A CONTEXTUALIST APPROACH TO FUNCTIONAL LOCALIZATION IN THE BRAIN

Daniel C. Burnston

Tulane University, Philosophy Department

Forthcoming in Biology & Philosophy. Please cite final version, DOI: 10.1007/s10539-016-9526-2

1. Introduction

All things being equal, we should want our theories to portray the objects of study as accurately as possible. This much sounds unobjectionable. However, certain kinds of properties that objects can exhibit seem to conflict with our norms for good theories, such as generality and predictive power. One such worrisome property is that of *context-sensitivity*. It is often assumed that context-sensitivity is incompatible with powerful theory, and that good theories should explain it away, rather than embracing it. Here I am interested in functional explanation in neuroscience, and specifically the question of whether functions can be *localized* to parts of the brain. I will defend a contextualist perspective on localization, and argue that it can produce good functional explanations, despite these concerns.

Functional localization is one instance of a general explanatory strategy in the life sciences: that of decomposing a system into parts, characterizing the parts functionally, and explaining the system's behavior in terms of those functions. Functional decomposition of this sort is an important explanatory aim (Bechtel & Richardson, 1993; Simon, 1962). However, localization and context-sensitivity have often been considered incompatible, and in recent discussions arguments to this effect have explicitly invoked theoretical norms. If context-sensitivity is the rule, the reasoning goes, we will not have theories of functional localization that can meet our epistemic goals. We had thus best hope that some kind of theory can be adduced that rules contextual variation out of our functional descriptions.

While often given in empirical contexts, these arguments are predominantly philosophical, since they rely on the norms of good explanation. The dialectic has analogs in other debates in philosophy, for instance in epistemology and the semantics of natural language (Preyer & Peter, 2005). Contextualists in these debates are generally inspired by data which they suggest is best read as evidence for contextual variation in the truth values of statements. Anti-contextualists (often called "minimalists" or "invariantists"), while their views differ in many ways, are very often motivated by the perceived negative epistemic consequences of adopting a contextualist position. Consider the following quote:

> "Semantics is a discipline that aims to characterize systematically certain features of linguistic expressions and to do so in a way that captures general truths about languages, and not just truths about particular speakers in specific contexts." (Cappelen & Lepore, 2005, p. 58)

Here, the assumption is that embracing contextualism about semantics will lead to bad theories—theories that are not "systematic" in the sense relevant for semantics. Instead, semantic theories should posit one univocal semantic value for each term/statement. Very similar statements can be found elsewhere (MacFarlane, 2009, p. 246; Stanley, 2005, p. 222).

The analogous debate regarding functional localization is between *absolutists* and contextualists. I use the term 'absolutism' to refer to a class of views which, while differing in their details, share a central assumption about the kinds of theories that can succeed at localizing function. Absolutists think that univocal function ascriptions should be given to individual parts of the brain (referred to henceforth as "areas") and that all of a given area's activities will be explicable as performing that single function.

> *Absolutism*: For a given neural area A, there is some univocal description D, such that D explains the functional role of A's activity whenever A functions.

Absolutism is anti-contextualist, because it says that in any context in which it functions, A performs the univocal function ascribed to it. Modern proponents of absolutism hold it *despite* seeming evidence of contextual variation, for broadly the same reasons as anti-contextualists in philosophy of language. The epistemic consequences of contextual variation are thought to be too dire; we should thus seek ways of describing functions such that contextual variation can be explained away. A variety of proposals are then put forth for how to do so—i.e., for how to structure D so that it does not need to vary as context changes (Anderson, 2010; Bergeron, 2007; Price & Friston, 2005; Rathkopf, 2013).

It is not my concern in this paper to flesh out the different possible varieties of absolutism. I describe several possible varieties, and argue against each of them, in other work (Burnston, 2015, Burnston, Forthcoming; cf. footnote 14). My goal here is to give a formulation of contextualism, which is only beginning to emerge as a realistic option (Klein, 2012; McIntosh, 2004), and to examine the epistemic consequences for adopting it. I will contrast contextualism with only one traditional framing of absolutism, as it has been employed in perceptual neuroscience, and suggest that the switch from absolutism to contextualism in this instance does not have the supposed negative effects. While my arguments in this paper will thus not tell against every variety of absolutism, they do suggest that one of the major *motivations* for absolutism in general—the idea that contextual variation is incompatible with good theory—is flawed.

The main move the contextualist must make is to reconsider the role that we expect functional theories to play in our overall understanding of the brain. The absolutist assumes that, in order to be successful, a theory must cover all cases. The contextualist data suggest that such a strong notion of generalizability is untenable. What the contextualist must offer is a way of understanding theoretical success that systematically meets a weaker standard without collapsing into a series of one-off explanations. I suggest that such a view can be found by considering how functional theories *shape ongoing investigation*, and reconceiving their success conditions in those terms, rather than attempting to see function ascriptions as exceptionless truths in the domain of neuroscience. Once this move is made, I contend, contextualism does not undermine epistemic success.

In section (2), I lay out the desiderata that I think most theorists would agree are reasonable for theories of functional localization, and how absolutism hopes to meet them. In section (3) I introduce a case study, the “middle temporal area” (MT) of the visual cortex, that is a traditional success case for absolutist-style function ascription.[1] In (4), however, I discuss evidence which I claim strongly suggests that this area’s functioning is context-sensitive. Then, in section (5) I argue that contextualism can give some reasonable answers about how to meet the desiderata in the case of MT. In section (6) I consider some outstanding issues for a contextualist approach, and section (7) concludes.

2. The Desiderata and Absolutism

---

[1] Since I only discuss one brain area, it is of course possible that other parts of the brain will be less amenable to contextualist analysis. McCaffrey (2015) suggests such a view in his “functional heterogeneity hypothesis.” My goal here is to articulate contextualism for perceptual brain areas, which are important test cases for functional localization. How far the view extends is an empirical question, as McCaffrey notes.

The function of a brain area is what it contributes to the functioning of the system of which it is a part, and thus to brain functioning in general.[2] When we functionally decompose a system, we attempt to understand how the causal organization of that system brings about phenomena or behavior of interest.[3] In the case of neuroscience, these include perception, memory, linguistic behavior, etc., and aspects thereof. This is at least a vaguely "mechanistic" approach to explanation, although I do not wish to commit here to a particular view of what mechanistic explanation entails.

Brain areas are individuated by a combination of morphological and connectivity properties. An area usually has a particular architecture, and its cells often share connectivity patterns with other parts of the brain (Felleman & Van Essen, 1991). According to absolutism, each part of the brain performs a univocal function, and the functioning of the whole brain in any context is to be understood in terms of combinations of the univocal functions performed by each of its parts.

The desiderata on theories of functional localization are as follows. To be clear, by "theory" here I mean a meta-theory that describes how particular functional theories of particular brain areas should be structured. These particular theories should then meet the desiderata in describing their respective areas.

- *Decomposition*: A good theory should be able to distinguish areas from each other on the basis of their functional properties.
- *Empirical adequacy*: A good theory should be compatible with the range of available data about the areas being studied.
- *Generalizability and projectability*: A good theory should give functional descriptions that apply over a range of instances of functioning. It should also provide functional explanations in particular contexts that are relevant to contexts not yet explored.

---

[2] Out of the extant philosophical options, this characterization most resembles "systems" accounts of functions (Cummins, 1975). As Craver (2007) notes, systems accounts normally must be supplemented with an account of causal relevance, to show which properties of the parts are actual components of the causal organization that produces the explanandum. In systems neuroscience, it is assumed that the electrical signaling activity of a part of the brain is its relevant causal contribution, and I will adopt this assumption.

