*Philosophy of Science* (2025), **92**, 185–203
doi:10.1017/psa.2024.50

PSA PHILOSOPHY OF SCIENCE ASSOCIATION
PHILOSOPHY OF SCIENCE

ARTICLE

# Context and Neural Function

David L Barack

Departments of Philosophy and Neuroscience, University of Pennsylvania, Philadelphia, PA, USA
Email: dbarack@gmail.com

(Received 02 April 2024; revised 26 July 2024; accepted 23 September 2024; first published online 28 October 2024)

## Abstract

The debate over context and neural function pits contextualists, who maintain that neural functions vary by context, against invariantists, who maintain that they do not. In this article, I defend a moderate position that permits some context sensitive yet invariant functions. I distinguish performance from competence as well as different types of contexts, accounting for many cases of variability by context. The discussion suggests changes in networks can sometimes change the competences of parts. I conclude that some context sensitivity of function is consistent with invariant functions whereas other changes in context imply changes in the competence to perform functions.

## 1. Introduction

Do the parts of the brain have functions? Some philosophers maintain that the parts of the brain cannot be assigned functions because of their complexity, dynamics, or network properties (Silberstein and Chemero 2013; Stinson 2016; Weiskopf 2016; Rathkopf 2018; Silberstein 2021). Others, in contrast, maintain that the parts of the brain can be assigned functions (Posner et al. 1988; Craver 2007; Bechtel 2008; Kanwisher 2010; Barack 2019b). There is overwhelming empirical evidence that the functioning of parts of the brain change as the context changes (Anderson 2010; Klein 2012; Pessoa 2014; Barack 2019a; Burnston 2021). The question of what to make of these changes has emerged as a major point of disagreement in the debate amongst those who maintain that the parts of the brain can be ascribed functions (Burnston 2021). According to contextualist theories of brain function, the ascriptions of functions to the brain's parts change as the context changes (Klein 2012; Rathkopf 2013; McCaffrey 2015; Burnston 2016b, 2016a; Viola 2021). Contextualists are opposed to invariantists, who think that the brain's parts have context-independent functions. In this essay, I develop a middle-ground position, intermediate between contextualism and invariantism.

This debate in the philosophy of neuroscience over context-sensitive function has run together important distinctions. The performance of a function must be distinguished from the competence to perform it. Parts of the brain might be invariantly competent to perform multiple functions but only perform one, or none, depending on the context. I also distinguish between two different types of contexts

© The Author(s), 2024. Published by Cambridge University Press on behalf of the Philosophy of Science Association.

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

implicit in function ascriptions: part-contexts and function-contexts. By using these distinctions to resolve many cases of the functions of parts seemingly varying across contexts, I will argue that the truth value of ascriptions to a part of the brain of the competence to perform a function can be invariant across a range of changes in part-context. As a result, I defend a position that is more invariant regarding neural function ascription than much recent work in the philosophy of neuroscience. Other changes in contexts, however, do change the competence of parts to perform functions. Whether a change in context results in a change in competence depends on how the context has changed. By specifying the different ways that contexts can change, a clearer picture of the context sensitivity of function emerges. Hence, I defend an intermediate position—while some changes in context result in changes in the truth value of ascriptions of the competence to perform functions, other changes do not. My aim in this essay is a positive one, to present my approach and defend it against objections.

On the view outlined herein, behavior can be explained by ascribing fixed functions to parts of the brain that are sensitive to context, vindicating much neuroscientific practice. Neuroscientists ascribe functions to parts of the brain despite how plasticity, development, the environment, and other factors impact which function a part performs. But they also describe changes in function following lesions or damage as well as following normal developmental trajectories. A framework for rationalizing these function ascriptions is missing from the literature, and my approach begins to fill that gap. Further, distinguishing performance from competence and between different elements of contexts provides an understanding of how context sensitivity and invariant function ascription are conceptually distinct. Finally, the grounds for cognitive neuroscientific generalizations are clarified: The invariance of the competence to perform some function across a range of contexts, even if the part does not perform it in a given context, supports the generalizability of neuroscientific function ascriptions.

I start with a general statement of the form of function ascriptions. I then present context sensitivity for function ascriptions in cognitive neuroscience, illustrating this sensitivity with the case of motion and depth perception in area MT. Next, intrinsicality of neural function is introduced and critiqued based on the competence-performance distinction. A problem arises from cases of contrary competences, which I illustrate with the case of pleasure generation in the Nucleus Accumbens. To address cases of contrary competences, two types of contexts are distinguished: function-contexts and part-contexts. This machinery is put to work to deconstruct cases of contrary competences. Besides part-context and function-context, function ascriptions also refer to goals, which I illustrate with the case of the Superior Colliculus. The ensuing discussion points to ways that context might vary which impact a part's competences, such as some changes in networks to which a part belongs. I conclude with a brief discussion of some upshots for the debate over function ascription in the brain.

## 2. Context and the form of neural function ascriptions

Ascriptions of neural functions to parts of the brain have the following (or analogous) general form: the function of a $\phi$-ing of part P in system S in context C is to $\Phi$ for goal

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

Ψ. In other words, P Φ's by ϕ-ing, and in context C this contributes to, aids in, or helps S's Ψ-ing. The parts of the brain can be neurons, coalitions of neurons, areas, networks, and more; in short, however the brain is to be physically decomposed.[1] These parts behave in some way, the ϕ-ing, and this behavior has the function to Φ. This formulation is close to Wimsatt's, who says that "[a]ccording to theory *T*, a function of behaviour *B* of item *i* in system *S* in environment *E* relative to purpose *P* is to do *C*" (Wimsatt 1972, 32). Wimsatt's *P* is my Ψ; his *C* is my Φ; the system variables are the same; his *i* is my P; and his B is my ϕ.[2]

In the preceding ascription schema, by "context," I refer to the physical environment in a broad sense around some part of the brain at a time. Context includes the environment around the system of which the brain part is a part (corresponding to Wimsatt's *E*), but also the system of which the brain part is a part (the variable S) and the immediate physical environment around the part in the system, such as its neuromodulatory or network properties, or the various inputs, control signals, and so forth to a part. The context can also include historical or statistical properties. These different elements of contexts will be important in the following text.

Neural functioning is widely acknowledged to be sensitive to context (Anderson 2010; Pessoa 2014; Barack 2019a; Burnston 2021). To illustrate this context sensitivity, consider the example of area MT, which contains neurons whose firing rate (the ϕ-ing) signals motion direction and speed ($\Phi_m$) (Britten 2004) as well as depth ($\Phi_d$) (Parker 2007). MT seems to be genuinely multifunctional, as the same neurons respond to information about motion and depth in the same way (the neurons maximally respond for specific depths or motion directions, with decreasing responses for stimulus values that are farther from the preferred ones; call this "tuned activity"). This prevents breaking up the area into smaller function-specific subregions or assigning different functions to different aspects of the neuron's behavior, and there are also no viable candidates for some appropriately granular univocal function (Burnston 2016a). Granted that the function MT performs depends on the context, how should we describe its function?

