

The View from Vector Space: an account of conceptual geography¹

Joshua Stein, New York University

Beneath the discussion of mental representation, this paper is about geography; it attempts to bring together two factions in California, the mind-as-joint-relation camp in northern California and the eliminative materialist camp in southern California. It emerges out of a pair of presentations given as an undergraduate that were received well by those in the mind-as-a-joint-relation camp, much to my surprise as someone with substantial eliminativist leanings. That experience led me to reevaluate the space between the factions; this paper is a result of that inquiry.

It is fortunate that the findings of this inquiry should be published in a neutral zone for those two factions and that much of this work was done while living in New York, far removed from the course of the conversation, admitting retrospection. On the other hand, the general removal of this retrospective discussion from the first-order discourse on which it comments requires a substantially expanded expository section.² The geographical and ideological distance between the eastern institutions where this paper was written and discussed led to a change in idiolect and grounding conceptual space, to make it accessible to those not familiar with the complexity of California's varied philosophical circles. To those reading in California, some of the conceptual exploration will seem rather odd, as my unpacking is responsive to those I have engaged with; I will not apologize for that, but ask for some patience.

The paper focuses on the account of mental representation put forward by by Teed

¹ Originally Published 2014. *Journal of Cognition and Neuroethics* 2.1: 71-91

² I am grateful to, among others, Arnon Cohen, Brent Kiou, Benedicte Veillet, and an anonymous commenter for clarifying the areas of the expository section necessary to make the discussion comprehensible.

Rockwell and Paul Churchland;³ it argues that, despite significant differences between the two accounts, there are important shared features that tie the two together. In tying the two positions together, it develops an argument for conceiving of mental space as vector space, and of neurological structures as a map of that space. Further, it illustrates the substantive philosophical and scientific success held by both theories e.g., the successful handling of learning histories and comparative psychology. I maintain that the differences between Churchland and Rockwell are neither entirely empirical nor philosophical; rather, the theories are successful in part because empirical differences bear on philosophical differences, and vice versa.

Finally, this discussion is a product of both accounts, and so finds itself aligned with one account or the other on particular issues. In the course of discussing some of the differences between the two accounts, I offer support one account, or subordinate claim within an account, over the other. I hope to leave you, the reader, with a something close to a complete picture of the synthesis of the views I have come away with, and an idea of where it fits into contemporary philosophical discussion.

1. Prefacing Churchland & Rockwell

Before moving into a positive account of the views presented by Churchland and Rockwell, it is helpful to characterize the position they emphatically reject; the more common account of mental content given in the literature on philosophy of mind from the 1960s onward.⁴

³ While large portions of this discussion are drawn from earlier work of both Rockwell and Churchland, the views are laid out fairly fully in their recent work. Churchland has discussed this approach for several decades, but most recently in *Plato's Camera* (2012). Rockwell lays out his account in *Neither Brain nor Ghost* (2007).

⁴ I am thinking here of the work following Chomsky, including work following Fodor on language of thought, laid out at length in Fodor (1975; 1981) and in brief in Ayede (2010). Fodor's view is notably folk-psychological, in the sense derided by Churchland and Rockwell, but there are various strains of the view that are not, e.g. those developed by Carruthers (2006) and Dennett (1986; 1998).

Characterizing this family of views, itself, could be an exhausting project, but I will just give a brief preface to the relevant facets of the literature. These are the ‘semantic content’ views, which maintain that the content of a representation is the truth-conditions for that representation, like content of a sentence in compositional semantics. The representation of my reclining on a brown coach is specified by conditions where "Joshua is reclining on a brown coach." is true and false. These are the conditions whether the representation is attached to a belief, desire, imagination, etc. The content of a representation are specified the same way as the contents of a linguistic token. Some members of this family take the account literally, suggesting mental content is given in an actual language of thought with definite syntactic structure, like Fodor’s *mentalese* (Fodor 1975; 1990). This family of views has been defended extensively, with interpretations ranging in empirical and conceptual commitment; many suppose something weaker than a language of thought, but maintain a semantic view of mental representation (Perry and Barwise 1983; Searle 1984; 1990; 1992; Stalnaker 1984). For a brief overview of the points of dispute of the language of thought hypothesis, see Ayede (2010, §5) bearing in mind that the views of Rockwell and Churchland are an idiosyncratic variety of what Ayede calls “functionalist materialism.” The semantic content family concerns what mental representations are (Siegel 2012, 7–8; 21–23), rather how they come to have the character they do (Prinz 2000), but positions on the latter bear on the former. Indeed, one purpose of the semantic contents views is to give an account of the formation of mental content, physically and computationally (Rey 1992, 289; Haugeland 1985). This is also the goal of the vector space view, as we will see.

There are a few reasons that Churchland and Rockwell reject the semantic content family of views, but one is particularly relevant to our discussion. The semantic content views attempt an assessment of mental content *per se*, independent of the physical structure that conditions its

instantiation. Rockwell writes,

“... advanced neuroscience will not just give us more information about what the brain does and how it does it. It could also end up eliminating the whole concept of brain, just as easily as it could eliminate any other concept originally derived from folk psychology.

Fodor, [the exemplar of the semantic content family] however, is quite explicit in this commitment when he claims that psychology must accept what he calls “methodological solipsism” (Fodor 1981). What he means by this is that mental states must be studied as an independent system that takes entirely within a brain, which can be understood without any reference to the outside world...

