Keeping the World in Mind

Introduction

This book records research that began with a simple question; "What are mental representations"? While the question is simple, asking it may seem puzzling. Mental representations have been alive and well in Anglophone philosophy for over 35 years. What could the problem be?

Mental representations, as commonly understood in such philosophy, can seem unproblematic. The inference from "Mary believes that her keys are on the table," to "Mary has a state with the content 'My keys are on the table'" may look obvious. However, they come with a considerable amount of theoretical debt. They standardly are thought to have content and aboutness, but how they acquire them and how they themselves enter into causal relations are unresolved questions. There are answers to these questions; in fact, there are too many, with nothing like a consensus in view.

In addition, recent research in cognitive neuroscience should lead us to question the idea that there are separable levels of explanation, such that explanations in terms of states with content have at least a conceptual isolation from explanations in terms of neural happenings. The now familiar idea that conscious thought is a small part of the brain processing that enables our intelligent actions starts to pose a dilemma. Either the unconscious processing is sub-personal and not really part of our sense of the personal, or the unconscious processing is explicable as operations over mental representations. The first gives us a very exiguous picture

of the personal level. The second, however, neglects the extent to which neural computations –

computations with neural states as values – are providing the successful explanations.

One response to problems concerning mental representations is to try to downplay them

or do entirely without them. Embodied cognition theorists such as Dreyfus(Dreyfus and

Wrathall 2006), Noe (Noë 2004), Gallagher (Gallagher 2005), Hutto (Hutto 2008), Chemero

(Chemero 2009), Wheeler (Wheeler 2005) and others have brought in notions from philosophers

and psychologists outside those traditionally considered in so-called analytic philosophy. In

doing so, they have stressed actions of the lived body in its environment as grounding our

understanding of cognition, even if, with the exception of Chemero, they have still found some

use for a notion of representation. Since neuroscience is replete with discussions of

representations, these philosophers tend to eschew neuroscience, or at least minimally engage

with it.

This alternative is unavailable to me. I have worked closely with neuroscientists and

neuropsychologists; the idea that we cannot learn a great deal about from such researchers about

human cognition seems to have many more counterexamples than one would want to count. One

area where the importance of neuroscience comes in concerns affordances, which are frequently

discussed in embodied cognition theories. One can just see, it is sometimes said, that an object

or setting affords one the opportunity, for example, to sit comfortably, to eat well, to rest a

burden, and so on. There are, however, aspects to the phenomena of affordances that are under

the purview of cognitive neuroscience, which for over a decade now has moved away from the

classical cognitive science of Marr's Vision into investigating the organism as shaped for

flourishing in its niche.

A neuroscientist may notice, then, that what objects afford one varies even for one individual. Osteoimmunology can tell us why a chair that once was so comfortable no longer affords one a good resting place; neuroscience attempts to explain not just what the resulting pain consists in but also how the change in what the world affords is recognized and acted upon. One experience after an accident reveals one's skeleton's weakness might lead one to see over-stuffed chairs as foes and not friends. What about us allows one-trial learning in this sort of situation? Cases such as this have been studied extensively, and we now have theories that promise to explain at least parts of very puzzling features of human life. One such feature, connected with the chair example, has been long recognized; instinct may win out over reason as a guide even to quite sophisticated parts of human life, as we will see in Chapters One, Nine and Ten. And neuroscience has a very good story to tell about how one instinctively updates oneself on many of the changing contingencies in one's environment. The story also may yield an explanation of why updating can radically fail in cases of cigarette smoking or risky behavior on the stock market, for example (Rangel, Camerer, and Montague 2008).

Though the work above was produced largely after the turn of the millennium, enough of the embodied cognition work was in place for me to realize in the late 1990's that it was focused on important facets of human cognition. Cognitive neuroscience, however, did not seem to me to hold we are in effect brains in vats; rather, it aimed to explain functioning well in one's niche. "Mental representations" remained a very contested area. Consequently, I used my library's main search engine to bring up all the uses of "represent" and its cognates it could find that were involved in theorizing about mind states. It was early enough in my university's acquisition of search engines that I got only 250 examples. Hence, the task of examining them all was manageable, though the articles cited still covered a very wide range of disciplines. That sum

was, of course, after I eliminated all the references to 'represent' in other contexts, particularly ones about representational government.

The results were very exciting and even astonishing. While philosophers standardly are agreed that representations have aboutness, a number of other disciplines appeared to attach a very different understanding to their use of "representation." Their use was instead at least very close to the Aristotelian-Thomistic idea that what is sensed and known in the world outside our heads gets realized inside our heads. As I saw when I went to a talk around that time, the experimenter who referred to a pattern of excitation in a monkey's brain as "the movement of the banana in the monkey's brain" might possibly be mean exactly what he was saying.

Making sense of the Aristotelian idea in modern terms was going to require a great deal of work, it quickly became apparent. How could a pattern of movement in the lab be a pattern of excitation in a brain? Further, if we could make sense of it on its own grounds, much work would remain to be done. For example, a significant amount of historical research had to be in error. Many historians of philosophy have been as captivated as everyone else by the idea of representing as having content that is about what is represented. Hence, it is not uncommon for them to find Aristotelian and many later uses of "represents" deeply mistaken. It looks as though the past philosophers were trying to explain content-based aboutness in distinctly different and indeed inadequate ways. The new thesis that our ancestors were actually explaining a quite different term started to look at least as promising. Maybe they were not so egregiously wrong.

Once one starts to think of alternative conceptions of how the mind relates to the world in perception, terms surrounding the standard use of "representation" take on a potent significance. Thus terms such as "content," "satisfaction-conditions" and "true" are seen as semantical terms, which come from an historically circumscribed theory of mind-states. That is, they come from a

linguistic model of the mind, which is hardly the only understanding of the mind that has been on offer during the history of philosophy.

