The Function of Perception

Peter J. Graham

1 Fitness, Functions, and Functional Analysis

I believe human perceptual systems—especially visual systems—have producing reliably accurate perceptual representations as a biological function (Graham 2010, 2012, 2014). I defend this against an argument from Tyler Burge. Burge argues that perceptual states cannot have representing accurately as a biological function, for there is a "root mismatch" between representational success and failure, on the one hand, and biological success and failure, on the other. Truth and accuracy are semantical, not practical, matters, and biology only cares about practical matters. Representational success and failure thus cannot be biological functions of any psychological state or system. Burge is not alone, as many have argued that truth and accuracy cannot be biological functions.[1]

I shall argue this isn't necessarily so. In the first Sect. 1 say a few words about biological functions before saying why, in the second, I think human perception has accurately representing the environment as a biological function. In the third and fourth I state Burge's case for thinking this isn't so. In the fifth I explain why Burge's grounds do not make his case and then in the sixth I critically examine an example Burge offers to buttress his case. In the seventh I say why the issue matters to Burge and why, even though I reject his argument, we are not at cross-purposes.[2]

There is an everyday sense of 'fitness' and a technical sense. I trust we would all like to stay fit—to stay in shape. And so many of us go to fitness centers to exercise and work out. That's the everyday notion of fitness. But it's not the sense in biology.

---

[1] E.g. Churchland 1987, Cruz and Pollock 2004, Plantinga 1993, Stich 1990.

[2] And since we are not at cross-purposes, there is always a chance I've misinterpreted his argument. This is especially true when interpreting a philosopher as subtle and sophisticated as Burge, who is often fighting on many fronts. And so I shall quote as extensively as the occasion demands.

P.J. Graham (✉)
University of California, Riverside, USA
e-mail: peter.graham@ucr.edu

A. Fairweather (ed.), *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of Science*, Synthese Library 366, DOI 10.1007/978-3-319-04672-3_2,

In biology, fitness is all about survival and reproduction; it's all about getting your genes in the next generation. Being "fit" in the first sense may contribute to being "fit" in the second, but not always. You may be in great shape without having many—if any—children, and you may be in terrible shape in the everyday sense but have more than your fair share of offspring.

The two main theories of biological functions connect functions with survival and reproduction. On the first, the function of a trait supervenes on its *propensity* to contribute to the fitness of its bearer. This theory "looks forward" to future fitness-enhancing effects. On the second, the function of a trait supervenes on its *past* contributions to fitness in ancestors undergoing natural selection, contributions that then partly explain via heredity why the trait exists in current organisms. This theory "looks backward."[3] Though I prefer the second, which one is correct does not matter for present purposes, as we'll see.[4]

To assign functions, biologists engage in what Robert Cummins calls "functional analysis" (1975). According to Cummins, a functional analysis explains how any system is able to produce an effect by *analyzing* the system. Suppose we want to know how a factory produces cars along an assembly line. An analysis breaks the faculty down into its parts and how they interact.

> Production is broken down into a number of distinct tasks. Each point on the line is responsible for a certain task, and it is the function of the workers/machines at that point to complete that task. If the line has the capacity to produce the product, it has it in virtue of the fact that the workers/machines have the capacities to perform their designated tasks, and in virtue of the fact that when these tasks are performed in a certain organized way—according to a certain program—the finished product results. Here we can explain the line's capacity to produce the product…by appeal to certain capacities the workers/machines and their organization into an assembly line. (Cummins 1975: 74)

Applied to biology, functional analysis is "essentially similar" (Cummins 1975: 74). Living organisms survive and reproduce. A functional analysis explains why. Start with the whole organism and then break it down into its major systems: digestive, circulatory, respiratory, reproductive, immune, nervous, and so on. They then break those into their components. The digestive system, for example, breaks down into the mouth, esophagus, stomach, liver, pancreas, intestines, and colon. Then break those down. The mouth, for example, includes saliva glands, teeth and tongue. The tongue in turn involves muscles, sensory receptors, and so on. Then explain how all the parts interact so as to contribute to fitness. The cells, by making

[3] The first, propensity theory is associated with Bigelow and Pargetter (1987) where functions are adaptive effects. The second, etiological theory is associated with Wright (1973), Millikan (1984), Neander (1991), Godfrey-Smith (1993) and others. These papers are anthologized in Buller (1999). For discussion of both theories and important elaboration of the etiological theory, see McLaughlin (2001). For more recent discussion, see Lewens (2004). For my preferred statement of the etiological account, see my 'Functions, Warrant, History' (2014).

[4] Burge agrees: "There are many explications of the notion of biological function. But the differences are not important for present purposes" (2010: 299). However, he seems to prefer the etiological account. See page 320, note 44.

up the muscles of the tongue, make it possible for the tongue to move food around our mouths, so that our teeth may masticate the food. The muscles in turn also assist in swallowing. Once broken down and swallowed, food passes through the esophagus to the stomach, where the stomach in turn processes the food. Each part has various capacities that contribute, through their role in the system, and the system's role in the whole organism as it interacts with other systems, to the ability of the whole organism to survive and reproduce in its natural habitat. Biological functions are then the capacities of the parts that contribute to fitness, the capacities that explain how the system is able to survive and reproduce in its natural habitat. As a result, biological traits often have more than one function. For they often contribute to survival and reproduction in many ways. The tongue helps us eat nutritious food. But it also helps us talk. Our hands also help us eat. But they also help us find food in the first place.

On the propensity theory, biological functions are the capacities of a trait that enter into a functional analysis of propensities to survive and reproduce (Lewens 2004). On the etiological theory, biological functions are the capacities of a trait that enter into a functional analysis of ancestor's survival and reproduction, and so enter into evolutionary explanations of the trait (Griffiths 1993). And so on either theory, when looking for functions look for the fitness enhancing capacities of the trait in a functional analysis of survival and reproduction.

2 The Biological Utility of Vision

What, then, is the biological function of vision? In the first chapter of his textbook *Vision Science*, Stephen Palmer asks what human vision is for.

