



Putting the Brakes on Enactive Perception

Jesse Prinz
Department of Philosophy
CB #3125, Caldwell Hall
UNC/Chapel Hill
Chapel Hill, NC 27599-3125
(c) Jesse Prinz 2006
jesse@subcortex.com

PSYCHE 12 (1), March 2006

KEYWORDS: Enactive perception, prosthetic vision, inverting lenses, motor deficits, change blindness, non-conceptual content, pictorialism

ABSTRACT: Alva Noë's *Action in Perception* offers a provocative and vigorous defense of the thesis that vision is enactive: visual experience depends on dispositional motor responses. On this view, vision and action are inextricably bound. In this review, I argue against enactive perception. I raise objections to seven lines of evidence that appear in Noë's book, and I indicate some reasons for thinking that vision can operate independently of motor responses. I conclude that the relationship between vision and action is causal, not constitutive. I then address three other contentious hypotheses in the book. Noë argues that visual states are not pictorial; he argues that all perception is conceptual; and he argues that the external world makes a constitutive contribution to experience. I am unpersuaded by these arguments, and I offer reasons to resist Noë's conclusions.

1. Introduction

Up until recently, cognitive scientists were happy to parcel up the mind-machine into neat parts. We had perception on one side, which is in the business of representing inputs from the external world. Then we had action, on the other side, which controls an organism's outputs, or behavior. Nestled between these "peripheral systems" when had central systems, which were presumed be the main engines of "cognition" or "thinking." Each of these systems was supposed to work independently, like separate committees in a great corporation, only vaguely away of what the others are up to. In cogsci lingo, each system was supposed to use proprietary rules and representations. Oh, how times have changed. We are now living in an era of border disputes. The orthodox divisions of the mind are being attacked. I have tried to join the front lines myself on occasion. I think the border

between perception and cognition needs to be renegotiated: thinking does not use a proprietary code; it redeploys representations used to perceive the world (Prinz 2002). I'm inclined to think that cognition also avails itself of representations used for action. To plan, for example, is to imagine acting without actually executing the imagined actions. On this view, there are no central systems. There are just peripheral systems that can be taken off line to simulate perceptions and actions in the absence of actual inputs and outputs. Thinking is just seeing with closed eyes and acting without moving.

Despite my skepticism about the traditional borders between central systems and peripheral systems, I'm a pretty old-fashioned philosopher. The view I defend is just a version of Classical British empiricism. I think cognitive science should follow the path of Locke and Hume. In recent years, a more radical critique of mainstream cognitive science has been under way. Members of this movement (an apt term) have set their targets on a different border. They think we should not draw sharp lines between perception and action, inputs and outputs. Seeing is a kind of doing, they say. Perception is enactive. Some members of this camp go even farther. They have launched a campaign against a border that most people consider impermeable: the boundary between mind and world. Once we collapse the boundary between perceiving and acting, the old Cartesian model according to which the mind is simply a model of external reality begins to look quaint. If seeing is a form of acting, then it is not principally a matter of representing anything; it is a matter of engaging with the environment. Within this paradigm, it becomes attractive to see the mind as emerging from interactions between internal processes and external entities. Mind doesn't supervene on the brain, as Cartesian materialists would suppose, but rather on a dynamical system containing brain, body, and things surrounding the body.

These radical new border disputes have been carried out by a number of very capable combatants. Soldiers in this battle include Rodney Cotterill, Susan Hurley, Sean Kelly, Kevin O'Regan, Pete Mandik, Evan Thompson, and, of course, Alva Noë. Noë has collaborated with half of these others, and he went to boot camp under the command of Daniel Dennett, who has been an important inspiration for many of those who have been waging war against traditional views in cognitive science. Noë has done a great service to his comrades by authoring a book that summarizes the case for the enactive view. *Action in Perception* is, in many respects, an excellent book. It is a captivating and sustained attack on deeply entrenched borders. Again and again, Noë invites the reader to reconsider fundamental assumptions. The book is engagingly written, and it is a wonderful model of how to do naturalistic philosophy: Noë moves effortlessly back and forth between empirical science and deep philosophical issues. He makes science relevant to philosophy and philosophy relevant to science. He generally resists dogmatism, and offers the enactive view as a plausible and provisional hypothesis, worthy of serious consideration. Rather than insisting that traditional cognitive science has been refuted, he suggests that traditional cognitive science has presupposed the integrity of borders that may not hold up. It is refreshing to read something that challenges us to step back and rethink much of what we have taken for granted. Noë's book does that.

But, after stepping back and after re-thinking, I find myself unpersuaded. Noë would like to undermine the border between action and perception and the border between mind and world. It would be exiting if he were successful, and the thrill of that

revolutionary outcome adds much to the rhetorical power of his book. Many people get swept up in enactive view, because it falls into that category of theories that would be very exciting if true. But we mustn't get seduced by exciting theories. The evidence proffered in favor of the enactive view falls far short of the mark. Despite being skeptical of traditional borders myself, I confess little temptation, on close evaluation, to follow Noë's siren-call. In this review, I will play the role of a reactionary. I will explain why I think we should put the breaks on enactive perception (see also, Prinz, forthcoming). I will focus my discussion on the question of whether there is a boundary between perception and action, but I will also have some brief remarks about several other themes in Noë's book.

2. Inactive Perception

It is a central theme of *Action in Perception* that perception and action are inextricably bound. Often this thesis is articulated using appealing catch phrases that add more charm than precision. We are told, for example, that perception is a "skillful activity," perception is an exercise of "practical knowledge," and perception is "mode of exploration of the environment." At this level of abstraction, defenders of traditional theories of perception need not disagree. In vision, for example, we deploy *skills* of pattern recognition; we know *how* to discriminate colors without knowing *that* they are constituted by thus and such physical magnitudes; and we *explore* the scenes by focusing, scanning, and surveying with our senses. Moreover, no one would deny causal interactions between action and perception: shifting our eyes can change what we see, and what we see can incite us to prepare different motor responses. The thesis becomes controversial when Noë says, "to perceive you must be in possession of sensorimotor bodily skill" (p. 11). Two things stand out about such claims. First, the term "sensorimotor" makes explicit reference to output capacities, and, second, the term "must" implies that the perceiving necessarily involves such capacities. I interpret Noë to mean the following: perception cannot take place (in us? in any naturally evolved creature?) without engaging the mental processes that underwrite our capacity to move. One way to make sense of this necessity thesis is to argue that there is a constitutive relation between perception and action. On this interpretation, the mental processes that enable us to act are literally components of the mental processes that allow us to perceive. I am fairly confident that this is the view that Noë wishes to defend. He often says that, in perceiving, we must use knowledge of sensorimotor contingencies, which implies, that perceptual mechanisms work by registering how stimulation will change when motor effectors are engaged.

