

The Role of Valence in Intentionality

David Leech Anderson
Department of Philosophy
Illinois State University, Normal, USA

Abstract

“Functional intentionality” is the dominant theory about how mental states come to have the content that they do. “Phenomenal intentionality” is an increasingly popular alternative to that orthodoxy, claiming that intentionality cannot be *functionalized* and that nothing is a mental state with intentional content unless it is phenomenally conscious. There is a consensus among defenders of phenomenal intentionality that the *kind* of phenomenology that is both necessary and sufficient for having a belief that “there is a tree in the quad” is that the agent be consciously aware of the meaning of “tree” and “quad”. On this theory, experiences with a valence – experiences like happiness and sadness, satisfaction and frustration – are irrelevant to intentionality. This paper challenges that assumption and considers several versions of “valent phenomenal intentionality” according to which a capacity for valent conscious experiences is either a necessary or a sufficient condition for intentionality (or both).

1. Introduction

One of the most important disputes in the philosophy of mind surrounds the question: *Is phenomenal consciousness necessary for intentionality?* The dominant theory of intentionality, *functional intentionality*, answers “no”. According to several popular versions of functional intentionality, a complex system will be in the intentional state of “believing that *p*” so long as some state of the system serves as a representation “that *p*” and is playing the belief-role in the overall economy of the system. For a state of the system to be a representation “that *p*” it must be that the semantic elements that constitute the representation are causally regulated by (or “track”) the objects and properties to which they refer.¹

Although functional intentionality continues to be the most often defended (or at least the most often presupposed) within both philosophy and cognitive science, it is not as dominant as it was twenty years ago. *Phenomenal intentionality* is an opposing theory which answers “yes” to

¹Two influential sources for functional theories of intentionality that rely on mechanisms of *tracking* are Fred Dretske’s (1981) *information-theoretic* account and Ruth Millikan’s (1984) *teleological* account. Many contemporary theories are built upon insights from these seminal works.

the question above and has been steadily growing in influence in the new millennium. Phenomenal intentionality holds that no matter how effective a system's tracking mechanisms and how flexible and adaptive its behavior, it will still lack genuine mental states with intentional content if it lacks phenomenal consciousness. To be in a phenomenally conscious mental state is for there to be *something that it is like* to be in that state. The phenomenal properties are how things *seem* to the agent, from her first person perspective.

While a complete theory of phenomenal intentionality must address a range of complex issues, there is value in giving a basic definition that captures the central intuition of the view without engaging in subtleties. One method for doing this is to define phenomenal intentionality by appeal to states that are paradigmatically mental and that uncontroversially possess intentional content. This is the strategy of Graham *et al.* (2007) as they make the following claim about paradigmatic mental states, with perceptual beliefs offered as an example (p. 470):

A mental state is phenomenally intentional in content just in case the intentional content of the state (*viz.* what it's about or represents or is directed at) is determined or constituted by its conscious or phenomenal character or what-it's-likeness alone.

They say that intentional content "presents itself immediately or directly to the subject" and that "the subject can *just tell* what he or she is thinking of" (p. 471). In what follows, I will use the expression "being consciously aware of X" to capture this immediate phenomenological awareness of what an intentional state is about.

It should be noted here, that Graham *et al.* (2007) are not merely making a claim about human beings and the form that phenomenal intentional content happens to take in the human case. They are making a claim about the very nature of intentionality. Having a conscious awareness of a specific content is not merely *sufficient* for intentionality; it is *necessary*. While there is room in this picture for unconscious brain mechanisms, any such content is derivative in that its intentional status comes not from its intrinsic properties but only from its integration within a phenomenally conscious system. The two claims about "consciousness of content" made by standard versions of phenomenal intentionality are as follows:

The Sufficiency Claim: For any mental state M, if the agent is consciously aware that the content of M is C, then M has non-derived content C.

The Necessity Claim: For any mental state, M, if the agent is not consciously aware (or potentially so) that the content of M is C, then M does not have non-derived content C.

It might seem inevitable that a theory of intentionality grounded in phenomenal consciousness would make the “conscious awareness of (a particular) content” the lynchpin of the theory, as reflected in the previous two claims. It is indeed true, I believe, that for *humans* the conscious awareness of what a mental state is about plays an essential role in determining the intentional content of that state. It does not follow, however, that this very same property is either necessary or sufficient for intentionality per se. We must take care that we are not so impressed with the characteristics of our own phenomenology that we fail to imagine other types of phenomenally conscious creatures whose intentionality might take different shape.

One particular prejudice we will be questioning is the assumption that each token mental state is a closed intentional unit and that for each token state it is the phenomenal properties of *that very state* (and none other) that are constitutive of its intentionality. Consider two examples: a perception and a belief. When I *see* the rhododendron in my backyard, there is a distinctive *sensory phenomenology* characteristic of perceiving this kind of green leaf and that kind of red blossom. And when I believe that rhododendrons are beautiful, there is a distinctive *cognitive phenomenology* characteristic of thinking about rhododendrons and of the concept, “beauty”. It may seem reasonable to assume that we need go no further than the intrinsic properties of *each token mental state* to find the phenomenal properties that make it a mental state with intentional content.²

This paper will challenge that assumption. A case will be made for the possibility of a creature that “believes that *p*” even though it is wholly incapable of being consciously aware that *p* as well as a creature that fails to “believe that *p*” despite the fact that it is consciously aware of *p*. The strongest conclusion to follow from these arguments is that the kind of phenomenology necessary for intentionality is the possession of conscious states *with a valence*, not states characterized by a *conscious awareness* of content.