The sensitivity of the function of parts to context can be described by the thesis that functions of parts of the brain vary relative to the context in which they function.

(Sensitivity): The function of a part P varies with respect to the context in which P functions.

[1] I assume that parts can be defined in some way independently of context. This could be challenged (see, e.g., Hutto et al. 2017), and so the decompositionalist is responsible for arguing that parts can be defined independent of context. My goal herein is not to make good on the context independence of parts but rather, granted parts, to explore the scope of context invariant ascription of context-sensitive functions.

[2] While Wimsatt explicitly couches his discussion in terms of teleofunctions, I will remain neutral on which sense of function is appropriate (e.g., causal-role function [Cummins 1975], selected-effects function [Neander 1991; Garson 2011], or some other sense). Similarly, I will remain neutral on the type of function; for example, some neuroscientists distinguish "behavior-function" from "operation-function" (Genon et al. 2018), and some philosophers distinguish "surface-function" from "deep-function" (Viola 2021), which may or may not map on to Ψ and Φ. The problem of context-sensitivity, though, may still arise for these types of function.

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

(Sensitivity) explains context sensitivity as variation in a brain part's function across contexts, and this variation could be described by referring to different contexts in the function ascription. For example, applying the preceding schema and assuming for simplicity that the goal Ψ is the same across functions, the function of tuned activity of neurons in MT in the primate in a motion context is to signal motion direction and speed for visual perception, and the function of tuned activity of neurons in MT in the primate in a depth context is to signal depth for visual perception. (Sensitivity) is accepted by the parties to the debate. Burnston defines context sensitivity as "[t]he function of a given P can change depending on events external to the part" (Burnston 2021, 746). Rathkopf says that "many neurofunctional hypotheses make essential reference to transient environmental contexts" (Rathkopf 2013, 2), consistent with (Sensitivity). Finally, de Wit and Matheson state that "context determines which computational operation that brain part implements in support of the behavior" (de Wit and Matheson 2022, 7). Granted the relevance of (Sensitivity), what are the implications for the debate between contextualists and invariantists?

The context sensitivity of function stated in (Sensitivity) seems to imply contextualism, that ascriptions of neural functions change depending on the context. Contextualism rejects intrinsicality, the thesis that the functions of brain parts are ascribed with limited reference to their context:

> (Intrinsicality): The function of a part P is ascribed with limited reference to the context in which P functions.

(Intrinsicality) is inspired by Burnston, who defines "intrinsic function" as "[t]he function of a part P is to be specified with limited reference to the system in which P functions" (2021, 745). McCaffrey defines a similar thesis as "the functions of neural structures can be characterized or understood in relative isolation" (McCaffrey 2023, 4). "Limited reference" or "relative isolation" means that parts of the brain are to be ascribed functions without regard to context except for reference to the resources that might be required for its functioning (Klein 2018) and background conditions like causal precursors (Burnston 2021). While MT's functioning requires oxygen and glucose, including them in the description of its function is permitted even though they are part of the context. In addition, developmental, genetic, and other causal precursors may also be included without violating (Intrinsicality) as regards the ascription of a brain part's functions.[3]

## 3. Competence, performance, and function ascriptions

(Sensitivity) and (Intrinsicality) are two main theses in the debate over context sensitivity of function. Contextualists like Burnston reject (Intrinsicality) and endorse

[3] There are narrow and wide readings of (Intrinsicality). The narrow reading is that the function is described with limited reference to context, such as claiming that the function of tuned activity of neurons in MT is to signal-in-context-$C_m$ motion direction and speed where $C_m$ is some motion stimulus context like moving dots. But then functions don't generalize across similar contexts and are individuated as finely as contexts. The wide reading is that the function ascription makes limited reference to context. I adopt the second reading herein.

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

(Sensitivity). To describe how function ascriptions can be invariant despite changes in contexts, I will consider a series of cases that illustrate overlooked distinctions in this debate. This will force revisions to both theses, with the result that versions of both are acceptable even if functions are context sensitive.

To start, (Intrinsicality) needs to be revised as an account of how functions for parts of the brain are ascribed. A part P can have function Φ consistent with (Intrinsicality) even if the function Φ of P changes by context because a part can have a function without performing it. Whether P performs Φ can depend on context; in some contexts, P may perform Φ, while in others, P may not. But these latter contexts do not imply that P does not have Φ merely because it does not perform it. Suppose, for example, that a monkey is presented with a stimulus that contains no visual motion. Then, tuned activity of neurons in area MT can't signal the speed and direction of the motion in the visual stimulus. However, this doesn't imply the region lacks that function.

This suggests two different intrinsicality theses, inspired by the competence-performance distinction from the philosophy of language (Chomsky 1965). The competence of speakers to speak a language is distinguished from their performances. Speakers might make errors when speaking, but this need not impugn their linguistic competence, which is based on the speaker's (implicit) knowledge of the language. I contend a similar distinction holds for the brain (cf. Millikan 1986, 59–60). A brain part may be competent to perform a function, in the sense of having the capacity to do so, without doing so, or even if performing the function should fail.

Drawing the competence-performance distinction for parts of the brain differs from the linguistic case in two important regards. First, unlike linguistic competence, the competence of brains parts to perform functions does not inhere in any kind of knowledge; rather, these parts have capacities that permit the performance of functions. Second, there is a difference in the source of the justification for assigning competence. In the case of speakers, linguistic competence is partly grounded in acceptability judgments, a source of norms unavailable in the case of neural functions. What could ground competence ascriptions to parts of the brain?[4]

I see the problem as one of specifying norms for the capacity of parts of the brain to perform functions.[5] A wealth of work in the philosophy of biology and mind on norms for function ascriptions can help here. There are at least four sources of such norms: evolution, development, learning, and control. Evolution in the form of selective pressures provides grounds for norms for functions because the present form of a structure is due to past behavior; it is because some part behaved in a certain way in the past that the part has the function it does even if that part never performs its function now (Millikan 1984). For example, suppose area MT has the function of signaling motion direction and speed because it was selected for such signaling. Then, area MT has the function to signal motion direction and speed even if,

[4] I thank an anonymous reviewer for pushing this objection and acknowledge there may remain concerns about the external standards relevant to such ascriptions.

[5] Note that function ascriptions are always made in some theoretical context, as pointed out in the function ascription schema in the preceding text, and so there will always be a theory-relative aspect of such ascriptions. This theory-relativity, however, need not call into question the resulting ascriptions or the theory independence of the norms.

