

Material Agency

Carl Knappett • Lambros Malafouris
Editors

Material Agency

Towards a Non-Anthropocentric Approach

 Springer

Editors

Carl Knappett
University of Toronto, Toronto,
Ontario, Canada

Lambros Malafouris
University of Cambridge
Cambridge, UK

ISBN: 978-0-387-74710-1 e-ISBN: 978-0-387-74711-8
DOI: 10.1007/978-0-387-74711-8

Library of Congress Control Number: 2008930774

© 2008 Springer Science+Business Media, LLC

“Where Brain, Body, and World Collide” reprinted with minor changes by permission of *Daedalus*, Journal of the American Academy of Arts and Sciences, from the issue entitled, “The Brain,” Spring 1998, Vol. 127, No. 2

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden. The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

springer.com

Contents

1	Where Brain, Body and World Collide	1
	Andy Clark	
2	At the Potter's Wheel: An Argument for Material Agency	19
	Lambros Malafouris	
3	Material Agency, Skills and History: Distributed Cognition and the Archaeology of Memory	37
	John Sutton	
4	The Actor-Enacted: Cumbrian Sheep in 2001	57
	John Law and Annemarie Mol	
5	Non-Human Agencies: Trees in Place and Time	79
	Owain Jones and Paul Cloke	
6	Intelligent Artefacts at Home in the 21st Century	97
	Richard Harper, Alex Taylor and Micheal Molloy	
7	In Context: Meaning, Materiality and Agency in the Process of Archaeological Recording	121
	Thomas Yarrow	
8	The Neglected Networks of Material Agency: Artefacts, Pictures and Texts	139
	Carl Knappett	
9	Some Stimulating Solutions	157
	Andrew Cochrane	
10	On Mediation and Material Agency in the Peircean Semeiotic	187
	Christopher M. Watts	

11 When ANT meets SPIDER: Social theory for arthropods 209
Tim Ingold

12 Agency, Networks, Past and Future 217
Sander E. van der Leeuw

Index 249

Contributors

Andy Clark

Department of Philosophy, Edinburgh University, Edinburgh, UK

Paul Cloke

Department of Geography, University of Exeter, Exeter, UK

Andrew Cochrane

School of History and Archaeology, University of Cardiff, Cardiff, Wales, UK

Richard Harper

Microsoft Research Cambridge, Cambridge, UK

Tim Ingold

Department of Anthropology, School of Social Science, University of Aberdeen, Aberdeen, UK

Owain Jones

Countryside and Community Research Institute, Dunholme Villa, The Park, Cheltenham, UK

Carl Knappett

Department of Art, University of Toronto, Toronto, Ontario, Canada

John Law

Centre for Science Studies and Department of Sociology, Lancaster University, Lancaster, UK

Lambros Malafouris

McDonald Institute for Archaeological Research, Cambridge University, Cambridge, UK

Annemarie Mol

Centre for Science Studies and Department of Sociology, Lancaster University, Lancaster LA1, 4YN, UK

Micheal Molloy

Microsoft Research Cambridge, Cambridge, UK

John Sutton

Macquarie Centre for Cognitive Science, Macquarie University, Sydney, NSW
2109, Australia

Alex Taylor

Microsoft Research Cambridge, Cambridge, UK

Sander van der Leeuw

School of Human Evolution and Social Change, Arizona State University,
Tempe, Arizona, US

Christopher Watts

Department of Anthropology, University of Toronto, Toronto, Ontario, M5S
2S2, Canada

Tom Yarrow

Department of Social Anthropology, School of Social Sciences, University of
Manchester, Manchester, UK

Material and Nonhuman Agency: An Introduction

C. Knappett and L. Malafouris

Material and nonhuman agency – surely this is a mistake? Is not agency a solely human property? How then can we devote a whole volume to a topic with such obviously shaky foundations? Certainly, the odds seem to be stacked against us when we think of agency as not only the capacity to act, but also the capacity to reflect on this capacity. A subject may feel his or her arm moving and recognise ‘ownership’ of that movement, but this is not necessarily the same as being able to reflectively understand that he or she is the cause or ‘agent’ of that movement (Gallagher 2007, p. 2). When agency is linked strictly to consciousness and intentionality, we have very little scope for extending its reach beyond the human.

Even those nonhuman entities that seem to threaten this neat equation most – let us say software agents or robotic agent-artefacts – do not get close to fulfilling these criteria of agency, even though they are closely modelled after the human (Suchman 2007; Sørensen and Ziemke 2007). In this view, then, ‘material agency’ is a secondary property, a mirage even, with agency (as consciousness and intentionality) still very much in human hands. Nonetheless, there is some sense, even an anxiety in folk psychology, that the autonomy and interactivity of such artefacts is a step towards agency. But achieving agency is from this perspective a question, essentially, of becoming human.

This human-centred view of agents and artefacts is not limited to those artefacts we design to be like agents. It extends to a much wider and more prosaic world of artefacts and matter, an environment of things that is conceived on our own terms, under our control and designed to serve. We do not give a second thought, on the whole, to chairs, mugs, steps, litterbins, wooden, ceramic, concrete or plastic: these objects are overlooked because we engage with them habitually and haptically every day. They would not serve our ends very well, if we could not overlook them. Designed to be secondary, they have to be secondary, forming the backdrop to our lives, of which we are of course the stars, the decision-makers, the agents. It is common sense that agency should be conceived anthropocentrically – how can it be otherwise? We are centre-stage in our lives, not these artefacts, however mundane, or indeed intelligent.

Although here glossed rather simplistically, this anthropocentric worldview means that the material or environmental counterpoints to human agency have

generally been given short shrift in scholarly discussion. Indeed, while agency is a much-debated theme across the social sciences, the terms of the debate have remained rather narrow, focussing overwhelmingly on the relationship between agency and structure (Emirbayer and Mische 1998, pp. 962–963). Arguments go back and forth over the degree to which agents, which may be, though certainly need not be, human individuals, are free to act in the world. This degree of freedom is socio-culturally mediated (Ahern 2001, p. 109), and thus varies considerably in different settings and societies. It is a fundamental question in how society is constituted, from the top-down or the bottom-up: does power lie with individuals, or with social institutions? This debate, often cast in terms of structure vs. agency, has seen attempts to bridge the divide, in the form of practice theory and structurationist ideas (Archer 1988; Bourdieu 1977, 1990; Giddens 1984). Such ideas are now familiar territory not only across much social science, but also in archaeology, with the work of Bourdieu and Giddens frequently cited (Barrett 2001; Dornan 2002; Gardner 2004; Johnson 2000; Smith 2001). Agency has become something of a buzzword in archaeological theory (Dobres and Robb 2000, p. 3), particularly in post-processual thinking, where it is a byword for individual rather than systemic power in the past.¹

Yet while the concept of agency is much contested, it is done so within the theoretical margins of a narrow anthropocentric perspective. This anthropocentric view of agency is based upon a general agreement about a single undisputable fact: that agency, in the real sense of the word, is a property of the human individual – “the only true agents in history are human individuals” (Giddens and Pierson 1998, p. 89). Whether this individual is conceived through a Cartesian or an existential lens makes no important difference. What is important is that when we speak about agents *proper*, we are referring to human individuals, and preferably of the modern Western-type. In short, agency is an attribute of the human substance.

