PSYCHE

Why Visual Experience is Likely to Resist Being Enacted

Pierre Jacob
Institut Jean Nicod,
CNRS/EHESS/ENS
UMR 8129
1bis, avenue de Lowendal
F-75007 Paris
France

jacob@ehess.fr

PSYCHE 12 (1), March 2006

KEYWORDS: enaction, knowledge of sensorimotor contingencies, change blindness, two-visual systems model of human vision, internal forward models of action

ABSTRACT: Alva Noë's version of the enactive conception in *Action in Perception* is an important contribution to the study of visual perception. First, I argue, however, that it is unclear (at best) whether, as the enactivists claim, work on change blindness supports the denial of the existence of detailed visual representations. Second, I elaborate on what Noë calls the 'puzzle of perceptual presence'. Thirdly, I question the enactivist account of perceptual constancy. Finally, I draw attention to the tensions between enactivism and two trends in cognitive neuroscience: the two-visual systems model of human vision and the theory of internal forward models of action.

Alva Noë's book, *Action in Perception*, is a sophisticated defense of the enactive conception of perceptual experience. As I understand it, the word "enaction" was introduced into cognitive science by the late Francisco Varela. According to the English Oxford Dictionary, the English verb "to enact" means both "to ordain" (or "to decree") and "to interpret a play". According to some of the most general versions of the enactive conception, perception is not something that happens to one, it is something one does. Perception is intrinsically active: it is the skillful bodily exploratory activity of one's environment. What one perceives—the content of perceptual experience—is determined by what one does (by what one is able to do, by what one would or is ready to do, or by what one knows how to do). One's perceptual experience arises from one's sensorimotor

knowledge, i.e., from one's practical or implicit knowledge of the sensory consequences of one's own movements.

It is not entirely clear which, if any, of the several meanings of the verb "to enact", the advocate of the enactive theory of perception wants to preserve. I, of course, agree that "perception" is the name of both a process (or an activity) and its product, i.e., the resulting experience. But I don't think that perception is something one does if this means that there is no relevant difference between perceiving a fact and either intending to act or acting. Intending is representing a goal or a non-actual state of affairs. Acting is turning a possible into an actual state of affairs. Unlike either intending or acting, perceiving is recording a fact (or a constituent of a fact). Furthermore, I am skeptical of the view that a proper account of perceptual processes—including bodily skills—can solve (or dissolve) all the phenomenological puzzles of perceptual experience. In particular, after reading Noë's book, I remain unconvinced by the enactive thesis that the content of one's perceptual experience of an object can be fully derived from one's implicit sensorimotor knowledge of the sensory consequences of one's own movements relative to this object.[1]

Action in Perception is a pleasure to read: it is replete with insightful arguments and interesting examples. I started reading the book with the suspicion that, like many other sensorimotor approaches to perception, the enactive conception is a form of behaviorism. After reading the book, I think that Noë convincingly rebuts the charge of behaviorism since his view is that my perceptual experiences arise, not from the lawful dependencies between my bodily movements and my sensory stimulations, but from my sensorimotor *knowledge* of the sensory consequences of my own movements.[2]

In a nutshell, Noë's overall project is to graft insights from James Gibson's ecological approach to visual perception onto John McDowell's view that the content of perceptual experience is thoroughly *conceptual*. I think that the attempt is quite interesting, but I fear that Noë successfully rebuts the charge of behaviorism only at the cost of "overintellectualizing" the content of perceptual experience. Here is how I shall proceed. First, I shall examine the question whether the empirical work on change blindness supports the denial of detailed internal visual representations of the world. Secondly, I shall explain why the book revolves around what Noë calls "the puzzle of perceptual presence". Thirdly, I shall address some objections to Noë's solution to the puzzle raised by McDowell's claim that we can embrace in thought all the different shades of color that we can visually experience. Finally, I shall object to Noë's enactive account of the phenomenon of perceptual constancy.[3] In so doing, I will draw attention to the tension between the enactive account of perceptual constancy and two significant trends in recent cognitive neuroscience: the two-visual systems model of human vision and the internal models theory of action.

