

## PHENOMENAL CONSCIOUSNESS AND THE ALLOCENTRIC-EGOCENTRIC INTERFACE

PETE MANDIK

*Department of Philosophy, William Paterson University, 300 Pompton Road  
Wayne, New Jersey 07470, U.S.A.  
(mandikp@wpunj.edu )*

**Abstract:** I propose and defend the Allocentric-Egocentric Interface Theory of Consciousness. Mental processes form a hierarchy of mental representations with maximally egocentric (self-centered) representations at the bottom and maximally allocentric (other-centered) representations at the top. Phenomenally conscious states are states that are relatively intermediate in this hierarchy. More specifically, conscious states are hybrid states that involve the reciprocal interaction between relatively allocentric and relatively egocentric representations. Thus a conscious state is composed of a pair of representations interacting at the Allocentric-Egocentric Interface. What a person is conscious of is determined by what the contributing allocentric and egocentric representations are representations of. The phenomenal character of conscious states is identical to the representational content of the reciprocally interacting egocentric and allocentric representations.

**Keywords:** Consciousness – Qualia – Representation – Philosophy – Neuroscience

### 1 Introduction

The purpose of this chapter is to sketch in a relatively small amount of space a relatively comprehensive theory of phenomenal consciousness, one that is both empirically warranted and philosophically productive. A theory of phenomenal consciousness must do two things: it must explain what makes a mental state a conscious mental state (as opposed to an unconscious mental state) and it must explain what it is for a conscious mental state to have phenomenal character, that is, the property of the state in virtue of which there is something it is like to be in that state.

The theory I shall sketch is what I shall call “The Allocentric-Egocentric Interface Theory of Consciousness” or “the AEI theory” for short. In brief, the theory posits that mental processes form a hierarchy of mental representations with maximally egocentric (self-centered) representations at the

bottom and maximally allocentric (other-centered) representations at the top. Part of what it means to be higher or lower in the hierarchy is to be further from or closer to the sensory and motor periphery of the nervous system. Focusing on the processing of sensory information, we can trace the path of information from relatively egocentric representations of the stimulus in sensation through stages of processing that increasingly abstract away from egocentric information and represent things in memory in an allocentric way. Further, we can note top-down effects from relatively allocentric representations high up in the hierarchy to egocentric representations lower in the hierarchy. I hypothesize that phenomenally conscious mental states are to be identified with states that are relatively intermediate in this hierarchy. More specifically, conscious states are hybrid states that involve the reciprocal interaction between relatively allocentric and relatively egocentric states: a conscious state is composed of a pair of representations interacting at the Allocentric-Egocentric Interface. Unconscious mental states are states that are either too high up or too low down in the hierarchy or are not engaged in the requisite reciprocal interactions. What a person is conscious of is determined by what the contributing allocentric and egocentric representations are representations of. The phenomenal character of these states is identical to the representational content of the reciprocally interacting egocentric and allocentric representations.

That, at least, is the theory in brief. The remainder of the paper will spell things out in a bit more detail. It is organized as follows. First I say a few more things about the way philosophers think of phenomenal consciousness and how this might relate to empirical work on consciousness. Then I dive into the exposition of AEI with an emphasis on empirical evidence for the theory. Finally I discuss philosophical ramifications of the theory.

## **2 Philosophy and Phenomenal Consciousness**

The phrase “phenomenal consciousness” involves a philosophical technical use of “consciousness” and it is best introduced by first noting some less technical uses of “consciousness” and the related term “conscious.”

The philosopher David Rosenthal has made some observations about different uses of the word “conscious” that are useful for gaining an understanding of what consciousness is supposed to be [1,2]. The first way in which uses of the word “conscious” may be distinguished is in terms of a distinction between the things that may be said to be conscious. The first sort of thing that we may say is conscious is a person or a creature. A person or

non-human creature that is awake and responsive is said to be conscious in this sense of the word “conscious” and Rosenthal labels this “creature consciousness”. A second sort of thing that we may say is conscious is a mental state of a person or a non-human creature. Many of us are familiar with the Freudian idea that some of our beliefs and desires are conscious while others are unconscious, and since the state of a creature is a very different sort of thing from the creature itself, the sense in which a state is conscious (“state consciousness”) is a very different sense of the word “conscious” than is the sense in which a creature is conscious (“creature consciousness”). Rosenthal further distinguishes the notions of creature consciousness and state consciousness in terms of the relative difficulty that theorists have had in understanding these types of consciousness. According to Rosenthal, being creature conscious amounts to no more than being awake and having mental states whereas state consciousness requires the satisfaction of additional criteria that distinguish conscious mental states from unconscious mental states [1].

A second kind of distinction that Rosenthal draws among uses of the word “conscious” distinguishes transitive from intransitive uses of the word [2]. The word conscious is used transitively when we speak of someone being conscious of something as when I am conscious of a buzzing insect that is pestering me. Intransitive uses of the word “conscious” are silent about whether the thing that is conscious is thereby conscious of something. Thus, both creature consciousness and state consciousness are instances of intransitive consciousness. We thus have on our hands at least three kinds of consciousness: creature consciousness, state consciousness, and transitive consciousness, the second and third of which will be especially important to the current discussion. Various theorists disagree as the ultimate natures of these three kinds of consciousness. In particular, they disagree as to how these three kinds might be explained in terms of one another.

I will briefly indicate a couple of positions to give a flavor of the issues. Rosenthal advocates an explanation of state consciousness in terms of transitive consciousness and an explanation of transitive consciousness in terms of mental representation. According to Rosenthal, a person is transitively conscious of  $x$  just in case they have a mental representation of  $x$ . Further, Rosenthal embraces as pre-theoretically intuitive the transitivity thesis whereby a person has a conscious mental state only if they are transitively conscious of that state [1]. It follows from these two points, then, that a person has a conscious state only if that person has a mental representation of that state.

