

EXPERIENCE WITHOUT THE HEAD

Alva Noë

To appear in *Perceptual Experience*, edited by
Tamar Gendler and John Hawthorne. (Oxford University Press).

DRAFT ONLY.

PLEASE DO NOT CITE WITHOUT PERMISSION

Some cognitive states — e.g. states of thinking, calculating, navigating — may be partially external because, at least sometimes, these states depend on the use of symbols and artifacts that are outside the body. Maps, signs, writing implements may sometimes be as inextricably bound up with the workings of cognition as neural structures or internally realized symbols (if there are any). According to what Clark and Chalmers [1998] call active externalism, the environment can drive and so partially constitute cognitive processes. Where does the mind stop and the rest of the world begin? If active externalism is right, then the boundary cannot be drawn at the skull. The mind reaches — or at least *can* reach --- beyond the limits of the body out into the world.

Can one extend active externalism to perceptual consciousness? There is a consensus that this question should be answered negatively.¹ The fact that we dream, and

¹ Dretske [1995] criticizes internalism about experience. An experiential state has a content only given the animal's history and relation to its environment. Dennett [1987] develops a similar position. Accounts such as these treat experiences as internal items whose nature is fixed by external, animal-world relations. They are externalist about the content of experience but internalist about the vehicles of experience [Hurley 1998]. On this sort of hybrid view, mental states are comparable to sunburns, as observed by Wilson

that neuroscientists can produce sensations by direct stimulation of the brain, shows that consciousness is a matter of what is going on in the head alone.

Or does it? The fact that some experiences can be produced by neural activity alone does not show that all experiences could be. Likewise, the fact that no change in experience can occur without changes in internal neural states does not show that neural states suffice for the occurrence of any experience. — We spend our lives in tight coupling with the environment (and other people). Why are we so confident that there could be a consciousness like ours independent of active exchange with the world? Why are we so certain consciousness depends only on our insides? Perhaps we are too hasty in dismissing externalism about perceptual experience itself.

In this paper I argue that the standard “consciousness is in the head” consensus rests on bad phenomenology. We think of experiences as like snapshots and we suppose that what is experienced, like the content of a snapshot laid out on paper, must be given all at once *in the head*. But experiences are not like snapshots. The content of perceptual experience is not given all at once the way the content of a picture is. In ways that I will try to explain, *what* we experience visually (for example) outstrips what we actually see. At least sometimes, the world itself can drive and so constitute perceptual experience. The world enters into perceptual experience the way a dance partner joins us in the dance, or – to change the image slightly – the way the music itself guides us.

[1997]: The sunburn (or mental state) is literally on the skin (or in the head), but its nature depends on its world-involving causal history. In this paper I advance a more thoroughgoing externalism about the very vehicles of experience.

1. A puzzle about perceptual presence

It is a basic fact about perception that solid, opaque objects, when seen from a vantage point, have visible and invisible parts [Koenderink 1984]. When you see a tomato, for example, you see its visible aspect. Euclid captured this thought when he wrote: “Nothing that is perceived is seen at once in its entirety.”

No surprise here. What could be more evident than that you can’t see the occluded portions of objects that you perceive? When you see a tomato, you can’t see its back. When you see a cat behind a picket fence, you only see, strictly speaking, those parts of the cat that show through the slats.

There is a way of thinking about experience – a reasonable way – according to which these Euclidean truisms can come to seem untrue. Yes, to use Euclid’s formulation, a perceived object is never perceived at once in its entirety. Nevertheless, one can hardly dispute that we take ourselves, when we see the tomato, or the cat, to have a sense of their presence – a *perceptual* sense of their presence – as wholes. In the case of the tomato, for example, you have a sense of the presence of a voluminous, ovoid, furrowed whole.

It is in defense of this reasonable way of thinking about perceptual experience that Thompson Clarke [1965] insisted that seeing is like nibbling. When you nibble a piece of cheese, you nibble it, the cheese, not merely a part of it; and so, when you see a tomato, you see precisely *it*, the tomato. It is only in special circumstances that it is correct to say, when you see a tomato, that you see only a part of it (just as, I presume Clarke would say,

it is only in special circumstances that it is correct to say, when you see a tomato, that you see *the whole of it*.)

P. F. Strawson [1979] is getting at a related idea – also in defense of this reasonable way of thinking about experience – when he claims that one *distorts* the visual experience of a tomato, one misdescribes or mischaracterizes it, if one describes it as a visual experience as of a tomato *part* (let alone as of tomato-like sense-data). The visual experience of the tomato, when one takes it at face value, presents itself to one precisely as a *visual experience as of the presence of a tomato*.

There is much to be said on behalf of these defenses of familiar beliefs about perception. It would be a mistake, however, to think that they give us reason to doubt the Euclidean observation stated at the outset. For the Euclidean observation is no less well entrenched in our ordinary thought and phenomenology. Both sides in this philosophical standoff – Euclid, on the one side, and common sense, on the other – show a tendency to lapse into dogmatism. The Euclidean insists, dogmatically, that when we take ourselves to perceive tomatoes, we “go beyond” what we really see. The reasonable philosopher, no less dogmatically, finds him or herself wanting to deny that our perceptual experience is confined by the limits of immediate perspective. “The plate doesn’t look elliptical, it looks round!”

