

Penultimate English language draft of Mandik, Pete. (in press). L'accomplissement Neural de L'objectivité [The Neural Accomplishment of Objectivity]. In: Elizabeth Ennen, Pierre Poirier, Luc Faucher, and Eric Racine (eds.) *Des Neurones à la Philosophie: Neurophilosophie et philosophie des neurosciences* [*Neurons with philosophy: Neurophilosophy and philosophy of the neurosciences*]. Paris: DeBoeck Université.

The Neural Accomplishment of Objectivity

Pete Mandik

1. Introduction

Philosophical tradition contains two major lines of thought concerning the relative difficulty of the notions of objectivity and subjectivity. One tradition, which we might characterize as “Cartesian”, sees subjectivity as comparatively less problematic than objectivity. On the Cartesian view, what we know best of all are the contents of our own minds and the major problematic is to pierce the veil of appearances and make contact with objective mind-independent reality. In contrast is a line of thought that reverses the order of difficulty. A pervasive materialistic and scientific mind-set takes objectivity as the unproblematic starting point. From this point of view, widespread through much of contemporary philosophy and especially explicit in the philosophy of mind, a world of physical, chemical, and biological events is taken as relatively given. The problematic here then is to make sense of any kind of genuine subjectivity within this physicalistic framework (See Mandik 2001 for an extended discussion of these points.)

One might expect neuroscientists and neurophilosophers alike to belong exclusively to this latter tradition, given their proclivity for seeing the mind as being intimately tied to, if not identical to, the brain—a physical thing presumably exhaustively describable in the objective idiom of physicalistic science. However, this is not so. Many practitioners of things neural count among adherents of what I have described as a Cartesian line of thought. This is especially clear when we recognize that the neural equivalent of the subjective/objective distinction is the egocentric/allocentric distinction. Egocentric representations, associated especially with activity in Posterior Parietal Cortex, code for things in “self-centered” reference frames. Allocentric representations (alleged by many to be involved in Hippocampal activity) in contrast, code for things in “other-centered” reference frames. Cartesians in neuroscience and neurophilosophy cast the egocentric as the relatively basic and unproblematic of the two sorts of neural representation. From this view, then, the allocentric is seen as especially difficult, and, under certain descriptions, impossible. My purpose in this paper is to review and ultimately counter this Cartesian line of thought.

It will be helpful to begin by briefly reviewing the distinction between objectivity and subjectivity. (For a longer review, see Mandik 1998.) Objective things exist independently of us, or more specifically, independently of our minds. Objective things are the things that would “be there anyway” that is, they would be there regardless of our representational capacities to perceive or think of them. In contrast, subjective things depend on our representational capacities. If beauty exists in the eye of the beholder and, as Hamlet says, “nothing is good or bad but thinking make it so”, then beauty, goodness, and badness are subjective. They are representation-dependent. There is, of course, a straightforward, indeed tautologous, sense in which representations themselves are representation dependent: representations exist if and only if representations exist. However, this does not mean that the idea of an objective representation is utter nonsense or that there is no sense in which representations can be anything but subjective. The sense in which it makes sense to call some representations “objective” and others “subjective” is the sense in which some representations are about objective things and others are about subjective things. Subjective representations, then, are egocentric representations. They are self-centered. They represent objects, properties, and relations that depend on the representing subject. In contrast, objective representations are allocentric or other-centered. They represent objects, properties, and relations that do not depend on the representing subject.

The purpose of this paper is twofold. First, I aim to sketch an account of the neural bases of objective and subjective representation. Second I aim to defend the account against a certain class of philosophical objections that dismiss the account as rendering objectivity too easy. Such objections conceive of objectivity as either a major cognitive accomplishment by a select few creatures or, at the extreme of difficulty, as an impossible ideal.

2. Egocentric representations

Let us tackle the topic of egocentric and allocentric representations by first starting with 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” (Levine and Shefner, 1991, p. 671). Hubel and Wiesel conducted classic early work on the receptive fields of various neurons in the 1950’s (Hubel and Wiesel 2001, p. 184). Retinal ganglion cells and neurons in the lateral geniculate nucleus have circular fields with either an excitatory center and an inhibitory surround 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.

Retinocentric representations are but just one kind of egocentric representation. Anderson (1995) describes 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.

Cells with retinocentric, head-centered, or body centered coordinates are not limited to brain areas discussed so far, but may additionally be found in cortical and subcortical regions such as the parietal lobes more generally, prefrontal cortex, and the superior colliculus (Andersen *et al.*, 1993; Galletti *et al.*, 1993).

The egocentric representations described above involve responsivity 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 motor neurons we may define “effective fields”: the region of space that a motor neuron effects (Mandik 2001, p. 191). Thus, for example, reach plans for arms might be encoded in egocentric coordinates. Indeed, Batista *et al.* (1999) find that reach plans are encoded in eye-centered coordinates.

