Manifolds, coordinations, imagination, objectivity

Rick Grush
Center for Semiotic Research, University of Aarhus, Denmark
rick@twinearth.wustl.edu
kultrg@cfk.hum.aau.dk

Draft, please do not quote or circulate.
Current word count (including notes and references): 15,427

> *Each of us distinguishes between himself and states of himself on the one hand, and what is not himself or a state of himself on the other. What are the conditions of our making this distinction, and how are they fulfilled? In what way do we make it, and why do we make it in the way we do?*
>
> Peter F. Strawson, *Individuals*

1. Preliminaries

Much current naturalistically inclined philosophy of mind has set up its defining problem in such a way as to echo Frege's puzzle of the informativeness of identity statements.[1] To take but one of dozens of examples, Fodor (1987) goes to some pains to ensure that his theory of content will entail that a cognitive token can have the content 'Jocasta' without having the content 'Oedipus' Mother', even though, in fact, Jocasta is Oedipus' mother. For Frege, of course, the parallel problem in linguistic semantics was the motivation for supplementing his theory of *Bedeutung* with a theory of *Sinn*. Part of the reason for this similarity between linguistic semantics and topics in the philosophy of cognitive science is historical, tracing back to a period during which the causal theory of names and theories which placed causation at the heart of epistemology crossed paths.[2] This interbreeding had the consequence that even well after the epistemological theories had given rise to theories of mental content, the latter are still largely concerned with getting ascriptions of mental

---

[1] Frege, 1952.
[2] See, e.g. Kaplan, 1969, Devitt (1974), Harmon (1977). Fred Dretske's classic *Knowledge and the Flow of Information* is a perfect example of this hybridization.

content to pass Frege's puzzle about the cognitive non-equivalence of coreferring representational vehicles.

There is nothing wrong with this task. Surely, something doing work similar to that done in linguistic semantics by Fregean sense will play a role in the ascription of content to cognitive tokens, so we can hardly fault such theorists for implicitly adopting it as a principle criterion to be met by fledgling theories of content ascription. What is to be faulted is the fact that, with very few exceptions, the naturalization of content has not seen fit to recognize anything beyond this neo-Fregean content-slicing as in need of investigation. And because in fact most naturalistic philosophy of mind has been obsessed with the project of naturalizing content, it has not advanced much beyond this content-slicing game. But there is another dimension to the content of our mental episodes which is not captured by such theories -- it is not even explicitly recognized. In this article, I will try to explain what this dimension of content is, and to lay the groundwork for a naturalistic theory of it.

Though I will put off for the moment the introduction of a name for this dimension of content, we can start getting an idea of what it is right away. To put it in a nutshell, there is a world for you in a way that there is not a world for an altimeter or a tree cross-section. You conceive of yourself as a subject with a point of view on a world which is largely independent of your experiencing it, and is 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, and your location and orientation in it. And you conceive of the world as consisting of objects, properties, etc., which are accessible from points of view other than your own. Much more will be said about this in the balance of this paper, but for now notice that this dimension to the content of our experience is plainly more basic and important than any dimension of content dealing with cows, horses, Homer, or Gödel. The capacity for

'outward' directed content is a logical precondition for even raising the question what the content of a referring thought is, let alone answering it. For, in the case of thought at least, the ability to entertain contents about objective extants *qua* objective is a logical precondition for entertaining a thought about any particular object *qua* external object. And it is equally obvious that this dependence is not symmetric. For instance, part of the content of a perceptually-based thought we might express with the words 'that cow is eating grass' is that the cow is an enduring object in a spatiotemporal realm that exists independently of my perception of it, and is the same cow that someone else might see from another point of view. A creature which could not conceive of an objective world of objects distinct from itself could not legitimately be credited (de dicto) with any thoughts about cows or horses on dark nights at all. But a creature could very well have outward-directed, or intentional (in Brentano's sense of the term) contents without being able to have any thoughts about cows.

The justification for calling this conceived subject/object divide part of the content of such a thought is quite straight-forward. The task of naturalizing the mind is the task of understanding how a physiological system supports a psychological system. And this involves the assignation of contents to physical states. One reason this needs to be done is that only so will it be possible to see how physical motions can be actions. Anything which plays an explanatory role *as a reason* (as opposed, perhaps, as a brute physical cause) why an agent acted in such-and-such a way must be a content. Failure to treat such reasons as a component of content to be naturalized with the rest of the components of content would render the aspiration of the naturalization of the mind unattainable, for there would be reasons to which an explanation of action would need to appeal, but which would be beyond, *ex hypothesi*, naturalistic explanation. All this should be uncontroversial. Part of the *explanation* of why you hide the cookies from me is that you represent the cookies as objects which are available to more than just your 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 perception of them, and hence as continuing to exist in the hiding place while you are not there. That is to say, *you represent the cookies as being independent of your point of view*. This is something not at all implied by simple sensorimotor engagement with cookies.

We can, for now, call this feature of content the 'subject/object divide'. It has been argued that a naturalization of the subject/object divide is impossible, if not an oxymoron. Three reasons for this intuition are first, the subject/object divide has often been confused with 'consciousness', or something similar, and this is often assumed to be unnaturalizable. Second, the subject/object divide is assumed to be a primitive feature of experience, and hence unexplainable. Both of these beliefs are wrong, at least given the sense of 'subject/object divide' I will be discussing. But given that we can leave consciousness out of the picture, and that the subject/object divide is the fruit of the interplay of a number of more basic cognitive mechanisms (the arguments for both of these are forthcoming), the naturalization of the subject/object divide would proceed in two steps. First, there would be an account of the various cognitive mechanisms whose joint interaction yields this divide. And second, there would be an account, in neurobiological terms, of each of these mechanisms and their interaction. This paper will address only the first of these two projects in any detail -- this is why I described this paper earlier as laying the groundwork for such a naturalization. But in the penultimate section, I will quite quickly highlight possible avenues of naturalization of these components via neurobiology. Clearly, then, this essay is aimed at the core part of what Kathleen Akins (1996) has dubbed the brain's *ontological project*.

The third and final reason that the naturalization of the subject/object divide is assumed to be impossible is that science, it is claimed, is necessarily a third person affair, while

subjectivity is essentially first person. But this argument is a bald nonsequitur whose continuing influence constitutes one of the true mysteries of philosophy. The task of science is to explain, not to recreate. It is no more an embarrassment to science that knowledge of neurobiology might not let the theorist 'know what it is like' to be the subject than is the fact that knowledge of black holes fails to create an event horizon around the theorist. The argument acquires whatever force it has only through the stratagem of couching the goings-on of the subject in terms of 'knowledge', and then claiming that this 'knowledge' cannot be had by a theorist through third-person science. In such shenanigans I must admit that I am unable to see anything other than a surprisingly seductive sophistry.

As a final preliminary, I should point out that for purposes of this paper, I will be using the terms *subjectivity* and *objectivity* in a constrained way. First and foremost, I do not intend to engage in old-fashioned prescriptive metaphysics, and hence I will not be undertaking an investigation into what objects *really* are, or what subjects *really* are, apart from what our common conceptual scheme says they are. Rather, I am interested in understanding how a system[3] comes to have a conception of an objective world, and a conception of itself as a subject with a point of view on that world. We are such systems, and hence this investigation concerns the cognitive machinery which accounts for this freature of our conception of things. Thus by 'subjectivity' I mean the property of conceiving of oneself as an entity with a point of view on an objective order. By 'objectivity' I mean the property of being represented as independent of the point of view of the representer. For instance, on this gloss, hallucinations are objective (i.e. have objectivity) if the representer does not realize that they are hallucinations.[4]

---

[3] I will use the term 'system' as a blanket term for a candidate psychologically describable entity. Thus humans, nervous systems, computers, and much else will count as a system.

