Cognitive Critique

SAVING TIME: HOW ATTENTION EXPLAINS THE UTILITY OF SUPPOSEDLY SUPERFLUOUS REPRESENTATIONS

JASON FORD

Department of Philosophy, University of Minnesota, Duluth, Minnesota

E-MAIL: jford@d.umn.edu

Accepted 26 November 2008

KEYWORDS

perception, attention, Alva Noë, enactive, sensorimotor, consciousness, David LaBerge, inattentional blindness, periphery, representation

ABSTRACT

I contend that Alva Noë's Enactive Approach to Perception fails to give an adequate account of the periphery of attention. Noë claims that our peripheral experience is not produced by the brain's representation of peripheral items, but rather by our mastery of sensorimotor skills and contingencies. I offer a two-pronged assault on this account of the periphery of attention. The first challenge comes from Mack and Rock's work on inattentional blindness, and provides robust empirical evidence for the semantic processing (and hence representation) of some wholly unattended stimuli. The second challenge draws on LaBerge's theory of attention to provide a substantial advantage to peripheral representations, saving time whenever we shift the focus of our attention to something which had been in the periphery, allowing us to respond to that thing more quickly than would be possible if Noë's account of perception were correct.

INTRODUCTION

The Enactive Approach to Perception (O'Regan and Noë 2001) is currently a hot topic in philosophy of mind. This theory holds that perception is an active skill, not a passive reception, and that much of our conscious experience (everything outside the focus of attention) is not grounded in representations of the items peripherally experienced, but by our mastery of "sensorimotor contingencies", knowing where and how to look for more information about the parts of our environment that are peripherally experienced. This is a significant consequence of the Enactive Approach, and I will use this claim to test the scope of the Enactive Account – to prove that it should not have the impact ascribed by its proponents (that the Enactive Account should ground a completely new approach to psychology and cognitive science[1]). I have two lines of argument, the first is based on Arien Mack and Irvin Rock's research on inattentional blindness. I contend that their results show that our brain creates representations that are far richer than Noë's account of perception will allow. The second line of argument is based on David LaBerge's model of attention. If I am right in extending LaBerge's model, then there is an evolutionary advantage to be had in producing rich peripheral representations, even if most of them never enter the focus of attention. Once that is done, however, I do not want to reject the Enactive Account outright – I believe that it has much to offer, once the more extreme claims made on its behalf are moderated. I will conclude by discussing the parts of the Enactive Account that I believe can be retained and incorporated into an account of conscious experience that does justice to attention.

Before I begin, however, I need to make several prefatory remarks about attention, its role in conscious experience, and the terminology that I will be employing. I hold that attention is the central structural feature of our conscious experience. I believe that careful attention to the way that attention shapes our conscious experience will allow us to solve some problems in both philosophy and psychology.

For a rough working definition of attention, I turn to James' *Principles of psychology* (1890, pp.381-382): "It [attention] is the taking possession of the mind, in a clear and vivid form, on one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence." Attention is whatever the thing you are most conscious of

right now has, and that the other things you are conscious of have less of. I'll call "the focus of attention" that part of your conscious experience where attention is most acute, and the rest of your field of consciousness "the periphery of attention". I contend that every moment of conscious experience is structured into a focus of attention and a periphery. There are varying degrees of salience within the periphery of attention, from nearly as strong as the focus all the way out to the dim edge of consciousness. Everything in the periphery of attention is part of our conscious experience, but we are most aware of what is in the focus of attention.

I also note, to avoid confusion, that the focus of attention and periphery of attention are distinct from the visual focus and visual periphery. The content of the focus of attention is whatever you are most aware of at any time (in any sensory modality), while the content of the periphery of attention would be all the rest of the things of which you are aware to some lesser degree. The content of the visual focus is the part of the world that is reflecting or producing the light that strikes the fovea, while the content of the visual periphery would be the parts of the world that are reflecting or producing the light that strikes the parafovea. We can distinguish the visual focus from the focus of attention by fixing our eyes on some object and then paying attention to the experience at the edge of our visual field. We don't do this very often, but the experiments to come will rely heavily on this distinction.

