Understanding is Simulating: a Defence of Embodied Linguistic Comprehension

Chris Letheby

Discipline of Philosophy

School of Humanities

The University of Adelaide

Submitted for the degree of Master of Philosophy

April 2013

CONTENTS

ABSTRACT

A topic of debate in current cognitive science is the nature of language understanding. One traditional view holds that we understand expressions of a natural language by translating them into an inner, abstract, symbolic Language of Thought. Recently, however, an increasingly plausible alternative has been proposed: that we understand natural languages by means of sensorimotor simulations of real-world objects and situations. This view is known as Embodied Linguistic Comprehension (ELC).

Much evidence has been found for ELC in such disciplines as psychology, linguistics, and neuroscience. However, the position faces several serious challenges. One is accounting for our comprehension of abstract terms, and other terms which refer to things beyond our own sensory experience. Other challenges include the productive and systematic nature of human thought, and difficult questions about how to interpret the relevant evidence.

This thesis is an exposition and defence of ELC. I review a representative sample of empirical data and major theoretical proposals, and then respond to objections. I argue that ELC is well-equipped to meet the challenges mentioned above. In particular, it has rich resources with which to account for abstraction, reference beyond a comprehender's own experience, productivity, and systematicity.

Responding to a recent challenge by proponents of a radical, anti-representational 'enactivist' theory of comprehension, I argue that ELC outperforms the enactivist view in accounting for the flexible and context-sensitive nature of language comprehension, and that rejecting mental representation is a costly and unnecessary step.

Perhaps the biggest challenge facing ELC at this point comes from powerful arguments purporting to show that the existing evidence is, at best, neutral between ELC and its rivals. I argue that, while the available evidence cannot rule out the existence of an abstract Language of Thought, we nonetheless have good reason to believe that sensorimotor simulation is a genuine constituent of all or most instances of comprehension, preserving the central point of the ELC proposal.

Declaration

For a thesis that does not contain work already in the public domain

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library catalogue and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Signature   Date....29/4/2013....

ACKNOWLEDGEMENTS

Thanks to Jon Opie and Matt Nestor for helping me find my way around a ground for thinking about the mind; to Andrew Gleeson for metaphilosophical challenges; and to John Mercier for early philosophical encouragement.

Thanks also to Denise Gamble for help and support throughout my candidature, and to Benjamin Bergen, Tyler Marghetis and colleagues for a copy of—and permission to cite—their forthcoming paper.

In the economy of understanding, words are merely money
JOHN HAUGELAND

1. Is Understanding Simulating?

Inquiring into Linguistic Comprehension

How do we understand words and sentences? This is an important and interesting question about language. In the context of cognitive science and the philosophy thereof, it amounts to a request for an explanation or theoretical account of a ubiquitous cognitive phenomenon: what are the mechanisms underlying linguistic comprehension?

To get a clear sense of what is at issue, consider the following. If you are a monolingual English speaker, something happens when you read the sentence 'the cat sat on the mat' which does not happen when you read the sentence 'l'uccello seduto nel nido'[1]. In each case you visually perceive symbolic forms. In the former case, your brain also—somehow—turns these symbols into meaningful thoughts[2]; in the latter case, it does not. My central question concerns what your brain does in the former, but not the latter, case—and how.

One possible answer is that your brain activates representations in its sensorimotor systems; among other things, representations of cats and of mats. Perhaps it even constructs a complex representation, or inner *simulation,* of a cat sitting on a mat. The details of this proposal clearly want development. But the basic claim that linguistic comprehension is fundamentally a process of accessing perceptual and sensorimotor representations to simulate denoted objects and described situations is one possible, and increasingly plausible, answer to the question about the mechanisms of comprehension. The idea is roughly that, when you

---

[1] Italian: 'the bird sat in the nest'.

[2] There is a risk of begging certain questions by couching the issue in terms of the brain turning symbols into meaningful thoughts, as will become apparent in chapter 6 when I discuss the enactivist theory of language comprehension. The least prejudicial way of putting the question is, perhaps: what important change or event happens to you *qua* cognitive being in the former, but not the latter, case—and how.

understand the word ‘cat’, you do so by undergoing some mental process (conscious or otherwise) which is similar in important respects to what happens when you perceive a cat. Following Daniel Weiskopf (2010) I will refer to this claim as *embodied linguistic comprehension* (ELC.)

My purpose in this thesis is to defend ELC. This thesis, therefore, offers an affirmative answer to the question: *is understanding simulating?* Unsurprisingly, however—given that this is first and foremost a work of philosophy—determining what those key terms *understanding* and *simulating* might amount to will prove to be no small matter.

ELC in Context

To understand ELC, it is important to understand the broad intellectual framework from which it has emerged. ELC has been proposed and developed by theorists working in a research programme in cognitive science known as ‘embodied cognition’. This research programme aims to establish the thoroughgoing dependence of higher cognitive functions generally—not just linguistic comprehension, but also conceptual reasoning and abstract thought—on the brain’s sensorimotor systems. Confusingly enough, however, the embodied cognition research programme is itself but one facet of a far broader and more eclectic movement within cognitive science which is *also* often referred to simply as ‘embodied cognition’. Clearly, this is a case of distinctions demanding to be drawn.

In chapter 2, below, I will undertake the required clarification, delineation, and scene-setting. The broader movement known as ‘embodied cognition’ also goes by other names, including ‘4E Cognition’--because its proponents, variously, claim that cognition is Embodied, Embedded, Extended, or Enactive (Menary 2010.) This is the name I will adopt for the movement, for two reasons: first, because it more accurately reflects the movement’s diverse nature, and, second, because it provides terminological clarity by distinguishing the broader movement from its narrower instance which is especially relevant to my project.

Having described how 4E Cognition arose as a critique of classical cognitive science and artificial intelligence, and having canvassed a few of its many manifestations, I will focus in more detail on one of those manifestations—the embodied cognition (henceforth EC) *research programme*. This is a sustained and somewhat concerted multidisciplinary endeavour to establish the claim that higher cognitive functions systematically exploit evolutionarily old and basic systems of perception and motor control. Compelling evidence is adduced from psychology, linguistics, neuroscience, and other disciplines to argue that abilities such as categorical inference, conceptual classification, abstract thought, and (of course) language comprehension are all subserved by sensorimotor resources.

Understanding the EC project of grounding cognition in perception and action is crucially important to the articulation and defence of ELC. Although it is helpful to be aware of the broader context constituted by the 4E Cognition movement, this context will rarely be directly relevant to the assessment of ELC itself. EC, however, is much closer to home. This is because ELC is basically just a special case of EC. That is to say, if we establish that higher cognitive functions in general are parasitic on perceptual and sensorimotor systems, then it seems likely that language comprehension will come along for the ride; it seems likely, in other words, that ELC will turn out to be true. So any evidence for the truth of EC, even if it does not directly concern language comprehension, constitutes at least indirect *prima facie* support for ELC. Similarly, any in-principle or *a priori* arguments concerning the truth or likelihood of EC will be relevant to ELC, too.

Defending ELC: How to Proceed

Once ELC has been situated in the broad context of 4E Cognition and the narrow context of the EC cognition-grounding project, in chapter 3 I will examine ELC itself in more detail. I will describe the two main theoretical developments of the claim: Arthur Glenberg's Indexical Hypothesis and Rolf Zwaan's Immersed Experiencer Framework. While differing significantly in detail and emphasis, both these views agree that language comprehension consists in the activation of sensorimotor representations to create inner simulations.

In chapter 3, I will also describe Lawrence Barsalou's influential theory of Perceptual Symbol Systems (PSS), an EC account of concepts which provides valuable resources for the defender of ELC. Barsalou develops the crucial notion of simulation in detailed and important ways, as well as addressing many classic objections to perception-based theories of conceptual knowledge. In the course of this tour through the two main versions of ELC and one of their chief sources of theoretical support, I will also describe some of the fascinating empirical evidence which has been gathered in support of EC and ELC.

In chapter 4, I will compare and contrast ELC with its main rival in cognitive science: the traditional *amodal* account of linguistic comprehension. This view—closely related to Jerry Fodor's (1976) Language Of Thought (LOT) hypothesis—claims that linguistic comprehension does not essentially involve sensorimotor representations, though it may have incidental causal connections with them (Adams 2010.) On this view, comprehension is fundamentally a process of manipulating abstract symbols in a separate, dedicated language faculty or conceptual system. These amodal symbols—'amodal' because they are separate and distinct in kind from representations in the sensorimotor modalities—are held to be the bearers of semantic content and the source of comprehension (Weiskopf 2010.)

My focus in chapter 4 will be the central motivations which have traditionally been given for adopting the amodal view. Many of these concern striking cognitive and linguistic phenomena such as productivity, systematicity, propositional content, abstraction, and the like, which have commonly been held achievable only by amodal means. However, Barsalou and others argue that these phenomena can, in fact, be accounted for by a modal theory—that is, by an EC or ELC approach (Barsalou 1999; Prinz 2002.) A central guiding question of chapter 4, therefore, will be: does the amodal view really have the advantages over ELC that are often claimed for it?

In chapter 5, I will address the converse to this question: Does ELC really have the advantages over the amodal view that are often claimed for it? This breaks down into two main parts. One of the key motivations for adopting ELC is the wealth of empirical data showing intimate and systematic connections between linguistic and conceptual cognition on one hand, and sensorimotor processing on the other. However, defenders of the amodal view respond that these data can all be accounted for within their framework, and that the data therefore do not give us reason to reject the amodal view in favour of ELC. I will examine this debate to determine what, if anything, can legitimately be inferred from the empirical data. One line of thought holds that the data establish merely causal connections between linguistic/conceptual cognition and sensorimotor processing (Adams 2010; Mahon & Caramazza 2008.) This issue goes deep, and I will revisit it in chapter 7.

The second advantage that ELC is alleged to have over the amodal view is that of solving the infamous Symbol Grounding Problem (SGP.) The SGP, at a very rough first pass, concerns how representations in a system can come to refer to things outside of the system. This has been seen as a problem for abstract symbol accounts like the amodal view because the representations they posit bear no intrinsic connection to their referents. Advocates of embodied approaches argue that they fare better in this respect because they are able to ground reference in sensorimotor representations, which are sometimes assumed to stand in non-arbitrary relations to the things they represent. Be that as it may, perhaps the biggest obstacle resolving the debate over the SGP is that of determining what exactly the problem is supposed to be—so I will devote a good deal of chapter 5 to such clarification.

After I have investigated whether or not the amodal view really has its alleged advantages over ELC, and vice versa, the main issue remaining concerns a quite different theory of language comprehension. The core ELC claim that comprehension consists in sensorimotor simulation has recently been challenged by advocates of enactivism, a radical approach to cognition which rejects the existence—or at least explanatory importance—of mental representations, and identifies cognitive processes with capacities for skilful environmental exploration and action (Van Elk et al 2010.) Chapter 6 consists of a detailed comparison and contrast of the 'cognitivist' or simulation version of ELC and its enactivist rival. It includes an exploration of the two main arguments which have been given for preferring the enactivist theory of comprehension, as well as a discussion of what seems to be a serious problem for the enactivist theory itself.

Conclusions to be Drawn

In chapter 7, I will take stock of the foregoing chapters, revisit the issue of empirical evidence for ELC, suggest some directions for future research, and set out the minimally qualified and specific version of ELC which seems defensible in light of my discussion.

My central conclusions will be as follows. First, that the most plausible and broadly applicable version of ELC is one which holds that episodes of simulation in Barsalou's technical sense are constituents of—even if not exhaustively constitutive of—all instances of language comprehension. Second, that comprehension is a graded phenomenon (one which admits of degrees) and recognition of this fact makes ELC much more defensible. Third, that the amodal view enjoys no advantage over ELC with respect to our comprehension of language about abstracta and things we have not experienced, because ELC has ample resources with which to account for this. Fourth, that the amodal view enjoys no advantage over ELC with respect to the productivity and systematicity of thought, because Barsalou's theory of PSS furnishes ELC with ample resources with which to account for these phenomena. Fifth, that ELC enjoys no advantage over the amodal view with respect to the SGP, properly understood, because any resources which are available to ELC for the solving of this problem are equally available to the amodal view. Sixth, that enactivist ELC enjoys no advantage over cognitivist ELC with respect to the comprehension of language about abstracta and things we have not experienced, because cognitivist ELC has ample resources with which to account for this. Seventh, that cognitivist ELC enjoys a significant advantage over enactivist ELC with respect to the flexible and context-sensitive nature of language comprehension, due to cognitivist ELC's emphasis on the dynamic and variable nature of sensorimotor simulation and enactivist ELC's antipathy towards mental representations as explanatory posits. Eighth, that—despite initial appearances—it is presently problematic to take the extant empirical evidence as support for ELC over the amodal view. Ninth, and finally, that friends of ELC need not be dismayed by my eighth conclusion, because there are several promising ways in which the case for ELC can be developed and strengthened. These include a parsimony argument and an alliance with the resemblance theory of mental representation.

In summary, I will conclude that cognitivist ELC, in the attenuated form I will develop, is a viable theory of language comprehension, even if much work remains to be done. Indeed, understanding *is* simulating—even if neither understanding nor simulating turns out to be precisely what we may have thought.

2. Embodied Linguistic Comprehension and Embodied Cognition

4E Cognition and the Critique of Classical Cognitive Science

Here is a story that has been told, more or less similarly, a number of times. Cognitive science and artificial intelligence got going in the mid-20th century as a consequence of the 'cognitive revolution', a paradigm shift in the study of the mind. For the first half of the century, scientific inquiry into human intelligence had been dominated by behaviourism, an approach which eschewed all mention of inner states or mechanisms. Motivated by the perceived unreliability of introspective methods and a desire to put psychology on a rigorous, objective, properly scientific footing, the behaviourists tried to account for all human intelligence in terms of conditioning: the implanting by learning of tendencies to respond in certain ways to certain kinds of stimuli. According to this vision, a complete psychology would make no reference to such seemingly intangible and unobservable mental events as beliefs, desires, experiences, or sensations, but only to patterns of stimulation and response (Gardner 1985.)

At least three events were crucial to the demise of behaviourism and consequent embrace of a new approach to the study of intelligence. These were: Alan Turing's (1936) formal analysis of the notion of computation; the development of the electronic digital computer; and a growing awareness within psychology and related disciplines of the inherent limitations of the behaviourist programme itself, articulated most famously in Noam Chomsky's (1959) devastating criticism of the behaviourist approach to language.

Turing's seminal work was originally motivated by investigations in logic which had little directly to do with either cognition or engineering. However, in the course of attempting to solve a theoretical problem concerning first-order predicate logic, he gave a thought experiment incorporating the specifications for an abstract 'machine' which turned out to form the basis of computer science (Clark 2001.) The Turing Machine, as it has become known, is a deceptively simple system consisting of three elements: a tape divided into squares; a read/write head capable of scanning a symbol from the current square, writing and/or erasing a symbol on the current square, and moving left or right along the tape; and a central processor which determines the next action based on the current input and the state set by the prior input (Turing 1936.)

Turing's purpose in describing the Machine was to give a precise, formal analysis of the intuitive notion of computability. On the intuitive notion, a mathematical function is computable just in case it can be worked out in a finite series of well-specified steps none of which requires any insight or creativity. Turing's famous thesis is that the class of functions computable in this sense is identical to the class of functions which can be solved by a Turing Machine, and hence that the definition of the Machine amounts to an exact specification of the conditions under which a function is computable (Copeland 1993.)

The relevance of this development to the study of the mind is that, prior to Turing's innovation, computing functions was an activity exclusively performed by human intellects. The description of a purely mechanical computing machine amounted to nothing less than a recipe for building a physical system apparently able to do something that could previously only be done by a mind. As Turing put it: "The idea behind digital computers may be explained by saying that these machines are intended to carry out any operations which could be done by a human computer" (1950, p.436.)

In particular, Turing's work suggested the possibility of constructing a symbol processing machine whose transformations from one state to another would be entirely determined by the syntactic properties—that is, the physical form or shape—of the symbols it trafficked in, but would nonetheless respect or preserve the semantic properties of those symbols. This insight into how to get meaning-respecting behaviour from a mere mechanism was pregnant with the thought that perhaps the meaning-respecting human mind just is such a mechanism—one built, of course, out of neurons, synapses, and the like (Clark 2001.)

In the wake of the remarkable theoretical work done by Turing (and others; e.g. Church 1936) came the construction of the first digital computers. The initial insight that the computation of mathematical functions could be mechanized and automated was one thing. However, matters were taken to a whole new level by the further insight that the symbols manipulated by a Turing-style device need not stand only for mathematical objects, but could potentially represent anything at all. This led to the construction of general-purpose, programmable digital computers: artefacts capable of solving indefinitely many well-specified problems by means of internal formal operations performed upon symbolic structures. In retrospect, it seems that these revolutionary machines were conceived of, designed, and created in perfect time for the position of scientifically-respectable-model-of-mind to be left vacant by the ignominious dismissal of behaviourism. By the middle of the century, Turing himself was defending the claim that there need be no contradiction or confusion involved in the idea of a machine that thinks (Turing 1950.)

This third development was a gradual process, but is often thought to be encapsulated by Chomsky's (1959) review of B.F. Skinner's (1957) book *Verbal Behavior*, the latter being a comprehensive application of the behaviourist framework to the specific phenomenon of linguistic ability. In brief, Chomsky argued—among other things—that the sophisticated nature of human language mastery far outstrips what can be accounted for solely in terms of environmental factors and simple principles of conditioning. To adequately explain verbal behaviour, he contended, it is necessary to posit rich and complex inner resources of exactly the kind eschewed by a behaviourist treatment.

The implications of such a claim for the study of intelligence are best expressed by Chomsky himself in his review of Skinner:

> The magnitude of the failure of [Skinner's] attempt to account for verbal behavior serves as a kind of measure of the importance of the factors omitted from consideration, and an indication of how little is really known about this remarkably complex phenomenon. (1959, p. 28.)

The "factors omitted", of course, were precisely the details of the cognitive resources whereby human organisms managed to achieve such sophisticated performance on the basis of the arguably meagre inputs available to a language learner. Chomsky's argument can be seen as a plea for the scientific study of intelligence to re-direct its attention to such factors, on pain of inability to yield any genuine insight into its subject matter.

In effect, then, the backlash against behaviourism breathed new life into the question: what is the structure and nature of the human mind such that it can do the remarkable things it can do? And at the same time, an undeniably appealing answer was suggested by the general purpose digital computer, a mechanical device capable of solving problems in many different domains by means of purely syntactically driven but nonetheless semantically sensible symbol manipulations. The answer, in Allen Newell and Herbert Simon's famous formulation, was that "a physical symbol system has the necessary and sufficient means for general intelligent action" (Newell and Simon 1976.) This conjecture has become known as the Symbol System Hypothesis (SSH; Copeland 1993.)

So, the question about the structure and nature of the mind was legitimized, and the SSH proposed as an answer. The beginnings in the 1950s and 60s of a research programme based directly on these two developments constituted the twin birth of (classical) cognitive science and artificial intelligence. These were theoretical and practical sides of a single coin: the effort to understand the human mind on the model of a digital computer or physical symbol system. This multifarious but unified research programme has since become known as Good

Old-Fashioned Artificial Intelligence, or GOFAI (Haugeland 1985); sometimes simply as 'classicism' (Horst 2011.) For convenience, I will use these two terms interchangeably.

For some time, this endeavour enjoyed considerable success and was accompanied by a sense of optimism. Computer programs were written which seemed to exhibit impressive signs of intelligent performance on certain tasks, such as problem solving (Simon & Newell 1958), playing games such as checkers (Samuel 1959), medical diagnosis (Shortliffe & Buchanan 1975), and natural language understanding (Winograd 1972.) But ultimately, a number of severe criticisms and seemingly intractable problems accumulated for the GOFAI programme.

One of these difficulties was the 'Frame Problem', which connects with deep issues concerning relevance. The problem is that of how a cognitive system, updating its inner model or representation of the world in real-time, knows what to check or update. Suppose a system, which harbours a complex internal model of its environment, performs some action or perceives some event. Since actions and events have effects, the system will need to update various parts of its model. For a GOFAI system, this means comparing various information-encoding symbol strings to the current state of the environment, and revising those symbol strings which no longer correspond to the environment.

But such a system cannot, for practical reasons, update every symbol string in its model every time some change occurs: that would be far too time-consuming and inefficient. Therefore, the system needs some method to determine, for any change that might occur in its environment, which parts of its model to update in response and which parts to leave alone, without considering every potential candidate. The problem of devising such a method is the Frame Problem, and it has been—along with other issues concerning relevance and knowledge organization—a key source of scepticism about GOFAI. Consider: the modus operandi of GOFAI systems is rule-governed symbol manipulation, so, it seems, any solution to the Frame Problem for these systems must be implemented as a set of rules codifying which kinds of symbol strings should be updated in the wake of which kinds of changes. But any such set of rules will certainly be subject to many qualifications and exceptions, necessitating a further set of higher-order rules regarding when the first-order rules do and do not apply; and so forth (Copeland 1993.)

Another influential criticism of GOFAI centred on the question of meaning, or semantics, and was brought to prominence by Searle's (1980) infamous Chinese Room Argument. (This is related—though in exactly what way is highly debatable—to the Symbol Grounding Problem, to be discussed in chapter 5.) Without delving into the details of this controversy, the problem, in a nutshell, is this: How do the symbols manipulated by a GOFAI system refer to things in the world? The question arises in part because these symbols, like the symbols of a natural language, are arbitrarily related to their referents. For example, in English, the symbol 'cat' could just as well have represented barking, canine animals, and the symbol string 'dog' could just as well have represented meowing, feline animals. There is nothing about either symbol which makes it especially well suited to denoting what it in fact denotes. In just this sense are the symbols manipulated by GOFAI systems also arbitrary (Haugeland 1985.)

In the case of public language, a plausible answer to the question is readily available: convention. The symbols of a language like English represent what they do because of tacit agreements, or intentions, or expectations, in the minds of those using them. But this answer cannot work for symbols in the mind, on pain of infinite regress. Of course, in the case of the actual GOFAI programs that were written and executed, it was clear that the assignment of symbol to referent stemmed from a decision on the part of the programmer (Winograd 1980.) But it was equally clear that this could not be the answer for the human mind, or any generally intelligent, genuinely cognitive system. Thus, the question stood open: if human minds, and minds more generally, are physical symbol systems—if the SSH is true—then how do their symbols mean anything? (Haugeland 1985; Shapiro 2007.)

The problems of relevance and knowledge illustrated by the Frame Problem, and the question of meaning, were among many criticisms of the SSH and the GOFAI programme. Others concerned the apparently significant differences between the workings of a physical symbol system and those of the human brain; others, the tendency of GOFAI research to focus on evolutionarily recent achievements such as rational thought and language use, while ignorant of how the older and more fundamental problems of perception and motor control had been solved by nature (Anderson 2003.) Meanwhile, writers such as Hubert Dreyfus (1972) argued on philosophical grounds that the SSH relied on dubious assumptions about the very nature of thought. Andy Clark articulates the general sentiment well:

> One possibility is that we simply misconstrued the nature of intelligence itself. We imagined mind as a kind of logical reasoning device coupled with a store of explicit data... In so doing, we ignored the fact that minds evolved to make things happen. We ignored the fact that the biological mind is, first and foremost, an organ for controlling the biological body... Minds are *not* disembodied logical reasoning devices. (1997, p.1)

Suffice it to say that serious scepticism about GOFAI and the SSH was in the air by the 1980s. Many researchers began to suggest that the classical programme was somehow incomplete, misguided, or fundamentally flawed and that a new paradigm— or at least significant change— in the study of the mind was required (e.g. Haugeland 1979; Winograd 1980; Searle 1980; Norman 1980; Lakoff and Johnson 1980a.) This ultimately led to the raising, throughout the late 1980s and the 1990s, of many different revolutionary banners in cognitive science and AI. Several proposals were made seeking to supplant GOFAI as the dominant approach to the scientific investigation of mental phenomena.

The first prominent aspect of this methodological explosion was the advent of connectionism, heralded as a neurally-inspired and (hence) more biologically plausible way of studying cognition, and one which might succeed where GOFAI had failed (Tienson 1987.) Instead of rule-based symbol systems, connectionists used (simulated) artificial neural networks to model cognitive tasks. These deceptively simple networks consist of units (or 'nodes') and connections between those units along which activation values are passed (Bechtel 1988.)

A typical connectionist network might have three 'layers' of units: an input layer, an output layer, and a 'hidden' layer. Suppose the network's task is to convert text to speech. In this case, it will need to receive textual information as input, and yield phonetic information as output. The textual information will be encoded in the input layer by means of activation values: each unit in the input layer will take some value, either binary or on a continuum. Next, the input units will transmit activation values to units in the hidden layer via their connections. Each input layer unit is connected to various hidden layer units, and each connection has a certain 'weight'. The activation value that a given hidden layer unit takes will be determined by its own threshold function plus the total activation it receives—and the latter will be determined by the activation values of the input layer units, plus the strengths or weights of the connections. In short, this process will be repeated from the hidden layer to the output layer, and, ultimately, if the network is successful, the resultant pattern of activation values at the output layer will encode the phonetic information corresponding to the textual input (Sejnowski & Rosenberg 1987.)

The foregoing is a very rough sketch, but it gives a sense of the first serious competitor to GOFAI in cognitive science. On the face of it, connectionism offers a very different vision of cognition to the SSH. Jerry Fodor once famously claimed that the symbolic model of the mind was "the only game in town" (1976, p.55) but by the late 1980s, it was apparent that this was no longer true. The rise of connectionism, however, was only the beginning. In its wake followed a huge variety of new positive proposals about how to gain traction on the vexing problems of intelligent thought and behaviour.

The publication of *The Embodied Mind* by biologist Francisco Varela, philosopher Evan Thompson, and psychologist Eleanor Rosch in 1991 may be said to mark the beginning of embodied cognition in the broadest sense—or, as I am calling it, 4E Cognition—as a self-conscious, revolutionary movement. This landmark work articulated a number of the key criticisms of the classical or GOFAI project, and drew on such diverse sources as Buddhist meditation, biology, and the phenomenological tradition in philosophy to propose a new paradigm: "enactive" cognitive science, which would emphasize the continuity between life and mind, and study cognition as first and foremost a biological phenomenon (Varela et al 1991.)

