

Self, world and space: The meaning and mechanisms of ego- and allocentric spatial representation.

Rick Grush
Department of Philosophy, and
Center for the Neural Basis of Cognition
University of Pittsburgh

Draft: Please do not quote from this draft. Rather, please see the updated and published version forthcoming in *Brain and Mind*.

Abstract: The problem of how physical systems, such as brains, come to represent themselves as subjects in an objective world is addressed. I develop an account of the *requirements* for this ability that draws on and refines work in a philosophical tradition that runs from Kant through Peter Strawson to Gareth Evans. The basic idea is that the ability to represent oneself as a subject in a world whose existence is independent of oneself involves the ability to represent space, and in particular, to represent oneself as one object among others in an objective spatial realm. In parallel, I provide an account of how this ability, and the mechanisms that support it, are realized neurobiologically. This aspect of the article draws on, and refines, work done in the neurobiology and psychology of egocentric and allocentric spatial representation.

Key words: Objectivity, spatial representation, Kant, cognitive maps, posterior parietal cortex

1. Introduction

Perhaps the deepest problem facing those who would provide a neurobiological account of the mind is the phenomenon of the mechanisms resulting in the subjective/objective distinction. This phenomenon is so pervasive and central to our own mental lives that it easily goes unnoticed -- in fact, it takes some work to show that there is a problem that needs to be solved.

With respect to our experience, we draw a distinction between ourselves as subjects of experience, and the world of which we have experience. We understand that our experiential

states, such as the seeing of a brown stick, is experience *of something* that exists independently of our experience of it. But we also realize that it is *experience* of the world, and not the world itself -- we can understand that although the stick in the water *looks* bent, it is *really* straight, and thus we cannot and don't simply equate 'what I experience' with 'the world'. The proverbial ostrich that seeks to deprive the predator of existence by hiding its head in the sand lacks the cognitive mechanisms that allow for the distinction between the *experiencing* self and the *experienced* world. We might, depending on the details of its experience, say that the ostrich has no conception of an objective world, or of objects that persist when unperceived. The bird attempts to exploit an "out of sight, out of world" rule, and this rule would only seem plausible to a creature that has not grasped the subjective/objective distinction, a creature that is unable to conceive of existence unperceived. But it is important to note that such a creature would equally lack a notion of itself as an experiencer. All there is is experience for such a creature. That it is experience that is had **by** something, *itself*, or that it is experience of something independent of itself, *the world*, is beyond its grasp.¹

You are more sophisticated than such a creature. You conceive of yourself² as a subject with a point of view on a world that you conceive as largely independent of your experiencing it, and as

¹ I don't mean to say that I think this is how things necessarily are with ostriches. But the (probably inaccurate) caricature of their conceptual scheme is familiar enough that I will exploit it for illustrative purposes.

² Phrases of the form 'a conceives of x as F' will see heavy use in this article. By it, I don't mean to imply that a must be employing concepts, nor that x's being F is something that a is consciously reflecting upon. I conceive of a chair as being capable of bearing weight when I move to sit on it, even if I do not consciously entertain any thoughts of the form 'This chair is capable of bearing my weight'. I may be thinking about something else entirely as I move to and sit on the chair. I might also experience the chair as being colored with thousands of shades of green, where each shade is part of the content of my visual experience of the chair, even if I lack many or all of the concepts required to specify these shades of green. In effect, 'a conceives of x as F' is shorthand for 'a has experience E such that a full explication of the content of E would include mention of the Fness of x.' By 'conceptual scheme'

largely independent even of your existence. You conceive of yourself as one object among many others in this spatially and temporally extended world, and you conceive of your experience as determined both by the state of the world, your location and orientation in it, and your perceptual receptivity to it. And you conceive of the world as consisting of objects, properties, etc., that are accessible from points of view other than your own. Part of the explanation of why you hide the cookies from me is that you represent the cookies as *objects* which are experienceable from points of view other than your own, specifically from my point of view. And you go back to the hiding place later because you represent the cookies as having an existence which is not dependent on your experience of them, and hence as continuing to exist in the hiding place while you are not there to perceive them. That is to say, you represent the cookies as being independent of your point of view, as being objective tenants of an objective world.

I must say a few words about my use of the terms 'objective' and 'subjective' in order to avoid confusion. On one common use of these terms, they mark a contrast between things that are dependent on being represented, or being experienced, (subjective) on the one hand, and things that are truly independent of all representation: something that exists, or is true, regardless of the experience or biases of the representer. This is a perfectly fine usage, but it is not quite the contrast I am after. The contrast I am interested in is between things that are represented as being dependent upon the representer (subjective), and things that are represented as being independent

I mean no more than 'the set of entertainable contents'. This leaves open the possibility that some entertainable contents may be nonconceptual. Those who don't worry about concepts and nonconceptual content may simply ignore these qualifications.

of the representer (objective). In my sense of objective, then, the large toothy spiders crawling up the LSD user's arm are objective, so long as the hallucinator thinks they are real. The hallucinator has the tools that the ostrich lacks, he is just misapplying them. I am interested in understanding the tools, not in the conditions of their correct application.

In this paper I outline a rather elaborate proposal concerning the cognitive infrastructure that supports this achievement. An outline of this outline is as follows. I first give some reason for thinking that the mechanisms hinge on spatial representation. The idea is that an understanding of the fact that I am a subject of experience in an independent world is made possible by my interpreting my experience as being from a particular point of view within an allocentric spatial realm that contains and is independent of that point of view (I will elaborate on this in a moment, and I will elaborate on what 'allocentric' means at length in Section 3). In Section 2, I provide an account of egocentric spatial representation, one motivated both philosophically and neuropsychologically. Though there are elements of this account that are novel, I don't think it will be very controversial. The controversial part of the proposal will be found in Section 3. There I first explain how an allocentric representation, together with a representation of egocentric space, make possible a grasp of the idea that what one has is a point of view on an objective order. I then argue that the allocentric spatial representation used is, to put it a bit bluntly, off-line egocentric spatial representation, or a representation from an imagined point of view -- an alter-ego-centric representation. I defend this way of looking at things philosophically, psychologically, and neuroscientifically. To the extent that this project is successful, we will

have in hand at least a rough draft of a naturalistic account one of the deepest, most subtle, and most elusive features of mentality.

But first, why tie the notion of the distinction between self and world to spatial representation?

The core of the idea is put nicely by Peter Strawson (1959):³

Roughly speaking, the crucial idea for us is that of a spatial system of objects, through which oneself, another object, moves, but which extends beyond the limits of one's observation at any moment, or, more generally, is never fully revealed to observation at any moment. This idea obviously supplies the necessary non-temporal dimension for, so to speak, the housing of the objects which are held to exist continuously, though unobserved; Thus the most familiar and easily understood sense in which there exist [entities] that I do not now [perceive] is this: that there are places at which those [entities] are [perceptible], but these are places at which I am not now stationed. (Strawson 1959, p. 73-4)

I think this idea is basically correct, but it will require considerable refinement, which will happen in the following sections, before anything serviceable emerges. We can restate the main idea: the subject/object distinction is the result of a cognizer's representation of space, such that i) the cognizer represents itself as being anchored⁴ somewhere in this space; ii) the cognizer represents this space as being independent of the cognizer (i.e. space is represented as independent of the cognizer's point of view)⁵; and iii) the cognizer understands that its experience is the joint

³ In Grush (in preparation) I examine Strawson's proposals concerning objectivity, as well as his philosophical ancestors and progeny (especially Kant and Evans) in much greater detail. There I develop and expand upon the Kant-Strawson-Evans line in a philosophical context, as well as relate that material to the sorts of empirical concerns touched upon in this paper.

⁴ I will use the expression 'x is anchored in y' to mean that x has both a location and orientation in y.

⁵ I suppose that this is the condition that the proverbial ostrich lacks. Not any cognizer that can engage in cognitive mapping will be able to conceive of the world objectively. See footnote 30 for a bit more discussion.

product of the state of the objective spatial world it is in, and the details of its point of view. So the idea is that if a cognizer can represent space, and can represent itself as being somewhere in this space, then it will thereby understand that it is one entity among others in a realm that exists independently of it, and the hope is that this will be sufficient to leverage a distinction between the subjective and the objective.

It will be objected that this initial analysis robs Peter to pay Paul, in that the problem of the source of the wherewithal to represent *things* objectively is addressed by an appeal to an unexplained ability to represent *space* objectively. In order for the above account to work, we must assume that the cognizer represents space as being something over and above its experience -- it must be able to represent itself as being in the space, for example. This objection must be allowed. However, progress has still been made, in that if we can give an account of the ability to represent objective space, then we will be in a position to provide an account of objectivity applicable to objects residing in that space.⁶

Thus even though James Bond is conceived as being in what we might call *a* space (how else can we imagine his exploits if not in a space?), James Bond is not in the real, objective space. This is to say no more than that real space is that which includes where I am. The Eiffel Tower is

⁶ This does not require reducing objects to mere spatial locations. See Yantis, 1993; Olson and Gettner (1996) for evidence to the effect that the objects are not represented simply as locations. And see Maguire et al (1998) for evidence that the 'allocentric' mapping of objects involves cortical regions over and above those used for mapping empty regions.

objectively real because it is spatially related to here -- it is possible to get to the Eiffel Tower from here by traversing a spatial trajectory. The same is not true of James Bond.

Our initial entry into the notion of the subject/object divide, or objectivity,⁷ has assumed that it relies on the ability to represent objective space. This involves three separate components: a capacity to represent an egocentric space, a capacity to represent a non-egocentric spatial realm, and the capacity to anchor the former in the latter. Each of these will require some elaboration. It is to the notion of egocentric space that we first turn.

