Available online at www.sciencedirect.com
SCIENCE DIRECT®
Language Sciences 26 (2004) 503–524
Language Sciences
www.elsevier.com/locate/langsci

Representation, levels, and context in integrational linguistics and distributed cognition

John Sutton *

Department of Philosophy, Division of SCMP, Macquarie University, NSW 2109, Australia

Abstract

Distributed Cognition and Integrational Linguistics have much in common. Both approaches see communicative activity and intelligent behaviour in general as strongly context-dependent and action-oriented, and brains as permeated by history. But there is some tension between the two frameworks on three important issues. The majority of theorists of distributed cognition want to maintain *some* notions of mental representation and computation, and to seek generalizations and patterns in the various ways in which creatures like us couple with technologies, media, and other agents; many also want to offer explanations at subpersonal levels which may undercut the autonomy of personal-level accounts. In contrast, dominant views in integrational linguistics reject all invocation of representation, resist the explanatory search for similarity across contexts and moments, and see linguistics as a lay discipline which should not offer explanations in terms alien to ordinary agents. On each of these issues, I argue that integrationists could move closer to the distributed cognition framework without losing the most important aspects of their view: integrationist criticisms of mainstream or classical theories can be respected while alliances with revised cognitivist views about representation, context, and explanation are developed.

Keywords: Integrational linguistics; Distributed cognition; Mental representation; Reduction; Context; Memory

* Tel.: +61 2 9850 8817; fax: +61 2 9850 8892.
E-mail address: jsutton@scmp.mq.edu.au
URL: http://www.phil.mq.edu.au/staff/jsutton.

0388-0001/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.langsci.2004.09.002

1. Introduction and preview

'Distributed Cognition' is an increasingly common label for a diverse set of views in the philosophy of cognitive science, views which are hard to render precise and which remain (to many working in the field) marginal and implausible. 'Integrational Linguistics' labels a movement in linguistics, originated by Roy Harris, also hard to render precise: as Joseph notes, may see it as 'merely negative, and relentlessly so, criticising virtually everything ever written about language without offering a coherent alternative vision or programme of work' (Joseph, 2003, p. 99). It has thus remained, despite more positive recent expositions (Harris, 1996, 1997, 1998), fairly marginal and implausible to most working in the field. Contributors to this special issue seek bravely to outline possibilities for fruitful interaction between these two marginal frameworks. In this paper I pin down specific points of contact and of conflict between them, aiming to work towards a more precise set of shared assumptions which might help challenge the marginality of the two frameworks in their respective intellectual and institutional domains. In keeping with the urge to increase the interdisciplinary diversity, I add memory to the mix, using the developmental psychology of autobiographical memory to exemplify the point that integrational constraints need not be violated by empirical psychology.

My presentation of Integrational Linguistics is hesitant. I rely particularly on Roy Harris' work of the last six or seven years, including his remarkable *Introduction to Integrational Linguistics* (1998), which (to the surprise of some of Harris' critics) lays out the programme 'in a systematic and orderly way' (Joseph, 2003, p. 100). Though troubled by Harris' tendency 'to present integrationism as monolithic' (Joseph, 2003, p. 108), I also draw heavily on Stephen Cowley's provocative and impressive recent attempts to bring Integrational Linguistics into contact with (the best bits of) the contemporary cognitive sciences (Cowley, 2002, 2004; Cowley and Spurrett, 2003). Integrationist literature is difficult to get into: while at its best Harris' writing is quirky and playfully elegant, it is not easy to track his use of idiosyncratic terms and his (paradoxical) desire to police their use. Few attempts to amend or extend integrationism have met with his approval. Even would-be sympathizers can be convicted of 'segregationism', of treating 'language' as an isolable system separate from social interaction, and as distinguishable from non-linguistic phenomena. Realizing then that my best efforts at avoiding segregationism may turn out to be in vain, that not all allies are welcome, I approach the task of relating – 'integrating' may be too hopeful—distributed cognition and integrational linguistics in a spirit of interested engagement. [1]

[1] The appropriate tactics here are similar to those needed by philosophers of cognitive science in dealing with diverse but related criticisms from Gibsonian, Wittgensteinian, and phenomenological points of view. While it is important to stress that 'cognitive science' is neither as monolithic nor as rigid (theoretically or institutionally) as these critics often suggest, a more useful response is carefully to pick out and work with the positive aspects of these alternative perspectives. I tried to do this for the case of Gibsonian and other criticisms of theories of memory in my *Philosophy and Memory Traces* (1998), chapters 15 and 16 (Sutton, 1998). From the critics' points of view, of course, this can look like attempted assimilation by the dominant paradigm: but the sociology of 'normal' cognitive science is much too complicated for that to be a simple process, and in fact most of the cognitive science discussed here is itself, as I have noted, far from mainstream.

I first offer bowdlerized sketches of some core commitments of both programmes in a way which highlights possible convergence (though further exposition will emerge as I put them together on specific issues). It is no part of my brief here to pin down precisely or to defend the central claims of the distributed cognition movement. Then I single out three issues on which Integrational Linguistics seems to be in tension with distributed cognition. In each case, it may just be that I have misunderstood the integrationist claims: but in each case I think that their motivating ideas can, to some extent, be salvaged without the extra negative claims which cause the conflict. At present, Integrational Linguistics is in some ways more radical than distributed cognition: my point is that these extra steps, which take it that much further away from the cognitive sciences, cause the trouble, and can (I suggest hopefully) be jettisoned without loss of the strong driving commitments.

This leaves, with luck, a modified position acceptable to both camps. To underline and exemplify this position, I finish with some brief remarks on one strand of current work in the interdisciplinary sciences of memory. It should be both theoretically possible and empirically fruitful to hold on to some of the integrationists' ideas without violating principles which enthusiasts for distributed cognition rightly want to retain.

2. Distributed cognition

With some inevitable loss of relevant diversity, it is possible to use 'distributed cognition' to bundle together some coalescing recent lines of research in philosophy of mind, developmental psychology, robotics, cognitive archaeology and anthropology, dynamical systems approaches to cognition, and human–computer interaction (among other fields). Key rallying-points for the messy movement include Ed Hutchins' fieldwork on the distribution of navigation processes across individuals, charts, and instruments in *Cognition in the Wild* (Hutchins, 1995), Andy Clark's synthesis and philosophical development of interdisciplinary work in *Being There: putting brain, body, and world together again* (Clark, 1997), and a paper Clark co-authored with David Chalmers defending the idea of 'the extended mind' (Clark and Chalmers, 1998). For convenience in this paper, the kind of 'Distributed Cognition' I will be comparing with Integrational Linguistics is roughly that defended by Clark, as extended further in his more recent writing (see for instance Clark, 2001a, 2003, 2004, and compare Rowlands, 2003, Chapter 9). [2]

On Clark's view, we are intricately psychologically tangled with, and our minds projected out into, a range of cognitive objects such as instruments, media, and other people. When (for example) an academic is writing a paper, or an artist is working

[2] Of course, there are many different routes to these ideas: among the relevant areas neglected in this presentation are related lines of thought about the evolution of cognition (Donald, 1991, 2000; Renfrew and Scarre, 1998), in the cognitive side of social anthropology (Strauss and Quinn, 1997), in sociology and philosophy of science (Latour, 1999), and in social ontology (Gilbert, 1989).

on an abstract artwork, the intelligent activity driving the process can include or span (as well as brain and body) a notebook, a sketchpad, scraps of paper with old notes, a friend at the other end of a phone call, various files on computer or on paper with records from different earlier stages, and so on. Humans, in general, are 'bio-technological hybrids … primed so as to participate in cognitive and computational architectures whose bounds far exceed those of skin and skull' (Clark, 2001a, pp. 138, 142). In certain circumstances, artefacts and other external structures are literally cognitive: notebooks, incised sticks, slide-rules, computer control-sticks, software agents, fingers, knots, rituals, monuments, roads, roadsigns, and landscapes can become components of an extended system (Haugeland, 1998). What makes this a view about distributed *cognition* is the point that such mind-tools are not simply cognitive commodities, for the use and profit of the active mind: rather, in certain circumstances, along with the brain and body interacting with them, they *are* the mind. For Clark, 'it is our basic *human* nature to annex, exploit and incorporate nonbiological stuff deep into our mental profiles' (Clark, 2003, p. 198).

