Beyond Componential Constitution in the Brain

Starburst Amacrine Cells and Enabling Constraints

Michael L. Anderson

Componential mechanism (Craver 2008) is an increasingly influential framework for understanding the norms of good explanation in neuroscience and beyond. Componential mechanism "construes explanation as a matter of decomposing systems into their parts and showing how those parts are organized together in such a way as to exhibit the explanandum phenomenon" (Craver 2008, p. 109). Although this clearly describes some instances of successful explanation, I argue here that as currently formulated the framework is too narrow to capture the full range of good mechanistic explanations in the neurosciences. The centerpiece of this essay is a case study of Starburst Amacrine Cells—a type of motion-sensitive cell in mammalian retina—for which function emerges from structure in a way that appears to violate the conditions specified by componential mechanism as currently conceived. I argue that the case of Starburst Amacrine Cells should move us to replace the notion of mechanistic componential constitution with a more general notion of enabling constraint. Introducing enabling constraints as a conceptual tool will allow us to capture and appropriately characterize a wider class of structure-function relationships in the brain and elsewhere.

Keywords
Componential constitution | Constitution | Constraint | Enabling constraint | Explanation | Functional levels | Levels | Mechanisms | Mechanistic explanation | Neuroscientific explanation | Spatial levels | Starburst amacrine cells | Structure function mapping

Author

Michael L. Anderson
michael.anderson@fandm.edu
Franklin & Marshall College
Lancaster, PA, U.S.A.

Commentator

Axel Kohler
axelkohler@web.de
Universität Osnabrück
Osnabrück, Germany

Editors

Thomas Metzinger
metzinger@uni-mainz.de
Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt
jennifer.windt@monash.edu
Monash University
Melbourne, Australia

1 Introduction

How, in the brain or any other system, does specific function arise from underlying structure? The question is a general one, and also in some sense a vague one, for it asks simultaneously about how structures shape events—generate causes—and also about what kinds of explanations one should aim for in neuroscience. Here I will focus on the second question in the hope of partially illuminating the first. One increasingly influential class of answers to this second question "construes explanation as a matter of decomposing systems into their parts and showing how those parts are organized together in such a way as to exhibit the explanandum phenomenon" (Craver 2008, p. 109; see also Craver this collection). This is an attractive idea as it is expressed, but what I hope to illustrate here is that the leading formalizations of this general idea (Craver 2008; Craver & Bechtel 2007) place overly restrictive conditions on good mechanistic explanation. In what follows, I lay out the norms of mechanistic explanation, as developed by Craver and Bechtel, and describe some cases that their model nicely cap-

tures. I then introduce the case of Starburst Amacrine Cells (SACs)—a type of motion-sensitive cell in mammalian retina. In SACs, and in the functionally coupled direction-selective ganglion cells, the function-structure relationship is hard to capture within the Craver/Bechtel mechanistic framework. I argue that we can better capture such cases by replacing the notion of mechanistic componential constitution with the more general notion of enabling constraints.

2 The requirements of mechanistic explanation

Craver (2008) sharply distinguishes between two traditions of understanding scientific explanation: reductive explanation and systems explanation. According to Craver, the first tradition accepts a version of the covering law model of explanation (Hempel 1965) whereby one explains regularities at a given level of organization by showing how these regularities (the laws describing events and their relations) can be derived from theories holding at lower levels. Put differently, one explains a phenomenon of interest by showing how it is to be *expected* based on the laws governing activity at lower levels of organization. This tradition is reductive because when such explanations are successful, one can strictly speaking *do without* the higher-level laws. However convenient they may be for understanding or predicting higher-level phenomena, the higher-level laws do not add, capture, or explain any facts that are not already contained in the lower-level laws. The lower-level laws are scientifically sufficient.

In contrast, in the systems tradition, a phenomenon of interest ψ exhibited by a system S is explained by identifying a set of component parts $\{X\}$ and showing how they are organized such that S ψ*s*. A systems explanation is similar to reductive explanation in that it too relies on the identification of levels of organization, since it requires identifying the parts of the system *S,* but, as I note below, it does not aim thereby at the reduction or explanatory absorption of one level by another. Craver & Bechtel write:

> In levels of mechanisms, an item X is at a lower level than an item S if and only if X is a component in the mechanism for some activity ψ of S. X is a component in a mechanism if and only if it is one of the entities or activities organized such that S ψ's. For that is what mechanisms are: they are entities and activities organized such that they exhibit a phenomenon. Scientists discover lower levels by decomposing the behavior of a mechanism into the behaviors of its component parts, decomposing the behaviors of the parts into the behaviors of their parts, and so on. (2007, pp. 548–549)[1]

As already noted, S is the system that ψs, or that exhibits phenomenon. It is, for instance, the car (S) that accelerates (ψ), and to explain car acceleration will require identifying the components $\{X\}$ that matter to S ψ-ing. To identify these components and their organization is to explicate the mechanism M that accounts for S ψ-ing. The target of mechanistic explanations of this sort is ψ: "mechanistic explanations are framed by the explanandum phenomenon" (Craver 2008, p. 121) and "[t]he explanandum of a mechanistic explanation is a phenomenon, typically some behavior of a mechanism as a whole" (Craver 2008, p. 139).

In mechanistic explanation, a given X is a component of the mechanism M if and only if it is one of the entities organized such that S exhibits some phenomenon ψ. So the engine, the accelerator, and the gas tank, but not the mudflaps or the windshield wipers are components of M that explain the car accelerating, even though these are *all* parts of the car S. In an

1 There is a terminological issue that needs to be raised at the outset to avoid confusion. Craver & Bechtel (2007; Craver 2008) usually, but not always, use S to refer to a *mechanism*. In contrast, I will always use S to refer to the *system* or entity exhibiting the explanandum phenomenon ψ, and I introduce the symbol M to refer to the responsible mechanism. I do this because M and S are clearly not identical. Moreover, they *are* (or at least appear to me) to be distinguished in this passage, at least on one reading. I think it is unfortunate that neither Craver nor Bechtel formally and consistently distinguish the system S and the mechanism M in their analysis, for reasons that will become clear at the end of this section. Here I'll attempt to faithfully capture the essence of the Craver–Bechtel mechanistic framework, were it to have included this important distinction.

ideal explanation, the mechanism defined by the parts $\{X\}$ will contain *all* and *only* the components relevant to S ψ-ing (see Craver 2008 for a discussion of constitutive relevance in this context). To identify the parts of M is thus to specify *both* a hierarchical and a functional relationship between M and its parts, and between M and S.

But although mechanistic explanation involves essential reference to hierarchical relationships between levels of organization, it is not thereby a species of *reductive* explanation because in a successful systems explanation nothing is rendered inessential or redundant. The phenomenon ψ is neither *derived* nor *derivable* from laws governing the parts of M; rather, the parts $\{X\}$ and their relationships simply *are* M, and together explain why S ψs. The explanatory relationship is not rational derivation, but functional composition: M is physically and functionally *constituted* by its parts, and S ψs in virtue of that constitution.

> Mechanistic explanations are constitutive or componential explanations: they explain the behavior of the mechanism as a whole in terms of the organized activities and interactions of its components. Components are the entities in a mechanism—what are commonly called 'parts'. (Craver 2008, p. 128)[2]

Given all this we can add one more criterion for a given X being a part of the mechanism M: each X must be not just a functional but also a *spatial* sub-part of M. As a component of M, X will be at a *lower level* than M, and *smaller than* M: "[b]ecause mechanisms are collections of components and their activities, no component can be larger than the mechanism as a whole, and so levels of mechanisms are ordered by size" (Craver & Bechtel 2007, pp. 549–550). Craver and Bechtel conclude: "[m]ost fundamentally, levels of mechanisms are a species of compositional, or part-whole relations" (Craver & Bechtel 2007, p. 550). In the overall framework developed by Craver and Bechtel, functional levels and spatial levels generally align.

Thus, although componential mechanistic explanations are not reductive, they generally *are* what one would call "bottom-up", or perhaps better in this context, "level-restricted": one explains the phenomenon ψ in S by reference to entities and relations at a lower level of organization, but never the reverse. In componential explanations of this sort, the intrinsic properties of and interactions between the mechanism's components account for a system's actions (where "intrinsic" means that such properties—such as the charge of an ion—are either basic to the entity or accounted for by reference to entities and properties at a still lower level of organization). Good mechanistic explanations on this view will not include references to unanalyzed properties of the whole S or M, its "shape" or overall organization, as the relations between the components $\{X\}$ at the lower level will already account for (in fact constitute) these.

This account of mechanistic explanation seems to me a clear and, indeed, compelling model of one kind of explanatory practice in the neurosciences. To satisfy the norms of mechanistic explanation, one must:

1. Identify the phenomenon of interest ψ
2. Identify the system S that ψs
3. Identify the relevant spatial sub-parts $\{X\}$ of M (and their relevant intrinsic properties)
4. Describe how the parts $\{X\}$ are organized such that S ψs

At least *prima facie*, a number of instances of successful (albeit incomplete) explanatory models in the neurosciences appear to neatly fit this description. Craver (2008) extensively discusses the mechanistic model of the action potential. Briefly, following the steps above:

1. The phenomenon ψ is the action potential, which consists of the rapid depolarization of neural cells from a resting membrane potential of approximately –70mV toward (and in many cases significantly exceeding) 0mV; an

2 Note that within this framework "componential mechanism", "constitutive mechanism", and "compositional mechanism" are synonymous.

equally rapid repolarization; a period of hyperpolarization, where the cell overshoots the normal resting potential; and a gradual return to the resting equilibrium (note that as even this simplified sketch illustrates, ψ will often be in and of itself complex, with many aspects that any adequate model must capture).
2. The system *S* that ψs is the neuron.
3. The parts in virtue of which *S* ψs include elements of the cell and its surrounding ionic milieu: positively charged K+ and Na+ ions; gated, ion-specific membrane channels; and the Na+/K+ pump.
4. Finally, the organization that explains ψ includes the following: The resting potential is in fact an equilibrium between two opposing forces: a chemical concentration gradient that pushes Na+ into the cell and K+ out of it, and an electrical gradient that pushes K+ into the cell, each maintained by the selective permeability of the cell to Na+ and K+. Na+ channels change their conformation in response to current flow (they are voltage-gated) such that they open to allow Na+ to flow into the cell. As Na+ flows into the cell this reduces the electrostatic pressure on K+, and opens voltage-gated K+ channels, allowing K+ to flow out of the cell. The net effect is to push the cell initially toward the electrochemical balance point for Na+, which is about +55mV. However, as the membrane potential drops, the Na+ channels close, thus slowing and eventually stopping the depolarization. The diffusion of K+ out of the cell combines with the activity of the Na+/K+ pump to repolarize the cell, which however overshoots the resting potential due to the fact that the K+ channels close later than the Na+ channels, thus allowing K+ to diffuse out of the cell for an extra millisecond or so during which the cell is hyper-polarized.