²It is important not to misunderstand the claim, here, that according to phenomenal intentionality it is the *intrinsic* or internal properties of a cognitive state that determine that something is a mental state with intentional content. This is not to say that phenomenal intentionality denies that *relational* or external properties ever play a role in determining the specific content of a mental state. It certainly does. But unlike functional “tracking theories” (for example), which make the right kind of causal connections – both internal to the system and externally related to the world – both *necessary* and *sufficient* for intentional content, for phenomenal intentionality they are neither. It is not necessary because a solipsistic mind causally isolated from all external objects will nonetheless have contentful mental states. Nor is it sufficient because the only external, relational properties with the capacity to influence the content of a mental state are those properties of semantic interest to phenomenally conscious minds.

2. Valent vs. Non-Valent Phenomenological States

The term, “valence” is commonly used in psychology to identify conscious states that are either positive or negative in character. Mental states with a positive valence include pleasure, joy, happiness, and satisfaction. Among those with a negative valence are pain, depression, sadness, and frustration. But what about everyday perceptions and beliefs about objects in the external environment? These are paradigmatic examples of intentional states. Do they have a valence?

There is some reason to think that they do, at least in human beings. Our sensory and cognitive states are often woven together with a rich valent phenomenology. In our own first-person experience, it is difficult (if not impossible) to separate the non-valent *visual* phenomenology of “the look” of the rhododendron blossom (its size, shape, and red color) from the valent *aesthetic* phenomenology of “the pleasure” we enjoy when observing it. And with respect to delicious tasting gustatory perceptions or disgusting smelling olfactory perceptions, the valence seems even more intrinsic to the intentional content.

Nonetheless, it is reasonable to believe that we have all had perceptual experiences wholly lacking in valence. (Or, if the valence is a dimension of the perceptual modality, then the proper description would be an experience of valence “0”.)³ Further, we know that phenomenally grounded visual processing and evaluative processing often take place in different parts of the human brain (Cela-Condea *et al.* 2011). Brain traumas can leave intact a patient’s capacity to recognize a loved-one’s face yet destroy the ability to feel the emotional connection that one previously felt when seeing the person. Capgras syndrome leaves people in the sad state of thinking their loved ones are imposters (Hirstein and Ramachandran 1997).

And yet, it doesn’t matter whether any human ever has had or could have a perceptual experience or a belief without valence. Even if not, it wouldn’t follow that the valent properties were an essential feature of the perception or the belief *qua* intentional state. The only question is: According to the standard theory of phenomenal intentionality, do paradigmatic cases of intentional mental states – like “*perceiving* a rock” or “*believing* that it will rain soon” – *require* valent phenomenological properties for them to possess the external-world content that they do? Presumably, the answer is “no”. Whatever other phenomenal properties a perceptual or doxastic state might have, the kind of property they possess in virtue of which they qualify as “perceptions” or “beliefs” is nothing other than the conscious awareness of a particular content, no part of which is *essentially* valent. (The one obvious exception are second-order beliefs about valent mental states; this special case will be discussed below.)

³This qualification is thanks to an anonymous reviewer.

If standard theories of phenomenal intentionality are true, then valence is not *necessary* for intentionality. And given the “necessity claim” (above), valent phenomenology alone is not sufficient for intentionality, assuming the valent features are not also accompanied by a conscious awareness of the content in question.⁴ The standard assumption, then, is that valent phenomenal consciousness is neither necessary nor sufficient for *perceptual* or *doxastic intentionality*. Before we consider how either or both of these claims might be challenged, let us first take a rough inventory of mental states, so that we can distinguish the ones that are necessarily valent from the ones that are not.

Although it seems quite possible that *beliefs* could be valence-free, there are many who would agree that the same cannot be said for *all* propositional attitudes. I am going to assume (the widely held view) that there are some propositional attitudes which are, by their very nature, valent. If a creature is going to *fear* that *p* or *delight* in *p*, it must be in a phenomenal state that has a negative and a positive valence, respectively. For someone committed to phenomenal intentionality, fear and delight are not merely functional states, nor can they be reduced to behavioral dispositions. For an agent to be in those states there must be *something that it is like* for the agent to “fear that *p*” and that something must be negatively valent. Likewise, for “delighting in *p*,” that something must be positively valent.

Figure 1 offers an inventory of some of the most important valent and non-valent phenomenological types. No authority is claimed for this classification system. It will fulfill its function if it provides a somewhat plausible way of sorting phenomenological types that are necessarily valent (labeled “valent”) from those that are not necessarily valent (labeled “non-valent”). If there is a taxonomy superior to this one, assume that taxonomy instead. The arguments to follow should go through with any (reasonable) classification system.

While there are many phenomenological types that are not included here (like proprioceptive, interoceptive, kinesthetic and somatic), the three listed in the chart (above) suffice for present concerns: perceptual, cognitive, and algedonic. *Perceptual (or sensory) phenomenology* includes the five sensory modalities, with the visual faculty being exemplary for the rich intentional content that it carries. *Cognitive phenomenology* features propositional attitudes – like belief, fear, desire, hope, etc. Both of the elements that constitute a propositional attitude – the propositional (or content) component and the attitudinal component – have their own phe-

⁴For those who think that *all* mental states have intentional content and, thus, that I am contradicting myself when I speak of “valenced phenomenology” which *lacks* a “conscious awareness of the content” – I must offer a promissory note, with more explanation to come. The point at issue here will ultimately be shown to be fully compatible with the view that all phenomenal states have intentional content.

