective, but to adopt a stance is to focus our attention on a particular *kind* of pattern, one that is most suitable for the phenomenon we want to explain. Our choice of stance (and thus, which patterns to focus on) also depends on the cognitive and material tools available to us—whether we have the right 'pattern detecting device'. For example, prior to the development of modern chemistry we were not able to identify molecular patterns, and prior to the cultural evolution of mindreading we might not have been able to identify the patterns focused on by the intentional stance.[5] We will return to this idea when we discuss 'pattern perspectivism' in Section 4.

Following Craver's proposal, we think there is a distinct 'mechanistic' stance that focuses on patterns that are not emphasised by the physical, design, or intentional stances. While Craver initially introduced the general idea by referring to the

[5] Indeed, if one takes Zawidzki's (2013) idea of 'mindshaping' seriously, these patterns might not even have existed yet, analogously to how synthetic chemicals display patterns that would not have existed unless we created them.

"mechanistic design stance" (2013: 156), suggesting a mere variation on the design stance, we believe it is more accurate to think of it as an independent kind of explanatory stance because it lacks distinctive features of the design stance and possesses unique features of its own. When we adopt the mechanistic stance, we view the behaviour of a system as the product of organised, interacting parts. In other words, we treat a system as if its behaviour is the result of organised activities resulting from parts and their interactions. Explanations from the mechanistic stance thus capture the operations of spatially and temporally organised components within a system that jointly constitute or produce the explanandum phenomenon (e.g. Craver 2007; Bechtel 2008). In turn, according to this approach, a system is a mechanism just in case describing its behaviour using the tools of mechanistic explanation gives us significant predictive and explanatory traction. Any physical system that is sufficiently organised, such that some of its microstructural parts can be grouped into macrostructural component types, is a potential target for the mechanistic stance. In other words, to be a legitimate target for the mechanistic stance, a system must be organised such that it is more than a mere aggregate, i.e. there must be some distinctive form of behaviour that emerges from the structure of its organised parts. Biological and artefactual systems will be obvious candidates for this kind of explanation, but so might certain non-biological natural systems, such as those studied by astrophysics (for further discussion, see Illari and Williamson 2012). An example of a system that would not be a suitable target for the mechanistic stance would be something that we typically treat as a mere aggregate, such as a homogeneous fluid, that would be better studied from the physical stance (i.e. using the tools of fluid dynamics).

We can draw out the distinguishing features of the mechanistic stance by comparing it to the physical and design stances. As we saw above, adopting the physical stance means focusing on the physical structure of a system and extrapolating, on the basis of physical laws, the future evolution of that system. By contrast, the mechanistic stance abstracts away from many low-level details about the chemical and physical structure of a system's constituents. Specifically, it abstracts away from microphysical details by 'clumping' microphysical parts into structured components. Meanwhile, the design stance requires no knowledge of the physical constitution underlying a system's behaviour. Predictions are based on assumptions of what the system was designed to do, or what its purpose is. By contrast, the mechanistic stance is concerned with the causal properties of a system's organised parts and how they interrelate. Moreover, the design stance requires invoking a designer (e.g., a human or divine actor), or a quasi-designer (e.g., natural selection) of the system in question. The mechanistic stance does not, and we take this to be what Craver means when he discusses the mechanistic design stance being "liberated from Dennettian associations with adaptationism and optimality" (2013: 156).

The mechanistic stance does, of course, involve ascribing functions to the system. However, such ascriptions are not bound by the assumption that the system in question was designed. Rather, they are bound by the assumption that the system and its parts are causally responsible for a phenomenon of interest (cf. Craver 2001 on causal role functions). The role of function ascription is then principally to pick out the causal contribution of a mechanism or component to the explanandum, among the many effects it produces. The sense of function inherent to the mechanistic stance, therefore, does not connote purpose or end-directedness, in any strong sense. Hence, the mechanistic stance does not require talk of 'proper functions' (derived from their design or otherwise). We acknowledge, of course, the many attempts by philosophers to provide accounts of proper functions for mechanisms (see e.g. Garson 2013; Maley and Piccinini 2017). Though we lack the space here to provide a full treatment of the literature, our present claim is that proper functions are not an essential element of the interpretive tools of the mechanistic stance, in contrast to the design stance. This puts our characterisation of the mechanistic stance in alignment with existing outlooks on mechanistic explanation which downplay the necessity of proper functions (e.g., Craver 2013; Dewhurst 2018a). One notable benefit of divorcing interpretation of a system from assumptions about its design is that it facilitates adopting the mechanistic stance towards natural non-biological systems, as Glennan (2017) and Illari and Williamson (2012) suggest.