[3] Functional decomposition is not the only explanatory goal for which we might discuss the parts of a system. We might, for instance, want to explain the evolutionary history of a brain area, or its morphological properties. But this kind of explanation is at least a primary goal of neuroscientists involved in pursuing localization.

Decomposition, as discussed above, is the general epistemic project of which localization is an instance, so any theory of localization should give principles for distinguishing parts of the brain in terms of how they differentially contribute to its functional organization. Empirical adequacy is straightforward; it is the idea that our theories, to the greatest extent possible, should explain all of the data about the objects being studied. Generalizability and projectability are the analogues of the "systematicity" discussed in debates about semantics. The idea here is that the most powerful theory is one that covers many cases, and moreover tells us something about cases we have not yet explored. I have left the notion of projectability (the 'relevant to') vague for the moment, since contextualists and absolutists must read it differently, as will become clear.[4] For now, suffice to say that both those theorizing about functional decomposition, and those interested in mechanisms, often tout the norm that theories should be predictive and general(Price & Friston, 2005; Craver & Darden, 2013).

The appeal of absolutism lies in its straightforward answers for how to meet the desiderata of decomposition and generalization/predictability. Absolutism is atomistic: it assigns each part of the brain a single function. One can then understand psychological phenomena as occurring due to the causal interaction of atomistic functions performed by distinct parts of the brain. Moreover, absolutism, if successful, would guarantee *very* powerful theories in terms of generalizability and projectability. Since it is built into the success conditions of functional localization on absolutism that functions must not vary with context, a successful absolutist account will be able to generalize across *all* instances of functioning, and will tell us what will happen in the contexts we've yet to explore. Absolutists often state this ideal explicitly. Consider:

> "[A function] is whatever single, relatively simple thing a local neural circuit does for or offers to all of the functional complexes of which the circuit is a part." (Anderson, 2010, p. 295)
>
> "It is most useful to label a region with a function that explains all patterns of activation" (Price & Friston, 2005, p. 268)

The norm being stated here is that our theories should be able to give a single function ascription that applies whenever an area functions, for instance in the distinct "functional complexes"

---

[4] I take it that these are three core desiderata for theories of functional localization. I don't claim that they are exhaustive, but they seem to capture relatively accurately the concerns that theorists have expressed.

alluded to by Anderson. This is a very strong notion of generalizability. Genuine context sensitivity would prevent the possibility of such function ascriptions, since what function is being performed would change with the context. Rathkopf (2013) thus argues that "context sensitivity systematically prevents [function ascriptions] from explaining neural structure" (p. 3) and prevents our theories from being "complete" (p. 10).[5]

Contextualism, in opposition to absolutism, stresses the *multi-functionality* of brain areas. While multifunctionality has been posited by a variety of theorists in both neuroscience and philosophy (McCaffrey, 2015; Poldrack, 2006), the best way of understanding and describing it is up for debate. Contextualism embraces multifunctionality in two ways. First, it says that the *contribution* made by a particular brain area to the functional organization of the brain, and hence to producing psychological phenomena, varies with context. Second, it proposes that the best way to account for this variation is to posit functional theories that index function ascriptions to contexts, and that have an *open* structure. Theories that have an open structure admit the adding of functional properties with exploration of novel contexts, and I will discuss how this works in detail in section 6.

"Context" can refer to either "internal" or "external" context. Internal context consists in the goings-on elsewhere in the brain that co-occur with a particular area's functioning (Klein, 2012, McIntosh, 2004). I will be primarily concerned with external context, which for my purposes consists in the perceptual and behavioral circumstances in which an area functions. In particular, I will discuss perceptual circumstances and tasks involving perceptual judgment, but these are only particular examples out of the many types of contexts an organism encounters. Internal and external contexts are closely related, however, since contextual variation of function with external context must be mediated by changes in internal context. I will discuss this further in section 5.

A caveat should be entered here. I will focus on only one neuroscientific method and one level of analysis: electrophysiological research into the functioning of brain areas, where these are construed as comprising groups of cells whose physiological activities together determine the area's functional output. Not only do a wealth of other techniques exist in neuroscience, but

---

[5] To be fair, Rathkopf is concerned with a particular *form* of functional explanation, namely descriptions that explain why a certain anatomical structure is present. He admits that context-sensitive descriptions might be useful for other purposes.

recently theorists have questioned whether the brain area is the right level of analysis, preferring instead to talk of cell populations or networks. While I think there is still merit in analyzing the functions of individual brain areas, I do not assume that this is the only or the most important level of organization in the brain. If the conceptual points made here are valid, then they should apply, *mutatis mutandis*, to other ways of carving up units and analyzing their functions. Importantly, theorists interested in both the network- and population-levels of analysis have recently invoked the notion that the functioning of their preferred units is context-sensitive (Pessoa, 2014; Rentzeperis, Nikolaev, Kiper, & van Leeuwen, 2014). To the extent that functional units at other possible levels of organization are context-sensitive in function, then the analysis I provide here can potentially be applied to them as well.

3. Absolutism and Perceptual Neuroscience

In perceptual neuroscience, absolutism has traditionally been implemented in the view that our intuitive perceptual attributives—for individual features such as color, motion, and shape, and for objects instantiating those features—have dedicated parts of the brain in which they are represented. The standard explanation of the mammalian visual system, the "modular functional hierarchy" (MFH) view, posits that each part of the visual cortex represents one distinct type of feature (or object) information. The overall causal organization of the visual system, as well as how perceptual phenomena are constituted in the brain, is explained in terms of the passing of representations from lower levels of the hierarchy to higher ones.

I use 'representation' here in the neuroscientist's sense of "carrying functionally significant information about." The standard means of attributing functions on the MFH view is to seek a privileged information-carrying correlation between a particular type of feature in the world and the physiological responses of a brain area. This is done with electrophysiological recordings from single cells or groups of cells as animals view scenes with particular types of features in them. On absolutism, each functional part of the visual cortex will represent one and only one type of feature or object information.[6] The individual cells or groups of cells within the

---

[6] Establishing a privileged correlation is by no means an easy matter. Perceptual neuroscientists have developed a variety of sophisticated strategies for arguing that one unique type of information is represented by a brain area in all cases, often despite seeming influence of other types on its responses. I present below the data for contextual variation that I take to best avoid the strategies, but I will not make the full argument here. I discuss the strategies in detail, and give arguments against them for these cases, in (Burnston, 2015).

area will be functionally specified as representing different aspects of this same feature. For instance, in area MT, particular cells represent motion in different directions and at different speeds. Since all of the cells represent aspects of the single preferred feature of the stimulus, motion, the area's overall function is defined as representing that feature.[7]

MT is a traditional, textbook success case for absolutist-style localization. MT is in the "dorsal stream" of the visual cortex, which is standardly construed as in charge of determining the locations of objects (Mishkin, Ungerleider, & Macko, 1983). MT's motion-detecting function within this stream was established using electrophysiology in monkeys (Zeki, 1974), and later with PET studies in humans (Zeki et al., 1991). In Zeki's (1974) pioneering study, monkeys viewed either moving elongated bars or colors within the "receptive fields" of individual MT cells. A cell's receptive field is the area of the visual field in which appropriate stimulation can cause the cell to respond. When motion was present in their receptive fields, individual MT cells responded to motion in particular directions. But *no* MT cells responded to particular colors in stationary stimuli. This supported the view that different areas are "functionally specialized to analyze different features of the visual environment" (Zeki, 1978, p. 243), such that "at every area a different type of information is analyzed" (p. 428). Maunsell and Van Essen (1983a) showed further that MT cells have distinctive speed preferences. The overall idea, then, is that in any instance of MT functioning, some subset of its cells will be carrying information about motion, namely the ones sensitive to the directions and speeds of whatever motion is in the stimulus. The MFH view contends that representing motion in this sense is the univocal function of MT.

