Getting over Atomism: Functional Decomposition in Complex Neural Systems

Daniel C. Burnston

Abstract

Functional decomposition is an important goal in the life sciences, and is central to mechanistic explanation and explanatory reduction. A growing literature in philosophy of science, however, has challenged decomposition-based notions of explanation. 'Holists' posit that complex systems exhibit context-sensitivity, dynamic interaction, and network dependence, and that these properties undermine decomposition. They then infer from the failure of decomposition to the failure of mechanistic explanation and reduction. I argue that complexity, so construed, is only incompatible with one notion of decomposition, which I call 'atomism', and not with decomposition writ large. Atomism posits that function ascriptions must be made to parts with minimal reference to the surrounding system. Complexity does indeed falsify atomism, but I contend that there is a weaker, 'contextualist' notion of decomposition that is fully compatible with the properties that holists cite. Contextualism suggests that the function of parts can shift with external context, and that interactions with other parts might help determine their context-appropriate functions. This still admits of functional decomposition within a given context. I will give examples based on the notion of oscillatory multiplexing in systems neuroscience. If contextualism is feasible, then holist inferences are faulty—one cannot infer from the presence of complexity to the failure of decomposition, mechanism, and reductionism.

- 1 Introduction
- 2 Atomism
- 3 Holist Inferences in Detail
- 4 Contextualism as an Alternative
- 5 Multiplexing and Contextualist Decomposition
 - **5.1** *Internal dynamics*
 - **5.2** *Dynamic interaction*
 - **5.3** *Network dependence*
- **6** Philosophical Upshot
 - **6.1** The scope and limits of mechanistic explanation
 - **6.2** *The context objection to reduction*
- 7 Conclusion

1 Introduction

Biological systems have parts, which can be distinguished spatially. When biologists functionally decompose a system, they attempt to determine how its distinct parts differentially contribute to its behaviour. Construed in this way, decomposition is fundamental to mechanistic explanation—explanation of the behaviour of a system in terms of its parts, the operations they perform, and their organization (Bechtel and Abrahamsen [2005]; Bechtel and Richardson [1993]; Machamer, Darden, and Craver, 2000). It is also central to post-Nagelian reductionist approaches to explanation. 'Explanatory' or 'mechanistic' reductions are those that explain the behaviour of a whole in terms of the behaviour of its parts (Bechtel and Hamilton [2007]; Sarkar [1992]; Wimsatt [1976]).

In this article I criticize a kind of inference that is often made by opponents of mechanistic and reductive explanation, which I call 'holist inferences'. Holists infer from the properties of context-sensitivity, dynamic interaction, and network dependence in complex systems to the failure of functional decomposition, and hence to the failure of mechanistic explanation and explanatory reduction. Often, holist inferences are accompanied by the view that we should switch to kinds of explanations that are taken to be non-mechanistic, particularly dynamical systems and graph theoretic approaches (Chemero and Silberstein [2008]; Silberstein and Chemero [2013]; Rathkopf [2018]).

I will argue that holist inferences are invalid, because they take an overly simplified form of decomposition as their target. I call this view of decomposition 'atomism'. Atomism holds that parts of systems should be functionally individuated according to what they do intrinsically, rather than what they do in interaction with other parts. While atomism is indeed falsified by complexity, there is a weaker, 'contextualist' approach to decomposition that is compatible with all of the properties. If this is the case, then one cannot infer from the presence of complexity to the failure of decomposition. I will discuss a case study from systems neuroscience that illustrates how the contextualist approach validates decomposition in the presence of complexity.

My primary aim is to critique holist inferences, but I also offer contextualism as a positive proposal for how mechanists and reductionists can embrace complexity. Many mechanists are not atomists, but the mechanist program as a whole has yet to agree on an alternative. Hence—as I will seek to establish exegetically—holists have continued critiquing atomism as a stand-in for the commitments of mechanistic and reductionist views. This has resulted in some disconnect in the literature, with proponents of mechanism claiming their views are compatible with complexity, and holists contending that they are not. I propose contextualist decomposition as an alternative core commitment for mechanistic and reductionist positions. Importantly, this clarifies a particular kind of explanatory target for these positions, namely,

understanding organizational principles that implement context-specific functional decompositions. The case study is offered as an exemplar of this kind of explanation. In addition, it shows how extant contextualist approaches to neural function (Burnston [2016a], [2016b]; Klein [2012]) can help make sense of massively multifunctional brain networks.

In Section 2, I will lay out atomism in more detail, and give some exegetical evidence that it is the position that critics of mechanism and reductionism attack. In Section 3, I will exposit two recent, sophisticated versions of holist inferences. Section 4 outlines the contextualist alternative. Section 5 then gives the case study from systems neuroscience, which is built around the notion of 'multiplexing' functions via signal modulation, and explains in detail how the framework embraces both complexity and decomposition. In Section 6, I will discuss the upshot of contextualist decomposition for questions about the scope and limits of mechanistic explanation and the feasibility of reductionism. Section 7 concludes.

2 Atomism

What I call 'holist' inferences move from the presence of contextsensitivity, interactive dynamics, and network-dependence in a system's functioning to the conclusion that its behaviour is not decomposable. A second inference is often made from here to the failure of mechanistic and reductive explanation, and/or to the presence of emergent properties. I am only targeting this type of inference in this article. I am thus not giving a positive argument in favour of mechanism or reduction, and will avoid getting bogged down in tricky details about how to define emergent properties (Boogerd *et al.* [2005]), or how 'strong' a reductionist thesis to embrace (Bickle [2006]; Gillett [2016]; Theurer [2013]). I will refer to any position based on these inferences as a 'holist' position, where holism is meant to contrast with explanations based on functional decomposition.

The main claim I will pursue is that there are two notions of decomposition: one that is indeed rendered glaringly false by complexity, and a weaker one that in fact incorporates complexity wholesale. The first conception I refer to as 'atomism', and is based around the notion of intrinsic function.

Intrinsicality (INT): The function of a part P is to be specified with limited reference to the system in which P functions.

The 'limited' in the above definition of INT is in reference to the fact that most atomist function ascriptions assume a set of normal background conditions (Klein [2018]) and causal precursors. So, one might have an atomistic view of the function of a particular gene—for instance, coding for a single trait—but the gene's exercising this function would of course depend on the presence of normal genetic transmission (causal precursor) and any transcriptional and epigenetic factors (background conditions) that

are required for the output of the gene to result in the trait. Similarly, a muscle fibre might function to move a hand, but its doing so requires the appropriate inputs and the actual presence of the hand, along with the appropriate downstream wiring. Atomism requires that the background conditions remain in the background. Changes in the background conditions might make the part unable to function, but they should not change what function it performs. If they did, then the (theory of the) part's operation would not be isolatable from the rest of the system in the way that atomism requires.

Atomism conjoins intrinsicality with a corollary about direction-ofexplanation in biological systems, which I call 'priority'.

Priority (PRI): The functional properties of the system, *S*, are to be explained in terms of the intrinsic functions of its individual parts.

Priority presents both an explanatory and methodological dictum. Explanations proceed from parts to wholes, and go only in this direction. As such, investigation should begin by first isolating parts and describing their functions, and only then describe how they interact to produce phenomena.

There are two implications of INT and PRI that I want to stress. First, they suggest a very minimal set of functions, or even a univocal function, for each part. This is not implied by INT on its own. In principle it is

possible that a part could perform multiple different functions purely based on its internal operations.¹ Multifunctionality is in tension with the conjunction of INT and PRI, however. Even if you could establish that a part could perform multiple distinct functions purely by looking at that part, it would require reference to properties of the surrounding system to say when and how those functions are implemented. So, INT and PRI together imply that there is a consistent function performed by each *P*.

Second, the combination of limited functions and directionality of explanation implies a relative lack of flexibility of explanations based on decomposition. Since each of the parts must be studied with limited reference to the system, the baseline resource that atomistic explanations start from is a list of the intrinsic functions of the parts. When one goes to explain a given system behaviour, one can at best recombine these functions in different ways. Since the list of intrinsic functions remains constant, there is a limit to how much the mechanistic organization of the system should change with context.

Atomism can perhaps best be read as a thesis about what mechanistic or reductive explanations should aspire to. On this reading, the best cases of these explanations assign a single function to each P, with limited reference to the system, and explanations will be less mechanistic—reductive the more reference to the system needs to be made. This 'more-

¹ Thanks to Colin Klein and an anonymous reviewer for pushing me to clarify these points.

or-less' conception of decomposition is often evinced in discussions of 'near decomposability' (Bechtel and Richardon [1993]; Simon [1962]).

Why would anyone hold atomism? First, it provides a methodological dictum that has in fact directed different stages of biological research, particularly early stages of research directed towards discovering relevant parts and exploring their capacities. Second, atomism offers particularly powerful generalizations—if we can find the univocal function of a series of parts, then the flexibility of the system consists only in how they are recombined, and we have fewer degrees of freedom in theorizing about the system. Regardless of its potential accuracy in describing some stages of research, or its theoretical appeal, I will argue that atomism should not be construed as definitive of decomposition.

The three properties that holists cite as incompatible with decomposition are:

Context Sensitivity: The function of a given *P* can change depending on events external to the part.

Interactive Dynamics: The function of a given P depends on temporally sensitive interactions between distinct quantities in or parts of the system.

Network Dependence: The function of a given P depends on the topological structure of the system.

All three properties of complexity are indeed incompatible with atomism. This is because they undermine the idea that one can fix a functional description in any given explanation without making reference to the broader system in which a part operates. Consider contextual variation first. Suppose that a given P's function changes with context. If that is the case, then we must figure out what context we are in when describing its function, and how that particular function is implemented in that context. A natural way of accounting for this change is that the part is interacting with other parts of the system and its environment in a new way, and these external factors can no longer be backgrounded or relegated to causal precursors of a univocal function. Moreover, the more continual and dynamic these interactions are, the more tightly coupled the behaviour of P will be with other components of the system, and the less isolable it will be. Finally, the more the behaviour of the system depends on global properties, such as those described by network theory, the less primary a description of the individual part will be in explaining system behaviour. These properties thus violate INT and PRI.

