Language Police Running Amok[1]

Justin M. Sytsma

In this article I critique Kathleen Slaney and Michael Maraun's (2005) addition to the ongoing philosophical charge that neuroscientific writing often transgresses the bounds of sense. While they sometimes suggest a minimal, cautious thesis–that certain usage can generate confusion and in some cases has–they also bandy about charges of meaninglessness, conceptual confusion, and nonsense freely. These charges rest on the premise that terms have specific correct usages that correspond with Slaney and Maraun's sense of everyday linguistic practice. I challenge this premise. I argue that they have not shown that there are such specific correct usages; and, further, that even if they had, they fail to justify that their definitions are the correct ones.

Kathleen Slaney and Michael Maraun (2005) make a recent addition to the ongoing philosophical attack on certain uses of language in the neurosciences. In *Analogy and Metaphor Running Amok*, they argue that neuroscientists commonly misapply terms that are specific to whole persons to their constituent parts, i.e. their brains or neurons. They hold "that the use of these and other similar literary devices in neuroscientific research sometimes leads to certain conceptual confusions and, thus, fails to aid in clarifying the nature of those phenomena they are intended to explain" (154). Slaney and Maraun's critique goes beyond this, however. They repeatedly argue that various uses of language in neuroscience are incorrect, fallacious, or nonsensical. Their arguments for such further claims are flimsy.

---

[1] I wish to thank Peter Machamer for many insightful conversations and comments on a previous draft of this paper. Furthermore, this article draws on a related critique of Bennett and Hacker's 2003 volume that was co-authored with him (Machamer and Sytsma, 2005).

To appear in *Journal of Theoretical and Philosophical Psychology*, 2007 (27:1)

Justin M. Sytsma
Department of History and Philosophy of Science
University of Pittsburgh
1017 Cathedral of Learning
Pittsburgh, PA USA 15260
Phone: (412) 624-5896
Email: jms124@pitt.edu

In this paper I will not go into the details of each of Slaney and Maraun's examples. Rather, my goal will be to draw out the general structure of their argument, articulate why it fails, and highlight this through a choice example–talk of "maps in/on the brain." The argument, most bluntly, is: (p1) a given term has a specific correct usage Slaney and Maraun know; (p2) this "correct" usage is transgressed in some neuroscientific work; as such, (c) the usage is incorrect and the claims made with these terms do not make sense. In Section 1, I target the first premise. I question both Slaney and Maraun's reasons for thinking that there are such singular, specific definitions for the correct use of terms and their reasons for thinking that they know them. I then consider the general form of their examples, or the so-called "mereological fallacy" in Section 2. Finally, in Section 3, I examine one of their examples in greater depth.

1. The Argument

It should, of course, be granted that neuroscientific usage *can* "result in misleading conceptual confusions… which impede rather than promote our understanding of the workings of the brain and mind" (154). The question is how endemic such confusions are and how we should understand them. Slaney and Maraun target particular usages held to be "common" (154) and argue that the confusion results from misuse of language by neuroscientists. They argue that the problem they are concerned with "stems mainly from a misunderstanding of particular concepts and of when it is appropriate to use one as a 'stand in' for another without serious conceptual distortions and/or confusions resulting" (154).

This "misunderstanding" suggests that there is a corresponding correct understanding. The two premises above are implied: Some terms have a specific correct usage that neuroscientists commonly misunderstand and transgress. The result is confusion. Elsewhere in the paper this is stated more bluntly. Paraphrasing Bennett and Hacker (2003), who give an extended attack on various uses of language in the neurosciences, Slaney and Maraun state that "for a hypothesis to be testable it must make sense, and in order for a hypothesis to make sense, its constituent terms must make sense and be used correctly" (169). Nonsense results from using terms in an incorrect way. Bennett and Hacker are even more blunt. They write that the "application of psychological predicates to the brain makes no sense" (72). Two questions should be raised: Are the bounds of sense really so clear? And, if so, do Slaney and Mauran correctly identify the boundaries? I charge that they fail to make the case for an affirmative answer to either question.