On the enactive view (a recent version of which I will discuss in chapter 6) cognition is identified with neither rule-governed symbol manipulation nor the parallel distributed processing of connectionist networks, but with "a history of structural coupling" between organism and environment (Varela et al 1991, p.206.) *The Embodied Mind* amounted to nothing less than a manifesto, urging cognitive science to rethink its foundations in light of a perceived gulf between computational theories of mind and the world of lived human experience (Varela et al 1991.)

The enactive proposal was but one of many. Robotics researcher Rodney Brooks (1991a) launched an influential assault on the traditional conception and role of mental representation in cognitive science and AI. To understand Brooks' challenge, first consider that the modus operandi of GOFAI systems, in general, is to take as input a description of some task or problem, convert that input into a symbolic representation of the problem, perform computational operations on the representation, and yield as output a solution to the problem. In the context of mobile robotics, this translates into what Brooks (1991b) calls the *sense-*

model-plan-act (SMPA) framework: the robot receives information about its environment via sensors, uses this information to construct a symbolic model or representation of the environment, uses the representation in turn to generate a plan of action, and then executes the plan.

Drawing attention to the typically poor performance of SMPA robots in navigating their environments in real time, Brooks proposed, and implemented, a new approach, which was—put simply—to do away with the middle two steps of the SMPA cycle. He and his colleagues designed and built robots which did not generate detailed internal models of their environment for use in planning, but which instead relied on direct couplings between sensing and acting. These minimal robots, it turned out, performed very well with such economical resources. One central moral Brooks drew from this was that explicit, detailed symbolic representation is an unnecessary and even counterproductive design feature when trying to get simple creatures to exhibit basic intelligent behaviour. He further argued that, since 'higher' cognitive activities like language use and explicit deliberation are the tiny tip of a gigantic evolutionary iceberg, the right way for AI and cognitive science to proceed is in a bottom-up fashion: First, we need to determine what is required for embodied beings to sense and act quickly and effectively in a changing environment, and then we can try to figure out how language and the other higher faculties might emerge from such fundamental capacities (Brooks 1991a.)

In one of his seminal papers, Brooks tells a delightful fable which is worth quoting in full, because it vividly articulates one common diagnosis of the problem with GOFAI:

> Suppose it is the 1890s. Artificial flight is the glamor subject in science, engineering, and venture capital circles. A bunch of AF researchers are miraculously transported by a time machine to the 1980s for a few hours. They spend the whole time in the passenger cabin of a commercial passenger Boeing 747 on a medium duration flight.

> Returned to the 1890s they feel vigorated, knowing that AF is possible on a grand scale. They immediately set to work duplicating what they have seen. They make great progress in designing pitched seats, double pane windows, and know that if only they can figure out those weird "plastics" they will have the grail within their grasp. (A few connectionists amongst them caught a glimpse of an engine with its cover off and they are preoccupied with inspirations from that experience.) (1991a, p.141.)

According to this way of thinking, GOFAI was led astray by its focus on symbol manipulation, a focus that amounted to seeking the essence of cognition in the inessential. Certainly, goes the thought, the design of digital computers was inspired by certain mental operations performed by human beings; but what sorts of mental operations? Symbol manipulation, as in explicit mathematical calculation and logical reasoning—which is a recent innovation, evolutionarily speaking, and perhaps marginal relative to the kinds of skilful environmental coping with which human and other brains were solely occupied for millions of years (Brooks 1991a.)

So, connectionism, enactivism, and Brooks' bottom-up, representation-lite, embodied-and-situated approach to intelligence each presented a serious alternative to GOFAI. Another such was the Dynamical Hypothesis—championed by Tim Van Gelder (1995; 1998)—that cognitive systems, fundamentally, are not computational but dynamical, and ought therefore to be studied using the mathematical tools of Dynamical Systems Theory. Meanwhile, articulating a growing appreciation of the ways in which cognitive creatures are "embedded" (Haugeland 1998) in their environments, Andy Clark and David Chalmers (1998) famously argued for the metaphysical thesis that the mind itself was not confined to the head. Under certain conditions, they claimed, objects external to a biological organism could literally be constituents of that organism's mental processes.

Much more could be said, but suffice it to note that, by the dawn of the 21st century, the eclectic and disorganised nature of the revolutionary movement led many to seek conceptual clarity. Did this vast proliferation of diverse methodological proposals and research programmes really amount to a cohesive movement at all, they wondered? Was it united by any positive claims or theses, or merely by the negative critique and rejection of GOFAI? And even if some of its positive claims were true, did they really have the paradigm-shifting implications claimed for them? (Wilson 2002; Shapiro 2007; Kiverstein & Clark 2009.)

Inspired by such musings, the label '4E Cognition' was coined to recognise the diversity of a 'movement' whose proponents claimed, variously, that the mind was Embodied, Extended, Enactive, and environmentally Embedded (Menary 2010.) Many commentators—including Margaret Wilson (2002), Michael Anderson (2003), and Larry Shapiro (2007)--have also suggested that the best way to engage with 4E Cognition is to identify the distinct positive claims being made by its proponents, and then consider each of these (insofar as possible) independently, asking what grounds we have for thinking they are true, and what follows if they are.

This seems like a sensible and fruitful way to proceed. In this spirit, then, I will for the most part ignore the Extended, Enactive, and Embedded strands of 4E Cognition research. My focus is on the claim that language comprehension is Embodied, in the sense of constitutively involving sensorimotor simulation. However, this claim has not arisen in isolation. It has been made by theorists working in a research programme which aims to establish that higher cognitive functions *in general* are embodied in this sense. This research programme, which has its home primarily in cognitive psychology and linguistics, I will call 'embodied cognition' (EC) and distinguish from the broader 4E Cognition movement of which it is part. Time, then, for a closer look at EC itself.

EC: Grounding (Higher) Cognition in Perception and Action

The diversity of the different proposals and research programmes which collectively comprise 4E Cognition is mirrored by a diversity of motivations. Although it is common for theorists in this area to emphasize one aspect or another of the critique of classicism outlined above, they differ in which aspects they emphasize and how they propose to address these concerns.

Insofar as the EC research programme is driven by negative or critical considerations, perhaps chief among them is the Symbol Grounding Problem. This, recall, is the question of how the symbols of a formal, rule-governed cognitive system can come to refer to things outside of the system; or, formulated differently, how a purely syntactic system can come to have semantic content. A common thought fuelling research into EC has been that establishing the thoroughgoing dependence of higher cognitive functions on sensorimotor systems would go a significant way toward solving this problem (Anderson 2003.) The suggestion seems to be that the structures used in abstract thought and reason, language use, and the like will be seen to have a non-arbitrary connection to their referents if it is established that they are essentially the same structures we use to represent those referents when experiencing them in direct, physical, embodied perception and interaction (Barsalou 1999.)

I will discuss the Symbol Grounding Problem in detail in chapter 5. Meanwhile, however, we can examine EC without worrying overmuch about whether it offers a solution to this problem or not. The key claim of EC is that (higher) cognition consists in embodied simulation, which essentially means the partial re-activation of sensorimotor states. As Margaret Wilson puts it:

> Many centralized, allegedly abstract cognitive activities may in fact make use of sensorimotor functions... Mental structures that originally evolved for perception or action appear to be co-opted and run “off-line”, decoupled from

> the physical inputs and outputs that were their original purpose, to assist in thinking and knowing. (2002, p. 633.)

Clearly, this is a significant and interesting claim in its own right. It is possible to consider the evidence for such a claim without the Symbol Grounding Problem or other critical considerations for motivation. Among others, linguist George Lakoff and philosopher Mark Johnson (1999) argue that this core EC claim—that abstract, conceptual, and off-line thought is subserved by mechanisms of on-line perception and action—is practically forced on us by convergent evidence from different areas of cognitive science, regardless of our prior theoretical interests or convictions.

At this point, it is natural to wonder what the evidence is which so impresses Lakoff and Johnson. One datum is the phenomenon of conceptual metaphor in linguistics. Conceptual metaphor, which has been extensively documented by Lakoff and Johnson (1980a, 1980b, 1999), refers to the existence of systematic structural mappings between abstract and concrete concepts. Studies in discourse analysis, etymology, and other subfields of linguistics reveal a pervasive tendency among human beings to instinctively deploy such metaphors as 'Purposes Are Destinations'. This metaphor maps the abstract notion of a purpose onto the concrete notion of a (spatial) destination, leading to such expressions as "He'll ultimately be successful, but he isn't *there* yet" (Lakoff and Johnson 1999.)

Lakoff and Johnson give detailed analyses of dozens of such metaphors, some of which are 'primary' or fundamental, and others of which are complex (built out of combinations of primary metaphors.) Examples include 'More Is Up' (as in "My income rose last year"); 'Understanding Is Seeing' (as in "I see what you're saying"); 'Ideas Are Money' (as in "He has a wealth of ideas"); and many more (Lakoff and Johnson 1980a.) Some of these metaphors—particularly primary ones—are found cross-culturally, while others are culture- or language-specific. Throughout their discussion, Lakoff and Johnson are at pains to emphasize their contention that conceptual metaphor is not merely about ways of speaking, but about ways of thinking. It is not mere terminology, but conceptual or inferential structure that is mapped from concrete to abstract domains. As Lakoff and Johnson put it when discussing the complex metaphor 'Love Is A Journey':

> [this] mapping does not just permit the use of travel *words* to speak of love. [It] allows forms of *reasoning* about travel to be used in reasoning about love. It functions so as to map inferences about travel into inferences about love... (1999, p.65.)

This means that, insofar as the abstract concept of a love relationship is structured on the concrete concept of a spatial journey, inferences which are valid concerning journeys will be mirrored by correspondingly valid inferences concerning love relationships. The idea is that our concrete concepts of objects, spatial properties, and the like are represented in our sensorimotor systems, and these structures are exapted for the purpose of abstract conceptual tasks. This, if true, is one way in which higher cognition might constitutively involve sensorimotor processing (Lakoff and Johnson 1999.)

The linguistic evidence for conceptual metaphor, then, is one reason for thinking cognition might be embodied in the EC sense. Some results from neuroscience have also been taken as supportive of the EC view. I will not discuss these in detail here. However, Gallese and Lakoff (2005) use findings concerning 'mirror neurons' and 'canonical neurons', as well as other neuroscientific data, to argue that the structure needed to process abstract conceptual inference is available in the brain's sensorimotor systems.

Much of the evidence for EC which will be directly relevant to ELC (recall: Embodied Linguistic Comprehension) comes from experimental studies in cognitive psychology. Systematic and intimate interaction effects have been shown to exist between conceptual and linguistic tasks, on the one hand, and sensorimotor variables, on the other.

For instance, in one study, subjects were presented with sentences on a screen describing various actions involving motion toward or away from their body. They then had to either move toward or away from their body to press a button indicating that they had read and understood the sentence. It turned out the subjects were significantly quicker to respond when the direction of motion implied by the sentence was the same as the direction in which they had to move to press the button. One possible explanation of this is that the process of comprehending sentences about motion actually engages motor control mechanisms in producing an inner simulation, thus interfering with subsequent motion production (Glenberg & Kaschak 2002.)

Many similar results have been found in other studies. To give just one more example, when subjects had first to read a sentence describing taking a particular orientation towards an object (e.g. *you are driving a car*) and then subsequently identify nouns presented to them as denoting parts of that object or not, they were quicker to respond correctly for object parts which would have been readily visible from the described perspective. Again, an obvious explanation is that their comprehension of the perspective-priming sentence involved constructing an inner perceptual simulation of taking the perspective in question (Borghi et al 2004.)

What I have just discussed is a brief but representative sample of the kinds of empirical evidence which the EC research programme has accrued. Meanwhile, work has progressed on theoretical and philosophical fronts, too. Psychologist Lawrence Barsalou (1999, 2003) has developed his theory of Perceptual Symbol Systems (which I will discuss in chapter 3), a detailed account of how concepts might be created, stored, and retrieved in the sensorimotor systems. Relatedly, philosopher Jesse Prinz (2002) defends Concept Empiricism, a position similar in spirit to the EC account, though differing in emphasis.

Armed with a basic understanding of the EC cognition-grounding project, and its place within the broader contexts of the 4E Cognition movement and cognitive science, it is time for a closer look at the particular instance of EC which is my concern: namely, the ELC project of grounding linguistic comprehension in sensorimotor simulation.

ELC as a Special Case of EC

As far as I am aware, two main versions of ELC have been articulated and defended in the cognitive science literature. These are the Indexical Hypothesis (IH) developed by Arthur Glenberg (Glenberg 1997; Glenberg & Robertson 1999) and the Immersed Experiencer Framework (IEF) due to Rolf Zwaan (2004.)

Although both these views claim that linguistic comprehension centrally involves perceptual simulation, they differ in detail. IH holds that linguistic comprehension is a three-stage process involving the mapping of words to corresponding (Barsalou-style) perceptual symbols, the subsequent derivation of affordances (roughly, perceived opportunities for organism-object interaction) from these perceptual symbols, and, finally, the creation of a simulated experience by combining object affordances as dictated by sentence structure (Glenberg & Kaschak 2002.) IEF, on the other hand, although relying on the notion of perceptual symbols, makes no explicit mention of affordances. IEF, like IH, understands comprehension as a three-stage process[3]; but its three stages are called 'activation', 'construal', and 'integration', and they do not correspond to the three stages identified by IH. In the first stage, various sensorimotor representations corresponding to a word or morpheme are activated; in the second stage, the representational activity is 'disambiguated' by context; and

[3] The term 'stage' is slightly misleading, as activation, integration, and construal are supposed to operate in an overlapping and recurrent fashion (Zwaan 2004.) It might be more accurate to call these *elements* of the comprehension process.

in the third stage, the disambiguated representations of objects and properties are combined to create a spatially and temporally coherent simulation (Zwaan 2004.)

The differences of detail between these versions of ELC are worth bearing in mind. My project is to see what the prospects are for defending the bare thesis of ELC: that language comprehension constitutively involves the activation of sensorimotor representations. I want to identify what—if any—is the most abstract, minimally qualified version of this claim which appears defensible in light of extant objections and considerations both theoretical and empirical. So if there is some specific claim made by IH or IEF which seems inessential to ELC itself, it will be possible to simply jettison that particular detail without needing to abandon the core position.

Of course, this does not mean that no further theoretical work will be required. The project just described amounts to identifying constraints on a plausible version of ELC. So, for example, if it turns out that most criticisms of ELC can be answered, but that its use of the notion of affordances is indefensible, then one constraint will be that the position is plausible provided some more defensible alternative to affordances can be recruited to do the same work or play the same explanatory role.

As well as emphasizing the relationship of these specific ELC theories to my more abstract investigation, I want to mention again the relationship of ELC to the broader EC project. In outlining their theories of language comprehension, people like Glenberg and Zwaan readily appeal to empirical findings and theoretical developments in the EC literature which point to the sensorimotor grounding of conceptual knowledge and abstract thought generally, not just language-specific processes. This seems like a reasonable strategy, granted only the crucial (and quite plausible) assumption that the mechanisms of linguistic comprehension and those of conceptual thought are intimately linked (Prinz 2002.) If understanding language involves having thoughts which feature the concepts corresponding to the words used, then certainly understanding language will involve (at least to some extent) whatever is involved in conceptual thought. I propose to accept this assumption, and, hence, its corollary: that direct evidence for EC constitutes indirect evidence for ELC. Every finding which lends support to the claim that sensorimotor processes are involved in conceptual cognition *ipso facto* lends support to the claim that such processes are involved in linguistic comprehension (though precisely how, and to what extent, remains to be seen.)

Conclusion

Now we know where we've come from, and have some idea where we're going. The eclectic movement now known as 4E Cognition arose from dissatisfaction with GOFAI and the classical project in cognitive science. The perceived inability of classicism to solve key problems led to a number of different research programmes, some more revolutionary than others. Paradigm-shifting implications aside, one of the most intriguing of these programmes is that which has come simply to be known as Embodied Cognition or EC, which seeks to establish that processes of higher, abstract and conceptual cognition take place by exploiting the brain's sensorimotor resources. The core of this approach can be loosely summarized as the claim that "thoughts comprise mental simulations of bodily experiences" (Casasanto 2009.)

Much empirical evidence has been gathered to support this core EC claim in such disciplines as linguistics, psychology, and neuroscience. A particular strand of EC research focuses on the claim that linguistic comprehension constitutively involves sensorimotor simulation. This claim is, of course, what I am calling ELC. The two main ELC theories are Glenberg's 'Indexical Hypothesis' (IH) and Zwaan's 'Immersed Experiencer Framework' (IEF), both of which appeal to the empirical and theoretical resources of the broader EC programme—particularly Barsalou's theory of conceptual knowledge known as Perceptual Symbol Systems (PSS.) I have suggested that this strategy in developing ELC is legitimate, given the plausible assumption that linguistic comprehension and conceptual thought are closely linked.

In the next chapter I will give a thorough exposition of the case for ELC. This will involve detailed descriptions of IH, IEF, and PSS, focussing in particular on points of comparison and contrast, as well as relationships of theoretical dependence. It will also involve a closer examination of the empirical evidence and other arguments cited in support of EC and ELC. It will emerge that there are many impressive and compelling reasons for concluding that understanding is simulating, before I move on to consider objections in subsequent chapters.

3. The Case for ELC

Perceptual Symbol Systems

As I have already remarked, ELC is situated firmly within the EC research programme in cognitive science. Probably the most thoroughly developed and oft-cited theoretical articulation of this programme is Lawrence Barsalou's (1999) theory of Perceptual Symbol Systems (PSS.) The two main versions of ELC which I will discuss below make liberal use of theoretical constructs from PSS, so an understanding of PSS is a prerequisite to a true understanding of the case for ELC. Let us, therefore, look more closely at PSS.

Barsalou's project is to give a schematic but comprehensive account of how such cognitive abilities as conceptual reasoning, categorical inference, planning, and (of course) language comprehension could be implemented in the brain's sensorimotor systems. He sets himself squarely in opposition to the classical view in cognitive science that such processes take place outside of the sensorimotor systems, in a separate, dedicated conceptual system using a representational code distinct from those deployed by the sensorimotor systems. Barsalou refers to this view in general as the *amodal view* of higher cognition, though its canonical articulation is Jerry Fodor's (1976) Language of Thought (LOT) hypothesis. To oversimplify slightly, then, PSS is offered as a rival hypothesis to LOT.

The widespread acceptance which LOT has enjoyed is largely due to influential theoretical arguments purporting to show that only such a view can account for certain striking features of human cognition, particularly those exhibited by the higher cognitive functions mentioned above. Those arguments will come to the fore in chapter 4, when I investigate the alleged advantages of the amodal view over ELC. For now, I mention them only to set the dialectical context for Barsalou's theory. His burden is to show that it is not necessary to postulate an amodal conceptual code, or LOT, by showing that sensorimotor resources alone can account for all the striking properties of higher cognition.

Barsalou begins this task by citing a well-established finding in psychology: during perception, attention can (and does) selectively fixate on distinct components of the perceptual scene—such as objects, shapes, properties, and relations—and those components on which it fixates are more likely to be stored in long-term memory from which they can later be retrieved. For instance, when visually perceiving a cat sitting on a mat, your attention might be drawn to the cat, and this will increase the likelihood that your visual representation of the cat will be stored and made available for subsequent use (Barsalou 1999.)

Suppose, then, that this has happened, and the following day you see a different cat—perhaps smaller, and of a different colour—sitting on a different mat. At this point, a match or similarity relation may be identified between your occurrent perception of today's cat and your stored perceptual representation of yesterday's cat. Despite the dissimilarities, they are classified as being instances of the same type. Now, due to the interaction of selective attention and long-term memory, your perception of today's cat is also stored for future use—at the same time probably modifying slightly your stored representation of yesterday's cat—and you now have a greater store of information which can be brought to bear on future cat-like perceptions. Moreover, if you find yourself wandering aimlessly around on the third day, hoping in vain to encounter another cat, it is possible that your stored representations of the first two cats could be re-activated, amounting to a simulated cat perception of sorts.

The foregoing is a simplistic sketch of Barsalou's account of concept formation. According to PSS, a concept consists of a mechanism called a *simulator*, the precise nature of which is not specified, but which serves to integrate information from various sensory and motor modalities derived from repeated experience of members of a certain category (of objects, properties, relations, etc.) Your concept of CAT, then, is a CAT simulator—and a CAT simulator is a mechanism which integrates stored visual, auditory, haptic, olfactory, gustatory, somatosensory and other kinds of information acquired through repeated encounters with cats (Barsalou 1999.)

Speaking generally, there are two main functions which a CAT simulator performs. One is on-line categorization and inference: if you have a sufficiently well-developed CAT simulator, then whenever you perceive a cat, a match will be established between your perception and the simulator, which constitutes a categorical judgement or classification to the effect that '*this* thing is a cat' (Barsalou 1999.) As a result of this classification, the wealth of information which your CAT simulator integrates will be made available to guide your actions in relation to the present cat; for instance, your simulator might recall stored memories of cats behaving pleasantly in response to pats or food, thus suggesting possible goal-conducive forms of behaviour in relation to this cat here.

The second function which a CAT simulator performs, and which is crucial to the aspirations of PSS to be a fully-fledged account of higher cognition, is that of underwriting off-line reasoning. If you are sitting in your armchair at home and wondering whether or not you should get a pet cat, your CAT simulator will be capable of recalling many sorts of cat-related experiences, including memories of what cats look, feel, and smell like, how they behave in relation to you, and how you typically feel in relation to them. This information will enable you to envisage or simulate (to a greater or lesser extent) what it would be like to have a cat, thus enabling your reasoning and decision-making processes in relation to the question (Barsalou 1999.)

According to PSS, then, concepts are simulators—mechanisms which integrate information from various modalities about some category of things, enabling on-line classification of and reasoning about members of that category, as well as off-line thought and planning about members of that category. This is the core of the theory. It is now time to fill in a few details.

One difference between PSS and the amodal view which Barsalou emphasizes is the fact that the representations posited by the former as the stuff of thought are both "modal and analogical" (Barsalou 1999 p. 578.) By "modal", he means that they are processed in systems whose function is representing information from the various sensory-motor modalities. By "analogical", he means that they stand (to some degree) in a relation of structural resemblance to the original modality-specific representations from which they were derived[4].

To get a sense of this difference, it helps to think about how the same thought might be represented in both systems. Returning to our earlier well-worn example, the thought that the cat is on the mat will be represented in the brain, according to the amodal view, by a symbol string in a linguiform representational code distinct from any codes used by the modalities. This string will probably contain symbols standing for CAT, ON, and MAT, combined in accordance with syntactic rules. (Recall the discussion in chapter 2 of GOFAI systems.)

According to PSS, that same thought will be represented in the brain by a schematic, simulated re-enactment of sensorimotor states typical of perceiving a cat on a mat. These will have been generated by mechanisms known as simulators, from stored traces of sensorimotor representations in long-term memory. Your thought that the cat is on the mat will be processed by the same systems that process your perceptions of cats on mats, and will stand in a relation of structural resemblance to those perceptions themselves.

At this point an objection along the following lines arises naturally: it doesn't feel like we think by simulating sensorimotor experience, at least not all the time. We are often not conscious of any such simulation; sometimes we seem to think in words, and sometimes we seem to think in wordless, imageless thought (Schwitzgebel 2008.)

---

[4] This seems to be the correct reading of Barsalou, at least. However, it is an interesting question exactly what various EC/ELC theorists mean when they talk about the 'analogical' or 'non-arbitrary' nature of sensorimotor simulations. This issue will come up again in chapters 5 and 7.

Barsalou's answer to this is twofold. First, there is his assumption that simulations or re-creations of sensorimotor states need not always be accompanied by conscious imagery in any of the modalities. The fact that we aren't conscious of imagining the smell of a rose, or its colour, or its texture, doesn't, on this view, show that schematic representations of such sensory experiences are not in fact being tokened in the various modality-specific systems (Barsalou 1999.)

Second, Barsalou makes a space in his theory for a kind of linguistic thought. However, he thinks that this, too, operates in the sensorimotor systems, using re-created perceptual information about words. The idea would be that your MAT, CAT, and ON simulators each possess, as well as traces of perceptions relating to their referents, traces of perceptions relating to the corresponding words. (Obviously this will not always be true, since we presumably have simulators for un-named categories.) So, for instance, your CAT simulator will integrate stored visual representations of the written word 'cat', auditory and motor representations of the spoken word 'cat', and so forth (Barsalou 1999.)

Speculatively, then, when we seem to be thinking in words, it may be that representations are being activated in multiple modalities, but are richer and more detailed in linguistic than non-linguistic information. Either way, Barsalou affords sensorimotor representations of words a role in governing the construction of multi-modal simulations. He thinks that our capacity to construct such complex simulations (i.e. thoughts) as we can piggybacks on the immense expressive and combinatorial power of natural language. He allows that a pre-linguistic perceptual symbol system would be capable of some productivity, but also holds that language helps a great deal. (Sensorimotor simulations of natural language will become very important in chapter 4.)

Another difference between PSS and its amodal rival which Barsalou emphasizes is the dynamic nature of the representations posited by the former. According to the amodal view, the same CAT symbol is tokened in each thought you have about cats (Zwaan et al 2002.) According to PSS, on the other hand, your CAT simulator—the mechanism Barsalou identifies as your CAT concept—changes, however minutely, with every new cat-encounter you have.

Another issue to mention in passing, which I will take up in more detail in chapter 4, is that of abstract concepts. Barsalou acknowledges that this is one of the main problems facing any sensorimotor or perceptual theory of higher cognition. Depending on what precisely one takes *abstract* to mean, the idea of a perceptual representation of an abstract concept might sound self-contradictory.