What would happen, then, if one tried to interpret the neuroscientific talk about Aristotelian representations in terms of a non-linguistic model of the mind? This book is a partial first attempt at answering that question. The first half of the book explicates Aristotelian representations and locates them historically, particularly in one prominent historical figure, Hume. The second half considers the implications of Aristotelian representations for some issues regarding concepts, perception, emotions and actions.

The book is, then, an experiment in looking at things differently. To the extent that it succeeds, it is analytic, as opposed to synthetic, in relation to developing a social, non-individual picture of the mind. The pairing here of synthetic and analytic is mine. It may be that any general statement of the distinction involves contentious divisions, though it can still be a useful way to regard a field. We can think of a synthetic investigation as an investigation into the phenomena, and most particularly for us, the ways in which our experience is embodied and embedded in a social life. An analytic investigation looks for a kind of foundation that will tell us what about us makes such a life possible.

We can allow that an analytic investigation can be scientific or, as in the present case, philosophical. The goals of such approaches can vary with individual investigators, but in my case the goal is to articulate a conception of human cognition influenced by cognitive neuroscience that does not subvert the sense that we are embodied creatures. The investigation also seeks to remind us that we are deeply social beings.

The division can be contentious into analytic and synthetic comes principally from the fact that the phenomenal aspects themselves may be considered foundational. So we will look at

two examples where it seems the phenomenal is considered foundation and yet where we might find openings for an analytic investigation. But we should first note that some foundational investigations can in effect cost us the phenomenal. Many embodied cognition theorists appear to agree that classic computational cognitive science tells us we might as well be brains in vats doing a lot of calculating. And this means that our embodiment is not essential; the deep sense that it is becomes invalid. Here the "classical" should be contrasted with current cognitive neuroscience, which is much more concerned with niche-flourishing.

Let us look at two examples where we might open up some space for a foundational investigation, though those providing the initial theses would not envisage such an investigation. The first example is Wittgenstein (Wittgenstein 1963), who took it to be a common view that language and meaning originate with private experience. Wittgenstein argued against such a view and maintained in contrast that language use is essentially a public practice grounded in the fact that members of one's community go on in the same way. Wittgenstein wanted to say that what went on in one's head would not explain language use, so heads might as well be full of straw.

Nonetheless, it seems consistent with most of what Wittgenstein says that we can ask what enables us to go on in the same way, as long as we do not in effect say that we are able to because we are all born with the same inner language. Among other things, no explanation is really achieved by making the outer into an inner phenomenon, Wittgenstein insisted. To uncover ways to describe what enables us to go on in the same way without positing an inner language is an analytic investigation, though definitely not Wittgenstein's.

Another case is the idea of "action-oriented representation," which is posited to capture the phenomenon of something's presence to us being a prompt for action (Wheeler 2005). The

cat is present as for petting, the glass of ice tea as for drinking, and the towel as for wiping. Situations or things can call for action.

In embodied cognition theories the notion of an action-oriented representation appears to be bedrock and does not itself receive an analytic investigation. The notion of representation appears to explain what they are, but now how we can have them. However, we can easily discern such an explanation being developed in cognitive neuroscience. There are investigations confirming the idea that experiences of rewards alter vision at a very early stage in visual processing. Thus a chair literally looks different to someone who sees it as a comfortable place to rest than it does to someone who does not (Shuler and Bear 2006). Further, according to an important theory, what affects vision also activates a relevant motor program (Montague 2007). One result of this is that the idea that perception itself prompts action may be not quite accurate. That is, some event such as a burst of dopamine may both alter one's vision and initiate the movement, so that the vision and the prompt are co-effects of a cause.

Eyeglasses change vision; in much more complex ways, interests in food and rest change vision. It is a prejudice, perhaps reinforced by classical cognitive science, to think that the changes that come from inside us must block or interrupt our access to the world in a way glasses do not. As Aquinas insisted, the sensible and intellectual species are not what we sense or know; rather, they are the means by which we sense and know. The investigations into the inner changes provide us with too rich a field of results to disregard them from the start.

In sum, this book is an analytic approach to our understanding of human cognition. It explores questions about the elements of cognition. Its central focus is on mental representation, and it is concerned with making possible explanations that leave in place the embodied and embedded phenomenology. Nonetheless, it is not part of the project, at least as I conceive it, to

try to rebut claims from those who are firm advocates of representations as standardly understood today. That is a thankless task, in my experience. In arguing for errors in extant causal theories of representational content one now counts as merely pointing out anomalies. In order to move beyond this stalemate, one should show that there cannot be a correct causal account or one should try to construct an alternative to causal theories. This books begins the latter task.

Chapter One: Introduction

1.   Introduction

In 1982 David Marr's book, *Vision,* appeared to describe permanent and important divisions in cognitive science. He maintained that there are three separate levels to investigating cognitive phenomena: the computational, the algorithmic and the implementation level. The computational level specifies the task to be performed, while the algorithmic level specifies the computational procedures used to accomplish the task. For example, the task might be to find the best among several alternatives, and the algorithmic explanation would specify formulas employed in calculating the decision's result.   The implementation level concerns the realization of the computations in the brain, but in 1982 many theorists did not even want to

conjecture about how the brain carried out the computations that thought supposedly consisted in.

Three decades later, the scene looks radically different as the values for computations in cognitive neuroscience are increasingly neural values. Relatedly, we have quite specific and empirically tested hypotheses about how the brain gives rise to kinds of behavior. We will take a specific example to illustrate how we can now characterize an episode in two quite distinct ways.

Suppose, for example, you are walking down a busy sidewalk in Soho, and you catch a glimpse of an old print in a shop window. You have had good luck in the past in finding wonderful old prints for reasonable prices, and you hurry over to the shop for a quick look, despite the fact that your companions are proceeding on to lunch. This descriptions of this event that we have used so far can provide us with a situation amenable to the standard account of reasons for action: You like buying old prints that are reasonably priced, and you think the shop may well have some.