Very often language can be confusing (to someone or some group) without thereby being meaningless or nonsensical. If I point at a china bowl and ask you to "hand me the dog dish," you might be confused. You might doubt that I *really* mean the china bowl, it seeming to you to be an odd choice for a dog dish. Nonetheless, your confusion in fact underlines that the statement is meaningful. It would be *unreasonable* for you to insist that "dog dish" simply *cannot* be used to refer to a china bowl. It is also possible for usage to be simply incorrect (by some technical standard) and yet meaningful: People often refer to chimpanzees as monkeys (they are not) or tomatoes as vegetables (in biological terms, they are fruits). While such usage might annoy me, I have no difficulty understanding what is meant.

More drastically, think about the children's cartoon *The Smurfs*. Imagine that Papa Smurf points to the china bowl and says "smurf me the smurf." Not only is the statement not nonsense, but anybody familiar with the cartoon would understand the intent. They would not be surprised when Hefty hands Papa Smurf the china bowl. Nonetheless, 'smurf' has a primary meaning that corresponds to the fictional little blue creatures that Hefty is an example of. This does not mean, however, that one can stop at this definition: One cannot claim that Papa Smurf spoke nonsense because ‘smurf’ means “fictional little blue creature” and not “china bowl.”

Even though we might feel that the Smurf’s language has run amok, nonsense, confusion, or meaninglessness does not result. In the same way, we should strongly doubt that the bounds of sense/nonsense are so black and white as to render common neuroscientific utterances meaningless. At the same time, the extended use of ‘smurf’ is not especially useful for most discourse (Papa Smurf could have simply said “hand me the china bowl”). In the next two sections I argue that this is not the case for the terms that Slaney and Maraun target: Not only is their inclusion in neuroscientific discourse meaningful, it is useful.

Furthermore, the ways in which Slaney and Maraun draw the supposed boundaries mask the ways in which the terms can be useful. The disservice that their definitions do to neuroscientific communication raises the issue of why we should accept Slaney and Maraun’s sense of the terms as normative. If the terms are meaningfully used in neuroscientific communications in ways that diverge from the definitions Slaney and Maraun rely on, then we either need to abandon the notion that there is some singularly “correct” definition, or we need convincing criteria for deciding which usage is legitimate and which has run amok. Such criteria are not in evidence in Slaney and Maraun’s article.

2. The Mereological Fallacy

Slaney and Maraun’s examples largely center on cases where “predicates usually reserved for humans are attributed to the brain” (169) or its parts. For example, talk of “maps in/on the brain” (156) runs afoul of the criterion that maps “are used by *people* for one purpose or another” (157). Likewise they hold that “humans, not brains, can be said to be capable of representing, and representations are created by language users" (159). In a similar vein, they question talk that "the central nervous system has its very own 'language' with which different structures within it may communicate" (163); that is, they complain that "the brain is ascribed with having language-using abilities" (169) when "humans are the only creatures who can intelligibly be said to use or to have them" (165).

These cases show a characteristic structure: Slaney and Maraun argue that terms that correctly describe whole humans are being misapplied to their parts. As such, we can take these examples as running afoul of the so-called mereological fallacy coined by Bennett and Hacker (2003).[2] The fallacy is to ascribe "to the constituent parts of an animal attributes that logically apply only to the whole animal" (73).[3] The articulation of this fallacy crucially involves the assumption that the terms for such attributes–primarily psychological predicates–only logically apply to whole animals. This is an instantiation of the first premise articulated above.

The use of "logically" here relates to Bennett and Hacker's goal of applying "the methods of connective analysis to the conceptual problems that characterize neuroscience at the point where it abuts on psychology" (378). Their guiding principle is that "the meanings of words are determined by their rule-governed use" (383): connective analysis draws out the connections between these rules. The mereological fallacy amounts to charging that the application of psychological predicates to people-parts breaks the rules. The result of this crime is nonsense.

Despite these charges, psychological predicates are often applied to entities besides whole humans (or whole animals) and it is ludicrous to think that such utterances are nonsense. Consider Bennett and Hacker's example of the psychological predicate "sees" (73): it is simply not the case that we only apply the term to whole animals. We can talk of a security camera

[2] See Machamer and Sytsma (2005) for an extended look at this work and critique of the mereological fallacy.
[3] Such a principle would seem to reject any use of metonymy; for example, it seems to make the sense of my mother's claim that I used to "give her lip" rather disturbing.

seeing a crook, or an electronic door *seeing* the customer. We might specify further–it is actually a part of the door, the sensor or electronic eye, which sees. Each of these phrases is not only meaningful but readily understood; their everyday "logic" does not appear to be solely the logic that Bennett and Hacker find. As such, the mereological fallacy stands in need of strong justification.