In addition to showing that particular information is represented in a physiological response, it is also important to show that that information has a functional effect on the system. Britten et al. (Britten, Newsome, Shadlen, Celebrini, & Movshon, 1996) developed a sophisticated way of doing so. They recorded from individual MT cells in monkeys while the monkeys performed a motion-judgment task, and showed that MT cell responses were *predictive* of the monkey's eventual decision. When MT cells with particular direction-preferences were

---

[7] Some areas are sub-divided into functionally specified parts. For instance, V1 is a single anatomical area, but comprises different parts specialized for representing wavelength, displacement, and oriented edges. V2 has a parallel internal organization. The hierarchy proposed by the MFH theory is posited to then separate these features into distinct areas corresponding to our standard perceptual attributives of motion, color, shape, etc. This sort of subdivision is entirely compatible with absolutism, so long as the distinct parts within the areas represent only one feature. Below I will explain that subdivision is not appropriate in the case of MT.

active, the monkey was significantly more likely to judge that the perceived motion was in the preferred direction. Thus, they concluded that MT signals contribute to the function of the visual system by conveying a motion representation.

The MFH theory ascribes different representational functions to the other parts of the visual cortex. Area V4, for instance, contains two distinct functional areas, one for representing object shape and another for representing color (Livingstone & Hubel, 1988). The functional representation at MT serves as input to further areas in the hierarchy in charge of representing more complex motion forms, and in charge of guiding eye movements. V4 provides input to areas in charge of object recognition. This kind of function attribution is extremely powerful, since it offers a clear picture of the representational and computational organization of the visual cortex. Moreover, it makes highly general predictions—for instance that in all contexts in which MT functions, it does so by conveying a motion representation. There is now significant evidence, however, that MT is *not* solely a motion-representing area. I discuss this evidence in the next section, and argue that it is best read as showing that what MT represents changes depending on the perceptual context.

4. Contextual Variation in MT Function.

A natural way to phrase a contextualist perspective in perceptual neuroscience is to reconsider the role that intuitive perceptual attributives play in function ascriptions. The MFH theory assumes that representations of perceptual features are psychological units that should each be instantiated uniquely somewhere in brain function. The contextualist perspective I will advocate argues that the standard attributives—features such as motion (displacement), depth (disparity), color (wavelength), brightness (luminance), shape, and the like—are best seen as *parameters whose variation determines the perceptual context*. All contextualist views require some notion of what changes when context changes. Contextualism in epistemology, for instance, is often cashed out via the idea that what changes with epistemic context is the *stakes* surrounding the knowledge claim being considered (DeRose, 1992). Linguistic contextualism takes things like speaker intentions, contrast classes, and the background knowledge of the listener to vary. On the variety of contextualism I am advocating, motion, depth, color, shape, etc., are all aspects that go into determining perceptual contexts. Individual contexts comprise particular values along these parameters, and their particular arrangements in the perceived scene. Looking for contextual variation in MT, then, requires looking for situations in which different combinations of parameters produce distinct functional contributions in its responses. This can occur in one of two ways. First, a change in external context could categorically modify the response to a given type of information. Second, new types of information might be represented depending on the context. MT exhibits both forms of contextual variation.

4.1. MT and Color

The traditional view of MT holds that it is "color-blind" (Livingstone & Hubel, 1988, p. 741)—that it does not have functionally significant responses to color. However, by the early 1990s some non-physiological results proved to be at least in tension with this assumption. Anatomical results suggested that there were connections between MT and areas more traditionally viewed as representing color. Moreover, a variety of psychophysical studies showed that humans have the ability to perceive motion defined *solely* by color (e.g., when the whole scene is equiluminant).[8] Dobkins and Albright (1994), and later Dobkins, Stoner, and Albright (1998), decided to test whether and in what way color information affects MT responses.

Importantly for the discussion later on, their initial study did *not* show an effect of color on what MT represents. They asked whether it was possible to perceive motion defined solely by color in a way that did not show any influence of color information *per se*. One such way would be to perceive the *border* between moving colors, and not the colors themselves. They constructed a red/green grating stimulus that moved in sequential time steps from left to right. This is shown in figure 1, below. In the figure, each row represents the entire grating at one time point.

---

[8] For a review of these results, see (Dobkins & Albright, 2004).

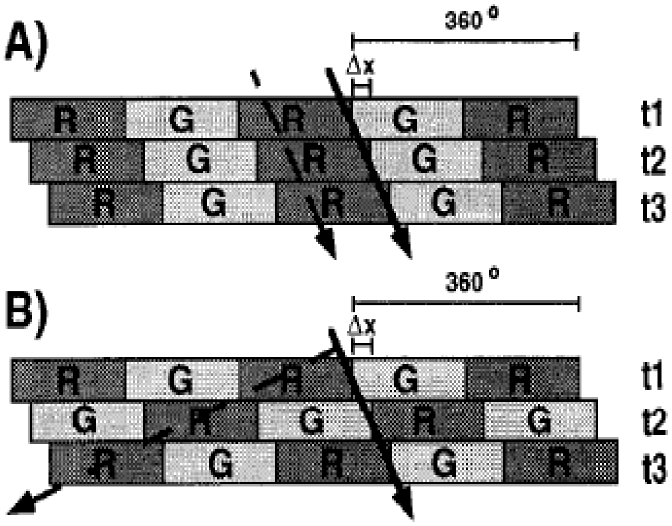

Figure 1. The stimuli from Dobkins and Albright's (1994) study. Patches labeled 'R' were red; patches labeled 'G' were green.

Normally, when a colored figure moves, the color-borders move in the same direction as the colors themselves, as shown in the top panel. However, Dobkins and Albright then varied the stimuli so that the colors in the grating would switch places at each time step, as shown in the bottom panel. In this stimulus, the direction of motion defined by color identity moves to the left (dashed line), and the border between colors moves to the right (solid line).[9] Dobkins and Albright showed that MT preferred the direction defined by the color border. Thus, color information *per se* was not influencing the response.

This result, however, led to a second study, which showed an effect on MT that could only be explained by a functional response to color identity. Dobkins, Stoner, and Albright (1998) constructed a stimulus consisting of two red/green gratings, with one moving across the other. Importantly, this stimulus can be perceived in one of two ways. It can be seen either as one grating moving across another, or as a unified plaid stimulus. The directions of these two motions will be different, as shown in figure 2 below.

---

[9] This is based on the attendant assumption that motion perception follows the border with the smallest displacement at each time step, which is independently well-established (Dobkins & Albright, 1994).

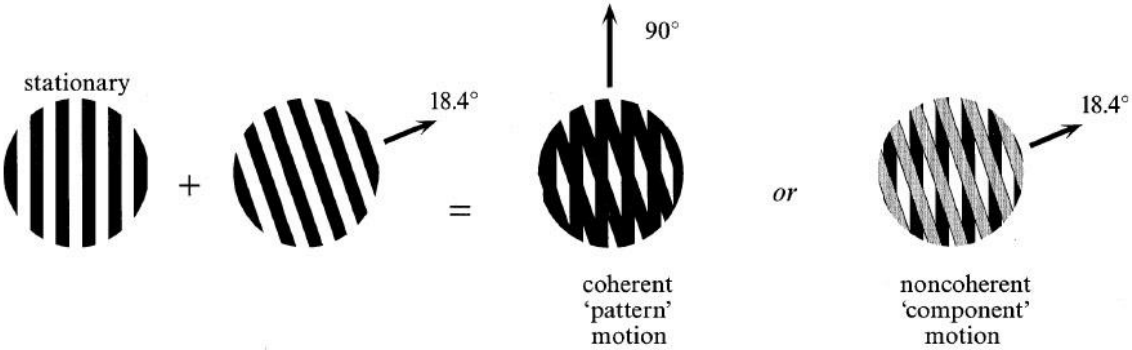

Figure 2. Component versus pattern motion in overlapping grating movement (the actual stimuli were red/green gratings).