While I can't fully canvass the relevant literatures here, I believe that atomism is standardly taken to be the core notion of decomposition in discussions surrounding mechanism and reduction. I'll here just present

some quotes that suggest this, and then discuss some specific examples of holist inferences in more detail in Section 3.

Consider the following quotes about functional decomposition:

A decomposable system is one in which each subsystem operates according to its own intrinsic principles, independently of the others. (Pessoa [2008], p. 155)

[A function] is whatever single, relatively simple thing a local neural circuit does for or offers to all of the functional complexes of which the circuit is a part. (Anderson [2010b], p. 295)²

These characterizations cite INT and PRI as definitive of decomposition. Now consider the following quotation that is critical of mechanistic explanation:

The guiding image of mechanisms as machinelike structures strongly suggests that they are made of discrete parts, each of which carries out a dedicated function. [...] the less easily the parts can be separated from one another while retaining their own functions, the further the system drifts away from being mechanistic. [...] if several functions are assigned to the same structural element, the specialization of functional parts that mechanism requires is violated. (Weiskopf [2016], p. 677)

² Anderson's view in his ([2014]) softens this notion of decomposition, and moves closer to the kind of contextualism I advocate. The relationship between my view and Anderson's current exposition of 'neural reuse' is too complicated to pursue in detail here.

This quote moves directly from a purported failure of intrinsicality to the failure of the mechanistic approach. Finally, consider some quotes about explanatory reduction:

Dynamic interactions can bring about qualitatively new behaviour in complex systems. This is precisely where prediction of system behaviour on the basis of simpler subsystems fails. We cannot predict the behaviour of the components within the entire system and so cannot predict systemic behaviour [from component behaviour]. (Boogerd *et al.* [2005], p. 156)

Reductive explanations in the life sciences exhibit three characteristics: first, they refer only to factors at a lower level than the phenomenon at issue, second, they focus on internal factors and thus ignore or simplify the environment of a system, and, third, they cite only the parts of a system in isolation. (Kaiser [2011], p. 453)

These quotes suggest that a reductive explanation must consist purely of conjunctions of intrinsic function ascriptions, thus reflecting INT and PRI. Each article then goes on to challenge reductionism on the basis of complexity. These quotes therefore express examples of holist inferences.

3 Holist Inferences in Detail

In a holist inference, one starts from the assumption that if a system is decomposable, then atomism must be true of that system. One then notes

the presence of complexity in that system, and argues that it falsifies decomposition. From the further premise that decomposition is the *sine qua non* of mechanistic and–or reductive explanation, one gets to the conclusion that mechanism and reduction must fail for that system. Since dynamical and–or network explanation do not fail for that system, we then get the further conclusion that these frameworks, rather than mechanistic ones, do explain the system. The atomistic assumption is often left implicit in these discussions, but it is revealed by the kinds of properties holists take to license the move away from decomposition.

Chemero and Silberstein (Chemero and Silberstein [2008]; Silberstein and Chemero [2013]) have recently advanced an influential holist argument, specifically in the neuroscience case. They argue that both mechanist and reductive approaches are firmly committed to decomposition, and that decomposition is not possible when properties of complexity are present. For them, complexity arises when system components dynamically interact, where some of these interactions are 'non-local' and are timing- and context-dependent. In these scenarios, they claim, one cannot isolate functions to distinct parts of the system, and thus the system's parts don't 'locally determine the capacities or properties' of the system as a whole. They state:

In order to get [...] reductionism off the ground, there have to be context-independent or invariant fundamental parts with intrinsic properties such as atoms or cells whose temporal evolution is governed by some fundamental context-independent [...] laws. (Chemero and Silberstein [2008], p. 14)

Systems biology and systems neuroscience contain robust dynamical and mathematical explanations of some phenomena in which the essential explanatory work is not be being done by localization and decomposition. More positively, the explanatory work in these models is being done by their graphical/network properties and the dynamics thereupon. (Silberstein and Chemero [2013], p. 960)

Hence, as befitting the holist inference schema above, Chemero and Silberstein argue that complexity undermines both reductionism and mechanism, and they offer dynamical and graph-theoretic description as a fundamentally different kind of explanation. The phenomena they take to be explained by these models—and hence not by decomposition—include 'cognitive capacities' and 'connectivity disorders' such as schizophrenia and autism.

Rathkopf ([2018]) argues similarly that in large organizations of interacting components, decomposition fails. He suggests that in these cases 'global patterns variables' are what explain the behaviour of the system, rather than descriptions of their individual components. The global pattern variables he cites are from graph-theoretic descriptions of networks. They include: 'characteristic path length', which is the average distance between any two nodes in a network; 'clustering coefficient', which is the tendency for two nodes connected to a common node to be

themselves connected; and 'betweenness centrality', which is the proportion of shortest paths between any two nodes in network that pass through a given node. Rathkopf argues that these descriptions 'collapse' the patterns of interactions between components into global explanatory features. He shows how 'small-worldness' (the combination of a high clustering coefficient and a low characteristic path length) has been used to explain the rate-of-transmission of diseases, and how betweenness centrality has been used to explain traffic patterns.

Rathkopf claims that in large interactive systems, it is the global pattern of organization that is explanatorily relevant, rather than anything about the individual nodes or their local interactions. When complexity is present, the system is 'not even close to being nearly-decomposable' and scientists are not 'free to develop a theory of each subsystem, and then compose the predictions of those theories in order to yield predictions about the behaviour of the whole' ([2018], p. 69; notice the atomism in the framing here). According to Rathkopf, complexity demands a 'completely distinct explanatory strategy' from decompositional explanation. He takes graph-theoretic explanation to be such a strategy, compared to decomposition approaches that 'strip away' complexity.

The holism that one ends up with as a result of these holist inferences is one in which a description that does not decompose the system, rather than one that does, is explanatory of the phenomenon. A few clarifications are in order. First, no holist thinks that biological systems are not

decomposable in any relevant sense. Holists of course admit that biological systems have parts, which are independently describable in some ways—one can talk about types of cells, neurotransmitters, spike patterns, and so on. Moreover, both Rathkopf and Chemero and Silberstein admit that some explanatory scenarios may be susceptible to decomposition. There may be systems with a small number of components or a particularly simple form of interaction for which decompositional approaches work. There may even be particular systems for which some behaviours involve simple interactions and other behaviours exhibit complexity. The first would be amenable to decompositional explanation; the second wouldn't be. So, while holists are pluralists to a degree—they posit a 'complementarity' (Silberstein and Chemero [2013], p. 960) or 'division-of-labor' (Rathkopf [2018], p. 75) between decompositional and holist explanations—they still see a situational opposition between decompositional and holist explanations. When a system behaviour exhibits complexity it requires a holistic, and not a decompositional, explanation.

Since I agree with holists that atomism is unsustainable for complex systems, I won't spend time here arguing against atomism as a regulative ideal. Rather, the question I will pursue is whether, given the failure of atomism, holist arguments in fact force the abandonment of decomposition. My goal is to prove that holist arguments are invalid. To do this, it needs to be shown that in the very kinds of scenarios in which

holist arguments are meant to apply, (i) complexity is indeed present and (ii) decomposition is possible and theoretically fruitful anyway.

In what follows I will offer an exemplar case that meets these requirements—explanation of cognitive capacities via multiplexing of function in the cortex. Cognitive capacities, recall, are one of the explananda that Chemero and Silberstein deny are susceptible to decomposition. And, as we will see, these explanations sometimes invoke pattern variables, and thus according to Rathkopf should be nondecomposable. I will argue that contextualist decomposition is still possible in these settings. If I am right about the case, then it shows that, in general, holist inferences are invalid—one cannot simply cite complexity in one's system or phenomenon of interest as a way of denying mechanistic or reductive explanations for that explanandum. This falls short of showing that all systems are decomposable, or that all explanations are mechanistic.³ Again, I am simply targeting holist inferences here, and offering contextualism as a potential strategy for mechanists and reductionists to pursue in other systems of interest.

³ For instance, there are added complications in thinking about system properties such as robustness (Huneman [2010]; Jones [2014]), about general 'design principles' (Brigandt, Green, and O'Malley [2018]; van Eck and Mennes [2018]), or about optimality tradeoffs (Rice [2015]). Since these frameworks are interested in system properties in general, they often intentionally abstract away from details about particular systems. There are also complicated considerations surrounding the explanatory relationship between mechanisms and models in cognitive psychology (Piccinini and Craver [2011]; Stinson [2016]; Weiskopf [2016]). Regardless of the overall merits of these discussions, I will argue that cognitive functioning in the cortex is not best explained in a holist way.

4 Contextualism as an Alternative

I call the alternative, non-atomistic form of decomposition 'contextualist', because it starts out from the notion that the function a given part might change depending on the context in which it operates. As discussed above, contextual variation is a wedge that once opened, immediately draws in the other properties. If the function of a part shifts with context, then a natural explanation for those shifts is that it is dynamically interacting with other parts of the system and the environment. Thus, contextualism abandons INT and PRI, which are definitive of atomism. Concomitantly, contextualism doesn't imply that we need a list of intrinsic functions that explains S's behaviour in any context. Quite the contrary: contextualism only suggests that we should be able to find differences between what the distinct parts of S are doing within a given context. Contextualism is thus fully compatible with multifunctionality, and is considerably more flexible than atomism in how it describes systems. It says that analysis starts from a context—we don't have to expect that the functional divisions made in one context will be the same as those made in another.

Contextualism is a significant departure from the recent history of function ascription in neuroscience, which has decidedly atomist leanings.⁴ Atomism attempts to describe each part of the brain as

⁴ As Mundale ([2002]) notes, there is also a long history of holist positions in neuroscience, ranging from equipotentiality across the whole brain to the possibility of some regional variation. There are also middle ground holist positions but, as Mundale ([2002], p. 316) states: 'Generally [...] holism does not accord well with functional

performing a univocal function, distinct from what other parts do. Consider the standard picture of the visual cortex. On an atomist view, visual cortex has the function of processing visual information. Each area in the visual cortex plays a particular role in this process, representing particular types of visual information at particular degrees of abstraction.