Trying to get at the periphery of attention directly is an impossible task, for when we try to grasp more about something in the periphery of attention, the only way that we can do so is by shifting the focus of attention to it. Doing that changes the experience as a whole, and shifts the character of that previously peripheral aspect of experience from peripheral to focal. We can never focus on the periphery of attention as peripheral, but I contend that it is a genuine component of our conscious experience, and that the periphery contributes to the experience of the content in the focus of attention. With those observations and distinctions in place, let us now turn to Noë's theory of perception.

THE ENACTIVE ACCOUNT

The main goal of the Enactive Account is to move our understanding of perception from a passive process to an active one. According to Noe:

> Perception is not something that happens to us, or in us. It is something we do. Think of a blind person tap-tapping his or her way around a cluttered space, perceiving that space by touch, not all at once, but through time, by skillful probing and movement. This is, or at least ought to be, our paradigm of what perceiving is. The world makes itself available to the perceiver through physical movement and interaction…. I argue that all perception is touch-like in this way: Perceptual experience acquires content thanks to our possession of bodily skills. What we perceive is determined by what we do (or what we know how to do); it is determined by what we are ready to do. In ways I try to make precise, we enact our perceptual experience; we act it out. (Noë 2004, p.1).

So far, the Enactive Account might only be bringing a hitherto unnoticed aspect of our experience to light, so that we might attend to the contribution that our sensorimotor skills make to our conscious experience. If that were all that Noë was doing, I wouldn't have any dispute with him, and in the end, I will argue that the Enactive Account should be preserved just as the claim that our sensorimotor skills make contributions to our experience that are often neglected.

Noë takes his account to be an alternative to the "snapshot conception" of perception (Noë 2004). According to the snapshot conception, we visually construct our entire visual field, in uniform detail from the edges of the visual periphery to the visual focus, by scanning our environment. That detailed picture is the sort of internal representation that Noë attacks at the outset. Again, if that were all that Noë was doing, I would have no dispute with him. As Noë points out, the picture model of perception is phenomenologically unmotivated, and is empirically undermined by the discovery of change-blindness and inattentional blindness. I am not defending the snapshot conception.

However, Noë goes much further and attempts to undermine most, if not all, of the representations[2] that might be thought to ground our conscious experiences. When it comes to the periphery of attention, Noë claims that our mastery of sensorimotor skills wholly constitutes the periphery of experience, as we see in the following passage:

> Surely it is a basic fact of our phenomenology that we enjoy a perceptual presence of at least some unattended

> features of the scene. So, for example, I may look at you, attending only to you. But I also have a sense of the presence of the wall behind you, of its colour, of its distance from you. It certainly seems this way… We must explain how it is that we can enjoy perceptual presence of unattended features of a scene.
>
> To begin to see… a solution to the problem of perceptual presence, consider as an example a perceptual experience such as you might enjoy if you were to hold a bottle in your hands with eyes closed. You have a sense of the presence of the whole bottle, even though you only make contact with the bottle at a few isolated points. Can we explain how your experience in this way outstrips what is actually given, or must we concede that your sense of the bottle as a whole is a kind of confabulation? (Noë 2002, p.8)

First, notice that what Noë calls "perceived but unattended", really belongs in the periphery of attention, if we adopt my way of describing conscious experience.[3] If the Enactive Account is correct, we don't really have peripheral experience, we have the experience in the focus conjoined with the implicit, skill-based knowledge needed to get more information, as we need it. According to Noë, sensorimotor knowledge explains how we can have the sense of perceptual experience that goes beyond our actual perceptions. This gives us the illusion of a richer stream of consciousness than is actually present, and is why Noë thinks that the experienced periphery poses a problem. Noë realizes that his theory requires radically different brain-level mechanisms for the focus of attention and the periphery of attention. He calls the need for a different account of the periphery "the problem of perceptual presence". Noë holds that our perceptual presence of the periphery is the result of a process of amodal completion (Noë 2004). In amodal perception, we fill in (complete) an incomplete representation, as with Noë's favorite examples – Kanisza triangles and the experience of seeing a cat behind a picket fence (we see the cat as a whole cat, not just the parts that are strictly visible). So, on the Enactive Account, our peripheral experience is very sparse, with our sense of a complete world available for our exploration created by amodally completing the rest of the visual field.