Importantly, which direction gets perceived depends on the colors in the stimulus. If the greens in both gratings are bright, and the reds dark (or vice versa), then the "pattern" motion of a single plaid will be perceived (the 90° movement in figure 2). If the brightness of the colors is mismatched, with the green bright in one grating and the red bright in the other (or vice versa), then "component" motion of one grating moving across another will be perceived (the 18.4° motion). MT exhibited this same variation. For instance, if the two greens in the stimulus were equally bright, MT signaled motion in the pattern direction. If the greens were mismatched, MT signaled motion in the component direction. This effect, however, depends on the *identity* of the colors—both of the *greens* must be bright in order to get the effect. As such, the responses in MT are not only correlated with motion information, they are also dependent on the colors in the stimulus. Dobkins, Stoner, and Albright interpret the effect as one in which MT's responses reflect *both* the direction of motion and the way that colors are segmented in the scene. Specifically, information about color-identity modifies, categorically, the *way* in which a motion signal is represented, namely as either pattern or component motion.[10] So, motion information can be represented differently depending on context. Notice that this effect does not occur in *every* context, but only in particular arrangements of color, luminance, and motion. I will expand on this point further below.

---

[10] One can compare this with non-categorical shifts, such as those posited by "gain-increase" mechanisms of attention, which suggest that top-down attention does not change what is represented, but simply heightens the response to a particular represented property (Treue & Martinez Trujillo, 1999).

Now, MT represents both component and pattern motion in situations not involving color, so at best color-segmentation contexts must be added to the list of contexts in which MT represents pattern motion. A hard-line absolutist response here would suggest that this influence of color does not undermine the absolutist account, because it is still *motion* which is ultimately represented in either case. I now discuss cases showing that MT represents distinct *types* of information depending on context, thus rendering the hard-line response unavailable.

4.2. MT and Coarse Depth

Binocular disparity is an important depth cue. When the eyes are focused at a particular point in depth, called the "fixation plane," light reflecting from an object will hit the eyes at two distinct points, and the angle between the object and the two points determines where in depth the object is. A closer object will have a larger angle; an object farther away will have a smaller one. "Absolute" disparity of this sort is shown in figure 3 below.

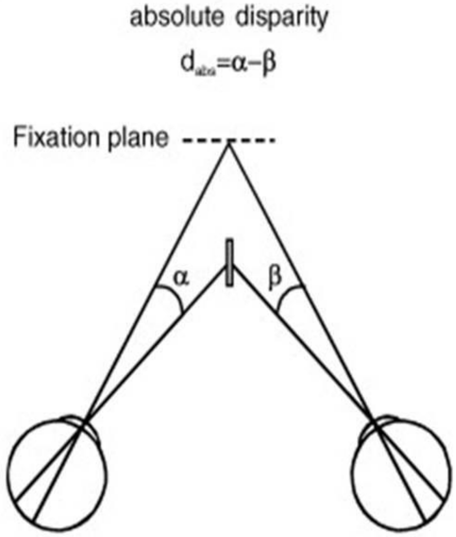

Figure 3. Binocular disparity. From Anzai & DeAngelis (2010). Moving the rectangle further from the viewer will result in a smaller angle. Disparity hence provides a depth cue.

While there were early suggestions that some MT cells respond to disparity (Maunsell & Van Essen, 1983b), the prevalence and functional import of these responses was not made clear until the late 1990s. DeAngelis et al. (DeAngelis, Cumming, & Newsome, 1998) showed that many MT cells had preferences for specific disparities, that these cells were regularly distributed throughout MT, and that they were organized in a map like fashion, with similar disparities represented in adjacent parts of MT. Moreover, they showed that the motion-selective and disparity-selective cells in MT overlapped—i.e., that the same cells which had the standard motion-responses in MT also had depth-responses. The study also gave an early indication that depth representations at MT are functional *independently* of motion. They had monkeys perform a depth-judgment task, where the goal was to guess the depth of an object as signaled by disparity. They then electrically stimulated MT cells with certain disparity preferences, and this biased the monkeys' judgments towards the depth entailed by that disparity. Importantly, this effect held when there was no motion in the stimulus. These combined results suggested that MT genuinely represents depth, and can do so independently of motion.

Later studies confirmed and extended these results. Uka and DeAngelis (2003) showed that roughly 94% of MT cells have significant disparity preferences, and that these preferences are *uncorrelated* with the cells' motion preferences (thus establishing that the two response properties are dissociable). This set of results strongly speaks against the temptation to try to save feature-specificity by further dividing up MT into smaller parts, each of which one might hope to be feature-specific (cf. footnote 9). Further, Palanca and DeAngelis (2003) showed that over a third of MT cells respond to depth in the *absence* of motion, in addition to their normal motion-signaling function. So, while there is *some* differentiation amongst MT cells in their relative responses to depth and motion, there are again not dissociable subsets of MT cells that uniquely represent depth and motion. MT as a whole comprises a group of cells that have both motion-and-depth representing functions, organized into distinct motion- and depth-representing maps. A number of further results established the relevance of depth representations at MT for depth perception. Using a methodology similar to the Britten et al. study discussed above, for instance, Dodd et al. (2001) showed that disparity signals in MT were predictive of depth judgments, concluding that "the perceptual role of MT is broader than previously assumed" (2001, p. 4820)—i.e., it represents depth in some contexts.

The dissociability of motion and depth representations in MT, as well as their independent contributions to motion and depth perception, present a compelling case for contextual variation. Unlike with the color case, there is no possible argument that disparity-processing in MT is solely for the purpose of specifying a motion representation, since depth-representation at MT occurs and is functional even in the absence of motion information. The overall picture, then, is that MT represents both motion and depth, depending on the context. If there is only motion in the stimulus[11], MT will represent motion; if only disparity, it will represent depth. If both motion and depth are present, then MT carries information about both of these properties. By the mid-2000s, a significant consensus had developed in the field that MT does indeed exhibit functional responses to both motion and depth. However, selectivity for absolute disparity is only relevant for *coarse* depth perception—for placing objects at rough levels in depth. *Fine* depth perception involves perceiving detailed depth properties, such as those exhibited by 3D objects oriented in depth. The consensus held that V4, not MT, represents this type of information. However, this assumption has recently been challenged by new explorations of particular stimulus contexts.

4.3. MT and Fine Depth

Fine depth can be signaled in a number of ways. One is by *relative disparity*--the difference in absolute disparity between two points in the visual field. Consider an object elongated in depth. Its closer and farther points will have two distinct disparities, and subtracting them will show the shape of the object as it is oriented in depth. Another is by motion. For instance, the different parts of a rotating object oriented in depth will appear to move at different velocities. Sanada et al. (Sanada, Nguyenkim, & DeAngelis, 2012) showed that, for stimuli exhibiting *tilt* and *slant*, defined *either* by motion or disparity, MT signals fine depth. Their stimuli are shown below, in figure 4.

---

[11] This would occur in a circumstance where the moving stimulus was precisely at the plane of fixation. See figure 3.