> [We] should ask what [visual perception] is *for*. Given it biological importance to a wide variety of animals, the answer must be that *vision evolved to aid in the survival and successful reproduction of organisms*. (1999: 5)

How does perception contribute to survival and reproduction? What role does perception—especially visual perception—play in a functional analysis of our ability to survive and reproduce? Palmer continues:

> Desirable objects and situations—such as nourishing food, protective shelter, and desirable mates—must be sought out and approached. Dangerous objects and situations—such as precipitous drops, falling objects, and hungry or angry predators—must be avoided or fled from. Thus, to behave in an evolutionarily adaptive manner, we must somehow get information about what objects are present in the world around us, where they are located, and what opportunities they afford us. All of the senses—seeing, hearing, touching, tasting, and smelling—participate in this endeavor.
>
> There are some creatures for which nonvisual senses play the dominant role—such as hearing in the navigation of bats—but for *homo sapiens*, as well as for many other species, vision is preeminent. The reason is that vision provides spatially accurate information from a distance…It gives a perceiver highly reliable information about the locations and properties of environmental objects while they are safely distant. (1999: 6)

Vision helps by accurately representing objects, properties and relations in the environment. Vision benefits humans because it produces accurate representations:

> Evolutionarily speaking, visual perception is useful only if it is reasonably accurate. … Indeed, vision is useful precisely because it is so accurate. By and large, *what you see is what you get*. When this is true, we have what is called veridical perception…This is almost always the case with vision. (1999: 6)

Palmer concludes:

> [The] evolutionary role of visual perception is to provide an organism with accurate information about its environment. (1999: 15)

Palmer is far from alone. Witness Andrew Parker:

> Today, vision is the most universally powerful sense in its impact on animal interactions and behavior. With the evolution of the first eye, the size, shape, color, and behavior of animals were revealed for the first time—the position and movement of animals could be accurately tracked. Hence, the introduction of vision can be considered to be the launch of the most powerful weapon on Earth…Since [the] first eye, vision has remained on Earth. Although only 6 of the approximately 37 animal phyla possess eyes, more than 95% of all species belong to these. Vision has been a powerful weapon and a successful innovation in the animal kingdom. (2010: 441)

Witness too Ludwig Huber and Anna Wilkinson of the University of Vienna:

> Perception is a universal phenomenon. It functions primarily as a means of allowing an organism to process changes in its external environment. Thus, perception has substantial survival value and can be observed in all living species. (2010: 401)[5]

So just as the heart contributes to survival and reproduction by pumping blood, and just as the lungs contribute to survival and reproduction by taking in oxygen and removing carbon dioxide, human perceptual systems contribute by

[5] Huber and Wilkinson continue: "The primary function of the brain is to compute dynamic, predictive models of the environment. Across the animal kingdom, organisms are able to rapidly evaluate their current situation and respond appropriately to it. This suggests that the perceptual constructions of the external world provide meaning or functional significance to object and situations. As humans, we perceive an object as having a particular shape or color and we perceive it as a dog, or tree (or whatever it is). Being able to identify objects as members of known categories allows the organism to respond to them in appropriate ways." (2010: 404) Hugh Foley and Margaret Matlin say in their textbook on sensation and perception that "Our senses evolved over time to enable us to succeed in responding to the environment…our senses are functional. We live in a physical world and our well-being is very much dependent on our ability to safely negotiate that world…For example, each sensory system serves to detect change in the world…As you can surely imagine, it is often vital to notice changes in the world ("that car is heading toward me")… Most of the time, our perceptions are sufficiently accurate to enable us to interact successfully with the world" (2010: 9–10). And John Frisby and James Stone write in their textbook on vision that by seeing "…we know what objects we are looking at…we are able to describe their various features—shape, texture, movement, size—or their spatial relationships one to another. Such abilities are basic to seeing—they are what we have a visual system for, so that sight can guide our actions and thoughts" (2010: 11).

reliably representing objects, properties and relations. Just as our teeth break down food for further processing, vision helps us identify food for consumption in the first place.[6]

3 The Root Mismatch

Burge argues this isn't so. Burge argues that perceptual systems do not, for they cannot, have accurate representation as a *biological* function.

Burge grants that perceptual systems and "some of their states" have biological functions. He even holds that "biological function is relevant to understanding both the content of perceptual states and their relation to actions that serve biological needs" (2010: 229).[7] But he denies that perceptual systems have producing accurate perceptual states as a biological function. For there is "a *root* mismatch" between *representational* success and failure and *biological* success and failure:

> Biological functions are functions that have ultimately to do with contributing to fitness for evolutionary success. Fitness is very clearly a practical value. It is a state that is ultimately grounded in benefit of its effects for survival and reproduction. Explanations that appeal to biological function are explanations of the practical (fitness) value of a trait or system. But accuracy is not *in itself* a practical value. (2010: 301)

Consider an accurate perceptual representation. Accuracy is a *semantic* relationship between representation and represented object. *As such, in itself*, accuracy contributes no good or benefit to the perceiving organism; this *semantic* fact is not a *practical* fact. And so *in itself, as such*, accuracy is not a biological good or benefit.

---

[6] Some people think the fallibility of perception—the possibility of perceptual illusion—undermines this conclusion. This, Palmer says, would be a mistake: "It is easy to get so carried away by illusions that one starts to think of visual perception as grossly inaccurate and unreliable. This is a mistake. As we said earlier, vision is useful to the extent that it is accurate—or, rather, as accurate as it needs to be. Even illusory perceptions are quite accurate in most respects. For instance, there really are two short horizontal lines and two long oblique lines [in a horizontal line drawing]. The only aspect that is inaccurately perceived is the single illusory property—the relative lengths of the horizontal lines—and the discrepancy is quite modest. Moreover, illusions such as these are not terribly obvious to everyday life; they occur most frequently in books about perception. All things considered, then, it would be erroneous to believe that the relatively minor errors introduced by vision overshadow its evolutionary usefulness" (1999: 8). Most perceptual errors, Palmer thinks, occur when the perceptual system is outside of normal conditions: "[P]erceptual errors produced by these illusions may actually be relatively harmless side effects of the same processes that produce veridical perception under ordinary circumstances" (1999: 9). So that under "most everyday circumstances…normal visual perception is highly veridical" (1999: 23–4).

