

Experience is not something we feel but something we do: a principled way of explaining sensory phenomenology, with Change Blindness and other empirical consequences.

J. Kevin O'Regan
Laboratoire de Psychologie Expérimentale,
Centre National de Recherche Scientifique, Paris, France

Talk given at Bressanone on 24 Jan 2001

ABSTRACT

Any theory of experience which postulates that brain mechanisms generate "raw feel" encounters the impassable "explanatory gap" separating physics from phenomenology.

A way round the problem is to postulate that experience is not something we feel, but something we do: a kind of give-and-take with the environment, analogous to the "feel" of driving a car. One consequence of such a "sensorimotor" theory of experience is that it provides a way of explaining the differences between seeing, hearing, touch, etc., which is more principled and has more explanatory power than Müller's notion of "specific nerve energy" or its modern counterpart, the notion of sensory pathways or cortical areas. The feasibility of sensory substitution is an empirically verifiable implication of this approach.

As applied to visual perception, a consequence of the sensorimotor approach is the idea that seeing does not consist in the creation of a "re-"presentation of the world inside the brain, but rather in knowledge that the outside world is immediately accessible through a flick of the eye or of attention, like an

"outside memory". The world-as-an-outside-memory idea has empirically verifiable consequences in the phenomenon of Change Blindness, among others.

If I show you a picture and make a change in it, like this, you immediately see the change. there's no problem, it pops out.

Change with no transient

But if I make the change at the same time as I put a slight flicker on the screen, like this, then in many cases you miss the change.

Flicker example: woman in kayak

Here's another example.

Couple at dinner

It's quite interesting that sometimes you can be looking directly at the change and still not see it. So here if I tell you to look at the man's nose, you'll be within a few pixels of the change, and yet often you won't see it: it's the bar in the background going up and down.

Here's another example

Notre Dame

Some changes are easier to see than others. But even a change which occupies a very large portion of the visual field may not be noticed if it's not part of what the picture would be said to be about. Here for example,

canalbridge

you would not normally say that the picture is about the reflection in the lake. But as soon as I tell you that, you see the change. Whereas here,

Milk glass

you probably consider the picture to be about a glass of milk, so its easy for you to see that that's what's changing, even though it's much smaller than the reflection in the lake was.

Change Blindness

- **Flicker**
 - Rensink, O'Regan & Clark, 1997; 1999
- **Eye saccades**
 - Currie, McConkie, Carlson-Radvansky & Irwin, 1995; McConkie & Currie, 1996
- **Blinks**
 - O'Regan, Deubel, Clark, Rensink, 1999
- **Film cuts, real life**
 - Levin & Simons, 1997
- **“Mudsplashes”**
 - O'Regan, Rensink & Clark (Nature, 1999)

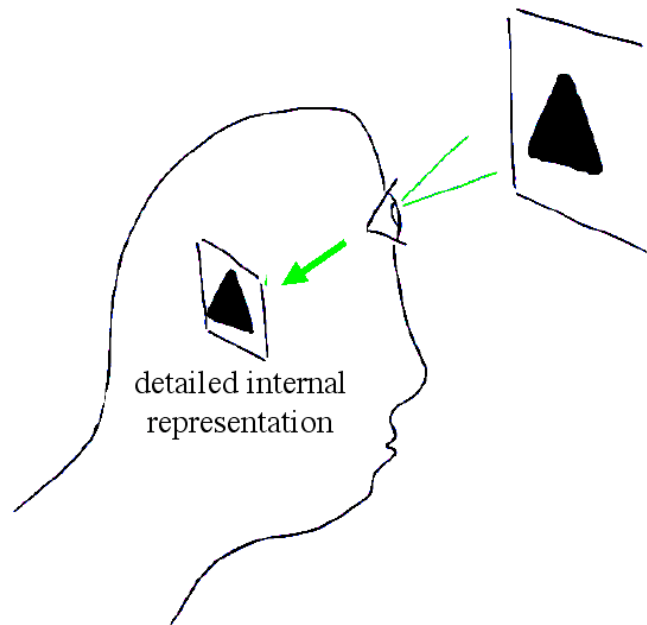
This phenomenon has been called change blindness, and has attracted quite a lot of attention over the last few years. There are a number of variations of the phenomenon. I've shown you the flicker paradigm, but the phenomenon has also been obtained with eye saccades, blinks, film cuts, and even in real-life situations.

A particularly interesting variant of the CB paradigm is the mudsplash paradigm.