There is another quite different way of describing your action. We understand now that on the neural level one important factor in many actions is that the dopamine burst which initially occurred when you got, for example, the reward of a good print now attaches to reliable signs of such prints. Further, it affects the perception of such signs; they are more visible and salient. Finally, it starts off the motor movement toward the pictures (Montague and Quartz 1999; Montague 2007).

We have, then, at least templates for fairly full explanatory accounts of two different sorts. The fact that this is so has, in many philosophers' eyes, been the source of problems. One particularly difficult problem has been articulated by Jaegwon Kim (Kim 1996, 2007, 2005, 2006), who has argued that in such cases as these, either beliefs and desires reduce to neural

states or they are explanatorily irrelevant. Relatedly, and equally seriously, the neural phenomena are at a sufficiently low level that it is doubtful they could constitute a basis for thinking as philosophers standardly conceive of it, as computations over mental representations. We are, rather, engaging computational neuroscience in which the values for the computations comes from neural features.

By the end of this book, we should have some good ideas about how to think about this situation, but it itself is not our topic. Rather, we will be looking at a way of thinking about human cognition that has made the situation problematic. One topic that will receive particular attention is that of so-called mental representations.

For now we will instead concentrate the virtues of having these two different kinds of explanation, and particularly the development of the neural explanation. Indeed, we might almost feel a sense of relief at having the general form of explanation for an exceptionally important feature of human life that shows up in many contexts when we perform rapid but purposeful actions. We evolved to be able to act quickly and successfully, and the neural story we have looked at underlies a large part of that ability. We now realize that those who lack the ability to respond quickly and successfully in many situations cannot be adequately compensated for by reasoning their way to an appropriate reaction. Deliberations tend to be unsuccessful when we need a quick (i.e., normal) reaction, and they may in fact cost us too much.

A shift toward considering how we are shaped to meet the demands for survival, reproduction and social interaction also brings out a change that has been in the making in neuroscience for over a decade; thus, much of the investigation in this book is influenced by the thought in the following important quotation (Montague and Quartz 1999)):

… Early investigators thought that the really important problem was to find the functions or computations being implemented by the brain independent of the specifics of their implementation using biological components. This view is now seen as impoverished because as structures constructed by evolution, most creatures are tightly woven into particular environmental and social niches, and are the 'answers' to manifold questions posed by their environs.

This quote rejects as impoverished what continues to be a main model in philosophy of mind, a model that once had few robust rivals for a scientifically acceptable approach to the mind. The model can be seen as a modern counterpart to Descartes' view of the mind, at least as it appears in his most famous work, the *Meditations*. On such a view, our essence as cognizers is to think, where thinking is conceived of on a model of mathematical computation. Further, the mind is whole and entire in itself; it is in important ways independent of its environment, including its social environment.

The quote refocuses our attention on other factors, our niches and the evolutionary history that produced our characteristics. That is, it invites us to see cognizers as products of changes responsive to requirements faced by their ancestors. When the creatures are human beings, this change in focus can be very dramatic, perhaps especially for philosophy of mind. Philosophy of mind is apt to concentrate on features that we think of as very likely to be distinctively human, such as truth-gathering and action planning. However, memory or sensory perception did not evolve to meet the human desire for the acquisition of truths. Even what we might call higher order cognitive functioning bears the marks of our evolutionary past; efficient cognitive function requires that a great deal takes place outside of our awareness, and is not the product of overtly deliberative reasoning.

Walking over to a shop is in fact an immensely complicated task, and much of it cannot be done through conscious deliberation; there are too many factors to cope with consciously. Rather, instinct is our guide, and the alternatives are selected outside of our awareness, perhaps by mechanisms that have the selection of one alternative wired into them. Further, an understanding that equips us for action may not meet other interests we have. Our capacity for survival depends on our missing more details than we may realize. As a consequence, the cognitive sciences are engaged in a quite radical revision of the idea that we are self-aware Cartesian egoes logically sorting through possible solutions.

> We all believe that we are capable of seeing what's in front of us, of accurately remembering important events from our past, of understanding the limits of our knowledge, of properly determining cause and effect. But these intuitive beliefs are often mistaken ones that mask critically important limitations on our cognitive abilities. . . . As we go through life, we often act as though we know how our minds work and why we behave the way we do. It is surprising how often we really have no clue. (Chabris and Simons 2010))

Vision and memory may be sufficient for successful action without necessarily supporting the mental abilities needed for learning much in human culture. Such recent conclusions may remind us of the twentieth century's discovery that human beings can use language meaningfully without being able to define their terms or even explain them well at all. For human beings, the surrounding culture and cultural codes may be more essential than a picture of thinking as computations in the brain supposes. Below we will look for ways in which human thought is covertly dependent on the community for its content.

The Montague-Quartz remark compares interestingly with the thought that life itself might be considered a logical category or that it is life, with the accompanying needs for maintenance, that introduces the special categories in terms of which we can understand mentality (Noë 2009; Thompson 2008; Foot 2001) . What is common to all such theorists is taking the biological status of the human being very seriously. Understanding human cognition on this option will involve understanding in the first place how it subserves demands on the creature's flourishing.

Philosophers stressing the biological status of human beings may not offer detailed biological descriptions. Instead they may employ a biological perspective as weak as Philippa Foot's in her recent discussions of virtues. The virtues, for Foot, are to be understood in terms of the well-being of particular biological creatures at a particular location in time. For Foot the ways in which being truthful aids one's flourishing is not itself a fact that can be understood only through specialized courses in biology. It is much less easy to given an account of how vision aids flourishing without any neuroscience, in part because human beings appear to have a strong inclination to false beliefs about vision, as we will see. However, whether one is talking about virtues or vision, what these theories eschew in contrast is an attempt to treat understanding moral choices or cognition as entirely about what is going on in one's head. As a result, these theorists tend to avoid according mental representations much of a role in their theories, since, as traditionally conceived, they encourage the idea that human cognition occurs in one's head.