Types of Phenomenology						
PERCEPTUAL	COGNITIVE			ALGEDONIC		
	<i>Propositional</i>	<i>Attitudinal</i>				
Sensory (Visual, tactile, etc)	Content (that <i>p</i>)	Doxastic (Beliefs)	Evaluative (Fears, Desires)	Pain & Pleasure	Emotions	Other Evaluative
NON-VALENT			VALENT			

Figure 1: Types of phenomenology; for details see text.

nominal character. A mental state’s *propositional phenomenology* is the qualitative character of the agent’s conscious awareness of the proposition (e.g. “that *p*”) that is the content of the mental state. A mental state’s *attitudinal phenomenology* is the qualitative character of the conscious attitude that one bears to the propositional content, attitudes like *believing*, *fearing*, and *desiring* (that *p*). The attitude of belief is not necessarily valent; many other attitudes are.

The third category, *algedonic phenomenology*, is of special importance to the current discussion. The term comes from the Greek $\alpha\lambda\gamma\omicron\varsigma$ ($\acute{\alpha}l\gamma\omicron\varsigma$, “pain”) and $(\eta\delta\omicron\nu\eta)$ ($\eta\delta\omicron\nu\eta$, “pleasure”). I will use it to refer to any phenomenally conscious state (*other than* propositional attitudes) that has a *valence*. (I use it in a way roughly similar to Kriegel 2015). Algedonic states include not only (1) bodily pain and pleasure, but also (2) emotions like embarrassment, happiness, and disgust, and (3) other positive or negative cognitive states like aesthetic pleasure or displeasure, existential angst or elation, and satisfaction or frustration. (Drawing a sharp line between valent propositional attitudes and algedonic states is hopeless because it is so difficult to say when a content is rich enough to qualify as “propositional”. I am assuming only a very rough boundary between them.)

This classification system has two broad phenomenological categories: (I) those phenomenological types that are, by their nature, valent, and include algedonic states and valent propositional attitudes, and (II) those phenomenological types that are not by their nature valent, which include perceptions and non-valent propositional attitudes.

3. Awareness Phenomenal Intentionality vs. Valence Phenomenal Intentionality

It is now time to consider an alternative version of phenomenal intentionality to the standard one. As a point of comparison, let us consider how each theory would handle the least controversial intentional states: Perceptions and beliefs about the external world. (Beliefs about mental

states – especially phenomenal states – would complicate things unduly at this stage.)⁵

With respect to this class of perceptions and beliefs about the external world, we can summarize, as follows, the received version of phenomenal intentionality together with its position on the role of valence as it applies to perceptual and doxastic states about the external world:

NS-Awareness Phenomenal Intentionality (or “*NS-awareness PI*”): For perceptions of and beliefs about the external world, a conscious awareness of the content of those states is *both necessary* (“N”) *and sufficient* (“S”) for genuine intentionality. It is wholly irrelevant (with respect to intentional content) whether the perception or the belief has a valence.

In contrast to this position, we will consider an alternative version of phenomenal intentionality according to which valence is relevant to intentionality. The rest of the paper will be devoted to the construction of such a view. The guiding intuition behind the view is sometimes expressed in general discussions of intentionality but it has rarely been advanced as an explicit *theory*. Bill Robinson (1992) has given the clearest expression of the position I have in mind, but it does fall short of a full-blown theory.

Robinson is an interesting participant in this discussion because he is, himself, committed to functional (not phenomenal) intentionality. That is, he does not think that phenomenal consciousness is necessary for genuine intentionality. Yet, he is not dismissive of phenomenal consciousness. On the contrary, he is an epiphenomenalist, a property dualist who believes that phenomenal properties exist and cannot be reduced to physical properties, yet that phenomenal properties have no causal powers and so do not figure in the explanation of any behavior. For Robinson, as with many contemporary epiphenomenalists, it seems reasonable that intentionality can be functionalized because, on this view, it is only the functional properties of cognitive states which feature in belief-desire explanations of behavior.

While Robinson has been an ardent defender of epiphenomenalism for decades, he did not come by his position without some internal conflict. In particular, he struggled with the competing claims of functional and phenomenal intentionality long before most people were taking phenomenal intentionality seriously. In his early book (Robinson 1992), he goes

⁵Beliefs about valenced phenomenal properties are an obvious exception to standard cases. It can reasonably be argued that one is incapable of understanding an intentional content that includes a *valenced phenomenal concept* if one is incapable of experiencing the valenced property that is being referenced. But this is an exception admitted by *any* theory of intentionality that recognizes a first-person requirement for fully grasping the meaning of terms referring to phenomenal properties, including modified versions of functional intentionality. This is not about valence in particular and it does not follow that intentionality, *per se*, is in any way sensitive to valence.

through a process (in his imagination) of constructing a robot, adding more complex features as he goes, and considering, at each step, whether the new addition confers on the robot anything genuinely *mental*.

In this process, his Robot “Frp” has reached a level of sophistication that makes the following behavior possible: When Robinson tells Frp to go to the pharmacy to get a prescription, it finds a map with which it successfully navigates, takes appropriate buses, finds the pharmacy, pays for the prescription, and brings the prescription back. Robinson claims that by virtue of this performance the words spoken by Frp (e.g. “pharmacy,” “prescription”) have genuine intentional content. Nonetheless, Robinson admits that when we ask the question, “Does the robot *understand* the meaning of the words?”, we could mean two different things because of an ambiguity in the term “understand”. He says (Robinson 1992, p. 54, my emphasis):

... we have a conflict because “understanding” is (naturally enough) used to cover two things. One of these is connection to sensation and feelings; one is connection to appropriate action. We could express the situation of Robot Frp by saying that it understands what it says and hears, but that its actions (including its utterances) ultimately *have no point for it*. We could also say, however, that it does not understand its words (meaning that they ultimately *have no point for it*) but that its words do stand for things to which it is connected both on the input side and in actions. How we put the matter is not so important; what is important is to maintain clear recognition of each of the two distinct characteristics that normal people have and whose possible separability gives rise to the puzzlement about understanding.