Such egocentric representations are thought to be wide-spread and figure heavily in perception and action. That is, they are discernable not only from neuroscientific points of view, but match up with much that may be discerned from a more overtly psychological perspective. A rich class of psychological examples comes from the literature on Gibsonian ecological approaches to psychological understanding. The gist of these examples will involve psychological states that have as their contents things that depend on the organism itself. For instance, if I see something as too large for me to eat, then there is a real sense in which the property I attribute to that object depends on me. Whether something instantiates the property of being too large for me to eat depends on facts about me—my size and metabolism etc. Gibson (1979) gives several examples in which the organism itself, or parts of the organism, or properties that depend on the organism figure in the contents of visual experience. In a discussion of these points, Bermudez (1998) calls such contents ‘self-specifying information’. Bermudez offers the following examples. The first example in which self-specifying information shows up in the contents of visual experience concerns the boundaries of the visual field. In the case of humans, the visual field contains less than half of the viewer’s visible environment. The boundary of the visual field constitutes a boundary between the seen and unseen world that is movable at will and supplies information to the viewer about her location in and trajectory through that world. Parts of the viewer’s body (especially the nose and cheeks) define the boundary of the visual field, and some parts, the nose in particular, is itself seen in the visual field. The occlusion of other seen objects by the nose supplies information about degrees of proximity of objects to the viewer. Another important proximity clue involves the amount of the visual field an object occupies. As objects move away from the viewer, they occupy less of the visual field. This fact allows for an interesting way of visually distinguishing between

parts of the subject's body and the rest of the seen world. The viewer's hand or foot can only move so far from the viewer, being typically (and hopefully) attached to the viewer, and thus there is a limit on how little of the visual field it can be restricted to. In contrast, other objects, being unattached, may move indefinitely far away from the viewer and thus occupy an indefinitely small portion of the visual field. This difference may also facilitate a clue as to what objects, though not a part of the organism's body, are at least attached to it or grasped by it. A stone held in my hand cannot occupy an indefinitely small portion of the visual field. Only if it is let go may it shrink into the distance. Another source of self-specifying information is optical flow. As a subject moves forward, the shapes in the visual field flow from the center to the boundaries. The center from which the flow emanates specifies the direction in which the subject is moving. These examples drawn from Gibson constitute rich sources of information available to the perceiving organisms. The step from the availability of this information to subjective representation is a small one. A common way of thinking of representation, especially in neuroscientific contexts, is that neural events represent environmental events in virtue of having the function of carrying information about the environmental events. The notion of information here is typically explicated in terms of causal dependence. The notion of function here is typically explicated in terms of natural selection. A neural event such as the firing of a detector cell carries information about an environmental event in virtue being caused by the environmental event. It has the function of carrying that information in virtue of being naturally selected to carry that information. If it has that function, then it not only merely carries information, but counts as a genuine representation. (See Mandik (2002, 2003) for further discussion of these issues.) If we suppose that organisms have states that function to carry that information, and that information includes information about the relation of the environment to the subject, then we have on our hands representations that are perspectival/egocentric/subjective.

All of the above examples of egocentric representations are instances of spatial representations. However, egocentricity may also be characterized for nonspatial representations. As discussed in greater length in Mandik 2001, representations of temperature may be regarded as egocentric, though this egocentricity is not characterized spatially.

A given temperature sensation does not just represent a temperature of a region on or near the skin but represents temperatures as being of varying degrees of hazard or harmlessness to the subject's tissues. The output of a thermoreceptor in response to a given temperature does not represent a given temperature *per se* but instead whether the given temperature is, e.g., too hot, too cold, or just right. The property of being too hot cannot be defined independently of answering the question "too hot for whom?" . . . Thermoreceptors include in the representational contents of their outputs relations that the temperatures bear to the representing subject, much in the way that retinocentric representations of spatial locations represent locations defined relative to the subject (Mandik, 2001 pp. 195 - 196).

The above remarks show that egocentric neural representations are thought to underlie many mental states, most obviously, those involved in perception, but also those involved in certain memories and intentions. However, there are reasons to think that 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 ourselves. 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 are more massive than electrons is not particularly a fact about me, in spite of the fact that I am chock full of neutrons and electrons. How can we account for this capability to have detached, objective mental states? More specifically, how can we account for this capability in neural terms? In other words, are there any genuinely allocentric representations?

3. Are there any genuinely allocentric representations?

The obvious place to look for proposals of the neural bases of objectivity is in neuroscientific literature on allocentric representations. The largest concentration of such work focuses on the hippocampus and nearby regions. The classic studies in these areas concern the spatial navigational capabilities of rats, especially comparisons of performance of rats with and without lesions to 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. Eichenbaum et al. (1990) performed experiments in which 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.

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. The main question addressed in this section is whether the hippocampus' contribution to spatial navigation utilizes genuinely allocentric representations or whether closer examination will reveal that the neural representations in hippocampus are ultimately egocentric.

One line of thinking that leads to viewing hippocampal representations as allocentric depends on first grasping a distinction between abilities that are perspective dependent and abilities that are perspective independent (Cussins, 1990). Cussins asks us to imagine two different abilities employed to arrive at some unseen destination in a city. The first, perspective dependent, ability would allow you to get to that destination, but only if you started at a particular point and then proceeded to follow a particular path. If your ability were maximally perspective dependent then you would be unable to arrive at the destination if you were to deviate from the original course or begin from a different starting location. In contrast, a perspective independent ability would allow you to get to your destination from any starting location.