[7] "An individual's perceptual capacities are individuated partly through causal and practical relations that the perceiver's perceptual system bears (normally in its evolutionary history) to elements in the environment" (Burge 2010: 256). "I believe that biological basic actions—eating, navigating, mating—along with whole animal biological needs figure epistemically and constitutively in background conditions for perception, representation, and empirical objectivity" (Burge 2010: 292). See also pp. 24, 69–71, 94, 211–15, 275–6, 319–20, 320–1, 324, 330–1, 345, 373.

> It is repeatedly said that the biological function of a sensory state [or perceptual representation] is to 'detect' [or accurately represent] the presence of some distal condition (perhaps a predator). Given this claim, any failure of correlation with the distal condition is in itself a biological failure at some level of explanation. But *in itself* detection [or accurate representation] does literally *nothing* to contribute to fitness…Being present when a certain condition obtains cannot *in itself* be a contribution to biological success…One cannot assimilate issues of accuracy and inaccuracy to issues of practical use. Functioning to be accurate is not *in itself* a biological function, at any level. Biological functioning is not a semantical matter. It is a practical matter, a matter of fitness for reproduction. (2010: 301, n. 17, n. 18)

Since accuracy *in itself* does literally *nothing* to further fitness, being accurate cannot be a biological function of a perceptual state. Burge thinks it is a just a mistake to attribute accuracy as a biological function to perceptual states.

What then is the connection between accuracy, on the one hand, and biological function, on the other? Burge says it lies in the *further effects* of the sensory or perceptual state. Concerning detection, Burge says:

> I do not doubt that biological functions can involve detection relations to distal conditions. I do doubt that biological functions, as ordinarily understood, ever reside strictly in detection by itself, or in mere correlation with distal conditions. A biologically more accurate description would be that the function is to initiate some sequence of states that ultimately issues in some response to the distal condition. Sensory states that are predator detectors, for example, have the biological function of initiating a chain of avoidance behavior, given further states and conditions, with respect to the predator. It is this initiation, not the detection per se, that contributes to biological success. (2010: 301)

Predator "detectors" do not have the function of detecting predators, but rather the function of initiating predator-avoidance behavior. For "detection" *in itself* has no practical significance, whereas avoiding a predator clearly does.

Concerning perception, Burge says:

> Although accuracy in perception[s]….usually contribute[s] to fitness, [accurate perceptions] are not in themselves contributions to fitness. When they do contribute, it is not the accuracy per se that makes the contribution. The tendencies of the state to produce efficient response to *need* or, more precisely, tendencies to produce evolutionary fitness—not the veridical aspects of the state—make the contribution. (2010: 302)

And so it's not accuracy per se or the "veridical aspects of the state" that helps the organism survive when it perceives its environment. Rather it is the further effects on behavior in the organism's environment that helps the organism survive.

Burge concludes:

> There is no question that biological structures that underlie perceptual and cognitive systems evolved and were selected for. These structures were selected for not because they are or underlie representational systems per se—systems for accurately representing the world (to within some degree of accuracy). They were selected for because they yielded results that were good enough to further fitness. Evolution does not care about veridicality. It does not select for verdicality per se. (2010: 302–3)

Palmer—and countless others in perceptual psychology and evolutionary science—has made a subtle error; confusing the biological utility of vision—a further practical effect of a perceptual state—with its representational accuracy—a semantical, non-practical relation between mind and world. The biological function of perception lies in its further practical effects; it cannot reside in its representational power.

4 The Argument

I grant that evolution does not care about veridicality per se, that nature does not select for truth and accuracy *as such*. I grant that semantical relations to the environment do not, in themselves, further fitness. Even so, I think perceptual states contribute to fitness by accurately representing the environment, and so have accurately representing the environment as a function; semantical matters are also sometimes practical matters. And so I think the argument I've just attributed to Burge makes a mistake.

To find the mistake, it will prove helpful to make the reasoning behind the argument explicit and fully general. Here's my interpretation:

1. Nature does not care about capacity F of trait T *as such*; F does not further fitness *in itself*.
2. F is a biological function of trait T only if nature cares about F *as such*, only if F furthers fitness *in itself*.
3. So F cannot be a biological function of trait T.

Of course trait T may have been selected for, or may contribute to fitness, and so may have other capacities as biological functions.

4. So if T has a biological function, it must reside in further effects or capacities of T, in the organism/natural habitat.
5. But to satisfy (2), those further capacities or effects of T must be ones that nature cares about *as such*; they must further fitness *in themselves*.

And so it's natural ask what capacities or effects of biological traits nature cares about *as such*. What capacities or effects of biological traits further fitness *in themselves*? What capacities or effects of biological traits further survival and reproduction *as such*?

Nature certainly cares about the capacity to survive and reproduce *as such*, and the capacity to survive and reproduce certainly furthers surviving and reproducing *in itself*. But this is trivial and non-explanatory. Functions, recall, are capacities that enter into a functional analysis of the organism's capacity to survive and reproduce, where the functional analysis *explains* how the organism is able to survive and reproduce in terms of the capacities of the parts and how they interact, given the organism's habitat. Nature may care about survival and reproduction *as such*, and so the capacity to survive and reproduce may meet the condition premise (2) lays down on biological functions, but since functions are *explanatory*, survival and reproduction are not the capacities we're looking for.

At this point the four Fs come to mind: Feeding, fleeing, fighting, and reproducing. For it seems empirically true that an organism can only survive if it eats nutritious food, flees from dangerous predators, successfully fights off real competitors, and finds fertile and cooperative mates. For these and related capacities seem empirically necessary for an organism to survive and reproduce. All known organisms need food. All known organisms need to avoid being eaten. All known organisms with competitors need to fight from time to time. And all known organisms that sexually reproduce need to find cooperative and fertile mates to reproduce their kind. Living organisms survive and reproduce by having their biological needs met, by feeding, fleeing, fighting and reproducing. These are all clearly practical matters, matters of great importance to survival and reproduction. And so it seems we have discovered four capacities that nature cares about *as such*, capacities that further fitness *in themselves*.