To stress the 'leakiness' of the human mind is to focus on its tendencies both to co-opt and to incorporate external resources. Cognitive processes sometimes constitutively involve multiple loops between brain, body, and world, where 'world' includes both the physical and the social environments with which embodied brains couple, the 'scaffolding' on which they lean. Cognitive states, like the processes in which they participate, are thus sometimes hybrid biological and non-biological states. So, in certain circumstances, things have a cognitive life (Sutton, 2002b). Claims like this can be put more or less strongly, and have methodological as well as metaphysical readings. [3] And of course there is much more to say about the relevant circumstances and conditions: what's critical is the context in which external resources are assimilated, parasitised, or internalised in some contingent 'dynamical singularity' spanning brain, body and world (Hurley, 1998a). Successful use of our various 'designer environments' counts as genuine distributed cognition when the external resources are used in recognizably flexible cognitive activity.

3. Integrational linguistics

Roy Harris writes (2000, p. 163)

> To make sense of any episode of human communication we have to recognize an integration of activities being carried out by particular individuals in a particular set of circumstances. Signs are created in the course of this integrational process.

[3] In particular, for present purposes they should not be read as imposing necessary conditions on cognition.

In Integrational Linguistics, then, analysis must focus on the context in which action-oriented symbolic activity arises. The integrationist urges us to focus on ordinary embodied skills and habits, on practical strategies of communication rather than on any set of context-free inner models of reality. Neither 'language' nor 'cognition', for Harris, is a system separate from communicative activity. The symbolic value of any sign emerges *only* in my reliance on it 'to integrate certain programmes of activity in my daily comings-and-goings' (Harris, 2000, p. 68) There is no abstract, permanent set of meanings and messages in either language or thought, prior to episodes of thinking and communicating.

In turn, Stephen Cowley offers a necessary integrationist account of the role of the brain, which should be attractive to the post-connectionist inspiration of much distributed cognition theorizing. Brains, for Cowley, are 'biosocial organs permeated by history' (Cowley, 2002, pp. 73, 75). Symbol-manipulation is unlikely to be either internal or innate in any interesting sense, but is rather an external or relational capacity learned in developing the capacities to do what feels right. What we call 'language' is then 'insinuated into developing neural organization as an individual exploits symbol-mediated activity to develop social skills and capacities' (Cowley, 2002, p. 85).

The implication for method in linguistics is that we should not seek to categorize an abstract and general 'language system', but instead investigate the complex and diverse practices that drive changes in 'contextualizing'. Contextualizing is the use of previous experience 'to integrate activities so that, in future, their effects are likely to benefit' the agent (Cowley, 2004, especially Section 2.3). Cowley puts this perspective into practice himself in strong empirical studies of (for example) conversational turn-taking and embodied interactivity in dialogue (see Cowley, 2002, p. 87–89) and by bringing an integrationist perspective to bear on developmental linguistics and psychology (Cowley, 2004). Microstudies of the properties of talk and silence in prosody, gesture, and facial expression can combine, from an integrational point of view, into a rich picture of the development of utterance activity, the contextual exercise of utterance capacity, and the propensity to exploit external symbols.

Already, with this bare thumbnail sketch, we can see a number of productive points of contact between Integrational Linguistics and distributed cognition. I point for now to four such convergences, on which both frameworks diverge from more mainstream perspectives in linguistics and cognitive science. The two frameworks share a stress on the *key role of context* in cognitive and linguistic activity, where 'context' includes social, environmental, bodily, and neural factors. On both views, the internal states and processes implicated in cognitive and linguistic activity are primarily *action-oriented*, set up for the integration of a range of embodied and interpersonal goals or processes. *External symbols*, in both frameworks, are not mere passive tools or supports for the active individual mind, but themselves play key roles in intelligent activity by transforming the tasks or the requirements of the situation. And theorists in both frameworks are starting to show that they have productive practical methodological and *empirical consequences*.

So far, so rosy. But, inevitably, there's a catch—or three.

4. Representations

The integrationist rejects a doctrine which Harris calls 'telementation'. Though it is not easy to render this doctrine precise, since Harris redescribes it as he ascribes it to most major figures in the history of linguistics and philosophy, the central telementationist idea is that language is for the transfer of thoughts from one mind to another (Harris 1996, pp. 114–115; 1998, pp. 22, 32 and *passim*). It is striking that Andy Clark's discussions of language also target this idea: public language does not just express pre-existing thoughts, but expands the realm of the thinkable, transforming rather than just reporting on cognitive processes (Clark, 1996, pp. 93–94). Clark calls his own position a 'supra-communicative' view of language (Clark, 1997, Chapter 10), and I have argued that it is a clear rejection of the 'expressivist' (or telementational) view of language (Sutton, 2002a, p. 380–384). [4] On 'classical' expressivist views in cognitive science and linguistics, the thoughts expressed in language are prior to, behind, and independent of that expression: the integrationist and the distributed cognition theorist disagree, seeing the role of language as more than that of a conduit or channel. A further point of contact with at least some distributed cognition theorists is in the integrationist's suspicion of any (innate or other) domain-specific mechanism or module of linguistic cognition.

Having shared enemies does not automatically make you allies: and in particular nothing I have said so far suggests agreement between the two frameworks in their positive views of the functions and nature of linguistic and other communicative activity. But there is a more immediate concern which can serve to introduce a first point of tension. For most distributed cognition theorists, questions about the nature of communicative activity and linguistic representation are one thing, while questions about the nature of *mental* representation are quite another matter. Distributed cognition, indeed, can be sensibly taken to be a view *about* (the nature and location of) mental representation (or at least the vehicles of mental representation—see below). But in Integrational Linguistics, the attack on telementational views of communication often seems to be taken as in itself an attack on mental representation too. Or at least, whatever the core reason, the integrationist is just as hostile to *any* invocation of mental representation as to the idea that the job of *language* is to represent thoughts.

