PHILOSOPHICAL PSYCHOLOGY, VOL. 12, NO. 1, 1999

Qualia, Space, and Control

PETE MANDIK

ABSTRACT *According to representionalists, qualia—the introspectible properties of sensory experience—are exhausted by the representational contents of experience. Representationalists typically advocate an informational psychosemantics whereby a brain state represents one of its causal antecedents in evolutionarily determined optimal circumstances. I argue that such a psychosemantics may not apply to certain aspects of our experience, namely, our experience of space in vision, hearing, and touch. I offer that these cases can be handled by supplementing informational psychosemantics with a procedural psychosemantics whereby a representation is about its effects instead of its causes. I discuss conceptual and empirical points that favor a procedural representationalism for our experience of space.*

Introduction

In this paper I critique recent representational theories of conscious experience, specifically, those offered by the philosophers Fred Dretske (1995), William Lycan (1996), and Michael Tye (1995). These authors argue that qualia—the introspectible properties of sensory experience—are exhausted by the representational contents of experience. I do not challenge this point here but instead take issue with what these authors say about the ways experiences get their representational contents [1]. That is, I take issue with these authors over the proper psychosemantics for the representational aspects of conscious experiences. These authors advocate an informational psychosemantics whereby a brain state represents one of its causal antecedents in evolutionarily determined optimal circumstances. I argue that such a psychosemantics may not apply to certain aspects of our experience, namely, certain aspects of our experience of space in vision, hearing, and touch. I argue that these cases can be handled by supplementing (but not replacing) informational psychosemantics with a procedural psychosemantics whereby a representation is about its effects instead of its causes. Motor commands are the prototypical instances of the kind of representation most amenable to a procedural psychosemantics. I argue for

Editor's note: This was the winner of the 1998 William James Prize of the Society for Philosophy and Psychology, awarded for the best paper by a graduate student presented at the Society's annual meeting. Beginning with this paper, winners of the James Prize will have the option of publishing a revised version of their paper in Philosophical Psychology.
Pete Mandik, Philosophy–Neuroscience–Psychology Program, Washington University in St. Louis, One Brookings Drive, St. Louis, MO 63130, USA; e-mail: pete@twinearth.wustl.edu. http://www.artsci.wustl.edu/~pjmandik/index.html

0951-5089/99/010047–14 $7.00

the plausibility of a procedural approach to spatial experience by examining the integral role that motor control plays in many of our experiences of space. I discuss conceptual and empirical points that favor a procedural representationalism for spatial qualia.

This paper is organized as follows. In the first section ("Qualia") I briefly discuss what I take qualia to be and the plausibility of representational accounts of qualia. In the second section ("Space") I argue for the existence of spatial qualia and the plausibility of procedural representationalism as applied to them. In the final section ("Control") I present conceptual and empirical considerations in favor of the case that some of our spatial qualia actually are the contents of procedural representations.

Qualia

Contemporary philosophers writing on consciousness are up to their necks in a morass of supposed synonyms. "Qualia", "consciousness," "experience," "subjectivity," and the horrendous "what-it-is-like" are tossed around in a word salad in which each term is treated as interchangeable with any other. My focus here is to get clear about qualia, and I start by insisting on the non-synonymy of the terms listed above [2]. Qualia are the introspectible properties of sensory experiences. They are not the experiences themselves, nor are they just any properties of the experiences, but only the introspectible ones. They are that in virtue of which we are able to discern by introspection similarities of and differences between our experiences. I may know by introspection that my visual experience of an orange is more like my visual experience of a grapefruit than my visual experience of a lime. I may also know by introspection that my visual experience of an orange is more like my visual experience of a lime than my experience of tasting a lime. My experiences also have properties that I may be unable to know by introspection, like whether my experience of pain is neurally instantiated by my c-fibers, or whether my visual experience of my coffee mug is the 51st visual experience of the mug that I have had today.

Representationalism is the view that qualia are representational contents. On this view, my visual experience of a red rose is a mental representation of a red rose—a representation of a thing as being red, as having a certain shape, a certain visual texture, etc. The introspectible properties of the experience are exhausted by the representational contents of the experience. Thus, the (introspectible) difference between my visual experience of a red rose and my visual experience of a white rose has to do with the fact that the former experience represents a red thing in my visual field and the latter experience represents a white thing in my visual field.