1.0 Egocentric space.

1.1 The skill theory of egocentric spatial representation.

In discussing the biological basis of spatial representation, it is useful to draw a distinction between (genuine) space and what I shall call a (mere) manifold. Genuine *space* is a multidimensional magnitude that has actual extension, it is a region within which things, and even oneself, can or could physically move about. A *manifold* is some range of discriminable stimuli or properties or locations that can be arranged along one or more dimensions (I will use the term 'manifold' in such a way that genuine space is a special kind of manifold, thus '*mere* manifold' is a manifold that is not genuinely spatial). For example, my ability to see the table near me, and the window farther away, exploits my ability to represent *space*. On the other hand, my ability to

hear middle C at 35dB and at the same time hear high A at 40dB is exploiting my ability to represent nonspatial manifolds. Though I represent pitches and volumes as being independently variable parameters of certain stimuli to which I am sensitive, and orderable along one or more dimensions, I do not represent them as being in themselves *spatially* located, or separated, or anything like that.⁸

The brain receives many channels of information that are dimensionally arranged: amount of pressure at a location on the skin, stretch of a certain muscle, intensity of light at some spot on the retina, exact location on the 2D retinal sheet that is being stimulated, pitch (or cochlear location of stimulation), volume, etc. etc. And efferents can be described as manifolds as well: the brain can release more or less of a certain hormone, can cause more or less tension on this or that muscle, etc. The big question is, why does the brain interpret, or experience, some of these manifolds as being genuinely spatial, and not others? I will argue that none of the more obvious or initially plausible answers to this question are adequate, and that a different answer, one I shall call the *skill theory*, is preferable. The four inadequate theories are: i) The *map theory*, which claims that topographic maps of egocentric space in the brain are necessary or sufficient for representing egocentric space (I will not discuss the map theory, since it is known to be inadequate -- there are no maps of egocentric space in the brain, and even if there were, such

⁷ From here on out, I will often use the term 'objectivity' as shorthand for 'the subjective/objective divide'.

⁸ Two remarks on this. First, I could of course represent such sounds as being located in space derivatively, if I represent them as being made by an object that is at a location in genuine space. Second, in making the distinction between genuine space and a mere manifold, I am -- as always -- distinguishing between how things are represented as being, and not how they really are. So on my usage, a brain in a vat that is interacting with a complex virtual-reality simulation will draw a distinction between genuine space and mere manifolds, even if in some sense it has no access to what we normally think of as real space. The brain might get it right, as presumably our brains -- housed

maps would fail to explain spatial content for the same reason that the receptive field theory will be shown to so fail, since maps are just topographically arranged sets of receptive fields); ii) the *receptive field theory*, which claims that neurons or neural groups that are selectively responsive to stimuli at certain locations in egocentric space do the job; iii) the *teleological theory*, according to which biological machinery carries spatial content if it has been evolutionarily selected for because its carrying this information has been selectively advantageous; and iv) the *motor theory*, which says that connections or associations with motor action imbue a sensory experience with spatial content. Along with showing why these are inadequate (though some have a measure of truth to them), I will articulate what I call the *skill theory* of egocentric spatial content.⁹ The skill theory claims that a creature represents egocentric space -- is able to grasp egocentric spatial contents -- in virtue of its mastery of a battery of sensorimotor skills.

I will begin with a discussion of a fascinating sensory substitution device, the *Sonic Guide*, whose purpose is to provide blind persons with a distance sense of use for spatial navigation and perception (for discussion, see Heil, 1987) It is a device worn on the head that transmits a continuous high frequency (inaudible) probe tone, and picks up echoes of that tone with a stereophonic microphone. The objects in the subject's vicinity reflect various components of the probe sound, in complex ways that depend on the size, distance, orientation, and surface

in the ways we call skills -- are representing real space in constructing representations of genuine space. But again, I am interested in the representational tools, not in the conditions of their correct application.

⁹ I develop this theory in more detail in Grush (1998), paying special attention to the work of the philosopher Gareth Evans, whose work was pivotal in the development of the theory. Crucial elements of the skill theory, including why skills are crucial, is to be found in this article.

properties of those objects. The guide takes these echoes, and translates them into audible sound profiles which the subject hears through earphones. There are several aspects to this translation.

1. Echoes from a distance are translated into higher pitches than echoes from nearby. E.g. as a surface moves toward the subject, the sound it reflects will be translated into a tone which gets lower in pitch.
2. Weak echoes are translated into lower volumes. Thus as an object approaches the subject, the subject will hear a tone which increases in volume (as the object gets closer, it will *ceteris paribus* reflect more sound energy, resulting in a stronger echo), and gets lower in pitch (because it is getting closer, as in (1) above). An object which merely grows, but stays at a distance, will get louder, but stay at the same pitch.
3. Echoes from soft surfaces, grass, fur, etc., are translated into fuzzier tones, while reflections from smooth surfaces, glass and concrete, are translated into purer tones. This allows subjects to distinguish grass from concrete, for example.
4. The left-right position of the reflecting surface is translated into different arrival times of the translated sound at each ear. (Note that it is not required that this coding exploit the same inter-aural differences which code direction in normal subjects. In fact, if the differences are exaggerated substantially by the guide, then one would expect a better ability to judge angle than we typically have.)

As John Heil (1987) describes it, the

...sonic guide taps a wealth of auditory information ordinarily unavailable to human beings, information that overlaps in interesting ways with that afforded by vision. Spatial relationships, motions, shapes, and sizes of objects at a distance from the observer are detectable, in the usual case, only visually. The sonic guide provides a systematic and reliable means of hearing such things.

I will suppose what seems to be plausible, that subjects who have been using the device for a while and are competent with it are actually *perceiving* the objects in their environment directly, rather than *reasoning out* what the environment must be like on the basis pitches and volumes.

This seems to be accepted by Heil who, in discussing the sonic guide and another sensory substitution device, the TVSS¹⁰, notes:

Devices like the sonic guide and the TVSS prove useful only after the sensations they produce become transparent. In employing a TVSS, for instance, one ceases to be aware of vibrations on one's skin and becomes aware, rather, of objects and events scanned by the camera. Similarly, successful use of the sonic guide requires one to hear things and goings-on rather than the echoes produced by the device. ...[children] less than about 13 months... do this quite naturally, while older children do so only with difficulty.

Let us take as an example a subject (call her 'Toni') who has been using the device from birth and is quite competent with it. For Toni, a pitch at 35dB at middle C has a clear egocentric spatial content -- of an object of such and such a size located *over there*. She doesn't hear the *sound* so much as here the *object*, one might say -- just as normal subjects don't see light so much as they see objects. Upon hearing such a sound, Toni would be prepared to immediately (non-inferentially) orient towards the object, point at it, or to try to hit it with a dart, in the same way our visual experience allows us to do these things. However, if you or I were to don the device and hear the same sound, middle C at 35dB, this would carry no spatial import for us. If asked to orient towards it, or hit it with a dart, we would be at a complete loss. How *would* you go about hitting middle C at 35dB with a dart?¹¹

¹⁰ The Tactile-Visual Sensory Substitution (TVSS) device is an array of tactile stimulators (small vibrators), worn by a blind subject on the stomach or back, which is driven by a video camera, typically worn on the head. Each vibrator in the grid is driven, or not, depending on the brightness of the video camera pixel to which it corresponds. See also Bach-Y-Rita (1972).

¹¹ Of course auditory phenomena have spatial import for us: we hear things as being to the left or right, here or there, for example. But what is at issue is that the features of the sound that *do* have spatial import for Toni, e.g. pitches and volumes, *do not*, by themselves, have any spatial import for us.

We can now see why the teleological and receptive field theories are inadequate. The teleological theory says that it is evolutionary pressure on certain biological states or structures that allows them to carry this or that content. But we may safely assume that Toni had no ancestors for whom there was evolutionary pressure for their auditory cortices, especially those areas dealing with pitch, to represent egocentric space. The *receptive field* theory says that egocentric spatial contents are carried by neurons or neural groups if they fire when and only when some stimulus is at a certain egocentric location. Though it is undeniable that such firing would carry *information* to the effect that some stimulus is at that location, information is not sufficient for content. Note that if you don't the guide there will suddenly be, in your auditory cortex, groups of cells that are egocentric receptive fields. The (distributed) group of cells that fires when and only when middle C at 35dB is heard will carry *information* to the effect that some object is at a given location in egocentric space. The problem is that even though that information is there, being carried by cell groups that are acting as receptive fields for that egocentric spatial location, *you grasp no egocentric spatial contents in hearing it* -- middle C at 35dB does not sound to you like an object *over there* in egocentric space. Nevertheless, to Toni it does.

Toni's brain is now set up so that a different set of manifolds is imbued with genuine spatial content. By this example I hope to have loosened the gripping intuition that there is something just naturally spatial about vision. I think that the normal sighted brain has to engage in a lot of work in order to imbue visual experience with spatial content, in the same way that Toni's has to do something in order to imbue auditory experience with spatial content. What is the difference between you and Toni in virtue of which her auditory experience is experience of an egocentric

space, while yours, in the required sense, is not (even when you are wearing the guide, thus guaranteeing that all the spatial information available to Toni is available to you)?

The most plausible initial answer is that Toni has learned to use the device, and you have not. This is correct, but what does 'learn to use' mean? The answer is that in Toni's case, but not yours, the auditory experience is such as to be able to noninferentially cue and guide sensorimotor activity -- her sensory (in this case auditory) and motor manifolds are coherently and appropriately coordinated so as to allow for the skillful execution of a range of activities in terms of a common reference frame.

Before this will make sense or seem plausible, some work needs to be done in explaining what 'coherently and appropriately' mean. Recall that I use the term manifold in a very inclusive sense, to cover ranges of *input* to the CNS, as well as ranges of *output*. In much the same way that one can specify a stimulus location by means of coordinates tied to the sense organ (retinal photoreceptor *a*, or tactile receptor *b*, where *a* and *b* are elements in ordered arrays), one can specify various aspects of motor activity in terms of coordinates as well. For instance, it is possible to specify the location of my hand relative to my torso by giving the angles of my shoulder, and elbow joints. Given that my shoulder has three (actually more than three, but let's keep it simple) and my elbow one degree of freedom, one can specify my hand position relative to my torso as a point in a four-dimensional joint-angle 'space' (= manifold). Motor commands to effectors governing my shoulder and arm are thus ways of effecting trajectories through this joint-angle manifold. Of course, one can also specify the forces applied to various joints as points in a

manifold as well -- they too are sets of ordered elements. As such, coordinating sensation and action can now be described more accurately as *coordinations* among manifolds.¹²

I will use the term 'coordination' to refer to an operation which establishes systematic relations between the elements in multiple manifolds. There are two types of coordination: coincidence-coordination (or c-coordination), and stabilization-coordination (or s-coordination). Two manifolds are c-coordinated if they have subparts which are identified. A simple example is that a map of California and a map of Oregon, so long as each includes at least a bit of the surrounding region, can be coordinated by identifying these regions -- the little bit of northern California on the southern end of the Oregon map is identified with the northern California on the California map, etc. One c-coordinates the two partial maps in order to construct a larger, higher order map. This higher order map may be virtual, in the sense that there is no need to actually physically abut the maps. The two component maps might even be at very different scales, and thus impossible to physically join so as to get a viable physical map.