Given the notion of perspective dependent and independent abilities, we can now begin to make sense of the notion of perspective dependent and independent representations or representational abilities. Imagine, then, that the representational repertoire of the perspective dependent urban navigator is something akin to the list of directions one might be told for getting to a party. They would be on the order of “starting at Moe’s tavern, go about five blocks down Green Street and turn left down the alley that has a blue van parked near its entrance”. Such a representation would be relatively useless to you if you started anywhere other than at Moe’s tavern. The representation employed by another navigator, however, might be something akin to a map of the entire city. This map represents all the relations that the destination bears to all other locations in the city, allowing the map user to get to the party from almost anywhere in the city. Let us contrast, then, the respective ways the list of directions and the map represent the location of the party. The perspective dependent representation of the party is a representation from the Moe’s tavern perspective. The map user, on the other hand, can represent the party from the Moe’s tavern perspective, but may also represent the location of the party from many other perspectives as well.

Thinking of hippocampal functioning as the use of a map as opposed to the list of indexical representations helps to see how the representations involved might be less egocentric than others insofar as they are less tied to particular perspectives. However, it is not clear that the metaphor of the map is sufficient to give a sense of anything genuinely allocentric in the sense of a representation that is genuinely non-egocentric. To see why, consider that in the literal viewing of the map, one adopts a point of view on a city, akin to, for instance, a view from up above the city: a bird’s eye view. Being a pictorial representation, it necessarily encodes information about where the viewer would have to stand in order to have an experience with content similar to the picture (Mandik 2001 pp. 186 – 188). The problem with the map metaphor, then, is that it does not give an example of a representation that abstracts away from points of view altogether.

However, the problem that plagues the metaphor may not necessarily apply to the proposal concerning hippocampal functioning. To attribute “cognitive maps” to creatures with intact hippocampi may not necessarily commit one to the postulation of a mental version of a literal map, where being a literal map involves being a visual representation of a spatial layout seen from some particular point of view or other.

To better assess whether the hippocampal representations subserving navigation are genuinely allocentric, it will help to consider a detailed model of hippocampal functioning. One especially prominent proposal is the slope-centroid system postulated by O'Keefe (1993, 1994). 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. For example, if the objects are arranged in a roughly ovoid shape, the slope would be the line connecting the oval's foci. 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, comparison 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. 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 (Taube, Muller and Ranck, 1990). The representation of speed is thought to be implemented by the frequency of sinusoidal oscillations of the hippocampal EEG called the "theta pattern" (O'Keefe, 1993).

The vectors that code for the location of the animal and other objects are thought by O'Keefe to be computed in the following manner. The vectors involve two quantities: the magnitude and the phase angle in a polar framework. Representation of magnitude is implemented by the number of spikes in a burst of neural activity. Representation of phase angle is implemented by the timing of the onset of the burst relative to the cycle of the theta pattern. Thus, according to O'Keefe, each pyramidal cell is a harmonic oscillator in which the amplitude of oscillation codes for the current distance to an object or location and the phase shift relative to the theta pattern wave codes for the angle of the vector in polar co-ordinates. Computations involving vector addition are accomplished by the addition of the vector's sinusoidal representations. Vector subtraction involves the addition of a sinusoid to the 180-degree phase reversal of the other sinusoid. (O'Keefe 1993, p. 55).

4. The Cartesian case against the allocentricity of the slope-centroid system

Grush 2001 questions whether the slope-centroid system utilizes genuinely objective spatial representations. Grush's arguments depend, in part, upon the introduction of a kind of representation intermediate between an egocentric representation and a genuinely allocentric representation that Grush calls "alter-ego centric". The basic idea of contrast between egocentric representation and alter-ego-centric representation might best be conveyed in terms of a difference between perception and imagination. So, in perception, I might see aspects of a certain landmark or building in relation to my actual location: for instance I can see the front of the building but not the top or the back. An example of an alter-ego-centric representation would be if one were to, without changing their actual location with respect to the building, imagine inhabiting an alternate position with respect to it by, for instance, imagining floating above the building. The imagined scene would still present things as they might appear in relation to me, and thus resemble egocentric representations, but in presenting counterfactual instead of actual relations things bear to me, they become, in Grush's phrase, *alter-ego-centric*. In contrast, then, full-blown allocentric representation would, I presume, contain absolutely no reference to either actual or counterfactual relations to the representing subject.

Grush's suggestion regarding the slope-centroid representational system is that it provides only alter-ego-centric representations of spatial position. Grush is of course live to one obvious objection to such a proposal. He writes:

how can this be compatible with the slope being defined by the objects in the environment? Simply because the slope may be no more than the preferred orientation in which the animal likes to populate its off-line egocentric space. It may be convenient, or perhaps efficient from a processing standpoint, when imagining a 2-D array of objects, to imagine them in such a way that their largest direction of spread lies on the left-right axis rather than the up-down axis. In other words, if one is going to adopt a surrogate point of view, from which to entertain the spatial relationships between objects in an alter-ego-centric reference frame, one needs to decide where to place this virtual point of view and how to orient it: if I am imagining being above San Diego so that I have a map-like view of it, shall I imagine being directly above downtown, with La Jolla in the periphery, or above La Jolla with downtown in the periphery? (Grush 2001, pp. 84-85).