[4] Plainly, then, the approach I am adopting is entirely **compatible** with the eliminativism of Paul Churchland (1979, 1989). That is, I am assuming that subjectivity and objectivity are the result of some sort of theory (where 'theory' is to be construed very broadly). Unlike Churchland I am not content to leave the matter there, for it strikes me that if the goal is to understand the functioning of the human brain, then we need to understand the nature of the theories that brains adopt, regardless of their truth or falsity. I think in many cases, including this one, much more of interest can be said than that 'the theory is a point in

This approach might require a bit of defense, because many will think that all and only things which 'really are' subjects could represent themselves as subjects. This intuition equates *is a subject* and *is conscious.* But this is simply a confusion. That consciousness is not sufficient for having subjectivity was demonstrated by Hume[5] long ago, and more recently by P.F. Strawson (1959). One can, perhaps with a little effort, imagine the mental life of a conscious being which did not conceive of itself as a subject, which had experience, but, in Strawson's phrase, had no use for a distinction between itself and its states on the one hand, and things not itself or its states on the other. This would be the *Ideal Humean Mind*, which was free of the errors and confusions that lead *us* to posit personal identity, objects, causality, and the rest. The question about *whose experiences they are* will arise for us, as theorists, but not for the Ideal Humean Mind itself. The fact that there very well could be a 'what it is like' to be a being which has no inkling that its experience is of a world, or that there is even a subject of experience at all, as Hume and Strawson have indicated in their own ways, should be enough to silence those who have assumed the equivalence. At a minimum it should take the burden of proof out of my cart and place it squarely in theirs. Consciousness is not sufficient for subjectivity. I believe that it is also not necessary, but I am not able to defend this claim here. The defense is implicit in what follows, and will be that the account of subjectivity I will give in this paper makes no appeal to consciousness. If that account is successful, then consciousness will be shown to be neither necessary nor sufficient for subjectivity.

So my approach is to dissociate consciousness and subjectivity, and this has a number of payoffs. For starters, it allows the two problems to be tackled separately. This is not only an advantage for me, but will be so even for philosophers whose intuitions are quite unlike

---

weight space'.
[5] Hume, David (1888).

my own. Those driven to a sort of pan-psychism[6] because they view consciousness (or experience) as a feature of the physical world have the task of distinguishing human consciousness from that of plants and salt crystals. But with a distinction between subjectivity and consciousness at hand, such philosophers can avoid either biting the 'hard' bullet (not distinguishing them) or making vague references to complexity. The way out would be to recognize that a human consciousness will be one of the few which is also subjective, in the sense I have described.

In my own case, the payoff is strategic: association of subjectivity with consciousness renders the former notion as vague, poorly understood and resistant to clarification as the latter. Dissociation is my way of trying to doff the proverbial cement galoshes. But I have a fallback position open to me if this dissociation is not allowed. This is that the balance of this essay can be considered as spelling out a feature of conscious beings which plays a role in their status as subjects. Thus a real subject, this fallback position continues, will be one which is organized as I will describe, and in addition is conscious (however that comes about). This is the minimal concession needed if in fact consciousness is a necessary condition for subjectivity. And with those few remarks, no more of consciousness.

Like any worthwhile divide, the subject/object divide has two sides -- in this case the subjective and the objective. I have already restricted my use of these terms to a descriptive, rather than prescriptive metaphysical, sense. Even so, some more words on what I am after are called for. It will be convenient to start with an exploration of the notion of an object. In one sense of the term, objects are entities over which it is possible to quantify, and which exhibit property bivalence. In a second sense of the term, objects are (roughly) entities which have an independent, more or less continuous existence, and which, though they typically can be experienced, are not themselves dependent on being experienced for their

---

[6] E.g. Chalmers, 1996.

existence -- I should like to sum this up by saying that objects are entities which are conceived of as being independent of the subject's point of view. It is this second sense which is the important one for understanding the subject/object distinction. For there is no shortage of entities which can be quantified over, and exhibit property bivalence, but which fail to be *objective* -- there are non-objective objects, one might say. Headaches, afterimages, and opinions, all fail to be objective, yet can obviously be quantified over, and exhibit property bivalence as well as rocks and rats. And their failure to be objective is not a result of their nonmateriality -- after all, shadows and magnetic fields can be objective. A good way to express the point is that objects, in the required sense of the term, exhibit a certain represented metaphysical distance from subjects of experience. It is something which a subject represents as being independent of its point of view (whether or not it really is so). Henceforth, when I use the term 'object', it is to be understood in this sense.

On the other hand, in order for a system to be *subjective*, it must represent itself as an enduring entity which has a position in and a point of view on an objective order, and the character of whose experience is determined in part by its position in that order.[7] As I have glossed them, objectivity and subjectivity are interdependent. The objective is that which is conceived of as independent of the subject, and the subject is that which is conceived of as having a particular continuous point of view on the objective. There is no circularity here, at least none vicious. All it means is that subjectivity and objectivity as described here form a local holism, as Peacocke (1992, p. 10) uses the term. Peacocke is interested in concept possession, but I think the point can be generalized to notions which may not be concepts. So for instance, Peacocke describes a concept F as 'That unique concept C to possess which a thinker must meet condition A(C).' A local holism is a set of concepts such that the possession conditions for any one are identical to the possession conditions for all. I.e.

[7] It will be noticed that I have used the usual spatial metaphors for describing the situation: the subject is 'in' an objective order, etc. But for the time being, let these be taken as metaphors, for the role that that spatial representation plays in objectivity has not yet been discussed.

Concepts F1, ..., Fn are those concepts C1, ..., Cn to possess which a thinker must meet condition A(C1, ..., Cn). Similarly, the claim I am here adopting is that subjectivity and objectivity, while not necessarily concepts, are nonetheless a local holism in that a system 'has' either, and hence both, only in virtue of satisfying some single possession condition. The interdependence of subjectivity and objectivity has lead Adrian Cussins (1992) to introduce the neologism *sojectivity* to capture this feature of our cognitive lives. It is the metaphysical distance which opens up when the world and the subject are separated.[8] I will adtopt Cussins' term. To put it in Peacocke's idiom, then, we are looking for the condition that must be met by a system in order for it to have sojective representational contents. Explaining this condition an a manner amenable to naturalization is this article's task.

Though there is considerable intuitive appeal to this interdependence, it has not gone unchallenged. Strawson, for example, thought that it would be possible for a system to conceive of a world as independent of its experience, and yet not conceive of itself as a subject. He felt that the notion of the self as a subject was borne of interpersonal interaction and communication. It will be seen later that Strawson's intuition depends on confusing two distinct but often conflated ideas -- objectivity and existence unperceived. These two are often confused because perception is typically given too much weight, at the expense of action, in defining a point of view. But in the next section we shall see that a point of view depends at least as much on action as it does on perception, and hence it is quite possible for entities to be conceived of as independent of perception, and yet still dependent on a single point of view. We shall explore this more in the next section.

But it might also be thought that one could have the notion of a subject without the notion of an objective order. Imagine not the Ideal Humean Mind, but rather the *Pure Cartesian*

---

[8] Though Cussins work has been crucially influential in my own thinking on these matters, and though I will adopt his terminology, I make no guarantees that my use of the term *sojectivity* will match his use in all respects.

Ego, for whom the limits of the universe are the limits of its own experience, thoughts, and memories, which has no use for the notion of anything other than itself, and which, unlike the Ideal Humean Mind, conceives of these experiences and thoughts *as its own*, as the thoughts and experiences of an enduring subject. But it is not evident that this is really coherent a possibility. Keep in mind that what is important for objectivity and subjectivity is not one's evidence for or belief in the truth of one's ideas about the objective order, or of oneself. Rather, it is the *capacity to conceive* of something independent of oneself which is at issue -- a quite different matter. Even in his most psychotic fantasies, Descartes never challenged the possibility of, as opposed to knowledge of, the objective.[9] The mischievous demon, for instance, is always a possibility, and has an existence conceived of as independent of the ego. Indeed, what import or content there could be to the self notion when there is literally nothing that can even be conceived of as non-self?

Let's initiate this investigation by considering the role that spatial representation plays in sojectivity. The conviction that sojectivity depends on the ability to represent space goes back at least as far as Kant, who tells us that "[t]hrough space alone is it possible that things should be outer objects to us." [A29][10] Exactly what is meant by 'representation' and what is meant by 'space' will emerge in the course of our investigation, but for now, we can get off the ground with the following passage from Strawson's *Individuals*:

> 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

[9] It is precisely on this point that the epistemological heritage of theories of content has made its deepest mark. Neither the epistemological worries of Descartes, nor those of Dretske, were concerned with objectivity. They assumed it.