The justification given rests on an appeal to authority. Slaney and Maraun (again following Bennett and Hacker) make much of a snippet from §281 of Wittgenstein's *Philosophical Investigations* (1958). They write that:

> Wittgenstein said: "Only of a human being and what resembles (behaves like) a human being can one say: it has sensations; it sees; is blind; hears; is deaf; is conscious or unconscious" (cited in Kenny, 1984, 125). Yet in neuroscientific research there are literally thousands of examples in which predicates such as "organize", "sort", "need", "dominate", and "control" are applied to the brain. (160)

The suggestion is that such usage is nonsense. Because of this Slaney and Maraun hold that "there is no obvious sense in which applying psychological predicates to the brain helps us better understand the mind" (160). Before turning to the brain/mind, note that each of these terms can be meaningfully and usefully applied to non-humans. We can meaningfully say that a crystal lattice *organizes*; that the machine at the bank *sorts* coins; that my car *needs* gas; that my heart *controls* the flow of blood through my body; and so on. Each of these sayings can be informatively descriptive and useful. In most cases we can readily rephrase each of the terms that Slaney and Maraun find suspect–we could say that having fuel is a necessary condition for my car to run and that it currently is low on fuel, for example–but there is no need to and brevity is often a virtue.

Our language is such that neuroscientists could likely rephrase the sentences to remove the psychological predicates whose use Slaney and Maraun challenge (although the onus is on them to show this). In a sense, then, perhaps the use of these terms does not help us to "better understand the mind" any more than talk of my car *needing* gas helps me to better understand that I need to stop at the filling station (since I could express the thought without the offending term). It is not the case, however, that "there is no obvious sense" in which the usage might improve understanding: "the car needs gas" is shorter, more concise, and because of that it is less likely to be misunderstood than the alternative statement. There is no *a priori* reason to expect anything different of neuroscientific utterances.

Slaney and Maraun appeal to, but do not explicate or defend, Wittgenstein. Moreover, they get him wrong. Wittgenstein did not claim that these terms only apply to humans, but says that they also apply to things that *behave like* them. The "behaves like," I argue, should be understood with regards to the term at issue: With regards to "sorting," for example, we can say that the machine at the bank behaves like a human (i.e. each sorts). If the terms at issue pick out behaviors, or doings, then it would seem reasonable to apply the term to any system that acts in those ways (whether a whole human or not). I see no reason to read §281 as denying this.

This is clear if the snippet that Slaney and Maraun quote is considered in context. Despite the strength of the language in §281, Wittgenstein did not hold that *secondary* applications of psychological predicates to things that do not behave like human beings are nonsensical. Rather, his point was that our primary use of such terms relates to human behavior. As such, in §282, he offered a contrast class for §281 that is rather different from the part/whole distinction: "We do indeed say of an inanimate thing that it is in pain: when playing with dolls for example. But this use of the concept of pain is a secondary one. Imagine a case in which people ascribed pain *only* to inanimate things; pitied *only* dolls!" He is not claiming that such ascriptions to dolls are *nonsensical*. While the secondary extension presupposes the primary usage, it does not follow that the extension is meaningless.

In these two sections, Wittgenstein emphasized action and drew a contrast with inanimate things that do not act. Brains, however, are not "inanimate," not clearly without behavior. Brains receive inputs that lead to series of neuronal firings, which activate further outputs. Such neuronal activity may warrant treating such firings as behavior (as what the neurons *do*). Thus, one way to make sense of attribution of psychological predicates to brains would be to consider their application to brain behavior (or what brains *do*–the regular changes they undergo and the direct effects they bring about). If psychological predicates may be used to describe brain behavior, then such uses are not nonsensical; further, they may be informative in a way that the attribution of such predicates to dolls, for example, are not.