Grush's argument in the quoted passage shows, at best, that alter-ego-centric representations are compatible with the slope specification output by the slope-centroid system. What they do not show, however, is that alter-ego-centric representations are employed within the slope-centroid system. Nor do they show that the slope-centroid system is devoid of genuinely allocentric representations.

Grush further argues that the representational system of the slope-centroid system would be quite useless if it was strictly allocentric.

. . . a purely structural description, sanitized of all egocentricity or POV effects, will be ambiguous between symmetric layouts, and thus presumably useless for guiding behavior or navigation. (p. 88)

Grush additionally provides positive evidence for thinking of the representations utilized in navigation as alter-egocentric. Grush notes that Posterior Parietal Cortex, widely implicated in the egocentric representation of space, when lesioned, results in disruptions of navigation skills “in some ways more severely than does damage to the hippocampus itself.” (p. 85). However, removing the creature’s eyes and/or feet may similarly rival hippocampal lesions in deleterious effects on navigational behavior. This just highlights the indisputable point that hippocampus is not the sole locus of navigational ability. It remains clear that hippocampus makes a contribution. Might the contribution be allocentric?

What Grush’s “Post parietal Cortex” argument shows is, at best, that egocentric and perhaps even alter-egocentric representations are necessary for navigation (I’m not saying the argument does show this, but only that if it showed anything it would be this.) What the argument doesn’t establish, however, is that egocentric and or alter-egocentric are alone sufficient (without some genuinely allocentric representations thrown in) for the sorts of navigational capabilities thought to be underwritten by the hippocampal based slope-centroid system.

Similar remarks apply to Grush’s arguments concerning the uselessness of genuine allocentric representations in navigation. One might grant that allocentric representations might be alone insufficient (without the help of some egocentric representations) for the accomplishment of certain navigational tasks, as in disambiguating between symmetric layouts. That is, one might grant this without thereby granting that there are no allocentric representations involved in the postulated navigational systems.

One set of issues that Grush’s remarks raise concerns the relative usefulness of allocentric and egocentric representations. Grush’s line of thought suggests that pure allocentric representations would be relatively useless. Consider: if you had a representation that there is a pot boiling over on a stove, you would be ill suited to do anything about it unless you knew where the pot and stove were relative to you. Is it in the room adjacent to the one you are in? Is it half way around the world? Plausibly, without some connection to egocentric representations, allocentric representation may be *relatively* useless. However, this does not mean that allocentric representations are *absolutely* useless. Why couldn’t it be useful to have a representation whose content was simply: the ball is in the box? Storing all sorts of extra stuff seems not worth it. Compare how much of your hard drive would be taken up by the text file containing the words “the ball is in the box” compared to a bitmapped photograph of a ball in a box.

It helps here to think of the representational mechanisms of organisms as having multiple distinct parts. Such a view need not commit one to the Fodorian theses of modularity (Fodor 1983) or Language of Thought (Fodor 1975). What such a view does commit one to is spelled out at some

length in Grush and Mandik (2002). For present purposes, the following points suffice. First, the various representational contents had by typical representing systems, such as human nervous systems, are *independently targetable*. The idea of independent targetability is that of the being able to represent different things in virtue of having distinct representational capacities. So, for instance, in imagining the moon to be a different shape than it is, one can represent it as being cubical, or oblong, all the while representing it as the moon. Conversely, one can keep the shape in question fixed while imagining it being had by objects other than the moon. One can shift from imagining the moon to be cubical to imagining what the Eiffel Tower would look like if distorted into a roughly cubical shape. Second, the most plausible hypothesis of how organisms have this capacity is in virtue of separate representational vehicles for the separate contents. That is, there are distinct physical processes internal to the organism in virtue of which the distinct contents are independently targetable.

With the idea of independent targetability in hand, we can see how the existence of allocentric representations in a cognitive system is perfectly consistent with the proposition that a cognitive system must have egocentric representations to behave meaningfully in the world. One need not, contra Cussins (1998), embrace the possibility of a purely allocentric cognitive system (“Fregean angels,” as Cussins’ calls them) in order to make sense of the notion of a purely allocentric representation. Since contents are independently targetable, the possibility opens up of a system having the capacity of independently targeting allocentric contents and egocentric contents. One possible implementation of such a system would be to have areas such as parietal cortex responsible for egocentric representations while hippocampal areas responsible for allocentric representations. Further still, slope-centroid may very well be considered as trafficking in such pure allocentric representations. Indeed Grush flirts with a similar view. . .

My proposal is that the hippocampus encodes, and remembers, the spatial relations between objects, perhaps only or primarily topographic information, in the form of many, perhaps partial and disunified representations of these objects as given in perception during exploration. This information can be accessed and processed by the parietal cortex, which creatively constructs an off-line alter-ego-centric space, and places the objects, including itself, in that space. (p. 85).