1. Does change blindness support the denial of detailed internal visual representations?

According to the classical computational representational paradigm, which has been criticized by supporters of the enactive approach, visual perception consists in the brain's computing and transforming a series of mental representations in response to visual inputs. Many supporters of the enactive conception have expressed the view that visual

perception does not require the construction of detailed representations of visual stimuli on the grounds that the world itself can serve, to quote the cognitive scientist Rodney Brooks and the vision psychologist Kevin O'Regan, as "its own best model" or as "an outside memory". As Noë puts it (p. 50), "there's no need to re-present the world on one's own internal memory drive. Off-loading internal processing onto the world simplifies our cognitive lives and makes good engineering and evolutionary sense". It is an interesting recurrent theme of Noë's book that our perceptual sense of the presence of details in a visual scene does not require the construction (or the postulation) of a detailed internal representation. We experience the detail of a visual scene as *virtual* or *accessible*. Our experience of the details of a visual scene does not arise from a detailed visual internal representation. The detail belongs so to speak to the outside world, it is not intrinsic to our visual representation of it. We implicitly know that the detail is available for us to retrieve if we need to.

The question I want to consider is: which aspects of the computational representational paradigm does the enactive approach exactly object to? When he considers Marr's computational research program (in chapter 1), it is not quite clear whether Noë wants to deny that "vision is a process whereby the brain produces an internal representation of the world" (p. 21) or whether he wants merely to deny that "vision is a matter of generating a detailed internal representation of the visual world on the basis of information available at the retina alone" (p. 22). Here, Noë slips from rejecting the view that internal *detailed* representations could be formed on the basis of information available at the retina alone, to denying internal representations at all. Now, one may certainly agree that in perceptual tasks, information available at the retina must be supplemented by information about the position of the eye in its orbit together with information about head-position without denying that vision involves the production of internal representations of the world.

Some enactivists claim that their own brand of anti-representationalism derives much empirical support from spectacular experimental work on change blindness and inattentional blindness by a number of perceptual psychologists including Kevin O'Regan among others.[4] The term "change blindness" refers to the surprising finding that, in many experimental circumstances studied in the lab, normal subjects may fail to notice large changes in a rich visual scene. It was first discovered that observers fail to notice changes in photographs when the changes are made during an eye movement. Interestingly, some of these initial experiments were anticipated, and their results were predicted, by Dan Dennett on the very last page (p. 468) of his famous 1991 book, *Consciousness Explained*. Then, using the so-called "flicker" paradigm, observers were shown repeated alternations between an original scene and a modified version separated by a brief blank display. Again, it took a long time until observers could detect the difference. Similar results were obtained for changes made during saccadic eye movement, eye blinks, so-called "mud splashes" and brief occlusions.

Some supporters of the enactive conception have claimed that the results on change blindness show that visual representations must be sparse, incomplete or altogether non-existent. Now, I am not so sure that this is what the experimental results prove. Nor is Noë by the way. At times, he claims that what change blindness shows is that "the representations needed to subserve vision *could* be virtual", not that there are *no*