What I am calling the problem of perceptual presence comes clearly into focus when we acknowledge that both sides in this standoff are, in a way, right.² The plate

² See Noë [2002]; O’Regan and Noë [2001] for more on the problem of perceptual presence.

looks round, and it looks elliptical. That is, we have a sense of the presence of the plate's roundness despite the fact that, plainly, it looks elliptical from here. And so for the tomato. When you see a tomato, you only see, strictly speaking, the visible face of the tomato, even though it is also true that you are visually aware of the presence of the parts of the tomato which you don't actually see.

How can both these facts about perceptual experience be, well, just that, *facts about perceptual experience*? How can it be true, as I think it is, that we are perceptually aware, when we look at a tomato, even of parts of the tomato which, strictly speaking, we do not perceive? This is the puzzle of perceptual presence: in what does our sense of the perceptual presence of a strictly unperceived features of the world consist?

Before proceeding, two warnings. First, the puzzle of perceptual presence should not be confused with another nearby and closely related problem, namely, the epistemological problem of *the given* in perception. Our problem is not with whether what is given provides sufficient rational basis for perceptual judgment, but rather with the question, *what is given*? Our focus is phenomenological, on the nature of perceptual content itself. The point is that what is given is, at least apparently, rife with conflict. The plate looks round *and* it looks elliptical from here; we encounter only the visible parts of the tomato *and* we take ourselves to be aware of the presence of its strictly unperceived parts. Or consider further examples: the two trees appear to be of the same size *even though* the nearer tree looks larger than the farther one; the wall appears to be a uniformly colored surface, *despite* the fact that this part of the wall is visibly brighter (where it falls in direct sunlight) than that part of the wall (which is cast in shadow).

Second, as the immediately preceding remarks indicate, the problem of perceptual presence is of surprising generality, comprising a range of perceptual phenomena not usually grouped together, including those already mentioned – occlusion shape, color constancy, apparent size – but also others as well (as we shall see).

2. First stabs at a solution

It may be tempting to bite the bullet and concede that we *don't really see* the whole tomato, or the roundness of the plate, or the whole cat, etc. We go *beyond* what is strictly given in an account of our experience when we in this way describe what we see. Our feeling that we see the whole tomato, say, is an illusion.

But this objection misses the point. The puzzle is not that it seems to us as if we see the whole tomato, when we don't really, or that we experience the color as uniform, when in fact it is nonuniform. This is the epistemological problem I mentioned in the last section. Our puzzle is that it seems to us *at once* as if we only see part of the tomato (evidently!) *and* as if the whole is perceptually present without being strictly seen. It seems to us as if we see the roundness of the plate *when* it looks elliptical. We take ourselves to sense the presence of a uniform color, even though the surface is dappled in light and thus variegated in apparent color. We take ourselves to be perceptually aware of features which, manifestly, we do not see and so feel no inclination to take ourselves to see.

Nor can it help us here to be told that although we don't *see* the hidden parts of the tomato, or the cat, we *infer* their presence. There is something to this line of thought,

no doubt. After all, we know what tomatoes and cats are, we have these concepts; we make use of these concepts in fleshing out or indeed in “cognitively filling in” what is given to us. I think this must be right; however it provides no solution to the problem of perceptual presence. It can’t be the whole story.

For what we want is an account not of our *thought* or *judgment* or *belief* that there is a whole tomato there, or a whole cat there, or a uniformly colored wall there. What we want is an account of our *perceptual sense of their presence*.

Crucially – and this is a phenomenological point – the cat seems present as a whole, *perceptually*. The voluminous tomato seems *perceptually* present. That is, we do not merely *think* that they are present. Indeed, this sense of perceptual presence does not depend on the availability of the corresponding belief.

Consider Figure 1, an illustration of Kanisza’s. It is natural to perceive it as depicting a rectangle partially occluding four disks. Our sense of the perceptual presence of the disks is not significantly altered by the explicit recognition that there are not really occluded bits present.

[Figure 1 About Here]

We don’t merely *think* the presence of the occluded bits: after all, they are, evidently, *not* present, they are blocked from view (or rather, not drawn); it *looks* as if they are blocked from view. We experience the presence of the occluded bits even as we experience, plainly, their absence. They are present *as absent*.

Psychologists call this phenomenon *amodal* perception: perception, but not in any modality. This can't be quite right either. The paradoxical quality is sharper. The phenomenon is best characterized as amodal *visual* perception, as a kind of seeing without seeing.

3. Presence as absence

Perceptual presence *in* absence – amodal perception -- is a widespread perceptual phenomenon. Consider some further instances.

Filling-in at the blind spot. In his recent textbook, Steven Palmer [1999] suggests that the brain fills in to make up for the gap or discontinuity at the blind spot. We know this, he explains, because of the results of demonstrations such as the following. Consider Figure 2. If you shut the left eye, and fixate the cross with the right eye, you can adjust the illustration of a broken line so that the break falls in the blind spot. (This will occur when the page is about a foot from the face.) When it does, we have the experience of an unbroken line. The experience must in turn be underwritten by a neural process whereby an internal representation which has a gap in it is filled in. As Palmer writes: “The line on the retina actually has a gap in it at the blind spot, but we experience it as complete and uninterrupted when the gap falls within the blind spot. The important point is that what we experience visually conforms not to the firing of retinal receptors, but to some higher level of neural activity” [617]. Neural processes of filling-in in a higher-level neural representation are what bridge the gap between low-level retinal input and experience.