However, an emergent suspicion of the humanistic determinations of agency can actually be traced quite far back, to the likes of Mauss (1954) and Heidegger (1977). Mauss’s seminal study, *The Gift*, illustrated the fluidity of the boundaries between persons and things and the capacity of the latter to embody and objectify, as well as produce, social consequences. This is a point also reflected in Munn’s observations on the Kula exchange system: “Men appear as the agents defining shell value, but without shells men cannot define their value. In this respect, shells and men are reciprocally agents of each other’s value definition”, (Munn 1983, p. 284). Indeed, as modern anthropology has plainly illustrated “commodities, like persons, have social lives” (Appadurai 1986, p. 3). The *enframing* that, according to Heidegger, characterises the attitude of the

¹ “[A]gency has become the buzzword of contemporary archaeological theory . . . a lingua franca – an ambiguous platitude meaning everything and nothing . . . there is little consensus about what “agency” actually means . . . nor has there been sustained consideration of basic methodological and epistemological issues as to make it applicable and appropriate to the premodern past” (Dobres and Robb 2000, p. 3).

Western individual towards the world as a “standing reserve” – a passive recourse to be controlled and manipulated for human ends – has no place and meaning in a number of ethnographic contexts with a very different understanding of what it is like to be a person, and what it means to engage the world (Heidegger 1977). For example, what could agency mean to a partible, composite and relationally constituted Melanesian person (Strathern 1988)? Clearly the idea of the isolated agent that acts upon the world, imposing shape and meaning upon inert matter, can hardly be accommodated in a Melanesian context where the categories of persons and things are inseparably distributed over biographical time and space.

Despite those examples, however, it has been primarily in the last two decades that the idea of decentralised agency has gained momentum across the social sciences. It is Actor-Network Theory (ANT) that has been particularly influential in this regard (Callon 1986; Law 1992, 1999, 2002; Latour 1994, 1999a, b, 2000, 2005). ANT can be defined as a semiotics of materiality that is symmetrical with respect to human and nonhuman agents (Law 1999, p. 4). Conceptualizing agency as variously distributed and possessed in relational networks of persons and things, ANT proposes that all entities participating in those networks should be treated analytically as of equal importance (Ashmore et al. 1994; Fuller 1994; Lee and Brown 1994). In other words, for ANT what we call actors or agents are essentially the products or effects of networks. That means that no primacy of the human actor – individual or collective – over the nonhuman actor can be accepted on a priori grounds. This may sound like yet another attempt to reconcile the two traditionally opposed poles of social theory, i.e. agency and structure, but in reality is something quite different. In drawing material things into the sociological fold the aim of ANT was not to overcome this contradiction but simply to ignore it, and develop what Latour calls a “bypassing strategy” (1999, pp. 16–17).

For example, to answer the question whether people or guns kill, we need to move beyond what is acceptable in either the materialist or the sociological account of activity (Latour 1999, p. 180). Both accounts start with essences, and essences result in antinomies, and antinomies are the reason that modernist theories fail to capture practice.

What does the gun add to the shooting? In the materialist account, *everything*: an innocent citizen becomes a criminal by virtue of the gun in her hand. The gun enables, of course, but also instructs, directs, even pulls the trigger . . . Each artifact has its script, its potential to take hold of passerby and force them to play a role in its story. By contrast, the sociological version . . . renders the gun a neutral carrier of will that adds nothing to the action, playing the role of a passive conductor, through which good and evil are equally able to flow (Latour 1999, p. 177).

What both accounts – materialist and sociological – fail to recognise is that agency “resides in the blind spot in which society and matter exchange properties” (Latour 1999, p. 190). Neither the isolated gun nor the isolated individual can bear the responsibility for the act of killing. The responsibility lies, on the one hand, in the way those two agents come together to construct a new hybrid

agent – the gunman – and on the other, in the socio-technical network that supports and makes possible such a meeting. Action involves a coalescence of human and nonhuman elements and as such the responsibility for action must be shared among them (Latour 1999, pp. 180–182). No distinctions between human and nonhuman entities can be sustained in terms of agency.

More will be said on ANT in subsequent chapters, such as those by Knappett and Watts. For similar perspectives, we might also consider Pickering's (1995) work on "the dance of agency" between humans and artefacts, Kaufmann's "la danse avec les choses" (1997), or Suchman's "sociomaterial agency" (2007). Cognitive science and philosophy have also advanced our understanding of the agency of artefacts in novel ways (see Hutchins 1995; Kirsh 1995; Clark 1997; Norman 1988). We might also very well turn to other domains such as human geography (Jones and Cloke 2002), political theory (Bennett 2004), economics (Lane and Maxfield 1997) and anthropology (Gell 1998; Layton 2003; Hoskins 2006). Many of these perspectives are more concerned with understanding agency as a situated process, rather than debating what or who is or is not an agent. This has, in part, been aided by a recognition that agency need not be coterminous with intentionality, which releases nonhumans into the process of agency. We hope to capture in this volume some of this move away from anthropocentric approaches that is happening across the social sciences. We also intend to show the worth of these new perspectives for archaeology, which has been slow on the uptake; despite the fact that the likes of Bourdieu and Giddens show little concern for material culture; much archaeological theory has remained faithfully wed to practice theory and structuration for 20 years. By using the term "material agency" we do not want to go to the other extreme and say that agency is material rather than human; it is more of a wake-up call, for social scientists and archaeologists, to encourage them to consider agency non-anthropocentrically, as a situated process in which material culture is entangled. Archaeology has the potential here to lead the way, as it deals with material culture "far more seriously and innovatively than do most social scientists" (Dobres and Robb 2000, p. 14).

Given this obvious focus on material culture in archaeology, and the acknowledgement that "material culture is clearly central to creating agents and expressing agency" (Dobres and Robb 2000, p. 14), it is surprising that the relationship between material culture and agency does not feature more prominently in current archaeological theory.² Strange as it might seem for a discipline that is in broad agreement on the "active" nature of material culture, archaeology remains, in our opinion, attached to an *anthropocentric* view of the world and by extension also of agency. The move towards a more active view of

² In a recent review, Dobres and Robb note that agency remains both "woefully under-theorised" and subject to remarkably "sparse methodological developments" (2005, p. 159). While Dobres and Robb stress in particular the need for robust methodologies in the study of agency (see also Joyce and Lopiparo 2005), we focus here more on the woeful under-theorisation.

material culture goes back 20 years or more; but this has been more a case of acknowledging the active rather than passive use of material culture by humans, rather than ascribing much dynamism to the artefacts themselves. What the active nature of material culture in its common usage seems to imply is essentially that *human* individuals, far from passively adapting to external systemic forces, are *actively using* material culture as an expressive symbolic medium for their social strategies and negotiations (Hodder 1982, 1986). In other words, the essence of the argument is that material culture may not simply reflect but also actively construct or challenge social reality, *on the necessary condition*, however, of human agency and intentionality. The above sounds too obvious to be wrong, and indeed this is precisely how material culture operates in many cases. However, this is only a part of the picture and, once you embrace it, leaves you with few chances to discover what the active nature of material culture really means.