We can further understand this difference between the mechanistic and other stances by returning to the notion of a pattern. In alignment with the broader stance framework, the mechanistic stance approach suggests mechanistic descriptions capture discernible regularities that allow for a compressed description of a system. Whilst the patterns recognised by the physical stance are those of the underlying physics, the mechanistic stance abstracts away from the physical microstructure of a system to the real patterns of its compositional structure, but (unlike the design stance) does not require that we posit any kind of designer or design process. Here we follow Illari and Williamson (2012) in claiming that while mechanisms are, at least in part, functionally individuated—they are individuated by their causal role relative to an explanandum—and function attribution plays an important role in mechanistic explanation, mechanisms are not functional in any strongly normative sense. Moreover, their identity is also fixed by specific details of their causal structure, insofar as two structurally distinct mechanisms that in some sense 'do the same thing' should nonetheless be considered to be separate mechanism types (we discuss this point in more detail below). Another way to put this is to say that while an explanandum phenomenon must be specified before we can identify a mechanism for that phenomenon, this does not imply that said mechanism was *designed* for producing that phenomenon, or that doing so is its *proper* function.

To illustrate the differences between the stances, take the example of a thermostat. From the physical stance, we can explain and predict the behaviour of a thermostat in terms of relevant physical laws and the chemical properties of its parts, e.g., we understand bimetallic strips in terms of the coefficient of thermal expansion of copper and steel and equations that describe the curvature of a bimetallic beam. This contrasts with understanding the thermostat from the design stance. Here, we explain and predict the behaviour of the system by understanding the thermostat's function, i.e., to sense temperature, with no consideration of its internal components and operations (though we could apply the same stance to parts of the system). Both approaches diverge from understanding the thermostat from the mechanistic stance. Here, we explain and predict the behaviour of the system in terms of interacting component parts and their operations. For example, in a standard two-wire thermostat we would appeal to the parts such as the control lever; bimetallic strip wound into a coil; flexible wire; fixed contact screw; and magnet. We would also describe operations performed by these parts—e.g., the movement of the coil clockwise as temperature decreases—and how they interact—e.g., the moving coil turns the connected lever which turns on a gas valve. Which of these stances is best suited to explaining the behaviour of a thermostat depends on the efficiency of the explanations it offers, or in Dennettian terms, the compressibility of the real patterns that it identifies and tracks. However, the relative value of efficiency also depends on one's explanatory goals. If your goal is simply to control room temperature, then the ability of the physical stance to describe the material properties of a thermostat is irrelevant, and the design stance is to be preferred, but if you are interested in the particular durability of your thermostat at extreme temperatures then the physical stance might provide information that is simply not available from the design stance. Different explanatory virtues like scope are also relevant here: a design stance explanation might cover thermostats in general, whereas mechanistic stance and physical stance explanations will differ depending on the particular constitution and organisation of different types of thermostat or individual thermostats. Hence, it is not *just* efficiency that matters—different stances provide qualitatively different kinds of information, and are thus more or less suitable for different explanandum phenomena.

The design stance and the mechanistic stance also individuate systems differently, and at different levels of grain. From the perspective of the design stance, we might describe a wide range of different systems as "corkscrews", without worrying about the specifics of how it is they remove corks from bottles. From the perspective of the mechanistic stance, however, these differences often matter, as componential and organisational details are also relevant to the identity of a mechanism. So, a traditional lever corkscrew and a modern electronic corkscrew would constitute different kinds of mechanism, even if from the perspective of the design stance they are both of the same kind, i.e. "corkscrews". A practical scientific case where the distinction between these two stances is relevant is the recent debate about the individuation of neural computations. Chirimuuta (2014) has argued that neural computations are individuated functionally, without any reference to fine-grained mechanistic details, while mechanistic accounts of computation would suggest otherwise (cf. Kaplan 2011; Kaplan and Craver 2011; Piccinini 2015; Craver and Kaplan 2020; Kaplan 2017; Dewhurst 2018b). Rather than attempting to find a definitive answer to this question, we could see it as a clash between two different stances or perspectives: Chirimuuta's *design stance* perspective is interested in individuating neural computations according to what they are *for* (in evolutionary or developmental terms), while the *mechanistic stance* perspective is instead interested in individuating them according to what they do and how they do it, i.e. their mechanistic composition and structure. These stances need not compete with one another, but instead offer distinct, complementary approaches to studying the same underlying system (in this case, the brain). In both cases, the function of 'computing' is attributed to the neural system, but from the mechanistic stance it plays a heuristic role in guiding the discovery of several distinct kinds of mechanism, whereas from the design stance it is definitional of the components being discovered, whose underlying physical structure is therefore less relevant.

In closing this presentation of the mechanistic stance, we want to draw attention to an underlying theme of the stance framework: different interpretive perspectives, such as the physical, design, and intentional stance, form distinct cognitive technologies that allow us to flexibly interact with systems with varying degrees of detail and abstraction, depending on our goals. The stances are marked by their different developmental histories. The intentional stance is mostly or wholly a product of evolutionary and ontogenetic development, whereas the physical stance has resulted, at least in large part, from cultural and technological progress. Building on this theme, we propose that the mechanistic stance represents another major epistemic innovation for humankind. Unlike the design stance, which understands systems 'teleologically'—in terms of their purpose given their design—the mechanistic stance allows for more flexible, nuanced interactions with systems by uncovering their parts, operations and the relations between those parts and operations. At the same time, it usefully abstracts away much of the fine-grained complexity of the physical stance. The mechanistic stance hits the sweet spot, for many practical purposes, between the coarse-grained design stance and the fine-grained physical stance. Of course, there are many cases where a mechanistic explanation is inappropriate. Our claim is simply that in some cases (paradigmatically, hierarchically composed and systematically organised systems) the mechanistic stance allows for greater explanatory purchase than either the design or physical stances.