[Figure 2 about here]

Dennett [1991] has criticized this style of reasoning on the grounds that one isn't entitled to assume that brain produces the filled-in percept by a neural process of perceptual completion. Perhaps, he suggests, the brain instead ignores the absence of information corresponding to the blind spot, thus giving rise to the gap-free percept without requiring the actual construction of a gap-free internal representation. This line of thinking relies on the Kantian point that representations needn't have the properties they represent the world as possessing.

It is striking that Dennett, no less than Palmer, takes for granted, as we might put it, that *it seems to you as if the line is filled in*. After all, it seems as if the line is unbroken. But this is a mistake, or rather, it is an equivocation. In the relevant sense of "the line looks unbroken," it isn't the case that the line looks unbroken. Perform the demonstration and pay careful attention to what you see. Notice, it does not seem to you as if you *look* at the break and see that it is filled in there, that the line is completed there. It is just that, from where you are (as it were, from where you orient your eyes), you don't perceive the break. But not perceiving the break is different from perceiving the absence of a break. We naturally *say* that we perceive the line as unbroken, but in saying this we are not committing ourselves to the proposition that we are in qualitatively the same state as when we actually visually study an unbroken line.

One proposal, advanced by Durgin, Tripathy and Levi [1995], is that the experience of the line as filled in at the blind spot is, phenomenologically speaking, like the experience of the far side of the tomato, or of the hidden portions of the disks in Kanisza's illustration. They are filled in, but not *as seen*. As support for this point, they invite us to cover the break in the line with your thumb. Doesn't the line look filled in, and indeed, as filled in *in the same way*? In this case, the *amodal* character of the filling in is evident.

This is a controversial issue and there are arguments ranging on both sides.³ My point is phenomenological, rather than empirical. The phenomenology of the experience of the filled-in broken line is not like that of the solid broken line.

Once we get clear about the phenomenology we can observe that even if Dennett's criticism of Palmer's argument is right, Dennett and Palmer share a commitment to a misdescription of what the experience of the line is like.

Color constancy. Color constancy is illustrated by such facts as that we do not experience a change in colors of surfaces as illumination changes. For example, when you go outside, you don't experience the color of your shirt as changing. Or consider the wall: we perceive it as uniformly colored even though it is brighter here than it is there.

³ See, for example, Ramachandran and Gregory [1991]. For a general discussion of issues in this vicinity, see Pessoa, Thompson and Noë [1998].

Standard thinking in visual theory would have it that the brain must figure out what color to represent the wall as having, or the shirt as having, despite substantial changes in the character of reflected light entering the eyes.

Any account of color must explain constancy, no doubt. But this framing of the problem of color constancy is unsatisfactory because it misdescribes the phenomenology. The fact is, the wall does differ perceptibly in color as the illumination differs. To match the color of different parts of the wall you would need different color chips. The problem of color constancy, then, ought better to be framed as a problem about perceptual presence. How is that we experience the presence of a uniform hue across the surface of the wall, despite the differences in color here and there? How can we explain our ability to experience the invariant color amidst so much color variation?

Peacocke [1983] has used color constancy to illustrate the difference between the *representational* content of an experience (how the world is represented by the experience), and the *qualitative* or *sensational* content of experience (what the experience is like apart from its representational features). The experience of the wall here and there are the same in their representational content, but they differ nonrepresentationally in their qualitative character. Sean Kelly [2001] has resisted this analysis, arguing that what color constancy reveals is that some aspects of the *representational* content of experience are nonconceptual. Kelly proposes that there is a qualitative difference between our experience of the two parts of the wall that does not correspond to a difference in *which* color concept correctly applies. The qualitative difference *is* representational, on Kelly's view (in contrast to Peacocke), but it is *nonconceptual*.

This is not the place to address these issues about representational and nonconceptual content. However, I think both these arguments go astray because they underscribe the representational content of our experience. We experience the parts of the wall as the same in color *and* as different in color (or perhaps: as the same in color but different in their apparent color). The presence of the invariant color is like the presence of the far side of the tomato, or the roundness of the plate, or the sameness of size of the near and far tree: it is presence in absence, or implicit presence.

Change blindness and the experience of detail. Perhaps the most fascinating phenomenon of perceptual presence is the visual experience of detail. First, some background.

Visual theory has tended to take as its starting point a way of thinking about seeing according to which visual experiences are like snapshots. The idea is that visual experiences represent the world the way pictures do – all at once, in sharp focus, from the center out to the periphery. This snapshot conception is captured in pictorial form by Mach’s famous drawing of the visual field; it is caricatured, we might say, in Gursky’s well-known photograph of a Los Angeles 99 Cents Shop. The central aim of visual theory, as it has been practiced for the last century, has been to understand how the brain gives rise to this sort of snapshot-like, richly detailed experience. The problem is hard for two reasons. First, it is *ill-posed*, i.e. the two-dimensional retinal projection does not uniquely determine a three-dimensional layout. Second, the retinal image itself is defective (distorted, gappy, of uneven resolution). How, on the basis of such an impoverished stimulus, do we come to enjoy richly detailed experiential snapshots? The orthodox strategy for answering this question is to hypothesize that the brain integrates

the information available in successive fixations to form a detailed internal representation, which then serves as the substrate of the experience. *What* we experience is *what* is represented in this internal representation. Vision, according to the orthodox view, is the process whereby this internal representation is constructed.