[10] It might be thought that Kant would deny that sojectivity was dependent on space, since inner sense is given through time alone. Perhaps for Kant, this line proceeds, a subject could conceive of itself as such without any spatial intuition, pure or otherwise. But it is not clear that this is the case, for Kant never envisaged the possibility of a subject lacking pure spatial intuition. It might be fair to say that for Kant, sojectivity is pure *a priori* exactly because space as a form of outer sense is non-optional, and *a priori* -- at least for us.

> exist continuously, though unobserved; it supplies this dimension for objects which are not themselves intrinsically spatial, such as sounds, as well as for objects that are. Thus the most familiar and easily understood sense in which there exist sounds that I do not now hear is this: that there are *places* at which those sounds are audible, 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 need considerable refinement before anything serviceable emerges. In particular, Strawson's formulation is saturated with the conflation of objectivity and perception-independence. I propose the following as our starting point, which already involves some refinements: sojectivity is the result of a system's representation of space, such that i) the system represents its point of view as being anchored[11] somewhere in this space, ii) the system represents this space as being independent of the system (i.e. space is represented as independent of the system's point of view), iii) at least some entities in this space are such that they are in principle accessible from many points of view, in particular points of view other than the system's own current point of view.

It might be objected that this initial analysis robs Peter to pay Paul, in that the problem of the source of the wherewithal to represent objects as independent of the point of view of the subject (or for short, POV-independent) is addressed by an undefended appeal to the ability to represent space as POV-independent. This objection must be allowed. However, progress has still been made, in that if we can give an account of the ability to represent POV-independent space, then we will be in a position to provide an account of sojectivity. Kant himself addressed this problem through a notion of 'synthesis'. I refer here not to the synthesis of empirical objects, but to the 'pure' synthesis of the forms of sensibility themselves, space and time. The imagination is active in this pure synthesis no less than in

[11] I will use the term 'anchored' to mean that something has both a location and orientation in some domain.

the synthesis of the empirical intuition of objects. In his deservedly famous commentary, Henry Allison explains that for Kant

> ... the imagination is required for the representation of time and space as they are described in the Transcendental Aesthetic. Let us first consider the case of time... ...each extent of time is represented as a determinant portion of a single all-inclusive time, which is itself characterized as an infinite magnitude. It follows from this that the actual awareness of a given portion of time... involves the awareness of it as a portion of this single time. As we have already seen in the case of space, however, this whole is not itself actually given in intuition as an object. In the case of time, we could say that it is given only one moment at a time.... I must be able to represent times that are not 'present', and ultimately the single time of which all determinate times are parts. This is what the imagination enables me to do. Similar considerations apply to space. (Allison 1983, p. 160)

So for Kant, the ability to represent objective space relies on two capacities: the capacity to have an immediately given space, a space which is given in intuition (like the single moment is given to intuition at a given time) and the capacity to represent this given space, with the aid of the imagination, as a determinate part of an objective spatial realm. For the sake of some consistent terminology, let us call the 'immediately given' space *egocentric* space, and the objective spatial realm *objective* space. 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; or to put it in Kantian terms, it is that objective space which one gets by (purely) synthesizing, with the imagination's aid, an infinite spatial realm anchored to *here*. And the Eiffel Tower is 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 sojectivity, has assumed that it relies on the ability to represent objective space. And if we follow Kant, which I suggest we do, at least for a while, then objective space is arrived at by using the imagination to anchor the immediately given space in an objective realm. This involves three separate

components: a capacity to represent an egocentric space, the capacity to conceive of 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 an egocentric space that we first turn.

2. Mach's Thesis

Our task now is to gain an understanding of egocentric spatial representation. Objective space will wait until the following section. An interesting perspective on the naturalization of spatial representation was articulated by Ernst Mach in *Space and Geometry* (1906). Mach argued that objective[12] space was an abstraction from 'physiological space' which was itself built out of atomic spatial information delivered by the senses and motor activity. This line proceeds in two stages. First are arguments to the effect that physiological sensation is intrinsically spatial in character. The second stage involves hypotheses about how objective space is extracted from this physiological space. The second stage of Mach's argument will not concern us, but the first stage runs something like this: from the observation that animals have sense organs composed of one- or two-dimensional arrays of sensory elements (such as tactile receptors in the skin, or photoreceptors in the retina), such that each element by itself can distinguish among some number of stimulus conditions, Mach concludes that such an arrangement would be well suited to providing two kinds of information, which might be called 'what' and 'where' information.[13] The 'what' information is provided by the specific response of the element is stimulated, regardless of which element it is (a single tactile receptor might respond in one way to light stimulation and in another to greater pressure), whereas the identity of which element in the array is responding, regardless of the character of the response, provides information about where

[12] Mach's term was not 'objective' but 'geometric', but by this he meant independent of any point of view, and so he clearly had in mind the same notion as do I.
[13] The distinction I am making here is not quite the same a that often made in neurobiology to distinguish between he so-called dorsal and ventral paths in the mammalian visual system.

this stimulus is located. Mach's conclusion is that "[t]he perfect biological adaptation of large groups of connected elementary organs among one another is thus very distinctly expressed in the perception of space."

Such an array of elementary sensory receptors makes available a 'physiological space' of a certain character -- visual space, tactile space, etc. About these Mach remarks:

> The physiological spaces of the different senses embrace in general physical domains which are only in part coincident. Almost the entire surface of the skin is accessible to the sense of touch, but only a part of it is visible. On the other hand, the sense of sight, as a telescopic sense, extends in general very much farther physically. ... Yet, loosely connected as the different space sensations of the different senses may originally have been, they have still entered into connection through association... (Mach (1906), p. 15)

The egocentric space of the organism, according to Mach, is built from the spaces given by sensory arrays and the behavioral actions appropriate to the associated stimuli. Space is thus given directly in sensory experience, and merely refined by coordinating the partial spaces of the senses.

There are two lessons to be drawn from Mach's position. First, his argument as stated works only by trading on an ambiguity in the term 'space'. In one sense of the term, space can be taken to refer to the three-dimensional external space of the extended physical universe, either in its objective form, or its egocentric form. In a second sense, 'space' refers to any set of ordered elements. To see the difference, note that the pitches to which the human auditory apparatus is sensitive are often arranged along a 1 dimensional continuum, from lower to higher. Indeed, receptors (hair cells) in the cochlea are arranged as a 1 dimensional line of hairs sensitive to different pitches.[14] As such, they comprise a physiological space by Mach's standards. Nonetheless, no one would mistake this 'space'

[14] Actually, only the inner hair cells are arranged in a more or less straight line. The outer hair cells are arranged in a somewhat broken wavy configuration. But the important point is that the hair cells occupy positions along the effectively 1 dimensional cochlea.

for space. What the sensory organs deliver are 'spaces' only in the second sense, not spaces in the first. One might object on Mach's behalf that in the case of audition, the analogy breaks down, because there is no split between 'where' and 'what' information as there is with the other sensory surfaces. But this is wrong. Different timbres, such as produced by a flute and a violin, can occupy any of the allowable pitches. As far as the physiology goes, there is no disanalogy between treating the timbre as the 'what' and the pitch as a 'where'.

We are evidently in need of some more discriminating terminology. Accordingly, I will use the term 'manifold' as a blanket term for a set of elements over which a distance measure is established. This includes, as a special case, the continuous ordered elements of locations which make up real space, but also includes the 1-D arrangement of pitches, etc. As such my usage is inspired more by the mathematician's use of 'manifold' than Kant's use. I thus leave the term 'space' to refer to real space, either objective or egocentric. With this distinction in hand, we can immediately see why Mach's argument fails. There is no more reason to interpret sensory information delivered from the two dimensional sheet of the retinae as spatial than there is to interpret different pitches as indicating different spatial locations. What we are given in sensation is, at best, manifolds. What we need is an explanation why some sensory manifolds give rise to experience interpreted as spatial, while others do not.