This is another aspect of what the integrationist calls 'the language myth', which arises through some combination of poor (Western) theorizing and ordinary integra-

[4] The label 'expressivism' is Gauker's (1999), and captures what Gauker had previously (1994) called 'the Lockean theory of communication'. Locke is one of Roy Harris' prime telementationists (Harris, 1996, p. 135; 1998, pp. 32–33). Peter Carruthers clearly lays out a range of options about this doctrine in current theory (Carruthers, 2002). Part of my argument in Sutton (2002a) was that expressivism (or telementation), which is a view about linguistic representation and communication, is logically independent of debates about whether or not thinking itself is linguistic in form (a doctrine I called 'lingualism', following John Preston (1997, p. 1)). My worry in this section is that the integrational linguist (deliberately, but in my view wrongly) is conflating the two issues.

tive activity. As well as 'telementation', the integrationist criticizes the 'fixed code' as mythical. The 'fixed code myth' is the idea that either language or thought is a system of determinate correspondences between form and meaning (Harris, 1996, p. 131; 1998, p. 22). (Harris accepts that telementation and the fixed code are logically independent aspects of the language myth, but suggests that in practice they are almost inevitably taken to be mutually supporting). Although the fixed code is sometimes presented simply as a view about *linguistic* determinacy (Harris, 1998, p. 32), it is also regularly framed as a thesis about *psychological* determinacy, and the rejection of fixed codes is taken to be a rejection of any psychological realism about representations, of what Cowley calls 'token-realism' (Cowley, 2002, pp. 76, 85–86). [5]

Here we have at least some confusion or, at worst, a serious problem. Although, as I will show in a moment, theorists of distributed cognition do have many different views on these points, what I am taking to be the mainstream position *requires* some notion of (weakened and contextualized) mental representation. Questions about the relations between form and meaning, or in contemporary jargon, between the vehicles of content and the contents they carry, should not be *assumed* to have either easy or determinate answers (Hurley, 1998a,b), for 'distributed cognition' *means* that the relevant vehicles can sometimes, partly, be outside the brain and body. [6] But that such questions are worth asking, that answers are worth looking for, and that there *are* both vehicles and contents for us to ask these questions about, are views widely, and in my view rightly, held among philosophers of distributed cognition. To show this, I will need to back up and say quite a bit more about the motivations for retaining mental representation in distributed cognition: but some quick conceptual geography will help first, to locate what I'm calling this 'mainstream' view among a range of alternatives. I want to distinguish this particular view, in particular, from the views of (actual or possible) enthusiasts for distributed cognition who couple this enthusiasm with *either* a classical *or* an eliminativist account of mental representation.

Firstly, on the one hand (on the right?), there are philosophers sympathetic to key parts of the distributed cognition framework who retain some version of a 'classical' view of mental representation. It's not *necessary*, that is, to reject the Language of

[5] So we can read Nigel Love's (2004) critique of Clark as claiming that, despite Clark's best efforts at genuinely contextualising both mind and language, Clark's position in the end attributes an unrealistic form of psychological determinacy to mental representations (Love, 2004). Love rightly notes that 'Clark's target is not the idea that brains represent *aspects* of a real, independent world, but rather the idea of those representations as inert and *action-neutral*': but the brunt of Love's argument is that there turns out to be no middle way here, and that even Clark's minimal invocation of mental representations collapses back into the fixed code: as Love puts it, Clark underpins his ideas about public language 'with the classical, codist view of private language'. See footnote 17.

[6] It's for this reason that Rowlands (2003, Chapter 9) *calls* the view 'vehicle externalism': this is what distinguishes distributed cognition from more common forms of semantic or social externalism, which is a view only about content and says nothing about the individuation or location of the vehicles of content.

Thought hypothesis (LOT) [7] to be sympathetic to distributed cognition: debates about whether the medium of thought is linguistic or not are independent of debates about whether and to what extent cognition extends into the environment. It is possible to argue, indeed, that the nature of computation within a classical account of mental representation *requires* distributed coupling with external symbol-systems (Wilson, 1994; Sterelny, 2000). [8] But while the fact that most adherents to the distributed framework reject LOT is contingent, there are, I will argue, sound reasons for them to do so and instead to retain revised, non-LOT-style conceptions of mental representation: their representations will not be atomic, independent, or context-free. Integrational linguists do not tend to distinguish between these different possibilities, or to acknowledge that not all accounts of mental representation are LOT-style accounts: so one question to put to them is whether their qualms about mental representations are *only* qualms about 'fixed codes' of the LOT style. Can the cognitive sciences invoke any mental representations which are *not* 'fixed codes'?

On the other hand (on the left?), there are philosophers sympathetic to distributed cognition who *do* reject mental representations wholesale. They are unwilling to postulate determinate contentful inner cognitive states which figure in computational cognitive processes. This 'eliminativist' strand in distributed cognition has representatives in robotics, developmental psychology, and dynamical systems cognitive theory (van Gelder, 1995). [9] The idea is that external representations alone do all the explanatory work, and that there's no need to talk of 'mental' representation at all: since in many cases 'the world is its own best representation' (Brooks, 1991), there's no need ever to postulate additional internal or spread representations (for an overview and discussion see Eliasmith, 1996). There's a superficial similarity here, then, with Integrational Linguistics, and a similar interpretive difficulty: are the representations being rejected simply LOT-style atomic, context-free representations? [10]

[7] In both of its main forms: Jerry Fodor's (1976) hypothesized Mentalese, an innate set of atomic constituents of thought which (re)combine according to combinatorial principles, and natural language as itself the medium of thinking (Carruthers, 2002).

[8] One reason that, in practice, few philosophers combine the Language of Thought (LOT) with Distributed Cognition may be that most adherents to LOT make a sharp distinction between *intrinsic* and *derived* content, by which all external symbols and resources have at most derived content. But again, such a distinction is not required by LOT: nor is it obvious that such a distinction in fact does rule out Distributed Cognition (Clark, 2004). And in one respect LOT-style theories sit *better* with Distributed Cognition than do connectionist theories: they see (or impose) a stricter symmetry or parity between the format and dynamics of the fixed stock of symbols in the brain, and (at least some) external symbols like notebooks. I argue in the rest of this section that connectionist-style theories, however, offer a better picture of the complementary relations between internal and external representations.

[9] It's historically awkward to use the term 'eliminativist' for these views, even though it is accurate, because it is apt to be confused with eliminativism about folk or common-sense psychology. But eliminativists in the latter sense, notably the Churchlands, are emphatically *not* eliminativists about mental representation (Sutton, 1999).

[10] This is a plausible interpretation at least of the strand in van Gelder's thought which focusses on clarification of the radically different (non-Language-of-Thought-style) properties of connectionist representations (van Gelder, 1991).

So issues about the existence and the format of mental representation are (and should be) distinct from issues about distribution. I am going to argue now for one particular combination of positions, but there are a cross-classifying range of coherent options here. It may be that neither debate is yet well formulated [11]: I suggest that the laudable motivations behind the integrationist's rejection of fixed codes do not in fact rule out revised accounts of mental representation, and offer a sketch of what such revised accounts might look like.