I think that representationalism goes quite a distance in solving certain puzzles about the relation of mind to brain. Leibniz, in his *Monodologie* (1714), imagined being miniaturized and strolling through a brain: Leibniz would not and could not, he assured us, find anything recognizable as an experience amidst the brain mechanisms that loomed around him (§17). While Leibniz took this consideration as an argument against the identification of experiences with brain process (or any other physical processes, for that matter), most of us know better than to commit Leibniz's fallacy. That experiences do not seem like brain processes, or vice versa, is insufficient to establish that experiences are not brain processes. However, despite his error, Leibniz put his finger on a powerful intuition—one that exerts widespread influence even today. The intuition is this: regardless of whether experiences *are* brain processes, they nonetheless do not *seem* like brain processes. Leibniz's problem is the problem of explaining this apparent difference.

Treating experiences as representations allows us to solve Leibniz's problem by appealing to a content/vehicle distinction. A miniaturized Leibniz seems unable to find experiences in the brain for the same reason that someone who cannot read English will not be able to extract the story of *Moby Dick* from even the most careful scrutiny of the paper and ink in a copy of Mellville's novel. Even the closest examination of a representation's vehicle need not yield any clues to the representation's content, for the properties had by the former need not be had by the latter, and vice versa. To think otherwise is to be surprised upon opening a copy of *The Red Badge of Courage* and discovering that it is not printed in red ink. Leibniz could not introspect that his experiences were realized by brain processes, nor could he, by examining brain processes, know what the introspectible properties were. The key insight of representationalism is that the properties of experiences that are and are not introspectible—the qualia and the non-qualia, as it were—are differences between representational contents and representational vehicles.

With the idea of a content/vehicle distinction in hand, we may go on to ask what, physicalistically speaking, are the contents and vehicles. The vehicles, most if not all agree, are brain state [3]. Which brain states they are is, of course, a tough question. However, an even tougher question is how to specify, in terms acceptable to physicalists, how brain states go about having representational contents. This is *the* problem of psychosemantics.

The favored psychosemantics of the representationalist is an informational psychosemantics, the key idea of which is that the representational content of brain states is something that *causes* brain states to be tokened in evolutionarily determined optimal conditions. The gross outlines of the story are as follows. When my nervous system is functioning as it should, my c-fibers (the neural correlates of pain experience) are caused to fire by damage to and disturbances of tissues in my body. That is, their proper function is to carry information about tissue damage/disturbance. Firing c-fibers *represent* tissue damage and disturbance and they may also *mis*represent tissue damage and disturbance. It is the proper function of my c-fibers to carry information about tissue damage and disturbance, and they have that function even in instances in which they are not carrying information about (that is, not caused by) tissue damage/disturbance. (This is analogous to the way that eyes have the function to see even in instances in which they no longer are capable of performing that function.) Typically, the teleological part of the story (the part about there being functions) gets analyzed etiologically (that is, historically) and the historical facts that matter are facts about evolution and natural selection. Thus, c-fibers are the result of a process of natural selection whereby creatures that did have c-fibers that fired in response to tissue damage/disturbance had a selective advantage over those creatures that did not. Activations of c-fibers represent those things that they were naturally selected to be caused by. I suppose the details of this kind of teleosemantics familiar enough not to need further unpacking here.

The part of the above story that I am most interested in at present is not the teleology part. Nor is it the etiology part. It is the information part—the part whereby representations are supposed to be about their *causes*. I want to contrast the informational story with what I shall call a procedural story: a story whereby representations can be about their effects instead of their causes. I will say more about the details of the procedural story in the next section. For now let me just flag the following. The informational story is plausible for *some* qualia—I just do not think that it will be the truth about *all* qualia. Some qualia, namely, some spatial qualia, will best be handled by a representational story that has procedural as well as informational components.

Before supplying more detail regarding procedural psychosemantics, I will say a few things about spatial qualia, like, that there are spatial qualia and which qualia are the spatial ones.