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

*per impossible*, the monkey is never in a visual context that contains motion. If the monkey is never in that context, then MT can never perform the function, but it still has the competence to do so. Development, construed as partly due to evolution and partly due to environmental and ontogenetic constraints, can also be a source of norms; as development unfolds, the behavior of parts can acquire capacities to perform functions. Learning is still another source of norms. When learning occurs, there is selection for a particular pattern of behavior that results from activity patterns in the brain. Once some behavior is learned, then a brain part that performs a function for that behavior can fail without calling into question the capacity to give rise to the behavior. So, learning involves the establishment of a capacity in a part of the brain. Finally, another source of norms is control. Some functions may be the result of soft assembly (Clark 2007), neural assemblies constructed online by the brain for the task at hand, or transient assembly (Anderson 2014), preexisting latent dynamics put to novel uses such as for new transformations of signals. Various control processes aim for parts of the brain to perform functions by constructing such assemblies (Burnston 2022). In these cases, the assembly comes to have the capacity to perform some transformation (say) even if it sometimes fails to do so. Control, then, is a further source for norms. In sum, there are numerous sources of norms for ascribing the competence to perform functions to parts of the brain, including evolution, development, learning, and control.

Two intrinsicality theses can be distinguished based on the competence-performance distinction. The first thesis is that parts of the brain perform functions that are ascribed with limited reference to their contexts:

> (Intrinsicality$_{perf}$): The performance of some function of part P is ascribed with limited reference to the context in which P functions.

Discussions of context sensitivity often intend this performance sense of intrinsicality. For example, Burnston says that "[c]ontextualism 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" (Burnston 2016b, 542; italics removed). McCaffrey says "[c]ontextualism ... proposes that intrinsicality is false: neural structures perform radically different functions in different contexts" (McCaffrey 2023, 9). And de Wit and Matheson say that the function "a region performs ... is determined by the behavioral and neural context in which the region finds itself, and can shift when the context changes" (de Wit and Matheson 2022, 7).

How, then, does (Intrinsicality$_{perf}$) hold up? While resources or causal precursors are permitted in the ascription of the function Φ that a part performs, other variables are often cited in Φ's performance. For example, consider inputs. Tuned activity of neurons in MT might signal depth when there is no motion input to the area, say when there is none in the visual scene. Does including inputs in the ascription of the function to MT violate the restriction to limited reference in (Intrinsicality$_{perf}$)? Yes, because the inputs to MT are essential to specifying which function it performs. MT is multifunctional and may signal motion when there is motion input with uniform depth, or depth when there is depth input without motion, and so inputs can't be excluded from the ascription of the performance of a function. But if the inputs are

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

elements of the context, are not resources or causal precursors (e.g., they are not excluded by the limited reference clause), and the ascription of the function that is performed requires reference to the input, then ($\text{Intrinsicality}_{\text{perf}}$) is false because reference is made to the context when ascribing to tuned activity of neurons in MT the performance of some function. While here inputs are described with reference to the environment, a similar point holds even if the inputs are mathematical or computational (as Rathkopf 2013 maintains) because the area still receives input determined by the context. There are also other variables besides input that impact whether some brain part performs some function (Wajnerman-Paz and Rojas-Líbano 2022), notably including control signals (Burnston 2021). Based on these considerations, I conclude that intrinsicality for performing functions is false.

The second intrinsicality thesis is that the competence of parts of the brain to perform functions are ascribed with limited reference to their contexts:

> ($\text{Intrinsicality}_{\text{comp}}$): The competence of part P to perform some function is ascribed with limited reference to the context in which P functions.

Because ($\text{Intrinsicality}_{\text{comp}}$) concerns the competence to perform a function, not the contingencies about performing it, a range of elements of contexts, possibly including inputs or control signals, that only bear on performance are less relevant to ascribing a function.

How does ($\text{Intrinsicality}_{\text{comp}}$) fare when tested by the case of MT previously mentioned? The case of MT is consistent with the thesis because MT can have the competence to perform the functions of signaling depth or signaling motion even if the region performs different functions depending on the input (say). Burnston addresses a similar dispositional objection to contextualist accounts of neural function (Burnston 2016b). The dispositional objection is 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" (Burnston 2016b, 14). Burnston responds that the objection "sits poorly with the explanatory aims of functional localization . . . . Recall that we are . . . attempting . . . to show how [MT's] functional properties contribute to overall brain function and to psychological phenomena. As such, when we ask what MT . . . contributes, our answer will have to vary depending on the perceptual and behavioral context" (Burnston 2016b, 14). Applied to the distinction between competence and performance, the contextualist can similarly retort that while MT may have different competences, it is the performance of a function that contributes to overall brain function and to psychological phenomena.

This reply fails for the same, well-known reason that appeals to performance typically fail in functional explanation: Functional explanations are generalizations in which parts can fail to perform functions even when they have them (a familiar point from the literature on teleofunctions; see e.g. Millikan 1984; Neander 1991). Certainly, during a given episode, the performance of some function may explain some behavior. But explanations in neuroscience, like much scientific explanation generally, focus not just on episodes but also generalities, which requires generalizing over the performance of functions in episodes. Enough failures of a part may constitute evidence that the part does not have the competence to perform some

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

function. Because tuned activity in area MT always fails to (say) signal touch, the part does not have the competence to signal touch. However, the cases I'm considering are not cases of typical failure like that, but rather cases in which brain parts have some success at performing functions and that permit some failures of performance. For example, in contexts of motion with uniform depth, area MT typically or often or usually signals motion with tuned activity, which permits some cases of failure to signal motion. To make those explanatory generalizations requires a way for parts to have functions despite the occasional failure, such as an appeal to the competence of parts.

The distinction between competence and performance also helps to reconcile context sensitivity, multifunctionality, and invariantism. Parts of the brain may be invariantly competent to perform multiple functions, but which functions are performed depends on the context. That is, the invariantist can endorse both (Sensitivity) and (Intrinsicality$_{comp}$). This shift to consider function competence allows for invariant ascriptions of different functions. On a function performance approach, a part performing different functions for different goals will have variable functions. Intuitively, however, this is problematic because the mere fact of filling different roles for different goals shouldn't imply variable functions; that is, multifunctionality seems to be conceptually distinct from the debate over context sensitivity of function. Some brain part might always perform the same function for a given goal, say, but perform different functions for different goals. Inferring variability in the function of a part based on performance overlooks the constant contribution the part could make to each goal. But on a function competence view, the inference to variability does not follow. A part P can have the competence to perform a function $\Phi$ for $\Psi$ and the competence to perform a function $\Phi'$ for $\Psi'$ in all contexts. That is, a part can have two functions ($\Phi$ and $\Phi'$) in all contexts because we are working with a competence-based notion of function. In the function ascription, the context is quantified over.

