# UC Merced

## Proceedings of the Annual Meeting of the Cognitive Science Society

**Title**

A New Theory of the Representational Base of Consciousness

**Permalink**

**Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 26(26)

**ISSN**

1069-7977

**Authors**

Brook, Andrew
Raymont, Paul

**Publication Date**

2004

Peer reviewed

# A New Theory of the Representational Base of Consciousness

**Andrew Brook** {abrook@ccs.carleton.ca}

**Paul Raymont** {praymont@ccs.carleton.ca}

Cognitive Science
Carleton University
Ottawa, ON Canada

## Abstract

Though we take mainly a philosophical approach, we hope that the results of our work will be useful to researchers on consciousness who take other approaches. Everyone agrees, no matter what their point of view on consciousness, that consciousness has a representational base. However, there have been relatively few well-worked-out attempts to say what this base might be like. The two best developed are perhaps the higher-order thought (HOT) and the transparency approaches. Both are lacking. Starting from the notion of a self-presenting representation, we develop an alternative view. In our view, a representation, a completely normal representation, is the representational base for not just for consciousness of its object (if it has one), but also itself and oneself as its subject. The unified picture of consciousness that results should assist research on consciousness.

## Introduction

Though we take mainly a philosophical approach, we hope that the results of our work will be useful to researchers on consciousness using other approaches. Current views on consciousness can be divided by whether the theorist accepts or rejects cognitivism about consciousness. Cognitivism is the view that consciousness is just a form of representation or a property of information-processing systems that have representations (e.g., focussed attention). Anti-cognitivists deny this, claiming that inverted spectrum and zombie thought experiments show that consciousness could change while everything cognitive or representational stays the same. Whatever, researchers on both sides of this fence agrees that consciousness has a *representational base.* Whether or not consciousness *simply is* representational or cognitive, it at least *requires* representation and cognition.

However, there have been few well-worked out attempts to say what this representational base might be like. The two best developed are perhaps the higher-order thought (HOT) approach, in which the representational base of consciousness is a thought directed at one's own psychological state(s), and the transparency theory, in which one's conscious states are said to be, not objects of representation, but things that one knows about by inference from consciousness of the world around one, one's body, and so on.

As we will see, both approaches are lacking. We then introduce a notion of a self-presenting representation and attempt to build a better alternative around this notion. On our view, representations are self-presenting but more than that. A single representation is the basis not just of consciousness of the world and itself but also of oneself as its subject. This notion leads to a unified picture of consciousness.

## The standard picture of representation

Many consciousness researchers accept the following as a principle of representation:

> RP: Representations represent something other than themselves and only something other than themselves.

If RP is right, our view is wrong and something like a HOT model or transparency model has to be right.

There are a host of problems facing HOT models (see Raymont, forthcoming). Perhaps the most serious arises from its separation of the representing state that confers consciousness from the state on which consciousness is conferred. The problem is that a representation can exist in the absence of its object. If so, a HOT that represents pain should be able to make things seem subjectively just as they would if one really were in pain – with no real pain. Rosenthal (1997, p. 744) at least bites this bullet but it is a pretty tough chew. Moreover, since in this case what is represented is not real, it is the representing state that has to be the conscious state. If so, there is nothing higher-order about consciousness. Indeed, the resulting conscious state would look remarkably like our self-presenting state.

So what about the transparency alternative? The basic idea behind transparency theory is that one is directly conscious only of what a representation represents and not the representation itself. We are conscious *via* representations, not *of* representations. Representations are transparent to us.

If so, consciousness of representing is an inference from the fact that we are conscious of what is represented. As Dretske puts it,

> You cannot represent something as F without, necessarily, occupying a state that carries the information that it is F (not G or H) that you are representing something as. [1995, p. 56]

All we know about our representing is what we can infer from how represented items appear. Dretske (p. 40) calls the resulting consciousness of our representing states *displaced perception*. The perception of an object is displaced by an inference onto the perception itself.

The transparency thesis faces some problems. First problem: when one is conscious of something by means of a given representation, one is thereby conscious of that repre-

sentation's content, i.e., of how that thing is represented by one's representation (visually, aurally, and so on). It is a *very* small step from this to consciousness *of* the representation. To take this step, various things may be needed but an inference from *what* one is representing does not seem to be one of them. One is conscious of not only *what* is represented but also *how* it's represented. The latter is an aspect of the representation itself.