With the basics of the Enactive Account in place, we reach the challenge I would like to focus on, in the form of a rhetorical ques-

tion that I will attempt to answer:

> But consider a more basic point: why should the brain go to the trouble of producing a model of the bottle when the bottle is right there in your hands and can serve as a repository of information about itself? All the information about the bottle you need is available to you in the world – you need only move your hands to gather it…
>
> The enactive, sensorimotor account explains how it can be that we enjoy an experience of worldly detail *which is not represented in our brains*. The detail is present – the perceptual world is present – in the sense that we have a special kind of access to the detail, an access controlled by patterns of sensorimotor dependence with which we are familiar. [Italics added] (Noë 2002, pp.10-11)

That last passage can easily be made into an evolutionary argument against the existence of peripheral representations[4]. Why should the brain bother to produce representations of all that peripheral stuff, when having the knowledge of where to look to get it will do the job just as well? Making all those peripheral representations, in all the sensory modes available to us, would take up a lot of brain processing capacity. If letting the external world serve as its own memory store will get the job done, with considerable savings in the brain's cognitive resources (which can be better employed elsewhere), then beings whose brains did not produce peripheral representations but relied instead on enactive perceptual mechanisms and sensorimotor contingencies would have a substantial advantage over their periphery-befogged brethren. Should we conclude that we are those creatures with the purported evolutionary edge provided by the enactive approach to perception?

I contend that we should not, and I have a two-pronged assault on Noë's position on peripheral experience. First, I will show that Mack and Rock's investigations into inattentional blindness provide very strong empirical evidence that we do have detailed representations of items in the periphery of attention. Second, my theory of how attention works (based on and extending the work of David LaBerge) will answer Noë's challenge and show that there is a substantial advantage to be had in representing peripheral items when compared with the Enactive Account of those peripheral items.

THE FIRST COUNTERARGUMENT: MACK AND ROCK'S EVIDENCE FOR UNATTENDED REPRESENTATIONS

In their *Inattentional blindness* (Mack and Rock 1998), Mack and Rock discovered that some objects presented outside the periphery of attention are actually processed up to the level of semantic content. That requires a detailed representation of the item, and if this happens even for things that are not experienced at all, then things that appear in the periphery would certainly also have representations (showing that they are, in fact, processed and represented in the brain).

Inattentional blindness is a fascinating phenomenon, and one that seems counter-intuitive to many. It is possible for an item to appear well within a person's visual field, and for that person not to see it. Indeed, if the focus of attention is directed into the visual periphery, then an item that appears in the visual focus may also go wholly unnoticed. Mack and Rock (along with a bevy of graduate students) conducted a series of experiments exploring this phenomenon, which I will briefly describe.

The general form of most of their experiments took a subject who was told to perform a task requiring attention (when shown a cross on a computer screen, the subject must judge which arm is longer, the vertical arm or the horizontal arm – the arms of the crosses vary in length, so sometimes this task will be easy and sometimes difficult). The subject performs this task a few times. Then, on the third or fourth trial, something extra is presented along with the cross, well within the subject's visual field (within 2.3 degrees of visual fixation, or at fixation if the distractor task is in the subject's parafovea). Surprisingly, many subjects do not report seeing anything in addition to the cross (Mack and Rock 1998). For some of the experiments, the cross is presented at the visual fixation point and the extra stimulus is presented in the parafovea. For the experiments that we will be focusing on, the cross is presented in the parafovea and the extra stimulus is presented at the point of visual fixation (making the failure to see it all the more surprising). For most of the stimuli (common words, shapes, colored spots) presented at the fixation point (with the cross in the parafovea), inattentional blindness (IB) averaged 61% (Mack and Rock 1998).

Mack and Rock found that some types of stimuli were seen under conditions that produced inattentional blindness in most

cases. These stimuli evidently have the power to attract attention even when we are consciously focusing on something else. These stimuli were: the subjects' own names, cartoon smiley faces, and large items (1 degree of the visual field or larger). The first two indicate semantic processing of unattended stimuli (spotting large items could easily be the result of an early-selection process).