detailed internal representations (p. 52). Let us see. First of all, in many experiments on change blindness, normal cues of change between two scenes are skillfully concealed from observers by experimenters. Secondly, as Dretske (2004) has recently argued, in the typical case, what observers fail to notice are the *differences* between two scenes or two displays that are produced by a change. Now, one cannot notice a difference between two distinct scenes perceived in succession unless one makes a comparison between the two successive scenes. Arguably, one could not compare at t_2 a scene perceived then with a scene perceived at t_1 unless one had available at t_2 a representation of the scene perceived at t_1. But as Simons and Rensink (2005) have recently observed, failing to perform such a comparison could occur in the presence of detailed internal representations of the stimuli for at least four separate reasons. (a) The detailed representation of the display perceived at t_1 could decay or fade away before comparison takes place at t_2. (b) The detailed representation of the display perceived at t_1 could be encoded in a neural pathway unavailable for comparison. (c) The content of the detailed representation of the stimulus perceived at t_1 could be encoded in a format unsuitable for comparison. (d) Although the detailed representation has not decayed until the comparison process takes place and although the representation is in the right pathway and in the right format, still the comparison process itself could fail for some other reason. For example, the comparison process could be too slow relative to the temporal window (the interval between t_1 and t_2) imposed by the experimenters on the task. To conclude: it is by no means obvious that experimental work on change blindness supports the weaker negative thesis that vision does not consist in forming internal detailed representations on the basis of retinal information alone, let alone the stronger negative thesis that vision involves no internal representation of the world at all.

2. The puzzle of perceptual presence

One of the major tasks—if not the major task—on the agenda of Noë's enactive conception of perception is to solve the puzzle of perceptual presence. Solving the puzzle of perceptual presence is the litmus test—or a condition of adequacy—for the enactive conception of perception. Roughly, the puzzle of perceptual presence is that in visually perceiving an object we are perceptually aware of unattended features of the object. This is true of perceptual tasks in both the visual and the tactile modalities, but I am not sure that it is true of either auditory or olfactory experiences.[5] The puzzle of perceptual presence has many different instantiations in visual tasks. In perceiving a single opaque three-dimensional physical object, whose back is hidden from me by its front surface, I am perceptually aware, not merely of the front of the object, but of the full object including its invisible back. In perceiving e.g., a tomato, I am not presented with the hidden surface of the tomato. Nonetheless, I enjoy a visual experience of the tomato as voluminous and three-dimensionally extended. In perceiving an object whose front surface is partly occluded by other objects, I am perceptually aware of the full object including the occluded parts of its surface. For example, in perceiving a cat partly occluded behind a picket fence, I have a perceptual sense of the presence of the full cat.

On the face of it, the enactive conception of perception provides a nice solution to the puzzle of the perceptual presence of both the voluminosity of a perceived tomato and the fullness of a partly occluded perceived cat. As I said, Noë takes the work on change blindness to show that the content of perceptual experience is *virtual*: our perceptual

sense of the presence of detail in a visual scene does not require the construction (or the postulation) of a detailed internal representation; we experience the detail of a visual scene as virtual or accessible. Our experience of the detail of a visual scene does not arise from a detailed visual internal representation. The detail belongs so to speak to the outside world, it is not intrinsic to our visual representation of it. We implicitly know that the detail is available for us to retrieve if we need to.

Now, the perceptual presence of the voluminosity of the perceived tomato or of the fullness of the perceived cat is virtual presence, which Noë also calls *amodal* perception, and about which he has very interesting things to say in connection with filling in the blind spot and with the perception of the virtual illusory contours of e.g., Kanizsa triangles. Our sense of the virtual presence of unattended features of a perceived object arises from our practical or implicit sensorimotor knowledge of the sensory consequences of our bodily movements. My sense of the voluminosity of the tomato arises from my practical sensorimotor knowledge that if I move my eyes and head, and if I grasp and manipulate the tomato with my hand, I shall be able to visually experience parts of the tomato which I am not presently experiencing. Similarly, my sense of the full cat partly occluded behind the picket fence arises from my practical sensorimotor knowledge that if I walk across the picket fence, I shall be able to see parts of the cat which are presently occluded and which I am not visually experiencing at the moment.