However, there is the problem that some ascriptions of competences to parts of the brain are contraries. Consider this striking example: the change in neural function depending on context in the Nucleus Accumbens (NAcc) (Berridge and Kringelbach 2015). The rostral (front) and caudal (back) ends of the NAcc have distinct functions in producing valenced behavior (in the parlance of the field, "desire" vs. "dread"). In between those two ends, the NAcc contains multiple sites that generate valenced behavior. Surprisingly, however, the valence of the behavior generated by some site in NAcc can change depending on the environment of the organism. To illustrate, rats prefer quiet, home-like environments to noisy, party-like ones (Reynolds and Berridge 2008).[6] Reynolds and Berridge injected an agent in the rat NAcc that disrupts excitatory signals and then placed them in either home, neutral, or party environments while measuring appetitive (eating and drinking) or fearful (defensive treading[7]) behavior. In the home environment, more sites toward the caudal end of NAcc increase their firing rate for appetitive behavior, whereas in the party environment, more sites toward the rostral end of NAcc increase their firing rate for fearful behavior. The increase in firing rate of a given site might correlate with

[6] Literally. The experimenters played Iggy Pop in the noisy environment.
[7] Defensive treading is when the rat shoves bedding toward the sides of the enclosure.

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

appetitive behavior in the home environment but with fearful behavior in the party one. As they put it, "[r]emapping by environmental ambience was powerful enough to completely switch the appetitive versus fearful valence of ... motivated behaviors generated at many intermediate locations between the rostral and caudal tips" (Reynolds and Berridge 2008, 425). The NAcc, then, contains sites between its rostral and caudal extent that produce valenced behavior, but because the valence (desire or dread) changes depending on the animal's environment, sites within NAcc change function depending on context. That is, the function performed by the site within NAcc changes in different contexts, while being competent to perform either function appropriate to the context.

The case of NAcc challenges those who endorse both (Sensitivity) and (Intrinsicality$_{comp}$). The problem is that unlike MT, sites within NAcc are ascribed the competence to perform functions that are in conflict.[8] To present the strongest form of this objection, let the function of producing desire ($\Phi$) be the negation of the function of producing dread ($\sim\Phi$), and assume the two functions contribute to the same goal, to generate valenced behavior. Applying the function ascription schema to the NAcc, the function of the increase in firing rate of neurons in a site in the NAcc in the home context in the rodent is to produce desire for the goal of generating valenced behavior. But also, the function of the increase in firing rate of neurons in the same site in the NAcc in the party context in the rodent is to produce dread for the goal of generating valenced behavior. These two ascriptions match the change in the function performed depending on context. However, because of the shift to ascribing the competence to perform functions, not just the performances, contexts are now quantified over; in all contexts, the function of the increase in firing rate in that site in the NAcc in the rodent is to produce dread for the goal of generating valenced behavior, but also, in all contexts, the function of the increase in firing rate of neurons in the same site in the NAcc in the rodent is to produce desire for the goal of generating valenced behavior. These last two ascriptions follow from the new focus on the competence to perform functions. Both functions, of producing desire and producing dread, are ascribed across contexts despite being contraries. Call this the ascription of contrary competences.

Now, the ascription of contrary competences could be seen as unproblematic because the performance of the functions occurs in distinct contexts. However, I've been urging a shift from performance ascriptions to competence ones, so falling back on differences in contexts of performance to distinguish the ascriptions undercuts the motivation for my view. Further, the competence ascription remains that the function of the increase in firing rate of part P (a site within NAcc) in the rodent across contexts is either to produce $\Phi$ (desire) and to produce $\sim\Phi$ (dread), which would be contradictory, or to produce $\Phi$ or $\sim\Phi$, which would be vacuous. This might ground an objection to the competence view, that it allows either incoherent

[8] There are differences in the inputs and control signals when NAcc generates dread or desire (Berridge and Kringelbach 2015), raising concerns that differences in NAcc activity between the two functions can be used to distinguish the ascriptions. However, as Berridge (2019) notes, sites in NAcc "can ... switch between positive and negative affective valence simply owing to a shift in the emotional ambience of the outside environment, even when they are in the same neurochemical state" (p. 229).

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

(if contradictory) or uninformative (if vacuous) function ascriptions. I call this the problem of contrary competences.

## 4. Contrary competences and function claims

Neuroscientists ascribe contrary functions to the same brain part, as the NAcc illustrates, and any theory of neural function ascriptions will need to provide an account of such contrary ascriptions. If we use the strategy of appealing to competence to explain the context sensitivity of functions, then we risk incoherent or uninformative function ascriptions. I will now distinguish between two types of context sensitivity to resolve the problem of contrary competences.

The (Sensitivity) thesis above says that the function of a part varies with the context. But there are at least two types of contexts that can vary, part-contexts and function-contexts. While these two contexts often go together, they can sometimes come apart because the two types don't pick out the same contextual elements, as I illustrate in the following text.

To start, the function of the part can vary by the context of the function. When Φ-ing is context sensitive, the truth value of the ascription "a function of the ϕ-ing of a part P is to Φ in context C for Ψ" may change as the context changes. So, context sensitivity may refer to changes in the context of the Φ-ing.

(Sensitivity$_{fun}$): The function of a part P to Φ for Ψ varies with respect to the context of the Φ-ing.

On this first thesis, whether a part is ascribed Φ depends on the context of Φ; that is, whether P counts as Φ-ing depends on the context of the Φ-ing. I will call this the function-context.

Second, the function of the part can vary by the context of the part. When the part is context sensitive, the truth value of the ascription "in context C, a function of the ϕ-ing of a part P is to Φ for Ψ" may change as the context changes. So, context sensitivity may refer to changes in the context of the part to which some function is ascribed:

(Sensitivity$_{par}$): The function of a part P to Φ for Ψ varies with respect to the context of P.

On this second thesis, whether P is ascribed Φ depends on P's context. I will call this the part-context. In the following, I will distinguish part-contexts from function-contexts in function ascriptions by denoting the part-context at the start of ascriptions and function-context after the statement of the function; that is, ascriptions will have the form "in part-context C the ϕ-ing of a part P has the function to Φ in function-context C′ for Ψ."

To illustrate this distinction, consider the example of a piston (taken from Klein 2012). Pistons can speed up a vehicle by compressing a fuel-air mixture so it detonates, providing power to the crankshaft. But pistons can also be used as brakes to slow down a vehicle, using power from the wheels to compress air in its cylinder. Consider the following ascription:

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

1) In the context of accelerating, the piston has the function of compressing air in the context of accelerating.

This ascription is false because in the context of accelerating, the piston's function is to compress the fuel-air mixture. Now consider:

2) In the context of accelerating, the piston has the function of compressing air in the context of braking.

This ascription is true. In the context of accelerating, when the car is speeding up, the piston is competent to perform a range of functions, including the function of compressing air when in the braking context (and so the function-context does not match the part-context).