This leads to the second lesson to be drawn from Mach, which is the idea that space, as perceptually experienced, is tied to the organism's *sensorimotor activity* in its environment. Recall that for Mach, the various sensory manifolds are coordinated by association. Mach clearly meant it to be the case that these associations are effected through motor engagement. It is through action that the various sensory manifolds are brought into appropriate coordination, because it is through action that the interrelations between, e.g.,

sight and touch are made manifest. This second idea provides a support which may bear the weight of Mach's conclusion even when his argument from physiological sensory arrays has crumbled -- Mach may be right, though for the wrong reasons. What distinguishes bare manifolds from space is that, as it were, real space is where one's sensory states and behavioral actions are appropriately integrated. A better way of putting the point, which I will call **Mach's Thesis** is this: *egocentric space is the represented manifold in which sensations and motor activities are coherently and appropriately coordinated*. For example, it doesn't matter if a creature locates its food as being at location x by sight, touch, hearing, or olfaction. If it is localized at location x, the same equivalence class of motor programs will be sufficient to obtain the food. Similar systematic coordinations between perception and action simply do not obtain for pitch.

Note that what I have been calling motor actions can be regarded as a different type of manifold. In much the same way that one can specify a sensory location by means of coordinates tied to the sense organ (retinal photoreceptor *a*, or tactile receptor *b*, where *a* and *b* are elements in an ordered array), 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 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). My 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.[15] As such, what Mach referred to as 'associations' between the physiological spaces can now be described more accurately as coordinations among manifolds. And

[15] Many manifolds will accordingly require a distance measure, but it need not be the case that the manifold is metric, because there is no reason to require that the distance measure be such as to allow the manifold to satisfy the triangle inequality.

Mach's Thesis is the statement that an egocentric space is a manifold which successfully coordinates sensory as well as motor manifolds.

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. This is the sort of coordination Mach was describing when he remarked that part of the body accessible to touch was also accessible to vision. Such areas of overlap establish a c-coordination between the two manifolds. Better examples of this will be along shortly.

An s-coordination is the establishment of a relationship between elements of more than one manifold which has the effect of stabilizing higher order features 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. The eye can rotate clockwise or counterclockwise. When a light appears on the screen, its direction relative to the eye cannot be determined by the location on the retina to which the light projects. This is because the same retinal position can correspond to many different points on the screen, depending on how the eye is oriented. But given both the location of the retinal projection, as well as the angle of the eye’s orientation, it is possible to determine the direction of the light. Consider Figure 1.

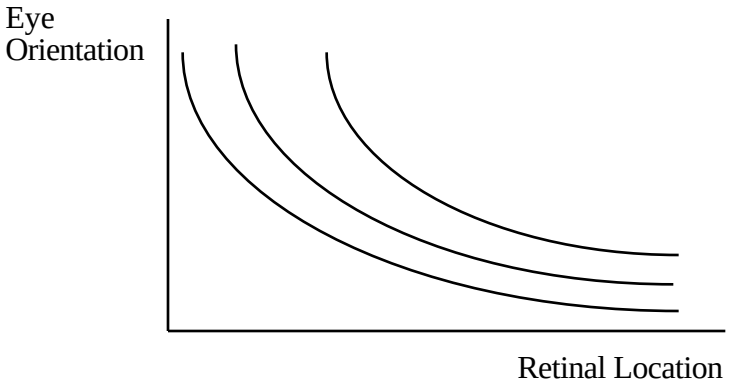

Figure 1

Here, retinal location 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 which all correspond to the same light location. By establishing these contours, a system is s-coordinating the two 1-D manifolds represented in the axes to stabilize a higher order 1-D manifold 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 which was selectively sensitive to points on a given contour would thereby be selectively sensitive to a given location. In this way the s-coordination of the manifolds stabilizes a representation of the possible locations of light sources.[16]

[16] I have argued elsewhere (Grush, 1997) that representations are possible only in off-line forward models. Thus what I say here might be thought to be in conflict with that position. There is an element of truth to this. But, the sorts of models which support representations can be used on line to guide action and process sensory information when playing a role in a Kalman filter. This was an architectural variant I did not discuss in that earlier article. Given the more sophisticated Kalman filter architecture, there is no conflict between the sorts of representations defended here, and the requirement, expressed in Grush (1997) to the effect that representations are only possible in forward models. For more detail on this, see Grush (in preparation).

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 also 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-dimensional 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 with its hand. The project of c-coordinating the region available to vision and the region available to tactile sensation requires a lot of work. Before they can be c-coordinated, each must be stabilized with respect to some common reference point, in this case the torso. To see what this means, consider the following. The bright object will project an image on 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 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. The wedge thus stabilized is not simply the wedge accessible to the retinae at any given moment, but rather the entire wedge anchored to the head which is accessible to the retinae, given the freedom of movement of the eyes. This is because every point in this entire wedge can be determined from the information available in the various manifolds. This 3 dimensional wedge is, as a set of ordered elements, a new manifold.

This visual wedge is of little use, however, because the head is free to move. This visual wedge needs to be stabilized with respect to the torso. This is done through stabilization with respect to the muscles which control the position of the neck, and mechanoreceptors and vestibular organs which give information about the head's orientation. 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 around 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 of the creature accessible to touch will be stabilized in an analogous manner, only though the s-coordination of rather different manifolds. When these two higher-order manifolds have been stabilized with respect to the creature's torso, it will then be possible to establish regions of overlap. These regions of overlap then underwrite the c-coordination of the higher order manifolds stabilized with repeat to the creature's torso: the 'visual space' and the 'tactile space'.

These examples have been painfully 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 interact so as to generate regions of sensorimotor stability. In all cases, what gets created is a manifold which stabilizes some range of sensorimotor activity, or which establishes regions of continuity between these stabilized manifolds. The end result of this process just described is the creature's *behavioral space*. Henceforth, I will use the term 'behavioral space' rather than 'egocentric space' because the former emphasizes the constitutive role of motor behavior in the stabilization of the space.

I should like to make it clear that the behavioral space thus made available is a representation, and is not the actual space surrounding the organism. Of course, the latter may be the 'referent' of the former, but that is a separate issue. I stress this because there is

a growing tendency among philosophers of psychology and cognitive science to off-load the work of cognition onto the world, a tendency associated with the labels ‘embedded cognition’ and ‘situated action’, among others. Such a philosopher might object to what I have said, claiming that it is the real space the system is in that coordinates perception and behavior, and that positing representations of this space is a manifestation of latent Cartesianism or unwitting allegiance to traditional artificial intelligence assumptions.

But such an objector would be making a fairly obvious error that can be exposed by cursory reflection. The space around the organism and the represented behavioral space can be doubly dissociated. First, virtual reality gear as well as dreams work because they induce the creation of a representation of behavioral space, regardless of how the system is actually situated in its ‘real’ environment. Second, various sorts of trauma to the central nervous system, notably lesions to the posterior parietal cortex of the right hemisphere, disrupt the representation of the behavioral space while obviously leaving the actual space surrounding the organism intact. If one thinks it is the actual space that does the work, then the debilitating effects of such lesions are difficult to account for. This leads to the antirepresentationalist’s final move, which is to claim that successful sensorimotor activity is a *context dependent skill*. Brain damage of the sort described is destroying the implementation base of a skill, the objection claims, and not any representational machinery. But this is an obfuscation. To be sure, there are context dependent skills: one must be in water (or some other suitable liquid) in order to swim. But one type of context dependent skill is representation-mediated activity. One can disrupt the success of such activity by compromising either the representations or the ‘context’ (what is represented). In either case, what one is doing is making it so that the representations misrepresent, and this *ceteris paribus* leads to failure. Merely labeling some ability a ‘context dependent skill’ shows nothing about its reliance, or lack thereof, on representations. The objector is making the same mistake as one who argues that the purpose of stop signs isn't to get cars

to stop, because their real purpose is to reduce the number of car collisions. In order to be worth anything, such an argument would require some independent reason to think that reduction of number of accidents is effected by some means other than getting cars to stop. Failure to show this leaves open the option that the both descriptions of the 'purpose' of stop signs are legitimate.