Obviously, this remains a sketch (see Craver 2008) or any basic neuroscience textbook for more detail), but it illustrates the main elements of a mechanistic explanation. The intrinsic properties, actions, and interactions of *M*'s spatial sub-parts together comprise the mechanism that allows *S* to ψ and thus explain how S ψs. One can likewise plausibly sketch the mechanisms that account for spatial long-term memory (e.g., the ability of an animal to return to some location in its environment) in terms of long-term potentiation of synapses in the hippocampus (Craver 2008), although it is worth noting that a more complete account of the functions of hippocampus will have some of the features I describe in 3 and 4 (Buckner 2010; Anderson 2015). Still, the fact that *some* explanations in neuroscience are like this is not under significant dispute.

But this brings us to the question of why I have distinguished *M* and *S* in my treatment. Because Craver (2008) does not formally distinguish these, he is never led to ask what the precise relationship is (or could be) between *M* and *S* (and between their respective parts). In fact, for Craver the symbol *S* usually (but not always) refers to what I have been calling *M*, and he frames his analysis of mechanistic composition entirely in terms of ψ and its mechanism. When he does mention the larger system it is generally to emphasize the fact that not every part of a system *S* is relevant to the mechanism in virtue of which it ψs. So what might the committed mechanist say about the relationships between *S*, *M* and {*X*}? One possibility is: all the parts {*X*} of *M* will be on a lower level than *S*. That would be in keeping with the level-restricted character of the framework, and its characteristic alignment between spatial and functional levels. It is certainly a feature of all the examples discussed in its support, including the model of the action potential outlined above. A slightly stronger possibility would be: all the parts {*X*} of *M* will be spatial sub-parts of *S*. I don't think anyone would or should endorse this stronger condition, but seeing why will be instructive, and will lead us to the reasons to reject the weaker formulation as well.[3]

3 On my reading, the framework developed in (Craver 2008) implicitly assumes the weaker condition, although most likely not the stronger one. But for my purposes here it is not crucial to pin this down. If the framework *does* assume the weaker condition, what follows should be read as arguing (contra this model) that there are systems for which functional

The immediate trouble with the stronger formulation is that it collides with a fact noted by Craver (2008), but not otherwise discussed: the mechanism that accounts for S ψ-ing may contain parts that are extrinsic to S (although not to M). For instance, in the mechanism for the action potential, the Na^+ and K^+ ions that are clearly part of M are (at least sometimes) extrinsic to S; and in embodied accounts of some cognitive processes like mathematics, the mechanism that accounts for a person (P) multiplying (ψ_m-ing) contains parts that are *always* extrinsic to P, such as pencil and paper (Clark 1997; see also this collection). These entities would arguably *not* be components of the systems that ψ, although they would be components of the mechanisms in virtue of which they ψ. At the very least, this suggests there are some details yet to be worked out about the necessary physical relationships between M and S that implement the hierarchical and functional relationships in virtue of which M can account for S ψ-ing. There will be (presumably rare) cases in which M and S are identical; cases such as the accelerating car where M contains only parts of S; and cases such as the action potential where M and S cross-cut one another, sharing some but not all of their parts.[4] There may also turn out to be cases in which they share no parts, perhaps because the parts of M and the parts of S are individuated by different criteria, or because S's ability to ψ is imposed by or inherited from an entirely extrinsic mechanism (indeed I'll discuss a potential instance of this class of cases later in the paper).

But distinguishing M and S in this way *also* allows one to ask whether all the parts of M need to be at a lower level than S. If not every X needs to be a spatial sub-part of S, then there is little reason to suppose that each X needs to be on a lower level than S, either. Indeed, I claim that in fact for some systems S the mechanism M will contain items that are neither intrinsic to *nor at a lower level than S*. For instance, I often use other people to help me remember things, in the easiest case by asking them to remind me at some future time. In such a case, this other individual is arguably part of the mechanism responsible for my remembering, but is certainly not for that reason on a lower ontological level than I am, qua remembering system. Moreover, as I will argue when looking at the case discussed below, some relevant parts of M (and certainly M itself) are at a *higher* organizational level than S. Now of course, Craver & Bechtel *define* the concept of lower level in terms of being a part of the mechanism: "an item X is at a lower level than an item S if and only if X is a component in the mechanism for some activity ψ of S" (2007, p. 548). I agree that this holds for the constitutive relationship between *mechanisms* and their parts. But it only holds for all systems S if we assume that all the parts of M are parts of S, and we have seen that this is not always the case. Thus although I think that Craver correctly analyzes the relationship between mechanisms and their parts in terms of constitution, I argue that the more capacious notion of *enabling constraint* better captures the relationship between mechanisms and the systems whose activities they enable.

In any case, with this as background, I now turn to the case of the SAC. In 3, I describe what we know about how the mechanisms in virtue of which the cell operates, and in 4 I discuss the implications of this case for componential mechanistic explanation.

3 Direction selectivity in SAC dendrites: Beyond componential constitution

Starburst Amacrine Cells are axonless neurons found in the retina of mammals and numerous non-mammalian species. Their morphology is planar, with multiple dendrites arrayed, as the

and spatial levels in fact dissociate. If it does not, then what follows should be read simply as offering an account of some of the possible functional relationships between mechanisms and systems, an issue not explored in the original analysis. Either path leads to the same recommended modification of the original model.

4 In the case of the action potential, one *might* mount the argument that the system that ψs is *strictly speaking* S + {the nominally non-S parts of M}, including the surrounding extracellular fluid. That would make M part of S in this case, but it is not clear to me that this move will be equally attractive in every such case, nor do I think the mechanist is *forced* to adopt this strategy.

name suggests, in a starburst pattern around the cell body (Figure 1).

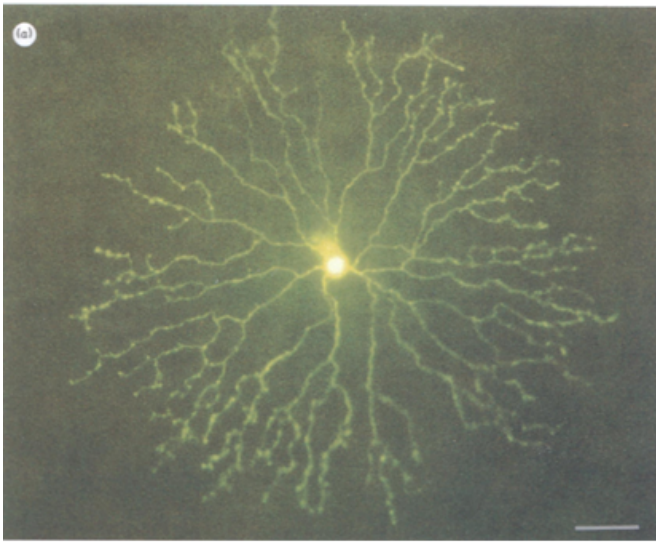

Figure 1: Micrograph of a Starburst Amacrine Cell. Calibration bar 50µm. Reprinted from Tauchi & Masland (1984).

SACs form dense, highly overlapping, co-fasciculating layers in the "on" and "off" levels of the inner synaptic layer of the retina, nestled physically and functionally between bipolar cells and direction-selective ganglion cells. Among the most numerous neural cells found in the mammalian retina, they represent a large proportion of the total neural volume in the eye; in the rabbit retina, for example, as much as six meters of SAC dendrites occupy each square millimeter of retinal surface—higher coverage than any other retinal cell by an order of magnitude (Masland 2005; Tauchi & Masland 1984; see Figure 2).

SACs are interesting for multiple reasons. Despite lacking axons, they synthesize and release both excitatory and inhibitory neurotransmitters (ACh (acetylcholine) and GABA (-Aminobutyric acid)) from the distal regions of their dendrites. Both the role and relative proportion of excitatory and inhibitory synaptic connections change over time. Cholinergic synaptic connections between neighboring SACs disappear over development, and GABAergic connections between SACs begin as excitatory but later become inhibitory. However, excitatory cholinergic synapses between SACs and ganglion cells remain (Masland 2005).

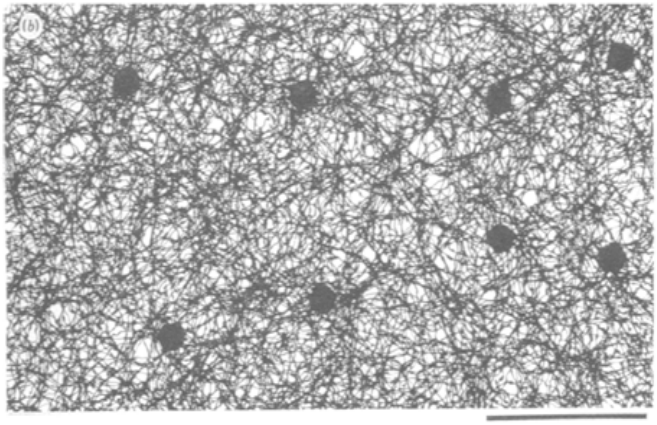

Figure 2: Depiction of the SAC network in peripheral retina. Calibration bar 50µm. Reprinted from Tauchi & Masland (1984).

Functionally, SACs play an important role in motion detection, and are part of the overall network for multiple uses including optokinetic eye movement and motion perception (Yoshida et al. 2001). In fact, each dendrite of the SAC acts independently of the others, and signals the presence of stimuli moving centrifugally, that is, from the cell body out in the direction of the signaling dendrite (Euler et al. 2002; see Figure 3). Put differently, each SAC dendrite is a directionally selective spatial sub-part of the overall cell, and this is the functional property that will interest us here. As with so much in the neurosciences, the mechanism that explains this function is complex and not fully understood. It is, however, possible to offer a sketch of it.

As mentioned above, SACs lie between bipolar cells and direction-selective ganglion cells. Bipolar cells thus mediate the initial stimulus such that a moving light causes them to fire in turn as the stimulus moves across the retina. The bipolar cells make excitatory synapses onto the SAC dendrites.[5] With these basic anatomical facts in view, we can turn to describing

5 In fact there are two classes of bipolar cells, "on" and "off", functionally differentiated by their disposition to respond to stimulus onset vs. stimulus offset—i.e., one responds to light and the other to dark—and anatomically distinguished by whether they synapse onto the "on" or "off" level of the inner synaptic layer (Figure 4). As the mechanisms for direction selectivity in SAC dendrites are the same regardless, I'll ignore this detail in what follows.

three different aspects of the overall mechanism for direction selectivity: wiring specificity between bipolar cells and the SAC dendrites; lateral inhibition between neighboring SACs; and active elements in the dendrites themselves.

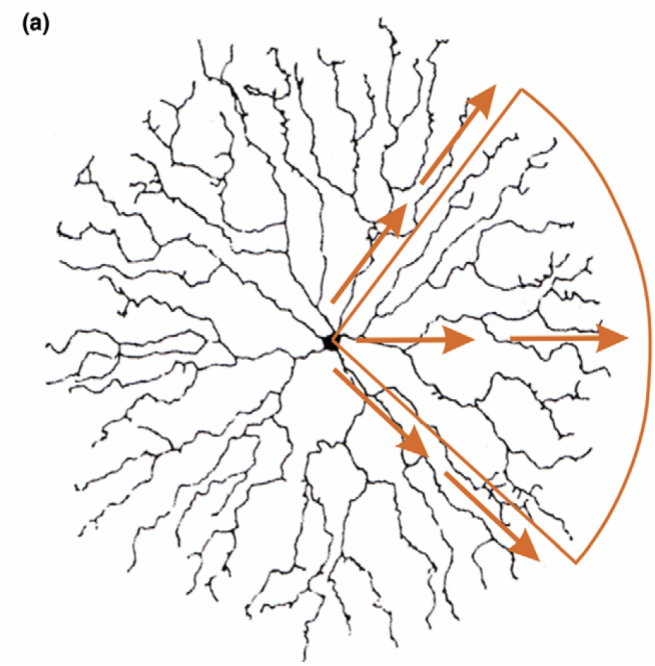

Figure 3: Depiction of direction selectivity in SAC dendrites. Reprinted from Masland (2005).