The psychological predicates that Wittgenstein mentions in §281 can be both meaningfully and usefully applied to brains and their parts. Consider "seeing" again. A person can evidence seeing by responding appropriately to visual cues. This notion is readily extended. Many systems respond in systematic, functionally relevant ways to visual stimuli. One way to capture this is through an extended use of the psychological predicate. If we are concerned with edge-detecting cells in the visual cortex, for example, we might say that a cell *saw* an edge in the visual stimuli presented when it fires; likewise, we might say that it did not *see* the edge if it did not fire. The cell's behavior, what the cell does, is enough to make the secondary application of the psychological predicate understandable. As Bennett and Hacker put it: "we say of an animal or a human being that they perceive something in their field of perception if, for example, they respond to what is visible… in appropriate ways" (81). The application of "sees" to the edge-detector cell follows this directly (assuming a functional understanding of "appropriate"), simply extending the usage to parts of animals in a straightforward way. Again, I see nothing in §281 that would deny this.

In fact, the selective use of §281 to support premise one, above, makes a rather un-Wittgenstinian point. A general point made repeatedly in the *Investigations* is that a term's use is "not everywhere circumscribed by rules" (§69); or, as he put it in *The Blue Book*: "We don't use language according to strict rules–it hasn't been taught to us by means of strict rules, either" (1965, 25). Bluntly, sense and nonsense are not so cut and dry in natural language as premise one supposes. Slaney and Maraun's approach, in Wittgenstein's words, is "to think that if anyone utters a sentence and means or understands it he is operating a calculus according to definite rules" (§81). Their argument supposes this and proceeds to argue that the calculus is violated resulting in nonsense or lack of meaning. Slaney and Maraun do not offer support for this view of language and it should not be assumed. Denying such a calculus, as the later Wittgenstein did, the presumption should be that there are ways to make sense of neuroscientific utterances, not to invent fixed definitions and call illegitimate anything that falls outside their scope.

My goal, here, is not to defend Wittgenstein's view of language; rather, it is to note that Wittgenstein does not lend authoritative support to Slaney and Maraun's claims. In fact, his views run contrary to premise one, or the claim that a given term has a specific correct usage. Nonetheless, we could disagree with Wittgenstein and still reject premise one: Supposing that there are definite rules, there is no reason to think that Slaney and Maraun get them right. In part, there is no reason to think this because Slaney and Maraun offer none: If they want to prescribe how our words are to be used, then the burden of proof is on them to first show what justification would look like here and then to provide that justification. In fact, both everyday and neuroscientific usage tells against them, as we saw for the example of "sees." I now turn to one of Slaney and Maraun's specific examples.

3. Example: Maps in/on the Brain

It is highly dubious that the general argument that Slaney and Maraun appeal to can be supported. Nonetheless, it is possible for linguistic usage to so transgress the bounds of typical usage that they become essentially meaningless (or useless for communication) in a general context. Without further explanation, a statement like "my coffee tastes isomorphic," for example, just doesn't have any obvious meaning (although this does not imply that it cannot be given some). Assessing whether this is at all common in neuroscience requires (charitably) considering specific neuroscientific usage. Slaney and Maraun begin with claims that "there are maps in/on the brain" (156). They then move to terms like organize, sort, etc., whose use they find to imply that "the brain has a will of its own" (160). Third, they look at terms that they feel suggest the confused claim that "the brain speaks to itself in its own language" (162). Finally, they consider talk that "the brain has, stores, and contains information" (165). I will briefly discuss the first of these interrelated targets.

The talk at issue for Slaney and Maraun involves expressions like neural maps (Whitaker, 1971), cortical maps (Calving and Ojemann, 1994), topological code (Young, 1987), retinotopic organization (Rosenzweig et al., 1996), and so on. Talk of maps and brains takes a number of forms. Brains have been mapped and maps made, atlases produced, that aid researchers in finding their way around actual brains. These are *our* maps (researcher's maps); what they are maps *of* are brains. The maps mark out different brain areas, and while one might question what the spatial demarcations demarcate and how they are arrived at, one can also question the demarcations for, say, a map of Europe (reflecting boundary disputes).