In contrast, theorists such as Dretske [3] and Tye [4,5] deny the transitivity thesis and instead embrace the transparency thesis whereby when one has a conscious state all one is conscious of is what the state is a representation of. A conscious state, according to Dretske, is a state in virtue of which one is conscious of something [3]. Thus, if one is conscious of a buzzing insect, one is thereby in a conscious state and need not additionally be conscious of that state. In spite of his disagreement with Rosenthal about state consciousness, he agrees that transitive consciousness is to be defined in terms of representation: being transitively conscious of  $x$  involves mentally representing  $x$ .

Note that the above discussion of creature consciousness, state consciousness, and transitive consciousness made no explicit mention phenomenal consciousness. The above points about state, creature, and transitive consciousness can all be made by calling our attention to the various ways in which people use the word “conscious” in non-technical discourse. The same cannot be said of so-called phenomenal consciousness. The term “phenomenal consciousness” is not only a technical term, but often those who introduce it into technical discourse give little indication as to how it involves a common sense use of the word “conscious.” Indeed, when people use the term “phenomenal consciousness” they are not much interested in a kind of consciousness distinct from the three already mentioned. They are interested instead in certain properties that conscious states allegedly have. These properties are referred to interchangeably as “phenomenal properties,” “phenomenal character,” or “qualia.” Two key theorists who discuss phenomenal consciousness and qualia are Block [6] and Chalmers [7]. Block characterizes phenomenal consciousness (P-consciousness) as follows.

P-consciousness is experience. P-conscious properties are experiential ones. P-conscious states are experiential, that is, a state is P-conscious if it has experiential properties. The totality of the experiential properties of a state are “what it is like” to have it. Moving from synonyms to examples, we have P-conscious states when we see, hear, smell, taste, and have pains. P-conscious properties include the experiential properties of sensations, feelings, and perceptions, but I would also include thoughts, wants, and emotions. (p. 230)

Chalmers writes that

a mental state is conscious if it has a qualitative feel — an associated quality of experience. These qualitative feels are also

known as phenomenal properties, or qualia for short. The problem of explaining these phenomenal properties is just the problem of explaining consciousness. This is the really hard part of the mind-body problem. (p.4)

Additionally, Chalmers writes that

what it means for a state to be phenomenal is for it to feel a certain way. ...in general, a phenomenal feature of mind is characterized by what it is like for a subject to have that feature. . . (p.12)

It is clear that the kind of consciousness that both Block and Chalmers are interested in is state consciousness, for all of their examples of consciousness are examples of conscious states. And what they find interesting about conscious states is, as Chalmers puts it, their “phenomenal properties, or qualia for short.”

What, then are qualia? This is of course a vexed issue, but we can nonetheless characterize some agreement among those who are fond of asserting their existence. Qualia are alleged to be (1) intrinsic properties of conscious states that (2) account for “what it is like” for the subject to be in that state and are (3) directly and fully knowable only by that subject. To further characterize qualia, we can start with (2) and work our way out to (3) and then (1). When I have a conscious mental state, e.g. a conscious perception of a red rose, there is presumably something it is like for me to be in that state, and for all I know, when you have a conscious perception of a red rose, what it is like for you is quite different from what it is like for me. For all I know what it is like for you to see red things is like what it is like for me to see green things, and vice versa. And, while you can tell me a bit about what it is like for you there’s a lot that you know that I can’t have access too. It is in this way that the characterization of “what it is like”-ness goes hand and hand with the claim that this whatever-it-is is only directly and fully knowable by the subject who has it. The allegedly intrinsic nature of qualia is hypothesized to explain why what it is like to be in a mental state is only directly and fully knowable by the subject of the state. The thought here is that if something is defined by relations it enters into, then we can describe it fully by describing the relations it enters into. Thus if there is something that we cannot describe fully by describing the relations it enters into then it is not defined by relations it enters into.

Many philosophers have correctly raised the question of whether there really are such properties as qualia, that is, whether there really are proper-

ties that satisfy all three conditions. One of the most questionable aspects of qualia concerns whether they are intrinsic. One can question whether there really are any properties that are intrinsic. Or one can grant that some properties are intrinsic, but question whether intrinsic properties are consistent with being the sorts of properties that account for what it is like to have a mental state or the sorts of states that account for whether they are directly and fully knowable by the subject. One sort of consideration against regarding qualia as intrinsic is that if qualia are definable in terms of being directly knowable by the subject, then that makes them extrinsic or relational after all, since qualia would be definable in terms relations to the subject and the various parts of the person in virtue of which the person has any knowledge at all.

Given the highly questionable status of the allegedly intrinsic nature of qualia, I will proceed with the following minimal account of qualia: qualia (phenomenal characteristics) are the properties of mental states in virtue of which there is something it is like to have them. This characterization thus leaves open whether such properties are intrinsic or fully and directly knowable only by the subjects of the mental states that have them. It is worth noting that in the characterizations of phenomenal consciousness from the quotations from Block and Chalmers given above, qualia were characterized solely in terms of this minimal characterization, that is, no explicit mention was made of qualia being either intrinsic or knowable fully and directly only by their possessors. On the account of consciousness I advocate below, qualia will not turn out to be intrinsic for they will turn out to be representational contents and representational contents are widely and correctly regarded as relational. Regarding the question of whether they are directly and fully knowable only by the subject, I address that question at length elsewhere and will thus devote no additional space to it here [8].

### **3 Neuroscientific Applications of the Concepts of Consciousness**

Let us leave these philosophical remarks aside for a moment and turn to discuss how one might apply these concepts of consciousness in empirical settings. Two useful kinds of cases to look at in this regard involve blindsight and motion induced blindness. I begin with blindsight.