First, the axonal projections of bipolar cells largely preserve the topography of their inputs, such that neighboring axons come from cells with neighboring inputs, and make neighboring synapses onto post-synaptic cells. What this arrangement means for SACs is that neighboring synapses on the dendrite are likely to come from neighboring bipolar cells, so that when a moving stimulus activates one cell, and then another immediately to its left (say), this will tend to activate a given synapse, and then another immediately to *its* left. Thus, in the case where such a stimulus moves along the direction of a dendritic process, the successive excitatory inputs to that dendrite will tend to reinforce (Demb 2007; Lee & Zhou 2006). This is an important part of the overall mechanism, but is not sufficient by itself to produce the observed directional selectivity, as these inputs would tend to reinforce even during centripetal motion, although this would result in a weaker response at the *distal* process of the dendrite (Hausselt et al. 2007).

Another important part of the mechanism for directional selectivity involves mutual inhibition between neighboring SACs (Figure 5). As a stimulus moves so as to stimulate the centrifugal dendrite of SAC1 (in Figure 5A), reinforcing inputs will cause the release of GABA onto the centripetal dendrite of SAC0, such that even when the light stimuli begins to excite the centripetal dendrite of SAC0, the leading inhibition dominates the signal. Similarly, as the stimulus moves to the *centrifugal* dendrite of SAC0, the successive excitatory inputs from the bipolar cells reinforce, and any inhibitory inputs from the neighboring SAC2 come too late. Moreover, SAC0 will largely inhibit SAC2's response (Figure 5B; Lee & Zhou 2006). An important element of this mechanism involves the relative time-course of ACh and GABA: ACh response from the bipolar cells ramps up and decays fairly quickly, while GABA response is relatively delayed and prolonged (Demb 2007). This temporal asymmetry helps ensure that when inhibition leads it dominates, and vice-versa. The distance between SACs also plays a role. The likelihood of synaptic connections between the distal portion of the dendrites of two SACs—where inhibitory connections are most effective—depends on the distance between the cell bodies. Cells that are very close together or very far apart will thus not mutually inhibit one another (Figure 5C).

Finally, direction selectivity depends upon properties of the dendrite itself. The dendrites are electrically isolated from one another, as a result of both overall cell morphology and the low impedance of the cell body. The uneven distribution of synaptic inputs and outputs also contributes: excitatory inputs from the bipolar cells are distributed along the length of the dendrite, but synaptic outputs are confined to the distal ends (as implied by the two aspects of the overall mechanism described above). A third, active aspect of the local dendritic portion of the mechanism appears to involve voltage-gated calcium channels. These channels lead to amplification of the ACh response beyond what the passive reinforcement caused by successive synaptic transmission from bipolar cells can account for (Hausselt et al. 2007).

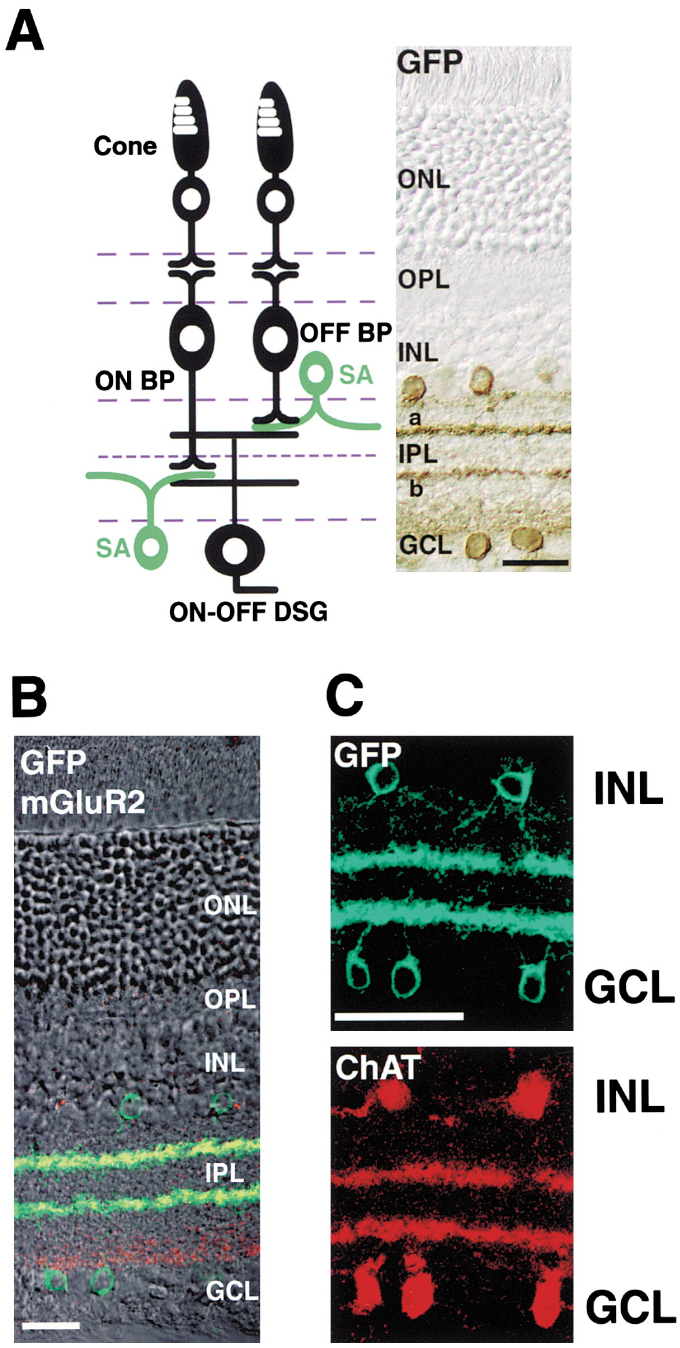

Figure 4: Schematic representation of the layered structure and synaptic relationships between bipolar cells and SACs. Reprinted from Yoshida et al. (2001).

All of these elements combine to produce the direction selectivity of the SAC dendrite. Bipolar cells successively synapse onto the dendritic process, resulting in passive reinforcement of excitatory input that preferentially promotes neurotransmitter release in response to motion in the centrifugal direction. Surrounding SACs selectively inhibit centripetal excitation, as a result of the different temporal activation profiles of GABA and ACh; the asymmetric distribution of input and output synapses; and the relative spatial placement of the SACs. And voltage-gated calcium channels in the dendrite actively amplify the centrifugal signal. Although this sketch leaves out many of the known details, and there remain many details still to be worked out, I believe it is sufficient to warrant the conclusion that this is (a) an instance of mechanistic explanation that (b) does not have the level-restricted character of the (canonical) mechanistic explanations laid out above. I spell out the reasons for this conclusion in the next section.

4 Constitution and constraint

We can most readily see why this case represents an interesting challenge for componential mechanism by fitting it to the four steps outlined in section 2, above.

1. Identify the phenomenon of interest ψ
2. Identify the system S that ψs
3. Identify the relevant spatial sub-parts $\{X\}$ of M (and their relevant intrinsic properties)
4. Describe how the parts $\{X\}$ are organized such that S ψs

The specific phenomenon of interest ψ_{ds} is direction selectivity or, more precisely, the release of neurotransmitter in and only in response to motion in a specific centrifugal direction. The system S_{ds} that exhibits ψ_{ds} is the dendrite of the SAC. It is also easy to say what the parts $\{X_{ds}\}$ of the mechanism M_{ds} are in virtue of which the dendrite ψ_{ds}s, and how they are organized. I have provided that sketch above. Finally, it seems right to say, following Craver (2008), that the relationship between M_{ds} and its parts $\{X_{ds}\}$ is one of componential constitution, such that all the parts $\{X_{ds}\}$ are at a lower level than M_{ds}, and together constitute M_{ds}. But now it gets interesting for componential mechanistic explanation as currently developed. For only some of the parts of M_{ds}—including the voltage gated calcium channels, and the input and output synapses—are at a lower (spatial) level than the

dendrite $S_{ds.}$ The inhibitory dendrites of the neighboring SACs are at the same level as S_{ds}, the bipolar cells and their spatial relations are arguably at a higher level than S_{ds} (although one might wish to screen these off as mere *inputs* to the mechanism), and the mechanism M as a whole in virtue of which S_{ds} $\psi_{\text{-ds}}$s is *certainly* at a higher level than, and is in no way a physical or functional component of S_{ds}.

I think this example demonstrates that not every mechanistic explanation will have the "bottom-up" or "level-restricted" character that the mechanism for the action potential has, where function is built entirely from the capacities of lower-level components and their interactions. In the SAC dendrite, we appear to have a case *not* of a system that ψs in virtue of the capacities and relations of its components (and that could in turn be thought of as a component supporting the activities of a larger functional system), but rather very nearly the reverse: a system that ψs in virtue of the properties of and interactions in the higher-level system of which it is a part. That is, the SAC dendrite is not functionally related to its surrounds as a component to a higher-level system; nor is the higher-level system related to the SAC dendrite as one of *its* components. Instead, I want to say that the higher-level mechanism M acts as an *enabling constraint* on S.

Before providing a bit more in the way of substantial analysis of the concept of an enabling constraint, let us pause to consider one way in which a supporter of componential mechanistic explanation might resist this conclusion by redefining the system S_{ds} to include the mechanism M_{ds}. I think this is not a viable option for a number of reasons. First, it would appear to violate standard usage: neuroscientists speak of direction-selective dendrites, and *not* of a directionally selective network spanning several retinal layers. The debate in the neuroscientific literature concerns *not* the definition of the direction-selective system, but the relative role of intrinsic and extrinsic mechanisms for dendritic direction selectivity in SACs (Hausselt et al. 2007; Lee & Zhou 2006).

Second, it appears that the mechanism as a whole is *not* direction selective. Any given SAC, for instance, and certainly the network as a whole, signals motion in *all* directions. Even if we restrict the definition of M_{ds} to the entities in virtue of which *one* particular SAC dendrite is directionally selective, the symmetry of the mechanism—the fact that SACs *mutually constrain* one another and the same bipolar cells synapse onto more than one SAC dendrite—strongly suggests that *very same mechanism* generates right direction selectivity in the rightward-reaching dendrite in SAC0, and left direction selectivity in the leftward-reaching dendrite in SAC2 (e.g., in Figure 4). The mechanism, that is, does not have the same direction selectivity as either of the dendrites. Rather, it's as if when you turn the crank one way (i.e., the stimulus moves one way) the mechanism produces one output; and when you turn it the other way, it produces the other output.