There are several possible understandings of the concept of abstraction. One intuitive sense is that involved in thinking about the idea of a cat in general, independently of the details of any particular cat, or thinking about a triangle in general, independently of the details of any particular kind of triangle. Barsalou, perhaps wisely, chooses to tackle the triangle first, and reasons as follows. If there are dedicated neural populations in the visual system coding for lines, vertices, angle, orientation, and the like, then it should be possible to token a representational state in the visual system which encodes the information about lines and vertices essential to triangularity in general, while remaining silent about any angles or orientation which would specify a particular kind or orientation of triangle (Barsalou 1999.)

This is one of the sensorimotor representations, of course, which Barsalou must think would be tokened without accompanying conscious imagery. The assumption that this can work is crucial for Barsalou's story. One of his chief selling points for PSS is its invulnerability to a key objection to perceptual theories of cognition: that they are committed to holding, implausibly, that every time we think about a category we must represent some particular instance of it or other. Simulation, for Barsalou, is a technical term, meaning the (more or less) schematic re-enactment of sensorimotor states under endogenous control, with or without conscious imagery (Barsalou 1999.)

Having described in some detail the theoretical tenets of PSS, it is now time to consider some of the empirical evidence cited in its favour. Of course, there are difficult questions about whether, and to what extent, various data actually support PSS against its amodal rival. However, I will defer such questions to chapter 5; for now, I will just recount the evidence and why it is supposed to favour PSS.

One piece of evidence which Barsalou cites is a finding from neuroscience: viz., that damage to modality-specific areas of the brain results in impaired conceptual processing of concepts whose exemplars are typically represented in that modality. For example, damage to the visual system results in problems performing conceptual tasks related to categories such as BIRD (Pulvermüller 1999.) Arguably, this should not be the case if conceptual cognition occurs in a system separate from the modalities, but it is exactly what we would expect if PSS is true (Barsalou 1999.)

A related neuroscientific result which Barsalou describes comes from neuroimaging studies of normal-brained subjects which show high levels of activity in modality-specific areas during conceptual tasks. For instance, visual areas are very active during tasks relating to such categories as animals, while motor and somatosensory areas are very active during tasks relating to such categories as tools (Pulvermüller 1999.) This is perhaps less decisive than the previous finding, but it seems *prima facie* to be an empirical prediction of PSS, and not of the amodal view, which is borne out (Barsalou 1999.)

Many findings from cognitive psychology have also been claimed to support PSS. The empirical literature demonstrating systematic interrelationships between conceptual tasks and sensory-motor variables has grown enormously in the last 15 years, and I cannot hope to cover all of it here. I will therefore settle for describing a representative sample of these fascinating findings.

One is that subjects who are asked specifically to perform conceptual tasks using mental imagery exhibit the same performance as those who are permitted to perform the same conceptual tasks using whatever method they select spontaneously. For instance, suppose two groups are asked, for a number of different concepts, to list features or properties which are true of that concept. One group is instructed to perform the task using mental imagery, while the other group is not instructed to use any particular method. When such experiments have been performed, it has been found that both groups give similar lists of features and perform similarly in other important respects (Barsalou et al 2003.) This result is predicted by PSS: if conceptual cognition operates in the normal case on sensorimotor simulation, then it is very similar to using mental imagery and should produce similar performance. On the other hand, if conceptual cognition is amodal in nature, then the two groups should show some interesting differences in performance.

Another finding is that subjects performing a conceptual property-verification task are slowed down by switching from one modality to another. For instance, suppose a subject has just verified that the auditory property *rustling* is (typically) true of the concept *leaves*. They will be quicker to respond on the next verification if it, too, is of an auditory property, such as *loud* being (typically) true of *blender*. On the other hand, they will be slower to respond if the next verification is of a non-auditory property, such as the gustatory property *tart* being (typically) true of the concept *cranberries*. An obvious explanation of this difference in reaction times is that the subjects' performance of this conceptual task is underwritten by simulations in their perceptual systems, and they are slowed down when they have to switch from one modality-specific system to another (Pecher et al 2003.)

I will describe more relevant empirical results in due course. In particular, studies pertaining to language comprehension will feature later in this chapter. However, it is clear that findings of the sort I have discussed, in conjunction with Barsalou's theoretical arguments, amount to a compelling case. PSS is a powerful and well worked-out theory of higher cognition and conceptual knowledge which cannot be ignored. Now that it is in place as a background framework, we can turn to the embodied theories of language comprehension which stand to it in a relationship of mutual support: Glenberg's Indexical Hypothesis (IH), and Zwaan's Immersed Experiencer Framework (IEF).

Before that, however, there is one more thing that needs to be mentioned in relation to PSS: namely, the term 'perceptual symbol' itself. Given that it constitutes most of the name of the theory, and is used liberally by theorists such as Glenberg and Zwaan, it is important to be clear what it means. This question is particularly confusing because some writers in the fields of cognitive science and artificial intelligence define the term 'symbol' in such a way that 'perceptual symbol' becomes almost oxymoronic. In this sense, a symbol is something which, by definition, stands in an arbitrary relationship to the thing it represents: no symbol is intrinsically better suited than any other to represent its particular referent. Symbols, on this construal, are primitive tokens in some computational system, the manipulation of which constitutes that system's computations, and which stand for their referents by virtue of convention, causation, or some other non-resemblance relationship. Such symbols are the stuff of thought, according to the amodal view (Harnad 1990.)

Since Barsalou is at pains to stress the analogical and non-arbitrary nature of the representations he posits as the stuff of thought, it is clear that he is not using the term 'symbol' in this strict sense. Rather, he is using it simply to mean a representation: any state of a system which stands in for something in the system's computations. Why, then, does he use the term 'symbol'? Largely for strategic reasons: he wants to emphasize that the restriction of the term to arbitrary, amodal, non-analogical tokens is unwarranted, because non-arbitrary, modal, analogical representations can play all the same roles, and do all the same work, which makes the term 'symbol' appropriate for their counterparts.

Very well, but what then are perceptual symbols? Well, they are not simulators, because although simulators are identified with concepts, they are not themselves the stuff of thought: rather, they are mechanisms with the capacity to generate various *representations* which are the stuff of thought[5]. Perceptual symbols are simulations: each of the indefinitely many more-or-less detailed cat simulations which your cat simulator is capable of generating is itself a perceptual symbol, a representation of a cat which stands in for that object in your mind's reasoning and cogitation. Perceptual symbols are the modality-specific representations of external stimuli which are stored by selective attention and re-activated in cognitive processes (Barsalou 1999.)

Time, now, to look at what role these symbols are claimed to play in understanding language, and how.

The Indexical Hypothesis

As well as Barsalou's theory of PSS, and a wealth of fascinating empirical evidence, Glenberg's IH is based in his broader views about the nature and purpose of cognition. These views intersect with what has become known as the study of situated cognition: that is, cognitive processes in the service of immediate, real-time action with the present environment (Wilson 2002.) Glenberg (1997) believes that the central purpose of higher cognitive processes such as memory, reason and language is, in short, action preparation. The function of cognition is to guide effective action.

It is important to note that this wider claim concerning the function of cognition is separable from ELC. The specific claim that ELC makes, that comprehension consists in sensorimotor simulation, might be true even if the wider teleological claim is false. If claims

[5] This is reminiscent of Cummins' (1996) insistence on the distinction between concepts, understood as knowledge structures, and the representations which comprise them.

about the action-guiding nature of cognition end up looking contentious, then that is not necessarily fatal to ELC: rather, it will amount to identifying constraints on a plausible version of the latter.

That aside, here is the IH story about comprehension. Comprehending a sentence, it says, is a three stage process: First, perceived words are indexed to perceptual symbols; second, affordances (to be defined below) are derived from those symbols; third, and finally, the derived affordances are combined or "meshed" to create a simulation of the state of affairs described by the sentence (Glenberg & Robertson 2000.)

Let us unpack these steps one at a time. The first step is the indexing of perceived words to perceptual symbols. This means that when you hear or read the sentence *the cat sat on the mat*, the first part of the comprehension process consists in the accessing or activation of sensorimotor simulations of, at least, CAT, MAT, SIT, and ON. The precise mechanism by which this occurs is not specified—a problem which besets much theoretical work in EC and ELC at the present time (Flusberg et al 2010; Pezzulo et al 2011.) However, presumably associative mechanisms combined with the perceptual encoding of words themselves play a substantial role. When you perceive the word 'cat', according to PSS, your recognition of it as that word and not another consists in a matching of your perception of the word to stored sensorimotor representations of the word. Assuming your stored representations of the word cat are integrated by a simulator mechanism with your stored representations of cats themselves, then your recognition and classification of the word will automatically tend to activate representations of its referents.

The second stage of the comprehension process consists of the derivation of affordances from the activated perceptual symbols. At a first gloss, the concept of an affordance—due originally to Gibson (1979)--is simply that of an opportunity for organism-object interaction. So chairs afford sitting to able-bodied humans, but not to elephants; tree tops afford eating to able-bodied giraffes, but not to humans; and so forth.

The existence and nature of affordances and our perceptions thereof has been a matter of some controversy. Gibson infamously claimed that we directly perceive affordances, with no inference or reasoning process required. Investigating this claim would take us too far afield. Happily, we can stay fairly neutral on this point. All that is required for Glenberg's analysis is that there genuinely are properties of objects which make them amenable to certain kinds of actions by certain kinds of animals, and not to others, and that we can perceive those properties and represent them in our sensorimotor systems. Whether our perception of these properties is conceived of as direct or indirect is neither here nor there.

That having been said, there are still two important points about affordances. One is that the term itself is ambiguous between properties of objects and our representations of those properties. Obviously, if we are talking about the cognitive system deriving affordances from perceptual symbols as part of the comprehension process, the latter sense is at play. The cognitive system is deriving representations of properties from representations of objects; it is certainly not deriving properties of objects from representations of objects (whatever that might mean.)

The final thing to be said is that if Glenberg’s analysis is to be plausible and general, there is a problem with relying on an egocentric notion of affordances. In order to comprehend the sentence 'the cat sat on the mat', what matters is representing that mats afford sitting *to cats*. Whether or not they afford sitting to me, the human language comprehender, is obviously beside the point. So, the affordances which we derive from perceptual symbols in comprehension must at least be understood to include any affordance of any action by any object to any creature. However, I think that even a non-egocentric notion of affordances is too narrow to do the required work. It is not affordances—that is, opportunities for organism-object interaction—which need to be derived from perceptual symbols. Rather, it is any opportunities for interaction between two objects, animate or otherwise. Justifying this claim takes us into a discussion of the third and final stage of the comprehension process posited by the IH: the combining, or meshing, of derived affordances to create a sensorimotor simulation of the described situation.

Glenberg and his colleagues would most likely not accept the modification I propose. The emphasis of the IH as they develop it is on comprehension as action preparation. Glenberg holds that "the world is perceived in terms of its potential for interaction with an individual's body" (1997, p.4) and that "the meaning of [a] situation consists of the set of actions available to the animal in the situation" (Glenberg & Kaschak 2002, p.558.) It is clear, therefore, that he is talking about affordances in the traditional sense of opportunities for action by an organism.

Consider, though, how the IH is applied to analysing results from empirical studies. Several experiments (e.g. Glenberg & Robertson 2000; Glenberg & Kaschak 2002; Glenberg et al 2008) have used a "sensibility judgement" paradigm (Shapiro 2011, p.106), in which a central task for subjects is to decide whether or not a given sentence is sensible, and make some prescribed response. Most of the sentences used in these studies describe actions, concrete or otherwise; e.g. the following pair:

(1a) After wading barefoot in the lake, Erik used his shirt to dry his feet.

(1b) After wading barefoot in the lake, Erik used his glasses to dry his feet.

(Glenberg & Robertson 2000, p. 384.)

The IH explanation of people's tendency to judge 1a, but not 1b, sensible goes as follows: The affordances of the objects in 1a, but not of the objects in 1b, can be combined or meshed into a coherent simulation. Shirts afford drying to wet humans; glasses do not. So what has occurred with the non-sensible sentence 1b is a failure in the third and final step of the comprehension process: meshing derived affordances into a coherent simulation (Glenberg & Robertson 2000.)

As soon as we go beyond explicitly action-describing language, however, this kind of analysis in terms of affordances will no longer work. Consider the following pair of sentences:

(2a) The coffee dripped into the cup.

(2b) The book dripped into the cup.

It is easy to believe that most people would judge 2a sensible and 2b not sensible, and this fact can be given a ready explanation which conforms to the structure of the comprehension process posited by the IH—but only if the notion of affordances is replaced, as I am suggesting, by the notion of *opportunities for interaction between two objects, animate or otherwise*. This is because the difference between coffee and books which makes one, but not the other, a plausible candidate for dripping into a cup has nothing to do with how people or other animals can interact with coffee and books, and everything to do with how cups can interact with coffee and books. The same point applies to our old friend *the cat sat on the mat*, and innumerable other simple sentences describing concrete situations which do not involve human action.

At this point, Glenberg and colleagues could still insist, on independent grounds, that affordances are important to the use of embodied simulations to guide action planning. Maybe after a coherent simulation of the coffee dripping into the cup is created, affordances are derived to generate an awareness of possible actions in relation to this situation. However, what is clear is that affordances cannot account for the important differences between our processing of sensible and non-sensible sentences where action is not explicitly involved. If the described situation does not consist of an action, then the situation itself must be coherently simulated before any action can be planned in relation to it.

The revision I am proposing is consistent with some of the central and distinctive emphases of the IH. Perhaps more than most EC theorists, Glenberg is explicitly concerned with the idea of simulations as analogical mental models, and the idea that what is meaningful, or comprehensible, or able to be simulated is determined by fairly literal relations between the shapes or structures of objects and actions as analogically represented in the sensorimotor systems (Glenberg 1997.) The idea of mental models, or situation models, will come up again in chapter 4, and the idea of analogical representation in chapter 7. For now, however, suffice it to note that the IH in outline looks like a potentially widely applicable model of comprehension—so long as step 3 involves the creation of a simulated model by meshing a broader class of structural relations between analogical representations than just affordances.

This leaves us with the following story of the comprehension process: words are indexed to perceptual symbols; structural relations are derived from those perceptual symbols; and the structural relations are combined to create a simulation.

One thing to note is that syntax is quite important to the final stage of the process, as is made obvious by a moment's reflection on the different simulations which would underwrite comprehension of the following two sentences:

[A] *Andy delivered the pizza to you*.

[B] *You delivered the pizza to Andy*.

(Glenberg & Kaschak 2002, p.560.)

The first two stages of comprehension—accessing perceptual symbols and deriving structural relations from them—would be exactly the same in the two cases, so clearly the difference is in the final stage, when these elements are combined, and clearly syntax must be doing a lot of work.

An objection may arise at this point that syntactic analysis cannot consist of sensorimotor simulation, and so the IH—or ELC more generally—cannot be the whole story about comprehension. There are two things to say in response to this. The first is that, despite the misleading oversimplification of the slogan ‘understanding *is* simulating’, ELC is not committed to holding that the comprehension process is exhausted by sensorimotor simulation. Rather, at least in the form I am defending, ELC claims that comprehension constitutively involves such simulation. This is consistent with it involving many other non-simulation processes too. So even if there are completely non-sensorimotor syntactic processing mechanisms, this is not a problem for ELC.

The second thing to say is that there are attempts in progress to account for syntactic processing in terms of sensorimotor simulation. In particular, the linguistic formalism known as Embodied Construction Grammar (Bergen & Chang 2003) attempts to ground grammatical constructions in sensorimotor schemas. I will not review this work here. For the purposes of this thesis—especially in light of the previous paragraph—I am content to treat the nature of syntactic processing mechanisms as a black box, although the effects of their outputs will be discussed somewhat below in the context of the IEF.

Now we have a basic understanding of IH as an account of the comprehension process. It claims, fundamentally, that understanding language consists of indexing words to perceptual symbols, deriving affordances (or structural relations) from those symbols, and meshing those affordances (or structural relations) to create a simulation of the described situation. It is also associated with the claim that the primary function of higher cognition is action preparation. For this reason, Glenberg and his collaborators have done extensive work focusing on the role of the motor system in comprehension.

One result they have found, much discussed in the literature on EC and ELC, is known as the Action-sentence Compatibility Effect (ACE). In brief, this is that merely understanding a sentence which implies motion of an object in one direction facilitates performing actual motion in the same direction, and interferes with performing action in the opposite direction. That is, actions are more quickly performed if they are *compatible* with (the same as) the direction of the motion described or implied by the sentence (Glenberg & Kaschak 2002.)

Subjects in the relevant study were asked to read a series of sentences, one at a time, on a computer screen, and decide whether or not the sentences made sense. They then had to press a 'yes' or 'no' button to indicate their answer. Some of the sentences presented implied motion towards the subjects' bodies, such as sentence [A] above, or *open the drawer*. Other sentences implied motion away from the subjects' bodies, such as sentence [B] above, or *close the drawer*. Some neutral sentences did not imply motion at all, while some nonsense sentences did not describe any coherent or imaginable state of affairs.

The action-related variable comes from the position of the 'yes' and 'no' buttons which the subjects had to press. Each subject would begin each trial with their hand resting in a position in between their 'yes' and 'no' buttons. Some subjects were in the 'yes-is-near' condition, in which their 'yes' button was nearer to their body than their hand, and their 'no' button further. This meant, of course, that they had to move towards their body to press 'yes' and away from their body to press 'no'.

Here is the striking result of the study: subjects in the 'yes-is-near' condition were consistently significantly quicker to press 'yes' for sensible sentences describing motion towards their bodies than for sensible sentences describing motion away from their bodies. The opposite result was found for subjects in the 'yes-is-far' condition (the nature of which I trust is obvious.) The theoretical inference which Glenberg and Kaschak draw is that comprehending the sensible sentences engaged the subjects' motor systems in simulating movement in the described direction, facilitating motion production in that direction—because the relevant motor programs were already active—and interfering with motion production in the opposite direction—because switching motor programs was required (Glenberg & Kaschak 2002.)

Several other studies conducted by Glenberg and his collaborators have found results congenial to the IH. For example, one study aimed to test competing predictions of the IH and two amodal accounts of meaning and comprehension, with respect to the role of affordances

(Glenberg & Robertson 2000.) Each of the amodal accounts in question—Latent Semantic Analysis (LSA; Landauer and Dumais 1997) and Hyperspace Analogue to Language (HAL; Burgess and Lund 1997)--endorses some version of the claim that the meanings of words and sentences are mathematically calculable functions of the co-occurrences of words with other words. There is no role given to sensorimotor representations of words' referents, nor to affordances derived from such representations.

Glenberg and Robertson's strategy was to find pairs of sentences so similar that they were assigned very close meaning values by LSA and HAL, but which differed significantly in the affordances they involved. One example of such a pair is the following:

[1] As a substitute for her pillow, she filled up an old sweater with leaves

[2] As a substitute for her pillow, she filled up an old sweater with water

(Glenberg & Robertson 2000, p.385.)

These sentences are given very similar numerical values by LSA, which amounts to the claim that they have very similar meanings and should be similarly understood by language comprehenders. On the other hand, the IH predicts significant differences in comprehension (and thus meaning) between the two sentences because of the difference in affordances. Because of the affordances of pillows, sweaters, leaves, and water, the situation described by [1] can readily be simulated, while the situation described by [2] cannot (Glenberg & Robertson 2000.)

Subjects in this study were presented with various sentences and asked to rate how sensible they were, on a scale from 1 to 7. As predicted by the IH, subjects assigned vastly different sensibility ratings to pairs of otherwise similar sentences which differed in their affordances. This has become known as the Affordance Compatibility Effect (Weiskopf 2010a.) This result seems strongly to favour the IH over its amodal rivals, since it is difficult to see why subjects should rate the sensibility of sentences like [1] and [2] so differently if their meanings are as similar as theories like LSA and HAL calculate.

Many other empirical predictions of the IH have been borne out (see, e.g., Glenberg & Robertson 1999; Kaschak & Glenberg 2000; Glenberg et al 2005.) The literature which it alone has generated is enough to establish that inquiry into ELC is a flourishing research programme. It is now time, however, to meet the IH's theoretical counterpart, the other main version of ELC—Zwaan's (2004) Immersed Experiencer Framework (IEF).

The Immersed Experiencer Framework

The IEF, like the IH, holds that linguistic comprehension is a three-stage process, the successful output of which is an inner simulation of a described situation. This much, the two accounts have in common. However, the three stages identified by the IEF—which are held to operate in an overlapping, rather than a strictly sequential, manner, and which do not correspond neatly to those posited by the IH—are called *activation*, *construal*, and *integration* (Zwaan 2004).

Each of the IEF’s three stages operates on a particular kind of linguistic unit, corresponding to a particular class of referents. Activation operates on single words or morphemes, corresponding to individual objects, properties, or actions; Construal operates on clauses or ‘intonation units’, corresponding to individual events; and Integration operates on entire discourses, corresponding to whole sequences of events. Already, it is clear that a central concern of Zwaan's is the dynamic and temporally unfolding nature of comprehension.

Considering an example in a coarse-grained fashion may be the clearest way to explicate the IEF's three stages. Take the following sentence, inspired by related examples in Zwaan (2004): ‘The eagle flew across the sky and then landed in its nest’. The first stage, Activation—which corresponds loosely to the Indexing stage of the IH—consists of the unconstrained activation of perceptual symbols corresponding to words or morphemes. Thus, upon reading the word 'eagle', traces of all kinds of previous sensorimotor representations of eagles are activated in a diffuse and non-specific fashion. Various different eagle-related representations will be activated, all to an equally low level.

Upon reading the word 'flew', the perceptual symbol for flight will be activated in a similar fashion (modulated, presumably, by the past tense form—a complication I will ignore here.) However, already, the second stage, Construal, will begin its work. This stage consists of the narrowing down or selection from the many possible representations of objects or actions to yield a relatively determinate, albeit schematic, representation of an event. So the activation of representations of flight constrains the activation of representations of eagles, narrowing them down to flying ones. The 'pull' exerted by the activation of flight-related representations will mean that of all the previously activated eagle representations, the flying ones will become activated to a higher level, thus being selected from among the available alternatives. A similar thing will happen with the next word, 'across', which will further constrain the representation by ruling out swooping or diving eagles; and so forth (Zwaan 2004.)

The third stage, Integration, is just the ordering of event simulations generated by Activation and Construal into temporally structured, coherent sequences. This is the process whereby the representation of the eagle not only flying, but flying *across*—and across *the sky*, to be precise—is joined with the subsequent representation of the eagle not only landing, but landing *in its nest*, to create a seamless and unified inner simulation of the sequence of events described by the sentence.

The IEF, then, views comprehension as a continuous, overlapping process in which perceptual and sensorimotor representations are accessed, constrained by further representations and other new information, and combined into simulations of denoted situations and (sequences of) events. As Zwaan (2004) puts it, "the comprehender is an immersed experiencer of the described situation, and comprehension is the vicarious experience of the described situation" (p. 36.)

A couple of qualifications should be made to this formulation. The first is that the word 'vicarious' is potentially misleading; there is no reason to suppose that a language comprehender simulating our example sentence would simulate it from the perspective of the eagle, or of any other object or entity in the situation. It is more natural to suppose that the comprehender would simulate perceiving the eagle's actions from an observer standpoint.

The second qualification is equally minor: Zwaan's use of the word 'experience' should not be taken too literally, in light of Barsalou's insistence—on good theoretical grounds—that sensorimotor simulation be understood as simply neural re-activation, with *or without* accompanying phenomenology.

As with the IH, some impressive empirical data have been adduced in support of the IEF. One study in cognitive psychology showed that language comprehenders are sensitive to differences between described situations in ways which allegedly can only be accounted for by appealing to sensorimotor simulation (Zwaan et al 2002.)

Subjects in this study, after reading sentences such as 'the ranger saw the eagle in the sky' and 'the ranger saw the eagle in its nest', were presented with a series of pictures, one at a time. For each picture, they had to indicate, as quickly as possible, whether the object it depicted was mentioned in the sentence or not.

The result was that subjects were quicker to respond correctly if the shape of the object in the picture matched the shape implied—but not explicitly stated—in the sentence. That is, after reading 'the ranger saw the eagle in the sky', subjects were quicker to respond correctly to a picture of an eagle with outstretched wings—as it would appear in the sky—than to a picture of an eagle with folded wings (Zwaan et al 2002.)

Zwaan et al. take this priming effect as evidence that the subjects generated a perceptual simulation of the eagle in the course of comprehending the sentence, because the information about its shape is not given by a purely semantic interpretation of the sentence. The fact that the eagle in the sky has spread wings only becomes obvious, they suggest, when one simulates or imagines (consciously or otherwise) what such an eagle would look like. Put differently, it is not obvious why we should expect this difference in response times if comprehending the sentence consists of activating amodal conceptual representations completely distinct from the sensorimotor modalities; whence the priming and interference effects, on such a picture?

Another psychological finding cited in support of IEF is closely analogous to the one just discussed, but concerns object orientation, rather than shape. This time, subjects had to read sentences such as 'John put the pencil in the cup' and 'John put the pencil in the drawer'. After each sentence, a picture of an object (e.g. a pencil) was presented, and the subject had to respond as quickly as possible to the question whether the depicted object had been mentioned in the sentence (Stanfield & Zwaan 2001).

Perhaps unsurprisingly by now, the result was precisely analogous to the one previous described: subjects were quicker to give correct responses when the depicted orientation of an object matched the object's orientation as implied by, but not explicitly stated in, the sentence. The explanatory inference, too, is the same: Stanfield and Zwaan attribute the different response times to subjects' having created, during the comprehension process, a perceptual simulation of the described situation. The effect found in relation to orientation and shape in these two studies has become known as the Appearance-sentence Compatibility Effect (Weiskopf 2010a.)

The IEF, then, is—like the IH—a version of ELC which describes comprehension as a three-stage process. In this case, the three stages operate in an overlapping and mutually modifying fashion. First, perceptual symbols representing objects and actions are activated by perceived words and morphemes. Second, mutual constraining of the perceptual symbols selects from among the many active possibilities to yield simulations of fairly determinate objects, properties, relations, and events, in accordance with the structure of clauses and intonation units. Third, the distinct events represented are integrated into coherent, temporally ordered sequences of events in accordance with the structure of the broader ongoing discourse. Once again, the result of successful comprehension is an inner simulation of what is being described by the comprehended sentence(s).