¹² I am describing the skills as involving sensation and action, but really a third type of information is involved as well -- postural. In this discussion I absorb postural signals, information about how the body, including its sense organs, are oriented, into the sensory and motor classes. I treat postural information separately and in more detail in Grush (in preparation).




Figure 1

An s-coordination is the establishment of a relationship between elements of more than one manifold which has the effect of stabilizing higher order elements that may themselves form an additional manifold. Consider first an extremely simple illustration of stabilization. In a two dimensional realm an eye is pointed at a screen (see Figure 1). The eye can rotate left or right. When a light appears on the screen, its location on the screen relative to the eye cannot be determined from the location on the retina to which the light projects (for simplicity I assume that all lights are equidistant, and hence that distance is not a factor in determining location). This is because stimulation at a given retinal position can correspond to many different points on the screen, depending on the eye's orientation. But given both i) the location of the retinal stimulation, as well as ii) the angle of the eye's orientation, it is possible to determine the location (= direction given the simplifications) of the light. Consider Figure 2. Here, location of retinal stimulation is given along the horizontal axis, and eye orientation is given along the vertical. Every point on the 2-D representation stands for an ordered pair of retinal location and eye orientation. The contours are sets of such ordered pairs each of which correspond to the same light location.

By establishing these contours -- that is, by becoming selectively sensitive to points that lie on a given contour -- a system can s-coordinate the two 1-D manifolds represented in the axes to stabilize a higher order 1-D manifold (the contour itself) corresponding to the position of the light source. This contour then can carry information about spatial location, in that any further mechanism of the system that is selectively sensitive to points on a given contour will thereby be selectively sensitive to a given location. In this way the s-coordination of the two lower-level manifolds (retinal location and eye orientation) stabilizes a higher order manifold, corresponding to the possible locations of light sources.

Now for a slightly less simplistic example. Imagine a creature with two eyes, each of which can move in the creature's head. The creature has an arm with shoulder and elbow joints which can bring the hand within the visual purview of the creature. There are a great many manifolds at work here. There is a 2-D manifold for each retina, a 2-D manifold for the orientation of each eye in the creature's head. There is a 2-D manifold for the orientation of the creature's head with respect to its body, and a 3-D manifold for the position of the creature's hand with respect to the torso (2 degrees of shoulder freedom, and one of elbow freedom). Let us also suppose that there is some bright object in front of the creature which it can sense visually as well as via touch. The project of c-coordinating the region available to vision and the region available to tactile sensation demands some sophisticated orchestration. Before they can be c-coordinated, some region of one must be identified with some region of the other. The only regions that will be stably identical are regions that are in a space centered on the creature's torso. All other regions will slide past and through each other as the creature moves this or that body part (for instance, as I move my head,

my craniotopic space slides through my torso-centered space). So each must be stabilized to a reference frame anchored to the torso, as follows. The bright object will project an image onto each of the creature's retinae. However, the location of the object relative to the creature's head cannot be determined on the basis of the representations in the retinal manifolds, because the eyes may move while the head and object remain stationary, and this will change the position of the image on each retina. But, if given access both to the retinal images and to the orientation of each of the eyes in the head, the position of the object relative to the head can be fixed. That is, by appropriately s-coordinating the retinal and eye position manifolds, one can stabilize a 3-dimensional wedge anchored to the head (I am simplifying the account, of course, as depth would also need to be determined in part on the basis of disparity -- but disparity can also be represented as a manifold). This 3 dimensional wedge is, as a set of ordered elements, a new manifold.

This visual wedge is of little immediate use, however, because the head is free to move -- something can remain at the same location in the wedge, while changing location relative to the torso and arms, if the head moves. This visual wedge needs to be stabilized with respect to the torso. This is done through stabilization with respect to a 2-D manifold that provides information about the orientation of the creature's head relative to its torso, information provided by the muscles which control the position of the neck, mechanoreceptors which give information about the head's orientation, efferent copies of commands to move the head, and vestibular information. When this visual scene wedge is s-coordinated with the position of the head with respect to the body, then it becomes possible to stabilize a visual region with respect to the

creature's torso. That is, for every pair (a, b) where a is an element in the visual wedge, and b is a specification of the orientation of the head with respect to the body, there is an associated point c , an element in a manifold anchored on the torso. The region accessible to touch will be stabilized in an analogous manner, only through the s-coordination of rather different manifolds -- limb joint angles, positions of stimulation on the skin surface, etc. When these two higher-order manifolds, visual and tactile, have been stabilized with respect to the creature's torso, it will then be possible to establish stable regions of overlap. These regions of overlap then underwrite the c-coordination of the higher order manifolds stabilized with respect to the creature's torso. The c-coordination of the 'visual space' and the 'tactile space' allows for the understanding that the *felt* thing and the *seen* thing seen are the *same* thing.

These examples have been irresponsibly brief, and have made many simplifications, but I hope they have been sufficient to make clear just what s-coordinations and c-coordinations are, and how they can interact so as to generate a unified region of sensorimotor stability from disunified and partial manifolds. In all cases, what gets created is a manifold that stabilizes some range of sensorimotor activity, or establishes regions of continuity between these stabilized manifolds. The end result of this process just described is the creature's *egocentric space*. The egocentric space is thus a virtual, high order 3-D manifold, anchored to the torso, which can (though it need not in every case) serve as a common frame for all the lower-order (sensory as well as motor) manifolds that are coordinated with it. (I call this manifold *implicit* because it is not given explicitly in any of the inputs, or outputs, nor is it explicitly represented by a separate map -- it is like the virtual map of California + Oregon mentioned earlier.)

There are three points worth mentioning before we move ahead. First, the egocentric space may very well require maintenance in order to remain stable, as when one gets a new pair of glasses, or more obviously when one wears prisms that alter the relationship between eye/head movement and change of retinal location. In such cases, it takes some time to re-coordinate the temporarily uncoordinated manifolds, and to regain stability. Second, it should be clear that once the behavioral space has been constructed, it is not tied exclusively to any modality. This follows almost by definition, because the egocentric space results from the coordination of different lower-level manifolds, and sensory modalities are bottom-level input manifolds. Third, the stabilizations required to create the egocentric space are stabilizations which depend on motor manifolds as much as sensory manifolds. Neither in isolation could possibly provide the materials for the construction of this space, as each provides the friction which the other needs. It is only because the organism is driven to coordinate sensation and action that a common, higher-order manifold must be established in relation to which both sensory and motor manifolds are ultimately stabilized. Without this imperative, there would be no reason to bother with construction, and without that fact of the interaction, there would be no means by which to learn the correct coordinations.

These considerations lead to what I will call the *coordination principle*:

The Coordination Principle: The content of an element e of manifold M is determined in part by the character of M , but also in part by the character of all manifolds c- or s-coordinated with M .

For example, if I have a map of California and another of Oregon, I can c-coordinate them, provided each shows at least a bit of the other. A person who is exploiting such a coordination can then see from the new, virtual map that Sacramento is closer to Medford than Portland, even though this fact is not represented on any of the maps individually. Or to put the point more accurately, given a coordination between the maps, the content of elements in the first map is influenced by that map's coordination with the second map. The first map now represents Sacramento as being south of Portland, for instance, which is something that neither map could do prior to their coordination. In the case of the egocentric space, the coordination principle entails that part of the content of, say, a visual stimulus is provided in part by how one would orient towards that stimulus (a motor action), and how one would move one's arm in order to bring the hand to that point, such as *THE THING GRASPABLE BY REACHING THUS*.¹³ Similarly, part of the content of a felt location is given by how one would visually orient to that location, and how the hand would look when the eyes are trained on it.

Let us now return to the case of Toni. It should be clear that what has effected the transition from her auditory sensations comprising mere manifolds (pitches, volumes) to providing genuinely spatial experience is her coordination of these auditory manifolds with the other sensory and motor manifolds at her disposal. The same *information* is available to Toni and to us when we wear the guide. The difference is the way this information is coordinated with the other sensory and motor manifolds. In Toni's case, the spatial import comes courtesy of the

¹³ I will use italicized all-caps expressions to indicate a nonconceptual content specification. That is, I am not claiming that the organism is thinking about grasping and reaching, or thinking that that object is graspable in such

coordination principle. Because she has s-coordinated and c-coordinated her auditory manifolds (tone, volume, etc.) with her other sensory and motor manifolds, it is now the case that part of the content of an auditory experience is provided by these motor and sensory manifolds. Part of the content of middle C at 35dB is *THE THING GRASPABLE THUSLY*, and *THE THING I COULD CENTER MY HEAD ON BY MOVING SO*, and *THE THING I COULD HIT WITH THIS DART BY THROWING IT ABOUT LIKE THIS*, and so on, for a potentially infinite number of possible sensory and activity possibilities. This is what 'learning to use the device' amounts to: coordinating the manifolds it delivers with the higher-order 3D manifold that is the behavioral space. This provides content for the elements in the manifold that is provided by this space (one hears middle C at 35dB as being *over there*), and the elements, or locations, in this space get their content from the entire web of sensory and motor manifolds that go into its construction, in accord with the Coordination Principle. The acquisition of these coordinations is part and parcel of the acquisition of sensorimotor skills: being able to move one's eyes towards a heard stimulus, to reach for a seen object, etc.

We are now in a position to see exactly why the skill theory is superior to the four alternatives. I have already addressed the map theory, the receptive field theory, and the teleological theory. What remains is the motor theory. The easiest way to see what is wrong with the motor theory is to consider an objection made to me by Pete Mandik (personal communication), who has interpreted an earlier expression of the Skill Theory as a sort of motor theory:

a way. We seldom reflect on such things when we grasp. Nevertheless, such expressions are the best means of characterizing the content of an experiential episode.

I suspect that what you're up to is sneaking in through the backdoor (action) what you claim won't fit through the front door (sensation). If the real spatial character of sensory input is mysterious, then the coordination of inputs with outputs seems equally so. Some outputs control the flailing of limbs through space, other outputs merely adjust the secretions of chemicals and whatnot. What makes some outputs spatial? Should be no more or less mysterious than the question of what makes some inputs spatial. The reason that coordinating inputs with activities seems to bestow genuine spatiality upon them is because prototypical instances of activity involve moving real bodies through real spaces. But nervous output is no more intrinsically spatial than is nervous input

Mandik is exactly right: neither sensation nor action, afferents nor efferents, carry any spatial content pinned to their sleeve (though some might carry spatial information, but that is a different matter, as we have seen). Spatial content does not fit in through either the front or back door. Rather, spatial content is built within, via the stabilization of a high-order 3 dimensional manifold from lower level manifolds, both sensory and motor. Only some of the efferents and afferents are coordinated with this manifold, because only some vary in ways which are coordinatable with it. Manipulations of glandular excretions are not, while manipulations of joint muscle tensions are, and retinal locations are, while pitches (at least in the normal case) are not.