This suggests a different way to illustrate the limitations of componential mechanism as formulated. Craver writes that the explanandum phenomenon ψ is "typically some behavior of the mechanism as a whole" (Craver 2008, p. 139), and he thus might insist, contra my way of formulating his framework in 2, that it is the mechanism M and not the system S that exhibits ψ. In this case, because I have agreed that the parts $\{X\}$ in fact constitute M, any conflict between functional and spatial levels disappears. But in the case before us it seems that the mechanism *responsible* for, say, rightward direction selectivity does not in fact *exhibit* rightward direction selectivity. So the functional puzzle reasserts itself in a different guise.[6]

One might nevertheless insist on distinguishing these mechanisms in subtle ways—perhaps M_{ds0} includes these synapses from bipolar cells, but not those synapses, while M_{ds2} includes those synapses but not these. I doubt whether this can work, because explaining direction selectivity in *either* direction will require reference to the excitatory inputs from bipolar cells to the centrifugal dendrite, and the inhibitory inputs from the overlapping centripetal dendrite, which are in turn a result of the excitatory inputs from the *very same* bipolar cells synapsing

6 Thanks to an anonymous reviewer for pointing out this way of expressing the matter.

onto the centrifugal dendrite. But let us take the possibility as granted. Then one seems forced to say something along the following lines: the mechanism as a whole ψs, but *signals* ψ-ing with the dendrite.

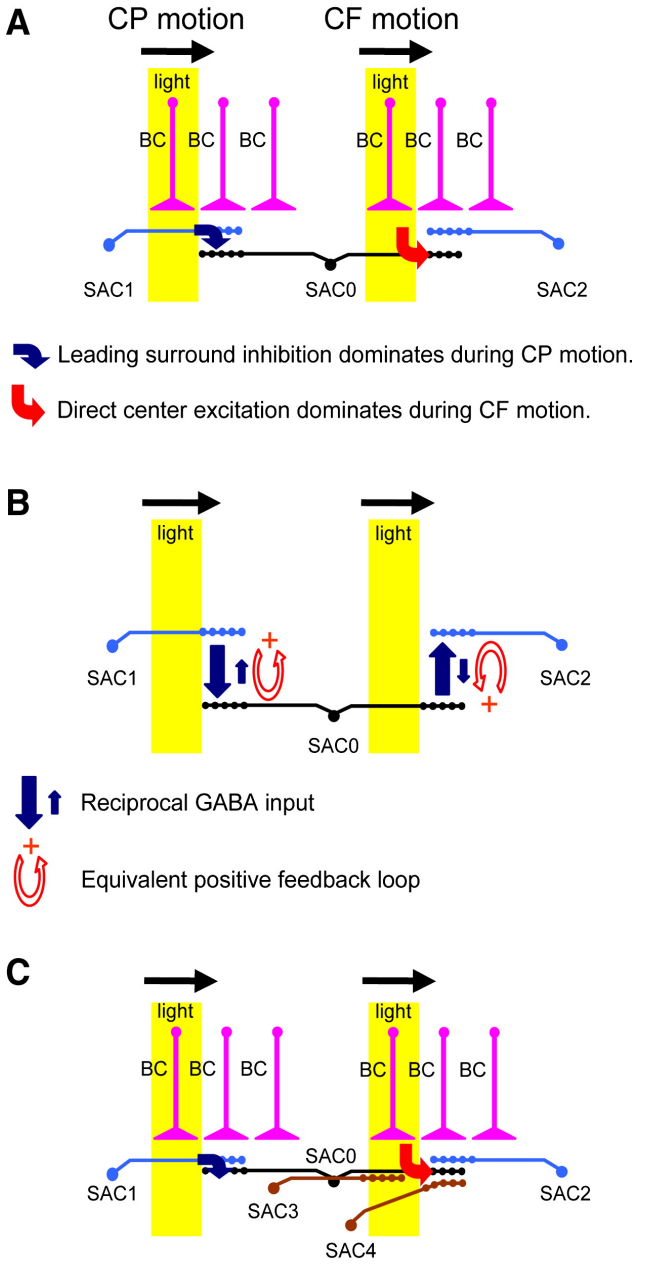

Figure 5: Lateral inhibition between neighbouring SACs contributes to direction selectivity in the dendrites. Reprinted from Lee & Zhou (2006).

Let us consider this possibility carefully. As I intimated above, scientists debate the relative importance of intrinsic and extrinsic mechanisms for dendritic selectivity in SACs. Hausselt et al. (2007) note that direction selectivity in SAC dendrites persists in the presence of GABA and glycene receptor antagonists, which would deactivate the portions of the normal mechanism that involve mutual inhibition between neighboring SACs. In these circumstances, one might argue that *only* the portions of the original mechanism *intrinsic* to the dendrite matter in the explanation of direction-selectivity, and in such a case it is clearly the dendrite that ψs. What shall we say, then, when we remove the antagonists from the system and reapply the same directional stimulus, resulting in neurotransmitter release from this dendrite? One option is: whereas before the dendrite ψ'd, now it merely signals the ψ-ing of the larger mechanism. But it seems clear to me that, if the dendrite can ψ, then adding network interactions that *aid and enhance* (that is, do not in any sense prevent) ψ-ing can hardly cause it to *not* ψ, but only signal ψ. This points to a fourth and final reason to reject the general move to extend the neural system S to include the mechanism M whenever it is (or contains entities that are) on a higher level than S: one would apparently need the ability to rigorously distinguish between ψ-ing and *signaling* ψ in an overall system where to ψ is generally also to signal it—that is, where signaling and doing are deeply intertwined. Thus, I believe we must insist: the dendrite ψs.

For all these reasons, I do not think it is wise to hold onto level-restricted explanations and componential composition by fiat. Instead, it is time to expand the scope of mechanistic explanation by considering the various ways in which systems S relate to the mechanisms M that enable their activities. I think the case of SACs is especially important because it illustrates one way in which local selectivity in parts of a network can be the result of the interplay of excitation and mutual inhibition between non-selective parts of that network, which is clearly something that we need to understand better if we are to accurately characterize the functional mechanisms at work in both small and large-scale brain networks (Anderson et al. 2013). But other structure-function relationships appear to call equally for a broader account of mechanistic explanation. For instance,

the direction-selective ganglion cell DSGC (Direction-Selective Ganglion Cell), mentioned briefly above, responds to stimuli moving only in its preferred direction (which of course varies cell-to-cell). In this case, there do not appear to be *any* intrinsic mechanisms for the direction selectivity of the DSGC. Rather, SAC dendrites selectively synapse onto DSCGs with preferred stimuli antiparallel to the SAC dendrite preference (Briggman 2011) thus suppressing responses to motion in the non-preferred direction. DSCGs seem to simply *inherit* their selectivity via their synaptic contact with SACs—and, in fact, elimination of SACs from the retina abolishes direction selectivity in DSCGs (Yoshida et al. 2001). Here I just don't see any case for a compositional relationship between the mechanism (or its parts) and the selective system. Instead, the relevant mechanism synapses onto the relevant system, and by suppressing a sub-set of its response tendencies, induces selectivity.

This brings us finally back to the notion of "constraint", which I think may help us understand the full range of mechanism/system relationships in the brain. The term constraint has been used in myriad ways in the literature on scientific explanation. In evolutionary biology, scientists refer for instance to stability constraints (Schlosser 2007) and both universal and local developmental constraints on evolvability (Maynard Smith et al. 1985). There are also law-like constraints on the possible states of physical systems generally (Lange 2011). None of these capture the sense of "constraint" that will be most helpful to us here.

One notion that gets us close is the idea of a "capacity constraint", that is, a limitation on the capacity of a process that might take the form of changing the relative probabilities of the range of possible process outcomes (Sansom 2009). This certainly has the right flavor, for in the mechanism under discussion above it appears that the excitatory and inhibitory interactions between bipolar cells and neighboring SACs bias the outcome of the dendritic processing of the moving stimulus. But insofar as a capacity constraint is generally conceptualized in terms of the reduction of some pre-existing whole ability—in Sansom's (2009) example, being handcuffed limits one's ability to move one's hands—this does not offer quite the right organizing frame for explanation in neuroscience.

The reason is that in the neurosciences we want to understand not just the capacities of entities, but how the structured interactions between entities give rise to *functions*, which are, crucially, *differential* and *differentiating* processes (that is, they differ from one another, and they differentiate between stimuli). Capacities in the sense of general powers (the capacity to generate an action potential, say) are necessary conditions for functions, but they are not yet functions; the DSGC is strictly speaking *non-functional* in the absence of SACs, even though it will continue to exercise its capacity to fire action potentials in response to inputs from bipolar cells. Constraints of the sort under investigation here serve to limit capacities, but in so doing they enable functions; they result in an *enhancement* (not a reduction) of the abilities of the system (and the organism).

For this reason I propose to analyze the general functional (and, crucially, *non-hierarchical*) relationship between mechanisms and systems in the following way: an *enabling constraint* is a relationship between entities and/or mechanisms at a particular level of description and a functional system at the same or a different level, such that the entities/mechanisms bias (i.e., change the relative probabilities of) the outcomes of processing by the system. Such enabling constraints offer necessary but not sufficient conditions for the instantiation of differential function in neural systems. Because enabling constraints are synchronic rather than diachronic, the idea shares the same explanatory advantage that the relation of constitution has over the relation of "causation" (when understood, e.g., as an event involving the transmission of some property, power, or conserved quantity from one entity to another). As Craver & Bechtel (2007) point out, such a conception of causation does not accomodate interlevel functional relationships well, because these are often synchronic and symmetric, whereas causa-

tion of this sort is temporal and asymmetric.[7] In addition, enabling constraints can be *mutual*, which gives the idea an advantage over both causation and constitution as an analysis of functional relationships in the brain.

> Enabling constraint $=_{\text{Df}}$ A physical relationship between a functional system S and entities {X} (and/or mechanism M), at the same or different level of description, such that {X} (and/or M) changes the relative probabilities of various possible functional outcomes of activity in S.

To understand function not just in systems like SAC dendrites and DSGCs, but also in the large scale networks that are partially constituted by the Transiently Assembled Local Neural Subsystems TALoNS (Transiently Assembled Local Neural Subsystems) crucial to the functioning of a dynamic brain (Anderson 2015), we need to accept that there is a broader range of relationships that mechanisms can have to functional systems, beyond componential constitution. Function in TALoNS results not from structured interactions between stable, autonomous low-level components, but rather from the interplay between the capacities of lower-level entities and higher-level network dynamics. That interplay, I argue, is best analyzed in terms of the mutual constraint that exists between bottom-up and top-down, feed-forward and feed-back mechanisms in the brain.

5 Conclusion

Although mechanistic explanation as developed by Craver & Bechtel (2007; Craver 2008) does seem to accurately characterize one kind of explanation in neuroscience, and one kind of functional arrangement in neural systems, I've argued here that the formulation is not wide enough to capture the variety of mechanisms in the brain. When we formally distinguish the system S from the mechanism M in virtue of which S exhibits the explanandum phenomenon Ψ, we see that although it seems correct to describe the relationship between M and its parts $\{X\}$ in terms of constitution, it will only sometimes be the case that S is (partially) constituted by $\{X\}$.