I call this myth... the myth of “the machine in the machine.” It is the basis for Fodor’s “language of thought” and any other theory of mind that holds that all we need to understand the mind is to open Skinner’s black box.” (Rockwell 2007, 10)

Churchland has a similar view (in an omitted portion of the passage above, Rockwell says that Churchland “almost breaks free”) in the consideration of a form of functionalism “construed broadly as the thesis that the essence of our psychological states resides in the abstract causal roles they play in a complex economy of internal states mediating environmental inputs and behavioral outputs” (Churchland 1989, 23). Churchland’s account of vector-spaces is offered to provide a very specific, neuroscience-driven account of how the environmental inputs realize complex representations that allow for sophisticated behavioral outputs.

1.1 Churchland’s account of vector-space & activation pattern

Churchland's alternative to the semantic content views is to account for mental states as neurological activation patterns (2002, 27–32; 1996, 21–36). The potential activation patterns in a given region of the brain map a vector space, and activation specifies the content of a tokened representation. The structure constitutes a state-space in which a given state of affairs can be represented. The neurological structure constitutes the standing possibility of possible representations i.e., the state-space underlying any and all possible representations, and when the perceptual system is stimulated in a certain way, those neurological structures activate in a

pattern specifying the content of the representation. Churchland arrives at something very much like an identity claim about the relation between mind and brain, where the vector space and the vector-map collapse together. This leads to the controversial claim that mental states simply *are* activation patterns, rather than discrete entities.⁵

Any sense in which there is some predicate content of the sort postulated by a classical content view is, then, capable of being given in terms of a vector coding (Churchland 1996, 29-30; 1989, 154-196), and content given in the vector space account is capable of being ascribed truth-conditions. What Churchland is really after is a story about how the neural structures realize rich representations. The general account is straightforward. He starts his own account,

“... at the bottom of several ladders, each of whose bottom rung is a population of *sensory* neurons—such as the rods and cones in the retina of the eye, the ‘hair cells’ in the cochlea of the ear, or the mechano-receptors within the skin. Each of these initial neuronal rungs projects a wealth of filamentary axonal fibers ‘upward’ to a proprietary second rung—to a receiving population of postsensory neurons... each *rung* of these ladders constitutes a unique cognitive canvas or representational space, a canvas or space with its own structured family of categories, its own set of similarity and difference relations, and its own peculiar take on some enduring aspect of the external world. What happens, as sensory information ascends a ladder, is its progressive transformation into a succession of distinct representational formats, formats that embody the brain’s background ‘expectations’ concerning the possible ways in which the world can be.” (2012, 35)

Consider one of Churchland’s favorite cases, facial perception. Churchland (*ibid*, 7–10 and 62–74) gives a fairly straightforward account of how neurological activation patterns specify the particular details of a face and allow for facial recognition within the state-space, showing

⁵ There are a number of important philosophical objections. The first is the standard line raised against an identity relation given by Chalmers (1997). I do not think that this objection turns out to be as devastating, even to someone who is sympathetic to much of Chalmers’ account, since we can say that the identity relation is not a matter of the intension, but a matter-of-factual claim about the relation in our world; this response requires additional fleshing out. The second objection comes from Bennett and Hacker (2003; 2007). I take this objection to be answered by methodological and conceptual criticisms advanced by Dennett (2007).

that the neural network is “trained to discriminate faces from nonfaces, male faces from female faces, and to reidentify the faces of various named individuals across diverse photographs of each” (ibid, 7). The multi-dimensional state-space realized by laddered neural networks allows for representation and, through weighting of certain collections of neurons throughout the system, memory of a complex array of faces. This offers a theory of facial perception and representation, and lends itself well to projects in mapping the relationships between neurological and psychological states.

The view can be extended even further, offering an explanation of the representation of salient social features, like facial expressions. When a subject sees someone experiencing distaste, there is the presentation of the visual stimulus of the face to the visual system, the categorization of the face as such in the activation pattern, and then the sequential activation of a series of regions and subregions associated with features of the facial recognition vector space i.e. the progression of rungs. Like recognition of a face as a particular person, recognizing the expression on someone’s face is a higher rung on that ladder. The behaviors associated with the ascription of a particular mental state are recognized by the activation of certain regions of the vector space, and those regions are connected to other regions related to other features of the cognitive architecture e.g. simulation of affect, or whatever consists mental state ascription; the cascades set off by the initial presentation of the visual stimulus to the visual system are sufficient to explain the progressive enrichment of the representation, e.g. the recognition of the face, the association of that face with a particular person, the ascription of mental states to that person, etc. The subject ascribes distaste based on the progression of the activation pattern representing the stimulus, explaining the defeasibility of judgments about mental states when certain features of the activation pattern are incidentally activated by a stimulus, like the

phenomenon of representing faces in Arcimboldo's paintings of non-faces.