Second problem: when one is conscious of *how* one is representing given contents (seeing them, imagining them, doubting them, remembering them), even Dretske allows that this knowledge does not come by way of an inference from representing something (1995, p. 57-8), though he does not say where it does come from. If transparency is not the case here, why anywhere?

Third problem: itches, pains and other bodily sensations. On the transparency view, feeling a pain, hurting, has to be nothing more than an inference from what the pain represents – some bodily damage or whatever. This is extremely implausible. (Dretske recognizes that he has a problem about pains, etc.: "this is a topic that I have neither the time nor (I admit) the resources to effectively pursue" [1995, p. 103].) All who accept RP face this problem.

Indeed, about the only phenomenon for which a transparency claim seems at all plausible is perception, especially vision, though even here transparency might have trouble with the difference between, say, seeing a corner and feeling it, in general with situations in which there is one content, two or more modes of representation. Put bluntly, the transparency thesis seems simply to be false, at least for most of consciousness.

As we said, if RP is right, something like a HOT model or transparency model are the only alternatives. If accepting RP leads to such problems, what happens if we deny it? It is not obviously true. Think again of pains and itches and other bodily sensations, not to mention feelings of tension and tiredness, mood states such as aimless anxiety or euphoria, and so on. What (other than themselves) are states like these about? Then there are altered states of consciousness. When consciousness is altered, what it is like to have those states certainly changes but there is no obvious candidate, other than the conscious states themselves, for what they are about.

There are two ways to reject RP. One would be to say that pains, etc., are not representations at all. This move is hopeless. In having pains (and mood states, altered states of consciousness, and so on), one is clearly aware *of something*, something about which a pain, for example, carries information. These are marks of representing. But what one becomes aware *of* in having a pain, what this state carries information *about*, is *itself*. Pains are representational by being *self*-representational; the qualities of which one becomes conscious by having these states are mainly qualities of the states themselves. So the other way of rejecting RP is to say that some representations at least are self-presenting.

Not only do we think that RP is false but our version of the opposing idea that representations can be self-present-ing is fairly radical. Here is how our story goes.

## Self-presenting Representations: The Representational Base of Consciousness

On our view, having a representation is all the representation that one needs to become conscious not only of what the representation is about and the representation itself, the standard view of self-presenting representations, but something more.[1] Consider the following sentence as uttered by MM:

1. I am reading the words on the screen in front of me.

Having the representation expressed by (1) can make MM conscious of the words on the screen, obviously. It can also make her conscious of the representation itself, that the words are being seen (not heard, imagined, touched, and so on). In addition, having this representation can make MM conscious of who is seeing the words, namely, herself. An ordinary single representation is all the representation that one needs to become conscious of that representation and also oneself as the subject of it. Let us call such a representation the *representational base* of becoming conscious of these items.

> **Representational base** – an act of representing that is all the representation that one needs to be conscious of it and of oneself as its subject

Almost any representation will do.

Imagining something unreal such as Pegasus will do just as well as perceiving an external object such as a computer screen. Indeed, even a representational state that had no object, and therefore could not make us conscious of anything other than itself, could still be the basis of becoming conscious of that state and of oneself. Moreover, a representation need not itself actually be recognized in order to provide a representational base for self-consciousness. Just recognizing what is represented in it would be an adequate representational base for one to be conscious of oneself (as conscious of that object).

Note carefully the term representational *base*. We are not saying that to have a representation is *to be* conscious of it. That would be a crazy view to hold. What we are saying is that having a representation provides everything *representational* needed to become conscious of having it and of oneself. Other things may be needed, too, shift of attention for example, or the conceptual resources to go from consciousness of something to consciousness of representing it.

Lest it be thought that the idea of a self-presenting representations is exotic, note that something as lowly as a bar code can be run quite a long way as an analogy. A bar code contains information about what it is 'about', usually the item's nature and price. But it also contains information about itself – a few of the bars are an integrity check on the

---

[1] Among others, James held the standard view, as more recently have Kriegel (2003) and Tye (2003).

bar code itself. And it contains information about the thing that has it – it is physically mounted on the thing that has it. How far the analogy can be run does not matter here. What matters is that even a representation as simple as a bar code can be self-presenting.