Recent work in perceptual psychology on scene perception challenges the orthodox conception of vision precisely by challenging whether experiences are snapshot-like in the way that orthodoxy has tended to suppose. If they are not, then we aren't saddled with the problem of explaining how the brain gives rise to picture-like experiences. For example, work on change blindness (and related phenomena such as inattentional blindness) weighs against the snapshot conception.⁴ Our success as perceivers depends on the fact that we are very good at noticing flickers of movement and other attention-grabbing concomitants of change. We spontaneously direct our eyes to these transients and so discover change as it happens. It turns out that if we are prevented from noticing the associated flickers, or if there are no flickers – because, say, the relevant changes are too gradual – we will remain unaware of the changes going on around us, even when they are large-scale and pertinent to our interests and background concerns. In one noteworthy recent demonstration, due to Kevin O'Regan, perceivers are shown a photograph of a Paris street scene. Over the seconds that you look at the picture, the color of a car prominently displayed in the foreground, changes from blue to red. Perceivers overwhelmingly fail to notice this change in color, even though the change is

⁴ For reviews of this and related phenomena, see O'Regan in press; Simons and Levin 1997, Simons 2000, Noë, Pessoa and Thompson 2000.

dramatic and occurs over a short period of time. When the color change is pointed out, perceivers laugh aloud and express astonishment that they could have failed to miss the change.

Change blindness does not demonstrate that we (ordinary perceivers) tend to overestimate what we see. Given normal circumstances, we *are* very good at noticing changes. What change blindness illuminates is the degree to which this ability is vulnerable to disruption. Importantly, the fact of change blindness does reveal that a certain theoretical account of what seeing is – the snapshot conception and the associated idea that seeing is a process whereby a detailed representation is built up corresponding to what is seen – must be wrong. We don't seem to have access to such a detailed internal representation when we contemplating our environments.

Scientists and philosophers have sometimes suggested that change blindness reveals that our visual consciousness is a kind of confabulation (a “grand illusion”). It seems to us – doesn't it? -- as if when we open our eyes we *see everything*, the whole scene, in sharp focus and uniform detail, right out to the edges of the visual field. We do not, however (as revealed by change blindness and many other experiments and demonstrations). Therefore, the visual world is a grand illusion. Perceptual consciousness is a confabulation.

But this confabulation hypothesis is wrong.⁵ Granted, we do not enjoy snapshot-like experiences, as orthodoxy had supposed and as change blindness demonstrates. But –

⁵ See the papers collected in Noë 2002a for discussion of this issue.

and this is crucial –it is not the case that it *seems* to us as if we enjoy such snapshot-like experiences. We take the world to be densely detailed, yes. But we do not take ourselves to represent all that detail in consciousness at a moment in time, in the way that a picture might represent that detail at a moment in time. The snapshot conception is no part of ordinary perceptual phenomenology.

A little consideration is enough to bring this out. Consider your current visual experience of, say, the view out your window. You no doubt have a sense of the scene outside as dense and rich in detail. If you pause to reflect, however, you will notice it is not the case that it seems to you, *now*, as if all that detail is seen by you all at once, in an instant, in sharp focus and high resolution. Some things are clearly in view, others are present only indistinctly as background elements, and some items are not really experienced at all. To bring detail into consciousness, it is necessary to probe the environment, by turning your eyes, and your head, by shifting your attention from here to there.

The suggestion that visual experience is a grand illusion rests on a misdescription of the character of our experience of seeing,

Why do we find change blindness so surprising? Why do audiences gasp with astonishment when presented with immediate evidence of their own change blindness? Dennett [2001, 2002] has pressed this point: surprise, he has suggested, is an indicator of a foiled epistemic commitment. The question is, *what* epistemic commitment is thus shown to be foiled? It is enough to explain the surprise, I would say, to point out that we don't realize quite how vulnerable to disruption our ability to detect change really is. This doesn't entail that we take ourselves to see everything all at once!

But there is stronger evidence that we are not committed to the snapshot conception. In daily life, we continuously move our eyes and head in order to get better looks at objects around us in the cluttered environment. Moreover, if you ask me to describe my room, I look around – first here, now there – I don't shut my eyes and reflect on my internal memory. Why are we *not* surprised by our need to constantly adjust and probe if in fact we take ourselves to carry around a detailed internal model with us? The absence of surprise here is a clear indicator of the absence of epistemic commitment to the snapshot conception.

Each of these phenomena illustrates the perceptual phenomenon of presence in absence. We have a sense of the presence of strictly unseen or unattended visual detail; we experience the presence of a uniform color despite the apparent difference in color across the surface of the wall; we experience the line as complete even though it is not the case that we actually see the break as filled in. In each of these cases, as with our examples of the tomato and the cat, what we experience visually goes beyond what we see (strictly speaking).