We can get at this by examining another reason why, in the distributed cognition framework, external processes and structures can be themselves (part of) cognitive processes and states. It is not that the outer resources—notebooks, sketchpads, and so on—have the same format and dynamics as the internal (neural and bodily) resources with which they are coupled in intelligent activity. Especially for connectionist sympathizers like Andy Clark (Clark, 1993), the internal aspects of such distributed systems and processes have quite different formats and dynamics from the typically more static and passive external symbol systems. Internal plasticity is the norm. In the brain's neural networks, information is 'stored' only superpositionally, in the enduring but modifiable connection weight matrices of the network. So biological memory is reconstructive, with recall in a context (for example) driven by the non-explicit information sedimented in the network through its experience. This is the connectionist equivalent or realization of Stephen Cowley's vision of organisms' brains sculpted by their history. 'Distributed cognition', then, is a convenient label for a set of sciences of the interface between plastic, permeable inner resources and more stable outer resources: crudely, it's just because there's relatively little fixity in the brain that we scaffold ourselves so thoroughly with artefacts and other people. So, as Clark puts it, the external media we use as our cognitive technologies are 'best seen as alien but *complementary* to the brain's style of storage and computation. The brain need not waste its time *replicating* such capacities. Rather, it must learn to interface with the external media in ways that maximally exploit their particular virtues' (Clark, 1997, p. 220). [12]

So distributed cognition still allows the brain to *have* a style of 'storage and computation'. Indeed there is no obvious reason to see the framework as incompatible with a generic computational theory of mind, provided computation is broadly construed: if computational processes are a particular variety of causal processes, those in which the transitions between states exhibit some kind or degree of semantic coherence, then both connectionist and distributed cognition

[11] The eliminativist-representationist front, for example, has only recently been clarified by Mike Wheeler's helpful decomposition of debates about internal representation into distinct, empirically-accessible claims about the particular organization of compartmentalized but communicating subsystems (Wheeler, 2001).

[12] This 'complementarity' argument for the extended mind should take precedence over the 'parity principle' (Clark and Chalmers, 1998). There's a helpful clarification of these points in Clark (1998), which is responding to O'Brien (1998).

frameworks allow for computation. [13] There's nothing in the idea of computation itself which specifies that only the brain can be the locus of computational cognition, or that the representations over which computations occur must be atomic and fixed. This does not *prove* that distributed cognition needs representation and computation. But it means that attempts to get by without these notions need strong alternative explanations of the various phenomena which have always prompted their invocation. It's extremely suggestive, for sure, and relatively neglected in cognitive science, that sometimes when we remember experiences in the distant personal past, or think about things which are absent or highly abstract, we can only do so (or can do so much more easily) in the presence of relevant external representations. [14] But we can also sometimes do these things without such scaffolding in our immediate present environment, [15] an ability which still seems on the face of it to require some primarily inner trace or representation (Clark, 1997, pp. 166–170). This is compatible with the point that the particular coalescence of internal resources driving the exercise of such a capacity may be temporary and context-specific: no extraction of a fully-formed item from long-term cold storage is required. Further, a number of recognizably 'distributed cognition'-style empirical programmes seem to require representational talk, applied equally to artefacts and to brains. Ed Hutchins (1995), for example, wants to be able to track the causal path of particular contents across the various vehicles—individual brains and bodies, machines and instruments, and so on—which can carry them in the course of complex navigation. It only makes sense to ask questions about transmission and potential distortion across media if we can trace a specific representation and its antecedents and descendants. [16] And finally, a core reason for Clark's rejection of 'telementational' ideas about the expressive function of language is that linguistic activity has significant effects on the mental representations of the language-user. If language is 'the ultimate artefact' (Clark, 1997, p. 193), it is because we and it have conspired, over both historical and developmental time, to gain us the capacity to think differently. For example, linguistic labels can turn thoughts or other fleeting mental phenomena into different kinds of object, increasing their stability, allowing us to consider, reconsider, reflect on, and re-use them: this, for Clark, is the basis of our capacities for self-criticism, and in turn such 'thoughts about thoughts' help us restructure the

[13] The fact that 'semantic coherence' and related notions like 'reason-respecting behavior' (Clark, 2001b, pp. 2–16) are themselves matters of degree does not rule out an objective account of computation which relies on them.

[14] It's not, in cases like these, that there are *no* inner traces, but that in the period before coupling with external triggers or cues, such traces only form memories which are (to greater or lesser extent) *incomplete.*

[15] And even when such scaffolding is present, it is natural at least to invoke internal representational differences to explain the psychological and behavioural differences between the system which, due to its history, is able to exploit it, and systems which cannot so exploit it.

[16] Compare also the epidemiology of representations developed by Dan Sperber (Sperber, 1996).

(physical and social) environment. [17] None of this could get going if there were no mental representations, distributed for sure but partly inner, for linguistic activity to influence.

Where does this leave the integrationist? Does Integrational Linguistics have arguments against even flexible, action-oriented, context-sensitive mental representations? Harris argues that there are no determinate relations between form and meaning for *any* sign: 'integrational theory treats the intrinsic indeterminacy of the sign as the foundation' (Harris, 2000, p. 72). And although Stephen Cowley does accept something called 'organic computation', which needs further explication, he denies that there are 'tokens of determinate value' which 'establish symbolic reference' (Cowley, 2002). There is no definite psychological typology: regularities are not to be found in the brain, but only in activities. It can only be a metaphor, for the integrationist, to say that tokens are produced and processed in the brain: to forget this, as in Cowley's view many theorists do, shows the lack of a concept of contextualization. For Harris, there are no sensible questions to be asked about the 'sameness' of a sign across contexts: what matters is only how activities are integrated on the basis of our uses of resources in the world.

Connectionist versions of distributed cognition theory do sometimes get close to this perspective, as they do to their shared ancestor, the considerations against a private language found in the later writings of Wittgenstein. But as with those considerations, the philosophers of distributed cognition want (for the reasons outlined above) to accept the underlying motivations without jettisoning all reference to mental representation. 'Tokens' do not have to be a fixed code or a permanent store. It may well be that often we create representations at the moment we need them: in autobiographical remembering, for example, we do not pull an intact experience out of a passive episodic memory system (Engel, 1999, p. 6). [18] The materials which drive such context-ridden processes of reconstruction operate in conspiracies of internal and external causes. There are real and difficult questions here about the identity of distributed implicit 'representations' (Ramsey, 1997): yet we want to be able to understand not only occurrent episodes of psychological activity, but also enduring or dispositional psychological states; and to be able to model (and refer in causal explanation to) the kinds of transforming, distorting, or generalizing processes which operate over time on such enduring implicit representations.

So theorists of distributed cognition can agree with Harris that a sign is not a 'discrete autonomous entity', if that means (as Harris takes it to) that it must be 'rather

[17] See also Millikan, 2001. This does not mean, however, that having 'thoughts about thoughts' turns the mind into a linguistic engine (or a 'Joycean machine'). Nigel Love (2004) on occasion reads Clark's remarks about the cognitive roles of, for example, 'the mental rehearsal of sentences' as if they explicated Clark's view of *all* mental representations: but of course, for post-connectionists like Clark, thought is not itself linguiform. This is particularly clear in Clark's exchanges with Daniel Dennett—see for example Clark (2002). For Love's further and more pertinent argument against even 'a realm of non-linguistic thoughts or ideas' see my comments in Section 6.

[18] In his account of a dispositional representation as a 'dormant firing potentiality' Damasio (1995, pp. 103–104) nicely compares the mythical town of Brigadoon which becomes visible only once every century: 'dispositional representations exist in potential state, subject to activation, like the town of Brigadoon'.

like a chess piece being returned to its box after the game, ready to come out again when next needed' (Harris, 2000, p. 82). They can also accept the caution that we should not assume in advance that representations, especially internal representations, *must* underlie intelligent and integrative activity. But they can rightly deny that there are no sensible questions to ask about signs across contexts, that it makes no sense to look for representations underlying integrative activity, and that we can achieve all our explanatory goals without invoking representations.