This formulation is a bit more perspicuous than many in the book, but it is still somewhat problematic. The claim that sense organs register sensorimotor contingencies implies that sense organs deliver information in conditional form: you will have sensation *S* if you make movement *M*. But this can't be right. Such conditional rules make explicit reference to sensations, and, in so doing, they imply that there is a way to individuate sensations independent of our movements. In order to capture the idea that perception (defined broadly enough to include sensation) necessarily engages motor processes, we cannot define the contents of perception as containing conditional relations between sensations and movements. The relationship between sensations and movements must be more intimate. Perhaps we should say that, for Noë, a sensation is constituted by a range

of possible movements. On this formulation, sensation is reduced to action. That has some very striking implications. First, it implies that there cannot be sensation without engaging processes involved in action; arresting action mechanisms should lead to “experiential blindness.” Second, it implies that similarities and differences between perceptions will always correspond to similarities and differences between the motor responses associated with stimulation to our sense organs.

Such radical conjectures carry a heavy burden of proof, and one of the great virtues of Noë’s book is that he makes a concerted effort to support his claims with empirical evidence. My chief complaint is that the evidence is unconvincing. Let me briefly review what I take to be the most promising empirical arguments presented in favor of enactive perception in *Action and Perception*. Further arguments have been put forward in some of Noë’s papers, and in work by other authors, including Susan Hurley’s excellent book, *Consciousness in Action*. I will restrict myself to the arguments in Noë’s book, which is the most recent sustained defense of the position.

One line of evidence for enactive perception comes from research on prosthetic vision. Bach-y-Rita developed an apparatus that allows blind people to detect optical information by converting the inputs from video cameras affixed to a pair of glasses into an array of electrical impulses transmitted to the torso or tongue. The electrical array is spatially arranged like a video monitor, and when subjects train on this apparatus, they learn to use electrical patterns to respond to distal objects as if they were actually seeing. Subjects can identify shapes of the objects around them, avoid obstacles, and discern when one object overlaps another. Usually, the sense of touch can only convey information about objects that make contact with the body, but, with Bach-y-Rita’s apparatus, subjects are able to use touch in the way that sighted subject use vision. Noë believes that these results support the enactive theory of perception. To make that case, he must rely two premises. First, he must show that the experience of using the apparatus is like vision, and second he must show that it takes on this visual quality in virtue of the fact that subject learn to associate its inputs with the kinds of motor responses that are usually reserved for vision. If learning new motor responses converts touch into a sense that is qualitatively like vision, then motor responses may be constitutive of sensory qualities.

I have little doubt that subjects master use of the apparatus by adjusting their behavioral dispositions; when their torsos are stimulated, they react by avoiding obstacles in their environment, not by flexing or contracting their abdominal muscles. But I seriously doubt that these subjects experience anything visual. Noë’s evidence for this claim is that the subjects begin respond to distal objects as if they were seeing them. But that is not sufficient evidence. After all, there are conditions under which we can use touch to sense objects that are not in contact with our bodies. When we tap an object with a cane we feel its shape and texture; when we drive, we feel the surface of the road; when we move towards a fan or a flame, we can sense our approximate distance; and so on. In each of these cases, the experienced qualities are distinctively tactile, not visual. I have not tried Bach-y-Rita’s apparatus, but none of the testimony that I have read persuades me that users experience anything visual. There is no reason to think that one sense begins to cause perceptual states that are qualitatively like another sense simply by conveying the same information. To give one more example, we all know that the sound

of a smoke alarm signifies the presence of smoke, but that does not mean that we smell smoke simply by hearing the alarm go off. Could a congenitally anosmic subject, deprived of olfaction, suddenly experience the smell of smoke after being trained to interpret the significance of a smoke alarm? Of course, a person who was not congenitally blind might begin to form visual imagery while using Bach-y-Rita's apparatus, through memory associations. If a congenitally blind individual reported visual qualities while wearing the apparatus, we would have little reason to believe the testimony, for such a person would not be able to recognize visual experiences as such. My best guess is that prosthetic vision devices simply allow subjects to make automatic inferences about where objects are located in space as a result of tactile information. Their spatial awareness may include visual imagery for subjects who have visual associations in memory, but for the congenitally blind, the awareness probably isn't visual at all. Touch systems, like sound, smell, and even taste, can encode spatial information.

Let me turn to a second line of evidence. Noë is impressed by research on adaptation to prism lenses. When subjects wear lenses that invert the visual field from left to right, they initially find it difficult to successfully reach for objects, avoid obstacles, and otherwise interact with the object that they see. But, after some training, they adapt, and their actions become quite successful. Noë claims that this supports the enactive view, but that conclusion hinges on a particular interpretation of the results. The enactive view entails that visual qualities depend on motor dispositions, so it is crucial for Noë's view that subjects' visual experiences transform as they learn to behave successfully with the lenses on. He takes this to be the case. He even goes so far as to say that subjects undergo "experiential blindness" as they abandon their usual motor responses to visual stimuli and search for new responses. Is Noë's interpretation correct?

The claim that subjects wearing inverting lenses undergo a period of experiential blindness is clearly false. Subjects never report being blind, and they seem to use visual information to correct their motor responses (e.g., when subjects' toss an object in the wrong direction, they correct for the observed error the next time around). Subjects who wear the lenses experience normal (but inverted) vision when they stand still. When they try to move around, they sometimes experience visual distortions, but those are very easily explained. It is highly likely that, when viewing a dynamic scene, we spontaneously form expectations about what will happen by generating visual images of the next instant. For example, if you are watching a ball fly, you visualize its trajectory. Now suppose the ball changes course in an unexpected way. That experience may be both surprising and disorienting. If you think the ball is moving away from you, and it's actually getting closer, you may be startled. Moreover, the image you have formed to anticipate the trajectory of the ball may co-occur with the actual experience of the ball, resulting in a double image. Now imagine that everything you are seeing is like this; all of your visual expectations are off. That would lead to an experience of a distorted unstable world. Objects would shift unexpectedly, shadowy double images would occur, objects that appear close would slip away. This is just what it's like to wear inverted lenses.