There are some recent exceptions that do take the activeness of artefacts themselves more seriously, but they are surprisingly few (Chapman 2000; Chapman and Gaydarska 2007; Gamble 2007; Olsen 2003; Shanks 1998; Webmoor and Witmore 2005). Perhaps we should look beyond archaeology, particularly given this volume's interdisciplinary breadth; if we turn to the emerging field of "material culture studies", does this fare much better? A good place to assess this is the *Journal of Material Culture*, founded in 1996. A search reveals the word "agency" mentioned in the texts of 111 papers over the last 10 years. While this might at first glance appear to represent a strong commitment to issues of agency across a range of approaches in material culture studies, there is a remarkable tendency in the vast majority of these cases to engage with agency in a very particular way, influenced predominantly by Alfred Gell's 1998 volume "Art and Agency". Thus even here the concept of agency is not subject to very much depth or breadth of investigation.

These modest advances notwithstanding, we believe that approaches to material agency would benefit greatly from a much broader interdisciplinary involvement. This is one of the main justifications for this volume.

We begin with a contribution from Andy Clark, one of the foremost philosophers of mind to have contributed to discussions of agency. In "Where Brain, Body and World Collide" (Chapter 2), which was originally published in *Daedalus* in 1998, Clark sets out a possible conceptual background and philosophical conception of self and the mind with which many of the independent perspectives and case studies presented in this volume can be anchored. Drawing on recent theoretical and experimental work ranging from monkey finger control and mirror neurons to interactive vision and robotics, Clark puts together a powerful argument for the integration of perception, cognition and action in the study of mind and agency. For him the agency of the material world is simply the natural consequence of a "deeply interanimated unity" between perception and action.

This view of the situated brain and the extended mind is then explored by Malafouris using the potter's wheel as his example (Chapter 3). Developing his

argument for Material Agency, Malafouris argues that the only available starting point and obligatory point of passage for studying the emergence and determination of agency is that of material engagement. As with many other dimensions of the human mind, agency and intentionality should be understood as distributed, emergent and interactive phenomena rather than as subjective experiences. The clay on the potter's wheel should not be construed as the external passive object of the potter's intentional states, but as a functionally co-substantial component of the intentional character of the potting experience.

Our trio of cognitive science approaches to agency is completed by John Sutton. His chapter "Material agency, skills, and history" (Chapter 4) explores the issues of agency, interactivity, skill and distributed cognition from a broader historic and diachronic perspective. How can we identify the significant dimensions of cross-cultural and historical variation that will enable the construction of better typologies of distributed cognitive systems? What might be the role of cognitive archaeology to this end? How can archaeology and anthropology broaden the empirical and theoretical horizons of the cognitive sciences? Taking issue with a recent paper by one of us (Malafouris 2004), and focusing on the agency of things in the context of memory, Sutton proposes a number of fruitful directions for interdisciplinary research and the cross-fertilisation of ideas.

The next three chapters present diverse interdisciplinary perspectives on material agency, ranging from sociology and human geography, through to economics and human-computer interaction. The first two chapters, those of Law and Mol, and Jones and Cloke, provide us with an important correction to the potential pitfall in a non-anthropocentric approach of focussing only on artefacts and technologies. They show us that "material" in this case means "nonhuman", with their case studies tackling sheep and trees, respectively. Law and Mol (Chapter 5) take a Cumbrian sheep as their prospective agent, specifically a Cumbrian sheep in the midst of the foot and mouth crisis of 2001. Situating their approach within the "material semiotics" afforded by ANT, they stress that entities give each other being: they both act and are "enacted" (hence their chapter title "The Actor-Enacted"). They allow for four versions of a sheep – as a veterinary, epidemiological, economic and farming entity. These entail multiple practices that enact the sheep, making for a "sheep multiple". However, an enacted sheep is not a passive sheep; but it is difficult in the English language to circumvent this active/passive dichotomy. Though difficult to imagine a simultaneously active-passive agent, this is precisely what is required if we are to understand sheep as actors in this scenario. Furthermore, the practices that enact sheep-actors form complex webs that merge, interfere and pull apart, dynamically and indeterminately; we would be better advised, say Law and Mol, to ask *what* is happening, not *who* has done it.

The chapter by Jones and Cloke (Chapter 6) also falls under the broad ANT umbrella, with its presentation of trees as "actants", and agency as a "hybridised" phenomenon. They note, however, that ANT, in its discussions of hybridity, has been "biased towards technological rather than organic nonhuman entities". Not only this, but when the organic is included, it ventures only

as “far” as animals. Thus Jones and Cloke set out to remedy this imbalance in material agency by looking at other kinds of organism, in this case trees. They take three “tree places” in and around Bristol (South-west England) and assess the agency of trees in the making of these places over time. In so doing they distinguish between four kinds of agency, as routine, transformative, purposive and non-reflexive action.

In the approach taken by Harper, Taylor and Molloy (Chapter 7), we see the kinds of contemporary and future human–computer interaction that might easily tempt us into assigning agency to artefacts that appear somehow “intelligent”. Yet the authors are quick to state that the objects in question, though dynamic and interactive, do not have intelligence. This they assign to humans alone. They show how surfaces and containers in the home have particular affordances that can be augmented with digital technologies; but surfaces remain just surfaces, and do not become artefacts of intelligence through digital augmentation.

In a third section of the volume, we have four chapters dealing with archaeological contexts, albeit from interdisciplinary perspectives. Yarrow’s contribution (Chapter 8) is actually more an ethnography of recording processes during archaeological excavation. He looks in particular at context sheets, typically thought of as passive records of archaeological features but which, Yarrow argues, do have an enactive role within wider networks which include both human and nonhuman actants. There are interesting resonances with some of Harper’s earlier work on the role of paper documents in offices.

Chapter 9 by Knappett picks up on the theme of networks of human and nonhuman actants, seeking to develop more systematic means for addressing the network properties of agency, paradoxically neglected in Actor-*Network* Theory. Using examples from the Aegean Bronze Age, he examines the interactions between different categories of “material culture” – artefacts, images and texts. Early Aegean scripts are to some extent imagistic, and indeed they include images of artefacts; and some texts take artefactual form, inscribed on various kinds of support (e.g. tablets, sealings and pots). By looking at these various connections, we can gain glimpses into how artefacts, images and texts had a networked presence in the world of the Aegean Bronze Age; and that this presence formed a kind of material agency within which humans were deeply entwined.