We will be stressing how much of understanding human cognition involves understanding the ways in which human beings are equipped to meet the demands of maintaining and replicating their lives. In Chapter Eight we will look quite specifically at the idea of functioning, and thesis that explaining well functioning is the first aim of a theory of

cognition. While the investigation is not itself an investigation in cognitive neuroscience, it is strongly influenced by such inquiry.

As we examine some of the foundations of philosophy of mind and cognitive neuroscience, we may want to look for "the mark of the mental," articulating which has been taken recently to be an important condition on the success of theories discussing the topics we will take up. The agreed-upon candidate so far appear to be mental representation as traditionally conceived (Adams and Aizawa 2008; Rowlands 2010). But such an answer problematically supposes that the mental nature of state is independent of the system in which it occurs. To ask for the mark of the mental independently of its context might make as much sense as to ask for the mark of a bit of swimming independently of its context.

2. Modeling our Cognitive Relation to our Environment

Since there are many different computational sciences these days – computational chemistry, for example – we should be sure we are clear about an important difference underlying the quote above. We can model chemical reactions by employing algorithms (or recipes) but we do not want to claim that the chemicals are using the algorithms. Thought may seem to us very different. In thinking, we might maintain, we use or follow the rules in a way quite different from that instantiated in a chemical system. More particularly, the algorithms apply because of the syntactic and semantic characteristics of the elements of thought. When someone talks about a computational theory of the mind, this is one of the claims we can expect.

The contrasting idea of embodied cognition has gained stature in recent years, but not all of its proponents will highlight the brain in the way the above quote does. This reluctance may be due to the wide acceptance had by the computational theory of the mind, which can be seen as

encouraging the approach labeled "impoverished," and as theorizing cognition as solely a brain activity. The key element of a computational theory of the mind that we will eventually focus on is the claim that a thought is a simple or complex mental representation possessing intentional content or aboutness, along with satisfaction conditions. The mind's cognitive relation to its environment consists, on views employing this model, in having states or processes that are about the environment.

Chemero (Chemero 2009), largely correctly I think, takes his radical embodied cognitive science to be the only sort that avoids representations with content; what this means is the almost all theorists of non-radical embodied cognitive science do admit what the introduction called 'Fodorian' representations, however attenuated the roles of such representations may be. My only qualification on Chemero's claim, I should say, concerns his idea that my theory posits representations with content. It does not.

These mental representations are 'Fodorian representations' in part because our first modern development of the idea of a mental representation is due to Jerry Fodor, I believe. However, that does not mean that "Fodorian representations" employed here means "representations in a language of thought." Even Fodor's account of intentional content, along with his views about how it is determined, are strenuously contested bu many proponents of mental representations. Still, it will be useful for Chapter Two to look at how far we can excise Fodor's view about mental representations from his view about the language of thought.

There is another model of the mind's relation to its environment and it is in cognitive neuroscience. The term "representation" also occurs in characterizations of this other model, and though these other representations are very different, it will be even more misleading to call them something else. So we will label them "Aristotelian Representations." Though they are the topic

of much in cognitive neuroscience and appear elsewhere in cognitive science, they are first described by Aristotle. In the few cases in which they are recognized in recent philosophy, they are dismissed as based on a wholly inadequate idea (Stump 2003; King 2007; Crane 2003; Shields 2007). It is interesting to note that modern cognitive scientists receive a similar criticism from some philosophers.. Thus, in his review of (Ramsey 2007), Hutto remarks (Hutto 2011):

> when [researchers in cognitive science] talk of representations these researchers really lack a proper understanding of their own theoretical commitments. They are misled and confused because of their attachment to a certain way of talking; a certain way of characterizing what goes on in the brain.

As a principle of exegesis, one needs to consider that when a large number of very distinguished philosophers and scientists over a spread of thousands of years are accused of employing a hopeless notion, it may be the accusers who have got it wrong.

Aristotelian representations play an essential role in a different model of the relation between cognizers and their environment. Fodorian representations are about the environment, while Aristotelian representations lack aboutness. Rather, Aristotelian representations represent by being copies, samples or examples; they are instantiations, exemplifications or manifestations. Were we practitioners of Aristotelian or Thomistic philosophy, we might be able to see immediately that these all can apply together, if not always entirely colloquially. Suppose, then, that we are trying to find something to represent the exact shade of mauve we want for a curtain. A color chip at the paint store might be a sample of just that color by instantiating that color; in more learned terms, the example exemplifies the color and manifests it. These are simply alternative ways of capturing the Thomistic notion that a particular accidental form is realized in

the sample. In more modern terms, we could say that the chip has the property of being that

particular mauve. 'Copy' is not discussed yet, and it may be that that terms carries the extra idea

that there is more than one example, along with a suggestion about how it was produced or the

intentions with which it was produced. Still, two samples exactly alike in a paint store would be

taken to be copies each of the other. And we do think of copying done on a cellular level

without literally ascribing intentions to the cells.

     A quick look a google result on the web reveals that 'example' and 'sample' are both

sometimes defined in terms of representing and being representative. Perhaps these associations

explain the use of 'represent' in ordinary English and in cognitive neuroscience that is strikingly

different from that in the usage of much contemporary philosophy. In any case, more formal

scientific speech can received the same sort of account. We can see the difference between the

Fodorian and Aristotelian senses of "represent" with the following comments in (Acebes et al.