What happens, then, when people adjust to the lenses? Noë would have us believe that their visual experiences transform as a result of their new motor dispositions. This is

that the enactive view requires, but notice that even if Noë is right about what happens, it wouldn't entail that the enactive view is correct. If motor dispositions alter vision, that does not prove that visual qualities are partially constituted by motor dispositions; they alternations might be causal. We know, for example, that vision can affect what we hear. In the McGurk effect, seeing the way a person's lips move influences the experience of the sound that we hear them producing. We don't conclude from this effect that auditory states are intrinsically visual, much less that hearing requires mastery of visual skills. Rather, we say that vision can influence responses in the auditory system. Likewise, if re-learning of motor skills alters vision when wearing inverting lenses, that wouldn't prove that vision is intrinsically motoric.

In any case, I seriously doubt that visual experiences are significantly altered after practice with inverting lenses. We might stop making erroneous visual predictions of the kind that I have been discussing, but I doubt that the inverted image generated by the lenses actually reverses back to the correct orientation. A more likely possibility is that vision remains inverted, and we adjust our motor responses accordingly. Consider, by analogy, the experience of combing your hair in a mirror. Mirror images appear in front of us, yet, to comb our hair we move our arms above our heads. This is a highly practiced motor response, but it does not transform the visual experience; the mirror image continues to look like it's in front of us! Likewise, if you practice drawing mazes using a mirror reflecting, you will become quite adept at it, but at no point does the image reverse. Or, to give one more example, Leonardo da Vinci wrote all his notes backwards, but scholars are accustomed to reading them. After practice, however, da Vinci's notes don't magically reverse and appear to be written in correct direction. Adaptation to inverting lenses seem to involve a process of motor training in which it becomes effortless and automatic to reach to the left when an object appears visually on the right. This, I suspect, is just like dragging a comb on one's head when we notice an offending stray hair in the reflection of a mirror.

The claim that lens adaptation leaves vision intact has been defended in detail by Harris (1965; 1980). Noë is aware of Harris's proposal, and he offers five remarks in reply. First, he says "I find it highly implausible that we have any grip whatsoever on spatial content apart from a perceiver's sensorimotor skills" (p. 94). This just begs the question, and, what's more, the visual modality offers a clear example of how space can be represented without a motor component. Visual images, like pictures, may have parts arrayed along spatial axes (compare, Kosslyn, 1994). Second, Noë suggests that Harris's position may actually be interpreted as an argument for the enactive view. Harris says that adaptation works by altering motor dispositions, rather than by altering the visual representation, but, if the enactive view is right, motor dispositions are constitutive of visual qualities, so Harris may be inadvertently supporting the enactive view by conceding motor adaptation. This reply is also question begging. Harris is proposing that we can explain adaptation without assuming that motor dispositions have any impact on vision. One cannot deny this possibility without argument. Third, Noë complains that Harris postulates an indefensible asymmetry between vision and motor responses. He thinks that Harris has no grounds for supposing that motor responses adapt to vision, rather than the reverse. But the evidence is on Harris's side. Vision is usually the dominant sense. We know, for example, that vision dominates hearing in the McGurk effect (McGurk and MacDonald, 1976). More to the point, studies of the brain structures

involved in prism adaptation suggest that they crucial changes take place in the cerebellum and parietal cortex (two regions associated with the body) and not in visual cortex (Martin *et al.*, 1996; Clower *et al.*, 1996). Fourth, Noë accuses Harris of making a homunculus fallacy. Harris assumes that the visual representation has an intrinsic left and right side, just like a picture, but there is no reason to think that the left and right side of a retinotopic brain state correspond to the left and right side of a visual percept. In response, Harris could point out that his arguments do not hinge on any facts about cortical retinotopy; there are phenomenal and behavioral reasons for thinking the visual percept does not change during lens adaptation. Moreover, if there were direct evidence that retinotopic representations remained unchanged during lens adaptation this would provide *prima facie* support for the claim that the visual percept doesn't change. After all, under ordinary circumstances, the spatial configuration of retinotopic configurations corresponds in systematic and predictable ways to the experienced percept (for example, retinotopically proximate cells correspond to components of a visual percept that appear proximate to the viewer). The brain does not contain an inner homunculus, but it seems to use adjacent regions of visual cortex to represent adjacent portions of the visual field. Finally, Noë complains that Harris is committed to the view that subjects who adapt to inverted lenses are deluded about the character of their own experience. Noë thinks that lens wearers will report that their visual percepts invert after they become accustomed to the lenses. But the reports of people who have trained with inverting lenses seem to be unclear on just this point. People get to a point where objects on the left once again seem to be on the left, but this seeming may not be visual. After habituation, we know to reach to left to pick up the beer on our left, but the beer may visually appear on the right. Mundane experiences with mirrors bear this out. When I see my reflection, the misplaced hair *seems* to be on top of my head, in a motor sense of seeming, but it visually appears to be a couple of feet in front of my head, in the reflection. I think Harris's view gets the phenomenology right.

As far as I can tell, Noë's empirical case for the enactive view rests quite heavily on his interpretations of lens adaptation and prosthetic vision. These interpretations are simply too contentious to support the theory. If there were no other evidence in favor of the enactive view, we should reject it out of hand. But I can identify four other arguments in his Noë's book.

One of these arguments involves image stabilization. Noë points out that a stimulus becomes invisible when it is stabilized on the retina, and that suggests that perception requires movement. This argument is far from decisive. We can see a stable stimulus initially. It disappears only after it remains stable for a period of time. Presumably this is because this visual system needs a way of distinguishing specs of dust, scratches, and blood vessels on the eye from objects in the environment. If vision is to deliver information about the outside world, it must ignore stimuli that move with the eye. In any case, blindness to stabilized images would show only that motion induced transformations can causally influence visual processing, not that such transformations are constitutive of visual states. The visual system also shuts down when we blink, and that certainly doesn't prove that visual states are partially constituted by motor states. Compare: the bulb in a motion sensitive light would flash when and only when there is movement, but its luminosity is not constituted by the movements on which it depends.