The experiment involving names (with the cross in the parafovea and the extra stimulus at the point of visual fixation) compared the rate of inattentional blindness for the subject's own name with that name with one vowel altered. That is, "Jason" would have been compared with "Jeson" (trials were also run with common words like "Time", and other types of stimuli). The subject's own name was almost always noticed (IB = .5%), but the slightly altered name was frequently missed (as were the other stimuli tested) (IB = 60%), this difference is highly significant ($X2$ (1df) = 13.8; $p < .001$). The whole word must have been processed (hence represented) in order for the real name to attract attention and be consciously experienced, where the misspelled name does not. If the unattended stimuli had only been processed up to the level of shape, then we would see a lot of false positive identifications (people seeing the slightly altered names as their names) (Mack and Rock 1998). Further, if Noë were correct about the mechanism that produces the periphery of attention, that it is primarily a product of amodal completion, then we should expect a goodly number of the subjects to complete the peripherally presented stimulus with their own names (attracting the focus of attention and entering conscious experience as a false positive result). Since that is not what the experiments showed, we have evidence that the periphery of attention is not the product of amodal completion. The power of the subjects own names to attract attention also shows that the periphery of attention is constructed of detailed representations (since that is the only way that correctly-spelled names could draw the subject's attention while misspelled names do not).

More evidence for the presence of semantic content in those representations that make it up to the threshold of attention, but which go unnoticed, comes from priming tests on the unattended stimuli (those which go wholly unnoticed by the subjects). In these, the targets (the extra stimuli) were short words, and the subjects were asked to perform a stem-completion task (in addition to the distraction cross task and being asked if they saw anything extra). If the target word were "Flake", then the stem to be completed would

be "Fla". Mack and Rock chose target words that were relatively uncommon, 4-5 letters, and for which the three-letter stem could be completed by a variety of other English words. They checked the frequency of the target word as a stem-completer among those subjects that were inattentively blind to it and compared it with the frequency found by asking random people to complete the stem with the first two English words that come to mind (the control group) (Mack and Rock 1998, pp. 175-182). They found that the inattentively blind subjects produced the target word significantly more often than the random controls did – the combined results from all the trials in this set of experiments was 34% of the 91 inattentionally blind subjects choosing the target word as one of their two stem-completers, compared with 6.4% of the 140 controls (Mack and Rock 1998, p. 183). Again, this result is highly significant.

To check whether the words were seen as meaningful words or as complex shapes (that is, to check for semantic processing), the stem completion task was replaced with a forced choice task among five pictures, each an object that began with the first few letters of the target word. So, if a subject was presented with the word cork, the choice of pictures was: a coat, corn cobs, a coffee cup, a couch, and a wine bottle with a cork. Out of 60 subjects, 29 (48%) failed to see the target word, and of those 29 people, 14 (48%) chose the picture named by the target word. This is very strong evidence of semantic priming, since the meaning of the target word would have to operate in order to impact the choice of pictures (as it did in almost half the subjects) (Mack and Rock 1998, pp.186-190), see Figure 1 below.

These priming experiments strongly indicate that the target words were perceived, and their meanings registered, but not consciously. Add that into the results that compared names with their near-misspelled versions and we see that Noë's bold claim (that our experience of the periphery is wholly constituted by our mastery of sensorimotor skills via amodal completion) must be false. Mack and Rock clearly demonstrated that even unattended stimuli were represented in the brain and processed all the way up to receiving semantic content that was sufficiently rich to allow comparisons from lexical stimuli (the words) to images (the pictures). If the wholly unattended stimuli are represented, how could we claim that stimuli in the periphery of attention are not represented? If that is sufficient to show that we really do form representations in the periphery of attention, I turn to a different aspect of Noë's argument. I

will show that there may be an evolutionary advantage to be had by forming peripheral representations.

Figure 1: Picture array for choices on the prime, "cork". From Mack and Rock, 1998, page 189 (where it is figure 8.4b), used with the permission of the author.