2011):

> Synapse loss correlates with cognitive decline in aging and most
>
> neurological pathologies. Sensory perception changes often
>
> represent subtle dysfunctions that precede the onset of a
>
> neurodegenerative disease.

Special cases aside, Fodorian representations refer to things other than themselves. Accordingly,

on the Fodorian reading the sensory perception changes are not the subtle dysfunctions; they

rather refer to them. On the Aristotelian reading, the sensory perception changes are part of the

subtle dysfunctions. The Aristotelian reading also seems the better one; it is very hard to

understand how we could sensibly take the standard sorts of philosophical theories of

representation to this sort of case. For example, it simply seems wrong, as far as we know, to

say that sensory degeneration has the function of indicating the fact a further dysfunction in the system. The teleological accounts will have trouble getting a foothold. 'Represent' here is like the 'represent' in "This paper represents her work at its best." That is, the paper is an instance of her best work; it is not about her best work, except perhaps when we a case of a scholar's paper that is about a past paper of hers that was one of her best.

To take another example, in their abstract, (Bach, Bayliss, and Tipper 2011) remark:

An important question for the study of social interactions is how

the motor actions of others are represented. Research has

demonstrated that simply watching someone perform an action

activates a similar motor representation in oneself.

As the examples in Chapters Three, Nine and Ten make clear, a motor representation of X's action is the action plan for that action realized in the observer. Motor representations partially prepare one for action; they are not about true or false of someone else's action.   Given two action plans are the same kind of plan if they share enough relevant properties, we can say then that X's action plan (as a type) is instantiated in Y's brain and so manifested in it. X's brain realizes that action plan. Y's brain, like X's, has an example of that action plan. If you want to see a sample of what X is doing, you could ask Y to do it; similarly, his brain has a sample of the action plan.

To summarize: A color sample might represent the color I want for a curtain, or a brave act might represent one's better character even though the color sample is not about me or the curtain, and the act is not about its agent. To take an example from Chapter Two below, the soul may be an example of, and so represent, one kind of substance. But though the soul is an example of a substance, on such philosophies, it is not about substances. Though Fodorian

representations need not be tied to a language of thought theory, they get situated in a core of notions that are linguistic. Thus, mental representations, have meaning, content, satisfaction condition, reference and so on. Aristotelian representations do not have these sorts of properties.

Let me stress that I will be following Chemero in trying to generalize about all theorists in recent philosophy who have any definite commitment to what representations are. Chemero and I roughly agree that they all hold that representations have content. Some of these philosophers may be hesitant to use "content," but nonetheless the representations have aboutness and satisfaction conditions. That is, they are about some bit of the world and they either fit it or fail to fit it. Thus, for example, we might understand "action-oriented representations" as Clark does (Clark 1997), in which case they have a declarative component and an imperative one. These representations say something like "This X is before you; use it." In such a case, the representations have two satisfaction conditions; they need to fit the world, but the world also needs to fit them, if both conditions are satisfied. If there is an X here before you, the first component is satisfied; if you use it, then you have fit the second component.

A philosopher may, such as Wheeler (Wheeler 2005, 2008), want to focus quite entirely on the action part of the representation. The representation still appears to have aboutness and satisfaction conditions. Rather like a shopping list, the action-oriented representation will only be satisfied by the environment's fitting it. Another alternative, which may also fit Wheeler and so other philosphers, would be to say that the philosopher's task is to specify simply how the representation fits into the phenomenology. That is, the representations are simply action-originating reactions and that they are so is not to be understood in the semantical terms of aboutness and satisfaction conditions. In this case, the theorist does not meet the condition I agree with Chemero on. We are not given an explanatory account of what representations are.

Someone proposing that Aristotelian representations can be used to explain cognition will be very receptive to this last way of introducing action-orienting representations, but Aristotelian representations have always been part of explaining how we have the phenomenology that we do. Hence, they are part of an explanatory picture, and not simply a description of the phenomenon.

It is also the case that when we fit action-oriented descriptions with human neurophysiology, they may have been incorrectly described, as we saw in the introduction. That is, our best theory now suggests that when we have a perception that seems to initiate an action, what is really going on is a boost of dopamine that enhances the perception and initiates the action. Thus the perceptual reaction and the action are effects of a joint cause.

3.   Language and Thought

Though the original developer of the language of thought hypothesis is Jerry Fodor (Fodor 1975), many factors before him helped put it in place. One of them, falling outside philosophy, is the theory of computation developed by Turing and realized in digital computers. This gives us a general picture of processing symbols according to algorithms, a picture that morphs easily into one of mental processes as computations over semantically contentful representations. And that provides a major piece of Fodor's language of thought hypothesis.

Other factors have helped. Among them are Davidson's work on action and belief explanation, and Chisholm's interpretation of Brentano (Brentano, Kraus, and McAlister 1973). Davidson's work legitimized causal accounts of beliefs and action more solidly than had been done before. Chisholm's interpretation introduced a conception of intentionality widely invoked

today. One unfortunate consequence of the interpretation is that a model of the mind under development since Aristotle effectively was denied entry into contemporary philosophical thought. In our terms, Brentano's 'Aristotelian Representations' became understood as 'Fodorian representations.' (We will see in Chapter Four that Brentano's own understanding may have been more mixed than is relevant here.) The resulting conception of the mind's cognitive relation to its environment is very different from that developed from Aristotle through Aquinas to Brentano. However, because of the misreading, Brentano's Thomistic use of the term "intentional" is now understood in a sense that makes it apt for defining Fodorian representations, even by those supposedly in his much older tradition, such as (Gilbert and Lennon 2005). That is, for Aquinas, outer features realized inwardly then have as such an "intentional" presence. The sense of "intentional" in use today is typically explained in terms of the semantical ideas of reference or extension, or what a term refers to.