Next, Noë cites examples of amodal completion in defense of his enactive theory of perception. When we see a partially occluded object, we can tell that the object continues beneath the occluder. We can also tell that a tomato facing us has a back. Noë claims that the absent elements are (a) part of the phenomenology and (b) that they are phenomenally available to us as a result of implicit motor knowledge. This suggests that motor dispositions make a constitutive contribution to experience. I don't think these examples support the enactive view. First of all, it's far from obvious that amodal completion is accomplished by the motor system. It could be explained entirely within the visual system, as a function of visual expectations. Second of all, it's far from obvious that the occluded elements are part of the phenomenology. We would be surprised if, when an occluder was lifted, the object behind it was incomplete. But the surprise only shows that we had an unconscious expectation, not that that we were, paradoxically, experiencing the hidden part. Third, even if Noë's motoric account of how we complete occluded figures is right, that does not prove that vision is partially constituted by motor responses. To the contrary, it is possible that vision leads us to generate motoric expectations that are completely separate from the visual representations that cause them. When I see ice cream, I imagine tasting it; that does not show that vision is intrinsically olfactory. Likewise, when I see a tomato, I imagine grasping it from all sides, but that does not mean that vision is intrinsically motoric.

Noë's next argument has to do with a related phenomenon. When we see an object, we see it from a particular vantage point, but we nevertheless recognize its intrinsic form. If we see a penny from an angle, we experience an ellipse, but we know, at the same time, that the penny is round. Noë explains this by supposing that, when we see an object we immediately impart our knowledge of how that object would look if we were to move around it. Seeing a penny as round intrinsically involves motor knowledge on this view. I find this unconvincing. Noë often conflates movement of an object with movement of an agent. Suppose he is right that discerning the intrinsic form of an object is a matter of knowing what it would look like from multiple angles. That does not show that such knowledge is motoric. After all, we can see an object from multiple angles if we remain stationary while the object moves. We may learn to understand objects by seeing them in motion. In principle, this can be done entirely within the visual system with no involvement of systems involved in motor control. We have good models to explain perception of form from motion (as in, e.g., point light displays). We also have good models of how the visual system recognizes the intrinsic form of objects without any motion, either of the object or of the agent. Biederman's (1987) geon theory, for example, explains recognition of intrinsic volumetric geometry via the identification of vertices that can be discerned in a two dimensional image. We also know, from cell recording in the visual system, that vantage-point invariant representations are extracted from vantage point specific representations during visual object recognition. Motor systems play no essential role in this process.

Another argument for the enactive view appears at the end of the book, and draws on work that Noë has done in collaboration with Susan Hurley. In a rich and interesting paper, Hurley and Noë (2003) draw our attention to the fact that there are two kinds of neural plasticity. In some cases, when one region of sensory cortex invades another, the resulting phenomenological qualities caused by activation in the invaded region correspond to the qualities associated with the invading region; they call this deference.

For example, when congenitally blind people read Braille, their tactile systems cause activation in areas of the brain normally associated with vision, but they experience the activations as touch. In other cases, the invaded region retains its original qualities and refuses to defer; Hurley and Noë call this dominance. For example, when cells that respond to tactile stimulation of the face invade brain regions associated with a limb that has been amputated, activation in the invaded region feels like sensation in the limb, not like sensations in the face (so-called phantom limbs). Hurley and Noë interpret these findings as evidence for the enactive view. They say that the only factor that determines whether the invaded region will defer or dominate is the motor dispositions that come to be associated with activations in that region; sensorimotor contingencies dictate phenomenology. I find this implausible. First, it may get the cart before the horse. Perhaps people respond to phantom-limb pains with limb directed movements because the sensations feel like they are arise from the limbs; Hurley and Noë would have us believe that the pains feel like they are arising from the limbs, because of how people are disposed to move. Second, if we train amputees to stroke their faces rather than their missing limbs when pain arises, the new motor behavior does not eliminate the experience of pain as coming from the missing limbs. Third, the factor that seems to determine whether a cell population will dominate or defer seems to be driven by the source of the inputs, not the outputs. In particular, dominance seems to occur when cells are invaded by other regions within the same sense modality, and deference seems to occur when the invaders come from other sense modalities. Hurley and Noë consider and reject this possibility, because they think there are counterexamples. In particular, they suggest that intermodal synesthesia is a case of dominance, rather than deference, contrary to the proposal I am considering. When the olfactory system causes synthetic tactile sensations, for example (as in the man who tasted shapes), the invading modality (olfaction) does not causes the invaded modality (touch) to relinquish its usually phenomenology. The problem with this counter-example is that synesthesia is not a case of neural plasticity; it is not a case of one sense invading another that is not being used. Rather it is a case of low-level causal links between two sensory systems that are intact. Nothing gets rewired in synesthesia; rather it is believed that cells linking the senses that are usually pruned in early development remain intact. The proposal that intermodal cases of cellular *invasion* cause deference remains plausible. No one knows why the brain defers in intermodal cases and dominates in intramodal cases, but one possible explanation is intermodal invasion occurs when an entire perceptual system (or large parts of it) are taken over by another perceptual system. With a global take-over, there may be sufficient rewiring to change the cellular dynamics in such a way as to change the qualitative character of resultant stimulation of those cells.

The final argument I will consider in favor of enactive perception is developmental. In 1963, Held and Hein performed a study on kittens which is interpreted by many to show that vision will not develop normally in the absence of action. Held and Hein raised two kittens in darkness, but allowed them to enter an illuminated room for a couple of hours a day. One kitten was allowed to road around in the room, and the other was suspended above the ground in a cradle that prevented the kitten from moving. The kitten in the cradle was harnessed to the mobile kitten, so it moved through the room, but it was unable control the movement itself. After being reared in this way, both kittens were freed. The mobile kitten had normal visual abilities, but the immobile kitten was

impaired. It was unable to use vision to direct its paw reaching behaviors and it was indifferent to visual cliffs. Noë suggests that this experiment supports the view that visual perception necessarily involves understanding of motor responses. I am unconvinced.