Interestingly, if the enactive conception of perception is a solution to the puzzle of perceptual presence, then our amodal perception of the occluded portions of the three disks in a Kanizsa triangle affords a model for our sense of the presence of the unattended features of either a perceived tomato or a perceived cat. If so, then two features of the enactive solution to the puzzle of perceptual presence are worth mentioning at this stage. First of all, there is, I think, an interesting difference between one's sensorimotor expectations relative to a tomato and one's sensorimotor expectations relative to a cat. In seeing a cat, but not a tomato, I expect it to generate sensory consequences from its own internally generated movements. Unlike the perceptual sense of the voluminosity of the tomato, the perceptual sense of a full cat that is partly occluded depends, not just on my knowing the sensory consequences of my own bodily movements, but also on my expectations of the sensory consequences of the cat's bodily movements. But Noë's treatment of the puzzle of perceptual presence fails, I think, to take these differences into account. Perhaps Noë thinks that the difference between my implicit knowledge of the sensory consequences of the cat's actions relative to me and my implicit knowledge of the sensory consequences of my own actions relative to the cat is immaterial. But if he does, then I think he is wrong, as I shall argue (towards the end of this paper).

Secondly, the question arises whether acceptance of Noë's enactive solution to the puzzle of perceptual presence meets the requirement (stated on p. 60) that what is wanted is not an account of our ability to think, believe or judge that unattended features of a perceived object are conceptually present, but an account of our *perceptual* sense of presence or that the phenomenon of perceptual presence is a genuinely perceptual phenomenon. The question is made more pressing by the fact that Noë subscribes to the conceptualist thesis that perceptual experience is intrinsically and thoroughly *thoughtful*. Indeed, the penultimate chapter of his book is devoted to providing challenging arguments for the thesis that the content of perceptual experience is fully conceptual and

to rebutting the arguments in favor of the distinction between the conceptual content of thought and the nonconceptual content of perceptual experience.

3. Noë on color concepts

As I said above, Noë's overall project is to offer further support for McDowell's (1994) conceptualist standpoint on the content of perceptual experience. Referring to our visual experience of colors, Noë fully endorses McDowell's well-known dictum that "we are equipped to embrace all the shades we can see in conceptual thinking" (cited by Noë on p. 190). His strategy is twofold. On the one hand, as Noë insightfully points out, acceptance of Evans' and other nonconceptualists' insistence on the belief *independence* of perceptual experience (in the case of e.g., visual illusions) does not commit one to the belief *indifference* of the content of perceptual experience. Only belief indifference, however, not mere belief independence, can serve the purpose of the distinction between the conceptual content of thoughts and the nonconceptual content of perceptual experiences. As Noë elegantly puts it on behalf of his conceptualist viewpoint, "the content of perceptual experience is conceptual not in the sense that it *is* judged, but in the sense that it *can be* judged" (p. 189). On the other hand, Noë rejects as "much too exalted" (p. 185) the conception of our own conceptual capacities according to which "all concept use must take the form of explicit deliberative judgment" (p. 199). For the purpose of deflating it, he urges a less demanding account of the cognitive conditions for the possession of some—e.g., sensorimotor—concepts according to which sensorimotor skills can count as a simple kind of concept possession.

No doubt, his account of the content of color experiences (in chapter 4), as arising from one's sensorimotor expectation that the apparent color of an object varies as one moves relative to the object, appears to be consistent with his deflationary account of the possession of sensorimotor concepts. However, I do not think that the account of the content of color experiences based on one's sensorimotor knowledge of the changes in the apparent color of objects consequent upon one's bodily movements can be squared comfortably with his further suggestion that the possession of color concepts amounts to the possession of what he calls "color-concept-formulae" (p. 194) or that color concepts are formal concepts in just the way numerical concepts are.

Anybody who accepts McDowell's idea that "we are equipped to embrace all the shades we can see in conceptual thinking" faces the challenge of specifying how we can be endowed with a concept for every possible shade of color to which we are perceptually sensitive. The analogy between color concepts and numerical concepts is meant to solve McDowell's puzzle. As Noë observes, our mastery of numerical concepts allows us to represent any of the infinitely many natural numbers. As he puts it (p. 196-7), "the symbolism we use to represent numbers is such that it contains within it recursive procedures for generating number names ad infinitum". Arguably, the contextual use of perceptual demonstratives together with the recursive resources of the grammars of natural languages may enable us to express demonstrative color concepts for an open ended set of color shades.