In this paragraph, Robinson seems intent on showing the intuitive appeal of both functional intentionality and phenomenal intentionality. This actually serves as a fairly plausible explanation for why John Searle’s infamous “Chinese Room argument” (Searle 1980) inspires such opposition among so many readers. Robinson, like Searle, is asking us what we mean by the word “understand.” When Searle asks if the Chinese Room *understands* Chinese, he thinks the only possible answer is “no”. Of course not all readers agree with him. Robinson has given a nice explanation for why people give two very different responses to the question. Those who are already drawn to functionalism interpret the word “understanding” in the *connection to appropriate action* way, and those already drawn to phenomenal intentionality interpret the term in the *sensation and feeling* way. This makes the Chinese Room more of a Rorschach Test – regarding what people see in the English word “understanding” – than a genuine argument. (This also explains why functionalists are so vitriolic in their disapprobation of Searle’s argument. Rorschach tests can’t be question-begging, but arguments can be.)

Notice the intriguing way that Robinson has chosen to describe the phenomenal properties (the “sensations and feelings”) that humans possess and that Frp lacks in virtue of which humans are capable of “understanding” in this stronger, more phenomenally-rich way. The property that he thinks is essential according to the phenomenal intentionality perspective is the property of “having a point”. Like the phenomenal properties of perceptual experiences – for example, a sunset “having a look” or a sandwich “having a taste” – something’s “having a point” is always indexed to an agent. Just as it is the way something looks to a *particular agent*, so too it is the point that something has *for a particular agent*. This is the essentially first-person character of all phenomenal properties.

A second thing that “having a point” has in common with perceptual experiences is that it is normally attributed to an object or state of affairs. When someone observes a sunset, the state-of-affairs that has a particular *look* to the agent is the projected light of the sun as seen through the atmosphere. But to what state-of-affairs does Robinson attribute the property of “having a point (for an agent)”? This requires some exegesis as the text is open to interpretation. Robinson uses the phrase twice in explaining what Frp lacks:

... its actions (including its utterances) ultimately have no point for it,

and

it does not understand its words (meaning that they ultimately have no point for it).

In the first usage, the kind of thing that Robinson says can “have a point” for an agent are *its own actions* with special mention of “utterances”. Verbal behavior is one particular kind of action. In the second sentence, it is linguistic items, “words,” which can have a point. It is not surprising that Robinson would affirm that it is things like “utterances” and “words” that can have a point for an agent. After all, we are considering competing theories of intentionality, and it is words and utterances which are paradigmatic examples of things that have cognitive content. If we were to assume that it is *only* words and utterances (and the like) that are capable of “having a point” for an agent, we would be left with little clue as to what *kind* of phenomenal property Robinson has in mind.

The only other thing we are told is that it involves “sensations and feelings”. That covers a lot of ground. A defender of the standard (Graham *et al.* 2007) version of phenomenal intentionality might argue that “having a point” is just a quirky, idiosyncratic way of referring to the *conscious awareness* of some content, that non-valent property which is common to both *cognitive* and *perceptual* phenomenology. That Robinson has something else in mind is evident from the first quotation. Robinson

says that the difference between a human being with a rich phenomenology and a conscious-less robot is that its actions *in general* – and not just its verbal behavior – “have a point” for it. Utterances are a very narrow, limited kind of cognitively generated behavior. It is a form of communication in which I seek to convey intentional content to another. For example, I want my audience to be consciously aware of the same proposition that I am currently contemplating. In such speech contexts, it is hardly surprising that cognitive phenomenology is primary. Yet *most* of my intentional states have nothing to do with communication and an approach that focuses exclusively on cognitive phenomenology will not generalize to all of our actions that “have a point”.

Robinson makes clear that he is not concerned here merely with cognitive or perceptual phenomenology of a non-valent type. He is concerned about Frp’s lack of *valent* phenomenological states, which would (presumably) include *algedonic phenomenology and valent propositional attitudes*. Here Robinson has his interlocutor give an argument that one might reasonably use in defense of a version of *phenomenal intentionality* that gives valence a more prominent place than on the received version (Robinson 1992, pp. 51f):

There is a structure and order to our behavior that depends, sooner or later, on consequences for sensations and feelings. Therefore, if we imagine a thing that does not have any sensations or feelings (as the robots in this and the previous chapter may not), we imagine a thing from which we have removed the ultimate purposes of action. There would be no real point for it to its doing anything. (There might be a point for its owner. For example, if some activity protected it, or fueled it, its owner would have an interest in its doing the activity.) If it went to the store, bought chops, cooked them, and consumed them, its motions would be all performance with a hollow core of pointlessness, no matter how perfect the exterior execution might be. Since there would actually be no point (for itself) to its actions, there is no way it could use an examination of its actions to understand their point. So, even if it could string together a series of sentences that expressed reasons for doing something, it would not really understand them, because it would not understand the last terms (that is, the ones for sensations and feelings) that describe what is ultimately the point of actions.