4.   Varieties of Representations

Readers need to be aware that the phrase "Aristotelian representation" can be a cause for confusion. A sign of the influence of Fodor's theories is that many philosophers today appear to think that representations by definition possess content or aboutness. One consequence is that Aristotelian representations can seem to be impossible by definition, since they do not possess content or aboutness. Given this problem, why call them "representations" at all? In fact, there are a number of reasons for doing so. First of all, that is what they are called in cognitive neuroscience. As we will see in Chapter Three and some following, there are clear cases where the Chisholmian account fails to fit the use in work in cognitive neuroscience. Secondly, the Aristotelian sense of representation is historically prior; for example, the Latin "representatio" in

Aquinas is Aristotelian and not Chisholmian. Finally, there are plenty of clear uses of "represent" and "representation" in Twenty-First Century English that are Aristotelian. While in recent philosophy "represent" as a somewhat technical term is understood in terms of content and satisfaction conditions, there is another use of "represent" with a quite different meaning. For example, we might ask whether certain medical procedures represent good value for money; if they do, then they are good value for money (Mihalopoulos et al. 2011). Morphologic and molecular data may suggest members of an investigated group represent a new family within a particular subclass (Liu et al. 2011). Epidermal Growth Factor Receptor (EGFR) may represent a biomarker for a certain disease (Cronin et al. 2011). In the latter two cases, the members are members of a new family and EGFR is a biomarker. Representing here is not to be understood in terms of intentional content and aboutness. And a philosopher may remark, "Hands and souls are parts of substances, although they represent different sorts of parts. A hand is an integral part…" (Stump 2003). That is, the hand is an example of a different sort of part; the hand is not about parts of the body, nor is it true or false.

There are a number of further questions about emotions that we will raise. Three particularly bear on the embodied cognition debate. The first concerns the vehicle of content; this issue is treated abstractly in Chapter Three, and it is given a more detailed treatment in Chapter Nine. The second is the idea of 'representation-hungry' representations that embodied cognition theorists something suggest must be Fodorian representations. I will argue that the premises for that conclusion are defective. The third question concerns whether representations can be decoupled from what they represent (Gallagher 2008; Wheeler 2008).

5. Looking Ahead

As we will see in the Chapters Two through Six, there is a lot of evidence that the representations of cognitive neuroscience are often Aristotelian, not Fodorian. If that is right, then we will have to revise a number of interpretations of the results of cognitive neuroscientific research. In explicating results in terms of Fodorian representations, philosophers seriously misunderstand them. Revising our analysis will affect our understanding of a wide range of mental phenomenon, including thinking, perceiving, feeling and acting. For example, the language-based model has led over two generations of theorists to think that belief and desire attributions are grounded in an agent's possession of contentful inner states. This is, I will be arguing, a very important if subtle mistake. It is the language-based model of the brain that has led us to think that neural activity and its causes give rise to semantic content. But, as we will see, a different model will push us toward locating the origin of semantic content in public language.

If we are to avoid starting with a language-based conception of the mind, we will have the interesting task of developing an account of Aristotelian representations. Two, Three and Six and Seven will be largely involved in this undertaking. Chapters Six and Seven are about the role of Aristotelian representations in conceptualized thought, while Chapters Two and Three provide a more basic introduction to them.

We will posit internalized language use for some cognition, but we will continue to reject the order of explanation Fodor posits. While, following Fodor, many philosophers take thought content to determine public content, we will understand meaning to originate in public language. The picture we will build will draw on work found in (Murphy 2002; Barsalou 2003; Barsalou 2008; Barsalou 2008; Barsalou 1999; Barsalou, Kyle Simmons, et al. 2003; Barsalou,

Niedenthal, et al. 2003; Dehaene 1997). These are issues developed in Chapters Six and Seven, where the discussion will be set in a context of a dialectic developing between Prinz and Machery (Prinz 2002; Machery 2009, 2010).

Chapters Four and Five are principally concerned with historical issues about Aristotelian representations. Hume figures prominently in this account, because he was the first philosopher to consider the dual component model we will use in Chapter Seven. That is, the account of human thought combines impressions and ideas that lack syntactic structure with the use of public language that provides it. It also helps us to appreciate Aristotelian representations if we see how different some of Hume's arguments look when understood in terms of them.

Chapters Eight through Ten will look at how human thought is a matter of coordination between our brains and our bodies in their environment. Among other things, the idea of flourishing and so functioning well is the target of investigation for a science of the mind. That this is so will, among other things, sharply divide a standard philosophical investigation of cognitive functioning from a cognitive neuroscientific one. Philosophy is interested in the division between, for example, seeing and not-seeing. Cognitive neuroscience, on the contrary, is concerned with what creates the distinction between seeing well and not seeing well. This distinction will make all the difference when we look at how to ground an embodied account of cognitive functioning. Conversely, an account that attempts to distill out what is common to all cases of seeing will end up with a very exiguous picture indeed. It is this very thin account that underlies the arguments for the internalism of cognitive states.

Philosophers influenced by Heidegger, Husserl or Merleau-Ponty might add a third model of the mind's cognitive relation to its environment; people cognize because they interact with the environment. It is important to see that this element is actually consistent with the

account being developed here. The idea that cognizing is action is at least accommodated by the

view that a science of cognition aims to explain functioning well in a niche. The fact that

functioning well is the basic explanadum does give a kind of foundational role to action.

Nonetheless, functioning well is a more general notion, and it allows us to connect action with

the broader categories of the maintenance and replication of life.