5. Reduction, levels, and subpersonal explanations

Roy Harris argues that the *mechanisms* by which I recognize and contextualize anything in a particular way in relation to certain activities should be investigated in the neurosciences, and not in linguistics. The possibilities and the limits of human communication, he argues, are set simultaneously by 'biomechanical', 'macrosocial', and 'circumstantial' factors. So in addition to the roles of context and of culture, the bodily capacities and state of the communicating agent are given their place: there is, however, no acknowledgement of an intermediate level of enquiry into cognitive or representational processes. To understand the integrationist programme, then, we need to see just what kind of 'biomechanical' factors are considered legitimate. To the physical properties of the human agent and of the symbolic media used by the agent Cowley adds the brain, which constrains language, but emphasizes that its use is in turn constrained by linguistic practice.

What then is the proper subject-matter of Integrational Linguistics? What are its relations, as a discipline and a mode of enquiry, with other apparently connected research programmes? Integrationists argue that their project has and should have a significant degree of autonomy:

> The business of integrational linguistics is not to be confused with neurophysiology, or speech pathology, or auditory phonetics, any more than watching cricket is to be confused with observing how bats are made or pitches laid. (Harris, 1998, p. 146).

The cricketing analogy suggests a peaceful disconnected coexistence of the activities mentioned: even if the processes and mechanisms studied in 'lower-level' sciences are necessary preconditions or grounds for the communicative activities studied by the integrationist, it would be absurd to think that understanding them would throw any light on the 'higher-level' phenomena. [19] Integrational Linguistics is about

[19] I'm going to use 'levels' talk in this section as if it was clear, which it is not. The kinds of looping and coupling effects dear to distributed cognition theorists do problematize the neat hierarchical pictures of levels which we have inherited from the days of the classical reductionist 'unity of science' doctrine. But though I cannot argue the case here, I believe that levels talk can in fact be rendered precise and defended within such more complex frameworks, along the lines developed for example by Carl Craver (2002): for a contrary view see Machamer and Sullivan, 2001.

embodied agents, not their components. Where machines compute, according to Cowley, living bodies contextualize. Not brains, but conscious human bodies are masters of symbolic reference, and the tasks those bodies engage in are separate from the neural instruments used (Cowley, 2002). Integrational Linguistics, then, is to operate at the personal level of explanation, which is autonomous from and irreducible to any subpersonal processes and mechanisms: this personal level, however, is not to be understood psychologically, but as the level of ordinary integrative agency in the natural and social world.

But in fact it is not clear just how neuroscientists should proceed even if they wanted to maintain the neat division of labour which Harris suggests. [20] In trying to deal with the mechanisms and processes underlying (for example) recognition and contextual action without violating integrationist principles, they would have to abjure all allegedly illegitimate psychological constructs, and studiously avoid treating their accounts of subpersonal processes as if (even collectively, even when set in context) they could explain personal action and contextualization. Unless the neurosciences are implausibly restricted to the study of basic anatomy (and perhaps not even then), functional hypotheses are required to get even structural inquiry going, or to identify and offer candidate taxonomies of the explananda.

Rather than taking this point into a discussion of general debates in philosophy of science about reduction and levels of explanation, I will again compare integrationist doctrine on these issues with a range of views in the literature on distributed cognition. Some philosophers see the idea that the vehicles of cognition are distributed across brain, body, and world as itself primarily pitched at a subpersonal level (Hurley, 1998a). But this is not the only coherent view. The personal/subpersonal distinction does not play a significant role, for example, for Andy Clark, who takes the distributed cognition framework to be applicable to the self or agent too. Notions like 'person', 'responsibility', and 'agency' are not protected from the potential leakage of mind into the world (Clark, 2003).

In some respects the resulting picture seems close to the integrationist's: although 'biomechanical' factors are reinterpreted as subpersonal cognitive factors, we are to investigate these in social and 'circumstantial' context, with no clear and discrete personal level in between. I am far from sure how to interpret integrationists on this point, but would like better to understand their resistance to this dissolution of the subject into the subpersonal and the social at once. It seems odd to say, with Cowley, that 'persons exploit brain dynamics to embody contextualizing activity': one wants to know who precisely is doing this exploiting, who is the user, and what is the nature of the user's knowledge of and access to such brain dynamics. It might be more consistent simply to say that contextualizing activity occurs, within a less individualist picture of personal identity.

This move would bring integrationism closer to one strand in the distributed cognition movement, which is to reject the idea of the person as prior to and

[20] Again, the same problem arises when Wittgensteinians, Gibsonians, or phenomenologists seek to distinguish purely scientific projects in the neurosciences from illegitimate, philosophically-tainted confusions of neuroscience and psychology. See for example Bennett and Hacker, 2003.

unencumbered by objects, tools, aids and so on. The autonomous individual as the possessor and independent user of artefacts, and as the owner and controller of memories and other mental states, is seen as an unfortunate residue of psychological internalism, which relies on an unrealistic picture of control (Clark, 2003). Instead there are simply brains and bodies interfacing (or contextualizing) variously with each other and with technologies and media in various contexts. When personal-level talk is solidified into an autonomous theoretical level, instead, it acts as an oddly static bottleneck between the dynamic processes studied in both subpersonal and social enquiry. [21]

I shouldn't give the impression that this is a mainstream perspective within distributed cognition. For the moment I want simply to point out that integrational linguists cannot easily go this way: they retain an autonomous personal level in being committed to 'lay-oriented linguistics'. For Harris, everyone is necessarily a linguist, because everybody deals with the contextual nature of communication, and thus has some grip on the biomechanical, macrosocial, and circumstantial factors which enable and constrain it. There can be no 'real' linguistic facts available to science alone, of which the competent lay speaker could be ignorant. A 'scientific linguistics' would be full of myths, arcane, and of no general interest, driven by an almost pathological urge to strip away the ordinary messiness of communicative activity in search of a less coarse and confusing pure or basic language underneath.

The idea that linguistics is and should be essentially lay-oriented seems at first bewildering, like the corresponding claim that common-sense psychological understanding of mind and activity is and should be the only kind of psychology. For integrationists, it is in mutual support with the attack on representations and on any postulated 'deep' mental processes and states which might be inaccessible to the subject. Lay-oriented linguistics, to make an initial point, will certainly be incompatible with almost any version of distributed cognition, since theories and particular explanations in the latter framework pay no particular respect, and offer no particular priority, to ordinary views about the mind, which are often marked with individualist assumptions. Though I will qualify this point at the end of this section, I want first to run through three problems, of increasing importance, with the focus on lay or common-sense understandings of mind and communicative activity.

If 'lay-oriented linguistics' were implemented literally, to the extent that whatever ordinary speakers think about language and communication is the end of the matter, then no-one could ever be wrong on these matters, and linguistics would be a mess of incompatible beliefs. Clearly this is not the idea: but it is hard to find principled reasons in the integrationist literature for taking some lay views and practices seriously and not others.

Harris accepts, of course, that much lay thinking about language is tainted by 'Western theories' of representation, and by the 'language myth'. But the integrationist somehow stands outside of the myth in order to criticise it, even if it is a lay myth and integrationism is essentially lay-oriented. There are threats of self-refutation

[21] Thanks to Don Ross for discussion on this point.

here: whatever kind of theory, or framework, or even attitude, Integrational Linguistics is meant to be, it is not easy to specify principled reasons to think that it alone can justifiably take a neutral view, untainted by the problematic aspects of lay thinking.