6. The explanatory capacities nature cares about *as such* are the four Fs (or other capacities at the very same level of explanation, capacities that nature clearly seems to care about *as such*, *in themselves*).
7. Given (2) and (6), capacity F of trait T is a biological function of T only if F is one of the four Fs.

In other words, the only biological functions of traits are feeding, fleeing, fighting and reproducing. Biological functions consist in the capacities or effects that most obviously serve survival and reproduction: finding *nutritious* food, finding *fertile* and *cooperative* mates, *successfully* fleeing from *dangerous* predators, fighting *effectively* with *competitors* for mates, food, shelter and so on. Biological functions consist in meeting or fulfilling biological *needs*. For these are all obviously—if not analytically—*practical* goods, goods that clearly contribute to, if not comprise, survival and reproduction.[8]

Applied to our question the consequence is clear: since nature does not care about representational accuracy *as such* (accurately representing the environment is not *the same as* [or at the same level as] eating nutritious food, fleeing from danger, fighting off a rival or predator), accurately representing the environment cannot be a biological function of perceptual states. Their functions lie in their further effects, in their further contributions to practical needs. And so the biological functions of perceptual states are to contribute to finding food, fleeing from predators, fighting off rivals and predators, and finding cooperative and fertile mates, and so on. Our perceptual systems were selected for because they were good enough to further fitness, not because they accurately represent the environment. Accuracy is not, for it cannot be, a biological function.

[8] There is some suggestive but inconclusive evidence that Burge identifies functions with needs, or the fulfilling of needs. On page 371 he says individuals fulfill "basic whole-animal functions" and on page 292 he calls eating, navigating and mating "biologically basic actions…along with whole-animal biological needs." And on page 94 he describes processes that "are ecologically relevant to the individual's basic functions—functions such as eating, navigating, and fleeing danger." Combined, passages such as these at least suggest a tendency to identify biological functions with capacities that obviously, if not constitutively, contribute to survival and reproduction.

5 The Mistaken Premise

Given the interpretation, it's pretty clear how the first premise supports the conclusion. But it's also pretty clear that the argument doesn't go anywhere without the second. I grant the first; I reject the second. I reject the claim that a capacity F of a trait T can be a biological function only if capacity F *in itself* furthers fitness. True, biological functions are contributions to fitness. True, the four Fs are empirically necessary contributions to fitness. Even so, they are not the only biological functions of biological traits. For the biological functions of traits are the capacities that enter into a functional analysis of how the trait contributes to feeding, fleeing, fighting and reproducing. The capacity of the trait that contributes to the four Fs is the function, even if the trait, *as such* and *on its own*, does nothing to further fitness. Biological functions of traits are capacities of the trait that *explains* how it contributes to fitness, where most traits only contribute to fitness as a contingent, empirically determined matter of fact, given their capacities, their role in the organism, and the broader environment, even though by themselves, all on their own, taken in isolation, they contribute literally nothing to fitness. Whether they contribute to fitness *as such* and *in themselves* does not matter. What matters is whether, as a matter of fact, they contribute to fitness given their role in the system and the system's role in the broader environment. Biological functions are explanatorily relevant capacities of traits that contribute to fitness, whether contingently or necessarily so; it does not matter whether they are means towards that end *as such*, *in themselves*; they need only be means to that end. The second premise is false.

Consider surface coloration. An organism's surface color *as such* and *in itself* clearly does not further fitness. Being red, white or blue isn't the same as eating nutritious food, or fleeing from a dangerous predator. Even so, surface coloration often makes a huge difference to fitness. Take a polar bear's white fur. Its function is to camouflage the bear as it stalks its prey. How does it do that? By matching the background snow. And so matching the background is functional for the bear. A brown polar bear would fail to stay hidden for very long; mismatching the background is obviously dysfunctional. Matching the environment then enters into a functional analysis of the bear's ability to survive and reproduce; matching the environment explain how it provides camouflage, which then in turns explains how it successfully stalks its prey, which then it turn explains how it gets enough food to eat. And so matching the environment is a biological function of the bear's fur, for matching the environment enters into a functional analysis of its ability to survive and reproduce.

Is matching the background environment a contribution to fitness *in itself*? No. Not at all. Camouflage is not the same thing as eating nutritious food, finding a mate, avoiding a predator, and so on. Camouflage *as such* is not a fitness enhancing effect; coloration *as such* is not a practical good. Camouflage *as such* is an *aesthetic*, not a practical, matter. Even so, the polar bear's white fur is supposed to camouflage the bear. Camouflage is often also a practical matter.

Surface coloration often makes a huge contribution to fitness in countless species, as a contingent matter of fact, even though it doesn't make a difference *as such*.[9] Some species identify mates by skin color. Some species hide from predators by skin color. Some species avoid detection as they stalk their prey by color. Surface color partly explains why these species survive and reproduce. The contingent, not-necessarily functional capacities of surface colors often enter functional analyses of how organisms survive and reproduce.

We can make the same points about accurate representations. Accuracy *as such* does not contribute to fitness. But does it follow that representational accuracy can never, in any circumstance, contribute to fitness? Does it follow that representational accuracy cannot enter into a functional analysis of an organism's ability to survive and reproduce? True, accuracy *as such* is not a practical good. Representational accuracy is not eating nutritious food, finding a mate, avoiding a predator, and so on. Even so, can representational accuracy make a contingent contribution to finding food, finding mates, avoiding predators, and so on?