The first thing to note about this passage is that it seems clear that the “sensations and feelings” that would give a *point* to its behavior are *valent* phenomenal states, like algedonic states and valent propositional attitudes. Our actions *have a point* for us because we *care* about them, they *matter* to us. And to say that “we care” or that “they matter” just is to say that the possible outcomes of our actions make us happy or sad, they give us pleasure or pain, or they are the occasion for great satisfaction

or frustration. Robinson says that nothing has a point for Frp, that no matter how “perfect” its performance, it lacks a phenomenally rich kind of understanding, not for lack of any cognitive phenomenology but for a lack of *valent* phenomenology, which Robinson seems to think is the more serious deficiency.

It is not yet obvious, however, how to take this idea of things *mattering* to an agent (their “having a point”) and turning it into a new version of phenomenal intentionality. Robinson has identified a salient distinction between functional intentionality and phenomenal intentionality that is not exclusive to speech acts but is present in all intentional acts. We have expressions that capture the pointlessness of an action that comes too late to accomplish its intended goal: “*There’s no point* in closing the barn door after the horse has bolted”. When we speak of the point of something it is usually in regard to the way it contributes to (or fails to contribute to) the satisfaction of a particular desire. For most of us, the *point* of spending many hours planting a flower garden in April is not just the digging and weeding and planting itself, but instead the aesthetic pleasure we enjoy in the summer when the flowers bloom. There would be *no point* (for most of us) to plant in April if we knew a flood in May would wash away all of our plants.

If we interpret Robinson’s words in this way, we can use the flower-planting case to draw the distinction between a human who possesses phenomenal intentionality and a conscious-less robot who lacks it. If robot Frp were to plant a garden it would have *no point* for it because Frp has no desires grounded in phenomenally conscious experience (call these “phenomenal desires”). It is incapable of experiencing aesthetic pleasure of a phenomenal kind and so its “perception” of the rose in July doesn’t have a *point* for Frp the way it would for us.

This is not to say, of course, that Frp couldn’t have *functional* desires. Frp’s primary behavioral imperative could be to plant flowers in April. In fact, Frp’s executive control mechanisms could have *all* of the functional properties that human gardeners’ phenomenal desires have. Frp’s flower-planting-behavior could be conditional upon first computing a low probability of a garden-wrecking flood in May. Frp might even say, “There is no point for me to plant roses in April because a flood is coming in May.” Frp could behave *as if* it mattered to it that there be blossoms in July because its behavior (verbal and otherwise) could be regulated by a *functional* desire for blossoms in July. But Robinson’s claim seems to be that no functional properties, no mere dispositions to behave – absent sensation and feeling – are sufficient for its actions to have a point for it. Since it has no phenomenal desires or any other valent states, none of its actions will have the kind of “point” for it that phenomenal consciousness makes possible for us.

Before turning to the task of applying these insights to a theory of

intentionality, we remind the reader, lest there be any confusion, that Robinson will not ultimately support any of these efforts. He is committed to *functional* intentionality and will reject all possible versions of phenomenal intentionality. He holds that Robot Frp's words have genuine intentional content by virtue of Frp's functional-behavioral properties alone. His insight was simply to draw attention to a family of phenomenal properties that contribute to their "being a point" for an agent. He acknowledges that these properties enhance the value of human existence but he ultimately concludes that these properties are not necessary for intentionality. So with one hand he honors Frp with the classification "intentional agent" and with the other he concedes (and almost bemoans) that Frp is plagued by "a hollow core of pointlessness". Robinson may be interested in where our road leads, but he will not be taking the journey with us.

Assuming that we have accurately captured Robinson's central point, how might this impact a theory of intentionality? There is both a stronger and a weaker position we could take. Let us begin with the stronger. On the strongest, priority is given to the following brief passage from Robinson:

... a thing [without] ... sensations or feelings [is] ... a thing from which we have removed the ultimate purposes of action. There would be no real point for it to its doing anything.

If there is "no real point for it to its doing anything," then this applies as much to its own (purported) intentional states, like perceptions and beliefs, as to anything else and the fact that none of Frp's own (candidate) intentional states *matter to it could* undermine their status as genuine mental states. But why exactly? The reasoning is this. Frp has no desire – no phenomenal (valent) desire – to believe what is true. Its own "beliefs" *don't matter* to it. Truth and falsity *don't matter* to it. Its own behavior *has no point* for it.

Of course you can still predict Frp's behavior by appeal to its functional beliefs (its information-bearing states) and its functional desires (its behavioral control conditionals), but this is nothing but the operations of a conscious-less mechanism. There exists no genuine *agent* with whom any intentional content could reside. If Frp *did* care, if things *did* matter to it, if it had a strong phenomenal desire to know truths about the world, and if it enjoyed rich algedonic responses to its own cognitive states as well as states of the world, then this valent perspective would constitute a locus of agency.

It is now time to consider the theory we will call *valence phenomenal intentionality*, which comes in different strengths. It can be advanced as a *necessary* condition, a *sufficient* condition, or both. We will leave until later the choice among these options and begin with the generic version.

Valence Phenomenal Intentionality (the generic version): For an agent to have genuine intentional states it is either *necessary* or *sufficient* or both that it has states with contents that *matter to the agent*. For something “to matter to an agent” it must be the case that the agent has valent states that are sensitive to the output of content-bearing states.

Robinson’s insight is of most interest if we interpret it as supporting some version of this position. Some might argue that we may have overreached. There is an alternative, weaker interpretation of the insight, which will be described only briefly. Earlier, the definition of *awareness phenomenal intentionality* focused exclusively on content about the *external world*. This was done to exclude mental states where the content of one’s thought is another (valent) mental state. It is now time to bring phenomenal concepts back into the picture.