We can roughly characterize blindsight as a condition following certain sorts of brain damage wherein subjects report a loss of consciousness in spite of the retention of visual ability. One source of difficulty in characterizing blindsight in a clear and uncontroversial way is that there are few if any clear and uncontroversial ways of characterizing the relevant no-

tions of consciousness and visual ability. Another source of difficulty is that, even when one is able to fix interpretations for “consciousness” and “visual ability” the data concerning blindsight shows neither a total loss of consciousness nor a total retention of visual ability. Many philosophers who delight in the superior chemical purity of thought experiments over real experiments have taken to discussing the conceptual possibility of the imaginary condition superblindsight wherein both loss of consciousness and retention of visual ability are total. I propose that in this case, if not in general, we should reject fake chemical purity as being as good as none whatsoever and thus seek to gain as much clarity as possible by examining real phenomena.

We can summarize blindsight in terms of the following questions: What are the lesions involved? What is the measure of retention of visual ability? The lesions involved are to primary visual cortex (area V1). In an early study of blindsight subjects were better than chance at moving their eyes to the location of a flash of light in a region of the visual field wherein they had reported not being able to see flashes of light [9]. Subsequent studies investigated the sorts of stimuli that blindsight subjects could respond to in spite of reporting no visual awareness, including stimuli characterized in terms of features such as wavelength [10,11] and motion [12]. Additional studies examined the way in which the presence or absence of consciousness could be indicated as in, for example, having the subject indicate by pressing one of two keys “whether he had any experience whatever, no matter how slight or effervescent” [13].

There are several natural suggestions of how to apply the concepts of consciousness from the previous section to the case of blindsight. The retention of visual abilities indicates that certain kinds of visual information are represented in the nervous system and that these representational states can guide certain behaviors. However, it seems clear that these representational states are not conscious states. Thus, first and foremost, the kind of consciousness that blindsight subjects seem to lack with respect to vision is state consciousness. Regarding the question of transitive consciousness, note that theorists that regard representation as sufficient for transitive consciousness will attribute transitive consciousness of items in the “blind” regions of the subjects’ visual fields. However, if we take subjects’ reports at face value, it seems that they lack transitive consciousness since they report not being conscious of the items in question.

One problem with blindsight is that the majority of the readers of the present paper are not themselves blindsight subjects and insofar as intro-

spection is an important source of information regarding the adequacy of theories of consciousness, the reader may find some of the remarks about blindsight difficult to evaluate. It will thus be helpful to look at empirical work on consciousness that the reader may have an easier time relating to.

Experiments concerning the phenomenon of motion induced blindness are very useful in this regard [14]. One way of eliciting the phenomenon of motion induced blindness is by having normally sighted subjects look at a computer screen that has a triangular pattern of three bright yellow dots on a black background. Moving “behind” the yellow dots is a pattern of blue dots. As subjects stare at the screen, fixating on the midpoint between the three dots, they report that one or more of the dots disappear. I’ve seen the effect several times myself and it is quite salient (readers are encouraged to search the internet for “motion induced blindness” and see for themselves). The yellow dots are not really disappearing from the screen although it looks as if they do. Further, there is evidence that the brain continues to represent the presence of the yellow dots and the “disappearance” is due to the representations changing from being conscious representations to being unconscious representations as opposed to being due to retinal suppression, sensory masking, or adaptation. Some of the evidence that the yellow dots are still represented in the brain includes the fact that the effect can be influenced by transcranial magnetic stimulation to parietal cortex (a relatively late stage of visual processing in the brain). Other evidence is that motion induced blindness is sensitive to object grouping so that, for example, when the stimuli are, instead of the three yellow dots, a pair of partially overlapping circles, one yellow and one pink, sometimes an entire circle will disappear leaving the other behind even though some of their contours are very close in the visual field. The brain mechanisms thought to mediate such object groupings are relatively late in the visual processing hierarchy. Thus, information concerning the yellow dots is represented at many levels of processing prior to consciousness.

To relate motion induced blindness to the kinds of consciousness described above, there are several compelling and plausible ways of describing the changes in consciousness. First, I think it is clear to anyone who has experienced the phenomenon that at one moment you are conscious of a yellow dot and then at another moment you are not. Since the consciousness in question here is “consciousness of” this is to describe the case in terms of transitive consciousness. Second, the combination of first-person introspective evidence and third-person empirical evidence indicates that in the course of a motion induced blindness experiment there is, under some



conditions, a conscious mental state and, under other conditions, an unconscious mental state. When it is obvious to subjects that they are in a state of perceiving the yellow dots this is a conscious state. When the dots disappear, even though there is a brain state that represents the yellow dots, it is not a conscious state. (It is accurate to regard this neural state as a mental state insofar as it is accurate to regard it as an unconscious perception of the yellow dots). Since phenomenal character attaches to states that are conscious, and phenomenal character is the property in virtue of which there is something it is like to be in that state, we are in a position to investigate what the properties are of that state that are relevant for determining what it is like for that person. More can be said later, but for now we can note that what it is like for that person is that it is like seeing a yellow dot appear and disappear. This suggests that at least part of what it is like to be in that state is determined by its representational content: the state represents the presence of a yellow dot. Returning to the question of what makes the state in question a conscious state, we can note one important similarity between various competing theories of state consciousness: namely that they explain state consciousness in terms relations between relatively low level states like sensations and relatively high level states like conceptual thoughts. So, for advocates of the transparency thesis (like Tye) what makes the representation of the yellow dot a conscious representation is that it is poised to interact with higher level conceptual states like beliefs and desires [4,5]. Many advocates of the transitivity thesis, like Rosenthal, will require that what makes the representation of the yellow dot conscious is not merely that it is poised to interact with higher level conceptual states, but further, that those conceptual states must be about — that is, representations of — the low level representation of the yellow dot [1,2].

I propose that these theorists are correct in trying to explain state consciousness in terms of interactions between high-level and low-level states. However, I want to argue ultimately that neither the transitivity thesis nor the transparency thesis is true. I favor a different account of the relative contributions of the various levels to state consciousness. In the remainder of the discussion, I will be primarily interested in visual consciousness, though I do intend the theory to apply to consciousness generally.

When we look at visual processing, we can characterize levels in a hierarchy of information processing. More specifically, we can characterize the levels in terms of how much the information has been processed, where the information is being processed, and what the nature of the processing is.