Zwaan (2004) cites many other empirical results in support of the IEF, including more from psychology, and others from disciplines such as neuroscience. However, we now have enough to step back from the details of the two theoretical versions of ELC and look at the big picture.

Considering ELC in General

Having looked at the IH and the IEF, as well as the evidence for them, in some detail, it is time for a little logical geography. Each of these theories is an explication of the central ELC claim that language comprehension constitutively involves sensorimotor simulation. Further, each of them holds that it is a process involving three distinct components or stages, although the components and stages they identify are different. With respect to their differences, one important question is precisely what 'simulation' amounts to.

For the purposes of the theory of PSS, recall, simulation is a technical term. It means the re-activation of more-or-less schematic sensorimotor representations, with or without accompanying phenomenology or conscious experience. In this sense, simply tokening, in isolation, a representation in the sensorimotor systems of the lines and vertices essential to being a triangle, without any representation of angles or orientation—and hence without phenomenology—counts as simulation. This is a minimal sense of the term.

Simulation, however, takes on more full-blooded meanings in the hands of Glenberg and Zwaan. Such claims as "language is a set of cues to the comprehender to construct an experiential (perception plus action) simulation of the described situation" (Zwaan 2004) make plain that they understand a simulation as something akin to a mental model of a state of affairs or sequence of events. Further, in light of the IH's emphasis on affordances, consider the following from Glenberg and Kaschak, part of which I quoted earlier:

> [T]he meaning of [a] situation consists of the set of actions available to the animal in the situation... language is made meaningful by cognitively simulating the actions implied by sentences...
>
> (2002 pp. 558-9.)

What is meant here is obviously something even more substantial—something along the lines of a mental model of an entire situation, *including* of the performance of various courses of action available in that situation.

We have, then, three possible senses that the term 'simulation' could take, in the core ELC claim. The weakest is the activation of more-or-less schematic sensorimotor representations. The next is the construction of a somewhat comprehensive model of a situation or sequence of events out of more-or-less schematic sensorimotor representations. Finally, the strongest sense is the construction of a somewhat comprehensive model of a situation or sequence of events, *including* affordances and actions, out of more-or-less schematic sensorimotor representations.

One of the reasons that this is important is that the plausibility and generality of ELC is likely to vary depending on what understanding of 'simulation' we adopt. This came out clearly in my earlier discussion of the limits of the classical notion of affordances in relation to the IH.

Also, for example, there may be some sorts of language—perhaps very abstract or technical language—the comprehension of which is difficult to account for if even a full-blown situation model must be involved, but which can be accounted for purely in terms of the activation of sensorimotor representations. In considering arguments for and against ELC, then, we can keep in mind these three possibilities about the meaning of the term 'simulation'.

Another dimension along which the strength of ELC could vary is helpfully identified by Weiskopf (2010a.) He notes that proponents of ELC could claim that sensorimotor simulation is constitutive of, or necessary for, all acts of language comprehension; more weakly, that it is constitutive of, or necessary for, only *some* acts of comprehension; or, more weakly still, that it is neither constitutive of nor necessary for any acts of comprehension, but that it plays a facilitating role in some such acts.

To my mind, it is questionable whether the last possibility is even a version of ELC at all. To claim that sensorimotor simulation aids some acts of comprehension, but is never constitutive of comprehension itself—that all genuine comprehension could potentially take place without any simulation—seems quite contrary to the spirit of the theory. Small wonder, then, that Weiskopf finds no conflict between this ‘weak’ sense of ELC and the amodal view of comprehension.

Be that as it may, the range of possibilities at our disposal now amounts to the following: understanding language always, or sometimes, constitutively involves, or is facilitated by, the activation of sensorimotor representations—perhaps to construct an entire situation model, perhaps also involving affordances and actions. In relation to the various arguments for and against ELC which we will encounter, we can always ask, firstly, what sense of 'simulation' is at play and to what extent comprehension is being identified with simulation in that sense.

Time, now, to look at the first set of arguments—those which aim to convince us that the amodal view is a better theory than ELC in some respect or other.

4. Alleged Amodal Advantages

Introduction

This chapter, as its name suggests, examines alleged advantages of the amodal view over ELC. Chief among these is the problem of accounting for our comprehension of abstract language. This is acknowledged as a problem for ELC by its proponents, as well as being pressed as an objection—both implicitly and explicitly—by its detractors. I will begin by presenting this problem in its various forms, and then discuss several attempts which have been made to account for the comprehension of abstractions in embodied terms. I will argue that these attempts, whatever problems they face in isolation, combine to create a compelling case for the claim that ELC is not at a disadvantage here. This argument will include a digression into issues relating to concept possession conditions, in which I will consider the possibility that friends of ELC might benefit from endorsing the Concept Pragmatism championed by Prinz and Clark (2004.)

The second alleged amodal advantage I will discuss is that of accounting for our ability to comprehend language whose meaning we seemingly cannot simulate with sensorimotor resources. This overlaps significantly with the problem of abstraction, but deserves independent discussion because it raises issues concerning the nature of our explanandum. It thus presents an opportunity to consider the question: what are we trying to explain when we try to explain linguistic comprehension? There is an influential view in the philosophy of language which holds that comprehension is knowledge of truth conditions. Weiskopf (2010a) uses this as a basis for criticizing ELC. My response to his arguments will be twofold. First, I will present some independent reasons for doubting that this analysis of comprehension is correct. Second, I will propose a dilemma about how truth conditions themselves are understood: on one horn, the truth conditions view will face infinite regress, while on the other, it will turn out to be compatible with ELC or something very like it.

The third alleged advantage of the amodal view concerns two distinct, but closely related, issues. As I have mentioned earlier, one of the strongest arguments for the amodal view or LOT in relation to conceptual thought generally has been its apparently unique ability to account for certain cognitive phenomena. Chief among these are the productivity and systematicity of thought. These issues do not pertain specifically to comprehension, so I will give them less attention than the others. However, since these theoretical arguments have been so influential, it is worth spending a little time to dispel lingering doubts that an embodied view of cognition in general is capable of accounting for such phenomena. In this section, I will rely chiefly on Barsalou's defence of PSS against amodal arguments.

The ultimate thesis of this chapter, then, is that the apparent advantages of the amodal view are merely that: apparent. ELC is, at least, not inferior to its rival when it comes to accounting for abstraction, comprehension of the unimaginable, productivity, and systematicity.

The Problem of Abstraction

The first thing to mention about the problem of abstraction is that the most prominent advocates of EC and ELC acknowledge it as a problem. Barsalou (1999) and Prinz (2002), champions of EC in psychology and philosophy respectively, both state in no uncertain terms that it is a real challenge for the embodied view to explain how we represent—using only sensorimotor resources—such concepts as TRUTH, IMPLICATION, NEGATION, LOVE, DEMOCRACY, ELECTRON, and so forth.

The problem is also raised sharply, whether explicitly or not, by critics of EC and ELC. For example, in discussing the IH, Fred Adams (2010) raises the possibility that it may be committed to a criterion of meaningfulness, such that only sentences which can be perceptually simulated are meaningful. In this case, suggests Adams, IH itself might be in deep self-refuting trouble. He asks: "can one perceptually simulate IH, itself? No! What are its affordances?" (p.624.) Now, I have argued elsewhere (Letheby 2012) that IH is not committed to a criterion of meaningfulness in the way Adams suggests. But this takes none of the bite out of his question. "[C]an one perceptually simulate IH", indeed? If not, then, even if IH is meaningful, it seems that—by its own lights—we cannot understand it. The reason, of course, why it is questionable whether we could simulate IH is that the following sentence is abstract if anything is:

> [IH] Language comprehension is a three-stage process of indexing words to perceptual symbols, deriving affordances from those symbols, and meshing those affordances to create a simulation.

Never mind criteria of meaningfulness; explaining how we could simulate the state of affairs described by [IH] in our sensorimotor systems is a challenge indeed.

The problem of abstraction also underlies an objection raised by Weiskopf (2010a.) He argues that there are pairs of sentences which do seem susceptible to sensorimotor simulation, but which have differences in meaning that could not be accounted for in sensorimotor terms. Consider, for instance, the sentences *the man stood on the corner* and *the man waited on the corner*. Obviously they both describe situations which we can imagine perceiving, and obviously they describe different situations: the second one contains information that the first does not. But it is not clear that this difference could be captured in sensorimotor terms. It seems that the simulations of both situations will be the same, and thus, that ELC cannot account for our undeniable ability to understand the difference. Why, though, does it seem that the difference cannot be captured in sensorimotor terms? I submit that it is simply because the difference between *standing* and *waiting* is an abstract one, in the sense of (arguably[6]) not being directly perceivable through any of the traditional sensory modalities.

Clearly, then, abstraction is a problem. But, inspired by the last sentence of the preceding paragraph, we might want to ask: what problem, exactly? What is meant by 'abstraction'?

There are at least two possible senses of the term[7]. The first is the sense of abstracting away from particular details, as in Barsalou's (1999) discussion of how we could represent a triangle in general. We achieve a certain measure of abstraction by only representing certain details—in this case, lines and vertices—and leaving others, such as angles and orientation, unspecified[8]. In this sense, most concepts—at least of kinds, rather than individuals—are abstract, in that they represent entire classes of things (or properties, or events, or relations) independently of the details of their instances.

The second sense of 'abstraction' is that to which I alluded in discussing the difference between standing and waiting. In this sense, abstract simply means something like *not perceivable by the senses*. Many concepts are abstract in this sense, some more clearly so than others. Possible examples include GOD, ELECTRON, TRUTH, JUSTICE, DEMOCRACY, LOVE, etc. This sense of abstraction is more obviously problematic for ELC. Even if the subtracting details strategy which Barsalou applies to the concept of triangle can gain a lot of mileage, it is hard to believe it can work to characterize such concepts as those I just listed. Clearly at least one more solution is needed.

Fortunately, many are in the offing. In their review of the literature on ELC and abstraction, Glenberg et al (2008) identify three distinct proposals. We can call these the Metaphorical Account, the Introspective Account, and the Action Schema Account. There are

[6] Arguably, because this may depend what view one takes about the phenomenon of 'mindreading' in social cognition; cf. Gallagher (2008.)

[7] Barsalou (2003) identifies six senses, but the fine grained distinctions involved will not be directly relevant here.

[8] I am grateful to Dr. Jon Opie for bringing this conception of abstraction to my attention.

two other proposals which I think are interestingly distinct from these three. I will call these other two the Situational Account and the Partial Dis-embodiment Account. Without further ado, then, let us take a look at each of these five proposals.

The Metaphorical Account has been given its fullest and most compelling exposition and defense by Lakoff and Johnson (1980a, 1999) and has also been pursued theoretically and experimentally by many others (e.g. Gibbs [1996], Boroditsky & Ramscar [2002], Casasanto [2009].) The basic idea is that abstract concepts are in some sense metaphorical extrapolations of concrete ones. There is much evidence from linguistics that people pervasively talk about abstract concepts in concrete terms. For instance, the metaphor MORE IS UP can be seen at work in such turns of phrase as *prices rose*, *stocks fell*, and the like. When we speak like this, we talk about the abstract concept of quantity in terms of the concrete, spatial concept of verticality.

Of course, talking is one thing, and thinking another. The Metaphorical Account claims that in discourse like this, how we talk genuinely reflects how we think. We do not merely talk about abstract concepts using concrete terms, but we actually think and reason about abstracta by tokening and manipulating representations of concreta in our sensorimotor systems (Lakoff and Johnson 1999.)

What is the evidence that warrants this inference from linguistic behaviour to cognitive process? For one thing, there is the fact that the inferential structure of abstract concepts seems to reflect that of the concrete concepts to which they are metaphorically mapped. Consider, for instance, the metaphor PURPOSES ARE DESTINATIONS, in which the abstract concept of a purpose is put in terms of the concrete concept of a (spatial) destination. Here, we find that inferences which are valid in the domain of destinations are mirrored by correspondingly valid inferences in the domain of purposes. For example, if X is travelling to destination D, then it follows that X has not yet reached D. Correspondingly, if X is working towards purpose P, then it follows that X has not yet achieved P.

Lakoff and Johnson (1980a, 1980b, 1999) in particular present an impressive array of evidence from different languages, cultures, times, and places, aimed at showing that such metaphors are real, ubiquitous, and genuinely cognitive rather than merely (superficially) linguistic. They describe in detail how complex metaphors such as A PURPOSEFUL LIFE IS A JOURNEY are constructed out of basic or primary metaphors such as PURPOSES ARE DESTINATIONS and ACTIONS ARE MOTIONS, many of which may be human universals (Lakoff and Johnson 1999.)

One prima facie difficulty for the Metaphorical Account is that there is obviously a difference between abstract concepts and concrete ones. It is all very well to say that the abstract concept of a purpose is metaphorically modeled on the concrete concept of a (spatial) destination, and represented—like the latter—in the sensorimotor systems; but what then is the cognitive difference between thinking *she has not yet reached her destination* and thinking *she has not yet achieved her purpose*? As Prinz (2002) puts the point: "metaphors leave remainders" (p. 171.) Call this, then, the Problem of the Remainder, because, given any abstract concept and any proposed concrete metaphorical basis for it, there will still be something of the abstract concept which remains unaccounted for.

So much for the Metaphorical Account. The second embodied story about abstraction is the Introspective Account, due to Barsalou (1999) in his original presentation of PSS. Being aware of the work of Lakoff and Johnson on conceptual metaphor, but believing—like Prinz—that it could not be the full story, Barsalou set out to give another account of abstract concepts in terms of perceptual symbols derived from repeated introspective experience. To explain this, it is probably best to look at a couple of examples. Therefore, I will now present Barsalou's treatment in these terms of the concepts TRUTH and NEGATION.

According to PSS, we can represent the proposition THE CAT IS ON THE MAT by combining our simulators for CAT, ON, and MAT to create a simulative model of the state of affairs of the cat's being on the mat. So far, so good. Having done so, we might then go into the room containing the mat and look at the mat to see if the cat is indeed on it. Suppose we see that it is. In that case, we now have two sensorimotor representations of situations: a simulation of the cat's being on the mat, and an occurrent perception of the cat's being on the mat.

What we also have now is a match, or mapping, between our perception and the simulation (which latter is our representation of a proposition.) We have a situation in which our off-line representation of the world matches our on-line perception of the world; we see that things are as we thought them to be. Suppose, now, that this cognitive event of a successful mapping can *itself* be stored in long-term memory, through the operation of selective attention, for later re-activation and use. Given repeated experiences of such successful mappings, we might develop a robust simulator mechanism capable of categorizing new cognitive events as belonging to the class or type of successful representation-to-world mappings. Under the influence of our linguistic community, this simulator mechanism might also come to integrate linguistic information from auditory, visual, or tactile modalities, such as representations of the words 'true' and 'truth'. Thus, we acquire a concept of TRUTH (Barsalou 1999.)

There is a certain intuitive plausibility to this account. Further, a corresponding account of NEGATION falls straightforwardly out of it. This is, of course, simply a simulator built up from repeated experiences of failed representation-to-world mappings, or of representation-to-world mismatches.

Barsalou applies this strategy to several other abstract concepts—notably, other logical ones such as IMPLICATION and DISJUNCTION. The extent to which his analyses of these various cases are convincing is open to debate. In any case, despite its merits, the main potential stumbling block for his Introspective Account seems to be its reliance on a questionable Inner Sense model of introspection.

In order to sustain the claim that we can develop simulators of such cognitive events as successful or failed simulation-to-perception mappings, one must also claim that introspection is a modality similar in the relevant respects to the various sensorimotor modalities. It must be a faculty which receives a channel of information from its domain (i.e. our own cognitive processes) and generates, on the basis of this information, on-line representations of the current state of affairs in that domain. Further, these on-line representations must be subject to the operations of selective attention, long-term storage, and recall, just as the five senses, proprioception, and the like are.

However, the view that introspection is a faculty similar in such important respects to the sensorimotor modalities is known as an Inner Sense theory of introspection or self-knowledge, and it is increasingly controversial in cognitive science and philosophy. Several thinkers—notably, Peter Carruthers (2011)—have been developing a case for the claim that our access to our own mental states is nowhere near as direct, reliable, or perception-like as such a model would suggest. Without venturing into the self-knowledge debate here, I will just note that the biggest potential stumbling block for the Introspective Account of embodied abstraction is that it seems committed to—and thus hostage to the increasingly doubtful plausibility of—an Inner Sense model of introspection.

Moving on, the third proposal concerning ELC and abstraction is what I have called the Action Schema Account, due to Glenberg and his colleagues (Glenberg and Kaschak 2002; Glenberg et al 2008.) This shares a large part of the spirit of the Metaphorical Account, in that it suggests that our understanding of (certain kinds of) abstract language operates via a metaphorical extension of our understanding of concrete language. In this case, however, what are allegedly re-used for abstract comprehension purposes are general schemas which enable our fluid performance of various kinds of physical actions.

Proponents of the Action Schema Account propose that we have 'action schemas' grounded in very basic patterns of bodily experience. For instance, we might have a schema for transfer of objects, consisting of the following parameters: kind of grip, determined by the object; level of force required, also determined by the object; and direction of motion, determined by the starting location of the object (oneself or another, in the case of giving and receiving respectively) and the destination of the object (another or oneself, in the case of giving and receiving respectively.) Research suggests that the development of such schemas underlies development of an understanding of transfer verbs such as *to give* (Glenberg et al 2008.)

Now, supposing that we have action schemas for object transfer, it seems reasonable to suppose, on ELC, that they might be involved in our comprehension of such sentences as *Andy delivered the pizza to you* and *you delivered the pizza to Andy*. This would provide a detailed and interesting explanation of the differential reaction times found by Glenberg and Kaschak (2002) in their study of sentence comprehension relating to object transfer (see discussion in section 2 of chapter 3 above.)

What is even more interesting, however, is that Glenberg and Kaschak found the same result in testing comprehension of sentences describing transfer of abstract objects. That is, reading and understanding the sentence *you radioed the message to the policeman* affected motor response times in the same way as did reading and understanding the sentence *you delivered the pizza to Andy*. And reading and understanding the sentence *Liz told you the story* affected motor response times in the same way as did reading and understanding the sentence *Andy delivered the pizza to you* (Glenberg & Kaschak 2002.)

The Action Schema Account of abstraction provides an attractive explanation of this result: viz., that the same action schemas which underwrite our comprehension of concrete transfer sentences are also implicated in our comprehension of abstract transfer sentences. Thus, when we understand *Liz told you the story*, we activate the same object transfer schema as we do in understanding *Andy delivered the pizza to you*, with the parameters being set as direction of motion towards the body.

One obvious shortcoming of this approach, at least in isolation, is its limited scope: clearly not all abstract language is comprehended using action schemas. What action schema would we recruit, for instance, to understand *physicists have demonstrated the existence of the Higgs Boson*, or *the concept of God is that of a being greater than which none can be conceived*? Perhaps these questions can be answered, but clearly they can also be multiplied, and it is doubtful that the Action Schema Account can generalize to all kinds and degrees of abstract conceptualization. The focus on *action* schemas specifically may reflect Glenberg's focus on the action-guiding aspects of language; Lakoff and Johnson (1999) discuss other kinds of sensorimotor schemas, particularly spatial ones, which may be recruited in the comprehension of abstraction.

Another shortcoming, which the Action Schema Account shares with the Metaphorical Account by virtue of their similarity, is the Problem of the Remainder. The proposal that we understand sentences about storytelling through the same processes whereby we understand sentences about pizza delivering strikes us as novel and provocative precisely because we know there is a big difference; we all understand that in one case a physical object is literally being transferred, and in the other case it is not. There must, therefore, be something going on other than the activation of an action schema when we comprehend *Liz told you the story*. This is not to say that we cannot specify what that something is. It is just to emphasize that the Action Schema Account cannot carry all the weight alone. Glenberg et al acknowledge this, commenting that the "various approaches to the embodiment of abstract language are not mutually exclusive; in fact, they may all be emphasizing different aspects of the same phenomenon" (2008, p.908.)

The fourth proposal to consider is the Situational Account. Like the Action Schema Account, this is not intended as a complete or comprehensive story about the embodied representation of abstract concepts. Rather, it is a proposal about one aspect of how such concepts might be represented and processed on an embodied view. The proposal, due to Barsalou and Wiemer-Hastings (2005), is that an important part of the content of abstract concepts consists of knowledge of the kinds of situations to which they apply.

Barsalou and Wiemer-Hastings review empirical evidence supporting the idea that situational content is central to both concrete and abstract concepts. For example, to have a full or adequate grasp of the concept HAMMER, it is not enough simply to possess visual, tactile, and proprioceptive representations of hammers themselves. One must also have some knowledge of the sorts of things people do with hammers, how they do them, in what circumstances, and the like (Barsalou & Wiemer-Hastings 2005.)

However important such situational content is to concrete concepts, Barsalou and Wiemer-Hastings argue that it is even more important to abstract ones. The previous paragraph notwithstanding, it is possible to represent the concept HAMMER—albeit in a severely limited fashion—simply by tokening straightforward perceptual representations of hammers. But it does not seem possible at all to represent the concept TRUTH simply by tokening straightforward perceptual representations of truths. Rather, one must represent a situation in which some claim is made and the world is arranged accordingly. It is only in the situation partly comprised by the relationship of claim to world that the concept of TRUTH becomes manifest. Thus, knowing the kinds of situations to which they correctly apply seems even more important to abstract concepts than to concrete ones.

This proposal, then, can perhaps be summarized by saying that the difficulty of representing abstract concepts in terms of simple perceivable features or performable actions can be mitigated by representing them, at least to a large extent, in terms of complex situations and situation types. As I have already mentioned, it is certainly not an adequate total theory of the embodied representation of abstraction. It also needs to be handled with care, lest it give rise to a Euthyphro-type dilemma: does abstract concept A apply to these situations because the situations independently have the property of being A-instances? Or do the situations have the property of being A-instances only because abstract concept A applies to them? The latter unpalatable horn only threatens if we suppose the Situational Account is being offered as a complete theory of abstract conceptual content, which is just another reason for emphasizing its lack of any such ambition.

The fifth and final proposal to consider is the Partially Dis-embodied Account, due to Dove (2011), with antecedents in the work of Barsalou (1999) and Louwerse (2011.) The core claim of this account is that the content of abstract concepts consists at least partially of their associative and inferential relationships with other concepts, encoded in sensorimotor representations of words. So, for example, part of the content of the concept TRUTH is its relationship of antonymy with the concept FALSITY, and this relationship is represented in terms of sensorimotor representations of words such as 'true', 'truth', 'false, and 'falsity', as well as their relationships with one another.

It is hard not to see this as a dramatic concession. After all, the claim that linguistic context is constitutive of conceptual content is at the heart of amodal theories of linguistic meaning like LSA and HAL—commonly viewed as antithetical to ELC (Glenberg & Robertson 2000.) The Partially Dis-embodied Account seems thus to amount to an admission that ELC is only true about certain kinds of language.

There is a sense in which this is true. To the extent that the Partially Dis-embodied Account is correct, we do not comprehend abstract sentences by constructing simulative situation models of the states of affairs they describe. But here we should recall the different possible ELC claims distinguished in the final section of chapter 3. It is not essential to ELC that the sort of simulation implicated in comprehension always amounts to the construction of a situation model. Rather, it can just be simulation in Barsalou's minimal sense: the tokening of more-or-less schematic sensorimotor representations under endogenous control. This is certainly still doing the work in comprehension, on the Partially Dis-embodied Account. And, importantly, what is *not* doing the work is any kind of amodal representational code: there is still no need to posit a LOT or anything like it. Sensorimotor representations can take us all the way, even if situation models cannot.

Dove (2011) coins the (hyphenated) term 'dis-embodied' to refer to representations which are modality-specific but whose semantic content is not a function of their modality-specific content. The hyphen in 'dis-embodied' is supposed to distinguish it from the term 'disembodied', which Dove reserves for amodal representations proper: symbols in a LOT. An example of a dis-embodied representation is an auditory representation of the word 'democracy'. This representation is couched in a modality-specific code, but—supposing it can represent, or stand in for, democracies in various cognitive processes—it has semantic or referential content which is arbitrary relative to its modality-specific content.

The Partially Dis-embodied Account, then, preserves the minimal core claims of EC and ELC: that all representations manipulated in higher cognition are couched in modality-specific codes, and that there is no LOT. However, it claims that modality-specific representations of natural language phonemes, graphemes and the like perform roles similar to those for which an amodal code has typically been postulated, representing categories arbitrarily related to their form—including, importantly, abstract ones.

There is much more which could be said about each of these five proposals concerning abstraction. Now, however, it is time to step back and consider what can be said in general terms about the ability to handle abstraction as an alleged advantage of the amodal view.

One question which merits attention, and to which I have alluded several times already, is whether a single, unitary, monolithic account of the representation of abstract concepts is necessary or desirable. Prinz (2005) mentions this question at the beginning of his articulation and defence of an embodied theory of moral concepts—a theory which, interestingly, does not fit neatly into any of the five embodied accounts of abstraction I have discussed. He contends that because abstract concepts themselves differ greatly and do not form a homogeneous class, there is no reason to suppose that a single story will be true about how they are all represented. Even if all representation and cognition operates on sensorimotor resources, different sorts of sensorimotor representations and combinations thereof might be used in different ways, by different processes, to represent different sorts of concepts.

This seems to me correct. Viewed in this light, the partial and incomplete nature of each of the embodied accounts of abstraction can be seen as a virtue. We should expect that various individually insufficient stories about the embodiment of abstraction can be told, amounting to a pool of resources which can be drawn on as the problem of abstraction is "dismantled piecemeal" (Prinz 2005, p.93.)

A second general question to consider concerns the nature of the explanandum: What are we trying to explain when we try to explain the representation of abstract concepts? Two possible opposing answers to this question can be found in the debate over Concept Rationalism and Concept Pragmatism (Fodor 2004; Prinz and Clark 2004.) These are two different doctrines about the conditions under which cognitive beings possess concepts.