There are at this point a number of points worth mentioning. First, the behavioral 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. It is, as Kant maintained, a resource which is used to ‘make sense of’ the deliverances of all the relevant senses. But we have gone beyond Kant, because we have seen exactly how and why motor action plays a constitutive role in the stabilization of spatial import. The stabilizations required to create the behavioral 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.

These considerations lead to what I will call 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.

The coordination principle entails that part of the content of, say, a point in visual space is provided in part by how one would orient (a motor action) towards that point, and how one would move one’s arm in order to bring the hand to that point, such as THE THING I COULD GRASP BY REACHING THUS (where the material in all caps is a

nonconceptual content specification). Similarly, part of the content of a felt location is given by how one would visually orient to that location. Something like the coordination principle is behind the thought of Evans when he writes:

> The subject hears the sound as coming from such and such a position, but how is this position to be specified? We envisage specifications like this: he hear the sound *up*, or *down*, *to the right* or *to the left*, *in front* or *behind*, or *over there*. It is clear that these terms are *egocentric* terms: they involve the specification of the position of the sound in relation to the observer's own body. But these egocentric terms derive their meaning from their (complicated) connections with the actions of the subject...
>
> Auditory input, or rather the complex property of auditory input which codes the direction of the sound, acquires a spatial *content* for an organism by being linked with behavioral output... (Evans 1985, p. 384-5)

There is abundant empirical evidence to support the connection between spatial representation and motor engagement. I will mention only two old, but rather famous examples here.[17] The first is Held and Hein's (1963) kitten experiments. Through a clever apparatus involving harnesses and small gondolas, they arranged it so that two kittens were each exposed to nearly identical visual input from birth. The difference between the two was that only one of the kittens actively explored its surroundings; the other was passively moved in an identical manner around identical visual stimuli. In terms of the present theory, only the first kitten was able to effect a coordination between its motor and sensory (in this case only visual) manifolds. The result was that only the first kitten developed normal 3-D visual perception. The second example is Bach-Y-Rita's (1972) visual prosthesis device. This apparatus was an array of tactile stimulators, worn by a blind subject on the stomach or back, which was driven by a video camera, typically worn on the head. The device worked marvelously well at letting blind subjects 'see' their surroundings, but only after an initial training period during which subjects were allowed to move around and actively explore their surroundings while wearing the apparatus. That is, for subjects whose camera

[17] More empirical evidence, including much which is quite recent, is cited and explained in Grush 1995, 1997, in preparation.

was not controlled by their own movements, either because the camera was on someone else's head, or mounted on a table, or because the subjects were not allowed to move at all, such subjects did not get to the point where they interpreted the tactile array as providing distal spatial information. Subjects who wore the camera, and were allowed to move, rather quickly got to the point where they non-inferentially interpreted the tactile stimulation as spatial perception of their surroundings in a quite appropriate manner.

Notice that the behavioral 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 is anchored to the torso, but whose vertical orientation is aligned with the gravitational field. The exact nature of this implicit anchor point of the behavioral space is less important that the fact of its existence. This implicit anchor of the behavioral space I will call 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. It's role as an anchor does not depend on restricting its extension to nothing. Regarding the modal metaphor, it should be clear that a point of *view*, as I have describe it, is not tied to a 'view' in the strictly visual sense, nor limited to sensation in general.

It can now be seen why objectivity requires more than existence unperceived, if objectivity means *conceived of as independent of the current point of view*. This is because the behavioral space which implicitly defines the POV is, as we have seen, not exhausted by

sensory input, but is constructed from these together with motor behavior. There are entities which are stabilized with respect to the POV, but not necessarily tied to perception. For example, when driving a car the gearshift and brake pedal is often not perceived when not being used, but they are represented in the behavioral space nonetheless: the brake pedal is right there [insert the motion of the foot to the left], and the gearshift right there [insert the characteristic motion of the right hand down to the side, hand formed such as to grasp the knob]. But an entity might very well be stabilized within the behavioral space without being conceived of as being independent of that space. We typically do conceive of stable entities in our behavioral space as being independent of our POV, but that is an achievement to be explained, not a counterexample to the claim.[18]

We now have an answer to the question what distinguishes a merely physiological manifold from a genuine spatial manifold. The answer is that a spatial manifold is one which is constituted by the appropriate coordination of motor and sensory manifolds. I am thus explicitly claiming that behavioral space is a separate manifold, separate from those of any of the sensory or motor manifolds, which is stabilized from these manifolds, and by the coordination principle, supplies these other manifolds with spatial import. It is, accordingly, a supra-empirical neurocognitive resource. The reason Bach-Y-Rita's subjects could learn to interpret the tactile array as spatial so quickly (typically on the order of a few hours) is that these subjects already had a behavioral space, as all blind people do. All they needed to do was to learn to coordinate the new tactile/visual manifold with this pre-existing spatial manifold, as opposed to trying to learn to coordinate it with the dozens of

[18] We can now see why Strawson in ‘Sounds’ made the mistake of claiming that one can have the notion of an objective particular without the notion of oneself as a subject. Strawson took himself to have provided the machinery for objective particulars when in fact all he provided was particulars independent of perception -- existence unperceived. That one can have a notion of unperceived existence without the notion of the self as a subject, even if true, is not an objection to the interdependency thesis. For Strawson, the subject emerges only when forced to interact with others. If this is so, this interaction also forces one to conceive of particulars as independent of the point of view. So even adopting Strawson’s moderate social constructivism does not force the abandonment of the interdependence thesis. But we shall see that the social element is unnecessary, and in fact is dependent upon more basic nonsocial capacities.

distinct sensory and motor manifolds with which they were already familiar. Whether the brain constructs this behavioral spatial manifold on the basis of experience together with simple associative learning rules, or is innately specified to produce it under the appropriate conditions, is a question I need not take a stand on here (but my suspicions are that there is a strong innate component). Once properly understood, this question becomes one for cognitive neuroscience to answer.

One might be tempted to stop at this point. The thought will be that the things that are stabilized in the behavioral space are objects, and the agent which is at the 'center' of this space, which occupies the POV, is a subject. If so, then we have provided an account of sojectivity with our account of behavioral space. But this idea stumbles shortly out of the gate. First, what we are after is the cognitive machinery which allows a system to conceive of itself as a subject with a POV, not the mechanisms which merely supply a point of view. All that has been done here is the latter. We have shown that a system has a POV in virtue of having an behavioral space. It just is the implicit nexus of behavioral efficacy constituted by the coordination of sensation and action.

The second reason we are not at a position where we can stop has already been foreshadowed a few paragraphs back. The 'objects' which are available in the behavioral space are not yet objects in the required sense. Recall, what we need to account for is the ability to conceive of something as the sort of thing which is independent of the current POV. Since the behavioral space provides only one point of view, simply having an behavioral space -- having a point of view -- cannot provide this ability.

The present account is thus in conflict with Daniel Lloyd's (198@) so-called *dialectical theory of representation*. Lloyd takes it that having access to an object via more than one sensory modality (or even sensory receptor!) is sufficient for objectivity -- presumably

because the object will be given from more than one point of view, so to speak. But if the present account is correct, then individual sensory modalities do not constitute points of view at all. A sensory modality of a system is either coordinated with the other sensory and motor manifolds of that system, or it is not. If it is, then it contributes to the same point of view in virtue of being coordinated with the same behavioral space. If it is not, then it provides no point of view at all. We might adapt Evans' famous phrase thus: for such systems, there is only one point of view, because there is only one perceptuo-behavioral nexus.

This section has not yet revealed the full structure of the mechanisms which make sojectivity possible. But two important things have been established: first, the nature of behavioral space has been analyzed in a way which both distinguishes it from nonspatial manifolds, and makes our understanding of it in terms of neural mechanisms unmysterious (this will be explored briefly in the final section). Second, we have had the opportunity to get clearer on what we are after, to distinguish *independent of perception* from *independent of the point of view*. Our next job is to explore 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 objects.