We have intentional states that include *phenomenal concepts*. I have a phenomenal conception of pain when I am thinking about it in terms of how it *feels*. An agent incapable of experiencing valent phenomenal state *V* will be incapable of having an intentional state that includes the phenomenal concept *V*. If we assume this is true, as do virtually all defenders of phenomenal intentionality and even some defenders of functional intentionality, then a theory of intentionality should make the following weak concession to valent phenomenal states:

The Phenomenal Concept Requirement: If an agent is to have an intentional state that includes a phenomenal concept, then the agent must have first-person experience of the phenomenal property picked out by the concept.

Note how weak a concession this is. Functionalists about intentionality who believe that robots have a complex set of intentional states can nonetheless admit that they are incapable of intentional states which include phenomenal concepts. This is of a kind with the admission that human beings do not have full mastery of the phenomenal concept of bat-echolocation, because we don’t know *what it is like* to “sonar” objects in our environment.

A weaker interpretation of Robinson’s original text might support only an expanded scope for phenomenal concepts so that they are invoked by many more terms in the English language than we might have initially supposed. In addition to obvious properties like the “taste of pineapple” or the “look of a red rose”, it may be that “gardening” and “making dinner” and many other terms for human behavior are subtly infused with meaning drawn from the valent phenomenal experiences that ultimately motivate those behaviors. If so, then a machine incapable of phenomenal

consciousness will not understand what gardening *is* even if it can distinguish cases of gardening-behavior as reliably as any human. On this weaker interpretation of Robinson's insight, it isn't intentionality *per se* that requires valent phenomenology but simply a *significant* percentage of terms in the English language. On this reading, a robot would have genuine intentionality, it would only suffer from a much reduced working vocabulary (Robinson 1992, pp. 51f, especially encourages this interpretation).

It is probably evident at this point that I am not ultimately in the business of "William Robinson exegesis". I will not speculate about whether Robinson intended his talk of "understanding" grounded in sensation and feeling to be interpreted as a nod towards valent *phenomenal intentionality* or only to an *expanded scope for phenomenal concepts*. Whether true to Robinson's intentions or not, we will focus on the stronger view and consider the various versions of *valent phenomenal intentionality*, because they offer a fascinating challenge to established orthodoxy. In the space remaining, we will consider two hypothetical cases which will allow us to compare, side-by-side, the virtues of the standard view and its new challenger.

4. Creatures without Valence vs. Creatures without Cognitive Awareness

In testing our considered judgments about which properties are and are not necessary for intentionality, it is standard practice to compare humans (with a rich phenomenology) to robots (with no phenomenology). The problem with these test cases, for present purposes, is that they provide no help whatsoever in discerning the relative importance of non-valent cognitive phenomenology as compared to valent phenomenology. Humans have both; robots have neither. What is required are cases where the two kinds of phenomenology *come apart*, with one creature having only the first type and another creature having only the second type. (Figure 2 shows the phenomenological capacities of the two creatures to be considered.)

4.1 Case 1: A valcogniphien Has Cognitive But No Valent Phenomenology

The standard view, *NS-awareness PI*, reduced to its simplest principle, holds that a "conscious awareness of the content" of what one is perceiving or believing is both necessary and sufficient for having intentional states of a perceptual or doxastic kind. Since valence is no part of this picture, let us imagine a creature with a brain *capable* of producing non-valent conscious awareness of content – of both a sensory and a propositional kind

		Types of Phenomenology						
		PERCEPTUAL	COGNITIVE			ALGEDONIC		
			Propositional	Attitudinal		Pain & Pleasure	Emotions	Other Evaluative
		Sensory (Visual, tactile, etc)	Content (that <i>p</i>)	Doxastic (Beliefs)	Evaluative (Fears, Desires)			
		NON-VALENT			VALENT			
Creatures	Human	✓	✓	✓	✓	✓	✓	✓
	Aval (w/o valence)	✓	✓	✓	X	X	X	X
	Algedon (w/o cognitive)	X	X	X	✓	✓	✓	✓
	Robot	X	X	X	X	X	X	X

Figure 2: Phenomenological capacities of the creatures considered in the text.

– and a brain *incapable* of producing valent experiences of any kind. Call this creature “Avalcogniphen” or “Aval” for short. The “a-val” captures the fact that it has *no* valent phenomenology; the “-cogniphen” reflects its capacity for cognitive phenomenology. (Henceforth, we will use “cognitive phenomenology” as an umbrella term to refer to any phenomenal state in which the intentional content of the state is determined by a conscious awareness of *what* is being perceived, believed, or otherwise cognized. This is broader than standard usage.)

We will consider two kinds of Avals. A “passive Aval” has no behavioral dispositions of any kind. Not only does it lack the kind of behavioral dispositions grounded in valent phenomenological states that motivate human behavior, but it also lacks the kind of purely mechanistic, behavioral control mechanisms that produce autonomous behavior in robots. It lacks the capacity for any self-generated behavior. Its sole capacities are that of a “cognitive observer”. It sits wherever it is placed, it eats (and doesn’t die) if someone feeds it, and it observes (in a phenomenal way) whatever part of the world is present to its sensory faculties. Since none of its mental states have valence, *nothing matters* to it. It *doesn’t care* about anything that it sees or believes. It doesn’t care where it is or where it goes.