Another homely example, a gauge, shows how our story goes.[2] A gauge presents information about something other than itself, namely, whatever it has the function of indicating information about. For example, an altimeter presents information about the distance to the earth's surface. However, an altimeter also presents information about itself, how far it is from earth, for example. And it is the gauge that presents this information about the gauge, not some higher-order gauge pointed at it! This is how we see conscious states. To be conscious of having a representation, all the representation that one needs is the state itself.

For the analogy of the gauge to be complete, the gauge would have to have one more function, It would have to represent the system that has it. Not a problem. Suppose that to provide information about altitude and itself, an altimeter has to port itself to a system. And suppose that to do so correctly, it has to recognize what sort of system it has been installed in. ('Ah, this is a Cessna Skylane.') Now we have at least a rough analogue of a representation presenting not just its object and itself but also its subject, the person who has it.

We don't have room to mount the full case for our notion of the representational base of consciousness but notice, if it is right, it would do some real work for us.

• It would entail that RP is false.

• It would entail that the idea behind transparency and displaced perception, the idea that we are not conscious of our own conscious states, is false and that the displaced perception move is unnecessary – the consciousness one has of one's own representations by having them is as direct and non-inferential as any consciousness of anything.

And,

• It would show that the HOT move is unnecessary. If a representation itself is all the representation we need to become conscious of that representation, there is no need for a higher-order representation of any kind.

## Global Representation

Our notion of the representational base gives a single, unified account of the basis of three forms of consciousness – consciousness of the world, consciousness of one's own states, and consciousness of oneself. This nicely unified account is progress. Progress – but not the whole story.

So far we have talked exclusively about individual representations as understood by the tradition. As Kant already knew, however, the representations that serve as the representational base of consciousness are usually much 'bigger' than individual representations traditionally conceived. Indeed, in complex beings like us, we do not believe that there are any individual representations as traditionally conceived; but we will not go into that here.

The representations serving as the representational base of consciousness usually have multiple objects and encompass multiple representations (as traditionally conceived). Let us call such a representation a *global representation.* In a global representation, one is conscious of many objects and/or many representations as traditionally conceived, and one is conscious of them as a single complex object and/or a single complex representation.

> **Global representation** – representing many objects and/or many representations as traditionally conceived as a single complex object and/or a single representation.

Our points about the representational base can now be made using this notion. A global representation is all the representation that one needs to be conscious not just of its complex object, if it has one[3], but also of the representation itself, and of oneself as the 'the single common subject' of the elements of this representation (Kant, 1781/7, A350).

The structure of a global representation is complicated. Suppose that at the same time as one has the representation expressed by (1),

1. I am reading the words on the screen in front of me,

one also has representations expressed by,

2. I am puzzled by your comments

3. I am enjoying the music I hear outside

4. I believe our agreement was to meet at 6:00

5. Yesterday I thought I understood Kant's notion of the object

6. I wish the world were a fairer place

There are three elements of (1)-(6) that could be united in a single global representation.

• One could be conscious of the various represented objects here as a single complex object.

• One could be conscious of the various ways in which these objects are being represented as a single complex representation.

And,

• One could be conscious of oneself, the subject, as the single common subject of the whole business.

The next question is, How does a global representation serve as the representational base of consciousness of its complex object, itself, and its subject, the person who has it?

## Joint Consciousness

Central to a global representation of objects is what we will

call *joint consciousness:*

> **Joint consciousness** – to be conscious of any of the objects of a global representation is to be conscious of other such objects.

It seems obvious that joint consciousness is a distinctive feature of a global representation. The notion of joint consciousness just stated clearly applies to consciousness of the world (plus such things as one's own bodily states), more exactly, to consciousness of intentional objects. What about consciousness of one's global representation and of oneself as its subject?

When one is conscious of representing in a global representation, here too there will be joint consciousness – to be conscious of some representings is to be conscious of others. And conscious of self? Here the idea that is plausible is a bit different. When one is conscious of oneself as subject of one bit of representing, one will usually be conscious of oneself as the subject of other representings, as their 'single, common subject', to use Kant's words again.