4 Three kinds of mechanism perspectivism

'Perspectivism', as intended here, principally concerns the nature of scientific knowledge. Though perspectivism and cognate terms have been used in importantly different ways, the underlying notion shared by all these is that scientific knowledge is essentially perspective-dependent, and in turn, scientific knowledge cannot be understood without recourse to the perspectives of agents and communities of agents. It is not our goal here to taxonomize the many different ways in which this broad sentiment can be interpreted, however, it will prove useful to note the different degrees to which scientific knowledge might be construed as perspectival. For example, we can imagine a spectrum at which one end lies an extreme kind of radical relativism. According to this relativism, the sole ingredient of scientific knowledge is the beliefs and assertions of the scientific community, coupled with widespread societal deference to such beliefs and assertions. At the other end of the spectrum lies an austere realism which holds that human perspectives play no role in scientific knowledge which is instead grounded solely in objective, agent-independent states-of-affairs. Somewhere in the middle of this spectrum lies 'perspectival realism' and related positions. Though a broad church itself, perspectival realists agree that (1) there are objective, perspective-independent states of affairs which science can have knowledge of, and (2) scientific knowledge of these states of affairs is shaped and constrained by perspectives. As this spectrum indicates, different theories can invoke different strengths of perspective-dependence, ranging from total relativism, through versions of perspectival realism according to which the objects in our scientific theories depend somewhat on our perspectives and finally traditional non-perspectival realism, according to which there is a single objective world and a single best way to model it.

One obvious sense in which mechanistic explanation and the mechanistic stance may be interpreted as perspectival is via their compatibility with previous expressions of scientific perspectivism, such as those due to Giere (2006) and Massimi (Massimi 2012, Massimi 2018a, Massimi 2018b; see also Massimi and McCoy's 2019 edited volume *Understanding Perspectivism*). These concern scientific knowledge in general, and not mechanistic explanation in particular. Perspectival realism, for instance, considers the limitations of all scientific instruments and models, from fundamental physics to cognitive science. At the same time, all proponents of mechanistic explanation highlight the important role played by the scientific community in selecting a phenomenon of interest, and the mechanistic stance approach clarifies how this selection process plays an important role in shaping the kind of scientific explanations that mechanisms feature in. Mechanistic explanation seems compatible with scientific perspectivism in this generic sense, but here we are interested in whether mechanistic explanation is perspectival in any more specific sense.

If it offered nothing more than logical consistency with more generic forms of scientific perspectivism, such as perspectival realism, the mechanistic stance approach would not be terribly exciting when it comes to evaluating the role of perspective-taking in mechanistic explanation. In what follows, we will consider three more significant senses in which mechanistic explanation could be considered to be perspectival, using the idea of the mechanistic stance as a starting point and unifying thread: phenomenon perspectivism, pattern perspectivism, and hierarchy perspectivism. Phenomenon perspectivism reflects the need for scientists to select the phenomenon which they adopt the mechanistic stance towards. Pattern perspectivism captures the role and limitations of an agent in choosing which patterns to focus on when providing a mechanistic explanation. We will suggest that neither of these are especially radical forms of perspectivism, and should both be palatable to any proponent of mechanistic explanation. A third kind of mechanistic perspectivism would have stronger implications for the hierarchical structure of mechanisms, via the constitutive role played by function attribution in mechanistic explanation. These three types of perspectivism are not exclusive nor necessarily exhaustive, but rather offer one typology that is useful for understanding the varieties of mechanistic perspectivism. Articulating these three senses in which one might understand perspectivism in mechanistic explanation is more constructive, we think, than providing a single, sweeping answer as to whether mechanistic explanation is perspectival, which mistakenly suggests the answer is unambiguous and one dimensional. To appreciate this, let's examine these three 'mechanistic perspectivisms' further.

Moving beyond the compatibility between the mechanistic stance and broader perspectivist theories within philosophy of science, what we refer to here as 'phenomenon perspectivism' reflects specific features of the mechanistic stance itself. The mechanistic stance involves adopting an epistemic stance towards a particular system, meaning scientists select the phenomenon they want to adopt the mechanistic stance towards. This suggests mechanistic explanation is dependent upon the explanatory choices made by scientists, i.e., it suggests a degree of perspective-dependency. Without the perspective-taking of scientists, there would be no mechanisms, the idea goes, for a system is only a mechanism to the extent it is interpretable via the mechanistic stance (much like a Dennettian claims that a system is only an intentional agent to the extent it is interpretable via the intentional stance). Amongst other things, this phenomenon-selection is vital for carving up the boundaries of a system which, as noted above, is embedded in a nested hierarchy of environments, each of which may bear on an explanation, depending on the explanandum. It is also important for setting the practical limit on decomposition. Complex systems may be decomposed further and further until we reach an 'absolute bottom', the level at which no further mechanistic decomposition can be given (and one must adopt the physical stance). However, for the purposes of giving an adequate explanation of a phenomenon, such complete decomposition is rarely necessary. For instance, a satisfactory explanation of the cardiovascular system does not typically require one to decompose down to the level of the cell nucleus, even though the latter is a good target for mechanistic analysis in its own right. In this way, the role of phenomenon selection is vital for understanding mechanistic explanation.