Giving a particular atomist description of a brain area's function involves settling on a particular spatial localization and applying a univocal functional description at the appropriate level of generalization. Area V1 is standardly subdivided into distinct populations representing orientation, wavelength, and displacement. Area MT is generally construed as a single unit, with the function of representing pattern motion on the basis of information about local displacement from V1. This function is distinct from areas further along the dorsal stream, which represent more complex motion patterns. V4 is standardly subdivided into two sections, one with the univocal function of representing natural colours, and the other for representing visual form (Van Essen and DeYoe [1988]).

Each atomist functional description is a generalization that is meant to capture a part's function as a whole, in a way that distinguishes it from other areas. A generalization pitched at too low a level would not accurately characterize the function. Arguing that MT processes direction

of motion would not account for the fact that it is sensitive to speed as well, while the pattern motion description does. Alternatively, a generalization pitched at too high a level would not distinguish MT from other areas—its function is more specific, say, than 'processing visual information', a description it shares with many other brain areas (Burnston [2016a]; Rathkopf [2013]. Giving these kinds of descriptions also involves determining causal precursors. For instance, MT requires information about displacement to be input from V1 before it can compute pattern motion.

The data that leads one away from atomist and towards contextualist views is data that speaks towards multifunctionality. If a brain area genuinely has multiple functions, and which function it plays varies with perceptual or behavioural context, then a contextualist rather than atomist description is appropriate for that area. Importantly, contextualist generalizations must operate at the same spatial and descriptive level as the atomist ones they are meant to replace.

Consider MT again. The last twenty years of investigation into MT have revealed that MT cells have detailed responses to several varieties of depth, as signalled by binocular disparity, and these responses are psychophysically important. MT responses to coarse depth (distance from the perceiver) are predictive of animals' depth judgements, similarly to how MT's motion responses are predictive of motion judgements, and these effects can be manipulated by microstimulation (Uka and DeAngelis

[2006]). Moreover, MT cells' depth and motion responses are statistically independent, meaning that MT can represent depth even in the absence of motion, or vice versa (for instance, a moving stimulus at the 'fixation plane', where no binocular disparity information is present; Palanca and DeAngelis [2003]). Moreover, and importantly, the vast majority of MT cells (94%) have both motion and depth-selectivity, and depth selectivity properties are organized into representational maps (in which cells with similar selectivity properties are close to each other) that cross-cut the organization of motion representation (DeAngelis *et al.* [1998]). In addition to coarse depth, MT cells also exhibit selectivity to fine depth features of tilt and slant, as cued by combinations of disparity and velocity cues.

These properties suggest that MT is genuinely multifunctional—its constituent cells have response properties to different types of information, and these lead to different generalizations about what the area is representing in different contexts. The genuine multiplicity of the functions speaks against a couple of responses atomists might make. First, it speaks against simply re-describing the univocal generalization to account for the new data. Whereas there is a natural generalization that groups leftward and rightward pattern motion under 'pattern motion', there are limited candidates to describe the combined motion, coarse depth, and fine-depth descriptions under a single generalization that would still distinguish MT from other visual areas (see Burnston [2016a] for

more thorough discussion). Moreover, since it is the entire area that exhibits these properties, one cannot apply what McCaffrey ([2015]) has called a 'subdivide-and-conquer' strategy, such as the one standardly used in V4, which breaks up a single area into multiple function-specific areas. Theorists have recently gotten quite liberal with subdivide-and-conquer strategies, even suggesting that functions can be spatially isolated to distinct but interdigitated populations (which Kanwisher [2010], calls 'archipelagos'), but if it is the very same population that is multifunctional, this strategy is off the table.⁵

A number of theorists have recently argued that contextualism and functional localization are not incompatible (Bechtel [2012]; Burnston [2016a], [2016b]; Khalidi [2017]; Klein [2012]; McIntosh [2004]; see Brigandt [2017] for contextualism applied to individuation of function more broadly). Different discussions of contextualism focus on different aspects of context and on different philosophical issues arising from the view. My purpose here is to apply contextualism to broader debates about decomposition rather than defend a particular version of it, so I will only

⁵ As an anonymous reviewer helpfully points out, one might attempt a kind of generalization about which individual cells or groups of cells could be context-sensitive, even while the area they constitute is uni-functional. There is some evidence that individual MT cells, for instance, can switch from representing component (local movement) to pattern motion depending on features of the stimulus, and one could try to interpret this as contextual variation at the lower-level and univocality at the higher (see Burnston [2016a], for detailed discussion of this case; for further discussion of complexities in MT cells' representation of motion, see Madary [2013]). This will not work for the depth case, however, since the entire area comprises distinct representational functions.

give rough outline of the view, drawing heavily on the formulation I've developed elsewhere (Burnston [2016a], [2016b]). Contextualism is an attempt to isolate functions to parts of the brain within contexts. It argues that the brain is functionally decomposable so long as (i) within a given context, we can specify what function a part of the brain is performing, and (ii) different areas differ in their functional 'profiles'—the way their functions vary as contexts change (Anderson [2014]; Scarantino [2012]). MT and V4 for instance, both represent a range of features in different contexts, but these features and contexts differ.

Contextualism faces epistemic challenges, particularly in how it produces useful generalizations and how it guides research. Contextualist function ascriptions are importantly 'open'—they admit continual amendment to function ascriptions as novel contexts are explored. The key to producing useful generalizations is to seek to demarcate the changes of context that produce a change in what a brain area does, versus the changes that do not. For instance, MT represents depth in both stationary and moving stimuli. Its depth-representing function is thus invariant across changes along the motion dimension, and so long as we are within any of these contexts, we can projectably say that MT will represent that

⁶ Of course, contextualism must also establish that it is the correct way to read the data. One alternative is to claim that evidence in favour of multifunctionality is only apparent, and that each brain area can be characterized as performing a variety of abstract signal transform (see, for example, Anderson [2010]; Bergeron [2007]; Rathkopf [2013]). Burnston ([2016a]) argues in detail that these views cannot capture the contextual variation in brain systems in a way compatible with atomism, and thus that there are non-eliminable contextual elements in the functioning of brain areas.

information (Klein [2012]). On the other hand, extant function ascriptions—whether starting from an atomist ascription or a contextually indexed one—guide research by providing defeasible pragmatic assumptions about what a brain area will do in an unexplored context. These assumptions invoke researchers to attempt to show changes in an area's behaviour in a novel context that cannot be accounted for under extant function ascriptions, before expanding their understanding of the area to include a new, contextually indexed function.

I have only focused on one particular brain area here, but as noted above, atomistic function ascriptions are very standard in neuroscience. The decomposition of the visual system into feature-specific areas is often seen as an exemplar for other sensory systems. It is also common, for instance, to divide the language system into production (Broca's area) and comprehension (Wernicke's) systems. Motor systems are standardly seen to operate on a representational hierarchy similar to the visual system (Grafton and Hamilton [2007]). More controversial recent proposals include the idea that there is a lateral-medial distinction in the prefrontal cortex for deliberative versus emotional decision-making (Cushman, Young, and Greene [2010]). But, as with MT, there is widespread evidence of multifunctionality in other areas of the brain, and this evidence exists from different measurement techniques and at different spatial scales. According to a meta-review of fMRI data by Anderson ([2010a], [2010b]), the average cortical area is implicated in no less than

nine distinct types of task.⁷ Electrophysiological advances have shown that cells across the frontal cortex and in perceptual areas, as well as in subcortical areas, are 'hyperdimensional'—they encode for a range of stimulus, response, and context parameters, depending in the task (Fusi *et al.* [2016]; Panzeri *et al.* [2015]). Again here, the data suggests multifunctionality both at the level of whole brain areas and at the level of individual cells, speaking against subdivide-and-conquer strategies.

In addition to being generally multifunctional, brain areas also exhibit high degrees of interactivity, and it has become clear that interactions between different parts of the brain are vital for fulfilling cognitive functions. I will give several examples below. These kinds of interactions are particularly important for the discussion of holist inferences. Contextualists argue, contra atomism, that contextual variation and multifunctionality are compatible with decomposition. The question then, naturally turns to how those different functions are implemented. What are the organizational features of the brain that implement contextually varying functions, and do these require resorting to a holist kind of view? Moreover, understanding these interactions is important for understanding how the brain itself is organized. If brain areas are multifunctional, then one of the major problems that the brain has to solve is how to solicit the

⁷ How these results turn out depends on how one parses cognitive domains across studies, and thus on adopting a cognitive ontology. Even across multiple ways of construing cognitive domains, however, brain areas tend to exhibit functional diversity in meta-analyses (see Anderson [2014] for review).

right functions in the right contexts. There thus need to be principles of control that implement the correct functional organization in the appropriate settings. In the next section I will explore a framework that has recently been developed in systems neuroscience, which has begun to address these problems. I will then argue that this proposal, while abandoning atomism, is compatible with contextualist decomposition.

5 Multiplexing and Contextualist Decomposition

The proposal I will pursue is that neural activity is 'multiplexed'. That is, physiological activity within a given brain unit—cell, population of cells, or brain area—is not dependent solely on its inputs, but is modulated by background variables. I will show examples where different patterns of modulation produce different kinds of functional contributions within the very same group of cells or brain area. This in turn shows how context-sensitive functions can be implemented.

Multiplexing is implemented by the relationship between action potentials of individual cells and the background 'local field potential' (LFP). The LFP is basically a summation of electrical activity in a particular part of the brain. It involves not only neurons firing action potentials, but also post-spike polarizations, subthreshold activity, and post-synaptic activation (Khanna and Carmena [2015]; Panzeri *et al.* [2010]). LFPs can be measured using implanted electrodes, as well as extra-cranially through EEG. While the basic units of LFP production are

networks of roughly 10,000 neurons (Canolty and Knight [2010]), EEG results show that there are coherent patterns of LFP activity at the level of entire brain areas and, as we will see later, of brain networks. LFPs within neural populations are oscillatory. They exhibit electrical activity at characteristic frequencies, which are described in terms of frequency bands—delta (1–4 Hz), theta (4–8 Hz), alpha (8–12 Hz), beta (12–30 Hz) and gamma (30 Hz and above).⁸ A given brain area can also exhibit multiple distinct frequencies within its LFP at a given time.