Now consider the case of NAcc again. We can hold fixed the part-context for evaluating "the function of the ϕ-ing of a part P is to Φ" while varying the function-context, and we can also hold fixed the function-context for Φ-ing while varying the part-context. For example, let the function-context vary. For some site in NAcc, we can wonder if increases in firing rate at that site has the function of generating desire in the home context. We can also wonder if increases in firing rate at that site has the function of generating desire in the party context. And so on. But we can also let the part-context vary. Then, we can wonder if in the home context increases in firing rate at that site has the function of generating desire. And we can wonder if in the party context increases in firing rate at that site has the function of generating desire. And so on. These two types of context sensitivity can be combined in single claims. We can wonder for some site in NAcc if in the home context increases in firing rate at that site has the function of generating desire in the home context. Or, if in the party context increases in firing rate at that site has the function of generating desire in the home context. Or, if in the home context increases in firing rate at that site has the function of generating desire in the party context. And so on.

Let's consider a site in the middle of the NAcc, restrict focus to the function of generating desire, and assume the home context and party context are exclusive. Suppressing the goal Ψ for clarity, there are at least four different ascriptions:

3) In the home context, increases in firing rate at the site in NAcc generates desire in the home context.

This is a true ascription.

4) In the home context, increases in firing rate at the site in NAcc generates desire in the party context.

This is a false ascription because in the party context, the site functions to generate dread.

5) In the party context, increases in firing rate at the site in NAcc generates desire in the party context.

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

This is false for the same reason as (4).

6) In the party context, increases in firing rate at the site in NAcc generates desire in the home context.

This ascription is true because even though the site in NAcc is in the party context, were the site in the home context, it would have the function of generating desire.

To elaborate, in the party context, when the animal is bombarded with noise, the site in NAcc will be competent to perform a range of functions. Is one of those competences the function of generating desire? Well, it depends. Consider (6). In the party context, the site does have the competence to perform the function of generating desire, when the site is in the home context. We set the context of the brain part to when the animal is surrounded by noise, and then we consider the context of the function. And that context (for this discussion) can be either when the animal is in its home cage or when it is surrounded by noise. We then set that function context to when the animal is in its home cage, which is counterfactual to the context set by the part. And, in that counterfactual context, increases in firing rate at the site does have the function of generating desire. Even though the site happens to be in the party context, nonetheless when in the home context the site still has the competence to perform the function of generating desire. Next, consider (5), and suppose we wonder if the site has the competence to perform the function of generating desire in the party context (and so the function context matches the context of the part)? The answer to that is clearly no—were the site to be placed in the party context, it would not have the competence to perform it. The increases in firing rate at in the site in the party context does not have the competence to generate desire; instead, it has the competence to generate dread.

Now, the truth values of (3)–(6) are determined by the function-contexts; while sometimes the function-context matches the part-context and sometimes not, it is the function-context that determines the truth of the ascription. So, it might be objected, what work is the part-context doing? In these four ascriptions, the change in part-context does not change the truth value of the ascription, which is determined by the function-context. In reply and as preview, this need not always be the case. Consider—counterfactually—the case in which a change in the part-context changes the competence of a site in the NAcc. Suppose that in the party context, the organization and connections of the NAcc change such that the increases in firing rate at sites in NAcc can no longer generate desire in any context. Then, (3) would be true but (6) would be false. This suggests a division of labor where part-contexts denote something like the mechanisms, pathways, and processes that endow a brain part with competences and function-contexts denote the other conditions under which competence is assessed. This also puts pressure on (Intrinsicality$_{comp}$), which I discuss in the text that follows. The key point presently is that two types of contexts are relevant to determining the truth value of function ascriptions, the part-context and function-context. In evaluating the truth of a function ascription, we can hold fixed the function-context and say whether the part is competent to perform the function. And that ascription might be true or false across all part-contexts, and so invariant.

Distinguishing these two types of contexts in function ascriptions resolves the problem of contrary competences. That increases in firing rate at the sites in NAcc in

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

one function-context generate desire and dread in another does not imply that the sites are ascribed contrary functions. Rather, what it means to ascribe the function of generating desire (dread) in the home context is different from what it means to ascribe the function of generating dread (desire) in the party context because ascribing the function of generating desire (dread) changes by context. That is, the competence ascription is that across contexts increases in firing rate at a site within NAcc in the rodent has the function of producing desire in the home context and producing dread in the party context. Because changing the function-context changes the function ascription, the two function ascriptions—the two ascriptions of what functions the sites of NAcc are competent to perform—are not incoherent or uninformative.

A simple objection to this account is that it implies radical contextualism about functions because what it is to function in some way changes from context to context. However, this is not accurate. The function ascription changes by function-context, but it is the same function across contexts. To ascribe to sites in NAcc the function of generating desire across some range of contexts is to claim that "the function of sites in NAcc is to generate desire" is true in those contexts. The function does not vary across those contexts even if the ascription does because of the different function-contexts.

A second challenge is posed by the ends of the NAcc. Unlike the locations in the middle, the caudal and rostral ends of the region do not change the sort of behavior that they generate; the caudal end generates dread and the rostral generates desire irrespective of the environment (home or party). What should we say of those sites? The form of the ascription—namely, "In context C increases in firing rate at the caudal site of NAcc has the function to generate dread in context C′" is true (*m. m.* for the rostral site)—is the same as for the sites in the middle of the NAcc, with open variables C and C′ unless quantified over or filled in. However, unlike the sites in the middle, the caudal site makes true a universally quantified sentence of the form "For all contexts C and all contexts C′, in context C increases in firing rate at the caudal site of NAcc has the function to generate dread in context C′," and similarly *m. m.* for the rostral site. The corresponding ascription "For all contexts C and all contexts C′, in context C increases in firing rate at a middle site of NAcc has the function to generate dread in context C′" is false. That's the difference between the ends and the middle of the NAcc. But note, this doesn't pose a challenge to variability in function-contexts—in fact, acknowledging function-context variability allows us to better specify the function of the ends of the NAcc.

There are several other advantages to distinguishing part-context from function-context. First, describing how context, part, and function interact is useful for understanding how some changes in context, such as changes in attentive state, do not imply changes in competence to perform functions whereas others, such as some following a stroke, do imply such changes. Second, disambiguating part-context from function-context rationalizes the statements and practices of working cognitive neuroscientists. Resolving this practice of function ascription rescues the statements and practices of working neurobiologists who otherwise could be convicted of making incoherent function ascriptions. Third, my account helps to make sense of how function ascriptions in neuroscience can generalize across contexts because brain parts can have competences that are exercised across contexts. Finally, fourth,

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

distinguishing part-context from function-contexts clarifies predictions. For example, Berridge and Kringelbach predict that "opioid recruitment in accumbens-pallidal hotspots … would plausibly generate pleasure 'liking'" (Berridge and Kringelbach 2013, 299). But this prediction assumes the system is not addicted to opioids, which can lead to changes in behavioral and accumbal responses to opioid administration (Koob 2020). That is, the prediction presumes a particular part-context, a nonaddictive one, and then describe what would occur following local opioid recruitment.