In addition, the Montague-Quartz quote we started with proposes we can see as a two-

fold revision in the older and impoverished view of cognitive neuroscience. The revisions will

be to both the explanans and the explanadum of traditional approaches. The cognitive activity to

be explained encompasses the actions and functions that are needed for flourishing. And the

causal account of these will have to include the many alignments and calibrations that make such

actions and functions possible.   For example, the brain does not just receive input; it receives

filtered and coordinated input, much of which serves the organisms efficient functioning. We

can then achieve – in a way that to our conscious minds is effortless – a unified perception of an

object, signals from which have taken varying times to travel from the sense organs to

associative areas of the brain.

Once we have completed the foundations presented in the first seven chapters, we will

pick up selected topics in the philosophy of perception, action and emotion. A series of articles

and books have attacked the idea that vision is essentially connected to bodily functioning; a

central target is Alva Noë (Noë 2004; Block 2005; Prinz 2006; Adams and Aizawa 2008; Rupert

2009). The account developed here provides a new and sustained response to those views hostile

to Noe. An account of seeing in terms of niche flourishing should make less compelling the

philosophical idea that only the brain is essential to seeing.

In Chapter Ten, we will look at the so-called folk psychological theory of action. It has proved enormously tempting to take the logical form of action explanation as a good guide to the causal ontology. Whether or not it is safe to read ontology off of logical form for most causal statements, it is certainly risky in the case of the mind and action. Among other things, it is questionable whether action explanations are causal statements. In addition, it is far from clear that the ontology proposed does in fact meet the folks' conceptions. Finally, cognitive neuroscience is uncovering very serious challenges to the philosophers' conception of action; it is, I shall argue, much less clear the challenge is really to the folks' thought.

There are also important issues about the reality of emotions, the topic of Chapter Nine. Is there some way in which they can be seen as a natural kind, or should they rather be regarded as more like artifactual kinds or social constructions? The distinction between Aristotelian and Fodorian representations becomes important here. As we investigate the relation between human emotions and emotions in near evolutionary ancestors, we see how to describe the immense difference language makes to human emotions. One response to an argument that emotions are not natural kinds (Griffiths 2004) says that the Aristotelian-representation part of emotions provides a natural kind that is homologous to the emotions of other animals. Language, however, can always introduce factors that take us away from the natural biological base. Griffiths, it should be noted, has since stop advocating his first argument regarding emotions and natural kinds in favor of a second (Griffiths and Scarantino 2009), but arguably the defense I offer against the first version also applies to the second.

6. In conclusion

This chapter introduces two very large themes. Both are Aristotelian, but both are important to understanding recent cognitive neuroscience. The first concerns explaining organisms, very much including ourselves, in terms of their aptness for flourishing in their environment. As this book progresses, perception, action and emotions should often look different when they are seen from this perspective.

The second addresses the question of how much cognition is in our brains. Clearly, our cognitive lives are made possibly by a vast amount of neural activity. Nonetheless, that does not means that the activity is one of computations over Fodorian representations, even when interpreted as minimally as they are here. Content may be less present in the brain than most theorists are inclined to say these days. Such a result is, following arguments should at least suggest, compatible with cognitive neuroscience. It definitely is not compatible with much in today's philosophy of mind.

The fact that philosophy of mind and cognitive neuroscience are in disagreement on such a basic topic makes the following investigation worth pursuing. It is appropriate to reiterate at this point an earlier warning: Given how "representation" has come to be used in recent philosophy, Aristotelian representations are going to look logically impossible. That should not, surely, be the case.

Bibliography

Acebes, A., A. Martin-Pena, V. Chevalier, and A. Ferrus. 2011. Synapse Loss in Olfactory Local Interneurons Modifies Perception. *Journal of Neuroscience* 31 (8):2734-2745.

Adams, Frederick, and Kenneth Aizawa. 2008. *The bounds of cognition*. Malden, MA: Blackwell Pub.

Bach, P., A. P. Bayliss, and S. P. Tipper. 2011. The predictive mirror: interactions of mirror and affordance processes during action observation. *Psychonomic Bulletin & Review* 18 (1):171-176.

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. 2008. Language and simulation in conceptual processing. In *Symbols, embodiment, and meaning*, edited by A. M. G. In M. De Vega, & A. C. Graesser. Oxford, UK: Oxford University Press.

Barsalou, Lawrence W. 1999. Perceptual symbol systems. *Behavioral & Brain Sciences* 22 (4):577.

Barsalou, Lawrence W. 2008. Grounded Cognition. *Annual Review of Psychology* 59 (1):617.

Barsalou, Lawrence W. . 2003. Situated simulation in the human conceptual system. *Language & Cognitive Processes* 18 (5/6):513-562.

Barsalou, Lawrence W., W. Kyle Simmons, Aron K. Barbey, and Christine D. Wilson. 2003. Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences* 7 (2):84.

Barsalou, Lawrence W., Paula M. Niedenthal, Aron K. Barbey, Jennifer A. Ruppert, and H. Ross Brian. 2003. Social Embodiment. In *Psychology of Learning and Motivation*: Academic Press.

Block, Ned. 2005. Action in Perception. *Journal of Philosophy* 102 (5):259-272.

Chabris, Christopher F., and Daniel J. Simons. 2010. *The invisible gorilla : thinking clearly in a world of illusions*. New York: Crown Publishers.

Chemero, Anthony. 2009. *Radical embodied cognitive science*. Cambridge, Mass.: MIT Press.

Clark, Andy. 1997. *Being there : putting brain, body, and world together again*. Cambridge, Mass.: MIT Press.

Crane, Tim. 2003. *The mechanical mind : a philosophical introduction to minds, machines, and mental representation*. 2nd ed. London ; New York: Routledge.

Cronin, James, Elizabeth McAdam, Antonios Danikas, Chris Tselepis, Paul Griffiths, John Baxter, Linzi Thomas, James Manson, and Gareth Jenkins. 2011. Epidermal Growth Factor Receptor (EGFR) Is Overexpressed in High-Grade Dysplasia and Adenocarcinoma of the Esophagus and May Represent a Biomarker of Histological Progression in Barrett's Esophagus (BE). *American Journal of Gastroenterology* 106 (1):46-56.