There is also talk of maps *in* the brain. Consider the following snippet from an introductory neuroscience text (Purves et al., 1997) discussing the somatic sensory cortex:

“Mapping studies in humans and other primates show further that each of these four cortical areas [Brodmann’s areas 3a, 3b, 1, and 2] contains a separate and complete representation of the body. In these somatotopic maps, the foot, leg, trunk, forelimbs, and face are represented in a medial to lateral arrangement” (159). What is meant by this? In particular, in what sense, if any, are we dealing with a map? Mapping studies are conducted that record “the approximate region of the human cortex from which electrical activity is recorded following mechanosensory stimulation of different parts of the body” (159). That is, the basic “intro to neuroscience” understanding is that (1) we can map areas of the brain on the basis of the stimulations that generate neural activity (here, Brodmann’s area 3b–the primary somatosensory cortex–is mapped in terms of responses to stimulation of different body parts). (2) The brain area is then talked about as giving a somatotopic map. This, I believe, is the move that Slaney and Maraun object to and the usage I will focus on.

Slaney and Maraun’s discussion of maps focuses on the everyday sort, such as topographical and geographical representations:

> A map in the usual sense of the term, is a representation, or symbolic description, *of* something; it can be scaled, distorted, incomplete, and can even contain errors, in which case those who attempt to use it would likely deem it a poor, or a not very useful, representation. This brings us to another important feature of maps: They are used by *people* for one purpose or another, an example being to find one’s way in unfamiliar territory. (157)

We can first ask why this is restricted to people. Often maps are artifacts; my map of Pittsburgh is printed on a large sheet of paper, for example. I could memorize some relevant bit of it, however, and then use my memory of it to navigate. I could also construct a map “in my head”–perhaps of the relation of my office to my home and alternative routes between the two. Non-human animals, of course, do not use maps like my paper map of Pittsburgh. Could they create and use maps similar to the one in my head? It at least seems possible that they could. Of course, our assessment of this will involve behavioral cues; nonetheless, these cues give us reason to think that some non-human animals do indeed use maps.

Consider Packard and Teather's 1998 work on rats (see Churchland, 2002, 276-277 for a further description). The rats are trained on a T-shaped maze in which the cheese is always to the left. They learn quickly enough where it is. They are then tested with the maze turned around: The cheese remains in the same spot in the room, but running left will not get the rats to it. Without further training, the rats now turn right, running directly to the cheese. The result indicates that it is not a simple conditioned response involved, but that the rats use a representation of the layout of the maze relative to the room. This should not be overly surprising. As Churchland notes: "Many animals, including dogs, horses, bees, and bears, exhibit behavior that shows good spatial representation, such as finding novel routes home" (277). It seems fair to talk of the animals as using/having maps; at least, they behave in essentially the same way I do when I use the map in my head.

The restriction to people, therefore, is at least somewhat dubious. What if we replaced it with a restriction to animals (or certain classes of animals)? The restriction is still dubious: Many machines can perform the same type of behavioral tasks. Consider what is perhaps the first such robotic system to do this, Shakey developed by the Stanford Research Institute starting in 1966:

> Shakey inhabited a set of specially prepared rooms. It navigated from room to room, trying to satisfy a goal given to it on a teletype. It would, depending on the goal and circumstances, navigate around obstacles consisting of large painted blocks and wedges, push them out of the way, or push them to some desired location. Shakey had an onboard black-and-white television camera as its primary sensor. An offboard computer analyzed the images and merged descriptions of what was seen into an existing symbolic logic model of the world in the form of first order predicate calculus. (Brooks, 1991, 1227-8)[4]

[4] See Dennett (1991, 85-95) for a more philosophical look at Shakey.

One thing that makes Shakey a particularly interesting case to look at is that a large collection of papers and technical notes about it are available (see Nilsson, 1984). That is, we have the information available to specify in excruciating detail what is involved in Shakey navigating through its world. In particular, we can specify exactly what is meant by saying that Shakey uses a map or an internal spatial representation of its world. The "world" is diagrammed below (adapted from Russell and Norvig, 1995, 361):