In sum, the invariantist can endorse context sensitivity and invariant function ascriptions. Further, by distinguishing part-context from function-context, they can endorse both (Intrinsicality$_{comp}$) and (Sensitivity$_{fun}$), thereby sidestepping problems from ascribing contrary competences.[9]

## 5. Goals, context, and competence

The foregoing discussion of counterfactual changes in the competence of NAcc points to a larger issue of which changes in context give rise to differences in the truth value of function ascriptions. While competences may be invariant to changes in inputs or control signals, other changes in context seem to violate (Intrinsicality$_{comp}$). Responding to some such cases will highlight the goal $\Psi$ in function ascriptions, whereas responding to others will require revising (Intrinsicality$_{comp}$).

Consider the superior colliculus (SC), a structure in the midbrain for orienting behaviors (Basso et al. 2021) that contains a topographic map of visual space and helps control both eye movements (Sparks 1986) and attention (Krauzlis et al. 2013). The SC is divided into superficial, intermediate, and deep layers, with neurons in the superficial layers showing responses to visual stimuli and neurons in the deeper layers showing motor responses (Basso et al. 2021). Increases in firing rate in intermediate layer visuo-motor neurons both correlates with eye movements and visual attention (Ignashchenkova et al. 2004; Ayar et al. 2023) and is causally implicated in both functions (saccades: Carello and Krauzlis 2004; McPeek and Keller 2004; attention: Müller et al. 2005; Lovejoy and Krauzlis 2010). The same intermediate SC neurons are hypothesized to underlie guidance of eye movements and changes in attention and are under the control of deep layer SC neurons (see Basso et al. 2021 for discussion).

For both eye movements and attention, and simplifying for this discussion, a given neuron in the SC increases in firing rate to help determine the endpoint of an eye movement and to help determine the locus of visual attention. Consider the original schema for function ascriptions again: the function of a $\phi$-ing of part P in system S in context C is to $\Phi$ for $\Psi$. Applying the schema to SC, a function of increases in firing rate of SC neurons in one context is to signal some location in the visual field for moving the eyes, whereas a function of increases in firing rate of SC neurons in another context is to signal a location for shifting attention. There may be some description that unifies these two functions of SC—such as, say, computing the relative value of orienting to a location in the visual field. However, the two types of

[9] The distinction between part-context and function-context also impacts the formulation of (Intrinsicality$_{comp}$). The relevant context for that thesis is part-context because the competence attaches to the part.

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

orientation are distinct, contributing to direction in the case of eye movements and contributing to divisive normalization in the case of visual attention, suggesting the goals may be distinct (cf. Barack 2019a). In what follows, I will assume the two functions subserve distinct goals.

This case of SC neurons is simply a case of multifunctional competences. For a given intermediate layer SC neuron, the neuron is competent to perform the functions of signaling locations in the visual field for eye movements or for guiding attention. But the case is helpful for highlighting the role of goals in function ascriptions. In ascribing these functions, there are differences in goal, that is, the Ψ. Goals can vary and that variability requires another variable besides part-context or function-context in the function ascription. So, for example and to make it fully explicit, the ascription may be: for some goal Ψ in some part-context C and some function-context C′, in C a ϕ-ing of part P has the function to Φ in C′ for Ψ. And that ascription of competence to perform a function might be true or false across all part-contexts, that is, all C, and so invariant.

While goals can vary without thereby causing problems for the invariantist, the case indicates a deeper issue. Some approaches to neural function take the network in which a part of the brain appears as relevant to ascribing functions to the part. As McIntosh puts it, "[c]ognitive operations are … the consequence of dynamic network interactions that depend on the processing demands for a particular operation" (McIntosh 2000, 864). More superficial neurons in SC are part of a network that helps move the eyes or the locus of attention. The contribution of increases in firing rates in those neurons to these goals involves many other parts, and the whole network determines which function is being performed; however, the part of that network that involves intermediate SC layers might invoke the same input, control signals, and signal processing. So, there might be some area A (such as thalamic nuclei or other regions), which sends inputs and control signals to the deep layer neurons in SC, which in turn control intermediate layer computations. Those in turn send outputs to some other area A′. But it is the interface between A, A′, and other parts of the brain that determines whether those increases in firing rate of SC neurons are for eye movements or for attention. That is, two different networks overlap in recruiting SC neurons and so the function of those neurons is determined by the network (cf. McIntosh 2000; Klein 2012; Stanley et al. 2019).

The case of variability across networks puts pressure on (Intrinsicality$_{comp}$), which states that the competence of a part is ascribed with limited reference to the context in which the part functions. Some changes in networks, which are elements of contexts, do not violate (Intrinsicality$_{comp}$). In those cases, the scientist or philosopher can allow networks to vary and remain committed to invariant functions. A different sort of change in network, though, would violate (Intrinsicality$_{comp}$). Some network might stand in relation to another network such that the part's membership in the first network prevents the part from being competent to perform some function that it is competent to perform when part of the second network. So, shifts in network membership are consistent with invariant competences in some cases but are inconsistent in others. In the case of SC, a change in network does not falsify either function ascription, but counterfactually, if the SC's being part of the attention network were to (say) sever the anatomical connection to the cranial nerves for eye movements, then the SC would no longer be competent to

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

help control eye movements, and that ascription would be false. Because changes in network can change the truth value of the ascription to brain parts of the competence to perform functions, a part's network needs to be included as a variable in the intrinsicality thesis. This suggests the following revision to (Intrinsicality$_{comp}$):

> (Intrinsicality$_{network\text{-}comp}$): In network N, the competence of part P to perform some function is ascribed with limited reference to the context in which P functions.

Any network variables that are part of the description of a context are intended to be fixed because the context lies within the variable's scope.

Parts of the brain appear in contexts with many different elements besides networks, including mechanisms (Wajnerman-Paz and Rojas-Líbano 2022), pathways, cascades, and processes (Ross 2021), neuromodulators (Marder et al. 2015), and presumably more (see also Burnston 2022). Which changes in context can modify or remove the competence to perform a function is presumably an empirical question, with some changes to these elements changing competences and others not. There are also historical or etiological elements of contexts (Khalidi 2017), the history of a part or system, and population elements of contexts (Viola 2021), the statistical reference class to which the system belongs, which may change competence ascriptions. To make good on the idea that the functions ascribed to parts of the brain vary by context, the key is to specify in which contexts (part-context or function-context) and for which goals function ascriptions vary and, for those contexts and goals, to further specify the changes in the elements (mechanisms, pathways, etiology, population, etc.) of the context with which the function varies. Variability in some of the elements of contexts may count toward variability in the part-context, specifically by violating (Intrinsicality$_{comp}$), but the contextualist can't simply assume so.