According to Concept Rationalism (or ‘Cartesianism’), defended by Fodor, a necessary and sufficient condition for possessing some concept C is the ability to “think about Cs ‘as such’” (2004, p. 29.) Defined negatively, this means that concept possession is not identified with any kind of ability, such as the ability to distinguish Cs from non-Cs, or to draw certain inferences about Cs. These are, of course, precisely the sorts of things with which Concept Pragmatism identifies concept possession: abilities of one sort or another, be they abilities of classification, of knowledge-deployment, or whatever.

This is not the place to adjudicate the debate between Rationalists and Pragmatists about concept possession. A few things, however, are worth noting. The first is that, insofar as the question is open, Concept Pragmatism seems like a natural ally of any embodied view of cognition, given the emphasis typically placed by the latter on the evolved, biological, and action-oriented nature of thought (Prinz and Clark 2004.) The second thing to note is that Concept Pragmatism looks like a far more suitable basis for an empirical science of concepts than Concept Rationalism, because such abilities as classification and knowledge-deployment are much easier to operationalize than the ability to think about Cs ‘as such’.

It is tempting, in light of this, to argue as follows: Proponents of ELC should also be Concept Pragmatists, and Concept Pragmatism eases the pressure on an account of abstract concept possession by more clearly specifying the explanandum in a way that takes much of the mystery out of the topic. If we adopt Concept Pragmatism, then it seems like we no longer have to tell a story about how people can manage, using only sensorimotor resources, to think about truth, negation, disjunction, democracy, electrons, and the like ‘as such’ (if that was in fact what we thought we had to do.) Rather, we now have to tell a story about how people can manage, using only sensorimotor resources, to distinguish truth from untruth, democracy from non-democracy, etc., or to deploy knowledge about the typical features of truths etc., or to draw sensible inferences about truths, etc. This leaves the cognitive operations underwriting such abilities rather under-specified, thus making the problem look refreshingly tractable.

However, there are two problems with this line of argument. One stems from the fact that, as Prinz and Clark point out, the line between Concept Pragmatism and Concept Rationalism is actually quite blurry. One can be something of a Pragmatist, holding that all and only those creatures able accurately to distinguish Cs from non-Cs (or whatever) possess the concept C, and yet maintain that as a matter of fact, such classificatory abilities are underwritten by the ability to think about Cs 'as such' (Prinz & Clark 2004.) In fact, I think this is the way we have to go if we want to endorse the cognitivist, representationalist form of EC and ELC which I am defending. To identify concept possession with various epistemic abilities and hold that those abilities are *not* underwritten by the ability to genuinely think about the referents of concepts is to flirt with a more radical enactivist or anti-representationalist strain of embodiment, which I will argue in chapter 6 is costly and under-motivated. So opting for Concept Pragmatism over Concept Rationalism does not obviously ease the burden on ELC proponents to explain, in sensorimotor terms, people's ability to think about abstracta as such.

The second problem with my strategy above is that this thesis is not about concept possession; it is about linguistic comprehension. It is true that I have been at pains to emphasize the intimate connection between conceptual cognition, on the one hand, and comprehension on the other, and the extent to which evidence for EC and ELC can be treated as interchangeable. However, it is clear that the two issues are nonetheless separate. Suppose I were to give, right now, a bulletproof, ironclad argument for the truth of some form of Concept Pragmatism. We would then know that any organism O possesses arbitrary abstract concept A just in case O is able to draw the right inferences (or whatever) about As. But we could still sensibly ask what goes on when O understands language about As.

We should be careful here, however. Perhaps a suitably adapted Pragmatist answer can be given to this question. Perhaps, when presented with a sentence S about As, O understands S just in case O is able to draw the right inferences from S, or something along those lines[9]. This would amount to a kind of Comprehension Pragmatism which would ease the explanatory burden on the ELC-proponent in the way originally intended, reducing the project of explaining abstract comprehension in sensorimotor terms to the project of explaining abstract inference-drawing (etc.) in sensorimotor terms.

Having reached this point, my commitment to a cognitivist or representationalist construal of ELC is again something of a stumbling block. Even if compelling philosophical analysis were to yield necessary and sufficient conditions of a pragmatist kind for either concept possession *or* language comprehension, there would still remain the question of how, or by what cognitive mechanisms, actual organisms satisfy those conditions. And I am committed to the view that these mechanisms are representational—they are mechanisms of simulation; in other words, of thinking. So concluding that O understands S just in case O can draw certain inferences raises the question of how O can draw those inferences, to which I must answer: because O has thoughts which accord with the (abstract) content of S. This, in turn, raises the question of how O manages to have thoughts which accord with the (abstract) content of S, which leaves me back where I started: attempting to account for our ability to represent abstract concepts—to *think about* abstracta—using purely sensorimotor terms.

There is one caveat to this, however. This is the idea that our ability to think about any given concept—and consequently to understand language involving terms for that concept—*admits of degrees*.

---

[9] This is very close to the enactivist proposal, to be discussed at length in chapter 6, that "language comprehension can be accurately described as the procedural knowledge how to respond in certain situations to specific utterances" (Van Elk et al 2010, p.4.)

Let me begin to unpack this by returning to the five embodied strategies for dealing with abstraction. These, recall, are: the Metaphorical Account; the Introspective Account; the Action Schema Account; the Situational Account; and the Partially Dis-embodied Account.

Now, suppose, that—as I have followed Prinz in arguing—there is no reason to suppose that any one of these stories will or should apply to all abstract concepts. Suppose, further, that—as I have followed Glenberg in arguing—the full story about any given abstract concept may well consist of some combination of two or more of these accounts.

In this case, when considering any abstract concept, such as TRUTH, we can start from the premise that its representation in a given human mind will consist of some combination of the following: a metaphorical mapping to one or more concrete concepts, encoded in sensorimotor representations; a stored stock of introspective representations of cognitive events, such as simulation-to-perception mappings; links to schemas for performing certain kinds of physical actions; a stored stock of schematic representations of different (kinds of) situations to which TRUTH is applicable or appropriate; and a stored stock of sensorimotor representations of the word TRUTH and its cognates, from various modalities, along with its contextual, inferential, and other relationships to many other words.

In light of this picture, two things seem abundantly clear. One is that this is a very rich stock of resources with which to account for the ability to perform various cognitive operations involving abstract concepts. Another is that possession of all these resources admits of degrees.

For example, an individual's concept TRUTH could be metaphorically mapped to only one concrete concept, or many. The concrete concept(s) to which it is mapped could contain more or less detail. It could contain many stored introspective representations, or few, or none. It could be linked to only one action schema, or many… and so forth.

To take an example: as a philosopher, my concept of VALIDITY is probably quite rich, containing much situational and linguistic information, in particular. I have a comparatively large body of knowledge about the sorts of situations to which VALIDITY is applicable, and its relationships (including definitional ones) to other words. In contrast, my concept of HIGGS BOSON is undoubtedly quite meager. It, too, contains situational and linguistic information, but relatively little; and due to my (introspectively apparent) ineptitude in conceptualizing of physics and similar topics in spatial or other similar terms, I would suspect that it is fairly bereft of metaphorical mappings.

To say that the ability to comprehend language about some concept C (and the corresponding ability to think about Cs) admits of degrees is not to deny clear-cut cases of comprehension and non-comprehension. For example, I definitely do not, at the time of writing, understand the word *calidad*[10], and I definitely do understand the word *tree*. The idea is, however, that of those words which I do understand, I understand some better, or more fully (or richly) than others.

This idea is helpful in tackling the problem of abstraction because it immediately suggests the plausible idea that we typically do not understand abstract ideas as well, or as fully, or as richly as concrete ones. Examples such as my minimal understanding of sentences about the Higgs Boson are probably pretty close to a limiting case of this putative phenomenon. I can understand language about the Higgs Boson, to a certain extent; and I can think about the Higgs Boson, to a certain extent. But explaining these abilities of mine is not too formidable a task, because, in each case, the extent is so limited that there is relatively little to explain!

This is the grain of truth to be gleaned from the initially promising Concept Pragmatist strategy: we should be wary of shouldering more of a burden than is necessary by trying to explain more than there actually is to be explained. ELC need only account for our ability to

[10] Spanish: 'quality', as I have since learned.

understand abstract language to the extent that we *can* understand it—and once it is admitted that understanding may admit of degrees of richness or fullness, it is not hard to see abstract language as a prime candidate for language which is typically less richly or fully understood than other kinds.

What can we now say about the claim that the amodal view has an advantage over ELC with respect to accounting for our ability to comprehend abstract language? Well, this conclusion is usually derived from the premise that ELC cannot account for this ability because it is limited to invoking only sensorimotor resources. I think we are now in a position to see that this is false. ELC has a wealth of promising strategies invoking only sensorimotor resources with which to account for the comprehension of abstraction, especially with the clearer conception of a gradable explanandum which I have just developed.

Of course, this does not show that the amodal view does not have some kind of advantage here. It just shows that such an advantage cannot consist in a mere ability to account for the comprehension of abstraction. It must consist in a superior ability to do so; there must be something about the amodal view which makes it better equipped to account for this than the ELC view, despite all the resources of the latter.

This seems implausible, however, in light of the Partial Dis-embodiment Account. Any special advantages which an amodal symbol system would confer with respect to the comprehension of abstraction could surely be mimicked by a system of sensorimotor representations of natural language, functioning effectively as an arbitrary semantic code despite its modality-specific nature. Further, as discussed earlier, this is not the substantial concession which it appears to be, because there is still no need to postulate an actual amodal code or LOT, and because embodied (as opposed to dis-embodied) representations are still doing the lion's share of the work in most instances of comprehension.

My final verdict, then, is that the amodal view does not enjoy any advantage whatever over ELC with respect to accounting for our ability to comprehend abstract language. With that conclusion established, it is now time to examine the second, closely related, alleged amodal advantage: that of accounting for our ability to comprehend language whose meaning we cannot simulate.

Comprehending the Unimaginable

This objection is often tangled up with the Problem of Abstraction, because there is considerable overlap between the two. It is pressed most clearly and distinctly, however, by Weiskopf (2010a) and can be summarized as follows: there is language which we can understand, but the meaning of which we cannot simulate; therefore, understanding is not identical to simulation; therefore, ELC is false.

There are at least three different putative varieties of language which we can understand but the meaning of which we cannot simulate. One is abstract language. Another is language describing things we have never experienced—e.g. sentences about dodos or dinosaurs, or, for physically disabled people, sentences about performing various kinds of physical actions. A third kind is nonsense sentences, such as Chomsky's famous "colorless green ideas sleep furiously" (1957, p.15.)

The claim that we can understand such sentences as Chomsky's may seem dubious. Weiskopf's grounds for this claim consist of an influential view in the philosophy of language that sentence understanding is identical to knowledge of truth conditions (Davidson 1967.) On this view, which Weiskopf calls a "*minimal* view of linguistic understanding" (2010a, p. 298), it is sufficient to demonstrate understanding of some sentence S to be able to infer from S what would have to be the case for S to be true. So, for example, we can infer that if Chomsky's sentence is true, then ideas “slept, and they were both colorless and green” (ibid.), and this demonstrates that we can understand Chomsky’s sentence. However, it seems clear, says Weiskopf, that we cannot create a sensorimotor simulation of the situation of colourless green ideas sleeping furiously; thus, here is a case of understanding without being able to simulate, which is all that is required to refute ELC’s identification of the former with the latter.

It is not difficult or distorting to reconstruct this argument as a claim that the amodal view enjoys an advantage over ELC. Given its identification of understanding with simulation, runs the objection, ELC cannot account for our (alleged) ability to comprehend such unimaginable sentences as Chomsky’s ‘green ideas’. Lurking not far from the surface, of course, is the thought that the amodal view *can* account for our ability to comprehend such sentences. If understanding language consists of tokening symbolic representations in a LOT, or something like it—representations which may or may not be incidentally associated with any sensorimotor content—then it seems clear that such representations could readily be tokened for the ‘green ideas’ sentence. It further seems clear that such truth-conditional inferences of the sort which interest Weiskopf could be drawn from them in a purely syntactic fashion, unhampered by any inability to construct a sensorimotor simulation.

There are a number of ways to respond on behalf of ELC. The first is just to recapitulate one of the strategies I used in dealing with the Problem of Abstraction: namely, to recruit the Partially Dis-embodied Account and its ability to simulate the properties of an amodal symbol system using sensorimotor representations of natural language. Even if it is conceded that the ability to compute truth-conditional inferences is sufficient for understanding, it does not follow that ELC—with its restriction to only sensorimotor resources—is unable to account for our understanding of the ‘green ideas’ sentence. We just need to claim that the computation of truth-conditional inferences is performed by the manipulation of dis-embodied (sensorimotor) representations of natural language.

Of course, this does concede something to Weiskopf: that the ability to construct simulations *in the sense of full-blown situation models* cannot be a necessary component of understanding. But I have already admitted this; one of the qualifications which ELC has accrued in the course of my discussion so far is that, in order to be generally applicable, the simulation with which it identifies comprehension must be understood in Barsalou's minimal, technical sense, meaning the re-activation of stored sensorimotor representations under endogenous control. So far, we have seen no reason to suppose that any kind of language comprehension—of abstract language, or of unimaginable-because-nonsensical green ideas type sentences—occurs in the absence of simulation, so construed.[11]

The second response to Weiskopf's argument involves questioning his claim that we do understand the green ideas sentence and its ilk. It is undeniable that we can draw truth-conditional inferences from it; and it follows from this, given the minimal view of linguistic understanding, that we understand it. But what reason do we have to accept the minimal view of linguistic understanding?

This rhetorical question could obviously be answered. There are deep and complex debates in the philosophy of language surrounding these issues. However, I think it is worth emphasizing that, if the minimal view and ELC are genuinely inconsistent with one another, then all the evidence we have for ELC is equally evidence *against* the minimal view. This point speaks to deep methodological issues concerning the relationship between *a priori* philosophical analysis, on the one hand, and empirical inquiry, on the other. As Prinz (2004) argues in the context of the philosophy of emotions, when we conduct scientific investigation into the nature of some putative phenomenon, we need to allow for the possibility that the

[11] One possible counter to this response is to claim that, even if sensorimotor representations of words are being tokened in the course of understanding the green ideas sentence, they are not doing the real work. Syntactic mechanisms—which, as I have admitted, may or may not be sensorimotor—are responsible for the computation of truth conditions.

Thus, the success of this first response to the green ideas objection may depend on the viability of an embodied theory of syntax. Apart from Embodied Construction Grammar, which I mentioned earlier, Alistair Knott has recently (2012) published a book-length treatment of an embodied interpretation of Chomskyan syntax—but I do not have space to address these issues here.

results of that investigation will unseat some of our pre-theoretic or, indeed, pre-scientific views about the nature of the phenomenon. If our best evidence indicates that paradigm cases of linguistic comprehension involve something richer than mere computation of truth conditions, then we must at least be prepared for the possibility of abandoning the view that comprehension consists of such computation.

The strategy of questioning Weiskopf's claim that we understand the 'green ideas' sentence can be taken in a slightly different direction, too. Recall the idea, discussed at length in the previous section, that comprehension admits of degrees. If I may indulge in a little intuition-pumping, I submit that it is eminently plausible to say we do not understand the 'green ideas' sentence to the same extent, or even in the same sense, as we understand the sentence 'the cat sat on the mat'. Suppose that, as Weiskopf believes, there is a wholly amodal (disembodied, not merely dis-embodied) conceptual system syntactically computing the inferential consequences of all sentences we process, and sensorimotor representations are only contingently activated as a causal consequence of *some* such computations. In this case, there would still clearly be a big difference between merely computing inferences from the 'green ideas' sentence, on the one hand, and on the other, computing inferences from 'the cat sat on the mat' while also simulating perceptions of a cat sitting on a mat. The best way I can think of to express this difference is that we would understand far more clearly *what the world would have to be like* for the cat sentence to be true than for the 'green ideas' sentence.

This last remark brings me to my final response to Weiskopf, which consists of a dilemma for the truth-conditional view of understanding which he advocates. Suppose we accept that understanding a sentence consists of knowing its truth conditions[12]. We may then ask in what manner those truth conditions themselves are known. On the first horn of the dilemma, they are themselves further linguiform entities. This seems to be what Weiskopf has

[12] Or satisfaction conditions, to include comprehensible non-declarative sentences.

in mind when he conceives of knowing truth conditions as the ability to draw inferences. But if our knowledge of the truth conditions of some sentence—which knowledge allegedly constitutes our comprehension of that sentence—consists in the tokening of further linguiform representations, then we can legitimately ask how those truth conditions *themselves* are understood. On this first horn of the dilemma, an infinite regress threatens.

The second horn, of course, involves denying that our knowledge of truth conditions consists in the tokening of further linguiform representations. In this case, we can legitimately ask in what it *does* consist. I submit that sensorimotor simulation is an eminently plausible answer to this question. If knowing the truth conditions of a sentence intuitively amounts to knowing what the world would have to be like for a sentence to be true, then it is difficult to imagine a more effective way of doing this than undergoing a simulated perception of the world being as the sentence describes.

I conclude that the ability to account for our comprehension of the unimaginable is not an advantage of the amodal view. Even if the minimal view of comprehension is correct and computation of truth conditions suffices for understanding, the qualified version of ELC which I am defending can explain such computation in terms of dis-embodied (sensorimotor) resources; no amodal code is necessary. Further, the empirical and theoretical case for ELC may give us good reason to reject the minimal view of comprehension. If not, and we accept that knowing truth conditions suffices for some degree or kind of understanding, there are clearly higher degrees and richer kinds of understanding which outstrip such mere inferential competence. Finally, the question of how truth conditions themselves are understood seems to lead either to infinite regress[13] or to compatibility of the minimal view with ELC.

[13] As Dr. Jon Opie has pointed out to me, a proponent of the minimal view could here respond by identifying knowledge of truth conditions specifically with an inferential *capacity*, as opposed to the tokening of a linguiform representing vehicle. I think this move is problematic because of conceptual considerations relating to the notion of understanding, which I will discuss briefly in my suggestions for future directions in chapter 7.

I turn, now, to the final alleged amodal advantage: the much-discussed productivity and systematicity of thought.

Productivity and Systematicity

As I have already mentioned, the issues to be discussed in this section tend not to arise in the context of ELC per se, but rather in the context of EC more generally. As a consequence, I will treat them rather briefly. However, given how influential they have been in debates about the necessity of postulating a LOT (Fodor & Pylyshyn 1988; Chalmers 1989) it is worth at least saying something, since I am defending a view according to which there is no need to postulate a LOT.

I will describe the problems one at a time. Firstly, productivity is the apparent ability of the human cognitive system to construct an indefinite number of complex representations using a finite stock of representational primitives and mechanisms for combining them (Prinz 2002.) It is probably best to illustrate this with an example, of a classic sort: Unless you have read this thesis before, then there is a good chance that you have never entertained the thought that rhinoceroses never ride shopping trolleys to the pub wearing purple jackets and polka dot ties. However, there is an equally good chance that you are perfectly capable of entertaining, and indeed endorsing, this thought.

The important point is that it does not stop there. You can just as well entertain the thought that rhinoceroses never ride shopping trolleys to the pub wearing purple jackets, polka dot ties, and tutus. Or you can entertain the thought that giraffes never ride shopping trolleys to the pub wearing purple jackets and polka dot ties. You can entertain, in fact, a seemingly indefinite number of novel, complex thoughts simply by combining members of the finite repertoire of concepts you already possess in indefinitely many novel, complex ways which are nonetheless governed by a finite repertoire of basic combinatorial principles or rules. (Propositions of mathematics provide another clear example.)

This is a first gloss of what is meant by the productivity of thought. Specifying exactly what it amounts to in a precise, formal fashion is no small matter. However, an intuitive grasp of the phenomenon is all we need to assess the question whether EC and ELC are capable of accounting for it. Before I argue that they are, however, I will briefly describe the systematicity of thought, because the same fundamental aspects of PSS allow it to account for both of these (not unrelated) phenomena.

Systematicity, like productivity, is difficult to characterize precisely. Imprecisely, it refers to the apparent fact that the ability to think certain thoughts is inescapably connected with the ability to think certain other, related, thoughts. To use a well-worn example, anyone who is able to think that *John loves Mary* is ipso facto able to think that *Mary loves John* (Fodor 2004.) You simply do not find people who have the ability to entertain one thought but not the other. Whether it is included in one's characterization of systematicity per se, or as part of an explanation of it, it seems natural to suppose that this is related to the way in which such complex thoughts as *John loves Mary* are constructed from their constituent parts. If we assume that the actual cognitive structure of such thoughts mirrors, to some extent, the structure of the language with which we express them, then it looks as though their meaning is a lawful function of the meanings of their constituents plus the ways in which those constituents are combined.

Languages both formal and natural are fairly clear examples of representational or communicative systems which exhibit productivity and systematicity in the senses I have described. Examples include spoken and written languages such as English, Hindi, and French; various sign languages; formal systems such as propositional logic and mathematics; and the symbolic encoding formats used by digital computers.

It has often been supposed that the productivity and systematicity of languages is a unique consequence of certain features which are essential to being a language—in particular, having a compositional syntax and semantics (Fodor 2008.) This simply means having a finite stock of primitive or atomic symbols plus a finite set of rules specifying the ways in which the primitives can be combined to create meaningful complexes, whose meanings are a joint function of their structure and the meanings of their constituents. The claim that productivity and systematicity can only be gotten from these quintessential properties of languages, plus the claim that thought is productive and systematic, together make a strong case for the claim that thought must take place in something very like a language in these key respects—hence, of course, the LOT (Fodor & Pylyshyn 1988.)

It is possible to question the extent to which, or the senses in which, thought is, in fact, productive and systematic (Cummins 1996.) However, we need not pursue that line of inquiry here. Productivity and systematicity are, of course, only issues for EC and ELC to address insofar as they are in fact features of human thought. So I will just give a brief overview of the reasons Barsalou gives for thinking that PSS can account for these phenomena.

First of all, productivity follows naturally from Barsalou's claims about how simulators interact to create simulations. Recall the schematic story about how we token the thought that the cat is on the mat. We have simulators for CAT, MAT, and ON, each of which integrates stored traces of perceptions of these categories from various modalities. These mechanisms, in combination with one another, can create a schematic multi-modal representation of a cat, and a mat, arranged in a way that exemplifies the spatial relationship of on-ness. This complex simulation constitutes our thought that *the cat is on the mat*.

Similarly, we could re-apply our simulators for MAT and ON to simulate another mat on top of the (unfortunate) cat. The same basic mechanisms which allowed us to entertain the thought that the cat is on the mat also allow us to entertain the thought that the mat is on the cat which is on the (other) mat. It is clear that this recursive capacity could be extended to allow us to generate an extremely large range of potential complex simulations from a finite stock of simulators. This, of course, is productivity.

Systematicity follows equally straightforwardly from the core properties of PSS. In order to simulate the cat's being on the mat, we must possess simulators for CAT, for ON, and for MAT. But if we have such simulators, then we are automatically equally able to simulate the mat's being on the cat. Likewise, to entertain the thought that John loves Mary, we must possess simulators for JOHN, LOVE, and MARY. But, of course, possessing such simulators automatically enables us to entertain the thought that Mary loves John. Thus, we have systematicity.

Much more could be said on this topic, but this is not the place. I take this brief discussion to have demonstrated that ELC and EC have, in the form of Barsalou's development of PSS, rich resources with which to account for the productivity and systematicity of human thought. Thus, the amodal view does not enjoy an advantage in ELC over this respect either.

To conclude this chapter, then, the alleged advantages of the amodal view do not survive scrutiny. A sufficiently rich and carefully qualified version of ELC is, at least, the equal of its rival in accounting for our comprehension of language about the abstract and the unimaginable, and in accounting for productivity and systematicity. In the next chapter, I turn to the converse question: does ELC really have the advantages which it is often held to have over the amodal view?

5. Alleged ELC Advantages

Introduction

In this chapter I subject to closer scrutiny the main advantages which ELC is claimed to have over the amodal view. There are two such advantages—one empirical, and one theoretical. The alleged empirical advantage of ELC simply consists in the wealth of evidence from such disciplines such as psychology, linguistics, and neuroscience which are held to support the embodied view against its rival. Meanwhile, the alleged theoretical advantage of ELC is its claimed ability to solve the much-discussed Symbol Grounding Problem (SGP), which many theorists think is deeply problematic for the amodal view (e.g. Barsalou 1999; Glenberg & Robertson 2000.)

I will discuss these two alleged advantages in turn. Several issues arise in connection with the inference from empirical data to embodied theory. One of these is the question of whether the data are genuinely inconsistent with the amodal view, as some claim. Another is the question of whether the studies relating to language comprehension specifically are sufficiently well-designed to identify facts about comprehension, rather than some other cognitive process. Yet another concerns the charge that the empirical data only support a conclusion about the contents of cognitive processes, while ELC is a thesis about representational vehicles. I will discuss each of these issues in turn, and argue that the inference from data to theory withstands such objections. None of these issues seriously damages the empirical case for ELC.

There is, however, one objection which is harder to answer. This is the contention made by Mahon and Caramazza (2008) that there is no sense in talking about 'the amodal view', because there are many different amodal theories of language comprehension—and, moreover, that some of these theories actually predict the same sensorimotor effects that ELC does. For the purposes of the present chapter, I will ignore this complication and assume (as do most parties to this debate) that we can legitimately contrast ELC with 'the amodal view' of comprehension. However, it should be borne in mind that the conclusions I reach by this method are tentative, and may need revision in light of Mahon and Caramazza's arguments, which I will discuss in chapter 7.

The SGP is a somewhat different story. It takes some effort merely to specify precisely what the problem is supposed to be. At a first pass, it is the problem of how the representations tokened in higher cognitive processes such as language comprehension manage to convey meaning about things in the world. But this first pass leaves open many important questions, such as whether the SGP is a problem about meaning or about understanding, and how it relates to philosophical debates about mental representation and content determination. After a close examination of these issues, I will conclude that the SGP does not pose any special difficulty for the amodal view, because—as Shapiro (2011) argues convincingly—amodal symbols can be grounded. The resources which the ELC theorist must posit to solve the SGP are equally available to the amodal theorist. At this dialectical point, a parsimony argument may hold some hope for the ELC proponent—but even this is by no means decisive, and in any case is, strictly speaking, a distinct argument from the SGP.