As both the semantic and vector space accounts note, there are a number of functions tied up in facial representation pertinent to the derivation and ascription of emotional responses. These features are often subtle and not obviously part of a semantic operation as they are not given as propositional. Particularly tricky for semantic content views are motor function and affective processes mapped to limbic activity (Gopnik and Seiver 2009; Hutto 2012), as well as visual attention to features of the face, which explain heterogeneous results for tests of facial recognition among subjects with Autism Spectrum Disorder (Hobson and Hobson 2012). Because these features require an account of embodiment, of the way that neurological features interact with parts of the peripheral nervous system and sensory organs, describing them as propositional attitudes is problematic.⁶ The semantic content views yield a sketch of social cognition where rendering a representation is a straightforward inferential process, distancing themselves from a discussion of the relevant physiological and behavioral differences in the subject. (We will revisit the importance of this difference when we talk about reasoning later.)

The major success of the vector space view comes directly from its rejection of semantic content views; by refusing to describe mental states as abstract entities with content given as truth-conditions, the vector space view can focus on embodiment when required. It can turn to the relation between sensory organs and the ladders of neurological structure that participate in representation, offering a mechanistic account of the content. If the content is handled more easily as an abstract object, we can render a heuristic account of that content which allows for that treatment. On the other hand, we can also provide a fine-grained account, a "highly specific

⁶ For discussion of social cognition in particular, see Goldman and de Vignemont (2009) and Niedenthal et al. (2005). For general discussion of the role of embodiment in perception and cognition, see Noë (2006) and Lakoff and Johnson (1999). The latter two, particularly Noë, are exemplary of what I call the "mind-as-joint-relation" camp.

and very different [unit] of representation” (Churchland 2012, 5) embodied within an organism embedded in an environment.

1.2 Rockwell’s account of mind as embodied and embedded

Rockwell’s account does not drift too far from Churchland’s, either in the spirit of characterizing the mental as embodied or in the vernacular used to tell the story. I trade on the linguistic similarities to reinforce the similarity of the accounts, but the deeper conceptual similarity can be demonstrated simply by looking at their goals. Churchland’s goal is to offer a fine-grained account of mental content based on the neurological structures of the organism itself. Rockwell is not so interested in this process of specification, but rather in acknowledging the role of the environment in instantiating mental states; he gives a conservative version of the claim, “even if there are or could be times when the brain is having experiences without receiving stimuli from the outside world, at least some of our experiences would not be recreated even if all of the appropriate neural activity were activated” (Rockwell 2008, 65). We need the peripheral nervous system and sensory organs to tell the story of our mental states. We need our embodiment to explain the mental representations that we have and we need the world to help bring about those representations.

I suggested above that Churchland’s account requires embodiment; this is true. The first step in the story that Churchland tells is at the level of sensory perception. But the goal of Rockwell’s account is to bring the embodiment of the organism into the foreground, to show that (both conceptually and empirically) we cannot tell the story simply by giving a description of the brain states, nor by accounting for abstract functions. Rockwell uses embodiment to emphasize the role of the world in facilitating mental states. Because we are organisms perceiving, the act of perception needs an object; it needs a world perceived. Another way, the sort of embodied

organism that we are must be embedded in a world.

Rockwell prefers the language of vector spaces (Rockwell 2008, 55) for talking about the specification of mental content. Though his language is ambiguous about the status of representation, and the use of the term "content," the view is interested in the same phenomena (Rockwell 2007). Because of the heavy emphasis on embodiment and embedding, Rockwell takes "representation" as the psychophysical states activated by sensory inputs and governing motor functions. This locates him firmly in the "mind-as-joint-relation" camp in northern California, among George Lakoff, Alva Noë, Bert Dreyfus and others.⁷ The entanglement of representation with other cognitive functions makes this view appealing, because the view demonstrates the continuity of direct perceptual experience (the stimulation of the receptor cells) with the peripheral and central nervous system. It does not require an attempt to discretely identify cognitive functions. Instead, it just takes "representations" to indicate the role of some set of functions within the system.

The neuroconstructivist view advanced by Rockwell is similar to the view of brain plasticity advanced by Churchland (1979), because both have an interest in the serviceability of epistemology. They want a theory of representation that can account for some representations as instances of knowledge. Both use "attractor" to indicate a series of vectors activated together, changing the vector space over a period of time (Churchland 1989, 207–208). This neurological and mental change is the basis for an account of learning. Like Churchland, this view of learning

⁷ Those who are particularly historically astute will note that some of the primary contributors to this tradition are not a part of the immediate geographical region, among them Edwin Hutchins (once colleague of Churchland at UC San Diego) and Mark Turner. The geographical sketch is a rough heuristic for identifying communities, recognizing the immediate influence of communities of philosophers on each other. Eliminativist discussions have played a greater role, generally, in southern California thinking on mind (even amongst sympathizers with embodied psychology, e.g. Antonio Damasio) and embodied cognition has been definitive of the discussion of philosophy of mind in much of northern California, especially around UC Berkeley.

allows Rockwell to tie representation into other cognitive functions, instead of giving an account of reasoning as the application of abstract functions to tokens of propositions, like the semantic content views do. For the eliminativist, we come away learning history as neurological change, creating simple causal stories about precisely how neurological structure changes between the initial state and the learned state; for the mind-as-a-joint-relation camp, the story admits of more conceptual complexity, because learning involves the presence of a certain sort of stimulus i.e., one that acts as justification, maintaining the philosophical difference between learning and the fostering of delusions. The two accounts experience tension (and the two advocates experience outright disagreement) in the desire of the eliminativist program to focus heavily on the brain-state and frustration with that approach within the mind-as-a-joint relation camp, but they can be conjoined so the fine-grained account of representation offered by Churchland fits neatly with the extension of the mental state into the body and world.