## 6. Conclusion

In this essay, I have evaluated the conflict between invariantists, who claim that parts of the brain have functions irrespective of context, and contextualists, who claim that parts of the brain have functions relative to a context. I have argued that both views can sanction some context-sensitive functions and defended a moderate position between invariantism and contextualism. Specifically, I argued that the truth value of ascriptions to have the competence to perform a function can be invariant across a wide range of changes in the context of the part to which the functions are ascribed but other changes imply a change in competences.

The first example of context-sensitive function ascription, the case of MT, does not pose a problem for my moderate position because the ascription may be about the part's competence, not performance. The second example of context-sensitive function ascription, the case of NAcc, also does not pose a problem for my moderate position despite the seemingly contrary competences because the ascription varies across function-contexts. The scientist or philosopher can adopt a view that permits function-context variability, and so context sensitivity, and yet invariant function ascriptions across part-contexts. Finally, the third example of context-sensitive

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

function ascription, the case of SC, highlighted how the function ascription varies across goals. My moderate position permits goal variability and yet invariant function ascriptions across part-contexts.

There remain challenging cases of genuine context-relativity of function, cases in which in one context, a part has some competence, but it lacks it in another. I briefly outlined how such a case might arise for changes in networks. In those changes in context, competences can be invariant, so long as the different networks don't change what a part is competent to do. Other network changes may change a part's competences. There are still other ways a part's competences may change by changing elements of contexts, such as after damage-related, learning, or developmental changes occur (De Brigard 2017). Hence, I defend a moderate position between invariantism and contextualism.

## References


Anderson, Michael L. 2010. "Neural Reuse: A Fundamental Organizational Principle of the Brain." *Behavioral and Brain Sciences* 33 (4):245–66. https://doi.org/10.1017/S0140525X10000853.

Anderson, Michael L. 2014. *After Phrenology*. Oxford University Press.

Ayar, Eve C., Michelle R. Heusser, Clara Bourrelly, and Neeraj J. Gandhi. 2023. "Distinct Context- and Content-Dependent Population Codes in Superior Colliculus during Sensation and Action." *Proceedings of the National Academy of Sciences* 120 (40):e2303523120. https://doi.org/10.1073/pnas.2303523120.

Barack, David L. 2019a. "Cognitive Recycling." *The British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axx024.

Barack, David L. 2019b. "Mental Machines." *Biology & Philosophy* 34 (6):1–33. https://doi.org/10.1007/s10539-019-9719-6.

Basso, Michele A., Martha E. Bickford and Jianhua Cang. 2021. "Unraveling Circuits of Visual Perception and Cognition through the Superior Colliculus." *Neuron* 109 (6):918–37. https://doi.org/10.1016/j.neuron.2021.01.013.

Bechtel, William. 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. Taylor & Francis.

Berridge, Kent C. 2019. "Affective Valence in the Brain: Modules or Modes?" *Nature Reviews Neuroscience* 20 (4):225–34. https://doi.org/10.1038/s41583-019-0122-8.

Berridge, Kent C., and Morten L. Kringelbach. 2013. "Neuroscience of Affect: Brain Mechanisms of Pleasure and Displeasure." *Current Opinion in Nneurobiology* 23 (3):294–303. https://doi.org/10.1016/j.conb.2013.01.017.

Berridge, Kent C., and Morten L. Kringelbach. 2015. "Pleasure Systems in the Brain." *Neuron* 86 (3):646–64. https://doi.org/10.1016/j.neuron.2015.02.018.

Britten, Kenneth H. 2004. "The Middle Temporal Area: Motion Processing and the Link to Perception". *The Visual Neurosciences*, edited by L. M. Chalupa and J. S. Werner, 1203–16. Cambridge, MA: Bradford Book, MIT Press.

Burnston, Daniel C. 2016a. "Computational Neuroscience and Localized Neural Function." *Synthese* 193:3741–62. https://doi.org/10.1007/s11229-016-1099-8.

Burnston, Daniel C. 2016b. "A Contextualist Approach to Functional Localization in the Brain." *Biology & Philosophy* 31:527–50. https://doi.org/10.1007/s10539-016-9526-2.

Burnston, Daniel C. 2021. "Getting over Atomism: Functional Decomposition in Complex Neural Systems." *The British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axz039.

Burnston, Daniel C. 2022. "Mechanistic Decomposition and Reduction in Complex, Context-Sensitive Systems." *Frontiers in Psychology* 13:992347. https://doi.org/10.3389/fpsyg.2022.992347.

Carello, Christopher D., and Richard J. Krauzlis. 2004. "Manipulating Intent: Evidence for a Causal Role of the Superior Colliculus in Target Selection." *Neuron* 43 (4):575–83. https://doi.org/10.1016/j.neuron.2004.07.026.

Chomsky, Noam 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.