The ultimate conclusion of this chapter, then, is that the extant empirical evidence does favour ELC over the amodal view, but that the SGP does not. A parsimony argument is a more promising line of theoretical inquiry.

Accounting for the Evidence

Recall the various empirical findings we have seen adduced in support of ELC. Merely understanding a sentence which implies motion relative to the body affects subsequent production of actual motion (Glenberg & Kaschak 2002.) Understanding a sentence which implies—but does not explicitly state—that an object has a certain position or orientation facilitates identifying the object when depicted with that position or orientation (Stanfield & Zwaan 2001.) Information about affordances which is not derivable from linguistic context predicts how sensible subjects will rate given sentences to be (Glenberg & Robertson 2000.) Reading category nouns activates areas of the brain implicated in perceptual and motor encounters with members of the relevant categories (Pulvermüller 1999.) This small and selective sampling of the wealth of data on offer leaves no doubt that inquiry into ELC constitutes a fruitful research program which has contributed enormously to our knowledge about the mind.

However, several theorists have recently sounded a note of caution. While not denying how impressive and exciting these empirical findings are, they are sceptical about the inference from such findings to the truth of ELC.

One such critic is Adams (2010) who notes that the systematic correlations discovered between language comprehension and sensorimotor processes are consistent with the claim that the latter are a mere causal consequence of the former, rather than a genuine constituent thereof. The fact that understanding language goes hand-in-hand with the activation of sensorimotor representations does not show that understanding language *is* the activation of such representations. So, he argues, the data do not pose a genuine threat to the amodal view, because they can be accounted for readily within that framework.

To a certain extent, Adams' point is well taken. Sometimes researchers sympathetic to ELC do claim that their results are inconsistent with amodal accounts of comprehension or higher cognition (e.g. Glenberg & Kaschak 2002.) This seems implausible and hyperbolic.

There is nothing inconsistent in the suggestion that language comprehension consists in the activation of amodal conceptual representations, which typically causes the activation of distinct but closely associated modality-specific representations. Of course, one would want the details of this story spelled out, but the same can be said about many extant ELC and EC proposals. So Adams is correct in this much: the empirical evidence adduced in support of ELC does not deductively, conclusively, or unequivocally demonstrate the truth of ELC or the falsity of the amodal view.

Once we set aside pretensions to conclusive falsification, however, the situation is somewhat different. The best kind of argument to make for ELC is an abductive one—that is, an inference to the best explanation. The way to do this is to enumerate all the available observations concerning language comprehension and then to attempt to show that ELC out-scores alternative explanations of these observations in terms of parsimony, consilience, fruitfulness, and other qualities of a good explanation. Once this sort of strategy is explicitly adopted, the prospects for arguing from empirical evidence to the truth of ELC look much better. I do not, in this thesis, provide a thorough and systematic abductive argument for the superiority of ELC to rival theories. And, as will be discussed in chapter 7, there are problems with inferring the truth of ELC from the evidence as things currently stand. However, while there may be problems with an abductive inference to the truth of ELC, the formal consistency of the evidence with alternative theories is not one of them.

As far as the prospects for an abductive argument are concerned, when it comes to the sorts of data we have been discussing, there seems to be a crucial difference between ELC and the amodal view. Granted, the latter can explain the various results post-hoc, but it is less clear that it can predict them in a principled fashion. This is, of course, exactly what ELC has done. Many of its predictions concerning the relationships between comprehension and sensorimotor processing have been confirmed empirically. As Barsalou (1999) argues, results which were predicted by one explanatory theory but can also be accommodated after the fact by a rival theory should not be considered to favour both theories equally, but should be considered to favour the theory which predicted them.

The second issue relating to arguing from data to theory is connected with Adams' suggestion that sensorimotor processing might merely be a causal consequence of higher cognitive processes such as language comprehension. Expanding upon Adams' arguments, Shapiro (2011) argues that the experimental designs used in some of the relevant studies are not sufficiently fine-grained to warrant the claim that the observed sensorimotor priming and interference effects occur as part of the comprehension process.

In particular, recall the experimental paradigm used in the study which showed the interaction between understanding motion-implying sentences and producing actual motion (Glenberg & Kaschak 2002.) Subjects had to read a sentence, then judge whether it was sensible or not, and subsequently move their hand toward or away from their body in order to make a 'yes' or 'no' response. Shapiro's suggestion is simple, and equally applicable to other studies which use this 'sensibility judgment' paradigm: we are not justified in concluding that the motor effects occur or are caused during the comprehension process, because they could just as well be a constituent or result of the process of making a sensibility judgment instead. Maybe subjects first understand the sentence, and then imagine or simulate what they have understood in order to decide whether it is sensible (Shapiro 2011.)

I think this is correct. Fortunately, Glenberg and his collaborators have since performed another study aimed at testing precisely this possibility, using more fine-grained methods from neuroscience. In this later study, subjects once again had to read sentences describing transfer of concrete objects, others describing transfer of abstract objects, and others not describing transfer at all. This much was the same as in the earlier study (Glenberg et al 2008.)

This time, however, Glenberg et al used a technique known as Transcranial Magnetic Stimulation (TMS) which can be used to evoke a response in the motor system, the modulation of which response can in turn be measured. This allows researchers to identify differences in modulation of the evoked motor response between stimuli.

In order to test the possibility that the motor system is only engaged *after* sentence comprehension, Glenberg et al designed the experiment so that the TMS pulse would be applied shortly after the presentation of the verb, while the subjects had not even completed *reading* the sentence. Obviously at this point they could not have finished comprehending it, so any motor response modulation measured must take place during, and not after, the comprehension process.

The results of the experiment confirmed the prediction made by ELC. Significantly greater modulation of the evoked motor response occurred during reading of the sentences describing transfer, both abstract and concrete, than during reading of the 'no-transfer' control sentences. This seems to settle the question pressed by Shapiro: the motor system is engaged during the comprehension process, and not merely during the post-comprehension process of making a sensibility judgment.

The third and final issue concerning data-to-theory arguments for ELC is raised by Weiskopf (2010a), who contends that some such arguments rely on a conflation of representational content with representational vehicles. This distinction—a mainstay of the philosophy of mind and cognitive science—is, intuitively, that between what is being represented and what is doing the representing. To give an example similar to Weiskopf's: consider a photograph of a cat sitting on a mat, the string of English words 'the cat is on the mat', and a line drawing of a cat sitting on a mat. Each of these is a different representational vehicle—a different thing which is representing something. But they all have the same content; they all represent the same thing (namely, the state of affairs of the cat's being on the mat.)

Without delving into technicalities, it seems plain that—as Weiskopf is at pains to stress—the content-vehicle distinction is a well-motivated and useful one. Why, then, does he think the arguments for ELC blur it?

His worries stem from Glenberg et al's use of the Affordance Compatibility Effect discussed in chapter 3. Recall, the Affordance Compatibility Effect is the finding that subjects reliably assign higher sensibility ratings to sentences describing scenarios which are afforded by the properties of the objects they describe than to sentences describing non-afforded scenarios. This holds true even when the two sentences are given very similar semantic values by amodal computational models such as LSA and HAL. Glenberg and Robertson (2000) argue that this is problematic for the amodal view because it gives us no understanding of why subjects' sensibility ratings would differ in accordance with affordance compatibility, while ELC—in particular the IH, with its emphasis on the derivation of affordances—explains (indeed, predicts) this result straightforwardly.

What, then, is Weiskopf's concern? Simply this: the empirical results only show that information about affordances is being mentally represented, but they do not show in what format or by what kinds of vehicles. This being the case, there is no reason to assume that they are being represented in modality-specific codes, as ELC claims. Information about what various objects afford could equally well be represented in an amodal LOT-style code and brought to bear on the tasks involved in these studies, with identical results. The studies, then, only establish the claim that content about affordances is represented by subjects, which falls far short of establishing the ELC claim that the representing vehicles involved are sensorimotor or modal in nature.

There are a few things to be said about this argument. The first, which may seem slightly pedantic but is nonetheless important, concerns Weiskopf's reference to it as his "second major argument against ELC" (2010, p.299.) This is misleading, since the point about the content-vehicle distinction is not, in any way, an argument against ELC. Rather, it is a criticism of one kind of argument *for* ELC. Even if totally successful, all it shows is that the Affordance Compatibility Effect cannot be used to support ELC. It leaves all the other arguments for ELC, as well as the internal coherence and plausibility of the position itself, untouched.

A second point to be made echoes one I made earlier about the relative degrees of support gained by, on the one hand, theoretical frameworks which predict empirical findings, and on the other, frameworks which merely accommodate them in an ad hoc manner. Glenberg and Robertson's (2000) study was specifically designed to test competing predictions of ELC and certain of its amodal rivals, LSA and HAL—which predicted that subjects should understand the paired afforded and non-afforded sentences very similarly. Setting aside those specific models, it is unclear on what basis the amodal view of comprehension generally would make a prediction that information about affordances would be accessed in comprehension. Granted, the empirical finding that information about affordances is represented in comprehension does not conclusively rule out the claim that such information is represented in an amodal format. However, given two alternative frameworks capable of explaining this finding, one of which predicted it and the other of which did not, then, as Barsalou (1999) urges, we should *ceteris paribus* prefer the predictively successful framework.

The final point I want to make about Weiskopf's criticism is that the possibility of affordances being represented in an amodal format does not generalise to the other kinds of sensorimotor information represented in comprehension. When the process of understanding language influences motor responses and performance in perceptual tasks, it is clear that the sensorimotor information being represented is being represented in the relevant modality-specific systems. In light of this, and insofar as we have good reason to think that affordances *are* represented in comprehension, it seems more parsimonious to suppose that they, too, are represented in modality-specific codes.[14]

Weiskopf's objection about the content-vehicle distinction, then, at most shows the untenability of one kind of argument for ELC—that which relies on the Affordance Compatibility Effect. It leaves ELC itself, and all the other arguments for it, untouched. Further, given the point about prediction vs. post-hoc explanation, it is not clear that it even deals a fatal blow to this one kind of argument.

I conclude that the objections I have discussed do not substantially harm the empirical case for ELC. Despite the formal consistency of the various empirical results with the amodal view of comprehension, when viewed from the point of view of an inference to the best explanation, they seem to favour ELC. I turn, now, to the major theoretical argument in favour of ELC: the infamous SGP.

The Symbol Grounding Problem

The SGP is presented by many theorists as a central motivation for research into ELC and EC, and as a strong reason for preferring these embodied frameworks to amodal views of higher cognition (e.g. Barsalou 1999; Glenberg & Robertson 2000; Anderson 2003.) Its canonical articulation is due to Harnad (1990) who introduces it via a re-worked version of Searle's (1980) famous Chinese Room thought experiment. The story goes roughly like this: Imagine landing in a foreign country, the language of which you do not speak. Disembarking at the airport, you

[14] Of course, we should note that the study in question utilises the same sensibility judgement paradigm which we have already seen is somewhat coarse-grained. This leaves open the possibility that affordances are not represented in comprehension, but only in the post-comprehension process of making a sensibility judgement. Unless and until some ingenious experimental design can discriminate between these two possibilities, the question remains open. But it is not really relevant here, since it amounts to a different criticism than Weiskopf's about the content-vehicle distinction.

look around and find yourself confronted with many signs bearing unfamiliar strings of unfamiliar symbols. Seeing your plight, a sympathetic local provides you with a dictionary in the local language. Opening this up, you look up some of the symbol strings and find, to your dismay, that they are defined only in terms of further unfamiliar symbol strings. Looking *those* up, in turn, you of course find the same thing again... and so on, ad infinitum. Never will this monolingual dictionary give you any understanding of the unfamiliar language. The moral of the story is that abstract symbols such as words cannot gain meaning merely through their relations with other abstract symbols: something more is needed to ground them, to establish their connection to extra-linguistic things.

Advocates of EC often take the SGP to demonstrate, or at least strongly suggest, a fatal flaw of the amodal view of higher cognition: that the abstract, language-like representations it posits as the stuff of thought cannot be meaningful because they are not grounded in our perceptual and bodily encounters with the things to which they refer. The embodied view is supposed to enjoy an advantage in this respect because the representations *it* posits as the stuff of thought are none other than stored traces of such encounters; the SGP does not arise because modal representations such as perceptual symbols are already grounded.

The first thing to ask about the SGP is how it relates to the philosophical problem of mental content. The way the SGP is often articulated in the literature on EC, it sounds like the exact same question: how do the representations manipulated in cognition get to refer to their specific referents? Of course, as Shapiro (2011) points out, there are many extant answers to this question in the philosophical literature. For instance, causal theories of representation hold that mental particulars refer to their referents by virtue of entering into the right kinds of causal relations with them (Dretske 1981) while resemblance theories of representation hold that some form of resemblance between mental particular and referent is what fixes the content of the former (O'Brien & Opie 2004.)

Given that the question of mental content has been so extensively explored by philosophers, the SGP presumably cannot simply be the same question re-stated as though it were a complete mystery. But sometimes this seems like what is going on. In particular, proponents of EC and ELC sometimes sound as though they are (i) assuming the truth of some kind of resemblance theory of mental content, and (ii) arguing that the embodied view is superior because it endorses such a theory while its amodal rival does not. Thus, Stanfield and Zwaan write:

> Perceptual symbol systems assume an analogue relationship between a symbol and its referent, whereas amodal symbol systems assume an arbitrary relationship between a symbol and its referent.
>
> (2001, p.153.)

In a similar vein, Glenberg et al write:

> The [amodal] symbols are arbitrary in the sense that there is no natural connection between the symbol and what it represents. For example, the word "chair" does not look, taste, feel, or act like a chair, and the word "eight" is not larger in any sense than the word "seven."
>
> (2005, p.116.)

These may simply be careless uses of language, but the distinctions involved are worth making, especially given the use of 'arbitrary' as a pejorative adjective in the literature on EC and ELC. The implication cannot be that mental particulars are incapable of representing categories or objects to which they stand in arbitrary (i.e. non-resemblance) relations. In order for this to be available as a premise in an argument for the embodied view, one would first have to establish the superiority of a resemblance theory of content. No such argument has been offered by any EC or ELC proponent, as far as I am aware, and certainly Harnad's popular thought experiment alone will not do the job.

Let us be clear, then: the issue is not one of content determination. A more plausible construal of the SGP, as Shapiro (2011) suggests, is that it is a problem about *understanding*, rather than about meaning. The question it is intended to raise is not about how the representations manipulated in thought acquire their specific contents. Rather, the question is about how the cognitive system manipulating the representations *knows* or *understands* what they refer to. This certainly seems to be what Glenberg and Robertson (2000) have in mind when they ask: “how does the system know what the symbols are about; that is, what it is thinking about?” (p.382.) We can now state the question more precisely. Assume the existence of a cognitive system manipulating symbols which represent things in the world, but which stand in arbitrary relations to those things. By what means can that system determine, for each symbol type it manipulates, what category of real-world things that symbol represents?

If this is the question, however, then it seems like a resemblance theory of representation has, in fact, been smuggled in through the back door. Consider: EC theorists argue that the amodal view falls foul of the SGP, while the embodied view does not, because the latter sees all cognition as consisting of the manipulation of representations which are modal in nature. But why does the SGP not arise for modal representations themselves? The implied answer seems to be: because they, unlike abstract symbols, are not arbitrary in the allegedly problematic sense. This suggests the following argument: a modal view of higher cognition solves the problem of mapping representations onto real-world referents because modal representations represent by some form of resemblance.

There is a different construal of the SGP as mapping problem, however, which does not rely on a resemblance theory of representation, even for modal representations. This is one which takes for granted the cognitive system's ability to map modal representations onto their referents, but then questions how a system of amodal representations could be mapped onto modal representations. This leaves open the question whether the representational codes used by the various modalities use any form of resemblance relation or not. On this reading, the embodied view's touted advantage of non-arbitrariness does not refer to relations between representations and things, but rather to relations between representations and representations—in particular, between the representations used in higher cognition and those used in perception and action. This seems like the most plausible resolution of the ambiguity surrounding the SGP.

So understood, however, it is doubtful that the embodied view enjoys any advantage with respect to the SGP. This is made clear by Shapiro (2011), who asks: why couldn't amodal symbols be grounded in the same way that perceptual symbols could be? The obvious reply—that perceptual symbols get grounding 'for free' because of their analogical relationships to perceptual states—will not work, for the reason that the simulators posited by Barsalou are multi-modal integration mechanisms. When you think about cats, or understand language about cats, according to PSS your CAT simulator recreates schematic representations—namely, simulations—of typical cat encounters in the various modalities. Of course, the auditory simulations it generates will bear an analogical relationship to your auditory perceptions of cats, because they are just re-enactments (albeit schematic) of those very perceptions. Similarly, the visual simulations it generates will bear analogical relationships to your visual perceptions of cats, etc. But there is no reason to suppose that your auditory representations of cats are analogically related to your stored visual, tactile, proprioceptive or other representations of cats. Assuming, plausibly, that the various modalities encode different information, in different formats, the very idea of a multi-modal simulator requires the existence of some mechanism distinct from the various stored modality-specific representations which can integrate them all and ensure that representations of the *same* category are re-created in all the various modalities. But what is this mechanism, if not an amodal symbol grounding device? When activated by a mere visual perception of a cat—say, in a photograph—your CAT simulator is, on this story, capable of mapping this to auditory, tactile, and other representations of cats, and *it does not do this by exploiting any kind of analogical relationship between its input and the representations onto which this gets mapped*. Simulators, then, are capable of symbol grounding in a way that may work equally well for abstract, arbitrary, amodal symbols.

In response to the suggestion that their need for multi-modal integration mechanisms commits them to the existence of amodal representations, embodiment theorists sometimes respond that these integration mechanisms are not amodal representations because they are not themselves the stuff of thought. Modal encodings are still what are doing the representational work; as Prinz says, "…occurrent cognitive activity and conceptual performance rely on activity within the modalities… If an amodal code exists, it works on credit rather than serving as the primary currency of thought." (2002, p.137.) Now, this is a perfectly acceptable response to the charge that EC theorists are forced to posit amodal symbols, and one which I will return to in chapter 7. But it will not work to deflect the present point, which is this: The amodal view is not disadvantaged with respect to symbol grounding because EC's symbol grounding (or multi-modal integration) resources could ground amodal symbols equally well. The representational status of the extra-modal integration mechanisms is irrelevant. If EC theorists must posit mechanisms for mapping information about categories across different forms of encoding, then a similar mechanism could map amodal information about categories onto modal information about them, thus grounding amodal symbols.

Another possible move for the EC theorist, at this point, is to appeal to a parsimony argument, along the following lines: It is true that amodal symbols could be grounded, given the machinery that EC/PSS already needs to posit. But since that machinery plus modal representations alone can do the work, why bother to posit extra (grounded) amodal symbols as well?

However, this response is inadequate in the present context, for two reasons. The first is that it misses the point: even if a parsimony argument could support EC over the amodal view, it is a change of subject. The parsimony argument does not establish any EC advantage with respect to the SGP.

The second reason that the parsimony argument may fail is that it is unclear that positing multi-modal integration mechanisms is more parsimonious than positing amodal symbols. Perhaps one could bolster this argument by showing that amodal symbols would require integration mechanisms as well. However, I will not investigate this strategy here[15].

For now, I conclude that ELC does not enjoy an advantage over its amodal rival where the SGP is concerned. The resources which ELC needs to posit in order to solve the problem could equally be posited to solve it by the amodal view.

However, as we saw earlier, the situation seems different when it comes to the empirical case for ELC. This appears[16] to constitute a genuine advantage. While the data can undoubtedly be accommodated by the amodal view, they provide much stronger support to the theory which predicted them before the fact and which continues to generate new, fruitful empirical predictions.

In the next chapter, I will turn from the dispute between ELC and its amodal rival to a different front. This is the clash between the relatively conservative, representationalist version of ELC which I endorse, and a more radical, 'enactivist' interpretation of the position, which has recently been advanced as a better explanation of the available evidence.

[15] I will mention it again briefly in chapter 7.
[16] Though, again, the arguments of Mahon and Caramazza (2008)—which I will revisit in chapter 7—should be borne in mind as a potential problem for this empirical verdict.

6. The Enactivist Challenge

Introduction

So far, this thesis has been considering the contrast between ELC and the amodal view—two very different accounts of linguistic comprehension. Despite their differences, however, there is a lot about which these two views agree. They are both committed to the claim that comprehension is constituted, at least partly, by the activation of mental representations of some kind, and they both hold that comprehension—like other cognitive functions—is wholly "intracranial", in Adams and Aizawa's (2001, p.56) term. That is, cognition takes place entirely within the bounds of "skin and skull" (Clark & Chalmers 1998, p.7.)

There is a more radical approach to cognition which denies both of these claims. This is the thesis of enactivism (mentioned briefly in chapter 2) which denies that cognition is best conceived of as operations on mental representations, and holds instead that it inheres in the dynamic relationships between an organism and its environment as the former acts skilfully in the latter (Varela et al 1991.) On this view, cognition is fundamentally knowledge-how, rather than knowledge-that; it consists of the exercise of practical abilities, rather than the representation of states of affairs (Van Elk et al 2010.) The contrast between enactivism and the version of ELC I have been defending is a stark reflection of the diverse nature of the broader 4E Cognition research programme, as outlined in chapter 2. Despite the fact that my version of ELC and the enactive approach both situate themselves under this banner, it may well be that ELC has more in common with the amodal view of higher cognition, when it comes to such fundamental issues.

Recently, a challenge has been mounted to ELC from proponents of enactivism. Van Elk et al (2010) argue that there are insurmountable problems facing a 'cognitivist' interpretation of the data concerning embodiment and language comprehension, and that these problems can and should be avoided by adopting an enactivist version of ELC instead.

In this chapter I respond to these arguments in detail. I will proceed as follows. In the second section, I will outline the enactivist view of cognition in general, including its intervention in the debate concerning the function of mirror neurons. Van Elk et al use this intervention as a model for their dialectical strategy concerning language comprehension. In the third and fourth sections, respectively, I will discuss the two arguments given for thinking that we should reject the cognitivist version of ELC in favour of the enactivist alternative. These relate to what Van Elk et al call the 'Simulation Constraint'—the question how we can understand language about things we have not experienced with our senses—and the 'Necessity Question'—the question whether activation in sensorimotor areas of the brain is necessary for understanding language.

I will argue that neither of these issues poses an insurmountable problem for the cognitivist version of ELC, and that it is not clear that the enactivist version of ELC enjoys any significant advantage with respect to these issues. Fundamentally, the claim that enactivist ELC has an advantage seems to rely on saddling cognitivist ELC with unnecessary assumptions about language—in particular, that there is a core meaning of words which is activated automatically regardless of context. I will argue that cognitivist ELC is just as well-placed as enactivist ELC to account for the context-sensitive nature of language, and, further, that enactivist ELC faces problems of its own. In the fifth and final section of the chapter, I will summarize and conclude the preceding discussion.

Enactivism and Cognitivism

The 4E Cognition movement, and the various claims known as 'embodied cognition', are often presented as a radical challenge to the orthodox practice and theoretical assumptions of cognitive science (Shapiro 2011.) Nowhere is this more accurate than when it comes to the enactivist view. Van Elk et al define enactivism as “the view that cognition emerges in the interaction between an organism and the environment, such that perception and action are co-constitutive of it” (2010, p.1.) Consider the narrow EC claim that activity in the brain’s sensorimotor systems is constitutively involved in higher cognitive functions. This is considered by some theorists to be quite radical (e.g. Adams 2010.) But it is positively conservative compared to the enactivist proposal that not just perceptual and motor activity in the brain but perception and action themselves are constitutively involved in cognition. On this view, the ability of organisms to orient themselves appropriately in relation to their environments and act skillfully cannot be said to be a causal consequence of those organisms’ cognitive capacities; rather, this ability *is* (at least partly) what those cognitive capacities are.

One of the most influential and detailed applications of the enactivist view is the sensorimotor theory of (visual) perception. On this view, perception is constituted by the active exploration of the environment, and the mastery of “sensorimotor contingencies”—reliable couplings in both directions between sensory input and action (O’Regan & Noë 2001, p.940.) One famous putative example of this is “presence-in-absence” (Noë 2009, p.472)—the phenomenon of looking at, e.g., a building from the front, only able to see the façade, and nonetheless having a sense of the presence or reality of the rest of the three-dimensional structure. This sense, according to the sensorimotor theory, is a function of the fact that the rest of the object is available to our sensory systems if we take certain actions. This is what is meant by a sensorimotor contingency.

Gaining a clear understanding of enactivism may be helped by considering what it is supposed to be an alternative to. Van Elk et al define cognitivism, the orthodoxy to which enactivism is meant to be a radical reaction, as “the theoretical approach that attempts to explain cognition in terms of the manipulation of discrete internal representations” (2010, p.1.)

This characterization, however, immediately raises questions. Whence the requirement that cognitivism be committed to the *discreteness* of representations? Certainly the *classical* cognitivist theory, the amodal view in the form of Fodor's LOT theory, holds that the mental representations manipulated in higher cognition are discrete. But there are other representational approaches to cognition which are consistent with cognition being entirely intracranial and ontologically distinct from skilful action, while not positing *discrete* representations. Connectionist frameworks, for instance, posit as representational vehicles either occurrent patterns of activation over a whole network of neuron-like units, or more durable patterns of connection strengths between those units (O'Brien & Opie 2002.) Computational models of phenomena relating to EC and ELC have been developed which utilize connectionist techniques; e.g. Flusberg et al's (2010) model of conceptual metaphor, and Joyce et al's (2003) model of visual cognition in Perceptual Symbol Systems.

These are thoroughly representational models of cognitive processing, consistent with a brain-bound ontology of cognition. But their representational vehicles are patterns of activation and patterns of connection strengths, neither of which are discrete phenomena. Not only are connectionist representations distributed across the whole network in this fashion, but they also exhibit superpositional storage—meaning that the very same vehicles are used to encode multiple items of information (Clark 1993.) This is about as far from a *discrete* representation as one could get.