Though important for fully understanding the nature of mechanistic explanation via the mechanistic stance, the perspective-dependency implied by phenomenon perspectivism is mild. It is often observed across diverse opinions within the new mechanism literature that selecting a phenomenon for explanation is vital for the process of mechanistic explanation. We suggest the phenomenon perspectivism inherent in the mechanistic stance is really just a stance-friendly way of understanding the common refrain that mechanisms are only ever mechanisms *for* a phenomenon. This is a relatively weak sense of perspectivism, principally because all positions in the new mechanism literature acknowledge that human epistemic interests determine the target phenomenon, and that this, in turn, affects how *we* carve up the relevant spatiotemporal regions of the world. However, those spatiotemporal regions may otherwise remain perfectly objective. In this sense, even an austere realist would likely accept this role for perspective-taking in mechanistic explanation.

It should be noted that Craver's (2013) perspectivism about mechanistic functions arguably qualifies as a kind of phenomenon perspectivism, given its attention to the role of explanandum selection in fixing a mechanism or component's function (though see the discussion on 'hierarchy perspectivism' below). We believe that Craver's emphasis on perspective-taking also serves a useful rhetorical purpose, namely, contrasting his theory of function with existing, naturalistic theories of 'proper function'. For Craver, function attributions capture the causal roles of mechanisms and mechanism components relative to a phenomenon of interest; function attribution plays an instrumental role in specifying the relevant causal aspects of a system. We take this absence of proper function for explanation to be a hallmark of the mechanistic stance, in contrast with the design stance. Nevertheless, as others have noted, the perspectivism inherent in this way of looking at functions is rather tepid to the extent that *once* one has selected a phenomenon of interest, the function of a mechanism or component is perfectly objective (e.g. Piccinini 2015: 143). Relative to the distribution of oxygen via the circulatory system, for example, the function of the heart is to pump blood, and for most scientific purposes this is the phenomenon we are interested in. Even though phenomenon perspectivism denies the existence of proper functions, in contrast with etiological or objective goal accounts, it does not deny that there is a fact of the matter about the function of some target system, relative to a phenomenon of interest. This kind of perspectivism is therefore not especially radical.

There is another, more interesting form of perspectivism suggested by a Dennettian stance approach to mechanistic explanation, which we will call 'pattern perspectivism'. The fundamental idea is that pattern-recognition itself implies a form of perspective-dependency. There are two components to pattern perspectivism:

First, during mechanistic explanation, an observer must choose to adopt the mechanistic stance from the range of possible stances available to them. The orderly arrangements in nature underlying a pattern exist objectively, in standard formulations, but an agent must deploy the right tools, usually from a range of options, to recognise those objective patterns. By adopting a stance, we also commit to certain norms of explanation, for example the design stance commits us to talk of proper functions whilst the mechanistic stance commits us to describing the causal structure of a system (cf. Kästner and Haueis 2019). The patterns we focus on when giving an explanation depend on these norms, and in this way, patterns *qua* feature of an explanation arise from the interplay between objective regularities and pattern detectors. As Dennett writes, "in the root case a pattern is 'by definition' a candidate for pattern recognition" (Dennett 1991a: 32; see also Haugeland 1998), not because the underlying regularity is observer-relative, but rather that for it to do any explanatory work it must be recognised by an observer.

Second, any pattern recognition system has material/cognitive constraints on which patterns it can detect, and so which patterns are in-principle recognisable by a system depend on its particular capabilities. Borrowing from Dennett, we can imagine a race of super-intelligent "Laplacian martians" who predict human behaviour in terms of micro-physical properties and physical laws, without ever treating them as agents. In failing to adopt the intentional stance—let's suppose due to hard wired biological constraints—such aliens would fail to detect one set of patterns that support predictions and generalisations. Similarly, we can suppose that the mechanistic stance is supported by a suite of evolved and culturally developed cognitive technologies. To this extent, mechanistic explanation depends on the particularities of human cognition that allow us to pick up on the associated set of patterns.

Like phenomenon perspectivism, pattern perspectivism captures an important role for the epistemic standpoint of agents in scientific explanation, by partaking in a two-way relationship with a pattern that would not be possible without adopting an explanatory perspective. By itself, a pattern does not explain or do anything, it simply offers a compressed description of a dataset that could be used by an epistemic agent when constructing an explanation. While phenomenon perspectivism reflects the need for a perspective in choosing what to adopt a stance towards, pattern perspectivism reflects the fact that the selection of a pattern to focus on is always made by a pattern detector with particular constraints.[6] In this sense, it says something stronger and more interesting about the role of agent perspectives in scientific knowledge, going beyond the mere choice of which phenomenon to explain, to the selection of which patterns to use to explain it. For example, given the target phenomenon of temperature regulation by a thermostat, we might choose to focus either on the interaction of the organised parts of the system (adopting the mechanistic stance) or on the particular thermal properties of the materials used in its design (adopting the physical stance). Each stance picks up on different patterns, and will be suitable for explaining different aspects of the target phenomenon. Furthermore, if we lack any knowledge of electrical engineering then the first stance might be unavailable, while if we lack the tools or expertise to measure thermal properties then the second stance would be unavailable. Nevertheless, in its standard formulation, such perspective-dependent choices are still dependent on the existence of objective, perspective-independent regularities (i.e. real patterns). As such, the perspective-dependency implied by pattern perspectivism says more about the role of perspectives in characterising mechanistic structures than about the nature of those structures themselves.