The questions of how much and where can be answered simultaneously by tracing the flow of visual information from the earliest stages of processing in the eyes through to the latest stages of visual processing in the cerebral cortex. More specifically, we can trace the flow of information from retinal ganglia through the optic nerve to the subcortical structures of the lateral geniculate nucleus. Next information is sent to the first stages of cortical processing in occipital cortex in the primary visual area (area V1). Later stages of cortical processing involve sending information along two branching paths [15]. The first is the dorsal stream that sends information from occipital cortex to posterior parietal cortex. The second is the ventral stream that sends information from occipital cortex to inferotemporal cortex. Still later areas of processing involve areas in frontal cortex [16] as well as in the hippocampus [15]. It is worth noting that the flow of information is not strictly feed forward from sensory input to the highest levels of brain processing but also includes many instances of feedback or back-projections of information being sent back from higher levels to lower levels [17].

The “where” and “how much” questions do not exhaust all there is to say about visual processing: there remains the question of what the nature of the processing is. I propose that one fruitful general way of understanding what is happening to visual information as it progresses through the levels of the processing hierarchy is that what begins as a relatively egocentric (self-centered) representation of visual information becomes increasingly abstracted and increasingly allocentric (other-centered) in the higher levels. We are to find the most egocentric visual representations in the lateral geniculate nucleus and also in the primary visual areas in occipital cortex. The most allocentric representations are found in frontal areas and hippocampus. Intermediate areas of cortical visual processing contain representations that are intermediate between being egocentric and allocentric. To get a clearer grasp of the proposal that visual processing can be characterized in terms of egocentric to allocentric transformations of represented information, it will be helpful to consider a more detailed discussion of egocentric and allocentric representations.

#### **4 Egocentric Representations and Allocentric Representations**

A useful starting place in characterizing egocentric representation is the notion of a receptive field. A good initial definition of “receptive field” is “area in which stimulation leads to response of a particular sensory neuron” [18]. Retinal ganglion cells and neurons in the lateral geniculate nucleus have circular fields with either an excitatory center and an inhibitory sur-

round or an inhibitory center and an excitatory surround. The locations of these fields are defined relative to retinal locations, that is, a particular cell in, e.g., lateral geniculate nucleus, is most responsive to a visual stimulus falling on a specific retinal location. The firing of such a cell is thus said to represent the location of a stimulus in a region of retinocentric space [19].

Retinocentric representations are the lowest level representations in a hierarchy from the most egocentric representations to the most allocentric. The progression of information up the hierarchy progressively abstracts away from the particularities of the maximally egocentric representations as in transformations from retinocentric to head-centered and body-centered representations. Such transformations involve neurons in area 7a of posterior parietal cortex. These neurons exhibit different responses depending in part on whether eye-position is fixed. When eye position is fixed, these neurons exhibit retinocentric receptive fields. However, when eye position is not fixed, stimulus of a given retinal region results in a neural response that varies linearly with eye position. Under these later conditions then, these neurons have a linear gain field defined over eye position. Response in normal conditions, then, is a product of retinal stimulus location and eye position resulting in a neuron tuned to a particular location in head-centered space [20].

The egocentric representations described above involve sensitivity to a spatial location relative to some part of the organism or the organism as a whole. However, egocentricity is not limited only to the responses of sensory neurons, but can be defined for motor neurons as well. For example, reach plans for arms are encoded in eye-centered coordinates [21].

There is more to our mental lives than can be accounted for by egocentric representations. Many of our thoughts have a detached or objective character that abstracts away from peculiarities about us. For instance, my knowledge that Pi is an irrational number is not in any obvious way about me, regardless of how irrational I might be. Similarly, my grasp of the fact that neutrons and protons are more massive than electrons is not particularly a fact about me, in spite of the fact that much of my mass and volume is determined by neutrons, protons, and electrons.

Our capability to have detached, objective mental states is grounded in our allocentric representations. Allocentric representations have been postulated to exist in frontal areas as well as in hippocampus. I here focus on research concerning the neural basis of allocentric hippocampal representations. The classic studies in these areas concern the spatial navigational capabilities of rats, especially comparisons of performance of rats with and

without lesions in hippocampus. One representative class of experiments concerns the performance of such rats in the Morris water maze. The Morris water maze consists of a container filled with water rendered opaque by the addition of milk powder. In typical conditions, rats swim in the water to goal locations consisting of platforms submerged deep enough to not be visible to the rats, but shallow enough to offer a place to rest and breathe without having to tread water. In one such study the water maze was set up such that rats had to swim to a platform rendered visible during training trials, but occluded by opaque water during testing trials. Orientation cues consisted of varied visual stimuli positioned around the maze. Intact and hippocampal damaged rats were trained to swim to the platform from a given start location. Test trials involved two general kinds of condition: one in which the starting position was the same as in the training trials and one with novel starting positions. In trials where starting positions were the same in test as in training, both intact and hippocampal damaged rats were able to swim to the platform. However, in trials where starting positions in the test differed from the training start positions, intact rats out performed hippocampal damaged rats. Hippocampal damaged rats took much longer to reach the platform, and in some cases never found the platform [22].

Results such as these have led to the hypothesis that the hippocampus functions in spatial navigation by supporting a cognitive map involving allocentric representations of the spatial layout of the creature's environment. Allocentric representations are implicated since the rats' navigational ability does not seem to be tied to any particular point of view or orientation within the environment.

One especially prominent proposal concerning how allocentric representations in the hippocampus underwrite successful navigation is the slope-centroid system [23,24]. The basic idea behind the slope-centroid system is that of a polar coordinate system based on the distribution of objects in the animal's environment. The centroid is the point at the center of the collection of objects in the environment. The slope is a line running through the longest axis of the collection of objects. Orientation within the environment is encoded in terms of angles relative to the slope. Position within the environment is encoded in terms of a vector defined by distance from centroid and angle relative to slope. Movements would be encoded in terms of vectors encoding direction and distance. As the animal moves around in the environment, the vector encoding movement is added to the vector for the current location resulting in a vector encoding the location expected at the end of the movement. Upon arrival at a goal location, com-

parison of sensory inputs to the representation of expected location results in a capability for mismatch detection allowing for continual correction and updating of the memory of the environmental layout.