Consider the abstract example of multiplexing given by Watrous *et al.* ([2015]) in Figure 1. Suppose that a single cell or group of cells exhibits the responses to 'light' and 'dark' stimuli shown in the bottom part of (a). The firing rate alone will not distinguish between light and dark. However, modulation by properties of the background LFP, including frequency, phase, and—or amplitude can tease apart the signals. Modulation occurs when spiking behaviour is indexed to certain properties of the LFP. Different combinations of such schemes are shown in the other panels, with the modulation schemes shown on the left and the signal differentiation in the heat maps to the right.

⁸ Where one draws the lines between frequency bands is somewhat arbitrary (these designations are from Canolty *et al.* [2010]). This arbitrariness makes no difference for the sort of coding principles I will discuss—what matters is that different frequencies are available for different functions via multiplexing. The dependence of function on LFP oscillation has been cited in Bechtel ([2012]) and Viola and Zanin ([2017]), but has not received thorough philosophical exposition. Haueis ([2018]) considers the possibility that oscillations in the cortex, particularly in the gamma band, have 'infrastructural' rather than cognitive functions, but as he admits, a general contextualism is compatible with oscillations playing both infrastructural and cognitive roles.

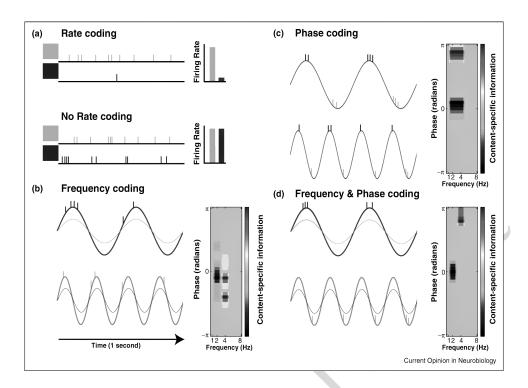


Figure 1. From Watrous et al. ([2015]). Reproduced with permission.

In panel (b) the signals are modulated by frequency. The dark signal is only affected by a low-frequency waveform, and the light one only by a higher frequency. The signals are not significantly modulated by phase, as is shown in the spread across the phase dimension of the heat map. Panel (c) shows phase modulation. Both firing patterns are modulated by two distinct frequencies, but at distinct phases of each—dark at the peak and light at the trough. Hence, frequency modulation does not differentiate them, but phase modulation does, as the heat map indicates. Panel (d) shows a combination of frequency and phase modulation, with dark signals grouped at the peak of the low-frequency waveform, while light signals are grouped at the trough of the higher frequency one. This further

differentiates the signals, again as the heat map indicates. Further, the schemes in (b) and (d) are also amplitude-specific, with modulation occurring only at a specific amplitude in each frequency.

The basic principle is abstractly similar to frequency modulation in wireless communication, where a content-signal is indexed to a background carrier frequency. Employing a metaphor from communication technology, Akam and Kullmann ([2014]) say that the background frequency is 'meta-data', which distinguishes different signals even if they spatially overlap, and allows for those signals to be selectively responded to for particular purposes.

Multiplexing aligns with contextualist function ascriptions when the background modulation changes what type of function a population or brain area is contributing to. This can be done, for instance, by changing what kind of information is being processed, as in the case discussed above, or by changing what functional network a brain area is part of, as we will discuss below. If multiplexing contributes to specifying what function a brain area is performing within a context, then there must be ways of setting up the appropriate multiplexing scheme for that context,

⁹ The term 'multiplexing' is occasionally employed in different ways in neuroscience, some of which are weaker than what I mean here. The term occasionally refers to distinct properties of a firing pattern (for instance, spike timing versus frequency) carrying different information (Harvey *et al.* [2013]). Akam and Kullmann's ([2014]) discuss multiplexing as primarily a means for signal routing. Neither of these uses imply a change in what is represented by the brain area across contexts. I will thus use the term to refer exclusively to the stronger, function-shifting role for multiplexing posited by Watrous *et al.* Thanks to Bryce Gessell for pushing me to clarify this point.

and for the rest of the brain to respond appropriately. Two principles that help implement the needed interactions on the multiplexing framework are 'phase synchrony' and 'cross frequency coupling'. Phase synchrony occurs when the LFPs in two distinct brain areas become synchronized. Cross-frequency coupling occurs within a brain area when multiple distinct phase components become coupled, usually with a lower-frequency component modulating a higher-frequency one. This can in turn affect cell firing. A common coding scheme is for lower-frequency activity, such as in the theta or beta band, to modulate gamma activity, which is shown to closely mirror spiking behaviour in populations of individual neurons (Lisman and Jensen [2013]).

This framework allows for a multi-scale account of how multiplexing contributes to functional modulation. If, for instance, A and B exhibit phase synchrony in their low-frequency LFPs, and at B the low frequency couples with high-frequency LFP components, this in turn provides a way for A to modulate the behaviour of cells in B. That is, the low-frequency interaction between A and B might provide the meta-data that modulates high-frequency spiking behaviour at B. This, in the abstract, is the framework I'll be analysing for the remainder of the article.

I have already argued above that context-sensitivity is not in principle incompatible with decomposition, due to the presence of a viable contextualist alternative. The question then turns to whether the other properties—dynamics and network dependence, in fact force a holist

interpretation. I will now look at three cases: dynamic variability within a brain area (Section 5.1), dynamic interaction between brain areas (Section 5.2), and dependence on network structure (Section 5.3), and argue that dynamics and network-dependence, far from being incompatible with decomposition, in fact show how context-sensitive functional decompositions are implemented in complex brain networks.

5.1 Internal dynamics

In a brilliant study, Canolty *et al.* ([2012]) measured activity from a group of cells as they contributed to multiple distinct tasks, in relation to the background LFP. They focused on motor cortex (M1) in monkeys. M1 is known independently to contribute to sophisticated behavioural movements of the limbs (Graziano [2006]). So, the first task the monkeys had to perform was to reach towards a target. The other task was a brain-computer interface task. In this kind of setup, an electrode is implanted into the brain, and the subject learns to control a computer cursor by intentionally manipulating the patterns of electricity on the electrode. Canolty *et al.* inserted the electrode over M1 cells involved in the reach task, so that the very same cells would contribute to both tasks. The question was how the same set of M1 cells might flexibly shift their contributions between the two contexts.

Canolty *et al.* measured how spiking in individual M1 cells varied in relation to the background beta LFP across tasks. What they found was

striking: even at the single cell level, firing patterns were modulated distinctly but reliably by the background LFP in the distinct tasks. Individual cells had particular patterns of modulation that depended on both the phase and the amplitude of the beta LFP. (The amplitude of the beta LFP varied during the course of the task, allowing the relative modulation of cells at different beta amplitudes to be measured.) But these patterns were different in the two distinct tasks—a given cell would exhibit a certain pattern of modulation by beta in the cursor task, then a distinct pattern in the reach task. So, even if the cells are active during both tasks, they can contribute to distinct behaviours depending on the task context. This is especially true at the population level. While different cells had different precise modulation patterns, the whole population would produce extremely different patterns in the two tasks, and these patterns were repeatable across different instances of the same type of task.

As Canolty *et al.* note, the combination of flexible but reliably invoked modulation schemes is just the kind of property that a multifunctional system needs to exhibit to produce the right functions in the right contexts. It is important to stress that these modulations are dynamical processes within the system—the change in function is due to the interaction of multiple temporally defined properties (firing and LFP), each of which vary over time, and which change their pattern of interactions as task context varies. So, if dynamic modulation of function were really

incompatible with decomposition, then this kind of result would be evidence against decomposition.

The results, however, are in no way incompatible with contextualist decomposition. Contextualism does not imply, as atomism does, that we could not map distinct functions, in distinct contexts, to the very same part. In each task context here, it is difficult to debate that the function of moving the arm or moving the cursor is being implemented by M1 cells. And the multiplexing framework provides explanatory principles for how those distinct functions can be implemented in distinct contexts. What we have here is a story about how M1 cells serve two distinct functions in distinct contexts, not one in which there are no specific function attributions that can be made to M1.

Now, this is not the whole story, since something must implement the change in modulation in M1 with changes in the task. The holist might suggest that the dependence of modulation on factors external to the population of cells means that the behaviour of the population is sensitive to the functioning of the whole system. We thus now turn to dynamic interaction and network dependence.

5.2 Dynamic interaction

At the outset, it is worth noting that the multiplexing framework opens up an extreme amount of flexibility in terms of the kinds of modulations that can be employed. There are at least three factors that contribute to this flexibility. Functions can vary by (i) the particular pattern of modulation at a given frequency; (ii) the particular frequencies that make up the coding scheme; and (iii) the particular brain areas involved. We just saw an example of the first type, in which different patterns of cell firing modulation by beta contribute to distinct functions. Beta, however, plays different roles in different networks—it also is involved in sensorimotor decision making, for instance, as we will see below. Alpha and theta similarly make diverse contributions depending on the functional network in which they are acting (Voytek *et al.* [2010]; Lisman and Jensen [2013]). And functional networks for given tasks can be differentiated according to which brain areas they comprise (for helpful discussion, see Stanley *et al.* [2019]). The examples discussed in this and the next section exhibit these principles of functional variability.