THE SECOND COUNTERARGUMENT: AN EVOLUTIONARY ADVANTAGE IN PERIPHERAL REPRESENTATIONS

Now I will take David LaBerge's account of how the brain processes the focus of attention, extend it to cover the periphery of attention and the unattended stimuli revealed by Mack and Rock, and argue that if this extended model of attention is correct, then peripheral representations provide us with a significant time-saving advantage in most of our perceptual processing.

According to LaBerge, in *Attentional processing* (LaBerge 1995) the focus of attention is produced when a triangular circuit

of feedback and feedforward connections is made between three parts of the brain: the prefrontal cortex (which controls the direction of the focus of attention), the thalamus (which enhances the firing rates to the needed threshold level), and whatever part of the cortex is representing the object in the focus of attention (vision, hearing, memory, etc.), see Figure 2 below. Several cycles through the circuit, with the thalamus providing continued enhancement, are required in order to reach the threshold of the focus of attention.

The three nodes in LaBerge's triangular circuit are:

The cortical areas that code for sense perceptions, objects and actions provide the expression of attention, which presents the "object" in the focus of attention.

The prefrontal cortex serves to direct and control the focus of attention.

The thalamus serves to enhance the activations at the sites of expression and to enhance the feedback and feedforward connections between the sites of expression and control.

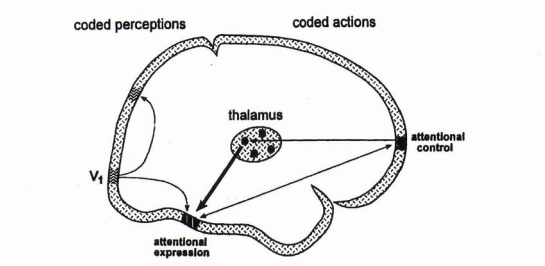

Figure 2: Schematic diagram of the three major (human) brain sites connected in an attentional triangular circuit that results in attentional intensification of activity in a cortical column cluster. The presumed control site in prefrontal cortex directly selects the site of attentional expression in the occipitotemporal cortex, and this control site also regulates the intensity of the cortical expression in the occipitotemporal cortex by means of the indirect connection through a thalamic nucleus. The posterior cortical areas serve mainly perceptions, while the anterior cortical areas serve mainly actions, and sites of attentional expression may exist in both areas. (LaBerge 1997, p.165, used with the permission of the author.)

To extend LaBerge's model to include the periphery of attention we need only make one additional supposition: that the same mechanism which produces the focus of attention is also responsible for the periphery of attention, acting at a lower level of thalamic enhancement (to create the "sense of perceptual presence", in Noë's parlance). Since there is a threshold level of neural activation needed for the focus of attention, we may suppose that there is a lower threshold needed for a representation to enter the periphery

of attention. Below that periphery threshold, representations could be processed for semantic content (as found by Mack and Rock) without entering consciousness at all. I would describe those representations as in the pipeline of attention.

If Noë is correct, then every time we shift the focus of attention to a new object, we must start the perceptual processing mechanism from scratch. If I am correct, then every time we shift the focus of attention to an object that had been in the periphery of attention, then most of the perceptual processing is already accomplished. The brain needs only to enhance the firing rate a bit more in order to move that representation up to the threshold of the focus of attention. Similarly, for unattended representations in the pipeline of attention, a bit of thalamic enhancement is all that is required for the representation to enter the periphery of attention, and from there into the focus of attention. How much time would we save? My best estimate is 150-200 milliseconds (from Thorpe, Fize and Marlot, *Speed of processing in the human visual system*, 1996). That may not sound like much, but if we gain that time in every instance where we respond to something that had been in the periphery of attention (and that is most of what enters the focus of attention), it would be a significant advantage indeed. Saving a few tenths of a second in perceptual processing also lets us respond to our environment a few tenths of a second faster, and that could well make the difference between survival and doom out on the veldt (or on today's freeways, just for instance).

COMPROMISE: WHAT IS RIGHT AND USEFUL IN THE ENACTIVE APPROACH TO PERCEPTION

Though I believe that Noë's claims regarding the representational commitments of the Enactive Approach are too strong, I would not want to reject Noë's Enactive Approach to Perception in its entirety. Perception really is a skill, it depends on our mastery of a lot of body-related "contingencies", and we often lose sight of that fact. I agree that our perceptions are constructed via active engagement with and exploration of the world, and I agree that perception is not snapshot-like.