Behind these first two concerns hides a deeper difference between integrational and mainstream linguistics and cognitive science, about method, explanation, and theory. Harris's ideal is a kind of linguistics which not only is not itself a natural science, but which has no particular relations with the natural sciences. This is partly because of his focus on context, the topic of the following section. For now it's worth pointing out that Clark's distributed cognition forcefully challenges the legitimacy of the kind of divisions between natural sciences and social sciences on which Harris relies. While Clark expects the distributed cognition framework to come up with surprising results which may not be easily assimilable into ordinary lay views, Harris's restriction of linguistics to (suitably purified versions of) common-sense understandings threatens to paralyze enquiry.

I see two possible ways out of this impasse. Perhaps Integrational Linguistics need not be so committed to being lay-oriented: and indeed there seems to be an impressive corpus of detailed theory already in place under the integrationist banner which goes far beyond anything lay speakers come up with in the course of ordinary thinking about communicative activity.

Or, to try to find some way better to respect the lay orientation of Integrational Linguistics, we could ask whether in fact there might not be significant strands of common-sense thinking about mind and language which are closer to the integrationist and distributed perspectives. Perhaps, outside highly selective and specific elite cultural contexts in the modern West, folk have long been and often still are happy to see language as deeply embedded in contextualizing activities, and cognition as distributed across bodies and world. Without assuming that the folk are always right, this line of thought would encourage us to seek closer contact with both phenomenological and common-sense thinking than has been traditional in mainstream linguistics and cognitive science.

6. Context and explanation

Both distributed cognition and Integrational Linguistics take context very seriously, and problems arise for both frameworks as a result. Attention to context seems to some to make the right kind of science impossible. Theorists of distributed cognition can just bite the bullet on this point, and in opting for slightly different kinds of science take on the burden of demonstrating the cash value of their vision. It is not clear, however, that the integrationist has the parallel resources to get out of jail here: in this section I show how the problem arises in each framework, and suggest some directions for possible integrationist responses.

Distributed cognition examines 'cognitive and computational architectures whose bounds far exceed those of skin and skull' (Clark, 2001a, p. 138). But such architectures include a daunting array of the social and technological systems with which

embodied brains can couple. What would cognitive science be like, how could it continue, if its objects include notebooks, sketchpads, and tattoos as well as brains? Critics of the framework complain that systems of brains coupled with cognitive tools, computing devices, or memory aids 'would seem to form such a motley collection that they will not form the basis for any significant scientific theorizing' (Adams and Aizawa, 2001, p. 63). Distributed cognition thus threatens to thwart cognitive science's connected quests for natural kinds and for disciplinary identity. If scrabble tiles, spice racks, cocktail glasses, slide-rules, incised sticks, shells, languages, rules and laws, moral norms, knots, codes, maps, diagrams, fingers, software devices, rituals, monuments, rhythms and rhymes, and roads (to take just a sample from the distributed cognition literature), can count as part of the legitimate subject matter of the sciences of mind, is not the project obviously absurd?

Clark's answer to this challenge has two strands (Clark, 2004). On the one hand, we should not work with an overly restricted or puritanical notion of scientific explanation; on the other hand, we should not rule out in advance the possibility that in fact there may be higher-level accounts which do find commonalities or patterns across the 'unscientific motley of capacities' (Adams and Aizawa, 2001, p. 62) exhibited by extended hybrid minds. Carefully contextualized accounts of particular forms of scaffolding and their cognitive roles need not descend into a nightmarish disarray of unrelated descriptions. Distributed cognition still encourages, of course, the study of the brain and its representational capacities as a key term in the cognitive relation: yet neuroscience must be coupled with the social sciences, history, media theory, and so on if we are to get at objective accounts of our cognitive hybridity. While we want to include treatment of both individual differences and historical differences in relation to the technologies and processes of distributed cognition, there is no principled reason to abhor theorizing about the dimensions on which such differences might be found.

Whether or not this strategy turns out to be appropriate in the case of distributed cognition, the same problem arises with even greater force in the case of Integrational Linguistics. Again, attention to context seems to rule out the search for patterns which characterizes even modest aims in explanation. For Harris, a sign 'has no continuous identity outside the contextualization that brought it into existence' (Harris, 2000, p. 82). Written signs have the same ephemerality as speech. The integrational constraints alter in every context, and there are always multiple contexts for any episode of integrative activities. The point is that contexts are not extrinsic to signs, and thus dispensable in analysis: rather they are built in to the act of use. [22] The net must be cast wide in addressing relevant contexts: even, for example, attitudes or implications about morality or authority can be built in to integrative activity. Harris offers an extended meditation on the contextual nature of writing your

[22] One epigraph to Harris's *Signs*, *Language and Communication* (1996) is from Giulini: 'The performance is ''this moment''. You can never do twice the same performance, thanks God'. Harris interprets claims or thoughts like this (if there can ever be two similar claims or thoughts on his view) as entailing that it makes no sense to ask whether and in what respects two performances are similar or different: lack of identity across contexts is taken to imply total dissimilarity across contexts.

signature, noting how different this activity is from merely copying it (Harris, 2000, p. 161–183). Since both writers and speakers create signs at a moment, Harris argues that there is no point in seeking a theory of signs-in-general, or of the nature of signs abstracted from specific uses. For Cowley, too, no utterance or other integrative activity is context-free, and every contextualization is necessarily unique; and for Love (2004), 'signhood is a transient, not a permanent, property of a sign'.

There are two kinds of ways to interpret this careful attention to context. I will recommend that it should be thought of as a regulating ideal for explanation, such that in principle the patterns and generalizations sought in our sciences of mind and language should have sufficient flexibility and richness to be applied in thick description of local historical or individual action. On the other hand, it could be thought of as an extremely pure form of empiricism, in which the uniqueness of particular constellations of events or contextualizing activities is treated as their total dissimilarity. At times, Harris's version of Integrational Linguistics seems to take the latter path: we are not even to seek patterns or generalizations across the different instantiations of signs in different contexts, we are not to theorize or try to offer explanation beyond the single case. This would make it hard to see how Integrational Linguistics can proceed or generate empirical programmes, and would raise questions about its utility in driving enquiry.

On this stronger integrationist view, only *occurrent* communicative activity is a proper object of study; there can be no investigation of any enduring or dispositional features of signs, and thus no patterns or similarities across different contexts of use to be investigated. This point underpins Love's critique (2004) of Clark: Love rejects the invocation of *any* mental representations, even 'a realm of non-linguistic thoughts or ideas', on the ground that such thoughts would have to be 'invariants that remain the same', definite entities, repeatable without loss of meaning, and unambiguously identifiable as *the same* across contexts. Love rightly points out that there are no such things, no fixed codes, no private languages: but he is wrong to think that Clark is invoking such mysterious context-free inner symbols. The fact that context-dependence renders the possibility of *identity* across contexts problematic does not, as Harris and Love assume, thereby rule out the possibility of (investigable) *similarities* across contexts.

To put the point another way, I would argue that context-dependence is itself a matter of degree. The extent to which meanings shift with contextual change is itself not constant across contexts: so our understanding of cases in which there *is* some degree of stability across contexts *requires* a framework in which the very idea of semantic stability across contexts makes sense. Indeed Harris's own practice implicitly depends on this point: he accepts that the media or vehicles of symbolic activity differ on a range of dimensions connected with semantic stability across contexts. The printed book, for example, is 'the least context-bound' of all forms of writing, approximating the kind of portable and autonomous set of representations which the 'language myth' erroneously assumes in advance. We should agree with Harris in rejecting *in general* the idea of 'medium-transferability', the strict identity or constancy of contents over different material instantiations: but, and as a consequence, we should expect to find different particular kinds of medium-transferability in

relation to the use of signs in different situations. Indeed the literature of distributed cognition includes well-developed frameworks for addressing the different properties of different external structures: in comparing internal and external symbolic structures (or 'exograms') in understanding the cultural evolution of cognition, for example, Merlin Donald picks out a number of key properties on which such structures can differ both from each other and over time (Donald, 1991, p. 315). Admittedly such attention to external structures alone needs to be supplemented by analyses of their uses, and by historically-specific narratives of technological and social change. But there is no reason to accept an inference from the uniqueness of contexts to their total incommensurability.