Following this comes a contribution on the intertwined agency of water and stone in Irish Neolithic passage tombs (Chapter 10). Andrew Cochrane argues that the imagery carved into the stone façades of these tombs was “stimulated” by solutions, i.e. when wetted by rainwater and perhaps other kinds of liquid. That these images might come to life when wet, stimulating new perceptions and experiences among their viewers, is very suggestive; we might see some parallels with the paper of Harper et al., in that the affordances of surfaces and substances might be exploited and augmented by designers and/or users. Once again, it may not be easy to locate agency in this scenario; and we might follow Law and Mol’s advice in asking *what* is happening in Cochrane’s stone-and-solution scenario rather than *who* has done it.

The final Chapter 11 in this archaeological section comes from Chris Watts, who grapples with the complexities of Peircean semiotics, or ‘semeiotic’, in developing what one might call a situated semiotic approach to material agency. His use of Peirce’s notion of “synechism” conveys very persuasively the way in which “people and things are conjoined through the principle of semiotic mediation”. Watts grounds his sophisticated theoretical arguments in a case study drawn from the Late Woodland period of southwestern Ontario; comparing Iroquoian with Western Basin ceramic assemblages, he is able to show how these traditions bring into being two quite different agentic networks.

The volume concludes with two sets of concluding remarks, by Tim Ingold (Chapter 12) and Sander van der Leeuw (Chapter 13), respectively. The first of these is an ingenious critique of the ANT approach as applied to questions of material agency. Ingold creates a metaphor whereby “ANT” is one kind of creature in the forest, and “SPIDER” is another, and they strike up a philosophical dialogue. We know that ANT stands for “Actor Network Theory”, but SPIDER is a new acronym devised by Ingold that stands for “Skilled Practice Involves Developmentally Embodied Responsiveness”. Whereas ANT sees agency as coming about through the networked interlinkage of diverse objects, as for example in an ant colony, SPIDER rather sees agency as emergent from the skilled action-perception of an organism that inhabits a particular milieu, as for example a spider-with-web. The skilled action-perception of an organism as it moves through an environment creates a kind of mesh, but this does not mean that agency is distributed evenly between the organism and its mesh (which could be a fish with water, a butterfly with air, a spider with web or a potter with clay). Ingold thereby provides us with a thought-provoking challenge to ANT, particularly its tendency to see agency as a distributed network phenomenon, without sufficient attention to either the role of bodily perception and movement in the creation of the network, or the different qualities of the entities that are implicated in this process. Ingold’s wider aim is to get us thinking in terms of “meshworks” rather than “networks” (see Ingold 2007).

Sander van der Leeuw’s concluding comments (Chapter 13) have a different flavour. He draws on his experience of long-term, deep-rooted interdisciplinary projects in assessing the overall impact of the papers in this volume, both for archaeology and the social sciences more broadly. At the same time, he considers their potential for the development of a new approach to invention and innovation, topics of particular interest to van der Leeuw. He differentiates between a priori and a posteriori perspectives on innovation, the former engendering a proactive approach and the latter a reactive one. Unfortunately, the tendency is towards the latter in much of the scientific writing on innovation, which means that the inventive creativity leading to innovation is poorly understood. Van der Leeuw thus takes this volume’s emphasis on relationality (established through network thinking) to rethink the relationship between invention and innovation, with reference to traditional pottery making in Mexico and the Philippines. He sees the potential for inventive action as a feature of the dynamic “network”, composed of people, things, objects and

contexts, within which a potter is situated. Invention is a local process, only involving a relatively restricted network, whereas innovation is more widespread, cascading through a broader network.

Finally, we should emphasise that the aim of this volume is not to establish “material” or “nonhuman” agency as some kind of sustainable alternative to the idea of human agency. Rather, our primary intention is to provoke debate, with the term “material agency” intended as a challenge to the anthropocentrism inherent in existing approaches to agency. Some of the contributions to this volume argue that it is correct to centre our understanding of agency on the human; we are not averse to this when it is clearly and explicitly argued, as it is here, rather than implicitly assumed. We have not been looking for a consensus, but have actively sought divergent views. Such divergence is all the more likely when a project cuts across disciplinary lines, as here in the apparently unlikely combination of computer science, cognitive science, philosophy, human geography, and sociology, as well as anthropology and archaeology. Nevertheless, we do find that there are many converging lines of enquiry too, which here we have only begun to broach, and which, we believe, merit much fuller exploration in the future.

References

- Ahern, L., 2001, Language and Agency, *Annual Review of Anthropology* 20: 109–137.
- Appadurai, A., 1986. Introduction: Commodities and the Politics of Value, In *The Social Life of Things*, edited by A. Appadurai, pp. 3–63. Cambridge: Cambridge University Press.
- Archer, M. S., 1988, *Culture and Agency: The Place of Culture in Social Theory*. Cambridge University Press, Cambridge.
- Ashmore, M., Wooffitt, R., and Harding, S., 1994, Humans and Others, Agents and Things, *American Behavioral Scientist* 37(6): 733–740.
- Barrett, J., 2001, Agency, the Duality of Structure, and the Problem of the Archaeological Record. In *Archaeological Theory Today*, edited by I. Hodder, pp. 141–164. Polity Press, Cambridge.
- Bennett, J., 2004, The Force of Things: Steps Towards an Ecology of Matter, *Political Theory* 32(3): 347–372.
- Bourdieu, P., 1977, *Outline of a Theory of Practice*. Cambridge University Press, Cambridge.
- Bourdieu, P., 1990, *The Logic of Practice*. Polity Press, Cambridge.
- Callon, M., 1986, Some Elements for A Sociology of Translation: Domestication of the Scallops and the Fishermen of St-Brieuc Bay. In *Power, Action and Belief: A New Sociology of Knowledge?*, edited by J. Law, pp. 196–223. Sociological Review Monograph, Routledge and Kegan Paul, London.
- Chapman, J., 2000, *Fragmentation in Archaeology: People, Places and Broken Objects in the Prehistory of South-eastern Europe*. Routledge, London.
- Chapman, J., and Gaydarska, B., 2007, *Parts and Wholes: Fragmentation in Prehistoric Context*. Oxbow Books, Oxford.
- Clark, A., 1997, *Being There: Putting Brain, Body and World Together Again*. MIT Press, Cambridge, MA.
- Dobres, M.-A., and Robb, J., 2000, Agency in Archaeology: Paradigm or Platitude? In *Agency in Archaeology*, edited by M.-A. Dobres and J. Robb, pp. 3–17. Routledge, London.