These ideas will be developed further below. First it will be useful to do some further ground clearing by considering a different line of argumentation against viewing slope-centroid representations as allocentric.

Another set of arguments against the allocentricity of slope-centroid come from the philosopher John Campbel (1993). Campbel construes the slope-centroid system as allowing the creature to “represent the vector from it to any arbitrary target recorded in the system.” (p. 81). Campbel argues, however, that this is insufficient to constitute genuinely allocentric representational capacities. According to Campbell, a criterion of allocentricity is that it “be possible to represent [spatial properties and relations] by having a disengaged picture of what is going on there”. (p. 81) However, a creature representing space using only the slope-centroid system

allegedly fails this criterion since “it cannot simply represent the vector between any pair of arbitrary places, regardless of whether or not it is thinking of itself as at either of those places” (p. 81). Campbell thus construes the slope-centroid system as having an indexical character. Because of this indexical character the creature “cannot represent the spatial relations between any two arbitrary places in its environment, irrespective of whether the organism itself is being supposed to be at one of them” p. 87.

I think that Campbell is mistaken here and his mistake depends on not adequately appreciating the slope-centroid’s use as a memory system. One thing Campbell seems to assume is that the tokening of representations in the slope-centroid system occurs only during its use in navigation. But this overlooks its use as a memory system. The relations between, say, the lamp and the box in the room are stored, and thus, represented, even though the creature is not currently in the room. They are stored then, even on occasions when the creature may have no idea where it currently is with respect to the room. So, for example, the fact that the ball is in between the box and the lamp is stored and thus represented even on occasions when the creature may have no idea where it is currently located with respect to the ball, box or lamp. The evidence that they are stored in an allocentric way is that they can be accessed from any arbitrary starting point when placed back in that room.

Consider these remarks in connection with the above discussion of the different sets of directions to the party. The behavioral evidence that the partygoer had indexical directions is the context dependence of his party finding ability. The one with the non-indexical representation was the one able to find his way from arbitrary starting places. It might be objected here that no partygoer can find the party without knowing where he himself is and further that this knowledge must itself involve an indexical mode of presentation. Thus, if a creature had only allocentric representations, then the creature would be unable to find its way. Note, however, in response to this objection, that I am not interested in defending the utility of having *only* allocentric representations. I am merely defending the utility of having allocentric representations in addition to egocentric representations. This defense is part of a larger argument that the hippocampal navigational system counts as a locus for where some of those allocentric representations might be.

As far as behavioral evidence goes, the ability to find the goal from arbitrary starting locations counts in favor of viewing the underlying representation as being allocentric. Of course, in the case of the partygoers, behavioral evidence is not all we have. We also can look at their lists of directions and see whether or not they are using indexical descriptions. Can we analogously look at the representations employed by hippocampus as hypothesized by the slope-centroid model?

Recall that in the slope-centroid model, the main things represented are place, heading and speed. Let us look at each to see how each is represented in the slope-centroid hippocampal system. Place is represented by pyramidal cell activity, with highest level of activity indicating that the location represented by that pyramidal cell is the same as the location where the creature currently is. The level of activity thus serves as a sort of “you are here” marker on a map. Of course, this description is making the pyramidal cell

representation sound indexical, and it might be, but note that other aspects of the pyramidal cell representations are non-indexical. Other locations are represented as well, for example, location b is represented and being in between a and c . And further, as mentioned earlier, the creature is capable of representing this in memory, regardless of whether it is at a , b , or c . Indeed, the creature can represent the between-ness relations of these three locations regardless of where in an indefinite number of locations the creature is with respect to them.

Recall that in the pyramidal cells, amplitude of oscillation codes for the current distance to an object or location and the phase shift relative to the theta pattern wave codes for the angle of the vector in polar co-ordinates. Again, considering these as components in a large scheme allows us to see how these components may be allocentric regardless of whether other components may be egocentric. Thus, for instance, a plausible hypothesis is that an amplitude of oscillation n codes for distance m regardless of which pyramidal cell the oscillation is occurring in and regardless of which objects, including the creature itself, are the relata in the distance relation. Similarly, a phase shift of p might represent an angle of q regardless of what objects are being localized in the scheme.

It is useful here to make note of a similarity of my proposal to the language of thought hypothesis (LOT). My intention here is not to endorse the hypothesis, but instead to construct an analogy. In a language of thought, like an ordinary spoken language, there are presumed to be components of thoughts each of which have their own meanings, analogous to words. So for instance, the thought expressed by “This cat now has fleas” is made out of parts, only some of which are indexical, namely “this” and “now”. “cat” and “fleas”, however are not indexical, at least not obviously so. The moral of this story is that representational systems have representational parts and the different parts may differ with respect to whether or not they are indexical parts or non-indexical parts.