[6] Potochnik & de Oliveira (Potochnik and Oliveria 2020) have recently made a similar suggestion in the context of cognitive scientific explanation, which they argue consists of a patchwork of different kinds of explanatory patterns, where "Which pattern is explanatory depends on both the cognitive phenomenon under investigation and the research interests occasioning the explanation" (ibid: 1).

There is a final form of perspectivism implied by the mechanistic stance, which we will call 'hierarchy perspectivism', that concerns the constitutive role played by function attribution in mechanistic explanation. Recall that function attribution is understood by Craver to be essentially perspectival (Craver 2013), insofar as the function of a system depends on the explanandum in question. Furthermore, the attribution of input/output functions, in Craver's (2013) sense, determines which causal relations in a target system are taken to be relevant for a given mechanistic explanation. Any given component of a target system will engage in myriad causal interactions with other components, sub-components, and incidental non-components, but only a tiny fraction of these interactions will be of interest from the perspective of a mechanistic explanation, i.e. those identified by the attribution of an input/output function to a component. We noted above that this itself does not imply an especially strong form of perspectivism. Importantly, however, these privileged causal relations subsequently determine the hierarchical structure of a mechanism, via the prohibition against inter-level causation articulated by Craver and Bechtel (2007). According to Craver & Bechtel, if two entities causally interact with one another, they must be at the same level of the (local) mechanistic hierarchy, in order to avoid issues to do with apparent cases of inter-level causation. Note that this is not meant to imply the existence of a *universal* hierarchy, of the kind proposed by Putnam and Oppenheim (1958), but merely a local hierarchy confined to the internal structure of a mechanism. The prohibition against inter-level causation is also not universally endorsed by mechanists (cf. Kaiser and Krickel 2017). However, we note it here to demonstrate the consequences of combining Craver's brand of perspectivism about function with Craver & Bechtel's understanding of causation within a mechanistic hierarchy: if function attribution plays a role in determining the relevant causal interactions for a given mechanistic explanation, and inter-level causation is prohibited, then function attribution also plays an important role in determining the structure of the mechanistic hierarchy. Thus, it seems Craver's function perspectivism implies a hierarchy perspectivism, with potentially strong implications for the observer-independence of mechanisms themselves, and not merely our models or descriptions of them.

A practical example of these consequences is the possibility of *cross-cutting* functional hierarchies, where two different decompositions of the same mechanistic system result in competing hierarchies, with token-identical entities appearing at different levels of each hierarchy. Dewhurst and Isaac (2020) discuss this issue in the context of the phenomenon of ephaptic coupling, compared with the more familiar phenomenon of synaptic transmission. Ephaptic coupling describes the effect on the firing rate of a single neuron that the electromagnetic field generated by another group of neurons can have (see e.g. Anastassiou et al. 2011; Chiang et al. 2019). In this case, the electromagnetic field and the affected neuron would have to be at the same level of the mechanistic hierarchy, as they are causally interacting, while the electromagnetic activity of the other neurons would be at a lower level, being constitutive of the field. However, it is also possible that the first neuron might be synaptically coupled with one of the other neurons, once again situated at the same level of the hierarchy due to this causal interaction. This would place the activity of the neurons generating the field both at the same level and at a lower level from the first neuron, leading to a conflict between the mechanistic hierarchies responsible for ephaptic coupling and synaptic transmission. Dewhurst and Isaac (2020) present this is a practical example of cross-cutting mechanistic hierarchies, which they argue poses a challenge to existing theories of realism about mechanistic hierarchies. The mechanistic stance approach, on the other hand, faces no such challenge, as the realism it endorses is a realism about *patterns*, which might uncontroversially cross-cut one another, rather than a realism about mechanistic hierarchies.