The skill theory dispenses with any simplistic reduction of the content of egocentric locations to motor commands or joint angles, or positions of stimulation on sensory arrays. Rather, it recognizes the entire web of coordinated sensory and motor manifolds, and principally the high-order virtual 3D spatial manifold constituted by this web, as the unsupplied supplier of spatial content. Of course, joint angles and motor commands are among the more important (being some of the foundational) elements of this web, as we have seen.

A final crucial point about the skill theory. When the appropriate manifolds have been coordinated, it is not required that one actually engage all or any of the coordinated manifolds in order for them to contribute to the content of some element of one of the manifolds. It can be part of the content of a point of light in my visual field that it is the sort of thing on which I could bring this or that motor action to bear, *even if I do not in fact engage in any such motor actions at that time*. The experience will cue a host of patterns of motor engagement, in that the organism will be poised to execute them if desired, and this is enough for purposes of the coordination principle. Presumably such coordinations might remain for some time even after complete paralysis.

We now have an answer to the question what distinguishes a merely physiological manifold from a genuinely spatial manifold. The answer is that a spatial manifold is one which is constituted by the appropriate coordination of motor and sensory manifolds -- it is a higher-order manifold within which the other sensory and motor manifolds (at least some of them) are coherently and appropriately coordinated. The content of this space is provided in a holistic and revisable manner, as a function of the stabilization of higher and higher order manifolds from lower-order manifolds ultimately given directly in sensation and action, in accord with the coordination principle.¹⁴ In the normal case, smell, taste, pitch, and anything else we might think of as a mere manifold are not coordinated in any such way with motor activity, because they are not coordinatable with such activity (and *mutatis mutandis* for certain kinds of non-spatial 'motor'

¹⁴ For some discussion of very similar proposal regarding the mechanisms which underwrite singular reference generally, see Cussins (1998).

actions). The pitch of a bird call, for example, is not normally something that varies in any systematic way with how you might move your head or arms. Apparatus like the sonic guide are constructed so as to allow certain things, such as pitch and volume, to be coordinatable with motor actions in such a way as to guide the execution of those actions. And when such coordinations are established, such experience has spatial import.

1.2 Neural Mechanisms and Psychological Consequences

This section will of necessity be briefer than I would like (see Grush (in preparation) for much more detail). I will discuss a very influential proposal for the neural mechanisms of egocentric spatial representation, and look at a few behavioral results, especially from hemineglect patients, for illumination.

1.2.1 Gain Fields¹⁵

¹⁵ I will be discussing the gain field proposal of Zipser and Anderson to the exclusion of other accounts, mostly because of space limitations, but also because it is a highly successful model. I discuss other proposals, such as Pouget's basis function model (Pouget @@, Pouget and Sejnowski 1997) in more detail in Grush (1999, in preparation). Here I will simply note that this model is not, contrary to first appearances, inconsistent with my skill theory proposal. This is because a fully general version (which the authors do not themselves present in Pouget & Sejnowski 1997) would require decomposition of all sensory and postural signals into a set of basis functions, of the form $B_i(S_1, S_2, \dots, S_a, P_1, P_2, \dots, P_b)$ (the authors only discuss a special case with one sensory and one postural signal, I am generalizing to a sensory and b postural signals). The smallest set of these adequate to underwrite motor control transformations generally will be one includes information about postural signals relative to the torso. Indeed, their proposal is quite consistent, in that this general basis set will be the common coding scheme for all sensory (all encoded as a single set of B_i 's) and motor (decoded from this single set of B_i 's via a different set of linear coefficients for each motor skill) throughputs.

One of the major successes of computational neuroscience is the Zipser and Anderson (1988) connectionist model of the response properties of neurons in the posterior parietal cortex, an area of the brain known to be crucial for the representation of egocentric space. They gave a connectionist network a problem that was similar to a problem that needed to be solved by the brain, and once the network had learned to solve it, they examined the 'response profiles', and compared them to the activation profiles of single neurons in the posterior parietal cortex. The match was fairly close, enough so to merit significant confidence that the solution the artificial network had found was the one that the biological system was in fact using.

The problem is one we have already seen: to take information about i) the location on the retina of a stimulus projection, and ii) the orientation of the eye, and from these determine the location of the stimulus relative to the head. There are many ways one can go about solving this problem. The one the Zipser and Anderson network found was that individual units took on the function of (linear¹⁶) gain fields. For each unit, there is a particular location on the retina at which the unit is most sensitive in that stimulation at that location drives the unit most strongly. As the stimulation moves from that location, the unit's response drops off in a Gaussian manner. Thus, each unit has a standard Gaussian retinal receptive field. But in addition, each unit has its response modulated by eye orientation in a linear fashion. That is, regardless of what the unit's activity would be as a result of the retinal projection, that response is attenuated in proportion to the angle of the eye's orientation away from a preferred orientation. The result is a Gaussian of

¹⁶ The fields are linear in this case, because the retina was only one dimensional. The 2D case, where there are 2 degrees of freedom for eye orientation, would be a *planar* gain field.

retinal location multiplied by a linear function of eye orientation, which means that the cell fires most strongly only with a particular combination of retinal position and eye orientation. And a particular combination of these two corresponds to a particular location relative to the head, as we have already seen (it would be one of the points in Figure 2). Of course, many different such combinations correspond to the same head-centered location, and so a large set of such units are required to represent head-centered location. And again, as we have seen, such units will all fall along a single contour as in Figure 2. A suitable collection of such contours provides coverage for the entire space.

The authors also suggest that the same process can yield torso-centered information, when the process is iterated to include information about the orientation of the head with respect to the torso. Subsequent experiments found cells with such properties (see e.g. Anderson et al. 1997).

As Anderson (1995) describes it:

Recent neurophysiological experiments suggest there exist intermediate and abstract representations of space interposed between sensory input and motor output. These intermediate representations are formed by combining information from various modalities. A head-centered representation refers to a coordinate system framed with respect to the head and is formed by combining information about eye position and the location of the visual stimulus imaged on the retinas. A body-centered coordinate representation is likewise achieved by combining head, eye, and retinal position information. ... There is reason to believe that the brain contains and uses all these representations. [p519]

If the gain field proposal is correct then the egocentric space is represented by the distributed set of cells whose joint response profiles signal positions in a space centered on the torso. Notice that such cells have both sensory and motor implications, since the limb's position is given by

joint angles that start with the torso.¹⁷ Notice that exactly the same coordination process is required for the kinematic specification of, e.g. hand position in egocentric space. The position of the hand in egocentric space cannot be determined by the angle of the elbow alone, since as the shoulder reorients, the same elbow angle can place the hand in different locations. But given both elbow and shoulder angle, the position of the hand relative to the torso, i.e. in egocentric space, can be determined. As I have already noted, it is only because creatures must coordinate sensation and action through sensory systems and motor systems that are all connected to the torso on a variety of movable platforms (neck, shoulders, etc.) that a higher-order space centered on that common platform is needed to coordinate them. This is exactly what the gain field proposal accounts for.

1.2.2 How many spaces?

Given what has been said so far to the effect that there is a single, unified egocentric space and the mechanisms of its representation, there are two sorts of concerns that might be raised. First, perhaps it would be better to speak of many partial *subegocentric* spaces being represented, as opposed to a single egocentric space. Second, perhaps there is more than one genuinely egocentric space at work. I will discuss this second issue in 1.2.3 and 3.2.

¹⁷ For more on the way in which motor action is implicated in the scheme, see Anderson et al. 1997. See also Grush (in preparation) for more discussion.

Both the skill theory and the neurobiology of egocentric spatial representation with which it is compatible analyze the mechanisms of spatial representation in ways that can be counterintuitive, especially since they do not appeal to topographic maps. This lack of maps together with the fact that many sub-egocentric spaces are represented has led some researchers to make some surprising claims. For instance, Stein (1991) claims that since there are no maps of egocentric space in the brain, it doesn't really represent egocentric space. And Farber et al. (ms.) claim that "What introspection presents as the "oneness" of spatial perception is undoubtedly illusory to some degree."¹⁸ Such remarks are puzzling at best. What has gone wrong?

It can be difficult not to fall into what philosophers call content/vehicle confusions. A representation is something -- some object or state (the vehicle) -- that is about, or means, or stands for, something else (the content). Ink spots on a page are one thing, the Eiffel tower is another. Nonetheless, some ink spots might *mean* the Eiffel tower. In fact, you have just seen an example of such ink spots. While few think that you need red ink to write about red things, or that you have to write in a very large script to describe large items, more subtle versions of content/vehicle confusions are seductive and difficult to resist. Spatial representation is a case in spades. From the fact that there is no spatially arranged topographic map in the brain, it does not (though it might seem to) follow that the brain doesn't represent egocentric space. And from the fact that the vehicles are partial, non-topographic, or distributed, it does not follow that the space represented is not unified. One might as well claim that because the novel itself is only black ink

¹⁸ Apart from these claims, both articles are quite interesting, an largely compatible with the position I am defending. If there were more space, I would sing their praises rather than simply focus on these minor complaints.

on white paper, *The Scarlet Letter* does not really concern a scarlet letter, and that the belief that it does is illusory; or that since the inscription 'one word' is itself two words, the notion that it means *one word* is illusory.

1.2.3. Hemineglect.

Lesions of the right posterior parietal cortex, usually resulting from infarct, often lead to impaired representation of the contralateral hemifield of egocentric space. The phenomenon is known as hemineglect, because though in some sense subjects can, under suitable circumstances, represent objects on the contralateral (left) side, such objects and locations are typically neglected. For example, if given a sheet of paper covered with short line segments, and asked to draw a short line through each of them (the line cancellation task), hemineglect patients will typically only cancel those lines on the right side.