It might be worth spelling out here why my view expressed here is not a species of the LOT defended by Fodor (1975). First, LOT is a hypothesis about propositional attitudes. Second, it construes the having of a propositional attitude that P as requiring the tokening of an inner representation, an inner sentence, with the content that P . Thus, one believes that George the cat is black in virtue of having an inner representation the content of which is that George the cat is black. Third, it construes the inner sentences as themselves being made of parts that individually have meaning. The resemblance between my account and LOT is limited to as resemblance to the third point: inner representations have parts that themselves are representations. Nothing I’ve said about hippocampal functioning commits me to attribute propositional attitudes to the hippocampus. Nor am I necessarily committed to the representations in hippocampus having a sentential structure, which I take to mean having a subject-predicate structure.

Another useful consequence of the analogy between my view and the language of thought hypothesis is that by moving away from a pictorial model of mental representation and toward a more linguistic one, it becomes even clearer how neural representations can manage to be purely allocentric. It is

important to guard against the line of thought that construes every representation as a view, that is, as a representation with pictorial perspective (Mandik 2001). Every picture is a view from somewhere, and if we construe allocentric representations as pictures with views from nowhere, then we have postulated something absurd. But words, unlike, pictures, need not encode anything about a point of view in any literal or non-trivial sense.

5. A more general view of objective and subjective neural representation

So far the line of argument in this paper has been as follows. First, I've countered arguments from Grush and Campbel that the hippocampus does not employ allocentric representations. Of course, merely exposing bad arguments against allocentric representations does little to prove that there actually are any allocentric representations in hippocampus. I suggested then, that the similarity in behavior between the various partygoers and the various nonhuman creatures suggests that analogous representations are used in each case. More specifically, behavioral evidence suggests that the rats exhibiting context independent behaviors are utilizing context independent representations.

However, the question arises of how, in general, to tell when a creature is utilizing allocentric representations. I argue that we can grant that creatures must utilize egocentric representations and that whatever allocentric representations they must have are used conjointly with egocentric representations. This involves conceiving of the organisms as having separate representational capacities: allocentric as well as egocentric. The question remains, however, of how to tell in general whether any allocentric representations are used.

Let us consider this question in the context of attributions of representations to the neural states of non-human creatures. Consider, for example, the considerations involved in attributing representations to the activity in the nervous system of a frog. Let us presume that with respect to frogs, flies are mind-independent: with respect to our science, at least, flies are frog independent. Let us also presuppose the old story about the fly detector (Lettvin *et. al.* 1959): that there is some event type that occurs in the frog's brain that is relatively well correlated with the presence of flies. Call that brain event *B*. Following Godfrey-Smith (1998, p.250) let us call the conditional probability that a fly is present given *B* 'reliability' and the conditional probability that *B* will occur given the presence of a fly 'sensitivity'. We can define the frog's hunger as the amount of time that has elapsed since it last ate a fly. Suppose, for the sake of argument, that upon observing the frog and its brain activity, we note the following. The frog's sensitivity to flies increases as a function of its hunger. The conditional probability that the presence of a fly will activate *B* is much higher when the frog is hungry.

Given these presumed facts let us question whether *B* really functions to represent flies (instantiations of the property of being a fly) or some other and perhaps mind-dependent (frog's-mind-dependent) property. *Ex hypothesi*, the property of being a fly is a frog-independent property. But consider the property defined as being a fly-when-a-frog-is-hungry (FWAFIH). FWAFIH is

pretty obviously frog-dependent. And while the frog is relatively sensitive to flies in virtue of brain activity B , the frog is even more sensitive to FWAFIH's. And arguably, what matters to frog survival is not so much the detection of flies simpliciter, but instead detecting flies when its hungry. That is, what matters more to the frog's survival than the detection of flies is the detection of FWAFIH's. Is it so large a leap to think that the frog is thus in the business of detecting FWAFIH's and not flies? I argue that this is indeed too large a leap. Having the function of detecting flies when a frog is hungry is distinct from the function of detecting the presence of fly-when-frog-is-hungry. But how are we to draw the distinction? How would one go about choosing between the hypothesis that the frog has neural representations of flies and the hypothesis that the frog has neural representations of FWAFIH's?

We can get a handle on this question by deciding an analogous question asked regarding scientific instruments instead of sensory systems. Consider the following question. How would one decide that an instrument, say a thermometer, was detecting an objective property temperature instead of the thermometer-dependent property *temperature-of-thing-thermometer-is-in/near*? The key to answering this question is by applying, in this context, a crucial distinction between the properties are that are explanatory, and the properties that need to be explained. If what needs to be explained is why a thermometer has a certain reading when it is in a substance, then the elements of an explanans that are restatements of what was said in the explanandum are strictly superfluous. Circularity is the death of explanation. "Grass is green because grass is green" is so obviously unexplanatory because it is so obvious that the properties picked out in the explanans are all and only the properties picked out in the explanandum. Molier's famous and funny example of explaining the fact that opium makes people sleepy in terms of positing a dormative virtue of the opium is both famous and funny because citing a dormative virtue of opium only restates what was known already: that opium makes people sleepy. What needs to be explained is the why the thermometer goes up when its placed in a substance. Invoking the property of temperature is a candidate explanation. But consider the property *temperature-of-substance-thermometer-is-in*. Note the overlap between the explanans and the explanandum. We already know that the thermometer is in the substance. Invoking the thermometer dependent property bears no explanatory load over and above that born by the invocation of the thermometer independent property. We are interested in the property that explains why, when the thermometer is in the liquid, the mercury goes up. We abstract away from the instrument. When you have a fever, you instantiate a property that is detected by the thermometer, and, let us grant, a property that depends on thermometers. But once we have the properties solely internal to the thermometer, and the properties external to the thermometer, no additional explanatory work is done by constructing properties that, while external to the thermometer, depend on the thermometer.