As an alternative to the relationship of componential constitution, I have offered the notion of an *enabling constraint* that can exist between a system and the mechanism(s) in virtue of which it has its various functions. SAC dendrites appear to have their function in virtue of the enabling constraints imposed by entities at the same and higher levels of organization; and DSGC function is enabled by the constraints imposed by the SAC dendrites. In neither case is it appropriate to describe the relationship between the mechanism M and the relevant system S in terms of constitution, nor are all (or, in the case of DSGCs arguably any) of the parts $\{X\}$ of M components of S.

Overall, I hope to have made the case that moving beyond level-restricted mechanistic explanation will allow us to better capture the variety of neural systems that emerge from the constant, constraining, biasing interplay between feed-forward, feedback, bottom-up, and top-down processes in the dynamic brain.

7 For instance, what explains why a neuron has a particular functional property cannot be an event involving the transmission of some property, power or conserved quantity from the parts of the neuron to the whole, because if causes must precede their effects, this would appear require that there be a time prior to which the neuron did not have the functional property conferred by its parts. Interlevel functional relationships do not generally appear to be temporal in this way. Rather, for Craver and Bechtel, what explains the functional property of the neuron is the way it is *constituted* by its parts. Enabling constraints are also synchronic in the relevant way, and so the view I am advocating here is also able to accommodate such cases of interlevel functional relationships.

References

Anderson, M. L. (2015). *After phrenology: Neural reuse and the interactive brain.* Cambridge, MA: MIT Press.

Anderson, M. L., Kinnison, J. & Pessoa, L. (2013). Describing functional diversity of brain regions and brain networks. *NeuroImage*, *73*, 50-58. 10.1016/j.neuroimage.2013.01.071

Briggman, K. L., Helmstaedter, M. & Denk, W. (2011). Wiring specificity in the direction-selectivity circuit of the retina. *Nature*, *471* (7337), 183-188. 10.1038/nature09818

Buckner, R. L. (2010). The role of the hippocampus in prediction and imagination. *Annual Review of Psychology*, *61*, 27-48. 10.1146/annurev.psych.60.110707.163508

Clark, A. (1997). *Being there: Putting brain, body and world together again.* Cambridge, MA: MIT Press.

——— (Ed.) (2015). Embodied prediction. In T. Metzinger & J. M. Windt (Eds.) *Open MIND.* Frankfurt a. M., GER: MIND Group.

Craver, C. F. (2008). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience.* Oxford, UK: Oxford University Press.

——— (Ed.) (2015). Levels. In T. Metzinger & J. M. Windt (Eds.) *Open MIND.* Frankfurt a. M., GER: MIND Group.

Craver, C. F. & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology & Philosophy*, *22* (4), 547-563. 10.1007/s10539-006-9028-8

Demb, J. B. (2007). Cellular mechanisms for direction selectivity in the retina. *Neuron*, *55* (2), 179-186. 10.1016/j.neuron.2007.07.001

Euler, T., Detwiler, P. B. & Denk, W. (2002). Directionally selective calcium signals in dendrites of starburst amacrine cells. *Nature*, *418* (6900), 845-852. 10.1038/nature00931

Hausselt, S. E., Euler, T., Detwiler, P. B. & Denk, W. (2007). A dendrite-autonomous mechanism for direction selectivity in retinal starburst amacrine cells. *PLoS Biology*, *5* (7), e185. 10.1371/journal.pbio.0050185

Hempel, C. G. (1965). *Aspects of scientific explanation.* New York, NY: Free Press.

Lange, M. (2011). Conservation laws in scientific explanations: Constraints or coincidences? *Philosophy of Science*, *78* (3), 333-352. 10.1086/660299

Lee, S. & Zhou, Z. J. (2006). The synaptic mechanism of direction selectivity in distal processes of starburst amacrine cells. *Neuron*, *51* (6), 787-799. 10.1016/j.neuron.2006.08.007

Masland, R. H. (2005). The many roles of starburst amacrine cells. *Trends in Neurosciences*, *28* (8), 395-396. 10.1016/j.tins.2005.06.002

Maynard Smith, J., Burian, R., Kauffman, S., Alberch, P., Campbell, J., Goodwin, B., Lande, R., Raup, D. & Wolpert, L. (1985). Developmental constraints and evolution. *Quarterly Review of Biology*, *60* (3), 265-287.

Sansom, R. (2009). The nature of developmental constraints and the difference maker argument for externalism. *Biology & Philosophy*, *24* (4), 441-59. 10.1007/s10539-008-9121-2

Schlosser, G. (2007). Functional and developmental constraints on life cycle evolution: An attempt on the architecture of constraints. In R. Sansom & R. Brandon (Eds.) *Integrating evolution and development: from theory to practice* (pp. 113-173). Cambridge, MA: MIT Press.

Tauchi, M. & Masland, R. H. (1984). The shape and arrangement of the cholinergic neurons in the rabbit retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, *223* (1230), 101-119. 10.1098/rspb.1984.0085

Yoshida, K., Watanabe, D., Ishikane, H., Tachibana, M., Pastan, I. & Nakanishi, S. (2001). A key role of starburst amacrine cells in originating retinal directional selectivity and optokinetic eye movement. *Neuron*, *30* (3), 771-780. 10.1016/S0896-6273(01)00316-6

Carving the Brain at its Joints

A Commentary on Michael L. Anderson

Axel Kohler

When neuroscientists explain the biological basis of a phenomenon of interest, they usually try to identify the parts of a system that seem do the relevant job, and propose a model of how those parts interact to produce the phenomenon. This mechanistic framework of explanation is widely used and has been investigated from a philosophical point of view by different authors. In his target article, Michael Anderson poses a challenge to the currently dominant version of mechanistic explanation as advocated, e.g., by Carl Craver. Taking empirical results and explanatory models from studies on retinal starburst amacrine cells as a starting point, Anderson suggests that the current framework for mechanistic explanation should be extended to include a differentiation between systems and mechanisms, which would allow more leeway in understanding processing in the nervous system. Mechanisms can then be seen to provide enabling constraints on the functioning of systems, where the mechanisms do not need to be subsumed under the system and do not even have to be on the same organizational level. Although Anderson's proposal is interesting and worth exploring, I am unconvinced that this extension conforms to real-world explanatory practice and/or is necessary for accommodating the understanding of direction-selectivity in the retina. I examine another sample of research on starburst amacrine cells, where the integration of empirical data and computational models shows that, on close inspection, it is distributed networks to which certain characteristics are ascribed—a situation that can be handled with the available tools of mechanistic explanation.

Keywords
Constitution | Direction selectivity | Enabling constraint | Enabling constraints | Mechanism | Mechanistic explanation | Motion processing | MT | Neuroscience | Neuroscientific explanation | Starburst amacrine cells | Top-down causation | V1

Commentator

Axel Kohler
axelkohler@web.de
Universität Osnabrück
Osnabrück, Germany

Target Author

Michael L. Anderson
michael.anderson@fandm.edu
Franklin & Marshall College
Lancaster, PA, U.S.A.

Editors

Thomas Metzinger
metzinger@uni-mainz.de
Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt
jennifer.windt@monash.edu
Monash University
Melbourne, Australia

1 Introduction

One of the dominant frameworks of explanatory practice in the neurosciences and the biological sciences in general is the model of mechanistic explanation proposed in its modern form by Bechtel & Richardson (1993) and recently extended by Carl Craver (2007). Mechanistic explanations describe entities and activities that together bring about a phenomenon of interest (Machamer et al. 2000). When we are interested in how vision works, for example, we try to localize the relevant parts of the brain, and identify components and their types of interactions in order to understand how we can see things (Bechtel 2008). This model of mechanistic explanation is thought to capture the dominant explanatory practice in the biological sciences (Bechtel & Richardson 1993), but normative claims are also made with respect to the adequacy of explanatory accounts. Craver (2007) proposes a number of constraints on constitutive mechanistic explanation in order to decide whether a mechanistic model is viable or not.

In his target article, Michael Anderson (this collection) takes current models of mechanistic explanation as a starting point for proposing an important extension of the existing accounts. In previous models, the system that

exhibits a phenomenon and the mechanism that explains the phenomenon were not separated. Sometimes parts of the system can be screened off with respect to the phenomenon at hand. The windshields of a car and its radio components are not really important in order to understand how it drives, for example. It's fine to say that the whole car drives, but that only the relevant components (engine, axles, tires) are doing the mechanistic work. Focusing on the essential components of a mechanism within a larger system is unproblematic. But Anderson worries about more complex cases in the neurosciences where the system displaying a phenomenon does not encompass the relevant mechanism producing the phenomenon and might not even be on the same level of description as the mechanistic components.

Anderson wants to demonstrate that componential constitution is not sufficient as a model of mechanistic explanation for the processing of directional selectivity in the retina. Mechanisms computing direction of motion are already available at the earliest stages of the visual hierarchy. The vital components of direction selectivity in the retina could be identified. In particular, in recent discussion starburst amacrine cells (SAC) have been viewed as a mechanistic substrate of motion processing. The SACs receive input from bipolar cells, which are not themselves directionally selective, and provide output to direction-selective ganglion cells (Zhou & Lee 2008). The SACs themselves seem to be the core component for retinal motion selectivity (Park et al. 2014; Yoshida et al. 2001).

Examining the current models of how direction selectivity is created in SACs, Anderson takes note of a discrepancy between how direction selectivity is mechanistically achieved and to which parts it is ascribed. He argues for a distinction between the system S that Ψs (that is, exhibits direction selectivity) and the mechanism M that accounts for S's Ψ-ing. For the case at hand, the SACs themselves or even just single dendritic compartments of SACs Ψ, but a much broader network of neighboring SACs and bipolar cells needs to be considered in order to provide a mechanistic account of SAC direction selectivity. Anderson proposes this distinction as an important extension of Craver and Bechtel's model of mechanistic explanation. This has two major advantages, according to Anderson: (1) there can be entities and actions that play a role for M, but are not necessarily parts of S. This allows a certain flexibility in defining the system that displays Ψ, while at the same time including all relevant components in the mechanistic account of S's Ψ-ing. (2) But if there are parts of M that don't need to be spatially subsumed under S, neither do they need to be at a lower level than S. So even the requirement of componential constitution might be relaxed to allow for higher-level mechanistic components that play an important role in S's Ψ-ing.

As an alternative account of the relationship between mechanisms M and the respective systems S, Anderson proposes that M acts as an enabling constraint on S:

> [A]n enabling constraint is a relationship between entities and/or mechanisms at a particular level of description and a functional system at the same or a different level, such that the entities/mechanisms bias (i.e., change the relative probabilities of) the outcomes of processing by the system. (this collection, p. 12)

In the case of retinal direction selectivity, the mechanistic interaction between neighboring SACs and BCs acts as an enabling constraint for the direction selectivity of a specific SAC dendritic compartment (i.e., the system).