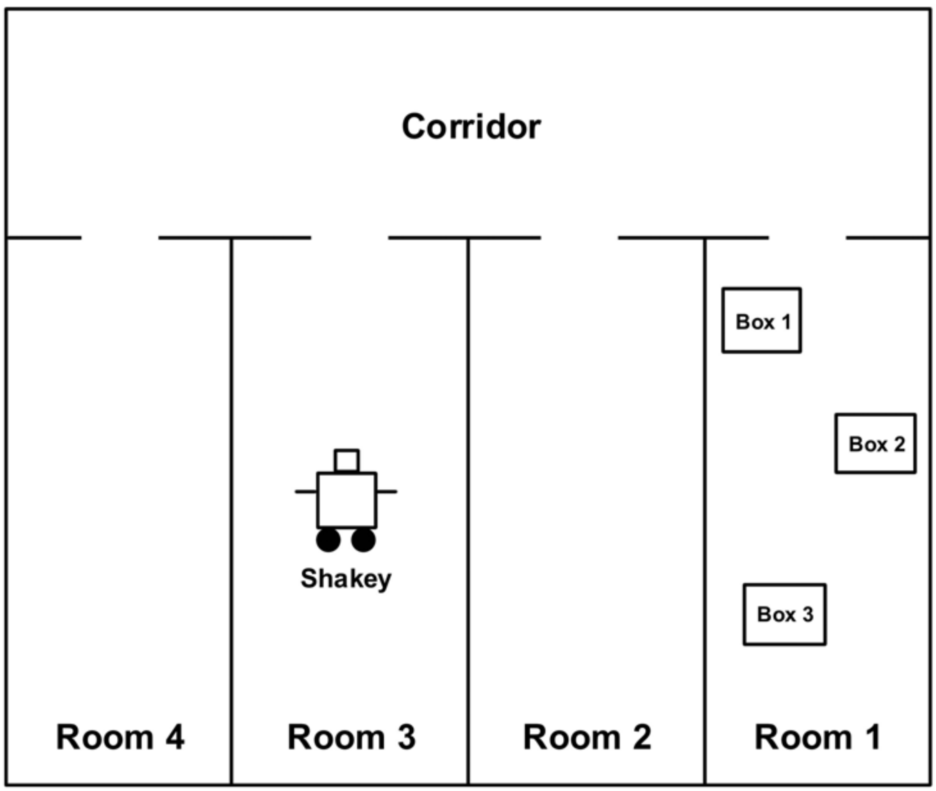

Although Shakey is described as using a "model" rather than a "map," the point is the same. This model is "the heart of the software that controls Shakey" (Nilsson, 1984, 6). Further, this is necessitated by the developers' goal of "enabling an automaton to function independently in realistic environments" (4)–if this environment was to be "extensive" and "non-trivial," they needed to focus on "methods for *efficient* internal representations of the world" (5). The method they came up with was to use a "single model for all [Shakey's] operations" (19); this was rendered as "a global data structure that can be accessed and modified by the other routines" (6).

The model was “a collection of predicate calculus statements stored as prenexed clauses in an indexed data structure” and utilizing “five classes of entities: Doors, wall faces, rooms, objects, and the robot” (19). For these, “all distances and locations are given in feet and all angles are given in degrees; the quantities are measured with respect to a rectangular coordinate system oriented so that all wall faces are parallel to one of the X-Y axes” (19). I will assume that this is enough detail to alleviate concerns that the use of “model” is vague, or that “the sense in which the term ‘representation’ is used is completely unclear” (Slaney and Maraun, 2005, 159; the charge is leveled against J. P. Frisby, as cited in Hacker, 1991).

Shakey is well described as using a map (its model). Furthermore, the map is *in* Shakey and we can specify exactly what it consists in. Shakey is an automaton, a mechanical system, such that the causal workings of Shakey’s parts explain how Shakey does what it does. These parts include its offboard computer and its physical memory. Shakey’s map *consists in* the states of some of this memory; it is *described in terms of* the data structure and predicate calculus statements that the engineers used to program it. That is, we can draw a distinction between Shakey’s model and the programmers’ descriptions of it. It would be a conflation, for example, to think that the diagram above is *in* Shakey *in that form*. Rather, Shakey is set up such that the information depicted in the diagram is encoded in states of some of Shakey’s parts such that the regular operation of those parts is causally involved in enabling Shakey to navigate its environment.

With this extended example as background, let’s return to Slaney and Maraun’s discussion of maps in the brain. “Map” doesn’t just refer to the everyday artifacts we typically think of, but also has a technical sense that Slaney and Maraun note: We can also speak of “the act of ‘mapping’, in which elements from one domain can be transformed into another (codomain) via some mathematical function” (157). They hold that this “may indeed be a legitimate aim of the neuroscientist" (157); for example, "it is reasonable to say… that the retina can be mapped onto the visual cortex" (157). Their concern is that this merely describes a relationship, however; that this "is *merely* to say that particular patterns of stimulation of the rods and cones of the retina are *predictive* of corresponding patterns of the firing of neurons in the visual cortex" (157; italics added).