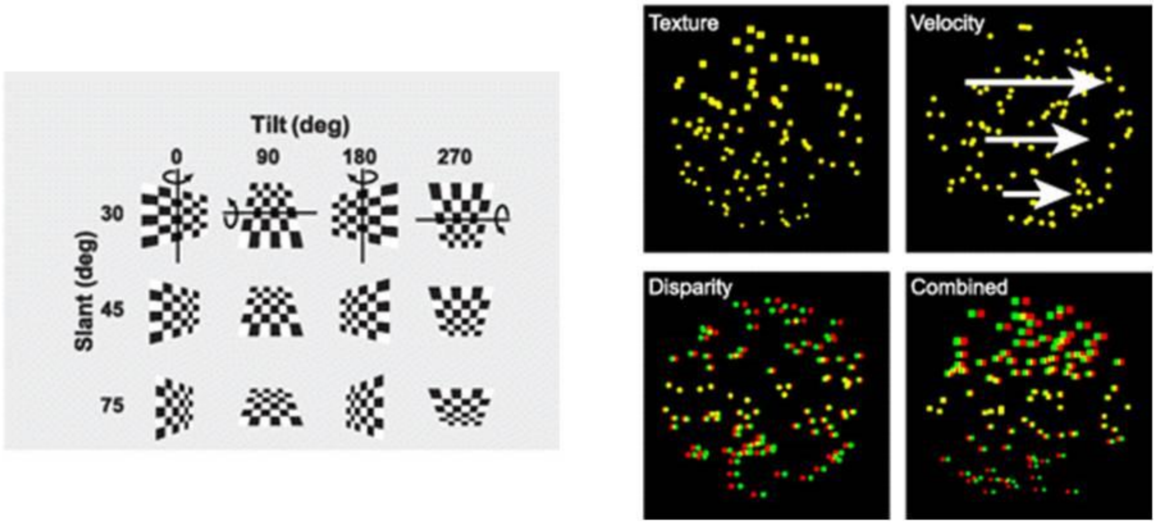

.

Figure 4. The tilt and slant stimuli of Sanada et al. (2012). Different tilts and slants are shown in the panel on the left. Implementations of particular tilts and slants via different cues are shown on the right.

MT cells, in the study, had specific preferences in their responses for particular tilts and slants. Again, these selectivities were in addition to, and distinguishable from, the cells' normal motion-receptive properties. Different cells, however, preferred tilt and slant cued by different types of information. Some responded to tilt and slant defined by only relative disparity. Some responded only to tilt and slant defined by relative velocity. Others responded to a combination of the cues. The conclusion, as with the original discoveries regarding motion, was that MT comprises cells with distinctive tilt and slant preferences. So, MT represents fine depth features in contexts involving tilted and slanted objects.[12]

4.4. Contextual Variation in MT

A contextualist approach to perceptual neuroscience takes perceptual attributives as parameters that vary to produce particular perceptual contexts. Evidence for contextual variation, then, would establish that what MT responses represent changes depending on the particular parameters in the scene, and their arrangement. This is exactly what the foregoing has shown.

---

[12] There have yet to be detailed studies, along the lines of Britten et al. (1996) or Dodd et al. (2001), establishing the functional use of this information. Given the progression of the field, however, and MT's detailed responses to tilt and slant, it is reasonable to expect that such studies will discover the functional usefulness of these responses.

When only motion is in the scene, MT represents its speed and direction. These motion representations can be categorically different—i.e., in terms of representing component or pattern motion—in situations where color information segments objects as joined or segregated. MT responds to coarse depth when it is cued by disparity information in the scene, and to the fine depth features of tilt and slant when those properties are cued by either relative disparity or relative velocity. Vitally, these informational responses are functionally relevant *for perception*. MT's varying responses affect what is in turn perceived, as measured in studies correlating physiological responses with perceptual judgments.

There are two kinds of arguments that might question a contextualist reading of this evidence, and thus perhaps motivate a return to absolutism. The first is logical. It says that we can define a single functional property, F, which is distinct from but extensionally equivalent to the full list of representational functions that MT plays. F, by definition, thus would not vary with context. The second is what we might call "dispositional." It suggests that MT always has the same set of dispositions—to represent motion, coarse depth, fine depth, etc. It is only that some of these dispositions get exercised in some contexts while others don't. Put another way, one might claim that MT just *always* represents every sort of information that it processes, but in some cases some of the parameters it represents take a null value (e.g., in stationary depth cases).[13]

Both of these moves are logically possible; however, each sits poorly with the explanatory aims of functional localization. Recall that we are not simply attempting to describe MT, but to show how its functional properties contribute to overall brain function and to psychological phenomena. As such, when we ask what MT does or contributes, our answer will have to vary depending on the perceptual and behavioral context we are discussing. This has important ramifications for understanding both how function is implemented in the brain, and how we need to investigate it. Suppose I am right that what MT represents varies depending on external context. We must expect that these different functions require the rest of the brain to

---

[13] Thanks to Nancy Cartwright for pushing me on the dispositional objection. There is in fact, a third strategy, which attempts to save absolutism by searching for a deeper functional principle that explains all of the specific functions that an area like MT performs. Several absolutists (including some of those cited in section 1) have argued that, rather than representing a specific type of information or contributing to a particular type of task, each neural area performs a particular *type of computation* in any context in which it functions. These views are still absolutist, in that they still posit univocal functions that are intended to account for all cases. I argue against this alternative form of absolutism in (Burnston, Forthcoming).

interact with MT differently as contexts vary. Several projects are beginning to investigate how the different contributions of particular brain areas are mediated by broader networks in particular contexts. In an early investigation, DeAngelis and Newsome (2004) studied cases in which both motion and depth signals are present in a scene. They were interested in how MT could underlie a motion-judgment task in these situations, where it was responding to both motion and depth despite the latter's irrelevance for the task. They showed a variety of possible "strategies," that could do this, including having only cells with weaker disparity preferences mediate the decision, or having cells with disparity selectivities corresponding to the depth of the moving stimulus do so. Intriguingly, they even suggested that individuals might learn different strategies in the course of learning the task. In either case, this shows the need for the broader network to account for contextual variation in MT responses in using these responses to produce behaviors.

Recently, theoretical advances in systems neuroscience have articulated more general mechanisms for contextually varying functional implementation, for instance through "multiplexing" (Watrous et al., 2015). Multiplexing is the differential modulation of particular neural signals by background activity, such as the local field potential (LFP). The LFP can be thought of as the sum of subthreshold activity in a population of cells, and recent results strongly suggest not only that a given spiking event can be modulated by this background activity, but can be modulated in different ways depending on behavioral circumstances. Multiplexing thus suggests that differently modulated signals can play different roles depending both on what information the signal carries and on the behavioral task. Canolty et al. (Canolty, Ganguly, & Carmena, 2012) describe the principle as one on which what function is performed in a particular context depends upon both internal and external factors. While it is early days for these projects, they explicitly recognize the need for the broader system to react differently depending on what information is being represented by a particular brain area in a particular context, and our empirical need to investigate how these processes work. Neither response on behalf of absolutism informs the need for this kind of theorizing. Moreover, the foregoing shows that it is not sufficient to call MT "multi-functional" and leave it at that. In order to understand how particular motion and depth perceptions come about, we have to understand both what information is being represented at MT in particular contexts, and how that information is processed by the visual system in the service of particular percepts and perceptual tasks.

In sum, what information is represented in MT responses varies depending on what kind of information is available, and these distinct representations contribute differentially to distinct perceptions and behavioral tasks. There is thus evidence of contextual variation in MT function. The remaining question is whether a theory that embraces this variation can meet the desiderata.

5. Contextualism

Contextualism builds contextual variation in at the ground floor of functional description.

> *Contextualism*: A neural area A exhibits an open conjunction of functional properties, D, where each conjunct in D is indexed to a type (or types) of context.