The second type, the “active Aval,” is just the same except that it has behavioral control mechanisms that produce autonomous movement. These mechanisms are accompanied by no phenomenal experiences of any kind, and *active Aval* has no phenomenal desires to do one thing rather than another. It has only conditional behavioral imperatives, like a robot, which determine how it behaves in its environment. None of its behaviors *matter to it* because it has no phenomenal response to any of its own

behavior. This means that most of its propositional attitudes function differently from ours.

“Beliefs” are similar for humans and Avals. The Aval has conscious awareness of both “the proposition” believed and “that” it believes it. Many other propositional attitudes, however, are valent for us, but are not even phenomenal for Aval. The Aval is aware of the proposition that is “desired”, or “feared”, but the “attitude” toward that proposition is a functional disposition wholly lacking in phenomenology. When Aval “fears that p ” it is disposed to aversive behavior when *consciously aware* that p (is true), but there is no valent phenomenology that accompanies that aversive behavior (although there is conscious awareness that the resulting behavior takes place). Like humans, it is consciously aware of what it perceives and believes and how it behaves, but it has no positive or negative phenomenal responses to any of it. It simply *doesn't care* what it perceives, believes, or does.

Our primary focus will be on the case of *active Aval* because its behavior will bear closer resemblance to human behavior, but the passive case is instructive because it shows how little is required to meet the requirements of *awareness PI*. According to the standard theory both active and passive Aval have genuine intentional states. Their perceptual and doxastic states meet all the same standards for intentionality that ours do. There is certain content (like valent phenomenal concepts) that is beyond their cognitive reach, but we have phenomenal concept limitations as well. According to *NS-awareness PI*, these limitations do nothing to threaten the Aval's general capacity for intentionality.

Valence PI challenges these assumptions. The generic version of this theory, defined above, did not choose between three different alternatives. Let us consider the two weakest, and therefore the most defensible, versions beginning with valence as a necessary condition:

N-Valence Phenomenal Intentionality. For an agent to have genuine intentional states it is *necessary* that it has mental states with contents that *matter to the agent*.

Remember, that for a content to “matter to” an agent it must be the case that the agent has positively and/or negatively valent states that are responsive to the content-bearing states. If *N-valence PI* is true then Aval lacks intentionality altogether. It has neither perceptions nor beliefs with genuine intentional content.

How plausible is this position? Defenders of *NS-awareness PI* will answer, “not very”! They will argue that once we say that Aval “is consciously aware of the meaning of p ”, it is simply a contradiction to deny that it *understands* the content “that p ” or to deny that it has a genuine mental state with the content “that p ”. So what that nothing *matters* to

Aval? That is a property wholly different from *understanding*, *thinking* about, or *believing*.

I admit that, in the end, I may agree with this reasoning and reject N-valence PI. However, some push-back is possible. Consider how diminished one's mental life is once all valence has been stripped from it. I suggest that such a creature has far more in common with a consciousnessless robot than it does with a normal human being. If we were to live in community with such creatures, knowing that they didn't care one way or the other about any of the contents with which they engaged, we might discover that we became wholly careless in our interpretation of their utterances.

Ultimately, we might feel little or no obligation to interpret the creatures' thoughts and utterances as having the content that awareness PI assigns to them. If this were to happen, it wouldn't exactly support the conclusion that Avals lack intentionality, but it might well have the consequence that we would judge that the intrinsic value of *non-valent intentionality* would be so diminished (as compared to the *valent* kind) that "having genuine mental states" would be of less consequence than what has traditionally been assumed by defenders of phenomenal consciousness.

4.2 Case 2: Algedon Has Rich Valent Experience and No Cognitive Phenomenology

Consider next the position that valent phenomenology is *not* necessary for intentionality but that it is *sufficient*:

S-Valence Phenomenal Intentionality: For an agent to have genuine intentional states it is *sufficient* that it has mental states with contents that *matter* to the agent.

To consider the plausibility of this theory, consider a second unique creature: Algedon (see Figure 2 above). Algedon is capable of a rich array of valent experiences – pains and pleasures, joy and sadness, aesthetic pleasure and disgust, despair and elation, and satisfaction and frustration in response to its level of goal-achievement. Things *matter* to Algedon. It cares about what it does, what it believes, how the world is, and much more. At least, that is one way of describing the situation. Settling on a proper description is complicated by the fact that Algedon has no cognitive or perceptual phenomenology. Its perceptual and cognitive capacities meet popular contemporary requirements for *functional intentionality*, but it has no "conscious awareness" of anything it perceives or believes.

Algedon has non-phenomenal, mechanistic sensory faculties that produce functional states that "carry information about" the environment

(Shannon and Weaver 1949, Dretske 1981), but those states have no phenomenal properties. We might imagine a cyborg with a properly functioning amygdala and the neural foundations for other valenced phenomenology but with all sensory and cognitive capacities replaced by silicon-based computational and sensory devices that (i) take in visual information from the environment with a digital camera, (ii) generate symbols which meet information-theoretic standards for being “representations” of the objects perceived, and (iii) write statements into a “belief box” that meet its justificatory standards – all *without conscious awareness* of the content of those states.

And yet, Algedon *cares about* the outcomes of those perceptual and doxastic computations, at least in this respect: If its belief box contains the sentence “my-dog-Scruffy-enjoys-being-my-pet”, that fact will produce a fully realized phenomenological experience of happiness; if it contains the negation of that sentence, it will produce phenomenal unhappiness. If the visual system takes in information about a sunset with a certain array of red-green-blue values, it will produce in Algedon strong feelings of aesthetic pleasure; other visual arrays will produce feelings of aesthetic disgust. Its rich algedonic phenomenology is sensitive to the purely functional and information-bearing properties of its computational states.