2. Points of Incompatibility between Churchland & Rockwell

The views, like the regions where they emerged, can appear to rather close to outsiders, but disparities should not be understated. They emerge out of very different approaches to philosophy, and are independently attenuated to conventional claims in the contemporary philosophy of mind. I will come to the former a bit later. My intention is to show that the general philosophical approach turns out to be largely irrelevant to the stories they want to tell about mental representation and, thus, to my re-appropriation of those stories. The burden is on me to show that programmatic baggage can be divorced from the account. With this in mind, a number of the important philosophical platitudes are immediately available for reflection and evaluation, and their role in these accounts helps establish the contrast that needs to be addressed.

2.1 neurocentrism & the empirical assessment of the peripheral nervous system

Consider the claim explicitly disputed by the “mind-as-joint-relation” philosophers like Rockwell; the claim that mind is a function of/caused by/identifiable with the brain. The mind is in, or emerges from, the space between our ears. I will call this view ‘neurocentrism.’ Rockwell says that neurocentrism is false; it is an open empirical and conceptual question as to whether the mind is even subdermal. Far from being an obvious platitude used to evaluate whether an account of mind is plausible, this claim faces a great deal of serious argumentation, especially on the basis that the peripheral nervous system and the sensory organs play an important role. Consider a counterexample to neurocentrism in sensorimotor coordination and embodiment. It seems that sensorimotor coordination occurs largely in the peripheral nervous system (Rockwell 2010, 734–744), in parts of the system that respond directly to stimuli rather than requiring input from the neurological system; e.g. a sharp or hot sensory input occurs in the hand, and the peripheral nervous system proper responds by contracting the arm to pull the hand away; the sensory input will still reach and be represented in the brain, but the actual task is being carried out in the peripheral nervous system and so a substantial portion of the mental state occurs in the arm, not in the brain.

While Churchland has argued that much of our sensory motor coordination is based in the central nervous system (Churchland 2002, 26–27; 1996, 91–96) both he and Rockwell take this to be largely an empirical issue. Which systems are responsive to stimuli that cause disorientation? Which are activated in the performance of certain tasks e.g., regaining balance after a slip? These questions are not a matter of introspection on the phenomenal character, but rather about the empirically assessment of the central nervous system in the performance of these tasks. “Mind-as-joint-relation” advocates, including Rockwell, take this as an empirical point in

favor of the notion that neurocentrism, as advocated by Churchland, is false.⁸

Why should we care about the acceptance or rejection of this platitude? Partly because it is not an accident that one account challenges a platitude while the other accepts it.

Neurocentrism is central to Churchland's views because of his account of developmental psychology and plasticity of mind, not simply as an incidental feature derived through the recitation of the platitude by his peers. Roughly, it is important to Churchland that the child's perception of the world be minimally mediated in order to lend itself to the reliable representation of the world that he is after in his epistemology of science. If embodiment plays a major role in facilitating the development of the brain, then this calls into question whether the development of perception and representation tracks the truth. Churchland is not enormously dogmatic in his espousal of the view, though, so long as complex representation develops as a result of the causal power of the world in a way inclined to track the truth of the situation, something that an advocate of the mind-as-joint-relation view can grant without problem.

2.2 Intersubstitutability of mental & neurological states

Another relevant conceptual issue is the disagreement about the non-intersubstitutability of mental and neurological states, where it is Rockwell who assumes the orthodox position. The mainstream view is simply that, whatever is constitutive of (or identifiable with) a mental state cannot be the neurological state. They are not the same. Therefore, in our description of a state, it is not acceptable, not sensible, to substitute a neurological state with a mental state. Churchland's denial of this claim is well known; as are the conceptual arguments supporting the mainstream

⁸ There are a number of arguments for extended mind. In the case of Rockwell (2010; 2007) and Noë (2010), this claim is taken to be empirically informed. By contrast, there are conceptual arguments for the position, most notably Clark and Chalmers (1998).

view.⁹ Rockwell denies intersubstitutability on the basis that mental states are a joint relation, and so brain states cannot be exhaustive. A mental state includes the neurological state, as well as the state of the peripheral nervous system, salient features of the world, etc.

This is a difference in the conventions governing the terms for mental states, in large part resultant from what they are choosing to focus on in their assessments. Churchland is focusing on the neurological structure, and so it is unsurprising that he advocates for using the language of mental states. Rockwell's entire account emphasizes the importance of the joint relation, and so the fact that he should be so insistent on rejecting intersubstitutability is unsurprising, because he wants to continue to emphasize the importance of embodiment and embedding, and accounts allowing for intersubstitutability generally ignore those features. The fix is simple. For the purposes of one discussion or the other, we might be inclined to note the intersubstitutability or not, but when we perform the substitution, it is always heuristic rather than an exhaustive description of the state. We can give an account of the activation pattern as an account of the mental state, but recognize that the account is not exhaustive because it does not include the full account of that state, including embodied features and the role of the environment.