Contextualism is inspired by the kinds of evidence given above, and explicitly relativizes function ascriptions to contexts. By the conjunction being "open," I mean that new conjuncts can be added. Contextualism thus admits that we may discover new functional properties when we study contexts that we have not yet investigated. I now turn to showing how contextualism can meet the desiderata.

5.1. Decomposition and Empirical Adequacy

Decomposition is the goal of distinguishing parts of the brain according to their functional properties. Contextualism posits that what distinguishes areas from each other is the particular conjunction of informational functions that they perform, *and the contexts in which they perform them*. MT, for instance, can be said to represent at least motion, coarse depth, and fine depth depending on the stimulus context, and to respond to color when that information helps segment a moving stimulus. Read in this way, there is ample evidence that decomposition can be met. I will here briefly compare current physiological data on MT and V4 to show that, even if they are not individuated by each responding only to a particular feature, they still can be distinguished from each other functionally.

The absolutist explanation of the visual system is based on proposing that each visual area represents information about a particular perceptual feature. I have argued already that MT does not meet this kind of description. Now consider the summary description of V4's physiological properties from Roe et al. (2012): "V4 comprises cells that exhibit diverse receptive field preferences related to surface properties (color, brightness, texture), shape (orientation, curvature), motion and motion contrast, and depth" (p. 1). Offhand, not only does V4not appear to be a feature-specific area, but the features it responds to have significant overlap with the featural responses of MT—including color, motion, and depth.[14] So, it appears that *if* particular features are what we use to individuate the functions of areas, MT and V4 will not be clearly dissociable. If one insists on individuating functions in this way, while ignoring context, and if it is shown that MT and V4 both "represent motion" and "represent depth" in some sense, then one will believe that decomposition between these two areas is impossible. But contextualism does not individuate in this way. If one distinguishes areas according to both the features represented and the contexts in which they are represented, then the problem evaporates.

Consider motion and color first. MT cells respond to motion in particular patterns, directions, and speeds. While V4 responds to motion as well, it does so in distinct contexts. For instance, a significant percentage of V4 cells respond to what is known as "kinetic edges." If a stimulus moves in such a way that its different moving parts form an edge, then V4 will represent this information. Importantly, MT does not carry information about kinetic edges, and V4 does not carry detailed information about, e.g., the speed of the stimulus. A similar story can be told for color. The contexts in which MT responds to color, so far as the current evidence goes, are extremely limited: the responses are tied to color-identity signals that segment moving stimuli. If this turns out to be the extent of color influence, then this means that the color responses of MT are restricted compared to V4 responses, which reflect categorical color assignments and exhibit color constancy. That is, V4 cells exhibit a range of responses to contexts involving color information that differs from those in which MT responds. The fact that color affects both MT responses and V4 responses does nothing to undermine a functional distinction between them.

Now consider depth. MT represents coarse depth information in both moving and stationary stimulus contexts. V4, so far as the current evidence shows, does not exhibit functional responses to coarse depth. Both MT and V4 *do* have fine depth responses, but unsurprisingly, the contexts differ. MT, as discussed above, represents the fine depth properties

---

[14] To actually establish that V4 is not a feature-specific area would require a thorough analysis of the evidence, along the lines of what I gave above for MT. I won't do so here; while I believe such an argument can be made, my primary interest in this section is to analyze the consequences for decomposition *if* both MT and V4 are genuinely context sensitive in their functioning.

of tilt and slant. V4 cells, however, in addition to their fine-depth responses, also respond to *curvature*, whereas MT cells do not (Connor, Brincat, & Pasupathy, 2007). As Roe et al. point out, V4's combination of curvature selectivity and relative disparity selectivity allows for the representation of protuberances—parts of an object that extend towards the perceiver. Since MT cells are not curvature-responsive, they are not particularly good for this aspect of perception, even if they are helpful for overall tilt and slant perception. The difference between the two in terms of their depth responses is *not* based on responding to fine depth or not responding to fine depth, but instead on the particular combinations of fine depth with other features—i.e., fine depth occurring in different contexts.

It should be noted here that, while part of the move involves getting more fine-grained about the informational functions being described (e.g., representing "tilt" rather than "fine depth"), these fine-grained descriptions are capturing *actual contextual variation* in the functional properties of the areas being analyzed. The same MT cells that represent tilt and slant also have speed and direction preferences, and will represent that information when no tilt or slant is being perceived. V4 cells with functional responses to protuberances will represent only their preferred curvatures when only curvature information is available, and so on. Contextualism easily captures these patterns of variation, and builds them directly into the notion of decomposition. Read in a contextualist way, not only is there evidence for the decomposability of the visual cortex, but the evidence is preponderant.

With decomposition established, empirical adequacy comes along in relatively short order. Notice, first, that contextualism *retains* the empirical success of the MFH theory, in a more flexible way that can accommodate new discoveries about contextual variation. No one should disagree that MT, for instance, is *a* motion area, an *important* area for motion, or even that it is *more* heavily motion responsive, in general, than V4. These claims are still valid, and contextualism provides an easy reading for them, namely as distinctions in the *quantity* and *diversity* of motion versus color contexts that produce functional responses in the respective areas. It is important for the empirical adequacy of contextualism that the standard results in favor of the MFH view can be brought along in this way. However, since contextualism posits an open conjunction of functional properties, it admits the possibility that exploration in new contexts can reveal new functional responses within these areas. Thus, new discoveries can be incorporated into the functional description of MT without sacrificing empirical adequacy.[15]

5.2. Generalizability and Projectability

Admitting contextual variation into function ascriptions is what gives contextualism the ability to meet the empirical adequacy and decomposition desiderata. However, it is exactly this admission that anti-contextualists of different sorts find so epistemically troubling, because they think that it sacrifices generalizability and projectability. If we admit that function varies with context, we also admit that new contexts may reveal new functions, thus limiting the generality and predictive power of our current theories. Rathkopf (2013) calls this "unboundedness," and it is captured by the "open" structure of function ascriptions that I have proposed.[16] The contextualist *cannot* embrace the ideal of complete generalizability, since the possibility of open-ended contextual variation is built in from the outset. What the contextualist must do is argue that contextualism supports alternative reasonable notions of generalizability and projectability that are useful to those attempting to understand functional decomposition in the brain. That is, if we don't simply *define* generalizability and projectability in absolutists terms, then we can have a good understanding of how they can be met that legitimizes the project of localization. The key, I suggest, is to modify what we expect from theories. Rather than insisting that they be complete in the absolutist's sense, they should instead provide function ascriptions that (i) do not collapse into one-off explanations, and (ii) *structure and guide investigation*. Instead of predicting every outcome in advance, theories should provide us with resources for addressing new contexts of investigation and, perhaps, for making new discoveries.

Interpreted in this way, the real worry about new contexts, vis-à-vis our current functional descriptions, is not that our current views might be proven *false* in new contexts, but that they are *irrelevant* in new contexts. If there is genuinely no projectability—that is, if every

---

[15] Importantly, this does not mean that contextualism is "unfalsifiable" in any interesting sense. Posits about particular conjuncts and how they differentiate between areas are highly falsifiable. See the discussion below.

[16] Strictly speaking, a contextualist view needn't be unbounded. Functional properties could be context-sensitive and conjunctive even if it were possible to give a complete list of the conjuncts. There are two reasons for embracing unboundedness. First, since it is the strongest form of contextualism, avoiding the epistemic worries evinced by absolutists for this kind of view should heavily lessen the motivations for denying contextualism writ large. Second, I agree with Rathkopf that once we open our function ascriptions to contextual variation, it will be hard to *know* that we have a complete theory. At several times during investigation of MT, consensus views of its function have been overturned, and it is part of the appeal of contextualism that it has the epistemic modesty not to rule out this possibility*a priori*. Thanks to Gary Ebbs for pushing me to clarify this point.

context will require a new functional description—then we will have to start from scratch when considering new contexts. This is the worry that the contextualist finds compelling. I thus propose a reading of projectability that solves it without being tied to completeness. This established, it will be easier to see how a contextualist understands generalizability.