First, of all, the immobile kitten was able to see when it was released. It walked around and responded to objects, albeit clumsily. Second, there is no reason to think that the immobile kittens' visual experiences were in any way abnormal. The fact that it couldn't direct its paws suggests that it could not successfully map visual stimuli onto motor commands, which is unsurprising, since it had no practice doing so. The fact that the kitten was indifferent to visual cliffs shows that the kitten did not have enough experience walking on edges to anticipate the bodily affordances of the visual world. This is also unsurprising. For all we know from the data, the two kittens had exactly identical visual experiences after they were freed, but only one kitten knew what all the visual experiences meant with respect to the body. It is also worth noting that the Held and Hein study was never replicated, the kittens recovered quickly, and studies of human infants with muscle atrophy show that when humans are prevented from moving in early development, there is no decrement in the visual comprehension of space (Rivière and Lécuyer, 2002).

I don't think any of the arguments in *Action In Perception* give us reason to think that perception is intrinsically motoric. In fact, I think there are many good reasons to reject this suggestion. The study of infants with muscle atrophy is one of many studies showing that damage to movement systems does not impair perception. Perception is not impaired by spinal cord injuries that cause paralysis, by paralysis of the eye muscles or brain structures that control them, by atrophy of motor cortex in Lou Gehrig's disease, by destruction of action-control centers in parietal cortex, or by destruction of brain areas that include mirror neurons in frontal cortex (which are presumably destroyed in many cases of Broca's aphasia) (see descriptions in Kandell *et al.*, 2000). Noë is not explicit about what brain structures are involved in the motoric component of perception (perhaps because he doesn't believe in neural localization), but it is certainly noteworthy that no motor deficits seem to undermine our ability to perceive. There are clear dissociations between perception and action. People with motor deficits can see the world, and people with perceptual deficits can act in it.

This is not to deny that there are some brain systems that serve the function of relating perceptual information to motor consequences. There may be bimodal cells in the dorsal visual stream, for example, that are used for the visual coordination of action. But this does not entail that all perception is essentially motoric. Indeed, the visual systems that are involved in guiding action are believed to work outside of visual awareness (Milner and Goodale, 1995). Noë believes that these findings have no bearing on the enactive view, but they seem to flatly contradict a central tenet of his approach: the assumption that visual awareness depends on motor dispositions. If the brain areas that are known (because of their behavioral consequences) to encode the motor consequences of visual stimuli are not implicated in visual consciousness, then there is no reason to think Noë's theory of consciousness is correct. Noë is forced to say that representations in the ventral visual stream are also involved in the coordination of action, but there is absolutely no evidence for this conjecture. All evidence implicates the dorsal stream. I

think the evidence from neuroscience provides an overwhelming case for the view that perception is not essentially linked to action.

There are also theoretical reasons to resist the enactive view. First of all, Noë is committed to the claim that all sensory qualities are distinguished by the associated motor dispositions. That strikes me as very far-fetched. Noë spends a chapter trying to convince us that colors can be distinguished by motor affordances, but this strikes me as a slow walk up a tall slippery hill. It's far from clear that there are different motor skills deployed when we look at a solid green color field and an equiluminous red color field. Noë has the burden of proof; he must show that there are motoric differences. Doing so is necessary, but not sufficient. After all, there is little reason to think that colors must be distinguished by their motor affordances when there are already differences in how colors are represented in the visual system. Why do we need motor information to distinguish colors when different colors elicit different cellular responses from the retina on into the brain? In addition, Noë himself points out that two color patches may appear identical, when viewed from one position at one time, and then different, when we change viewing conditions. If experience of color were determined by motor affordances, such patches should appear different. Noë sometimes tries to support the enactive view by saying it is necessary to explain why colors seem constant across viewing conditions. But color constancy, like shape constancy, is believed to be resolved within the visual system using purely perceptual machinery. Good biologically-grounded models exist to explain these effects, and there is no reason for motor responses to get involved.

The enactive view becomes even less plausible when we move beyond vision. Consider two perfumes: they may smell different even if they do not have different consequences for action (especially if they are equally appealing). Do we sniff different smells differently? Two high-pitched tones sound different even if their effect on action is the same—even, for example, if both fall far out of our range of vocal imitation. Two bowls of tepid water feel different, even if neither makes me more inclined to withdraw my hand. And so on (see Jacob, this issue, for more on this theme). Perhaps these sensations afford different responses, but Noë must argue for that, and he must show that these differences are *constitutive* of the differences in perceptual qualities, rather than consequences of the differences in perceptual qualities.

In sum, I don't see a plausible case for the enactive view. The evidence offered in its defense falls far short of the mark, and there is considerable reason to be skeptical. Indeed, the dissociation between perception and action may have been the most important evolution step in the development of intelligence. Simple organisms have responses automatically linked to the stimuli that impinge on them. There is no choice, no flexibility. By severing the link between perception and action, an organism can bring multiple contextual factors to bear before responding. I don't take this evolutionary story as a serious argument against the enactive view; Noë ends his book with an equally speculative evolutionary story of his own. But it does show, from a design perspective, why it might be advantageous to keep perception and action apart.

3. Do We See in Pictures?

The arguments thus far are intended to raise doubts about what I take to be the central thesis of Noë's book: the claim that perception necessarily involves action. I think the

border between perception and action can fully withstand Noë's attack. Perception is inactive, not enactive. I could conclude here, but *Action in Perception* is filled with rich and interesting treatments of other philosophical topics. I cannot comment on these exhaustively or adequately in this review, but, in these final sections, I want to register misgivings about a few other contentions in the book.

Let me begin with Noë's critique of the hypothesis that we see by constructing mental pictures. There is certainly a sense in which Noë's position is right. We do not have color filled pixels in the head, much less an inner observer to discern all the parts. But this view is not seriously defended by anyone. Even Kosslyn (1994), who is famous for defending a "pictorial" view of visual imagery, is careful to avoid a regress of homunculi and tubes of mental paint. The claim that vision produces pictures is best understood as the claim that the parts of a visual representation function as if they were spatially aligned; the brain accesses representational components that correspond to adjacent parts of a visually perceived scene as if they were adjacent. The pictorial view is also associated with the conjecture, though it need not be, that perceptual representations are rich. Like a photograph of scene, defenders of mental pictures often assume that vision picks up detailed information from every corner of the visual field.