- Dobres, M.-A., and Robb, J., 2005, "Doing" Agency: Introductory Remarks on Methodology, *Journal of Archaeological Method and Theory* 12(3): 159–166.
- Dornan, J. L., 2002, Agency and Archaeology: Past, Present, and Future Directions, *Journal of Archaeological Method and Theory* 9(4): 303–329.
- Emirbayer, M., and Mische, A., 1998, What is Agency? *American Journal of Sociology* 103(4): 962–1023.
- Fuller, S., 1994, Making Agency Count: A Brief Foray into the Foundations of Social Theory, *American Behavioral Scientist* 37(6): 741–753.
- Gallagher, S., 2007, The Natural Philosophy of Agency, *Philosophy Compass* 2 (j.1747-9991.2006.00067.x).
- Gamble, C., 2007, *Origins and Revolutions: Human Identity in Earliest Prehistory*. Cambridge University Press, Cambridge.
- Gardner, A., editor, 2004, *Agency Uncovered: Archaeological Perspectives on Social Agency, Power and Being Human*. UCL Press, London.
- Gell, A., 1998, *Art and Agency: Towards a New Anthropological Theory*. Clarendon Press, Oxford.
- Giddens, A., 1984, *The Constitution of Society: Outline of the Theory of Structuration*. Polity Press, Cambridge.
- Giddens, A., and Pierson, C., 1998, *Conversations with Anthony Giddens: Making Sense of Modernity*. Stanford University Press, Stanford, CA.
- Heidegger, M., 1977. The Question concerning Technology and Other Essays. New York (NY): Harper and Row.
- Hodder, I., 1982, *Symbols in Action. Ethnoarchaeological Studies of Material Culture*. Cambridge University Press, Cambridge.
- Hodder, I., 1986, *Reading the Past*. Cambridge University Press, Cambridge.
- Hoskins, J., 2006, Agency, Biography and Objects. In *Handbook of Material Culture*, edited by C. Tilley, W. Keane, S. Küchler, M. Rowlands and P. Spyer, pp. 74–84. Sage, London.
- Hutchins, E., 1995, *Cognition in the Wild*. MIT Press, Cambridge, MA.
- Ingold, T., 2007, *Lines: A Brief History*. Routledge, London.
- Johnson, M., 2000, Conceptions of Agency in Archaeological Interpretation. In *Interpretive Archaeology: A Reader*, edited by J. Thomas, pp. 221–227. Leicester University Press, London.
- Jones, O., and Cloke, P., 2002, *Tree Cultures: The Place of Trees and Trees in Their Place*. Berg, Oxford.
- Joyce, R. A., and Lopiparo, J., 2005, PostScript: Doing Agency in Archaeology, *Journal of Archaeological Method and Theory* 12(4): 365–374.
- Kaufmann, J.-C., 1997, *Le coeur à l'ouvrage: théorie de l'action ménagère*. Nathan, Paris.
- Kirsh, D., 1995, The Intelligent Use of Space, *Artificial Intelligence* 73: 31–68.
- Lane, D. A., and Maxfield, R., 1997, Foresight, Complexity and Strategy. In *The Economy as an Evolving Complex System 2*, edited by W. B. Arthur, S. N. Durlauf and D. A. Lane, pp. 169–198. Addison-Wesley, Reading, MA.
- Latour, B., 1994, Pragmatogonies: A Mythical Account of How Humans and Nonhumans Swap Properties, *American Behavioral Scientist* 37(6): 791–808.
- Latour, B., 1999a, *Pandora's Hope: Essays on the Reality of Science Studies*. Harvard University Press, Cambridge, MA.
- Latour, B., 1999b, On Recalling ANT. In *Actor Network Theory and After*, edited by J. Law and J. Hassard, pp. 15–25. Blackwell, Oxford.
- Latour, B., 2000, The Berlin eyK or How to do Words with Things. In *Matter, Materiality and Modern Culture*, edited by P. M. Graves-Brown, pp. 10–21. Routledge, London.
- Latour, B., 2005, *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press, Oxford.
- Law, J., editor, 1992, *A Sociology of Monsters: Essays on Power, Technology and Domination*. Routledge Sociological Review Monograph, London.

- Law, J., 1999, After ANT: Complexity, Naming and Topology. In *Actor Network Theory and After*, edited by J. Law and J. Hassard, pp. 1–14. Blackwell, Oxford.
- Law, J., 2002, *Aircraft Stories: Decentering the Object in Technoscience*. Duke University Press, Durham.
- Layton, R., 2003, Art and Agency: A Reassessment, *Journal of the Royal Anthropological Institute* (n.s.) 9: 447–464.
- Lee, N., and Brown, S., 1994, Otherness and the Actor Network, *American Behavioral Scientist* 37(6): 772–790.
- Malafouris, L., 2004. The Cognitive Basis of Material Engagement: Where Brain, Body and Culture Conflate. In *Rethinking Materiality: The Engagement of Mind with the Material World*, edited by E. DeMarrais, C. Gosden and C. Renfrew, PP. 53-62. McDonald Institute Monographs, Cambridge.
- Mauss, M., 1954. *The Gift*. London: Cohen and West.
- Munn, N., 1973. *Walpiri Iconography*. Ithaca and New York: Cornell University Press.
- Norman, D. A., 1988, *The Psychology of Everyday Things*. Basic Books, New York.
- Olsen, B., 2003, Material Culture after Text: Re-Membering Things, *Norwegian Archaeological Review* 36(2): 87–104.
- Pickering, A., 1995, *The Mangle of Practice: Time, Agency and Science*. Chicago University Press, Chicago.
- Shanks, M., 1998, The Life of an Artifact, *Fennoscandia Archeologica* 15:15–42.
- Smith, A. T., 2001, The Limitations of Doxa: Agency and Subjectivity from an Archaeological Point of View, *Journal of Social Archaeology* 1(2): 155–171.
- Sørensen, M. H., and Ziemke, T., 2007, Agents Without Agency? *Cognitive Semiotics* (special issue), 102–124.
- Strathern, M., 1988, *The Gender of the Gift*. University of California Press, Berkeley, CA.
- Suchman, L., 2007, *Human-Machine Reconfigurations: Plans and Situated Actions, 2nd edition*. Cambridge University Press, Cambridge.
- Webmoor, T., and Witmore, C., 2005, *Symmetrical Archaeology*, Metamedia, Stanford University @ traumwerk.stanford.edu:3455/Symmetry/home, accessed 26-11-06.

Chapter 1

Where Brain, Body and World Collide

Andy Clark

Introduction The brain fascinates because it is the biological organ of mindfulness itself. It is the inner engine that drives intelligent behaviour. Such a depiction provides a worthy antidote to the once-popular vision of the mind as somehow lying outside the natural order. But it is a vision with a price. For it has concentrated much theoretical attention on an uncomfortably restricted space; the space of the inner neural machine, divorced from the wider world which then enters the story only via the hygienic gateways of perception and action. Recent work in neuroscience, robotics and psychology casts doubt on the effectiveness of such a shrunken perspective. Instead, it stresses the unexpected intimacy of brain, body and world and invites us to attend to the structure and dynamics of extended adaptive systems – ones involving a much wider variety of factors and forces. Whilst it needs to be handled with some caution, I believe there is much to be learnt from this broader vision. The mind itself, if such a vision is correct, is best understood as the activity of an essentially *situated* brain: a brain at home in its proper bodily, cultural and environmental niche.