Yes. Representational accuracy often makes a huge contribution to fitness in countless species. We all know this. Some species rely on perception to identify mates. Misrepresenting a predator as a mate can bring your life to an early end. Some species identify and flee or hide from predators by first accurately representing them as predators or as danger. Misrepresenting predator as prey can be just as bad or worse as misrepresenting a predator as a mate. Organisms rely on their perceptions to navigate their environments. Accurate representations are better guides. Just as white fur helps the bear because white matches its environment, accurate perceptions help countless creatures because accurate perceptions match their environments. The accuracy of perceptual representations—especially visual representations in humans—plays a role in the functional analysis of how organisms with perceptual systems are able to survive and reproduce. Getting it right often contributes to fitness, as a contingent, empirically determined matter of fact, in countless creatures with perceptual systems. Just take away accuracy but leave everything else intact and see what happens. Would you rather walk towards a cliff with accurate, or inaccurate, representations as your guide? If you find yourself at all puzzled by this, re-read the second section, including the notes.

Burge says "there is no question that biological structures that underlie perceptual" systems underwent natural selection. He says they were not selected because they "underlie representational systems per se" but rather they were selected because they further fitness. "Evolution," he says "does not care about veridicality" per se. But evolution also does not care about coloration per se; it does not care about pumping blood per se; it does not care about sharp teeth or long legs per se; it

[9] We can imagine cases where coloration is completely irrelevant to survival and reproduction. Think of animals in lightless caves. These animals do not use vision to identify anything, and so they do not use color to identify food, mates, predators, etc. Nor do they use skin color to avoid predators or to avoid detection by their prey. Their color makes no difference whatsoever to their chances for survival and reproduction, both in their current environment and in their evolutionary history. In such a case, color makes no contribution whatsoever to fitness. *A fortiori* it makes no contribution whatsoever to fitness *as such*.

does not care about oxygen diffusion or photosynthesis per se. Evolution only cares, per se, about contributions to fitness and reproduction. It does not follow from any of this that evolution did and does not care, *as a matter of fact*, about coloration, pumping blood, oxygen diffusion, sharp teeth and long legs. It does not follow from any of this that coloration, pumping blood, oxygen diffusion, sharp teeth and long legs cannot enter into the functional analysis of an organism's ability to survive and reproduce. And so it does not follow from the fact that evolution does not care about veridicality per se that it does not care about veridicality as a contingent, empirically well-established matter of fact. All the point shows is that if accurate representations did not contribute to fitness, nature would not have cared about them. But since they do, nature cares.[10]

I think the tendency to infer from the fact that evolution does not care about a capacity *as such* to the conclusion that the capacity cannot be a biological function results from not thinking through the functional analysis of the trait in the overall economy of the organism and the its natural habitat. If biological functions are the capacities and effects of traits that contribute to meeting needs, albeit contingently given their role in the organism in its natural habitat, then many capacities are biological functions even if they don't contribute to fitness *as such, in themselves*. To suppose otherwise is to *identify* biological functions with needs, and thereby *exclude* the capacities or effects of the traits of the organism that, often as a contingent matter of fact, *explain* how those needs are met. Then only the four Fs would fall under the category of function. The function of the heart would not be to pump blood, but only to assist in fleeing, feeding, fighting. The function of the kidneys would not be to remove wastes. The function of the eyes would not be to see. Though biological functions are necessarily *associated* with survival and reproduction—with meeting practical needs—it does not follow the biological functions are restricted to those capacities that contribute to need *as such*.[11]

[10] Burge says predator detectors have the function of "initiating a chain of avoidance behavior with respect to the predator." But the organisms *succeeds* at avoiding the predator by first detecting it; it relies on detection of the predator to avoid the predator. No detection, no initiation of avoidance behavior. And so detection enters into the functional analysis of how the organism avoids predators. Detection is not epiphenomenal when explaining fitness.

Burge says accurate perceptions are not *in themselves* contributions to fitness. Burge says the "tendencies of the state to produce" evolutionary fitness and "not the veridical aspects" of the perception "make the contribution" to fitness. True, the tendencies of the state to produce fitness make the contribution to fitness; that's trivial, but for the same reason non-explanatory; we want explanations of how traits contribute to fitness. And it's pretty clearly true that "veridical aspects" often partly explain why perceptual states contribute to fitness. It is true that without the further effect on behavior the perceptual state would not contribute to fitness. The perceptual state does not contribute to fitness all on its own, *by itself* or *as such*; further behavior is required too. But it is also true that that the behavior contributes to fitness partly because guided by an accurate perceptual state. The accuracy of perceptual states is not epiphenomenal in the explanation of how perceptual systems contribute to fitness. Getting it right often matters.

[11] It may be helpful to sketch out the following in the margins, just to the right. At the top of the paragraph write 'survive and reproduce'. Then write down 'the heart' at the bottom. Then write the 'four Fs' just under 'survive and reproduce'. Now think about how the heart contributes to the four Fs, and so to survival and reproduction, and write them in as well. You will end up writing down

I agree that nature per se only cares about fitness-enhancing traits. But nature cares, as a contingent matter of fact, about countless capacities of traits for their contingent, matter of fact contributions to fitness. It is because they contribute to fitness that nature cares about them. So from the fact that nature does not care about a particular capacity or effect per se shows nothing about whether nature, as a contingent, empirical matter of fact, cares about that capacity or effect. Nature may not care about veridicality per se, but for all that it may care about veridicality a great deal. Veridicality can be, and surely is, a *biological* function of many of our perceptual states and perceptual systems. The second premise is just false.

6 Burge's Example

This concludes my discussion of Burge's argument. Burge also provides an example to make his point. And since examples are sometimes more compelling than arguments from general principles, it would be wrong not to discuss his case. Burge's example purports to show that misrepresentation does not entail failure of biological function, so that correct representation is not a biological function.

Burge imagines a creature like a rabbit that relies on a detection mechanism to avoid predators. Such mechanisms are often unreliable, for false positives ("danger is present" when there is nothing to fear) outnumber true positives ("danger is present" when it's time to run). Burge further imagines that every triggering increases strength and agility: being frightened spurs the exercise required to stay in shape (like having a workout buddy that drags you to the gym everyday, or an alarm clock that reminds you it's time to go to the gym). And so every triggering

---

the steps in a functional analysis of the heart's contribution to survival and reproduction. In my sketch I wrote 'pump blood' just above 'the heart'. Moving up, I wrote 'moving blood and other stuff through the organism'. I then wrote 'and so assist the organism in fighting diseases, providing energy to the muscles, removing dangerous wastes, etc.' I then wrote 'and so contribute to digestion, locomotion, cognition, etc.' And by doing all of that it contributes to fleeing, feeding, fighting and reproducing. Voila, a crude functional analysis of the heart's role in the organism's ability to survive and reproduce.