4. **Presence as access**

Do you take yourself, when you open your eyes and look, to be aware of the whole scene before you, in sharp detail, all at once? The correct answer to this, as we have seen, ought to be: yes and no. Yes, in so far as you take yourself to have a sense of the presence of a richly detailed world. But *no* in so far as it does not seem to you as if you actually see every bit of detail. There is no such thing as seeing all the detail at once, just as there is no such thing as seeing the tomato from all sides all at once.

Phenomenologically, the world is given to perception *as available*.

To solve the problem of perceptual presence – comprising as it does a broad range of phenomena – we need to make explicit this feature of the relevant phenomenology.

We visually experience the scene before us as densely detailed without seeing all that detail, just as we visually experience the tomato as voluminous and three dimensional even though we don't see all of it. The presence of the detailed environment –of the occluded parts of the tomato, of the uniform color of the wall's surface – consists, then, not in our feeling of immediate contact with those features, but in our feeling of *access* to those bits of detail. The detail is present *now*, though absent (unseen, out of view, partially occluded, et cetera), because we *now* possess the skills needed to bring the relevant features into view.

The scene is present to me now as detailed, even though I do not now see all the detail, because I am now able – by the exercise of a repertoire of perceptual skills – to bring the detail into immediate perceptual contact. For example, I need but move my eyes, or move about, or direct my attention here or there to bring the relevant detail to focus. The detail is present because it is, as it were, *within reach*.

The basis of our feeling of access is our possession of the skills needed actually to reach out and grasp the relevant details. We are familiar, as a general rule, with the ways our sensory experience changes as we move. Moving the eyes, blinking, turning the head, moving the body – all this produces familiar kinds of sensory change. Familiarity with the ways sensory stimulation changes as we move is the ground of our perceptual access. Perceivers *know how* to gain access, to make contact, with the environment around them.

This is the key to the problem of perceptual presence: our sense of the perceptual presence of the detailed world does not consist in our representation of all the detail in consciousness *now*. Rather, it consists in our access *now* to all of the detail, and in our knowledge (itself practical in character) that we have this access. This knowledge takes the form of our comfortable mastery of the rules of sensorimotor dependence that mediate our relation to the world, to the surrounding detail, the cat, the wall. My sense of the presence of the whole cat behind the fence consists precisely in my knowledge, my implicit understanding, that by a movement of the eye or the head or the body I can bring bits of the cat into view that are now hidden. And so for the tomato: My relation to the strictly unseen portions of the tomato is mediated by familiar visual laws of sensorimotor dependence. The presence of the tomato to me as a voluminous whole consists in my knowledge of the sensory effects of my movements in relation to the tomato.

5. Intramodality and intermodality

Presence is explained in terms of access. The *modality* of presence – whether the presence is perceptual or merely *thought* (as it were), or whether it is visual, or tactile, say – is explained by the different *kinds of access* required, and by the different sorts of skills needed to secure access.⁶

⁶ Here I build on the idea presented first in O'Regan and Noë 2001. See also Noë 2002 and Noë, forthcoming.

Contrast the felt presence of peripheral detail in the visual scene with the sense of the presence of the far side of the tomato. In both cases, the feature in question is strictly unperceived, but is sensibly felt to be present. However, the features are present in different ways. In the one case, detail is grasped or obtained by the mere flick of the eye, or by turning the head to the left or right, towards the periphery. Your sense of its presence now is a confidence that *by doing those things* you can obtain the detail. In the other case, you must move around the tomato, to bring the other side into reach. To the qualitative differences in the felt presence of the features, there corresponds the different things we need to do to attain those features, and the different sensorimotor skills we need to deploy.

Of course, it is true that we have access to more than we take ourselves to experience perceptually. Not everything that is present is perceptually present. The room next door feels present, for example, but it doesn't seem as if I *see* it. It is merely *thought* to be present. That it is not felt to be present *perceptually* comes out in such facts as that I can jump up and down, turn around, turn the lights on and off, blink, and so on, and it makes no difference whatsoever to my (visual) sense of the presence of the room next door. My relationship to the room next door — however strongly I believe or know or assume or feel that it is present — is not a perceptual relation. My relation to hidden parts of the cat, or to the far side of the tomato *is* perceptual, even though I don't actually see these items. For my relation to these items is mediated by patterns of sensorimotor dependence. My relation to these items is *affected* by bodily movements.

You may wonder whether this can be quite right. First, my visual relation to the room next door is also mediated by patterns of sensorimotor dependence. Moving over to

and through the door will produce changes in my sensory stimulation. The differences in my relation to the strictly unperceived but perceptually present and to that which is simply strictly unperceived is one of degree. Second, is it really true that I am perceptually linked to the unseen bits of the tomato or cat?

To reply: In general, sensorimotor dependencies can be characterized as having two important features. The first of these, which I have been emphasizing, is that they are *movement-dependent*. The slightest movements of the body modulates my sensory relation to the object of perception. But they are also *object-dependent*. Suppose I am looking at you and someone off to the side gets up to leave. In normal circumstances I will notice the movement off to the side and turn my eyes to it. Part of what my sense of the perceptual presence of the periphery of my visual field consists in is just this fact that movements there grab my attention. (This explains our sense of the unboundedness of the visual field.)⁷

It is true, then, then there are movements of my body that will bring me into visual contact with the room next door, but it is generally not the case that movements or changes in the room next door will produce sensory changes in me. The sensorimotor contingencies mediating me and next door are not *object-dependent*.