2.3 The meta-philosophical difference

These differences between the two accounts raise an issue; how are we to reconcile two accounts so different in certain parts of the story that they tell? The position that mental states compose a joint relation between brain, body, and world cannot be cleanly excised, nor can the claim that the underlying neurological state is an exhaustive explanation of a given mental state.

⁹ See Kripke's (1980) and Bennett and Hacker's (2003; 2007) arguments against mind/brain identity theory. We might suppose that arguments against identity of mental states with qualia (Chalmers 1997; Jackson 1998) or intentionality (Searle 1980; 1983) might apply, but those cases are somewhat more contentious, because Churchland can deny that those phenomena occur as characterized in the argument, and that the actual phenomena are conceptually reducible. The general accounts are a much more explicit challenge.

This difference matters a great deal, especially if spatial perception (and, derivatively, representation) is a candidate for embodiment in the way sensorimotor coordination is. All phenomena in the vector space account of representation require an answer as how they are embodied before being able to tell the story. I maintain the best solution here is to use criteria shared by both philosophers, and they will tend to converge on the same answers, at least around the account of representation.

It is important to note the substantive meta-philosophical divide between the two accounts, in order to get at how we can justify (independently) a form of pragmatism from each position. We begin with the realism of Churchland and the pragmatism of Rockwell, with a pair of methodological choices far from exhaustive of our inclinations. In these cases, certain resolutions are satisfying for both camps, though there may be some cases where no such resolution can be reached. What I need is a set of shared criteria that allow for the adjudication of a substantial number of these cases. The shared set of criteria is addressed in Rockwell (1995), what he calls "pragmatic pluralism." This occurs as a criticism of Churchland, based on Churchland's acceptance of pragmatic pluralism, but failure to be sufficiently pluralist about folk psychology. In "pragmatic pluralism," we entertain various theories and evaluate them based on a set of pragmatic criteria, entertaining theories that function with respect to some criteria in certain situations, but not dispense altogether with theories that fail to meet other criteria.

In these shared criteria, we find a strong prospective solution to the issue; with respect to accounts where the mind-as-a-joint-relation theory seems pragmatically viable, (e.g. epistemological issues, reference of mental representation, etc.) we are welcome to employ the account. In cases where we want a fine-grained account, where particular features of the neural map are described, we find the eliminativist account more valuable because of the focus on the

relation between the activation pattern and the state space. Those on either side might argue that there are simply not cases where they are incapable of telling the relevant sort of story; after all, surely an eliminativist can give an account of a particular epistemological state and someone with a mind-as-a-joint-relation position can give a fine grained account of a particular mental state. However we are still inclined to prefer one position or other for offering such an account, to create emphasis, to avoid a problem, and the pragmatic pluralism accepted by both Rockwell and Churchland entitle us to do so. The appropriate solution is to treat the philosophical disputes as points of dispute about which we are non-committal, and qualify the account.

3. Reasoning as functions performed on representations

Perhaps the most contentious part of this account is the notion that this has something of interest to tell us about reasoning, and features of cognition generally understood as non-representational; this is partly because it is a departure from the excursus on Churchland and Rockwell and a venture into a different part of the sandbox. I no longer limit myself to contending that the activation pattern is constitutive of a representation, and its occurrence in a particular region of the vector-map is sufficient for an entire story. Residual of Churchland's account is the process by which the representation is modified, and it is useful to talk about that process; its variance across populations of individuals is of interest to developmental and comparative psychology, as well as the general assessment of cognitive variation in cases like Autism Spectrum Disorder.

As Churchland and Rockwell note, there are regions of the vector-map corresponding to features of a given representation, and an activation pattern in that region of the map corresponds to a representation of a particular region in the vector space. This description accounts for the rich representations that interest us. In the case of someone who has a lesion and becomes

agnosiac, the lesion precludes the activation of the region of the vector-map corresponding to the feature lost by the agnosiac patient, and this account seems exhaustive. However, there are cases where we recognize the cognitive impairment as the inability of the brain to perform a function. Cases of agnosia might be given as a lost ability (e.g. the inability to see faces) though the presence of this locution inclines us to talk about the loss of a cognitive function. We can tell the story of a cognitive power (“being able to see faces”) or of a representational function; it makes no difference. The cognitive capacity of facial recognition given in vector spaces, Churchland notes,

“... embodies an elegant solution to a chronic problem that bedevils the classical or lingua-formal approach to abductive inference generally... What will count, for a given individual in a given situation, as the *best* of the explanatory takes or interpretations available, is a global function of *all* (or potentially all) of one’s current background information. But short of conducting an exhaustive search of all one’s current background beliefs and convictions—an effectively impossible task, at least in real time—how can such an evaluation ever be relevantly or responsibly performed? ... [citing Fodor 2000] the problem here at issue takes center stage as the potential Grim Reaper for the hopes of the classical Computational Theory of Mind...

Where, if anywhere, does the network’s acquired background information reside? Answer: in the acquired weight configuration of the 327,680 synapses meeting the second rung, and in the further weight configuration of the 327,680 synapses meeting the output layer—roughly two-thirds of about a million information-bearing connections, even in a tiny model network.” (Churchland 2012, 68)

In the case of facial perception, the cognitive powers associated with social cognition and facial recognition are functions being performed on representation in virtue of the background information provided by previous events altering the weights of the relevant collection of synapses. The loss of some portion of those synapses then restricts the possibilities for representation and so coincides with the loss of cognitive function. Now, it seems unlikely that either Churchland or Rockwell would ever deny that some things are cognitive functions, like the working memory and prediction, but it is clear on this account that these cognitive functions are,

in fact, parts of the representational process given in terms of vector spaces. The purpose, here, is to illustrate cases where something can be described both as a region of the space and as a cognitive function, like the ascription of mental states. Representational capacities are themselves mental functions.