For at least three related reasons, however, I think that the analogy between color concepts and numerical concepts is flawed. First of all, as Noë himself acknowledges (pp. 197-8), it is likely that such demonstrative color concepts will not allow much by way of

recognition of a color shade over time. But they should. Now, if demonstrative color concepts do not allow recognition over time, in what sense are they genuine *color concepts* at all? Secondly, whereas the application of a perceptual demonstrative to a color is highly contextual, the contribution of a numerical concept to a numerical thought is as context-independent as human thought is ever going to get. Finally, whereas arithmetic possesses its own recursive resources for constructing the representation of any natural number, the only recursive procedures available for constructing color concepts are provided either by the grammatical resources of natural languages or by the syntactic resources of the language of thought, if there is a language of thought. In either case, the recursive resources are not specific to color concepts and/or color thoughts.

4. Can perceptual constancy be enacted?

As I said above, the litmus test for Noë's enactive conception of perception is that it ought to solve every instance of the puzzle of perceptual presence. It is, I think, fair to say that much of the burden of Noë's whole book is to pave the way for the claim that every instance of the problem of perceptual constancy is best framed as a version of the puzzle of perceptual presence for which the enactive conception of perception provides a solution.[6] Now, the problem of perceptual constancy arises for the visual experience of size, shape and color. An object looks to be the same with respect to its size, shape and color even though the way the size, the shape and the color of the object look are different in different circumstances (as the distance between the object and the observer changes, as the illumination conditions change, and so on). As the circumstances change, the apparent size, the apparent shape and the apparent color of an object change. Even though its "perspectival" properties (or P-properties, as Noë calls them) change according to the circumstances, nonetheless the object's actual shape, its actual size and its actual color look to be the same (perceptual constancy). As Noë is fond of emphasizing, "perceptual experience is radically ambiguous" (p. 34). Perceptual content has two dimensions: a factual dimension and a perspectival dimension. Perception is a way of representing and keeping track of both the actual properties of things (the way things are) and the perspectival properties of things (our relations to things). "Perception is thus world-directed and self-directed." (p. 168)[7]

On the face of it, the enactive conception of perception does provide an elegant solution to each version of the puzzle of perceptual presence raised by the perceptual constancy of the size, the shape and the color of an object. Consider shape. From a particular location, the circular plate appears elliptical. The plate's actual shape is circular, but its apparent, perspectival or P-shape is elliptical. According to the enactive conception, our experience of the actual circularity of the plate arises from our sensorimotor knowledge of the set of transformations of the P-shapes of the plate determined by our own bodily movements relative to the plate. As Noë repeatedly puts it, we experience the circularity of a circular plate in virtue of experiencing its P-shape from a particular location: "the actual shape and size are invariants we encounter when we explore visual variation produced by movement" (p. 85). Our visual experience of the actual shape results somehow from our experiences of the different P-shapes of the object together with our sensorimotor knowledge of the set of transformations of our visual experiences of the different P-shapes produced by our bodily movements relative to the plate. The same account applies to the perceptual constancy of colors:

> […] color perception and shape perception are on a par. You experience the roundness of the plate in the fact that it looks elliptical from here and that its elliptical appearance changes (or would change) in precise ways as your relation to the plate […] changes. In exactly this way, we experience the color of the wall in the fact that the apparent color of the wall varies as lighting changes. We are able to experience the actual color of the object as, so to speak, the condition which governs or regulates the way these changes unfold (p. 128).