There are two issues that merit discussion in the present context. The first concerns something that hemineglect can show us about the representation of egocentric space. My brief gloss (like most glosses on hemineglect) described it as an impairment of egocentric spatial representation. The details are less clear, though. Subjects will do poorly on the line cancellation task even if the entire sheet is in the right half of their egocentric space. This seems to suggest an object-centered hemineglect: the left side of the sheet of paper is neglected, regardless of where it is located in egocentric space. Furthermore, and more to the present point, even in circumstances where it

seems to be half of egocentric space that is neglected, it is unclear if relevant framework is aligned with the body, the environment, or the gravitational field (see e.g. @@@). This suggests that there may in fact be more than one egocentric space represented by the brain, which means that the account I provided in this section was at best a simplification. But this does not threaten the core of the account, since regardless of how many such spaces there are, there will all be coordinated with each other in such a way as to provide for a unified realm supporting behavior.

Second, the hemineglect phenomenon raises a host of tangled issues that I won't even try to disentangle here. The behavioral manifestations of hemineglect can suggest that it is best thought of as an impairment of *attention management*, such that the subject is unable to focus attention on, or move attention to, the left half of whatever the current locus of attention is, rather than, or in addition to, an impairment of spatial content per-se.¹⁹ Indeed, it suggests that the representation of egocentric space and of whole objects must be intact, at least in some sense, for it seems possible to systematically ignore half of things if one is capable, at some level, of representing the things as wholes: for it is only if there is, somewhere, a representation of the whole that it will be possible to systematically ignore just half of it.²⁰

The skill theory that I have been pushing in this article would claim that there is a very close tie between spatial sensorimotor skills and spatial contents. Of course, this leaves open questions

¹⁹ See Stein (1991) for an expression of this view.

²⁰ A simple metaphorical illustration would be this: I choose some number between 1 and 100, and without telling you what number I have chosen, your task is to count upward from 1 by 1s until you reach half my number. If you managed to stop at the number that was half of the one I chose, then this would be a lucky shot in the dark. The

concerning what exactly a sensorimotor skill is, what counts as an implementation of it, and when this implementation is impaired (does strapping my arms to my side destroy my reaching skills? amputating my arms? damaging my spinal cord? lesioning primary motor cortex? lesioning PPC? Or more generally, what is the distinction between having a skill, and being able to execute it?) It also leaves open the question whether attention management is a type of sensorimotor skill, or more specifically, if spatially-directed attention management is a spatial-content-bestowing skill.

I don't have any answers to these questions sufficiently coherent to put forth even as speculations. I will simply register them as topics raised by the current account (and indeed any account), to be addressed in future work.

1.3 Subjectivity, Objectivity, and the Point of View

Notice that the egocentric space is possible only through the implicit adoption of some position and orientation as an origin or center. The coordination of retinotopic stimulus location and eye position to stabilize object location stabilizes that location in a coordinate frame anchored to the *head*. Coordinations with further manifolds leads to a coordinate frame anchored to the *torso*.

And coordinating this with vestibular information gives rise to a hybrid manifold which is of great use: a manifold centered on the torso, and whose rotational orientation in the horizontal plane is

ability to consistently succeed at this task requires that SOME information about the number I choose be available, even if only subconsciously.

anchored to the torso, but whose vertical orientation is aligned with the gravitational field. The exact nature of this implicit anchor point of the egocentric space is less important than the fact of its existence. This implicit anchor of the egocentric space is the point of view, or POV. Though in so doing I wish to flag that this term employs two metaphors, one spatial and one modal. The spatial metaphor is that of a point in a spatial realm. While I think the spatiality is called for, I would like to stress that the 'point' need not be a 0-dimensional point, but may be spread over some region of space, such as the torso or entire body, perhaps. Its role as an anchor does not depend on restricting its extension to nil. Regarding the visual metaphor, it should be clear that a point of view, as I have described it, is not tied to a 'view' in the strictly visual sense, nor limited to sensation in general.

We are now in a position to make a distinction that is rarely made, a distinction *between being conceived of as objective* and *being conceived of as capable of existing unperceived*. If a creature has a fully coordinated egocentric space, it will be able to conceive of some item as being (continuing to exist) *there*, even if it is not currently being perceived (*there* might be a location just behind the head, a location one could grasp at, or look at if one turned around). If one makes the common mistake of defining a point of view exclusively by what is currently *sensed* there will be an inclination to confuse 'objective' with 'perception independent', because in such cases something will be independent of the current point of view if it is independent of sensation. This is a mistake because, as we have seen, a point of view is co-defined with the egocentric space, which is provided as much by action as it is perception. Given this, things that are not perceived, but are nevertheless conceived as *the thing right there* (where the material in italics has as its

primary content a way of grasping or orienting) will not necessarily be conceived of as objective. Since the content of the object's location is still given exhaustively in terms relating to the creature's own perception and action, there is no need that the creature represent this space, or the things in it, as having an existence which transcends the immediately available possibilities for perception and action.

Failure to appreciate this point has greatly hindered progress on our understanding of the abilities that underwrite objectivity. In *Individuals*, Strawson (1959) does not clearly distinguish them (though one can read the distinction as being implicit in the distinction between reidentifiable particulars and objective particulars), and Evans (1985a), in commenting on Strawson's chapter, embraces the confusion outright. Furthermore, in developmental psychology, experiments testing a child's ability to know that something under a rag or behind a screen is still there have been taken to be testing the child's ability to conceive of objects as having an independent, or objective existence. In this regard the proverbial ostrich has very inadequate conceptual tools indeed, as it really does seem to have existence tied exclusively to sensory perception and not the egocentric space. Perhaps the lesson is rather that the description of the proverbial ostrich's cognitive life can't be right. Surely, for example, the ostrich continues to run from the predator when the predator is behind it.

Given this, I will now restate the goal of this study in more adequate terms. Many organisms are such that they have a point of view -- their brain stabilizes sensation and action by means of a coherent 3 dimensional space centered stably on a POV. But organisms of a more sophisticated

sort, in addition to *having* a point of view, understand the fact that what they have is a point of view, and that this point of view is one of many possible points of view on the same stuff, the same world. They also understand that their point of view is itself located in the world, like anything else. They realize that the character of their experience is largely shaped by their point of view. In order to conceive of something as objective, one must conceive of it as independent of one's experience, that is, independent of one's point of view. Not just independent of (sensory) perception.

At best, all I have so far accomplished as an account of the less sophisticated mechanisms, those that supply an egocentric space, or a point of view. We have seen that a system has a POV in virtue of having an egocentric space. It just is the implicit nexus of sensorimotor efficacy constituted by the coordination of sensation and action. In the next section, I begin the account of the more sophisticated abilities -- the mechanisms which allow a system to represent its own point of view *as such*, and to make sense of points of view other than its own, which nonetheless are points of view on the same world.

2.0 Allocentric Space.

The issue of allocentric spatial representation is one of the thorniest issues in neuroscience, psychology and philosophy. Not the least reason for this thorniness is that the term itself is often used in different, sometimes inconsistent, ways. So in section 2.1 I will run through some

of the ways it is used, and indicate the one which I will be interested in. Section 2.2 will show how an allocentric representation of the sort in question addresses the problem of subjectivity and objectivity. Section 2.3 then runs through the neurobiological and psychological evidence in support of the proposed analysis.

2.1 What is allocentric space?

Etymologically the expression *allocentric* means 'centered on an (or the) other'. It turns out, though, that there are at least four different ways that this expression is commonly used, and at least three of them are consistent with the etymological meaning. To differentiate them, we will consider the following example: I am in a room that is empty except for an upright statue of an elf, a lamp that is situated above the statue, and Jones, who is lying down at one side of the room. Now we will consider different spatial descriptions or frames that might be employed in describing or representing these spatial relations.

1. Egocentric Space. I will suppose that it is unproblematic that a description of the locations of the three entities as *the statue is over there*, *the lamp is up there*, and *Jones is right here* are all egocentric, when the *theres*, *ups*, and *heres* are understood correctly. They all involve taking myself as the origin of a space within which the objects are located, and taking my body and behavioral possibilities as defining axial asymmetries (up vs. down, left, vs. right, ahead vs. behind), and also taking myself as the reference point from which these objects are located.

2. Egocentric space with a non-ego object reference point. I might say (or think) that the statue is to the left of Jones, while still using an egocentric space. That is, I assume a space with myself at the origin, and myself as setting up, in the usual way, the axes of left-right, above-below, front-back, etc. I then use one object in this space as a reference point for locating another object in this space: locate Jones, and then move to the left (where *left* means *left as constituted in my egocentric space*), and there you will find the statue. This is a tricky case, because it can seem like an object centered reference frame, but really it is not. Notice that the statue may not be to Jones' left at all, it may be in front of or behind or even to the right of Jones, as judged from Jones' egocentric space.

3. Object-centered reference frames. While an object may serve as a reference point within an egocentric space, as in the last example, it might also serve as the origin of an object-centered space. Consider Jones' thought that the lamp is above the statue. Since Jones is lying down, the lamp is really to the left of the statue in his own egocentric space. Jones is using a space that is centered on, *and* whose axes are provided by, the object itself. In this case, *above* means *above from the point of view of the statue of the elf*, not just *above the statue as a reference point*. Such a reference frame might be centered on any object or person (other than oneself, of course).

Confusions between (2) and (3) type cases are likely to occur for objects that have no natural orientation or asymmetries, such as a ball. When one thinks or says "the rock is next to the ball" one could be i) using the ball as a mere reference point in one's egocentric space, together with one's ability to register contiguity, to arrive at the thought that the ball is next to the rock; or ii) one could be adopting a spatial framework truly centered on the ball.

4. Virtual points of view. Another option similar to the last would be to center a space at some location where there is no object. This is sometimes called a 'neutral' perspective (see e.g. Franklin et al, 1992). For instance, I might represent the locations of the four objects in the room (myself being one of them) as they would be seen from a point 10 or so meters above. Most maps are somewhat stylized versions of such a representation, abstracting away from particularities that create parallax, for instance. But they still have an orientational structure, left-right and up-down axes. They are identical to the 'object centered' reference frames, except that there happens to be no actual object or person that is constituting the point of view.

5. 'Objective' or 'nemocentric' maps. The most puzzling proposal is the 'view from nowhere'. Some researchers, mostly philosophers, have hankered after some kind of truly viewpoint-independent notion, especially those interested in objectivity. The fear is that any viewpoint is a subjective viewpoint, and hence not an objective viewpoint. Though the expression 'allocentric' is often used to mean such a representation, I will use the expression 'nemocentric' (meaning *centered on nobody*). For a number of reasons, I will not concern myself with nemocentric representation. I will go into more detail as to why in Section 3.