The posterior parietal (PPC) and prefrontal (PFC) cortices are important for a variety of functions—they are sometimes known as the 'frontoparietal control network' (Westphal *et al.* [2017]). In a simple sensorimotor task, for instance, a subject must map a sensory input, such as correlated movement in a dot-motion stimulus, to a particular motor response, such as a button press. Siegel *et al.* ([2011]) used EEG to measure patterns in LFP activity in humans when they made decisions about motion direction at close-to-threshold levels of discriminability. They showed that during the initial stimulus and the actual movement, gamma synchrony increases and beta synchrony decreases in extrastriate

cortex and in motor cortex, respectively. In the interval between the presentation of the stimulus and the required motor response, beta synchrony increases in the frontoparietal control network. Vitally, the strength of frontoparietal beta phase synchrony predicts whether the subject is correct in the task, but does not predict the particular decision that is made (Siegel *et al.* [2011]).

Siegel et al.'s interpretation of the results is that the beta phase synchrony in frontoparietal regions links up sensory information and motor commands in the relevant regions (which exhibit beta synchrony during the delay, but at a lesser degree than frontoparietal regions), and maintains the sensory evidence necessary for the decision during the delay period. The description here, as the authors themselves stress, is of a dynamic pattern of activity across these brain areas, which is underlain by both feed-forward and feedback connections between the areas (Siegel et al. [2012]). But the results still support a functional decomposition of the network in the decision task. The gamma frequency activation at the two regions during stimulus onset and motor response support the view that MT underlies the motion perception, and premotor cortex the movement preparation. The frontoparietal network exhibits particular frequency- and time-specific engagement in this process, which leads to the hypothesis about its evidence-maintaining role—a different role than is posted for MT or motor cortex.

It is also important to note that as task demands change, more areas can interact with this system, but it ways that are still functionally decomposable. If the task explicitly requires spatial attention, further phase synchrony is implemented in the gamma band between an overlapping, but non-identical network including the frontal eye fields, the PPC, and visual cortical areas (Siegel *et al.* [2012]). While the PPC is involved in both aspects of the task, distinct frequencies modulate these distinct contributions, and different overall interactants take part, as per the flexibility discussed above. And as discussed in the last section, these changes can be produced reliably as task demands change. So, even dynamic interaction between areas, and dynamic changes in what brain areas are included in a network, are compatible with decomposition on the multiplexing—contextualist picture.

5.3 Network dependence

The frontoparietal network is also important for episodic memory, now in conjunction with the medial temporal lobe (MTL), which includes the hippocampus, as well as the entorhinal cortex and parahippocampal gyrus (Watrous and Ekstrom [2014]). Episodic memory, however, involves different coordinating LFP frequencies from sensorimotor decision making, and intriguingly these frequencies can be used to dissociate between distinct memory tasks. Ekstrom *et al.* ([2013]) trained human subjects in a virtual reality environment—their task was to drive a virtual

'friend' around from store to store in a virtual neighbourhood, picking them up and dropping them off. After experience with the environment, subjects were then asked one of two distinct types of question about relationships between stores: which ones were closer to each other (spatial question), or which ones were visited in which order (temporal question). In two separate studies, one with intracranial EEG and one with fMRI, Watrous and colleagues measured phase synchrony and 'functional connectivity'—the coactivation of areas in particular conditions.

Phase synchrony was shown between the MTL, PPC, and PFC in both spatial and temporal recall, but at different frequencies—delta (1–4 Hz) for spatial memory, and high theta (7–10 Hz) for temporal. So, even for very closely related tasks involving similar brain areas, the frequency of background LFP phase synchrony can discriminate between behavioural conditions. This supports the kind of analysis of dynamical modulation given in the last section. However, these were not the only areas to show phase synchrony in these conditions, and it is well established that successful episodic memory recall involves greater interconnection across a wide brain network compared to non-successful recall (Westphal *et al.* [2017]). To explore this connectivity further, Watrous and colleagues analysed the data in terms of network measurements. In particular, they looked at the degree of each area, which is the number of other areas with which each area was synchronized or co-active, and at betweenness

centrality, which measures the number of shortest paths between other areas that connect through the area of interest.

Recall Rathkopf's variety of holist inference from Section 3. Rathkopf thinks that one is forced into the non-mechanistic (non-decompositional) side of the explanatory division-of-labour by complexity, and that the explanatory relevance of pattern variables is evidence that that divide has been crossed. Degree-distribution across the network and betweenness centrality are precisely the kind of global pattern variables Rathkopf is interested in, and thus his argument predicts that decomposition should be unfruitful in these contexts. This does not, however, match the way that the research proceeded, or the kinds of explanatory hypotheses the researchers offer.

Watrous and colleagues' functional connectivity measurements showed that both synchrony and connectivity, measured by the overall degree measurement in the network, were increased during correct as compared to incorrect trials, as were both node degree and betweenness centrality in the 'hub' regions of PFC, PPC, and MTL. The MTL had similar degree and betweenness measures in each type of task, and these were the highest measures in the network. Intriguingly, however, there were also differences between the tasks. Prefrontal areas showed greater values for network measures in the temporal context, and parietal in the spatial context. Moreover, while there was significant overlap between the networks, it was not complete. Both tasks activated a wide frontal network

in the superior and medial frontal gyri, but different parietal areas—the precuneus versus the inferior parietal lobule—were involved in the network in spatial versus temporal tasks, respectively. This matches prior MRI results implicating these sub-areas of the PPC in these types of task (Watrous *et al.* [2013]).

While Rathkopf is right that the pattern variables are playing an important explanatory role here, the case is not one in which decomposition fails. Two distinct frequencies of cross-frequency coupling are involved in the two distinct tasks, and while there is overlap in the networks, that overlap is not complete. Moreover, and importantly, the researchers take the MTL's centrality for each task to indicate that it plays a distinct functional role from other areas. Watrous and Ekstrom ([2014]) propose that the MTL is in charge of coordinating a widespread low-frequency oscillation that is appropriate in the context, and that cross-frequency coupling in cortical areas implements item-specific reactivation.

To summarize: Rathkopf's version of holist inference posits a distinction between system behaviours that are explained by pattern variables and those that are decomposable. The studies reviewed here employ the kinds of variables that Rathkopf cites to help understand the distinction between types of episodic memory. However, these measurements, which establish the centrality of the MTL in both contexts, and the respective centrality of the prefrontal and parietal for specific tasks, inspire an attempt at functional decomposition rather than a denial

of it. And it is worth noting that this extends to other settings where centrality is important. Siegel *et al.* ([2012]), in summarizing the studies from the last section, suggest that the PPC is a central hub in both attentional and decision networks, although again at distinct frequencies. But as discussed above, these networks are not non-decomposable.

I thus conclude that the contextualist approach is compatible with all three properties of complexity that motivate holist inferences. If this is true, then those inferences are invalid. I now proceed to analyse how this result intersects with extant debates in the literature.

6 Philosophical Upshot

6.1 The scope and limits of mechanistic explanation

As I noted in the introduction, many mechanists and reductionists are not atomists. Mechanists have argued that mechanism identity is explanandum relative (Craver [2007]), that mechanism description should be sensitive to explanatory context (Craver [2001]), and for a variety of ways in which dynamical and network-based models of systems can interact productively with mechanistic research. Bechtel and Abrahamsen ([2009], [2010]; Bechtel [2016]) in particular have argued that understanding dynamics, 'recomposing' the mechanism, and situating the mechanism in the environment (Roe and Baumgartner [2017]) are all important for

understanding how a mechanism produces a phenomenon. ¹⁰ In discussion of reductionism, Wimsatt ([2006]) has suggested that inter-level correspondences within mechanisms will be heavily context-dependent. Bechtel and Abrahamsen ([2008]) suggest that even within a reductionist framework, one must occasionally resort to higher-level descriptions to understand the organization one hopes to reduce. Bechtel and Levy ([2016]; Bechtel [2017]) have recently sought a substantive notion of mechanism that avoids implications about rigid organization. According to their 'mechanism 2.0', distinct mechanisms may overlap, mechanisms may gain or lose parts over time, and mechanisms may exist only transiently.

Why, then, has failure of atomism been taken by holists to equate to failure of mechanism? I suggest that the likely reason is the continuing influence of the more-or-less notion of decomposability inspired by taking atomism as a regulative ideal (see Section 2). If mechanists and reductionists are committed to atomism being 'close enough' to accurate, then widespread complexity really should undermine mechanism and reductionism, and loosening these views to accommodate complexity weakens them to such an extent that they lack distinctive content.

¹⁰ The literature on the relationship between dynamical, network, and mechanistic explanations is waxing voluminous, and I can't do full justice to it here. On dynamics, see (Kaplan and Craver [2011]; Levy [2013]; Ross [2015]). On networks, see (Bechtel [2015], [2017]; Levy and Bechtel [2013]; Craver [2016]; Green *et al.* [2017]; Matthiessen [2017]).

Nowhere is this kind of sentiment more clear than in current debates over the 'scope and limits' of mechanistic explanation (Halina [2017]). Some, particularly Weiskopf ([2016]) and Moss ([2012]), argue that mechanistic theorizing is deeply beholden to the 'machine analogy', according to which neatly separable parts engage in simple interactions like those in everyday machines. Along these lines, Silberstein and Chemero ([2013]) suggest that one can save mechanism in complex systems only by abandoning the idea that decomposition is the key commitment of a mechanistic approach. Rathkopf ([2018)] argues that the mechanistic commitment to decomposition is what underlies the division-of-labour he proposes. Giving up or weakening decomposition would, it seems, bring together mechanistic and network explanations in an unprincipled manner, by fiat.

Some mechanists respond to these kinds of concerns by eschewing the machine analogy (Craver [2007]; Bechtel and Abrahamsen [2012]). Others, however, seem to have been at least tempted by the move to abandon the need for decomposition as a distinctively mechanistic commitment (Stinson [2016]; Zednik [unpublished]). I agree with the critics of mechanism that decomposition should be seen as the core commitment of mechanistic and reductionist approaches, and thus that abandoning this commitment would deflate the view. But I of course disagree that this distinctive commitment must be read in atomist terms. What mechanists and reductionists need is a well-fleshed-out alternative to

atomism that retains decomposition as a central commitment while also explicitly embracing complexity. I have offered contextualism as a general alternative principle, and the multiplexing framework as an exemplar. The multiplexing framework shows how neuroscientists have sought principles that decompose functions even in cases where dynamic interaction and network-mediation are abundantly in evidence. I have argued that contextualism captures the kinds of decompositions at work in these explanations.