For the sake of argument, I am willing to grant Noë that the content of visual experience is twofold: "the plate looks circular (it really does) *and* it looks elliptical from here (it really does). As Kelly (2004) argues convincingly, however, I don't think that the plate can *simultaneously* look circular and elliptical. The wall looks to be uniform in color across its surface *and* it appears brighter, where it falls in direct light" (p. 164). But I have serious doubts about the view that we experience the actual shape and the actual color of an object in virtue, or as a result, of our experience respectively of its P-shape and of its P-color together with our sensorimotor knowledge of the dependencies of our experiences of P-properties upon our bodily movements. In other words, I have serious doubts about the enactive account of the phenomenon of perceptual constancy. For one thing, in spite of Noë's explicit disclaimer (on p. 85), his enactive account does seem to turn the experience of perceptual constancy into the conclusion of a plain inferential process—whether this process is thought of as a plain inferential process at the personal level or as a computational process at the sub-personal level. Noë claims (p. 85) that "my experience of the circularity *just is* my experience of the variation in its perspectival shape" (*ibid*.). This does fit with Noë's claim that visual experience is thoroughly conceptual and thoughtful. But it also does make it seem as if one thinks or concludes that the plate is circular from the belief that it looks elliptical from here.

In fact, I think both parts of the enactive account of perceptual constancy are open to doubt. I cannot accept the priority given to the representation of the P-properties of objects over the representation of their actual properties. Nor can I, for two complementary reasons, fully accept the idea that perceptual experience depends on our sensorimotor knowledge of the sensory consequences of our own bodily movements.

First of all, focusing on shape for the sake of concision, I agree that a circular plate may look circular and that it may look elliptical from here, at different times. But the reason I do not think that my visual experience of the circularity of the plate can be derived in any intelligible sense from its looking elliptical from here is precisely that, whereas the plate just looks circular, it looks elliptical *from here*. Now, I may have been brainwashed by John Perry's (1986) seminal essay on unarticulated constituents. But it strikes me that the judgment (or the thought) that the plate looks elliptical from here is more complex, not less complex, than the judgment (or the thought) that it looks circular. Arguably, the latter thought reflects (or arises from) my immediate perceptual experience of the plate. Unlike the latter, the former includes an explicit constituent referring to my current location. The former thought would simply not arise were I unable to reflect on my spatial relation to the plate. I can see the plate and experience its circularity. I can also reflect upon (or think about), but I cannot see, my spatial relation to the plate. Reflection

upon the spatial position of the plate relative to me, however, is required for me to become aware of the way the circular plate looks from here, i.e., of the P-shape of the plate relative to my current position. It is unlikely that one *generates* the less complex judgment about actual shape from the more complex judgment about how P-shape looks from here if 'generating' means moving from simple to complex according to some compositional principle. So I agree that perception is two-dimensional over time, but I do not think that the perceptual representation of the actual non-relational properties of things (perceptual constancy) can arise from the perceptual representation of the perspectival properties of things, which are relational properties.[8]

Secondly, the question arises whether the enactive conception of perception is really compatible with the evidence in favor of the two-visual systems model of human vision, according to which humans can visually process objects for two distinct purposes: for perception and for the visuomotor control of action.[9] Noë quickly dismisses the question (e.g., on pp. 12, 19) by distinguishing the claim, which he takes to be *constitutive* of the enactive conception, that visual perception is made possible by the perceiver's sensorimotor "practical grasp of the way sensory stimulation varies as the perceiver moves" from a putatively distinct presumably empirical claim, which he rejects, according to which "perception is *for* acting or for guiding action". No doubt, the two-visual systems hypothesis is an empirical hypothesis that should not be treated as an infallible dogma.