The most straightforward move by proponents of existing models of mechanistic explanation, as Anderson (this collection) also notes, would be to claim that the differentiation of system and mechanism is vacuous. Only the mechanism as a whole can do the work. Even in complex cases, one just has to pick out the right subparts of the network (specific synapses, specific compartments of neurons) that together produce the phenomenon of interest. Anderson provides a number of arguments against this way of extending the concept of mechanism/system, which I would like to briefly summarize:

1. Neuroscientists just don't talk about complex directionally selective networks, but about the direction selectivity of certain dendritic branches.
2. The mechanism as a whole does not display a specific direction selectivity (it is not rightward-selective etc.), it only contributes to the specific selectivity in the respective SAC dendrites. The mechanism contributes to different kinds of selectivities in different dendrites.
3. Making fine-grained distinctions between subparts (synapses, axon branches, dendrites etc.) of the very same neurons that contribute to different directional selectivities is implausible.
4. When the whole network is said to be direction-selective (i.e., it Ψs), what about the dendrite itself? Is it supposed to only signal direction selectivity (signal Ψ-ing)? It is unlikely that a clear distinction between Ψ-ing and signaling Ψ-ing can be made.

The aim of this commentary is twofold. First, I would like to argue that the described cases can be handled by current models of mechanistic explanation when one considers the options of reconstituting the phenomena and top-down causation. Second, using another example of research on SACs, I would like to show that the straightforward ascription of direction selectivity to the SAC dendrites is at least debatable. When looking at how empirical results are often integrated with computational models of direction selectivity, it becomes clear that those phenomena can only be understood by considering the distributed nature of the involved networks.

2 Reconstituting the phenomena and top-down causation

Anderson proposes a separation between systems and mechanisms. No matter whether the system is constrained to be a dendritic compartment or whether it is extended to encompass all mechanistically relevant parts, there are tools available to describe the respective situation. The mechanistic model does not necessarily consider systems in isolation from the environment or surrounding processes. Even if the system is defined as the dendrite only, factors influencing dendritic processing as well as the embedding of the system in the overall economy, its organization, have to be considered in order to arrive at an understanding of the system's functioning (Bechtel 2008, pp. 148–150). On the other hand, I would like to argue that we have good reason to extend the boundaries of the system to encompass all the contributing parts. This is a situation in which the original ascription of a function to a system part has to be revised to accommodate new findings. This process is termed *reconstituting the phenomena* by Bechtel & Richardson (1993). Although direction selectivity was thought to be bound to or even intrinsically generated in SAC dendrites, it turns out that the system can only be understood in combination with other neural elements that vitally contribute to the mechanism in question.

One advantage that Anderson suggests comes with the differentiation of system and mechanism is that mechanistic components can then be set at a different level of organization than the relevant system. The SAC dendrite is at a lower level compared to the input from bipolar cells and the network structure (bipolar cells and neighboring SACs) that enables SAC direction-selectivity. But once the question of how exactly we should carve up the brain into systems and mechanisms has been answered, I don't think that complex inter-level relationships are much of an issue for mechanistic accounts. They can be easily accommodated within the framework of top-down causation proposed by Craver & Bechtel (2007). They suggest that any reference to inter-level interactions can be analyzed in terms of within-level causal relationships between parts of entities, where parts and entities are related in a constitutive fashion and entities can be located on different levels. Emphasizing the fact that complex inter-level interactions often need to be considered in order to offer adequate explanatory accounts in neuroscience is important, but it is not outside the scope of current models of mechanistic explanation.

3 Systems and mechanisms for direction selectivity

Since the processing of direction selectivity in the retina is currently a very active research field, there is substantial controversy concerning the relevant entities and activities that contribute to the mechanism, as Anderson points out in his target article. Some accounts focus on local processes within the SAC dendrites themselves (Hausselt et al. 2007), while others draw a broader picture of a multi-component process, where the exact arrangement of cell types and their compartments is vital for direction selectivity (Lee & Zhou 2006). For our purposes here, I would like to use a most recent update on SAC function offered by the group working with Sebastian Seung. The group uses high-resolution electron-microscopy images of brain tissue to reconstruct complete brain networks on a cellular level. Apart from trained reconstruction experts, the project also makes use of so-called "citizen neuroscientists"—volunteers who contribute to the reconstruction process through an online platform that employs gaming features to guide and motivate the community effort (http://www.eyewire.org).

In their study, Seung and colleagues used images from the mouse retina to analyze SAC circuitry. They took a closer look at the exact wiring between bipolar cells (BCs) and SACs (Kim et al. 2014). BCs provide input to SACs, but do not show any directional selectivity by themselves. The main point of the article is to show that different BC subtypes display different patterns of connectivity with SACs. By analyzing branch depth and contact area, they could show that one subtype (BC2) has mainly connections close to the soma, while another subtype (BC3a) has more connections far from the soma in the outer parts of the dendrites. Importantly, the BC subtypes, in turn, have different intrinsic visual response latencies. BC2 seems to lag BC3a by 50ms and more. It can be shown that the differential connectivity patterns and the divergent latencies add up to produce selectivity for a preferred direction of movement going out from the soma on the respective dendrite in accordance with empirical results.

What is important about the paper is not just the main result itself. Any empirical observation may be overruled in the (near) future. So it is not particularly relevant whether these exact cell types and this exact type of wiring is vital for the phenomenon at hand. What I found intriguing in this study, however, was how the relevant mechanism was described and how the data were integrated with a computational model of direction selectivity, reflecting a recent trend in the neurosciences to combine biological and computational perspectives in explanatory accounts. It shows how neuroscientists pick out the relevant parts of a system that contribute to a specific phenomenon in question. The proposed computational model (Fig. 1a; Kim et al. 2014) maps the biological entities onto specific parts of the computational circuit. The output element at the lower part of the figure is the SAC. The input stems from BC2 (left) and BC3a (right); their respective response properties are captured as delay values and sustained vs. transient response types. The circuit combines elements of classical models of direction selectivity, the Reichardt (Fig. 1b) and the Barlow-Levick detectors (Fig. 1c). Clearly, the direction selectivity cannot be attributed to any one of the system components in isolation. Mechanistic accounts and the corresponding computational models both point to the whole complex of entities as the relevant system that achieves directional selectivity.

In its computational abstraction, the model can be thought of as a canonical system of directional selectivity. Similar models have also been applied to different hierarchical levels of neural processing and different species. For example, mechanisms of directional selectivity have been studied for a long time in the fly visual system. With very different neural elements and wiring, a system of interconnected neurons achieves directional selectivity with response properties closely resembling the Reichardt-type of motion detector (Borst & Euler 2011). Again, only the combination of elements from different processing stages succeeds in delivering direction selectivity as a system. On a cortical level, direction selectivity has been first described for complex cells of the

primary visual cortex (V1) in the seminal work of David Hubel & Torsten Wiesel (1962). Without offering a quantitative computational model, they nevertheless suggest a hypothetical connectivity pattern between different cell types that might underlie the observed responses to moving patterns in complex cells (Hubel & Wiesel 1962, Fig. 20). The model shares features with other motion detectors; a mapping between components is possible.

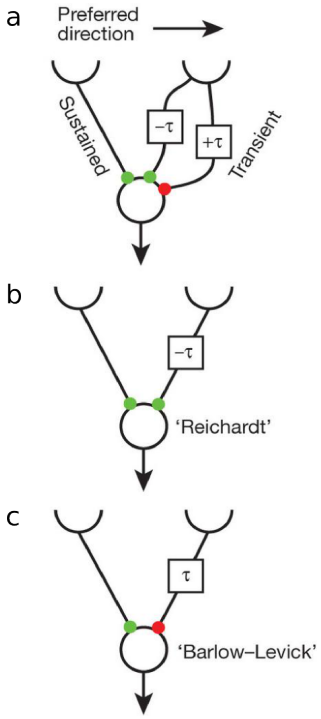

Figure 1: Computational models of direction selectivity (a) The selectivity of SACs described in Kim et al. (2014) can be modeled with a computational framework using a combination of sustained and transient response properties as well as excitatory and inhibitory connections. The displayed wiring would lead to direction selectivity for rightwards motion. The proposed model can be considered to combine previous classical models of direction selectivity, the Reichardt detector (b) and the Barlow-Levick model (c). Green dots indicate excitatory and red dots inhibitory synapses. '-τ' indicates a temporal lead and '+τ' a temporal lag. Reprinted by permission from Macmillan Publishers Ltd: Nature (Kim et al. 2014), copyright (2014).

When it comes to motion selectivity in the brain, one of the most intensively studied cortical areas is the middle temporal (MT) region. The region was first described in the macaque (Dubner & Zeki 1971; Zeki 1974) and owl monkeys (Allman & Kaas 1971). The human homolog, the human MT complex (hMT+; Tootell et al. 1995; Zeki et al. 1991), turned out to be a collection of areas with related response properties (Amano et al. 2009; Kolster et al. 2010). Again, to understand the direction selectivity of MT, it is necessary to consider the cooperation of cells in MT and the input processing stages, mainly from V1. This cooperation and the need for an integrated perspective is emphasized in empirical studies (Saproo & Serences 2014) as well as computational models of MT functioning (Rust et al. 2006). Only the V1-MT system as a whole is understood to deliver motion selectivity as output of the MT stage.

But in terms of the role of MT in motion processing, a case could be made in support of Anderson's suggested distinction between a system that exhibits a certain selectivity and the mechanism that produces this selectivity. The apparent locality and modularity of motion processing in MT is based on very selective deficits in patients with lesions in and around MT (Zeki 1991; Zihl et al. 1983). And stimulation of MT with transcranial magnetic stimulation (TMS) in healthy participants leads to selective deficits in motion perception (Beckers & Hömberg 1992; Beckers & Zeki 1995; Hotson et al. 1994; Sack et al. 2006). In a recent study, patients undergoing brain surgery near MT could be investigated with electrical stimulation (Becker et al. 2013). Only stimulation of MT and a related area nearby, MST, led to an inability to perform a simple motion-detection task, a rather specific result concerning the relevance. Results of that kind drive the intuition that the system that is responsible for motion perception, independent of any cortical areas that might mechanistically contribute to the processing chain leading up to MT (like V1), are localized in MT.

Lesion and other interference studies (e.g., with TMS) are suggestive, but there are also well-known difficulties with interpreting the results. Lesions mostly affect larger parts of the brain and are rarely limited to a single cortical site. As such it is often hard to identify the actual parts of the complex brain networks that

are affected. The advantage of stimulation techniques is that the interference is temporary and can be precisely targeted on a specific location. But, given the rich connectivity structure of neural networks, stimulation effects can be seen even in remote target sites (Bestmann et al. 2004; Sack et al. 2007). In addition, TMS studies have shown that activity of MT might not even be sufficient for conscious motion perception without the involvement of V1 (Pascual-Leone & Walsh 2001; Silvanto et al. 2005). There are also further empirical as well as philosophical reasons for rejecting the claim that motion perception can be attributed to MT in a stringent fashion (Madary 2013), which I won't discuss here.[1]

So while at first glance MT is a very strong candidate for straightforward and very local attribution of function, it seems again that the relevant system is more appropriately described on a network level. The tendency to see system parts as vital for a function may also stem from the limitations of our employed methods. Lesion cases and interference techniques are commonly interpreted as being informative about the relevant gray-matter structures that are affected by the lesion or stimulation. But there is evidence that interference with white-matter connections between network parts can be even more incapacitating than gray-matter damage. It has long been known that frontoparietal areas are implicated in a deficit of visuospatial attention called *neglect*. But very recently Thiebaut de Schotten et al. (2005, 2011) revealed that the properties of fiber connections between frontal and parietal sites are most predictive of visuospatial processing capacities, and that their electrical stimulation leads to severe deficits. Transferring this insight to the case of MT, we simply have most direct access to the cortical gray-matter centers involved in motion processing, and since they are vital components of the system, this also leads to corresponding deficits when they are affected or stimulated. But this might conceal the fact that motion selectivity is a product of a wider network that crucially depends on integrated processing for proper functioning.