Space

In discussing early versions of the present paper, I have encountered people who thought it weird to say that there were *spatial* qualia. They granted that there were color qualia and odor qualia, but not *space qualia*. As I see it, if there are any qualia at all, then it is very likely that there are spatial qualia. If qualia are the introspectible properties of sensory experience and insofar as we can introspect that we experience objects as having spatial properties, then we have spatial qualia. There is no reason to allow that there are qualia in virtue of which I am seeing something as green while ruling out that there are qualia in virtue of which I am seeing something as square. The introspectible differences between seeing a square and seeing a triangle are spatial qualia differences. The introspectible similarities between seeing a square and feeling a square are spatial qualia similarities. I move on to catalog a few varieties of spatial qualia.

Our senses of taste and smell seem not to have much to do with the spatial properties and relations of physical objects. In marked contrast, our senses of hearing, touch, and—most of all—sight present us with richly structured spatial manifolds. We hear sounds as coming from various directions in relation to our bodies. And most notably and importantly for fans of stereophonic recordings of music, we can hear sounds as coming from either the right or left sides of the spaces our bodies occupy. Our senses of touch and sight allow us to discriminate the shapes of objects, as well as detect the various distances between each of the objects [4].

One particularly noteworthy feature of vision is that it functions as a "distance" sense: the properties sensed via that modality are presented as being located at various distances from the surfaces of our bodies, even though the sites of transduction are located in our retinas. In contrast to our retinas, our tongues present their stimuli (flavors, temperatures, and textures) as occurring right at the site of transduction. Interestingly, our sense of touch, while typically not a distance sense, can sometimes function that way. The most familiar example of this is the way that, when one end of a stick is held in the hand and the other end tapped on the ground, the felt taps are, as it were, "projected" to the end of the stick—we feel the taps at the far end of the stick, not in the hand grasping the other end. An example of tactile projection familiar to automobile drivers is the way that bumps and slick spots in a road are felt in the tires of the automobile and not in the parts of the driver's body that are the sites of contact with the rest of the car.

Perhaps the most intriguing example of such sensory projection is given in experiments Paul Bach-y-Rita (1972) performed in developing prosthetic vision devices for the blind. The devices typically consist in a camera worn on the subject's head which sends low resolution video signals to a 16-by-16 or 20-by-20 array of tactile stimulators worn on the subjects back. After only a few hours of training with the device, subjects could utilize the tactile stimulation fed to the surface of their skin by the camera to recognize distal faces and objects and read printed words that the camera was focused on.

Arguably, the trained subjects' experiences manifested many of the qualia of vision [5]. Just as a person's tactile point of view can extend to the tip of a walking cane, the point of view of the trained subjects shifted from the sites of stimulation on their back to the point of view of the camera. Bach-y-Rita reports an occasion in which the zoom control of a subject's camera was activated without warning by an experimenter, causing the image in the camera to loom suddenly. The subject reacted by raising his arms to his head and lurching backwards (Bach-y-Rita, 1972, pp. 98–99). Although the tactile stimulus array was located *low* on the subject's *back*, the subject's behavior indicates that he located the percept *high* and in *front* of him. And in looming, the objects seen through the camera were seen by the subject as rapidly approaching him, indicating that the device was functioning as a distance sense for the trained subject.

Thus completes my brief catalog of the relevant and interesting spatial qualia. I want to return now to the topic of representationalism.

Representationalism about conscious experience gets its strongest foothold for any case in which we can point to a distinction between vehicle and content—between the experience itself and that which the experience is an experience *of* [6]. This is why, I offer, orgasms and moods seem to some to be especially difficult for representationalism [7]. One may be skeptical, then, of representationalism's ability to handle all kinds of qualia. For my purposes here, all that matters is the plausibility of representationalism for some qualia, namely, spatial qualia.

The content/vehicle distinction gets a pretty good grip when we consider our sensory experiences of space. The experiences themselves, most will suppose, are (or at least token identical to) states of the brain. I can experience something as being far from me without the experience itself being far from me. I can experience something as being much larger than me without the experience itself being larger than me. The clearest way to explain the difference is by saying that experiences of spatial properties and relations represent those properties and relations and need not have those properties and relations in order to do so. The introspectible properties of spatial experiences—spatial qualia—are thus the contents of representations.