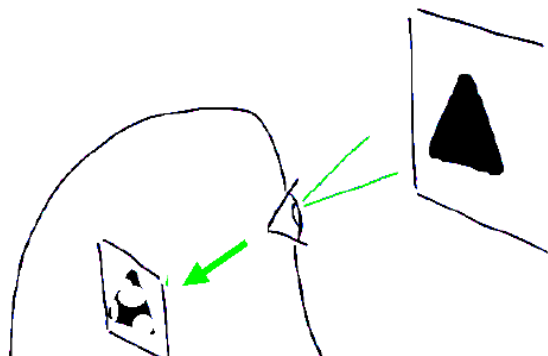
mudsplash example


The reason this is interesting is that the mudsplashes are positioned in such a way so as to not cover the change location. That way it can't be argued that the reason you miss the change is that it is somehow masked or wiped out by any kind of superposition with flicker.

It's clear that there is something very shocking about all these experiments: they seem to suggest that our internal representations of the outside world, instead of being very detailed and rich...




are actually rather sparse...So what's going on here? How can we have the impression of richness in the world if there is no richness in the head?






There are some other examples of cases where we have the impression of richness, but where, in another sense, we know that actually the information available is rather sparse. One such example is the case of the blind spot. There is a place at the back of the eye where the optic nerve and the blood vessels that irrigate the eye come into the eyeball. At that point there can be no photoreceptors, meaning that there is a position in the visual field of each eye where we are effectively blind.



You can convince yourself that this blind spot really exists, by holding an orange at arm's length and closing one eye. You'll find that there's a position where the whole orange suddenly disappears, engulfed in the blind spot. So here is another example of a case where the internal

information about the visual field is lacking a big chunk, and yet we have the impression that everything out there is perfect.

Another illustration concerns color vision. It turns out that whereas in central vision we have good color vision, as close as a couple of degrees into peripheral vision, color processing is exceedingly poor. You can appreciate this if you do the following experiment

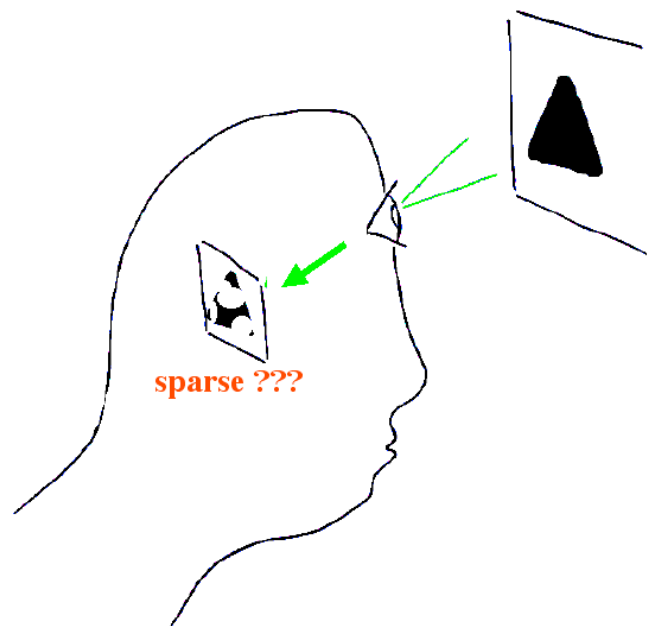


You ask a friend to stand in front of you, and you look directly at her nose, without moving your eyes. She takes a bunch of colored pencils and brings them in slowly from the side. You try to say what colors the pencils are, and in what order they are arranged. You'll find that you're quite incapable of doing this task, even if she holds the pencils as close as right next to her ear. You may be able to see a red pencil, but you'll certainly miss most of the colors, and you definitely will get the order wrong.

Again: here we have another kind of evidence that certain aspects of vision provide very poor information. Yet we see the world as perfectly

colorful, and we see no fading of colors in peripheral vision, and no mixing up of their orders.

I've shown you several examples of where we can prove that the information available in the brain about the visual field is very sparse, and yet we have the impression of a perfect visual world. How can we have the impression of richness when the internal representation is sparse?