Something similar could be said about Shakey: There is a *relationship* between the patterns of inputs to its television camera and the states of its memory. This is not *merely* a *predictive* correlation, however. There is a *causal* relation between the two that legitimates talk of Shakey having an internal map or representation. In similar fashion, neuroscientists have detailed how the "stimulation of the rods and cones" *causally produce* the "corresponding patterns of the firing of neurons in the visual cortex." The neurons involved carry out a "mapping"–they can be described as *computing* a mathematical function via a series of causal changes that generate the pattern of firings in the visual cortex from the pattern of stimulation on the retina. If the pattern of activation is thought to be causally/mechanistically related to organismal behavior, then it is typical to talk of the neuronal activation as a representation. If the pattern of activation in the visual cortex is involved in causally enabling the animal to appropriately navigate through or interact with its environment then we can meaningfully say that it is a map. This is exactly what is meant when we talk of Shakey having a map of its world; it is absurd to claim that such talk is meaningless or useless in the case of the robot/memory and correspondingly strains credulity to make such claims in the case of the animal/brain.

Slaney and Maraun's deeper concern is that the supposed "map metaphor" is often "indicative of a particular theoretical perspective on the part of the researcher, namely, the *representational theory of mind*" (158). They hold that talk of brains representing, or having internal representations, is nonsense. This gets into some tricky waters. Much work in philosophy of neuroscience is involved in working out what is meant by such representation talk. (See, for example, Part V of Bechtel et al., eds., 2001). A perusal of this literature will, I hope, show just how shallow and ill-considered Slaney and Maraun's critique is; for example, their conclusion that "the claim that 'there exists a symbolic description in the brain' has no *sense*, and so would not be amenable to empirical verification (or falsification)" (160) simply fails to engage with the literature or to attempt to understand what such talk might mean.

Their argument is that:

> For something to be a representation, the concept denoting it must have a rule-governed use; that is, there must be a correct, and so also incorrect, way of using the concept. It must have a grammar, or criteria within which its applications are bounded, and it may only be used by symbol-employing creatures. It is unintelligible to suggest that the brain is capable of representing anything at all, let alone sensory stimuli, and that the mind has access to such representations. Humans, not brains, can be said to be capable of representing, and representations are created by language users. (159)

This argument is, in fact, far more confusing than most of the examples that Slaney and Maraun critique. The relationship between representations and concepts (as well as the meaning of "concept" itself) is especially unclear: It is not clear why we should think that being a representation depends on there being a concept that denotes it and which has a rule-governed use. A photograph is surely a representation. Does its being so depend on a concept denoting it? (If so, which one?) It is also unclear that concepts (as opposed to languages) have grammar. This will depend on the theory of concepts you adopt (for example, if concepts are analogous to words in a "language of thought"). Further, it is unclear why brains lacking such concepts (supposing that they do) must lack the representations that the concepts denote. I suspect, however, the confusion involved is due to Slaney and Maraun failing to recognize the distinction between a regular process being *describable by* rules and that process being governed or *produced by* rules.[5]

[5] As Wittgenstein writes: "We must distinguish between what one might call 'a process being in accordance with a rule', and, 'a process involving a rule'" (1965, 13).

Mechanical systems can be highly regular and this regularity can be described through rules (or through concepts whose use we specify in terms of rules); but, this does not necessarily mean that the system's behavior is *governed by* those rules. Of course, what neuroscientists are trying to do is to describe the workings of brains; they use *their concepts* to give the descriptions. Brains need not be "rule-governed," merely causally regular. It is this causal regularity, again, that is at the heart of representation talk in neuroscience (just as it is in robotics). This talk is not nonsense, as we have seen; it is both meaningful and useful.

While Slaney and Maraun give definitions that the neuroscientific usage sometimes "transgress," they have not shown that their definitions are the right ones. They haven't even clearly stated or defended what being "right" would mean here. What we are left with is the fact that representation talk is both useful and meaningful in neuroscience. A similar story could be easily told for each of the phrases that Slaney and Maraun object to above, although I won't bore the reader with the details. Until the objectors meet *their* burden of proof, there is no reason to.