However, I am not convinced that Noë can reconcile the enactive conception with the evidence for the two-visual systems hypothesis by conferring to his own appeal to one's sensorimotor knowledge of the sensory consequences of one's own actions the status of a constitutive claim. Nor is it always clear that Noë really rejects the view that perceptual experience guides action, as when he writes (e.g., on p. 90) that "there is no question that experience can and does guide movement, but it is not the business of the enactive view to emphasize this humdrum fact". The evidence for the two-visual systems hypothesis is precisely what would lead one to doubt at least the strong assertion that the function of perceptual experience is to guide movement. On several occasions, Noë endorses consequences of the view that the function of visual perception is to serve action. For example, in the footsteps of Poincaré's famous view that "to localize an object simply means to represent to oneself the movements that would be necessary to reach it" (quoted p. 75), Noë seems to take the view that localizing an object in some egocentric frame of reference (centered on the axis of the agent's body) is a necessary condition for one to experience its P-properties. For example, he writes:

> the plate looks elliptical to me because, to indicate its shape, I can (and indeed, in some sense, *must*) move my hand in a characteristic manner. That is, to experience a thing as elliptical is precisely to experience it as occupying a particular region in one's egocentric, sensorimotor space […] the P-size of an object is given precisely as that region to which I, for example, would point or reach (etc.) if I wished to direct myself to the object (p. 89).

On the one hand, Noë's account of the P-shape of the plate smacks of behaviorism: my hand movements may reveal or manifest my visual experience of the P-shape of the plate (e.g., to others), but the former does *not* explain the latter. On the other hand, the human visual system can represent the spatial location of an object relative to several different frames of reference: e.g., relative to an egocentric frame of reference or relative to an allocentric frame of reference (centered on some constituent of the visual array).[9] An object's location must presumably be represented in an egocentric frame of reference for the agent to act upon it (e.g., to reach and grasp it). But it is by no means obvious that an object's location must be represented in an egocentric frame of reference for an observer to experience either the object's P-shape or its P-size.

Finally, the question arises whether the content of perceptual experience can arise from one's sensorimotor knowledge of the sensory consequences of one's own bodily movements with respect to perceived objects. The reason the question does arise is that much work in the cognitive neuroscience of action and the study of human motor cognition shows that when we perform a voluntary movement, we do anticipate and, on this basis, we discard the sensory consequences of our own movements. According to the so-called "internal (forward) models" theory of voluntary action, the human motor system is richly endowed with predictions of the sensory consequences of one's own bodily movements.[10] Each time, my motor intention to act sends a motor command to some of my muscles, it also sends a so-called "efference copy" to some central monitor for planning, controlling and checking the execution of the action. Since no efference copy is being sent to my central monitor when I perceive an action executed by another, the mechanism of sending an efference copy is thought to contribute importantly to the sense of agency, i.e., to the sense of my being the agent of my own action. So far so good. But there is another side to the hypothesized contribution of the efference copy mechanism, which is that when they arise from my own actions, sensory stimuli are predictable and, therefore, not worth paying perceptual attention to. Unlike sensory stimuli that do *not* arise from my own actions, those that do need not be perceptually processed. So the efference copy mechanism contributes both to one's sense of agency and to decreasing the computational load of the perception of sensory events that arise from one's own actions, by attenuating their sensory consequences.[11]

Contrary to the enactive view that my sensorimotor knowledge of the sensory consequences of my own bodily movements is the main source of the phenomenal content of my perceptual experiences, the internal models theory of action suggests that because they are predictable, the sensory consequences of my own actions should not give rise to perceptual experiences endowed with much phenomenal content. This is why I earlier said that there is a difference between my sensorimotor expectations relative respectively to a cat and to a tomato. Given the enactive theorist's commitment to emphasizing the contribution of one's own actions to the content of one's perceptual experiences, I thought the enactive theorist should face the tension between his own account and the internal models theory of action.[12]

References

Blakemore, S.-J., Wolpert, D. & Frith, C.D. (2000). Why can't you tickle yourself? *NeuroReport*, 11, R11-R16.

Dennett, D. (1991). *Consciousness Explained*, London: Allen Lane, The Penguin Press.