Noë rejects this pictorial view, but his arguments are not always compelling. For example, he says the pictorial view is threatened by the fact that we saccade several times a second, and that the visual apparatus is not uniformly colored or focused. But these discoveries do not undermine the view that perceptual representations are spatial or rich. Saccades may help us discern many details, which would be lost if our eyes remained in one place, and variations in focus or color can either be captured by pictorial elements that fade and blur or can be eliminated by visual mechanisms that sharpen and colorize. No facts about the cell distribution in the retina can rule either of these possibilities out.

Noë also tries to undermine the pictorial view by pointing out that each eye has a blindspot, where the optic nerve makes contact with the retina. He says that slaves to the pictorial view of seeing are led to postulate unnecessary mechanisms for filling in these gaps. Noë is right to think that filling in is unnecessary. As Dennett (1991) has remarked, we mustn't confuse the absence of representation with a representation of absence. If the visual system is not picking up any information in the blind spot, that doesn't mean it is picking up a gap. Therefore, we shouldn't expect to see two holes on either side of our visual fields. But the defender of pictorial vision does not necessarily predict that blindspots will cause visual holes. Vision can be rich and spatial without being committal about every region of space in front of the viewer. Only a defender of inner homunculi should expect visual gaps to be discernable from within. The visual percept exhausts the visible world; there is no viewer to see its contours.

Thus, if there are no mechanisms for filling in the blind spot, the defender of pictorial vision should not blush. As it happens, however, there *are* mechanisms for filling in the blindspot. In a series of ingenious experiments, Ramachandran and Gregory (1991) have shown that the early visual system actually uses information around the blindspot to fill in the region that receives no inputs from the outside world. In one experiment, subjects stare at a uniform "snow" screen on a television monitor, and then the monitor is switched to a broadcast a solid color field. For a brief moment, subjects report seeing snowy patches in the color field in the positions of the blind spots. These

findings are not essential to support the pictorial theory of vision, but are embarrassing for opponents of pictorialism. If you think that visual representations are sparsely detailed, rather than rich, you should not expect the visual system to generate representations corresponding to small, unimportant regions outside the center of focus.

Noë's most important objection to the pictorial theory of visual perception is designed to undermine the hypothesis that visual representations are richly detailed. He invokes the phenomenon of change blindness. When subjects see one scene followed by a slightly different scene, they often fail to notice the change. Noë interprets this finding as demonstrating that the visual system extracts only a sparse representation that leaves many details out. Another possibility, however, is that the visual system generates very rich representations, but doesn't compare every detail. On the view, the visual system registers changes in the scene, but not the fact that a change has taken place. The differences are all encoded, but no comparison system picks up on them. This interpretation is consistent with the pictorial theory of vision, and Noë must argue against it to show that his interpretation is right.

Curiously, Noë does not argue against the alternate interpretation, and, in a concessive passage, he acknowledges evidence that strongly supports that interpretation and undermines his suggestion that vision leaves out many details. In particular, he cites a study by Simons and Levin in which subjects fail to notice that a basketball is replaced by volleyball. When forced to guess afterwards, however, subjects are accurate about what kind of ball they initially saw. This suggests that the neglected information was internally represented. Noë concedes the point. Subsequent studies have added further support. In one study by Mitroff *et al.* (2004), subjects are shown two arrays of object sequentially, and they fail to notice that some objects in the initial array have disappeared; but, when given a forced choice test afterwards, they are better than chance at guessing which objects were in the initial scene. In another study, Silverman and Mack (2001) have shown that, when subjects fail to notice which letters change in briefly displayed arrays, the unnoticed letters prime subsequent performance on a letter completion task. Once again, this supports the view that visual representations are rich, rather than sparse. We just don't keep track of all the changes in the details. This supports a pictorial view of vision and it undermines the uses that Noë makes of the change blindness results throughout his book (see pp. 62, 191, and 218).

4. The Contents of Experience

In one of the more philosophically nuanced sections of the book, Noë defends the claim that perceptual content is conceptual, rather than non-conceptual. He reviews some of the arguments for non-conceptual content and offers plausible replies. For example, he rejects the suggestion that non-human animals perceive the world using non-conceptual resources, because he thinks we have no reason to deny that such creatures have concepts. I found much of this discussion compelling, but at a crucial juncture, I would want to abandon Noë's ship and embrace non-conceptual content.

Noë rejects Evans's (1982) suggestion that perceptual contents must be non-conceptual because perceptual representations are richer than conceptual ones. I think this is the most compelling argument for the conclusion that perception is not always conceptual. It is a well-known fact that our powers of perceptual discrimination outstrip

our powers of perceptual recognition. There are nearby shades of color that we can discriminate even though, on a subsequent recognition test, we would have no idea which of the two we had seen on an earlier occasion.

Noë has two reasons for rejecting this argument. First, he endorses John McDowell's (1994) suggestion that we may be able to form demonstrative concepts corresponding to every discriminable shade of color. According to McDowell, when we keep a color in view, we can internally point to it, and think, a color like *that*. McDowell thinks that such internal pointers are concepts. I find this move unconvincing for two reasons. First, it seems to me that, before we can point to a color experience and refer to it, there has to be an experience there to point at. The pointing seems to be a separate mental act that occurs after the color is already being perceived. If so, perception of colors must initially be unmediated by concepts. Second, I don't think we should grant such internal pointers the status of concepts. The term 'concept' is a technical term, of course, and theoreticians can define it as they will. I stipulate that it is a necessary condition on concepts that they can be tokened endogenously, under the control of an organism (Prinz, 2002). Representations that depend on the presence of exogenous inputs don't qualify. There is, at least, an important difference between the concepts that we can freely and spontaneously use in thought, and the perceptual representations that we hold in place. That's a distinction worth marking with the word "concept." McDowell and Noë might reply by saying that we can make spontaneous use of a demonstrative; we can choose to mentally point to a color or not. True enough, but what is spontaneous in this case is the empty demonstrative pointer, not the color to which we point. Our capacity to point may be conceptual, but the compound state comprising an inner pointer and a color experience is, in part, non-conceptual; the color cannot be re-generated at will.