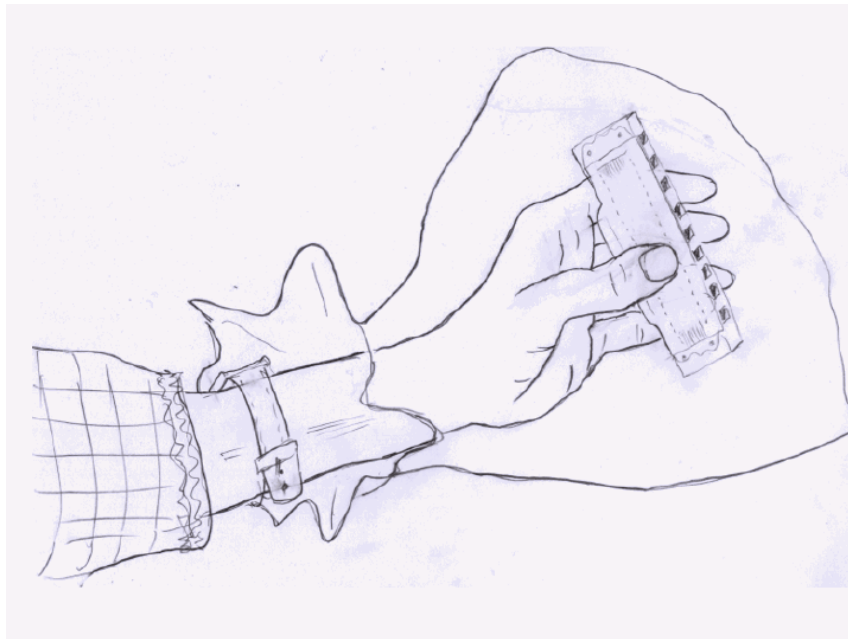
I suggest that a possible solution is the idea of what I call the World as Outside Memory. The idea is that to get the impression of richness, there's actually no need for the richness to be in the head. What has to be in the head is merely algorithms or recipes for getting **at** the information in the world.

Such algorithms we have, in the form of movements of the eyes or shifts of attention. If we're interested in some detail of the visual scene, we simply need to move our eyes or our attention to that detail, and it is immediately

available.

Instead of storing all the information about the outside world in the brain, we use the outside world as an external memory storage. Thus, we get the impression that we're seeing **everything** there is to see in the visual field, because if we so much as faintly wonder whether we're actually seeing something, we turn our eye (and our attention) to that thing, and it becomes available for processing.

Perhaps you've played the children's game in which someone puts a household object like a cork, or a potato, or a harmonica into a bag, and you put your hand into the bag and try to figure out what the object is.



At first you feel this or that texture on the tips of your fingers. You have no idea what the object is. But suddenly you have a kind of "Aha" experience. Suddenly you feel you're no longer touching bits of texture on the ends of your fingers, but you're holding a whole object: it's a harmonica. And it's ALL there at once, even though you're in fact only touching a few parts of

it. It's not just that you **know** it's a whole harmonica, you actually **feel** it's a whole harmonica

The reason you have the feeling of touching the **whole** harmonica is that you KNOW that IF you were to move your fingers this way, then you would get THIS feeling, and if you moved them that way, you would get THAT feeling. You feel familiar with all the things you can do with your fingers at this moment. It is the knowledge of feeling at home with the possible things you can do, and with the resulting changes in the feelings on your fingertips, that gives you the impression of having the whole harmonica in your hand.

Why we think we see everything

- Seeing is having access
- Refrigerator light analogy (N. Thomas)




Extended to the domain of vision, this analogy suggests how it could be possible to have the distinct feeling of seeing whole objects and whole scenes, even though only a minute part of the scene is actually available for visual processing at any moment. The impression of seeing **everything there is to see** in the visual field is therefore a sort of illusion...

It's a bit like the light in the refrigerator:


The light always seems to be on: you open the door, the light is on. You close it. You quickly open the door again to check: yup, the light is still on. You have the illusion that the light is on all the time, but you have to keep checking to convince yourself.

Let me go on to show you some of the consequences of this way of thinking for scene perception.




This picture shows the path that the eye of one observer took while searching for large changes which occurred every time he blinked.

Here is an example of a scan path for another picture.






This scan path corresponds to a few seconds of exploration, and you would think that since the observer was actively looking for a change, he would look around in the picture in a rather systematic way, covering all the picture elements. But if you look at what the observer did over the next few tens of seconds, this is what you find.



It looks like the observer is just going round and round in circles. In fact this kind of behavior is typical of what people do when they look at pictures. Only a fairly limited number of positions are directly fixated by the eyes, and they are fixated repetitively. Why is this?