3. Evans' Thesis

Much of Gareth Evans' later work addressed the nature of objectivity, spatial representation, and their inter-relations. In *The Varieties of Reference*, Evans makes the following remarks on the relation between egocentric (= behavioral) space, and objective space (I hope I may be forgiven the use of such a long quote):

> We now have to inquire what makes such Ideas of places in egocentric space adequate Ideas of positions in public space... Any subject who is able to think objectively about space -- any subject who can be credited with a cognitive map of any region -- must know what is involved in ... imposing his knowledge of the objective spatial relations of things upon an egocentric space. Someone who has a cognitive map of Oxford, for example, must be able to contemplate the imposition of this map in the course of his travels (perhaps in a very dense fog). 'If I am here, midway between Balliol and the Bodleian, then that must be Trinity, and so the High must be down there.' In such a situation, one may have to choose between several ways of effecting a coincidence between egocentric space and one's conception of objective space. Each way of effecting a coincidence would generate hypotheses about what one should be able to observe if oriented in this or that direction, and what one would observe if one moved in this or that direction. At the same time, of course, each way of effecting the coincidence would entail an identification between every discriminable point in egocentric space and some point in objective space...
>
> Such thought presupposes the ability to represent the spatial world by means of a cognitive map. But nothing that the subject can do, or imagine, will entitle us to attribute such a representation to him if he cannot make sense of the idea that *he* might be at one of the points representable within his map. We say that the subject thinks of himself as located in space (in an objective world that exists independently of him, and through which he moves); only if this is so can the subjects egocentric space be a space at all. But what does this thinking of himself as located mean except that the subject can in general regard his situation 'from the objective point of view'? And this means that in general he has the ability to locate his egocentric space in the framework of a cognitive map. (Evans 1982, p. 162-163)

What is crucial for Evans is ‘effecting a coincidence’ between one’s egocentric space and a cognitive map. Evans is nowhere entirely clear about what he means by a cognitive map, but we can get the mileage we need with the following example, taken from Evans (1985):

> If a, b, and c are envisaged to lie upon a straight line, when what is envisaged is an instance of the simultaneous concept -- a concept whose most direct application lies in a presentation of the three elements a, b, and c together -- then a, b, and c must thereby be conceived to exist in exactly the same way. If a, b, and c are believed to exist in such an arrangement when b is perceived and a and c are not, then a and c are conceived to exist, though not perceived, in exactly the same sense in which b, now perceived, exists.

In this example, there is an initial representation of the spatial relations between a, b and c (suppose that these are statues in a park that can be recognizable by sight). Call this representation the cognitive map. Suppose further that I can locate a and b in my behavioral space, perhaps because they are directly in front of me and in sight, a is just ahead and to

the left, and b is just ahead and to the right (I am here departing from the details of Evans' example). If my map of a, b and c lying on a straight line is brought into coordination with the a and b as represented in my behavioral space, then I will be in a position to locate c in my behavioral space even though I cannot perceive it; it is *just over there to the right*.

It is clear that what Evans means by 'effecting a coincidence' is a case of what I have been calling a c-coordination. It involves establishing identities between elements in one manifold (the behavioral space) and elements in another (the cognitive map). Among the elements identified is the point of view. The nexus which defines the behavioral space is identified with one of the myriad points on the map. Upon such a c-coordination, and by the coordination principle, the elements in the cognitive map inherit the locatability of elements in the behavioral space. Moreover, the elements in the behavioral space inherit the properties of objects in the cognitive map. What needs to be established is that the cognitive map represents objects and space as objective.

This is easy enough to do. Because the cognitive map represents points of view as locations in the space (e.g. the familiar 'YOU ARE HERE' arrow), there is nothing special about any of the points of view as such. The anchor of the system's own point of view has no special status *in the map* -- it is one point among many. This, together with the fact that the map is unaffected by altering the location of the anchor, guarantees that the elements within the map are thus represented as being independent of this point of view, because those elements and their relations remain the same no matter where one places the 'YOU ARE HERE' indicator. This brings us to what I will call **Evans' Thesis**:

The elements in the behavioral space achieve objective import when the behavioral space, *B*, and thus the elements in it, are c-coordinated with a spatial manifold, *A*, that is itself independent of *B*, and capable, in principle, of supporting arbitrary c-coordination with *B*.

Moreover, this very c-coordination involves placing the POV, which is only implicit in B, explicitly in A, and thus provides the necessary tools for conceiving of the POV *as a point of view* on an order independent of that point of view.

The remaining puzzle is that it is not initially obvious where the cognitive map, or as I shall call it, the allocentric manifold (*A* in the statement of Evans' Thesis above), comes from. The manifolds whose coordination results in the behavioral space are all immediately available to the system, as bare manifolds, in perception or action. But the allocentric manifold, by definition, is not. To address this, we return finally to Kant. Recall that for Kant, it is the imagination which allows one to represent objective space. I propose a particular reading of this, which is that the imagination provides a representation of space via the adoption of an imagined point of view other than the imaginer's actual point of view, such that the imaginer's actual point of view is one of the points explicitly represented in the imagined point of view.

It might be thought that this is not what Kant had in mind. And perhaps it isn't. But weather or not it is what Kant meant, I think it is correct. In addition, I think there is at least some reason to attribute this reading to Kant. The alternate reading of Kant would be that the synthesis of objective space and time proceeds by, as it were, appending more and more space and time to that bit of it given in intuition. For instance, while sitting at the corner cafe, I can project my attention in this or that direction in order to 'imagine' an indefinitely large spatial expanse projecting out from *here*. This idea has appeal, but it is not sufficient for objective space, because all one is doing is expanding the behavioral space. The appended spatial locations are represented as being *way way over in that direction,* for example, where the directions and distances represented are still given in egocentric terms, and hence are not objective.

Bigger blemishes become visible on this alternate reading when we turn our attention from space to time. In the first place, thinking of times in the distant past or future obviously relates them to the present, which is specifiable only from the behavioral anchor. Furthermore, in the case of time, there is textual evidence that Kant had in mind the adoption of an imagined viewpoint, rather than the imaginative augmentation of the actual behavioral space(time). Time, for Kant, could only be represented as a spatial magnitude: a line. This is clearly not the way time is given *to us* in intuition -- appending more future times or more past times to *now* does not result in an intuited line. In thinking of time as a line, we adopt the imagined point of view of a being which can ***intuit*** time as a spatial magnitude. It is the coordination of this 1-dimensional manifold with *time as we experience it* which provides *time for us* with objective import.

Similarly, it seems reasonable to suppose that in imagining objective space, we are imaginatively adopting another point of view on some spatial realm or other, and coordinating the space represented from this other point of view with our behavioral space. This other point of view might be a possible point of view, such as a view of a city from high above, or it might not, as with the point of view of a being which has all of space given to it in intuition without its point of view being located in that space -- the spatial intuition of God.

There are two final, and brief, detours before we can proceed to the next section. The first has to do with an Evansian objection to the position I have been defending.

> It is often said that in such thinking we are taking the third-person, or God's-eye, point of view, but... I reject this way of looking at the matter. This formulation expresses ideal verificationism; whereas in fact the thinking is truly objective -- it is from no point of view. (Evans 1982, p. 152)

First, I should point out that in making this objection, Evans is creating an inconsistency in his later philosophical views. In ‘Molyneux’s Question’, an article which is a direct inspiration of my own thinking on these topics, Evans argues that the ability to represent spatial relations as genuinely spatial depends on relating them to complex behavioral dispositions. We have already seen some of Evans’ remarks on this topic.