The psychosemantic question as applied to spatial representation is thus the question of how brain states come to represent spatial properties. How can a brain state that is not eight feet wide come to represent something as being eight feet wide? How can a brain state that is not a large triangular thing off to my right come to represent something as being a large triangular thing off to my right? I proceed now to describe procedural psychosemantics in general and then move on to describe procedural psychosemantics as applied to spatial representation.

An informational representation is about its causes and a procedural representation is about its effects. Prototypical instances of these two kinds of representation are neural representations in sensory cortex and motor cortex, respectively. Activations of motor representations in the brain eventuate in the sending of efferent signals to the effector organs. Sensory representations are activated by afferent signals from sensory organs. Informational representations are individuated by the flow of causation from the outside in, whereas procedural representations are individuated by the flow of causation from the inside out.

For a linguistic analogy to sensory and motor representations, consider the difference between observation statements and commands. The semantic content of an observation statement is its truth conditions and a reliable observer is one who is caused by the observed object to make true reports about the object. (Compare "The litmus paper is red" and the litmus paper's being red.) Instead of truth conditions, the semantic contents of commands are their success conditions. Successful commands serve to bring about the state of affairs that comprise their success conditions. (Compare "Clean your room!" and a clean room.) The aptness of this analogy is evident in the frequency with which efferent signals are equated with commands in the scientific and philosophical literature [8].

The proper psychosemantics for motor representations will be a procedural psychosemantics. The big question is, however, what could this have to do with spatial representation?

Consider how informational and procedural psychosemantics may each deal with the problem of spatial representation. I propose to pursue this by way of a thought experiment. The thought experiment is, without a doubt, both fanciful and highly artificial, as is typical of thought experiments. It is also quite illuminating, so I beg the readers' indulgence for the moment.

Imagine a creature living in an environment that consists of a smooth, flat, almost featureless, plane. The plane is topographically quite uninteresting, but it will allow our creature an extremely simple form of locomotion. The creature's means of locomotion consists in a pair of tank treads; thus, I will hereafter call the creature "Tanky." Imagine the plane that Tanky traverses as a grid of squares. The width of each square is equal to the distance Tanky travels is a straight line if his tread turns exactly one full rotation. Just like a tank, Tanky can go straight forward and backward by turning both of his treads in concert and turn to the left or right by turning his treads in opposition. Most of the square regions in Tanky's planar world are devoid of anything noteworthy. However, a few are good places for Tanky to be (we can imagine the good squares being covered with a nutritious chemical) and a few are bad places for Tanky to be (we can imagine the bad squares being covered with a toxic chemical). Tanky is equipped with a sensory receptor on his underside so that when he is in a square, he can detect whether it is nutritious, toxic, or neutral. A nutritious region may be eventually drained of its nutritious chemical by Tanky, and he then will leave to find a new one. The longer Tanky spends in a toxic region, the more deleterious it is for his health: thus, it will be in his best interest to leave as soon as possible, and avoid future encounters with toxic regions.

We now have the gist of Tanky's environment and his basic survival goals. Everything that we need to know about Tanky's outer anatomy is exhausted by mentioning his two motor organs and his one chemoreceptor sensory organ. We must turn now to Tanky's inner anatomy. We are especially interested in finding out how it is that Tanky knows where he has been and figures out where he is going. The way Tanky knows whether he is in a region that is good, bad, or neutral is in virtue of afferent signals from his chemoreceptor. If we give Tanky the representational resources to have sensory experience and memory, then the psychosemantics for his chemical representations will clearly be informational. Let us turn from his chemical representations to his spatial representations. Consider how Tanky might go about figuring out how far it is between, say, here and the next nutritious square (his next "square meal" one might say).

If Tanky has been traveling in a straight line, passed over a toxic square, and traveled through four neutral squares before encountering a nutritious square, how might he come to know this? Let us consider two possible design solutions to this problem, one informational, and the other procedural [9].

The informational solution is as follows. First, in addition to Tanky's chemical receptor, he has sensory organs in his tank treads that send afferent signals to Tanky's brain indicating whether the tread is turning forward or back, and how fast it is turning. Second, note that Tanky's motor commands to his effector organs are not accompanied by efference copies telling him what signal he sent. (An efference copy is a copy of the signal sent to the effector organs, but serves not as a trigger of the effectors but instead as a record of what commands were issued and when.)