This hierarchy perspectivism also has implications for the debate between 'ontic' and 'epistemic' conceptions of explanation. According to the ontic conception, it is mechanisms in the world themselves that should be considered explanatory of some phenomenon. On this side of the debate is Craver, who writes "[o]bjective explanations are not texts; they are full-bodied things. They are facts, not representations" (Craver 2007, 27). By contrast, according to the epistemic conception it is the scientific representations of causal structures (our models and theories) which do the explaining. As Bechtel puts it, "[e]xplanation is fundamentally an epistemic activity performed by scientists" (Bechtel 2008, 18). A complete exploration of the ontic/epistemic debate is not possible here, but we wish to note a few points pertaining to the present discussion. First, we follow Illari (2013) in diagnosing the underlying disagreement in the contemporary debate as one fundamentally concerning norms of explanation. In a nutshell, the disagreement lies in whether mechanistic explanation does and should aim to uncover objective causal facts or provide pragmatic understanding of a phenomenon. Second, we agree with Kästner and Haueis (2019) that both sets of norms play an important role in mechanistic explanation: scientists must recognise epistemic norms of intelligibility but the ontic adequacy of these norms, in turn, depends on the objective causal structure of the entities in question. Third, reinforcing Kästner & Haueis's claim, we think a Dennettian stance approach poses a problem for an austere ontic approach, insofar as it suggests epistemic considerations are baked into our conception of mechanisms. While there are objective regularities 'out there' in the world, any individuation of mechanisms depends on explanandum selection, pattern recognition and which causal relations are privileged when determining the hierarchical structure of a mechanism.

An austere ontic approach suggests a robust notion of what a mechanism *really* is, independent of any epistemic consideration, such that it could serve 'ontically' as a scientific explanation. Craver's commitment to the ontic account of mechanistic explanation is therefore hard to square with his function perspectivism, insofar as the latter suggests some acceptance of the idea that the causal structure of a mechanism might depend somewhat on our explanatory perspective, in which case it becomes unclear which particular mechanistic structures are meant to serve as the grounds for ontic explanations. However, the mechanistic stance approach does potentially offer an alternative, less austere version of the ontic account, as it comes with a Dennettian reconception of what it means for something to be 'real'. Recall that within the stance framework, to be 'real' at a higher level of abstraction is just to be a real pattern (cf. Andersen's 2017 "pattern ontology"), i.e. to be a recognisable structure that provides more information at some level of grain. Function attribution could be perspectival, in the sense Craver suggests, and play a constitutive role in determining the hierarchical structure of a mechanism, without ruling out the mechanism being 'real' in the Dennettian sense. In this way, the mechanistic stance can accommodate a kind of perspectival realism about mechanistic structures, understood as real patterns, retaining some of the flavour of the ontic account while accommodating the relevant epistemic factors.

In this section, we presented three types of perspectivism that are suggested by a mechanistic stance approach to mechanistic explanation. Phenomenon perspectivism reflects the emphasis that is placed in the Dennettian stance framework on agents choosing to adopt a stance towards a particular system. We claimed this amounts merely to a stance-appropriate way of framing the familiar and uncontroversial point that selecting a phenomenon of interest is an important step in constraining the spatiotemporal regions and processes that are relevant to an explanation. Having said this, one potentially controversial area where the mechanistic stance conflicts with many commonly held theories is that of function ascription. In keeping with Craver's 'perspectivism about function', the mechanistic stance suggests functions are no more than causal roles of mechanisms and components relative to an explanandum which one is adopting the mechanistic stance towards in order to explain. Pattern perspectivism reflects the interesting role of a pattern detecting system in defining the role of patterns in explanation. This is a significant consequence for a Dennettian approach to mechanistic explanation, however, it is also one that remains grounded in the relationship between agents and dependable regularities in nature which, in the standard formulation, are objective and perspective-independent. Put otherwise, in our estimation, neither phenomenon nor pattern perspectivisms will especially worry those realists who are wary of scientific knowledge appearing too perspective-dependent. Hierarchy perspectivism, on the other hand, potentially threatens the objectivity of mechanistic structures themselves, posing a particular challenge for a purely ontic account of mechanistic explanation. Nevertheless, when situated properly within the mechanistic stance approach, which offers a form of realism understood in terms of real patterns, it becomes possible to see how even hierarchy perspectivism about mechanisms could be compatible with a kind of perspectival (and potentially ontic) realism about the hierarchical structure of mechanisms.

We take the analysis offered here to be indicative of a useful strategy for evaluating the perspectival credentials of mechanistic explanation more generally. Instead of, or at least before, passing final judgement on whether mechanistic explanation is or is not perspectival, we would do well to uncover the importantly distinct ways in which perspectives play a role in mechanistic explanation. In our estimate, this is likely to produce a more nuanced view of the conceptual landscape.

Moving forward, we believe a stance approach promises to shed light on other areas of debate surrounding mechanistic explanation. For example, the mechanistic stance could provide the resources for a useful distinction between two kinds of error in formulating putative explanations, namely, pattern misdetection and missing patterns. This is pertinent, for instance, in making sense of historical errors in medical diagnosis, such as understanding harmful constructs like drapetomania or hysteria, both in terms of implicating non-existent pathological mechanisms suggested by the alleged mental illness and omitting actual mechanisms (both physiological and social) that would explain behaviour (cf Pöyhönen 2013; Potochnik and Oliveria 2020).[7] This could perhaps be elucidated in terms of what Dennett has elsewhere called "lovely" and "suspect" qualities (1991b), with the former corresponding to real patterns and the latter

[7] We would like to thank an anonymous reviewer for drawing our attention to this point.

to *merely* perspectival patterns, i.e. patterns whose existence depends entirely on projections made by an observer, without any basis in the underlying structure of the world. While we have emphasised the perspectival nature of mechanistic stance explanation, it is importantly still dependent on the existence of real patterns, and therefore a successful mechanistic explanation should track lovely rather than suspect properties.