So, if the cognitivism vs. enactivism dichotomy is meant to be mutually exclusive and exhaustive, and the former posits representations while the latter does not, then it cannot be true that cognitivism is committed to the discreteness of representations. Perhaps the most illuminating contrast is to be found in the answers given by the two positions to the following two questions. First, is action a constituent of cognition? Second, does cognition involve internal representations of the external world? Enactivism answers yes to the first question, no to the second; cognitivism answers no to the first question, and yes to the second. So, enactivism is the claim that cognition consists in appropriate sensorimotor couplings of an organism to its environment, while cognitivism is the claim that cognition consists in some kind of operations performed over some kind of internal mental representations of an organism's environment.

In proposing their enactivist version of ELC, van Elk et al (2010) draw inspiration from the enactivist interpretation of neurological evidence concerning mirror neurons. Briefly, mirror neurons are a class of neurons which fire during both performance of certain actions by us and our observation of performance of those same actions by others (Gallagher 2005.) So, for example, a given neuron might fire in my brain whenever I activate a precision grip *or* observe someone else activating a precision grip (Gallese & Goldman 1998.)

One interpretation of the available information concerning mirror neurons is that they enable social perception by subserving a simulation process (*ibid*.) The idea is that when we see someone performing some action, we are able to understand their behavior and intentions thanks, at least in part, to a mirror neuron-based simulation of performing that same action ourselves[17].

However, the simulation interpretation of the mirror neuron evidence is controversial. One objection which has been raised stems from the observation that higher-level intentions are under-determined by 'motor intentions'. That is, simulating the activation of a precision grip can only help us to understand that someone is attempting to grasp an object of a certain size; it cannot help us understand *why* they might be doing that, what pragmatic or other signals their action might be intended to convey, and so forth (Van Elk et al 2010.)

In light of this and other worries about the simulation interpretation, an enactivist theory of the social-perceptual function of mirror neurons has been proposed. This involves re-

[17] Another line of thought holds that mirror neurons play an important role in learning, specifically by imitation (Gallese & Goldman 1998.)

conceiving what the phenomenon of social perception is. In line with the core enactivist claim that cognition is constituted by skillful action, this theory claims that social perception is constituted, at least in part, by *social interaction* (De Jaegher 2009); it is not a mere matter of interpreting sensory input and generating simulations or other representations on that basis, but rather it consists of couplings between sensory stimulation and the right kinds of motor outputs. It is worth quoting van Elk et al on this point:

> Mirror neurons…should be interpreted as contributing to the processing of the perceived behavior of others for the direct purpose of social *interaction*. The idea here is to think of social perception as an enactive process involving sensorimotor skills and not as mere sensory input processing. This idea is borrowed from enactive theories of perception according to which perception involves active engagement with the world rather than mere passive reception of information from the environment…
>
> (2010, p.3)

If we consider the dialectical structure of the mirror neuron debate, it is not hard to see why it is an inspiring model for van Elk et al. It begins with empirical evidence showing the activation of sensorimotor areas during a particular cognitive process. A theoretical account of this evidence is proposed according to which the cognitive process in question is underwritten by sensorimotor simulation. However, serious difficulties are raised for the simulation account, opening the door for an alternative, enactivist theory, which promises to account for the sensorimotor activation data in a way that avoids the difficulties of the simulation account—viz., by identifying the cognitive process in question with (among other things) actions of the entire organism.

In this spirit, van Elk et al propose an enactivist account of language comprehension. On this view, comprehension consists in skills instantiated in sensorimotor couplings. It is not *knowledge that* certain words or sentences mean such-and-such, but rather *knowledge how* to generate appropriate behavior in response to linguistic input. To use their example, when you are sitting in a restaurant after having finished your meal and a waiter asks you if you are alright, your understanding of this utterance consists in your ability to respond along such lines as: *yes, thank you, I'm fine; can I get the bill?* Once again, a relatively lengthy quote is warranted:

> Learning how to understand language is learning how to couple specific linguistic inputs to specific actions. These actions may be immediate but they may also be in the more distant future (e.g., as in understanding the sentence "the election will be on May the 5th"). They may also be only "virtual" in the sense that understanding an utterance only involves being disposed to act in certain ways given certain circumstances… In short, on an enactivist account, language comprehension can be described as procedural knowledge—knowledge *how*, not knowledge *that*—that enables us to interact with others in a shared physical world.
>
> (2010, pp.4-5.)

This is enough, for now, to have a fairly clear sense of the alternative conception of comprehension which van Elk et al think is demanded by the data. Time, now, to look at the problems which they think face a cognitivist version of ELC.

The Necessity Question

The first difficulty which van Elk et al raise for cognitivist ELC is the question whether activation in modality-specific brain regions is necessary for language comprehension. This is primarily an empirical problem; a potential falsifier of cognitivist predictions. Van Elk et al argue that cognitivist ELC predicts two things: firstly, that damage or disruption to sensorimotor processing should cause corresponding deficits in comprehension, and, secondly, that comprehension should always be associated with rapid and automatic activation in motor areas. However, they say, the empirical verdict on these questions is still inconclusive. The data on conceptual and comprehension-related deficits stemming from damage to sensorimotor areas is equivocal, and awaits further research; and the same is true of the data concerning rapid motor activation during comprehension (particularly, during comprehension of action verbs—a case much discussed by van Elk et al.) According to them, this constitutes, at present, a lacuna in the empirical case for cognitivist ELC, which could, depending on the outcome of future inquiry, amount to a fatal flaw (Van Elk et al 2010.)

The first thing which can be said in response to this argument is that we need to be careful about what exactly cognitivist ELC predicts. Remember that words and sentences are comprehended via the construction of multi-modal simulations, which may be more or less schematic, and incorporate different amounts of information from different modalities. Particularly, in accordance with Dove's (2011) Partial Dis-embodiment Account of the comprehension of abstraction, the representation of linguistic context in the form of simulations of written and spoken words can also play a role in comprehension. Thus, what kind of deficits cognitivist ELC predicts should result from damage to sensorimotor areas is a delicate question. Suppose, since birds are primarily encountered visually, that the bulk of the work in comprehending language about birds is done by the generation of visual simulations. This does not mean that auditory simulations, motor simulations, and dis-embodied linguistic simulations are not also playing a role. Therefore, if a subject's visual cortex is damaged, we should expect that their ability to comprehend language about birds will be somewhat impaired. However, it is not clear that we should predict it will completely fail. Consistently with my arguments in chapter 4 about abstraction and degrees of comprehension, it is possible that such a subject will understand language about birds less richly or fully, but still to a certain extent, and perhaps that dis-embodied linguistic simulations will pick up some of the slack left by the absence of visual simulations.

The second way of responding to this argument—strictly speaking, an elaboration and extension of my first response—concerns the issue of rapid and automatic activation of motor areas, in particular during the reading of action verbs. The issues involved here will be treated in more depth when I discuss the context-sensitive nature of language below. For now, suffice it to say that the multi-modal nature of the simulation posited by ELC proves, once again, to be an explanatory advantage. To be sure, we should expect that in the vast majority of cases, comprehension of verbs such as 'kick' and 'take' would be primarily underwritten by motor simulations. But, of course, we do not have only motor experiences of kicking and taking. Most of us also have stored visual perceptions of these kinds of actions, as well as linguistic information, and possibly, to a lesser extent, information from auditory and other modalities. Also, according to Barsalou's theory of PSS, context is an important factor in determining exactly what kind of simulation is generated of a given category on a given occasion. So we may reasonably expect that the kinds of sensorimotor information activated, and the strength of activation, will vary across occasions as well as individuals in the comprehension of a given word or sentence.

My third and final response to the argument from the Necessity Question is to ask how, exactly, adopting enactivist ELC is supposed to help. The suggestion seems to be that enactivist ELC would withstand the potential finding that motor information is only sometimes activated during comprehension, or by the potential finding that sensorimotor lesion patients can still understand language, because it is well-equipped to account for these findings by its commitment to language as a flexible and context-sensitive phenomenon. Consider:

> An advantage of an enactive approach is that it allows for the fact that language comprehension is a context-bound phenomenon that is dependent on the relation between the organism and the context in which the organism is acting. Cognitivist embodied approaches often make the implicit assumption that there is a core meaning of words that can be specified in terms of a specific representational vehicle. More specifically, cognitivist embodied approaches to language processing seem to imply that the sensorimotor representations that are activated in association with the processing of words occur relatively fast, automatic [sic], and in a bottom-up fashion.
>
> (van Elk et al 2010, p.5.)

The two assumptions mentioned should be carefully distinguished, however. It is true that cognitivist ELC endorses the claim that the initial activation of sensorimotor representations in comprehension occurs "relatively fast, automatic[ally], and in a bottom-up fashion". But this does not, by any means, imply the existence of a "core meaning of words that can be specified in terms of a specific representational vehicle".

Attention to the details of both Zwaan's theory of IEF and Barsalou's theory of PSS can help us understand how the one claim does not entail the other. Recall, from chapter 3, that IEF claims comprehension consists of three overlapping and mutually modifying stages: Activation, Construal, and Integration. Activation involves the diffuse and under-specified tokening of various representations associated with a given word or morpheme, while in Construal, greater specificity is achieved through the imposition of constraints by the wider linguistic context—i.e. the word, clause or sentence in which the initial word or morpheme is embedded. So, for instance, when understanding the sentence *the ranger saw the eagle in the sky*, on first reading the word *eagle*, all sorts of eagle-related representations from different situations and modalities will be accessed in a preliminary fashion. Subsequently, the activation of representations of *in the sky* will serve as an attractor which will preference representations of eagles with outstretched wings, resulting in those being more strongly activated, and other representations of eagles with folded wings losing activation strength.

This shows that the assumption of “relatively fast, automatic, and… bottom-up” activation of sensorimotor representations in comprehension does not commit one to the existence of a core meaning of words. The stored representations associated with any given word are many and varied, and which are activated on a given occasion will depend on many factors, including linguistic and extra-linguistic context. The latter is especially important, because it is so rich and varied. It includes many factors emphasized by enactivism, such as social context, goals, and the like. Consider just one simple example: a smoker, trying desperately to quit, who reads or hears the word ‘smoke’. I conjecture that such a person would activate many more cigarette-related representations (as opposed to, say, bonfire-related representations) than another person reading the same word—because of the fixation (conscious or otherwise) on an object of addiction which often happens during quitting attempts.

All this amounts to pointing out that if a common network of perceptual symbols is used in all the higher cognitive functions, including planning, deliberation, and so forth, then the total context biasing the activation which occurs during a given instance of comprehension can be extremely complex and multifaceted. The IEF’s ‘first stage’, Activation, never occurs in a vacuum; there is always some prior cognitive context which will bias the interpretation of a given word in a certain direction—a direction which may be subsequently altered.

Barsalou, in his development of the theory of PSS, also strongly endorses the idea that the construction of multi-modal simulations—and, hence, the comprehension of language—is highly flexible and context-dependent:

> Because a perceptual symbol is an associative pattern of neurons, its subsequent activation has dynamical properties. Rather than being reinstated exactly on later occasions, its activations may vary widely… Different contexts may distort activations of the original pattern, as connections from contextual features bias activation toward some features in the pattern more than others. (1999, p.584.)

In light of these explicit commitments of theories such as Zwaan's and Barsalou's, it is somewhat puzzling why van Elk et al should think that cognitivist ELC is unable to account for the context-sensitive nature of language comprehension. Its ability to do so is touted by its proponents as one of its major advantages over the amodal view, in the analysis of results from experiments which demonstrate the contextually-determined nature of the sensorimotor representations activated in comprehension. For instance, remember the study in which subjects were quicker to verify a horizontal pencil than a vertical one, after reading 'John put the pencil in the drawer', and vice versa after reading 'John put the pencil in the cup' (Stanfield & Zwaan 2001.) Obviously the pencil is being represented differently in comprehending these two different sentences, and the difference in representations of the meaning of this word is determined by context. The fact that cognitivist ELC predicts this contextually-determined difference in simulation in a principled fashion is the entire point of the study. Cognitivist ELC is extremely well equipped to account for the context-sensitive nature of language comprehension.

Moreover, while it does not, strictly speaking, relate directly to the Necessity Question, there is another issue about the flexibility of language comprehension which is worth mentioning. This issue seems far more difficult for the enactivist proposal. It is simply that there are many occasions in which we seem to comprehend language without exercising any abilities to respond, in the ordinary sense of physical or verbal action. It is possible to attenuate the notion of skilful action to account for such cases, and to characterize the relevant abilities as dispositions, but if one takes this path, then it starts to look inevitable that comprehension is underwritten by the tokening of mental representations of the world, even if one wishes to define it in other terms.

To see this, recall this passage from van Elk et al, quoted earlier in this chapter:

> Learning how to understand language is learning how to couple specific linguistic inputs to specific actions. These actions may be immediate but they may also be in the more distant future (e.g., as in understanding the sentence "the election will be on May the 5th"). They may also be only "virtual" in the sense that understanding an utterance only involves being disposed to act in certain ways given certain circumstances…
>
> (2010, pp.4-5.)

Consider briefly the specific sorts of actions to which one might couple the linguistic input *the election will be on May the 5th*. They are innumerable and extremely varied; they are verbal, physical, and cognitive. One might respond *May the 5th* if asked when the election will be; one might respond *the first week of May* if asked in what week the election will be. One might go to a polling booth on May the 5th, if one is a conscientious citizen, or go to the pub on May the 5th, if one is not. One might make a note in one's diary to update one's address on the electoral roll before it is too late, or make a mental note to remind one's 18-year-old son to enroll to vote, or—depending on what sort of parent one is—make a mental note to take one's 18-year-old son to the pub on May the 5th.

Obviously I could enumerate such actions indefinitely, but I take it that is not required to make the point. The diverse and open-ended nature of the arbitrarily temporally and spatially distant actions which could constitute a comprehender's total set of appropriate responses to the sentence about the election is surely not determined by any simple or straightforward sensorimotor coupling. It is undeniable that when someone hears and understands the election sentence, some durable change occurs in their cognitive profile: some item of information is stored which is capable of interacting with other items of information to produce sensible responses. Until a viable alternative (which is difficult to even imagine) is presented, the most parsimonious conclusion is that the item of information which is stored upon comprehension of this sentence is some kind of structure (discrete or distributed) which, in some format or other, *represents the world as being such that the election will be on May the 5th*.

Instances of comprehension which do not demand an immediate response cry out for a representational analysis. This is made clear by the claim above that the appropriate responses "may also be only "virtual" in the sense that understanding an utterance only involves being disposed to act in certain ways given certain circumstances" (van Elk et al 2010, p.4.) What are "virtual" actions, if not representations of actions? It is logically coherent, of course, to identify comprehension with the disposition to act in certain (appropriate) ways, etc., and to treat the internal neural or cognitive mechanisms underwriting this disposition as a black box. But it is difficult to see how this amounts to anything more than an ill-motivated return to behaviourism. This is the same issue which posed problems for the Concept Pragmatist-inspired analysis of comprehension: once one has philosophically identified some phenomenon with a disposition, or ability, or whatever, one can still sensibly ask how that disposition or ability is realized, and this seems to lead rapidly and inexorably to the manipulation of mental representations in the brain.

I conclude, then, that the Necessity Question does not pose any special problem for cognitivist ELC, or furnish us with any reason to reject this view in favour of the enactivist alternative. The cognitivist view is very well-equipped to account for the flexible and context-sensitive nature of comprehension, while the enactivist view faces its own problems about accounting for comprehension which causes complex dispositions to act in various ways best explained in representational terms. I turn, now, to the second argument for the superiority of enactivist ELC to cognitivist ELC: the Simulation Constraint.

The Simulation Constraint

According to van Elk et al, while the Necessity Question is a contingent, empirical matter, the Simulation Constraint is a “more principled problem” for cognitivist forms of ELC (2010, p.2.) It amounts to the question how we can understand language describing things which we have not experienced. This includes abstract language, as well as concrete language referring to things which we simply happen not to have encountered personally.

It is easy to see why the Simulation Constraint might be thought to constitute an advantage for enactivist ELC; it is a problem which, on that view, simply does not arise. Because enactivism does not conceive of cognition in terms of representation, but rather in terms of the ability skillfully to respond to stimuli, the enactivist account has no need to explain how we can represent or simulate the meaning of language which goes beyond our own experience. All it needs to explain is how we can learn to respond appropriately to receiving such language as inputs.

However—and this will come as no surprise to the attentive reader—the Simulation Constraint is no reason to reject cognitivist ELC, because cognitivist ELC has ample resources with which to answer the question it poses. I examined these resources in detail in chapter 4, when I discussed the alleged advantages of the amodal view relating to abstract and other language describing things beyond our own sensorimotor experience. The Simulation Constraint is just that same problem, being offered now as an alleged advantage of the enactivist view over cognitivist ELC. My response now is the same as it was then. Cognitivist ELC can account for our ability to comprehend language about things we have not experienced by deploying the combined resources of the Metaphorical Account, the Introspective Account, the Action Schema Account, the Situational Account, and the Partially Dis-embodied Account.

At this point, one might wonder whether the Simulation Constraint does not still constitute some kind of advantage of the enactivist proposal. Even though cognitivist ELC has resources to account for it, the thought runs, perhaps we should nonetheless prefer a theory which does not need to appeal to five different accounts in order to explain comprehension, let alone introducing the idea that comprehension admits of degrees, and other such complications.

My response to this line of thought is twofold. Firstly, until enactivist ELC is developed in more detail, we have no reason for thinking that its story will be any less complex or multi-faceted. According to enactivist ELC, comprehension is a skill. Therefore, given only the plausible assumption that skillfulness admits of degrees, it will also entail that comprehension admits of degrees. Further, suppose I am right about direct, straightforward sensorimotor couplings being insufficient to account for the diverse and open-ended nature of the sets of skilful responses required even by a simple sentence such as *the election will be on May the 5th*. In this case, a fully worked-out enactivist account of comprehension will have to specify the nature of the mechanisms involved in generating such sets of responses; and there is certainly no *a priori* reason to think that such mechanisms will be more simple than the five different accounts of comprehension recruited by the cognitivist version of ELC which I am defending.

Secondly, even if enactivist ELC enjoys some advantage with respect to the Simulation Constraint, it is hard to imagine that this will outweigh the disadvantage which accrues to it by virtue of its rejection of mental representations. Of course, one can reasonably quarrel about the nature of such representations, and that is what I have been doing for most of this thesis. But when we consider the undeniable human ability to generate indefinitely many kinds of behavior which relate sensibly to the content of sentences they have heard and understood decades earlier and continents distant, we should be dismayed by the prospect of attempting to explain this ability by abandoning the idea that our brains—somehow or other—genuinely represent the world as being the way sentences describe it to be.

Conclusion

Neither the Simulation Constraint nor the Necessity Question shows that we should abandon cognitivist ELC in favour of enactivist ELC. The argument from the Simulation Constraint can readily be answered by pointing to the resources developed in chapter 4 to account for the comprehension of abstract and other language describing things we have not experienced. And the Necessity Question is only a problem for cognitivist ELC if one saddles the latter with assumptions about language which are directly contradictory to its stated principles. Cognitivist ELC is very well-equipped to account for the flexible and context-sensitive nature of language; perhaps better equipped than the enactivist view, when we consider the difficulties engendered by rejecting mental representations as explanatory posits. Therefore, we should retain the idea that language comprehension consists of the construction of sensorimotor simulations.

7. Conclusion and Future Directions

The ELC vs. Amodal Debate

It is now time to return to our original question: how do we understand words and sentences? We have seen many good reasons to believe the answer that our brains perform the task of converting perceived symbolic forms into meaningful thoughts by generating sensorimotor simulations.

In chapter 2, we saw that, without ignoring its historical antecedents, this claim can and should be distinguished from various others made by theorists in the 4E Cognition movement. We can hold that understanding is simulating without abandoning mental representations; without explaining cognition purely in dynamical systems terms; without holding that the mind extends beyond the brain; and so forth. The task of assessing the arguments and evidence for and against ELC was made much more tractable by considering the position on its own merits.

Chapter 3 gave us a good sense of the strength and richness of the theoretical and empirical case for ELC. Barsalou’s theory of PSS, in particular, provides an indispensable theoretical background, including its technical definition of simulation as the activation of more-or-less schematic sensorimotor representations—with or without accompanying phenomenology—and its characterization of simulators as mechanisms integrating our knowledge about categories from various modalities. Meanwhile, empirical findings from neuroscience show close relationships in both directions between performance in conceptual tasks and activity in modality-specific brain regions, while psychological experiments reveal systematic interrelationships between performance on conceptual tasks and sensorimotor variables.

With the framework of PSS established, we examined the details of two different ELC theories. The first of these was Glenberg's IH, with its emphasis on affordances and action (which I argued needs refinement) and its three stages of indexing words to perceptual symbols, deriving affordances from those symbols, and meshing the affordances into a simulation. The second was Zwaan's IEF, with its three overlapping, mutually modifying stages of Activation, Construal, and Integration, reflecting the dynamic and temporally unfolding nature of comprehension. Consistently encouraging results have been found by behavioural studies aimed at testing these theories—among them the Action, Appearance, and Affordance Compatibility Effects.

In the final section of chapter 3, I distinguished two different dimensions along which interpretations of ELC could vary. When we disambiguate the claim that 'understanding is simulating', we can construe the term 'simulating' in three different possible ways. It could minimally mean the activation of sensorimotor representations, as in Barsalou; or it could mean the construction of a full-blown situation model, as in Zwaan; or it could mean the construction of a full-blown situation model incorporating information about affordances and actions, as in Glenberg. Further, however we construe 'simulating', we can ask—as Weiskopf points out—whether simulating in the relevant sense is constitutive of all or some instances of comprehension, or a mere unnecessary aid to all or some instances.

With this logical geography in mind, chapter 4 took us on a tour of alleged advantages of the amodal view. I argued that ELC has ample resources to account for the comprehension of abstract language, in the form of the Metaphorical Account, the Introspective Account, the Action Schema Account, the Situational Account, and the Partially Dis-embodied Account. Although adopting Concept Pragmatism does not ease the explanatory burden on the ELC proponent in the way one might think, the related strategy of incorporating the assumption that comprehension admits of degrees does help. It makes the problem of abstraction even more tractable, yielding the conclusion that the amodal view enjoys no advantage in this respect.

As regards the comprehension of the unimaginable (or un-simulable), many of the same resources can be recruited to account for this as for the comprehension of abstract language. Further, the minimal or truth-conditional view of comprehension involved in this objection does not constitute an advantage of the amodal view, because there are good reasons to think that comprehension is a richer phenomenon than this, *a priori* considerations notwithstanding. Another reason why the minimal view cannot help the amodal theorist here is that it faces a dilemma. On one horn, an infinite regress threatens, with language understood in terms of linguiform truth conditions which must themselves be understood, and so on. On the other horn, comprehension of truth conditions grounds out in something non-linguiform, such as sensorimotor simulation. Finally, any purely syntactic computation of truth conditions can be accounted for by Dove's Partial Dis-embodiment proposal, according to which sensorimotor representations of natural language mimic many of the properties of an arbitrary symbol system without needing to actually postulate an amodal code or LOT.

The final argument of chapter 4 was essentially a recapitulation of part of Barsalou's defence of PSS. The amodal view does not enjoy an advantage with respect to the productivity and systematicity of thought, because these fall naturally out of the properties of simulators as defined by Barsalou. Indefinitely complex simulations can be constructed by recursively applying simulators to one another, and possession of the simulators required to entertain any complex thought such as *Mary loves John* automatically provides the ability to entertain systematically related complex thoughts. In short, I argued that the amodal view simply does not have the advantages which are claimed for it over ELC and EC, because the latter two views are very well-equipped to account for our comprehension of the abstract and the unimaginable, as well as productivity and systematicity.

In chapter 5, we looked at the two main alleged advantages of ELC over the amodal view: the empirical data concerning embodiment, and the Symbol Grounding Problem (SGP). I here endorsed Barsalou's argument that, while the amodal view may well be able to account for the empirical results, they should be taken as favouring ELC because it predicts them in a principled fashion. However, I then argued that the SGP, properly understood, does not constitute an advantage of ELC as is often claimed. This is because all the resources which proponents of EC and ELC deploy to solve the SGP could, as Shapiro argues, equally well be used for the grounding of amodal symbols. An open possibility is that a parsimony argument could favour ELC over the amodal view, but this remains to be developed, and depends upon the claim that postulating multi-modal integration mechanisms is more parsimonious than positing grounded amodal symbols.

At this point, the scales look tilted in favour of ELC purely by the weight of the empirical evidence. The two rival accounts of language comprehension look pretty evenly matched from a theoretical perspective—both can equally well account for comprehension of abstraction, productivity, systematicity, symbol grounding, and the like. But predictions of the embodied view have been borne out again and again, in a convergent way, in psychology, neuroscience, and linguistics. Given the equality of the two views in other respects, the strength of the empirical case seems to be a very good reason for believing that ELC is the correct account of comprehension.

Before concluding, I will briefly revisit the arguments purporting to show that the empirical evidence does not favour ELC over the amodal view. Deeper consideration of these arguments shows that inferring the truth of ELC from the extant data may not be as straightforward as Barsalou's response suggests. This has implications for the next steps that should be taken in the debate over ELC.

Theory Comparison Issues and Directions for Future Research

Recall the argument Adams (2010) gave in relation to the data: The data cannot be taken as a straightforward confirmation of ELC or as a falsification of the amodal view because they are consistent with the amodal view. It is coherent, he argued, to claim that comprehension consists of the tokening of amodal symbols, and that any accompanying sensorimotor processing is a mere contingent causal consequence, subsequent to and distinct from the comprehension process itself.

In response to this argument, I conceded that such a claim is indeed coherent and that the data therefore do not conclusively falsify the amodal view or prove the truth of ELC. However, I suggested that the best way to argue for ELC is by an inference to the best explanation, and echoed Barsalou's contention that ELC gains a significant advantage over the amodal view by predicting the empirical results in a principled way, while the amodal view merely accommodates them post-hoc.