On the procedural solution, in contrast, Tanky's motor commands *are* accompanied by efference copies. Further, unlike the informational solution, on the procedural solution, Tanky would *not* have sensors in his tread indicating rate and direction of tread turning.

On both of the solutions we may assume that Tanky does not slip and slide in his travels, and no one ever picks him up and carries him to a different location. All of his travels across the plane are due to self-initiated turns of treads with good traction. On *both* solutions, then, Tanky has all of the efferent and afferent connections to his world to give his memory and computational modules what they would need to compute his previous, current, and possible future locations. On both solutions, Tanky has the representational resources for representing a nutritious and a toxic region as being, say, four tank tread distances away from each other.

The difference between the two solutions hinges on the different semantics of Tanky's spatial representations. On the informational solution, Tanky's spatial representations are triggered by afferent signals from his tread sensors. On the procedural solution, Tanky lacks tread sensors and his spatial representations are instead triggered by efference copies. On the informational solution, the spatial representations represent tank tread speed and direction in virtue of being *caused by* activity in the treads. On the procedural solution, the spatial representations represent tank tread speed and direction in virtue of *causing* activity in the treads [10].

The Tanky thought experiment shows that under certain conditions, procedural as well as informational representations can serve the needs of spatial representation. Informational representation does not *have* to be the only game in town. The thought experiment makes salient the possibility of the procedural solution to the problem of spatial representation. The thought experiment reveals the happenings in a possible world. The next question is: how is it done in the actual world? What solution did Mother Nature opt for us and our evolutionary relatives? Below I argue that the procedural solution is realized in actual cases. The points I raise below are conceptual and empirical considerations showing that motor differences—that is, efference differences—can give rise to differences in spatial qualia.

Control

The first consideration in favor of the behavioral constituency of perceptual space is a conceptual consideration due to George Pitcher (1971). Pitcher asks whether it is conceivable that the perceived directions from which sounds originate are not intrinsically tied to our dispositions to behave in ways appropriate to the perceived directions of the sound. As Pitcher points out, if it is so conceivable,

> … then it ought to be logically possible for someone to hear the direction from which a certain single sound (a bird-call, for example) is coming, and yet for him not to know in what direction he must point (or walk) if he is to point (or walk) in the direction from which the sound is coming, not to know in what direction must look if he is to look in that direction … and so on for all related abilities. (p. 189)

But such a possibility seems not to make sense. We just cannot imagine hearing a sound coming from a particular direction while simultaneously not knowing how to orient (or point, or walk …) toward the direction from which the sound seems to originate.

Similar considerations can be brought to bear on the visual apprehension of distributions of features in our visual fields. Imagine seeing a red stripe across your visual field such that its endpoints lie in the far left and right of your visual periphery. It makes little sense to suppose that one could perceive the endpoints of the stripe as occupying these regions of egocentric space without knowing how to orient to bring one or the other into one's fovea.

The case for the behavioral constituency of tactual–kinesthetic perception is even easier to make than for audition and vision. As Gareth Evans (1985) points out, it is quite obvious that when a blind man uses his hand to perceive the spatial relations that parts of a chair bear to each other, the content of his experience is "partly determined by the disposition he has thereby exercised—for example, that *if* he moves his hand forward such-and-such a distance and to the right he will encounter the top part of a chair" (p. 389).

Our sensory awareness of spatial locations is not simply a matter of an egocentric coordinate system based on the arrangement of our own body parts; instead it is a function of what we are disposed to do with those parts. As brought out nicely by Charles Taylor, up is not simply where our head is (as is evident when lying down), but instead a function of "how one would move and act in the field" (1979, p. 154).

Motoric output plays a constitutive role in the projective phenomena I had mentioned in connection with Bach-y-Rita's subjects. A necessary condition on being able to "see through" the camera-driven tactile array is that the subject be allowed to exert control over the inputs to the camera (Bach-y-Rita, 1972). Consider the following points in favor of this thesis.