There are thus three major representational components of the slope-centroid system: the representation of place, the representation of heading, and the representation of speed. The hippocampal implementations of these representational capacities are postulated to be the following. Place representations are thought to be implemented by pyramidal cell activity, with highest level of activity in a pyramidal cell corresponding to the animal's current location, irrespective of the animal's heading. Representation of heading is thought to be implemented by activity in cells in nearby brain regions with activity corresponding to the direction the animal is facing irrespective of the animal's location within the environment [25]. The representation of speed is thought to be implemented by the frequency of sinusoidal oscillations of the hippocampal EEG called the "theta pattern" [24].

## 5 Locating Consciousness in the Allocentric-Egocentric Hierarchy

It is reasonable to ask where in the processing stream conscious states arise. I turn now to considerations that we should regard conscious states as residing neither at the highest most allocentric levels nor at the lowest most egocentric levels. Instead, visual consciousness resides at an intermediate level. We can arrive at this conclusion by first noting that neither fully allocentric representations nor fully egocentric representations are ever conscious states.

Purely egocentric representations are not sufficient for conscious states. Egocentric representations count among the most basic and primitive forms of representations. For example, the kinds of spatial representations that arguably underwrite taxes (movement toward or away from a stimulus) in organisms as simple as the nematode worm *C. Elegans* (a creature with a nervous system of only 302 neurons) represent spatial distances and directions in egocentric terms [26,27]. While such creatures are complex enough to support egocentric representations, few theorists would regard them as complex enough to support phenomenal consciousness.

Another consideration against thinking that egocentric representations are alone sufficient for conscious states comes from the case of Milner and Goodale's patient DF, a victim of carbon monoxide poisoning that resulted in bilateral lesions to lateral occipital cortex. DF's lesions gave rise to

visual form agnosia, a condition in which DF reports being unable to see objects, especially aspects of objects concerning their shape or form. In brief, patient DF seems not to be visually conscious of the form of objects. Nonetheless, it can be demonstrated that in spite of this lack of visual consciousness she is able to make use of certain unconscious representations of visual information about the form of objects in order to guide her actions. One demonstration of DF's condition is her performance on a task in which she had to put a card into a slot that could be variously oriented. For a given orientation of the slot, DF was able to orient the card correctly and post into the slot. DF's performance on this task was about as good as normal subjects. However, when asked to not put the card into the slot but instead to merely report on the orientation of the slot by holding the card in a comparable orientation, DF's performance was quite poor compared to normal subjects. The aspect of DF's performance that is relevant for our purposes is that while DF was not conscious of the orientation of the slot, her successful performance on the task indicates that her nervous system had unconscious egocentric representations of the slot. That she was able to correctly post the card into the slot demonstrates not only that she had representations of the orientation of the slot, but also that the orientation was represented relative to her hand and was thus egocentric [15].

One final consideration against regarding consciousness as purely egocentric involves making note of how frequently conceptual knowledge can affect what it is like to have various conscious experiences. For example, a pattern of black splotches on a white page can suddenly resolve as an image of a dog for someone who has a concept of a dog. The concept of a dog and the conceptual knowledge of what dogs are involve allocentric representations. The categorical knowledge that dogs are furry need not encode any information about the current relations of any dogs to oneself. Nonetheless, the way in which conceptual knowledge can be brought to bear on perceptual experience shows that conscious experience is not solely a matter of egocentric representation.

Just as consciousness is not solely egocentric, neither is it solely allocentric. One consideration in favor of this view is that visual consciousness is perspectival in the sense of embodying a pictorial perspective. The different features that characterize perspective in paintings and photographs also characterize a key feature of visual consciousness. For example, the visual perception of a row of three houses, like a picture, contains information about the locations and distances relative to the viewer in a way that the purely allocentric thought that there are three houses does not [8].

Another reason for believing that purely allocentric representations are insufficient for phenomenal consciousness comes when we realize that many propositional attitudes lack phenomenal character. Consider, for example, one's belief that  $\pi$  is an irrational number. It is implausible to suppose that this thought has any particular phenomenal character associated with it. As Jackendoff suggests, any apparent phenomenal character of the thought actually is the phenomenal character of associated sensory imagery, not of the thought itself [28]. And leaving the question of phenomenal character to the side momentarily, we can note that one can have a belief without it being a conscious belief. For example, the reader may have believed for many years that  $\pi$  is an irrational number, but this fact was probably not in the forefront of the reader's consciousness until this paragraph. Thus one can have allocentric representations (in this case, representations of  $\pi$  and irrationality) without those representations thereby being conscious. This is not to say that allocentric contents can never enter into consciousness. Indeed, the point of the above discussion concerning the influence that conceptual knowledge of dogs can have on visual perceptions of dogs was to demonstrate that allocentric contents do enter into conscious experience. The point here is that conscious experience is never solely allocentric.

Since conscious experience is never solely allocentric or solely egocentric, it is never to be found at either end of the allocentric-egocentric hierarchy. This lends plausibility to the hypothesis that conscious states are to be identified with representations at an intermediate level of the processing hierarchy. Additional evidence comes from research on the neural correlates of consciousness in binocular rivalry. In binocular rivalry research, human and animal subjects are presented with contradictory stimuli to their eyes, such as horizontal stripes to the left eye and vertical stripes to the right eye. While two stimuli are presented, both stimuli do not enter into the conscious percept but instead compete in the following way. At one moment the subject will see only the vertical stripes and at another moment the subject will see only the horizontal stripes. Neuroscientific investigations look for which neural activations seem most closely associated with the conscious percept. Monkeys can be trained to indicate which of the two stimuli they are aware of at any given time and single cell recordings can indicate whether activation in a cell is correlated with the conscious percept. Logothetis [29] found that among monkey cortex cells associated with the conscious percept, 90% were in inferotemporal cortex whereas only 40% were in extrastriate cortex (regions of cortex adjacent to area V1).