Dretske, F. (2004). Change blindness, *Philosophical Studies*, 120, 1-18.

Jacob, P. & Jeannerod, M. (2003). *Ways of Seeing, the Scope and Limits of Visual Cognition*, Oxford: Oxford University Press.

Jeannerod, M. (2003). The mechanism of self-recognition in humans, *Behavioral Brain Research*, 142, 1-15.

Kelly, S.D. (2004). Perception, action, and the constancies (unpublished paper read at the Institut Jean Nicod on December 10, 2004).

McDowell, J. (1994). *Mind and World*, Cambridge, Mass.: Harvard University Press.

Milner, A.D. & Goodale, M.A. (1995). *The Visual Brain in Action*, Oxford: Oxford University Press.

Noë, A. (2004). *Action in Perception*, Cambridge, Mass.: MIT Press.

O'Regan, J.K. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness", *Behavioral and Brain Sciences*, 24, 939-1031.

Perry, J. (1986). Thought without representation, in Perry, J. (1993) *The Problem of the Essential Indexical and Other Essays*, Oxford: Oxford University Press.

Simons, D.J. & Rensink, R.A. (2005). Change blindness: past, present, and future, *Trends in Cognitive Sciences*, 9, 1, 16-20.

Tsakiris, M. & Haggard, P. (2004). Experimenting with the acting self, *Cognitive Neuropsychology*, 21, 1-19.

Wolpert, D. (1997). Computational approaches to motor control, *Trends in Cognitive Sciences*, 1, 209-216.

Notes

1. Incidentally, although Noë does not restrict his discussion to visual perception, he devotes far more space to visual perception than to perception in other modalities. Furthermore, his discussion of visual perception is restricted to the perception of physical objects at the expense of the perception e.g., flashes, shadows, holes, events or actions, about which it is not clear what the enactive conception would say.

2. On one specific occasion (p. 89), to which I shall come back towards the end of the paper, however, Noë does lapse into behaviorism.

3. While reading and thinking about Alva Noë's book, I heard and read an unpublished paper by Sean Kelly, which he delivered at the Institut Jean Nicod in the Fall 2004. I found Kelly's paper quite helpful.

4. See e.g., D.J. Simons & R.A. Rensink (2005) for a review of work on change blindness.

5. Noë does not say anything about either audition or olfaction. Nor will I.

6. But as I will argue, I think that the enactive conception fails to solve the problem of perceptual presence raised by perceptual constancy.

7. Noë's two-dimensional approach to the content of perceptual experience allows him to make an interesting contribution to issues raised by the causal theory of perception (pp. 169-75).

8. Noë concedes this point when he writes (p. 166): "There can be little doubt that we do not normally reflect on apparent shapes, sizes, and colors when we look around". My target is the view that my conscious visual experience of the 3-D shape of an object derives from my conscious visual experience of its P-shape. Given Noë's conceptualist assumptions, the view I criticize is that my thought that an object exemplifies its 3-D shape derives from my thought that it exemplifies a variety of P-shapes. But of course I am not arguing against the view that the visual processing of the 2-D shape of an object contributes to the sub-doxastic computational process whereby I become visually aware of the object's 3-D shape.

9. See Milner & Goodale (1995) and Jacob & Jeannerod (2003) for a slightly different account.

10. In contemporary philosophy and cognitive science, the word "egocentricity" is sometimes used to capture the context-dependency of demonstrative and indexical thoughts (e.g., expressed by "here"), which are genuine conceptual representations. It is also sometimes used to refer to a property of a particular kind of frame of reference relative to which the visual system represents the spatial position of an object.

11. I draw upon numerous sources. See e.g., Blakemore et al. (2000), Jeannerod (2003), Tsakiris & Haggard (2004) and Wolpert (1997).

12. Apparently, this is why one cannot tickle oneself (see Blakemore et al., 2000).

13. Thanks to Tim Bayne and Anne Tüscher for comments on these comments.