Bach-y-Rita notes that the major portion of the first few hours of the subject's training is occupied by learning the techniques of camera manipulation, including controlling the operation of the zoom lens, the aperture and focus, and directing the camera towards regions of the subject's immediate environment (1972, pp. 3–4). Bach-y-Rita further notes that subjects with a high degree of manual dexterity acquire the ability to see through the prosthetic vision devices more quickly than those with a low degree of manual dexterity (p. 7).

In discussing the case of the subject who raised his hands to protect his head in response to an unanticipated activation of the camera's zoom control, Bach-y-Rita writes:

> The startle response described above was obtained by a well-trained subject who was accustomed to have camera movement, zoom, and aperture under his control. (1972, p. 99)

Further, Bach-y-Rita writes, when subjects receive tactile array inputs from a *static* camera, they

> … report experiences in terms of feelings on the skin, but when they move the camera their reports are in terms of externally localized objects. (1972, p. 99)

The subject's being able to control the camera seems to be a necessary condition for the tactile stimuli on the back to become transparent and to enable the device to function as a prosthetic distance modality. Indeed, this is the hypothesis advocated by Bach-y-Rita:

> [E]xternal localization of percepts depends critically on such movements and … a plausible hypothesis is that a translation of the input that is precisely correlated with self-generated movement is the necessary and sufficient condition for the experienced phenomena to be attributed to a stable outside world. Conversely, in the absence of such a correspondence, the origin is perceived as being within the observer. (1972, pp. 99–101)

I must note, however, that I differ from Bach-y-Rita in that I deny that exercising motor control over the inputs is *sufficient* for the projection phenomena discussed. I wish only to claim that the motor control is necessary. Consider, in this regard, a datum presented by Bach-y-Rita: when a subject trained to see through the tactile array of the prosthetic vision device later scratches his back, he does not "see" the scratches (1972, p. 32). But, nonetheless, in scratching his back, he is in control of the tactile input. Thus, while necessary for instantiating the sensory projection essential to functioning as a distance sense, exerting motor control over one's sensory inputs seems not to be sufficient.

Plausibly, the necessity of motor control for sensory projection can be extended beyond the prosthetic vision case to the other instances of tactile projection I discussed above. Consider, for example, the case of the stick held in the hand. When you hold the stick in the hand and tap and rub the end over textured surfaces, you feel the sensations at the end of the stick. Imagine, however, that the same information was delivered through the stick to the hand, but without your being able to control the rubbing and tapping of surfaces against the end of the stick. Imagine that you are blindfolded and strapped into a chair with your hands and arms restrained. A stick is placed in your hand and an experimenter taps and rubs the end of the stick with, say, a rough rock. It is not implausible to suppose that your tactile sensations would *not* project to the end of the stick, but instead would be perceived as occurring entirely within your hand.

If the considerations examined above are correct, then it follows that the qualities of phenomenal consciousness associated with stimuli in various modalities are determined not only by the nature of the information transduced by the nerve-endings in the sensory organs, but also by what type of subsequent motor activity that information is employed in. In the case of Bach-y-Rita's prosthetic vision experiments, the difference between perceiving skin stimulation as disturbances on the skin and seeing through the disturbances to a display of objects located in a three-dimensional egocentric space entailed differences in the kind and degree of motor control that the subjects were able to exert on the dynamics of the sensory input.

For another line of evidence linking efference and spatial qualia, consider what happens when one attempts to move one's eye when the eye muscles have been paralyzed. Gallistel (1980, p. 175) describes the following interesting phenomenon. A person may have their eye muscles paralyzed by curare. If a person with curare-induced paralysis in her eyes attempts to shift her gaze to the left, the whole world will appear to her to have jumped to the left. This change in the visual scene is not due to any actual movement of the eye, nor has there been any change in the pattern of light on the retina. Gallistel (1980, p. 175) writes that what has happened is that the typical expected association of efference copy and afference has been interrupted.

I interpret the case of the paralyzed eye as follows. Usually when one moves one's eyes to the left, the visual image moves to the right with respect to the eye. This typical association between the efferent output and afferent input is used to construct a representation of the world as a stable scene—even though the visual image is moving to the right with respect to the eye, the world is not seen as leaping to the right. But when the efferent command to shift is not accompanied by the afferent signal indicating the change of light on the retina (flowing to the right), the whole world is perceived not as stable, but instead as leaping in the direction indicated by the efferent command to shift to the left. Again, efference has had an effect on the introspectible properties of spatial experience.