Noë's second reason for concluding that perception is conceptual is that he thinks the motoric knowledge that we bring to bear when perceiving qualifies as conceptual knowledge. I reject this suggestion. First, as I have argued, I don't think perception necessarily involves such motoric responses, and so Noë's argument could not be used to show that perception is always conceptual. Second, I reject Noë's proposal that we regard the motor knowledge that he has in mind as conceptual knowledge. The abilities we apply in scanning the surface of an object with our eyes, for example, are largely sub-personal. We could not exercise these abilities by act of will. I don't think they satisfy the condition of spontaneity or endogenous control. Nor do they satisfy other conditions stipulated by other authors to be necessary for concepts. It's far from obvious that we can recombine our motor skills in a way that would satisfy Evans's Generality Constraint, and it's also far from obvious that motor skills are governed by norms—a condition on concepts that has been defended by McDowell. Could someone hold me accountable for not staring at a red patch in the right sort of way? Do I feel compelled to give reasons for my saccades? Is there no such thing as a private disposition to grasp?

Like any theoretician, Noë is entitled to define concepts as he likes, and I see no reason for denying that some concepts are realized by motor skills. But the kinds of motor dispositions that he postulates in his theory of perception fail to satisfy standard criteria on concepthood in the literature, and therefore, his insistence that perceptual is conceptual cannot be regarded as incompatible with the views of those who have defended non-conceptual content. This strikes me as a verbal dispute won by stipulation.

To have a substantive debate about non-conceptual content, we should agree to define the key terms in the same way.

5. Where is the Mind?

Let me turn, finally, to Noë's provocative concluding chapter. In my introduction, I mentioned that there is a traditional border in cognitive science that could be disputed only by someone inclined towards brazen radicalism. That is the border between mind and world. Most cognitive scientists think that mental states, especially conscious mental states, supervene on brain states. Conscious states may have their content in virtue of being causally related to the world in a certain way, but the qualitative character of a perceptual experience depends exclusively on the brain: if one changed the environment and kept the brain constant, the character of the experience would remain unchanged. Noë wants to question this assumption. He views it as a dogma in cognitive science, and seriously entertains the idea that some conscious experiences supervene on interactions between mind and world. The position is interesting, of course, and the chapter is entertaining, but, ultimately, the arguments fall short, and I think that Noë's position unwittingly plays into the hands of dualists, mystics, and those who are disposed to be skeptical about science. This is not Noë's intention. He is a card-carrying materialist, and he is right to think that neuroscientists should pay more attention to how the environment interacts with the brain. But the metaphysical conclusion—wide supervenience—strikes me as gratuitously radical and almost certainly false.

Noë offers both positive and negative arguments for the thesis that some experiences supervene on mind-world relations. The positive arguments are used to marshal evidence in favor of his view, and the negative arguments are used to deflect standard reasons that are given for supposing that experiences supervene exclusively on the brain.

On the positive side, Noë invokes the experience of presence in absence, as in cases of amodal completion. He suggests that these experiences include bits of the external world—in particular, I assume, the parts of an occluded object that we don't see are present to us in virtue of being in the surrounding environment. Should this lead us to say that conscious experiences supervene on the world? Absolutely not. First, in some cases, that occluded bits of objects are not parts of phenomenal experience at all. Remember, absence of representation is not representation of absence. In these cases, there is no correlate for the phenomenal experience of an occluded part, because there is no such experience. Second, in other cases, we become phenomenally aware of occluded bits by either visualizing them, or, to follow Noë's own account, by registering possible motor interactions with them. Noë could only defend wide supervenience if he had a further argument to show that these aspects of experience (visual and motoric) do not supervene on the brain. He has no such argument. The mere fact that phenomenology includes features hidden from view certainly doesn't show that phenomenology includes the world. Third, if the world really were a constitutive part of our phenomenal experiences, then the bits of objects that are hidden from view should not be phenomenally "present as absent;" they shouldn't be absent at all. Phenomenal experience should be as richly detailed as those aspects of the environment to which we have access. Elsewhere, Noë is at pains to deny that phenomenology is richly detailed.

Without further argument, his wide supervenience thesis is at odds with his critique of phenomenal richness. To escape this problem, Noë needs to say that the bits of the world that contribute to phenomenology are absent yet present. I honestly don't know what this means.

Noë also tries to provide positive evidence for his wide supervenience thesis by arguing that we would not have phenomenal states with the rich character of those found in perception if our sense organs were not making contact with the world. If this were true, it would show only that the world is a causal precondition for having some phenomenal experiences. It would show only that the brain is incapable of entering certain configurations without external stimulation. For all I know, that is true, but it wouldn't support wide supervenience. To support wide supervenience, Noë should show that, when we keep the brain fixed and change the environment, there can be changes in experience. He attempts no argument of this kind.

Why think that experience supervenes on brain states? Isn't this just a dogma? No; it's the general consensus after centuries of argument, reflection, and science. People who believe in an immortal soul would love it to be false, and scientists have had to work hard over many, many years to support the conclusion that experiences are in the head. Noë considers and rejects two arguments for the brain view. First, he considers the argument that experiences must be in the head because they can be triggered by stimulating the brain, in the absence of any other environmental change. His reason for rejecting this view is that such experiments have only ever generated "very simple sensations such as the illusion of the presence of flashes of light" (p. 211). This is completely false. Penfield's classic studies showed that stimulation could generate vivid experiences of songs, conversations, views from childhood windows, and detailed, polysensory episodes from the biographical past. Subjects report that these episodes are much more detailed than ordinary images and they are more like reliving the events. Noë also complains the brain stimulation does not show that brain events are *sufficient* for experience. To make this point he draws an analogy: a car's running engine is necessary for driving, but it would be sufficient only under certain environmental conditions, as when the car is right side up. But I fail to see the point of the analogy. Brain stimulation establishes sufficiency; it shows that rich perceptual experiences can arise in the absence of related environmental factors. To show that brain stimulation is not sufficient, for experience, Noë would have to show that some environmental variable is making a substantive contribution in these cases. What variable could that be? Noë suggests that it might be the scientist and his electrode, but it's hard to take that suggestion seriously. After all, the same scientist, Penfield, can cause different experiences by stimulating different brain areas. Does Penfield transform into the supervenience base of one experience on one occasion and another experience on the next? Would it matter if Penfield were replaced by another scientist or a machine? Of course not! The only variable that varies directly to the character of the experience is the neurons that are active in the brain.