I can incorporate the main features of the Enactive Account (moderately interpreted) into my attention-based model of consciousness, and if Noë loosens his restrictions on the sorts of de-

tailed representations that can be produced in the brain, he can incorporate my conclusions as well. So, hopefully, this will be more of a cooperative venture than a direct conflict. Both of us are trying to draw attention to features of our conscious experience that have been neglected. The Enactive Account can and should be conjoined with my Attention Model to include peripheral representations in the explanation of the sense of perceptual presence. Our sensorimotor skills operate in conjunction with peripheral representations (not alone) to produce our conscious experience, structured by attention into a focus and a range of degrees of peripheral experience.

ENDNOTES

[1] For instance, "What perception is, however, is not a process in the brain, but a kind of skillful activity on the part of the animal as a whole. The enactive view challenges neuroscience to devise new ways of understanding the neural basis of perception and consciousness," (Noë 2004, p.2).

[2] It is my sincere hope that the dispute between Noë and myself does not boil down to a terminological dispute over what "representations" are. That would certainly make Noë's Enactive Approach less interesting, and undercut his bold claims to a revolution in psychology in a very unsatisfying way. So, in what follows, I will endeavor to show that the very things that Noë denies the existence of ("peripheral representations" – but the argument should remain whatever one chooses to call them) have empirical support for their existence and could serve an evolutionarily useful purpose.

[3] In fact, I think that many problems in contemporary philosophy of mind would dissolve if we did not treat everything outside the focus of attention as if it were "unattended" or even "unconscious". See Ford (2005); Ford and Smith (2006); and Ford (2008) for more instances of this neglect of the periphery of attention and the problems it causes.

[4] I am dealing with Noë's claims about the periphery of attention rather than the focus of attention because he is unambiguous about the periphery (ironically enough). He consistently claims that the periphery of attention is solely constituted by our mastery of sensorimotor skills and contingencies, where he is not consistent in his claims about the focus of attention. At some places, Noë does also claim that our focal experience is wholly constituted by our sensorimotor skills (e.g., Noë 2004, p.2), but in other places, he claims that our sensorimotor skills only constitute part of the content of our focal experience (Noë 2004, p.25).

ACKNOWLEDGMENTS

I would like to thank Wade Savage, Mark Newman, David Cole and Orlin Vakarelov for their thoughtful comments on earlier drafts of this paper. I am also indebted to the comments and questions raised by audiences at the University of Minnesota, Duluth, at the 2007 meeting of the Minnesota Philosophical Society at Inver Hills Community College, at the Center for Cognitive Sciences at the University of Minnesota, Twin Cities, and at the Pacific Division of the American Philosophical Association, 2008. This paper was written with the research support of the College of Liberal Arts of the University of Minnesota, Duluth, and the support of a Chancellor's Grant from the University of Minnesota.

REFERENCES

Ford J (2005) The attention model of consciousness. UMI dissertation. ProQuest, Ann Arbor, MI

Ford J (2008) Attention and the new sceptics. J Consciousness Studies 15(3):59-86

Ford J, Smith DW (2006) Consciousness, self and attention. In: Kriegel U, Williford K (eds) Consciousness and self-reference. MIT Press, Cambridge, MA

James W (1890) The principles of psychology. Harvard University Press, Cambridge, MA, 1981

LaBerge D (1995) Attentional processing: the brain's art of mindfulness. Harvard University Press, Cambridge, MA

LaBerge D (1997) Attention, awareness, and the triangular circuit. Conscious Cogn 6:149-181

Mack A, Rock I (1998) Inattentional blindness. MIT Press, Cambridge, MA

Noë A (ed) (2002) Is the visual world a grand illusion? Imprint Academic, Charlottesville, VA

Noë A (2004) Action in perception. MIT Press, Cambridge, MA

O'Regan JK and Noë A (2001) A sensorimotor theory of vision and visual consciousness. Behav Brain Sci 24:939-973

Thorpe S, Fixe D, Marlot C (1996) Speed of processing in the human visual system. Nature 381:520-522