In sum, I think that close inspection of how direction selectivity is investigated and treated in neuroscientific research is in disagreement with Anderson's arguments (1) and (3). Although it is true that investigators sometimes refer loosely to local elements as displaying a certain characteristic, the corresponding detailed and extended accounts of direction selectivity give credit to the distributed nature of the relevant systems that figure in explanations. Even considering the case of conscious motion perception, it is unclear whether the presumed locality of motion representation stands up to stringent tests. Rather, it seems to be a case of localized interference with a distributed system where damage to vital hubs leads to fundamental deficits.

4 Conclusion

In this commentary, I have defended the claim that the current tools of mechanistic explanation are sufficient for accommodating the explanatory goals in current neuroscience, particularly in the special case of direction selectivity in the retina and other neural systems. A closer look at explanatory practice shows that, in representative cases of empirical research, models of direction selectivity have to take a number of components in a distributed network into account in order to provide a full-fledged description of the relevant processes. On the philosophical side, the conceptual tools of "reconstituting the phenomena" (Bechtel & Richardson 1993) and "top-down causation" (Craver & Bechtel 2007), offered by existing models of mechanistic explanation, might be sufficient for capturing the problematic cases to which Anderson (this collection) points.

On the other hand, Anderson's proposal (this collection) to extend existing models of mechanistic explanation with the notion of enabling constraints is very interesting and might offer an avenue to more nuanced mechanistic

1 Madary (2013) uses two sets of empirical results to show that representation of motion cannot be ascribed to MT simpliciter. One is the recent emphasis on spontaneous activity making significant contributions to the state of sensory systems—they add content referring to the attentional or sensorimotor state of the organism to input-derived sensory representations. The other demonstrates that in MT specifically, the response properties of cells can be quite variable and are not consistently related to perceptual content only.

descriptions of systems in their contextual embedding. In almost all relevant cases in neuroscience research, there are various external factors influencing the workings of a system, and it is often difficult to draw clear boundaries between vital and non-vital, but nevertheless highly influential system components. Anderson's framework would offer a viable solution for handling those modulatory constraints. Resolving this debate will also depend on a clear conception of how the entities that display a certain phenomenon are best identified and described.

References

Allman, J. M. & Kaas, J. H. (1971). A representation of the visual field in the caudal third of the middle temporal gyrus of the owl monkey (Aotus trivirgatus). *Brain Research*, *31* (1), 85-105. 10.1016/0006-8993(71)90635-4

Amano, K., Wandell, B. A. & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, *102* (5), 2704-2718. 10.1152/jn.00102.2009

Anderson, M. L. (2015). Beyond componential constitution in the brain: Starburst Amacrine Cells and enabling constraints. In T. Metzinger & J. M. Windt (Eds.) *Open MIND*. Frankfurt a. M., GER: MIND Group.

Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience.* New York, NY: Taylor & Francis.

Bechtel, W. & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as scientific research strategies.* Princeton, NJ: Princeton University Press.

Becker, H. G. T., Haarmeier, T., Tatagiba, M. & Gharabaghi, A. (2013). Electrical stimulation of the human homolog of the medial superior temporal area induces visual motion blindness. *Journal of Neuroscience*, *33* (46), 18288-18297. 10.1523/JNEUROSCI.0556-13.2013

Beckers, G. & Hömberg, V. (1992). Cerebral visual motion blindness: Transitory akinetopsia induced by transcranial magnetic stimulation of human area V5. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, *249* (1325), 173-178. 10.1098/rspb.1992.0100

Beckers, G. & Zeki, S. (1995). The consequences of inactivating areas V1 and V5 on visual motion perception. *Brain*, *118* (1), 49-60. 10.1093/brain/118.1.49

Bestmann, S., Baudewig, J., Siebner, H. R., Rothwell, J. C. & Frahm, J. (2004). Functional MRI of the immediate impact of transcranial magnetic stimulation on cortical and subcortical motor circuits. *European Journal of Neuroscience*, *19* (7), 1950-1962. 10.1111/j.1460-9568.2004.03277.x

Borst, A. & Euler, T. (2011). Seeing things in motion: Models, circuits, and mechanisms. *Neuron*, *71* (6), 974-994. 10.1016/j.neuron.2011.08.031

Craver, C. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience.* Oxford, UK: Oxford University Press.

Craver, C. & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, *22* (4), 547-563. 10.1007/s10539-006-9028-8

Dubner, R. & Zeki, S. M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*, *35* (2), 528-532. 10.1016/0006-8993(71)90494-X

Hausselt, S. E., Euler, T., Detwiler, P. B. & Denk, W. (2007). A dendrite-autonomous mechanism for direction selectivity in retinal starburst amacrine cells. *PLoS Biology*, *5* (7), e185. 10.1371/journal.pbio.0050185

Hotson, J. R., Braun, D., Herzberg, W. & Boman, D. (1994). Transcranial magnetic stimulation of extrastriate cortex degrades human motion direction discrimination. *Vision Research*, *34* (16), 2115-2123. 10.1016/0042-6989(94)90321-2

Hubel, D. H. & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, *160*, 106-154.

Kim, J. S., Greene, M . J., Zlateski, A., Lee, K., Richardson, M., Turaga, S. C., Purcaro, M., Balkam, M., Robinson, A., Behabadi, B. F., Campos, M., Denk, W., Seung, H. S. & EyeWirers, (2014). Space-time wiring specificity supports direction selectivity in the retina. *Nature*, *509* (7500), 331-336. 10.1038/nature13240

Kolster, H., Peeters, R. & Orban, G. A. (2010). The retinotopic organization of the human middle temporal area MT/V5 and its cortical neighbors. *Journal of Neuroscience*, *30* (29), 9801-9820. 10.1523/JNEUROSCI.2069-10.2010

Lee, S. & Zhou, Z. J. (2006). The synaptic mechanism of direction selectivity in distal processes of Starburst Amacrine Cells. *Neuron*, *51* (6), 787-799. 10.1016/j.neuron.2006.08.007

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

Madary, M. (2013). Placing area MT in context. *Journal of Consciousness Studies*, *20* (5-6), 93-104.

Park, S. J. H., Kim, I.-J., Looger, L. L., Demb, J. B. & Borghuis, B. G. (2014). Excitatory synaptic inputs to mouse on-off direction-selective retinal ganglion cells lack direction tuning. *Journal of Neuroscience*, *34* (11), 3976-3981. 10.1523/JNEUROSCI.5017-13.2014

Pascual-Leone, A. & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science*, *292* (5516), 510-512. 10.1126/science.1057099

Rust, N. C., Mante, V., Simoncelli, E. P. & Movshon, J. A. (2006). How MT cells analyze the motion of visual patterns. *Nature Neuroscience*, *9* (11), 1421-1431. 10.1038/nn1786

Sack, A. T., Kohler, A., Linden, D. E., Goebel, R. & Muckli, L. (2006). The temporal characteristics of motion processing in hMT/V5+: combining fMRI and neuronavigated TMS. *NeuroImage*, *29* (4), 1326-1335. 10.1016/j.neuroimage.2005.08.027

Sack, A. T., Kohler, A., Bestmann, S., Linden, D. E., Dechent, P., Goebel, R. & Baudewig, J. (2007). Imaging the brain activity changes underlying impaired visuospatial judgments: simultaneous fMRI, TMS, and behavioral studies. *Cerebral Cortex*, *17* (12), 2841-2852. 10.1093/cercor/bhm013

Saproo, S. & Serences, J. T. (2014). Attention improves transfer of motion information between V1 and MT. *Journal of Neuroscience*, *34* (10), 3586-3596. 10.1523/JNEUROSCI.3484-13.2014

Silvanto, J., Cowey, A., Lavie, N. & Walsh, V. (2005). Striate cortex (V1) activity gates awareness of motion. *Nature Neuroscience*, *8* (2), 143-144. 10.1038/nn1379

Thiebaut de Schotten, M., Urbanski, M., Duffau, H., Volle, E., Levy, R., Dubois, B. & Bartolomeo, P. (2005). Direct evidence for a parietal-frontal pathway subserving spatial awareness in humans. *Science*, *309* (5744), 2226-2228. 10.1126/science.1116251

Thiebaut de Schotten, M., Dell'Acqua, F., Forkel, S. J., Simmons, A., Vergani, F., Murphy, D. G. & Catani, M. (2011). A lateralized brain network for visuospatial attention. *Nature Neuroscience*, *14* (10), 1245-1246. 10.1038/nn.2905

Tootell, R. B., Reppas, J. B., Kwong, K. K., Malach, R., Born, R. T., Brady, T. J. & Belliveau, J. W. (1995). Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *Journal of Neuroscience*, *15* (4), 3215-3230.

Yoshida, K., Watanabe, D., Ishikane, H., Tachibana, M., Pastan, I. & Nakanishi, S. (2001). A key role of starburst amacrine cells in originating retinal directional selectivity and optokinetic eye movement. *Neuron*, *30* (3), 771-780. 10.1016/S0896-6273(01)00316-6

Zeki, S. M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *Journal of Physiology*, *236* (3), 549-573.

——— (1991). Cerebral akinetopsia (visual motion blindness): A review. *Brain*, *114*, 811-824. 10.1093/brain/114.2.811

Zeki, S., Watson, J. D. G., Lueck, C. J., Friston, K. J., Kennard, C. & Frackowiak, R. S. J. (1991). A direct demonstration of functional specialization in human visual cortex. *Journal of Neuroscience*, *11* (3), 641-649.

Zhou, Z. J. & Lee, S. (2008). Synaptic physiology of direction selectivity in the retina. *Journal of Physiology*, *586* (Pt 18), 4371-4376. 10.1113/jphysiol.2008.159020

Zihl, J., von Cramon, D. & Mai, N. (1983). Selective disturbance of movement vision after bilateral brain damage. *Brain*, *106* (2), 313-340. 10.1093/brain/106.2.313

Functional Attributions and Functional Architecture

A Reply to Axel Kohler

Michael L. Anderson

In his commentary (Kohler this collection) on my target article (Anderson this collection), Axel Kohler suggests that componential mechanism (Craver 2008) in fact suffices as a framework for understanding function-structure relationships, even in complex cases such as direction selectivity in Starburst Amacrine Cells. Here I'll argue that while Kohler is correct that the framework *can* accommodate such cases, this approach misses an opportunity to draw important distinctions between what appear to be different sorts of relationships between functioning systems and the mechanisms in virtue of which they function. I tentatively suggest further that the avenue that one prefers may turn on whether one expects the functional architecture of the brain to be primarily componential and hierarchical (Craver 2008; this collection) or typically more complex than that (Pessoa 2014).

Keywords
Constitution | Direction-selective ganglion cells | Enabling constraint | Explanation | Hierarchy | Levels | Mechanisms | Mechanistic explanation | Neuroscientific explanation | Starburst amacrine cells | Structure function mapping

Author

Michael L. Anderson
michael.anderson@fandm.edu
Franklin & Marshall College
Lancaster, PA, U.S.A.