Software

Humans, dogs, ferrets – these are, we would like to say, mindful things. Rocks, rivers and volcanoes are not. And no doubt there are plenty of cases in between (insects, bacteria, etc....). In the natural order, clear cases of mindfulness always involve creatures with brains. Hence, in part, the fascination of the brain, understanding the brain looks crucial to the project of understanding the mind. But how should such an understanding proceed?

An early sentiment – circa 1970 and no longer much in vogue – was that understanding the mind depended rather little on understanding the brain. The brain, it was agreed, was in some sense the physical medium of cognition. But

A. Clark
Department of Philosophy, Edinburgh University, Edinburgh, UK
e-mail: Andy.Clark@ed.ac.uk

everything that mattered about the brain *qua* mind-producing engine turned not on the physical details but on the computational and information-processing strategies that the neural stuff (merely) implemented. There was something importantly right about this view, but something desperately wrong as well.

What was right was the observation that understanding the physical workings was not going to be enough; that we would need also to understand how the system was organised at a higher level in order to grasp the roots of mindfulness in the firing of neurons. This point is forcefully made by the cognitive scientist Brian Cantwell Smith who draws a parallel with the project of understanding ordinary computer systems. With respect to, for example, a standard PC running a tax-calculation program, we could quite easily answer all the “physiological” questions (using source code and wiring diagrams) yet still lack any real understanding of what the program does or even how it works.¹ To really understand how mental activity yields mental states, many theorists believe, we must likewise understand something of the computational/information-processing organisation of the brain. Physiological studies may contribute to this understanding. But even a full physiological story would not, in and of itself, reveal how brains work *qua* mind-producing engines.²

The danger, of course, was that this observation could be used as an excuse to downplay or marginalise the importance of looking at the biological brain *at all*. And so, in the early days of cognitive science, it was common to hear real neuroscience dismissed as having previous little to offer to the general project of understanding human intelligence. Such a dismissal, however, could not long be sustained (for a powerful critique, see Reeke & Edelman (1988)). For although it is (probably) true to say that a computational understanding is in principle independent³ of the details of any specific implementation in hardware (or wetware), the project of *discovering* the relevant computational description (especially for biological systems) is surely not.

One key factor is evolution. Biological brains are the product of biological evolution and as such often fail to function in the ways we (as human designers) might expect.⁴ This is because evolution is both constrained and liberated in ways we are not. It is constrained to build its solutions incrementally via a series of simpler but successful ancestral forms. The human lung, to give one example, is built via a process of tinkering (Jacob (1977)) with the swim bladder of the fish. The engineer might design a better lung from scratch. The tinkerer, by contrast, must take an existing device and subtly adapt it to a new role. From

¹ See B.C. Smith (1996), p. 148. Note that Smith’s worry, at root, concerns the gap between physiological and semantic or intentional questions.

² Thus we read, for example, that computational approaches make possible “a science of structure and function divorced from material substance [that] ... can answer questions traditionally posed by psychologists” (Pylyshyn (1986, p. 68).

³ See, for example, David Marr’s (1982) distinction between the levels of computation, algorithm and implementation.

⁴ See, for example, Simon (1982), Dawkins (1982), Clark (1997) Chapter 4.

the engineer's ahistorical perspective, the tinkerer's solution may look bizarre. Likewise, the processing strategies used by biological brains may surprise the computer scientist. For such strategies have themselves been evolved via a process of incremental, piecemeal, tinkering with older solutions.

More positively, biological evolution is *liberated* by being able to discover efficient but 'messy' or unobvious solutions: ones which may, for example, exploit environmental interactions and feedback loops so complex that they would quickly baffle a human engineer. Natural solutions (as we will later see) can exploit just about any mixture of neural, bodily and environmental resources along with their complex, looping and often non-linear interactions. Biological evolution is thus able to explore a very different solution space (wider in some dimensions, narrower in others) than that which beckons to conscious human reason.

There are, of course, ways around this apparent mismatch. Computationalists lately exploit so-called "genetic algorithms"⁵ that roughly mimic the natural process of evolutionary search and allow the discovery of efficient but loopy and interactive adaptive strategies. Moreover, hard neuro-scientific data and evidence is increasingly available as a means of helping expand our imaginative horizons so as to better appreciate the way real biological systems solve complex problems. The one-time obsession with the "software level" is thus relaxing, in favour of an approach which grounds computational modelling in a more serious appreciation of the biological facts. In the next section, I highlight some challenging aspects of such recent research – aspects that will lead us directly to a confrontation with the situated brain.

Wetware and Some Robots

Recent work in cognitive neuroscience underlines the distance separating biological and 'engineered' problem solutions and displays an increasing awareness of the important interpenetration – in biological systems – of perception, thought and action. Some brief examples should help fix the flavour.

As a gentle entry point, consider some recent work on the neural control of monkey finger motions. Traditional wisdom depicted the monkey's fingers as individually controlled by neighbouring groups of spatially clustered neurons. According to this story, the neurons (in Motor Area 1, or M1) were organised as a 'somatotopic map' in which a dedicated neural sub-region governed each individual digit, with the sub-regions arranged in latero-medial sequence just like the fingers on each hand. This is a tidy, easily conceptualised solution to the problem of finger control. But it is the engineer's solution, not (it now seems) that of Nature.

⁵ For a review, see Clark (1997) Chapter 5.

Marc Schieber and Lyndon Hibbard⁶ have shown that individual digit movements are accompanied by activity spread pretty well throughout the M1 hand area, and that precise, single digit movements actually require *more* activity than some multi-digit whole hand actions (such as grasping an object). Such results are inconsistent with the hypothesis of digit-specific local neuronal groups. From a more evolutionary perspective, however, the rationale and design are less opaque. Schieber (1990, p. 444) conjectures that the basic case, from an evolutionary perspective, is the case of whole hand grasping motions (used to grab branches, to swing, to acquire fruits, etc.) and that the fundamental neural adaptations are thus geared to the use of simple commands which exploit inbuilt synergies⁷ of muscle and tendon so as to yield such coordinated motions. The ‘complex’ coordinated case is thus evolutionarily basic and neurally atomic. The ‘simple’ task of controlling, for example, an individual digit represents the harder problem and requires more neural activity, namely the use of some motor cortex neurons to *inhibit* the synergetic activity of the other digits. Precise single digit movements require the agent to tinker with whole-hand commands, modifying the basic synergetic dynamics (of mechanically linked tendons, etc.) adapted to the more common task.

Consider next a case of perceptual adaptation. The human perceptual system can, we know (given time and training), adapt in quite powerful ways to distorted or position-shifted inputs. For example, subjects can learn how to coordinate vision and action while wearing lenses that invert the entire visual scene so that the world initially appears upside down. After wearing such lenses for a few days, the world is seen to flip over – various aspects of the world now appear to the subject to be in the normal upright position. Remove the lenses and the scene is again inverted until readaptation occurs.⁸ Thach et al. (1992) used a variant of such experiments to demonstrate the motor-specificity of some perceptual adaptations. Wearing lenses which shifted the scene *sideways* a little, subjects were asked to throw darts at a board. In this case, repeated practice led to successful adaptation,⁹ but of a motor-loop specific kind. The compensation did not “carry over” to tasks involving the use of the non-dominant hand to throw, or to an underarm variant of the visual overarm throw. Instead, adaptation looked to be restricted to a quite specific combination of gaze angle and throwing angle: the one used in overarm, dominant-hand throwing.