5 Conclusion

The new mechanism literature provides a substantial framework with which to understand the nature of scientific explanation in the biological and cognitive sciences and beyond. It has often been claimed that agent perspectives play some vital role in mechanistic explanation but there is little consistency or sustained analysis of mechanism's perspectivist credentials. This paper went some way to address this by beginning to develop a novel 'mechanistic stance' understanding of mechanistic explanation. According to this Dennettian approach, mechanistic explanation involves adopting a stance that lies between the more familiar physical and design stance, and which possess distinct features of its own.

The mechanistic stance is of independent interest, but it also provides a lens through which to assess the role of agent perspective in mechanistic explanation. We presented three ways in which one might construe the perspective-dependent nature of mechanistic explanation. Beyond embedding the mechanistic stance in a broader perspectivist theory in philosophy of science, 'phenomenon', 'pattern', and 'hierarchy' perspectivism picks up on specific features of the mechanistic stance. We do not wish to make any conclusive statement here about which of these perspectivisms should be endorsed by proponents of mechanistic explanation, although we do think that while both phenomenon and pattern perspectivism should be palatable to almost all mechanists, hierarchy perspectivism offers both a more controversial and yet potentially more interesting interpretation of the hierarchical structure of mechanisms.

Compliance with ethical standards

Conflict of interest We have no conflict of interest to declare.

Ethical approval Our research did not require ethical approval.

Informed consent Our research did not require informed consent.

References

Anastassiou, C., Perin, R., Markram, H., & Koch, C. (2011). Ephaptic coupling of cortical neurons. *Nature Neuroscience, 15*, 217–223.

Andersen, H. K. (2017). Patterns, information, and causation. *The Journal of Philosophy, 114/11*, 592–622.

Bechtel, W. (2008). Mental mechanisms: Philosophical perspectives on cognitive neuroscience. Lawrence Erlbaum Associates.

Bechtel, W. (2019). Resituating cognitive mechanisms within heterarchical networks controlling physiology and behavior. *Theory & Psychology, 29*(5), 620–639.
Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science C, 36*(2), 421–441.
Chaitin, G. J. (1966). On the length of programs for computing finite binary sequences. *Journal of the ACM, 13*(4), 547–569.
Chiang, C., Shivacharan, R., Wei, X., Gonzalez-Reyes, L., & Durand, D. (2019). Slow periodic activity in the longitudinal hippocampal slice can self-propagate non-synaptically by a mechanism consistent with ephaptic coupling. *The Journal of Physiology, 597*(1), 249–269.
Chirimuuta, M. (2014). Minimal models and canonical neural computations. *Synthese, 191*, 127–153.
Coelho Mollo, D.C. (Forthcoming). Against Computational Perspectivism. *British Journal for the Philosophy of Science* online first.
Craver, C. (2001). Role functions, mechanisms and hierarchy. *Philosophy of Science, 68*, 31–55.
Craver, C. (2007). *Explaining the brain*. Oxford: OUP.
Craver, C. (2013). In Huneman (Ed.), *Functions and Mechanisms: A Perspectivalist View. Functions: Selection and Mechanisms*. Dordrecht: Springer Netherlands.
Craver, C., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy, 2*, 547–563.
Craver, C., & Kaplan, D. (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.
Dennett, D. C. (1991a). Real patterns. *The Journal of Philosophy, 88*(1), 27–51.
Dennett, D. C. (1991b). Lovely and suspect qualities. In E. Villanueva (Ed.), *Consciousness*. Atascadero: Ridgeview.
Dewhurst, J. (2018a). Computing mechanisms without proper functions. *Minds and Machines, 28*, 569–588.
Dewhurst, J. (2018b). Individuation without representation. *The British Journal for the Philosophy of Science, 69*, 103–116.
Dewhurst, J. & Isaac, A.M.C. 2020. Mechanistic hierarchy realism and function Perspectivalism. Preprint available at http://philsci-archive.pitt.edu/16935/. Accessed 19 Dec 2020.
Dupré, J. (1993). *The disorder of things: Metaphysical foundations of the disunity of science*. Cambridge: Harvard University Press.
Garson, J. (2013). The functional sense of mechanism. *Philosophy of Science, 80*(3), 317–333.
Giere, R. (2004). How models are used to represent reality. *Philosophy of Science, 71*(5), 742–752.
Giere, R. N. (2006). *Scientific perspectivism*. Chicago: University of Chicago Press.
Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science, 69*(S3), S342–S353.
Glennan, S. (2017). *The new mechanical Philosophy*. Oxford: OUP.
Hardcastle, V. (1999). Understanding functions: A pragmatic approach. In Hardcastle (Ed.), *When Biology Meets Philosophy*. Cambridge: MIT Press.
Haugeland, J. 1998. “Pattern and being”. In *Having thought. Essays in the metaphysics of mind*. Cambridge: Harvard University press.
Illari, P. (2013). Mechanistic explanation: Integrating the ontic and epistemic. *Erkenntnis, 78*, 237–255.
Illari, P., & Williamson, J. (2012). What is a mechanism? *European Journal for the Philosophy of Science, 2*(1), 119–135.
Kaplan, D. (2011). Explanation and description in computational neuroscience. *Synthese, 183*, 339–373.
Kaplan, D., & Craver, C. (2011). The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science, 78*(4), 601–627.
Kaplan, D. 2017. “Neural computation, multiple Realizability, and the prospects for mechanistic explanation.” In *Explanation and Integration in Mind and Brain Science*, ed. Kaplan. Oxford: OUP.
Kaiser, M., & Krickel, B. (2017). The metaphysics of constitutive mechanistic phenomena. *British Journal for the Philosophy of Science, 68*, 745–779.
Kästner, L. (2018). Integrating mechanistic explanations through epistemic perspectives. *Studies in the History and Philosophy of Science, 68*, 68–79.
Kästner, L. & Haueis, (2019). Discovering patterns: On the norms of mechanistic inquiry. Erkenntnis.
Kellert, S., Longino, H. & Waters, C.K.. 2006. “The Pluralist Stance.” In *Scientific Pluralism*, eds. Kellert, Longino, & Waters. Minnesota Studies in the Philosophy of Science 19. Minneapolis: University of Minnesota Press.
Kolmogorov, A.N. 1965. “Three approaches to the quantitative definition of information.” *Problems of Information Transmission*, 1 /1: 1–7.
Ladyman, J., & Ross, D. (2007). *Every thing must go: Metaphysics naturalized*. Oxford: OUP.