We can easily imagine an analog to the paralyzed eye phenomenon in Tanky. In the procedural version of Tanky, he may construct a representation of his world based on expectations of which afferent signals would be paired with which efference copies. We could paralyze Tanky's treads so that a motor command to go forward half a square is not accompanied by a change in the information presented to his chemoreceptor. A command to move forward, when not accompanied by the backward flow of the pattern of chemicals across the chemoreceptor surface, would result in the perception of the world lurching forward.

A last empirical consideration I offer in favor of the proceduralist hypothesis is Held and Hein's (1963) famous kitten experiment. In this experiment a pair of kittens were placed in an environment controlled to make the visual input to each kitten as similar as possible, while allowing only one kitten to have control of the pattern of information coming into its eyes. The environment was an upright cylindrical volume with vertical stripes on the walls. In the center of the cylinder was a column. On top of the column was a horizontal bar that was attached at its center to a spindle. The horizontal bar could be turned around. Each kitten was attached to opposite ends of the bar. One kitten, the active kitten, was attached to the bar by means of a harness so that as the kitten walked around the cylindrical room, the bar would turn on the spindle. The other kitten, the passive kitten, was restrained in a gondola that hung from the other end of the bar. As the active kitten made the apparatus turn, the passive kitten was moved through the room at the same speed as the active kitten. When the kittens were not in the cylindrical room attached to the apparatus, they were both free to move about but only in the dark. After spending some time being reared in these conditions, the kittens were subjected to tests of their vision. The results were that only the active kitten developed normal visual perception. For example, only the active kitten avoided a visual cliff (that is, it avoided walking on a piece of glass suspended a considerable height over a surface below it)—a behavior typical of kittens its age.

We can easily imagine constructing Tanky-style systems that would replicate results analogous to those in the kitten experiment. Imagine two procedural-style Tanky creatures attached to the spindle, with one being allowed to actively move while the other was dragged along (with efferent signals to its treads blocked). Visual stripes on a wall would be replaced with chemical squares on the floor. Only the active tank would be able to develop certain efferent/afferent associations, associations like those discussed above in connection with the case of the paralyzed eye. The results of the kitten experiment show yet another connection between efference and perception, and the Tanky analog to the kitten experiment helps show the point of entrance for a procedural psychosemantic interpretation of the kitten experiments. The passive tank would be subject to certain chemoreceptive inputs but be unable to correlate those with the effects that motor activity has on the pattern of inputs. Thus, the passive tank would be unable to represent the spatial locations of the sources of the chemical stimuli.

One may question the relevance of the kitten experiment to a paper on the introspectible properties of sensory experience. It is not clear how we could have any idea as to what, if anything, the kittens introspect. At least in the discussion of the paralyzed eye and the prosthetic vision devices, human subjects were involved, and we had their introspective reports as evidence of efferent-sensitive changes in spatial qualia. The kitten experiment is relevant, then, insofar as it may be considered as evidence that similar results would be obtained if humans instead of kittens were subjected to similar conditions. Such an experiment conducted on human children would obviously be unethical. So we can never be *certain* what would result from such an experiment. However, the kitten data, considered in light of the paralysis and prosthetic vision data, lend further support to the proceduralist hypothesis about the nature of the representations in the experience of spatial properties.

Conclusion

In this paper I have argued that if there are any qualia, then there very likely are spatial qualia. I have argued that if one is a representationalist about any qualia, then one should be a representationalist about spatial qualia. After arguing these points I offered a procedural psychosemantics as an alternative to the informational psychosemantics typically advocated by representationalists. I introduced the thought experiment of Tanky to urge the intuitive plausibility that a procedural psychosemantics could have anything to do with the way an organism represents spatial properties and relations. In the last section of the paper, I discussed conceptual and empirical considerations that procedural representations actually are involved in some of our experiences of space. Note that I have not argued that all sensory representations must have procedural semantics. Nor do I think that all spatial representations necessarily have procedural semantics. I have argued that the informational story cannot be the *whole* story, and I have sketched a positive alternative—the procedural alternative—to handle those cases in which the informational story does not apply.