## 6 The Allocentric-Egocentric Interface Theory of Phenomenal Consciousness: Empirical Evidence

According to the AEI theory, not only are conscious states to be identified with representations at the intermediate level of the egocentric to allocentric processing hierarchy, they are, more specifically, to be identified with representations for which there is a mutual influence between egocentric and allocentric representations. That is, conscious states are hybrid representations in which there is both-bottom up influence of egocentric representations on allocentric representations and top-down influence of allocentric representations on egocentric representations.

Evidence for the reciprocal interaction between egocentric and allocentric representations comes from multiple sources. Already noted was the way in which conceptual knowledge can influence the nature of a perceptual experience. Additional evidence comes from studies of the relative contributions of low and high levels of the processing hierarchy conducted by Pascual-Leone & Walsh [17]. They applied precisely timed pulses of trans-cranial magnetic stimulation to different regions of visual cortex so as to test which areas seemed necessary for a conscious percept. In particular they looked at the relative contributions of area V1 and the relatively higher-level adjacent area known as MT or V5. Activity in neither area was alone sufficient for a conscious percept (a perception, in this case, of a moving stimulus). The conscious percept arose only when information was allowed to feedback from MT to V1.

We can relate the allocentric-egocentric interface proposal to the phenomenon of motion induced blindness. As already mentioned above, parietal areas-relatively intermediate in the processing hierarchy-are implicated in the phenomenon. Especially noteworthy are the contributions of relatively allocentric representations to the phenomenon. Bonnef and Cooperman [30] investigated what frames of reference seemed most implicated in the motion induced blindness and found that head-centered and object-centered mechanisms are involved in the disappearance effect.

Another promising line of evidence concerning the role of higher-level processes concerns the processes implicated in the kinds of learning that seem to involve consciousness. For example, there is evidence from fear conditioning studies that trace learning but not delay learning depends on consciousness. In the trace learning, there is a time gap between the conditioned stimulus and the unconditioned stimulus and in delay learning the two stimuli overlap. Additionally, it has been suggested that trace but not delay learning depends critically on hippocampus and certain prefrontal



structures [31].

The question arises of whether Milner and Goodale's dual systems theory of vision is inconsistent with the AEI theory of consciousness. One way of seeing a tension between the two accounts involves reading Milner and Goodale's view as the hypothesis that consciousness arises only in ventral stream processes and never in dorsal stream processes whereas the AEI theory allows that consciousness (at least sometimes) involves parietal processing. Two main points need to be made to ward off any threat that might be posed by Milner and Goodale's account. First we need to see that parietal areas do indeed sometimes get implicated in conscious states. The second point is to give an account of what distinguishes the occasions in which parietal processing affects consciousness and when it does not. Regarding the first point, it has already been noted that motion induced blindness may be modulated by transcranial magnetic stimulation of parietal areas. Further, parietal activity is implicated in conscious motor imagery [32,33]. Regarding the distinction between conscious and unconscious parietal activity, the distinction can be drawn as follows: direct projections from parietal areas to pre-motor areas do not result in conscious states, whereas projections from parietal areas to pre-motor areas via prefrontal cortex do give rise to conscious states [34]. This fits nicely with the Allocentric-Egocentric Interface theory given the role frontal cortex plays as a high-level area of visual processing implicated as a locus of allocentric representations.

## 7 Philosophical Implications of the Theory

In this section I turn to spell out some of the philosophical implications of the Allocentric-Egocentric Interface theory of phenomenal consciousness. It will be useful to first relate the empirical theory to the philosophically motivated terminology described in the second section of this chapter. Recall that the basic gist of the theory at hand is an account of state consciousness. A conscious state is comprised of two mutually interacting representations (via feed-forward and feed-back connections), one of which is relatively more allocentric and the other more egocentric. With respect to the question of transitive consciousness, the question of what a person is thereby conscious of in virtue of having a conscious state is that they will be conscious of whatever the allocentric and egocentric representations involved are representations of. What one is conscious of is always the content of both egocentric and allocentric representations. Regarding the question of phenomenal character, the account that emerges is that phenomenal character is identical to the representational contents of the implicated repre-

sentations. Further details concerning these issues will come to light in the following discussion of how the AEI theory compares to some of its main philosophical competitors. I turn now to discuss comparisons between AEI and, on the one hand, the First-Order Representational theories of consciousness favored by proponents of the transparency thesis and, on the other hand, Higher-Order Representational theories favored by proponents of the transitivity thesis.

First-order representational theories of consciousness attempt to explain consciousness in terms of first-order representations, that is, representations that, whatever they represent, do not represent other representations. Higher-order representations are representations of representations. A representation of a first-order representation is a second-order representation and a representation of one of these is a third-order representation and so on. First-order representationalists are especially fond of the transparency thesis that when one has a conscious experience, all that one is conscious of is what the experience is an experience of. First-order representationalists presuppose the transparency thesis as intuitively obvious and utilize it to justify their claim that conscious states have only first-order representational content and that phenomenal character is identical to the representational content of these first order states. If the transparency thesis turns out to be false then first-order representationalism turns out to be false for the following reason. If one were able to be conscious of something other than what an experience is an experience of (like vehicular properties of the state itself) then what it is like to be in a conscious state (a conscious state's phenomenal character) is something other than the just the first-order representational content of that state. Phenomenal character would thus include either the vehicular properties of experiences or the representational contents of the higher-order states in virtue of which one is conscious of the experiences.