4. Conclusion

The bulk of Slaney and Maraun's critique of the use of language in neuroscience rests on their pretending that particular terms have exact definitions that they are in a position to articulate for us. From this premise they argue that neuroscientific usage that transgresses these boundaries is confused. Confusion, however, is not best conceived in absolute terms. Much neuroscience *is* confusing to the uninitiated; *some* is also confusing to the initiated. The primary target of actual neuroscientific work is other neuroscientists or related scientists. If particular terminology is confusing to them, then it stands in need of clarification.

The point is that confusion is relative to a community. Most of the terms that Slaney and Maraun target presumably confuse them, but this is simply not to say that the use is confused in any absolute sense. Bluntly, most of the usage they target is not necessarily confused (and certainly not confused to the point of nonsense). Neuroscientific usage might not follow the definitions that Slaney and Maraun claim to find in their everyday sense of the terms, but it does not need to. Language is not so rigid as that. While Slaney and Maraun might aspire to be language police, such policing is inappropriate. It is inappropriate because we have no reason to think that language is "everywhere circumscribed by rules" (Wittgenstein, 1958, §69), and no reason to think that it needs to be. Even if it were, Slaney and Maraun give us no reason to think that the rules would be solely and exclusively those that they claim to find in popular usage.

References

Bechtel, W., Mundale, J., & Mandik, P. (editors) (2001). *Philosophy and the Neurosciences*. Oxford: Blackwell.

Bennett, M. R. & Hacker, P. M. S. (2003). *Philosophical foundations of neuroscience*. Oxford: Blackwell.

Brooks, R.A.(1991). Intelligence Without Reason. *Proceedings 1991 International Joint Conference on Artificial Intelligence (IJCAI 91)*, 569-595.

Churchland, P. S. (2002). *Brain-Wise: Studies in Neurophilosophy*. Cambridge: MIT Press.

Dennett, D. (1991). *Consciousness Explained*. Boston: Little, Brown and Company.

Hacker, P. M. S. (1991). Seeing, representing and describing: An examination of David Marr's theory of vision. In J. Hyman (Ed.), *Investigating psychology: Sciences of the mind after Wittgenstein* (pp. 119-154). London: Routledge.

Kenny, A. (1984). *The legacy of Wittgenstein*. Oxford: Basil Blackwell.

Machamer, P. K. & Sytsma, J. M. (2005). Neuroscience & the Nature of Philosophy. *Iride*, 46, 495-514.

Marr, D. (1981). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.

Nilsson, N. (1984). *Shakey the robot*. Technical Note 323, SRI International, Menlo Park, California.

Packard, M. & Teather, L. (1998). Double dissociation of hippocampal and dorsal-striatal memory systems by post-training intra-cerebral injections of 2-amino-phosphopentanoic acid. *Behavioral Neurosciences*, 111, 543-551.

Purves, D., Augustine, G. J., Fitzpatrick, D., Hall, W. C., Lamantia, A., McNamara, J. O., Williams, S. M. (editors) (1997). *Neuroscience*. Sunderland, MA: Sinauer Associates.

Rosenzweig, M. R., Leiman, A. L., & Breedlove, S. M. (1996). *Biological psychology*. Sunderland, Mass.: Sinnauer Associates, Inc.

Russel, S. J. & Norvig, P. (1995). *Artificial intelligence: A modern approach*. Upper Saddle River, NJ: Prentice Hall.

Slaney, K. L. & Maraun, M. D. (2005). Analogy and metaphor running amok: An examination of the use of explanatory devices in neuroscience. *Journal of Theoretical and Philosophical Psychology*, 25(2), 153-173.

Whitaker, H. A. (1971). *On the representation of language in the human mind*. Edmonton: Linguistic Research, Inc.

Wittgenstein, L. (1958). *Philosophical Investigations, Third Edition*. Upper Saddle River, NJ: Prentice Hall.

Wittgenstein, L. (1965). *The Blue and Brown Books*. New York: Harper & Row.

Young, J. Z. (1987). *Philosophy and the brain*. Oxford: Oxford University Press.