My account of allocentric space takes it to be a space in the senses of 3 and 4 above, which is the literal sense of allocentric -- centered on another. Such a space is a space that has another object, person, or perhaps just location as its origin, and this space has its axial asymmetries (right-left, up-down, front-back) established from that point of view.

2.2 The objectivity thesis.

In any case, it should be common ground, even for people who hanker after cognitive maps of the nemocentric (5) sort, that maps are non-egocentric in that they are not centered on the representer's actual location. So I will try to make the points I need to make for now in terms which neutralize the distinctions between (3), (4) and (5). I will speak, for now, simply of a non-egocentric map, or NEM. Suppose that I have an NEM which represents the spatial relations between three objects, *a*, *b* and *c*, and that according to the NEM they are arranged in a straight line (assume that these are statues in a park that can be recognized by sight). Suppose further that I can locate *a* and *b* in my egocentric space, perhaps I wake up from a drunken stupor, and they are directly in front of me and in view: *a* is just ahead and to the left, and *b* is just ahead and to the right. If the NEM of *a*, *b* and *c* lying on a straight line is brought into c-coordination with *a* and *b* as represented in my egocentric space, then I will be in a position to locate *c* in my egocentric space even though I cannot perceive it -- even though *c* is not at all part of my current experience from my point of view --; it is *just over there to the right*. Thus, such a c-coordination will allow me to populate my egocentric space with entities that had no previous residence there.

By the same token, this c-coordination also allows me to place things in the *map* that were not previously *there*, such as the bag of loot that is clearly visible at the foot of *b*, yet was not indicated on the map. Importantly, among the elements which will be located in the map via this c-coordination is *the point of view itself*. The nexus which implicitly defines the egocentric space

is identified with one of the points (or regions -- recall we do not need to restrict the POV to an extensionless point) on the map. If *a* is *to the right*, and *c* is *to the left*, and *b* is *right here* -- I have moved to the bag of loot (a few more journal subscriptions for the department library, of course) -- then I will have placed myself in the map, between *a* and *c*, just at *b*.

Now upon such a c-coordination of the NEM and the egocentric space, and by the Coordination Principle, the elements in the NEM inherit the content provided by the egocentric space, most pertinently in this case locatability (one can thus point to, or walk towards, the unseen statue). Moreover, the elements in the egocentric space inherit content made available through the NEM. What I need to establish is that the NEM represents objects and space as *objective*, as independent of the point of view of the representer.

This is easy enough to do. Because the NEM represents the point of view as a location in the space (e.g. the familiar 'YOU ARE HERE' arrow), there is nothing special about the point of view in the NEM's representation. My own point of view *has no special status in the NEM*. The other objects in the map are unaffected by altering the location of my POV, or by removing it altogether. Another way to put this is that the elements within the NEM are represented as being independent of my point of view (in exactly the same way that any two items in the NEM are conceived as being independent of each other). This point deserves a lot of emphasis: *because the actual point of view has no special status in the NEM, and because the entities represented in the NEM are such that their being represented in the NEM is independent of the details or existence of the point of view's location in the NEM, the NEM thus provides these objects and locations with*

content that is not tied to the current perceptuo-behavioral exigencies of the representing system. It thus provides for an understanding of the fact that such objects are genuinely independent of the (actual) point of view. This brings us to what I will call the *Objectivity Thesis*:

Objectivity Thesis: The elements in the egocentric space achieve objective import when the egocentric space, *E*, and thus the elements in it, are c-coordinated with a spatial manifold (the NEM in the example), *A*, that is itself independent of *E*, and capable, in principle, of supporting arbitrary c-coordination with *E*. Moreover, this very c-coordination involves placing the POV, which is only implicit in *E*, explicitly in *A*, and thus provides the necessary tools for conceiving of the POV as a point of view on an order that is independent of that point of view.

Compare the objectivity thesis to a recent gloss by John O'Keefe concerning the objectivity of a cognitive map (we will return to O'Keefe's model in more detail in 2.3):

By objective, I mean a representation which is conceived as independent of the organism in terms of either its actions or its purposes. The environment is conceived as a stable entity through which the animal moves, and which exists independently of the animal's motivations, goals or desires. [O'Keefe, 1994, p.39]

It should be pointed out that, just as in the case of the manifolds whose coordination results in the egocentric space, it is not necessary that one actually engage in the project of coordinating the egocentric and allocentric manifolds in order for elements in the egocentric space to be conceived of as objective. The fact that the organism has the ability to effect such a coordination is sufficient to imbue these elements with the requisite content. Once such a skill has been mastered, it allows me to experience my egocentric space as the sort of thing which can be coordinated with another point of view, even if it is not actually on any given occasion so coordinated. This is entirely parallel to the fact that I need not actually engage any motor

programs pertinent to a given stimulus in order for the experience of that stimulus to be infused with content derived from those cued-but-not-executed skills, provided the skills are in place.

The c-coordination of the egocentric space, with its merely implicit inclusion of the point of view as a special region *unlike* any other (the limit of the experienced world rather than anything in it), and the map, with its representation of the point of view as being entirely explicit and *just* like any other run-of-the-mill object or location, provides, via the coordination principle, for the unique, almost contradictory content of the self-notion: as something at once very special and unlike anything else, while at the same time entirely unspecial and exactly like everything else.

2.3. The cognitive map(s).

The psychological and neurobiological literature on cognitive mapping is rather large, and I won't pretend to do any sort of justice to it here. Rather, I will focus on the basic background, one good model, and a few interesting psychological and neurobiological facts, and try to show that they are all consistent with, and may even shed some additional light on, the ideas I have proposed.²¹

The additional light will be to the effect that the cognitive map is implemented not as a nemocentric representation, but rather as an alter-ego-centric representation. Though I think the details are a bit more complicated than this, the fundamental idea is that cognitive mapping

²¹ I go into more detail, in all aspects of the psychological and neurobiological work on cognitive maps, including discussion of more models, in Grush (in preparation). Here I merely draw attention to some of the more detailed work in footnotes where possible.

exploits mechanisms that create an 'imaginary' egocentric space,²² with a concomitant placing of a virtual alter-ego as the 'viewer' of this space.²³

First the background. In the thick of an era that strove to understand psychological phenomena in behavioristic terms, Edward Tolman, his students, and some 'underpaid research assistants' (see Tolman 1948) pushed the idea that rather than executing spatial behaviors by means of S-R mechanisms, maps 'get established' in the rat's brain during spatial learning, and these maps are employed in representational processes that intervene between sensory intake and motor output. In the course of subsequent decades, the relevant psychological question stopped being "Are there cognitive maps?" and became "What are cognitive maps like, how do they work, what are their properties?"²⁴

²² For example, I can imagine what it would look (and even feel) like to get up out of this chair and walk into the next room. I can even imagine turning around at the door, and looking at the chair in which I now sit. I can imagine what it might look like to see, from the doorway, myself sitting in this chair typing away. In such cases, I am imaginatively relocating my POV to somewhere it is not. Such imaginings provide me with a manifold -- in this case a representation of objects in a region together with the spatial relations obtaining between them -- that is not provided from my actual current point of view. At least some of these objects can be the same as those objects that are given to me in my actual behavioral space, such as this chair and this table. And this is all that is needed to get the Objectivity Thesis up to speed. This will allow for a c-coordination between my behavioral space and the representation provided by the imagination, and this allows me to conceive of the things I am currently experiencing as being independent of my (current, actual) point of view. And that is what we have been after all along. The Objectivity Thesis works even when the 'allocentric' manifold that includes the current point of view is itself just another from-some-point-of-view representation. This is just one example, though. More detail will follow.

²³ As I said, the details are more complicated. I will mention only three here. First, in calling this a process of imagination, I do not intend to claim that it is (necessarily) consciously available, nor that it is done via an act of will as opposed to automatically. See Farrell & Robertson (1998) for evidence suggesting that such 'imaginative' processes are automatic, and not always consciously accessible. Second, I should emphasize that the imagery involved is not *visual* imagery, but rather amodal, or supramodal spatial imagery, the sort of imagery one would expect to be generated by the off-line operation of the posterior parietal cortex, which deals with amodal spatial representation. See Luzziatti et al. (1998) for some discussion of this issue. Third, as a point of clarification, I would like to say that rather than thinking of cognitive mapping as a type of imagery, I would rather say that imagery, mapping, perception, and memory are all applications of a more general reconstructive representational operation that employs forward models in Kalman-filter-like processes. See Grush (in preparation) for more discussion of this.

²⁴ But the "Are there any?" question is not yet dead. See Bennett (1996).

The neurobiological landmark was the discovery by O'Keefe and Dostrovsky (1971) of 'place cells' in the rat hippocampus -- cells that would fire strongly when and only when the rat was at a given place in a known environment. In *The Hippocampus as a Cognitive Map*, O'Keefe and Nadel (1978) claimed to have found Tolman's maps in the hardware of the brain -- a finding consistent with the long-known fact that damage to the hippocampus severely impairs the sorts of spatial abilities the cognitive maps are supposed to support. But the question of the structure of the map was still very much unknown. O'Keefe and Nadel proposed that it was a Euclidean map. O'Keefe has more recently championed an updated and more specific version of this, the 'slope-centroid' model (for more detail, see O'Keefe 1994).²⁵ On this model, upon entering a new environment, the animal spots the various objects in the surroundings, and determines two things on the basis of their distribution: i) a centroid -- roughly a center of mass or spot that is in the center of everything; and ii) a slope, or a line that runs through the centroid at an orientation that is parallel to the direction in which the objects are most widely spread: roughly the first principle component. So for example if the objects are arranged in an oval distribution, the centroid would be at the center of the oval, and the slope would run through the long axis.

Once the centroid and slope have been established, all objects in the environment can have their locations coded as a vector whose length is distance from the centroid, and whose angle is its angular displacement from the slope -- in effect, polar coordinates with the centroid as origin and the slope as the reference for angular displacement. The animal's own location can also be coded

²⁵ I will not discuss any other models, such as those that claim the cognitive mapping mechanism consists of orientation-specific local views or 'snapshots' (see McNaughton, Leonard and Chen, 1989)

in such coordinates, of course (and so the cells that code for the animal's current vector would be its 'place cells'). Furthermore, the direction and distance from the animal to any entity in the environment can be found by simply subtracting the animal's vector from the object's vector.