On the issue of context as well, then, the possibility of developing a shared vision for Integrational Linguistics and distributed cognition depends on the integrationists dropping some more extreme glosses on their important insights. To ensure that the appropriate care for particularity is not translated into a ban on explanation or on the search for specific similarities across occasions, integrationists could reinterpret existing empirical work in their tradition as already investigating different continua along which to understand particular kinds of contextualized signs.

7. Cash value: memory development

I have suggested that the most valuable insights of the integrationist tradition, in both its critical and constructive modes, can be retained while dropping the more extreme claims which put it in tension with the distributed cognition framework. I want to finish with brief remarks about one strand of current interdisciplinary enquiry into memory, which demonstrates how psychological enquiry into cognition and language might respect the integrationists' insights. This work in the developmental psychology of autobiographical memory also shows the intellectual and institutional possibilities for integration of these insights with work in neighbouring disciplines: this openness to related but differing theory might also serve as a useful model for Integrational Linguistics.

This developmental work on early autobiographical memory relates to an increasing consensus in recent cognitive psychology about the reconstructive nature of remembering. As I've characterized this elsewhere (Sutton, 2003; compare Schacter, 2001), this involves three main strands which should each be compatible with integrationist thinking: a search for multicausal accounts of any episode of remembering, a stress on the plasticity of any 'storage' system (whether internal or external), and a focus on the key explanatory role of context, especially the contexts of recall. [23] In the resulting general vision of adult human memory, remembering is seen

[23] In relation to the concerns about representation discussed in section 4 above, it is worth underlining that this post-connectionist perspective on dynamic internal processes still includes space for non-occurrent memory representations (many 'stored' or embodied in the same set of weights, for example), which are more like embodied dispositions or habits than 'fixed codes'.

as an integrative activity in the present rather than an automatic extraction of a permanent trace from cold storage. This opens the way to a potential general integration of social and interpersonal analyses of memory with mainstream psychological accounts of processes within the individual.

Social-interactionist research in the developmental psychology of memory exemplifies precisely these themes. Like much work in distributed cognition, this tradition relies on Vygotskian notions of scaffolding: just as in motor development, where children initially rely on external aids before coming to embody relevant skills by themselves, so in cognitive development the particular paths by which capacities are acquired and honed may leave traces as the socially-available resources are idiosyncratically internalised. Although children talk about the past almost as soon as they begin talking, adults initially provide much of the content and the structure of their references to the past. From early abilities to capture generic regularities, temporal sequences, and routines of action, children gradually develop more sophisticated abilities spontaneously to remember *particular* episodes in their personal past. Memory sharing practices, usually at first initiated by adults, encourage the child's acquisition of the notion that there can be different perspectives on the same once-occupied time, so that the past becomes a potential object for negotiation, shared attention, and discussion (McCormack and Hoerl, 1999).

While there is much uncertainty about the necessary conditions for the development of these full autobiographical memory abilities, and increasing agreement that integrative multi-factor accounts are required, the social-interactionist research of Katherine Nelson, Robyn Fivush, and others offers a powerful angle on the process (see Nelson and Fivush, 2000 for a review and survey). Parental and cultural models and examples for the recounting of past events act as scaffolding on which children start to hang their own memory narratives. They may then, to differing degrees, internalise the forms and narrative conventions appropriate to their context. The framework has produced rich empirical evidence for individual, gender, and cultural differences (see Sutton, 2002a for review and discussion). What's worth noting here is that this work is not necessarily anti-cognitivist, and does not require neglect of either brain processes or the internal aspects of representational processes: indeed recent work suggests that there may be quite distinct pathways to similar developmental outcomes in autobiographical memory, with some children leaning more heavily on the cultural and linguistic resources of their local narrative environment, and others relatively shielded from these influences (Harley and Reese, 1999).

Language can thus play a key role in shaping autobiographical memory both in particular contexts and over time, through its effects on a representational system. Without assuming that language is itself the medium of cognition or of autobiographical memory, we can see ways in which local narrative resources help to 'freeze' thoughts about the past, allowing children to develop the perspective-switching abilities which ground our understanding of the past. So the case of autobiographical memory development shows how to see strong contextual factors operating often *by way of* their effects on the child's representational capacities and their exercise.

8. Conclusions

I've offered sketches of the distributed cognition and Integrational Linguistics frameworks so as to highlight some shared intellectual motivations and assumptions, then pinpointed three potential areas of conflict or difficulty. Where integrational linguists are sceptical about any invocation of mental representations, I've suggested their criticisms should be applied only to rigid atomistic conceptions of internal representation, and that they should accept the more flexible picture of representation developed by post-connectionist theorists of distributed cognition. Where integrationists seem committed to the autonomy of personal-level explanation, so that neuropsychological accounts are excluded from their 'lay-oriented' disciplines, I have suggested that this uneasiness about cross-level explanation and about the subpersonal is dispensable, and has significant costs. And where integrationists' admirable commitment to context makes it hard to get an explanatory programme going, we should accept the inevitability of multiple contextualization and still get on with the work of seeing which contexts matter, and why, rather than remaining squeamish about the whole project of generalizing across contexts. Attention to the sculpting of brains by history and context which integrationists like Cowley, at least, envisage, which allows for substantial individual and cultural differences within a broad developmental system, does not require us to jettison mental representation, cross-level influences, or explanation.

If this is the cost of an alliance with the distributed cognition framework, the integrationist may indeed be tempted to refuse the offer. If distributed cognition is still taken to be burdened by the many myths of mainstream linguistics and cognitive science, integrationists may decide to remain free from the proffered integration. I have argued that this continuing isolation and insulation is both unnecessary and undesirable. The future development of both frameworks in an integrative rather than combative spirit is certainly not necessary, but might be of considerable mutual advantage.

Acknowledgments

Many thanks to Maria Trochatos for research assistance. I'm deeply grateful to David Spurrett and Stephen Cowley for suggesting the topic and setting up the possibility for discussion of it, and to all participants in the Mind and World conference in Durban. Comments by Akeel Bilgrami, Wayne Christensen, Laurence Goldstein, Nigel Love, Richard Menary, Don Ross, and Mike Wheeler were particularly helpful. Conversations with Karen Salmon have aided my understanding of the developmental research mentioned in Section 7.

References

Adams, F., Aizawa, K., 2001. The bounds of cognition. Philosophical Psychology 14, 43–64.
Bennett, M.R., Hacker, P.M.S., 2003. Philosophical Foundations of Neuroscience. Blackwell, Oxford.