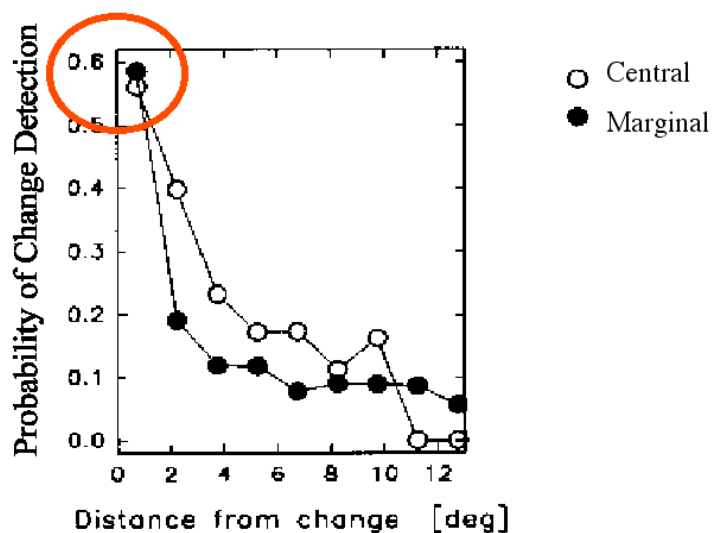
Under the "World as outside memory" point of view I'm sketching here, this can be explained. It could be that seeing a picture is not: accumulating information into an internal representation, but rather: checking that you have access to the

things that the picture is about. If you think the picture is about a couple having dinner, then seeing the picture involves making sure that those things that you think the picture is about, are really there. The eye will therefore go round and round checking.



Let me go back to the experiment where we measured eye movements while people were looking for changes in pictures.

O'Regan, Deubel, Clark & Rensink, 2000



We looked at the probability of detecting the change as a function of the position of the eye. We found, as you might expect, that the further your eye was from the change location, the smaller the probability of detecting the change was. You can see this from the fact that the graphs drop lower and lower as we go to the right, corresponding to greater eccentricities. Don't worry about the fact there are two curves, these just correspond to two different kinds of changes we used.

But one very surprising thing is visible in this graph. The probability of detecting the change when you are looking directly at the change, that is the leftmost point of the graph, is less than 60%.


That is, in more than 40% of the cases when the eye was looking directly at the change, it is not seen!

This fact is coherent with the approach I've been putting forward. According to this, when something falls on your retina, or when your visual system processes something, that does not necessarily mean you see it. Seeing only occurs when you are currently exercising your mastery of the sensorimotor contingencies associated with that thing, sort of "manipulating" it with your eyes. When I look at an object, I can be conscious of any number of its aspects: its color, its identity, its background, its position, etc. I would say that only the aspect that I'm currently checking on is actually being seen. So other aspects, even if they are being directly looked at, will not be seen.






Ambiguous figures and figure-ground competition actually provide examples illustrating this. Here, you can be fixating on the white nose and not see the black nose, even though it is in the same location.



Haines, 1991. Ames Res. Center, NASA

A similar rather shocking finding was reported by Haines, at the NASA Ames Research Center in California. He had commercial airline pilots land a 727 in a flight simulator, using a heads-up display of certain instruments on the windscreen. On certain landing approaches, Haines suddenly

superimposed a stationary small aircraft right in the middle of the runway. He expected pilots to immediately abort their landing approach. However 2 out of 8 pilots simply blithely landed through the obstructing airplane. When shown a video of what they had done, the pilots were shocked and incredulous and noted that they should perhaps resign from commercial flying.



Here is yet another example: you can stare at this for minutes and still think it says "The illusion of seeing". But actually it doesn't; it says: The illusion of of seeing.

Slow change demos

Here is another example of where you can be looking directly at the change and not see it. This picture is changing. Except it's changing very slowly. The change is rather large: see if you can find it. Dan Simons at Harvard has also been doing experiments with slow changes like this.

The point is, seeing is mentally manipulating some aspect of the scene. If no attention-grabbing visual transient brings your eye or your attention

on to some area of the picture, you won't see it at all, let alone see it change.

- **Inattention blindness**

- Neisser
- Mack & Rock
- D. Simons

Studies like these are part of a growing literature on what's called inattention blindness: Ulrich Neisser was one of the first to look at this, but Mack & Rock have just published a book on the question. Dan Simons has recently done some other beautiful experiments showing your eye can be very close to some totally obvious thing in a picture and yet not see it.

THE RICHNESS OF VISION

**SEEING IS: knowing you have access
to visual information**

In summary up to now, I've suggested the possibility that the feeling we get of seeing everything in the visual field doesn't require us having an internal representation of everything. It suffices to have immediate access to the information in the external world, which acts somewhat like an external memory store.