Contexutalist projectability is based on what I call "minimal hypotheses." A minimal hypothesis is basically a null hypothesis that is invoked when investigating function in new contexts. The idea is that in ongoing investigation one starts from an already accepted functional description, D, which has proven adequate for explaining the function of area A in some contexts (e.g., motion perception in MT). The default assumption is that the minimal hypothesis will continue to hold, unless results are adduced which cannot be interpreted according to it. While this is a simple enough idea, the importance of the minimal hypothesis is that it *sets the conditions* for whether one should modify D, and thus helps to determine *what kind of contexts one should investigate*. I think this idea has implicitly been at work in the advances in understanding MT function outlined in section 3, and that it can serve as the basis for a robust notion of projectability. In ongoing investigation, D can be modified one of two ways, corresponding to the two kinds of contextual variation discussed above. One can add a contextual index to a current conjunct via discovering a contextual difference in when that conjunct is instantiated. Or, one can add an entirely new conjunct, also indexed to specific contexts.

The motivation for Dobkins and Albright to undertake the study of color responses in MT was converging evidence from anatomy and psychophysics suggesting that color information might be able to affect MT responses. Hence, they undertook to find out whether there was genuine influence of color on MT, and if so, what sort. Their two studies were implicitly guided by the null hypothesis that MT only responds to motion, and never to color. They investigated two types of contexts, one in which color border could determine the relevant direction of motion, and one in which only color identity could. If color only ever affected MTvia the establishment of a color border, then there is no reason to incorporate color into the set of features to which MT responds. Put more generally, a particular feature parameter P never affects the area in question, or affects it in a way compatible with the minimal hypothesis, then no modifications involving P must be made to D. However, if the parameter affects responses in a context that is not so interpretable, one must incorporate P as part of the functional description of A in that context. Out of this process we do not come up with an attribution on which color *as such* affects MT, but one on which color identity in *specific contexts*—namely as a cue for object segmentation—influences MT in specific ways. In contexts where color identity segments moving stimuli into a unitary object rather than two separate ones, MT switches from representing component to pattern motion. So, color-segmentation contexts are one type of context to which pattern motion representation in MT is indexed. As discussed in the previous section, the specificity of the contextual attribution is part of what distinguishes MT from other areas that respond to color identity. Equally importantly, the researchers were *not* wandering in the dark in exploring the new contexts. The minimal hypothesis shaped the successive stages of experimental manipulation as well as the interpretation of the results. The result was scientific progress.

The color case exhibits the first type of contextual variation—MT motion responses can be categorically modified depending on information about color. However, one can tell a similar story about the depth case, which is one in which *different sorts* of information are represented depending on the context. The key set of manipulations and results regarding absolute depth involved manipulating disparity and motion in a variety of contexts, to show that MT responses to depth are *dissociable,* both informationally and functionally, from its responses to motion. This involved looking at stationary depth contexts, as well as depth contexts at a variety of motion directions and speeds. These manipulations produced the results that the vast majority of MT cells have coarse depth responses, that these responses are uncorrelated with their motion responses, and that a large number of MT cells respond to depth independently of motion. The manipulations were done specifically with the idea of manipulating parameters in such a way that MT responses would have to be interpreted as genuine depth responses, not motion responses.

Once these results became widely accepted, the field moved towards a conjunctive interpretation of MT function, including both motion and coarse depth in its functional description, but not including fine depth. In exploring whether this consensus was correct, the Sanada et al. study was designed to test whether responses to fine depth stimuli could be uncovered that were not interpretable in terms of coarse depth response. Since tilt and slant are inherently fine depth features, finding significant MT responses to these stimuli would not be interpretable in terms of coarse depth. Again, the discovery of responses to fine depth in these specific contexts helps distinguish the fine depth responses of MT from those of V4.

The final point to make, then, about the methodology employing minimal hypotheses is that it works in precisely the same way whether one is starting from a univocal function ascription or a conjunctive one. In either case, one has a known set of function ascriptions, and uses them as null hypotheses that suggest possible manipulations of contexts and ways and to interpret the evidence. From the contextualist perspective, not only is embracing a conjunctive view of function not disastrous for the progression of a science, but conjunctive theories are not significantly different from univocal ones in terms of their ability to shape investigation.

This view provides the resources for answering how contextualist views can be generalizable. Generalizability is a theory's ability to account for a range of instances. But unless generalizability *must* be understood in terms of the ideal of completeness, the current functional understanding is perfectly general, both as a whole and in terms of the individual conjuncts. Contextualism can perfectly well incorporate the fact that "MT represents motion" is a powerful generalization about its function. Further, MT represents coarse depth in *both* stationary and moving stimuli. The conjunct thus covers a wide range of potential instances, including multiple types of context, namely stationary ones and ones involving motion. None of the individual conjuncts is *completely* generalizable, of course, but insisting that a single ascription be completely generalizable, at this point, would simply be a restatement of absolutism, and thus beg the question. Contextualism asserts that the level of generality of each of the conjuncts in the theory is precisely a matter for empirical investigation—discovering the *limits* of generalizability for a certain conjunct is just as important an advance as discovering that it holds in many instances.

It is true that the contextualist, like contextualists in other debates, changes the standards for successful theorizing about the explanandum. And like other anti-contextualists, absolutists are unlikely to be satisfied with the result. A knockdown conceptual argument one way or another is unlikely to be forthcoming. What I have attempted to argue here is that there is a valid contextualist interpretation for how to achieve the desiderata, and thus that contextualism can resist the claim that it has disastrous epistemic consequences. If the argument is a good one, then a major motivation for denying contextualism, in this case at least, loses its bite.

6. Contextualism and the Progress of Neuroscience

One worry one might have is that I have relied too strongly on the notion of intuitive attributives as determiners of contexts, despite the fact that I have drawn on the practice of perceptual neuroscience in doing so. I have not, for instance, given a principled criterion of context individuation. Moreover, one might question the ecological validity of describing contexts in this way.[17] These concerns are important. I want to suggest, however, that they are not problems for contextualism per se, but rather point out avenues for theory development within the contextualist perspective.

To start off with, the demand for a single principled criterion of context individuation is too stringent. Recall that the parameters are only, in a sense, ingredients for determining contexts. They are not the contexts themselves. When perceptual neuroscientists investigate function, what they in effect do is look for differences in context—i.e., different combinations of the parameters—that make a difference in both the physiological responses under study and the perceptual outcomes for the organism. So, the relevant contexts are discovered, sometimes with considerable ingenuity, over the course of investigation. This is a relatively good description of the progress of MT research, but it resists an ultimate, single formulation for how to individuate contexts. Moreover, for different parts of the brain subserving different functions (e.g., more "cognitive" temporal and frontal areas), different sets of parameters and different types of contexts will need to be described. Attempting to provide a single criterion for context-individuation thus seems inadvisable overall, but lacking one does not prevent us from understanding how functional discoveries are made.