The entering wedge of a case against the transparency thesis begins by noting the way that allocentric representational content, especially conceptual content, can influence what it is like to have a particular experience. So, for example, what it is like to look at a ladybug and conceive of it as an example of *Hippodamia convergens* is, intuitively, quite different from what it would be like to conceive of it as one's reincarnated great-great-grandmother. This in and of itself is not a threat to the transparency thesis, since representations of ladybugs, *Hippodamia convergens*, reincarnation, and great-great-grandmothers need not be anything other than first-order representations. However, the possibility that conceptual content can en-

ter into what it is like to have an experience opens up the possibility that higher-order conceptual content can enter into what it is like to have an experience.

To spell out the possibility just described it will be helpful to spell out some of the conditions sufficient for violating the transparency thesis. First-order representationalists read the transparency thesis as saying that when an experience is conscious one can only be conscious of what the experience is an experience of and thus one cannot be conscious of the experience itself. The question arises, of course, of what it means to be conscious of the experience itself and here the answer is best understood if we grasp the distinction between representational contents and representational vehicles. We can illustrate the distinction in terms of an analogy concerning non-mental representations. The English sentence “an orange cat is on a red mat” has as its representational content an orange cat’s being on a red mat. Vehicular properties of the sentence would include what font and ink color it is printed in. Even though the sentence represents an orange thing being on a red thing, the sentence itself — the representational vehicle — need not be printed in either red or orange ink. Returning to the topic of mental/neural representations, consider that a conscious experience can have as its content an orange cat being three feet away while the experience itself — the representational vehicle — is a state of the nervous system and is thus neither orange nor three feet away (nor is it a cat). Returning to the transparency thesis with the content/vehicle distinction in hand we can see that part of what it is denying in denying that we are conscious of experiences themselves is that it is denying that we are conscious of vehicular properties (on the assumption, of course that the vehicular properties of a representation just are whatever properties it has other than its content). However, if we combine the suggestion that conceptual contents enter into the phenomenal character of experience with the suggestion that some of our concepts can be concepts of certain vehicular properties, the suggestion presents itself that we can indeed be aware of properties of experiences themselves. Thus, following Churchland’s suggestion of the possibility of the direct introspection of brain states [35], if a person had the conceptual knowledge that consciously perceiving motion involved reciprocally influencing activity in areas V1 and MT, and acquired the skill of being able to automatically and without conscious inference apply that conceptual knowledge to experience, then that person would be able to be conscious of the vehicular properties of that experience. One consequence of this view that concerns phenomenal character is that when brain states are directly

introspected it is not the vehicular properties of experiences that contribute to phenomenal character but instead the representational content of the introspective states (which, of course, represent vehicular properties) that contribute to phenomenal character.

The above remarks spell out the falsity the transparency thesis in terms of the fact that when we have conscious states we are sometimes able to be conscious of the states themselves. This, however, is not to endorse the transitivity thesis that requires that we are always conscious of our conscious states. Indeed, I believe the transitivity thesis to be false. The Allocentric-Egocentric Interface theory is an empirically warranted theory which is logically consistent with the falsity of the transitivity thesis. Further, there are philosophical reasons for being suspicious of the transitivity thesis.

First off, according to advocates of the transitivity thesis it is supposed to be intuitively obvious that it is a requirement on having a conscious state that one is conscious of that state [1,36]. If the transitivity thesis is true it should be obviously incorrect to say of a state that it was conscious before any one was conscious of it. However, if we consider a particular example, it seems that the transitivity thesis is not obviously correct (which is not, of course, to say that it is obviously incorrect). Consider, for example, how one might describe what happens in motion induced blindness experiments when the yellow dots pop into and out of consciousness. It seems equally plausible to say either (1) that first the perception of the yellow dot becomes conscious and then you become conscious of your perception of the yellow dot or (2) the perception of the yellow dot becomes conscious only if you also become conscious of your perception of the yellow dot. If the transitivity thesis were pre-theoretically obvious, then option (1) would be obviously incorrect and (2) would be obviously correct. However, since neither (1) nor (2) seem obviously correct (or obviously incorrect), the transitivity thesis is not pre-theoretically obvious.

A second consideration that casts suspicion on the transitivity thesis concerns how easily we can explain whatever plausibility it has without granting its truth. We can grant that the transitivity thesis may seem plausible to very many people but explain this as being due to the fact that counterexamples would not be accessible from the first-person point of view. If we ask a person to evaluate whether the transparency thesis is true, they will call to mind all of the conscious states of which they have been conscious. But this can not constitute conclusive proof that conscious states are necessarily states that their possessor is conscious of. Consider

the following analogy. Every tree that we have ever been aware of is, by definition, a tree that we have been aware of. But this is not due to the definition of being a tree, but only due to the definition of being aware of it. The fact that every tree that we are aware of is a tree of which we have been aware cannot constitute proof that trees are essentially the objects of awareness or that no tree can exist without our being aware of it. By analogy we should not conclude from our being conscious of all of our conscious states that we have been aware of from the first-person point of view that all conscious states are *necessarily* states that we are conscious of. We should instead view our first-person access to conscious states as a way of picking out a kind of state that we can further investigate utilizing third-person methods. The description “states we are conscious of ourselves as having” thus may be more profitably viewed as a contingent reference fixer of “conscious state” that leaves open the possibility that it is not part of the essence of conscious states that we are conscious of them. Instead, the essence of conscious states is that they are hybrid representations that exist in the allocentric-egocentric interface.

### Acknowledgements

This work was supported in part by grants from The National Endowment of the Humanities and The James S. McDonnell Foundation’s Project for Philosophy and the Neurosciences. I am grateful to audiences of versions of this material at the University of Quebec in Montreal; the Heron Island, Australia meeting of the McDonnell Project of Philosophy and Neuroscience; the City University of New York Graduate Center Cognitive Science Symposium and Discussion Group; the William Paterson University Department of Psychology; and the ZiF conference on Endophysics, Time, Quantum and the Subjective. I am especially grateful for feedback from the following individuals: Ralph Abrams, Kathleen Akins, Omid Bayati, Ned Block, Mike Collins, Luc Faucher, Thomas Hogan, Doug Meehan, Rachelle Nashner, Anatoly Nichvoloda, David Rosenthal, Alex Vereschagin, and Josh Weisberg.