PROBLEM:

If seeing is like memory, why doesn't it feel like memory?

MUST EXPLAIN

**ONGOINGNESS CONTINUOUSNESS
PRESENCE REALNESS**

But you might object that there's still a problem. If seeing is supposed to be like memory (be it an external memory), because it is so immediately accessible, why doesn't it have the same phenomenology as memory. When I remember my grandmother, I can immediately conjure up information about any aspect of her I wish. So this aspect is immediately accessible. But I don't have the same feeling of presence or realness as I

do when I really see her. When I really see her, I have the feeling of an ongoing, continuous presence which is lacking in the case of memory. This needs to be explained.

Take again the example of the refrigerator light. By opening the refrigerator door I can immediately check whether the light is on, and I get the impression that it is indeed continuously on. But I somehow don't feel it is continuously on in that same, acute way that I feel the presence of an object that I'm seeing directly. Why is that?

I think the answer is to do with two things. I call one "bodiliness", and the other "grabiness".

The illusion of seeing continuously

- **bodiliness**
 - tight link to body motions
- **grabiness**
 - transients
 - incontroversibly grab attention



Bodiliness is the fact that in vision, the things you do in order to get information are very closely linked to minute and even unconscious bodily actions: The slightest twitch of an eye muscle allows you to change from one point of the scene to another. A small head movement or body movement modifies what you see.

By virtue of this bodiliness, the outside world is intimately linked to you, almost as though it was

part of your own body. I suggest that this makes seeing more real, more felt and ongoing than the memory of your grandmother, which doesn't budge when you move around.

Now let me talk about grabbiness.

We all know that if there's a sudden flicker in the visual field, we can't help but immediately look at it. This is because there exist mechanisms in the first stages of the visual system designed to detect fast transitions in local luminance, and which incontrovertibly grabs your attention. Movement detectors are examples of such transient detectors.

It could be that this grabbiness of sudden events constitutes a second factor which contributes to the feeling of continual presence, and ongoingness of visual perception. Grabbiness makes it seem like we have continual "tabs" on everything that's going on in the visual field, and gives us the illusion of seeing things continuously, because if anything should change we're immediately informed.

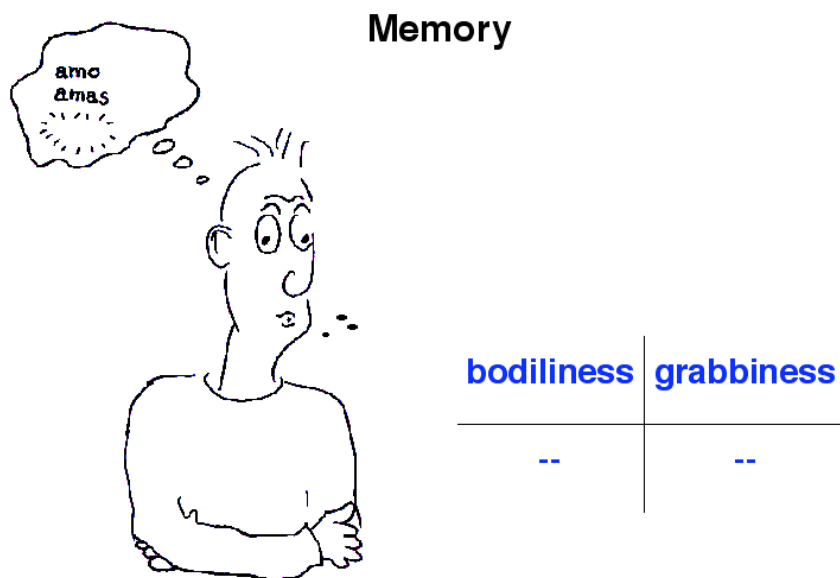
Memory and Vision

- **Immediate access to information**
 - memory: in brain
 - vision: in world
- **Difference in quality**
 - bodiliness/grabbiness

Let me summarize again: I've been suggesting that, contrary to our intuitions, memory and visual experience might actually be one and the same kind of thing: both involve knowing ways of getting at information: in one case the information is in the brain, in the other case it's in the outside world.

In the case of memory of latin verbs, for example, I know I can recover the conjugation of a particular verb by attending to that verb. Analogously, in the case of seeing, I know I can recover information about some object in the scene by paying attention to it.

The explanation for the difference in the quality of the experience we get from our memory for latin verbs and the experience we get from seeing, could be due to the amount of bodiliness and grabbiness that was involved.