The analogy between mental representations and measuring instruments leads to another way of conceiving of the attributions of objective contents. Common scientific practice has it that lines of converging evidence are indicative of the truth of a hypothesis. The more lines of evidence that

converge, the more likely the hypothesis. This point holds not only at the level of theory, but at the level of data. How does one go about deciding that a putative datum is not just an artifact? For example, if using one kind of microscopy, one observes a bulge in a cell wall, one may wonder whether that bulge is just an artifact of the particular preparation. One tries multiple preparations to see if the bulge appears in more than one of them. (For an extended discussion of precisely this sort of issue as it arises in cell microscopy, see Bechtel 1994.)

Why is it that multiple lines of converging evidence are taken as indicative of objectivity? Because the agreement of instruments only makes sense on the supposition that they are in the business of detecting properties independent of the instruments. Likewise, using one instrument to calibrate another (and how else can instruments be calibrated?) only makes sense on the supposition that the two different instruments are tapping into something independent of the two instruments. If what each instrument was in the business of detecting was a property dependent on that instrument, then each separate instrument would be detecting its own separate property and there would be no common measurement that they could be said to be agreeing on (and there would be no room for the notion of calibration to get a grip).

This argument is distinct from the common one that treats the objectivity of what is measured as an inference to an explanation of the agreement of the instruments. The common argument assumes agreement, and then posits the objectivity of what is measured as an inference to a best explanation of why there is agreement. What I am proposing is different from this common argument. What I am proposing is that the very notion of agreement itself presupposes objectivity (in the sense of instrument-independence of the properties measured). If what is measured is instrument dependent, then the notion of agreement cannot grab hold because where there are different instruments, then there are different instrument-dependent properties that each is tracking.

Thus we have a sketch on how one might go about deciding in the favor of the objectivist option for attributing contents to the states of the frog. The sketch shows how notions of objectivity hook up with notions of consensus or agreement without being identical to those notions. Indeed, part of the notion of agreement between, say, measuring instruments presupposes a metaphysically objective thing for them to be in agreement about.

To tie this back into hippocampal function, consider the following. If a creature is able to represent the thing to the right as being the goal and on other occasions, represent the goal as being off to the left, then it must have some representation of the goal that is detached from any particular egocentric representation. If the creature has merely goal-to-the-right and goal-to-the-left, there is no sense in which they can count as agreeing in any way. There is no sense in which they are about the same thing. They cannot have as their content one and the same goal. In order for two representations to be in agreement there has to be the capacity of representing the goal as an independently targetable content from right or left.

Indeed, some researchers regard what I have been calling “independent targetability” as a feature of hippocampal representations more

generally. Cohen and Eichenbaum (1993) have advanced a hypothesis of hippocampal function whereby it is not so much and exclusively spatial navigation organ as much as it is a locus of certain flexible representations more generally. The idea here is illustrated in the following study conducted by Eichenbaum et al (1989). The researchers showed that rats with hippocampal lesions can be trained to prefer certain stimuli, but the preferences will not be exhibited in contexts that differ from the initial training conditions. The rats demonstrate their conditioned preference by moving toward the target odor. Consider that a hippocampal damaged rat, when presented with two odor sources, *A* and *B*, can be trained to prefer *A* over *B*. This rat will demonstrate this preference even after a sustained delay period. However, if the rat is presented with the preferred odor *A*, along with some odor *N*, instead of the non-preferred odor *B* that accompanied *A* in the learning trials, the rat will demonstrate no preference whatsoever. Healthy rats do not exhibit such a lack of preference—their preference is demonstrated in novel situations (situations in which the odors accompanying the target *A* are different than the accompanying odor in the learning trial). Thus, the representation of *A* had by the hippocampal intact rats is more flexible and less context dependent than that of the hippocampal damaged rats.

6. Conclusion

In this paper I have presented a sketch of how one might attribute genuinely allocentric representations to neural states, especially those found in the hippocampal system. Central to my arguments is the view that representational systems consist of distinguishable representational parts. These parts may differ with respect to whether they have egocentric or allocentric representational contents. This view is no less true when applied to systems such as the hippocampal based slope-centroid system thought to underlie spatial navigational abilities. Once it is granted that creatures represent, and that there are organism independent aspects of the creatures' environments, then, contra the Cartesians, there is no *special* problem in positing that some of a creature's representations are about organism-independent features of the creature's environment. The Cartesians are correct insofar as they argue for the utility of egocentric representations. However, they err if this insight leads them to reject allocentric representations as useless, or worse, impossible. On the contrary, allocentric representations are neither useless nor impossible. Thus, the Cartesian impulse, which has caused so much trouble elsewhere in philosophy, should be avoided in neuroscience as well.