Brooks, R., 1991. Intelligence without representation. Artificial Intelligence Journal 47, 139–159.
Carruthers, P., 2002. The cognitive functions of language. Behavioral and Brain Sciences 25, 657–674.
Clark, A., 1993. Associative Engines. MIT Press, Cambridge, MA.
Clark, A., 1996. Linguistic anchors in the sea of thought. Pragmatics and Cognition 4, 93–103.
Clark, A., 1997. Being There: Putting Brain, Body, and World Together Again. MIT Press, Cambridge, MA.
Clark, A., 1998. Author's response. Metascience 7, 95–104.
Clark, A., 2001a. Reasons, Robots, and the extended mind. Mind and Language 16, 121–145.
Clark, A., 2001b. Mindware: An Introduction to the Philosophy of Cognitive Science. Oxford University Press, Oxford.
Clark, A., 2002. Minds, brains and tools. In: Clapin, H. (Ed.), Philosophy of Mental Representation. Clarendon Press, Oxford.
Clark, A., 2003. Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence. Oxford University Press, Oxford.
Clark, A., 2004. Memento's Revenge. In: Menary, R. (Ed.), The Extended mind. Ashgate, Aldershot. Forthcoming.
Clark, A., Chalmers, D., 1998. The extended mind. Analysis 58, 7–19.
Cowley, S.J., 2002. Why brains matter: an integrational perspective on the Symbolic Species. Language Sciences 24, 73–95.
Cowley, S.J., 2004. Contextualizing bodies: human infants and distributed cognition. Language Sciences 26 (6), 565–591.
Cowley, S.J., Spurrett, D., 2003. Putting Apes (body and language) together again. Language Sciences 25, 289–318.
Craver, C., 2002. Interlevel experiments and multilevel mechanisms in the neuroscience of memory. Philosophy of Science (Supplement) 69, S83–S97.
Damasio, A., 1995. Descartes' Error: Emotion, Reason, and the Human Brain. Picador, London.
Donald, M., 1991. Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition. Harvard University Press, Cambridge, MA.
Donald, M., 2000. The central role of culture in cognitive evolution. In: Nucci, L. et al. (Eds.), Culture, Thought, & Development. Erlbaum, New Jersey, pp. 19–38.
Eliasmith, C., 1996. The third contender: a critical examination of the dynamicist theory of cognition. Philosophical Psychology 9, 441–463.
Engel, S., 1999. Context is Everything: The Nature of Memory. W.H. Freeman, New York.
Fodor, J.A., 1976. The Language of Thought. Thomas Crowell, New York.
Gauker, C., 1994. Thinking Out Loud: An Essay on the Relation between Thought and Language. Princeton University Press, Princeton, NJ.
Gauker, C., 1999. Language and Thought. Available from: <http://host.uniroma3.it/progetti/kant/field/lat.htm>.
Gilbert, M., 1989. On Social Facts. Princeton University Press, Princeton, NJ.
Harley, K., Reese, E., 1999. Origins of autobiographical memory. Developmental Psychology 35, 1338–1348.
Harris, R., 1996. Signs, Language and Communication: Integrational and Segregational Approaches. Routledge, London and New York.
Harris, R., 1997. From an integrational point of view. In: Wolf, G., Love, N. (Eds.), Linguistics Inside Out: Roy Harris and his critics. John Benjamins, Amsterdam/ Philadelphia, pp. 229–310.
Harris, R., 1998. Introduction to Integrational Linguistics. Pergamon/ Elsevier, Oxford.
Harris, R., 2000. Rethinking Writing. Athlone Press, London.
Haugeland, J., 1998. Mind embodied and embedded. In: Haugeland, J. (Ed.), Having Thought: Essays in the Metaphysics of Mind. Harvard University Press, Cambridge, MA, pp. 207–237.
Hurley, S., 1998a. Consciousness in Action. Harvard University Press, Cambridge, MA.
Hurley, S., 1998b. Vehicles, contents, conceptual structure, and externalism. Analysis 58, 1–6.
Hutchins, E., 1995. Cognition in the Wild. MIT Press, Cambridge, MA.
Joseph, J., 2003. Orthodox Unorthodoxy (review of Harris 1998). Language Sciences 25, 99–109.

Latour, B., 1999. A collective of humans and nonhumans. In: Latour, B. (Ed.), Pandora's Hope: Essays on the Reality of Science Studies. Harvard University Press, Cambridge, MA, pp. 174–215.
Love, N., 2004. Cognition and the Language Myth. Language Sciences 26 (6), 525–544.
McCormack, T., Hoerl, C., 1999. Memory and temporal perspective: the role of temporal frameworks in memory development. Developmental Review 19, 154–182.
Machamer, P., Sullivan, J., 2001. Levelling Reduction. Pitt-London Workshop in the Philosophy of Biology and Neuroscience. Available from: <http://philsci-archive.pitt.edu/archive/00000400/>.
Millikan, R., 2001. The language-thought partnership: a bird's eye view. Language and Communication 21, 157–166.
Nelson, K., Fivush, R., 2000. Socialization of memory. In: Tulving, E., Craik, F.I.M. (Eds.), The Oxford Handbook of Memory. Oxford University Press, Oxford, pp. 283–295.
O'Brien, G., 1998. The mind: embodied, embedded, but not extended – review of Clark 1997. Metascience 7, 78–83.
Preston, J., 1997. Introduction: thought as language. In: Preston, J. (Ed.), Thought and Language. Cambridge University Press, Cambridge, pp. 1–14.
Ramsey, W., 1997. Do connectionist representations earn their explanatory keep? Mind and Language 12, 34–66.
Renfrew, C., Scarre, C. (Eds.), 1998. Cognition and Material Culture: The Archaeology of Symbolic Storage. Macdonald Institute, Cambridge.
Rowlands, M., 2003. Externalism: Putting Mind and World Back Together Again. Acumen, Chesham.
Schacter, D.L., 2001. The Seven Sins of Memory: How the Mind Forgets and Remembers. Houghton Mifflin, Boston and New York.
Sperber, D., 1996. Explaining Culture: a Naturalistic Approach. Blackwell, Oxford.
Sterelny, K., 2000. Roboroach, or, the extended phenotype meets cognitive science. Philosophy and Phenomenological Research 61, 207–215.
Strauss, C., Quinn, N., 1997. A Cognitive Theory of Cultural Meaning. Cambridge University Press, Cambridge.
Sutton, J., 1998. Philosophy and Memory Traces: Descartes to connectionism. Cambridge University Press, Cambridge.
Sutton, J., 1999. The Churchlands' Neuron Doctrine both cognitive and reductionist. Behavioral and Brain Sciences 22, 850–851.
Sutton, J., 2002a. Cognitive conceptions of language and the development of autobiographical memory. Language and Communication 22, 375–390.
Sutton, J., 2002b. Porous Memory and the Cognitive Life of Things. In: Tofts, D., Cavellero, A., Jonson, A. (Eds.), Prefiguring Cyberculture. Power Publications and MIT Press, Sydney and Cambridge, MA, pp. 130–141.
Sutton, J., 2003. Memory. In: Zalta, E.N. (Ed.), The Stanford Encyclopedia of Philosophy (Summer 2003 Edition). Available from: <http://plato.stanford.edu/archives/sum2003/entries/memory/>.
van Gelder, T., 1991. What is the D in PDP? a survey of the concept of distribution. In: Ramsey, W., Stich, S.P., Rumelhart, D.E. (Eds.), Philosophy and Connectionist Theory. Lawrence Erlbaum, Hillsdale, NJ, pp. 33–59.
van Gelder, T., 1995. What might cognition be, if not computation? Journal of Philosophy 92, 345–381.
Wheeler, M., 2001. Two threats to representation. Synthese 129, 211–231.
Wilson, R.A., 1994. Wide Computationalism. Mind 103, 351–372.