Memory for latin verbs has no bodiliness and no grabbiness: no bodiliness because my body motions dont affect the availability of latin verbs in my memory; and no grabbiness because

changes in my memory dont attract my attention -- for example, if a word dropped out of my memory overnight, no bell whistles in my mind to tell me.

Seeing

bodiliness	grabbiness
+++	+++



On the other hand seeing involves a lot of bodiliness and a lot of grabbiness: the slightest twitch of an eye muscle changes my visual input, and any change of the input attracts my attention.

	bodiliness	grabbiness
memory	--	--
wealthiness	+	-
driving	++	-
seeing	+++	+++

"raw feel"

↓

So we see that memory and seeing are situated at two ends of a continuum of bodiliness and

grabbiness. As expected therefore, memory has nothing like what one would want to call a "raw feel". Seeing on the other hand has lots of raw feel.

It's interesting now to reflect on whether there are intermediate cases.

Consider for example the experience of being wealthy. Like seeing, wealthiness is a form of knowledge about accessibility. It has more bodiliness than memory of latin verbs however because wealthiness consists in the expectations that when I do certain things with my body, I expect certain results. For example, I ask the bank manager to give me the money, and he does. However the things you do to get the results are things which are not very intimately linked to your slightest body motion. Thus I dont want to give wealthiness full marks on bodiliness.

Wealthiness on the other hand is not grabby: unless I have a particularly conscientious bank manager or stock broker, when my account is depleting or the market is down, nothing advises me of this fact. No bell rings in my mind. But we see that on the scale of raw feel, because it has a bit of bodiliness, wealthiness gets better grades than latin verbs. This explains why people do say, sometimes: I **feel** wealthy.

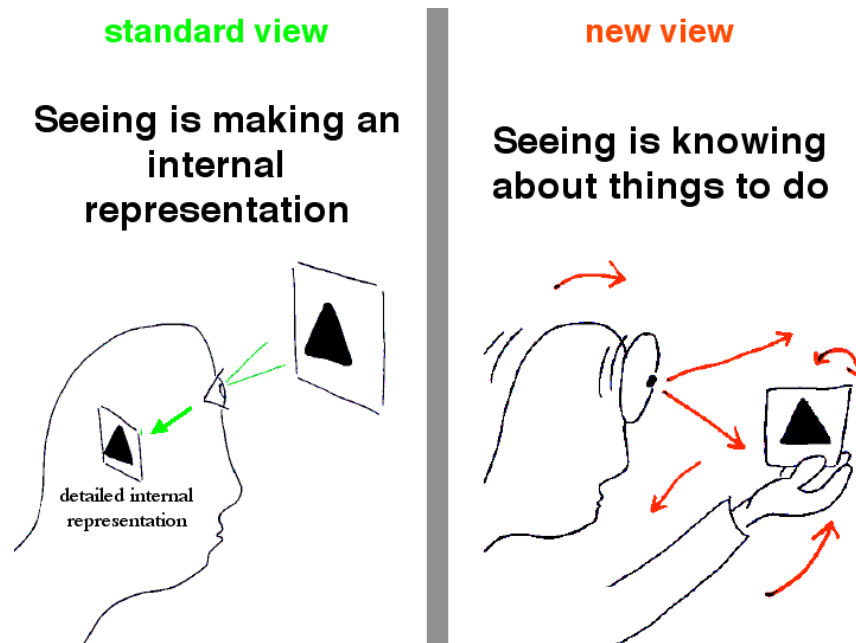
Consider now driving a car. Here bodiliness is a bit more intimate than for wealthiness: the slightest twitch of my foot on the accelerator or my hand on the steering wheel have effects on the car and thereby on the sensory input. But what about grabbiness. It's true that as I drive my attention can be automatically attracted to various things happening, but it's not ever attracted to the driving experience itself. Thus I would say that driving has no grabbiness.

Comparing with wealthiness and memory of latin verbs we see that because of the extra bodiliness there might be a bit more "feel" to driving. I think again that this corresponds to people's intuitions: there is such a thing as a "feel" to driving a car.

I must say that I find very promising this use of the concepts of bodiliness and grabiness to make a classification of the amount of "feel" that a mental state or activity possesses. I think there may be a simple way to extend this approach to include emotions and pain.


Some philosophical remarks

Let me now stand back and make some remarks about the way of thinking that I've sketched here.




Under the standard view, seeing consists in making an internal representation of the outside world. Under the new view, seeing consists in knowing different things you can do, and knowing the changes that these things will produce in your sensory input.