Acknowledgments

I owe special thanks to Rick Grush and Tadeusz Zawidzki for stimulating discussions on the material in this paper. For helpful and detailed comments on previous drafts I thank Alan Belasco, Chase Wrenn, the members of Washington University's Philosophy–Neuroscience–Psychology Research Group and the rest of the audience of an ancestor of this paper that was presented at a PNP Works-In-Progress Session in the spring of 1997, and audiences of presentations of this material at the 1998 meetings of the Society of Philosophy and Psychology and the Southern Society of Philosophy and Psychology.

Notes

[1] In Mandik (1997) I raise doubts about the plausibility of representationalism as applied to *all* qualia.

[2] "What-it-is-like" will have no further use or mention here. As for "subjectivity" I discuss that at length elsewhere (Mandik, 1998, in preparation). For what it's worth, I take subjectivity to be a property of certain representations, but not all subjective representations are experiences. My key concern in this section is to observe a many-way distinction between experience, the properties of experience, the introspectible properties of experience (qualia), and the introspected properties of experience. Due to limitations of space, I can do little here beyond simply asserting the distinctness of these notions. I refer the reader to Lycan (1996) for an excellent treatment of these separate notions.

[3] One may prefer to say not that experiences are brain states, but instead that *our* experiences are *token identical* to brain states. One may want to say this because one may think that computers and other things without brains have experiences. For the purposes of this paper, that would be fine with me.

[4] For further discussion of spatial and non-spatial experience, see Mandik (1999).

[5] Although it is highly implausible that they replicated *all* of the qualia of vision, since color qualia arguably are missing in prosthetic vision. Whether prosthetic vision has enough of the qualia of natural vision to really merit being called "vision" is not an issue I am interested in.

[6] For further discussion of what the distinction between content and vehicle is and is not, see Grush and Mandik (in preparation).

[7] See Tye (1995), Lycan (1996) and Mandik (1997) for discussions of these kinds of difficult cases.

[8] See e.g. Grush (submitted) and Gallistel (1980).

[9] It is not my intention to claim that these two are the only possibilities. My point here is that what I am calling the informational solution is not the only one and that the procedural solution is a viable alternative.

[10] I can, of course, retell the story to incorporate all of the teleology and etiology one may desire, but I will not do that here.

References

BACH-Y-RITA, P. (1972). *Brain mechanisms in sensory substitution*. New York and London: Academic Press.

DRETSKE, F. (1995). *Naturalizing the mind*. Cambridge, MA: MIT Press.

EVANS, G. (1985). Molyneux's question, In Antonia Phillips (Ed.) *The collected papers of Gareth Evans* (pp. 364–399). Oxford: Clarendon.

GALLISTEL, C. (1980). *The organization of action: a new synthesis*. Hillsdale, NJ: Lawrence Eribaum Associates.

GRUSH, R. (submitted). Manifolds, coordinations, imagination, objectivity.

GRUSH, R. & MANDIK, P. (in preparation). Anti-anti-representationalism.

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

LEIBNIZ, G. (1714). *Monadologie*. Hamburg: Meiner.

LYCAN, W. (1996). *Consciousness and experience*. Cambridge, MA: MIT Press.

MANDIK, P. (1997). Review of *Ten problems of consciousness: a representational theory of the phenomenal mind*, Michael Tye. *Philosophical Psychology*, *10*, 127–129.

MANDIK, P. (1998). Objectivity without space. In R. GRUSH (guest editor) *The Philosophy of Gareth Evans: Special Issue of the Electronic Journal of Analytic Philosophy*. At http/www.phil.indiana.edu/ejap/.

MANDIK, P. (in preparation) The philosophy and neuroscience of subjectivity. In W. BECHTEL, P. MANDIK, J. MUNDALE & R. STUFFLEBEAM (Eds) *Philosophy and the neurosciences: a reader*. Oxford: Basil Blackwell.

Pitcher, G. (1971). *A theory of perception.* Princeton, NJ: Princeton University Press.
Taylor, C. (1979). The validity of transcendental arguments. *Proceedings of the Aristotelian Society*, *79*, 151–165.
Tye, M. (1995). *Ten problems of consciousness: a representational theory of the phenomenal mind*. Cambridge, MA: MIT Press.