Now consider the tomato in front of me. True, I am unlikely to notice an ant crawling on the far side of the tomato. After all, the ant is hidden by the tomato. On the

⁷ This distinction between movement- and object-dependence corresponds to but differs in significant ways from the distinction between bodilienss and grabbiness first presented in O'Regan and Noë 2001 (and developed in a series of unpublished papers by O'Regan, Myin, and Noë).

other hand, the slightest movement of the tomato will grab my attention and may bring me into contact with now occluded parts of the tomato.

For our relation to an object to be perceptual, it must be mediated by patterns of sensorimotor dependence which are both movement-dependent and object-dependent.

This account of perceptual presence provides the basic elements, I think, for a more full-blown account of perceptual content. For the question, in what does my sense of the unseen part of the tomato consist is equivalent to the question, in what does my experience of the tomato's shape consist.

Consider a simple case of shape perception:

A cube has six sides; there are twelve edges and eight vertices. You can never see more than three sides from a single point of view. As you move with respect to a cube, its aspect changes dramatically. Sides come into view while others disappear. Any movement determines a set of changes in perceived aspect; any set of changes in perceived aspects determines equivalence classes of possible movements.

When you see the cube from a particular vantage point, you encounter its aspect from that vantage point. When you experience an object as cubical on the basis merely of its aspect, you do so because you bring to bear, in this experience, your sensorimotor knowledge of the relation between changes in cube aspects and actual and possible movement. To experience the figure as a cube, on the basis of how it looks, is to understand *how* its look changes as you move.

Similar points, *mutatis mutandis* go for different shapes and different perceptible properties.

This idea was anticipated by Merleau-Ponty [1945/1962: 82], who wrote:

I know that objects have several facets because I could make a tour of inspection of them, and in that sense I am conscious of the world through the medium of my body.

And a similar thought has been articulated by Poincaré:

“To localize an object simply means to represent to oneself the movements that would be necessary to reach it.”

Experiencing the ovoid character of the tomato depends on one’s implicitly grasping the sensory effects of movement in relation to the tomato. To experience the tomato as voluminously ovoid is to experience it as providing the possibility of a range of movement-induced sensory changes. The possible pattern differs depending on the experienced shape of the perceived object.

6. Enactive externalism and virtual presence

With this account of perceptual presence on the table, let’s return to the question of externalism. Is experience in the head?

The content of a perceptual experience is not given all at once the way the content of a picture is given in the picture all at once. What we experience *now* goes beyond what we represent *now* in consciousness. Detail, three-dimensionality, color, are present in experience not *as represented*, but rather *as available*. In this sense, experience has this content only *as a potentiality*. The content is accessible, thanks to the perceiver’s

embedding in the world, and the perceiver's possession of the skills – sensorimotor, perhaps also conceptual – needed to gather the content.

This is the first point I want to make in support of the idea that the body and the world enter inextricably into the making of experience. There is a sense, then, in which the content of the experience is not *in the head*. But nor is it *in the world*. Experience isn't something that happens in us. It is something we do: it is a temporally extended process of skillful probing. The world makes itself available to our reach. The experience comprises mind and world. Experience has content only thanks to the established dynamics of interaction between perceiver and world.

In defense of this idea, consider the proposal that the world is present in experience *virtually*, the way information from a remote server is present on your desktop. The world is present virtually thanks to the way we are bound to it, in bodies with the right sort of networked connections. The flick of the eye, the turn of the head, the movement of the body, brings us the detail we need as we need it. The world is present virtually thanks to our online, dynamic access to it.

This metaphor invites an objection though. All that is present in your computer, *really*, is what is already downloaded. Information on the network is accessible, but it isn't really present. The illusion of presence depends only on the current state of your local machine. Isn't the same true of perception? The content of your current experience is determined by your current brain-state. Crucially, this brain-state includes all the needed information about the sensory effects of movement. Anyone with a brain state identical to yours would have the sense of presence of the same variety of features, even if their environment were radically different than yours!

This is a compelling line of objection. Here is how I think we should respond to it.

First, phenomenologically speaking, virtual presence is a kind of presence, not a kind of non-presence or illusory presence. Recall, it doesn't seem (for example) as if you can actually *see* the partially occluded bits of the cat. It only seems as if they are present in that by movements we can bring them into view. Virtual presence is all the presence we need, phenomenologically. Crucially, virtual presence can be explained given a conception of the perceiver as embodied and as situated in and coupled with an environment that affords possibilities of exploratory movement.