Though unnatural at first sight, the new view has an interesting advantage.



In neuroscience today, one of the problems people are grappling with is to try to understand how a physical entity like a brain can give rise to something like the feeling of seeing, which is patently not physical.

Some as yet unknown mysterious, possibly even nonphysical mechanism has to be postulated to instill experience into the brain.




But under the new view, the problem disappears, because experience is not in the brain at all.


It's in the doing of the exploration, and in the knowledge of the things that will change as you explore.

But under the approach I'm suggesting, seeing is actually **not** something ongoing. Thus, the problem of finding a mechanism to generate it in the brain disappears. Instead of the role of the brain being to generate the **experience** of seeing, the role of the brain simply becomes that of generating the **exploratory activity** which underlies the seeing, and that of **holding the knowledge** of current possibilities for action that underlies seeing.

Still, there is an objection that may be disturbing you. You could say, ok, seeing is a thing we do... I see the red cup when I'm going about checking I have access to it by the flicks of my eye movements. OK. But what about when I finally **am** actually looking directly at the red cup. I now have red stimulation on my retina. Now surely there must be something that takes that takes red stimulation on my retina and causes me to experience the redness. We seem to be back to the situation of having to explain how brain activation can engender experience.

But it seems to me that it's possible to escape from the difficulty even for the raw sensation of redness.






“Red” is the way red things change the light

(Broackes, 1992)

Consider looking at a piece of red paper. Depending on whether you turn the paper so that it's yellowish sunlight or bluish skylight, or reddish incandescent light that is reflected off the paper, the spectrum of light being sensed by the eye is quite different. I suggest you see the paper as red when the laws that are obeyed by the changes in incoming spectrum are typical of redness. Thus, red is not a pattern of excitation caused by incoming light, but knowledge about the laws that the excitation obeys when you move the paper around.

Non-uniformity of color sampling



Another fact about red has to do with the way the eye samples color. At the center of the retina, color information is readily available, being sampled by retinal cones sensitive to long,

medium and short-wavelength light, symbolized by colored dots in the slide. But the density of the cones falls off quite rapidly, so that the nature of the neural stimulation that arises from looking straight at a red surface is quite different from that obtained by looking at the surface in peripheral vision. here, there are many more rod photoreceptors, symbolized by black dots, not sensitive to different colors. I suggest that the quality of red is NOT just the particular combination of long, medium and short wavelength stimulation, but also the way in which the stimulation changes as you move your eye on and off the red object.

Seeing Red

knowing that sensorimotor contingencies typical of red are currently being obeyed.

Like the experience of seeing everything, the experience of seeing red then is also a kind of knowledge: knowledge that the appropriate contingencies between sensory input and motor actions are currently applicable.

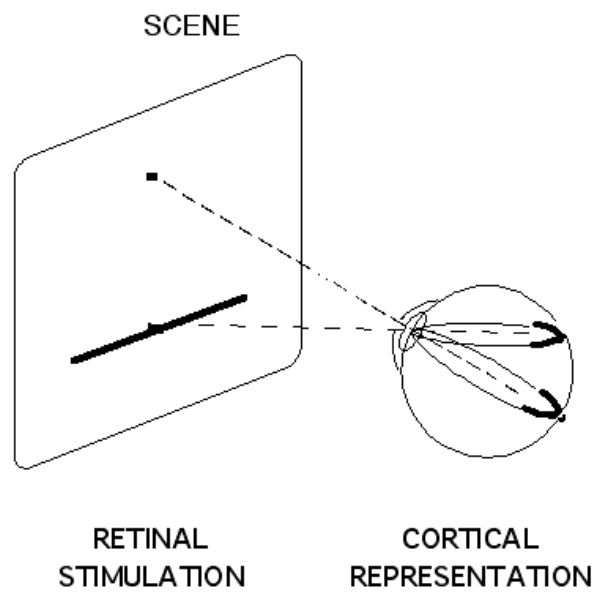
The notion of sensorimotor contingency can be generalized to cover not just the sensation of red, but, I suspect, all aspects of vision, both general and particular. For example the fact that the retinal image essentially goes blank when we


blink, or shifts in lawful ways when we move our eyes, or has an expanding or contracting flow field when we move our heads backwards and forwards, are facts about vision in general.