A similar line of reasoning applies to the worry about ecological validity. Using the parameters to investigate contexts is entirely compatible with there being a norm that the particular contexts we consider should be ecologically valid ones. Now, one might worry that the particular contexts I have discussed are not ones that obtain in real perceptual situations. I think this claim is overstated: it seems that there *are* perceptual circumstances in which we perceive stationary objects at particular depths, moving objects segmented primarily by color (e.g., in low-contrast environments), kinetic edges, protuberances, etc. This is certainly not to say that these contexts are *all* the relevant contexts, or that these parameters are all of the possible relevant ones. There are certainly many real-world situations—ones involving multi-

---

[17] Thanks to Ben Sheredos and Olivier Morin for pushing me on the ecological validity objection.

modal information, locomotion, etc.—for which these particular analyses would be insufficient. But that is no problem for contextualism. It just reiterates the fact that our current explanations are only partial. Consider multi-modality for instance. While I have assumed that MT is a *visual* area in describing its function, nothing about contextualism precludes it from occasionally contributing to multi-modal perceptual processing. Given the growing recognition of widespread multi-modal processing in the brain (Stein & Stanford, 2008), such a preclusion would be unwise. Not only does contextualism countenance the possibility, it provides normative guidelines for how to structure investigation with minimal hypotheses: given a putative influence of multi-modal information on MT, one should exhaust purely visual interpretations of its responses before expanding its function ascription to count it as a genuinely multisensory area (in certain contexts, of course).

To summarize: I admit that full answers to these issues await development within the contextualist perspective. They do not, however, pose conceptual problems at the outset. Neuroscience is an evolving discipline, and contextualism has the right kind of open-endedness to both structure investigation and incorporate new results. The details await determination.

7. Conclusion

The ideal of completely generalizable theories is a soothing one, epistemically speaking. But it is possible that certain systems just are not organized in a way that admits of its fulfillment. If we are studying one of those systems, or if the evidence strongly suggests that we are, it would be a shame if completeness were the only theoretical norm that could guide us in our deliberations. It would also be a shame if the failure of completely generalizable theory meant we needed to abandon other cherished epistemic aims, for instance understanding complex systems in virtue of the functions of their assorted parts. I have suggested that a fully developed contextualist view gives a reasonable way to fulfill the desiderata on theories of localized function, and does so in a way that is open to and informs new advancements in neuroscience. For studying an object as complex as the brain, that may have to suffice.

References

Anderson, M. L. (2010). Neural reuse: A fundamental organizational principle of the brain. *The Behavioral and Brain Sciences, 33*(4), 245-266; discussion 266-313.

Anzai, A., & DeAngelis, G. C. (2010). Neural computations underlying depth perception. *Current Opinion in Neurobiology, 20*(3), 367-375. Bechtel, W., & Richardson, R. C. (1993). Discovering complexity: Decomposition and localization as scientific research strategies: Princeton, NJ: Princeton University Press.

Bergeron, V. (2007). Anatomical and functional modularity in cognitive science: Shifting the focus. *Philosophical Psychology, 20*(2), 175-195.

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., & Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience, 13*, 87-100.

Burnston, D. C. (2015). *Perceptual Context and the Nature of Neural Function*. Doctoral dissertation, University of California, San Diego.

Burnston, D. C. (Forthcoming). Computational neuroscience and localized neural function. *Synthese*.

Canolty, R. T., Ganguly, K., & Carmena, J. M. (2012). Task-dependent changes in cross-level coupling between single neurons and oscillatory activity in multiscale networks. *PLoS Computational Biology, 8*(12), e1002809.

Cappelen, H., & Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*: John Wiley & Sons.

Connor, C. E., Brincat, S. L., & Pasupathy, A. (2007). Transformation of shape information in the ventral pathway. *Current Opinion in Neurobiology, 17*(2), 140-147.

Craver, C. F., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*. Chicago: University of Chicago Press.

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

DeAngelis, G. C., Cumming, B. G., & Newsome, W. T. (1998). Cortical area MT and the perception of stereoscopic depth. *Nature, 394*(6694), 677-680.

DeAngelis, G. C., & Newsome, W. T. (2004). Perceptual “read-out” of conjoined direction and disparity maps in extrastriate area MT. *PLoS Biology, 2*, e77.

DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research, 52*(4), 913-929.

Dobkins, K. R., & Albright, T. D. (1994). What happens if it changes color when it moves?: the nature of chromatic input to macaque visual area MT. *The Journal of Neuroscience, 14*(8), 4854-4870.

Dobkins, K. R., & Albright, T. D. (2004). Merging processing streams: Color cues for motion detection and interpretation. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1217-1228). Cambridge, MA: MIT Press.

Dobkins, K. R., Stoner, G. R., & Albright, T. D. (1998). Perceptual, oculomotor, and neural responses to moving color plaids. *Perception, 27*, 681-709.

Dodd, J. V., Krug, K., Cumming, B. G., & Parker, a. J. (2001). Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *The Journal of Neuroscience, 21*(13), 4809-4821.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex, 1*(1), 1-47.

Klein, C. (2012). Cognitive Ontology and Region- versus Network-Oriented Analyses. *Philosophy of Science, 79*(5), 952-960.

Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science, 240*(4853), 740-749.

MacFarlane, J. (2009). Nonindexical contextualism. *Synthese, 166*(2), 231-250.

Maunsell, J. H., & Van Essen, D. C. (1983a). Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *Journal of Neurophysiology, 49*(5), 1127-1147.

Maunsell, J. H., & Van Essen, D. C. (1983b). Functional properties of neurons in middle temporal visual area of the macaque monkey. II. Binocular interactions and sensitivity to binocular disparity. *Journal of Neurophysiology, 49*(5), 1148-1167.

McCaffrey, J. B. (2015). The Brain's Heterogeneous Functional Landscape. *Philosophy of Science, 82*(5), 1010-1022.

McIntosh, A. R. (2004). Contexts and catalysts: A resolution of the localization and integration of function in the brain. *Neuroinformatics, 2*(2), 175-182.

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in Neurosciences, 6*, 414-417.

Palanca, B. J. A., & DeAngelis, G. C. (2003). Macaque middle temporal neurons signal depth in the absence of motion. *The Journal of Neuroscience, 23*(20), 7647-7658.

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

Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences*, *10*(2), 59-63.

Preyer, G., & Peter, G. (2005). *Contextualism in philosophy: Knowledge, meaning, and truth*: Oxford University Press.

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

Rathkopf, C. A. (2013). Localization and intrinsic function. *Philosophy of Science, 80*(1), 1-21.

Rentzeperis, I., Nikolaev, A. R., Kiper, D. C. & van Leeuwen, C. (2014). Distributed processing of color and form in the visual cortex. *Frontiers in Psychology, 5,* 1-14.

Roe, A. W., Chelazzi, L., Connor, C. E., Conway, B. R., Fujita, I., Gallant, J. L., . . . Vanduffel, W. (2012). Toward a unified theory of visual area V4. *Neuron, 74*(1), 12-29.

Sanada, T. M., Nguyenkim, J. D., & DeAngelis, G. C. (2012). Representation of 3-D surface orientation by velocity and disparity gradient cues in area MT. *Journal of Neurophysiology, 107*(8), 2109-2122.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 467-482.

Stanley, J. (2005). Semantics in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 221-253). Oxford: Clarendon Press.

Stein, B. E., & Stanford, T. R. (2008). Multisensory integration: Current issues from the perspective of the single neuron. *Nature reviews. Neuroscience, 9*(4), 255-266.

Treue, S., & Martinez Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature, 399*(6736), 575-579.

Uka, T., & DeAngelis, G. C. (2003). Contribution of middle temporal area to coarse depth discrimination: comparison of neuronal and psychophysical sensitivity. *The Journal of Neuroscience, 23*(8), 3515-3530.

Watrous, A. J., Fell, J., Ekstrom, A. D., & Axmacher, N. (2015). More than spikes: common oscillatory mechanisms for content specific neural representations during perception and memory. *Current opinion in neurobiology*, *31*, 33-39.

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

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

Zeki, S. M. (1978). Functional specialisation in the visual cortex of the rhesus monkey. *Nature*, *274*(5670), 423-428.