But there is more here for mechanists—in particular, the multiplexing field is now pursuing the kind of multilevel decompositions and interventions that many take to be the hallmark of mechanistic explanatory schemata. While the studies I've cited are generally correlational, neuroscientists have begun to (i) pursue lower-level explanations of how LFPs, cross-frequency coupling, and phase synchrony are enacted, and (ii) intervene at multiple levels to show causal relevance of multiplexing for phenomena of interest. It has been shown that different LFP waveforms exhibit differences both in the interlayer patterns of connectivity between cortical cells, as well as different neuromodulatory patterns (Womelsdorf *et al.* [2014]). Hypotheses have also been stated to address the links by which phase coupling between areas occurs. Geib *et al.* ([2015]), for instance, propose that the caudate nucleus provides a mediating connection between MTL and frontal areas that underlies functional

connectivity, and the thalamus seems to be important for connecting up cortical areas more broadly (Hwang *et al.* [2017]).

Lastly, it is possible to intervene on LFPs in several ways to test their causal contribution to phenomena (Woodward [2010]). Spellman *et al.* ([2015]) showed that optogenetic interruption of hippocampal-prefrontal phase synchrony inhibits memory encoding in rats. Polania *et al.* ([2015]), applied transcranial alternating current stimulation to impose antiphase oscillations between the medial frontal and parietal cortices, which interrupted value-based choices in human participants. The multiplexing framework thus supports the kinds of multilevel, interventionist investigation typical of mechanistic research, in addition to successfully decomposing function, albeit in a non-atomist way.

6.2 The context objection to reduction

The 'context objection' to reductionism states that context-sensitivity of function within a mechanism is incompatible with reductionism. Delehanty ([2005]) has given an elegant response to this objection, based on the idea that we can expand the mechanism being analysed to include whatever contextual influences are relevant to explaining the system's behaviour. I here argue that contextualism has more resources than Delehanty takes advantage of, and thus that contextualism can avoid further objections by Hütteman and Love ([2011]).

Suppose we have an explanandum E and a purported mechanism M responsible for it. However, in the course of investigation we find that M is systematically affected by some part of the system outside of it in E contexts, and that without this influence we cannot explain why it produces E. Delehanty's solution is that we expand the mechanism for E to include M+n, where n is the set of interactions that, in addition to M, are cited in the explanation. So long as the explanation invoking M+n still consists in a description of distinct parts and their interactions, the explanation is still reductive. As Hütteman and Love helpfully summarize, the idea is simply to reduce the context as well.

Hütteman and Love give two considerations against the contextualist argument. First, they suggest that the move to expand M is illicit, since rather than providing a reductive explanation for the behaviour of the mechanism we set out to analyse, we have shifted to explaining a wholly different mechanism, namely, M + n. Second, they suggest that even if the expansion of M is granted, the contextualist may face a pernicious regress, where the mechanism must be continually expanded to account for temporally prior causal influence. I claim that contextualism has the resources to resist the first consideration, and that the second clarifies an explanandum for contextualism rather than being an objection against it.

The problem with the debate about context and reduction is that it is still framed in terms of atomsim, and Delehanty's approach does not do enough to abandon that framing. According to Delehanty, we start out

with a pre-existing mechanism and then 'add on' the components needed to account for its behaviour in the context. This invites the response from Hütteman and Love, because it starts from the standpoint that a given mechanism's behaviour is what's at stake, and that contextual influence is only an add-on to that mechanism. It thus seems like something is 'left out' of the new explanation—namely, the original explanandum, the behaviour of M. This is not the way the contextualist should go. Instead, the contextualist should lean on the flexibility described in the previous sections to suggest that mechanism identity *simpliciter* is dependent on context. Recall that the idea of a consistent mechanistic organization is an implication of atomism, not contexutalism. According to the full-throated contextualist, we aren't beholden to a previously specified mechanistic organization when context changes—we are not attempting to explain M +n but N, the mechanism that is relevant in the explanatory context. Hence, there is no residual explanandum left unexplained.

One might worry that this is too permissive, that there must be stricter conditions on mechanism identity, lest we be left with a hodgepodge of single cases.¹¹ The contextualist response to this is two-fold. First, the need to index mechanism identity to context doesn't mean that there is no generalizability to particular mechanistic descriptions (see Section 4). For instance, there seem to be consistent, discoverable organizations to the mechanisms underlying sensorimotor decision making, attention, episodic

¹¹ Thanks to Alan Love for pressing me to consider this objection.

memory, and adaptation after feedback (Section 5). The point of the contextualist position is that these generalizations must be discovered. We don't know that a given mechanistic organization will remain constant across contexts just because we've isolated it in one context.

The regress worry is a real one. But as Hütteman and Love admit (Love, personal communication), it is an empirical worry. It might be the case that some systems are so causally integrated that there is no place one can stop the potential regress—each mechanistic organization might depend so far back in a causal chain for its identity that there is no way to genuinely set it apart from the surrounding system. Stated as such, the regress objection is not an in-principle one, but a challenge to show, in a particular system, how mechanism identity can be established within a context, as separable from prior contexts.

There are two potential avenues of answering this challenge within the multiplexing framework, at least as it's developed so far. The first is that there may be dynamical properties of the system that provide 're-sets' for its behaviour. For instance, Canolty *et al.* ([2012]) suggest that periods of low beta-power in the course of their experiments might serve to interrupt a given multiplexing scheme and ready the system for the next task. Second, and more broadly, there are intriguing results suggesting that areas of the so-called 'rich club' network—high-degree nodes that are also intimately connected—oscillate at an extremely slow frequency, and change their patterns of functional connectivity between rest and tasks.

One way to read the function of this network is as a higher-order system that mediates competitions between other networks, in terms of which will become active in the given context (Senden *et al.* [2018]; Burnston and Haueis [unpublished]). Which sub-network is activated is reliable across tasks, and thus the rich club mediation may contribute to reliable, but context-sensitive, organization. The two mechanisms—re-sets and competition—both suggest natural points at which mechanism organization shifts. Once we figure out how the context shift is mediated, we can then focus on decomposition within that context. So, while both of these proposals need fleshing out, they illustrate the kind of explanation needed in order to avoid the regress worry.

7 Conclusion

I hope to have established that certain common inferences in the philosophy of science are mistaken, and that dialectics surrounding mechanistic and reductive explanation in biology and neuroscience have become muddied as a result. The multiplexing framework I've articulated will not be applicable to all complex systems we might want to analyse—if complexity and decomposition are indeed compatible, we can expect that different systems will have very different ways of implementing decomposition. Hopefully, this account might serve as an exemplar for analysis of other systems, though the specifics will differ. If my arguments are on track, then philosophers of science should be spending their time

trying to articulate the organizational principles that implement decomposition in context-sensitive, dynamically interactive networks, rather than using this complexity to deny decomposition altogether.

Acknowledgements

Many people contributed to the lengthy developmental phase of this article. I'm grateful to Krista Perks for first introducing me to the multiplexing literature years ago. Very early forms of this material were presented to thoughtful audiences at Georgia State University, the British Society for the Philosophy of Science conference in Cardiff, and the International Society for History, Philosophy, and Social Science of Biology meeting in São Paulo, all in 2016. In 2018, a pre-submission version of the article was given in the "Neural Mechanisms Online" talk series. I'm grateful to Fabrizio Calzavarini and Marco Viola for the invitation and for giving comments, and to Phillip Haueis, Joe McCaffrey, Charles Rathkopf, Ernesto Restrepo, and Michael Silberstein for subsequent stimulating discussion. Since then, I've received very helpful written comments from Bryce Gessell, Colin Klein, Alan Love, and Jason Winning, and benefitted from discussion with Colin Allen, William Bechtel, Ingo Brigandt, Mazviita Chirimuuta, and Natasha Navejar. The penultimate version was shared with my colleagues at the Center for Philosophy of Science at the University of Pittsburgh in spring of 2019. Janella Baxter, Tomasz Bigaj, Chunghyoung Lee, Edouard Machery,

Antonella Tramacere, Naftali Weinberger, Barb Wimsatt, Bill Wimsatt, and Liying Zhang provided extremely helpful discussion and comments during the finishing stage. Julie Crow provided vital help in preparing the final version of the manuscript. Lastly, three anonymous reviewers for BJPS provided careful and constructive comments on the article, and responding to them has made it much better.

Department of Philosophy
Tulane Brain Institute
Tulane University
New Orleans, LA, USA
Dburnsto@tulane.edu

References

- Akam, T. and Kullmann, D. M. [2014]: 'Oscillatory Multiplexing of Population Codes for Selective Communication in the Mammalian Brain', *Nature Reviews Neuroscience*, **15**, pp. 111–22.
- Anderson, M. L. [2010a]: 'Neural Reuse: A Fundamental Organizational Principle of the Brain', *Behavioural and Brain Sciences*, **33**, pp. 245–66.
- Anderson, M. L. [2010b]: 'Cortex in Context: Response to Commentaries on Neural Reuse', *Behavioural and Brain Sciences*, **33**, pp. 294–313.
- Anderson, M. L. [2014]: *After Phrenology: Neural Reuse and the Interactive Brain*, Cambridge, MA: MIT Press.
- Bechtel, W. [2012]: 'Referring to Localized Cognitive Operations in Parts of Dynamically Active Brains', in A. Raftopoulos and P. Machamer (eds), Perception, Realism, and the Problem of Reference, Cambridge: Cambridge University Press, pp. 262–84.
- Bechtel, W. [2015]: 'Can Mechanistic Explanation Be Reconciled with Scale-Free Constitution and Dynamics?', *Studies in the History and Philosophy of Biological and Biomedical Sciences*, **53**, pp. 84–93.
- Bechtel, W. [2016]: 'Mechanists Must Be Holists Too! Perspectives from Circadian Biology', *Journal of the History of Biology*, **49**, pp. 705–31.
- Bechtel, W. [2017]: 'Analysing Network Models to Make Discoveries about Biological Mechanisms', *British Journal for the Philosophy of Science*, **70**, pp. 459–84.