This gives us access not just to a useful locution, but also puts us in a position to characterize behaviors associated with the loss of cognitive functions.¹⁰ In giving an account of mental states, the functional account helps establish the relation between the neurological function of the region of the vector-map and the relevant set of behaviors correlated with that map. It helps us to tie together the mind and world in a way friendly to their conceptual employment in those cases, without taking away that fine-grained assessment that emerges from the vector-space account of mental representation.

3.1 Applications to comparative psychology

One promising empirical project, though I will only preface it here, is to develop a way of talking about heterogeneity in representational vector spaces, as opposed to talking about the presence or absence of a particular faculty in a given individual. Consider "theory theory" in developmental and comparative psychology (Tomasello 1999a; 1999b; Trevarthen 2012; Gopnik 2009). These accounts, whether given in terms of attentional features or theory-of-mind, treat social cognition not as a mode of representation, but as a function performed on certain mental representations. In showing the way that functions operate in the vector-space account, we see

¹⁰ One of the problems in characterization and ascription of mental states is the role of behavior. Part of the value of the mind-as-a-joint-relation account is that it acknowledges this difficulty as central in most areas of philosophy of mind, refusing to give in to a characterization of behaviorism or psychologism. The eliminativism advanced by Churchland becomes ultimately psychologistic primarily by virtue of its neurocentrism. i.e. Acts of mental state ascription are primarily about internal states. If that neurocentrism turns out to be empirically defeated, then it seems reasonable to suppose that behavior will play more of a role.

that this literature is not lost, but impacted by the difference in approach. The alternative is to recognize a difference in the vector space, the space of possible representations, conditioning the differences in mental representation and cognition in these cases. We shift talk about social reasoning to a set of relations between regions representing features (e.g. vector-spaces for facial representation, spatial representation, temporal representation, etc.) treating some functions as enriching the representation in different ways e.g., processing representations about affect. These processes can, but need not, be given a functional locution, despite being a part of the representational process, as characterized above. What is the upshot?

The upshot is similar to the pragmatic pluralism in the discussion above. There are instances where different approaches turn out to be valuable, either for giving an account of the mind as a computational entity or for looking at correlated neurological structures in the neurological map. There are instances where macrostructural comparisons that look at a set of representational regions need to be more coarse-grained than the vector-space account will allow.¹¹ However, the vector-space account can serve to correct our locution even in these cases. When we talk about the differences and similarities in neurological and mental states between two individuals (or two populations) we often fall into the trap of talking about one group as possessing a particular cognitive capacity and another group lacking the capacity e.g., "Horses do not have color vision." "Sparrows cannot introspect." etc. Where there is a major structural difference, these claims reflect the state of things, but in many of the interesting cases, it leads us to dramatically overstate the differences, or even misrepresent them.

¹¹ Carruthers' (2006) revision of massive modularity is a good candidate for such an account. It provides a compelling and empirically viable approach to discussing differences in the cognitive architecture of various species on the basis of comparison of modules (redefined to be less objectionable than Fodor's initial account) that seem commensurable with this account. Treating the populations of neurons as systems themselves allows Carruthers to keep an eye on the proverbial forest instead of getting lost in the complexity of the neurological microstructures.

One of the major points in comparative psychology is whether the "theory-of-mind" is present in chimpanzees or not, as though this has a "yes or no" answer. It is either there or it is not, supposing that a substantially weakened "theory-of-mind" would not count. This leads to extended and frustrated discussions in the literature in psychology and primatology about whether the evidence supports one conclusion or another; see ongoing disputes between Povinelli (et al. 1997; et al. 1992) and Leavens (2012; et al. 2005; et al. 2004) *inter alia*. On the one hand, Povinelli maintains that the theory-of-mind is not present in chimpanzees in the same way that it is present in young children; on the other hand, Leavens maintains that the theory-of-mind is present.¹² This appears to be a straightforward empirical issue, and is treated as such by psychologists, but it is conceptually more complicated when we peer into the particulars. After all, Povinelli's view is not that there are no social cognitive capacities at work in the chimpanzee, and Leavens' view is not that the social cognition of adult chimpanzees is the same as the social cognition of humans. Rather, the issue is that there are substantive cognitive differences in the chimpanzees, and the extent of those differences can be assessed in part through an analysis of the social behavior of the chimpanzees.