I call all of the contributions at all of the different levels "vertical" functions of the trait. And so there are biological functions at the highest level and the lowest and all the levels in between. It's a biological function of the heart to contribute to survival and reproduction, to contribute to fleeing and feeding, to bring oxygen to the brain and take wastes to the kidneys, ….. and to pump blood by beating regularly. It contributes to survival and reproduction by contributing to meeting needs, and it contributes to meeting needs by driving circulation of blood and oxygen, among other things, through the body, and it does all that by pumping blood. Though pumping blood—or pumping anything at all—is not *as such* a practical good (taken in isolation pumping fluid is but a mechanical property) it's pretty clearly, as a contingent matter of fact, a practical good. If it followed that F is not a function of a trait because nature does not care about F per se, then pumping blood is not a function of the heart. That cannot be right (Graham 2011, 2012; Fodor 1998).

contributes to fitness, for it keeps the animal in tip-top shape, so the animal is more likely to avoid predators when really present. So when the mechanism fires and there's no predator to avoid—in a case of representational error—there would be "no biological sense in which the mechanism failed to fulfill a biological function…The biological function is to contribute to a fit response to the predator—which entails contributing to avoiding predators" which is exactly what this inaccurate perception does (2010: 302).

Though Burge does not do so, we can put the example in terms of a functional analysis of survival and reproduction. How does an inaccurate perception contribute to fitness? The inaccurate perception spurs exercise, which contributes to strength and agility, which contributes to its capacity to evade predators, which contributes to survival and so to fitness. Spurring exercise figures in a functional analysis of the creature's capacity to survive and reproduce, and so is a function of the detection mechanism. And in the very case Burge imagines, the inaccurate perception fulfilled that function. Burge claims the case shows that perceptual states do not have the biological function of representing accurately, for even though the creature misrepresented its environment there is no biological sense in which the mechanism failed to fulfill a biological function. Representational error without biological error entails that representational success and error is not a species of biological success and error.

I do not think the example works. Burge has overlooked the possibility that the mechanism has more than one function. Spurring exercise may be one function, representing danger another. Many traits have more than one function: think of how your tongue helps you eat as well as talk, the way your hands help you communicate, eat, fight, climb, and so on. From the fact that a trait fulfills one function nothing automatically follows about whether it succeeded or failed in fulfilling other functions. The creature's mechanism may have the function of spurring exercise (so as to run quickly from predators) as well as accurately representing the presence and location of predators (so as to run at the right time in the right way from predators). As long as accuracy plays a role in the functional analysis of the danger-detection mechanism's contribution to fitness, the danger-detection mechanism would have accurately representing the presence of danger as a biological function.

And surely accuracy plays a role. Imagine that the detector failed to represent the presence of danger when danger was present. Then the animal would be in big trouble indeed. Or imagine that though it correctly represented danger, it represented it in the wrong location. The animal might then run into the open arms of its predator. Being full of strength and agility wouldn't help at all.

Danger detectors in many animals have, I believe, the biological function of detecting danger, for detecting danger—even if they are not very reliable at it—plays a role in the functional analysis of how the detector contributes to the capacity of the animal to survive and reproduce. From the fact that they sometimes or even usually misrepresent, or from the fact that there are cases where misrepresentation has very little costs, or from the fact that the device might contribute to fitness in other ways and so have more than one biological function, it does not follow that when they misrepresent with no obvious immediate costs, that there is "no biological sense in which the mechanism failed to fulfill a biological function." If representing accurately is a biological function of the detection device, then every representational error is also a biological error, even if there are no obvious or immediate biological costs.

Burge seems to be reasoning as follows. Suppose the exercise of an avoidance mechanism, in each and every case, increases strength and agility, and so increases the effectiveness of predator avoidance behavior. Suppose it fires on an occasion when danger is not present. Then, Burge concludes, because the device contributed to fitness, there was no biological sense in which it failed to fulfill a biological function. To see that something has to be wrong with Burge's example, consider an analogous case. Suppose the exercise of a sperm producing device increases, in each and every case, its own health and vitality, and so on average it would produce more sperm over time, and so it would fertilize more eggs. Suppose it fires on an occasion and fertilizes no eggs on that occasion. Then, by parallel reasoning, we should be entitled to conclude that there was no biological sense in which the device—and the sperm it produced—failed to fulfill a biological function. But that, of course, is absurd. From the fact that a device may fulfill one biological function on an occasion, nothing follows about whether it succeeded or failed in fulfilling its other biological functions. Burge has not imagined a case where representational accuracy plays no role in contributing to fitness, and so he has not imagined a case where representational accuracy is not, or cannot be, a biological function of a representational state.[12]

Burge's example exploits a fact about functions that is worth making explicit: the biological functions of traits are not always reliable capacities or effects. The trait may have a function that it only fulfills once in a blue moon. Ruth Millikan uses the example of sperm to make this point (1984). The biological function of sperm is to fertilize eggs. However, the vast majority of sperm never come close to an egg. The biological function of an item is what it does *often enough* to contribute to fitness, even if it hardly ever does. Nature settled on a mechanism for reproduction—sperm and egg—where countless sperm are produced for every egg. As a result, though each and every sperm is supposed to fertilize an egg—that is its function—nature is perfectly okay with nearly every sperm failing to fulfill its function; a male can reproduce if only one of millions of its sperm fulfills its function. For sperm, success

[12] The function of triggering exercise is, I think, a decoy. It's there to get us to agree that the device contributed to fitness, despite the error. But imagine a case where the animal doesn't need to exercise to stay in shape. Or imagine a case where the animal doesn't get any "fitter" in the colloquial sense from exercise. Or just imagine that the detector mechanism isn't there in the animal because it helps the animal stay in shape. In all of these cases, there would be no fitness-enhancing benefit to running away when there is no danger to runaway from. And so imagine cases where triggering exercise doesn't enter into the functional analysis of the animal's ability to survive and reproduce, and so isn't a biological function of the device. And so when the animal misrepresented the presence of danger and sprinted away, there would be a clear sense in which the device failed to fulfill a biological function.

once in a blue moon is success often enough.[13] Though *effective*—sperm do indeed fertilize eggs—they are not very reliable.