- Bechtel, W. and Abrahamsen, A. [2005]: 'Explanation: A Mechanist Alternative'. *Studies in History and Philosophy of Biological and Biomedical Sciences*, **36**, pp. 421–41.
- Bechtel, W. and Abrahamsen, A. [2008]: 'From Reduction Back to Higher Levels', in B. C. Love, K. McRae and M. Sloutsky (*eds*), *Proceedings of the Annual Meeting of the Cognitive Science Society*, Austin, TX: Cognitive Science Society, pp. 559–564.
- Bechtel, W. and Abrahamsen, A. [2009]: 'Decomposing, Recomposing, and Situating Circadian Mechanisms: Three Tasks in Developing Mechanistic Explanations', in H. Leitgeb and A. Hieke (eds), Reduction and Elimination in Philosophy of Mind and Philosophy of Neuroscience, Frankfurt, DE: Ontos, pp. 173–86.
- Bechtel, W. and Abrahamsen, A. [2010]: 'Dynamic Mechanistic Explanation: Computational Modeling of Circadian Rhythms as an Exemplar for Cognitive Science', *Studies in History and Philosophy of Science*, **41**, pp. 321–33.
- Bechtel, W. and Abrahamsen, A. A. [2012]: 'Thinking Dynamically about Biological Mechanisms: Networks of Coupled Oscillators', *Foundations of Science*, **18**, pp. 707–23.
- Bechtel, W. and Hamilton, A. [2007]: 'Reduction, Integration, and the Unity of Science: Natural, Behavioural, and Social Sciences and the Humanities', in T. A. F. Kuipers (*ed.*), *General Philosophy of Science: Focal Issues*, Amsterdam: Elsevier, pp. 377–430.
- Bechtel, W. and Richardson, R. C. [1993]: *Discovering Complexity*, Princeton, NJ: Princeton University Press.
- Bergeron, V. [2007]: 'Anatomical and Functional Modularity in Cognitive Science: Shifting the Focus', *Philosophical Psychology*, **20**, pp. 175–95.
- Bickle, J. [2006]: 'Reducing Mind to Molecular Pathways: Explicating the Reductionism Implicit in Current Cellular and Molecular Neuroscience', *Synthese*, **151**, pp. 411–34.
- Boogerd, F. C., Bruggeman, F. J., Richardson, R. C., Stephan, A. and Westerhoff, H. V. [2005]: 'Emergence and Its Place in Nature: A Case Study of Biochemical Networks', *Synthese*, **145**, pp. 131–64.
- Brigandt, I. [2017]: 'Bodily Parts in the Structure–Function Dialectic', in S. Lidgard and L. K. Nyhart (eds), Biological Individuality: Integrating Scientific, Philosophical, and Historical Perspectives, Chicago, IL: University of Chicago Press, pp. 249–74.
- Brigandt, I., Green, S. and O'Malley, M. A. [2018]: 'Systems Biology and Mechanistic Explanation', in S. Glennan and P. Illari (eds), The Routledge

- *Handbook of Mechanisms and Mechanical Philosophy*, New York: Routledge, pp. 362–74.
- Burnston, D. C. [2016a]: 'Computational Neuroscience and Localized Neural Function', *Synthese*, **193**, pp. 3741–62.
- Burnston, D. C. [2016b]: 'A Contextualist Approach to Functional Localization in the Brain', *Biology and Philosophy*, **31**, pp. 527–50.
- Burnston, D. C. and Haueis, P. [unpublished]: 'Evolving Concepts of "Hierarchy" in Systems Neuroscience'.
- Canolty, R. T., Ganguly, K. and Carmena, J. M. [2012]: 'Task-Dependent Changes in Cross-Level Coupling between Single Neurons and Oscillatory Activity in Multiscale Networks', *PLoS Computational Biology*, **8**, p. e1002809.
- Canolty, R. T. and Knight, R. T. [2010]: 'The Functional Role of Cross-Frequency Coupling', *Trends in Cognitive Sciences*, **14**, pp. 506–15.
- Chemero, A. and Silberstein, M. [2008]: 'After the Philosophy of Mind: Replacing Scholasticism with Science', *Philosophy of Science*, **75**, pp. 1–27.
- Colombo, M. and Weinberger, N. [2018]: 'Discovering Brain Mechanisms Using Network Analysis and Causal Modeling', *Minds and Machines*, **28**, pp. 265–86.
- Cushman, F., Young, L. and Greene, J. D. [2010]: 'Our Multi-system Moral Psychology: Towards a Consensus View', in J. M. Doris and the Moral Psychology Research Group (*eds*), *The Moral Psychology Handbook*, Oxford: Oxford University Press, pp. 47–71.
- Craver, C. F. [2001]: 'Role Functions, Mechanisms, and Hierarchy', *Philosophy of Science*, **68**, pp. 53–74.
- Craver, C. F. [2007]: *Explaining the Brain*, New York: Oxford University Press.
- Craver, C. F. [2016]: 'The Explanatory Power of Network Models', *Philosophy of Science*, **83**, pp. 698–709.
- DeAngelis, G. C., Cumming, B. G. and Newsome, W. T. [1998]: 'Cortical Area MT and the Perception of Stereoscopic Depth', *Nature*, **394**, pp. 677–80.
- Delehanty, M. [2005]: 'Emergent Properties and the Context Objection to Reduction', *Biology and Philosophy*, **20**, pp. 715–34.
- Fusi, S., Miller, E. K. and Rigotti, M. [2016]: 'Why Neurons Mix: High Dimensionality for Higher Cognition', *Current Opinion in Neurobiology*, **37**, pp. 66–74.

- Geib, B. R., Stanley, M. L., Wing, E. A., Laurienti, P. J. and Cabeza, R. [2015]: 'Hippocampal Contributions to the Large-Scale Episodic Memory Network Predict Vivid Visual Memories', *Cerebral Cortex*, **27**, pp. 680–93.
- Gillett, C. [2016]: *Reduction and Emergence in Science and Philosophy*, Cambridge: Cambridge University Press.
- Glennan, S. [2002]: 'Rethinking Mechanistic Explanation', *Philosophy of Science*, **69**, pp. S342–53.
- Grafton, S. T. and de C Hamilton, A. F. [2007]: 'Evidence for a Distributed Hierarchy of Action Representation in the Brain', *Human Movement Science*, **26**, pp. 590–616.
- Graziano, M. [2006]: 'The Organization of Behavioural Repertoire in Motor Cortex', *Annual Review of Neuroscience*, **29**, pp. 105–34.
- Green, S., Şerban, M., Scholl, R., Jones, N., Brigandt, I. and Bechtel, W. [2017]: 'Network Analyses in Systems Biology: New Strategies for Dealing with Biological Complexity', *Synthese*, **195**, pp. 1751–77.
- Halina, M. [2017]: 'Mechanistic Explanation and Its Limits', in S. Glennan and P. Illari (eds), The Routledge Handbook of Mechanisms and Mechanical Philosophy, London: Routledge, pp. 213–24.
- Haueis, P. [2018]: 'Beyond Cognitive Myopia: A Patchwork Approach to the Concept of Neural Function', *Synthese*, **195**, pp. 5373–402.
- Huneman, P. [2010]: 'Topological Explanations and Robustness in Biological Sciences', *Synthese*, **177**, pp. 213–45.
- Hüttemann, A. and Love, A. C. [2011]: 'Aspects of Reductive Explanation in Biological Science: Intrinsicality, Fundamentality, and Temporality', *British Journal for the Philosophy of Science*, **62**, pp. 519–49.
- Hwang, K., Bertolero, M. A., Liu, W. B. and D'Esposito, M. [2017]: 'The Human Thalamus Is an Integrative Hub for Functional Brain Networks', *Journal of Neuroscience*, **37**, pp. 5594–607.
- Jones, N. [2014]: 'Bowtie Structures, Pathway Diagrams, and Topological Explanation', *Erkenntnis*, **79**, pp. 1135–55.
- Kaiser, M. I. [2011]: 'The Limits of Reductionism in the Life Sciences', *History and Philosophy of the Life Sciences*, **33**, pp. 453–76.
- Kaiser, M. I. [2015]: *Reductive Explanation in the Biological Sciences*, Dordrecht: Springer.
- Kanwisher, N. [2010]: 'Functional Specificity in the Human Brain: A Window into the Functional Architecture of the Mind', *Proceedings of the National Academy of Sciences USA*, **107**, pp. 11163–70.

- Kaplan, D. M. and Craver, C. F. [2011]: 'The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective', *Philosophy of Science*, **78**, pp. 601–27.
- Khanna, P. and Carmena, J. M. [2015]: 'Neural Oscillations: Beta Band Activity across Motor Networks', *Current Opinion in Neurobiology*, **32**, pp. 60–7.
- Khalidi, M. A. [2017]: 'Crosscutting Psycho-neural Taxonomies: The Case of Episodic Memory', *Philosophical Explorations*, **20**, pp. 191–208.
- Klein, C. [2012]: 'Cognitive Ontology and Region- versus Network-Oriented Analyses', *Philosophy of Science*, **79**, pp. 952–60.
- Klein, C. [2018]: 'Mechanisms, Resources, and Background Conditions', *Biology and Philosophy*, **33**, available at <doi.org/10.1007/s10539-018-9646-y>.
- Levy, A. [2013]: 'What Was Hodgkin and Huxley's Achievement?', *British Journal for the Philosophy of Science*, **65**, pp. 469–92.
- Levy, A. and Bechtel, W. [2013]: 'Abstraction and the Organization of Mechanisms', *Philosophy of Science*, **80**, pp. 241–61.
- Levy, A. and Bechtel, W. [forthcoming]: 'Towards Mechanism 2.0: Expanding the Scope of Mechanistic Explanation'.
- Lisman, J. E. and Jensen, O. [2013]: 'The Theta-Gamma Neural Code', *Neuron*, 77, pp. 1002–16.
- Machamer, P., Darden, L. and Craver, C. F. [2000]: 'Thinking about Mechanisms', *Philosophy of Science*, **67**, pp. 1–25.
- Madary, M. [2013]: 'Placing Area MT in Context', *Journal of Consciousness Studies*, **20**, pp. 93–104.
- Matthiessen, D. [2017]: 'Mechanistic Explanation in Systems Biology: Cellular Networks', *British Journal for the Philosophy of Science*, **68**, pp. 1–25.
- McCaffrey, J. B. [2015]: 'The Brain's Heterogeneous Functional Landscape', *Philosophy of Science*, **82**, pp. 1010–22.
- McIntosh A. R. [2004]: 'Contexts and Catalysts: A Resolution of the Localization and Integration of Function in the Brain', *Neuroinformatics*, **2**, pp. 175–82.
- Moss, L. [2012]: 'Is the Philosophy of Mechanism Philosophy Enough?', *Studies in the History and Philosophy of Biological and Biomedical Sciences*, **43**, pp. 164–72.
- Mundale, J. [2002]: 'Concepts of Localization: Balkanization in the Brain', *Brain and Mind*, **3**, pp. 313–30.