In the initial discussion of Churchland, I gave a brief sketch of how social cognition regarding mental state ascription can arise from the facial-recognition system. We can see that an account of the socially rich representation can be given in the vector space account, but consider that we could give two parallel analyses; a generic analysis for humans and a generic analysis for chimpanzees. In comparing the two analyses, what we find is not a comparison in terms of

¹² David Leavens has written to me to clarify that I mischaracterized his view in the initial publication of the paper; Leavens denies that either children or chimpanzees have a theory of mind in the sense advocated by Povinelli. Rather, the point in the cited passage is that Povinelli et al. uses a double-standard in evaluating chimpanzees and children; they are more equal than Povinelli et al. argues, but don't have a theory of mind. I apologize for the misrepresentation. (Added 1/31/2016)

straightforward presence or absence of a particular mental function like the ascription of a complex mind, but rather the finer differences in the social cognition of the two communities, in how the enriched representations tie in to the affect of the individual and how the affect manifests itself behaviorally. These differences are cognitive manifestations of neurological and physiological differences; some of these differences will result from the environment in which the individual members of the community develop and some will be the result of genetic disposition.

The concerns expressed by Leavens about the developmental differences and the role of environment in facilitating psychological difference is a point of open empirical discussion that lends itself well to the literature on embodiment and environmental embedding invoked by Rockwell. Given the claim that a particular feature of the environment e.g., growing up in a community that uses some linguistic communication, is relevant for the development of certain capacities, we use empirical discussions in developmental psychology to assess impact on human and chimpanzee cognition; then account for the relevant cognitive similarity and dissimilarity at various stages of development. We maintain the more fine-grained analysis per Churchland, but bringing the causal power of the world and the embodiment of the individual in facilitating the development of social cognition into view per Rockwell. The vector space approach outlined above allows for a treatment of the cognitive psychology of both communities in a fine-grained way, engaging with causal structures that interest developmental and comparative psychologists.

The purpose of the discussion is to turn us towards the different cognitive features, which allows us to maintain the empirical discussion that Povinelli and Leavens and co. intended to have, and it still about the same cognitive explananda as before, though framed differently. It deprives us of the more coarse-grained straightforward denial or acceptance, instead working

from the shared assumption that some social cognition is present, that social cognition is substantively different, and that it is the structure of those differences that is of interest to comparative psychology. We deprive the discussion of little more than hot rhetoric, moving away from an all-or-nothing ascription of function in favor of a comparative analysis.

4. closing notes

The stated purpose of the discussion of representation was to map the space between two positions, but the ground we have covered is rather large. As we interest ourselves in the account of representation offered by Churchland and Rockwell, it is difficult to separate that account out from its implications for the philosophy of mind and psychology. Drifting into a commentary on the consequence of such a theory is inevitable, though perhaps I have been a bit more ambitious than is advisable.

The purpose, though, is to illustrate a theory in the philosophy of mental representation. Prefacing the bridges into psychological discussions is important, as the adjacent discussions are a place where the vector space theory comes together with the embodiment and environmental embedding of the organism, demonstrating the synthesis. The commentary on developmental psychology is, at least for Churchland's purposes (and mine), more valuable than the attempt at conflict resolution. There are a number of important conceptual and empirical issues tied up in the general story of representation, and even those who dismiss the accounts given here can make use of how a dismissal positions them on a number of issues; what is at stake in the rejection of the view.

In the course of this discussion I have addressed substantial empirical and conceptual disputes between Churchland and Rockwell, in order to bring together their accounts of vector space. The philosophical differences, both in their disagreement on issues in the philosophy of

mind and epistemology, are reconcilable, as are the programmatic differences that inform the animus of the conflict. I have shown that the first set of differences are likely subject to resolution either through empirical work, or some simple conciliatory conceptual work, and that the larger conceptual conflict does not serve as a major block, because they have shared criteria for the evaluation of many of the claims under consideration. Following that, I addressed the implications of the shared account of representation for a few areas in the philosophy of mind, illustrating with the question of social cognition in comparative and developmental psychology.

Most importantly, I have illustrated the role of reasoning in representation on the vector space view, where reasoning is the enrichment of representation by moving up the ladder of the activation pattern. This models the explanation of the representation of straightforward states, like the presentation of color and the contours of faces, as well as the problematic inferential states that have frustrated prominent philosophical accounts, including and especially the semantic content views I drew on for contrast. The complexities of distinguishing faces and non-faces and ascribing mental states to other individuals can be realized through extrapolation of the vector space account; the account can be fleshed out further by extending the account of mental states out into the perceptual organs, accounting for the origination of the activation pattern in the joint between the body and the world.

The conclusion emerges from that final look at the prospective role for the vector-space account in understanding cognitive psychology. The moderate and naturalized stories offered by the vector space account of representation suggest this approach is promising for empirical psychological work. While, for Churchland and Rockwell, the concepts function as a basic philosophical foundation to address conceptual issues, they can be pressed into service in favor of the resolution of scientific problems, like those in comparative and developmental psychology,

and developed into more sophisticated scientific devices in the course of that discussion. This approach promises a valuable asset in a complex and conceptually difficult workspace.