https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

Clark, Andy. 2007. "Soft Selves and Ecological Control. Distributed Cognition and the Will: Individual Volition and Social Context." Cambridge, MA: MIT Press, 101–22.
Craver, Carl F. 2007. *Explaining the Brain*. Oxford: Oxford University Press.
Cummins, Robert. 1975. "Functional Analysis." *Journal of Philosophy* 72 (20):741–65.
De Brigard, Felipe. 2017. "Cognitive Systems and the Changing Brain." *Philosophical Explorations* 20 (2):224–41. https://doi.org/10.1080/13869795.2017.1312503.
de Wit, Matthieu M., and Heath E. Matheson. 2022. "Context-Sensitive Computational Mechanistic Explanation in Cognitive Neuroscience." *Frontiers in Psychology* 13:903960. https://doi.org/10.3389/fpsyg.2022.903960.
Garson, Justin 2011. "Selected Effects and Causal Role Functions in the Brain: The Case for an Etiological Approach to Neuroscience." *Biology & Philosophy* 26 (4):547–65. https://doi.org/10.1007/s10539-011-9262-6.
Genon, Sarah, Andrew Reid, Robert Langner, Katrin Amunts, and Simon B. Eickhoff. 2018. "How to Characterize the Function of a Brain Region." *Trends in Cognitive Sciences* 22 (4):350–64. https://doi.org/10.1016/j.tics.2018.01.010.
Hutto, Daniel D., Anco Peeters, and Miguel Segundo-Ortin. 2017. "Cognitive Ontology in Flux: The Possibility of Protean Brains." *Philosophical Explorations* 20 (2):209–23. https://doi.org/10.1080/13869795.2017.1312502.
Ignashchenkova, Alla, Peter W. Dicke, Thomas Haarmeier, and Peter Thier. 2004. "Neuron-Specific Contribution of the Superior Colliculus to Overt and Covert Shifts of Attention." *Nature Neuroscience* 7 (1):56–64. https://doi.org/10.1038/nn1169.
Kanwisher, Nancy. 2010. "Functional Specificity in the Human Brain: A Window into the Functional Architecture of the Mind." *Proceedings of the National Academy of Sciences* 107 (25):11163–70. https://doi.org/10.1073/pnas.1005062107.
Khalidi, Muhammad A. 2017. "Crosscutting Psycho-Neural Taxonomies: The Case of Episodic Memory." *Philosophical Explorations* 20 (2):191–208. https://doi.org/10.1080/13869795.2017.1312501.
Klein, Colin. 2012. "Cognitive Ontology and Region-Versus Network-Oriented Analyses." *Philosophy of Science* 79 (5):952–60. https://doi.org/10.1086/667843.
Klein, Colin. 2018. "Mechanisms, Resources, and Background Conditions." *Biology & Philosophy* 33 (5–6): 36. https://doi.org/10.1007/s10539-018-9646-y.
Koob, George F. 2020. "Neurobiology of Opioid Addiction: Opponent Process, Hyperkatifeia, and Negative Reinforcement." *Biological Psychiatry* 87 (1):44–53. https://doi.org/10.1016/j.biopsych.2019.05.023.
Krauzlis, Richard J., Lee P. Lovejoy, and Alexandre Zénon. 2013. "Superior Colliculus and Visual Spatial Attention." *Annual Review of Neuroscience* 36. https://doi.org/10.1146/annurev-neuro-062012-170249.
Lovejoy, Lee P., and Richard J. Krauzlis. 2010. "Inactivation of Primate Superior Colliculus Impairs Covert Selection of Signals for Perceptual Judgments." *Nature Neuroscience* 13 (2):261–66. https://doi.org/10.1038/nn.2470.
Marder, Eve, Marie L. Goeritz, and Adriane G. Otopalik. 2015. "Robust Circuit Rhythms in Small Circuits Arise from Variable Circuit Components and Mechanisms." *Current Opinion in Neurobiology* 31:156–63. https://doi.org/10.1016/j.conb.2014.10.012.
McCaffrey, Joseph B. 2015. "The Brain's Heterogeneous Functional Landscape." *Philosophy of Science* 82 (5): 1010–22. https://doi.org/10.1086/683436.
McCaffrey, Joseph B. 2023. "Evolving Concepts of Functional Localization." *Philosophy Compass:* e12914. https://doi.org/10.1111/phc3.12914.
McIntosh, Anthony R. 2000. "Towards a Network Theory of Cognition." *Neural Networks* 13 (8–9):861–70. https://doi.org/10.1016/S0893-6080(00)00059-9.
McPeek, Robert M., and Edward L. Keller. 2004. "Deficits in Saccade Target Selection after Inactivation of Superior Colliculus." *Nature Neuroscience* 7 (7):757–63. https://doi.org/10.1038/nn1269.
Millikan, Ruth G. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
Millikan, Ruth G. 1986. "Thoughts without Laws; Cognitive Science with Content." *The Philosophical Review* 95 (1):47–80. https://doi.org/10.2307/2185132.
Müller, James R., Marios G. Philiastides, and William T. Newsome. 2005. "Microstimulation of the Superior Colliculus Focuses Attention without Moving the Eyes." *Proceedings of the National Academy of Sciences*

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press

102 (3):
524–29. https://doi.org/10.1073/pnas.0408311101.
Neander, Karen. 1991. "Functions as Selected Effects: The Conceptual Analyst's Defense." *Philosophy of Science*:168–84. https://doi.org/10.1086/289610.
Parker, Andrew J. 2007. "Binocular Depth Perception and the Cerebral Cortex." *Nature Reviews Neuroscience* 8 (5):379–91. https://doi.org/10.1038/nrn2131.
Pessoa, Luiz. 2014. "Understanding Brain Networks and Brain Organization." *Physics of Life Reviews* 11 (3): 400–35. https://doi.org/10.1016/j.plrev.2014.03.005.
Posner, Michael I., Steven E. Petersen, Peter T. Fox, and Marcus E. Raichle. 1988. "Localization of Cognitive Operations in the Human Brain." *Science* 240 (4859):1627–31. https://doi.org/10.1126/science.3289116.
Rathkopf, Charles A. 2018. "Network Representation and Complex Systems." *Synthese* 195 (1):55–78. https://doi.org/10.1007/s11229-015-0726-0.
Rathkopf, Charles A. 2013. "Localization and Intrinsic Function." *Philosophy of Science* 80 (1):1–21. https://doi.org/10.1086/668878.
Reynolds, Sheila M., and Kent C. Berridge. 2008. "Emotional Environments Retune the Valence of Appetitive Versus Fearful Functions in Nucleus Accumbens." *Nature Neuroscience* 11 (4):423–25. https://doi.org/10.1038/nn2061.
Ross, Lauren N. 2021. "Causal Concepts in Biology: How Pathways Differ from Mechanisms and Why It Matters." *The British Journal for the Philosophy of Science.* https://doi.org/10.1093/bjps/axy078.
Silberstein, Michael. 2021. "Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences 2.0." *Neural Mechanisms: New Challenges in the Philosophy of Neuroscience*:363–93.
Silberstein, Michael, and Anthony Chemero. 2013. "Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences." *Philosophy of Science* 80 (5):958–70. https://doi.org/10.1086/674533.
Sparks, David L. 1986. "Translation of Sensory Signals into Commands for Control of Saccadic Eye Movements: Role of Primate Superior Colliculus." *Physiological Reviews* 66 (1):118–71. https://doi.org/10.1152/physrev.1986.66.1.118.
Stanley, Matthew L., Bryce Gessell, and Felipe De Brigard. 2019. "Network Modularity as a Foundation for Neural Reuse." *Philosophy of Science* 86 (1):23–46. https://doi.org/10.1086/701037.
Stinson, Catherine. 2016. "Mechanisms in Psychology: Ripping Nature at Its Seams." *Synthese* 193 (5):1585–1614. https://doi.org/10.1007/s11229-015-0871-5.
Viola, Marco. 2021. "Beyond the Platonic Brain: Facing the Challenge of Individual Differences in Function-Structure Mapping." *Synthese* 199 (1–2):2129–55. https://doi.org/10.1007/s11229-020-02875-x.
Wajnerman-Paz, Abel, and Daniel Rojas-Líbano. 2022. "On the Role of Contextual Factors in Cognitive Neuroscience Experiments: A Mechanistic Approach." *Synthese* 200 (5):402. https://doi.org/10.1007/s11229-022-03870-0.
Weiskopf, Daniel A. 2016. "Integrative Modeling and the Role of Neural Constraints." *Philosophy of Science* 83 (5):674–85. https://doi.org/10.1086/687854.
Wimsatt, William C. 1972. "Teleology and the Logical Structure of Function Statements." *Studies in History and Philosophy of Science Part A* 3(1). https://doi.org/10.1016/0039-3681(72)90014-3.

**Cite this article:** Barack, David L. 2025. "Context and Neural Function." *Philosophy of Science* 92 (1):185–203. https://doi.org/10.1017/psa.2024.50

https://doi.org/10.1017/psa.2024.50 Published online by Cambridge University Press