The next question. How could a global representation serve as the representational base of *joint* consciousness of representings and of oneself as their *common* subject? It seems plausible to hold that only certain elements of a global representation are needed for one to have an adequate representational base for consciousness of a single representing and of oneself as its subject. One will be conscious of doing an act of seeing by doing that act of seeing, whatever the rest of one's current global representation is like. And one will conscious of oneself *seeing*, which makes it likely that the act of seeing is the basis of this consciousness of oneself. If so, what is the representational base of *joint* consciousness of objects, of being *jointly* conscious of various representings, of being conscious of oneself as the *common* subject of a number of acts of representing?

Here we can only sketch what would have to happen for these forms of joint consciousness to occur. In the same way as an act of representing is the representational base for consciousness of that act and of oneself, it will have to be the representational base for consciousness of doing *other* acts of representing and of oneself as their *common* subject.

Moreover, the relationship will be symmetrical. One is not 'located' at any given representing. When consciousness of a representing carries one to consciousness of other representings, one has all the representings concerned equally. So it would be better to say that representings in general are the base for consciousness of representings in general. It is not any given representing that is the base of consciousness of oneself as common subject but representings in general.

Thinking of representations in the traditional way, it is hard to make sense of what we just said. But that may be the fault of the traditional conception. There is another way to think about representation within which what we just said makes fine sense.

## Structure of a Global Representation

Instead of trying to make sense of the base of joint consciousness of representing and of oneself as subject in terms of relations *among* representations, let us try out the idea that these forms of consciousness are something that obtains *within* a single complex representation.

Our test case will be a person seeing something, hearing something, and tasting something, all parts of a global representation and global object, where the person is jointly conscious of the objects, the representings, and herself as subject. To be the base for such consciousness, how are the three acts of representing brought together in a global representation? Here are three possibilities:

1. The three acts and their objects become the object of a fourth, higher-order representation.

2. The three acts and their objects become parts of a single subsuming representation.

3. While their contents are taken up in a global representation, the three acts of representing do not survive even as parts of this state, though their objects remain distinct. They become three (for the moment let us call them) *modalities* of a single representation.

How might (3) work?

Consider what happens if one goes from a situation, at time $t$, that has objects $o1$ and $o2$ to a situation that has $o1$ but not $o2$. We could try to capture the change in two ways. We could say that where once there were two representations, $r(o1)$ and $r(o2)$, which were bundled in a mental structure of some kind, $[r(o1) \& r(o2)]$, we now have only the one representation, $r(o1)$; one representation has been dropped from the bundle that existed at $t$. Or we could say that there was just one representation, $r(o1 \& o2)$, at $t$ and it has been replaced by another single representation, $r(o1)$.

The second view is simpler, since it does not involve postulating representations as parts of an encompassing representation. According to it, at $t$ there was one representation that had a complex content. The content was complex because it had multiple contents, $o1$ and $o2$, as its parts. If it were a conscious representation, what would make it one representation is that to be conscious of any of its objects by means of it is to be conscious of other of its objects, too. Here, the part-whole relation obtains among objects, but there is no parallel multiplicity of representational states. On this approach, unlike the first, the representing state does not have 'smaller' or less complex representing states as parts.[4]

On this picture, globality obtains *within* a representation but not *among* representations. It does obtain among objects. Returning to the joint consciousness condition, where to be conscious of one thing is to be conscious of others, on our picture joint consciousness of objects can be present or absent but it is trivially present in a global representation,

---

[4.] If the representing state is a brain state, then it will have parts, but these parts will not be representing states.

simply because a global representation is a single representing state, r($o1$ & $o2$). For conscious representing states of of this form, to be conscious of $o1$ by means of this state is also, by that very same act, to be conscious of $o2$. One representational state-token provides consciousness of both $o1$ and $o2$.