Something of the neural mechanisms of such adaptation is now understood. It is known, for example, that the adaptation never occurs in patients with

⁶ Schieber & Hibbard (1993).

⁷ The notion of synergy aims to capture the idea of links that constrain the collective unfolding of a system comprising many parts. For example, the front wheels of a car exhibit a built-in synergy which allows a single driver ‘command’ (at the steering wheel) to affect them both at once. Synergetic links may also be learnt, as when we acquire an automated skill, and may be neurally as well as brute-physiologically grounded. See Kelso (1995), pp. 38, 52.

⁸ For a survey of such experiments, see Welch (1978).

⁹ In this case, *without* any perceived shift in the visual scene.

generalised cerebellar cortical atrophy, and that inferior olive hypertrophy leads to impaired adaptation. On the basis of this and other evidence, Thach et al. speculated that a learning system implicating the inferior olive and the cerebellum (linked via climbing fibres) is active both in prism adaptation and in the general learning of patterned responses to frequently encountered stimuli. The more general lesson, however, concerns the nature of the perception-action system itself. For it increasingly appears that the simple image of a general purpose perceptual system delivering input to a distinct and fully independent action system is biologically distortive. Instead, perceptual and action systems work together, in the context of specific tasks, so as to promote adaptive success. Perception and action, on this view, form a deeply inter-animated unity.

Further evidence for such a view comes from a variety of sources. Consider, for example, the fact that the primate visual system relies on processing strategies that are not strictly hierarchic but instead depends on a variety of top-to-bottom and side-to-side channels of influence. These complex inner pathways allow a combination of multiple types of information (high-level intentions, low-level perception and motor activity) to influence all stages of visual processing. The Macaque monkey (to take one well-studied example) possesses about 32 visual brain areas and over 300 connecting pathways. The connecting pathways go both upwards and downwards (e.g., from V1 to V2 and back again) and side-to-side (between subareas in V1) – see, for example, Felleman and VanEssen (1991). Individual cells at ‘higher’ levels of processing, such as V4 (visual area 4) do, it is true, seem to specialise in the recognition of specific geometric forms. But they will each also respond, in some small degree, to many other stimuli. These small responses, spread unevenly across a whole population of cells, can carry significant information. The individual cells thus function not only as narrowly tuned single feature detectors but also as widely tuned filters reacting to a whole range of stimulus dimensions (see Van Essen & Gallant 9 (1994)). Moreover, the responses of such cells now look to be modifiable both by attention and by details of local task-specific context (Knierim & Van Essen (1992)).¹⁰

More generally, back projecting (corticocortical) connections tend, in the monkey, to outnumber forward ones, that is, there are more pathways leading from deep inside the brain outwards towards the sensory peripheries than vice versa.¹¹ Visual processing may thus involve a variety of criss-crossing influences which could only roughly, if at all, be described as a neat progression through a lower-to-higher hierarchy.

Such complex connectivity opens up a wealth organisational possibilities in which multiple sources of information combine to support visually guided

¹⁰ In a somewhat related vein, Caminiti et al. (1990) show that the directional preference of individual cells encoding reaching movement commands varies according to initial arm position. See also discussion in Jeannerod (1997), Chapter 1.

¹¹ Though much of the connectivity is reciprocal. See Van Essen & Anderson (1990), Churchland et al. (1991), p. 40.

action. Examples of such combinations are provided by Churchland, Ramachandran and Sejnowski (1994) who offer a neurophysiologically grounded account of what they term “interactive vision.” The interactive vision paradigm is there contrasted with the approaches that assume a simple division of labour in which perceptual processing yields a rich, detailed inner representation of the visual scene which is later given as input to the reasoning and planning centres which in turn calculate a course of action and send commands to the motor effectors. This simple image (of what roboticists call a “sense-think-act” cycle) is, it now seems, not true to the natural facts. In particular:

1. daily agent-environment interactions often do not require the construction and use of detailed inner models of the full visual scene;
2. low-level perception may “call” motor routines that yield *better perceptual input* and hence improve information pick-up;
3. real-world actions may sometimes play an important role in the computational process itself and
4. the internal representation of worldly events and structures may be less like a passive data-structure or description and more like a direct recipe for action.

Evidence for proposition 1 comes from a series of experiments in which subjects watch images on a computer screen. Subjects are allowed to examine an on-screen pictorial display. Then, as they continue to saccade around the scene (focusing first on one area, then another) small changes are made to the currently unattended parts of the display. The changes are made during the visual saccades. It is an amazing fact that, for most of the part,¹² quite large changes go unnoticed: changes such as the replacement of a tree by a shrub, or the addition of a car, deletion of a hat, etc. Why do such gross alterations remain undetected? A compelling hypothesis is that the visual system is not even attempting to build a rich, detailed model of the current scene but is instead geared to using frequent saccades to retrieve information *as and when it is needed* for some specific problem-solving purpose. This fits nicely with Yarbus’ classic (1967) finding that the pattern of such saccades varies (even with identical scenes) according to the type of task the subject has been set (e.g., to give the ages of the people in a picture, to guess the activity they have been engaged in, etc.). According to both Churchland et al. (1994) and Ballard (1991), we are prone to the illusion that we constantly command a rich inner representation of the current visual scene because we are able to perform fast saccades, retrieving information as and when required. (An analogy:¹³ a modern store may present the illusion of having a massive amount of goods stocked on the premises, because it always has what you want when you want it. But modern computer ordering systems can automatically count off sales and requisition new items so

¹² The exception is if subjects are told in advance to watch out for changes to a certain feature. See McConkie (1990), Churchland et al. (1994).

¹³ Thanks to David Clark for pointing this out.

that the necessary goods are available just when needed and barely a moment before. This fine-tuned ordering system offers a massive saving of on-site storage whilst tailoring supply directly to customer demand.)

Contemporary research in robotics avails itself of these same economies. One of the pioneers of “new robotics,” Rodney Brooks (see e.g., Brooks, 1991) coined the slogan, “the world is its own best model” to capture just this flavour. A robot known as Herbert (Connell, 1989), to take just one example, was designed to collect soft drink cans left around a crowded laboratory. But instead of requiring powerful sensing capacities and detailed advance planning, Herbert got by (very successfully) using a collection of coarse sensors and simple, relatively independent, behavioural routines. Basic obstacle avoidance was controlled by a ring of ultrasonic sound sensors that brought the robot to a halt if an object was in front of it. General locomotion (randomly directed) was interrupted if Herbert’s simple visual system detected a roughly table-like outline. At this point, a new routine kicks in and the table surface is swept using a laser. If the outline of a can is detected, the whole robot rotates until the can is centred in its field of vision. This simple physical action simplifies the pick-up procedure by creating a standard action-frame in which a robot arm, equipped with simple touch sensors, gently skims the table surface dead ahead. Once a can is encountered, it is grasped, collected and the robot moves on. Notice, then, that Herbert succeeds without using any conventional planning techniques and without creating and updating any detailed inner model of the environment. Herbert’s ‘world’ is composed of undifferentiated obstacles and rough table-like and can-like outlines. Within this world the robot also exploits its own bodily actions (rotating the ‘torso’ to centre the can in its field of view) so as to greatly simplify the computational problem involved in eventually reaching for the can. Herbert is thus a simple example both of a system that succeeds using minimal 12 representational resources and one in which gross motor activity helps streamline a perceptual routine (as suggested in proposition (2) above).