This scheme can seem like a parade case 'object-centered' or 'environment-centered' representation. And it might be. But note that as described it is also compatible with being an alter-ego-centric reference frame with an object as reference point (the type (2) case). Recall that in (2) an object is used as a reference point, but the axes are established by the creature's own egocentric space. My suggestion is that the slope-centroid model might well be a (2) case, with the space axes being established by an alter-ego reference frame rather than the egocentric frame. But 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? And shall my virtual body be aligned with the coastline, or perpendicular to it? It would make sense to locate

myself so that the majority of what I want to represent is in the center of the space, in order to reduce the amount of processing that has to be performed in the periphery. And if there is an difference in the representational loads that can be carried by the up-down and left-right axes, then let that determine the orientation. It is known, for example, that in humans the visual field (as well as the 'visual field' of visual imagery, see Farah et al. (1992)) has a larger extent in the left-right direction than the up-down direction. Given this, it would seem to make sense to align my imagined location so that left-right falls along the direction of greatest spread of objects.

Though it may be compatible with the slope-centroid model, what positive evidence is there that the representation used is alter-ego-centric? Again, space limitations force me to be sketchy (and again, more detail can be found in Grush (in preparation)), but here is a speculative proposal.

First, it is uncontroversial that the PPC is primarily involved in the representation of egocentric space. Given this, it is interesting to note that lesions to the PPC also severely disrupt navigation skills, in some ways more severely than does damage to the hippocampus itself.²⁶ It is also known that the hippocampus is in the business of memory (not just spatial memory), and also the memory and learning of structural information rather generally (Eichenbaum, 19@@, Cohen and Eichenbaum, 19@@).

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

²⁶ See, for example, Hyvarinen, 1982; Kolb & Walkey, 1987; Kolb, 1990. See also Ammassari-Teule et al. (1998) for an interesting discussion of the contributions of hippocampal and PPC regions.

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.²⁷ The information that goes into the construction of the map at any given time will typically be a function of both memory of spatial relations encountered before as well as spatial relations between objects that are currently perceived.²⁸ But the constructed map is not a mere collection of such partial information, but rather a coherent reconstruction in another form, perhaps something like the centroid-slope representation proposed by O'Keefe.

Compatible with this proposal is the fact that the memory and processing of remembered spatial layouts, even large ones that presumably require the construction of cognitive maps, are orientation-dependent. For instance, Roskos-Ewoldsen et al. (1998) found clear orientation effects for both small and large spatial layouts, using both artificial as well as natural spatial configurations. The idea that this alter-ego-centric reference frame is linked to the imagination is supported by the discovery of 'mental rotation' latency effects in aligning one's actual heading with the orientation of the map (Farrell and Robertson, 1998).

3.0 Discussion

²⁷ Compare this to the proposal of Petrosini et al (1998, p.208): "Thus it is possible to postulate that the parietal cortex sets up a coordinate system representing space within which it locates objects identified, but not located, by the inferior temporal cortex." This topic is discussed in more detail in Grush (in preparation).

²⁸ I note that my proposal is similar to that of Poucet (1993), though he couches it in terms of the hippocampus providing topographic representations, and the posterior parietal providing metric ones.

3.1 Why is this 'objective' or 'allocentric'?

As I noted in section 2.1, some people seek a sort of representation that is what they might call 'truly objective', or from no point of view. There I passed on the issue, but now I will briefly take it up. One expression of this hankering comes from John Campbell:

A first interpretation is that what is wanted [for objective representation -- RG] is a kind of top-down view, so that we can think in terms of a kind of aerial photograph, and then one has to identify oneself as one of the people in the photo. But that would not be enough, for it would only give the viewpoint of the photographer, and we need a picture of the world that is objective, in that it is not from any viewpoint at all. It is a picture not drawn from any perspective. A natural reaction to this demand is to move the photographer further away from the earth, but of course no finite amount of movement will be enough. And once he is at an infinite distance, what will he be able to photograph? What would you expect to be able to discern from nowhere? (Campbell, 1994, p. 6)

It seems to me that the antidote to Campbell's argument is to simply realize that in saying that an objective representation is one which is not tied to any point of view, taking, so to speak, the narrow scope interpretation of 'any' is sufficient. In fact, this is stronger than is necessary, as all that is required *for the sense of objectivity that I am after* is that things are experienced as being independent of the current point of view. It is compatible with this that the experience might be dependent upon a mere handful of possible points of view. It is an unnecessarily large step (for my purposes, anyway) from *independent of the current point of view*, to *independent of all possible points of view*. And the proposal so far accommodates this, because obviously the imagination is free to roam over many different possible points of view, each of which can be coordinated with the egocentric space (this is the force of the 'in principle arbitrary' condition in the statement of the Objectivity Thesis). This ability to freely effect coordinations with alternate

allocentric viewpoints provides the notion that what is in the egocentric space is available from other points of view, and indeed (potentially) from any point of view.²⁹ I confess that I am unable to discern any advantage in thinking of an objective representation as one that is *from no point of view* over thinking of it as being one that is *compatible with all points of view*. And thus I am also unable to discern any advantage in Campbell's broad-scope interpretation of 'not any', over the narrow-scope interpretation. In all fairness, Campbell may be discussing some other sense of objectivity. If so, fine, but then the stated objection is not an objection to my account. If there is some form of nemocentric representation used by our brains, or needed for more esoteric purposes than those I am concerned with, fine. But for my purposes such representations are fifth wheels.

My position here is compatible with that of O'Keefe, who, in discussing animals more sophisticated than the rat, speculates that:

... in other, more cognitively developed, animals the point of view could be uncoupled from the location of the animal's body and imagined to move to other parts of the environment. This would be accomplished by uncoupling the movement of the focus from within the map from the animal's actual movements with the development of a mechanism for the suppression of the motor signal to the musculature. These intended or anticipated movements would now afford the animal a means for representing the environment from all possible perspectives or -- what is equivalent -- from no particular perspective. (O'Keefe, 1994)

²⁹ Of course, in order to use a map, the organism must be able to represent objects in a manner that allows for them to be recognized from multiple orientations. This should not be confused with representing them as objective, however. Having a representation of an object such that one can recognize it from multiple viewpoints is not the same as representing an object as the sort of thing that is available from more than one viewpoint. The proverbial ostrich has the first but not the second -- presumably it can recognize the predator as a predator from any, or at least

Leaving aside the fact that I have argued that O'Keefe's own centroid-slope model is a sort of off-line representation, his point here is entirely compatible with the objectivity thesis, and also an expression of a narrow-scope interpretation of objectivity.³⁰

Second, I'm never quite sure what such nemo-centric representations are supposed to be. It can't be simply a spatial description that makes no explicit reference to an ego, because such descriptions can still be egocentric, for example the (2) cases mentioned in section 2.1, where an object is used as a reference point. Such relations as 'next to' and 'between' can be egocentric in this sense.

Furthermore, 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.³¹ Consider the following suitably sanitary structural description: Line segments DE, DF, and DG are all mutually perpendicular, and 1 meter in length. This is fine, and seems to be a POV independent description. But notice that exactly because it is independent of any point of

very many, orientations. I discuss this distinction, with special attention to orientation-invariant recognitional capacities, in Grush (1999, in preparation).

³⁰ O'Keefe's suggestion that the ability to go completely offline is correct though, but because of space limitations, I have not been able to explain why. In brief, the objectivity thesis requires that the organism represent the region in the map as being independent of the current POV, and this seems to require that the organism has the ability to create a map in which the organism itself is not located, in order to give content to the notion that the map represents things that are truly independent of the organism itself. Because in the rat, O'Keefe reasonably presumes, the animal's mapping machinery is hardwired in such a way that the rat is itself always in the map, its map seems to lack this dimension of content. This might also allow for the idea that the proverbial ostrich, though it can use a cognitive map to navigate, might lack the full-blooded sense of objectivity, as will any cognizer whose cognitive map is hardwired in such a way as to always include its own location. I explore this more in Grush (in preparation).

³¹ I am bringing up considerations similar to those Kant marshaled against the Leibnizian view of space, as purely relational (Kant, @@). As far as I can tell, the best one can do in trying to create a nemo-centric representation of space would be to mention only the relations between the objects themselves, and we can call this a structural description.

view, it is applicable to incompatible configurations. Imagine DE going *up* from D to E, and DF going *left* from D to E. Now, DG can go either *towards you* or *away from you*, and still be compatible with the description. And the two configurations cannot be rotated into each other (at least not in three dimensions).³²

These remarks have been brief, but I hope at the very least to have shown that an account of the objective/subjective distinction need not appeal to nemo-centric representations, and that alter-ego-centric representations are compatible with some of the best models and evidence coming from the neurobiological and psychological literature.

3.2 The experience of unity and coherence.

Perhaps the dominant theme of this discussion has been the coordination of many many different manifolds to construct a higher-order spaces (egocentric or objective). So the question naturally arises: How can something unified and coherent emerge from so many partial, disunified representations, especially when there is no explicit map with which they are all coordinated?

The image one can get is that of a room full of maps: cities, states, countries, which jointly depict

³² For those familiar with vector cross-products, DG can be either $DE \times DF$ or $DF \times DE$. One might think that a structural description that included something like $DG = DE \times DF$ would be unambiguous and purely structural. But it is only unambiguous given a background convention, the 'right-hand rule', for determining the direction of vector cross-products, and application of this convention depends on orientational axes that are external to the configuration itself (a person's left and right hands, for example). The gambit to the effect that one can disambiguate by appeal to relations to something external to the configuration itself is correct. But it does not solve the problem, because either this something will be a point of view, in which case we don't have a non-nemo-centric representation after all,

the entire surface of the earth with a reasonable amount of overlap and differences of detail -- with no single, unified map of the world given at all. A cartographer in such a room would certainly be aware of all the component maps, and would only be able to reconstruct the world map -- to find a route from Madrid to India -- through the execution of a degree of *ad hoc* work. Why is this not how it is with us? Why are we typically aware of one unified space, not aware of all the partial component spaces? (Try training your attention on a point in head-centered space, say about 4 feet up and to the right, as you walk about. If you're like me, you can do it, but the point is always conceived of as being in, and sweeping through, a more fundamental, given, unified space.)

The question why we experience space as a single, unified objective thing, and not as a jumble of partial spaces, is a deep and revealing one. The first step to an answer lies in the realization that we are interested in contents (not vehicles). We have already remarked on this issue in Section 1.2.2. The fact that the vehicles are disunified, partial, and whatever else does not affect the contents they carry, and in particular does not entail that the contents need be disunified, partial, or whatever.