Noë also considers an argument from dreaming. We have rich experience when we are asleep, and our eyes are closed. It is also worth noting that people in sensory deprivation chambers experience richly detailed hallucinations. This seems to show that experiences supervene on internal events, not external events. Noë replies by conceding

that some experiences supervene on the brain, while insisting that others depend on interaction with the environment. To make this case, he contends that dreams have very different experiential content than waking life. But this response is not successful. The differences between dream experiences and waking experiences are not generally differences in content, but differences in narrative order. The experiences of color and shapes in dreams are often just like those in waking life. Some people have richly detailed dreams, and eidetic dreamers claim that, for them, dreaming is just like seeing. Of course, dreams have fanciful and incoherent plots, but that shows that rational thought depends on being awake, not that qualitative character depends on external inputs.

Furthermore, I don't see the allure of Noë's hybrid position. If he concedes that some experiences supervene on brain states, why not admit that all do? What is supposed to be the principled distinction between the aspects of experience that supervene on brain states and those that don't? It can't possibly be the case that experiences of physical interaction require actual interaction with the world, because we are often active in our dreams. Nor could it be that experiences of occluded objects supervene on real objects in the world, because occlusion occurs in dreams. Noë must tell us the difference. He should convert his wild speculations about wide supervenience into testable hypotheses.

Let me add, finally, that the supervenience thesis enjoys further support for extensive research in neuroimaging. Every aspect of experience, from illusory contours to motion illusions, from phantom limbs to diffuse pains, can be correlated with some neuronal response. There are cells that respond to every feature we perceive, and, in many studies the time course of the neuronal activation correlated with the time course of the perceptual experience. The search for neural correlates of experience has been a hugely productive research program. We have learned a tremendous amount about how perceptual systems process information by studying the brain, and, crucially, nothing seems to be left out. Neuroscience shows every promise of being deliciously complete. Each nuance of experience has a correlate, and each correlate can be found when we look for it.

Many philosophers would balk at this claim. Neuroscience is incomplete in one crucial respect: it can't explain why all these neural correlates feel the way they do. It can't even explain why they feel like anything at all. Perhaps this is the missing element that makes Noë feel like he has to go outside the head. If so, it's a fool's errand. For this deep epistemic problem of how physical states can be phenomenal experiences is in no way ameliorated by broadening the supervenience base. Just as it's hard to understand why brain states feel a certain way, it's hard to understand why brain states together with bits of the external environment would feel a certain way. As Noë realizes, wide supervenience will not help solve the hard problem of consciousness. So why go outside the head? To take Noë's invitation seriously, we would have to be shown an example of an aspect of experience that has no inner correlate. No such example has been given.

6. Conclusions

Action In Perception is an enjoyable, well-written, and highly provocative contribution to the philosophy of mind. It is a worthwhile read, and it may provoke fruitful debates. But the arguments in the book are not powerful enough to threaten the borders they seek to demolish. The border between action and perception is unscathed by Noë's attacks, and

so is the border between mind and world. I think we have a lot to learn from the enactive approach. We should bear in mind that we perceive the world in order to respond successfully to it. Perception is, in this respect, for action. Cognitive science should pay attention to the environment. Living organisms have bodies and are situated in the world. I have little doubt that, as cognitive scientists pay more attention to these humble facts, their theories will be improved, and perhaps a few cherished assumptions will be abandoned. But it would be an unfortunate step backwards if we embraced Noë's claim that perceptual systems are constituted by mechanisms that control action, and it would profoundly retard progress if we stopped searching for the neural correlates of experience. We should give action its due, but let's not get carried away.

References

- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Clower, D. M., Hoffman, J.M, Votaw, J.R., Faber, T.L., Woods, R.P., and Alexander, G.E. (1996). Role of posterior parietal cortex in recalibration of visually guided reaching. *Nature*, 383, 618–621.
- Dennett, D. (1991). *Consciousness explained*. Boston, MA: Little, Brown & Co.
- Evans, G. (1982). *The varieties of reference*. Oxford: Clarendon Press.
- Harris, C. S. (1965). Perceptual adaptation to inverted, reversed, and displaced vision. *Psychological Review*, 72, 419-444.
- Harris, C. S. (1980). Perceptual plasticity in the human adult. In C. S. Harris (Ed.), *Visual coding and adaptability* (pp. 95-149). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Held, R. and Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 56, 872-876.
- Hurley, S. and Noë, A. (2003). Neural plasticity and consciousness. *Biology and Philosophy*, 18, 131-168.
- Kandell, E.R., Schwartz, J.H., and Jessel, T.M. (2000). *Principles of neural science*, 4th ed. New York: McGraw Hill.
- Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge: MIT Press.
- Martin, T.A., Keating, J.G., Goodkin, H.P., Bastian, A.J., and Thach, W.T. (1996). Throwing while looking through prisms. I. Focal olivocerebellar lesions impair adaptation. *Brain*, 119, 1183-1198.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
- McGurk, H. and MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264, 746-748.

- Milner, D. and Goodale, M. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- Mitroff, S. R., Simons, D. J., and Levin, D. T. (2004). Nothing compares two views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics*, 66, 1268-1281.
- Prinz, J.J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.
- Prinz, J.J. (forthcoming). Is consciousness embodied? In P. Robbins and M. Aydede (Eds.) *Cambridge Handbook of Situated Cognition*. Cambridge: Cambridge University Press.
- Ramachandran, V.S. and Gregory R. (1991). Perceptual filling in of artificially-induced scotomas in human vision. *Nature*, 350, 699-702.
- Rivière, J., & Lécuyer, R. (2002). Spatial cognition in young children with spinal muscular atrophy. *Developmental Neuropsychology*, 21, 273–283.
- Silverman, M., and Mack, A. (2001). Priming from change blindness [Abstract]. *Journal of Vision*, 1, 13a.