Obviously, I think Evans is right about this, but if so, then it is not possible to represent spatial relations save as relations between things in a (perhaps imaginary) behaviorally fertile arena. And this just means representing them from a point of view. The idea that there could be a representation of spatial relations as genuine spatial relations represented from no point of view is plainly inconsistent with Evans’ own views on spatial representation. Moreover, I think it is possible to drop the ‘no point of view’ requirement while doing justice to Evans motivations for adopting that requirement. These motivations are to avoid ideal verificationism, the position that the *meaning* of statements about, e.g. distance places or microscopic objects, is to be analyzed as the content of the experience of a possible, *ideal*, being that could experience the places, entities and properties which such statements concern. This is avoidable because in claiming that objective space is an imaginatively constructed manifold -- an imagined experience -- one is not thereby forced to the conclusion that the *meaning* of the element-representations in this way is exhausted by the content derived from these imaginings. Recall, for example, the coordination principle. Thinking of a location as objective involves thinking of it as playing a part in some imagined experience from another point of view, but also necessarily involves thinking of it as being related to *here* (via coordination with the behavioral space). It may involve much else as well. That such imaginings play a foundational role in objective thought no more entails ideal verificationism than the fact that the violin is a crucial element in Stravinski’s *Petrushka* entails that *Petrushka* is a violin solo.

The second detour is the redemption of an IOU issued in the opening section. There we saw that Strawson thought that a system could have the capacity to represent entities as objective even though it did not conceive of itself as a subject of experience.[19] This latter capacity, Strawson argues emerges from social interaction (1959, p. 84-85, 94-98).

There are two problems with Strawson's position. First, as we have seen, the system which Strawson describes in 'Sounds' does not in fact have the capacity to conceive of particulars as objective. At best it can conceive of them as not being dependent upon being perceived. But these sound particulars are still within Hero's behavioral space, because the Master Sound is immediately given to hero in experience, and is coordinated with Hero's other experiential manifolds. That is, because the master sound is given to Hero in experience, there is no need for Hero to conceive of the dimension defined by that pitch spectrum of that sound as independent of his point of view. In effect, it brings the entire auditory world within Hero's POV.

Second, and more importantly, it should be clear that the ability to imaginatively adopt another viewpoint is a logical precondition for engaging in interpersonal interaction , as such, at all. The opposite position, advocated by Strawson and many others, maintains that one achieves the ability to conceive of other viewpoints (and thus the ability to recognize that one's own point of view is just another point of view), from interpersonal or social interaction. But this gets things hopelessly the wrong way around. A system that has no capacity to conceive of other points of view will, in virtue of this very fact, be barred from conceiving of interactions with other subjects as anything other than encounters with objects in its behavioral space. Perhaps they will be objects that make strange noises, perhaps they will be objects that move about in ways that are a challenge to predict, but they will be no more than objects for all that. And all the noisy and unpredictable objects in

[19] ref.

Heaven and Earth don't add up to a single alternate point of view. Unless, of course, the system's proprietary cognitive apparatus already has the capacity to imaginatively entertain other viewpoints, including other spatial viewpoints. Such a system will have the tools to hypothesize (perhaps a hypothesis forced on it by innate mechanisms) that some of these noisy, unpredictable things have a point of view -- e.g. *that one over there is experiencing the situation like this* (where *this* stands for some imaginatively created experience or other). Much cognitive benefit surely comes from such interpersonal interaction, but the bare ability to concieve of one's own point of view as such, and other possible points of view, is not one of these benefits -- it is the capital which makes these benefits possible.

I will close this section with the following remark. We can now see that the construction of objective space is much like the construction of the behavioral space. Each involves the coordination of experiential manifolds. The difference is that the construction of the behavioral space involves the coordination of sensory and motor manifolds given directly in experience, while objective space is constructed from the coordination of the behavioral space with a manifold not given in experience at all, but rather provided by the imagination. I take this to be a vindication of Kant's position on the matter, even if he would not have expressed the details in quite this way.

4. Remarks on naturalization

It is obvious that an adequate account of sojectivity would cover much more than what has been addressed here, and would do so in more depth. Nonetheless, I am confident that I have at a minimum drafted a preliminary plan for the form such an account will take.[20] But

[20] I am uncertain of the relationship between my labors here and those of Cussins (1992), who also provides a sketch of an account of sojectivity, in addition to coining the term. Cussins' approach and my own are quite different, in part owing to my emphasis on neurobiological realization, and Cussins' insistence on maintaining a symmetric metaphysics. At a minimum I can say that it is not obvious to me

this is only half of the goal I have set myself. Then remainder concerns a naturalistic account of these mechanisms, in this case in neurobiological terms. Anything approaching an adequate exposition of these neurobiological mechanisms is clearly beyond the scope of this paper, but it will at least be possible to run through an annotated laundry list of these mechanisms, in order to provide some substance, however thin, to my claims of neurobiological implementation.

Imagination. I have elsewhere argued extensively[21] that the neural basis of imagination is neural circuitry that recapitulates the dynamics of overt activity. To take a simple example: during motor control, certain areas of the brain send motor commands to the musculoskeletal system (MSS), and on the basis of those commands, proprioceptive signals are generated and sent back to these control centers, in part to provide these control areas with information about the progress of the MSS. There is evidence to the effect that there are neural circuits that mimic the dynamics of the MSS: that is, these circuits, which I call *emulators*, take a copy of the command sent to the body, and process it so as to produce a signal which is a prediction of the proprioceptive signal which the real MSS will produce on the basis of those same commands. Such emulators facilitate fast movements, as they are positioned to provide crucial feedback faster than that feedback is available from the periphery. But if run off-line, that is, if driven by a motor signal which is suppressed from being sent to the MSS, this emulator produces internally generate proprioceptive information -- a.k.a. motor imagery.

The application of the strategy of emulation is not limited to motor imagery. Visual imagery is also plausibly the result of the operation of emulators of the motor visual loop. Given a current visual presentation, such as a retinal image, and the current motor command (such as a step forward, or a saccade to the left), the next visual presentation is in part predictable

---

that our accounts are incompatible.

[21] Grush 1995, 1997.

(an enlargement of the retinal image, or a shift of the image to the right). Neural mechanisms can learn these regularities and exploit them to generate visual imagery.[22] There is evidence that many of the cortical sensory areas subserve this emulation function in that they can be taken off-line, and driven directly by mock motor commands. This however is no mean feat, as the cortical mechanisms which subserve sensory processing must have special circuitry and learning opportunities for them to be run off-line in this manner -- not any old system with a sensory processing component will be able to do this.

In the case of the behavioral space, matters are no different. Because the behavioral space is constructed from coordinations with motor manifolds, it is possible for the entire behavioral space to be taken 'off-line' and run via efferent copies. Something like this no doubt happens in dreams, but it can also be effected without dreaming -- you can imagine what it would look and feel like to walk around your desk and look at it and everything on it from a different angle. When this happens, one is exploiting the coordinations between motor and sensory manifolds to imaginatively move the behavioral space, or even create a new imagined behavioral space, by driving the behavioral space with efferent copies of motor commands which are suppressed from acting on the real motor apparatus. In this way, experiential structures can be created which are not linked to the real behavioral space and the point of view it supports. These counterfactual experiential structures just are allocentric manifolds.

S-coordinations. The posterior parietal cortex, particularly on the right hemisphere, seems to be engaged in exactly the sort of s-coordinations described in Section 2, and exactly for the purpose of building a coherent behavioral space.[23] For instance, there are cells in this area which are selectively responsive to locations relative to the head. These cells combine

---

[22] My discussion of the neurobiological implementation of imagery is torturously brief. Much more detail on this proposal can be found in Grush, 1995, in preparation. Also see Kosslyn, 1994, Mel, 1986.
[23] See Anderson, et all, @@@, @@@, @@@.

information about retinal location of a stimulus with information about the orientation of the eyes in the head in order to stabilize a reference frame anchored to the head. One plausible mechanism, advocated by Anderson et al (19@@) is that such cells implement gain fields: each cell responds to stimuli at a certain retinal location, but also has its activity modulated by eye orientation. The result is that the cell fires most strongly in response to a particular combination of retinal location and eye orientation, and such particular combinations signal orientations *with respect to the head*. Other groups of cells are involved in responding selectively to locations in a reference frame anchored to the torso, by being sensitive to head position relative to the torso in addition to retinal position and eye orientation. These cells are also implicated in motor action, such as reaching behaviors,[24] which again is to be expected, as reaching behaviors require that the goal of the movement be stabilized with respect to the torso. Auditory and vestibular inputs to this area are also well documented.[25]