Li, M. and Vitányi, P. 2008. An introduction to Kolmogorov complexity and its applications. Texts in Computer Science, Vol. 9. Springer, New York.
Machamer, P., Darden, L., & Craver, C. 2000. "Thinking about Mechanisms." *Philosophy of Science*, 67/1: 1–25.
Maley, C. & Piccinini, G. 2017. "A unified mechanistic account of teleological functions for psychology and neuroscience." In Kaplan (Ed.), *Explanation and integration in mind and brain science*. Oxford: OUP.
Massimi, M. 2012. "Scientific Perspectivism and Its Foes." *Philosophica*, 84: 25–52.
Massimi, M. 2018a. "Four Kinds of Perspectival Truth." *Philosophy and Phenomenological Research*, 96/2: 342–359.
Massimi, M. (2018b). Perspectival Modeling. *Philosophy of Science, 85*(3), 335–359.
Massimi, M., & McCoy, C. (2019). *Understanding Perspectivism*. Routledge.
Millhouse, T. Forthcoming. "Compressibility and the Reality of Patterns." *Philosophy of Science*.
Mitchell, S. D. (2009). *Unsimple truths: science, complexity and policy*. Chicago: University of Chicago Press.
Piccinini, G. (2007). Computation Mechanisms. *Philosophy of Science, 74*, 501–526.
Piccinini, G. (2015). *Physical Computation*. Oxford: OUP.
Potochnik, A. & de Oliveria, G.S. 2020. "Patterns in Cognitive Phenomena and Pluralism of Explanatory Styles." *Topics in Cognitive Science*, 12/4: 1306–1320.
Pöyhönen, S. 2013. "Carving the mind by its joints. Natural kinds and social construction in psychiatry". In Talmont-Kaminski K. & Milkowski M. (eds.), *Regarding the Mind, Naturally: Naturalist Approaches to the Sciences of the Mental.* Cambridge Scholars Press. pp. 30–48.
Putnam, H. & Oppenheim, P.. 1958. "Unity of Science as a Working Hypothesis." In Feigl et al. (eds.), *Minnesota Studies in the Philosophy of Science, vol. 2*. Minneapolis, MN: Minnesota University Press.
Robbins, P. & Jack, A.I. 2006. "The Phenomenal Stance." *Philosophical Studies*, 127/1: 59–85.
Ross, D. 2000. "Rainforest realism: A Dennettian theory of existence." In Ross, D., Brook, A., & Thompson, D. (eds.), *Dennett's Philosophy: A Comprehensive Assessment*. Cambridge, MA: MIT Press.
Shagrir, O. & Bechtel, W. 2017. "Marr's Computational Level and Delineating Phenomena." In *Explanation and Integration in Mind and Brain Science*, ed. Kaplan. Oxford: OUP.
Solomonoff, R.J. 1964a. "A formal theory of inductive inference. Part I." *Information and Control*, 7/1: 1–22.
Solomonoff, R.J. 1964b. "A formal theory of inductive inference. Part II." *Information and Control*, 7/2:224–254.
Suñé, A. & Martinez, M. 2019. "Real patterns and indispensability." *Synthese* online first.
Vermass, P.E., Carrara, M., Borgo, S., & Garbacz, P. 2011. "The design stance and its artefacts". *Synthese*, 190: 1131–1152.
Wimsatt, W.C. 1994. "The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets." *Canadian Journal of Philosophy Supplement*: 207–74.
Zawidzki, T. (2013). *Mindshaping*. Cambridge: MIT Press.