However, in the course of making a similar argument to Adams', Mahon and Caramazza (2008) explicitly and forcefully argue that the amodal view does not merely accommodate the results post-hoc. They point out that there are many extant amodal theories of language comprehension, and so there is never any such thing as 'the' unique amodal prediction. Rather, different amodal theories will yield different empirical predictions, so it always needs to be specified which theories are being compared. Having made this point, Mahon and Caramazza sketch the outline of an amodal theory of conception and comprehension which they say *does* predict the empirical findings in a principled way. In brief, their theory—which they call the hypothesis of "grounding by interaction" (2008, p.67)—claims that abstract, symbolic, amodal representations do exist, and constitute the representational vehicles accessed in comprehension and conception. However, it also claims that these amodal representations serve, like Barsalou's simulators, to integrate information about categories from different sensorimotor modalities, and to activate this sensorimotor information on specific occasions for functional purposes. This theory holds that the sensorimotor processing which occurs during comprehension is the result of activation spreading from amodal conceptual representations to sensorimotor regions.

The existence of a coherent theory which holds that comprehension consists in the tokening of amodal symbols, but which also predicts systematic two-way interactions between comprehension and sensorimotor processing, is troubling for my argument. It raises the possibility that the kind of empirical evidence I have discussed—on which my entire case now rests—does not support an inference to the truth of ELC at all.

There are four things I want to say about this development. First, even if the extant evidence does not discriminate between ELC and amodal theories, I have still shown at least that ELC is at no theoretical *disadvantage* with respect to its rivals. Systematically detailing the rich resources with which it can account for the comprehension of abstraction and other such phenomena, as I did in chapter 4, is a major step forward in the development of the embodied approach to language. It may be that further such steps will ultimately tip the scales in favour of ELC on theoretical grounds, even if the issue cannot be settled on the basis of the current evidence.

Second, one theoretical move which might tip the scales is the development of a parsimony argument for ELC, as I briefly suggested in chapter 5. Barsalou (1999) and Prinz (2002) deploy parsimony arguments for EC/ELC, and their efforts could be extended. In particular, as I mentioned in chapter 5, this would require exploring the issue of the multimodal integration mechanisms which ELC is forced to posit. If amodal theories are forced to posit distinct integration mechanisms in addition to their amodal symbols, then ELC—which must posit integration mechanisms, but not symbols—may be a more parsimonious theory. This is one line of inquiry which could be pursued.

Third, some of the ELC studies explicitly aim to compare and contrast it with specific amodal theories, and the results of these would seem immune from Mahon and Caramazza's arguments. For example, Glenberg and Robertson (2000) set out to test the IH against the statistical amodal models LSA and HAL. They derived certain predictions about meaning and sensibility judgements from LSA/HAL, and distinct predictions from the IH. They then ran the experiment and found that the results were as predicted by the IH. Of course, it is possible to dispute the details of Glenberg & Robertson's derivation of predictions, or interpretation of their results (Burgess 2000.) My point is simply that the Mahon and Caramazza arguments do not preclude the possibility that existing empirical evidence shows specific ELC theories to be superior to specific amodal theories.

If this much has been shown, then at least one thing which has been achieved by research into ELC is the identification of constraints on a plausible amodal theory of comprehension. The amodal approach may not have been shown to be untenable in general, but certain forms of it may have. This thought leads directly into my fourth and final point, which is this: one fruitful direction for the debate over ELC may be away from the question 'can we do without postulating an amodal code?' and toward the question 'what positive contribution might sensorimotor simulation make to comprehension?'

A similar suggestion has recently been made by Marghetis and colleagues (Marghetis et al, under review.) Starting from the possibility that the current evidence does not discriminate between modal and amodal theories of comprehension, they argue that no evidence ever could so discriminate, because the modal/amodal dichotomy is ill-conceived and unsustainable. They argue that there is no viable characterization of what it is for a representation or system to be modal, and that we should therefore abandon questions of representational code or format in favour of questions about the functional roles of neural circuits: "Which neural circuitry, performing which functions for perception and action, also serves what functional role(s) in what higher cognitive processes (if any)?" (*ibid*, p.4.)

I agree with the spirit of this proposal. The way to move forward is definitely to start asking more positive and functional questions about the mechanisms of comprehension. But problems with the 'modal vs. amodal' question should not lead us to eschew all questions of representational format. In particular, while Marghetis et al may be correct that the modal/amodal distinction "dissolves at the level of neurons and their functional properties" (*ibid*, p.20) there is a distinction of representational format which does not: the symbolic/analogue distinction.

Very roughly, symbolic representing vehicles are those which stand in arbitrary relations to what they represent, and analogue representing vehicles are those which do not. Analogue representations stand in relations such as structural resemblance to their referents (Copeland 1993; O'Brien & Opie 2004.) The debate about the vehicles of comprehension might well be re-cast in analogue vs. symbolic terms, with advocates of ELC being, as much as anything, concerned with the potential importance of analogue representation to understanding.

Some ELC theorists explicitly hold that the perceptual symbols underwriting comprehension are analogically related to their real-world referents (e.g. Glenberg 1997, Stanfield & Zwaan 2001.) Others are more cautious, claiming that perceptual symbols are analogically related to the sensorimotor states from which they arise, but refraining from claiming that the latter stand in any analogical relations to their referents (Barsalou 1999.)

Long out of favour, the idea that analogicity or resemblance might be the basis of mental representation has recently been making a comeback in philosophy (e.g. O'Brien & Opie 2004; Bartels 2006; Isaac 2012.) When discussing the Symbol Grounding Problem in chapter 5, I commented that advocates of ELC cannot assume the truth of a resemblance theory of mental representation in arguing for their position. That much is true; but perhaps they can argue for the truth of such a theory, as others have. ELC and the resemblance theory may well be natural allies. Perhaps, with a little more theoretical work, advocates of ELC can make a case along roughly these lines: resemblance is the (sole) basis of original meaning or non-derived mental content; sensorimotor states represent by resemblance, while any dis-embodied or disembodied symbols which exist do not; so, any symbols which exist cannot be meaningful except derivatively, by virtue of their connections with analogue (sensorimotor) representing vehicles.

This line of thought suggests that the truth in ELC is not the rejection of symbols but the insistence that symbols without sensorimotor content are blind. As my epigraph from Haugeland says, words [symbols] are “merely money” and must be cashed out in goods which our cognitive systems can trade directly with the environment: sensorimotor states encoding the structure or shape of the environment itself and our actions in it. This would apply equally to disembodied or dis-embodied representations. Of course, this is reminiscent of Prinz’s insistence that “if an amodal code exists, it works on credit” (2002, p.137.) If this line of thought can be made to work, it will have the consequence that understanding is still simulating, *even if there is a LOT*[18].

There are other positive ideas about the contribution of simulation to understanding which might be developed, too. Haugeland (1979) and Winograd (1980) developed prescient critiques of symbolic accounts of natural language comprehension, and these critiques are pregnant with suggestions about what sort of features a plausible model of natural language comprehension must have. They emphasize such aspects as the holistic nature of

[18] This echoes the point I made in discussing Weiskopf’s ‘minimal view’ of comprehension in chapter 5: even if there is a wholly disembodied conceptual system computing truth-conditional inferences from sentences, there is still surely a big difference between merely computing the truth conditions of Chomsky’s ‘green ideas’ sentence and computing the truth conditions of ‘the cat sat on the mat’ while also simulating the kind of perceptual state we would enjoy if things were thus.

understanding, the role of common sense, the social and pragmatic aspects of language, and the phenomenology of comprehension:

> Much of what we recognize as "making sense" is ... about some (possibly unique) circumstance or episode, which a longer fragment leads us to "visualize." Introspectively, it seems that we imagine ourselves into the case, and then decide from within it what's plausible. Of course, *how* this is done is just the problem.
>
> (Haugeland 1979, p.624.)

I cannot develop these ideas here. But I think they show that Mahon and Caramazza's arguments should be viewed more as a challenge than an objection. The empirical evidence may not clearly support ELC in the latter's current state of theoretical development. However, we can be justifiably optimistic that deeper examinations of the parsimony argument, the resemblance theory of mental representation, and other theoretical issues will serve to clarify and strengthen the case for simulation as a central constituent of comprehension.

Conclusion

It may seem that we have a tie between cognitivist ELC and the amodal view, theoretically and empirically. However, I have proposed several promising ways in which the tie might be broken in favour of ELC, and suggested that this leaves us still in a position to tentatively endorse the theory (while acknowledging the work that remains to be done.) In that case, it is reasonable as things stand to believe that language comprehension is not mere activation of disembodied LOT symbols—or, for that matter, skilful action (capacities) and sensorimotor couplings—but the activation of sensorimotor representations to amount to a form of internal simulation.

Before reflecting a little more broadly on this conclusion, we should return to the logical geography I presented at the end of chapter 3. Cognitivist ELC claims that understanding is simulating; but simulating in what sense? And 'is' in what sense?

Firstly, the relevant sense of simulation must be Barsalou's minimal, technical understanding—at least, if 'is' denotes a universal relationship of necessity or constitution This follows directly from my heavy reliance on the Partially Dis-embodied Account in defending ELC's ability to handle abstraction and other objections. Probably, when we comprehend (to whatever extent we do) extremely abstract language, we are not constructing a full-blown situation model—not even one *sans* affordances. But I think we are warranted in concluding that whenever we comprehend any language, to however great or small an extent, a necessary constituent of that comprehension process is the activation of more-or-less schematic sensorimotor representations.

It is also reasonable to think that when we comprehend language more fully or richly, especially language describing concrete objects and situations, the construction of situation models often occurs—sometimes incorporating information about affordances. Given that the richness of the comprehension is partly a function of the level of detail being simulated, it seems as though the construction of such situation models, in these kinds of cases, is a genuine constituent of understanding. So, all instances of understanding are instances of simulating, at least in the minimal sense. And many instances of understanding are partially constituted by—not merely causally related to—instances of simulating in one or another of the richer senses.

Now, let us briefly return to my central conclusion: that cognitivist ELC is a better explanation of the data than either the amodal view or enactivist ELC. The broader significance of this is worth considering. This conclusion is in the spirit of Prinz and Barsalou's (2000) chapter, 'Steering a Course for Embodied Representations'. The course to be steered runs between two untenable extremes. One of these—the amodal view—amounts to explaining away the intimate relationships between conceptual and linguistic thought, on the one hand, and sensorimotor processing, on the other, as a mere causal, contingent epiphenomenon, and claiming that the stuff of thought is a representational code completely distinct and divorced from all of our perceptual and bodily encounters with the world and with natural languages. The other untenable extreme—the radical thesis of enactivism—amounts to denying that our skilful, flexible, and open-ended engagement with the world is best explained by our ability to represent the world as being certain ways rather than others, and by our ability to simulate objects, properties, and events in their absence.

ELC, along with the simulation-based 'cognitivist' form of EC more generally—as embodied in Barsalou's theory of PSS and Prinz's Concept Empiricism—allows us the best of both worlds. It does justice to the rich and deep connections between thinking and doing, without rejecting the immeasurable explanatory purchase to be gained by positing the existence of mental representations. Understanding is not doing; it is thinking. But thinking is simulating, and is inextricably tied to our nature as situated, acting beings.

Bibliography

Adams, F. (2010). Embodied cognition. Phenomenology and the Cognitive Sciences 9(4), 619-628.

Adams, F. & K. Aizawa (2001). The bounds of cognition. Philosophical Psychology 14(1), 43-64.

Anderson, M.L. (2003). Embodied cognition: a field guide. Artificial Intelligence 149, 91-130.

Barsalou, L.W. (1999). Perceptual symbol systems. Behavioral and Brain Sciences 22(4), 577-660.

Barsalou, L.W. (2003). Abstraction in perceptual symbol systems. Philosophical Transactions: Biological Sciences 358(1435), 1177-1187.

Barsalou, L.W., W.K. Simmons, A.K. Barbey & C.D. Wilson (2003). Grounding conceptual knowledge in modality-specific systems. Trends in Cognitive Sciences 7(2), 84-91.

Barsalou, L.W. & K. Wiemer-Hastings (2005). Situating abstract concepts. In D. Pecher & R.A. Zwaan (eds.) Grounding Cognition: The Role of Perception and Action in Memory, Language and Thinking. New York: Cambridge University Press.

Bartels, A. (2006). Defending the structural concept of representation. Theoria 55, 7-19.

Bergen, B.K. & N. Chang (2005). Embodied construction grammar in simulation-based language understanding. In J. Östman & M. Fried (eds.) Construction Grammars: Cognitive grounding and theoretical extensions. Amsterdam: John Benjamins.

Borghi, A.M., A.M. Glenberg & M.P. Kaschak (2004). Putting words in perspective. Memory & Cognition 32(6), 863-873.

Boroditsky, L. & M. Ramscar (2002). The role of body and mind in abstract thought. Psychological Science 13(2), 185-189.

Brooks, R. (1991a). Intelligence without representation. Artificial Intelligence 47(1-3), 139-159.

Brooks, R. (1991b). New approaches to robotics. Science 253(5025), 1227-1232.

Burgess, C. (2000). Theory and operational definitions in computational memory models: a response to Glenberg and Robertson. Journal of Memory and Language 43(3), 402-408.

Burgess, C. & K. Lund (1997). Modelling parsing constraints with high-dimensional context space. Language and Cognitive Processes 12, 177-210.

Carruthers, P. (2011). The Opacity of Mind. Oxford: Oxford University Press.

Casasanto, D. (2009). Embodiment of abstract concepts: good and bad in right- and left-handers. Journal of Experimental Psychology: General 138(3), 351, 367.

Chalmers, D.J. (1989). Why Fodor and Pylyshyn were wrong: the simplest refutation. Proceedings of the Twelfth Annual Conference of the Cognitive Science Society, 340-347.

Chomsky, N. (1957). Syntactic Structures. The Hague: Mouton.

Chomsky, N. (1959). Verbal behaviour by B.F. Skinner. Language 35(1), 26-58.

Church, A. (1936). A note on the Entscheidungsproblem. Journal of Symbolic Logic 1(1), 40-41.

Clark, A. (1993). Associative Engines: Connectionism, Concepts, and Representational Change. Cambridge, Mass.: MIT Press.

Clark, A. (1997). Being There: Putting Brain, Body and World Together Again. Cambridge, Mass.: MIT Press.

Clark, A. (2001). Mindware: An Introduction to the Philosophy of Cognitive Science. New York: Oxford University Press.

Clark, A. & D. Chalmers (1998). The extended mind. Analysis 58(1), 7-19.

Copeland, J. (1993). Artificial Intelligence: A Philosophical Introduction. Oxford: Blackwell.

Cummins, R. (1996). Systematicity. The Journal of Philosophy 93(12), 591-614.

Cummins, R. (1996). Representations, Targets, and Attitudes. Cambridge, Mass.: MIT Press.

Davidson, D. (1967). Truth and meaning. Synthese 17, 304-323.

De Jaegher, H. (2009). Social understanding through direct perception? Yes, by interacting. Consciousness and Cognition 18(2), 535-542.

Dretske, F. (1981). Knowledge and the Flow of Information. Cambridge, Mass.: MIT Press.

Dreyfus, H.L. (1972). What Computers Can't Do. New York: Harper & Row.

Dove, G. (2011). On the need for embodied and dis-embodied cognition. Frontiers in Psychology 1(242), 1-13.

Flusberg, S.J., P.H. Thibodeau, D.A. Sternberg & J.J. Glick (2010). A connectionist approach to embodied conceptual metaphor. Frontiers in Psychology 1(197), 1-11.

Fodor, J. (1976). The Language of Thought. Sussex: Harvester.

Fodor, J. (2004). Having concepts: a brief refutation of the twentieth century. Mind and Language 19(1), 29-47.

Fodor, J.A. (2008). LOT 2: The Language of Thought Revisited. Oxford: Clarendon.

Fodor, J.A. & Z.W. Pylyshyn (1988). Connectionism and cognitive architecture: a critical analysis. Cognition 28(1), 3-71.

Gallagher, S. (2005). How the Body Shapes the Mind. Oxford: Clarendon.

Gallagher, S. (2008). Direct perception in the intersubjective context. Consciousness and Cognition 17(2), 535-543.

Gallese, V. & A. Goldman (1998). Mirror neurons and simulation theory of mind-reading. Trends in Cognitive Sciences 2(12), 493-501.

Gallese, V. & G. Lakoff (2005). The brain's concepts: the role of the sensory-motor system in conceptual knowledge. Cognitive Neuropsychology 25, 1-25.

Gardner, H. (1985). The Mind's New Science: A History of the Cognitive Revolution. New York: Basic Books.

Gibbs, R.W. Jr. (1996). Why many concepts are metaphorical. Cognition 61, 309-319.

Gibson, J.J. (1979). The Ecological Approach to Visual Perception. Boston: Houghton Mifflin.

Glenberg, A.M. (1997). What memory is for. The Behavioral and Brain Sciences 20(1), 1-55.

Glenberg, A.M. (2007). Language and action: creating sensible combinations of ideas. In G. Gaskell (ed.) The Oxford Handbook of Psycholinguistics, 361-370. New York: Oxford University Press.

Glenberg, A.M., D. Havas, R. Becker & M. Rinck (2005). Grounding language in bodily states: the case for emotion. In D. Pecher & R.A. Zwaan (eds.) Grounding Cognition: The Role of Perception and Action in Memory, Language and Thinking 115-128. New York: Cambridge University Press.

Glenberg, A.M. & M.P. Kaschak (2002). Grounding language in action. Psychonomic Bulletin & Review 9(3), 558-565.

Glenberg, A.M. & D.A. Robertson (1999). Indexical understanding of instructions. Discourse Processes 28(1), 1-26.

Glenberg, A.M. & D.A. Robertson (2000). Symbol grounding and meaning: a comparison of high-dimensional and embodied theories of meaning. Journal of Memory and Language 43(3), 379-401.

Glenberg, A.M., M. Sato, L. Cattaneo, L. Riggio, D. Palumbo & G. Buccino (2008). Processing abstract language modulates motor system activity. Quarterly Journal of Experimental Psychology 61(6), 905-919.

Haugeland, J. (1979). Understanding natural language. The Journal of Philosophy 76(11), 619-632.

Haugeland, J. (1985). Artificial Intelligence: The Very Idea. Cambridge, Mass.: MIT Press.

Haugeland, J. (1998). Mind embodied and embedded. In J. Haugeland, Having Thought: Essays in the Metaphysics of Mind. Cambridge, Mass.: Harvard University Press.

Harnad, S. (1990). The symbol grounding problem. Physica D 42(1-3), 335-346.

Horst, S. (2011). The computational theory of mind. In E. Zalta (ed.) The Stanford Encyclopedia of Philosophy (Spring 2011 Edition), accessed 18 Jan 2013. <http://plato.stanford.edu/archives/spr2011/entries/computational-mind/>.

Isaac, A.M.C. (2012). Objective similarity and mental representation. Australasian Journal of Philosophy, forthcoming.

Joyce, D.W., L.V. Richards, A. Cangelosi & K.R. Coverntry (2003). On the foundations of Perceptual Symbol Systems: specifying embodied representations via connectionism. In F. Detje, D. Dörner, & H. Schaub (eds.),The Logic of Cognitive Systems: Proceedings of the Fifth Annual Conference on Cognitive Modeling, 147-152. Bamberg: Universitätsverlag Bamberg.

Kaschak, M.P. & A.M. Glenberg (2000). Constructing meaning: the role of affordances and grammatical constructions in sentence comprehension. Journal of Memory and Language 43(3), 508-529.

Kiverstein, J. & A. Clark (2009). Introduction: mind embodied, embedded, enacted: one church or many? Topoi 28, 1-7.

Knott, A. (2012). Sensorimotor Cognition and Natural Language Syntax. Cambridge, Mass.: MIT Press.

Lakoff, G. & M. Johnson (1980a). The metaphorical structure of the human conceptual system. Cognitive Science 4(2), 195-208.

Lakoff, G. & M. Johnson (1980b.) Conceptual metaphor in everyday language. The Journal of Philosophy 77(8), 453-486.

Lakoff, G. & M. Johnson (1999). Philosophy in the Flesh. New York: Basic Books.

Landauer, T.K. & S.T. Dumais (1997). A solution to Plato's problem: the Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. Psychological Review 104(2), 211-240.

Letheby, C. (2012). In defence of embodied cognition: a reply to Fred Adams. Phenomenology and the Cognitive Sciences 11(3), 403-414.

Louwerse, M. (2011). Symbol interdependency in symbolic and embodied cognition. Topics in Cognitive Science 3(2), 273-302.

Mahon, B.Z. & A. Caramazza (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. Journal of Physiology-Paris 102(1), 59-70.

Marghetis, T., T. Davenport, R. Metusalem & B. Bergen (under review). From format to function: embodiment and the functional roles of neural circuits.

McCarthy, J. & P.J. Hayes (1969). Some philosophical problems from the standpoint of artificial intelligence. Machine Intelligence 4, 463-502.

Menary, R. (2010). Introduction to the special issue on 4E cognition. Phenomenology and the Cognitive Sciences 9(4), 459-463.

Newell, A. & H.A. Simon (1976). Computer science as empirical inquiry: symbols and search. Communications of the ACM 19(3), 113-126.

Norman, D.A. (1980). Twelve issues for cognitive science. Cognitive Science 4(1), 1-32.

Noë, A. (2009). Conscious reference. The Philosophical Quarterly 59(236), 470-482.

O'Brien, G. & J. Opie (2002). Radical connectionism: thinking with (not in) language. Language and Communication 22(3), 313-329.

O'Brien, G. & J. Opie (2004). Notes towards a structuralist theory of mental representation. In H. Clapin, P. Staines & P. Slezak (eds.) Representation in Mind: New Approaches to Mental Representation, 1-20. Amsterdam: Elsevier.

O'Regan, J.K. & A. Noë (2001). A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24, 939-1031.

Pecher, D., R. Zeelenberg & L.W. Barsalou (2003). Verifying different-modality properties for concepts produces switching costs. Psychological Science 14(2), 119-124.

Pezzulo, G., L.W. Barsalou, A. Cangelosi, M.H. Fischer, K. McRae & M.J. Spivey (2011). The mechanics of embodiment: a dialog on embodiment and computational modelling. Frontiers in Psychology 2(5), 1-21.

Prinz, J.J. (2002). Furnishing the Mind: Concepts and their Perceptual Basis. Cambridge, MA.: MIT Press.

Prinz, J.J. (2004). Gut Reactions: A Perceptual Theory of Emotion. New York: Oxford University Press.

Prinz, J.J. (2005). Passionate thoughts: the emotional embodiment of moral concepts. In D. Pecher & R.A. Zwaan (eds.) Grounding Cognition: The Role of Perception and Action in Memory, Language and Thinking, 93-114. New York: Cambridge University Press.

Prinz, J.J. & L.W. Barsalou (2000). Steering a course for embodied representation. In E. Dietrich & A. Markman (eds.) Cognitive Dynamics: Conceptual Change in Humans and Machines, 51-77. Cambridge, Mass.: MIT Press.

Prinz, J. & A. Clark (2004). Putting concepts to work: some thoughts for the twentyfirst century. Mind and Language 19(1), 57-69.

Pulvermüller, F. (1999). Words in the brain's language. The Behavioral and Brain Sciences 22, 253-336.

Samuel, A.L. (1959). Some studies in machine learning using the game of checkers. IBM Journal of Research and Development 3, 210-229.

Schwitzgebel, E. (2008). The unreliability of naive introspection. Philosophical Review 117(2), 245-273.

Searle, J. (1980). Minds, brains, and programs. The Behavioral and Brain Sciences 3(3), 417-457.

Sejnowski, T.J. & C.R. Rosenberg (1987). Parallel networks that learn to pronounce English text. Complex Systems 1(1), 145-168.

Shapiro, L. (2007). Symbolism, embodied cognition, and the broader debate. In M. de Vega, A.M. Glenberg & A.C. Graesser (eds.) Symbols and Embodiment: Debates on Meaning and Cognition, 57-74. New York: Oxford University Press.

Shapiro, L. (2007). The embodied cognition research programme. Philosophy Compass 2(2), 338-346.

Shapiro, L. (2011). Embodied Cognition. Oxford: Routledge.

Shortliffe, E.H. & B.G. Buchanan (1975). A model of inexact reasoning in medicine. Mathematical Biosciences 23(3-4), 351-379.

Simon, H.A. & A. Newell (1958). Heuristic problem solving: the next advance in operations research. Operations Research 6(1), 1-10.

Skinner, B.F. (1957). Verbal Behavior. New York: Appleton-Century-Crofts.

Stanfield, R.A. & R.A. Zwaan (2001). The effect of implied orientation derived from verbal context on picture recognition. Psychological Science 12(2), 153-156.

Tienson, J.L. (1988). An introduction to connectionism. Southern Journal of Philosophy 26(supplement), 1-16.

Turing, A.M. (1936). On computable numbers, with an application to the Entscheidungsproblem. Proceedings of the London Mathematical Society 2(42), 230-265.

Turing, A.M. (1950). Computing machinery and intelligence. Mind 59(236), 433-460.

Van Elk, M., M. Slors & H. Bekkering (2010). Embodied language comprehension requires an enactivist paradigm of cognition. Frontiers in Psychology 1(234), 1-9.

Van Gelder, T. (1995). What might cognition be, if not computation? The Journal of Philosophy 92(7), 345-381.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. The Behavioral and Brain Sciences 21, 615-665.

Varela, F., E. Thompson & E. Rosch (1991). The Embodied Mind: Cognitive Science and Human Experience. Cambridge, Mass.: MIT Press.

Weiskopf, D.A. (2010). Embodied cognition and linguistic comprehension. Studies in History and Philosophy of Science 41(3), 294-304.

Wilson, M. (2002). Six views of embodied cognition. Psychonomic Bulletin & Review 9(4), 625-636.

Winograd, T. (1972). Understanding Natural Language. Edinburgh: Edinburgh University Press.

Winograd, T. (1980). What does it mean to understand language? Cognitive Science 4(3), 209-241.

Zwaan, R.A. (2004). The immersed experience: toward an embodied theory of language comprehension. The Psychology of Learning and Motivation 44, 35-62.

Zwaan, R.A., R.A. Stanfield & R.H. Yaxley (2002). Language comprehenders mentally represent the shapes of objects. Psychological Science 13(2), 168-171.