The interactive vision framework envisages a more elaborate natural version of this same broad strategy, namely the use of a kind of perceptuo-motor loop whose role is to make the most of incoming perceptual information by combining multiple sources of information. The idea here is that perception is not a passive phenomenon in which motor activity is only initiated at the endpoint of a complex process in which the animal creates a detailed representation of the perceived scene. Instead, perception and action engage in a kind of incremental game of tag in which motor assembly begins long before sensory signals reach the top level. Thus, early perceptual processing may yield a kind of proto-analysis of the scene, enabling the creature to select actions (such as head and eye movements) whose role is to provide a slightly upgraded sensory signal. That signal may, in turn, yield a new proto-analysis indicating further visuo-motor action and so on. Even whole body motions may be deployed as part of this process of improving perceptual pick-up. Foveating an object can, for example, involve motion of the eyes, head, neck and torso. Churchland et al. (1994, p. 44) put it well: “watching Michael Jordan play basketball or a group

of ravens steal a caribou corpse from a wolf tends to underscore the integrated, whole-body character of visuomotor coordination.” This integrated character is consistent with neuro-physiological and neuro-anatomical data which show the influence of motor signals in visual processing. There are – to take just two small examples – neurons sensitive to eye position in V1, V3 and LGN (lateral geniculate nucleus), and cells in V1 and V2 that seem to know in advance about planned visual saccades (showing enhanced sensitivity to the target¹⁴).

Moving on to proposition (3) (that real-world actions may sometimes play an important role in the computational process itself), consider the task of distinguishing figure from ground (the rabbit from the field, or whatever). It turns out that this problem is greatly simplified using information obtained from head movement during eye fixation. Likewise, depth perception is greatly simplified using cues obtained by the observers own self-directed motion. As the observer moves, close objects will show more relative displacement than farther ones. That is probably why, as Churchland et al. (op.cit., p. 51) observe, head bobbing behaviour is frequently seen in animals: “a visual system that integrates across several glimpses to estimate depth has computational savings over one that tries to calculate depth from a single snapshot.”

And so to proposition (4): that the neural representation of worldly events may be less like a passive data structure and more like a recipe for action. The driving force, once again, is computational economy. If the goal of perception and reason is to guide action (and it surely is, evolutionary speaking), it will often be simpler to represent the world in ways rather closely geared to the kinds of actions we want to perform. To take a simple example, an animal that uses its visual inputs to guide a specific kind of reaching behaviour (so as to acquire and ingest food) need not form an object-centred representation of the surrounding space. Instead, a systematic metrical transformation (achieved by a point-to-point mapping between two topographic maps) may transform the visual inputs directly into a recipe for reaching out and grabbing the food. In such a set-up, the animal does not need to do any computational work on an action-neutral inner model as to plan a reaching trajectory. The perceptual processing is instead tweaked, at an early stage, in a way dictated by the particular use to which the visual input is dedicated. This strategy is described in detail in P.M. Churchland’s (1989, ch. 5) account of the “connectionist crab,” in which research in Artificial Neural networks¹⁵ is applied to the problem of creating efficient point-to-point linkages between deformed topographic maps.

In a related vein, Maja Mataric of the MIT Artificial Intelligence Laboratory has developed a neurobiologically inspired model of how rats navigate their

¹⁴ See Churchland et al. (1994), p. 44, Wurz & Mohler (1976).

¹⁵ For an accessible introduction, see P.M. Churchland (1995).

environments. This model exploits the kind of layered architecture¹⁶ also used in the robot Herbert. Of most immediate interest, however, is the way the robot learns about its surroundings. As it moves around a simple maze, it detects landmarks which are registered as a combination of sensory input and current motion. A narrow corridor thus registers as a combination of forward motion and short lateral distance readings from sonar sensors. Later, if the robot is required to find its way back to a remembered location, it retrieves¹⁷ an inter-linked body of such combined sensory and motor readings. The stored ‘map’ of the environment is thus immediately fit to act as a recipe for action, since the motor signals are part of the stored knowledge. The relation between two locations is directly encoded as the set of motor signals that moved the robot from one to the other. The inner map is thus *itself* the recipe for the necessary motor actions. By contrast, a more classical approach would first generate a more objective map which would then need to be *reasoned over* in order to plan the route.

The Mataric robot (which is based on actual rat neurobiology – see McNaughton & Nadel (1990)) and the connectionist crab exemplify the attractions of what I call ‘action-oriented representations’ (Clark, 1997, p. 49): representations that describe the world by depicting it in terms of possible actions.¹⁸ This image fits in nicely with several of the results reported earlier, including the work on monkey finger control and the motor loop specificity of ‘perceptual’ adaptation. The products of perceptual activity, it seems, are not always action-neutral descriptions of external reality. They may instead constitute direct recipes for acting and intervening. We thus glimpse something of the shape of what Churchland et al. (1994, p. 60) describe as a framework that is “motocentric”; rather than “visuocentric.”

As a last nod in that same direction, consider the fascinating case of so-called “mirror neurons” (DiPellegrino et al. (1992)). These are neurons, in monkey ventral premotor cortex, that are action-oriented, context-dependent and implicated in both self-initiated activity and passive perception. They are active both when the monkey observes a specific action (such as someone grasping a food item) and when the monkey performs the same action, where sameness implies not mere grasping but the grasping of a food item (see also Rizzolatti et al. (1996)). The implication, according to the psychologist and neuroscientist Marc

¹⁶ This is known as a ‘subsumption’ architecture, because the layers each constitute a complete behaviour-producing system and interact only in simple ways such as by one layers subsuming (turning off) the activity of another or by one layer’s co-opting and hence ‘building-in’ the activity of another (see Brooks (1991)).

¹⁷ By a process of spreading activation amongst landmark encoding nodes – see Mataric (1991).

¹⁸ Such representations bear some resemblance to what the ecological psychologist J.J. Gibson called “affordances,” although Gibson himself would reject our emphasis on inner states and encodings. For an affordance is the potential of use and activity that the local environment offers to a specific kind of being: chairs afford sitting (to humans) and so on. See Gibson (1979). The philosopher Ruth Millikan has developed a nice account of action-oriented representation under the label ‘pushmipullyu representation’ – see Millikan (1995).