Dehaene, Stanislas. 1997. *The number sense : how the mind creates mathematics*. New York: Oxford University Press.

Dreyfus, Hubert L., and Mark A. Wrathall. 2006. *A companion to phenomenology and existentialism*, *Blackwell companions to philosophy*. Malden, MA: Blackwell Pub.

Foot, Philippa. 2001. *Natural goodness*. Oxford: Clarendon.

Gallagher, Shaun. 2005. *How the body shapes the mind*. Oxford ; New York: Clarendon Press.

Gallagher, Shaun. 2008. Are Minimal Representations Still Representations? *International Journal of Philosophical Studies* 16 (3):351-369.

Gilbert, Paul, and Kathleen Lennon. 2005. *The world, the flesh and the subject : continental themes in philosophy of mind and body*. Edinburgh: Edinburgh University Press.

Griffiths, Paul E. 2004. Is Emotion a Natural Kind? In *Thinking about feeling : contemporary philosophers on emotions*, edited by R. Soloman: Oxford University Press.

Griffiths, Paul E, and Andrea Scarantino. 2009. Emotions in the Wild. In *Cambridge Handbook of Situated Cognition*, edited by P. Robbins and M. Aydede. Cambridge, UK: Cambridge University Press.

Hutto, Daniel D. 2008. *Folk psychological narratives : the sociocultural basis of understanding reasons*. Cambridge, Mass.: MIT Press.

Hutto, Daniel D. 2011. Representation Reconsidered. *Philosophical Psychology* 24 (1):135-139.

Kim, Jaegwon. 1996. *Philosophy of mind*, *Dimensions of philosophy series*. Boulder, Col.: Westview Press.

Kim, Jaegwon. 2005. *Physicalism, or something near enough*, *Princeton monographs in philosophy*. Princeton, N.J.: Princeton University Press.

Kim, Jaegwon. 2006. *Philosophy of mind*. 2nd ed. [Boulder, Colo.]: Westview Press.

Kim, Jaegwon. 2007. Causation and Mental Causation. In *Contemporary debates in philosophy of mind*, edited by B. P. McLaughlin and J. D. Cohen. Malden, MA: Blackwell Pub.

King, Peter. 2007. Rethinking Representation in the Middle Ages: A Vade-Mecum to Mediaeval Theories of Mental Representation. In *Representation and objects of thought in medieval philosophy*, edited by H. Lagerlund. Aldershot ; Burlington, VT: Ashgate.

Liu, Weiwei, Zhenzhen Yi, Xiaofeng Lin, and Khaled A. S. Al-Rasheid. 2011. Morphologic and Molecular Data Suggest that Lynnella semiglobulosa n. g., n. sp. Represents a New Family within the Subclass Choreotrichia (Ciliophora, Spirotrichea). *Journal of Eukaryotic Microbiology* 58 (1):43-49.

Machery, Edouard. 2009. *Doing without concepts*. Oxford ; New York: Oxford University Press.

Machery, Edouard. 2010. Précis of Doing without Concepts. *Behavioral & Brain Sciences* 33 (2/3):195-244.

Mihalopoulos, Cathrine, Theo Vos, Jane Pirkis, Filip Smit, and Rob Carter. 2011. Do indicated preventive interventions for depression represent good value for money? *Australian & New Zealand Journal of Psychiatry* 45 (1):36-44.

Montague, P. Read , and Steven R. Quartz. 1999. Computational Approaches To Neural Reward And Development. *Mental Retardation And Developmental Disabilities Research Reviews* 5:86–99.

Montague, Read. 2007. *Your brain is (almost) perfect: How we make decisions*. New York: Penguin Group.

Murphy, Gregory L. 2002. *The big book of concepts*. Cambridge, Mass.: MIT Press.

Noë, Alva. 2004. *Action in perception*, *Representation and mind*. Cambridge, Mass.: MIT Press.

Noë, Alva. 2009. *Out of our heads : why you are not your brain, and other lessons from the biology of consciousness*. 1st ed. New York: Hill and Wang.

Prinz, Jesse. 2006. Putting the Brakes on Enactive Perception. *Psyche* (12).

Prinz, Jesse J. 2002. *Furnishing the mind : concepts and their perceptual basis*. Cambridge, Mass.: MIT Press.

Ramsey, William. 2007. *Representation reconsidered*. Cambridge, UK New York: Cambridge University Press.

Rangel, Antonio, Colin F. Camerer, and P. Read Montague. 2008. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience* 9 (7):545-556.

Rowlands, Mark. 2010. *The new science of the mind : from extended mind to embodied phenomenology*. Cambridge, Mass.: MIT Press.

Rupert, Robert D. 2009. *Cognitive systems and the extended mind*. Oxford ; New York: Oxford University Press.

Shields, Christopher John. 2007. *Aristotle*, *Routledge philosophers*. London ; New York: Routledge.

Shuler, Marshall G., and Mark F. Bear. 2006. Reward Timing in the Primary Visual Cortex. *Science* 311 (5767):1606-1609.

Stump, Eleonore. 2003. *Aquinas*, *Arguments of the philosophers*. London ; New York: Routledge.

Thompson, Michael. 2008. *Life and action : elementary structures of practice and practical thought*. Cambridge, Mass.: Harvard University Press.

Wheeler, Michael. 2005. *Reconstructing the cognitive world : the next step*. Cambridge, Mass.: MIT Press.

Wheeler, Michael. 2008. Minimal Representing: A Response to Gallagher. *International Journal of Philosophical Studies* 16 (3):371-376.

Wittgenstein, Ludwig. 1963. *Philosophical investigations*. English text reprinted. ed. Oxford: Blackwell.