Does Algedon meet the requirements of *S-Valence PI*? Algedon has valent phenomenological states which are sensitive to its changing computational states which reliably track states of the environment. So there is at least a sense in which those information-bearing states *matter* to Algedon. Nonetheless, a defender of *NS-awareness PI* will likely argue that the computational states, unaccompanied by any conscious awareness of what those states *mean*, does not rise to the level of genuine “content”. The claim would be that Algedons may have content-less *raw feels*, which are evoked by certain syntactic objects, which bear causal relationships to states of the environment, but that does not mean they actually *believe* that the train has left the station or that it sees the tree in the quad.

Much of the force of this argument undoubtedly derives from the central role that “conscious awareness of content” plays in our own mental life. And yet, we must be careful that the harsh judgment that Algedons lack intentionality is not motivated by a purely abstract commitment to a philosophical theory that ignores the relevance of the value we would naturally place on the inner mental life of the Algedon. If we were to actually live in a community populated by large numbers of Algedons, we might be surprised to discover that we unself-consciously empathized with them because of the obvious richness of their valent mental states and that we spontaneously interpreted their mental states as having genuine intentional content.

5. Tentative Conclusions

The strongest conclusion we might draw from our two cases is that Avals do not have intentionality and that Algedons do. That would support *NS-valent P* and reject both elements of *NS-awareness PI*. A weaker but more defensible option is to say that neither cognitive nor valent phenomenology is *necessary* for intentionality but that either one is *sufficient*. On this view, *S-awareness PI* and *S-valent PI* are both true. This is the theory one would hold if both Avals and Algedons were judged to be genuine intentional agents.

If this were true, there would be two different ways for phenomenal consciousness to ground intentionality: (1) by having phenomenal states that carry content via a conscious awareness of that content or (2) by having valent phenomenal states that elevate information-carrying states to the level of intentional content. Humans would have a bounty of riches, meeting both conditions.

One final reflection. It is one thing to advance theories of intentionality in the abstract. It is another thing to observe actual members of a linguistic community populated by diverse creatures and see what judgments each makes regarding the intentionality of the others. Since we do not yet live in such a community, we must ask instead: To what extent would creatures (like ourselves) with a full complement of both cognitive and valent phenomenal experiences *respect as intentional agents* creatures with only valent phenomenology (Algedons) and creatures with only cognitive phenomenology (Avals)?

To raise this question is to think of intentionality not as a theoretical designation or a natural kind but as a social practice. If this distinction can be born out, then we might distinguish between intentionality *de jure* and intentionality *de facto*. In a previous paper (Anderson 2012), I began to explore possible connections between the moral status of a cognitive agent and its status as an intentional agent. These considerations provide further support for the view that in a society populated with diverse creatures, a race of Algedons might earn a kind of *respect* as intentional agents that exceeded that enjoyed by a race of Avals. Is that relevant to a theory of intentionality or not? The standard response would probably be “no”! One goal of this paper has been to explore reasons for thinking the correct answer might be “yes”.

Acknowledgments

I am grateful for helpful comments on the topics discussed in this paper from John Barker, Jeff Berends, Daniel Breyer, Tom Buller, Chris Horvath, Gina Schouten, Todd Stewart, Harry Deutsch, Pat Francken,

Jim Swindler and the participants in “The Reality of Consciousness” conference in Chicago – Harald Atmanspacher, Luis H. Favela, Georg Franck, Scott Jordan, Lana Kühle, Bill Seager, Serife Tekin, and Jeff Yoshimi. Detailed comments by two reviewers were especially helpful. Special thanks go to Charles Siewert who has, more than anyone else, focused attention on the moral aspects of valent consciousness and who challenged my earlier work on this topic and helped to motivate this paper.

References

- Anderson D.L. (2012): Machine intentionality, the moral status of machines, and the composition problem. In *Philosophy and Theory of Artificial Intelligence*, ed. by V.C. Müller, Springer, New York, pp. 321–333.
- Cela-Condea C.J., Agnati L., Huston J.P., Mora F., and Nadal M. (2011): The neural foundations of aesthetic appreciation. *Progress in Neurobiology* **94**(1), 39–48.
- Dretske F. (1981): *Knowledge and the Flow of Information*, MIT Press, Cambridge.
- Graham G., Horgan T., and Tienson J. (2007): Consciousness and intentionality. In *The Blackwell Companion to Consciousness*, ed. by M. Velmans and S. Schneider, Wiley, New York, pp. 468–484.
- Hirstein W. and Ramachandran V.S. (1997): Capgras syndrome: A novel probe for understanding the neural representation of the identity and familiarity of persons. *Proceedings of the Royal Society B* **264**, 437–444.
- Kriegel U. (2015): *The Varieties of Consciousness*, Oxford University Press, New York.
- Millikan R. (1984): *Language, Thought, and Other Biological Categories*, The MIT Press, Cambridge.
- Robinson W. (1992): *Computers, Minds and Robots*, Temple University Press, Philadelphia.
- Searle J. (1980): Minds, brains and programs. *Behavioral and Brain Sciences* **3**, 417–457.
- Shannon C.E. and Weaver W. (1949): *The Mathematical Theory of Communication*, University of Illinois Press, Champaign.
- Siewert C. (1998): *Significance of Consciousness*, Princeton University Press, Princeton.

Received: 24 September 2016

Revised: 31 December 2016

Accepted: 28 January 2017

Reviewed by Vincent C. Müller and another, anonymous, referee.