Many predator detectors work like this, where the representation of danger is not very reliable; it often represents the presence of danger when there is nothing to fear. Though *effective*—they "fire" almost every time danger is present and so keep the organism safe from harm, or at least give the animal a fighting chance—they frequently fire when danger isn't present, and so are not very reliable. Nature has settled on such a way of avoiding predators because false negatives ("there is no danger present; I'm safe" when danger is lurking) are so much worse than false positives ("danger is present, run!" when there's nothing to fear). If the animal overestimates the chances of danger and runs away at the slightest sign, it will effectively avoid predators when they are present, even if it frequently runs away when, in fact, it is perfectly safe. That's why the detector is *effective* (when danger is present it usually says it is) but *unreliable* (most of the time it's mistaken and there is nothing to fear). Burge's example exploits the fact that false positives are often pretty cheap. But the low cost of false positives does not diminish the high cost of false negatives. And it is the cost of false negatives, as well as the low cost of false positives, that explains why nature settled on an unreliable, but nevertheless effective, danger detection device. Accurate detection obviously matters—it explains why the device is effective—even if the device isn't very reliable. Nature settled on an unreliable but effective device, effective because accurate often enough. Most of our perceptual states and systems, however, are not like this. Most are reliable, and contribute to fitness by being reliable. Unreliable danger detectors are the exception that, so to speak, proves the rule.

7 At Cross Purposes?

I've critically discussed Burge's argument and his supporting example at length because the issue matters to me. But why does the issue matter so much to Burge?

The answer involves one of the main themes of his book, the distinction between sensation and sensory systems, on the one hand, and genuinely perceptual systems, on the other. Burge notes the popularity of "Sensation and Perception" as a textbook title, but laments the lack of a good account of the difference. And so in his book he sets out to provide one. But his account, as we will see, comes under threat from "deflationary" accounts of perceptual representation, accounts that effectively reduce perception to sensation where sensation in turn reduces to biological function. Burge then uses the

[13] Not so for the human heart. Not only must the heart pump blood once in a while to contribute to fitness, it must pump blood all the time. Unlike sperm, "often enough" for the heart is all the time. When I've made this point before, I've said that the heart, unlike sperm, not only has a certain effect as its function—pumping blood—the heart has producing that effect *reliably* as its function; the heart isn't just supposed to pump blood, its supposed to pump blood reliably (Graham 2012). Most organisms with hearts can survive and reproduce only if their hearts pump blood all the time. I believe most of our perceptual systems have reliably representing as a function.

premise that nature does not care about accuracy as such to block the reduction. That's (at least one reason) why Burge cares about the issue.

Here's Burge's account of sensation. According to Burge, non-perceptual sensory discrimination—sensation—involves functional information carrying. Information carrying is a broadly law-like correlation between a property of the signal and property of the source. For example, the rings of a tree (the signal) carry information about the age of the tree (the source), because of a law-like correlation between the number of rings in a tree and the age of a tree. Or, to take another example, a ringing doorbell (the signal) carries the information that someone is at the door (the source), because of a law-like correlation between ringing doorbells and visitors at the door. It's been widely recognized that our sensory systems carry information in this sense. Just as iron responds to the presence of oxygen, so too our skin responds to the presence of hot and cold temperatures, to light and dark illumination, and so on. When we touch something with our hands our sensory transducers respond to the change in shape and texture. When light enters our retinas the pattern of light absorbed causes regular changes in our visual system. Changes in the world cause changes in our retina, which in turn cause changes in our visual systems, all in a law-like way, such that changes in us (the signals) carry information about changes in our environment (the source). Our sensory organs—eyes, skin, ears, nose, tongue—not only carry information about our surrounding environment, they are *supposed* to, unlike iron and the rings of a tree. Carrying information, Burge holds, is one of the biological functions of our sensory organs: sensory states involve sensitivities to the environment that are "biologically functional for the individual" (2010: 293); the sensory systems of organisms are systems that are supposed to carry information (2010: 317).

Perception, however, differs. For functional information carrying is "pre-perceptual" and does not require that animal be able to perceptually represent its environment. Plants are sensitive to changes in their environment and respond in functionally useful ways. Likewise bacteria are sensitive to light, oxygen, and magnetic fields, and respond appropriately in turn. But plants and bacteria do not perceive. They do not genuinely represent distal objects, properties and relations.

"Deflationary" accounts of perceptual representation would fail to mark this difference between sensation and perception. Deflationary accounts seek to reduce perceptual representational content to the information that a state or system is supposed to carry. If a state is supposed to carry the information that the source is F, then the state, on these accounts, represents the source as F. If the source is F, then the state has fulfilled its biological function. If the source is not F, then the state has failed to fulfill its biological function. Deflationary accounts of representation then identify perceptual accuracy with fulfillment of biological function to carry information and identify perceptual error with failure to fulfill this function. On these accounts representational success is necessarily and essentially biological function fulfillment, and representational failure is necessarily and essentially failure to fulfill a biological function. Such accounts are associated with the work of Fred Dretske, among others.

Burge accepts such accounts for sensory registration. He rejects them for perception, as accounts of genuine objective sensory perceptual representation as of particulars in the distal environment. These accounts are deflationary, according to Burge, for they do not involve any genuinely psychological terms—iron carries information and nature is full of functions in non-psychological beings—and the accounts apply equally to the non-perceptual sensory systems of mollusks, paramecia and worms. They wrongly assimilate sensory *perception* to mere sensory *registration*. And so when Burge argues that nature does not care *as such* about accuracy, he's largely out to undermine deflationary views of representation that assimilate representation to biological function.