Consider an act of reference as an analogy. Suppose I refer to Toronto. Scarborough is part of that city; it is a part of the thing to which I referred. It does not follow that my act of referring to Toronto contains a numerically distinct reference to Scarborough. It was of course *possible* for me to refer simply to Scarborough, and thus to refer to part of the thing to which I actually referred, but it does not follow that I actually did so. The mere fact that Scarborough is part of the thing to which I referred does not entail that a reference to that borough figures as part of my act of referring. Similarly, the mere fact that $o1$ is part of ($o1$ & $o2$) does not entail that a distinct representation of $o1$ must figure as part of my act of representing ($o1$ & $o2$).

It may be objected that in advancing these observations and contrary to our intentions, we actually *provide the resources* for showing that a subject's global representation will contain parts that are themselves representations. Didn't we ourselves just say that representations are individuated by their objects? This suggests that for each object that we can individuate, including objects that are parts of a global object, there will be a corresponding representation individuated by that object.

This objection does not work. If we have relied on a claim that each representation is individuated by an object, we have not said how objects individuate a representation. What if only some objects individuate a representation? Let us introduce the idea of a *global object.*

**Global object** – a group of represented objects that is the single complex object of a global representation

If we now say that only a global object individuates at least a conscious representation, we are clear of the objection.

To capture these claims about singularity of representation, let us generalize (3),

3. While their contents are taken up in a global representation, the three acts of representing [viz., seeing something, hearing something, and tasting something] do not survive even as parts of this state, though their objects remain distinct. They become three modalities of a single representation.

into (3′),

3′. A global representation at a given moment is a single representation not made up of multiple distinct representations and it has a complex object.[5]

Now that we have said what we mean by (3) and (3′), what about (1) and (2)? Though most philosophers hold to (1) or (2), they do so uncritically and seldom offer any

support for the views. We do not know of any good argument for either one of them. To see where those adopting (1) or (2) might go wrong, recall the representation, r($o1$ & $o2$).

Suppose that this representation occurs at $t$. Supporters of (1) and/or (2) may confuse the *untokened* (at $t$) representational type that would take part of r($o1$ & $o2$)'s object (*viz.* $o1$) as its object, with a representational token, r($o1$), held to be *part* of r($o1$ & $o2$) and to have one of the latter state's objects as its sole object. There is no reason to suppose that there actually is such a distinct representational state nestled within r($o1$ & $o2$). There could have been such a token at $t$, but the mere fact that we can entertain such a possibility – that is, the mere fact that we can think of an instantiation at $t$ of that type – is no reason to conclude that there actually exists a token of that type at $t$.

Conclusion: (3′), the idea that at a given moment a global representation is not a group of representations but a single representation with a complex object, is a perfectly coherent point of view, one well supported by examples and analogies. Before we leave it, however, there are some objections that we need to answer. One of them is so obvious that it has probably occurred to most readers already.

## Single Representation View: Objections

The obvious objection is this. In a global representation, we are conscious, within a single unified representation, of several sensory modalities. The phenomenal field is polymodal: it involves tactual data, visual data, auditory data, and so on.

(1) or (2) would try to account for the contributions of the different modalities to consciousness by saying that there are several distinct representations here – visual representations, auditory representations, and so on – that come together as parts of an encompassing global representation. On this view, the cognitive system constructs a variety of representations in different modalities. These representations are not simply superseded by a global state that combines their informational contributions in one representation. Instead, they are preserved as distinct representations within it. The polymodal complexity of the resulting global state is due to the presence in it of this range of representations.

Against this, we offer the following picture: we do not have several visual, aural, etc. representations, not conscious ones anyway. Rather, the information we represent is *formatted* visually, aurally, etc. The cognitive system receives some information in a visual format (reflecting, perhaps, the wave length of incoming energy) or tagged as visual, some formatted or tagged as aural, and so on. When this information appears in a global representation, its modality appears with it. But there is just one representation, the global representation. In this one state, diverse bits of information are formatted in a variety of modalities.

This view has the twin virtues of adequacy and parsimony. The onus rests with proponents of the more complicated view, in which a global representation is held to be an assemblage of other representations, to show that our account has failed to account for something. It is difficult to

---

5. James (1890, vol. 1, pp. 145-61) held something like this view.

see what that could be.