Second, and this is perhaps the most important idea in this paper, experiential presence is virtual *all the way in*. This is an important disanalogy with the computer case. Consider the tomato again. You see the facing side. You can't see the far side, but you have a perceptual sense of its presence thanks to your practical grasp of sensorimotor patterns mediating your relation to it. The rear side is present virtually, but the facing side is present *simpliciter*. Notice, however, that you do not, as a matter of fact, have the *whole* of the facing side of the tomato in consciousness all at once. The facing side has extent, and shape, and color, and you can't embrace all this detail in consciousness all at once, anymore than you can embrace the whole detailed scene. This is clear to careful consideration. Take a tomato out. Look at it. Yes, you have a sense that the facing side of the tomato is all there, all at once. But if you are careful you will admit that you don't actually experience even every part of its visible surface all at once. Your eyes scan across the surface and you direct your attention to this or that. Further evidence is provided by change blindness. As mentioned above, the very color of the object you are

staring at can change right before your eyes without your noticing it, so long as you are not attending to the color itself!

What this shows is that you cannot factor experience into an occurrent and a merely potential part. Pick any candidate for the occurrent factor. Now consider it. It is structured too; it has hidden facets or aspects too. It is present only in potential.

The point here is not that one can only attend to a small number of features at an instant in time, although this is widely agreed to be true [Pylyshyn; Sperling 1960]. The point rather is this: a perceptual experience doesn't analyze or break down into experiences of atomic elements, or simple features. Experience is always a field, with structure, and because of this you can never comprehend the whole field in a single act of consciousness. Something always remains present, but out of view. All you can do is serially run through features. But the moment you stop and try to make *that* feature the sole object of your consideration – *this shade of red*, for example – it exceeds your grasp. This is true even of a *Ganzfeld*. Suppose you are in a giant grey fog. Nothing visually distinguishes here from there. And yet, you are not given the greyness around you as a simple property. There's the color, which is spread out in space, but there's also the texture.

Qualities are available in experience as possibilities, as potentialities, but not as completed givens. Experience is a dynamic process of navigating the pathways of these possibilities. Experience depends on the skills needed to make one's way.

The upshot of this is that there is no basis, in phenomenology, for thinking that what is given now, to me, as present in my consciousness, is ever enough to account for the character of my current conscious experience. My phenomenal experience expands

my immediate horizons and takes me beyond myself to the world. This sounds paradoxical, but it is not. Presence in absence, I have tried to show, is a pervasive feature of our perceptual lives.

“Does experience supervene on internal states of the brain?” The correct answer to this ought to be: “maybe”. I have argued that what we experience outstrips what is represented in consciousness. This does not entail that what we experience outstrips what is represented in our brains. However, it does remove the major theoretical obstacle to entertaining the possibility that experience might supervene not on the brain, but rather on brain-animal-world systems. *It is an empirical question whether our brains can do the work needed to enable us to enact our virtual worlds.* It is a mistake – a prejudice really – to think this question has already been settled.

“Wouldn’t my neural duplicate have the same experiences as me?” Perhaps, but your neural duplicate would almost certainly be embedded in and interacting with a duplicate of your environment. What else could explain the neural identity? (This point has been discussed by Hurley [1998]; see also Hurley and Noë 2003.) “But wouldn’t my neural duplicate have the same experiences as me *whatever differences there might be between its environment and mine?* This is the force of the claim that experience supervenes on the brain.”

Again, this is an empirical question. If, as I have suggested, experience may make use of online, dynamic access to the world, then it comes down to the question whether brains can suffice without active reference to the world itself. There is some preliminary

evidence that they cannot. This is one way to interpret the change blindness findings. We don't have a handy world-model in the head.⁸

A reasonable bet, then, at this point, is that *some* experience, or some features of some experiences, supervene on the brain, but that full-blown, mature human experience does not. This is supported by such facts as that, as of now, neuroscientists have been able to produce only relatively primitive experiences by direct stimulation of the cortex [Koch in press]. More importantly, it is just not clear, given the virtual character of perceptual content, why an internal representation would be any better than access to the world itself. This harkens back to Wittgenstein's point that anything a picture in the head could do could be done by a picture held in the hand. We go a step further: why need a picture at all? The world is right there, after all. We are *in the world*. We have the skills needed to enact our perceptual experience.⁹

The computer comparison may help again: my desktop may have enough memory to allow me to download the entire online version of *The NY Times* at one go. No need to make use of dynamic, networked access. But maybe I don't know how to do that, or

⁸ Hurley and Noë [2003] further argue that a range of phenomena of neural plasticity (what they call *cortical deference*) can best be explained by supposing that neural activity participates in dynamic sensorimotor loopings that involve distal objects, the animal's sensory periphery, brain and nervous system, and also the animal's skills of movement.

⁹ This idea that we can forego representation by making use of the presence of the world was perhaps first suggested by Minsky [1985], and, independently, by Dreyfus [1972/1999]. In recent discussion, it was hit upon by Kevin O'Regan [1992] (who suggests that the let the world serve as an external memory store) and Rod Brooks [1991] (who suggests that the world can serve as its own best model).

maybe my desktop isn't big enough to allow such a big download, or maybe doing that would interfere with computer's performance in other ways. Or perhaps, although I *could* do that, I choose not to because doing so would mean I would miss out on the constant updates that are provided online. Clark and Chalmers' point was not that the mind *must* extend beyond the limits of the head, but that there was no theoretical obstacle to thinking that it *could* do so, for some range of tasks.