### References

1. Rosenthal, D. (1993) Thinking that one thinks, in M. Davies and G. Humphreys (eds.), *Consciousness*, Basil Blackwell, Oxford, pp. 197–223.
2. Rosenthal, D. (1993) State consciousness and transitive consciousness, *Consciousness and Cognition* **2**(4), 355–363.
3. Dretske, F. (1995) *Naturalizing the Mind*, MIT Press, Cambridge (MA).

4. Tye, M. (1995) *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*, MIT Press, Cambridge (MA).
5. Tye, M. (2000) *Consciousness, Color, and Content*, MIT Press, Cambridge (MA).
6. Block, N. (1995) On a confusion about a function of consciousness, *Behavioral and Brain Sciences* **18**(2), 227–288.
7. Chalmers, D. (1996) *The Conscious Mind*, Oxford University Press, New York.
8. Mandik, P. (2001) Mental representation and the subjectivity of consciousness, *Philosophical Psychology* **14**(2), 179–202.
9. Pöppel, E., Held, R., and Frost, D. (1973) Residual visual functions after brain wounds involving the central visual pathways in man, *Nature* **243**, 295–296.
10. Stoerig, P., and Cowey, A. (1989) Wavelength sensitivity in blindsight, *Nature* **342**, 916–918.
11. Stoerig, P., and Cowey, A. (1992) Wavelength discrimination in blindsight, *Brain* **115**, 425–444.
12. Weiskrantz, L. (1995) Blindsight: not an island unto itself, *Curr. Directions Psychol. Sci.* **4**, 146–151.
13. Weiskrantz, L. (1996) Blindsight revisited, *Curr. Opin. Neurobiol.* **6**(2), 215–220.
14. Bonneh, Y., Cooperman, A., and Sagi, D. (2001) Motion induced blindness in normal observers, *Nature* **411**(6839), 798–801.
15. Milner, A., and Goodale, M. (1995) *The Visual Brain in Action*, Oxford University Press, New York.
16. Olson, C., Gettner, S., and Tremblay, L. (1999) Representation of allocentric space in the monkey frontal lobe, in N. Burgess, K. Jeffery, and J. O’Keefe (eds.), *The Hippocampal and Parietal Foundations of Spatial Cognition*, Oxford University Press, New York, pp. 359–380.
17. Pascual-Leone, A., and Walsh, V. (2001) Fast backprojections from the motion to the primary visual area necessary for visual awareness, *Science* **292**, 510–512.
18. Levine, M., and Shefner, J. (1991) *Fundamentals of Sensation and Serception* (2nd edition), Brooks/Cole, Pacific Grove (CA).
19. Hubel, D., and Wiesel, T. (2001) Brain mechanisms of vision, in W. Bechtel, P. Mandik, J. Mundale and R. Stufflebeam (eds.), *Philosophy and the Neurosciences: A Reader*, Basil Blackwell, Oxford, pp. 179–198.
20. Andersen, R. (1995) Coordinate transformations and motor planning in posterior parietal cortex, in M. Gazzaniga (ed.), *The Cognitive Neurosciences*, MIT Press, Cambridge (MA), pp. 519–532.
21. Batista, A., Buneo, C., Snyder, L., and Andersen, R. (1999) Reach plans in eye-centered coordinates, *Science* **285**, 257–260.
22. Eichenbaum, H., Stewart, C., and Morris, R. (1990) Hippocampal representation in spatial learning, *Journal of Neuroscience* **10**, 331–339.

23. O'Keefe, J. (1993) Kant and the sea-horse: an essay in the neurophilosophy of space, in N. Eilan, R. McCarthy and B. Brewer (eds.), *Spatial Representation: Problems in Philosophy and Psychology*, Blackwell, Oxford, pp. 43–64.
24. O'Keefe, J. (1994) Cognitive maps, time and causality, *Proceedings of the British Academy* **83**, 35–45.
25. Taube, J., Muller, R., and Ranck, J. (1990) Head direction cells recorded from the postsubiculum in freely moving rats. Description and qualitative analysis, *Journal of Neurosciences* **10**, 420–435.
26. Mandik, P. (2002) Synthetic neuroethology, *Metaphilosophy* **33**(1/2), 11–29.
27. Mandik, P. (2003) Varieties of representation in evolved and embodied neural networks, *Biology and Philosophy* **18**(1), 95–130.
28. Jackendoff, R. (1987) *Consciousness and the Computational Mind*, MIT Press, Cambridge (MA).
29. Logothetis, N. (1999) Vision: a window on consciousness, *Sci. Am.* **281**, 69–75.
30. Bonnef, Y., and Cooperman, A. (2003) Motion induced blindness is affected by head-centered and object-centered mechanisms, *Journal of Vision* **3**(9), 221.
31. Carter, R., Hofstotter, C., Tsuchiya, N., and Koch, C. (2003) Working memory and fear conditioning, *Proc. Natl. Acad. Sci.* **100**(3), 1399–1404.
32. Decety, J., Perani, D., Jeannerod, M., Bettinardi, V., Tadary, B., Woods, B., and Mazziotta, J. (1994) Mapping motor representations with PET, *Nature* **371**, 600–602.
33. Grafton, S., Arbib, M., Fadiga, L., and Rizzolatti, G. (1996) Localization of grasp representations in humans by PET: 2. Observation compared with imagination, *Experimental Brain Research* **112**, 103–111.
34. Boussaoud, D., di Pellegrino, G., and Wise, S., (1996) Frontal lobe mechanisms subserving vision-for-action versus vision-for-perception, *Behav. Brain. Res.* **72**, 1–15.
35. Churchland, P. (1979) *Scientific Realism and the Plasticity of Mind*, Cambridge University Press, Cambridge.
36. Lycan, W. (1995) *Consciousness and Experience*, MIT Press, Cambridge (MA).