Commentator

Axel Kohler
axelkohler@web.de
Universität Osnabrück
Osnabrück, Germany

Editors

Thomas Metzinger
metzinger@uni-mainz.de
Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt
jennifer.windt@monash.edu
Monash University
Melbourne, Australia

1 Introduction

In my target article (Anderson this collection), I argued that the complexity of the function-structure relationships that give rise to direction selectivity in Direction-Selective Ganglion Cells Direction-Selective Ganglion Cells (DSGCs) and in the dendrites of Starburst Amacrine Cells Starburst Amacrine Cells (SACs) represent a challenge to componential mechanism as currently formulated (Craver 2008). First, I argued that distinguishing between the system S that exhibits the target phenomenon ψ, and the mechanism M in virtue of which it ψs allows one to paint a more nuanced picture of the various ways entities can be organized so as to give rise to observed function. Second, I suggested that the function-structure relationships in these particular cases appeared to violate the bottom-up hierarchical assumptions at the center of the componential mechanistic framework, which requires that the components of M in virtue of which a system exhibits ψ are at a lower level of organization than S. In the cases under

discussion, I argued that some parts of the mechanism in virtue of which SAC dendrites function are at a *higher* level of organization than the dendrite, and that parts of the mechanism in virtue of which DSGCs function are at the *same* level. Moreover, I noted that in neither of these cases were all the entities that constituted M constitutive parts (components) of S.

To accommodate such cases, I recommended extending the notion of mechanistic *constitution* with the notion of an *enabling constraint*: mechanisms, we should say, enable function in systems by changing the relative probabilities of functional outcomes of activity in S. I suggested that this change would allow us to more accurately characterize the variety of structure–function relationships in the brain (and in other complex systems). However, in his commentary on my article (Kohler this collection), Axel Kohler argues that such an extension is unnecessary, for in fact the componential mechanistic framework of Craver and Bechtel (Craver 2008; Craver & Bechtel 2007) can accommodate these cases.

Kohler is correct. The extension is strictly speaking unnecessary, and componential mechanistic explanation can offer one plausible characterization of function-structure relationships in these cases. In fact, it is probably the case that *no* example or set of examples *ever* forces one to give up on an explanatory framework (certainly not one as well-motivated and useful as componential mechanism). What examples such as these *can* do, however, is illuminate the potential *advantages* of a new approach, and I would like to use the opportunity offered by this reply to reiterate what I take some of those advantages to be.

2 Three possible system-mechanism relationships

In my target article (Anderson this collection) I suggested that once one distinguishes between the system S that ψs and the mechanism M in virtue of which it does so, it is easy to see that there are three possible relationships between M and S. First, the components of M can all also be components of S, such that M is a relevant sub-component of S. Let's call this relationship R1. A relationship of type R1 obtains between the drive-train of an automobile and the automobile as a whole. Second, (R2), M and S can be identical. I can't think of an uncontroversial example of this relationship, and imagine that such a case is relatively rare. Third and finally, (R3), M and S can cross-cut in various ways, sharing some but not all of their parts. In my view, for instance, it is the neuron the fires an action potential, but not all of the entities that comprise the mechanism for generating action potentials are also part of the neuron. For example, the ions in the extracellular fluid that are crucial for establishing the membrane potential are not part of the neuron, although they are clearly part of the mechanism. Similarly, I argued in my target article that in the case of direction-selectivity in SAC dendrites, although it is the dendrite itself that is directionally selective, many of the parts of the relevant mechanism are not in fact parts of the dendrite. Moreover, in the case of DSGCs, the cell and the mechanism in virtue of which it is direction-selective share at most *one* part: the synapse between the SAC dendrite and the DSGC.

One advantage of making these distinctions, I believe, is that it allows one to see quite clearly when top-down constraints are responsible for function, as I argued is the case for direction selectivity in SAC dendrites. But Kohler suggests that appearances may be misleading here. In fact, he argues, we should "reconstitute the phenomenon" by recognizing that the relevant direction-selective system is *not* the SAC dendrite, but is rather the dendrite + the non-dendritic elements of the mechanism, including other SACs. This larger system can be then be treated within the standard framework of componential mechanism. We can call this approach to addressing these sorts of cases "the Kohler strategy".

As I noted in my target article, the Kohler strategy is certainly open to the mechanist. It does, however, have the following effects. First, it tends to make the systems of the brain to which functions are attributed relatively *larger* and more diffuse, which arguably reduces preci-

sion. Second, it would in effect turn all apparent instances of R3 into instances of R2.[1] I noted above that I thought the class of R2 would be small. If I am right about the prevalence of R3 functional relationships in the brain, then this strategy would make R2 very large. But it would do so essentially by legislation, as a way of preserving the universal applicability of the componential mechanist framework. How forced this appears will depend on how closely one believes the guiding assumptions of that framework match the architectural facts of the brain. We will return to this last point after reviewing some of the considerations that appear to favor the Kohler strategy.

3 Motivations for the Kohler strategy

Kohler maintains that actual scientific practice in fact supports the Kohler strategy. Exhibit A in his argument is a recent article (Kim et al. 2014) detailing part of the mechanism for visual motion detection. Kohler reproduces a figure depicting their model, and argues that the inclusion of the distributed network in the model suggests that the authors are strictly speaking attributing function to the whole system as depicted:

> Although it is true that investigators sometimes refer loosely to local elements as displaying a certain characteristic, the corresponding detailed and extended accounts of direction selectivity give credit to the distributed nature of the relevant systems that figure in explanations. (Kohler this collection, p. 6)

I agree that this is one possible interpretation of the practice. But here is another: these scientists are distinguishing between the system that exhibits the phenomenon and the mechanism that produces it, and are open to different sorts of relationships between them. Consider the following from the paper Kohler discusses:

> Research on [the visual detection of motion] has converged upon the SAC. An SAC dendrite is more strongly activated by motion outward from the cell body to the tip of the dendrite, than by motion in the opposite direction. *Therefore an SAC dendrite exhibits DS*, and outward motion is said to be its 'preferred direction'. Note that it is incorrect to assign a single such direction to a SAC, because each of the cell's dendrites has its own preferred direction. DS persists after blocking inhibitory synaptic transmission, when the only remaining inputs to SACs are BCs, which are excitatory. As the SAC exhibits DS but its BC inputs exhibit little or none, *DS appears to emerge from the BC–SAC circuit*. (Kim et al. 2014, p. 331; emphases added)

Far from seeming loose, the attribution of direction-selectivity to the dendrite appears to me clear and precise. Moreover, note that in the final sentence quoted above, the attribution of direction-selectivity to the cell is reinforced, even in the context of a reference to the mechanism as the "BC-SAC circuit". Indeed, I would argue it is natural and permissible to gloss the last clause in the following way: "DS *in the dendrite* appears to emerge from the BC-SAC circuit." On this reading, of course, the authors of this article would be proposing an R3 functional relationship such that parts of the mechanism are on a higher level of organization than the system exhibiting the phenomenon.

That these authors are open to R3 relationships of various sorts appears to be reinforced by a line later in the paper:

> Previous research suggests that On–Off direction-selective ganglion cells *inherit their DS from SAC inputs* owing to a strong violation of Peters' rule. (Kim et al. 2014, p. 335; emphasis added)

Here again we see the same pattern: a clear attribution of direction-selectivity *to the DSGC* in

1 Actually, there are some questions here, for there seem to be *obvious* instances of R3 with which the mechanist is and should be entirely comfortable, e.g., the neuron and the mechanism of the action potential. So presumably this strategy would be employed *only* when the relationship appeared to violate the "lower-level entity" constraint. I've not the space to pursue this further here, so will note only that *selectively* pursuing the Kohler strategy would need separate justification.

the same sentence as a reference to the distal mechanism (the SACs), in the context of what is obviously an R3 relationship between system and mechanism. Thus, while I agree that the Kohler strategy is viable, I don't see that consideration of scientific practice forces us to adopt it, or even necessarily favors it.

So what might be other reasons for adopting the Kohler strategy over extending mechanism to include enabling constraints? As I mentioned at the end of the previous section, the matter might come down to how closely one thinks the architectural facts about the brain match the guiding assumptions of the componential mechanist framework. If one expects that the brain is at root a decomposable or nearly-decomposable system of well-defined interacting components, then componential mechanism does indeed seem like a very appropriate framework for capturing at least the majority of its functional relationships (with the few exceptions to be dealt with perhaps as secondary elaborations or special cases). If, however, one takes seriously the notion that the brain is a massive network marked by multiple, nested, cross-cutting, dynamic hierarchies interacting in bottom-up, top-down, feed-forward and feedback fashions (Pessoa 2014), then one might wish for some of the explanatory flexibility that the notion of enabling constraints appears to offer. I, of course, am in this latter camp (Anderson 2015).

4 Conclusion

As Kohler correctly points out, it is possible to accommodate these complex cases of function-structure relationships within the componential mechanistic framework, by reconstituting the phenomenon and ascribing function to the whole mechanism that produces it. I have tried to indicate what I think some of the costs are to the Kohler strategy, including an apparent conflation of R2 and R3 functional relationships and a potential loss of grain in our ascriptions of function to structure. For some, paying these costs will be preferable to the proposed alternative, which might appear to require the admission of spooky top-down causes into our ontology.

For those who instead want to maintain the greater attributional specificity that appears to conform to scientific discourse, and in the current case to explain direction selectivity *in the SAC dendrite*, then I would argue that the most promising strategy is to recognize the ways in which functional parts (including networks) can impose constraints on other functional parts, at whatever relative level of organization. Adopting this strategy will of course focus attention on the nature of these constraints, whether bottom-up, top-down, or synpedionic. I would hope that the careful study of such R3 relationships as those showcased here would result in a better understanding of the varieties of causal interactions in complex systems.

References

Anderson, M. L. (2015). *After phrenology: Neural reuse and the interactive brain.* Cambridge, MA: MIT Press.

——— (2015). Beyond componential constitution in the brain: The case of starburst amacrine cells. In T. Metzinger & J. M. Windt (Eds.) *Open MIND.* Frankfurt a. M., GER: MIND Group.

Craver, C. F. (2008). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience.* Oxford, UK: Oxford University Press.

——— (2015). Levels. In T. Metzinger & J. M. Windt (Eds.) *Open MIND.* Frankfurt a. M., GER: MIND Group.

Craver, C. F. & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, *22* (4), 547-563. 10.1007/s10539-006-9028-8

Kim, J. S., Greene, M. J., Zlateski, A., Kisuk, L., Richardson, M., Turaga, S. C., Purcaro, M., Balkam, M., Robinson, A., Behabadi, B. F., Campos, M., Denk, W., Seung, H. S. & the EyeWirers, (2014). Space-time wiring specificity supports direction selectivity in the retina. *Nature*, *509* (7500), 331-336. 10.1038/nature13240

Kohler, A. (2015). Carving the brain at its joints: A commentary on Michael L. Anderson. In T. Metzinger & J. M. Windt (Eds.) *Open MIND.* Frankfurt a. M., GER: MIND Group.

Pessoa, L. (2014). Understanding brain networks and brain organization. *Physics of Life Reviews*, *11* (3), 400-435. 10.1016/j.plrev.2014.03.005