Cognitive maps. That areas of the brain, notably the hippocampus, acts in some respects as a cognitive map has been known for some time. For example, O'Keefe and Nadel (1978) found that there are individual calls, called 'place cells' in the rat hippocampus which fire when and only when the rat is in a specific location in its environment. As the rat moves about, different place cells become active. But this sort of 'map' does not provide what is required, because it does not allow the animal to represent its point of view as such, it merely provides information about where the point of view is located. For instance, if the rat is at location L and then issues a motor command which takes it to location M, the map will represent this change accordingly, by altering which place cells are firing. At no point in this process is the capacity to imaginatively disengage the point of view needed. But O'Keefe (1994) makes a speculative claim about the ability of such maps to be run off-line, in much the way I have described the off-line running of the behavioral space:

[24] ref. of review article.
[25] refs @@@

> I have speculated that in [animals more cognitively developed than the rat] 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. (1994, p. 40)

O'Keefe's description is not quite correct, for merely allowing the 'you are here' arrow to be imaginatively moved does not result in the representation of the environment from a different perspective. Representing the environment from another perspective requires taking the POV off-line as well as the 'you are here' element in the allocentric manifold. But his speculation is crucial nonetheless, because when the POV is taken off-line (as described above) the same suppressed motor commands which drive it can also drive the location indicator in the map, the place cells. It is this coordination of the moving POV and the moving place indicator which, via the coordination principle, supplies objectivity. This is what Evans took such pains to express. It is the ability to run both spatial representations off-line, as well as to coordinate them wether on- or off-line that provides objectivity. It allows, via the coordination principle, the elements in the map to be concieved of as real, because they are linked to *here* (recall James Bond and the Eifel Tower), and allows the elements in the behavioral space to be conceived of as objective, because they are available from other perspectives.

C-coordinations. C-coordinations involve the establishment of identities between elements of different manifolds, or more generally, between different features in distinct feature manifolds. For example, in order for me to conceive of the building in front of me as objective, I must establish an identity between the building represented in my behavioral space and the building represented in my allocentric representation. This problem has been addressed in the neurobiological literature under the heading of the *binding problem*. For example, the mammalian visual system processes different features of visual stimuli, such

as color and motion, in different cortical areas. While this distribution of labor has advantages, is introduces the problem of keeping track of which features belong to the which objects. This problem is most clear when one considers the visual processing of two objects, say a stationary blue one and a moving red one. If the features are processed in different locations, how is the brain to keep track of the fact that *red* and *moving* are to be identified, and not *blue* and *moving*? Much evidence suggests that this is done via temporally binding the responses of neurons representing features that belong together. To illustrate with a recklessly simplified example, neural pools representing RED and MOVING LEFT could not only be firing, but firing in synchrony, say at t = 5ms, 15ms, 25ms, etc., while the neural pools responsible for processing BLUE and STATIONARY would fire are 10ms, 20ms, 30ms, etc., while pools that process GREEN, YELLOW, MOVING RIGHT, etc., would be quiescent. Each phase would carry the information about a single entity, in that all features associated with that object would be represented by neural groups firing at that phase. This is one manifestation of c-coordination, in that elements in the color manifold are being identified with elements in the spatial/movement manifold. But obviously the same mechanisms could handle c-coordinations of other sorts as well. For instance the neural pools representing the building in your behavioral space and those representing the building in your imaginatively constructed map could fire in synchrony when the cognitive system represents them as *the same building*.

5. Conclusion

Any honest assessment of the degree to which I have clarified all the 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 distinction between objective *things* and objective *facts.* I have also said very little about the different sorts of imagination required for the allocentric viewpoint, such as the ability to imaginatively assume the position of a being which can intuit time as a spatial line. I have said nothing about the conceived appearance/reality distinction which seems to be an integral part of our own conceptual scheme. I have said nothing about the role that our notions of causality, or our naive theories of perception play in out self conception.[26] A full catalogue of the deficiencies of the theory I have developed would tax 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 objective order. This ability is arguably the starting point for many of the other issues just mentioned, and it also seemed to be a place where an empirical beachhead might be established. But regardless of the success of my own theoretical venture in the matter, I hope at a minimum to have helped bring to the fore the importance of a cluster of issues to which few engaged in the naturalization of the mind have paid sufficient attention.

Acknowledgments:

[26] John Campbell's (199@) *Past, space and self* is a penetrating analysis of many of these additional features that I have had to ignore, though Campbell unconcerned about neurobiological implementation.

I would like to thank the Center for Semiotic Research at the University of Aarhus, Denmark, and the Danish National Research Counsel for supporting this research, and Per Aage Brandt for his beneficence in securing this support.

References:

Akins, Kathleen (1996) Of sensory system and the 'aboutness' of mental states. *Journal of Philosophy* 93(7):337-72

Allison, Henry (1983) *Kant's transcendental Idealism: An interpretation and defense*. New Haven: Yale University Press.

Bach-y-Rita, P. (1972). *Brain Mechanisms in Sensory Substitution.* New York and London, Academic Press.

Chalmers, David (1996) *The Conscious Mind*. Oxford: Oxford University Press.

Churchland, Patricia, and Terrence Sejnowski (1992) *The Computational Brain*. Cambridge MA: MIT Press.

Churchland, Paul (1979) Scientific Realism and the Plasticity of Mind. Cambridge UK: Cambridge University Press.

Churchland, Paul (1989) *A neurocomputational perspective*. Cambridge MA: MIT Press.

Cussins, Adrian (1992) Content, embodiment and objectivity: the theory of cognitive trails. *Mind* 101(404):651-688.

Devitt, Michael (1974) Singular Terms. *Journal of Philosophy* 71:183-205.

Dretske, Fred (1981) *Knowledge and the flow of information.* Cambridge MA: MIT Press.

Evans, Gareth (1982) *The Varieties of Reference*. John McDowell, ed. New York: Oxford University Press.

Evans, Gareth (1985) Molyneux's Question. In Gareth Evans (1985) *The Collected Papers of Gareth Evans*. London: Oxford University Press.

Fodor, Jerry (1987) Psychosemantics. Cambridge MA: MIT Press

Frege, Gottlob (1953) On Sense and Reference. Reprinted in Moore, Ed. (1993) *Meaning and Reference*. Oxford: Oxford University Press.

Grush, Rick (1995) *Emulation and Cognition*. Doctoral Dissertation, University of California, San Diego. UMI.

Grush, Rick (1997) The Architecture of Representation. *Philosophical Psychology* 10(1):5-23.

Grush, Rick (in preparation) *The Neural construction of mind, language and reality. Volume I: Representation, perception and objectivity*.

Harmon, Gilbert (1977) How to use propositions. *American Philosophical Quarterly* 14: 173-176.

Held, R. and Hein, A. (1963) Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology* 56(5):872-876.

Hume, David (1888). *A Treatise of Human Nature*. L.A. Selby-Bigge. Oxford: Clarendon Press.

Kant, Immanuel (1929) *Critique of Pure Reason*. Trans. Norman Kemp Smith. New York: St. Martin's Press

Kaplan, David (1969) Quantifying In. In Donald Davidson and Jaakko Hintikka, eds. *Words and Objections*. Dordrecht: Reidel.

Kosslyn, Stephen (1994) *Image and Brain*. Cambridge MA: MIT Press.

Lloyd, Dan (1989) *Simple Minds*. Cambridge MA: MIT Press.

Mach, Ernst (1906) *Space and Geometry*. LaSalle: Open Court.

Mel, Bartlett (1986) "A connectionist model for learning 3-d mental rotation, zoom and pan" In *Proceedings of the 8th Annual Conference of the Cognitive Science Society*. Hillsdale NJ: Lawrence Earlbaum Associates.

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.

O'Keefe, John, and Lynn Nadel (1978) The Hippocampus as a cognitive map. Oxford: Clarendon.

Peacocke, Christopher (1992) *A Study of Concepts*. Cambridge MA: MIT Press.

Smith, Brian Cantwell (1995) *On the origin of objects*. Cambridge MA: MIT Press.

Strawson, Peter F. (1959) *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.