Bibliography

- Ayede, Murat. 2010. "The Language of Thought Hypothesis." in *The Stanford Encyclopedia of Philosophy*.
- Barwise, Jon, and John Perry. 1983. *Situations and Attitudes*. Cambridge, MA: MIT Press.
- Bennett, Maxwell and Peter Hacker. 2003. *The Philosophical Foundations of Neuroscience*. New York: Wiley Blackwell.
- with John Searle and Daniel Dennett. 2007. *Neuroscience and Philosophy*. New York: Columbia University Press.
- Carruthers, Peter. 2006. *The Architecture of the Mind*. Oxford: Oxford University Press.
- 1996. *Language, Thought and Consciousness*. Cambridge: Cambridge University Press.
- Chalmers, David. 1997. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Churchland, Paul. 2012. *Plato's Camera: How the Physical Brain Captures a Landscape of Abstract Universals*. Cambridge, MA: The MIT Press.
- 2002. "Outer Space and Inner Space: The New Epistemology." *Proceedings and Addresses of the American Philosophical Association* 76(2): 25-48.
- 1996. *The Engine of Reason, the Seat of the Soul*. Cambridge, MA: The MIT Press
- 1989. *A Neurocomputational Perspective*. Cambridge, MA: The MIT Press.
- 1979. *Scientific Realism and the Plasticity of Mind*. Cambridge, UK: Cambridge University Press.
- Clark, Andy and David Chalmers. 1998. "The extended mind." *Analysis* 58 (1): 7-19.
- Dennett, Daniel. 2007. with Peter Hacker, Maxwell Bennett, and John Searle. *Neuroscience and Philosophy*. New York: Columbia University Press.
- 1997. *Brainchildren: Essays on Designing Minds*. Cambridge, MA: The MIT Press
- 1986. *Content and Consciousness*. New York: Routledge.
- Fodor, Jerry. 2000. *The Mind Doesn't Work that Way: The Scope and the Limits of Computational Psychology*. Cambridge, MA: MIT Press.
- 1990. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- 1981. *RePresentations: Philosophical Essays on the Foundation of Cognitive Science*. Cambridge, MA: The MIT Press.
- 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press
- Goldman, Alvin and Frederique de Vignemont. 2009. "Is social cognition embodied?" *Trends in Cognitive Science* 13(4): 154-159.
- Gopnik, Alison and Elizabeth Seiver. 2009. "Reading Minds: How Infants Come to Understand Others." *Zero to Three* 30(2): 28-32.
- Haugeland, John. 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.
- Hobson, Peter and Jessica Hobson. 2012. "Joint Attention or Joint Engagement? Insights from Autism." In *Joint Attention*, edited by Axel Seemann, 115-136. Cambridge, MA: The MIT Press.
- Hutto, Daniel. 2012. "Elementary Mind Minding, Enactivist-Style." In *Joint Attention*, edited by Axel Seemann, 307-342. Cambridge, MA: The MIT Press.
- Jackson, Frank. 1998. "What Mary Didn't Know." in *Mind, Method and Conditionals*. London: Routledge.
- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lakoff, George, and Mark Johnson. 1999. *Philosophy in the Flesh*. New York: Basic Books.

- Leavens, David. 2012. "Joint Attention: Twelve Myths." In *Joint Attention* edited by Axel Seemann, 43-72. Cambridge, MA: The MIT Press.
- 2005. J.L. Russell and W.D. Hopkins. "Intentionality as measured in the persistence and elaboration of communication by chimpanzees." *Child Development* 76: 291-306.
- 2004. W.D. Hopkins and R.K. Thomas. "Referential communication by chimpanzees." *Journal of Comparative Psychology* 118: 48-57.
- Noë, Alva. 2010. *Out of Our Heads*. New York: Hill and Wang.
- 2006. *Action in Perception*. Cambridge, MA: Bradford
- Niedenthal, Paula, Lawrence Barsalou, Piotr Winkielman, Silvia Krauth-Gruber, and François Ric. 2005. "Embodiment in Attitudes, Social Perception, and Emotion." *Personality and Social Psychology Review* 9(13): 184-211.
- Povinelli, Daniel, J.E. Reux, D.T. Bierschwale, A.D. Allain, and B.B. Simon. 1997. "Exploitation of pointing as a referential gesture in young children, but not adolescent chimpanzees." *Cognitive Development* 12: 424-461.
- Povinelli, Daniel, K.E. Nelson and S.T. Boysen. 1992. "Comprehension of role reversal in chimpanzees: Evidence of empathy?" *Animal Behavior* 43: 633-640.
- Prinz, Jesse. 2000. "The Duality of Content." *Philosophical Studies* 100(1): 173-189.
- Rey, Georges. 1992. "Sensational Sentences Switched" *Philosophical Studies* 68: 289-319.
- Rockwell, Teed. 2010. "Extended Cognition and intrinsic properties." *Philosophical Psychology* 23(6): 741-757.
- 2008. "Dynamic Empathy: A new formulation for the simulation theory of mind reading." *Cognitive Systems Research* 9: 52-63.
- 2007. *Neither Brain nor Ghost: A Nondualist Alternative to the Mind-Brain Identity Theory*. Cambridge, MA: The MIT Press.
- 1995. "Beyond Eliminative Materialism: Some unnoticed implications of Paul Churchland's Pragmatic Pluralism." Unpublished.
- Searle, John. 1992. *The Rediscovery of Mind*. Cambridge, MA: MIT Press.
- 1990. "Is the Brain a Digital Computer?" *Proceedings and Addresses of the APA* 64(3).
- 1984. *Minds, Brains and Science*. Cambridge, MA: Harvard University Press
1983. *Intentionality*. Cambridge, UK: Cambridge University Press.
- 1980. "Minds, Brains, and Programs." *The Behavioral and Brain Sciences* 3(3): 417-424.
- Siegel, Susanna. 2012. *The Contents of Visual Experience*. Oxford: Oxford University Press.