Suppose that one is consciously representing things both visually and aurally. What accounts for something there being aural? This question must be answered either in terms of *how* the representation represents or in terms of *what* it represents. There are no other ways to specify the modality of a representation. But a global representation can incorporate either of these ways. There can be a diversity in what it represents; it can represent many things having many properties. And it can represent these contents in a variety of ways, visually, aurally, and so on.

In support of this contention, note that we are *compelled* to postulate a plurality of representations only when there is a particularly strong sort of incompatibility in how or what a cognitive system represents, that is, only when the system represents in ways that exclude one another. For example, theorists tend to attribute distinct representations (not all of which are held to be conscious) to a system in such cases as binocular rivalry, or to explain the fact that consciously seeing the Necker Cube in one way precludes consciously seeing it in the alternative way.

These cases and all others that we've been able to think of are not problems for us. To be a counter-example to our single representation/complex object picture, the elements would have to be: all conscious, and all available simultaneously to one conscious subject. The cases we've just considered do not meet these conditions.

In the Necker cube case, the conflicting representations are successive. In binocular rivalry, the representations may be simultaneous but one is not simultaneously conscious of them (Baars 1988, pp. 82-3; see p. 126). So even when we are driven to multiply representations in a subject at a time, the results are not a problem for our view.

Notice that cases such as binocular rivalry and the Necker cube arise *within* a modality. It is not clear that there are any strong incompatibilities across modalities. Certainly there aren't many. In the absence of strong incompatibilities, nothing compels us to posit *more than one* representation to make a place for different perceptual modalities.

Second objection. Think of a picture of a car in front of a house. It is plausible to say that this picture includes a picture of the car and a picture of the house, that the bigger picture of the car and house together contains little pictures of the car and of the house.

Is it so clear, though, that the picture of house and car together literally contains several distinct *pictures*, one for each item depicted? The belief that it does threatens to introduce an implausible multiplication of pictures. If one can discern ten thousand blades of grass in front of the house, then there would have to be ten thousand pictures in the larger picture. Clearly, at some level of decomposition, we stop positing a distinct picture for each part of the content that we are able to distinguish. Why not stop at the whole picture and say that it is the only picture, with no smaller pictures in it?

Third objection: If conscious representation consists of one big, non-compositional representation, how are acts of judging particular bits of content, forming beliefs about things based on particular bits of content, possible? Response: In the same way as information in the complex global object of a global representation can come 'marked' with various modalities (aural, visual, etc.), particular bits of the information in a complex global object can enter into particular information-processing activities: judging, remembering, and so on. These *activities* do not need to merge into a single representation and they can pick and choose information to work on *ad libitum*.

Even after responding to these objections, the support we have offered for (3′) is not decisive. It is strong, however, certainly strong enough to justify acceptance of the view.

Our theory that a single global representation is the representational base of consciousness has real potential. As we have seen, it provides a unified account of three major kinds of consciousness, consciousness of world, one's own representations, and oneself as subject. Though we can't show this here, it also opens the way to a nice account of: the unity of consciousness; the special features of consciousness of self; and the subject of consciousness (Brook and Raymont, forthcoming).

## References

Baars, B. 1988, *A Cognitive Theory of Consciousness.* Cambridge: Cambridge University Press.

Brook, A. and P. Raymont. forthcoming. *A Unified Theory of Consciousness.* Cambridge, MA: MIT Press.

Dennett, D. 1978e. Mind writing and brain reading. In *Brainstorms.* Montgomery, VT: Bradford Books, pp. 39-50.

Dretske, F. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press/A Bradford Book.

Kant, I. (1781/1787) *Critique of Pure Reason* (trans. P. Guyer and A. Woods). Cambridge: Cambridge University Press, 1997. (References are in the pagination of the 1st (A) or 2nd (B) editions.)

Kriegel, U. 2003. Consciousness as intransitive self-consciousness. *Canadian Journal of Philosophy* 33:103-132.

James, W. 1890. *Principles of Psychology.* London: Macmillan.

Raymont, P. submitted. From HOTs to self-representing states

Rosenthal, D. 1997. A theory of consciousness. In N. Block, O. Flanagan, and G. Gűzeldere, eds. *The Nature of Consciousness.* Cambridge, MA: MIT Press.

Tye, M. 2003. *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press/A Bradford Book.