Acknowledgements

This work was supported in part by grants from the McDonnell Project in Philosophy and the Neurosciences as well as William Paterson University. I thank Rick Grush for many stimulating discussions on topics related to this paper.

Bibliography

- Andersen RA, Snyder LH, Li C and Stricane B (1993). Coordinate transformations in the representation of spatial information *Current Opinion in Neurobiology* 3, 171-176
- Andersen, R.A. (1995). Coordinate transformations and motor planning in posterior parietal cortex, in M.S. Gazzaniga (ed.), *The Cognitive Neurosciences*, MIT Press, pp.519-532.
- Batista AP, Buneo CA, Snyder LH and Andersen RA (1999). Reach Plans in Eye-Centered Coordinates *Science* 285, 257-260
- Bechtel, W. (1994). Deciding on the data: Epistemological problems surrounding instruments and research techniques in cell biology. *PSA 1994*. Volume 2. pp. 167-178
- Bermudez, J. (1998). *The Paradox of Self-Consciousness*. Cambridge, MA: MIT Press.
- Campbell, J. (1993). "Physical objects in spatial thinking" In Eilan, N., McCarthy, R., and Brewer, B. (eds.) 1993. *Spatial Representation: Problems in Philosophy and Psychology*. Oxford: Blackwell.
- Cohen, N. and Eichenbaum, H. (1993) *Memory, Amnesia, and the Hippocampal System*. Cambridge, MA: MIT press.
- Cussins, A. (1990). The connectionist construction of concepts. In M. Boden, ed *The Philosophy of AI*. Oxford University Press.
- Cussins, A. (1998). *The Electronic Journal of Analytic Philosophy, Special Issue on the Philosophy of Gareth Evans*. <http://ejap.louisiana.edu/EJAP/1998/cussins98.html>
- Eichenbaum, H., Stewart, C., & Morris, R. (1990). Hippocampal representation in spatial learning. *Journal of Neuroscience*, 10, 331-339.
- Eichenbaum, H., Matthews, P., & Cohen, N. J. (1989). Further studies of hippocampal representation during odor discrimination learning. *Behavioral Neuroscience*, 103 1207-1216.
- Fodor, J. A. (1975). *The language of thought*. New York: Crowell.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press/Bradford Books.
- Galletti C., Battaglini P.P., Fattori P. (1993). Parietal neurons encoding spatial location in craniotopic coordinates. *Experimental Brain Research*, 96: 221-229
- Gibson, J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Godfrey-Smith, P. (1998). *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.
- Grush, Rick (2001). Self, world and space: on the meaning and mechanisms of egocentric and allocentric spatial representation. *Brain and Mind* 1(1):59-92.
- Grush R.; Mandik P.(2002). Representational parts. *Phenomenology and the Cognitive Sciences*, vol. 1, no. 4, pp. 389-394.
- Hubel, D. H. and T. N. Wiesel (2001). Brain mechanisms of vision. In Bechtel, William, Mandik, Pete, Mundale, Jennifer, and Stufflebeam, Robert.

- (Eds.) (2001). *Philosophy and the Neurosciences: A Reader*. Oxford: Basil Blackwell.
- Lettvin, J., Maturana, H., McCulloch, W., and Pitts, W. (1959). What the frog's eye tells the frog's brain." *Proceedings of the IRF*.
- Levine M.W. and Shefner, J.M. (1991). *Fundamentals of sensation and perception, 2nd ed.* Pacific Grove, CA: Brooks/Cole.
- Mandik, Pete. (1998) Objectivity Without Space. *The Electronic Journal of Analytic Philosophy, Special Issue on the Philosophy of Gareth Evans*. <http://ejap.louisiana.edu/EJAP/1998/mandik98.html>
- Mandik, Pete. (2001) Mental Representation and the Subjectivity of Consciousness. *Philosophical Psychology* 14 (2): 179-202.
- Mandik, P. (2003). "Varieties of representation in evolved and embodied neural networks." *Biology and Philosophy*. 18(1): 95-130.
- Mandik, Pete. (2002). Synthetic Neuroethology. *Metaphilosophy*. 33 (1-2): 11-29. Reprinted in *CyberPhilosophy: The Intersection of Philosophy and Computing*, James H. Moor and Terrell Ward Bynum, (eds.), Oxford: Blackwell, 2002.
- O'Keefe, J. 1993. Kant and the sea-horse: an essay in the neurophilosophy of space. In Eilan, N., McCarthy, R., and Brewer, B. (eds.) 1993. *Spatial Representation: Problems in Philosophy and Psychology*. Oxford: Blackwell.
- O'Keefe, John (1994). Cognitive maps, time and causality. *Proceedings of the British Academy* 83:35-45. Reprinted in Peacocke, Christopher (1994, ed.) *Objectivity, simulation, and the unity of consciousness*. Oxford: Oxford University Press.
- Taube, J.S., Muller, R.U., and Ranck, J.B. (1990). Head direction cells recorded from the postsubiculum in freely moving rats. Description and qualitative analysis. *Journal of Neurosciences*, 10: 420-435.