But this cannot be the whole answer, because there is reason to think that many of these partial spaces are *represented* -- that the brain actually carries *contents* that are retinocentric, craniocentric, object centered, etc. Clever psychological experiments and phenomena that surface

or it will not be, in which case we can add this something and its purely structural relations to the original configuration, and thus create a new, larger, ambiguous structural description.

upon damage to the brain show that this is the case. The harder question is, why are only some of these contents that the brain clearly carries available to us consciously and reflectively -- available at the *personal level*, and not merely subpersonal contents? These partial spaces, whatever they are (partial topographic maps, partial subpersonal spatial contents) are, for some reason, transparent to us. Recall Heil's remark (section 1.1) that in order for the sonic guide to work, the tones and volumes must become transparent -- the subject must pay attention not to them, but to the objects they are signifying. For us (at the personal level of consciously available content), all the partial maps are transparent -- we 'see through' them to the single objective world they signify. But *why* is this so? Why are we aware only of a single, unified, objective space? Surely the demands to have a single space for the programming and assessment of sensorimotor activities can be met while also having conscious access to the other partial spaces. So why don't we? One can only speculate.

So here is an admittedly sketchy and probably unconvincing speculation.³³ The question of why only certain contents are available at the personal level is closely interwoven with the question what the personal level is to begin with. If we think of the personal level as a homunculus, then there are two questions: what is the homunculus, and why does the homunculus have access only to this information and not that? This would suggest answers in the form of limits to attention or working memory -- in effect, we hypothesize that the homunculus has a desk only big enough to carry those contents the exclusive possession of which we are trying to explain. For various

³³ In Grush (in preparation) I develop this obviously Kantian speculation in much more detail, removing much of the sketchiness and with any luck at least some of the unconvincingness.

reasons (its homuncularity being not the least among them), I don't like this line. Rather, perhaps we can exploit the more esoteric idea that the personal level *just is* that which is primarily aware of an objective world.

What on earth does THAT mean? tk

On this account, the cognitive machinery that makes possible representational access to an objective world (including the transparification of the machinery itself), of necessity also creates the ego, the POV, as that which is aware of this world, and as that to which the machinery is transparent -- in the same sense that the inside and outside of a circle are both created by drawing a circle. The neurocognitive construction of the objective world and the neurocognitive construction of the (personal) subject of experience are not two separate constructions, but rather different aspects of the same construction.

4.0 Conclusion

Any honest assessment of the degree to which I have clarified all the philosophical details of subjectivity and objectivity would be rather harsh. There is much about the nature of these notions which I have not addressed, and for which it may not be clear that the theoretical machinery I have developed will even be capable of addressing. To illustrate: I have said nothing about the conceived distinction between mental states and physical states, I have said nothing about our conception of the minds of others, I have said nothing about the relation between

objective *things* and objective *facts*. I have said nothing about the conceived appearance/reality distinction that seems to be an integral part of our own objective conceptual scheme. I have said nothing about the role that our notions of causality, or our naive theories of perception, or our memory, play in our self-conception. A full catalogue of the deficiencies of the theory I have developed would burden patience.

While such lacunae speak to the shortcomings of this article, they speak with a louder voice to the difficulty and complexity of the problem. But standing slack-jawed at the challenge of such complexity and the risk of failure it presents is hardly productive. I have accordingly endeavored to address these issues by concentrating on what strikes me as one of their roots -- the ability of a system to represent itself as having a point of view on an order that is independent of its point of view. This ability is arguably the starting point for many of the other issues raised, such as the appearance/reality distinction, and it also seems to be a place where an empirical beachhead might be established.

Acknowledgments:

I would like to thank many people for comments on earlier drafts of this article. These include Adrian Cussins, Pete Mandik, Ingar Brinck, the class participants of a graduate seminar on 'Subjectivity, objectivity, and their neurobiological foundations' which I taught at the Center for Semiotic Research at the University of Aarhus, Denmark, in the Fall semester of 1997, as well as audiences at the philosophy departments of the University of Connecticut and the University of Pittsburgh, where I presented earlier versions of this paper. I would also like to thank the Center for Semiotic Research at the University of Aarhus, Denmark, and the Danish National Research Counsel for a research fellowship, during which the initial outlines of this research were begun, and Per Aage Brandt for his beneficence in securing this support.

References:

- Amassari-Teule, M., E. Save, B. de Marsanish and C. Thinus-Blanc (1998). Posterior parietal cortex lesions severely disrupt spatial learning in DBA mice characterized by a genetic hippocampal dysfunction. *Behavioral Brain Research* 95:85-90.
- Andersen, Richard A. (1995) Coordinate transformations and Motor Planning in Posterior Parietal Cortex. In Gazzaniga (ed.) *The Cognitive Neurosciences*. Cambridge, MA: MIT Press.
- Anderson, Richard A., Lawrence H Snyder, David C. Bradley, and Jing Zing (1997). Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience* 20:303-330.
- Bach-y-Rita, P. (1972). *Brain Mechanisms in Sensory Substitution*. New York and London: Academic Press.
- Bennett, A.D.T. (1996). Do animals have cognitive maps? *Journal of Experimental Biology* 199:219-224.
- Berkeley, George (1988) *The Principles of Human Knowledge*. (Roger Woolhouse, ed.) London: Penguin Books.
- Bermudez, Jose Luis (1997). *The Paradox of Self-Consciousness*. Cambridge, MA: MIT Press.
- Campbell, John (1994) *Past, Space and Self*. Cambridge, MA: MIT Press.
- Cussins, Adrian (1998) Subjectivity, Objectivity and Frames of Reference in Evans's Theory of Thought. In *The Philosophy of Gareth Evans: Special Issue of the Electronic Journal of Analytic Philosophy*, Rick Grush, guest editor. At <http://www.phil.indiana.edu/ejap/ejap.html>
- Evans, Gareth (1982). *The Varieties of Reference* (edited by John McDowell). Oxford: Clarendon.
- Evans, Gareth (1985a) Things without the mind. In Gareth Evans (1985) *The Collected Papers of Gareth Evans*. London: Oxford University Press.
- Evans, Gareth (1985b) Molyneux's question. In Gareth Evans (1985) *The Collected Papers of Gareth Evans*. London: Oxford University Press.

- Farah, M., M. Sos, and R. Dasheiff (1992). Visual angle of the mind's eye before and after unilateral occipital lobectomy. *Journal of Experimental Psychology: Human Perception and Performance* 18:241-246.
- Farber, Ilya, Will Peterman, and Patricia Smith Churchland (ms). The view from here: Spatial representations.
- Farrell, Martin J., and Ian H. Robertson (1998) Mental rotation and the automatic updating of body-centered spatial relationships. *Journal of Experimental Psychology: Learning, Memory and Cognition* 24(1):227-233.
- Franklin, N., Tversky, B., & Coon, V. (1992). Switching points of view in spatial mental models. *Memory & Cognition*, 20, 507-518.
- Grush, Rick (1998). Evans on spatial content. In *The Philosophy of Gareth Evans: Special Issue of the Electronic Journal of Analytic Philosophy*, Rick Grush, guest editor. At <http://www.phil.indiana.edu/ejap/ejap.html>
- Grush, Rick (1999). Manifolds, coordinations, imagination, objectivity. Part One: Kant, Strawson, Evans, and other philosophical motivations. Part Two: Ego- and allo-centric representation in the brain. Available at <http://www.pitt.edu/~grush/papers/manifolds/>.
- Grush, Rick (in preparation). *The Machinery of Mindedness*.
- Heil, John (1987) The Molyneux Question. *Journal for the Theory of Social Behavior*. 17:227-241.
- Hyvarinen, J. (1982). *The parietal cortex of monkey and man*. Berlin: Springer Verlag.
- Kolb, B. (1990) Posterior parietal and temporal association cortex. In B. Kolb and R.C. Tess (eds.) *The cerebral cortex of the rat*. (pp. 459-471). Cambridge, MA: MIT Press.
- Kolb, B. and J. Walkey (1987). Behavioural and anatomical studies of the posterior parietal cortex in the rat. *Behavioural Brain Research* 23:127-145.
- Luzziatti, C., T. Vecchi, D. Agazzi, M. Cesa-Bianchi and C. Vergani (1998). A neurological dissociation between preserved visual and impaired spatial processing in mental imagery. *Cortex* 34:461-469.
- McNaughton, B.L., Leonard, B., and Chen, L.L. (1989). Cortical-hippocampal interactions and cognitive mapping: A hypothesis based on reintegration of the parietal and inferotemporal pathways for visual processing. *Psychobiology* 17:236-246.

- Maguire, E.A., C.D. Frith, N. Burgess, J.G. Donnett, and J. O'Keefe (1998). Knowing where things are: parahippocampal involvement in encoding object locations in virtual large-scale space. *Journal of Cognitive Neuroscience* 10(1):61-76.
- O'Keefe, John and J. Dostrovsky (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. *Brain Research* 34:171-5.
- O'Keefe, John, and Lynn Nadel (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon.
- 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.
- Olson, Carl R., and Sonya N. Gettner (1996). Brain representation of object-centered space. *Current Opinion in Neurobiology* 6:165-70.
- Petrosini, Laura, Maria G. Leggio and Marco Molinari (1998). The cerebellum in the spatial problem solving: a co-star or a guest star? *Progress in Neurobiology* 56:191-210.
- Poucet, Bruno (1993). Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review* 100(2):163 - 182.
- Pouget, Alexandre, and Terrence Sejnowski (1997). Spatial Transformations in the Parietal Cortex using Basis Functions. *Journal of Cognitive Neuroscience* 9(2):222-237.
- Roskos-Ewoldsen, Beverly, Timothy P. McNamera, Amy L. Shelton, and Walter Carr (1998). Mental representations of large and small spatial layouts are orientation dependent. *Journal of Experimental Psychology: Learning, Memory and Cognition* 24(1):215-226.
- Stein, J.F. (1991) Space and the parietal association areas. In Jacques Paillard, ed. *Brain and Space*. Oxford: Oxford University Press.
- Strawson, Peter F. (1959) *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.
- Tolman, E.C. (1948). Cognitive maps in rats and men. *Psychological Review* 55:189-208.
- Yantis, Steven (1993) Stimulus driven attention capture. *Current Directions in Psychological Science* 2(5):156-161.
- Zipser, David, and Richard A. Andersen (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331(6158):679-684.