Sensorimotor contingencies of SEEING

blink:	big change
move forward:	expanding flow
turn sideways:	shifting flow
cover ears:	nothing
cover eyes:	big change


A fact that's more specific, characteristic of straight lines, for example, is the fact that when you move your eyes along them, nothing much happens to sensory input, whereas when you move your eyes across them, sensory input changes more drastically.





In sum then, it could be that the experience of seeing derives from being familiar (in the sense one is familiar with practical know-how) with a wide variety of sensorimotor contingencies related to the way the visual apparatus samples the environment.

**Seeing is:
Knowing the sensorimotor contingencies**



It could be that we feel we are seeing at this moment, when we know (in a practical way) that all these contingencies are currently applicable. The experience of seeing would then not be generated by activation of a brain mechanism. It would be constituted by knowledge that if you do certain things, certain things will happen to sensory input.

Müller's specific nerve energy and

sensory substitution

Now I'd like to discuss another consequence of what I've been saying, which concerns the quality of different sensory modalities.

Quality of sensory modalities

- **Old view:**
 - Müller's specific nerve energy
 - Cortical maps, neural pathways
- **New view:**
 - Different things you do

Everyone agrees that the qualitative nature of experience in one sensory modality is quite different from the experience in another modality: hearing feels quite different from seeing, which feels different from taste and touch... The explanation for this has remained problematical ever since Johannes Müller at the end of the last century had suggested that different neural pathways might have what he called different "nerve energies".

On the other hand a natural and principled approach to the problem might be available if we adopt the view that seeing is a kind of knowledge about what happens when you do certain things. Driving a car feels different from driving a truck or riding a bicycle, because it involves doing different things. Similarly, seeing feels different from hearing, tasting, touching, because it too involves doing different things.

Examples of sensorimotor contingencies

	SEEING	HEARING
blink:	big change	no change
move forward:	expanding flow	increasing amplitude
turn sideways:	shifting flow	asynchrony
cover ears:	nothing	big change
cover eyes:	big change	nothing

For example, we know we're seeing when we know that: if we blink, the sensory input changes drastically; if we move forward, there is an expanding flow field; if we move our eyes, there is a translating flow field; if we block our eyes with our hands the visual field is obscured; on the other hand, if we block our ears with our hands, nothing much happens.

On the other hand we know we're hearing if: when we blink or move our eyes nothing much happens; if we move forward the intensity of incoming stimulation obeys an inverse square law; if we move our head the asynchrony and spectrum of the input changes in certain characteristic ways; if we block our eyes with our hands, nothing much happens; if we block our ears with our hands, the intensity changes in a certain way.

Conclusion

Sensorimotor theory of visual experience

(O'Regan & Noë, Behavioral and Brain Sciences, under review)

<http://nivea.psvcho.univ-paris5.fr>

**Seeing is having immediate access
Knowing sensorimotor contingencies**

- **illusion of seeing everything**
- **illusion of continuousness**
- **bodiliness, grabbiness**

In conclusion then, the approach consisting of taking the rather counternintuitive stance of saying that seeing is not something ongoing, but rather a form of knowledge, like memory, has at first been hard to swallow.

It led us to postulate that the impression we have of seeing everything in the visual field is actually a sort of illusion, generated by the immediate availability, by a mere flick of the eye or of attention, of visual information.

It also led us to postulate that the impression of continuousness or ongoingness of vision was also an illusion. I suggested that the concepts of bodiliness and grabbiness might account for why we have this illusion.

Now this approach seems at first rather hard to swallow, it has some very interesting advantages.

Sensorimotor theory of visual experience
(O'Regan & Noë, Behavioral and Brain Sciences, under review)
<http://nivea.psychu.univ-paris5.fr>

Advantages

- **experience not generated by brain**
 - in exercising knowledge of sensorimotor

contingencies

- **explains quality of sensory modalities**
- **explains ineffability**
- **explains "feltness" of mental states**

The approach puts experience in the doing of exploration, rather than in the brain. That way we escape from the problem of having to find a brain mechanism that generates experience.

The approach explains in a principled way the differences in sensory qualities of the different sense modalities. I haven't had time to mention how it accounts for the ineffability of sensations.

Finally the approach provides a neat classification of the phenomenology of certain mental states like memory, wealthiness, and sensation.

If you've been interested in this way of thinking about vision, you may like to look at the paper we have under review in BBS, where with Alva Noë we've more carefully laid out the relation of this work to the problem of consciousness and qualia, and where we've tried to show how the view allows us to bring conveniently together strands of empirical research not only from Change Blindness, but also from a variety of other domains: among them sensory substitution, synesthesia, and sensory adaptation.