It always seemed that there were obstacles to thinking that consciousness could so extend beyond the limits of the skull. "Gosh darn it, experience just feels like it's in the head." But this is bad phenomenology, I have argued, and it is probably bad science. As we have seen, detail may be present in consciousness only virtually. We thus open up the possibility of an account of (for example) the perceptual experience of detail that is consistent with its not being the case that that detail is represented at once in the head. Although of course it *could* be. The upshot is not that experience is without the head, but that it might be. The world is safe for an externalism that allows that we enact perceptual content by the exercise of sensorimotor skills over time.

Figure 1




Figure 2



Figure 1 Caption

Due to Kanisza

Figure 2 Caption

Filling in at the blind spot

Shut your right eye and fixate the cross with your left eye. Adjust the distance of the book from your eye. At one point the gap in the line on the right falls within the blind spot. What do you experience when that happens?

Works Cited

- Brooks, R. A. (1991) Intelligence without reason. Proceedings of the 1991 International Joint Conference on Artificial Intelligence: 569-595.
- Clark, A. 1999
- Clark, A. and Chalmers, D. J. [1998] The extended mind. *Analysis* 58, 1: 7-19.
- Clark, T. [1965] Seeing surfaces and physical objects. In M. Black (ed.) *Philosophy in America*. Ithaca: Cornell University Press.
- Dennett, D. C. [1987] Error, evolution and intentionality in *The Intentional Stance*. Cambridge: MIT.
- Dennett, D. C. [1991] *Consciousness and Explained*. New York: Little Brown.
- Dennett, D. C. [2001] Surprise, surprise. [Commentary on O'Regan and Noë]. *Behavioral and Brain Sciences* 24, 5.
- Dennett, D. C. [2002] How could I be wrong? How wrong could I be? *Journal of Consciousness Studies* 9, 5-6: 13-16.
- Dreyfus, H. L. [1972/1999] *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, MA: The MIT Press.
- Dretske, F. [1995] *Naturalizing the Mind*. Cambridge, MA: The MIT Press.
- Durgin, F. H., Tripathy, S. P., & Levi, D. M. (1995). On the filling in of the visual blind spot: Some rules of thumb. *Perception*, 24(7), 827-840.
- Hurley, S. L. [1998] *Consciousness in Action*. Cambridge, MA: Harvard.
- Hurley, S. and Noë, A. [2003] Neural plasticity and consciousness.
- Kelly, S. D. [2001]. The non-conceptual content of perceptual experience: Situation-dependence and fineness of grain *Philosophy and Phenomenological Research* lxii, no. 3: 601-608.
- Koch, C. [in press] *The Question for Consciousness*.
- Koenderink, J. J. [1984]. The internal representation of solid shape and visual exploration. In L. Spillmann & B. R. Wooten (Eds.), *Sensory experience, adaptation, and perception. Festschrift for Ivo Kohler* (pp. 123-142). Hillsdale, New Jersey: Lawrence Erlbaum.
- Merleau-Ponty, M. (1945/1962) *The phenomenology of perception*. Trans. Colin Smith. London: Routledge Press, 1962. [French original: 1945.]
- Minsky, M. (1985) *The Society of Mind*. New York: Simon & Schuster.
- Noë, A. [2002] Is the visual world a grand illusion. *Journal of Consciousness Studies*. 9:5/6: 1-12.
- Noë, A. [2002a] *Is the Visual World a Grand Illusion?* Thorverton: Imprint Academic.
- Noë, A. Pessoa, L., Thompson, E. [2000]. Beyond the grand illusion: what change blindness really teaches us about vision.

O'Regan, J.K. [1992] On the real mysteries of visual perception.

O'Regan, J. K. [in press] Change blindness in *Encyclopedia of cognitive science*, Nature Publishing Group.

O'Regan, J. K. and Noë, A. (2001) A sensorimotor approach to vision and visual consciousness. *Behavioral and Brain Sciences* 24/5: 883-975.

Merleau-Ponty, M. (1945/1962) *The phenomenology of perception*. Trans. Colin Smith. London: Routledge Press, 1962. [French original: 1945.]

Palmer, S. (1999) *Vision: From Photons to Phenomenology*. Cambridge, MA: The MIT Press.

Peacocke, C. [1983] *Sense and Content*. Oxford: Oxford University Press.

Pessoa, L, Thompson, E. and Noë, A. [1998] Finding out about filling in: a guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences* 21/6: 723-802.

Poincaré

Ramachandra, V. S. and Gregory, R. L. [1991] Perceptual filling in of artificially induced scotomas in human vision. *Nature* 350, 6320: 699- 702.

Simon, D. S. [2000] Current approaches to change blindness. *Visual Cognition* 7: 1-15.

Simons, D. S. and Levin, D. T. Change blindness. *Trends in Cognitive Science* 1/7: 261-267.

Sperling, G. [1960] The information available in brief visual presentations. *Psychological Monographs*, 74.

Strawson, P. F. (1979) Perception and its objects. In G. F. MacDonald (ed.) *Perception and Identity: Essays Presented to A. J. Ayer with His Replies*. Ithaca: Cornell University Press. Reprinted in Noë, A. and Thompson, E. (eds) *Vision and Mind: Selected Readings in the Philosophy of Perception*. Cambridge, MA: The MIT Press.

Wilson, R. A. [1997] *Cartesian Psychology and Physical Minds: Individualism and the Science of the Mind*. Cambridge: Cambridge University Press.