- Palanca, B. J. A. and DeAngelis, G. C. [2003]: 'Macaque Middle Temporal Neurons Signal Depth in the Absence of Motion', *The Journal of Neuroscience*, **23**, pp. 7647–58.
- Panzeri, S., Brunel, N., Logothetis, N. K. and Kayser, C. [2010]: 'Sensory Neural Codes Using Multiplexed Temporal Scales', *Trends in Neurosciences*, **33**, pp. 111–20.
- Panzeri, S., Macke, J. H., Gross, J. and Kayser, C. [2015]: 'Neural Population Coding: Combining Insights from Microscopic and Mass Signals', *Trends in Cognitive Sciences*, **19**, pp. 162–72.
- Pessoa, L. [2008]: 'On the Relationship between Emotion and Cognition', *Nature Reviews Neuroscience*, **9**, pp. 148–58.
- Piccinini, G. and Craver, C. [2011]: 'Integrating Psychology and Neuroscience: Functional Analyses as Mechanism Sketches', *Synthese*, **183**, pp. 283–311.
- Polania, R., Moisa, M., Opitz, A., Grueschow, M. and Ruff, C. C. [2015]: 'The Precision of Value-Based Choices Depends Causally on Fronto-parietal Phase Coupling', *Nature Communications*, **6**, p. 8090.
- Rathkopf, C. A. [2013]: 'Localization and Intrinsic Function', *Philosophy of Science*, **80**, pp. 1–21.
- Rathkopf, C. [2018]: 'Network Representation and Complex Systems', *Synthese*, **195**, pp. 55–78.
- Rice, C. [2015]: 'Moving beyond Causes: Optimality Models and Scientific Explanation', *Noûs*, **49**, pp. 589–615.
- Roe, S. M. and Baumgaertner, B. [2017]: 'Extended Mechanistic Explanations: Expanding the Current Mechanistic Conception to Include More Complex Biological Systems', *Journal for General Philosophy of Science*, **48**, pp. 517–34.
- Ross, L. N. [2015]: 'Dynamical Models and Explanation in Neuroscience', *Philosophy of Science*, **82**, pp. 32–54.
- Salazar, R., Dotson, N., Bressler, S. and Gray, C. [2012]: 'Content-Specific Fronto-parietal Synchronization during Visual Working Memory', *Science*, **338**, pp. 1097–100.
- Sanada, T. M., Nguyenkim, J. D. and DeAngelis, G. C. [2012]: 'Representation of 3-D Surface Orientation by Velocity and Disparity Gradient Cues in Area MT', *Journal of Neurophysiology*, **107**, pp. 2109–22.
- Sarkar, S. [1992]: 'Models of Reduction and Categories of Reductionism', *Synthese*, **91**, pp. 167–94.

- Scarantino, A. [2012]: 'Functional Specialization Does Not Require a One-to-One Mapping between Brain Regions and Emotions', *Behavioural and Brain Sciences*, **35**, pp. 161–2.
- Schedlbauer, A. M., Copara, M. S., Watrous, A. J. and Ekstrom, A. D. [2014]: 'Multiple Interacting Brain Areas Underlie Successful Spatiotemporal Memory Retrieval in Humans', *Scientific Reports*, 4, p. 6431.
- Senden, M., Reuter, N., van den Heuvel, M. P., Goebel, R., Deco, G. and Gilson, M. [2018]: 'Task-Related Effective Connectivity Reveals That the Cortical Rich Club Gates Cortex-Wide Communication', *Human Brain Mapping*, **39**, pp. 1246–62.
- Siegel, M., Warden, M. R. and Miller, E. K. [2009]: 'Phase-Dependent Neuronal Coding of Objects in Short-Term Memory', *Proceedings of the National Academy of Sciences*, **106**, pp. 21341–46.
- Siegel, M., Donner, T. H. and Engel, A. K. [2012]: 'Spectral Fingerprints of Large-Scale Neuronal Interactions', *Nature Reviews Neuroscience*, **13**, pp. 121–34.
- Siegel, M., Engel, A. K. and Donner, T. H. [2011]: 'Cortical Network Dynamics of Perceptual Decision-Making in the Human Brain', *Frontiers in Human Neuroscience*, **5**, available at <doi.org/0.3389/fnhum.2011.00021>.
- Silberstein, M. and Chemero, A. [2013]: 'Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences', *Philosophy of Science*, **80**, pp. 958–70.
- Simon, H. A. [1962]: 'The Architecture of Complexity', *Proceedings of the American Philosophical Society*, **106**, pp. 467–82.
- Spellman, T., Rigotti, M., Ahmari, S. E., Fusi, S., Gogos, J. A. and Gordon, J. A. [2015]: 'Hippocampal–Prefrontal Input Supports Spatial Encoding in Working Memory', *Nature*, **522**, pp. 309–14.
- Sporns, O. [2011]: *Networks of the Brain*, Cambridge, MA: MIT Press.
- Stanley, M. L., Gessell, B. and De Brigard, F. [2019]: 'Network Modularity as a Foundation for Neural Reuse', *Philosophy of Science*, **86**, pp. 23–46.
- Stinson, C. [2016]: 'Mechanisms in Psychology: Ripping Nature at Its Seams', *Synthese*, **193**, pp. 1585–614.
- Theurer, K. L. [2013]: 'Compositional Explanatory Relations and Mechanistic Reduction', *Minds and Machines*, **23**, pp. 287–307.
- Uka, T. and DeAngelis, G. C. [2006]: 'Linking Neural Representation to Function in Stereoscopic Depth Perception: Roles of the Middle Temporal Area in Coarse Versus Fine Disparity Discrimination', *The Journal of Neuroscience*, **26**, pp. 6791–802.

- van Eck, D. and Mennes, J. [2018]: 'Mechanism Discovery and Design Explanation: Where Role Function Meets Biological Advantage Function', *Journal for General Philosophy of Science*, **49**, pp. 413–34.
- DeYoe, E. A. and Van Essen, D. C. [1988]: 'Concurrent Processing Streams in Monkey Visual Cortex', *Trends in Neurosciences*, **11**, pp. 219–26.
- Viola, M. and Zanin, E. [2017]: 'The Standard Ontological Framework of Cognitive Neuroscience: Some Lessons from Broca's Area', *Philosophical Psychology*, **30**, pp. 945–69.
- Voytek, B., Canolty, R. T., Shestyuk, A., Crone, N. E., Parvizi, J. and Knight, R. T. [2010]: 'Shifts in Gamma Phase-Amplitude Coupling Frequency from Theta to Alpha over Posterior Cortex during Visual Tasks', *Frontiers in Human Neuroscience*, 4, p. 191.
- Watrous, A. J. and Ekstrom, A. D. [2014]: 'The Spectro-contextual Encoding and Retrieval Theory of Episodic Memory', *Frontiers in Human Neuroscience*, **8**, available at <doi.org/10.3389/fnhum.2014.00075>.
- Watrous, A. J., Fell, J., Ekstrom, A. D. and Axmacher, N. [2015]: 'More than Spikes: Common Oscillatory Mechanisms for Content Specific Neural Representations during Perception and Memory', *Current Opinion in Neurobiology*, **31**, pp. 33–9.
- Watrous, A. J., Tandon, N., Conner, C. R., Pieters, T. and Ekstrom, A. D. [2013]: 'Frequency-Specific Network Connectivity Increases Underlie Accurate Spatiotemporal Memory Retrieval', *Nature Neuroscience*, 16, pp. 349–56.
- Weiskopf, D. A. [2016]: 'Integrative Modeling and the Role of Neural Constraints', *Philosophy of Science*, **83**, pp. 674–85.
- Westphal, A. J., Wang, S. and Rissman, J. [2017]: 'Episodic Memory Retrieval Benefits from a Less Modular Brain Network Organization', *Journal of Neuroscience*, **37**, pp. 3523–31.
- Wimsatt, W. C. [1976]: 'Reductive Explanation: A Functional Account', in R. S. Cohen, C. A. Hooker, A. C. Michalos and J. W. Van Evra (*eds*), *PSA* 1974, Dordrecht: Reidel, pp. 671–710.
- Wimsatt, W. C. [2006]: 'Reductionism and Its Heuristics: Making Methodological Reductionism Honest', *Synthese*, **151**, pp. 445–75.
- Wimsatt, W. C. [2007]: Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality, Cambridge, MA: Harvard University Press.
- Womelsdorf, T., Valiante, T. A., Sahin, N. T., Miller, K. J. and Tiesinga, P. [2014]: 'Dynamic Circuit Motifs Underlying Rhythmic Gain Control, Gating, and Integration', *Nature Neuroscience*, **17**, pp. 1031–9.

Zednik, C. [unpublished]: 'Are Systems Neuroscience Explanations Mechanistic?', available at <philsci-archive.pitt.edu/10859/>.