It is not my goal to defend these views, or to advance the general philosophical outlook from which they arise. It is not my goal to "naturalize" perception or epistemology in the sense of 'naturalize' Burge intends (2010: 296–8). It is not my goal to dispute Burge's distinction between sensation and perception or to dispute his account of perception. Nor is it my goal to replace the explanatory enterprise of perceptual psychology with the explanatory enterprise of evolutionary biology. My purposes are not at odds with Burge's.

Burge thinks that perceptual psychology and evolutionary biology ask different questions. I agree. Perceptual psychology asks how veridical and illusory perceptual representations of a distal environment are formed from proximal sensory registration on sensory transducers. Evolutionary biology asks about origins and fitness enhancing effects of perceptual systems and perceptual states. Psychology and biology ask different questions and offer different explanations about overlapping subject matters. Compare physiology and evolutionary biology. Physiology asks about the biochemistry and functional role of organs within an organism and evolutionary biology asks about origins and fitness enhancing effects of those organs within the organism and its natural habitat. They ask different questions and offer different explanations about overlapping subject matters, without the former reducing to the latter. My goal is to defend a claim about the fitness enhancing effects of perceptual states, not to reduce the very nature of perceptual states and perceptual representations to fitness enhancing effects.

Even so, for the sake of argument I can accept that representational content does not reduce to biological function. I can accept that "perceptual accuracy does not necessarily and constitutively contribute to biological" success. I can accept Burge's claim that "functioning in interacting successfully with respect to a beneficial or detrimental distal condition is not the same thing as accurately detecting the condition" (Burge 2010: 302).

Though I critically discussed Burge's example, we can imagine another Burge-inspired case where a perceptual state misrepresents without failure of biological function. Imagine an animal with veridical perceptions but the animal does not use them to control behavior in any way. Perhaps the creature over evolutionary time has become immobilized, has no predators, receives nutrition like a plant, and reproduces asexually. In such a case, we can imagine that the perceptions play no role in the functional analysis of the organism's ability to survive and reproduce, either in its current propensity to survive and reproduce or in its evolutionarily recent past. Perhaps its perceptual system is a vestigial, non-functional trait. And so on both accounts of functions, its perceptions would lack a biological function. *A fortiori* its representational successes and failures are not "fulfillments or frustrations of biological functions" (Burge 2010: 308). Unlike our perceptual states that contribute to fitness by accurately representing our environment, its perceptual states make no contribution at all.

And so I am not at odds with Burge's opposition to deflationary accounts of perceptual representation. Nor I am at odds with his main premise that nature does not care about truth and accuracy *as such*. Even so, our perceptual systems have the contingent, empirically established biological function of producing reliably accurate perceptual systems. Or so I have argued.[14]

References

Bigelow, J., and R. Pargetter. 1987. Functions. *Journal of Philosophy* 84: 181–196.
Buller, David. 1999. *Functions*. Albany: SUNY Press.
Burge, Tyler. 2010. *Origins of objectivity*. New York: Oxford University Press.
Churchland, Patricia. 1987. Epistemology in the age of neuroscience. *Journal of Philosophy* 84: 544–553.
Cruz, Joseph, and John Pollock. 2004. The chimerical appeal of epistemic externalism. In *The externalist challenge*, ed. R. Schantz. Berlin: de Gruyter.
Cummins, Robert. 1975. Functional analysis. *Journal of Philosophy* 72: 741–764.
Fodor, Jerry. 1998. Is science biologically possible. In *Critical condition*. Cambridge: The MIT Press.
Foley, Hugh, and Margaret Matlin. 2010. *Sensation and perception*. Boston, MA: Allyn & Bacon.
Frisby, John, and James Stone. 2010. *Seeing: The computational approach to biological vision*. Cambridge, MA: The MIT Press.
Godfrey-Smith, Peter. 1993. Functions: Consensus without unity. *Pacific Philosophical Quarterly* 74: 196–208.
Graham, Peter J. 2010. Testimonial entitlement and the function of comprehension. In *Social epistemology*, ed. A. Haddock, A. Millar, and D. Pritchard. Oxford: Oxford University Press.
Graham, Peter J. 2011. Does justification aim at truth? *Canadian Journal of Philosophy* 41: 51–72.
Graham, Peter J. 2012. Epistemic entitlement. *Nous* 46: 449–482.
Graham, Peter J. 2014. Functions, warrant, history. In *Naturalizing epistemic virtue*, ed. A. Fairweather. New York, NY: Cambridge University Press.
Griffiths, Paul. 1993. Functional analysis and proper function. *British Journal for the Philosophy of Science* 44: 409–422.
Huber, Ludwig, and Anna Wilkinson. 2010. Evolutionary approach. In *Encyclopedia of perception*, ed. E. Goldstein. Thousand Oaks, CA: Sage Publications.
Lewens, Tim. 2004. *Organisms and artifacts: Design in nature and elsewhere*. Cambridge, MA: The MIT Press.
McLaughlin, Peter. 2001. *What functions explain*. Cambridge: Cambridge University Press.
Millikan, Ruth Garrett. 1984. *Language, thought, and other biological categories*. Cambridge, MA: The MIT Press.
Neander, Karen. 1991. The teleological notion of 'function. *Australasian Journal of Philosophy* 69: 454–468.

[14] I benefited from comments from Zach Bachman and Colleen Macnamara on previous drafts. I presented an earlier version at the Academica Sinica in Taipei, Taiwan. I am grateful for comments on that occasion. And I am deeply grateful to Tyler Burge for getting me interested in these topics and for his continued encouragement and support.

Palmer, Stephen. 1999. *Vision science: From photons to phenomenology*. Cambridge: The MIT Press.

Parker, Andrew. 2010. *Encyclopedia of perception*, ed. E. Bruce Goldstein. Thousand Oaks, CA: Sage Publications.

Plantinga, Alvin. 1993. *Warrant and proper function*. New York, NY: Oxford University Press.

Stich, Stephen. 1990. *The fragmentation of reason*. Cambridge, MA: The MIT Press.

Wright, Larry. 1973. Functions. *The Philosophical Review* 44: 409–422.